Technical Infrastructure Architecture Deck

EQ Creator **Eunice Quezon** Created **Apr 05, 2025, 01:27** Last updated **May 07, 2025, 11:24**

Technical Architecture Document: AI-Augmented System Transformation

Overview

This document outlines the current technical landscape at BBB and the transformation toward a unified, AI-augmented architecture. It is intended for technical teams, architects, data engineers, and cross-functional stakeholders who are contributing to the planning and development of the future-state platform.

The proposed system emphasizes scalable infrastructure, intelligent automation, unified data layers, and Retrieval-Augmented Generation (RAG) for contextual, explainable AI outputs.

1. Current System Workflow (As-Is)

Fragmented and Manual Workflow Landscape

Observations

- Siloed systems across HR, Innovation, Core Services, Account Management, Sales, and Marketing
- Manual inputs via Outlook, Monday.com, Aircall, RingCentral, and HubSpot
- Disparate tools and disconnected integrations (e.g., NMI, Gravity Payments, Mandrill)
- Survey and response processing via InMoment and Email Services

Technical Challenges

- Redundant data entry and duplicate work
- Lack of centralized data synchronization
- Fragmented customer and internal touchpoints
- High dependency on human input for execution

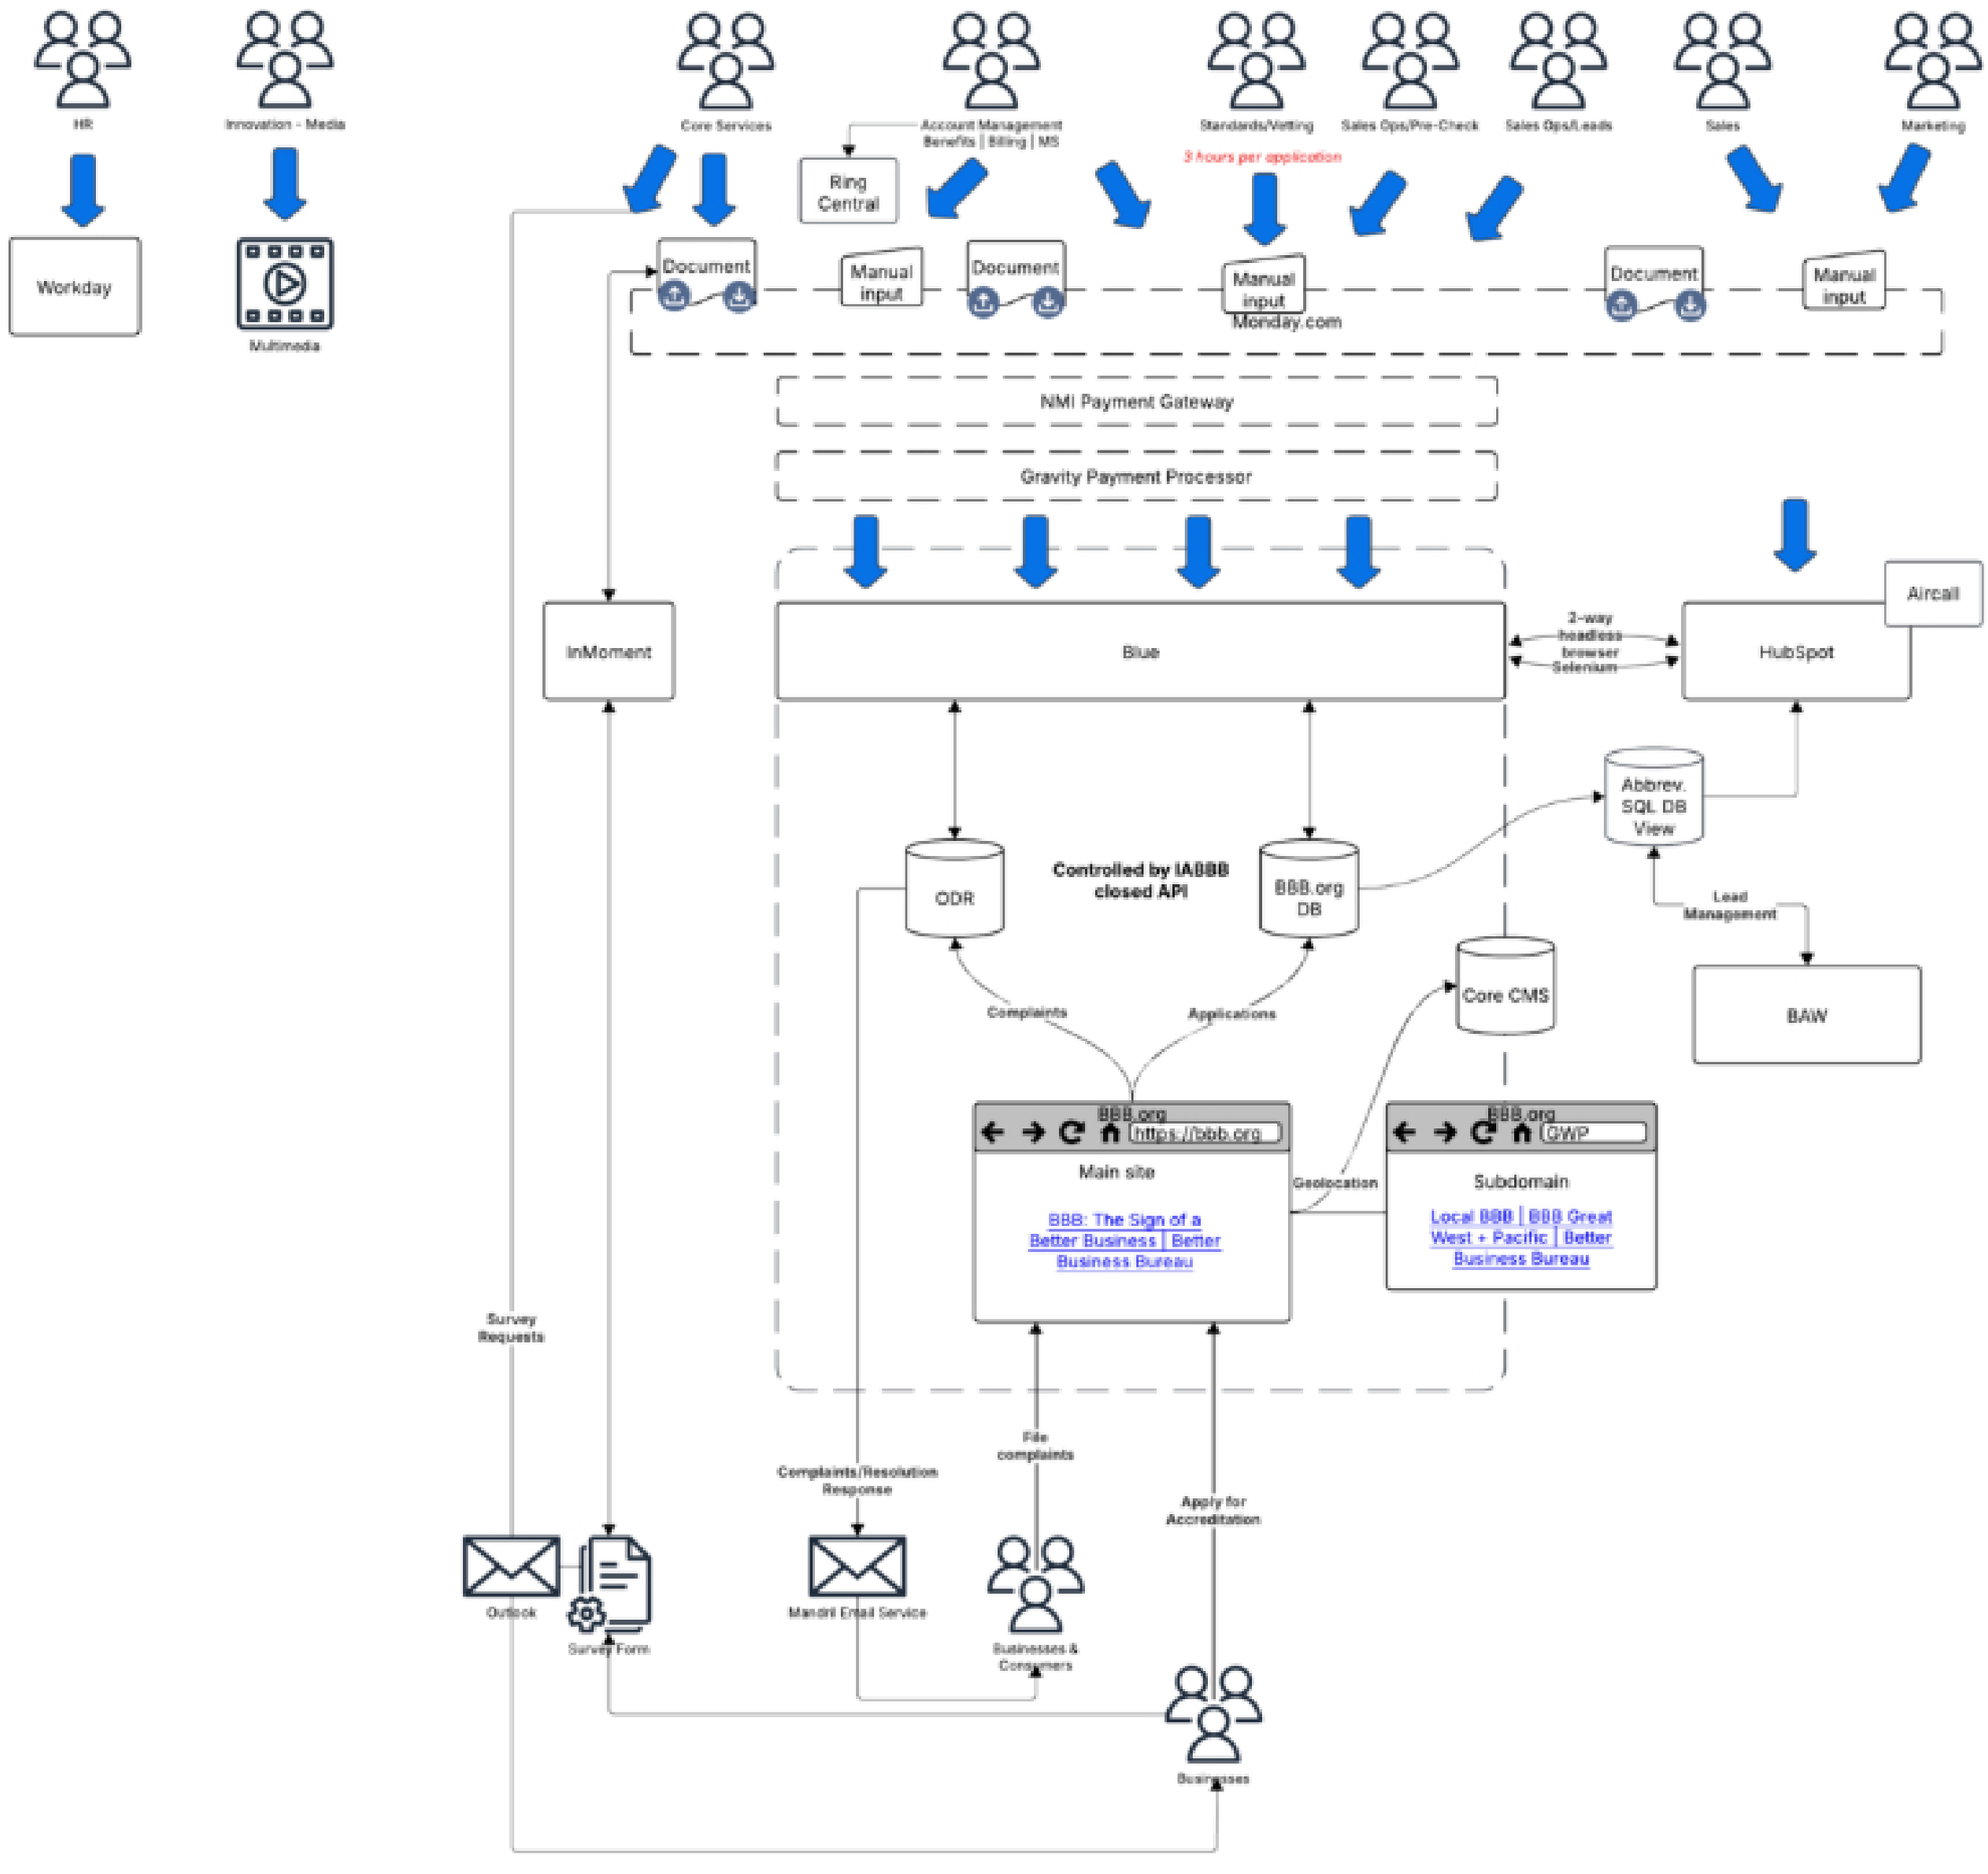

2. Current Data Model (As-Is)

Segregated Data with Limited Semantic Integration

Observations

- Data resides in spreadsheets, Blue SQL DB, Azure SQL DB, and HubSpot
- Analytics tools (Power BI, HubSpot Analytics) operate in silos
- No consolidated semantic layer or centralized processing

Technical Risks

- Incomplete data = inaccurate insights

- Poor lineage tracking & trust issues
- Difficult to enable machine learning/AI with inconsistent structure

