Arya.ai
an aurionpro company

Making AI Agents Work in Enterprises

Exploring a Novel Protocol for Integration and Agent Orchestration

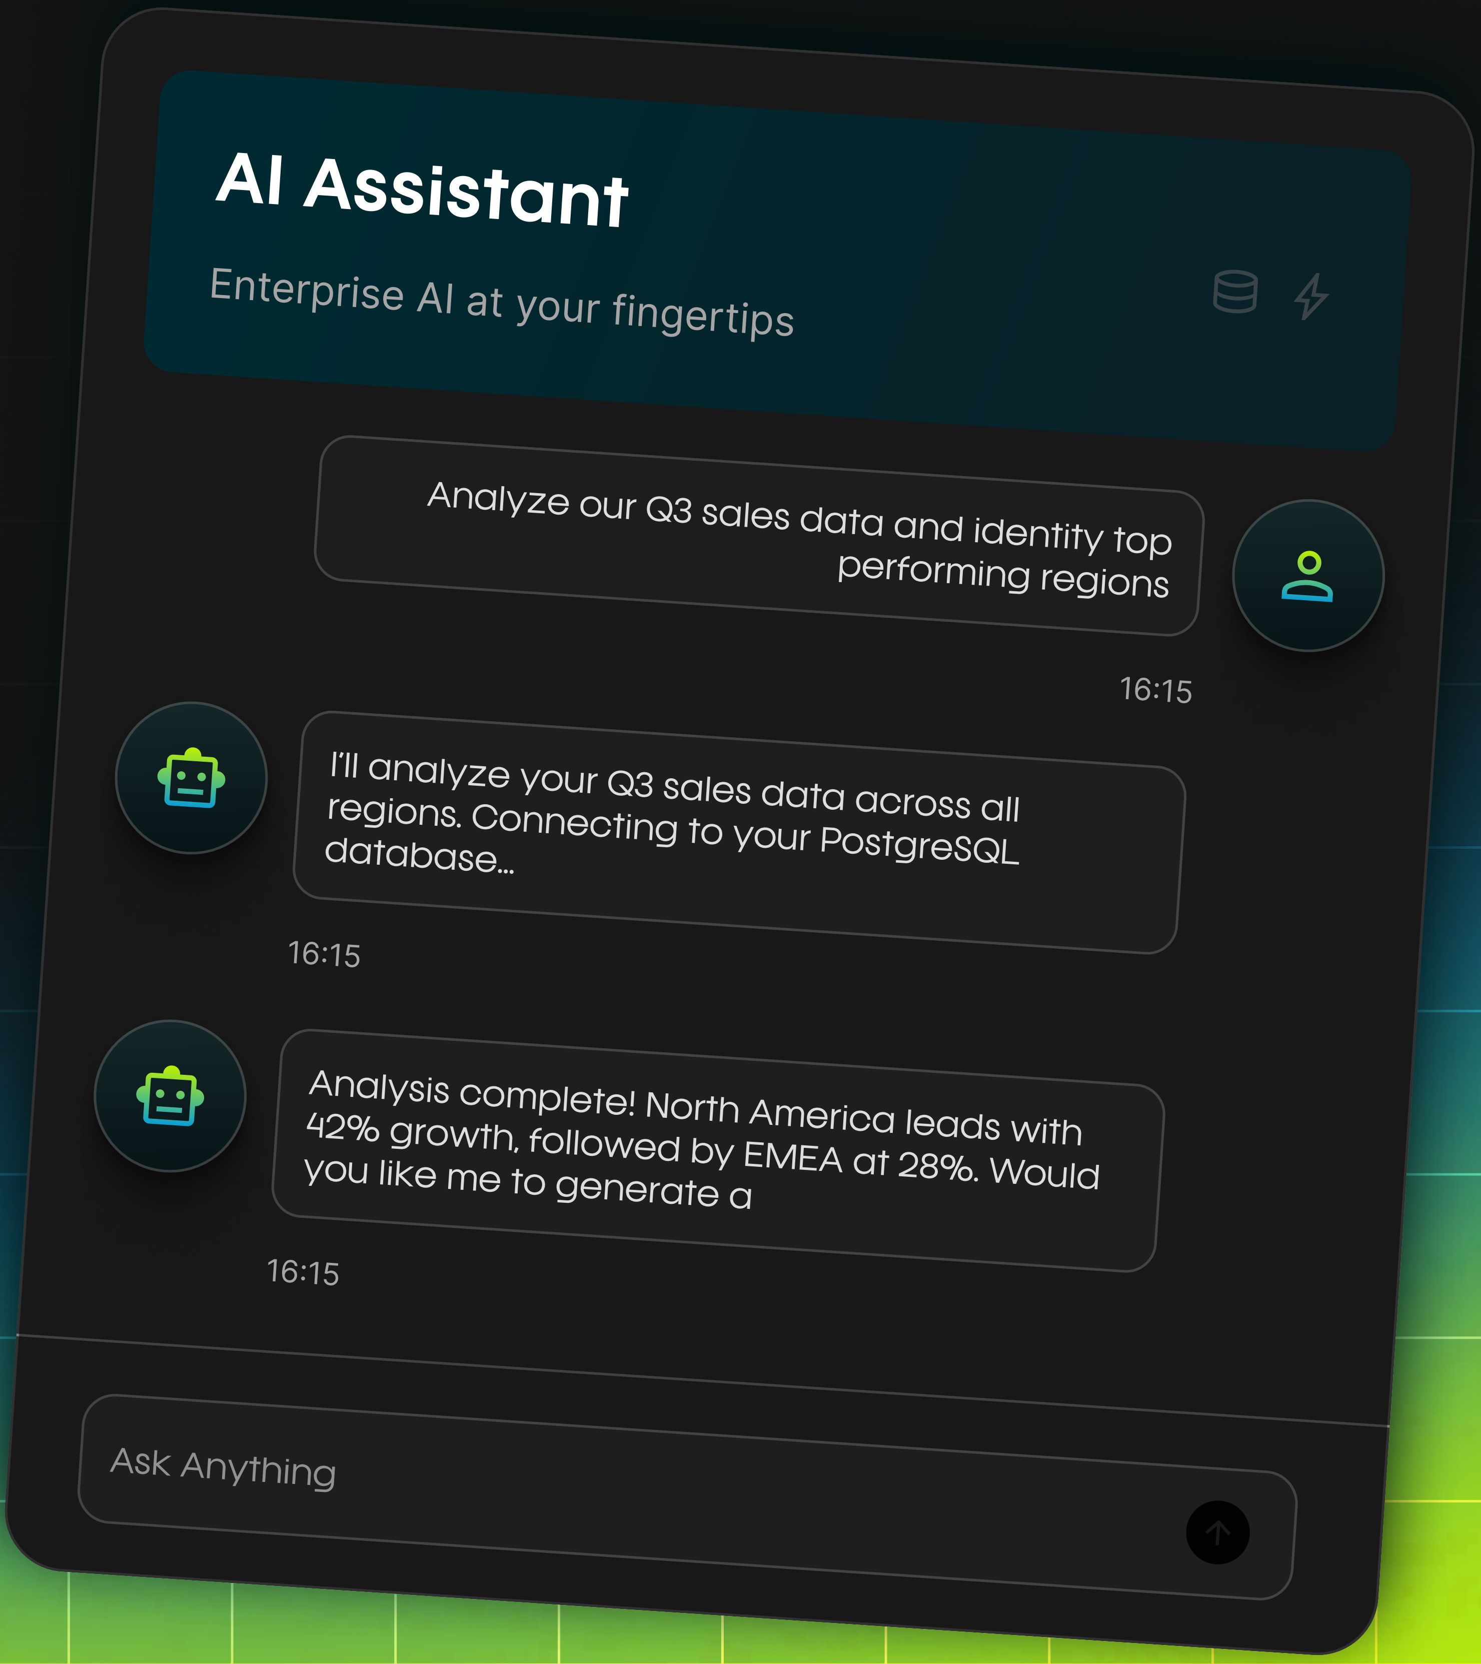

AI Assistant

Enterprise AI at your fingertips

Analyze our Q3 sales data and identity top performing regions

16:15

I'll analyze your Q3 sales data across all regions. Connecting to your PostgreSQL database...

16:15

Analysis complete! North America leads with 42% growth, followed by EMEA at 28%. Would you like me to generate a

16:15

Executive Summary

AI agents tend to lose context across extended workflows, and that is where they falter in real-world applications. What if we could use a new protocol to provide AI agents the ability to connect with the right sources to obtain context and never lose sight of the task at hand?

The Model Context Protocol (MCP) is an open standard designed to unify the way AI models connect with external data sources, services, and tools. Originally introduced by Anthropic, MCP has rapidly gained support from industry leaders.

Here, we explore MCP's value proposition, technical architecture, enterprise benefits, implementation best practices, and real-world use cases, culminating in recommendations for organizations seeking to modernize their AI strategy.

Background

Enterprises today are inundated with AI research breakthroughs but struggle to translate innovation into production. Legacy systems remain fragmented, and bespoke integrations hinder reuse. AI models often lack structured access to the data and services they need to deliver value, resulting in siloed pilots rather than scalable deployments.

The Integration Challenge

- Multiple Data Silos: CRM, ERP, analytics platforms, and custom databases create disparate endpoints.
- Custom Connectors: Each system requires bespoke API adapters, increasing development time and maintenance overhead.
- Security & Compliance: Sensitive data must be handled under strict governance, yet many AI integrations rely on third-party proxies or manual processes.
- Agility Gap: Rapid innovation cycles are stifled by monolithic integration projects that take months to complete.

From Arya.ai's Desk

Accessing additional data points and external services for AI agents has always been difficult, especially in an enterprise setting. Adding these abilities and providing agents with access to additional tools required a lot of work. Model Context Protocol standardizes this integration by negating the need for adding bespoke connectors, costly rewrites, and security compromises.

Yet standardization alone isn't enough. Workflows in an enterprise are complex, where one agent alone cannot bear the burden. We need to devise agentic workflows where multiple agents work in tandem and each agent draws context from external sources when needed.
So you need to orchestrate these agents on one platform. That is why we built Weave. MCP gives agents the standardized plumbing, and Weave provides the automation brain. Together, empowering organizations to deploy sophisticated, multi-agent AI solutions at scale.

