

Deloitte.



AMBIENT AGENTS IN FINANCIAL SERVICES
DELOITTE'S PERSPECTIVE ON THE
FUTURE, POWERED BY GOOGLE CLOUD

INTRODUCTION

WHAT ARE AMBIENT AGENTS?

Ambient agents are a new breed of intelligent digital assistants designed to operate within the fabric of digital environments. Unlike traditional bots that require explicit user commands or interactions, ambient agents are

“invisible helpers” - they observe, learn, and act in the background, often anticipating needs before users are even aware of them.

Imagine a scenario in a large financial institution where thousands of invoices arrive daily via email. Traditionally, employees would manually review these emails, extract invoice data, categorize payments, initiate approval workflows, and update accounting systems. With ambient agents, this process is transformed, routine tasks are taken off employees' plates as a team of agents continuously monitors incoming emails, automatically extracts relevant data from invoices, categorizes payments based on predefined rules and historical patterns, and initiates the appropriate approval

workflows. The user is only prompted when their input is truly necessary, such as for final approval or exception handling.

But ambient agents aren't just limited to back-office operations. They can be embedded in customer-facing applications, monitor user behavior and proactively surface relevant information, offers, or alerts. For example, an ambient agent in a mobile banking app might detect unusual spending patterns and proactively notify the customer or automatically categorize transactions for easier budgeting.

This shift from reactive to proactive, from explicit to implicit, marks a significant evolution in the design and deployment of digital assistants. Ambient agents are not just tools; they are intelligent collaborators that enhance productivity, reduce friction, and enable more human-centric experiences.

To realize their full potential, ambient agents must be designed and governed according to a set of foundational principles. These principles ensure that ambient agents are effective, trustworthy, and aligned with organizational goals.

KEY FEATURES OF AMBIENT AGENTS



EVENT-DRIVEN

They respond to triggers within digital environments, such as the arrival of an invoice or detection of unusual activity, rather than waiting for user input.



CONTEXT-AWARE

By integrating data from multiple sources, they understand the user's context, such as preferences, historical behavior, and real-time signals to deliver relevant and timely assistance.



PROACTIVE

Instead of reacting to commands, ambient agents anticipate needs, automate routine tasks, and provide recommendations or alerts without being prompted.



HARMONIOUS INTEGRATION

They are embedded within existing workflows and systems, minimizing disruption and maximizing adoption.



PERSONALIZED

Actions and recommendations are tailored to individual users or scenarios, enhancing relevance and engagement.



SECURE AND COMPLIANT

They operate with stringent privacy and security standards, maintaining detailed logs and audit trails to support compliance.

WHY DO AMBIENT AGENTS MATTER?

The real challenge in enterprise AI adoption isn't just having powerful models or infrastructure, it's integrating AI so that it truly augments daily business operations. Most traditional AI agents or large language models (LLMs) require clear, manual input through chat interfaces or dashboards; they respond to explicit requests and remain dormant otherwise. This user-driven approach can limit their reach and scalability in complex environments.

Ambient agents change this paradigm. Instead of waiting for a user's command, ambient agents are designed for continuous, unified integration with enterprise systems and workflows.

They operate invisibly in the background, monitoring data flows and business events as they happen, automatically triggering intelligent actions without disrupting established background processes. In this sense, ambient isn't strictly a new category of agent, but describes a mode of operation where the agent works unobtrusively and persistently, rather than only in response to foreground tasks. This ambient state allows the agents to learn and adapt passively over time, compounding their impact with every cycle.

Some of the key impacts of ambient agents include:

CONTINUOUS OPERATION

Running persistently in the background, ambient agents monitor and act on events as they happen.

WORKFLOW EMBEDDING

Integration into core business processes allows for automation, optimization, and compliance without disruption existing practices.

SCALABILITY

Ambient agents can handle large volumes of tasks and adapt to changing business needs without additional manual intervention.

SMOOTH USER EXPERIENCE

By reducing cognitive load and automating repetitive tasks, they allow employees and customers to focus on high-value activities.

INNOVATION ENABLEMENT

Automating routine operations creates space for experimentation, new product development, and the exploration of emerging business models.

COMPOUNDING VALUE

Each run helps refine policies and improve signal quality, driving ongoing performance gains.

