



**AXOMA**  
Centralized. Secure. Intuitive

Playbook

# Introducing Axoma

Axoma is an enterprise-grade Generative AI Accelerator Platform that empowers organizations to rapidly explore, validate, and scale AI use cases with confidence, governance, and speed. It acts as a strategic enabler for enterprises looking to adopt GenAI without the traditional complexities of long development cycles, high experimentation costs, and fragmented technology choices.

Designed to integrate seamlessly into diverse IT ecosystems, Axoma provides the foundational capabilities required to transform enterprise knowledge, workflows, and services through AI-driven automation and intelligence. Its modular architecture, flexible deployment options, and strong governance standards make it suitable for both cloud-first and highly regulated industries needing on-premise deployments.

## Axoma Vision

The vision behind Axoma is to accelerate AI adoption through a structured, repeatable, and scalable approach. Enterprises often struggle with identifying high-value use cases, validating early ideas, and converting prototypes into production solutions. Axoma bridges this gap by offering:

- Pre-built use case accelerators
- A model-agnostic, cloud-agnostic foundation
- A governed experimentation environment
- Rapid deployment capabilities

Axoma helps enterprises move from “idea to impact” faster by enabling a Quick Try → Quick Fail → Quick Scale strategy. This ensures innovation occurs rapidly while minimizing risk and maximizing business value.

# Core Capabilities

Axoma unifies the essential components needed for enterprise AI adoption:

## **i) Application Templates**

Ready-to-deploy templates for cross-industry use cases including customer support, operations automation, HR workflows, finance copilots, network management, and more. These apps significantly reduce time and cost for proof-of-concept development.

## **ii) Enterprise Knowledge Fabric**

A sophisticated ingestion and RAG (Retrieval-Augmented Generation) framework that allows secure integration of enterprise data from documents (pdf, docx, doc, text).

## **iii) Agent Builder**

Visual building blocks, agent builder allow teams to prototype and deploy AI agents quickly without needing extensive coding skills and just with help of tools.

## **iv) Workflow Builder**

Using Low Code, No Code build a workflow for automating the user processes. The workflows can be scheduled and can be integrated with other agents to perform some specific tasks.

## **v) Security, Governance & Control**

Axoma ensures enterprise-grade trust with features like role-based access, audit trails, model usage monitoring, content filtering, and custom guardrails.

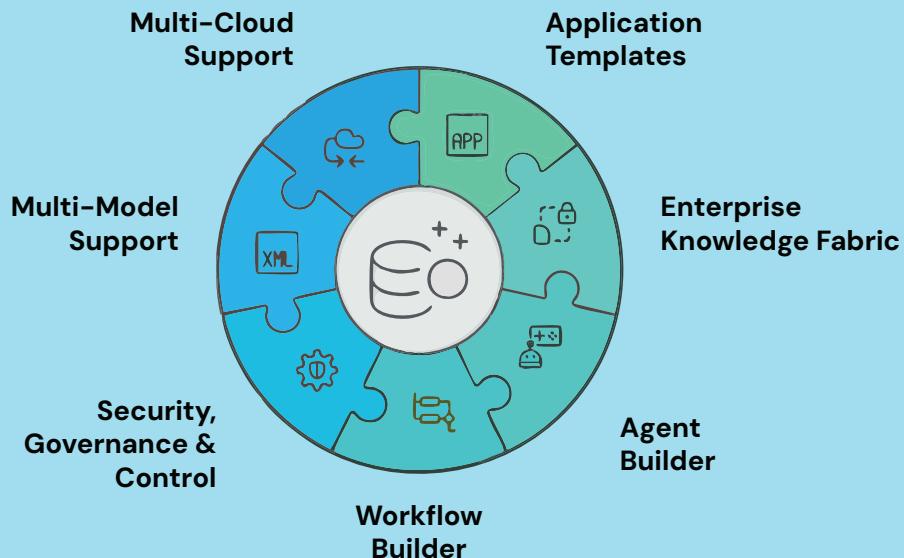
## **vi) Multi-Model support**

Axoma is fully agnostic, supporting OpenAI, Azure OpenAI, Gemini, Llama, and custom LLMs.