Deekshith Marla

FOUNDER

ARYA.AI

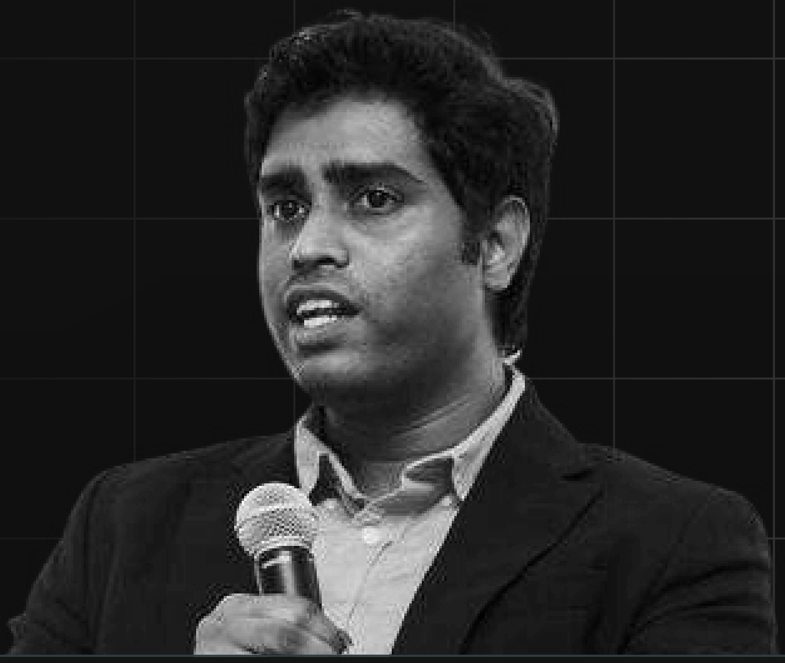

01

What is MCP?

MCP solves a core challenge: connecting LLMs and AI agents to external data sources, tools, and services.

MCP standardizes how AI models connect to external sources, enabling them to:

- Fetch information
- Interact with APIs
- Execute tasks beyond their built-in knowledge

The Model Context Protocol (MCP) is an open standard that provides a unified way to connect LLMs with various data sources and tools. Think of MCP as the "USB-C port" for AI applications. Instead of custom interfaces for each system, it offers one standardized protocol for plugging in databases, services, and software APIs.

Originally pioneered by Anthropic in 2024, MCP has quickly gained industry traction as a key enabler for building AI agents and complex workflows. Major tech players like OpenAI have embraced MCP, signaling that it is on its way to becoming a new standard for context integration in AI systems.

What Does MCP Mean for Enterprises?

For enterprises, MCP's promise of standardized AI integration is especially compelling. Organizations manage vast amounts of sensitive data across siloed legacy systems under strict regulatory oversight. MCP matters to enterprises because it can improve control, traceability, and performance in how AI systems access and use data.

In a compliance-heavy environment, having a consistent protocol means AI-driven solutions can be deployed with greater confidence, knowing that every interaction between an AI model and enterprise systems is logged and structured. Ultimately, MCP enables enterprises to leverage advanced AI capabilities (from customer-facing chatbots to back-office analytics) in a more secure, auditable, and efficient manner.

02

Strategic Advantage for Enterprises

Greater Agility in Innovation

MCP makes it much easier to develop and iterate on AI solutions. Teams can plug in new data sources or tools rapidly without rebuilding integrations.

This agility means a bank or an organization could prototype a new AI-driven service (like a mortgage chatbot or a fraud detection agent) in weeks rather than months.

The ability to "swap in" better models or additional data sources on demand gives organizations a flexible, modular AI architecture that evolves with their strategy.

Enhanced Data Security and Control

While integrating systems, MCP follows best practices to keep data within an enterprise's control.

All data retrieval and actions occur through servers that an organization can host in its secure environment, rather than sending sensitive information through third-party platforms.

This setup means enterprises can enforce their security protocols (encryption, access controls, monitoring) around the MCP connections.

Operational Efficiency Gains

By enabling AI to directly interface with multiple systems, tasks that used to require manual data gathering, entry, and cross-checking can be done in seconds.

This leads to faster turnaround times and lower operational costs. For instance, financial reporting that once took weeks of collating spreadsheets can be generated by an AI in real-time, since MCP feeds the AI model all necessary data, rules, and context without human handoffs.

Innovation and New Service Enablement

MCP doesn’t just make existing processes faster – it opens the door to new capabilities and services. Organizations can combine AI models with their rich data stores to create offerings that were previously not feasible.

For example, a bank might launch a personalized financial coach that uses an LLM connected via MCP to a customer’s transaction history, budgeting tools, and market data feeds to provide tailored advice in real time.

This encourages experimentation and adoption of the best technologies over time. In short, MCP helps firms modernize their services and differentiate themselves while maintaining the governance needed in a regulated industry.

