ENTERPRISE DATA & AI PLATFORM

SkyIntelligence

From Data to Intelligent Action

Ontology-Driven · GraphRAG & KAG · Agentic AI

MCP + A2A · Sovereign AI

Executive & Technical Briefing | 2026

The Enterprise Data & AI Gap

Enterprises invest billions in data platforms but still struggle to bridge the gap from data infrastructure to autonomous, governed AI decisions at scale.

01 Data-Rich but Decision-Poor

Invested in data lakes & warehouses, but analytics stop at dashboards — no semantic layer for AI-driven decisions

02 AI Disconnected from Data Platform

AI pilots run in silos, disconnected from enterprise data — results can't be trusted, audited, or scaled

03 No Unified Data + AI Governance

Separate data & AI governance tools create compliance gaps, cost sprawl, and shadow AI

04 AI Hallucination Risk

LLMs generate plausible but incorrect outputs — no grounding layer to verify facts against enterprise knowledge

05 AI Black-Box Decisions

AI outputs lack explainability, traceability, and audit trails — unacceptable for regulated enterprise decisions

ENTERPISE DATA PLATFORM & AI GAP
BROKEN DATA FLOW
GAP
DATA WAREHOUSE & SILOS
AI INTELLIGENCE

The Corporate Brain

An intelligent enterprise nervous system that perceives, understands, reasons, decides, and acts — with full governance and sovereign AI.

Memory
Skybase + Agent Memory

Cognition
SkyOntology + GraphRAG & KAG

Perception
Multimodal Processing

Action
Agents + MCP + A2A + UI Auto

Governance
AI Gateway + Guardrails + AI Ops
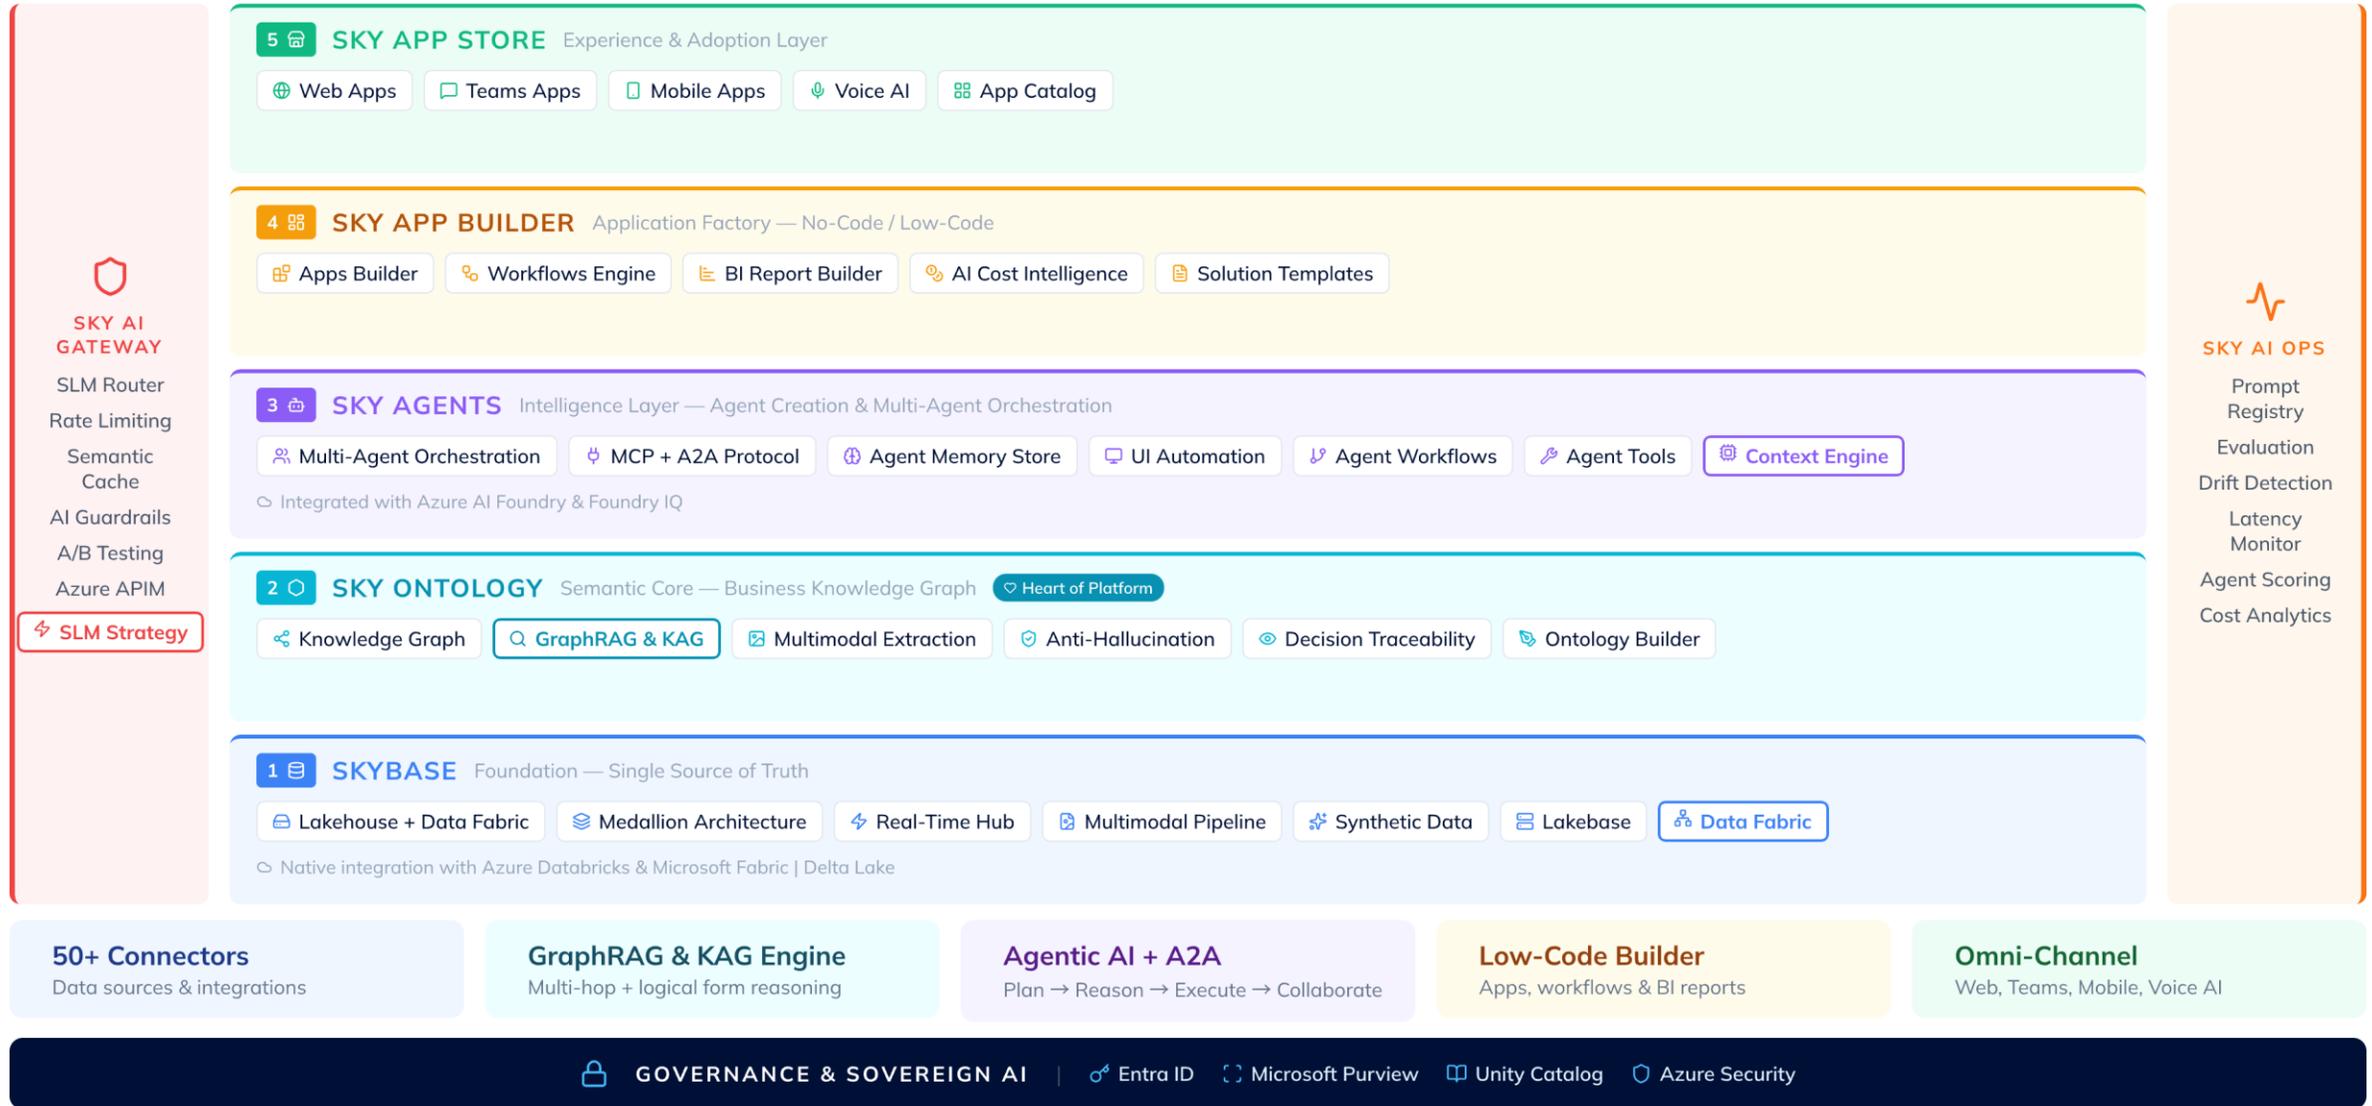
Attention
Sky Context Engine

