## 1. Solution Overview

**Agent Analyst** is an enterprise-ready, AI-native solution designed to analyze data across both modern data platforms (like Databricks, Snowflake, Azure Synapse) and legacy enterprise systems (such as SAP, Oracle EBS, Dynamics 365). It uses a modular multi-agent architecture orchestrated via an intelligent Orchestrator Agent built on **Python A2A SDK** and communicates through the **A2A (Agent-to-Agent) Protocol**. The system delivers insights, recommendations, and narrative summaries directly through natural language interfaces, empowering non-technical users while ensuring data privacy, security, and Responsible AI compliance.

## 2. Problem Statement

Enterprises face significant challenges in leveraging data from legacy systems alongside modern platforms. These include:

- Fragmented data silos across ERP, CRM, and data lakes
- Lack of user-friendly tools to query or interpret data
- Manual, error-prone, and time-consuming analysis
- Inability to enforce responsible AI principles in LLM-based analysis
- No unified system to provide real-time insights and root-cause analytics

Agent Analyst bridges these gaps by providing a centralized, autonomous, and intelligent data analysis platform.

## 3. Solution Detail

Agent Analyst operates as a multi-agent intelligence layer over enterprise systems and data platforms. At its core is an **Orchestrator Agent** built using **Python A2A SDK**, which controls and manages specialized remote agents through the **A2A protocol**. These remote agents include:

1. **Data Access Agent** – Interfaces with MCP servers to securely query ERP, CRM, and lakehouse platforms.
2. **Metadata & Ontology Agent** – Maps and resolves schema differences, aligns data fields to enterprise ontology.

3. **Semantic Interpretation Agent** – Translates natural language queries into formal analytical intents.
4. **Query Planning & Execution Agent** – Designs and executes optimized data retrieval and transformation workflows.
5. **Data Analysis Agent** – Conducts statistical analysis, trends, root cause detection, forecasting, etc.
6. **Narrative Agent** – Converts analytical outputs into natural language narratives or dashboard-ready summaries.

Each agent is designed to operate independently but in coordination, with tasks delegated and sequenced by the Orchestrator Agent. The agents are deployed in a scalable containerized environment using **Azure Kubernetes Service (AKS)** and rely on supporting services such as Redis, Cosmos DB, and Azure AI Search for memory and knowledge retrieval.

The solution also embeds a full observability and evaluation layer using Azure Monitor, Prompt Flow, and Data Explorer. Prompts, outputs, errors, and performance metrics are continuously logged, evaluated, and made available to admins and developers. With integration into Microsoft's Responsible AI framework, all outputs are monitored for fairness, bias, safety, and auditability.

This architecture allows enterprise users to ask high-level analytical questions in natural language, and receive validated, explainable, and context-aware responses — whether the source data resides in SAP, Salesforce, Snowflake, or elsewhere. Agent Analyst bridges the technical-business gap, driving faster insights with governance and scale.