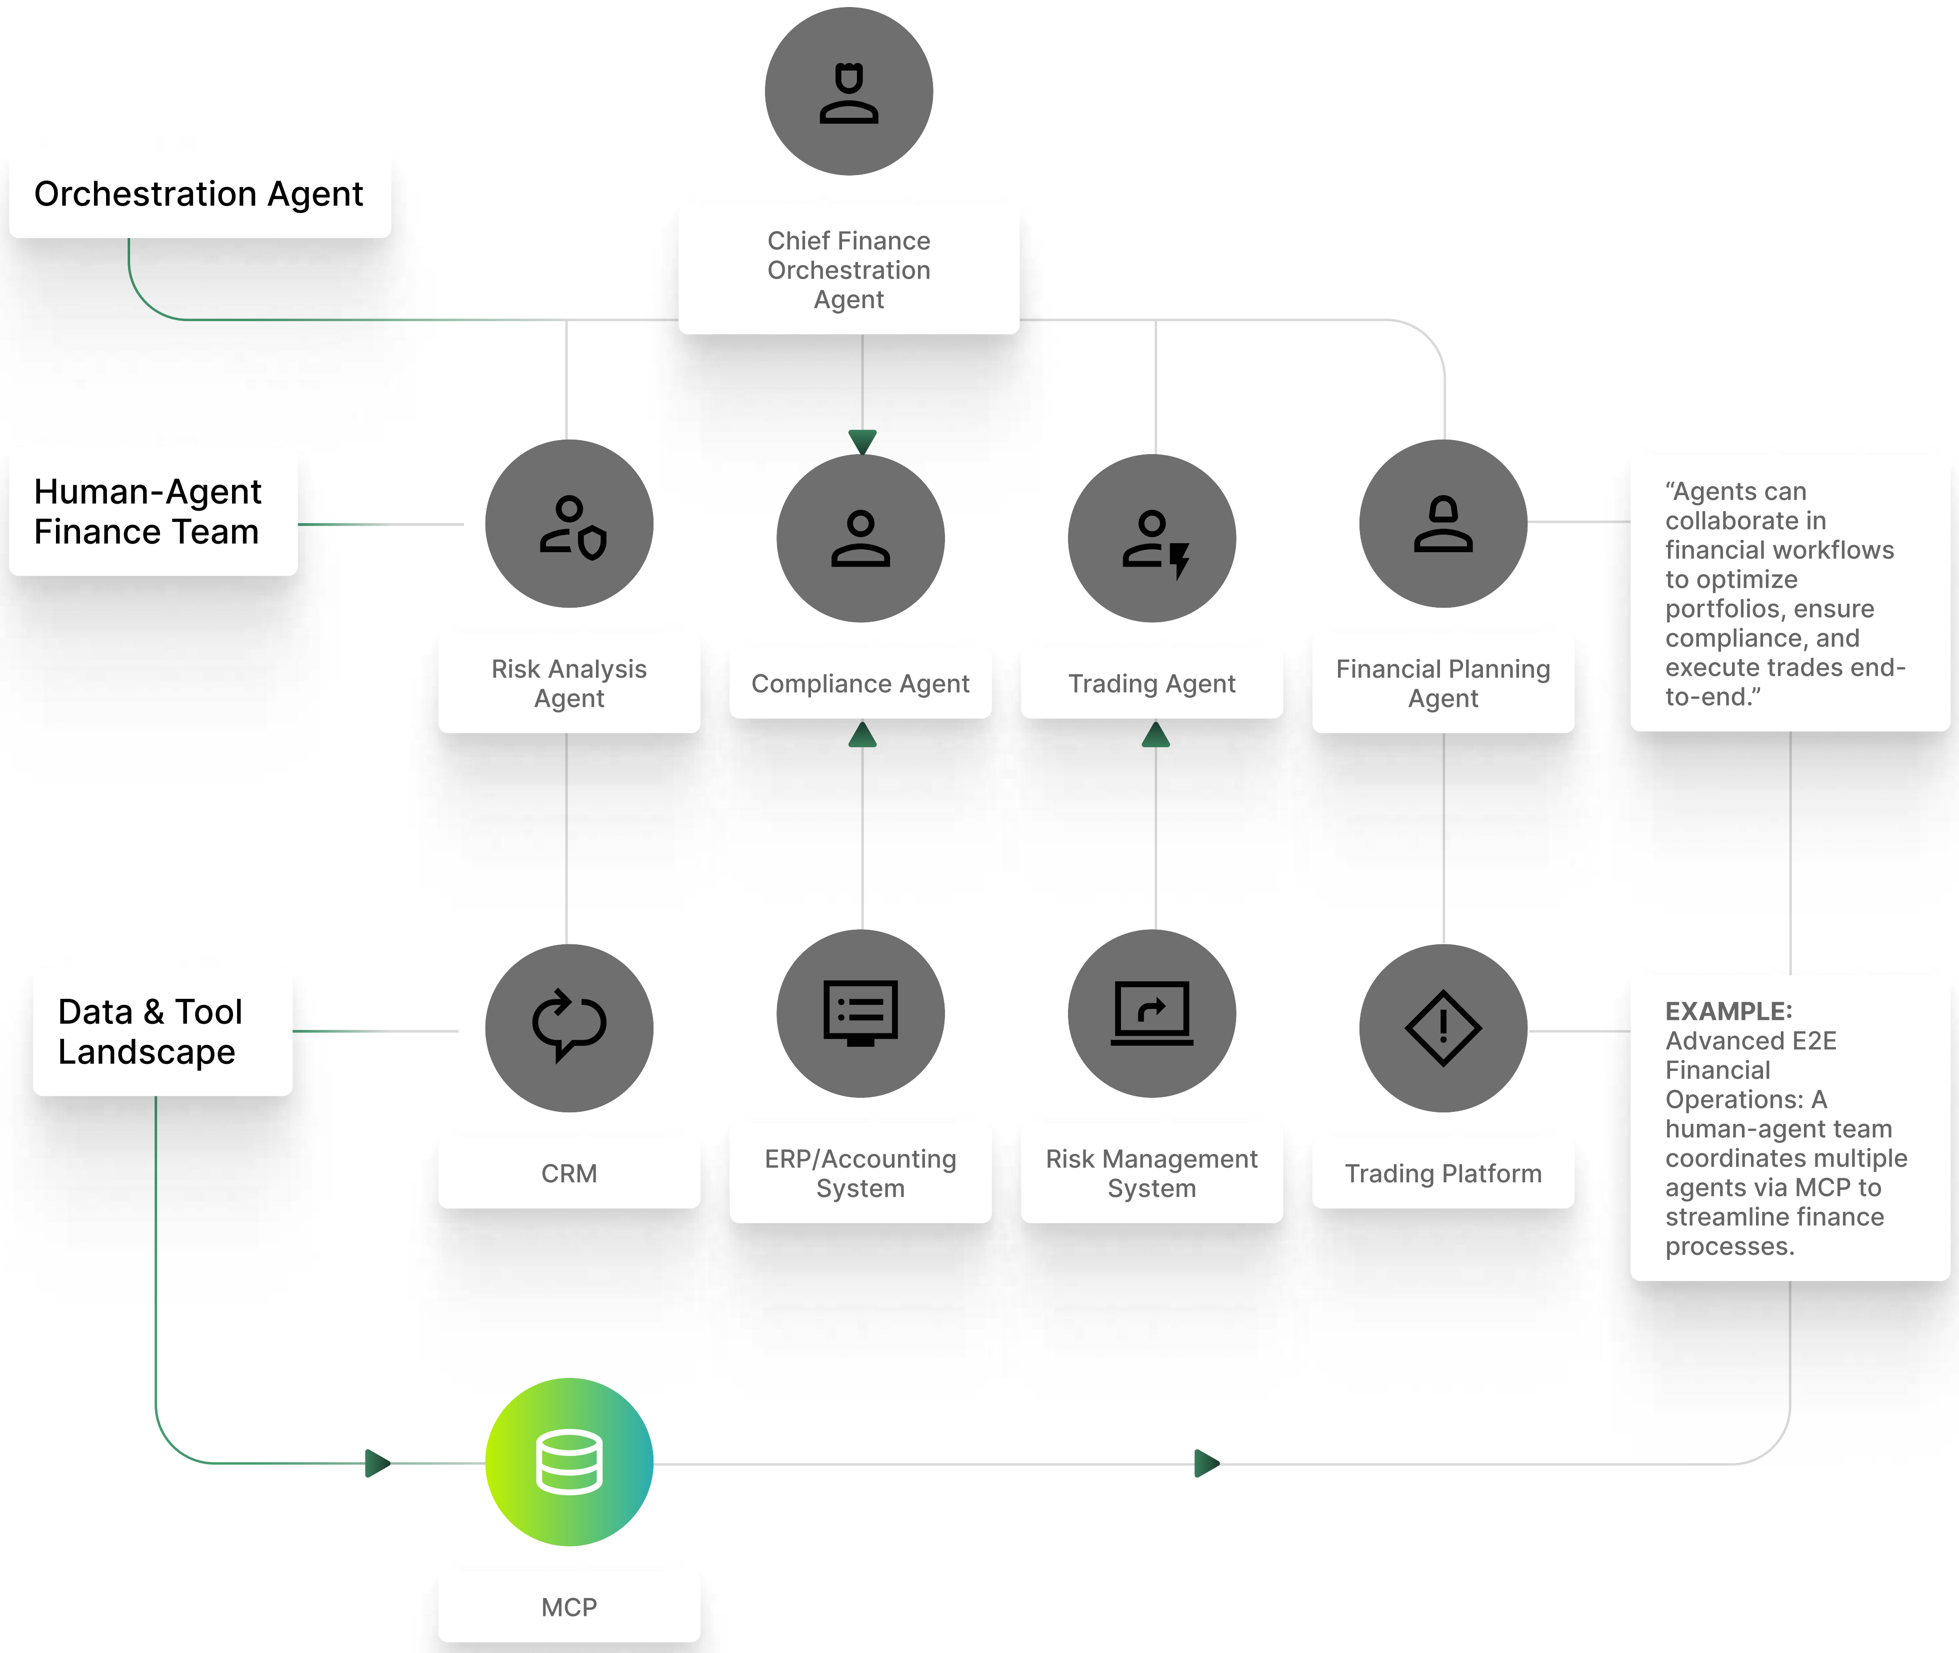

“Agents can collaborate in financial workflows to optimize portfolios, ensure compliance, and execute trades end-to-end.”

EXAMPLE: Advanced E2E Financial Operations: A human-agent team coordinates multiple agents via MCP to streamline finance processes.

03

Agent-to-External Source & Agent-to-Agent Architecture

Agent-to-External Integration

Connection: Agents connect to external sources through Model Context Protocol.

Security Enforcement: MCP's built-in authentication (OAuth2/JWT), authorization (RBAC/ABAC), encryption, and rate limiting guard every interaction.

