

saguna consulting»



# The AI Aperture

Opening the portal for **people** and businesses to explore what AI can unlock, today and *beyond*.

**Our AI solutions turn that view into action.**

# Table of contents

|    |                  |
|----|------------------|
| 03 | Introduction     |
| 04 | Our AI solutions |
| 05 | AI Advisory      |
| 08 | AI Architecture  |
| 11 | AI Engineering   |
| 14 | AI Foundry       |
| 17 | AI Factory       |
| 20 | AI Governance    |
| 23 | Conclusion       |

## Are the possibilities bigger than you think?

What if you could uncover  
AI opportunities you  
haven't considered yet?

What if you could turn AI  
from an idea into  
something your **people**  
actually use?

What if you could turn AI  
possibilities into  
capabilities that keep  
evolving with your  
business?

What if you could see  
beyond what AI can do  
today?

\_\_\_\_\_ This is what the AI aperture reveals and what our AI solutions make possible. \_\_\_\_\_

AI is no longer limited to a few use cases or isolated experiments. It is creating new ways to work, build, serve the **community**, make decisions, and imagine what a business can become.

At Saguna Consulting, we call this The AI Aperture- a way of looking beyond what AI does today and exploring what it could make possible tomorrow. What you see is only the beginning. The possibilities extend far beyond what AI can do today. Our AI solutions turn that view into action. From identifying where AI can make the biggest difference to building, integrating, and scaling it across the business, we help turn possibilities into something real.

**Possibilities go where you take them.**

# Our AI solutions

Six solutions covering the full AI lifecycle, from first assessment to enterprise-wide deployment.

AI  
Advisory

1

AI  
Architecture

2

AI  
Engineering

3

AI  
Foundry

4

AI  
Factory

5


AI  
Governance

6

# AI Advisory

Organizations can struggle to define a clear AI direction, prioritize high value opportunities, or translate AI initiatives into scalable business outcomes.

Our AI Advisory model helps you get there through readiness workshops, ROI-based use-case scoring, and clear roadmaps that turn priorities into a phased path to enterprise-wide execution.



## Four pillars of our AI advisory services

### AI Maturity & Opportunity Mapping:

Assess AI readiness using maturity scorecards, data-quality audits, and infrastructure diagnostics, then rank use cases by feasibility, ROI, and business impact.

### Use Case Prioritization:

Identify high-value GenAI and SLM-based use cases via capability mapping, ROI modeling, and value-stream analysis to prioritize investment and impact.

### Pilot to Production:

Move validated use cases from POC to production using MVPs, RAG pipelines, agentic orchestration, and reference architectures built on real enterprise data.

### Data & Platform Foundations:

Build scalable foundations on cloud-native architectures, vector databases, data lakehouses, and LLM orchestration to support future workloads.



# AI Architecture

Complex technology environments can become fragmented when cloud, infrastructure, data, and applications evolve without a unified architecture.

Our AI Architecture model gets you there through blueprint design, cloud-native patterns, and phased modernization plans across compute, data, applications, and operations for resilient, long-term AI adoption.



# Four pillars of our AI architecture services

01/

## **Architecture strategy:**

Define future-ready architectures using microservices, API gateways, and event-driven design for scalable, resilient security.

02/

## **Cloud & infrastructure architecture:**

Design hybrid, edge, and on-premises architectures on Kubernetes and a data mesh model for scalable, sovereign-compliant environments.

03/

## **Infrastructure & compute:**


Architect GPU-accelerated infrastructure using NVIDIA clusters, high-density centers, and low-latency networking for Enterprise AI.

04/

## **Secure & optimized architecture:**

Design governed architectures using Terraform, Prometheus observability, FinOps tooling, and MLOps/LLMOps workflows.

Figure 2- How companies best describe their current stage of cloud journey  
Source- Accenture | Cloud outcomes: the race to cloud



## What you'll achieve

### » Hybrid cloud architecture design

Build a resilient hybrid footprint using workloads that run seamlessly across cloud, edge, and on-device inference, with sovereign AI compliance built in automatically.

### » Governed enterprise data foundation

Get started with a governed data backbone through knowledge graphs and vector database layer turning scattered enterprise data into trusted, queryable insight without compromising compliance.

### » Modern cloud-native applications

Leaner, faster applications via legacy monoliths becoming container-native microservices that deploy in minutes and flex automatically with demand.

# AI Engineering

AI software development can become complex as repetitive engineering tasks, testing demands, and technical challenges increase across the lifecycle.

Our AI Engineering model gets you there via AI copilots, automated pipelines, and agentic workflows across development, quality, and operations, accelerating delivery and building resilient software.



# Four pillars of our AI engineering services

01/

**AI native software engineering:**

Embed GenAI copilots across the lifecycle, from requirements and architecture to CI/CD pipelines and automated testing to accelerate delivery.

02/

**Agentic engineering:** Deploy multi-agent orchestration systems, such as LangGraph, AutoGen, and CrewAI, that plan, reason, and coordinate tasks across CI/CD toolchains.

