Solution Overview

Agent SRE is an AI-powered observability and incident management platform built on Microsoft Azure. It leverages a LangGraph-based multi-agent architecture to autonomously monitor, diagnose, and resolve infrastructure issues across hybrid and multi-cloud environments. Integrated deeply with Azure Kubernetes Service (AKS), Azure Monitor, Azure OpenAI Service, and Azure Event Grid, it delivers predictive monitoring, automated remediation, and intelligent decision-making to ensure operational continuity, SLA compliance, and reduced human intervention.

Problem Statement

Modern enterprises face increasing operational complexity due to:

- Distributed systems and multi-cloud deployments
- Alert fatigue and noisy monitoring systems
- Fragmented visibility across infrastructure stacks
- Manual, error-prone incident response processes
- Inability to scale observability securely and compliantly

Traditional monitoring tools lack the intelligence and automation capabilities to manage today's performance, scalability, and reliability needs.

Solution Detail

Agent SRE delivers proactive and autonomous reliability engineering via Azure-native capabilities:

Core Features

- Autonomous Incident Detection & Remediation: Multi-agent system analyzes telemetry in real-time and executes automated remediations.
- Automated On-Demand Code and Configuration Generation: Automatically generate and deploy code or configurations as required, enabling seamless updates and supporting continuous deployment—all without manual intervention.
- AI-Powered Prompt-Driven Workflow Management: Utilize AI-generated prompts to orchestrate workflows on demand, accelerating task execution and enhancing overall operational efficiency.
- Proactive, AI-Driven Configuration Management: AI agents automatically manage and update configurations across environments, ensuring compliance and eliminating the need for manual intervention.

Technical Architecture

Agent SRE is designed with a modular, microservices architecture deployed on Azure Kubernetes Service (AKS), orchestrated using Azure-native serverless components.