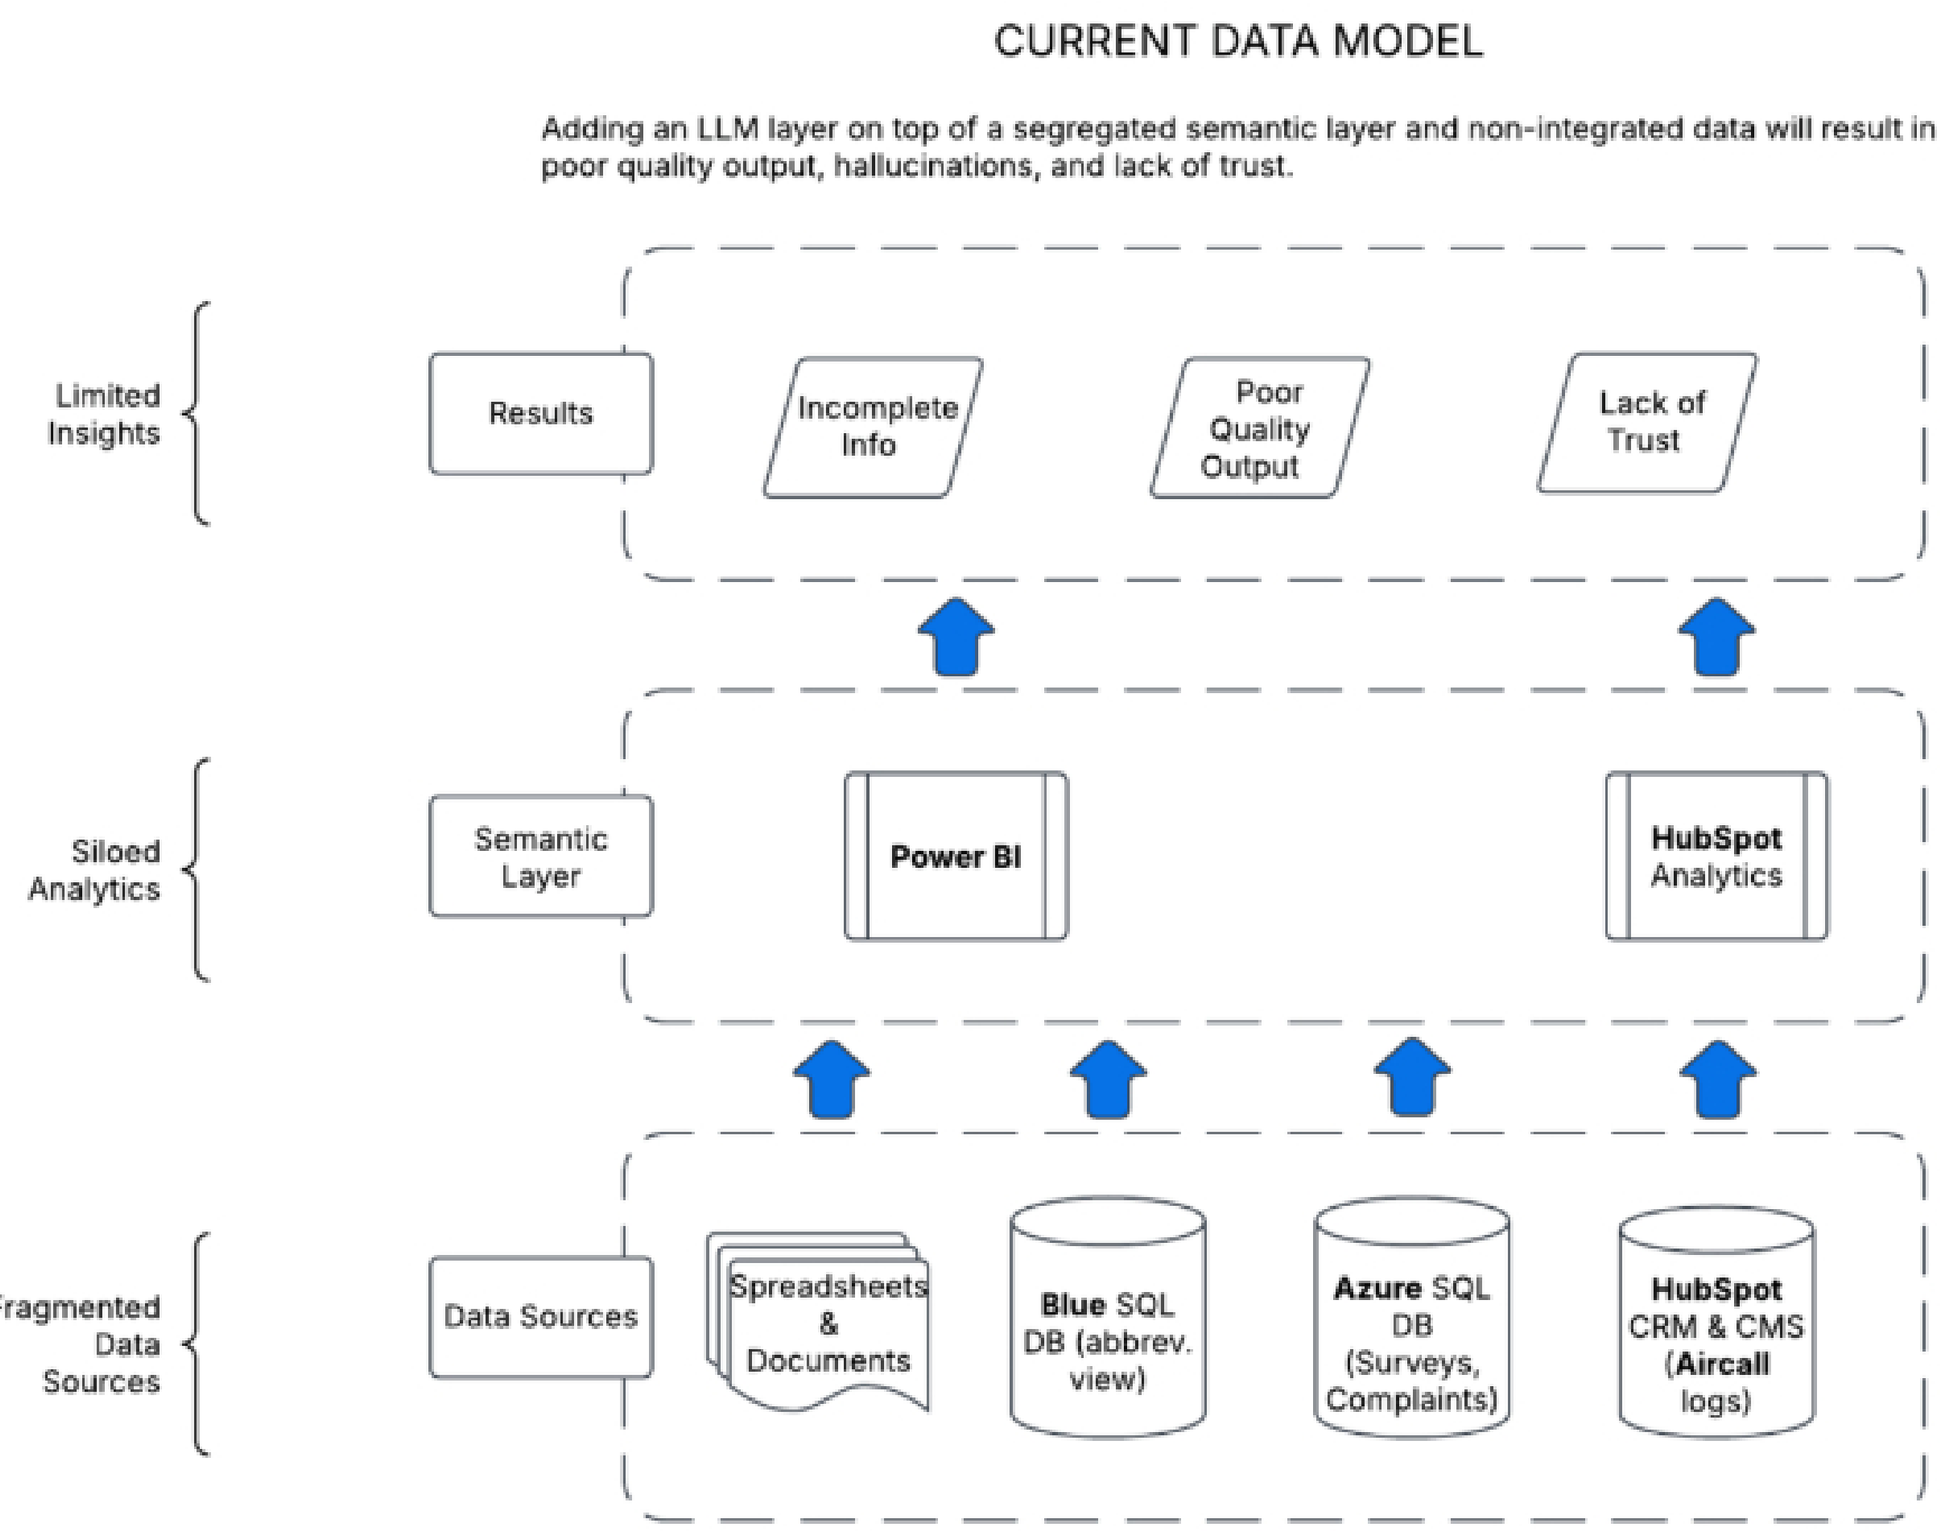

3. Target Data Model (To-Be)

Unified, Explainable Semantic Data Architecture

Features

- Snowflake: Acts as both data lake and warehouse with semantic modeling
- Pinecone: Embedding-powered semantic search
- Intelligence Layer: Agentic RAG layer to interpret tasks and orchestrate AI queries

Benefits

- Real-time, explainable, high-quality outputs
- Stronger AI integrations and knowledge retrieval
- Supports intelligent agents and automated workflows

