FROM PIE IN THE SKY TO ROI

GETTING A.I. DEPLOYMENT RIGHT IN THE FINANCIAL SERVICES INDUSTRY

JULY 2026

QUINLAN &ASSOCIATES

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The financial services industry is entering a new phase of artificial intelligence ("A.I.") adoption, with global financial institutions ("FIs") investing an estimated USD 76 billion in A.I.-related initiatives in 2025, a number we forecast to reach USD 132 billion by 2030.

Despite vast investments and institutional ambition, most financial institutions struggle to translate A.I. initiatives into measurable business value and clear returns on investment. In many cases, deployment activity has outpaced the development of the organisational foundations required to support it. Strategy is frequently backward engineered around available technology, risk is considered only after design and development have progressed, and existing operating models are not equipped to govern, implement, and scale A.I. effectively.

To overcome these barriers, financial institutions should adopt an integrated roadmap across three interconnected layers:

- **Strategise:** A.I. should be treated as an enabling capability to deliver the enterprise strategy rather than as a standalone technology agenda. Institutions should balance outside-in awareness of emerging technological possibilities with an inside-out assessment of strategic priorities, customer journeys, and operational workflows. This requires systematically identifying friction points across existing processes, determining where A.I. can generate meaningful value, and applying non-negotiable eligibility thresholds before prioritising viable use cases. Beyond improvements to existing processes, there is a broader opportunity for institutions to fundamentally reimagine workflows around A.I. capabilities, which can potentially unlock new ways of delivering value.
- **Govern:** Governance should be established before material resources are committed. Traditional approaches that treat governance as a downstream approval process are proving insufficient for dynamic A.I. environments. Financial institutions should implement systematic eligibility gates to exclude use cases that do not meet minimum viable requirements. Governance must also extend beyond use-case approval into the underlying workflow, defining how A.I. operates within controlled boundaries through runtime controls, human checkpoints, and more. When designed effectively, early governance can serve not only as a compliance safeguard, but also as a competitive edge for more effective and confident deployment by establishing clear boundaries.
- **Enable:** Once strategic priorities and permissible guardrails have been established, institutions must develop the technological and operational capabilities required for execution. From a technology perspective, this includes selecting an appropriate model architecture and developing the surrounding A.I. engineering capabilities, including prompt engineering, context management, and harnesses to deliver reliable, secure, and cost-effective outcomes. Enablement also requires a fundamental shift in people, roles, and ways of working. As A.I. becomes embedded across institutional workflows, frontline staff will increasingly direct A.I. outputs, middle managers will orchestrate mixed teams of humans and A.I. agents, and senior leaders will be responsible for setting strategic direction, boundaries, and organisational design. A.I. transformation therefore cannot rely solely on a traditional top-down technology rollout as meaningful value is often identified closer to the frontline, where real domain knowledge and operational realities are most visible.

Ultimately, the quality of transformation will depend on how effectively institutions codify human expertise into the context, rules, controls, and validation logic that guide A.I.-enabled workflows. Technology can automate and augment execution, but it cannot independently determine how an institution should operate, which risks it should accept, or what constitutes an appropriate outcome. Those decisions must remain grounded in institutional strategy, domain expertise, and clearly assigned accountability.

As financial institutions continue to advance their A.I. ambitions, particularly towards more autonomous and agentic models, they should conduct a structured readiness diagnostic covering strategy, governance, technology, people, and processes. This diagnostic can identify critical gaps, prioritise interventions, and establish a practical 12 to 24-month roadmap for scalable adoption. By strengthening these foundations before accelerating deployment, institutions can shift from fragmented experimentation towards controlled, repeatable, and measurable value creation.

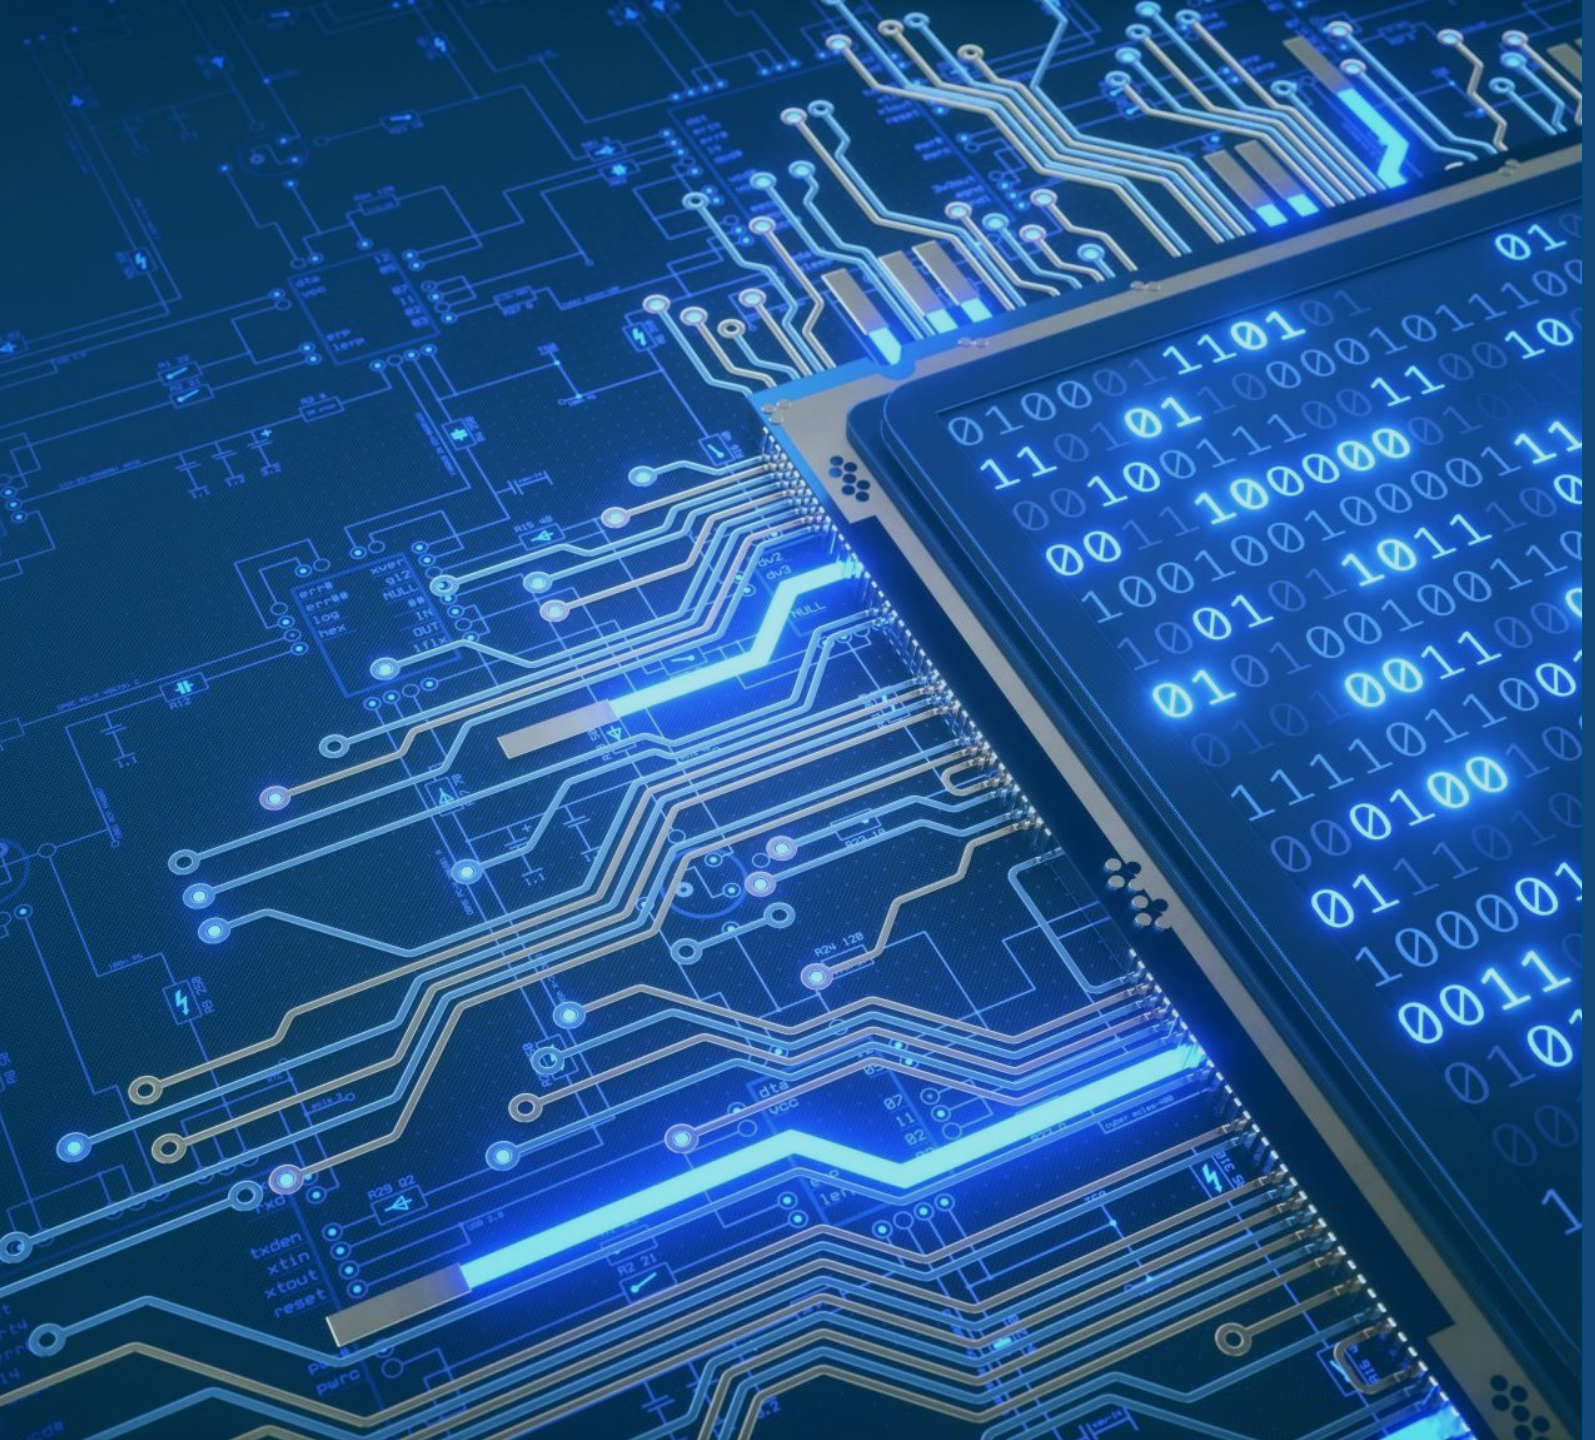

SECTION 1

LARGE INVESTMENTS, DISAPPOINTING RETURNS

QUINLAN &ASSOCIATES

QUINLAN &ASSOCIATES

GLOBAL A.I. SPENDING ACROSS THE FINANCIAL SERVICES INDUSTRY

The global financial services industry invested USD 75 billion in A.I.-related initiatives in 2025, with this figure projected to reach USD 132 billion by 2030; APAC commitment and allocation to A.I. continue to lag their peers in North America and Europe

Global A.I. Spending Size

Financial Services, USD Billion, 2025-30E

1.2 – 2.7x

Expected increase in A.I.-related budget allocation over the next five years

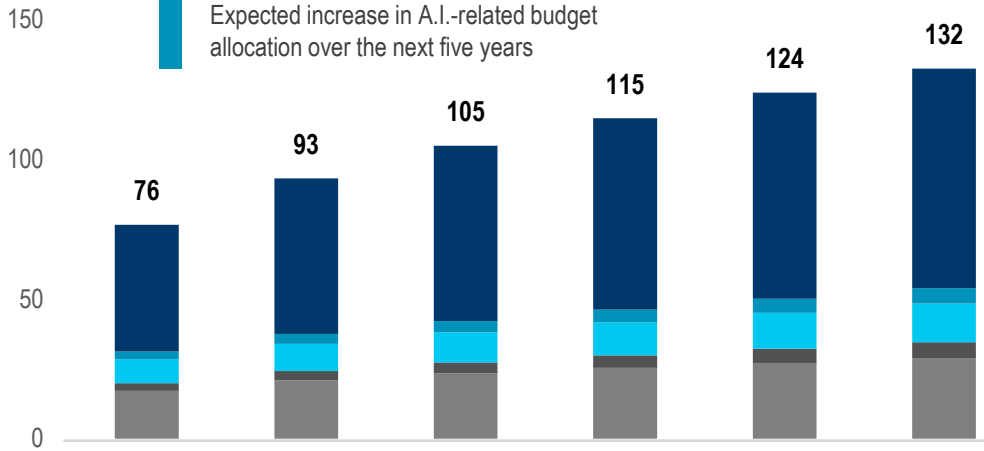

	2025	2026E	2027E	2028E	2029E	2030E	CAGR
Banking[1]	45	56	63	69	74	79	12%
Sell-side Markets[2]	3	3	4	5	5	5	15%
Buy-side Markets[3]	9	10	11	12	13	14	10%
Market Infra & Payments[4]	3	3	4	4	5	6	16%
Insurance	17	21	23	25	27	29	11%
Total A.I. Spending	**76**	**93**	**105**	**115**	**124**	**132**	**12%**

A.I. Spending Allocation

% A.I. of IT Expenditure, Average, 2025

BANK OF AMERICA

- **Total IT Budget:** USD 13.5 billion
- **A.I. Allocation[5]:** USD 4 billion

J.P.Morgan

- **Total IT Budget:** USD 20 billion
- **A.I. Allocation:** USD 2 billion

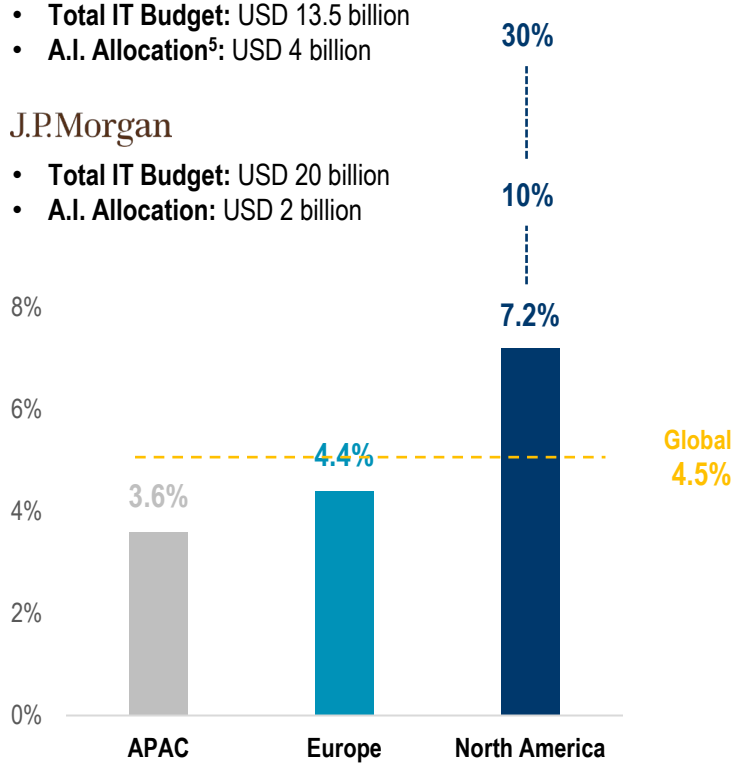

[1]Retail, commercial, and corporate banks; [2]Investment banks and brokerages; [3]Asset managers, wealth managers, private banks, and high frequency trading firms; [4]Securities exchanges, digital assets exchanges, and payment providers; [4]New initiatives, including AI

Source: Public disclosures (e.g., annual statements, financial positions, etc.), Expert interviews, Oxford Economics, Coalition Greenwich, Nvidia, Quinlan & Associates estimation

A.I. DEPLOYMENT AREAS

Across banking, capital markets, and insurance, A.I. is increasingly viewed as a source of business value across both external client-facing propositions and internal operating functions