## **vii) Multi-Cloud Support**

Axoma can be deployed on AWS, Azure, GCP, or in On-Prem Infra.

## Axoma's Comprehensive AI Solutions



## Axoma-Powered GenAI Use Cases Across Business Functions

### Telecom Use cases

#### Low

1. Customer Support Query Assistant
2. Knowledge Base Search Assistant
3. Internal Policy Summarizer
4. Plan Recommendation Bot

#### Medium

1. Contact Center Agent Copilot
2. Complaint Category Classification
3. Internal Reporting Automation
4. Field Operations Knowledge Assistant

#### High

1. Predictive Ticket Resolution
2. Fraud Detection & Explanation Engine
3. Network Incident Triage Assistant
4. Sales Advisor Copilot

#### Complex

1. 360° Customer Intelligence Engine
2. Network Planning & Optimization Assistant
3. Autonomous Network Operations Copilot

## Banking Use cases

### Low

1. Knowledge Base Search Assistant
2. Product Comparison Assistant
3. Policy & Compliance Summarizer
4. Internal Knowledge Search

### Medium

1. Loan Application Summarizer
2. Dispute Resolution Assistant
3. Investment Product Assistant

### High

1. AML / Fraud Case Explanation Engine
2. Intelligent Regulatory Reporting Assistant
3. Credit Underwriting Copilot

### Complex

1. 360° Customer Intelligence Engine
2. Autonomous Compliance & Regulatory Engine
3. Fraud Prevention & Autonomous Monitoring Platform
4. Enterprise Knowledge Graph + GenAI Reasoning

## Retail Use cases

### Low

1. Product Description Generator
2. Customer FAQ/Support Assistant
3. Store Policy Summarizer
4. Review Summarization Assistant
5. Knowledge Search Assistant

### Medium

1. Customer Service Agent Copilot
2. Inventory Assistant
3. Product Tagging & Classification
4. Loyalty & CRM Campaign Assistant

### High

1. Personalized Product Recommendations
2. Demand Forecasting Assistant
3. End-to-End Product Planning

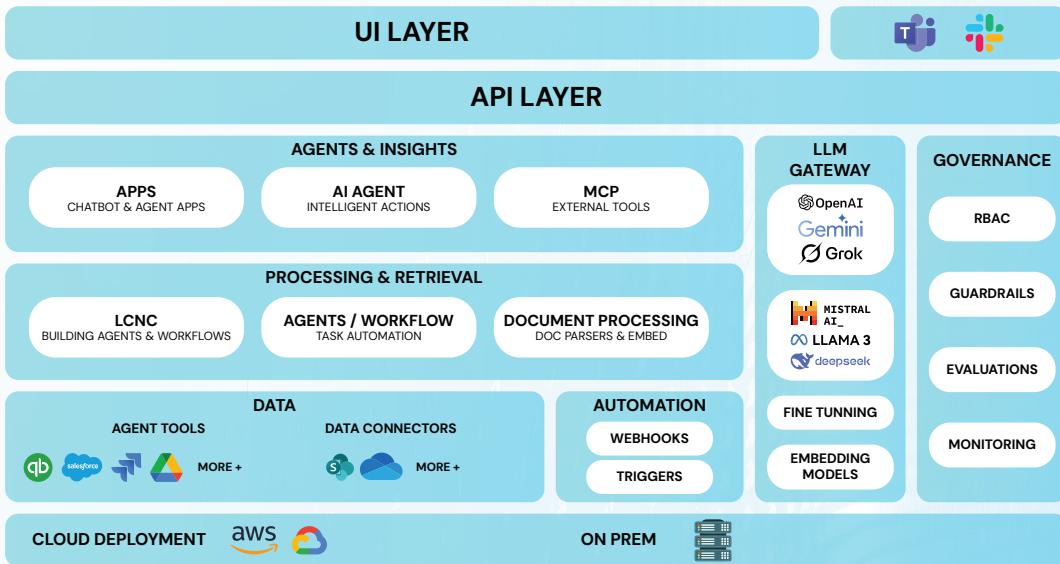
### Complex

1. 360° Customer Intelligence Engine
2. Dynamic Pricing & Offer Optimization Engine
3. End-to-End Assortment Optimization Platform