Event Hooks: Agents subscribe via MCP channels (webhooks or messaging streams) for real-time triggers.

Agent-to-Agent Communication

Connection: Agents dynamically discover and delegate tasks to specialized sub-agents (e.g., a Customer Onboarding Agent invokes KYC, Document Verification, and Risk Assessment Agents in sequence).

Orchestration Layer: An Agent Router consults the MCP Service Registry to locate peers and forward context-rich payloads.

Understand the Key Components Involved:

By enabling AI to directly interface with multiple systems, tasks that used to require manual data gathering, entry, and cross-checking can be done in seconds.

This leads to faster turnaround times and lower operational costs. For instance, financial reporting that once took weeks of collating spreadsheets can be generated by an AI in real-time, since MCP feeds the AI model all necessary data, rules, and context without human handoffs.

Agent Router

Routes requests among agents based on capability and load.

Service Registry

Stores metadata on available agents, endpoints, and versions.

Security Agent

Centralizes MCP-level authentication, authorization, and audit logging.

Connector Agents

MCP adapters for external data sources, systems, and tools.

Business Value of the Architecture:

By orchestrating AI agents through the Model Context Protocol, your organization unlocks rapid time-to-value, where AI workflows can be shipped in weeks. The built-in security across every step, from discovery to external calls, is encrypted and fully auditable. We can also add, swap, or extend specialized agents on demand.