Current and Expected Deployment Areas

By Sector, 2025

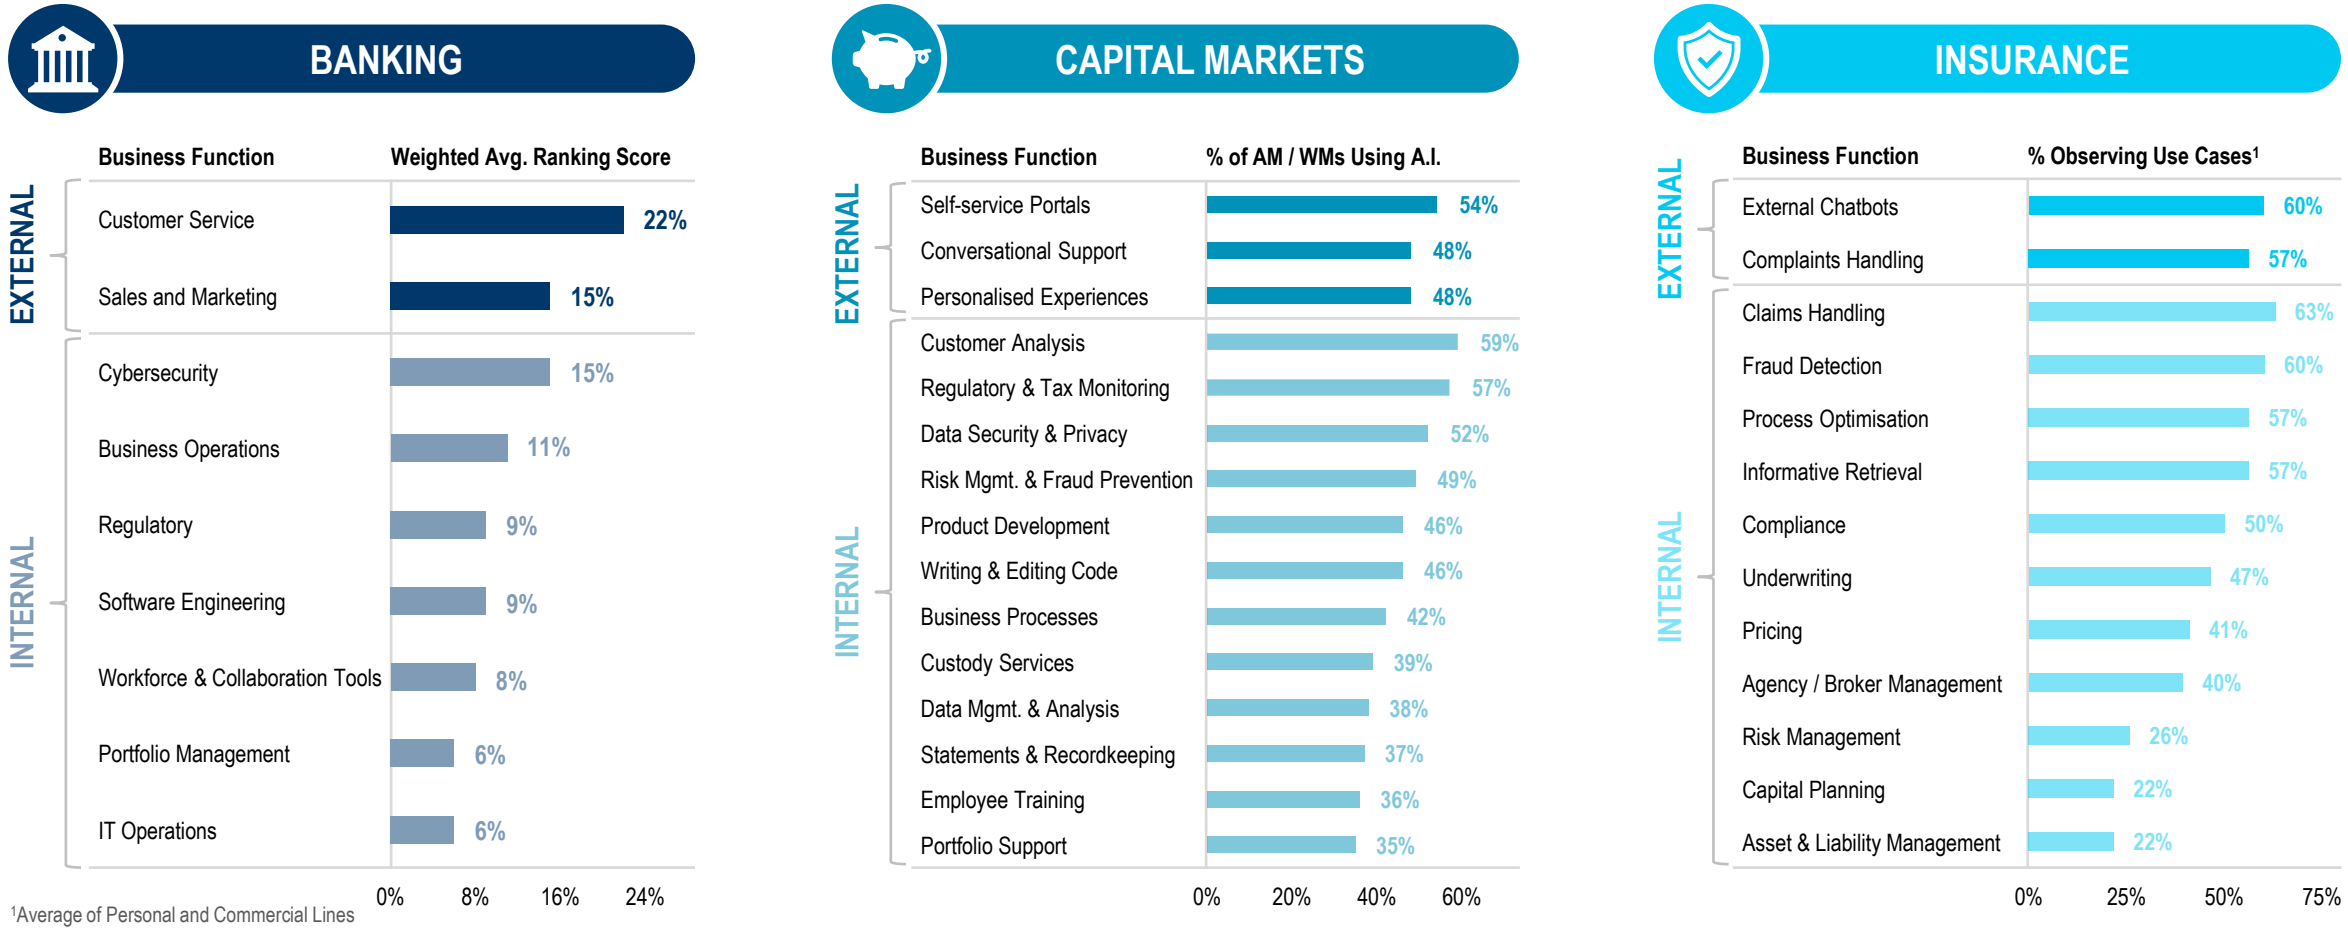

[1]Average of Personal and Commercial Lines

Source: Grant Thornton, Infosys, International Association of Insurance Supervisors, Quinlan & Associates analysis

STATE OF A.I. ADOPTION IN FINANCIAL SERVICES

While financial institutions are accelerating A.I. deployment, adoption remains concentrated in assistive use cases, including chatbots and copilots, while agentic A.I. has yet to achieve widespread adoption

Types of A.I. Applications

By Level of Complexity & Adoption

	Chatbot	Retrieval Assistant	Copilot	A.I. Agent	Agentic Operating System
Description	A rule-based or A.I.-powered conversational interface that handles inbound customer queries through text or voice	An A.I. tool that ingests large volumes of proprietary documents, data, and surfaces relevant information on demand	An A.I. assistant embedded into the business applications professionals use to better perform daily tasks with enterprise context	An A.I. system that can interpret a defined goal, plan the required steps, and execute tasks on behalf within predefined constraints	A coordinated network of specialised A.I. agents handling complex, multi-step workflows end-to-end across functions
Adoption Status	**Deployed Nearly Everywhere** The world's 10 largest banks by assets have all deployed A.I. chatbots	**Becoming Table Stakes** Idea generation, data processing, and market analysis are top 3 adoption areas for AMs[1]	**Most Institutions are Here** Microsoft's introduction of Copilot in Nov 2023 marked the beginning of accessible Copilots	**High Intent, Limited Deployment** 42% of finance professionals use or assess agentic A.I. capabilities	**Aspirational** Following single-agent pilots, leading firms are exploring the adoption of multi-agent system
Early Adopter[2]	ally BANK. **2015** Ally Assist	STATE STREET **2017** Quantextual Idea Lab	Swiss Re **2024** Copilot rollout	Allianz **2025** Project Nemo	BNY **2026** Multi-agent System
Other Examples[2]	Schroders Schroders Go TIGER BROKERS TigerGPT GEICO Kate	J.P.Morgan Contract Intelligence Vanguard Client Summaries AVIVA Medical Summary Tool	DBS CSO Assistant Aladdin by BlackRock Auto Commentary 中信证券 CITIC SECURITIES Copilot	BANKOFSINGAPORE SOWA Goldman Sachs Devin	AIG Agent Orchestration Layer

[1]Asset managers; [2]Examples are non-exhaustive

Source: Public disclosures, AITechtonic, FinText, Microsoft, NVIDIA, OpenAI, Reinsurance News, Quinlan & Associates analysis

WIDELY ADOPTED YET CHALLENGING TO CAPTURE VALUE

More than 70% of traditional financial institutions struggle to translate A.I. investment into measurable value and clear ROI, often because strategy is backward engineered around technology, risk is considered too late, and the operating model is not fit for purpose

Adoption Status of A.I.
2026, n=149, Traditional financial institutions across all sizes

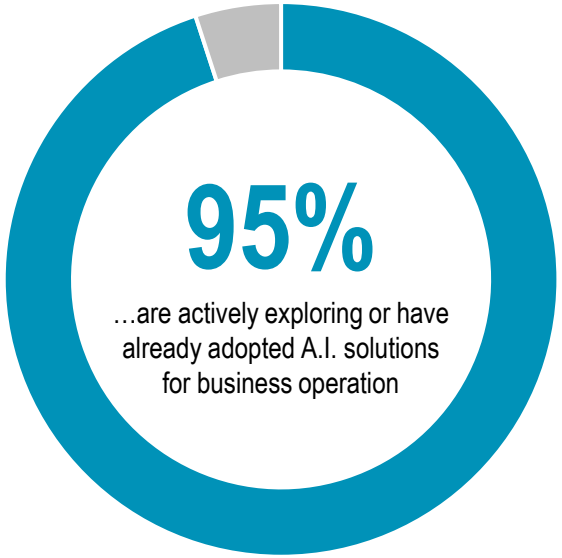

Most financial institutions are actively adopting A.I. Solutions

VS.

Challenging to Capture A.I. Value
2026, n=101, Traditional financial institutions across all sizes

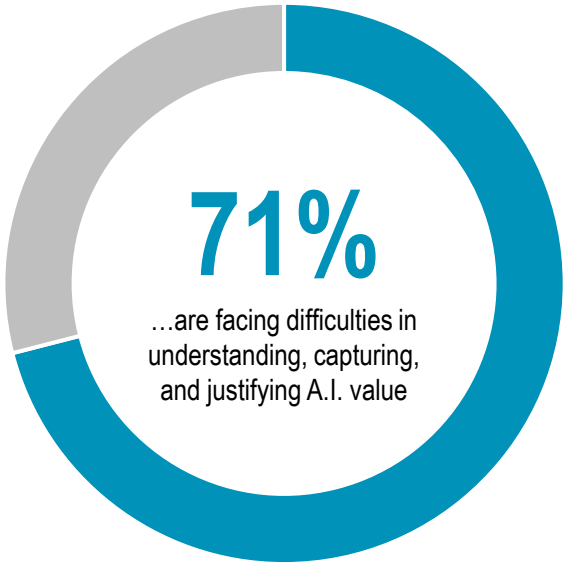

Yet more than 2/3 are having difficulties in capturing the A.I. Value

1

Strategy is backward engineered

Organisations often select A.I. tools before defining the business problems, workflows, and outcomes they should address. Strategy is then adapted to fit the technology, rather than technology being selected to support business priorities

2

Risk is treated as an afterthought

Risk is often assessed after use cases and solutions have been selected, rather than being embedded into their design. This creates late-stage remediation, inconsistent controls, and unclear ownership across the A.I. lifecycle

3

Operating models are not fit-for-purpose

Existing operating models often lack the proper governance, measurement frameworks, workforce capabilities, and change-management disciplines required to scale A.I. As a result, A.I. fails to be embedded across the organisation

Source: Cambridge Centre for Alternative Finance, Quinlan & Associates analysis

1. STRATEGY IS BACKWARD ENGINEERED

The rush to deploy A.I. in response to 'fear of missing out' can generate visible momentum and media attention, but initiatives launched without strategy-led design often fail to improve underlying business performance

Problem
Backward Engineering

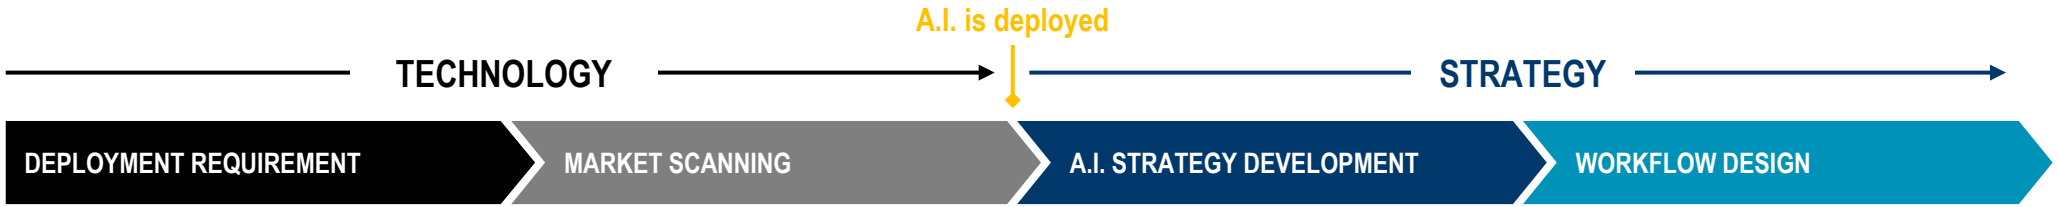

Current State (Technology First)	DEPLOYMENT REQUIREMENT	MARKET SCANNING	A.I. STRATEGY DEVELOPMENT	WORKFLOW DESIGN
Questions	**How quickly can we buy and deploy A.I. solutions to avoid falling behind?**	What A.I. tools and capabilities are currently available in the market?	**How should our A.I. strategy fit around the technology we have deployed?**	**Which workflows can be automated, augmented, or redesigned using A.I.?**
Driver Behind the Weak Starting Point	Competitive anxiety creates urgency to deploy quickly, often before the organisation has assessed value, feasibility, operational readiness, or control	Market hype and vendor activity create pressure to scan available tools before the organisation has clearly defined the business problem it wants to solve	Once technology is selected, the strategy is reverse-engineered, creating a risk that the organisation adapts its business needs to technology capabilities	The organisation risks redesigning the workflow around the technology, rather than selecting technology to support the workflow that needs to be redesigned
Result	**Pressure for "More A.I." vs. Scalable A.I.** Senior management is under pressure to deliver rapid A.I.-driven ROI, but often lacks visibility into scaling constraints → **Pressure translates into demand for more projects and faster delivery, equating activity with impact, flowing downwards**		**Technology defines business needs** Organisations deploy technology before defining the business problem, target workflow, or desired outcome → **Strategy and processes are redesigned around the A.I. tools, creating fragmented solutions that optimise existing activities rather than fundamentally improving business performance**	

Source: Quinlan & Associates analysis

2. RISK IS TREATED AS AN AFTERTHOUGHT

As LLM-based solutions ingest increasingly sensitive data and assume greater autonomy (e.g., recommending actions or making decisions on users' behalf), risk must be embedded early to avoid control gaps, costly remediation, and project cancellation

Technology Deployment Journey
Problem

IDEATION *Map out business requirements, which are used to determine the core functional requirements of the prospective solution*	**Proposition Development**	Determining the specific value-add of the suggested solution
	Features Identification	Laying out features of the solution that address the specific business problem
	UI / UX Mapping & Verification	Conceptualising the user experience and aligning it with end-users
VALIDATION *Demonstrate that the solution concepts have the potential to be applied and carried out in a real-world situation*	**Feasibility Assessment**	Understanding capability gaps that are required to achieve the future state
	Resource Requirement	Proposing timeline, budget, and human resources required for the initiative
	Risk Identification	Identifying potential risks that may arise prior to development and implementation
DEVELOPMENT *Carry out the development of functionalities and test it out with the end users to collect and incorporate feedback, iteratively*	**Prototype Development**	Building out core functionalities that aim to address the business needs
	Agile Feedback Loop	Seeking and incorporating the end-users' feedback under a controlled environment
	System Integration	Integrating with existing systems and deploying the solution

The Reality of A.I. Adoption
% of A.I. Initiatives[1], Dec 2025

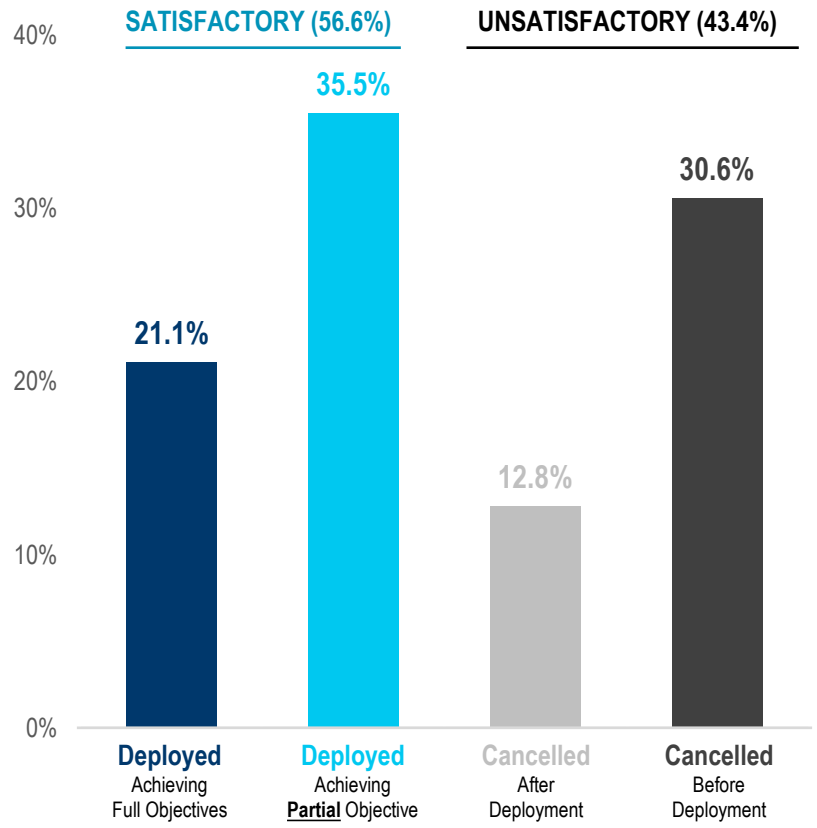

[1]Out of 25,289 A.I. initiatives developed by 400 surveyed banks; [2]1,649 business leaders worldwide across 14 sectors, including 15% in banking and 4% in insurance

Source: Celonis, Infosys, Quinlan & Associates analysis

QUINLAN &ASSOCIATES

3. OPERATING MODELS ARE NOT FIT FOR PURPOSE

Across the A.I. lifecycle, different blockers emerge at different stages, but they consistently trace back to one structural root cause: the operating model is no longer equipped to support A.I. adoption at scale

A.I. Lifecycle / Value Chain

Operating Model-related Issues, Illustrative

	DATA INGESTION	WORKFLOW DEFINITION	MODEL SELECTION	SOLUTION DEPLOYMENT	OUTPUT EVALUATION	SCALING OPERATION
Common Problems Among Institutions	**Fragmented and Unreliable Inputs** A.I. systems may ingest inconsistent, outdated, or over-permissioned data across repositories, resulting in hallucination risks	**Persistent Workflow Inefficiencies** Organisations fail to unlock the full potential of A.I., as workflows are inefficiently adapted around A.I. capabilities (vs. redesigned to address business needs)	**Trade-off in Capability vs. Explainability** Default to less capable, pre-approved models, prohibiting the use of outperforming, frontier models at the outset due to model risk requirements	**Fragmented A.I. Deployment** When capability is handed to different functions, teams assemble their own agents, resulting in duplicative efforts and fragmented deployment	**Inconsistent A.I. Output Quality** A.I. outputs are generated efficiently but lack business value due to inconsistent evaluation processes and unclear quality benchmarks	**Limited Enterprise-Wide A.I. Adoption** A.I. tools remain underutilised across the organisation, preventing widespread adoption and limiting the return on investment
Current State of Operating Model	**Poor Data Governance** × Customer data sitting in multiple repositories across business lines (e.g., private banking and wealth) × Inconsistent ownership and classification of data	**Technology-first Approach** × A.I. is deployed as a solution looking for a problem, resulting in workflows being reshaped around A.I. vs. being purpose-built around workflow requirements	**Lack of Robust Evaluation** × Model selection is driven primarily by risk avoidance and binary approval mechanisms × The same model is evaluated and applied the same way across use cases	**Decentralised Build** × A lack of central coordination where development and deployment of A.I. solutions across functions are done in isolation × Limited interoperability	**No Shared Standards** × Lack of standardised workflows, human review checkpoints, and use-case specific KPIs to assess and validate A.I. outputs	**Limited Capability Building** × Existing training and change management efforts have not evolved alongside changing employee roles × Staff do not have the skills and confidence to use the A.I. tools

Current operating models are no longer fit to support A.I. adoption at scale – addressing these problems is therefore important for:

- Faster Issue Resolution
- Reduced Operational, Regulatory, and Model-related Risks[1]
- Improved Compliance
- Accelerated Time to Market

[1]Includes hallucination risks, data leakage, etc.