## Platform Architecture, Value & Positioning

- **How Axoma Work**

Axoma's platform is structured as a layered architecture to support rapid experimentation and enterprise-grade scaling:



- **UI Layer (User Interaction)**

The UI Layer provides a clean, intuitive, and enterprise-ready interface for interacting with Axoma's capabilities. It enables users to engage with agents, trigger workflows, visualize insights, manage configurations, and monitor performance all through a secure, responsive, and role-based experience. Designed for simplicity and scalability, the UI Layer ensures that even complex AI interactions feel seamless for business, IT, and governance users.

- **API Layer**

The API Layer acts as the connectivity backbone of Axoma, enabling smooth integration with enterprise systems, databases, operational tools, and external AI services. It exposes secure, modular APIs that allow organizations to trigger agents, fetch insights, push data, and embed Axoma capabilities into existing applications. With standardized protocols and high scalability, the API Layer ensures Axoma fits naturally into any enterprise architecture whether cloud, on-premise, or hybrid.

- **Agents and Insights Layer**

Axoma's Agents are autonomous, domain-specific AI workers that execute business tasks end-to-end across enterprise applications. They handle

everything from data retrieval and workflow orchestration to decision-making and action execution reducing manual effort, improving accuracy, and enabling continuous automation across IT, business, and governance functions.

- **Processing and Retrieval Layer (Agents, Workflow, Vector Search & Knowledge)**

The Processing & Retrieval Layer is Axoma's core engine that unifies low-code/no-code orchestration, workflow automation, document processing, and enterprise-grade retrieval. It allows teams to rapidly build and iterate AI use cases by designing workflows with minimal engineering effort while automating complex business processes end-to-end. The document processing component converts unstructured content such as PDFs, text and docx, structured data using OCR and LLM-based extraction. Combined with vector search and RAG pipelines, this layer ensures every AI interaction is grounded in accurate, real-time, domain-specific knowledge.

- **Data Layer (Integration Hub)**

The Data Layer in Axoma provides a unified foundation for connecting, managing, and leveraging enterprise data across structured systems, unstructured repositories, and external APIs. Its library of pre-built connectors enables seamless integration with databases, CRMs, ERPs, cloud storage, etc ensuring fast onboarding and minimal engineering effort. Complementing this, Axoma's tools enable agents to interact directly with these applications by executing actions such as querying data, updating records, triggering workflows, or retrieving domain-specific information. Together, the tools and connectors ensure agents can operate reliably across the enterprise ecosystem while maintaining clean, governed, and real-time data access.

- **Automation Layer**

The Automation Layer in Axoma enables real-time, event-driven intelligence across enterprise workflows. Triggers allow Axoma to respond instantly to system events such as data updates, ticket creation, form submission, or user actions initiating agents, workflows, or business logic without manual intervention. Webhooks extend this capability by allowing external applications to push events directly into Axoma, enabling seamless integration with CRMs, ERPs, operational tools, and custom applications. Together, triggers and webhooks ensure Axoma remains tightly synchronized with enterprise processes, enabling continuous automation and faster decision-making.

- **LLM Gateway**

The LLM Gateway in Axoma serves as the unified control layer that manages all interactions with open-source and proprietary language models, giving enterprises full flexibility in how they deploy and scale AI. It supports seamless integration with leading commercial models, as well as on-premise or cloud-hosted open-source models, ensuring organizations can choose the right model for each use case. The gateway also manages embeddings and fine-tuning pipelines, enabling teams to customize models with domain knowledge while maintaining governance and version control. With built-in usage monitoring, rate limiting, and cost budgeting, the LLM Gateway ensures high performance, predictable spend, and secure AI operations across the entire enterprise.

- **Governance Layer**

The Governance Layer in Axoma ensures every AI interaction is secure, compliant, and aligned with enterprise standards through a unified framework that spans access, safety, quality, and oversight. Role-Based Access Control (RBAC) enforces granular permissions across users, agents, workflows, and data, while guardrails apply policy-based constraints to prevent undesired model outputs or actions.

