Calibo

A business-led AI execution framework

for trusted, production-grade use cases

Enterprises today operate in **highly complex, siloed environments**. As scale and operational complexity grow, **traditional big-bang approaches to business innovation break down**. Fragmented data, disconnected systems, and manual workflows block end-to-end execution, slow critical decisions, and undermine confidence in outcomes. Despite continued investment in digital and AI initiatives, organizations face a **persistent loss of productivity and operational efficiency**.

To address this, businesses need an **experimentation environment** that allows them to break complex problems into **bite-sized units**, engage directly with data and AI, and then **converge these into workflows that drive adoption and productivity**. Rather than tackling enterprise complexity monolithically, successful organizations start from execution reality—how work actually happens across siloes, documents, and handoffs.

Workflows are decomposed into **decision-centric use cases** with clear ownership and measurable outcomes. Execution begins with **end-to-end digitization**, making data AI-ready and reusable. AI and GenAI **augment human decision-making**, improving speed and consistency while retaining accountability. **GenAI-powered smart interfaces** enable collaboration and adoption. As use cases mature, they reconnect into a governed execution model, enabling **trusted, production-grade AI delivery in weeks, not months**.

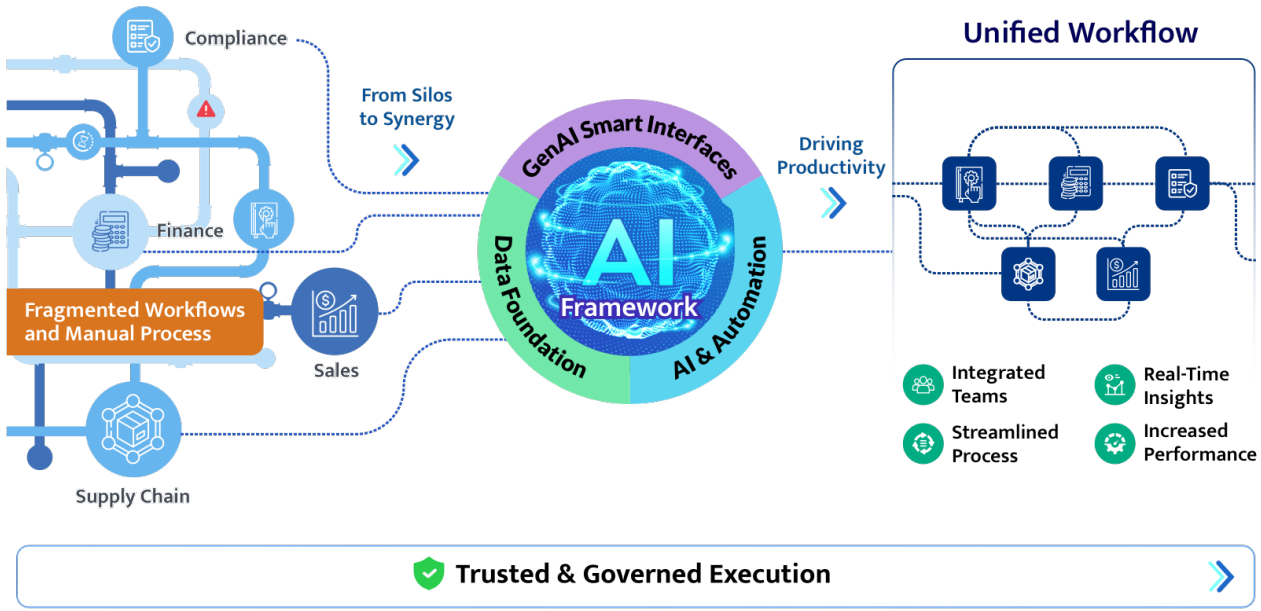

- Multi-model deployment architecture validated across Calibo-managed and customer-hosted environments
- Proven scalability through auto-scaling, self-healing services, and load-tested pipelines
- One-click / self-service deployment capabilities designed for restricted enterprise access models
- Built-in observability covering performance, cost, and operational health
- Web & API penetration testing
- Infrastructure security testing
- Architecture, documentation and stakeholder reviews

Each use case demonstrates how complex enterprise workflows are transformed by:

- Decomposing siloes into bite-sized, decision-centric use cases
- Digitizing end-to-end workflows, including unstructured data
- Augmenting human decision-making rather than replacing accountability
- Leveraging GenAI-powered smart interfaces to reduce cognitive load
- Reconnecting optimized siloes into a collaborative execution environment
- Embedding governance, explainability, and auditability by design

The framework provides shared data foundations, reusable AI services, and intelligent interfaces that allow these use cases to move from experimentation to trusted, production-grade deployment.