Source: Quinlan & Associates analysis

BUILDING BLOCKS FOR A.I. DEPLOYMENT

To overcome these barriers and realise value at scale, financial institutions must address three interconnected layers that collectively form a practical roadmap for A.I. adoption: strategise, govern, and enable

Building Blocks
Strategy & Operating Model

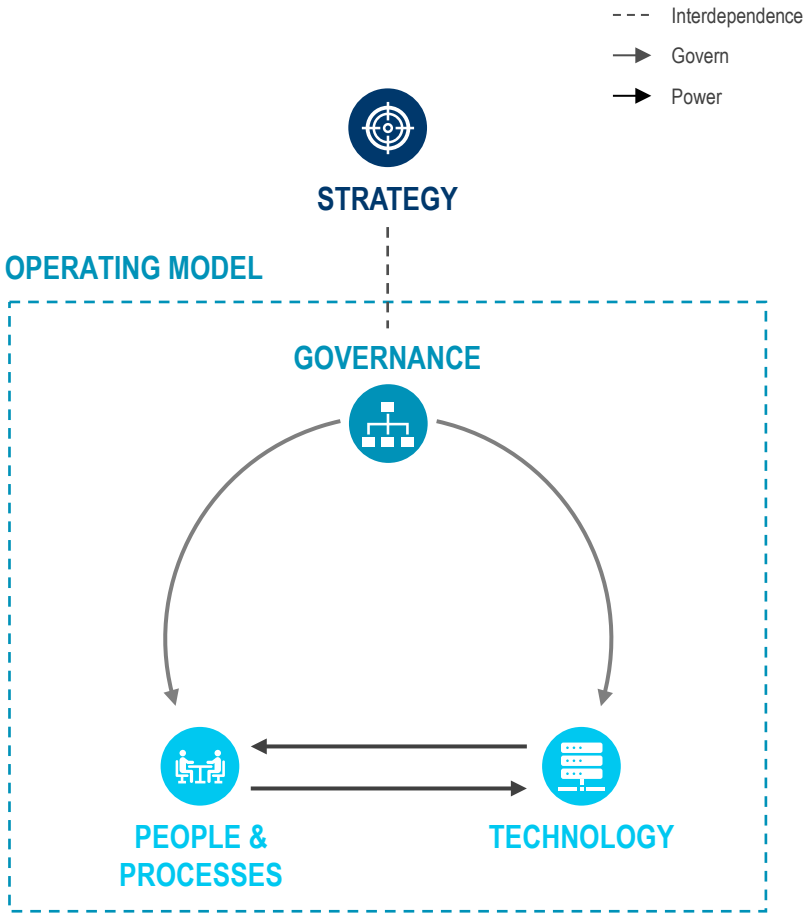

STRATEGISE
Define Objective

Is our A.I. strategy and plan surrounding the use cases and models applied to the right problems, and do we have a clear line of sight from use cases to business outcomes?

DESCRIPTION

- Translate overarching business objectives into a prioritised portfolio of use cases
- Identify use cases through workflow-led assessment
- Adopt a structured prioritisation framework rather than relying on ad-hoc selection

GOVERN
Establish Controls

Is A.I. pre-selected and operating within defined risk parameters and delivering reliable, high-quality outputs by following appropriate controls?

- Establish gating criteria (e.g., data readiness and risk thresholds) prior to use case prioritisation
- Centralise governance backed by a systematic playbook with clear accountability
- Redesign controls for A.I.-enabled workflows and technology operations, executed by teams
- Embed KPIs across the end-to-end A.I. lifecycle

ENABLE
Drive Better Usage

Do we have the appropriate capabilities in place to adopt A.I. at scale and in a way that sustains value over time?

- Select the right models and approaches for each use case, not just the most capable
- Build A.I. literacy across the organisation, from leaders to front-line staff, to reflect their evolving roles
- Establish ways of working that embed A.I. and governance into day-to-day processes, not as an add-on

Source: Quinlan & Associates analysis

SECTION 2

STRATEGISE

QUINLAN &ASSOCIATES

DEFINING THE A.I. STRATEGY

A.I. value is maximised when organisations balance outside-in awareness of emerging possibilities with a rigorous inside-out approach that embeds A.I. into core workflows and strategy-led priorities

A.I. Strategy
Double-edge Perspectives

Ideal State (Strategy First)	STRATEGY-FIRST APPROACH	WORKFLOW DESIGN	USE CASE PRIORITISATION	MARKET SCANNING
Key Questions	**How should A.I. support the organisation's strategy and strengthen its operating capabilities?**	**Which workflow steps should be automated, augmented, redesigned, or kept human-led?**	**Which use cases are worth pursuing, based on measurable value, feasibility, and strategic relevance?**	**Which available A.I. tools and capabilities are best suited to enable our strategic and tactical priorities?**
Why This Is a Better Starting Point	Starts with the organisation's strategic priorities, not the technology market, ensuring A.I. is positioned as an "enabling capability" that allows the organisation to achieve its intended business outcome	Forces the organisation to understand the workflow before deploying A.I and to avoid rushing into deployment, ensuring each use case is clearly defined (e.g., judgement-based, risk-sensitive, etc.)	Ensures the organisation focuses on use cases that are valuable, implementable, and strategically relevant, rather than spreading effort across attractive but low-impact A.I. experiments	Allows A.I. tools to be evaluated only after the organisation has clarified its strategy, automation needs, priority use cases, and tactical implementation requirements
	DEFINE BUSINESS OBJECTIVES	**TRANSFER DOMAIN EXPERTISE**	**PLACE STRATEGIC BETS**	**SELECT THE RIGHT TECHNOLOGY**

START WITH STRATEGY…THEN REDESIGN THE WORKFLOW

Ensure that A.I. adoption is tied to the overarching strategy first and foremost, and review existing workflows to spot operational friction (i.e., strong automation targets), applying A.I. as a process engineering lever

→ ***Tailored opportunities to pain points and value creation***

BE INFORMED BY THE MARKET…NOT BE DRIVEN BY IT

Look both within and outside the financial sector, defining what is technically and commercially possible by leveraging what is already proven in the market

→ ***Short time-to-market based on best practices***

This ensures that market-proven A.I. capabilities are selectively embedded into workflows and operating models to generate actual value

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

STRATEGY-FIRST APPROACH

Financial institutions should first determine whether and where A.I. supports the overarching business strategy, establishing clear directional intent before identifying use cases where adoption can create meaningful value

Strategic Alignment

Existing Strategy and A.I. Capabilities (Illustrative)

Type of Institutions	RETAIL / COMMERCIAL BANKS	ASSET / WEALTH MANAGERS	SECURITIES BROKERAGES	INSURANCE FIRMS
Strategic Objective *(Illustrative)*	**Rationalise the Footprint** Streamline the bank's structure and operations by exiting non-core markets and service lines	**Position for Wealth Transfer** Ensure smooth asset transition and reinvent current workflows to align with the values and behaviours of the next generation	**Deepen Retail Penetration** Deepen monetisation avenues with retail clients by reinventing the "do-it-yourself" approach to investment	**Expand Protection Services** Introduce innovative products to protect clients and their assets from a range of emerging risks

	IDENTIFY EMERGING RISKS	DEVELOP BUSINESS CASES	DESIGN PRODUCT & PRICING	OFFLOAD AGGREGATE RISKS	LAUNCH & MARKET PRODUCTS	PROCESS FUTURE CLAIMS
Journey	Pinpoint market gaps and test if the new risks are insurable	Build formal business cases to establish the target market	Set premium pricing and write the policy guidelines	Manage the insurer's exposure and obtain regulatory approvals	Train agents and brokers for marketing and process claims	Set up robust workflows to review and pay out claims
Relevant A.I. Capabilities *(Example)*	**Horizon Scanning** A.I. can continuously scrape unstructured data (e.g., litigation, research) for signals of new risks	**Persona Testing** A.I. can simulate client personas and stress-test if and how the product may be exploited	**Policy Drafting** A.I. can draft precise definitions and exclusions before comparing them with existing policies	**Pitch Creation** A.I. can create collateral, such as impact projects, that are presented to reinsurers and regulators	**Sales Enablement** A.I. chatbots and employee copilots can inform customers and agents more intuitively	**Pre-processing** A.I. can automate document processing and provide preliminary opinions on each claim
Alignment Assessment	Highly Aligned	Somewhat Aligned	Highly Aligned	Somewhat Aligned	Highly Aligned	Highly Aligned

Source: Quinlan & Associates analysis

WORKFLOW DESIGN (1/2) – REVIEW EXISTING PROCESSES

Within the relevant functional areas, institutions should identify operational frictions and bottlenecks across end-to-end workflows, establishing a structured basis for applying A.I. and redirecting resources towards higher-value activities

Use Case Development

A.I. Deployment

	Description	Implications of Inaction
DEFINE	**STANDARD OPERATING PROCEDURE / WORKFLOW** Map the end-to-end workflow across business lines, including inputs, outputs, decision points, approvals, controls, and hand-offs	*Automating fragmented, undocumented processes, resulting in inefficiencies and weakened controls*
ASSESS	**MANUAL GRIND AND PAIN POINTS** Identify workflow bottlenecks and manual-intensive activities by assessing frequency, effort, turnaround time, rework, error rates, and control sensitivity	*Deploying A.I. based on assumptions or hype rather than real pain points and business needs that generate value*
MATCH	**RISK SUITABILITY ASSESSMENT** Evaluate whether each activity is suitable for automation, augmentation, or continued human execution, taking into account regulatory requirements and model risk / impact	*Applying A.I. to unsuitable use cases where risks exceed potential value / benefits (e.g., compliance, op. risks, etc.)*
REDESIGN	**A.I.-INTEGRATED WORKFLOW** Revisit and update the end-to-end workflow to incorporate A.I., defining new roles, controls, escalation paths, and oversight as needed	*Resources intended to be 'freed up' end up correcting A.I.'s mistakes due to an unclear, outdated process*

Manual vs. High-Value Tasks

Illustrative, Front Office and Back Office

FRONT OFFICE

(Private Banker)

× Client Meeting Notes	✓ Needs Discovery
× Portfolio Summaries	✓ Tailored Advisory Provision
× Product Information Retrieval	✓ Client Relationship Management
× Internal Documentation	✓ Commercial Prioritisation
× Routine Enquiry Response	✓ Estate Planning Coordination

Free up Resources for →

BACK OFFICE

(HR Recruiter)

× CV Screening Support	✓ Candidate Interview
× Candidate Outreach Draft	✓ Workforce Planning / Advice
× Interview Scheduling	✓ Hiring Judgement
× Feedback Summaries	✓ Talent Event Participation
× Candidate Sourcing	✓ Candidate Relationship Building

Free up Resources for →

× **Manual Tasks Led / Augmented by A.I.** ✓ **Higher-Value Human-led Tasks**

Source: Quinlan & Associates analysis

WORKFLOW DESIGN (2/2) – EXPLORE NEW WAYS OF WORKING

While adopting a workflow-led approach is essential, ongoing developments in A.I. capabilities also present financial institutions with a unique opportunity to completely reimagine their workflows

A.I. Adoption Pathways
Possible Pathways

START (REQUEST)
The user submits a specific business request

1 ROUTE
The request is captured, logged, and assigned to the appropriate team

2 PROCESS
Specialist addresses the case by putting together data and applying judgement

3 VALIDATE
The proposed output is checked for completeness, accuracy, and eligibility

4 APPROVE
An accountable decision-maker provides final authorisation and sign-off

FINISH (DELIVERY)
The user receives the outcome

PATHWAY 1: AUGMENTATION
Supporting Human Processes

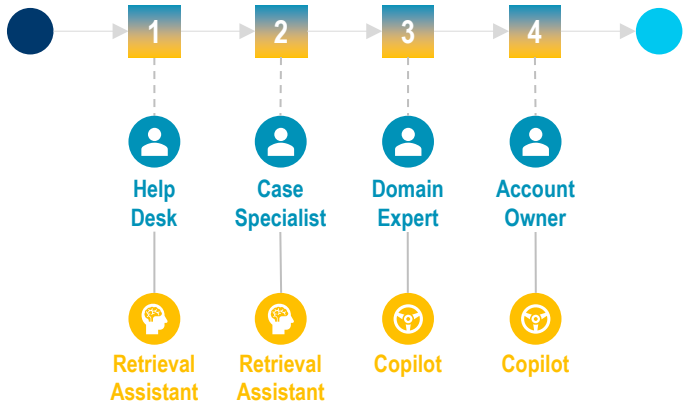

PATHWAY 2: OPTIMISATION
Automating Human Processes

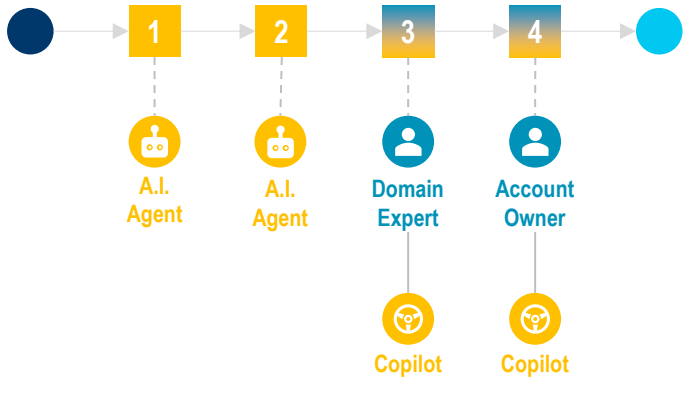

PATHWAY 3: RE-INVENTION
Reimagining the Entire Processes

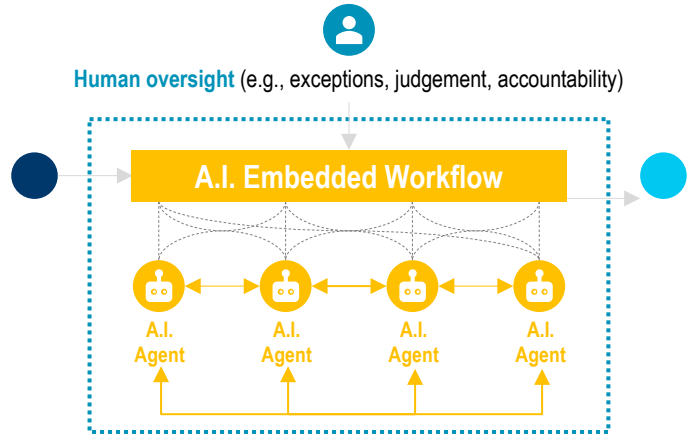

Source: Quinlan & Associates analysis

USE CASE PRIORITISATION

A minimum eligibility threshold ensures that institutional resources are focused on viable and compliant use cases, while a robust prioritisation framework determines the optimal delivery sequence to maximise impact