ENTERPRISE
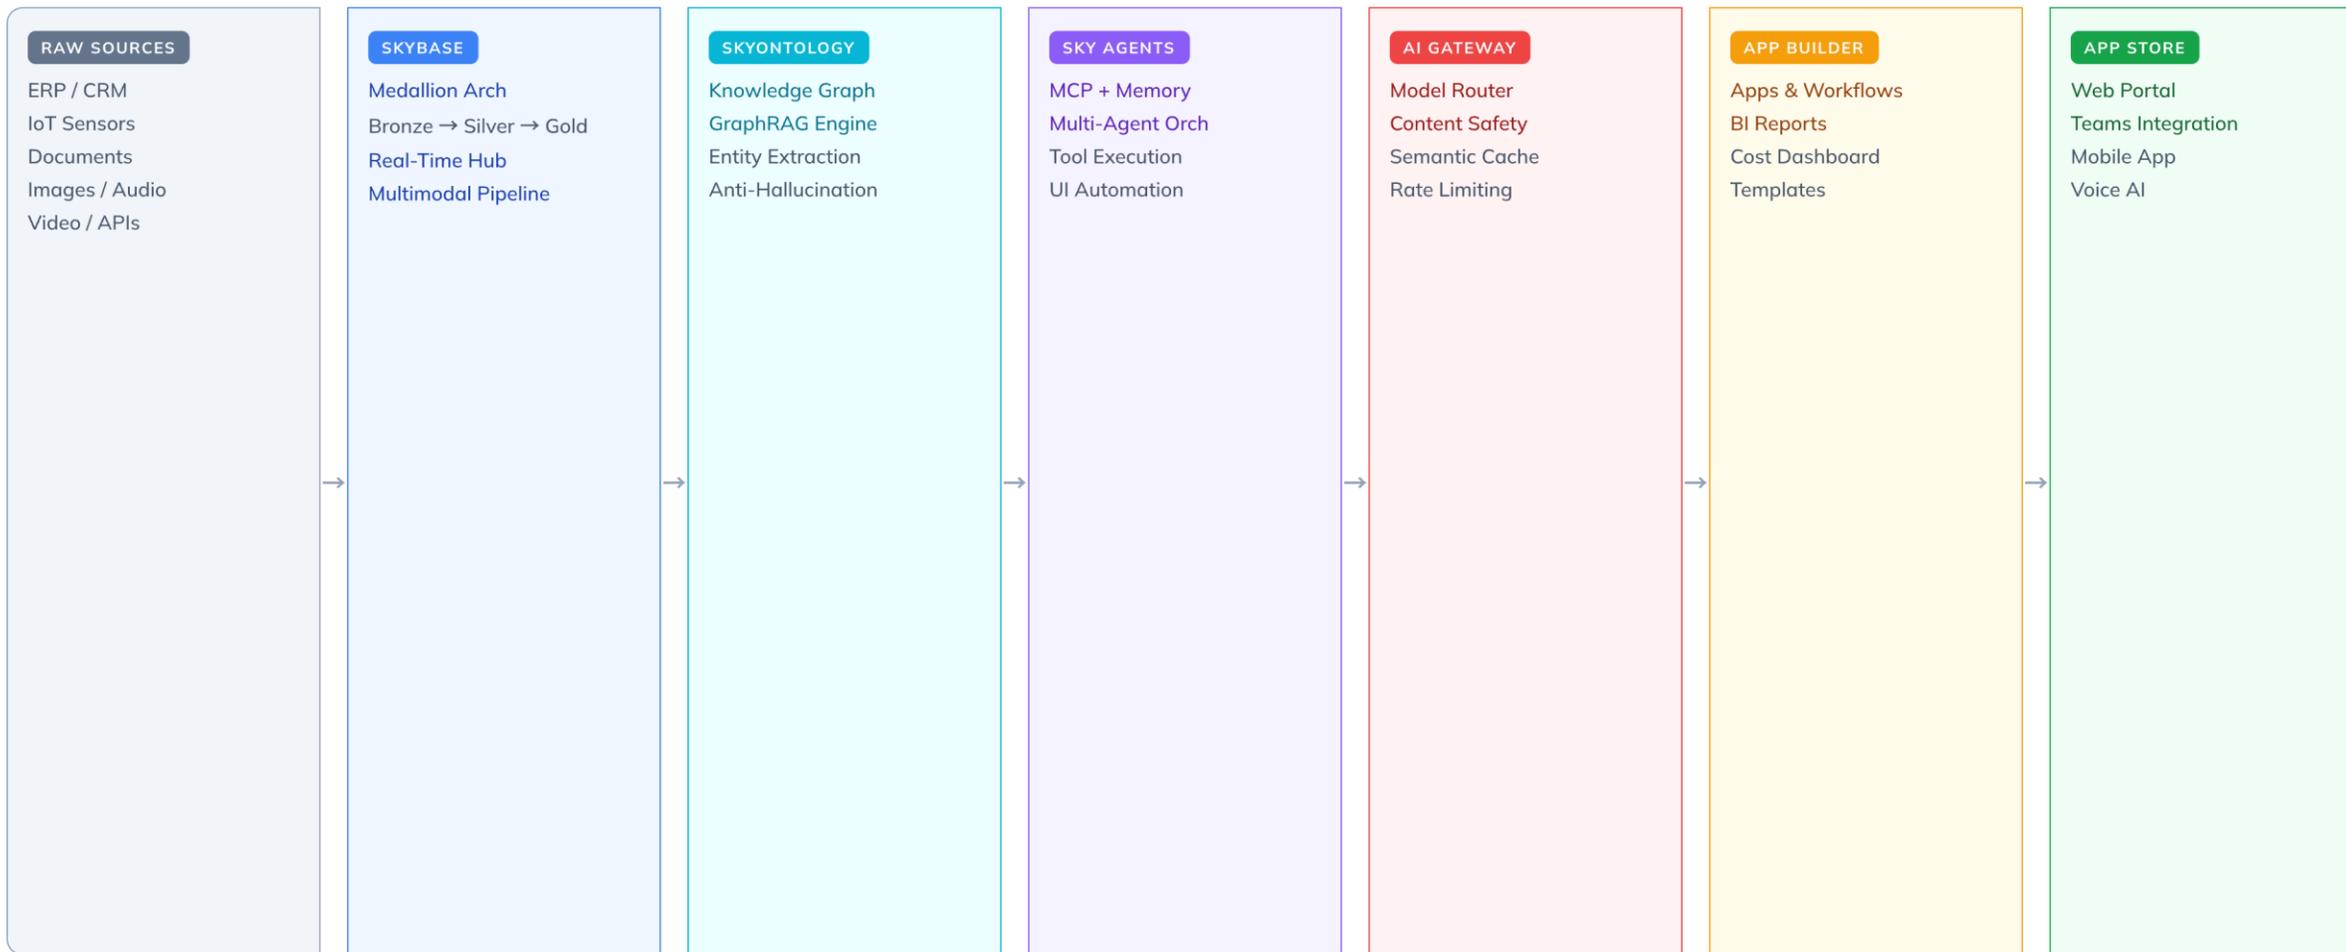TECHNOLOGY

SkyIntelligence - 5-Layer Enterprise AI Stack

Modular architecture from data foundation to AI-powered applications — with cross-cutting governance, SLM cost optimization, programmable guardrails, and sovereign AI deployment.