All use cases are built on a common set of principles enabled by the framework:

Business-led execution
AI aligned to decisions and outcomes

Decompose complexity
Small, high-impact use cases over monoliths

Digitization before intelligence
Strong data foundations first

Human-in-the-loop by design
Especially for regulated or critical flows

Smart interfaces drive adoption
GenAI as the interaction layer

Governance enables scale
Trust unlocks production deployment

Reuse over reinvention
Shared pipelines, models, and interfaces

Applied consistently, these principles enable organizations to transition from:

Fragmented siloes	Connected, collaborative execution
Manual effort	High productivity workflows
Isolated analytics	Increased adoption and trust
AI pilots	Governed, scalable production systems

This preface sets the context for the use cases that follow, each illustrating how the framework enables end-to-end AI execution across industries and domains.

USE CASE 1

AI-Augmented Customer Due Diligence & Risk Profiling

(Wealth Management / Regulated Financial Services)

Use Case Overview

Wealth onboarding operates across parallel and highly regulated siloes—relationship management, compliance, risk, and operations. Customer Due Diligence (CDD) and Risk Profiling run as separate processes but jointly determine onboarding approval, product access, and exposure limits. Manual document handling, rule-based scoring, and sequential approvals make onboarding slow, costly, and difficult to scale for complex clients.

This use case demonstrates how fragmented DD and EDD workflows are decomposed into bite-sized Data-AI use cases, digitized end-to-end, and reconnected through GenAI-powered smart interfaces into a collaborative, explainable, and governed onboarding execution model.

Business Problem

- Manual, document-heavy DD and EDD processes
- Rule-based risk scoring with high false positives
- Delayed escalation and senior approvals
- High operational effort for UHNW and complex clients
- Strong regulatory expectations for explainability and auditability

Success Metrics

- **Time to onboard**
- **Drop-off rate**
- **Cost per onboarding**
- **Compliance exceptions**

Objective

- Accelerate low-risk onboarding
- Improve accuracy of customer risk assessment
- Reduce manual effort in EDD
- Maintain strong regulatory defensibility

Key Stakeholders & Actors

Relationship Managers | Internal Audit | Compliance & AML Teams | Risk Management | Data Science & ML Engineering | Operations

Business Understanding

Breaking Siloes into Bite-Sized Use Cases

- Identity verification and KYC
- Sanctions and adverse media screening
- Customer risk scoring
- Enhanced Source of Wealth analysis
- Ongoing behavioral and risk drift monitoring

Data Understanding

- Customer identity and financial documents
- Sanctions and watchlists
- Media, legal, and regulatory data
- Transaction and behavioral data

Data Preparation

- OCR and document intelligence
- Entity resolution and relationship mapping
- Feature engineering for risk indicators
- Privacy masking and lineage tracking

AI / ML Modelling Approach

- ML-based customer risk scoring
- NLP-based adverse media analysis
- Graph analytics for ownership discovery
- Behavioral anomaly detection

GenAI-powered interfaces enable compliance teams to review risk narratives conversationally and trace decision rationale.

Evaluation & Governance

- Explainable AI for all risk decisions
- Human-in-the-loop approvals for EDD
- Full audit trails and regulatory reporting

Deployment & Operating Model

- Parallel DD and risk profiling workflows
- Real-time EDD trigger detection
- Scalable deployment across client segments

Monitoring & Continuous Improvement

- Continuous risk drift detection
- Periodic DD / EDD refresh cycles
- Model retraining aligned to regulatory change

Business Outcomes

- Same-day to 1–2 day DD completion
- 5–10 day EDD vs 30–60+ days traditionally
- Increased continuous adoption by 45%
- Strong regulatory defensibility

USE CASE 2

AI-Driven Demand Management

(Retail/ Consumer Industries)

Use Case Overview

Retail enterprises operate within deeply siloed planning environments where merchandising, supply chain, store operations, and finance optimize independently. As assortment complexity and omnichannel demand increase, manual forecasting processes fail to scale and limit end-to-end visibility.