Methodology Overview

Illustrative

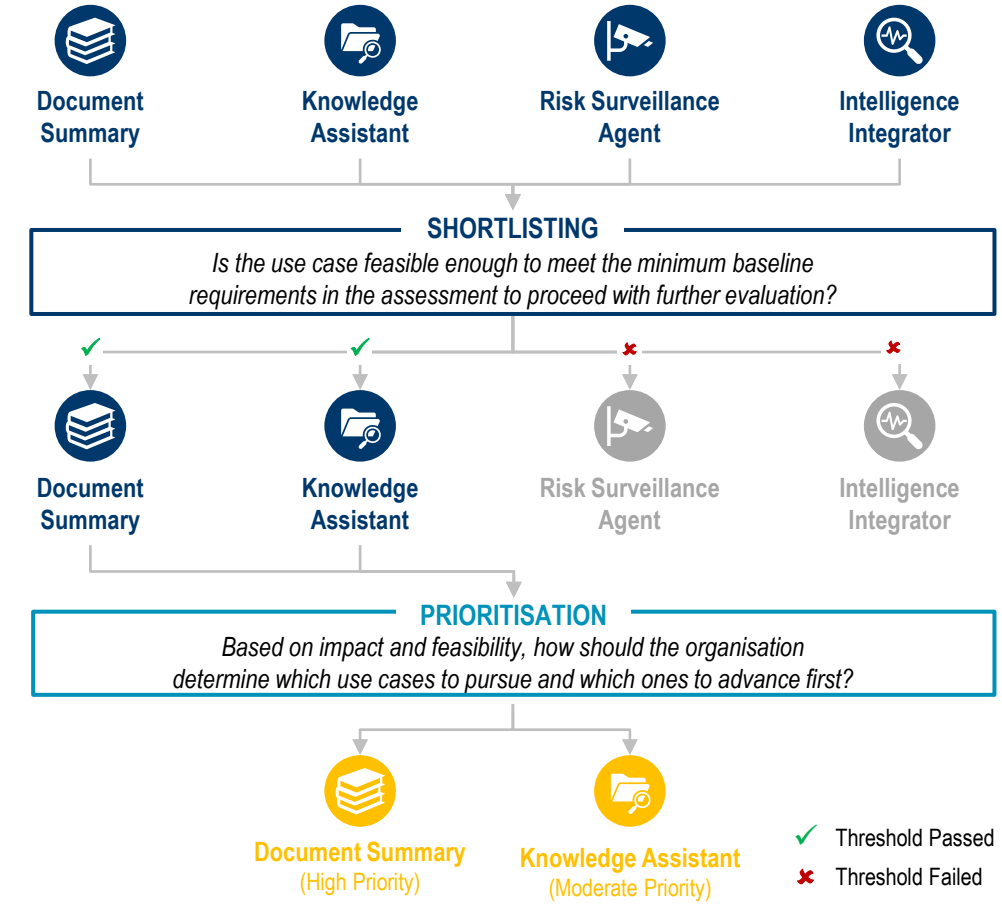

	1	2
Stages	**SHORTLISTING*** **(Eligibility Gate)**	**PRIORITISATION (Delivery Ordering)**
Description	A pass / fail screening stage that filters out use cases that are not viable for a financial institution to pursue	A subsequent comparative assessment stage that ranks shortlisted use cases to determine what to do first, next, and later
Sample Criteria	• Data readiness • Risk sensitivity threshold	• Value impact magnitude • Value impact reach • Time-to-market
Goal	Ensures only feasible and compliant use cases enter the pipeline, **working on what is actually buildable and permissible**	Ensures that financial institutions are building towards what matters first to **maximise impact and momentum**

Using a minimum shortlisting threshold ensures organisations focus their time and effort on viable and compliant use cases, while a robust prioritisation framework optimises delivery sequencing to drive maximum impact

*The concept of eligibility gate is further elaborated on Slide 21 as part of the "Govern" section

Source: Quinlan & Associates analysis

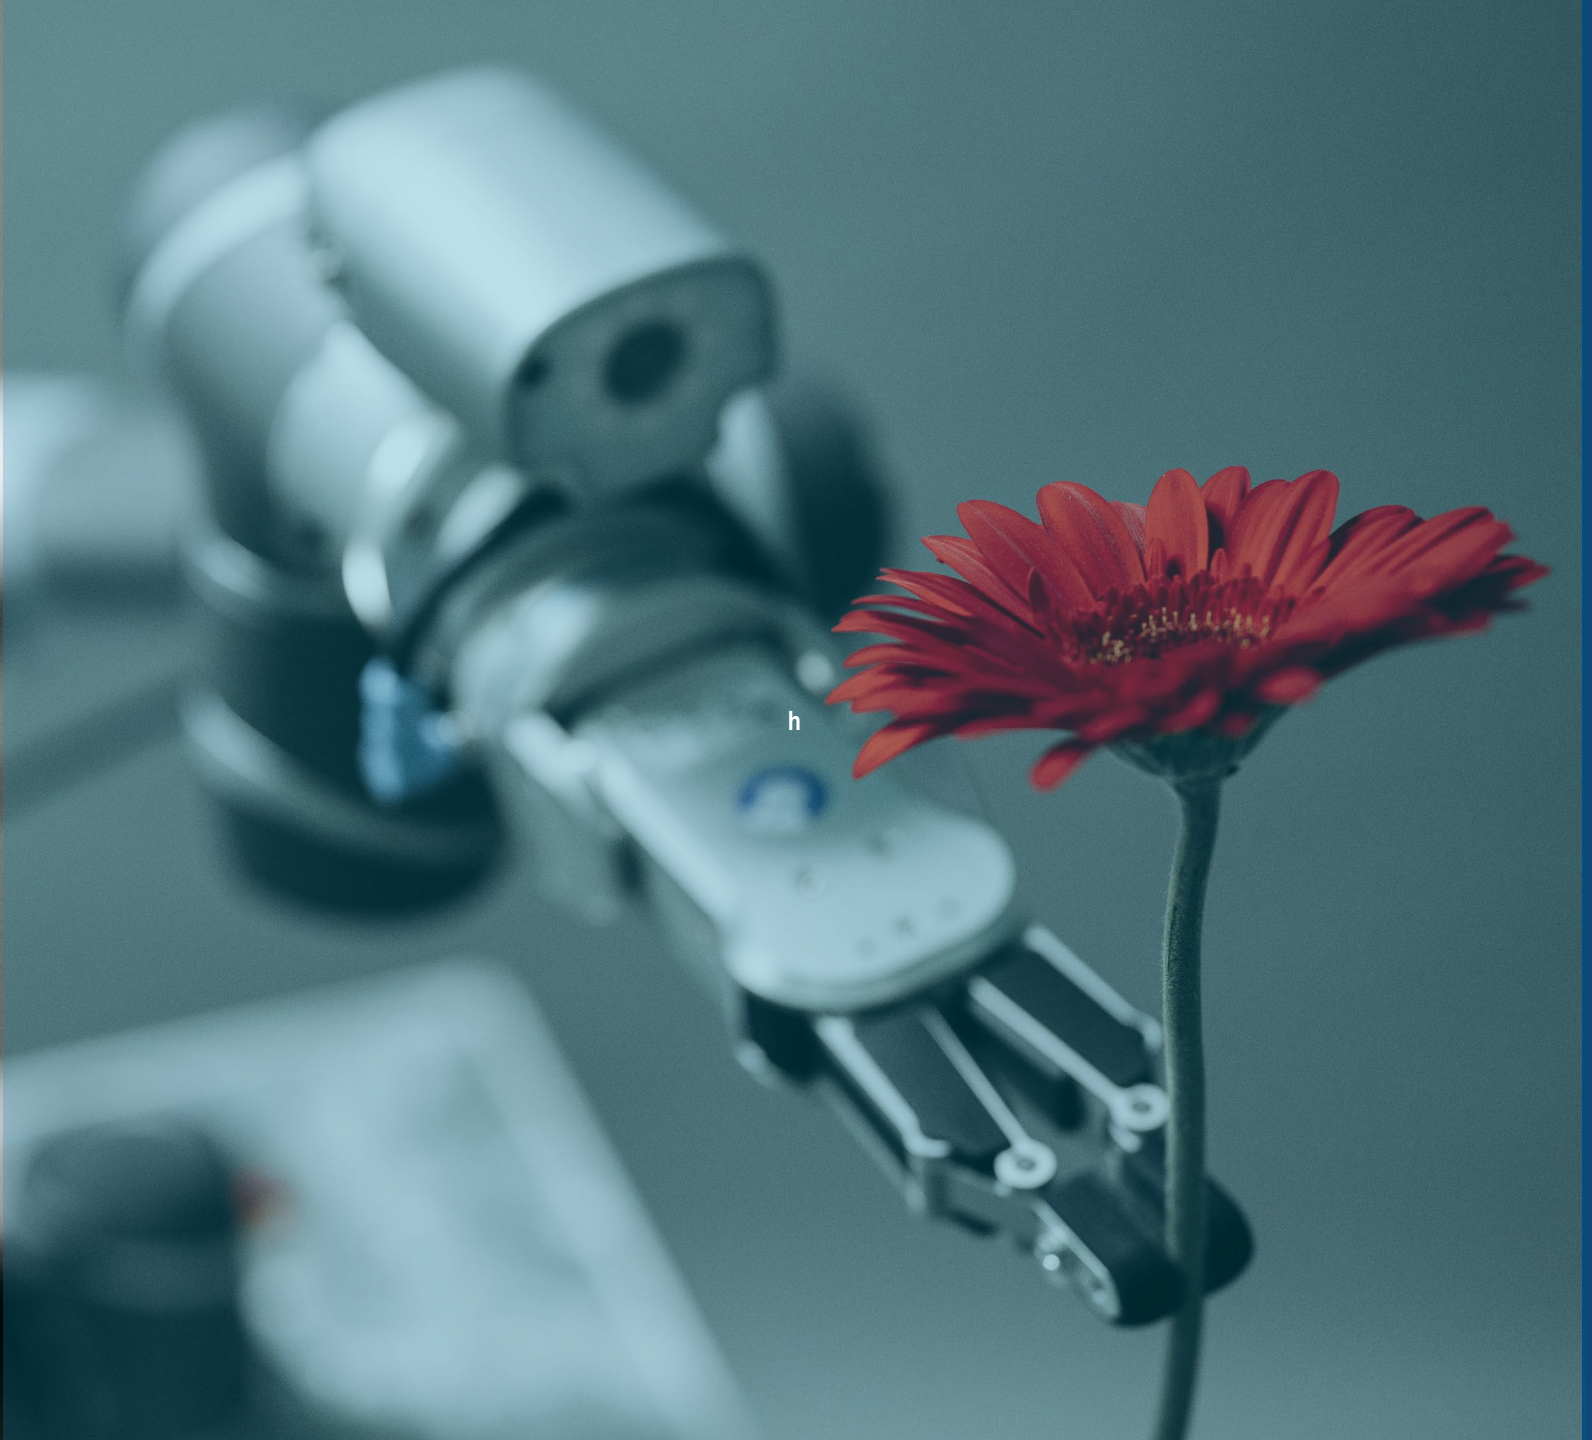

SECTION 3

GOVERN

QUINLAN
&ASSOCIATES

A.I. GOVERNANCE FRAMEWORK

As the A.I. landscape evolves and industry practices remain in flux, establishing governance provides both a compliance safeguard and a source of competitive advantage

Governance Considerations

Key Factors

	Current State	Target State
1 STRUCTURE	**DECENTRALISED** A.I. governance is currently dispersed across functions with limited decision power, resulting in fragmented enforcement and oversight	**CENTRALISED** A centralised decisioning layer with real authority (e.g., central coordination body) is needed to ensure firm-wide visibility, coordination, and enforcement of A.I. priorities
2 ASSESSMENT	**INCONSISTENT JUDGMENT** A.I. initiatives are still assessed case-by-case without a consistent evaluation logic, leading to subjective sign-offs and unclear outcomes	**SYSTEMATIC PLAYBOOK** A standard playbook should guide a central governing body in systematically assessing use case feasibility (e.g., risk assessment) and impact using objective criteria
3 ACCOUNTABILITY	**UNCLEAR ROLES** Ownership across A.I. governance is often unclear, with weakly defined RACI structures and inconsistent escalation pathways	**ENFORCEABLE ACCOUNTABILITY** Clear decision rights must define who proposes, validates, and approves, ensuring governance is not just advisory but enforceable

Systemic Playbook (Illustrative)	Use Case Charter
STEP 1: Intake Is the use case sufficiently defined for formal assessment?	• Strategic relevance & business problem • Ownership & stakeholders • Scope & intended users
STEP 2: Eligibility & Prioritisation Should the institution invest further efforts in developing the use case?	• Data availability assessment • Risk and control viability • Feasibility & impact assessment
STEP 3: PoC Approval Should controlled experimentation begin?	• PoC hypothesis and scope • Test design and success criteria • Pre-approved data & risk restrictions
STEP 4: Pilot Approval Should the solution be tested in a live under controlled environment?	• Validated functional solution • Pilot operating model • Pilot control framework
STEP 5: Production Approval Is the solution ready for business-as-usual deployment?	• Risk and control approval • Business-as-usual operating readiness • Production resilience and monitoring
STEP 6: Scale Approval Should the solution be expanded across the organisation?	• Adoption and operating evidence • Realised business benefits • Scalability assessment
STEP 7: Periodic Review Should the use case continue, be modified, restricted, or retired?	• Performance and risk outcomes • Economic and operational impact • Continued relevance and fitness

Elaborated on the next slide

Source: Quinlan & Associates analysis

ELIGIBILITY BEFORE PRIORITISATION

As part of a systematic A.I. playbook, financial institutions should establish non-negotiable eligibility gates for initial shortlisting, only applying feasibility and impact criteria to prioritise use cases that have passed those gates

Shortlisting and Prioritisation Framework

Illustrative

		ELIGIBILITY GATE: INITIAL SHORTLISTING		DELIVERY ORDERING: PRIORITISATION OF SHORTLISTED USE CASES	
	Criteria	**Threshold**	**Rationale**	**Weight**	**Rationale for Weighting**
FEASIBILITY	**Data Readiness:** Is there sufficient data available, and is it structured and of quality to train and run an A.I. model?	>80%	…of the required data is digitised, structured, and owned. Without this, no use case can be built, so the threshold is a hard pass / fail check	15%	Foundational requirement, as poor data quality often prevents A.I. use cases from progressing beyond concept
FEASIBILITY	**Risk Sensitivity:** Does the use case touch a process that is subject to regulatory scrutiny and pose material risks?	100%	…of explainability and audit trail requirements can be met, which is a gated requirement for regulated financial institutions	15%	Critical for financial institutions, where customer, conduct, regulatory, model, cyber, and reputational risks are material
FEASIBILITY	**Resource Needs:** How much monetary and human resources are needed for the proposed use case?	If a use case does not pass the eligibility gate, it should not proceed to further assessment, ensuring that time and resources are not spent on unsuitable opportunities		10%	Important for feasibility, but should not outweigh strategic relevance, risk, or value creation
FEASIBILITY	**Time-to-Market:** Can a working prototype or pilot be delivered quickly and seamlessly?			5%	Early wins build confidence, but speed must remain balanced against governance and control
FEASIBILITY	**Technical Readiness:** Could the system (e.g., integration, security, and controls) facilitate this use case?			10%	Existing infrastructure, integration, security, and controls can materially delay or block implementation
IMPACT	**Strategic Alignment:** Does the initiative directly support the organisation to achieve statutory objectives?			15%	Ensures A.I. initiatives support business strategy, operating priorities, and value creation
IMPACT	**Value Impact Magnitude:** How large is the expected benefit (e.g., regulatory quality, faster turnaround)?			15%	Directly measures expected benefits, including revenue, cost, productivity, risk, and customer outcomes
IMPACT	**Value Impact Reach:** How widely shared is the use case across functions and divisions?			10%	Scalable use cases can generate broader enterprise-wide benefits across functions, products, and geographies
IMPACT	**Change Readiness:** How open and willing are the broader team to adopt A.I.-assisted ways of working?			5%	Important for adoption, but better addressed through training, change management, and workflow redesign
	TOTAL WEIGHT			100%	

Source: Quinlan & Associates analysis

WORKFLOW-LEVEL CONTROL MEASURES

Financial institutions should also extend governance to the workflow level by defining runtime controls, human checkpoints, escalation triggers, resilience mechanisms, and recovery procedures that establish the operating boundaries for implementation

Control Measures
Key Elements, Illustrative Example

	Description
WORKFLOW CONTROLS	Govern how each A.I. model **operates at runtime**: • Permitted input types • Output formats and post-processing validation • Checkpoint trigger conditions (e.g., accuracy threshold breach) and action taken (e.g., human review, hold)
ESCALATION CONTROLS	Define what happens when **something goes wrong**: • Human intervention triggers and exception handling (e.g., responsible person, timeline, and authority) • Risk-based escalation paths, routing different severity levels to different reviewers
RESILIENCE CONTROLS	Ensure **contained damage and smooth recovery** when an A.I. model fails / produces harmful outputs: • Kill switches and its fallback procedures • Recovery procedures (e.g., resolution evidence, authority to reinstate, and monitoring)

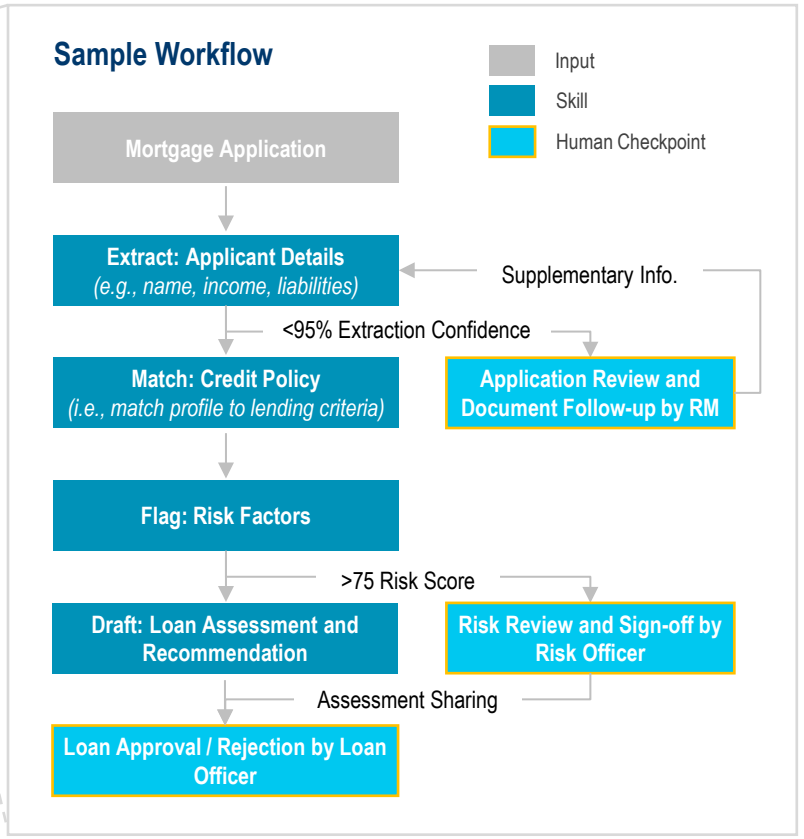

HUMAN CHECKPOINT

Defined points in the workflow where human intervention is required, whether it be for approval, escalation, or fallback

Best Practices:

- ✓ Adopt a "design for failure" approach (e.g., log for compliance, escalation triggers, etc.)
- ✓ Define clear trigger conditions for intervention (e.g., risk score, confidence threshold, etc.)
- ✓ Ensure right-sizing of human involvement