5 🏬 SKY APP STORE — Experience & Adoption Layer

- 🌐 Web Apps
- 💬 Teams Apps
- 📱 Mobile Apps
- 🎤 Voice AI
- 🔲 App Catalog

4 🔲 SKY APP BUILDER — Application Factory — No-Code / Low-Code

- 🔲 Apps Builder
- 🔗 Workflows Engine
- 📊 BI Report Builder
- 🔗 AI Cost Intelligence
- 📄 Solution Templates

3 🤖 SKY AGENTS — Intelligence Layer — Agent Creation & Multi-Agent Orchestration

- 👥 Multi-Agent Orchestration
- 🔀 MCP + A2A Protocol
- 🌐 Agent Memory Store
- 🖥 UI Automation
- 🔀 Agent Workflows
- 🔧 Agent Tools
- ⚙️ Context Engine

☁ Integrated with Azure AI Foundry & Foundry IQ

2 ◯ SKY ONTOLOGY — Semantic Core — Business Knowledge Graph ♡ Heart of Platform

- 🔗 Knowledge Graph
- 🔍 GraphRAG & KAG
- 🖼 Multimodal Extraction
- 🛡 Anti-Hallucination
- 👁 Decision Traceability
- 🔗 Ontology Builder

1 🗄 SKYBASE — Foundation — Single Source of Truth

- 🗄 Lakehouse + Data Fabric
- 🗂 Medallion Architecture
- ⚡ Real-Time Hub
- 🗄 Multimodal Pipeline
- ✨ Synthetic Data
- 🗄 Lakebase
- 🔗 Data Fabric

☁ Native integration with Azure Databricks & Microsoft Fabric | Delta Lake

SKY AI GATEWAY
- SLM Router
- Rate Limiting
- Semantic Cache
- AI Guardrails
- A/B Testing
- Azure APIM
- ⚡ SLM Strategy

SKY AI OPS
- Prompt Registry
- Evaluation
- Drift Detection
- Latency Monitor
- Agent Scoring
- Cost Analytics

50+ Connectors
Data sources & integrations

GraphRAG & KAG Engine
Multi-hop + logical form reasoning

Agentic AI + A2A
Plan → Reason → Execute → Collaborate

Low-Code Builder
Apps, workflows & BI reports

Omni-Channel
Web, Teams, Mobile, Voice AI

🔒 **GOVERNANCE & SOVEREIGN AI** | 🔑 Entra ID 🔲 Microsoft Purview 📖 Unity Catalog 🛡 Azure Security
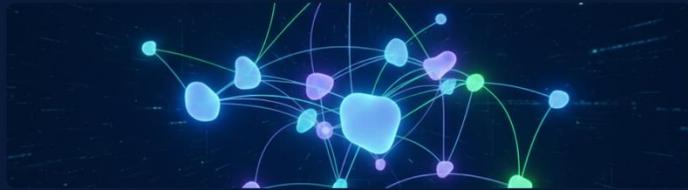
End-to-End Data Flow

From raw data sources to AI-driven business outcomes — through all 5 layers + 2 cross-cutting capabilities of the SkyIntelligence Platform
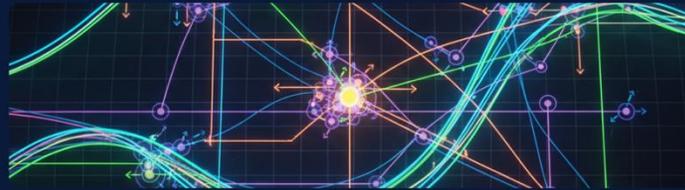
RAW SOURCES

ERP / CRM
IoT Sensors
Documents
Images / Audio
Video / APIs

SKYBASE

Medallion Arch
Bronze → Silver → Gold
Real-Time Hub
Multimodal Pipeline

SKYONTOLOGY

Knowledge Graph
GraphRAG Engine
Entity Extraction
Anti-Hallucination

SKY AGENTS

MCP + Memory
Multi-Agent Orch
Tool Execution
UI Automation

AI GATEWAY

Model Router
Content Safety
Semantic Cache
Rate Limiting

APP BUILDER

Apps & Workflows
BI Reports
Cost Dashboard
Templates

APP STORE

Web Portal
Teams Integration
Mobile App
Voice AI

SKY AI OPS — Monitoring Span Across All Stages   Prompt Registry   Evaluation   Drift Detection   Cost Analytics   Latency Monitor

Key Differentiators

Ontology-Centric AI

Grounded in structured business knowledge — not just vector search. Eliminates hallucination by design.

GraphRAG & KAG

Graph traversal + KAG logical form reasoning for superior multi-hop accuracy — query decomposition and ontology-guided inference.
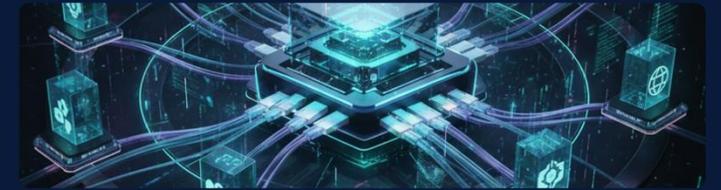
MCP + A2A Open Ecosystem

MCP for AI-tool integration + A2A for agent-to-agent collaboration. Dual open standards, no vendor lock-in, full governance.

Intelligent Data Fabric

Lakehouse core + Data Fabric semantic layer (SkyOntology) + Data Mesh domain ownership — complement, not replace your data platform.

Multimodal AI

Process documents, images, audio, and video — all linked to the unified knowledge graph.

Governance & Sovereign AI

End-to-end governance with programmable guardrails, SLM cost optimization, and full sovereign AI — data, model, compute & governance sovereignty in customer tenant.