## 4. Technical Architecture

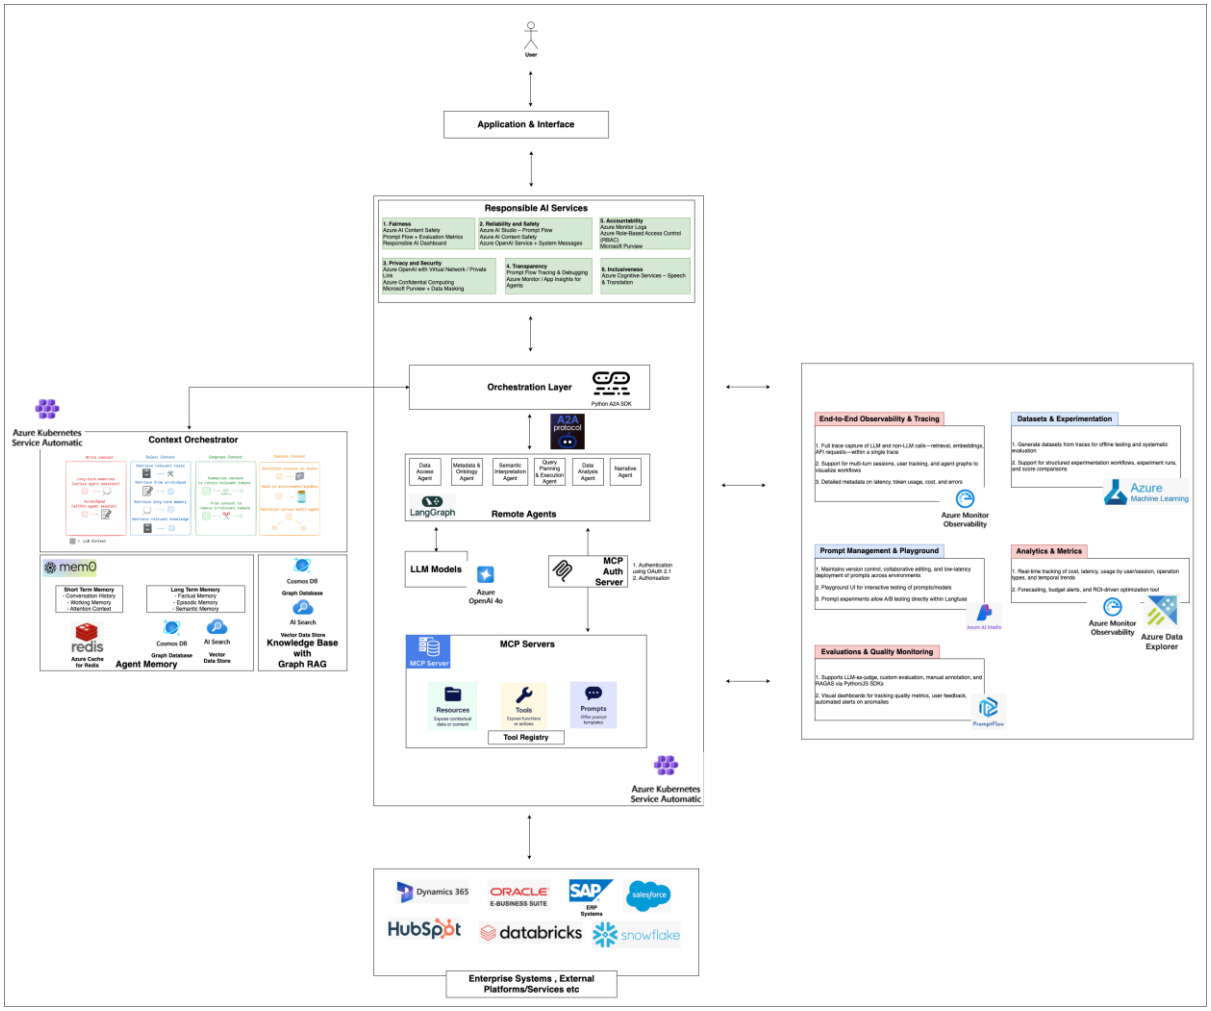


The architecture includes the following layers:

- **User Interface Layer**: Chat-based or dashboard UI for user interaction
- **Responsible AI Layer**: Services for fairness, safety, privacy, and transparency
- **Orchestration Layer**: Includes Orchestrator Agent using **Python A2A SDK** and A2A protocol for agent communication
- **Context Orchestrator**: Handles session context, memory (Redis, MemQ), and knowledge base (Cosmos DB, Azure AI Search)
- **Remote Agents**:
    - Data Access Agent
    - Metadata & Ontology Agent
    - Semantic Interpretation Agent
    - Query Planning & Execution Agent
    - Data Analysis Agent
    - Narrative Agent
- **MCP Servers**: Act as secure data and tool connectors to enterprise systems

- **Observability Stack**: Azure Monitor, Prompt Flow, Data Explorer, Application Insights
- **Deployment Layer**: Azure Kubernetes Service (AKS) for scalability and fault tolerance

## 5. Key Components

1. **Orchestrator Agent**: Built using **Python A2A SDK**, it governs the execution of all remote agents, manages routing, error handling, and state via A2A protocol.
2. **Remote Agents**:
    a. **Data Access Agent**: Interacts with MCP servers to access ERP, CRM, and data lakes.
    b. **Metadata & Ontology Agent**: Handles schema alignment, entity resolution, and business semantic mapping.
    c. **Semantic Interpretation Agent**: Converts user queries into analytical intent, grounding with metadata.
    d. **Query Planning & Execution Agent**: Generates execution workflows and submits them to MCP for fulfillment.
    e. **Data Analysis Agent**: Performs statistical, causal, and predictive analysis using structured inputs.
    f. **Narrative Agent**: Generates text-based summaries, visual explanations, and BI-compatible responses.
3. **Context Orchestrator**: Maintains conversational state, session history, and memory context.
4. **MCP Servers**: Handle secure access to external enterprise systems, abstracting integration complexity.
5. **Responsible AI Services**: Azure AI Content Safety, Prompt Flow Evaluators, Fairlearn toolkits, etc.