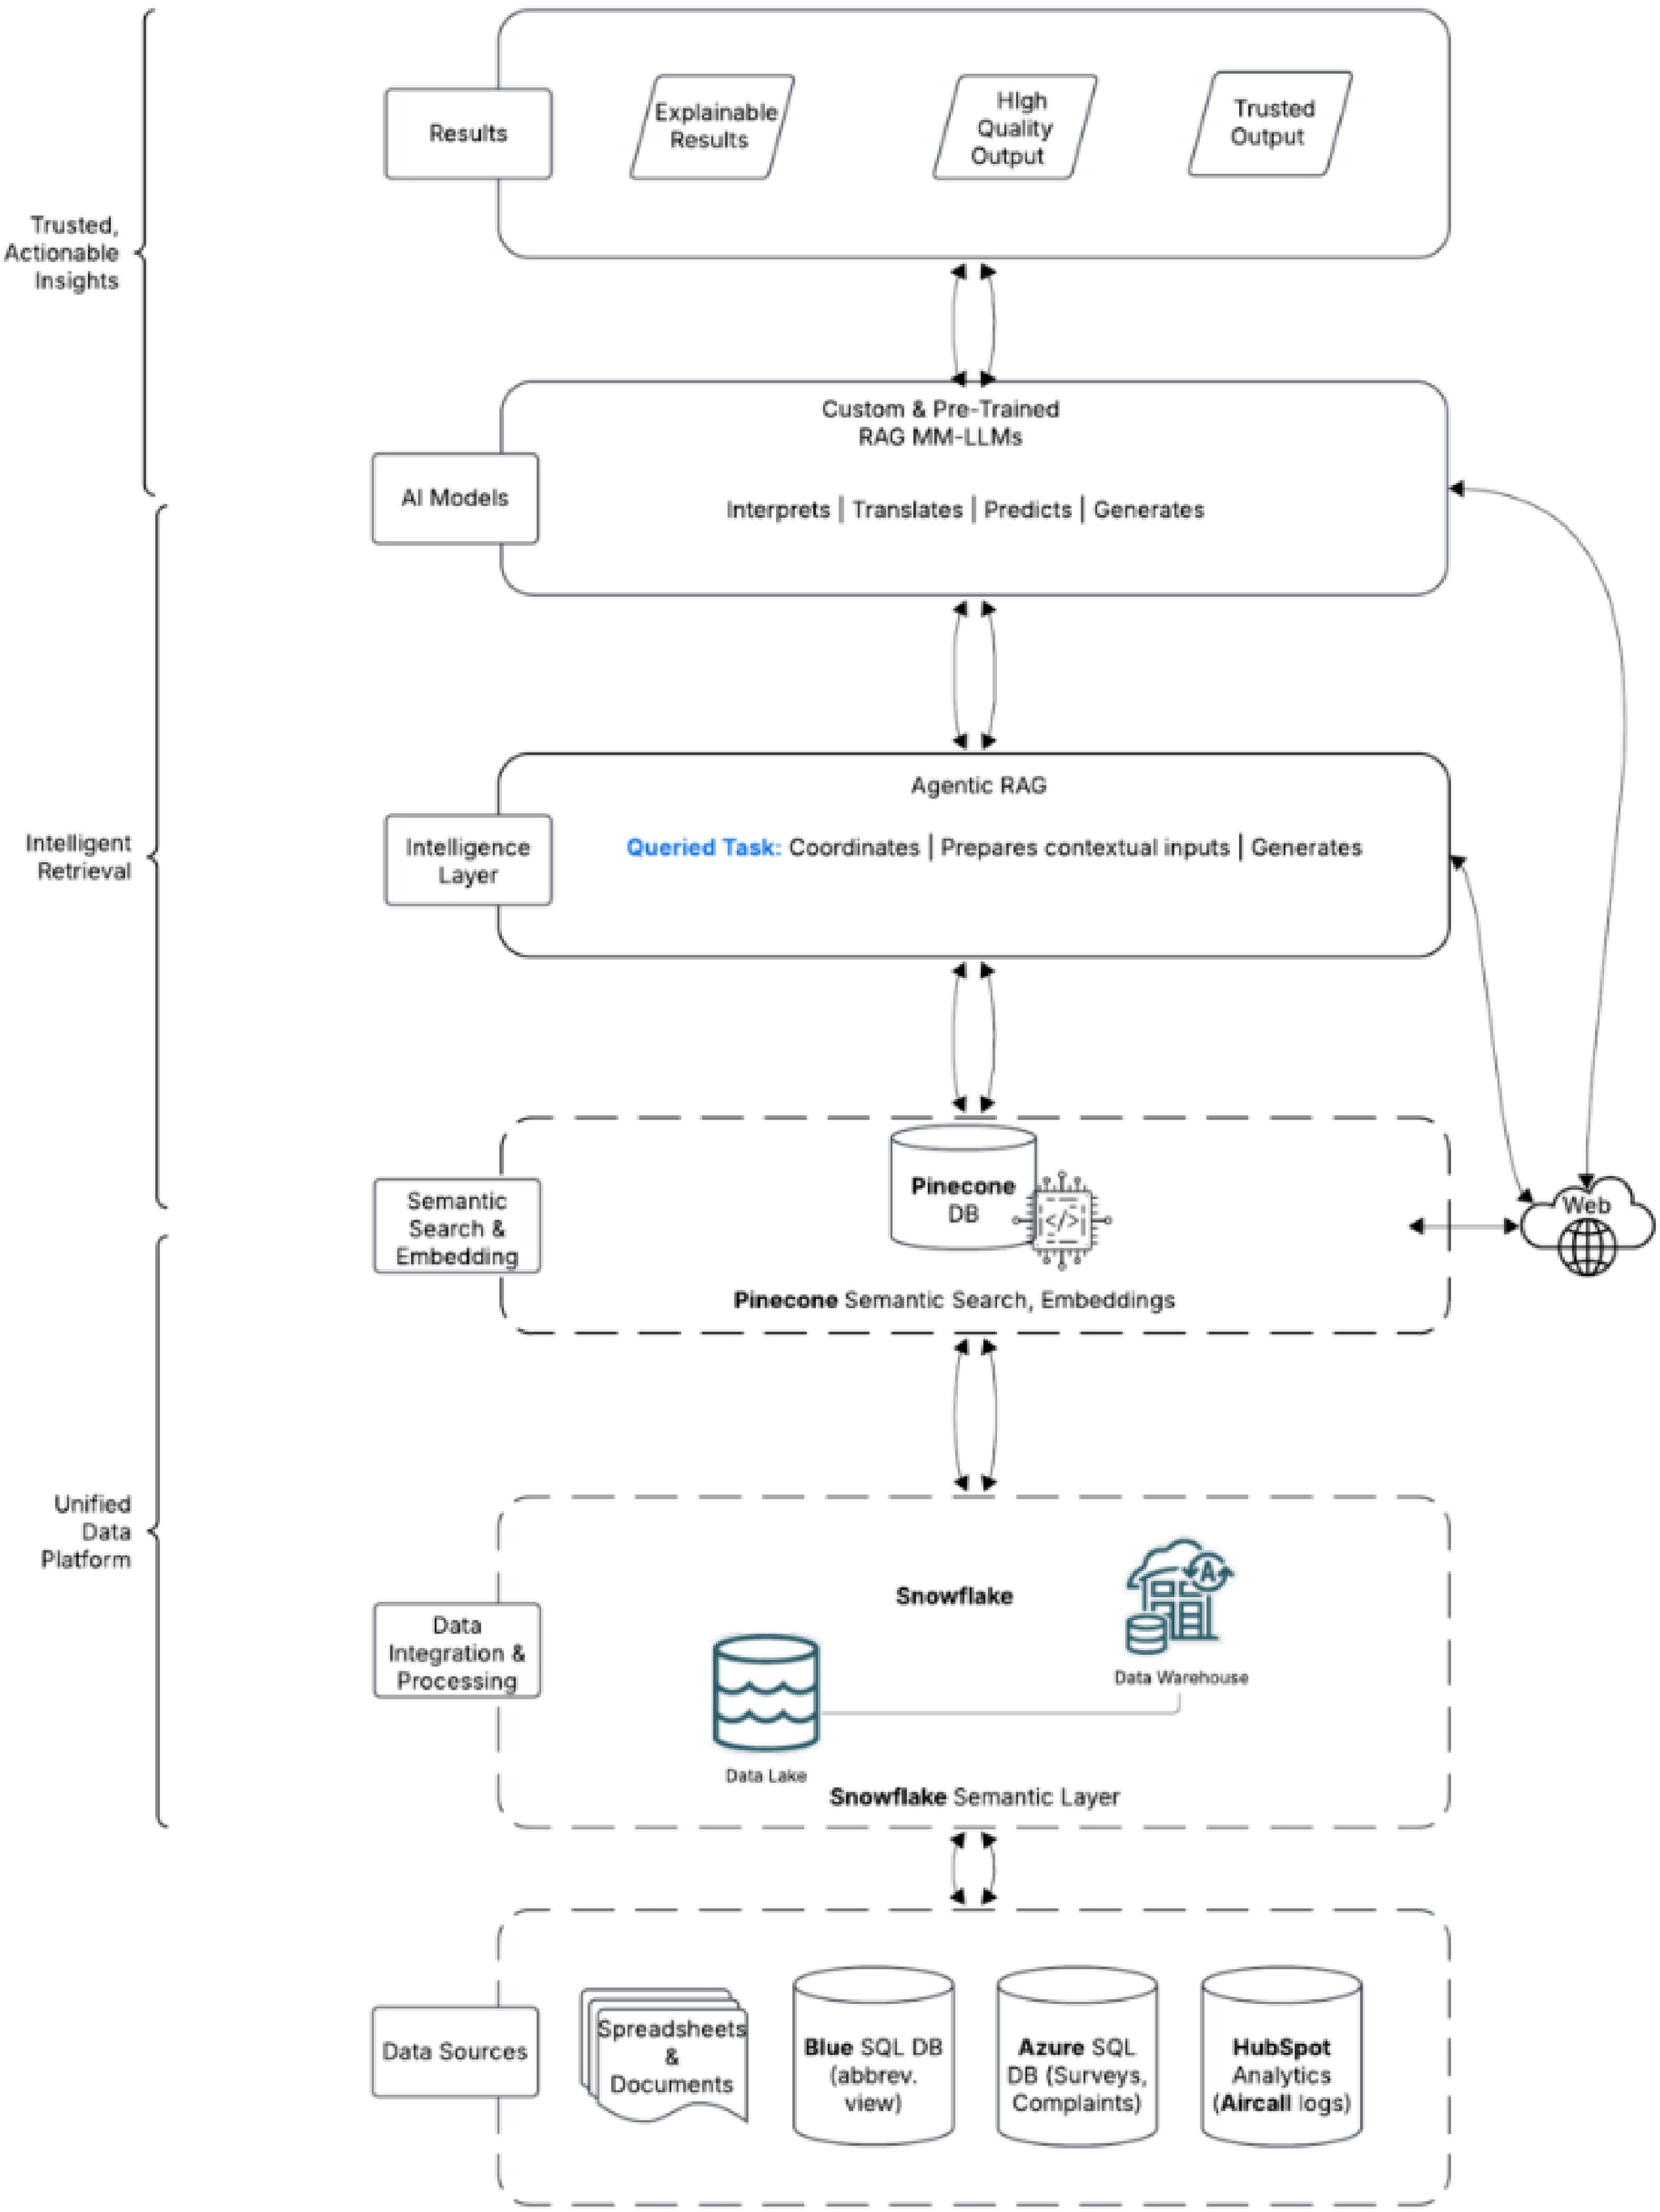

4. Target System Workflow (To-Be)

AI-Driven, Department-Specific Automation

Features

- Department-specific agents (e.g., AI Sales Agent, AI Marketing Agent, etc.)
- Orchestrator Agent coordinates workflows across roles
- IPA/RPA for back-end integrations with Blue, IBM BAW, HubSpot, etc.

Impact

- Reduced time-to-action
- Higher autonomy for departments with less manual input
- Unified experience across departments