Skybase — Architecture

Single Source of Truth — Lakehouse core + Data Fabric (SkyOntology semantic layer) + Data Mesh domain ownership. Native Databricks & Microsoft Fabric integration for ingesting, cleansing, and unifying structured, unstructured, and multimodal data at enterprise scale.
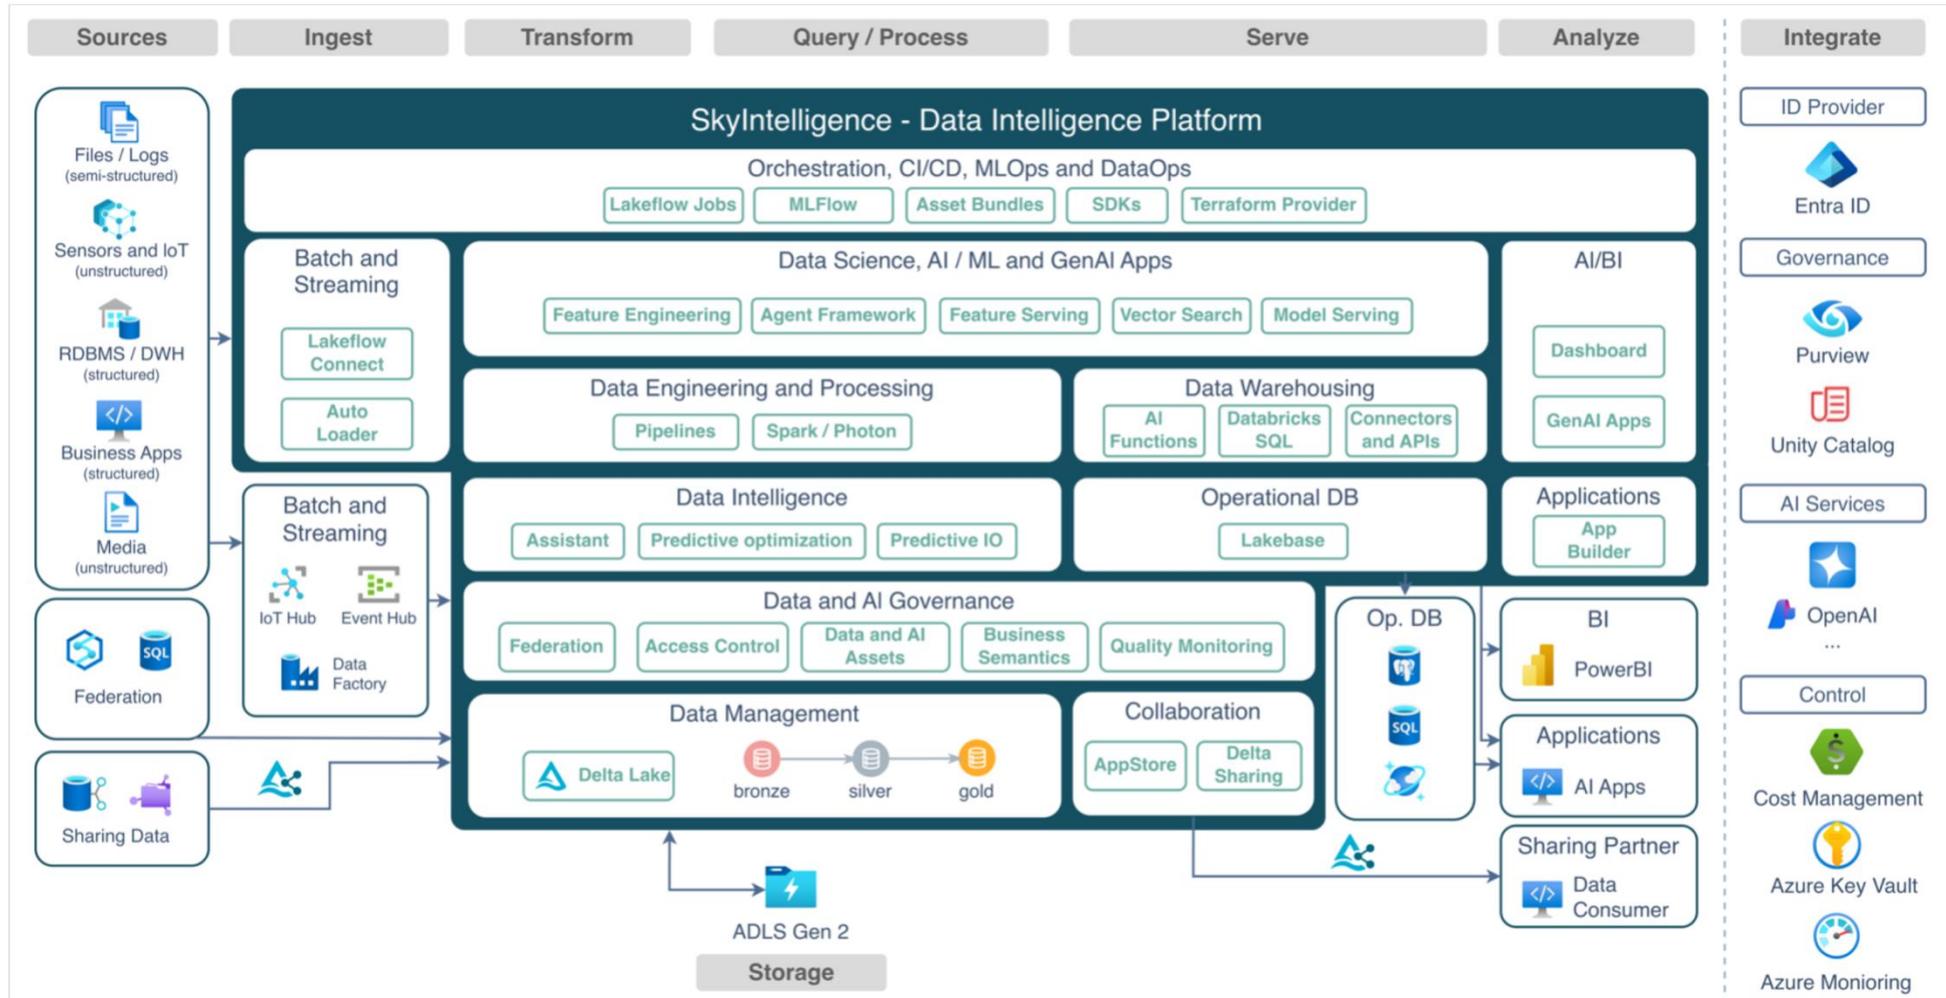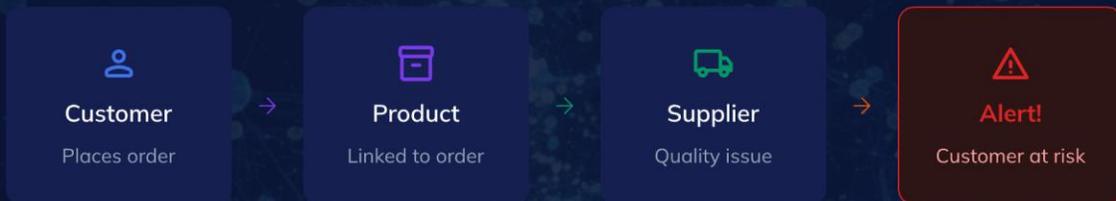
SkyOntology + GraphRAG & KAG

Ontology-grounded AI with KAG logical reasoning — eliminates hallucination and delivers superior retrieval accuracy

Multi-Hop Reasoning Flow

Entity → Relationship → Traversal → Insight

Customer		Product		Supplier		Alert!
Places order	→	Linked to order	→	Quality issue	→	Customer at risk

💬 *"Which customers are affected by Supplier X's quality issue?"*

🔍 Traditional RAG

1. Vector search: "Supplier X quality"

2. Returns docs mentioning Supplier X

3. Misses: Customer → Order → Product → Supplier links

✕ Incomplete — only finds directly mentioned customers

⚛ GraphRAG

1. Find Supplier X in Knowledge Graph

2. Traverse: Supplier → Products supplied

3. Traverse: Products → Orders → Customers

✓ Complete — finds ALL affected customers via relationships

KEY VALUE

Zero Hallucination
Ontology-constrained generation ensures every answer is grounded in verified enterprise data

Multi-Hop Reasoning
Query decomposition + logical inference across complex entity relationships

Single Source of Truth
Structured knowledge graph unifies all business entities & relationships for AI

Retrieval Accuracy Improvement

+58% vs traditional RAG (KAG-enhanced)

CAPABILITIES

Knowledge Graph · GraphRAG · KAG Engine

Anti-Hallucination · Hybrid Search
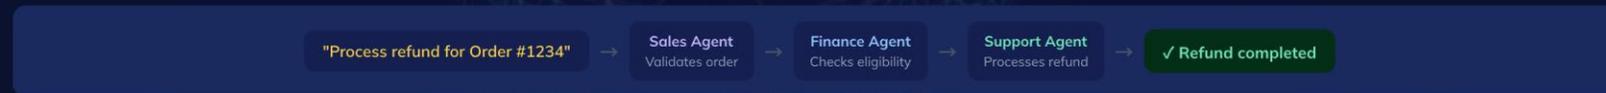
Sky Agents + MCP & A2A Architecture

Autonomous AI agents with MCP (agent-to-tool) + A2A (agent-to-agent) dual-protocol — orchestrate workflows, connect tools, and collaborate across enterprise boundaries with full governance

Agent Orchestration Engine

Plan → Reason → Act → Learn

Planning
Task decomposition & strategy

→

Reasoning
Context analysis & decision making

→

Execution
Tool calls via MCP & actions

→

Learning
Memory store & improvement

A2A IN ACTION — Agent-to-Agent Collaboration Example

"Process refund for Order #1234" → **Sales Agent** Validates order → **Finance Agent** Checks eligibility → **Support Agent** Processes refund → ✓ Refund completed

MCP Connections: GitHub | Slack | Salesforce | ServiceNow | Databases | SAP | Oracle | Custom Tools

KEY VALUE

Build Once, Run Anywhere
MCP for tools + A2A for agent collaboration — no custom integrations needed

Gets Smarter Over Time
Persistent memory at session, user & org levels — agents learn from every interaction

Legacy System Access
Browser & desktop automation lets agents work with systems that have no API