Axoma's evaluation and monitoring capabilities continuously measure model accuracy, performance, drift, and behaviour, ensuring reliability throughout the lifecycle. Comprehensive auditing tracks every decision, interaction, and workflow execution, enabling full transparency and regulatory readiness across all deployments.

## Developer Personas & Role Responsibilities

This section provides an overview of the user personas that interact with Axoma within an enterprise environment, outlining how different technical and operational roles engage with the platform. It explains the responsibilities and access levels associated with each persona, how they contribute to configuring, managing, and consuming AI capabilities, and how Axoma supports collaboration across teams. By defining these personas clearly, the playbook helps organizations understand how Axoma fits into their existing workflows, who will use the platform at various stages, and how responsibilities are distributed to ensure secure, scalable, and efficient adoption.

Persona	Role Details	What They Do
<b>Super Admin</b>	AI COE Lead Architects, Platform Engineers, DevOps, Security Teams	Platform governance, security, infra, global settings, enterprise-wide integrations
<b>Admin</b>	Full-stack devs, ML engineers, integration devs, data engineers, Product Owner, Analysts, Functional Managers	Build & configure use cases, workflows, agents, connectors, tenant settings
<b>User</b>	Business User, End Customer	Run workflows, interact with agents, validate results, consume insights

## Scalability and Deployment

Axoma is engineered as a fully scalable platform that can adapt seamlessly

to enterprise workloads, from initial pilots to large-scale production deployments. Its architecture supports horizontal scaling through autoscaling policies that dynamically adjust compute resources based on real-time demand, ensuring optimal performance without over-provisioning. Resource configurations—including CPU and memory utilization targets can be fine-tuned through load testing, enabling precise control over system efficiency at scale. Axoma also supports robust operational reliability through readiness and liveness health checks, ensuring smooth rollouts, self-healing behaviour, and consistent uptime across all stateless services. This scalable and resilient design allows organizations to confidently expand their AI use cases while maintaining stability, performance, and operational excellence.

## Why Enterprises Choose Axoma

### **Accelerated Time-to-Value**

Axoma reduces the typical AI lifecycle from 6–12 months to just 12–16 weeks, thanks to pre-built accelerators, reusable components, and guided workflows.

### **Reduced Experimentation Cost**

Its controlled experimentation environment ensures rapid validation at minimal cost, reducing wasted development cycles.