This use case illustrates how planning siloes are decomposed into granular Data-AI use cases, digitized using standardized data pipelines and GenAI-enabled smart interfaces, and reconnected into a collaborative planning and execution ecosystem.

Business Problem

- Fragmented planning across stores, products, and channels
- Manual heuristics unable to scale
- Poor cross-functional visibility
- Low trust in centralized forecasts

Success Metrics

- **Stock-out rate**
- **Fill rate**
- **Forecast accuracy (WAPE, bias)**

Objective

- Improve forecast accuracy and fulfillment
- Reduce stock-outs
- Shift planners from manual creation to decision orchestration

Key Stakeholders & Actors

Retail Planners & Business Users | **ML Engineers** | **Data Scientists** | **MLOps / Platform Teams**

Business Understanding

Silo Decomposition into Granular Use Cases

- Product–store level forecasting
- Promotion impact analysis
- Allocation and replenishment optimization

Data Understanding

- Historical sales and inventory
- Product and store hierarchies
- Promotion and seasonality data

Data Preparation

- Standardized ingestion pipelines
- Feature engineering for demand signals
- Data contracts and lineage

AI / ML Modelling Approach

- Multiple model experimentation
- KPI-aligned model selection
- Explainability for planner trust

GenAI-enabled interfaces allow planners to interact conversationally with forecasts and trade-offs.

Evaluation & Governance

- Baseline vs AI performance comparison
- Bias and drift detection
- Governance-aligned approvals

Deployment & Operating Model

- Automated MLOps pipelines
- Batch and real-time inference

Monitoring & Continuous Improvement

- Drift monitoring
- KPI alerts
- Feature and use case expansion

Business Outcomes

- ~30% reduction in stock-outs
- 98% order fulfillment
- Significant planner productivity gains
- Trusted, governed planning intelligence at scale

USE CASE 3

Agentic AI-Enabled External Manufacturing & Operations Integration

(Manufacturing / Life Sciences / Industrial)

Use Case Overview

Manufacturing organizations operate across complex internal and external siloes, including plants, quality, regulatory, supply chain, and a broad partner ecosystem. External partners operate with heterogeneous IT stacks and document-heavy processes, creating fragmented visibility and manual execution.

This use case demonstrates how manufacturing and partner siloes are decomposed into bite-sized operational Data-AI use cases, digitized using reusable data foundations and GenAI-powered agentic workflows, and reconnected into a collaborative, governed manufacturing execution environment.

Business Problem

- Partner data locked in PDFs and disparate systems
- Manual transcription and reconciliation
- Delayed operational visibility
- Low trust in execution metrics