*An enforceable control mechanism with **clear accountability***

Source: Quinlan & Associates analysis

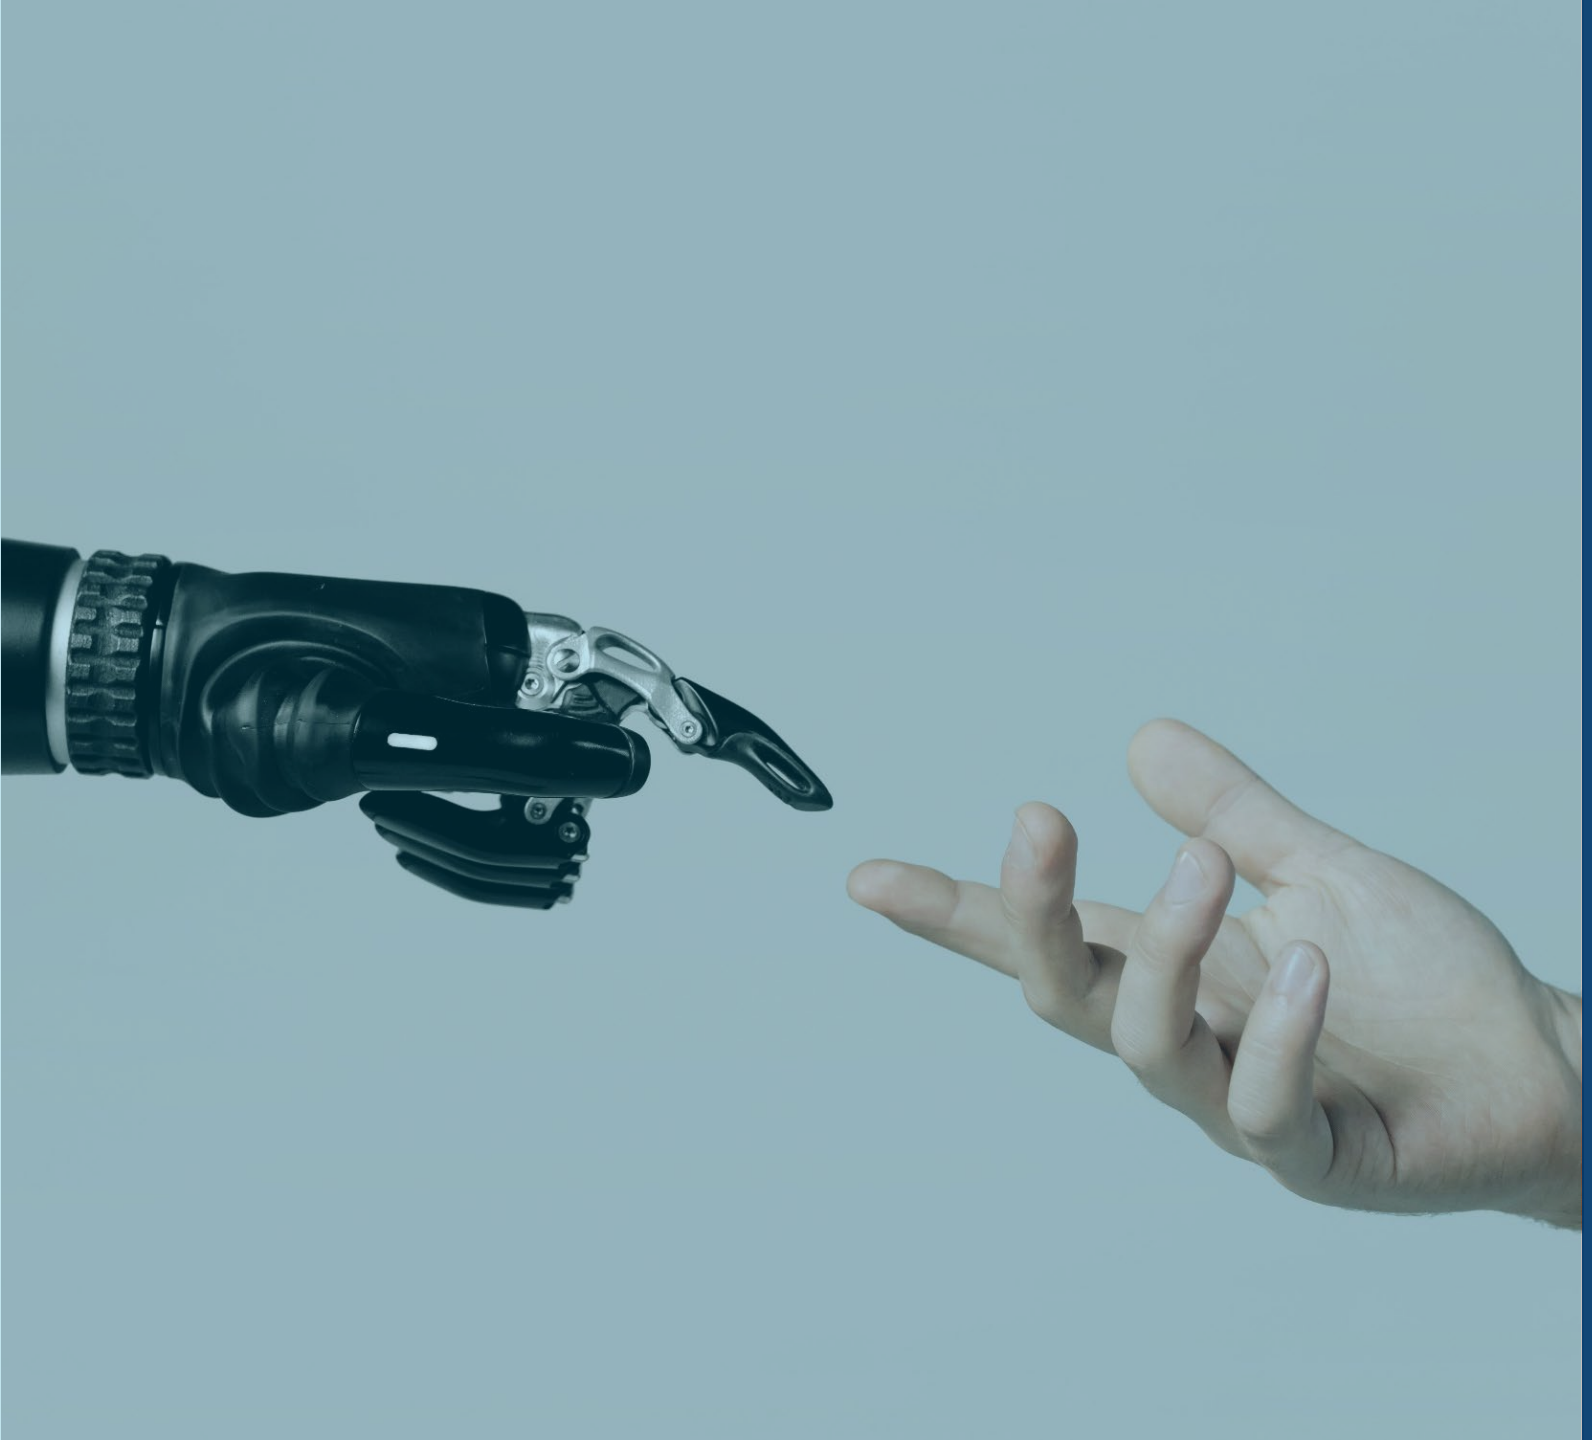

SECTION 4

ENABLE

QUINLAN &ASSOCIATES

SECTION 4-1

TECHNOLOGY

MODEL DEPLOYMENT APPROACH

Once governance establishes permissible operating boundaries, financial institutions must determine whether to use frontier, open-source, or open-weight models, or adopt a multi-model architecture supported by an LLM orchestration layer

Deployment Approach

Frontier vs. Open Source vs. Hybrid

THIRD-PARTY DEPLOYMENT

(Closed-source Frontier Model)

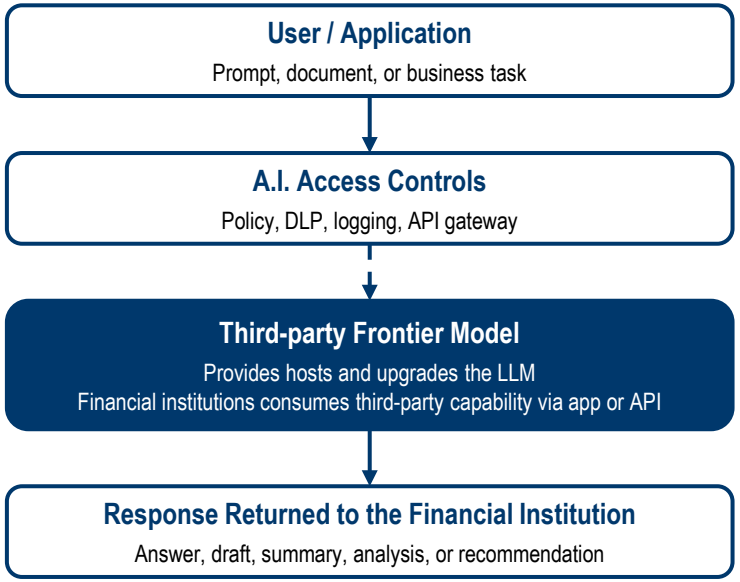

Third-party managed LLMs accessed through web, application, and/or API channels, allowing financial institutions to access leading model capabilities without directly operating the underlying model or infrastructure

SELF-HOSTED DEPLOYMENT

(Open Source / Weighted Model)

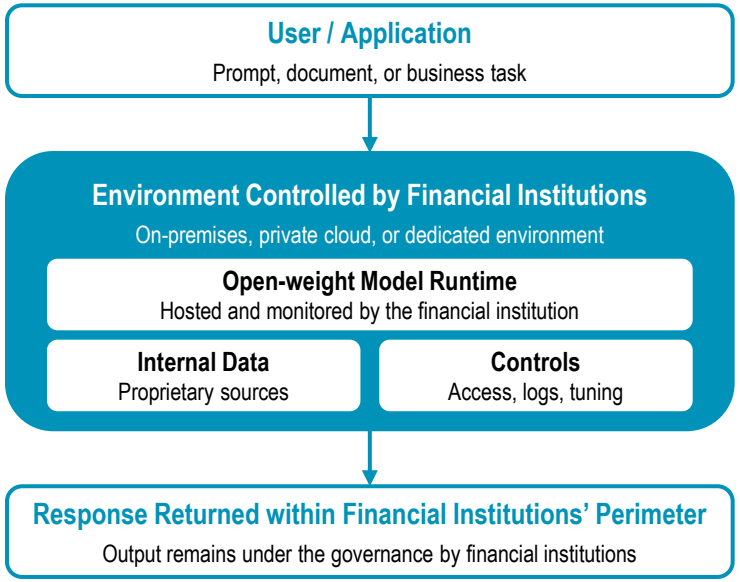

LLM is controlled by the financial institutions and deployed inside on-premises or private cloud environments, enabling tighter control over hosting, data handling, access permissions, finetuning, and infrastructure integration

HYBRID DEPLOYMENT

(Multi-model Architecture with Orchestration Layer)

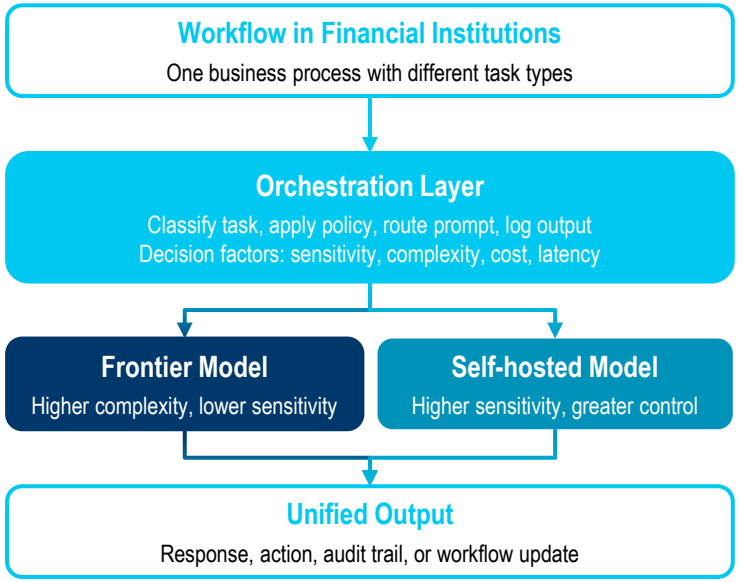

An orchestration approach that routes different tasks or workflow steps to the most suitable model environment based on sensitivity, performance, cost, governance, and operational requirements

Source: Microsoft, Google, Quinlan & Associates analysis

MODEL SELECTION

When selecting an appropriate LLM, a structured scoring methodology across key decision criteria should be employed, adjusting the weighting of each criterion to reflect the needs and sensitivity of different business units

	Model Selection Methodology Key Criteria and Assigned Weights (%)	Frontier LLM (Closed Source)	Self-hosted LLM (Open Weight/Source)	Front-office (Weight)	Back-office (Weight)
BUSINESS	**Data Privacy:** How sensitive are the data inputs and conversations processed by the LLM?	■□□□□	■■■■■	20%	5%
	Domain Finetuning: To what extent should the LLM be fine-tuned for the financial institution's business-specific requirements?	■□□□□	■■■■■	20%	5%
	Counterparty Risks: How critical is it for the financial institution to control the model lifecycle and steer clear of vendor lock-in?	■■□□□	■■■■■	10%	5%
MODEL	**Capability:** How important is long-context understanding and advanced reasoning for achieving quality outcomes?	■■■■■	■■■□□	20%	20%
	Latency: Is there a hard requirement for response time, throughput, uptime, and fallback?	■■■■■	■□□□□	10%	5%
IMPLEMENTATION	**Time-to-Market:** What is the financial institution's internal timeline for the LLM to be deployed into production environments?	■■■■■	■■□□□	5%	20%
	Total Cost of Ownership: What is the financial institution's appetite to shoulder the upfront and ongoing cost (e.g., API calls)?	■■■□□	■■■□□	5%	30%
	Operational Readiness: Does the financial institution have internal capabilities (e.g., infrastructure and talents, etc.) to support usage?	■■■■■	■■□□□	5%	10%
SUITABLE MODEL				**Self-hosted LLM**	**Frontier LLM**

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

FOUNDATIONS OF A.I. ENGINEERING

The success of A.I. deployment depends not only on the model itself, but also on the underlying model-level harness engineering, which must be robust to ensure model capability is reliably translated into production-grade outcomes

A.I. Solution Pillars
Key Components

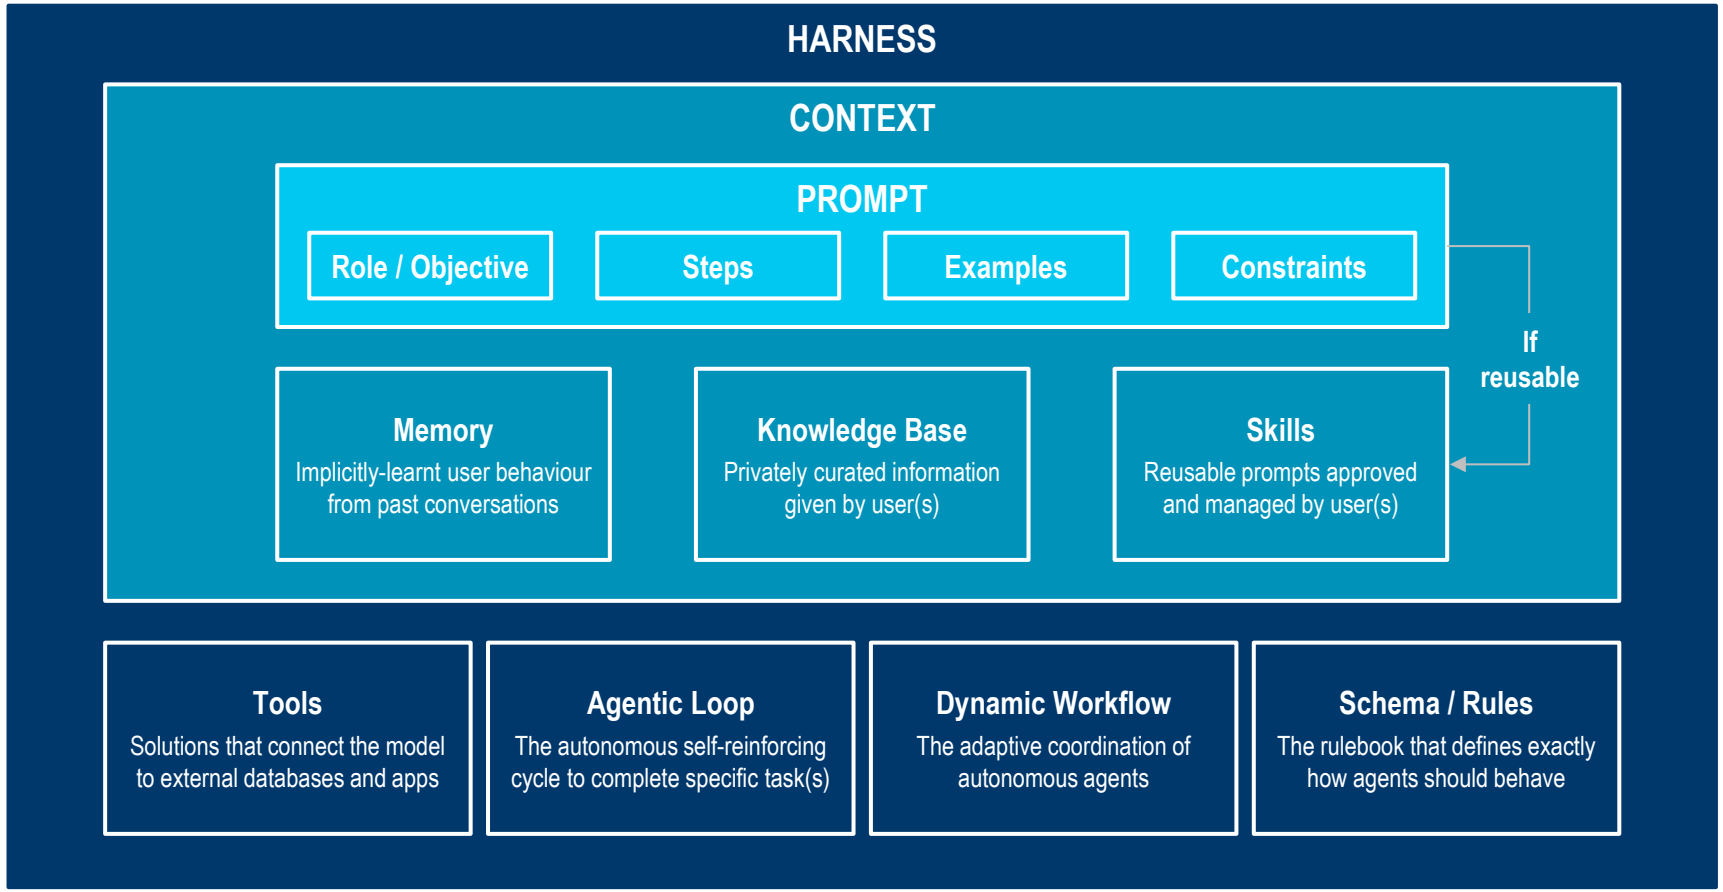

PROMPT ENGINEERING (2022)