03/

**AI assisted development:** Apply Copilot-style LLMs and synthetic data generation to code generation, automated refactoring, and testing, cutting repetitive work and boosting velocity.

04/

**AI product & platform engineering:** Build intelligent software platforms using agentic AI, Kubernetes-based microservices, and AI-powered orchestration engines for secure, production-grade solutions.

Organizations often face increasing engineering complexity as development, testing, and operations become more demanding across the software lifecycle.

We embed AI and agentic workflows across the engineering lifecycle automating repetitive tasks, accelerating software delivery, enhancing quality and engineering intelligence, and enabling scalable, resilient applications.

## What you'll achieve

### » **AI powered development acceleration:**

Ship faster, safer using fine-tuned code-gen agents and copilots cut manual engineering effort across the lifecycle, compressing release cycles without sacrificing quality.


### **Intelligent quality & system resilience:**

» Fewer defects through AI observability into every test run catches regressions before release, predicting failure points manual QA typically misses.

### **Intelligent software products & platforms:**

» Smarter platforms by default via multimodal, agent-orchestrated products ship with decision intelligence built in, giving you a durable market edge.

AI-enabled engineering outcomes:  
Productivity and efficiency gains



Faster information retrieval was recorded.

Figure 3

Source: Intellias | AI enabled engineering services

# AI Foundry

Organizations often struggle to connect intelligent agents across business processes, systems, and workflows, limiting their ability to scale agentic capabilities effectively.

Our AI Foundry model gets you there through a shared agent framework, API integrations, and orchestration layers that build, connect, and govern agentic workflows across the enterprise.



# Four pillars of our AI foundry services

01/

**Agentic AI application development:** Build and deploy enterprise-grade agentic apps using LLM orchestration frameworks, combining conversational UX, real-time data, and autonomous execution.

02/

**Custom & pre-built AI agents:** Provide ready-to-deploy LLM & SLM agents, domain-specific agents built with RAG pipelines for enterprise workflows and needs.

03/

**Multi agent orchestration:** Develop multi-agent systems using LangGraph and AutoGen, where specialized agents collaborate, delegate tasks, and coordinate complex workflows.

04/

**Enterprise AI integration:** Integrate AI agents with enterprise systems via MCP (Model Context Protocol), event streaming, and iPaaS connectors for seamless, low-code adoption.

**AI Foundry is built for the gap between experimentation and scale.** Most organizations don't need more AI experiments. They need a way to turn those experiments into systems that can actually work together.

AI Foundry brings agents, data, enterprise applications, and workflows into one connected environment, creating the foundation for intelligent operations that can move from prototype to production and from individual use cases to enterprise scale. It is where ideas become working systems.

# 15%

of organizations  
have scaled  
orchestrated, cross-  
functional multi-  
agent adoption

Source: Deloitte — AI Agents are Only the  
Beginning: AI Readiness Gap

## What you'll achieve

### » Enterprise agentic AI applications:

Agents that actually collaborate using domain-specific agents connected via an agent-to-agent (A2A) protocol coordinate complex workflows autonomously, without constant human handoffs.

### » Intelligent workflow automation & integration:

Workflows that actually connect through agents link directly to your existing systems and data sources for true end-to-end automation.

### » Reliable AI agent operations & governance:

Agents you can actually trust using real-time monitoring and guardrails keep every agent reliable and responsible at scale, with human oversight always available.

# AI Factory

AI Factory environments can become fragmented when legacy data, platforms, applications, and infrastructure operate in silos, limiting scalability and increasing operational complexity.

Our AI Factory model gets you there by modernizing data pipelines, migrating to cloud-native platforms, and automating operations to enable scalable, reliable enterprise AI.



# Four pillars of our AI factory services

01/

**Data modernization & readiness:**

Modernize fragmented data using profiling, ETL/ELT pipelines, data-quality rules, and governance catalogs to build trusted, AI-ready foundations.

02/

**Cloud data modernization:** Transform enterprise data using lift-and-shift migration, cloud-native lakehouses, and containerized platforms to improve scalability, cost efficiency.

03/

**Data engineering & integration:** Build modern pipelines using Airflow, Spark, and dbt, with lakes and hubs to accelerate processing and time-to-insight.

04/

**AI ready data platforms:** Enable AI and analytics using Snowflake, Databricks lakehouses, real-time streaming, and model quantization to accelerate AI/ML adoption and decisions.

## Market overview

|                           |                    |
|---------------------------|--------------------|
| <b>Study period</b>       | <b>2020 - 2031</b> |
| Market size (2026)        | USD 381.5 Billion  |
| Market size (2031)        | USD 782.40 Billion |
| Growth rate (2026 - 2031) | 15.45%             |
| Fastest growing market    | Asia- Pacific      |
| Largest market            | North america      |
| Market concentration      | Medium             |




Figure 4- AI factory infrastructure market | Market size in USD billion

Source: Modor intelligence

## What you'll achieve

### » Trusted AI ready data foundation

Modernize fragmented data using profiling, ETL pipelines, and governance catalogs. Create Snowflake and Databricks lakehouses that enable trusted, scalable AI, analytics organization-wide.