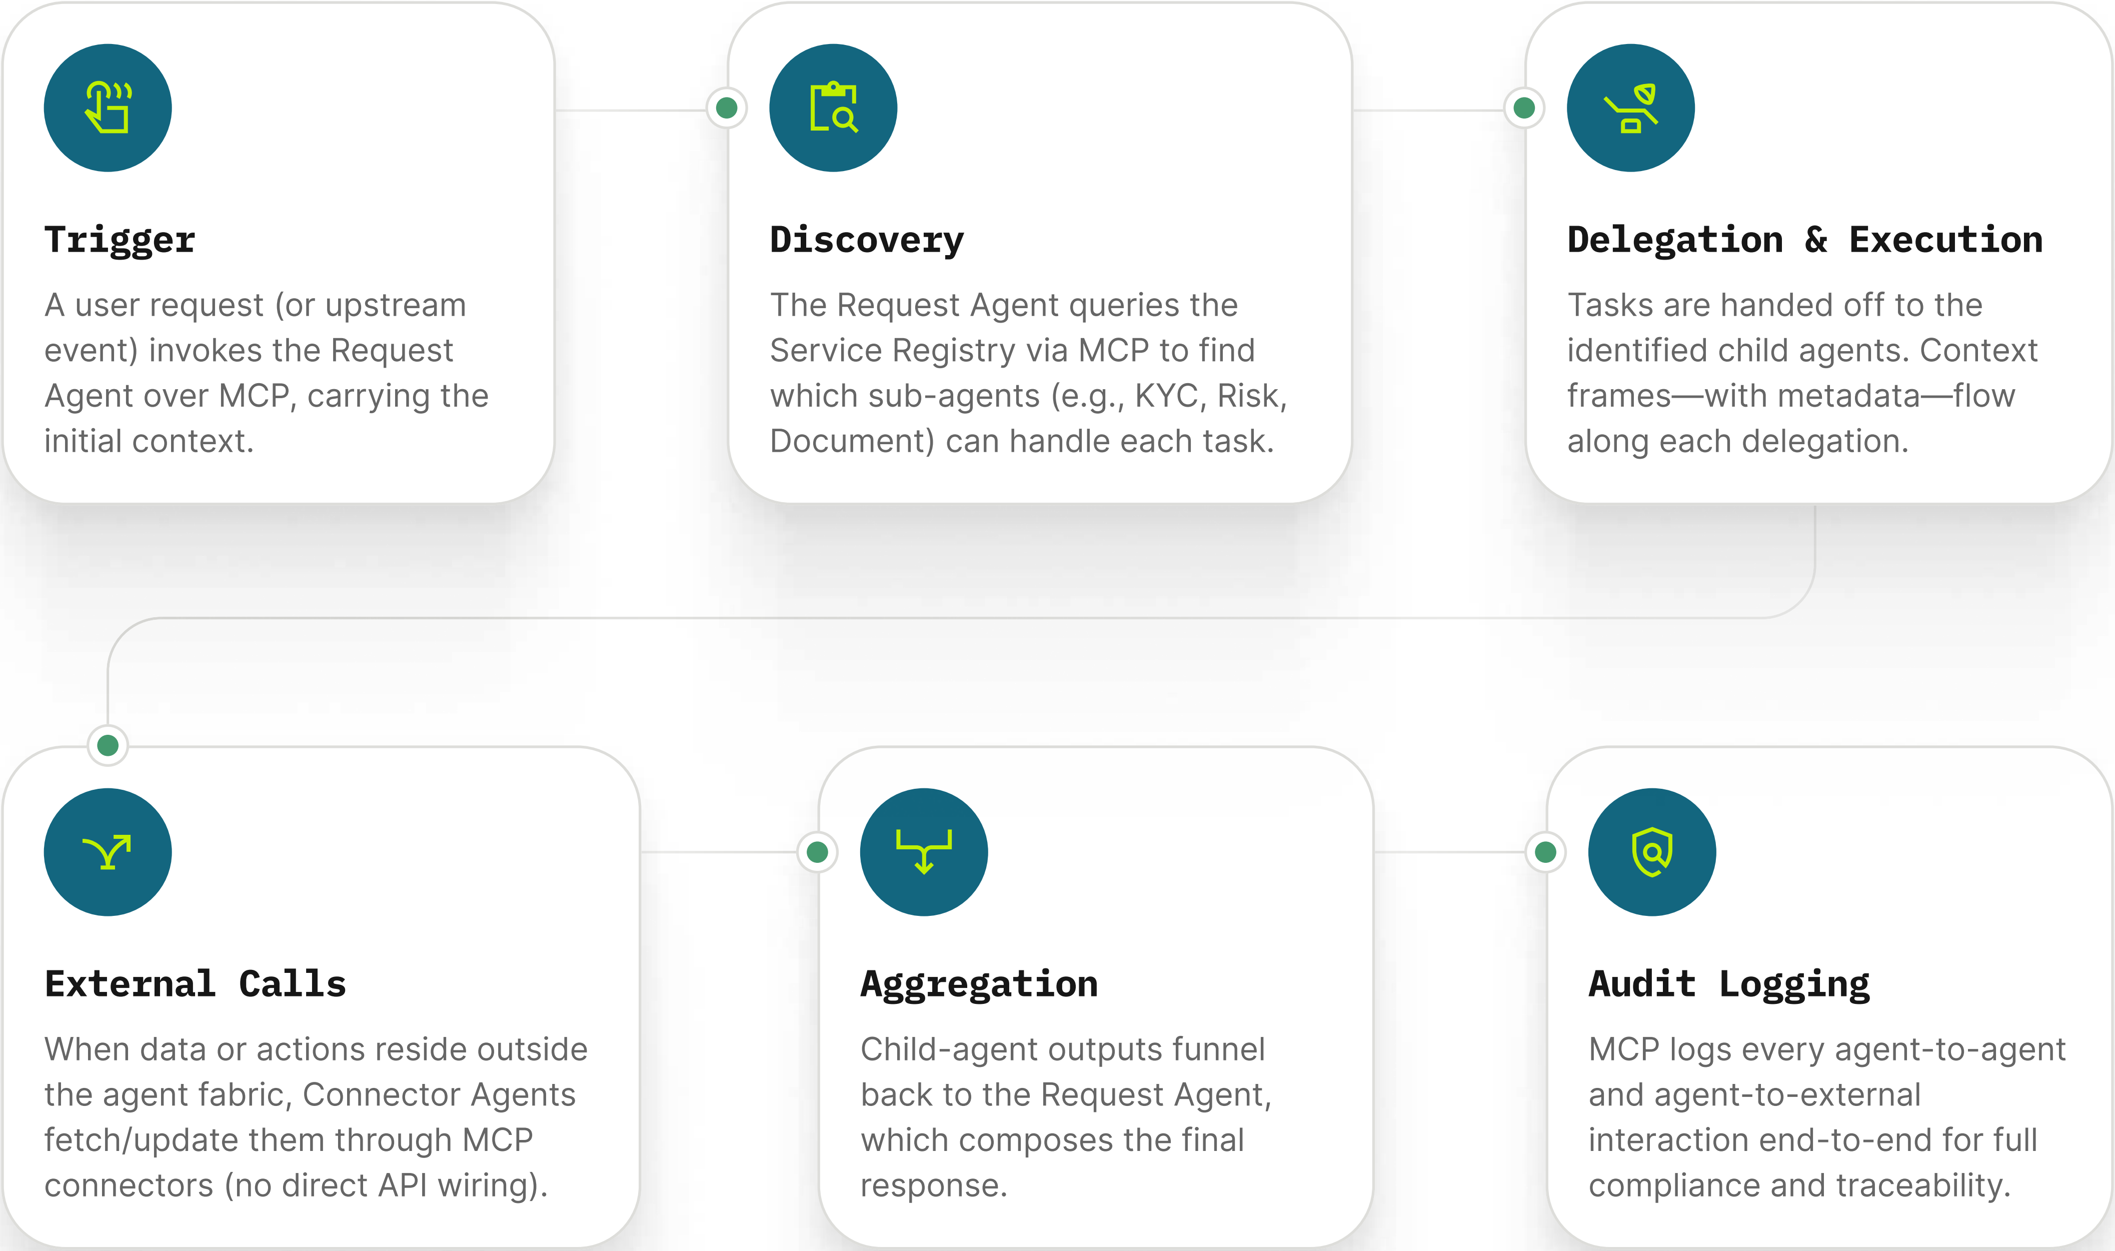

Trigger

A user request (or upstream event) invokes the Request Agent over MCP, carrying the initial context.

Discovery

The Request Agent queries the Service Registry via MCP to find which sub-agents (e.g., KYC, Risk, Document) can handle each task.

Delegation & Execution

Tasks are handed off to the identified child agents. Context frames—with metadata—flow along each delegation.

External Calls

When data or actions reside outside the agent fabric, Connector Agents fetch/update them through MCP connectors (no direct API wiring).

Aggregation

Child-agent outputs funnel back to the Request Agent, which composes the final response.