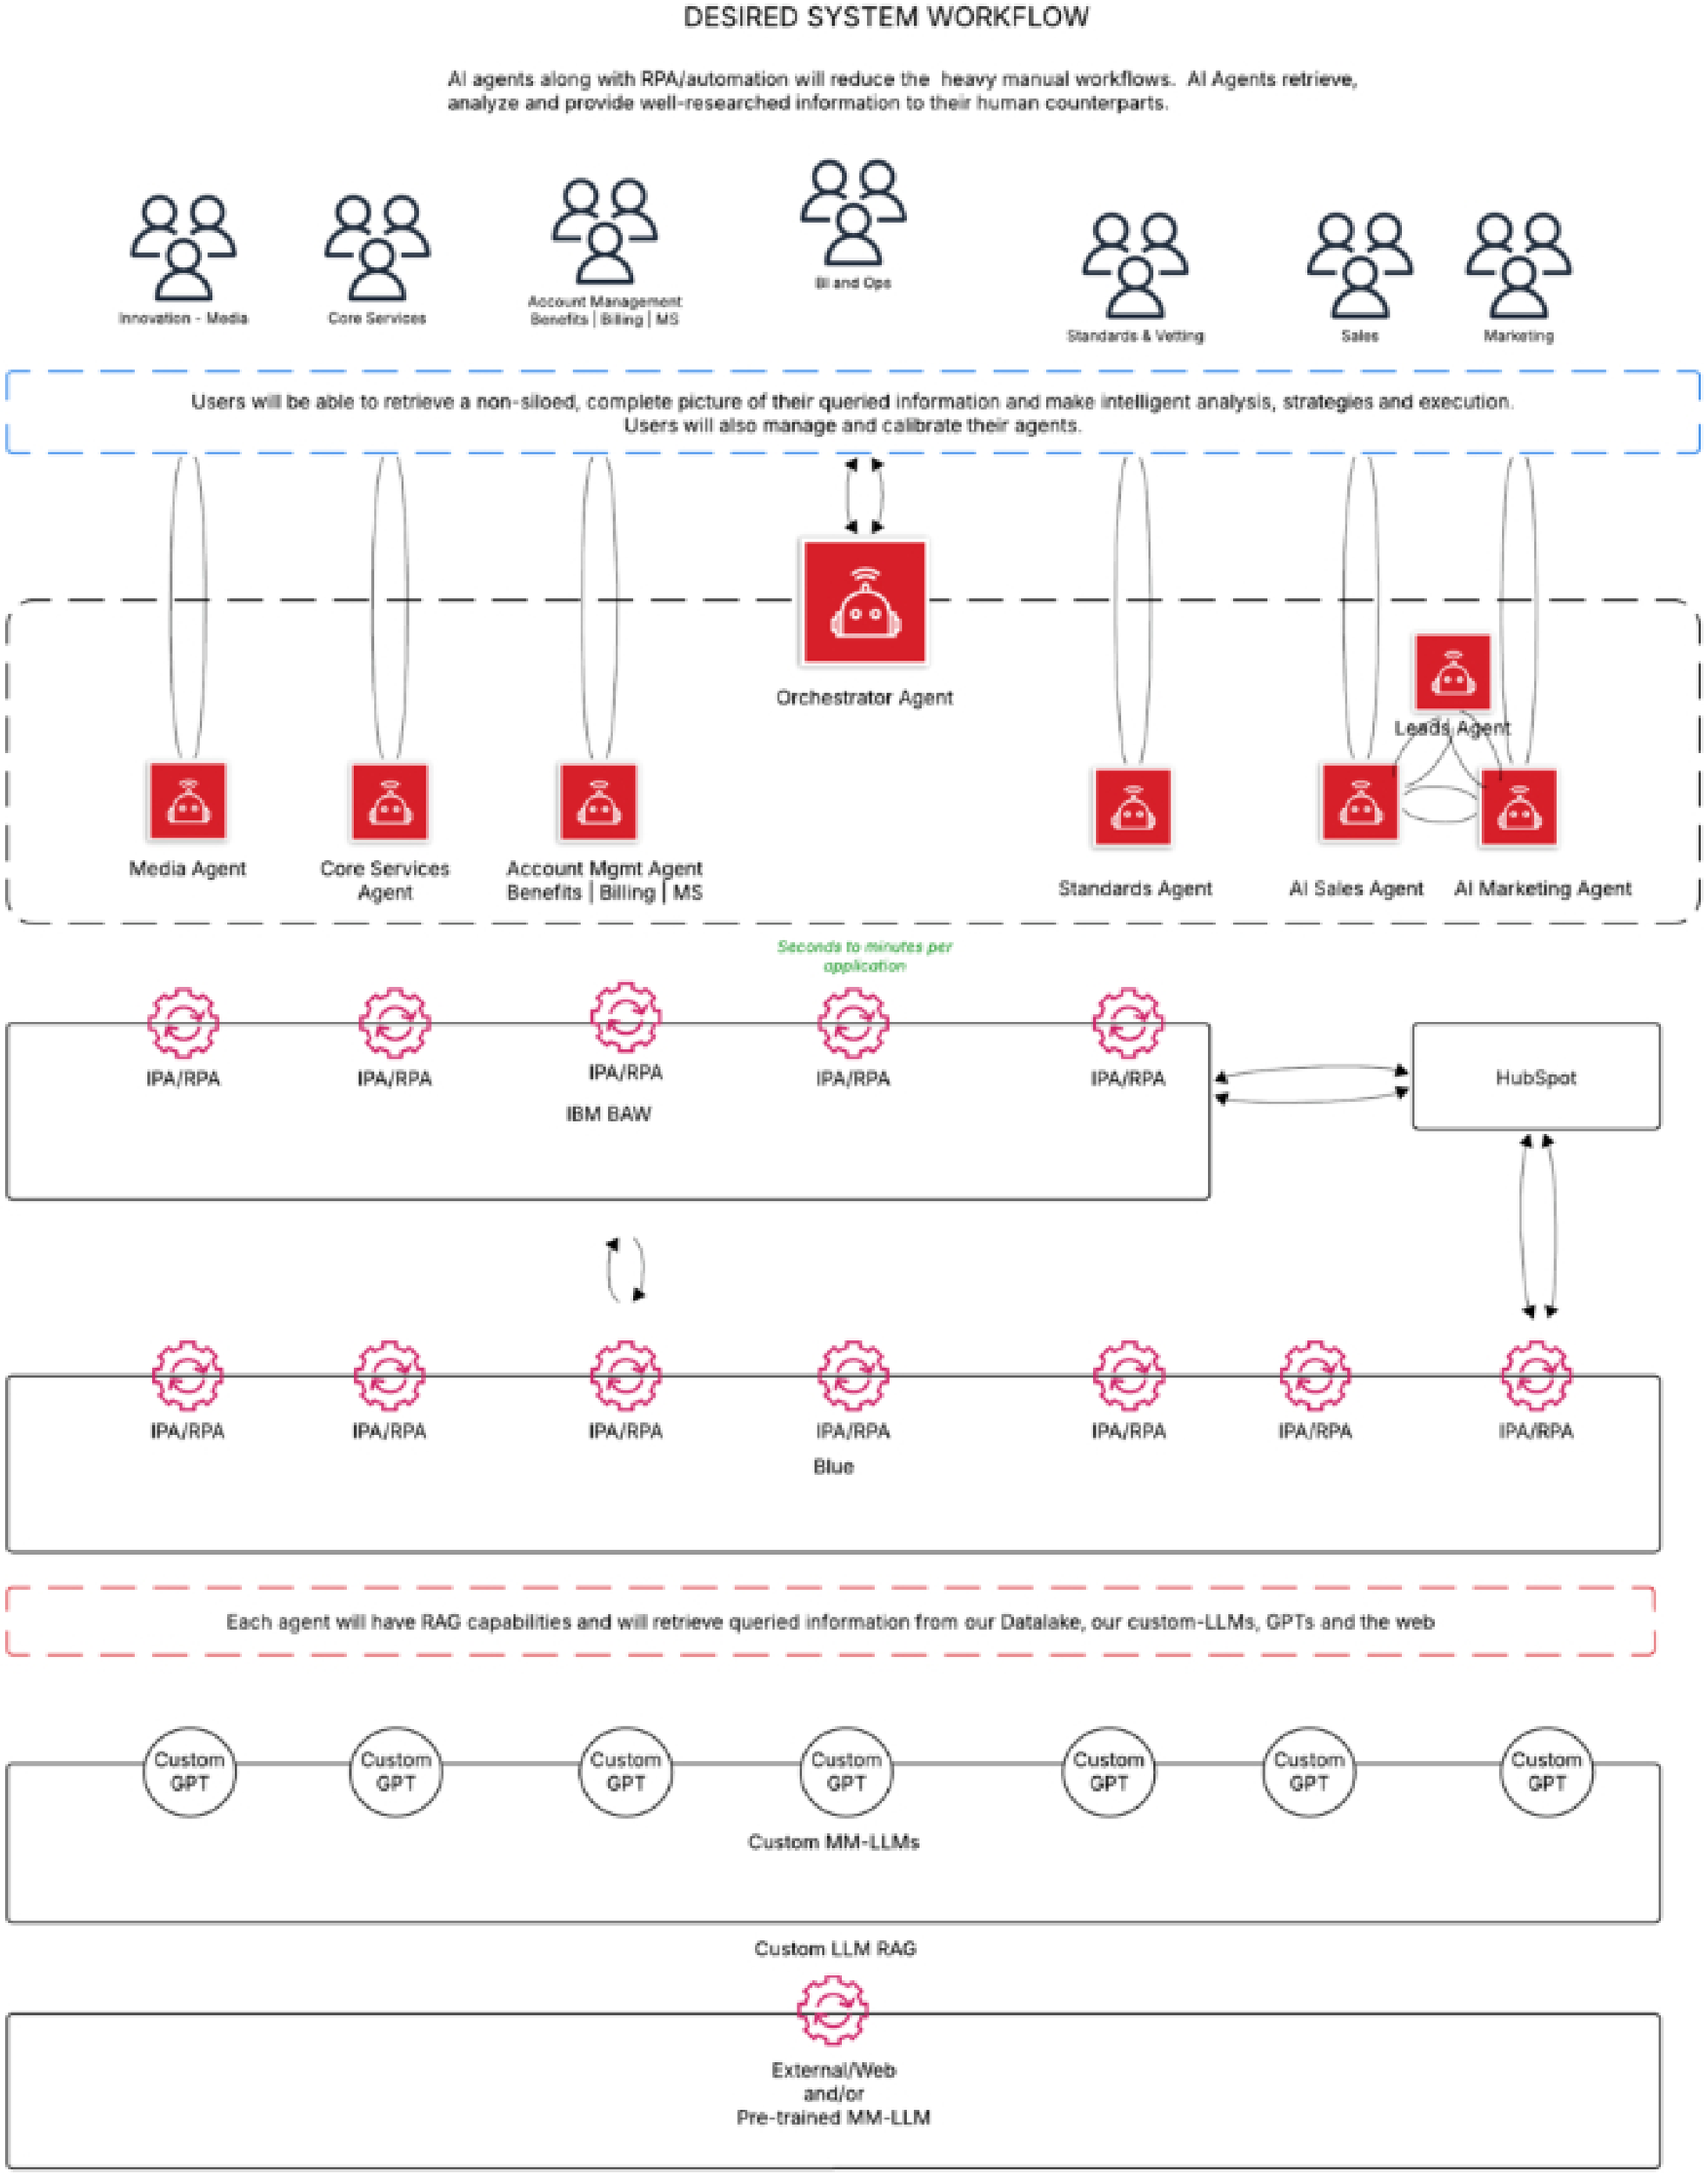

BAW-Focused:

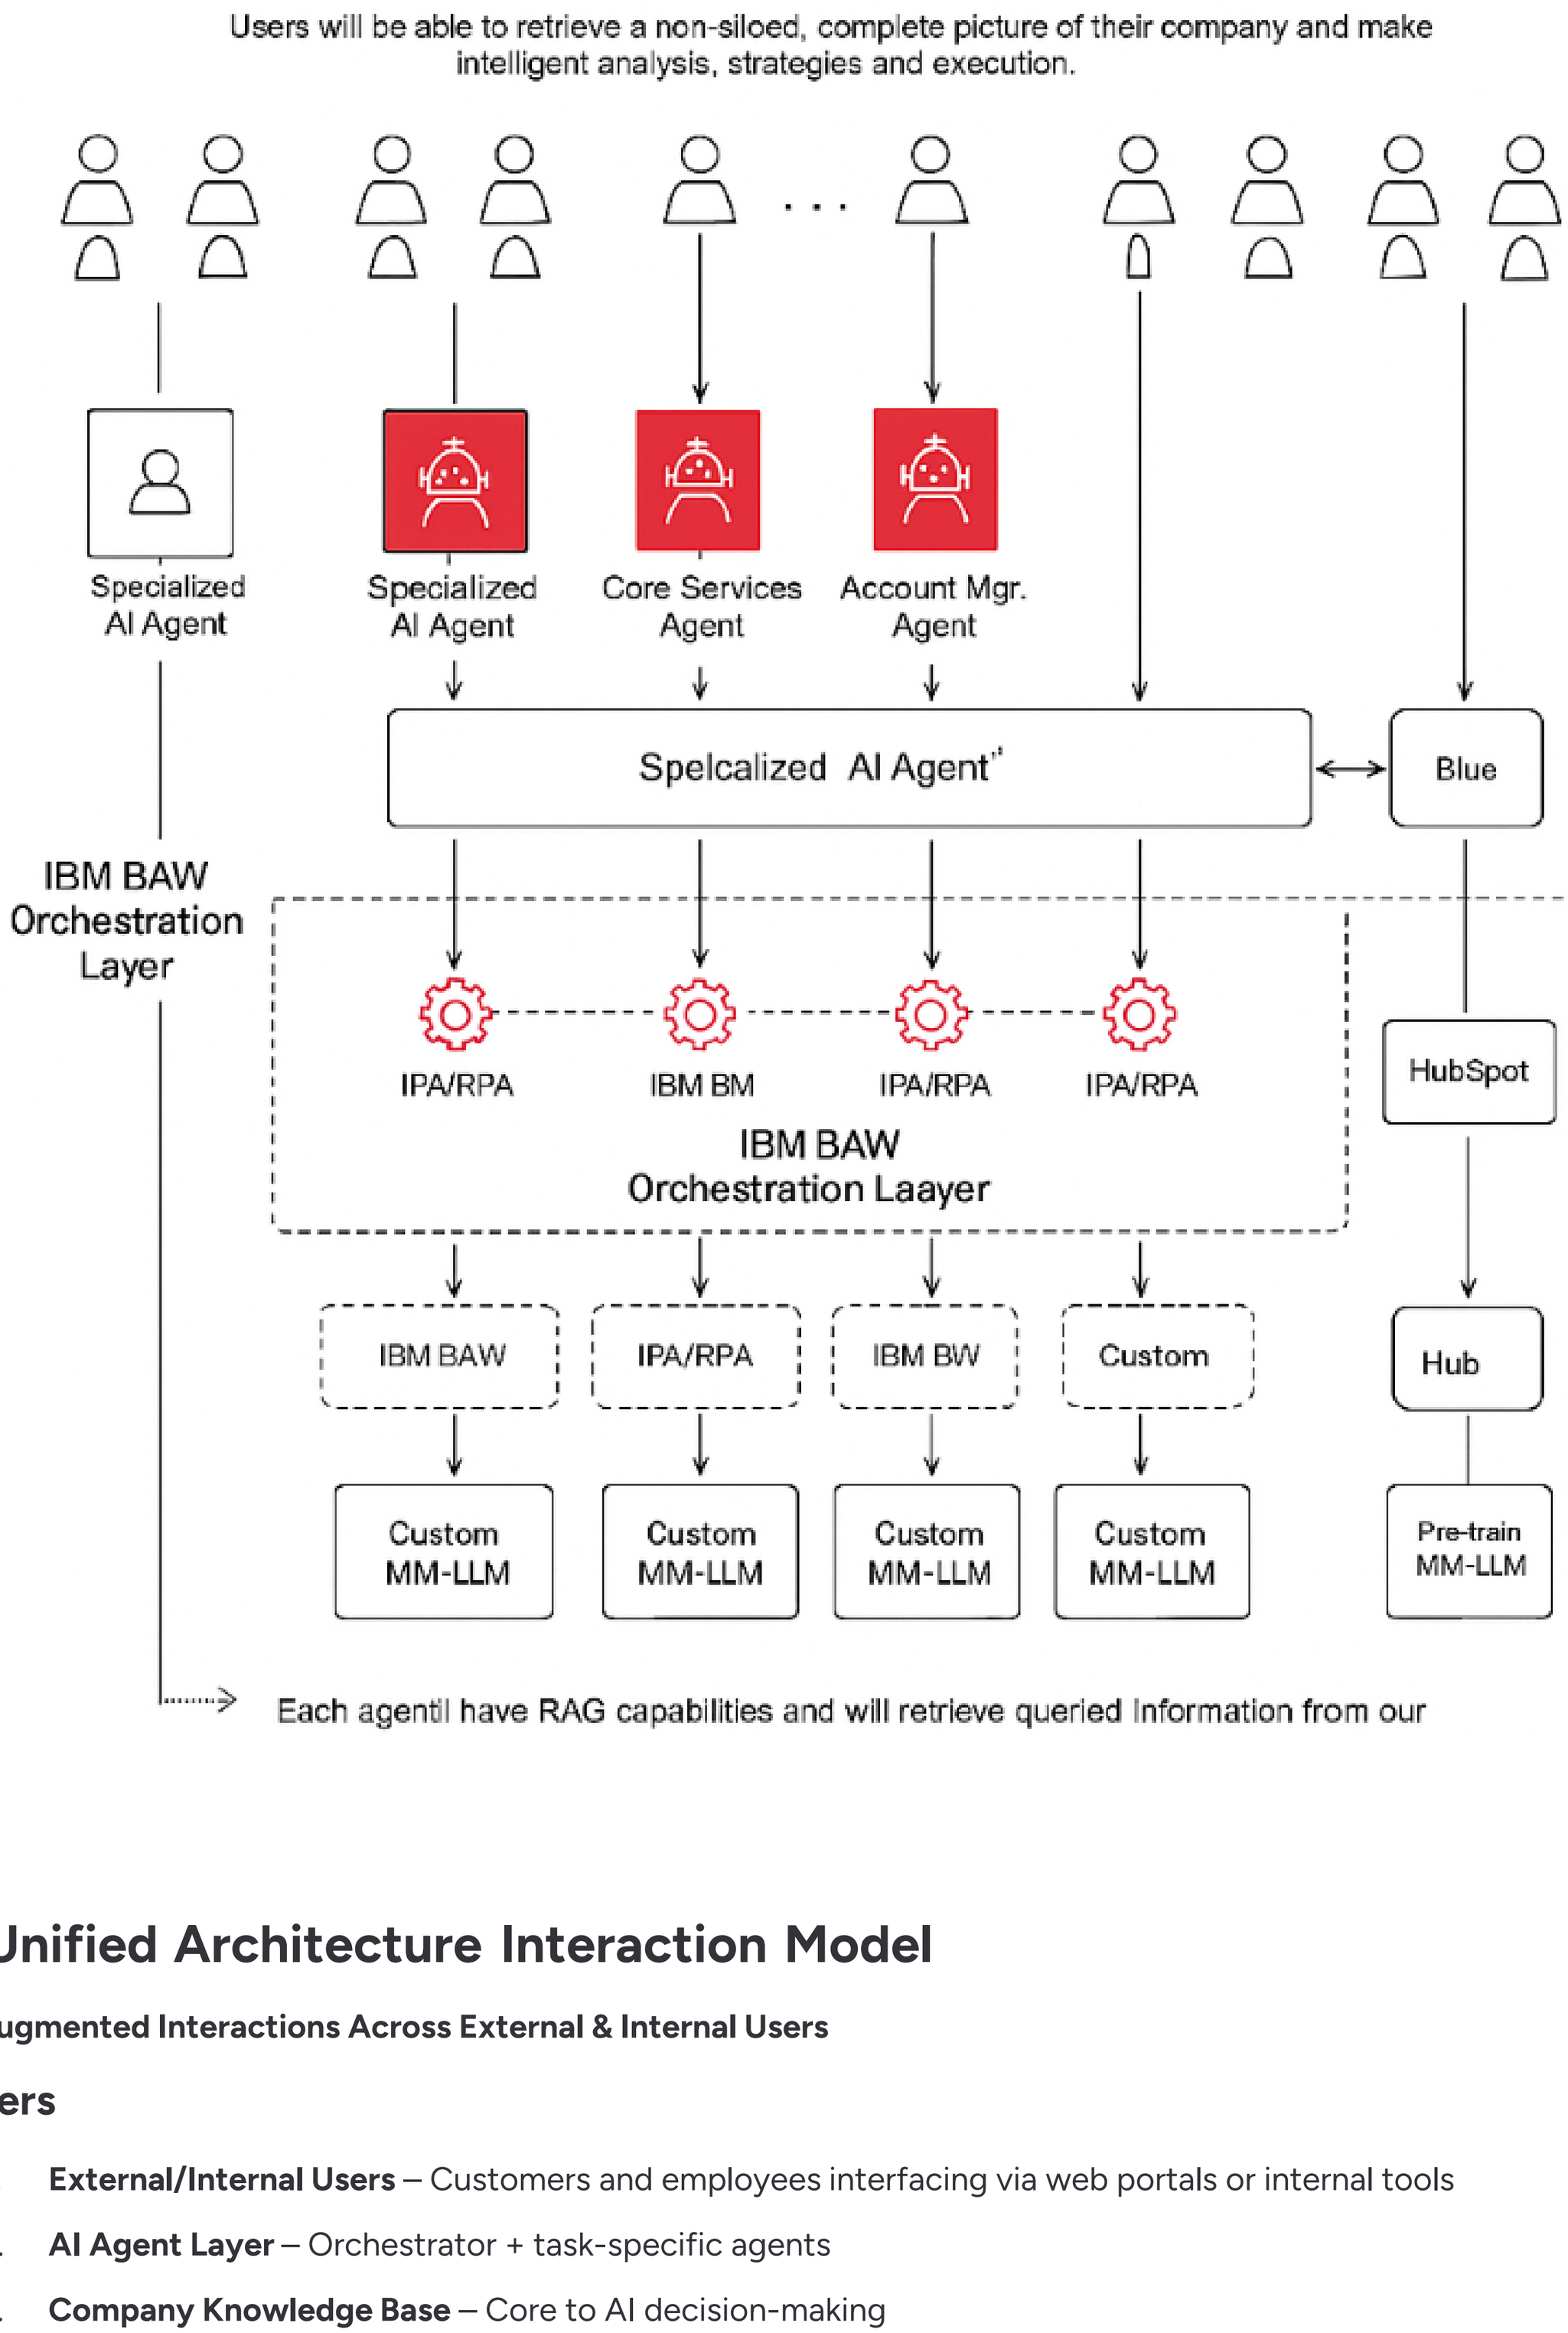

5. Unified Architecture Interaction Model

AI-Augmented Interactions Across External & Internal Users

Layers

1. **External/Internal Users** – Customers and employees interfacing via web portals or internal tools
2. **AI Agent Layer** – Orchestrator + task-specific agents
3. **Company Knowledge Base** – Core to AI decision-making

4. **Intelligence Layer** – Powers contextual query interpretation (Agentic RAG)
5. **Data Integration** – Connects semantic search (Pinecone) and structured data (Snowflake, SQL DBs)