PRODUCTIVITY GAINS

Automating time-consuming processes frees up human capacity for strategic decision-making and innovation.

PERSONALIZATION

Leveraging contextual and historical data, ambient agents tailor actions and recommendations to individual users and scenarios.

RISK MANAGEMENT AND COMPLIANCE

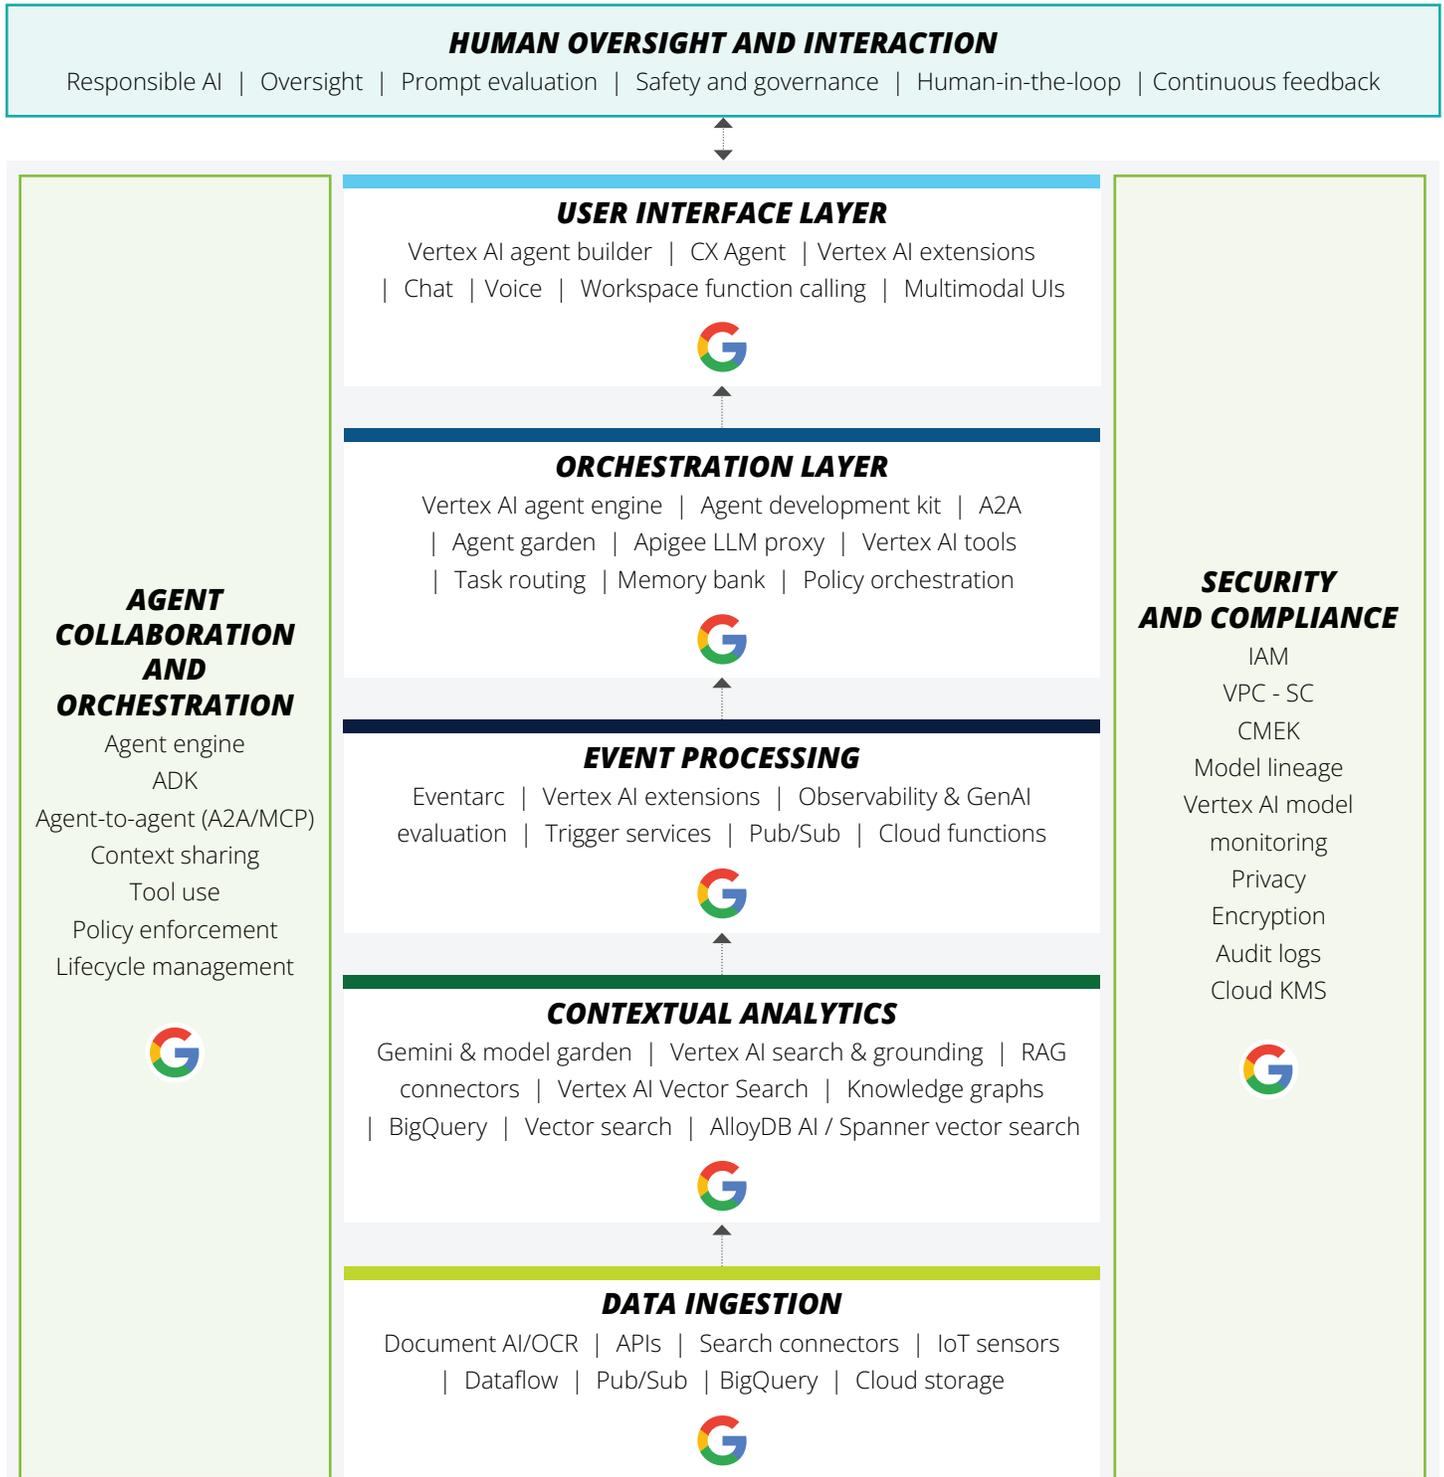
Proactive monitoring, automated compliance checks, and detailed log support auditability and regulatory requirements.