Audit Logging

MCP logs every agent-to-agent and agent-to-external interaction end-to-end for full compliance and traceability.

04

Why not reverse engineer it?

The last few years were all about business leaders exploring the prospect of integrating AI into “systems.” But MCP reverse engineers the entire thing. It inverts the traditional model-first approach—rather than shoehorning AI into each bespoke application, it treats every database, API, and service as a modular context endpoint that can be plugged directly into the AI.

In doing so, MCP deconstructs monolithic integrations into a standardized protocol layer, allowing teams to expose discrete system capabilities on demand. The result? Rapid, maintainable AI deployments that evolve with your infrastructure—no more fragile point-to-point connectors, just a USB-C-style port for any tool or data source you need.

Why Enterprises Need Standardization for MCPs?

The need for standardization arises from the inefficiencies of previous processes.

What it used to be.

- When the business team requested a new workflow, the IT team faced extensive efforts to build it, often involving complex API integrations.
- This not only duplicated efforts but also created potential inconsistencies, as each integration was bespoke.
- For instance, connecting to a new database might require weeks of development, delaying business agility.

What it is now.

- With MCP, the process shifts to a one-time IT effort to integrate the MCP client, after which business teams can design or modify workflows using the standardized interface.
- This minimizes IT involvement, empowering business teams to be more agile and reducing overall workload.
- The client-server architecture facilitates this, with MCP servers exposing data sources and services.

05

Introducing Weave: A Multi Agent Orchestration Platform by Arya.ai

Weave by Arya.ai is an intuitive platform for enterprises to leverage GenAI by connecting with pre-trained domain modules and 100+ external applications. You get end-to-end control, security, and performance in one seamless platform.

Deploy in Days

Skip lengthy training; drop in the modules you need.

Lock in Accuracy

Domain-specific validation reduces errors and "hallucinations."

Scale with Confidence

Mix, match, and extend modules as your needs evolve—no code rewrites.

Pre-Trained 100+ Domain Modules (APEX)

Each module combines domain rules, data checks, and proven workflows so you hit the ground running:

Finance & Compliance

- **Bank Statement Analyzer**

  Automatically parses statements, categorizes transactions, and highlights anomalies.

- **Invoice Processor**

  Extracts invoice data, validates line items, and reconciles purchase orders.

- **Risk-Scoring Engine**

  Applies configurable credit and fraud risk rules in real time.

Privacy & Security

- **PII Masking**

  Locates and redacts personally identifiable information in documents.

- **Invoice Processor**

  Differentiates live captures from static images for secure authentication.

- **Deepfake Detector**

  Uncovers manipulated audio and video through forensic analysis.

Insights & Automation

- **Sentiment Analysis**

  Delivers channel-wide feedback scoring to track customer mood.

- **Contract Clause Reviewer**

  Flags non-standard clauses and ensures regulatory compliance.

- **Supply-Chain Parser**

  Converts bills of lading into structured feeds for logistics workflows.

...and many more covering healthcare diagnostics, legal discovery, HR onboarding, marketing analytics, and beyond.

External Integrations

Seamlessly connect your MCP client & server applications to MongoDB, PostgreSQL, Kite, and 100+ other on-premise or cloud systems, ensuring real-time data flow, end-to-end security, and centralized governance across all your AI pipelines.

Github
Notion
GitLab
MongoDB
Snowflake
PostgreSQL

About Arya.ai

(AN AURIONPRO COMPANY)

Arya.ai is an enterprise AI solutions provider assisting organizations in unlocking the power of AI. We are working on bridging the gap between advanced technology and real-world challenges. Our offerings empower enterprises with solutions to integrate, manage, and scale AI across a range of functions.

Our Vision

We envision a future where AI is not a siloed experiment but a seamless extension of every business process. Arya.ai aims to provide an intelligence layer that can automate any critical workflow while maintaining rigorous security and complying with regulations.

What We Do

Apex (Pre-Trained AI Models)

Access a library of 100+ ready-to-use AI models ranging from deepfake detection to finance statement analysis.

Prism (Domain Specific Model)

An intelligence layer that unlocks the power of Generative AI to automate tasks, enhance collaboration, and drive productivity—seamlessly integrated into your workflow

Weave (Agent Orchestration Platform)

Simplify AI deployment where Weave acts as a central hub that routes AI requests to the appropriate model or data source, orchestrating multiple agents on one platform.

Autonomous Finance

Empower your finance function with a suite of AI-powered decision engines that autonomously manage tasks such as onboarding, underwriting, cash flow forecasting, etc.

Learn more about how Arya.ai can catalyze your AI transformation at https://arya.ai.

Authors

Kushagra Bhatnagar

HEAD OF RESEARCH

Kushagra Bhatnagar is responsible for the technical direction and foundational research for Arya.ai’s platforms. He concentrates on the architecture of multi-agent systems and the protocols required for secure, high-performance communication between AI models and enterprise data sources. His work ensures the company’s products are built on a stable and technically sound framework.

Deekshith Marla

FOUNDER

Deekshith Marla directs the company's vision for enterprise AI. His work is centered on removing the persistent barriers between advanced AI capabilities and their practical implementation within complex corporate environments. He is responsible for the strategy that guides the development of scalable, secure, and integrated AI solutions for the firm’s clients.