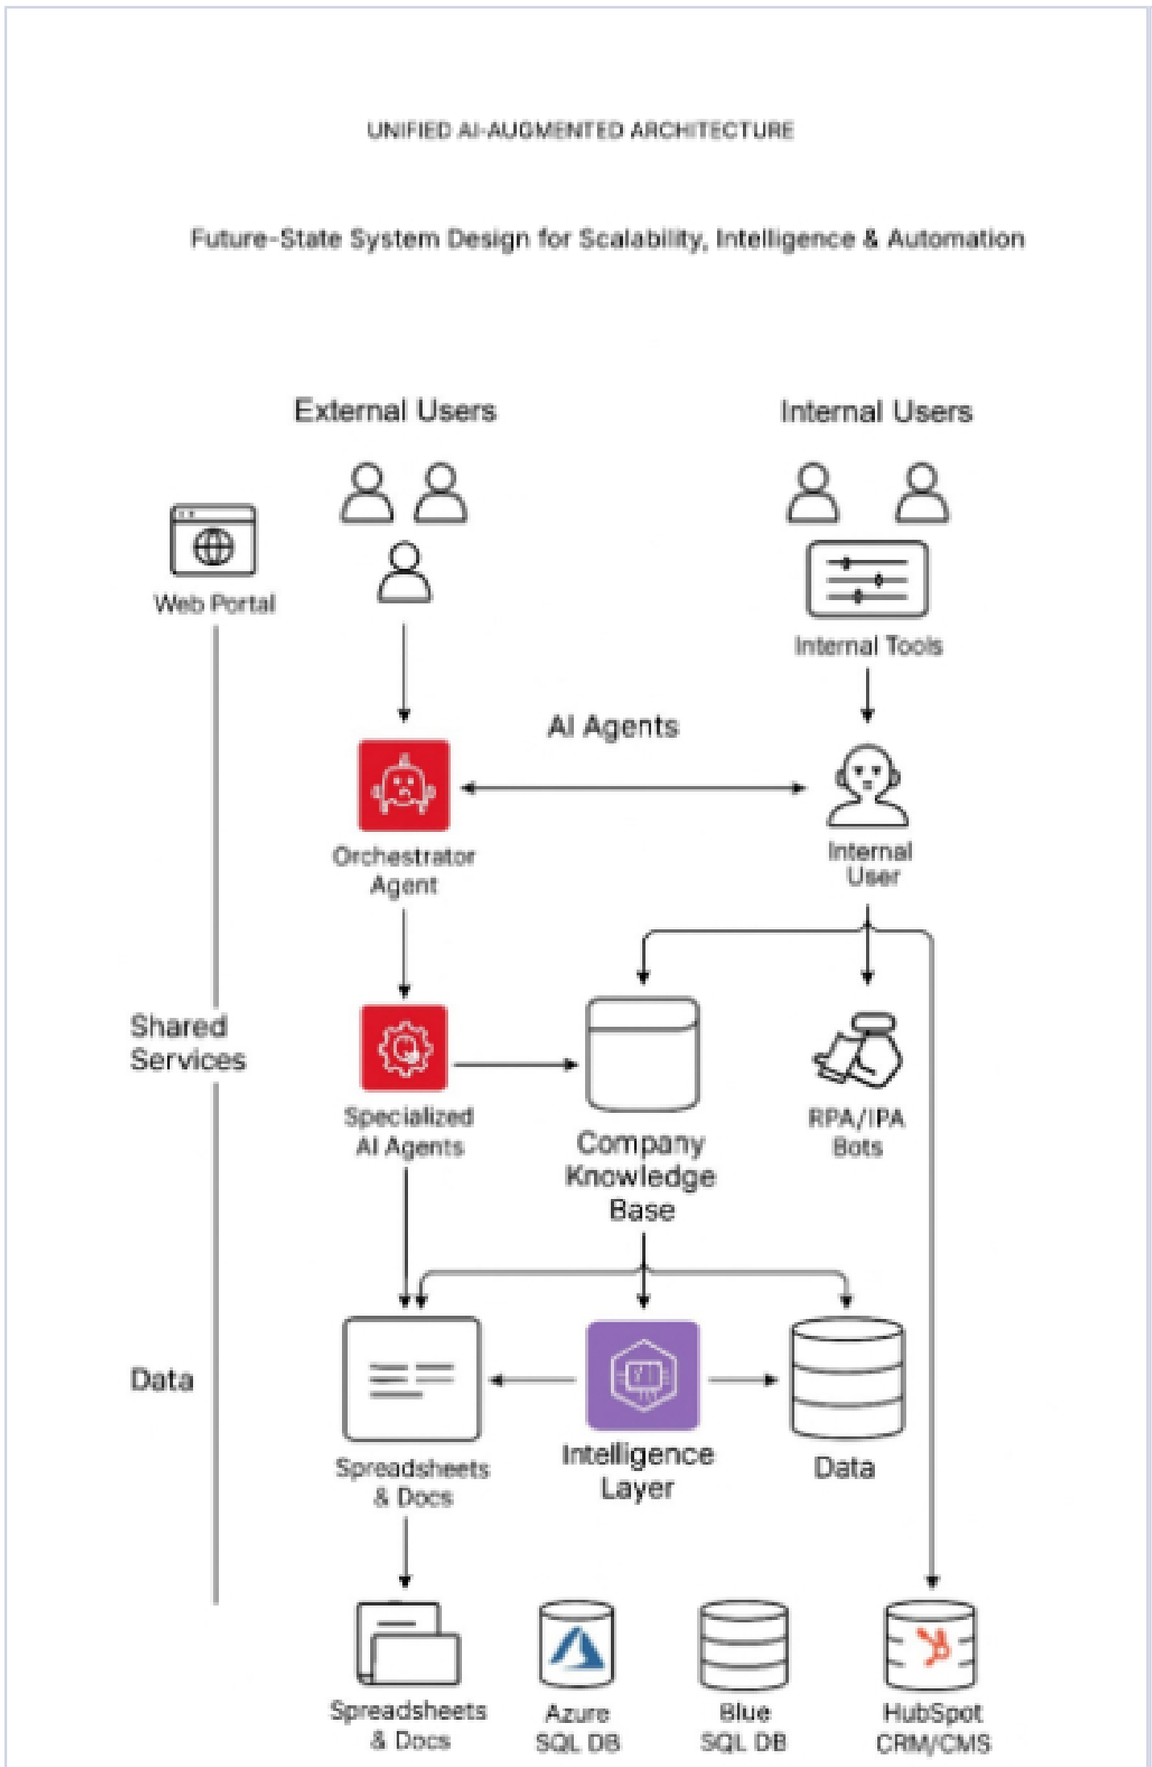

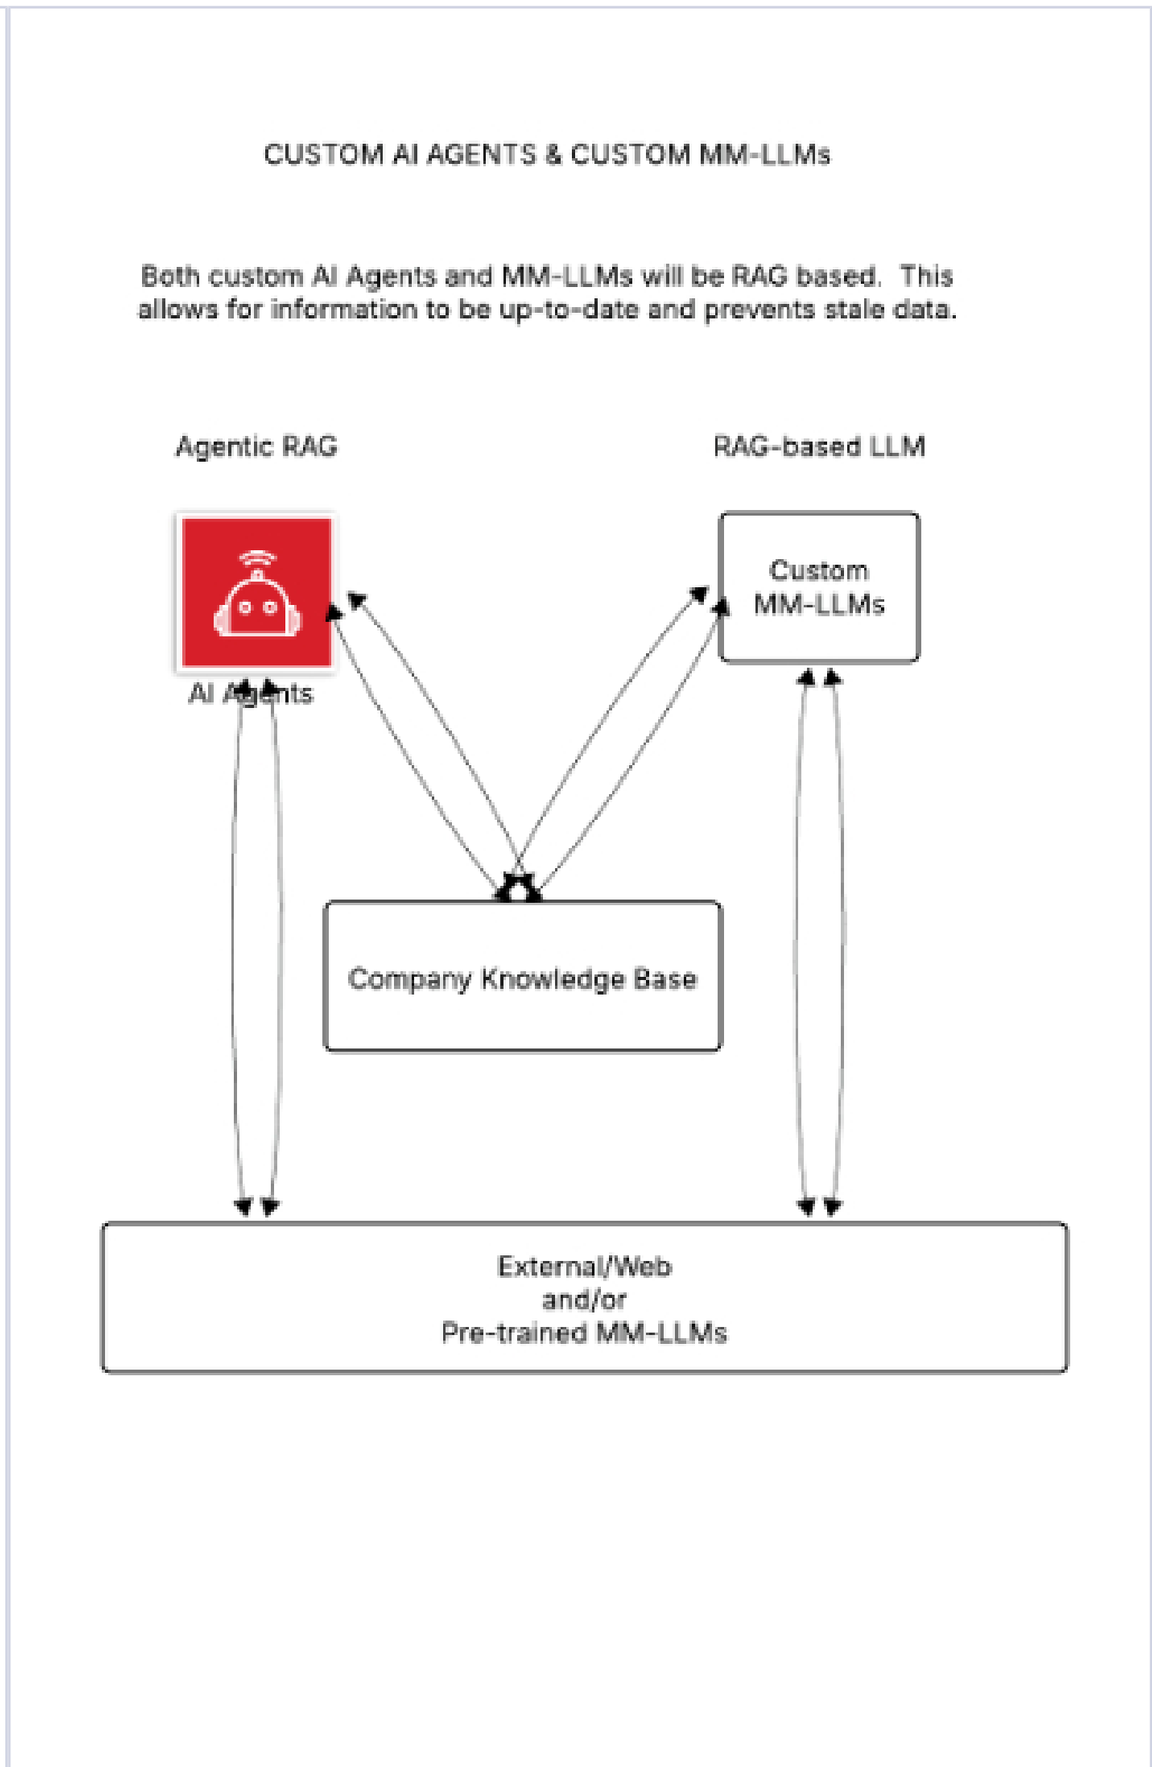

6. RAG Technical Architecture

RAG Orchestration for Real-Time Enterprise AI

Components

- **Inference Layer**: MM-LLMs, GPTs, and AI Agents
- **RAG Layer**: Retrieval, Context Augmentation, Generation
- **Vector Store**: Pinecone for fast semantic similarity
- **Data Sources**: Snowflake, CRM (HubSpot), SQL, Web, Docs

Governance Focus

- Privacy, Relevance, Explainability, and Quality Control
- KPIs: Sub-1s response times, 95% augmentation quality, 3TB/day ingestion