END-TO-END AUTOMATION AND PROCESS ACCELERATION

By connecting seamlessly to enterprise systems and data sources, ambient agents enable automated, closed-loop workflows that reduce manual handoffs and enhance data flows across systems.

ARCHITECTURE: WHAT POWERS AMBIENT AGENTS?

The architecture of ambient agents is purpose-built for the complexity of modern digital ecosystems, blending advanced AI-driven automation with structured human oversight.



VPC - SC: Virtual Private Cloud-Service Controls
CMEK: Customer Managed Encryption Keys
Cloud KMS: Key Management Services

MCP: Multi-Context Protocol
ADK: Agent Development Kit
RAG: Retrieval-Augmented Generation

 Solution and accelerators by Google

KEY ARCHITECTURAL COMPONENTS

DATA INGESTION

Ambient agents begin with real-time, multi-source data acquisition to ensure a holistic, up-to-date environment. BigQuery, Cloud Storage, Document AI/OCR, API connectors, search connectors, IoT sensors, and streaming via Dataflow and Pub/Sub enable consistent aggregation, ingestion, and normalization of diverse data streams.

CONTEXTUAL ANALYTICS

Advanced analytics and semantic understanding set the foundation for intelligence. Gemini, Vertex AI Model Garden, Vertex AI Search and Grounding, RAG connectors, and the Vertex AI Vector Search interpret data, infer context, connect to knowledge sources (including Knowledge Graphs), and provide Retrieval-Augmented Generation (RAG) for deeper, contextual insights.

EVENT PROCESSING

To enable instant responsiveness, the architecture leverages a robust, event-driven backbone. Eventarc, Pub/Sub, Vertex AI Extensions, Cloud Functions, Trigger Services, and comprehensive observability along with GenAI-specific evaluation power real-time, scalable event detection and processing, coordinating triggers with downstream AI capabilities.

ORCHESTRATION LAYER

This layer manages complex workflows and integrates with core systems (e.g., banking, ERPs). Vertex AI Agent Engine, Agent Development Kit (ADK), A2A, Apigee LLM Proxy, Agent Garden, Task Routing, Memory Bank, and Policy Orchestration manage workflows and assign tasks, enforce policies, and maintain shared memory.

AGENT COLLABORATION AND ORCHESTRATION

Agent Engine, ADK, Agent-to-Agent (A2A/MCP) protocols, Context Sharing, Tool Use, Policy Enforcement, and Lifecycle Management enable autonomous agents to communicate, share context, leverage tools collaboratively, and coordinate as teams. This layer supports real-time sharing of state and capabilities, orchestrating collective intelligence and unlocking automation patterns such as swarming, escalation, and dynamic role-based delegation helping agents adaptively solve challenges that exceed the capabilities of any one agent working alone.

USER INTERFACE LAYER

Ambient agents keep users informed and empowered with light, context-aware, human-centric interfaces. Vertex AI Agent Builder, CX Agent, Vertex AI Extensions, plus chat, voice, Workspace function calling, and multimodal UIs, present timely notifications, actionable insights, and approval requests, embedding agent interactions directly in the flow of work.

SECURITY AND COMPLIANCE

Robust security measures such as identity management, encryption, audit logging are essential, especially for sensitive financial data. IAM, VPC-SC, CMEK, Model Lineage, Vertex AI Model Monitoring, Cloud KMS, privacy, encryption, and audit logs combine for end-to-end protection, governance, and transparency, especially critical in sectors with high compliance requirements.

THE ROLE OF HUMAN INSIGHT

Despite their autonomy, ambient agents are designed to augment and not replace human judgment. Responsible AI methodologies, prompt evaluation, safety and governance frameworks, and human-in-the-loop practices help ambient agents operate within set guardrails, with meaningful oversight and continuous feedback. Human insight is critical for:

OVERSIGHT

Defining operational guardrails, approving critical actions, and reviewing exceptions

TRAINING

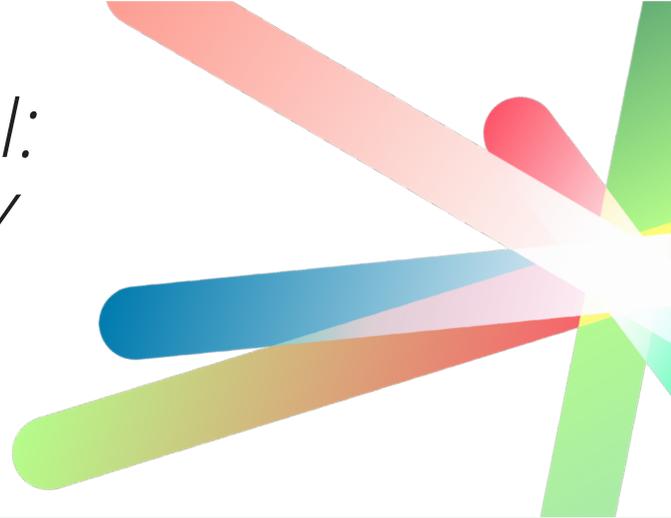
Providing feedback to refine AI models and ensure continuous learning

ETHICS AND TRUST

Ensuring decisions are ethical, transparent, and compliant with regulations

In highly regulated sectors like financial services, this partnership between humans and ambient agents is essential for building trust and ensuring responsible innovation.