## 6. Integration Points

- **Enterprise Systems**: SAP, Oracle EBS, Dynamics 365, Salesforce, HubSpot
- **Data Platforms**: Azure Synapse, Databricks, Snowflake
- **Authentication/Authorization**: Azure AD, MCP Auth Server
- **Monitoring & Evaluation**: Azure Monitor, Prompt Flow, Azure Data Explorer
- **Prompt Management**: Azure AI Studio with Prompt Flow and evaluation datasets

## 7. Use Cases

- Revenue breakdown and margin analysis across geographies
- Supply chain forecasting and inventory optimization
- Customer churn detection using CRM and support logs
- Root cause analysis of operational bottlenecks
- Financial anomaly detection and audit trail generation
- Generative executive summaries for monthly performance reports

## 8. Customer Pain Points Addressed

- Siloed data across legacy and modern systems
- Slow manual analysis cycles
- Lack of transparency in LLM-driven analytics
- Over-reliance on technical teams for insights
- Difficulty ensuring compliance and audit readiness

## 9. Industry-Specific Applications

- **Manufacturing**: Vendor delays, parts shortfalls, quality issue trends
- **Retail**: Product category performance, inventory forecasting
- **Financial Services**: Budget tracking, regulatory audit insights
- **Healthcare**: Claims anomalies, service coverage analysis
- **Telecom**: Network issue resolution, service desk optimization

## 10. Sample Customer Journey

1. User asks: "Why did our North America revenue drop in Q2?"
2. Semantic Interpretation Agent identifies the query scope.
3. Metadata Agent aligns "revenue" and "Q2" with ERP data schema.
4. Query Planning Agent creates a retrieval plan via MCP.
5. Data Access Agent pulls data from SAP and Snowflake.
6. Data Analysis Agent finds root cause (e.g., vendor delays).

7. Narrative Agent responds: "Q2 NA revenue dropped due to a 3-week delay from Vendor ABC, impacting SKU performance."
8. User follows up with: "What's the forecast for Q3 based on current trend?"

## 11. Technical Requirements

- Azure subscription with access to OpenAI and Prompt Flow
- AKS cluster to deploy agents
- Azure Cache for Redis for Agent memory layer
- Cosmos DB and Azure AI Search for KB access
- MCP connectors or APIs for system integration
- Azure Monitor and Data Explorer setup

## 12. Security Architecture

- **MCP Auth Server**: Centralized authorization and role-based control
- **Private Networking**: VNet and Private Link usage for secure access
- **Data Governance**: Azure Purview for lineage and access policy
- **Confidential Execution**: Optional Confidential Computing for PII-sensitive workloads
- **Audit Logs**: Full traceability via Azure Monitor and App Insights

## 13. Performance Considerations

- Modular remote agents allow parallel execution
- Memory-first design reduces repeated retrieval latency
- Load-aware orchestration for scalability
- Caching layer to speed up repeated business questions
- Real-time tracing to detect and mitigate slow data sources

## 14. Tools and Azure Services Used

- Azure OpenAI
- Prompt Flow & Azure AI Studio

- Azure Monitor, Data Explorer, App Insights
- Azure Purview, Confidential Computing, Private Link
- Azure Kubernetes Service
- Cosmos DB, Redis, Azure AI Search

## 15. Users of Agent

- Enterprise Analysts
- Financial Controllers
- Sales and Operations Managers
- CX and Support Leads
- Data Engineers and Architects (admin roles)

## 16. Dependencies

- MCP platform for data connectivity
- Pre-built agent prompts and planning logic
- Integration with internal and external identity providers
- Azure cloud infrastructure with observability stack configured

## 17. Key Benefits and Differentiators

- Specialized modular agents with defined roles
- Fast and explainable insights over legacy and modern systems
- First-class Responsible AI support with evaluations and safety controls
- Natural language interface tailored for business roles
- Plug-and-play deployment model for Azure-native environments

## 18. Value Proposition

**Agent Analyst transforms how enterprise users access insights from complex systems.** It allows decision-makers to interact with data through natural language, get reliable and explainable analysis, and close the gap between raw data and business

outcomes. The agent framework is modular, secure, observable, and scalable, offering future-ready intelligence for modern enterprises.

## 19. Conclusion

Agent Analyst is a transformative enterprise analytics solution designed for organizations that rely on legacy ERP systems, modern cloud data stacks, and diverse operational tools. By combining intelligent agent-based orchestration, secure connectivity via MCP, and Azure's Responsible AI ecosystem, it empowers enterprises to democratize insights and accelerate decisions.