The natural language instruction or question to trigger a specific response

→ *What does the A.I. model need to do?*

CONTEXT ENGINEERING (2024)

The surrounding context and data to ground the A.I.'s response in relevant facts

→ *What does the model need to know?*

HARNESS ENGINEERING (2026)

The infrastructure around a model to automate, control, and validate its end-to-end execution

→ *How should A.I. perform the task?*

Source: AWS, Microsoft, Google, IBM, Quinlan & Associates analysis

QUINLAN &ASSOCIATES

A.I. ENGINEERING

Across the different generations of LLM-based solutions being explored, multiple A.I. engineering components must be appropriately configured to ensure that the underlying technology performs reliably within defined governance boundaries

A.I. Solutions Evolution

Applicability of Prompt, Context, and Harness Engineering

✓ Fully Applicable - Applicable to Some Extent × Not Applicable

FIRST GENERATION → SECOND GENERATION → THIRD GENERATION →

	CHATBOT	RETRIEVAL ASSISTANT	COPILOT	A.I. AGENT	AGENTIC A.I. SYSTEM
	Conversational A.I. interface where the user prompts the model to generate a specific output, such as answering a question, drafting content, and summarising information	A.I. assistant connected to external knowledge sources and/or internal databases, which retrieves and references relevant data before generating a response	A.I. assistant embedded within end users' workflows and applications, supporting them with context-specific tasks within their real working environment	A goal-based software that can interpret an objective, plan steps, call tools, monitor progress, and execute multi-step tasks within end-user-approved parameters	System of one or more A.I. agents that pursues defined goals with limited human supervision, often across multiple tools, workflows, data sources, and approval layers
PROMPT					
Role / Objective	✓	✓	✓	✓	✓
Steps	✓	✓	✓	✓	✓
Examples	✓	✓	✓	✓	✓
Constraints	✓	✓	✓	✓	✓
CONTEXT					
Memory	✓	✓	✓	✓	✓
Knowledge Base	×	✓	✓	✓	✓
Skills	×	×	×	✓	✓
HARNESS					
Tools	-	-	✓	✓	✓
Agentic Loop	×	×	×	✓	✓
Schema	×	×	×	✓	✓
Orchestration	×	×	×	×	✓

Low — LEVEL OF AUTONOMY — High

Source: AWS, Microsoft, Google, IBM, Quinlan & Associates analysis

QUINLAN &ASSOCIATES

PROMPT ENGINEERING

To promote consistent and reliable outputs across the organisation, financial institutions should define standardised prompt templates that set out objectives, examples, and step-by-step guidance for employees using GenA.I.-powered applications

Key Components

Well-configured Prompt

SET OBJECTIVES

Designate the model's role and give an overview of the task, as well as relevant contexts to ensure the model captures task-specific nuances

INSTRUCT STEPS

Detail each step involved, including any constraints that the model is prohibited from doing while carrying out the task

INCORPORATE EXAMPLES

Give specific examples of expected outputs help the model learn about the task and the complexity, tone, and concision of expected outputs

USE DELIMITERS (e.g., <>, ```, “”, etc.)

Specify where inputs are inserted in the prompt to help model identify user inputs (e.g., company profiles), especially when they are long

STRUCTURE OUTPUTS

Specify the response with a fixed syntax to parse the output into be more readily available formats to be fed into the following processes

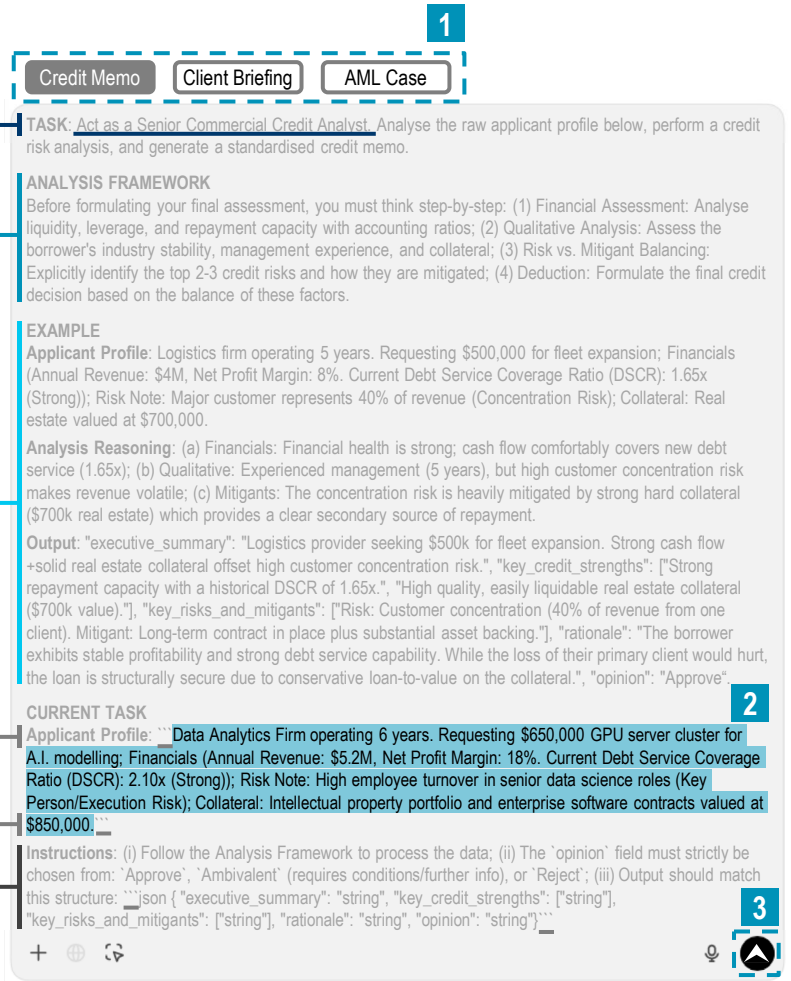

Employee Workflow

1 Select Task

Choose the task for which predefined templates have been created with all key elements to be loaded in the chatbox

2 Input Information

Provide case-specific information for the model to carry out the specific task, such as client profile and financial position

3 Generate Output

Request the model to generate outputs following the format and steps as instructed in the template

Source: CoreWeave, Comet, Quinlan & Associates analysis

QUINLAN &ASSOCIATES

SKILLS (1/2) – OVERVIEW

Skills are reusable and scalable procedural prompts that can be used by humans on demand or automatically invoked by the relevant A.I. agent based on specific configuration

Lifecycle of Skill

Origination, Storage, and Usage

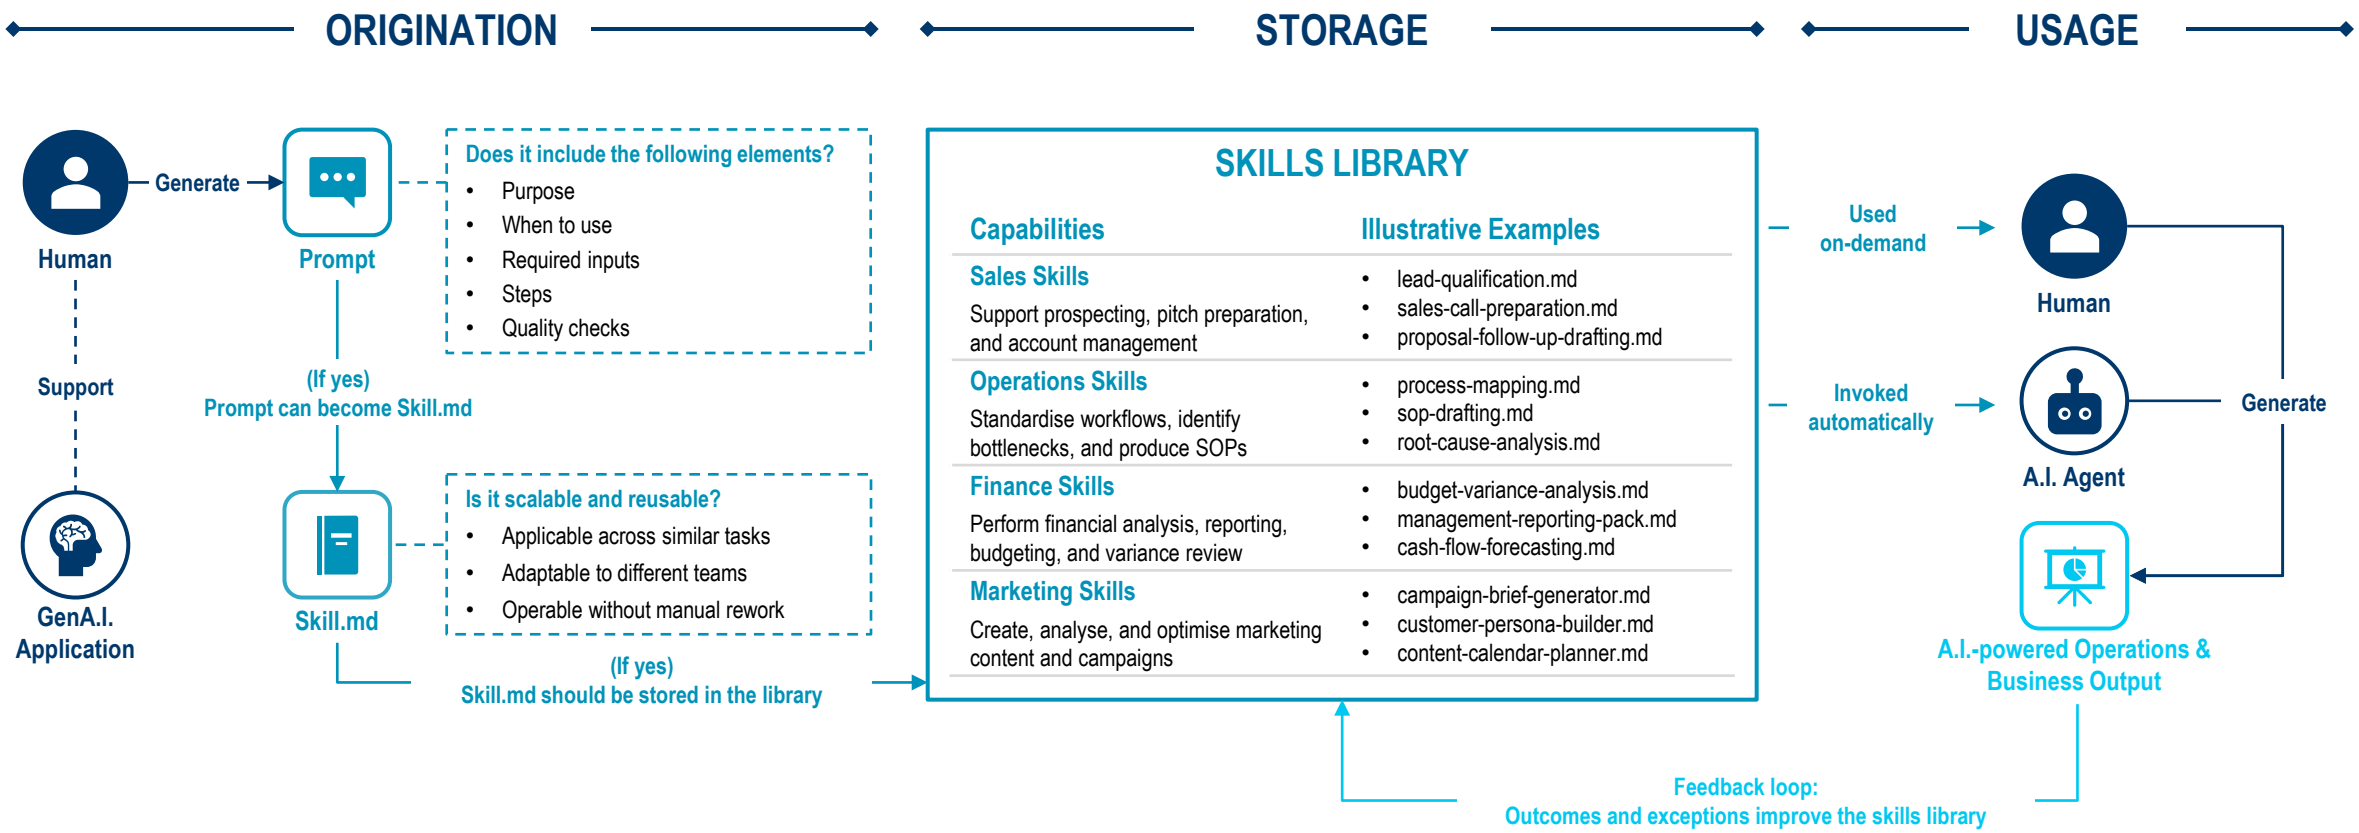

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

SKILLS (2/2) – COST & PERFORMANCE IMPACT

On average across the model lineup, skills increase per-trial token consumption, while reducing per-trial cost in practice as a greater share of tokens shifts towards cheaper input-side overhead and improvements in first-pass accuracy with fewer retries

Cost & Performance Improvement

Skills vs. Without Skills, Performance-focused vs. Efficiency-focused Model

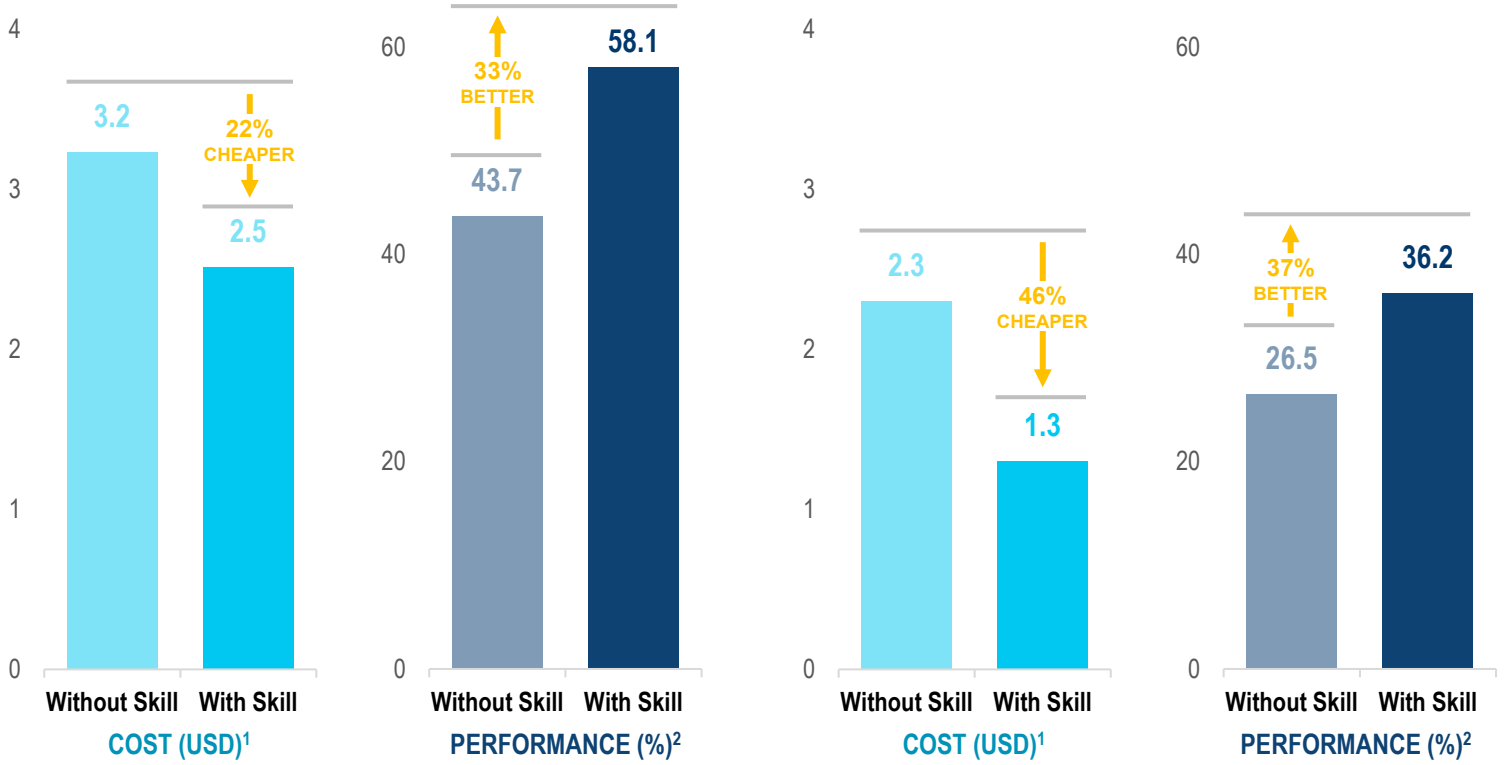

PERFORMANCE-FOCUSED MODELS
(GPT-5.5, Opus 4.8, and Gemini 3.1 Pro)

EFFICIENCY-FOCUSED MODELS
(GPT-5.4 mini, Sonnet 4.6, and Gemini 3.1 Flash Lite)

[1]Price per 1 million tokens; [2]Task-macro success rate by agent configuration

Source: SkillsBench, Quinlan & Associates analysis

Key Reasons

Token Economics

OBJECTIVE FUNCTION

The curated skill provided the model with the code template or algebraic structure across steps / calls that it can then attempt to optimise