Enterprise-Grade Security
Entra ID auth, RBAC, audit logging & policy gates on every agent action

CAPABILITIES
Multi-Agent | MCP + A2A | Agent Memory
Computer Use | Governance | Human-in-Loop

Sky AI Gateway — SLM Strategy, Safety & Cost Control

Intelligent routing with SLM plan-and-execute pattern, programmable guardrails, semantic caching, and real-time cost visibility — all in one gateway

Gateway Request Pipeline

Example: Sales agent asks complex question → GPT-5.3 plans response → Domain SLM executes at 1/10th cost → Cached for next similar query

REQUEST SOURCES | Sky Agents | App Builder | External APIs

Model Router
Route to optimal model per task

Semantic Cache
Deduplicate similar queries, save cost

Safety Filter
Block harmful & policy violations

Rate Limiter
Token quotas & budget enforcement

A/B Testing
Split traffic to test model quality

⚡ LIVE FLOW | "Summarize Q3" → Safety ✓ → Cache miss → GPT-5.3 plans → SLM executes → ✓ Cached & served

38% cost saved | **Per-Agent** budget alerts

KEY VALUE

90% Cost Reduction
Frontier LLMs plan complex tasks, domain SLMs execute routine work at 1/10th the cost

Smart Caching
Semantic cache deduplicates similar queries — reduce latency and cost automatically

Zero Trust Security
Every request validated: jailbreak detection, PII redaction, prompt injection & policy enforcement
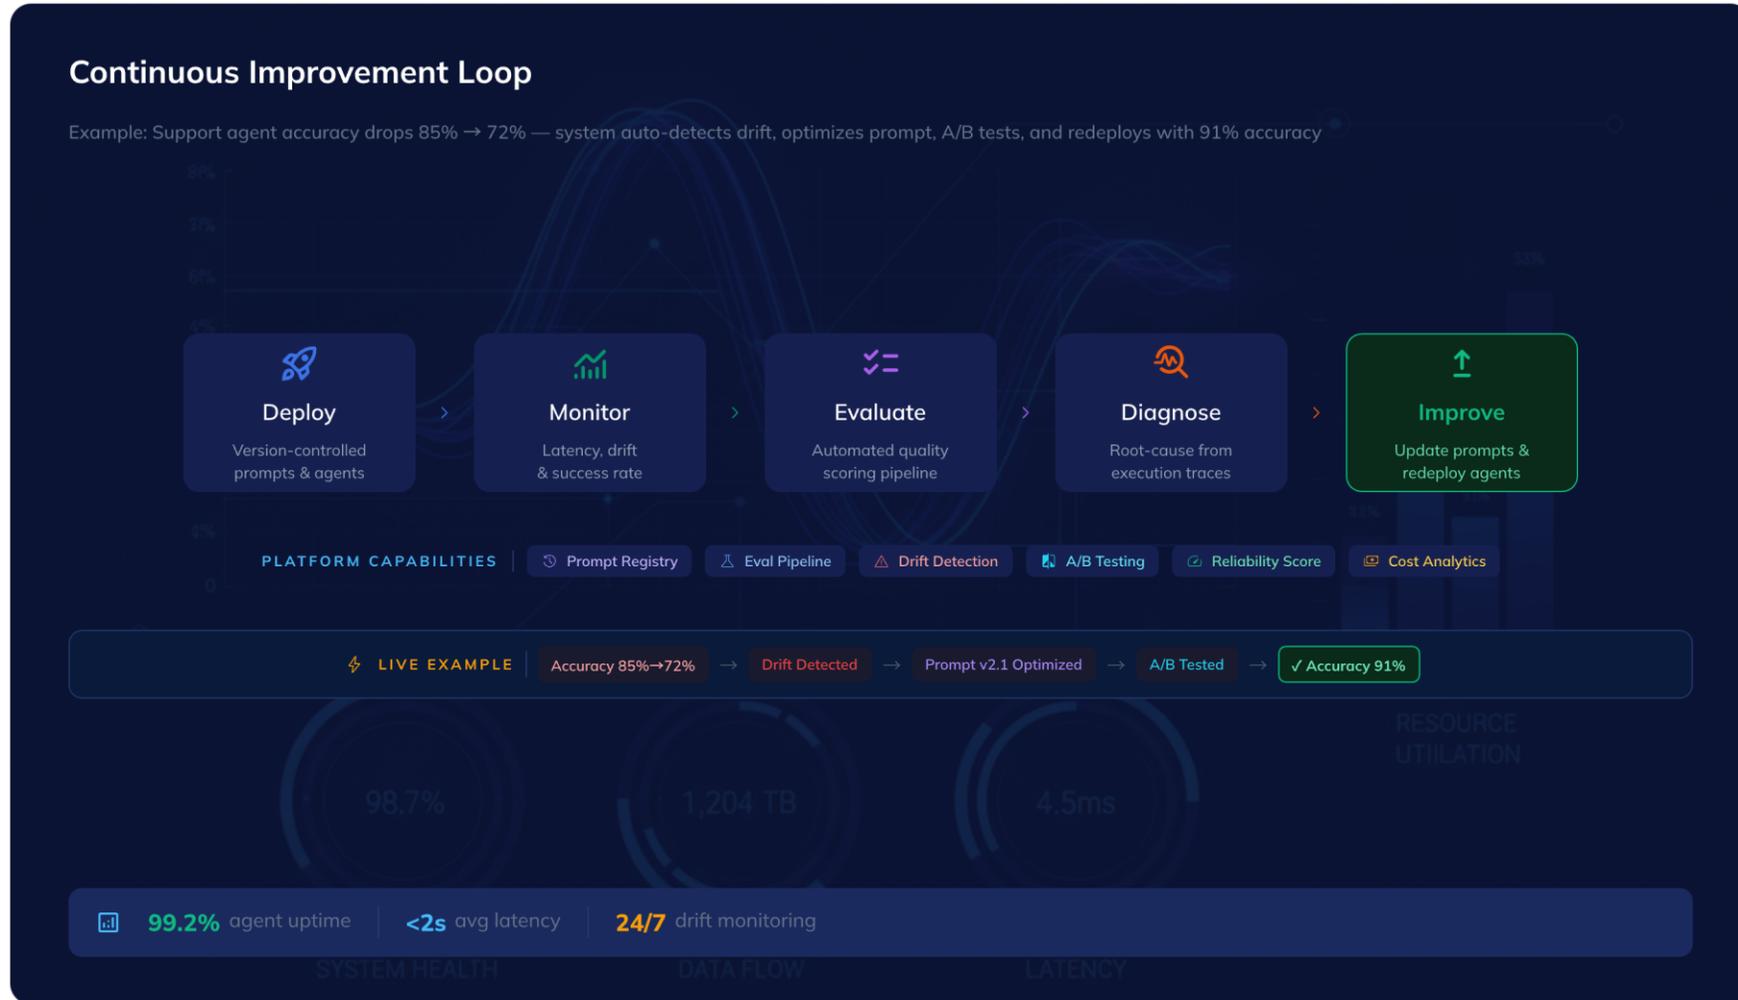
Full Cost Visibility
Per-agent, per-workflow cost analytics with budget alerts & optimization recommendations

PIPELINE STAGES

Content Safety   Semantic Cache   Rate Limiter
Model Router   A/B Testing   Observability