### » Modern data pipelines

Build data pipelines, lakes and hubs using Airflow, dbt automation and GraphRAG-powered analytics turn raw data into decisions instantly, across every business unit.


### » Legacy modernization

Modernize legacy applications and mainframes using LLM-assisted code translation and automated monitoring. Optimize environments using Prometheus monitoring and automated pipelines for improved performance.

# AI Governance

AI governance becomes complex as organizations manage multiple models, risks, regulations, and **stakeholders** across the AI lifecycle.

Our AI Governance model gets you there through risk-scoring frameworks, automated LLM evaluation, and centralized AI gateways that continuously keep systems secure, compliant, transparent, and responsibly managed at enterprise scale.



## Four pillars of our AI governance services

### AI Principles & Accountability:

Define AI governance principles and operating models using model registries, policy-as-code, and centralized AI inventories for accountability.

### AI Risk Management:

Map AI systems by risk tier using automated risk-scoring models, bias-detection algorithms, and continuous policy engines to evaluate fairness, security, and compliance.

### Responsible AI Testing:

Integrate LLM eval frameworks like RAGAS and DeepEval, plus guardrails-as-code, into the CI/CD pipeline to test bias, fairness, and trustworthiness.

### Multi-Model Governance:

Scale and govern multiple LLMs through API gateways (Kong, Apigee) and model routing, with centralized RBAC controls for security.

The AI governance market is projected to grow at a 36% CAGR through 2033, driven by tightening regulations and rising demand for transparency, bias mitigation, and risk management.

This growth signals that governance is no longer optional – it's foundational to scaling AI responsibly.

Estimated AI governance  
market expansion

# 36%

CAGR, 2026–2033


Source: Grand View Research

## What you'll achieve

- » **AI governance & risk framework:** Achieve governance from day zero through a live model registry which establishes accountability automatically, replacing manual risk assessments and static spreadsheets.
- » **Responsible AI testing & compliance:** Integrate automated red-teaming and bias-eval pipelines test every release continuously, flagging drift before it becomes real risk.
- » **Scaled AI model governance capability:** Govern multiple LLMs through Kong and Apigee API gateways with model routing. Centralized routing governs the fine-tuning vs. RAG tradeoff per use case, letting teams manage responsible AI independently.

# Making AI work takes more than AI

AI is only one part of the equation. Making it work takes the right capabilities around it.



**We bring the whole picture together.**

Strategy, **people**, technology, systems, and AI.  
What businesses need to move from possibility to implementation.

# From AI assessment to enterprise transformation

## 1.

### **Assess & Understand**

Evaluate your AI readiness- infrastructure, data maturity, organizational alignment. Map opportunities and gaps to build a realistic, prioritized roadmap.

## 2.

### **Design & Build**

Engineer custom AI solutions with governance built in. Move beyond pilots, design agentic workflows, integrate systems, and establish responsible AI practices from day one.

## 3.

### **Scale & Evolve**

Deploy AI across enterprise operations with observability and compliance. Continuously monitor, adapt, and expand, turning AI into competitive advantage at scale.

**Most organizations get stuck between assessment and scale. The difference between pilots and production is **partnership**.**

## About Saguna Consulting

Saguna Consulting is an AI-first consulting and technology company that builds products, delivers technology, and **partners** with organizations to transform ideas into durable, scalable outcomes rooted in real-world execution. Human-first in our approach, outcome-led in our delivery, and technology-enabled at every step, we work closely with **teams** to understand their problems, sharpen their vision, and move forward with confidence. From shaping early concepts to executing production-ready platforms, we help clients build solutions that are useful, usable, and built to last, always with shared ownership and collective momentum.

We collaborate with product-led teams, founders, and enterprises, engaging deeply as advisors, builders, and long-term **collaborators**. Sometimes this takes the form of traditional consulting; other times, it means aligning through equity-based collaboration, sharing not just work but conviction and outcome. Our multidisciplinary community bridges strategy, design, engineering, and delivery to ensure technology accelerates impact. *Forward together*, and AI-first by design, we help organizations deliver meaningful products and solutions that endure."

CTA:

 [sagunaconsulting.com](https://sagunaconsulting.com)

 [connect@sagunaconsulting.com](mailto:connect@sagunaconsulting.com)

 [saguna consulting](https://www.linkedin.com/company/saguna-consulting)

Disclaimer- This document is provided for general informational purposes only and reflects information available at the time of preparation. It does not constitute legal, financial, or professional advice. Saguna Consulting disclaims liability for decisions made based on this content.

Any third-party names or marks referenced remain the property of their respective owners and do not imply endorsement. Some visual elements may be illustrative or AI-generated.

@ 2026 All Rights Reserved Saguna Consulting Services LLC  
Saguna Consulting Services India Pvt. Ltd (wholly owned subsidiary)