→ *Focused intelligence = less waste / attempt*

STRUCTURED PROCEDURE

Skills break a complex task into structured steps, increasing tokens from context-reading and acting but fewer total trials are need to achieve outcome

→ *Higher success rate with fewer trial counts*

REASONING / OUTPUT CAPABILITY

While skills introduce additional low-cost input tokens (e.g., workflows, templates), it reduces expensive output, retries, and reasoning loops

→ *Less costly reasoning & correction cycles*

The skills context allows models to reach a solution with significantly less explanatory work

QUINLAN &ASSOCIATES

KNOWLEDGE BASE AND MEMORY

The right context engine should be driven by the intended use case, with RAG excelling at retrieving and grounding responses in the document repositories, while LLM Wiki can embed institutional tone and know-how, making it a powerful source of differentiation

Context Framework

RAG vs. LLM Wiki

User | Knowledge Base / Source / A.I. Model | Output | Manual Flow | Automated Flow | ✓ High | - Moderate | × Low

Retrieval-Augmented Generation ("RAG")

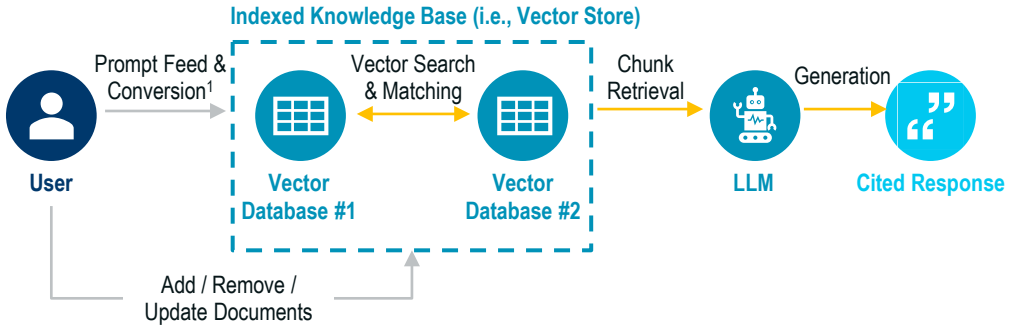

Large Language Model ("LLM") Wiki

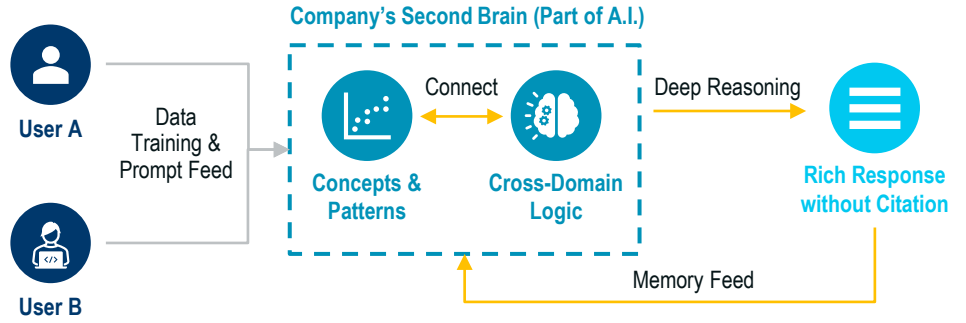

Ease of Setup and Maintenance	×	Self-managed deployments require chunking logic, an embedding model, a vector database, and ongoing retrieval evaluation and monitoring	✓	No complicated back-end infrastructure is required and users can easily drop files into context and edit them directly
Knowledge Depth & Reasoning Quality	-	Reasoning quality matches the underlying model, but synthesis is bounded by retrieval: the model reasons only over the top-ranked chunks returned per query	✓	The model sees the entire corpus at once with no chunking losses, enabling cross-document synthesis and nuanced reasoning
Knowledge Freshness & Update Latency	-	Source documents can be updated and re-indexed incrementally, with pipelines routinely automated on an event-driven or scheduled basis	-	Updates take effect instantly once the markdown file is edited, but past interactions (if not well governed) can silently propagate user errors
Protection from Hallucination Risk	✓	Answers are grounded in retrieved passages and can carry citations back to source, making claims auditable	-	In-context content provides a degree of grounding, but the model can confidently generate results that falls outside the wiki

[1]Conversion to numerical embeddings in order to run similarity search, also known as a vector

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

AGENTIC LOOP

When the components of harness engineering work together, they enable the agentic loop: the core mechanism by which Agentic A.I. iteratively executes, evaluates, and refines its own outputs autonomously to achieve defined objectives

Agentic Loop

Diagram

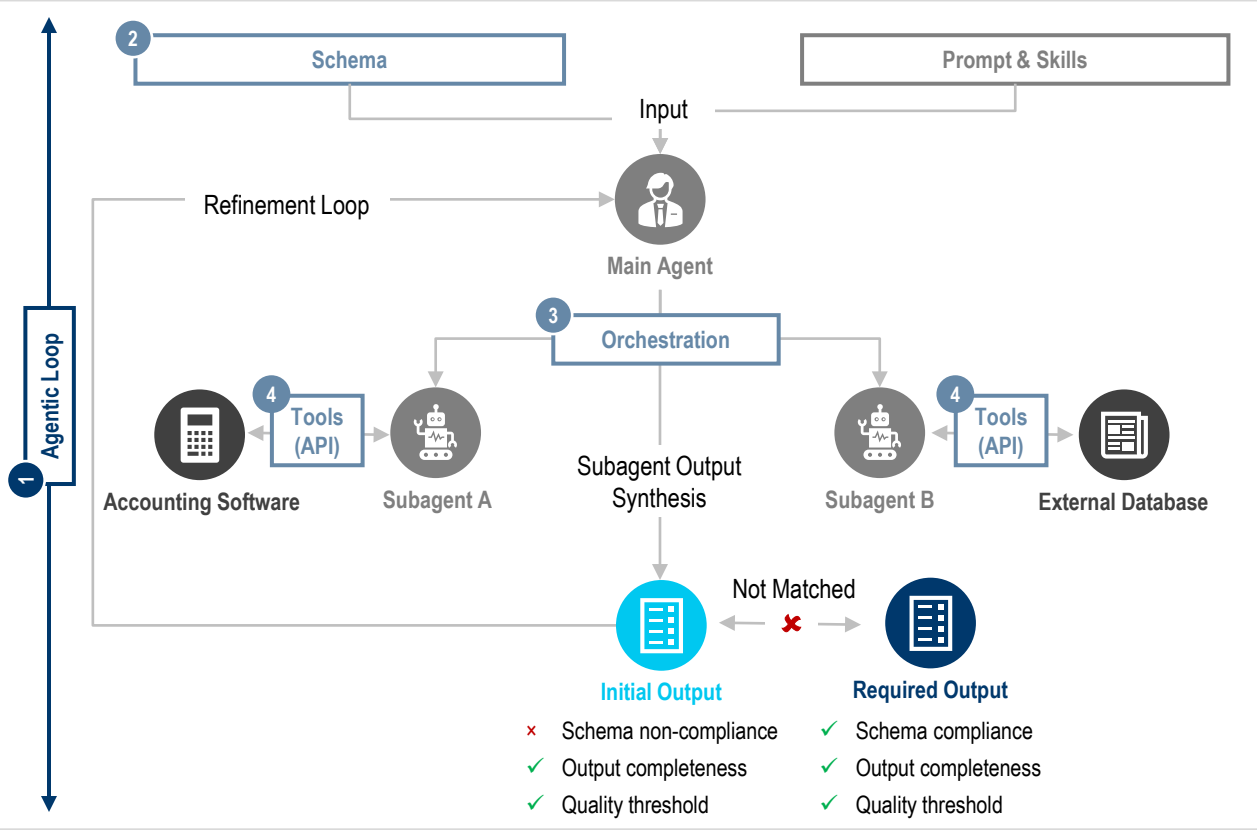

1

Agentic Loop

The autonomous cycle that the model executes to complete a specific task (including observing and matching output with the required conditions to instigate a refinement loop, if needed), all without human intervention

Powered by

2

Schema

The rulebook that defines the plan of attack and target required output that guides how agents should behave

3

Orchestration

The coordination of autonomous agents that reason and act with specific control logic and workflow sequencing

4

Tools

Solutions (e.g., API and MCP) that connect the model to external databases and applications

Source: Quinlan & Associates analysis

CONSEQUENCES OF NOT GETTING IT RIGHT

Without effective engineering, A.I. solutions can erode the value of institutional investment, resulting in economic waste, unreliable outputs, and weakened system integrity, resulting in heightened operational, legal, and reputational risks

Key Consequences
Use of A.I. in Other Industries

1 ECONOMIC WASTE

Runaway Cost from Context Rot
Without disciplined skill and workflow design, A.I. systems reprocess bloated context at every stage and consume far more tokens

Lost Staff Time from Manual Rework
Vague skill definitions and undocumented failure modes renders outputs unusable, forcing employees to re-run the same task or redo it manually

2 UNRELIABLE OUTPUTS

Hallucination and Factual Errors
Without skills in cross-verifying facts, users may accept A.I.-generated fabrications, leading to compliance, legal, or technical issues

Poor Judgement and Reasoning
Outputs may fail to reflect sound judgment, domain expertise, or organisational intent due to flawed decision logic / workflow

3 COMPROMISED INTEGRITY

Data Breach
Users frequently leak proprietary information into the A.I. open context if restrictions and data boundaries are not clearly defined in the prompt

External Manipulation
Without strict isolation layers, an A.I. can process malicious instructions (prompt injections) hidden within external sites, compromising systems

J.P.Morgan

Erosion of returns from ungoverned A.I. Usage

J.P. Morgan's Chief Data and Analytics Officer noted that some employees are incurring A.I. token expenses exceeding their salaries, raising skepticism about the true efficacy and return on A.I.

Klarna

Poor judgement in over-automated A.I. support

Klarna's A.I. assistant applied the same automation logic across all query types and hence could not distinguish routine requests from high-stakes cases such as fraud or disputed charges

Community Bank

Sensitive customer data exposed via A.I. tool

A US bank employee uploaded customer data, including Social Security numbers, into an unsanctioned A.I. application, triggering severe regulatory and operational consequences

Source: Forbes, Fortune, Quinlan & Associates analysis

SECTION 4-2

PEOPLE & PROCESSES

TEAM SETUP

Every role is becoming more managerial in nature, as team members direct and review A.I. outputs, managers orchestrate mixed human and A.I. teams, and leaders steer institutions through a transformation they may not yet fully understand

Team Structure

Traditional Team vs. A.I.-enabled Team

TRADITIONAL TEAM SETUP

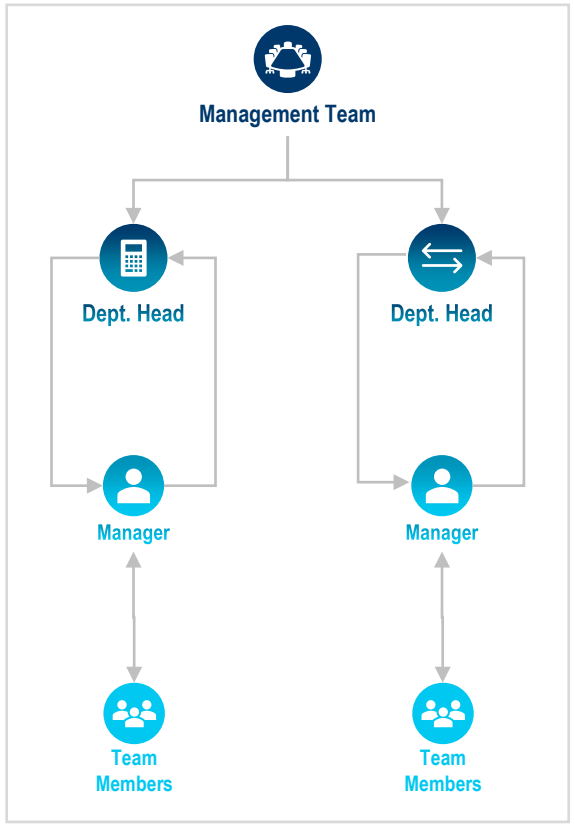

A.I.-ENABLED TEAM SETUP

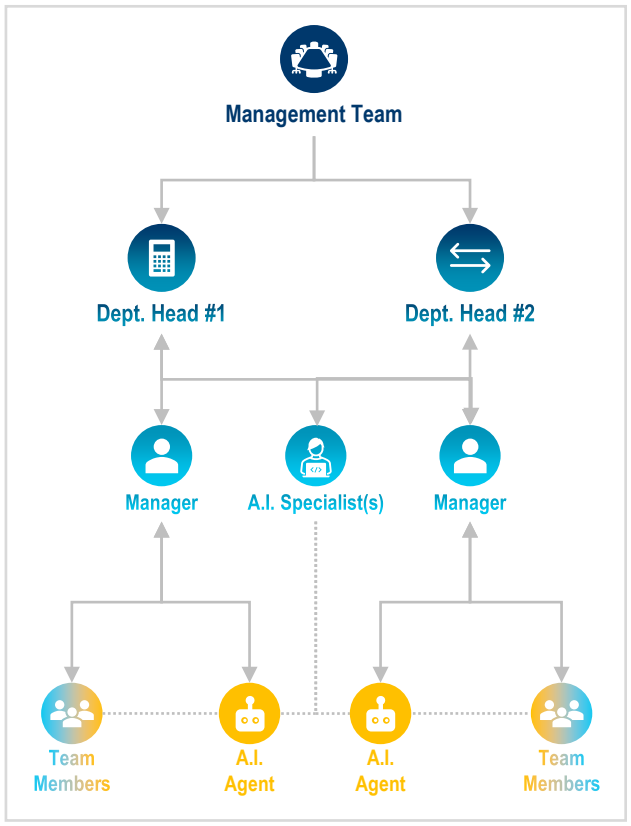

Senior Leaders · Middle Managers · Frontline Staff · Artificial Intelligence

SENIOR LEADERS

TRUST STEWARD

Senior leadership sets direction and relies on technical teams to recommend and manage technology decisions, given their domain expertise

DIAGNOSTICIAN

Senior leadership is accountable for how A.I. is used across the organisation, the tone on responsible A.I. use, and leading teams who are often more A.I.-capable

MIDDLE MANAGERS

COORDINATOR

Middle management allocates work to team members based on capacity and skills, reviews outputs, and manages performance

HUMAN / A.I. ORCHESTRATOR

Workflow design is not a core mgmt. skill, and middle management owns morale, ensuring team members feel secure, valued, and clear on changes in roles

FRONTLINE STAFF

EXECUTION SPECIALIST

Team members act as skilled contributors, producing output directly and reporting to the team heads, with their performance measured

QUALITY ASSURER

Team members decide what to prompt, review A.I. output critically, and make adaptations, along with adding a judgment layer that may be missing in A.I.

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

NEW ROLES IN THE A.I. AGE

Senior leaders set strategic direction and investment boundaries, middle managers translate those priorities into workflow-level redesign, and frontline staff determine what works in day-to-day execution

Updated Roles
Key Questions and Purpose

UPDATED ROLES IN THE ORGANISATION	KEY QUESTION TO ADDRESS	PURPOSE
SENIOR LEADERS *[DIAGNOSTICIAN]* Define the strategic objective, allocate resources, and set risk boundaries, but should not be expected to diagnose every workflow-level inefficiency	• What business outcome are we trying to achieve?	✓ Ensures A.I. starts with a business objective, not a tool
	• What is our level of commitment and preparedness?	✓ Clarifies budget, sponsorship, and risk appetite
	• What does success look like?	✓ Defines measurable KPIs (e.g., revenue uplift, cost reduction, better client experience)
	• What should not be compromised?	✓ Sets non-negotiables (e.g., client trust, compliance, confidentiality, quality control, etc.)
MIDDLE MANAGERS *[ORCHESTRATOR]* Translate senior-level priorities into workflow redesign by identifying where processes break down, where A.I. should be embedded, and where human control must remain	• Which end-to-end workflow is underperforming?	✓ Identifies friction points (e.g., approval delays) of the specific workflow to redesign
	• Which steps should be automated, augmented, or left human-owned?	✓ Prevents blunt automation of tasks that require judgement
	• Where should human review and approval sit?	✓ Builds control into the redesigned workflow
	• Where do exceptions, workarounds, and edge cases occur?	✓ Ensures the redesigned workflow handles real-world complexity
FRONTLINE STAFF *[QUALITY ASSURER]* Provide the operational reality by showing where work actually breaks, where A.I. helps or fails, and what must be refined before the workflow can scale	• What actually happens day to day?	✓ Captures operational reality rather than assumed process design
	• Which steps are most repetitive, manual, or frustrating?	✓ Reveals high-potential use cases for A.I. deployment
	• Where is A.I. output useful, unreliable, or unusable?	✓ Calibrates the quality and practicality of the A.I. deployment
	• What needs to be fine-tuned before the workflow can scale?	✓ Creates the feedback loop for adoption and continuous improvement

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

KEY CONSIDERATIONS FOR SENIORS LEADERS

As A.I. blurs the boundaries between strategy, operations, and technology, senior leaders must understand the underlying shifts, build first-hand capabilities, and actively lead the resulting organisational redesign