AMBIENT AGENTS IN FSI: HARNESSING INDUSTRY TRENDS WITH GOOGLE CLOUD SOLUTIONS



CURRENT TRENDS IN FSI AND GOOGLE CLOUD

The financial services industry (FSI) is rapidly advancing its use of AI, automation, and cloud technology. Google's Gemini Enterprise empowers financial institutions to automate complex tasks such as fraud detection and compliance, by orchestrating AI agents that can autonomously analyze data and trigger actions. These agents communicate and collaborate seamlessly, enabling

efficient information sharing and coordinated responses. Google Cloud provides the infrastructure and tools necessary for financial services to build and deploy these intelligent agents, enhancing efficiency and decision-making in areas like risk management and customer service. Below are a few key trends:

 <p>HYPER AUTOMATION End-to-end financial workflows powered by AI-driven automation, reducing manual effort and boosting operational efficiency</p>	 <p>EMBEDDED INTELLIGENCE Real-time, automated decision making within core systems through seamless AI integration across business functions</p>
 <p>SECURITY AND PRIVACY Comprehensive security, encryption, and access controls safeguarding AI deployments in sensitive financial environments</p>	 <p>REGULATORY COMPLIANCE Proactive compliance risk identification and audit trail generation supported by AI-powered monitoring and analytics</p>
 <p>DATA-DRIVEN PERSONALIZATION Highly tailored product recommendations, fraud alerts, and communications delivered through advanced analytics and AI models</p>	 <p>MULTI-AGENT COLLABORATION Networks of specialized AI agents interacting and sharing insights to address complex challenges and optimize outcomes</p>
 <p>EXPLAINABILITY AND AUDITABILITY Transparency and traceability for automated decisions achieved through explainable AI solutions supporting regulatory trust</p>	 <p>MULTIMODAL INTERFACES Seamless, user-friendly engagement across chatbots, voice assistants, and digital portals for both customers and staff</p>

**WHAT'S HAPPENING
IN THE INDUSTRY**

TREND

**THE FUTURE:
HOW AMBIENT AGENTS CAN
TAP INTO THESE TRENDS**

Institutions are scaling automation with Purpose-Built ML and Google Cloud AI Platform, streamlining onboarding, document review, and transaction processing for greater speed and compliance.



**HYPER
AUTOMATION**

Ambient agents orchestrate and optimize complete workflows autonomously by handling exceptions, learning from outcomes, and driving operational efficiency end-to-end.

Enhanced security with Google Cloud Identity, data encryption, and granular access controls enables safe and compliant deployment of AI and automation solutions.



**SECURITY
AND PRIVACY**

Agents proactively enforce policies, detect and respond to threats, and adapt security protocols to new risks and regulations, all at machine speed.

Banks use BigQuery, Vertex AI, and real-time analytics for tailored offers, fraud alerts, and predictive segmentation driven by deep customer insights.



**DATA-DRIVEN
PERSONALIZATION**

Agents synthesize real-time signals and context from multiple data sources to deliver hyper-personalized experiences everywhere a customer interacts.

Explainable AI tools on Google Cloud provide transparency into automated decisions, helping satisfy regulator and client demands for trust and clarity.



**EXPLAINABILITY
AND AUDIBILITY**

Agents generate real-time, detailed, and accessible explanations and audit logs for every decision, boosting compliance and stakeholder confidence.

Firms embed Google Cloud-based AI in payments, lending, and risk decisioning using APIs and event-driven architectures to respond instantly to transactions and events.



**EMBEDDED
INTELLIGENCE**

Agents autonomously monitor events and activate intelligent actions across systems, ensuring seamless, adaptive, and timely decisions without manual intervention.

Continuous risk monitoring and automated reporting are powered by Google Cloud's security and compliance frameworks, though critical review still requires manual input.



**REGULATORY
COMPLIANCE**

Agents enable real-time compliance checks, autonomous reporting, and immediate flagging of risk, minimizing reliance on manual reviews and audits.

Institutions are deploying specialized agents for fraud, customer service, and compliance, coordinating via Google APIs and shared data to resolve complex cases.