Sky AI Ops — Production-Grade LLMOps

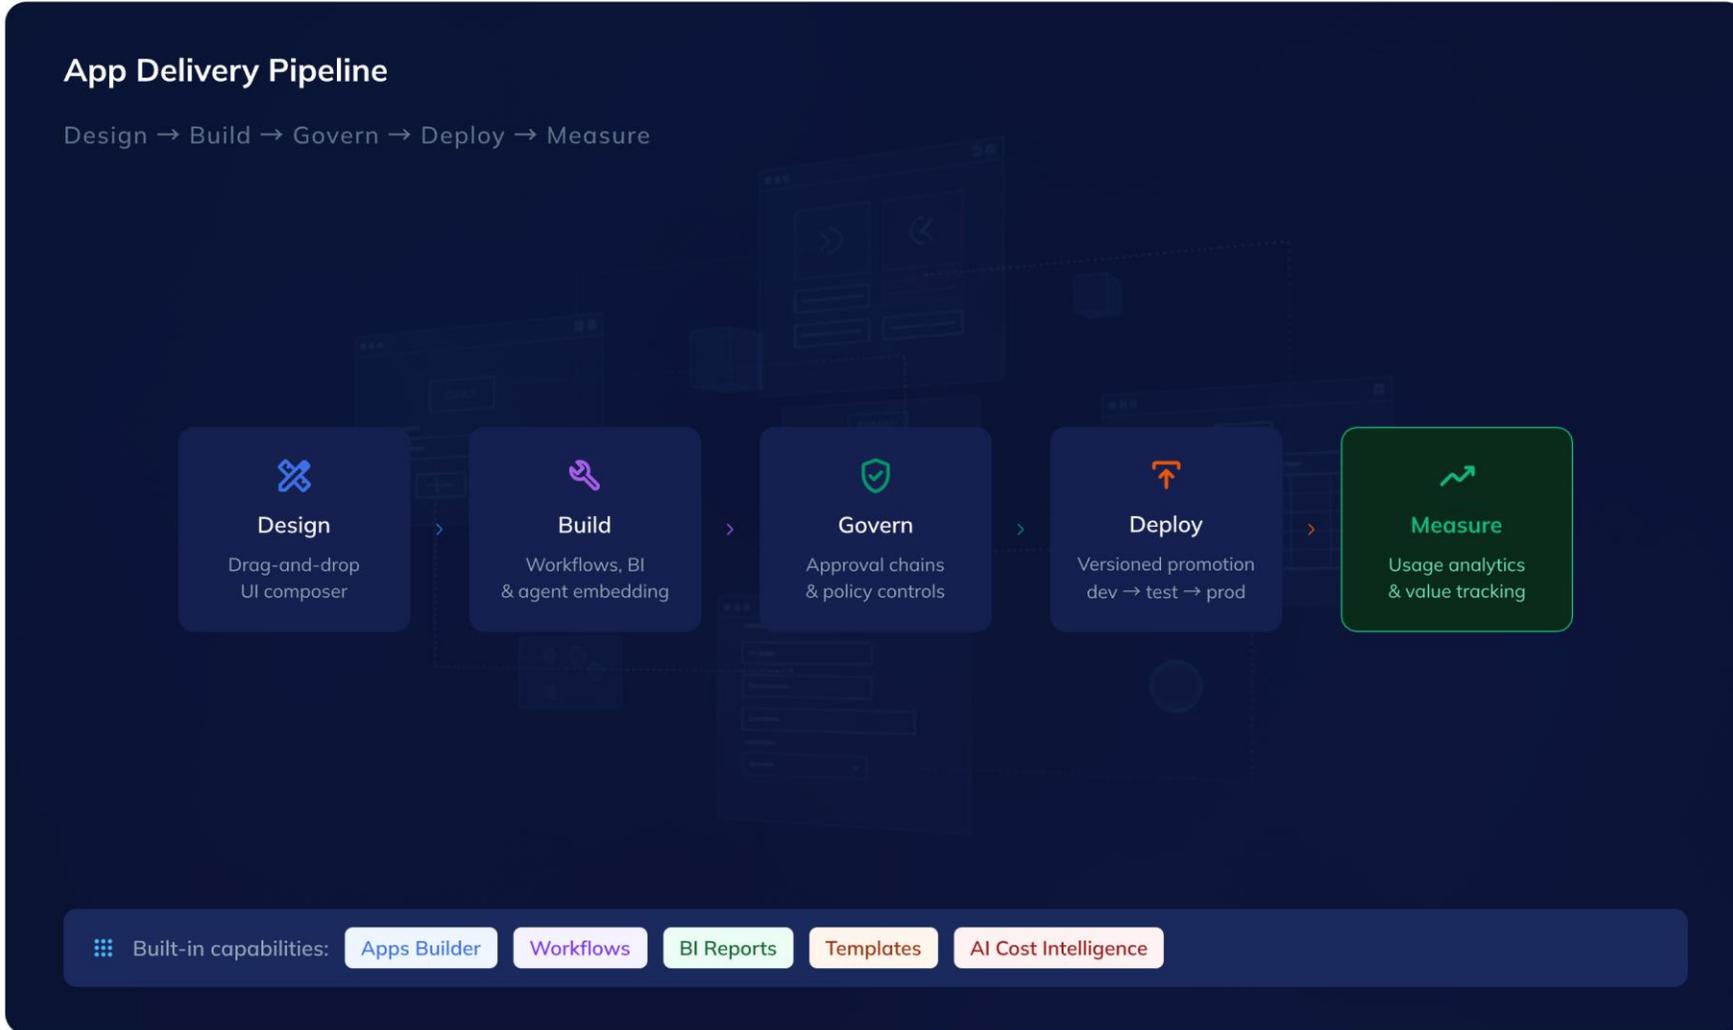End-to-end observability for AI agents — prompt versioning, evaluation, drift detection, and incident analysis

Continuous Improvement Loop

Example: Support agent accuracy drops 85% → 72% — system auto-detects drift, optimizes prompt, A/B tests, and redeploys with 91% accuracy

Deploy
Version-controlled prompts & agents

Monitor
Latency, drift & success rate

Evaluate
Automated quality scoring pipeline

Diagnose
Root-cause from execution traces

Improve
Update prompts & redeploy agents

PLATFORM CAPABILITIES   🕘 Prompt Registry   ⚗ Eval Pipeline   ⚠ Drift Detection   📊 A/B Testing   🔲 Reliability Score   💳 Cost Analytics

⚡ **LIVE EXAMPLE**   Accuracy 85%→72% → Drift Detected → Prompt v2.1 Optimized → A/B Tested → ✓ Accuracy 91%

RESOURCE UTILIZATION

98.7%   1,204 TB   4.5ms

📊 **99.2%** agent uptime | **<2s** avg latency | **24/7** drift monitoring

SYSTEM HEALTH   DATA FLOW   LATENCY

KEY VALUE

🛡 **Proactive, Not Reactive**
Drift detection catches quality degradation before users notice — automated alerts, not manual reviews

✓ **Automated Quality Gates**
Every response scored for groundedness, relevance & coherence — no manual spot-checking

🐷 **Full Cost Visibility**
Per-agent, per-workflow token cost analytics with budget alerts & optimization recommendations

🚀 **Zero-Downtime Optimization**
A/B test prompt improvements in production without disrupting live agents — deploy winners automatically

CAPABILITIES
Prompt Registry   Eval Pipeline   Drift Detection
A/B Testing   Reliability Score   Cost Analytics

Sky App Builder — No-Code AI Apps

Drag-and-drop application factory — from prototype to production app with built-in governance, BI, and AI cost intelligence

App Delivery Pipeline

Design → Build → Govern → Deploy → Measure

Design
Drag-and-drop
UI composer

Build
Workflows, BI
& agent embedding

Govern
Approval chains
& policy controls

Deploy
Versioned promotion
dev → test → prod

Measure
Usage analytics
& value tracking

Built-in capabilities: Apps Builder | Workflows | BI Reports | Templates | AI Cost Intelligence

---

Drag-and-Drop Apps Builder
Visual interface for creating apps with reusable UI components

Workflows Engine
Visual design of business logic and agent orchestration

BI Semantic Report Builder
Native BI dashboards — reduces third-party license dependency

Reusable Solution Templates
Domain starter kits for service ops, sales, risk & compliance

AI Cost Intelligence Dashboard
Per-agent, per-use-case cost tracking with budget alerts & FinOps

Embedded Agent Experiences
Native copilots & agent actions inside forms, reports & screens

Sky App Store — Enterprise AI Delivery

Omni-channel distribution — Web, Teams, Mobile & Voice AI — with role-based targeting, governed release, and adoption analytics

Omni-Channel Delivery & Adoption Loop

Publish → Target → Deliver → Adopt → Improve

Publish	Target	Deliver	Adopt	Improve
Catalog listing & approval gates	Role-based audience targeting	Web, Teams, Mobile & Voice	Usage analytics & value tracking	Feedback loop to ontology & agents

OMNICHAANEL DELIVERY

Channels: **Web Apps** **Microsoft Teams** **Mobile Apps** **Voice AI**

Enterprise App Catalog
Centralized marketplace for approved workflows, agent apps & BI by domain

Role-Based Targeting
Right app to right audience by function, geography & authority

Governed Release Management
Staged rollout, approval gates & rollback controls for safety

Voice AI Channel
Voice-first experiences for call center copilots & field operations

Adoption & Value Analytics
Track active users, cycle-time impact & value realization per use case

Feedback-to-Improvement Loop
User feedback feeds back to ontology, agents & app design

Sovereign AI Deployment Model

Sovereign AI / PaaS-First
Full deployment in customer Azure tenant — data, model, compute & governance sovereignty

Data & Model Sovereignty
Customer owns all data, fine-tuned models & inference — zero cross-tenant exposure

Entra ID + RBAC
Enterprise identity & access

Full Audit Trail
Every action logged & traceable

Compliance Ready

ISO 27001   SOC 2   GDPR   HIPAA   EU AI Act   PCI DSS

CYBER GUARDIAN
SOLUTIONS

Sovereign AI by Default

All data, models, inference & governance remain within the customer's Azure tenant — full AI sovereignty without additional architecture or configuration

Full Data Sovereignty
All enterprise data — structured, unstructured, embeddings, model outputs & AI decision logs — remains within customer Azure tenant and chosen geographic region

Model Sovereignty
Deploy models within customer Azure subscription — private endpoints, self-hosted open-source models & customer fine-tuned SLMs. Zero external API calls for sensitive workloads.