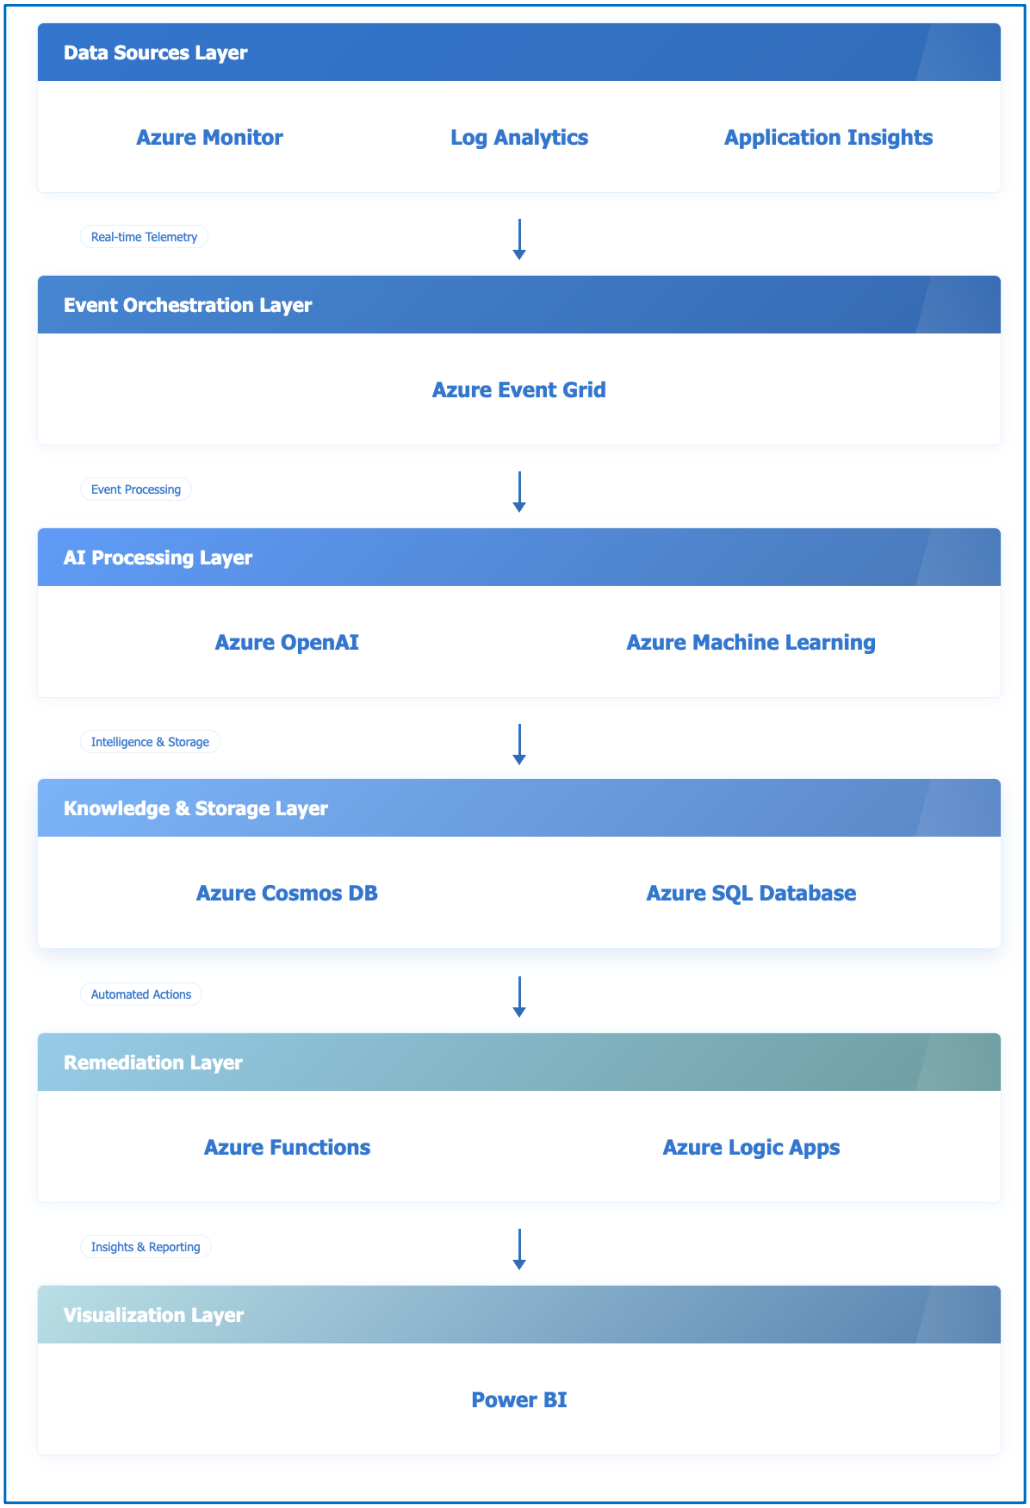

Architecture Highlights:

- Ingests real-time telemetry from Azure Monitor, Log Analytics, and Application Insights
- Uses Azure OpenAI and Azure Machine Learning for anomaly detection and agent reasoning
- Stores incident history and knowledge graphs in Azure Cosmos DB and Azure SQL Database
- Remediation actions via Azure Functions and Logic Apps
- Event orchestration with Azure Event Grid
- Unified dashboards built with Power BI

Key Components

Component	Azure Service
Multi-Agent Framework	Azure OpenAI + LangGraph on AKS
Telemetry Ingestion	Azure Monitor, Log Analytics, App Insights
Knowledge Graph & Runbooks	Azure Cosmos DB + Azure Cognitive Search
Remediation Engine	Azure Functions + Logic Apps
Workflow Orchestration	Azure Event Grid + Azure Automation
Observability 3.0 and Dashboard	Power BI
Secure Credential Handling	Azure Key Vault
Ticketing & Collaboration	Microsoft Teams, ServiceNow

Integration Points

- Monitoring: Azure Monitor, Prometheus, App Insights
- CI/CD: Azure DevOps, GitHub Actions
- Collaboration: Microsoft Teams
- Ticketing: ServiceNow, Jira
- Notification: Azure Notification Hubs, PagerDuty

Use Cases

- Proactive Incident Prevention
- Autonomous Resolution
- Alert Correlation
- Cross-Cloud Observability
- Compliance Enforcement

Customer Pain Points Addressed

- MTTR reduced by 85%
- Alert fatigue reduced by 92%
- Downtime costs reduced by $1.8M annually
- Compliance prep time cut from weeks to days

Industry-Specific Applications

- E-Commerce

- FinTech
- Healthcare
- Telecom

Sample Customer Journey

Scenario	Before Agent SRE	After Agent SRE
Incident Resolution	Manual triage; 3.5-hour MTTR	Predictive scale-up avoids incident entirely
Alert Volume	5 redundant alerts for one issue	One contextual alert with root cause and fix applied
Collaboration	Manual Slack/Teams update	Auto-notification and ticket update in Teams

Technical Requirements

- Azure Kubernetes Service (AKS)
- Azure Monitor & Log Analytics enabled
- Azure Functions and Event Grid access
- Integration with Teams, ServiceNow, etc.
- Historical telemetry and runbooks for model training

Security Architecture

- Zero Trust Architecture using Azure AD and Entra ID
- Encryption at Rest and In Transit using Azure Key Vault and TLS 1.2+
- Private Networking via VNet, NSGs, and Azure Firewall
- Credential Management using Azure Key Vault
- Audit Trails using Microsoft Defender for Cloud and Log Analytics

Compliance Features

- SOC 2, ISO 27001, PCI-DSS
- Tokenized access to sensitive data
- SSO and MFA with Azure Entra ID

Performance Considerations

- Inference agents: D-series or NCasT4_v3 VMs
- AKS nodes: Autoscaling enabled
- Pre-index runbooks
- Stream telemetry ingestion
- Partition ML pipelines

Tools and Azure Services Used

- Infrastructure: AKS, Azure Bicep, ARM Templates
- AI/ML: Azure OpenAI, Azure ML, LangChain
- Observability: Azure Monitor, App Insights, Log Analytics
- Automation: Azure Functions, Logic Apps, Event Grid
- Security: Key Vault, Entra ID, Defender for Cloud
- Data: Azure Cosmos DB, Azure SQL, Azure Data Lake
- Dashboards: Power BI

Users of Agent SRE

- Site Reliability Engineers (SREs)
- DevOps and Platform Engineers
- Cloud Infrastructure Teams
- IT Compliance & Governance Analysts
- Engineering and Operations Leadership

Dependencies

- Access to telemetry and historical incident data
- Azure DevOps or GitHub for CI/CD integration
- Documented runbooks
- Integrated ticketing and notification systems

Key Benefits and Differentiators

- 85% reduction in MTTR
- 92% drop in alert fatigue
- Predictive remediation of 78% major incidents
- Azure-native + multi-cloud ready
- Human-in-the-loop agent governance
- Scalable, serverless architecture
- Continuous learning and feedback loops

Value Proposition

Agent SRE brings the power of autonomous AI operations to Azure environments. It transforms traditional monitoring into proactive, intelligent observability—ensuring maximum uptime, reducing operational burden, and enabling secure, scalable incident response at enterprise scale.

Conclusion

Agent SRE is the future of observability and reliability engineering on Azure. Powered by multi-agent intelligence and the Azure ecosystem, it moves beyond dashboards and alerting—delivering autonomous operations with unparalleled efficiency, resilience, and compliance readiness.