**MULTI-AGENT
COLLABORATION**

Ambient agents dynamically collaborate, share expertise, and tackle multi-faceted issues as unified teams, accelerating resolution and learning in real time.

Institutions use CX Agent, Contact Center AI, and omnichannel portals to deliver services across chat, voice, and digital channels, though experiences may be fragmented.



**MULTIMODAL
INTERFACES**

Ambient agents ensure unified, context-aware support across every channel, carrying understanding and history seamlessly for frictionless, consistent interaction.



PREREQUISITES FOR REALIZING AMBIENT AGENTS: BUILDING THE FOUNDATION

While the vision of ambient agents is compelling, successful implementation requires addressing several foundational prerequisites. These technical and organizational requirements serve as the essential framework that enables truly autonomous, context-aware agents to operate effectively in enterprise environments.



KNOWLEDGE MODELS AS ORGANIZATION MAPS

The foundation of effective ambient agents lies in comprehensive knowledge models that serve as organizational maps, constraining and guiding agent behavior within appropriate boundaries. These knowledge models could encompass process documentation, domain expertise, regulatory frameworks, etc. On platforms like Google Cloud, these knowledge models are often implemented as Enterprise Knowledge Graphs (EKG), allowing agents to dynamically access and reason over up-to-date information and relationships across the organization.



LEGACY SYSTEM MODERNIZATION AND THE STRANGLER PATTERN

Many legacy applications in financial institutions were not designed for event-driven architectures, presenting a significant implementation challenge. These systems often rely on batch processing, scheduled jobs, and manual triggers that are incompatible with the real-time, event-driven nature of ambient agents. Rather than attempting wholesale system replacement, organizations can adopt the strangler pattern to incrementally transition to ambient agent capabilities.



COMPREHENSIVE AUDITABILITY THROUGH EVENT SOURCING

Given the autonomous nature of ambient agents and the regulatory requirements of financial services, comprehensive auditability becomes paramount. Traditional logging approaches are insufficient for the complex, multi-agent interactions that characterize ambient systems.



REAL-TIME DATA INFRASTRUCTURE

Ambient agents require low-latency data pipelines capable of processing and routing events in milliseconds. This necessitates robust streaming architectures using technologies like Google Cloud's Dataflow and Pub/Sub.

USE CASES FOR AMBIENT AGENTS IN FINANCIAL SERVICES

Building on these emerging trends in AI and automation within financial services, organizations can now explore how ambient agents can deliver tangible business value. By orchestrating networks of specialized AI agents on Google Cloud, firms can automate complex, multi-step processes while maintaining

robust human oversight. The following use cases illustrate how leading financial institutions can leverage Google Cloud-powered agentic AI to transform core operations, enhance compliance, and elevate customer experiences.

USE CASES	GOOGLE CLOUD SOLUTIONS THAT WOULD BE CRITICAL	AGENTS THAT COULD BE INVOLVED	HUMAN OVERSIGHT AND INPUTS REQUIRED
PAYMENT ECOSYSTEM AUTOMATION	Agent Payments Protocol (AP2), Document AI, Vertex AI*, Pub/Sub, BigQuery, Cloud Functions	Invoice Agent, Payment Agent, Approval Agent, Reconciliation Agent	Oversight of exceptions, review of flagged transactions, audit process
LOAN ORIGINATION AND SERVICING	Document AI, Vertex AI*, BigQuery, AutoML, Cloud Functions	Application Intake Agent, Document Verification Agent, Credit Risk Agent, Approval Agent	Oversight of edge cases, final credit decisions, customer communication
ANTI-MONEY LAUNDERING (AML) AND FRAUD MONITORING	Anti Money Laundering AI, Vertex AI*, BigQuery, Pub/Sub, Looker, Security Command Center	Transaction Monitoring Agent, Pattern Detection Agent, Case Management Agent	Oversight of escalated cases, confirmation or dismissal of alerts
CUSTOMER ONBOARDING AND KYC	Document AI, Vertex AI*, BigQuery, Cloud Functions, Identity Platform	Data Extraction Agent, Verification Agent, Compliance Agent	Oversight of discrepancies, onboarding approvals, compliance checks
CLAIMS PROCESSING (INSURANCE)	Document AI, Vertex AI*, Pub/Sub, BigQuery, Cloud Functions	Intake Agent, Assessment Agent, Validation Agent, Payout Agent	Oversight of complex claims, high-value payouts, customer support
REGULATORY REPORTING	BigQuery, Vertex AI*, Pub/Sub, Looker, Cloud Functions	Data Aggregation Agent, Validation Agent, Submission Agent	Oversight of final reports, sign-off, regulator engagement
PERSONALIZED FINANCIAL ADVISORY	Vertex AI, BigQuery, CX Agent, Looker, Cloud Functions	Portfolio Analysis Agent, Risk Assessment Agent, Customer Interaction Agent	Oversight of recommendations, compliance review, customer escalation