Compute Sovereignty
Leverage Azure sovereign cloud offerings — control where AI compute physically resides. Azure Government, regional zones & dedicated infrastructure.

Governance Sovereignty
All governance through customer's Entra ID, Purview & Unity Catalog. No SkyIntelligence admin dependency outside customer control plane.

ARCHITECTURE

Private Endpoints	Customer-Managed Keys	Managed Identity
Azure Private Link	Azure Key Vault	Zero-secret RBAC

KEY VALUE

No Data Exfiltration Risk
All data stays within customer's Azure tenant — zero external API calls for sensitive workloads

No Vendor Lock-In
Customer owns all data, models & infrastructure — switch or disconnect anytime

Full Regulatory Compliance
Pre-aligned with EU AI Act, GDPR, HIPAA, PCI DSS, FedRAMP & Japan APPI out of the box

Customer Retains Full Control
All governance via customer's Entra ID, Purview & Key Vault — no SkyIntelligence admin dependency

SOVEREIGNTY PILLARS
Data   Model   Compute   Governance

REGULATORY ALIGNMENT

⊘ EU AI Act	⊘ Japan APPI
⊘ PCI DSS	⊘ FedRAMP / SOC 2
⊘ GDPR	⊘ HIPAA / HITRUST

Cross-Industry Enterprise AI

Finance & Banking

Intelligent Fraud Triage

AML/KYC automation

Risk scoring agents

Energy & Utilities

Predictive Asset Maint.

Demand forecasting

Grid operations AI

Retail

Customer 360

Next-best-action

Omnichannel analytics

Manufacturing & Automotive

Quality intelligence

Root-cause analysis

Supplier management

Healthcare

Clinical documentation

Patient flow optimization

Care coordination AI

Japan's Top-5 Super General Contractor

$25B infrastructure leader — 130+ years building Japan's most iconic bridges, tunnels & skyscrapers — now transforming enterprise operations with AI

| **60%**	**1**	**5**
Faster AI use case delivery	Single Source of Truth platform	AI use cases in development

- DynamicERP: unified project mgmt, HR, AI customer care, cost analytics & 2D/3D geospatial mapping
- Replacing fragmented legacy systems with a single platform
- Strategic shift from siloed point solutions to full-platform AI transformation across the entire enterprise

PRODUCTION
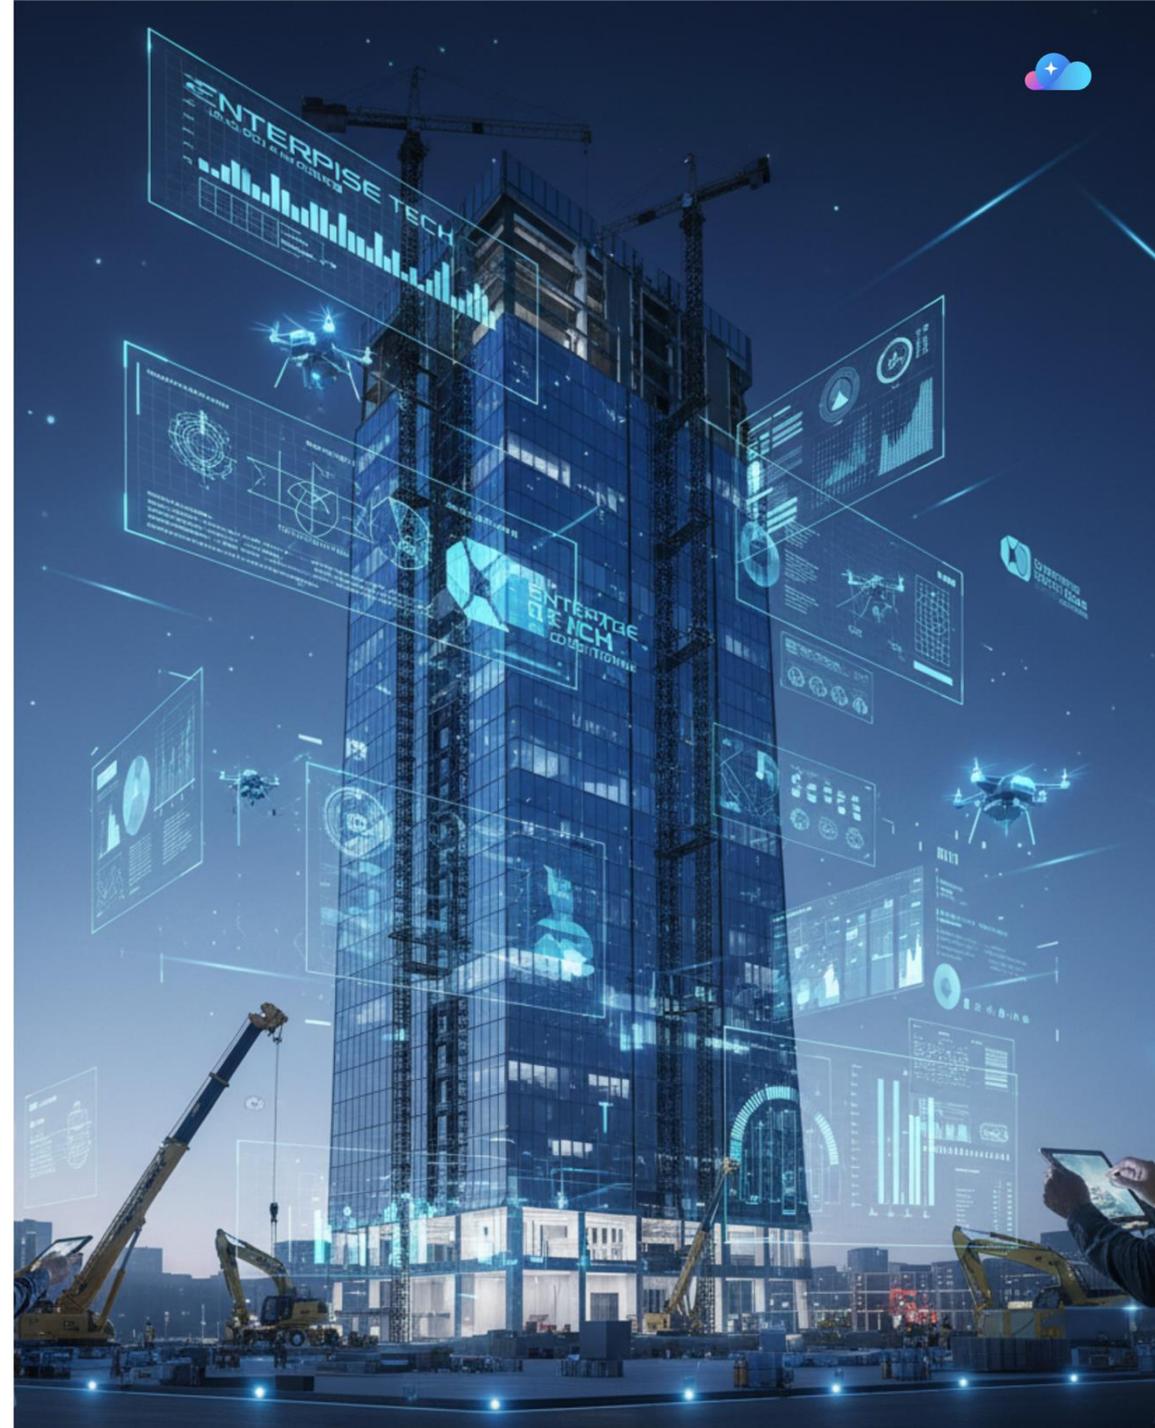
Top Retail Gas & FMCG Enterprise