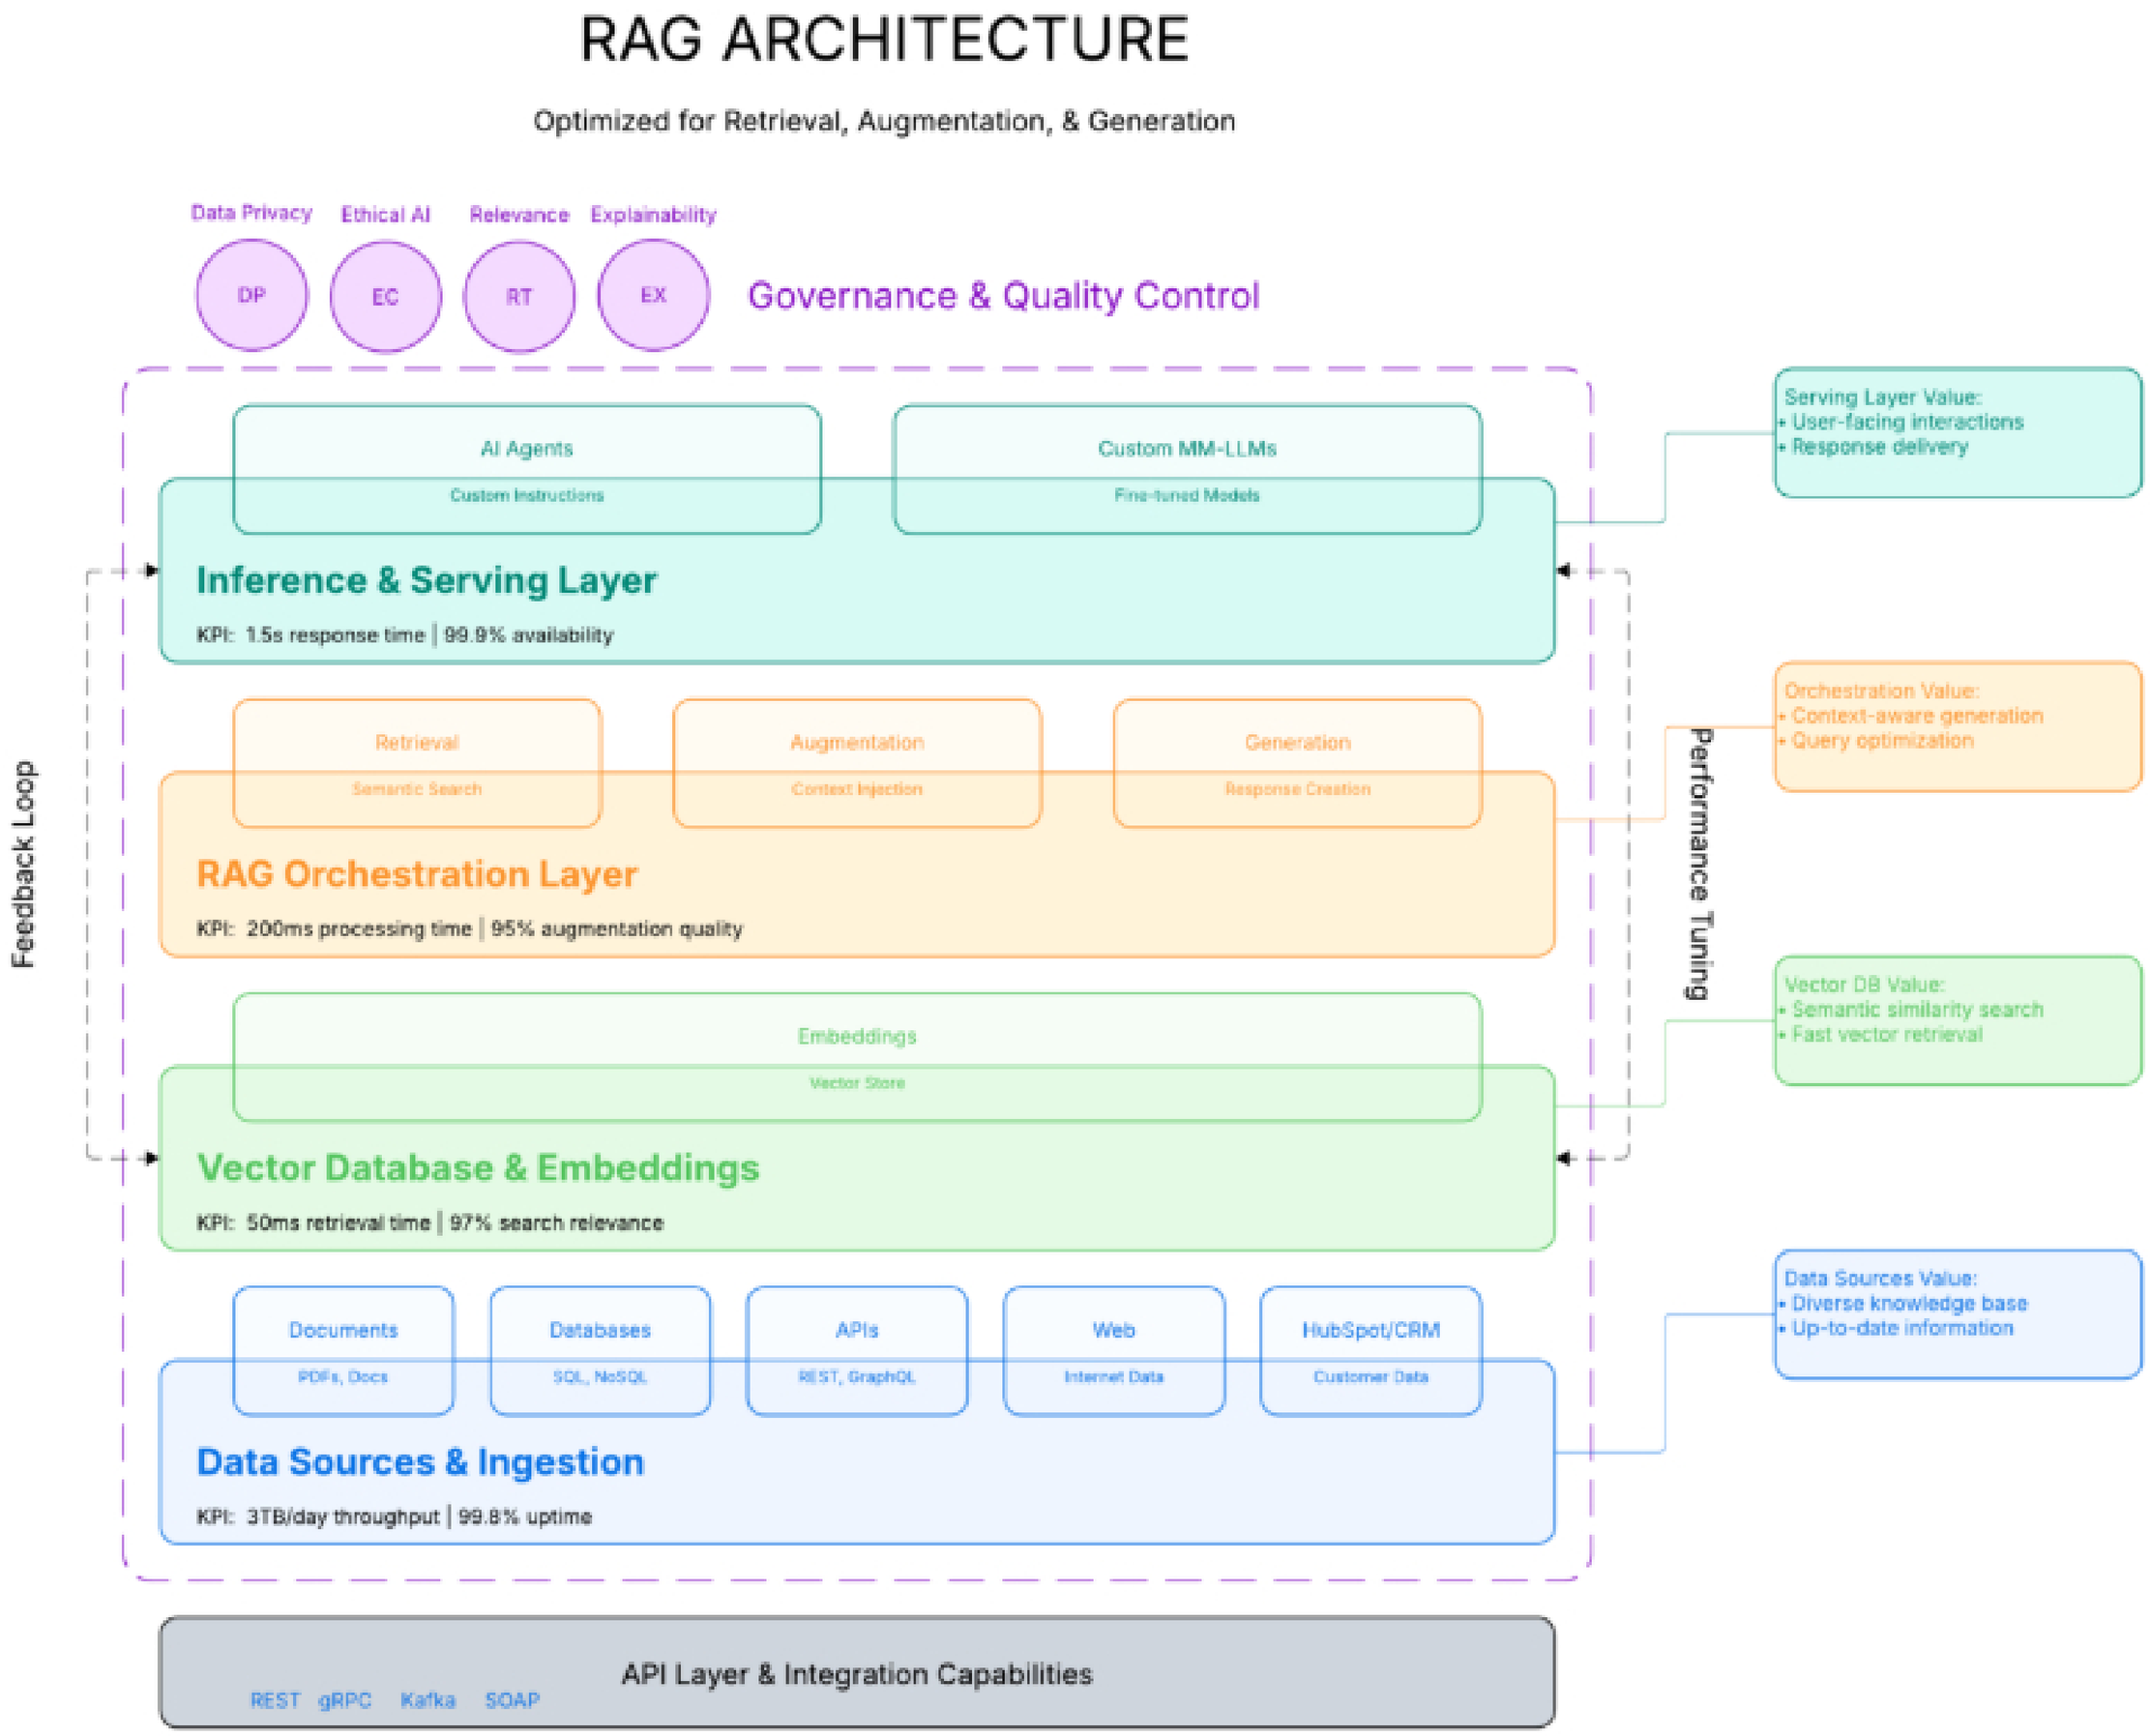

7. Summary and Next Steps

Enabling Intelligent Transformation

Summary

- Current systems are siloed and manual
- Future state enables intelligent agents, unified data, explainable AI
- RAG architecture allows scalable AI interactions using trusted data

Action Items

- Finalize unified data schemas
- Define and build ETL/real-time data pipelines
- Implement Pinecone + Snowflake stack
- Pilot AI agent(s) by function