### **Designed for Enterprise Scale**

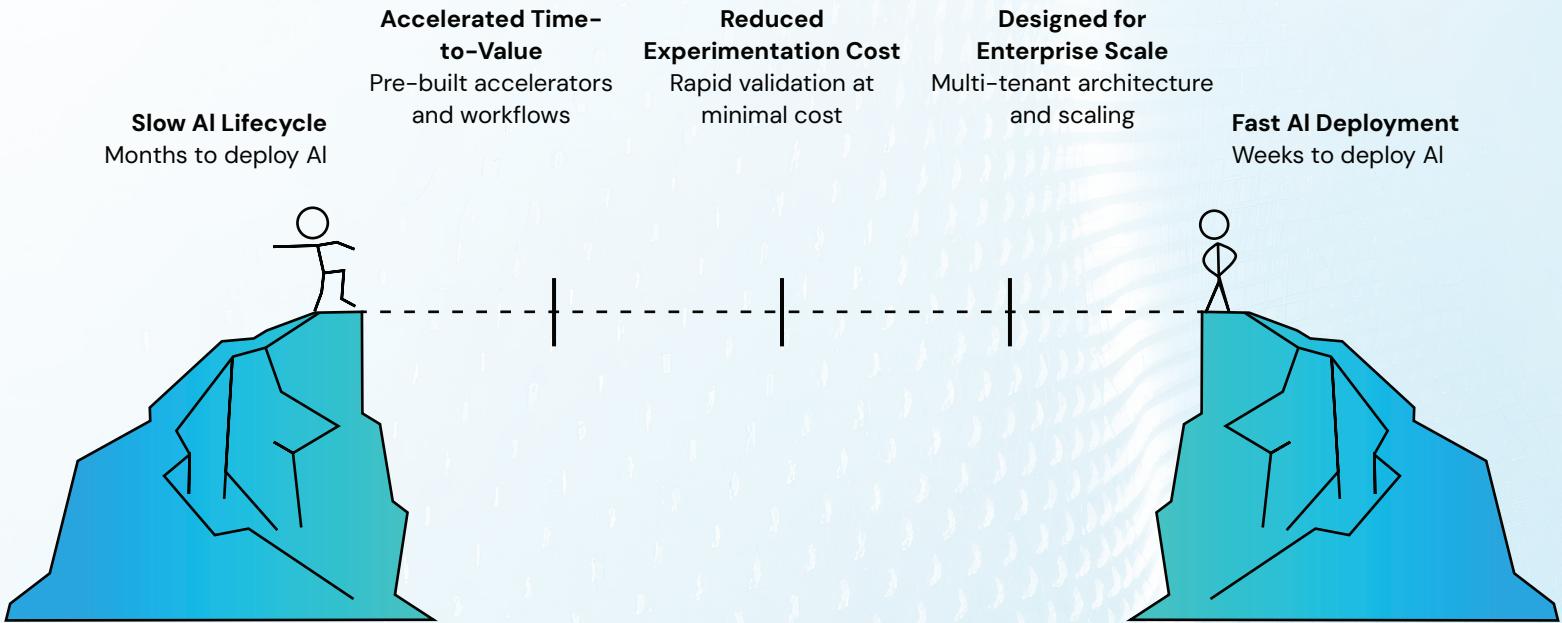
With multi-tenant architecture, horizontal scaling, and workload distribution, Axoma grows effortlessly as enterprise adoption expands.

### **Safe and Governed AI Adoption**

Security remains central from encrypted data pipelines to role-based controls and model governance.

## Flexible Deployment Options

Axoma supports cloud, hybrid, and fully on-prem deployments to align with enterprise security policies, data residency laws, and compliance needs.



## Strategic Positioning: Axoma as an AI Accelerator

Axoma is not just a platform it is a strategic accelerator for enterprise AI transformation. Its combination of domain-specific knowledge, strong governance, rapid prototyping, and scalable deployment framework makes it ideal for:

- Enterprises starting their GenAI journey
- Organizations seeking a governed sandbox environment
- Companies wanting to validate ideas quickly
- Large enterprises scaling AI across business units
- Regulated sectors requiring secure on-prem AI

Axoma enables organizations to innovate boldly, experiment safely, and scale confidently, making AI adoption predictable, cost-effective, and strategically impactful.

# Axoma Pricing Options

Axoma offers a flexible, scalable pricing model designed to support organizations at every stage of their AI adoption journey. The Starter plan is ideal for teams beginning their AI experimentation, providing essential platform capabilities to prototype use cases and validate value quickly. The Growth plan is built for organizations looking to scale AI across multiple functions, offering advanced automation, higher usage limits, premium connectors, and expanded governance features. For large enterprises with complex requirements, the Enterprise plan delivers unlimited scalability, dedicated infrastructure options, enhanced security and compliance controls, and customized support to meet industry-specific needs. Together, these plans ensure that organizations can adopt Axoma at their own pace while maintaining predictable cost and enterprise-grade performance.

	<b>Starter</b>	<b>Growth</b>	<b>Enterprise</b>
<b>Users</b>	10 Users	25 Users	100 Users
<b># of Agents</b>	5 Agents	15 Agents	40 Agents
<b># of Docs</b>	50 Docs	100 Docs	200 Docs
<b># of Tools</b>	5 Tools	15 Tools	50 Tools
<b># of Apps</b>	Unlimited	Unlimited	Unlimited
<b># of Tenants</b>	Unlimited	Unlimited	Unlimited
<b>Security Group</b>	Unlimited	Unlimited	Unlimited
<b>Usage</b>	Unlimited	Unlimited	Unlimited
<b>Support</b>	Standard Support	Standard Support	Priority Plan Available
<b>Pricing</b>	\$999 / Month	\$1999 / Month	Contact Us