Leading Vietnamese gas retail & consumer goods distributor — full platform transformation with SkyIntelligence

95K+
Active Customers

500K+
Orders / Year

- Replaced entire legacy BI system with Sky BI — unified reporting across all business units
- Sales team can chat real-time with data for customer care insights — no SQL, no waiting
- Full ERP & CRM replacement — rebuilt on SkyIntelligence for modern data-driven operations

PRODUCTION

SkyTrans — AI Translation at Scale

Enterprise translation application built natively on the full 5-layer SkyIntelligence stack

10K
Monthly active users

$700K
Saved in H2 2025

5B
Characters projected 2026

PRODUCTION **$2.5M PROJECTED SAVINGS 2026**

Vietnam's #1 Retail & Pharmacy Chain

2,700+ stores nationwide — processing billions of transactions yearly across electronics, pharmacy & healthcare

Month→Day BI report delivery	**60%** Faster new use case development	**10M+** Customers served nationwide	**5B+** Transactions processed per year

- Sales & customer care teams chat directly with real-time data — instant insights without manual reports
- Full SkyIntelligence stack: Skybase → SkyOntology → Sky Agents → Sky App Builder → Sky BI
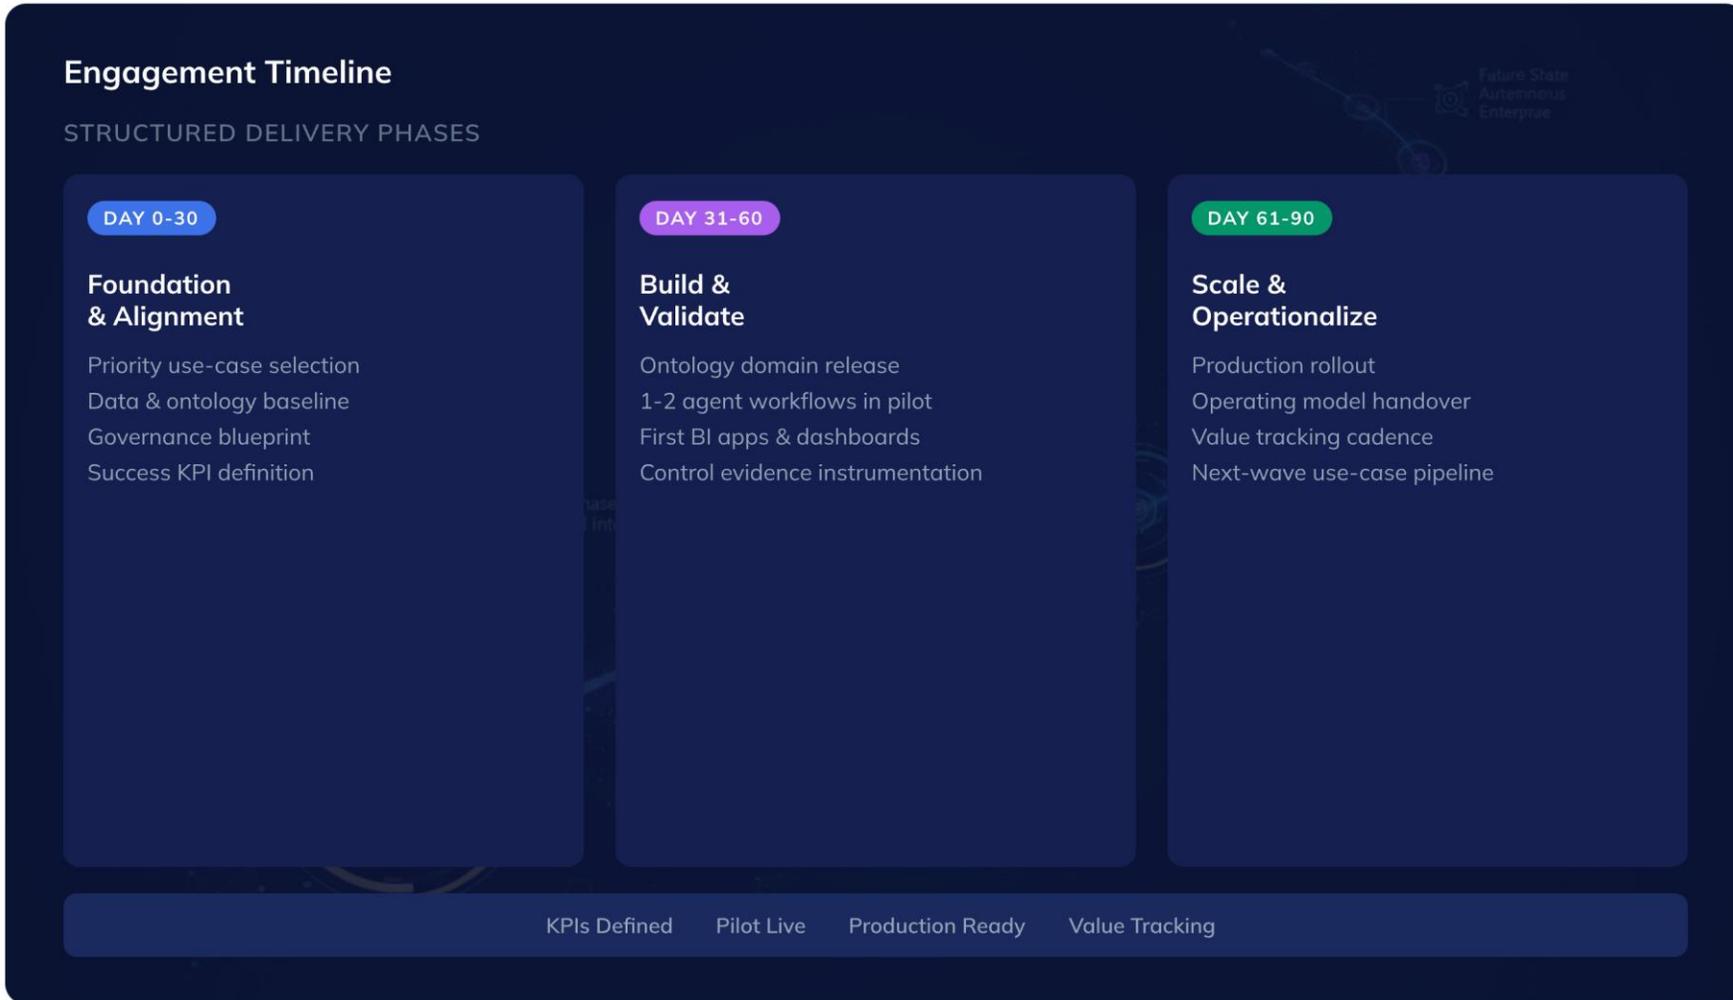
MVP IN DELIVERY

30-60-90 Day Program

A structured engagement model to accelerate time-to-value — from strategic alignment through production rollout.

Engagement Timeline

STRUCTURED DELIVERY PHASES

DAY 0-30

Foundation & Alignment

Priority use-case selection
Data & ontology baseline
Governance blueprint
Success KPI definition

DAY 31-60

Build & Validate

Ontology domain release
1-2 agent workflows in pilot
First BI apps & dashboards
Control evidence instrumentation

DAY 61-90

Scale & Operationalize

Production rollout
Operating model handover
Value tracking cadence
Next-wave use-case pipeline

KPIs Defined   Pilot Live   Production Ready   Value Tracking

Key Deliverables

Use-Case Prioritization

Identify and rank AI use-cases by business impact, feasibility, and data readiness

Ontology & Data Foundation

Enterprise knowledge graph with domain-specific ontology baseline and data integration

Agent Workflows in Pilot

1-2 production-grade AI agent workflows validated with real users and business data

BI Apps & Dashboards

First business intelligence applications with real-time analytics and executive dashboards

Production Rollout

Full production deployment with governance, operating model handover, and value tracking cadence

6-12 Weeks to First Value
vs. 4-9 months industry average

Let's Build Your Corporate Brain

01 Discovery Workshop

02 Priority Use-Case Selection

03 90-Day Program Proposal