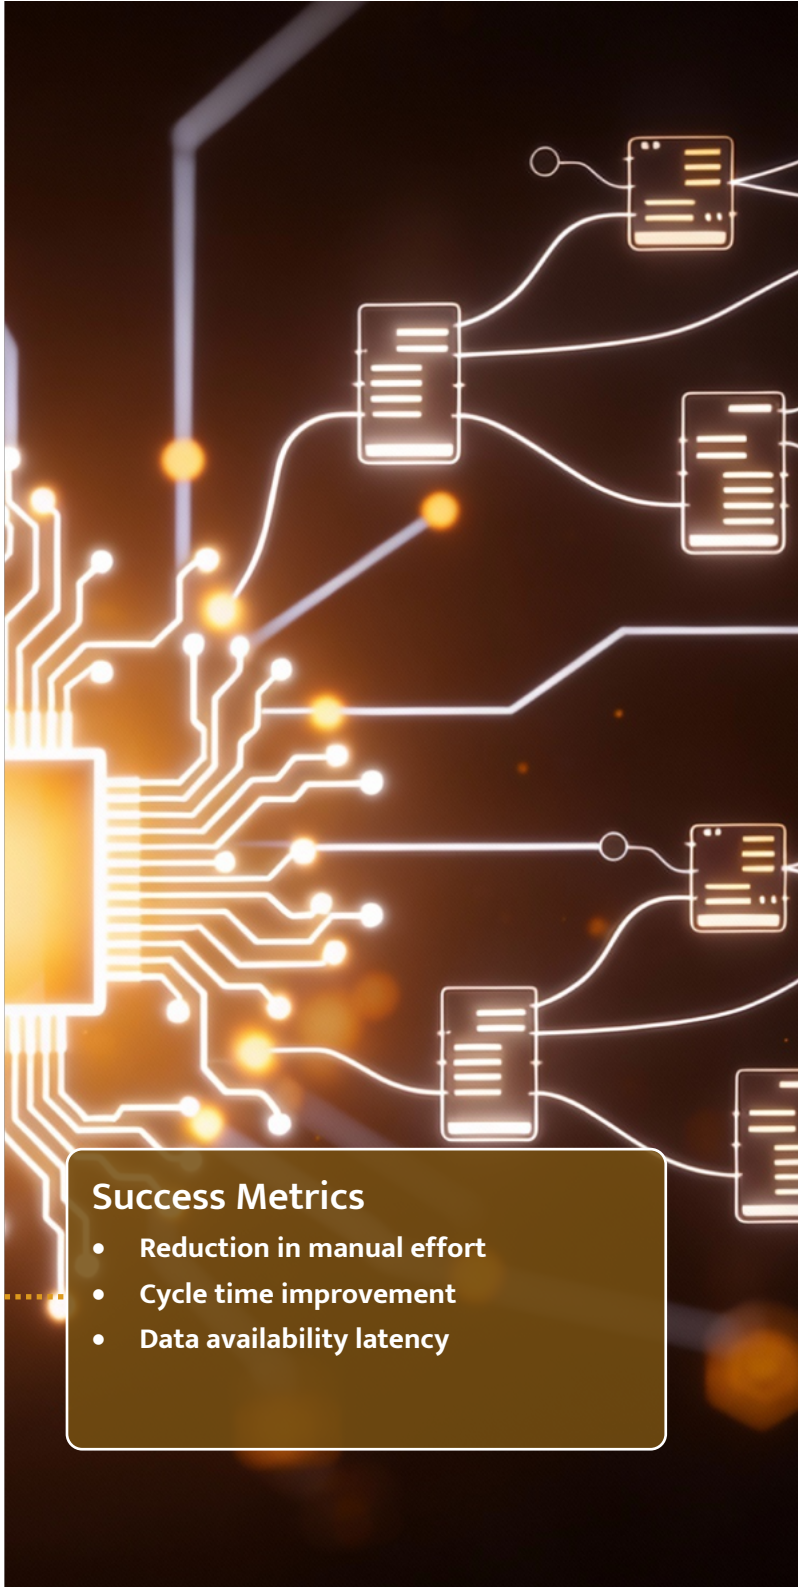

Success Metrics

- **Reduction in manual effort**
- **Cycle time improvement**
- **Data availability latency**

Objective

- Create near real-time operational visibility
- Reduce manual document effort
- Enable trusted, governed execution across partners

Key Stakeholders & Actors

Manufacturing & Operations Leadership

Quality & Regulatory Teams

Data / AI Practitioners

Supply Chain & Planning

Business Understanding

Breaking Manufacturing Siloes into Use Cases

- Partner document ingestion and extraction
- Batch status interpretation
- Deviation and exception analysis
- Operational Q&A for planners and quality teams

Data Understanding

- External partner documents and attachments
- Internal manufacturing and quality systems
- Planning and demand signals

Data Preparation

- Reusable ingestion pipelines
- SSOT-aligned data assets
- Data harmonization and lineage

AI / ML Modelling Approach

- GenAI-based document extraction
- Agentic workflows for reasoning and validation
- Human-in-the-loop checkpoints

Evaluation & Governance

- Sprint-level business validation
- Audit-ready outputs and approvals

Deployment & Operating Model

- Incremental, PI-based delivery
- Reuse-first integration strategy

Monitoring & Continuous Improvement

- Continuous onboarding of new partners
- Expansion to multi-entity workflows

Business Outcomes

Near real-time manufacturing visibility

Higher adoption through automated AI

Scalable, governed control-tower execution

www.calibo.com