* Vertex AI includes Google Cloud solutions like Agent Builder, Agent Runtime, Memory Bank, Enterprise Search, Tool Call Orchestration, Agent Development Kit (ADK), etc.

RISKS AND CHALLENGES

While ambient agents offer transformative potential, their adoption in financial services is not without risks and challenges. Addressing these proactively is essential for successful deployment and long-term value realization.

01

DATA PRIVACY AND SECURITY

Ambient agents often require access to sensitive financial and personal data. Ensuring robust data protection, encryption, and compliance with regulations such as GDPR, CCPA, and industry-specific standards (e.g., PCI DSS) is paramount. Institutions must implement strong identity and access management, regular security audits, and real-time monitoring to prevent unauthorized access or data breaches.

02

EXPLAINABILITY AND TRANSPARENCY

As ambient agents make increasingly complex and autonomous decisions, providing transparency and explainability is critical. Stakeholders, including regulators, customers, and internal auditors must be able to understand how and why decisions are made. This requires the integration of explainable AI (XAI) techniques and the maintenance of comprehensive audit logs.

03

INTEGRATION COMPLEXITY

Many financial institutions operate with legacy systems and fragmented technology stacks. Integrating ambient agents into these environments can be complex, requiring careful planning, robust APIs, and middleware solutions. Ensuring interoperability and minimizing disruption to existing workflows are key challenges.

04

CHANGE MANAGEMENT AND ADOPTION

The introduction of ambient agents represents a significant change for employees and customers. Effective change management including communication, training, and support are essential to drive adoption and realize the full benefits of automation. Organizations must also address cultural resistance and ensure that employees see ambient agents as collaborators rather than threats.

05

BIAS AND FAIRNESS

AI models, if not carefully designed and monitored, can perpetuate or even amplify biases present in historical data. This is particularly sensitive in areas such as lending, fraud detection, and AML, where biased decisions can have serious ethical and legal implications. Continuous monitoring, bias audits, and human oversight are necessary to ensure fairness and compliance.

06

REGULATORY AND ETHICAL CONSIDERATIONS

Regulatory requirements in financial services are constantly evolving. Institutions must ensure that ambient agents comply with all relevant laws and guidelines and be prepared to adapt to new regulations as they emerge. Ethical considerations, such as customer consent and data usage, must also be addressed transparently.

Looking ahead, ambient agents aim to redefine the financial services landscape by harnessing advanced AI and cloud technologies to streamline processes and deliver more personalized, intuitive experiences.

Realizing this vision requires not just technical modernization, but also robust governance and proactive change management. While ongoing attention to privacy, integration, and ethics is

crucial, effective human oversight and thoughtful design can help institutions unlock new levels of efficiency, resilience, and innovation, ensuring they thrive in an evolving digital future.

AUTHORS

Shailender Sidhu

Principal
Deloitte Consulting LLP
ssidhu@deloitte.com

Prakul Sharma

Principal
Deloitte Consulting LLP
praksharma@deloitte.com

Naveen Balawat

Senior Manager
Deloitte Consulting LLP
nbalawat@deloitte.com

Drew Davidson

Senior Manager
Deloitte Consulting LLP
ddavidson@deloitte.com

Jatin Balani

Manager
Deloitte Consulting LLP
jbalani@deloitte.com

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte provides industry-leading audit and assurance, tax and related services, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's approximately 457,000 people worldwide make an impact that matters at www.deloitte.com.