Key Considerations

Cheat sheet

	THE SHIFT: WHAT IS AHEAD?	**THE IMPACT:** WHY DOES THIS MATTER TO YOU?	**YOUR MOVE:** WHAT SHOULD YOU DO NEXT?
STEP 1 UNDERSTAND	• **Answering → Acting:** A.I. is no longer merely answering your questions, it acts on your behalf • **Adoption → Scale:** Everyone now uses A.I., but only a few convert usage into value through workflow redesign • **Specify & commission → Sketch & prototype:** Business can directly pilot test its own solutions without waiting for IT team	• **The emergence of new risk category:** A system that acts can now do the wrong thing, not just say it • **The race has started:** Having A.I. is no longer a differentiator; advantage now sits in addressing the adoption-to-value gap • **The new bottleneck:** From IT capacity to your own ability to articulate, test, and iterate ideas	• **Supervise by design:** Treat A.I. agents as digital workers: for each one, define who oversees them and how they are audited • **Fewer deeper bets:** Place capital on deployments that redesign whole workflows with clear business case • **Rewrite the IT contract:** Business builds and iterates; technology certifies, validates, and hosts
STEP 2 SPONSOR	• **Arm's-length sponsorship → First-hand experience:** You can no longer sponsor A.I. credibly from a distance • **Endorsing → Judging:** The scarce executive skill becomes judging the output quality for verification and calibration • **Speed → Fit:** Value no longer comes from adopting fast, but from systematically screening where A.I. fits your business	• **The credibility gap:** Sponsors who have never built anything misjudge the potential, the risk, and the speed of deployment • **Error is a design parameter:** Inaccuracy is a commonly reported A.I. output failure, presenting real business risks • **Focus on high-value work:** Judgement-heavy synthesis pay best; deterministic high-volume work often does not	• **Build one thing:** Create one working prototype from your own work in a sandbox before asking others to do the same • **Verify, then trust:** Adopt structured verification habits: demand sources and pressure test cases • **Rank your P&L:** Screen your own function and produce a ranked use-case shortlist before committing any budget
STEP 3 LEAD	• **Forecasting → Option management:** Given the pace of A.I. development, organisational strategy remain open questions • **Sequential approvals→ Risk-tiered gating:** Low-risk gates are automated, human review is reserved for high risk • **Buying software → Buying capability:** You are no longer procuring tools, you are buying capability that gets outdated fast	• **Single scenario does not work:** Betting on a single rigid strategy and architecture will lead to irreversible consequences • **Traditional controls measures are outpaced:** Your control framework was built for human-paced change, not for A.I. • **"Lock-in" is the new risk:** Long commitments are dangerous while capability and pricing shift this fast	• **Stay reversible:** Hold options: configurable human gates, reskilling over restructuring, audit-by-design from day one • **Redesign, no bypass:** Classify use cases by risk, automate the low-risk, and concentrate human review on the high-risk • **Move away from traditional cycle:** Transition from annual SaaS subscription to tailored solution leveraging LLM

Source: Quinlan & Associates analysis

TRADITIONAL APPROACHES NO LONGER WORK

A.I. transformation cannot rely on a top-down model in which management selects software and pushes it across the organisation, as meaningful value is identified closer to the frontline, where workflows, pain points, and domain knowledge are most visible

Traditional vs Modern Approach

Why it fails and works

TRADITIONAL TOP-DOWN APPROACH

A predefined software solution is selected first, then the organisation attempts to standardise work around the tool

SENIOR LEADERS

- Approve a technology solution based on strategic ambition, vendor claims, and pressure to modernise

MIDDLE MANAGERS

- Convert the selected system into an implementation plan, often with limited ability to reshape the underlying workflow

FRONTLINE STAFF

- Are expected to adopt a fixed workflow, even when the tool does not reflect exceptions, judgement, or informal workarounds

✘ WHY THIS FAILS

Technology defines the redesign boundary, causing the organisation to optimise around system configuration rather than workflow requirements

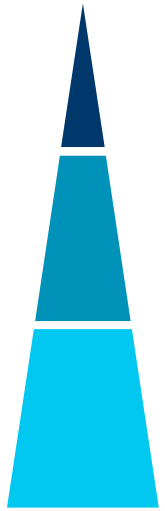

MODERN BOTTOM-UP APPROACH

A.I. act as flexible workflow discovery and augmentation tools, helping teams expose, test, and codify how work should change

FRONTLINE STAFF

- Use A.I. to surface recurring pain points, draft outputs, test prompts, and reveal where work is repetitive or judgement-heavy

MIDDLE MANAGERS

- Translate frontline insight into A.I.-embedded workflow designs, control points, data requirements, and scalable use cases

SENIOR LEADERS

- Prioritise, fund, govern, and scale the use cases that are strategically relevant, feasible, and measurable

✓ WHY IT WORKS

The workflow defines the technology need, allowing A.I. to be embedded where it improves execution, decision quality, and scalability

Source: Quinlan & Associates analysis

IMPORTANCE OF DOMAIN EXPERTISE FROM FRONTLINE STAFF

As A.I. applications progress from simple workflow automation to more advanced A.I.-enabled redesign, transformation quality depends on how effectively human expertise is codified into context, rules, controls, and validation logic

Complementary Relationship

Human Domain Expertise & A.I. Execution Capabilities

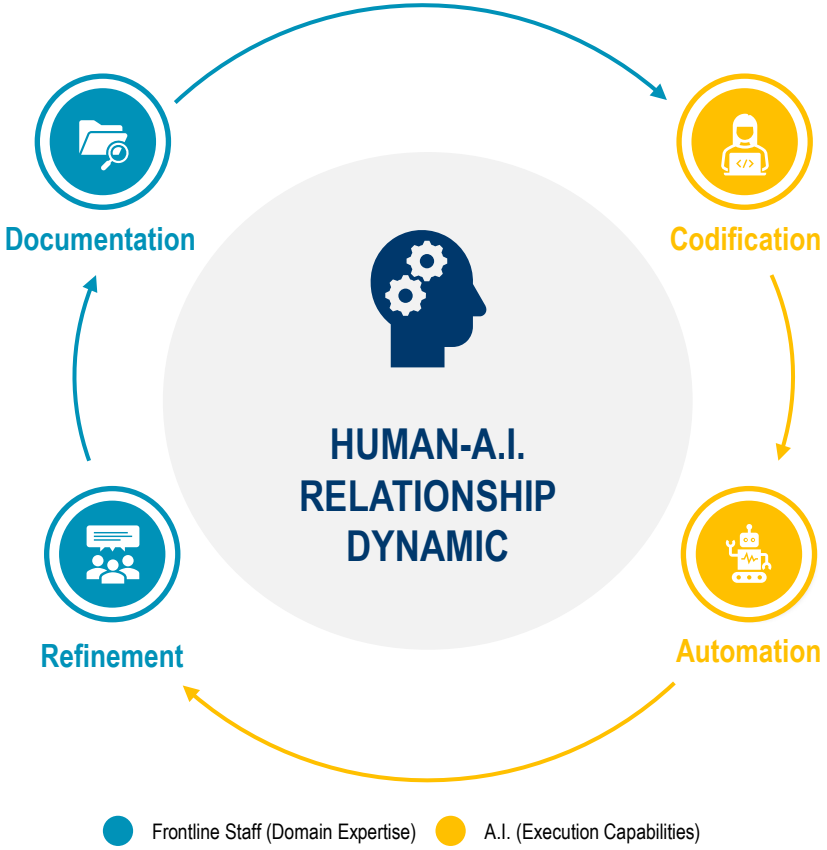

Types of Activities	Description	Examples	Goals
1. Documentation (By frontline staff)	Human domain expertise is captured from documents, procedures, examples, business rules, and operational experience	• Business intent • Real-life workflow and example • Judgement criteria	**Raw expertise base that explains what matters and why**
2. Codification (By A.I.)	A.I. helps structure, clarify, test, and convert domain expertise into repeatable workflow logic	• Codified SOP • Dynamic workflow development • Decision rules & control thresholds	**Codified logic that can be consistently applied and governed**
3. Automation (By A.I.)	A.I. applies the codified expertise to automate workflow steps within defined permissions and boundaries.	• Retrieval • Drafting • Analysis and routing	**Faster and more consistent execution of repeatable manual work**
4. Refinement (By frontline staff)	Human experts review outputs, correct errors, handle exceptions, and refine the workflow based on real-world performance	• Output review • Exception handling • Rule refinement	**Improved expertise base and stronger next-cycle automation**

Source: Quinlan & Associates analysis

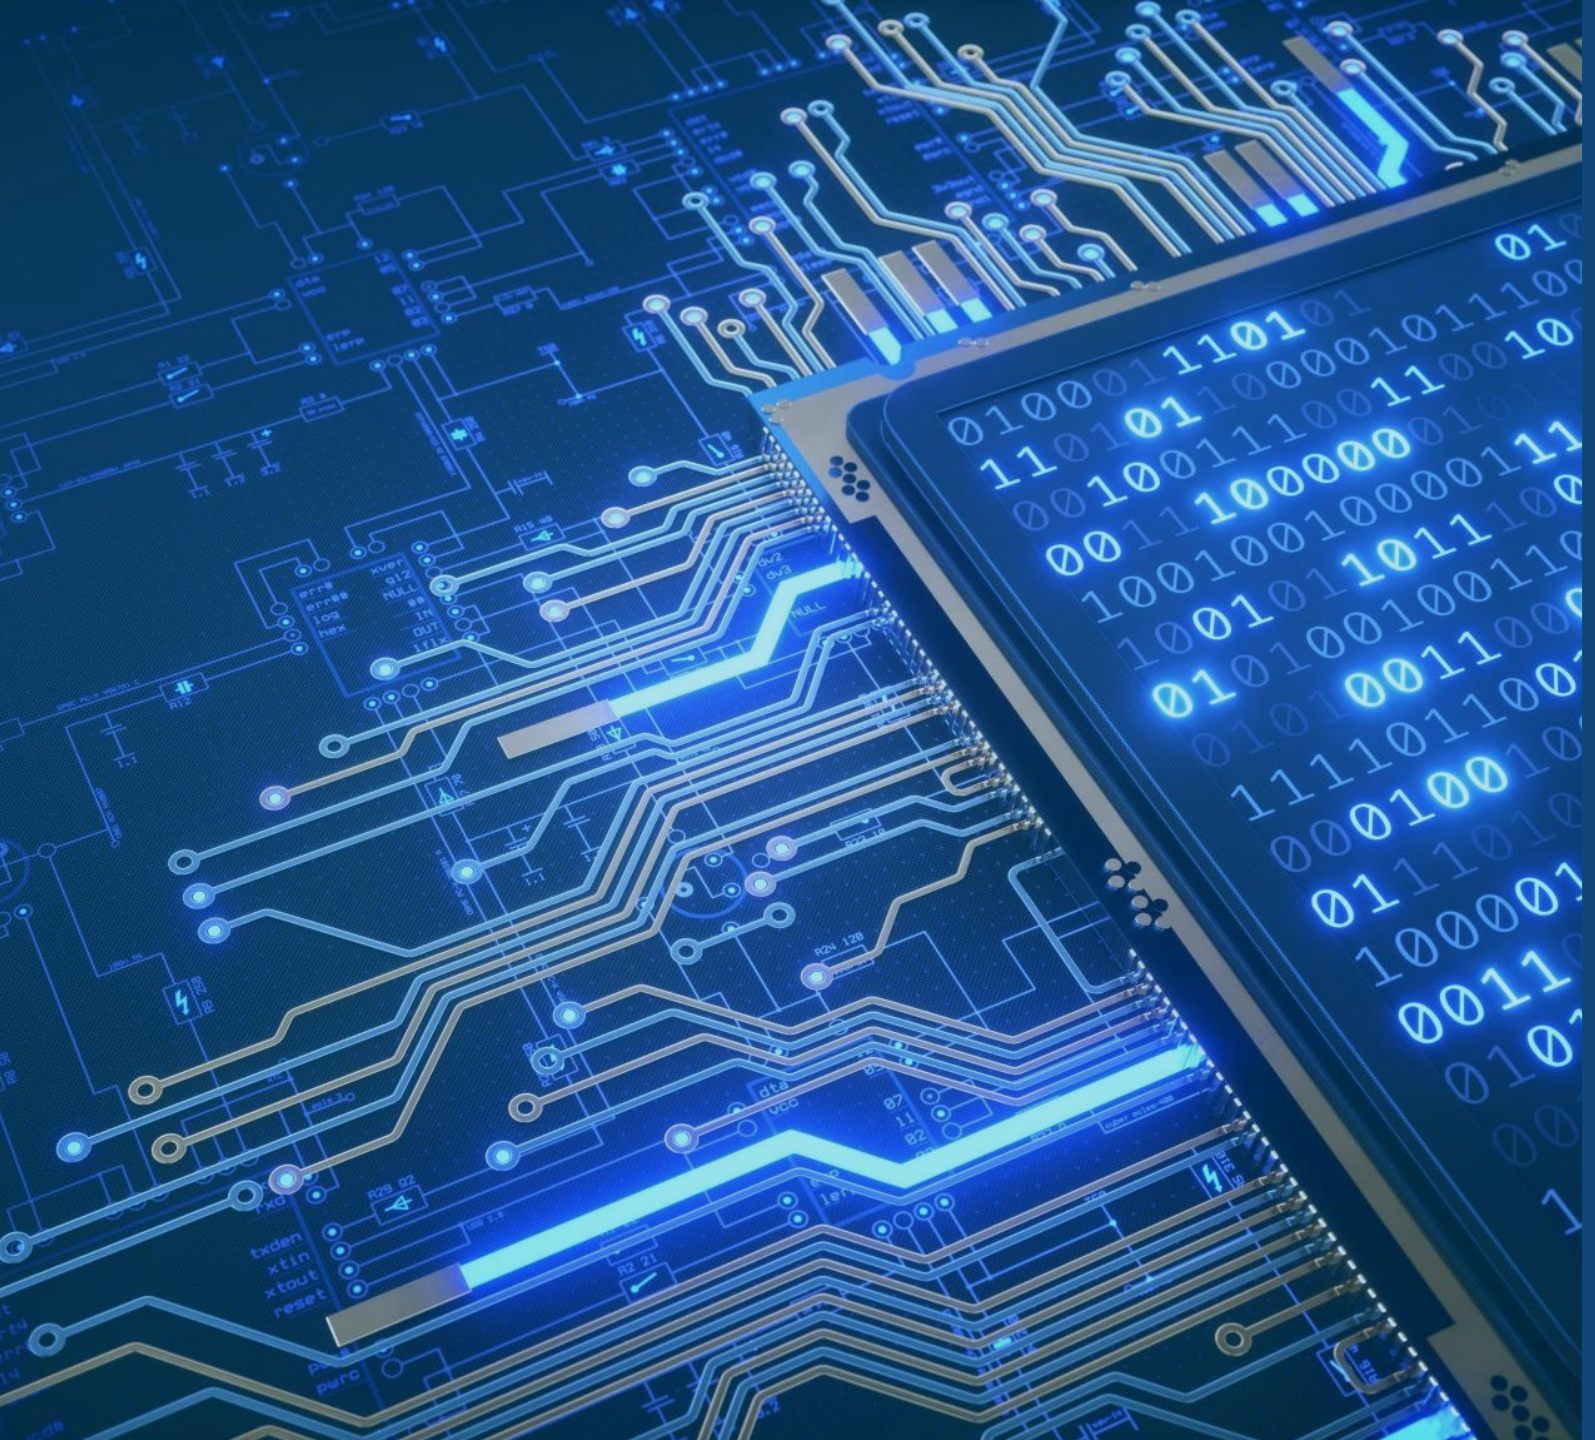

SECTION 5

HOW WE CAN HELP

QUINLAN
&ASSOCIATES

HOW WE CAN HELP

Before pursuing further A.I. initiatives, particularly agentic A.I., financial institutions should conduct a structured readiness diagnostic to identify gaps, prioritise interventions, and establish a 12–24-month roadmap for scalable adoption

How We Can Help
End-to-End Support

UNDERSTAND STATE → **BUILD FOUNDATION** → **ACTIVATE MODEL**

Best Practice Identification	Internal Diagnostic Assessment	Strategy Development / Enhancement	Target Operating Model ("TOM") Design	Implementation Roadmap Creation	Execution Support
What can we learn from others?	***Where are the critical gaps?***	***What should we prioritise?***	***What needs to evolve?***	***How do we get there?***	***How do we make it happen?***
*Evaluate **leading industry practices** across the financial services sector while referencing key **regulatory standards and emerging supervisory expectations** for A.I.*	*Assess the state of A.I. maturity across strategic and operational dimensions within the institution to **identify gaps / whitespaces** and **inform initial recommendations***	*Define / refine the A.I. strategy, and prioritise **A.I. use cases** based on key criteria, and define the approach to develop those use cases via model selection*	*Design a **future-ready operating model** that enables scalable and safe A.I. adoption and prepare any **supporting documentation**, based on agreed recommendations*	*Develop an **implementation roadmap** based on TOM & prioritised use cases, sequenced for controlled delivery, capability readiness, and managed dependencies*	*Provide **hands-on execution support** to translate the roadmap into tangible outcomes, ensuring initiatives are delivered on time, at quality, and achieve the intended impact*
• Peer mapping • Competitive benchmarking (e.g., A.I. use cases, governance structures, operating model, etc.) • Regulatory review across key operating jurisdictions • Trend analysis and case study development	• Use case mapping • Data, model, and tooling review • Governance and workflow assessment • Gap, duplication, and bottleneck identification and analysis • Severity assessment and prioritisation	• Business case development • Use case prioritisation (e.g., by value, feasibility, and risk) • Deployment approach (e.g., frontier, open-source, or hybrid) • Service model selection (e.g., build, buy, or partner) • (If partner) Potential partners evaluation and shortlisting	• Governance framework design (e.g., data governance, decision-making, coordination, etc.) • Improved operational workflow / rulebook documentation (e.g., role-based prompt templates, reusable A.I. skill library, human checkpoints, etc.) • Technology design	• Phased transformation roadmap • Use case rollout roadmap • Milestone and KPI definition • Stakeholder mobilisation plan (including resources, ownership, and internal and external communications)	• Leadership development • Employee training programme facilitation • PMO setup and execution tracking • PoC / prototype development • MVP build support • Model / A.I. solution testing support

*Collaborate with **different functions / divisions** and facilitate **senior stakeholder alignment via workshops** to validate findings, surface priorities, and agree on recommendations to drive A.I. to scale across the organisation*

Source: Quinlan & Associates analysis

QUINLAN &ASSOCIATES

STRATEGY WITH A DIFFERENCE

Website www.quinlanandassociates.com

Email enquiries@quinlanandassociates.com

Telephone (+852) 2618 5000

Address Level 20, One International Finance Centre
1 Harbour View Street, Central
Hong Kong

CONTACT US

Benjamin Quinlan
CEO & Managing Partner
T: +852 2618 5000
E: bquinlan@quinlanandassociates.com

Justin Chung
Associate Partner
T: +852 2618 5000
E: jchung@quinlanandassociates.com

Jeanny Ang
Engagement Manager
T: +852 2618 5000
E: jang@quinlanandassociates.com

Moanna Tang
Senior Consultant
T: +852 2618 5000
E: mtang@quinlanandassociates.com