MCP Dev Summit Mumbai

Operationalizing MCP:

Security, Control planes and Risk Governance

AGENDA

- Why MCP changes enterprise risk
- Registry at the core of control plane
- Identity as the building block
- Capability graphs
- Risk tiering and lethal trifecta
- MCP proxy pattern
- MCP Risk ⟶ Enterprise Architecture Mapping

SPEAKER

Sagar Dashora

AI/ML Lead, JPMorgan Chase

AAIF WG member -

- Security & Privacy
- Identity & Trust
- Governance and Compliance

WHY MCP CHANGES ENTERPRISE RISK

- Identity
- Trust
- Authorization
- Data Protection
- Governance
- Auditability

- Traditional API Risk Model: Application ⟶ Known APIs
- MCP Risk Model: Agent ⟶ Discovers Tools ⟶ Invokes Tools ⟶ Combines Results ⟶ Makes Decisions
- Key risks: dynamic discovery, agent autonomy, cross-system execution, emergent attack paths
- Risk surface is now a graph of agents, tools, identities, and runtime decisions.

MCP THREAT MODEL

- Token Mismanagement & Secret Exposure
- Privilege Escalation via Scope Creep
- Tool Poisoning
- Software Supply Chain Attacks & Dependency Tampering
- Command Injection & Execution
- Intent Flow Subversion
- Insufficient Authentication & Authorization
- Lack of Audit and Telemetry
- Shadow MCP Servers
- Context Injection & Over-Sharing

REGISTRY CENTRIC CONTROL PLANE

The directory isn't just a catalog — it's the building block of control plane that enforces governance, identity, and security across the MCP ecosystem.

- Prevents *rogue MCPs* and *shadow deployments*
- Ensures *ownership verification* and *security classification*
- Enables *policy enforcement* and *identity management*
- Integrates with enterprise CMDB for auditability and governance

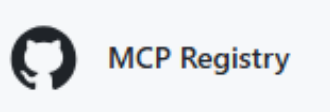

EVOLUTION OF AN ENTERPRISE REGISTRY

The maturity of an MCP platform is determined by how high up the stack governance decisions are made.

What exists

Discovery and Registration

- Inventory
- Metadata
- Discovery API
- Versioning

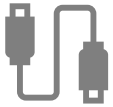

Who can connect

Connectivity Control

- Identity / authentication
- Traffic policies
- Proxy and Gateway
- Secure communications
- Service mesh (Istio, Linkerd)

What is allowed

Governance and Policy

- Policy Engine
- Risk Tiering
- Capability Graph
- Agent Scopes

IDENTITY MANAGEMENT

Service Identity

MCP trusts the platform

Individual agents do not have separate identity

Simplest operational model

Difficult to enforce per-agent authorization

Agent Identity

Every agent authenticates independently

MCP authorizes individual agents

Fine-grained access control

Requires pre-defined logical agent identities

Workload Identity

Identity tied to runtime execution

Implemented through mTLS or workload certificates

Strong cryptographic trust

More operational complexity

IDENTITY MANAGEMENT – DELEGATED IDENTITY

Token Forwarding

Simple implementation

No token exchange required

Anti-pattern in most cases

OBO Token Exchange

Supported by major IdPs

Fine-grained scopes and consent

Strong auditability

Token propagation across multiple hops can become complex

DELEGATED IDENTITY WITH CONTROL PLANE BOUND SCOPES

If we could have an overlap of human identity and agent identity

- Every actor requests OBO(on-behalf-of) tokens
- ID provider generates token with claims of original id and actor.
- Tokens with overlapping scopes from both actor and id
- ID provider and registry planes are synced

✓ Preserves full accountability chain

✓ Better auditability than token forwarding

- Requires changes in IDP level
- Need sync between IDP and control plane

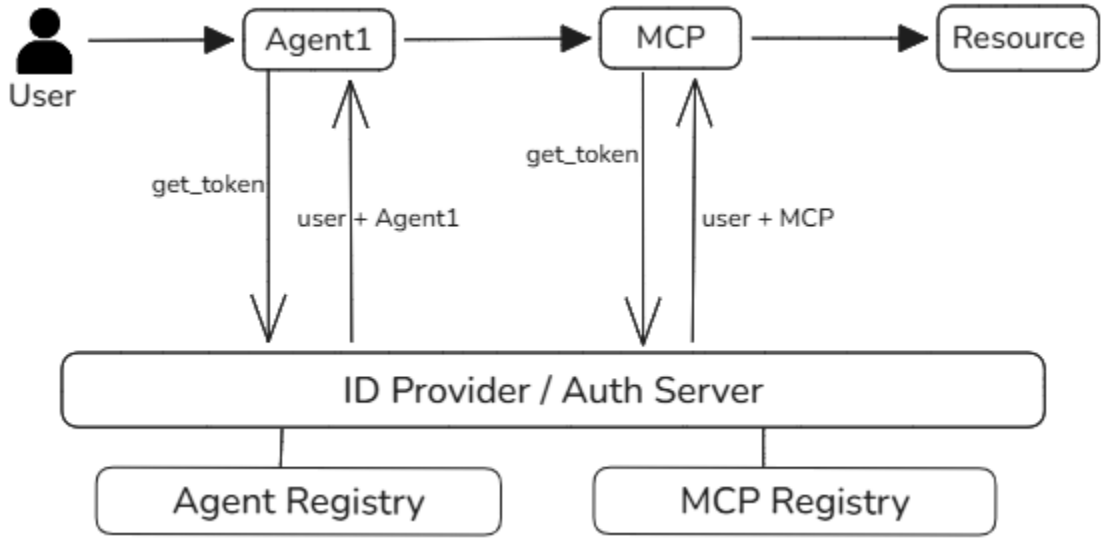

CAPABILITY GRAPHS

An agent's security boundary is determined by the tools and capabilities available within its universe.

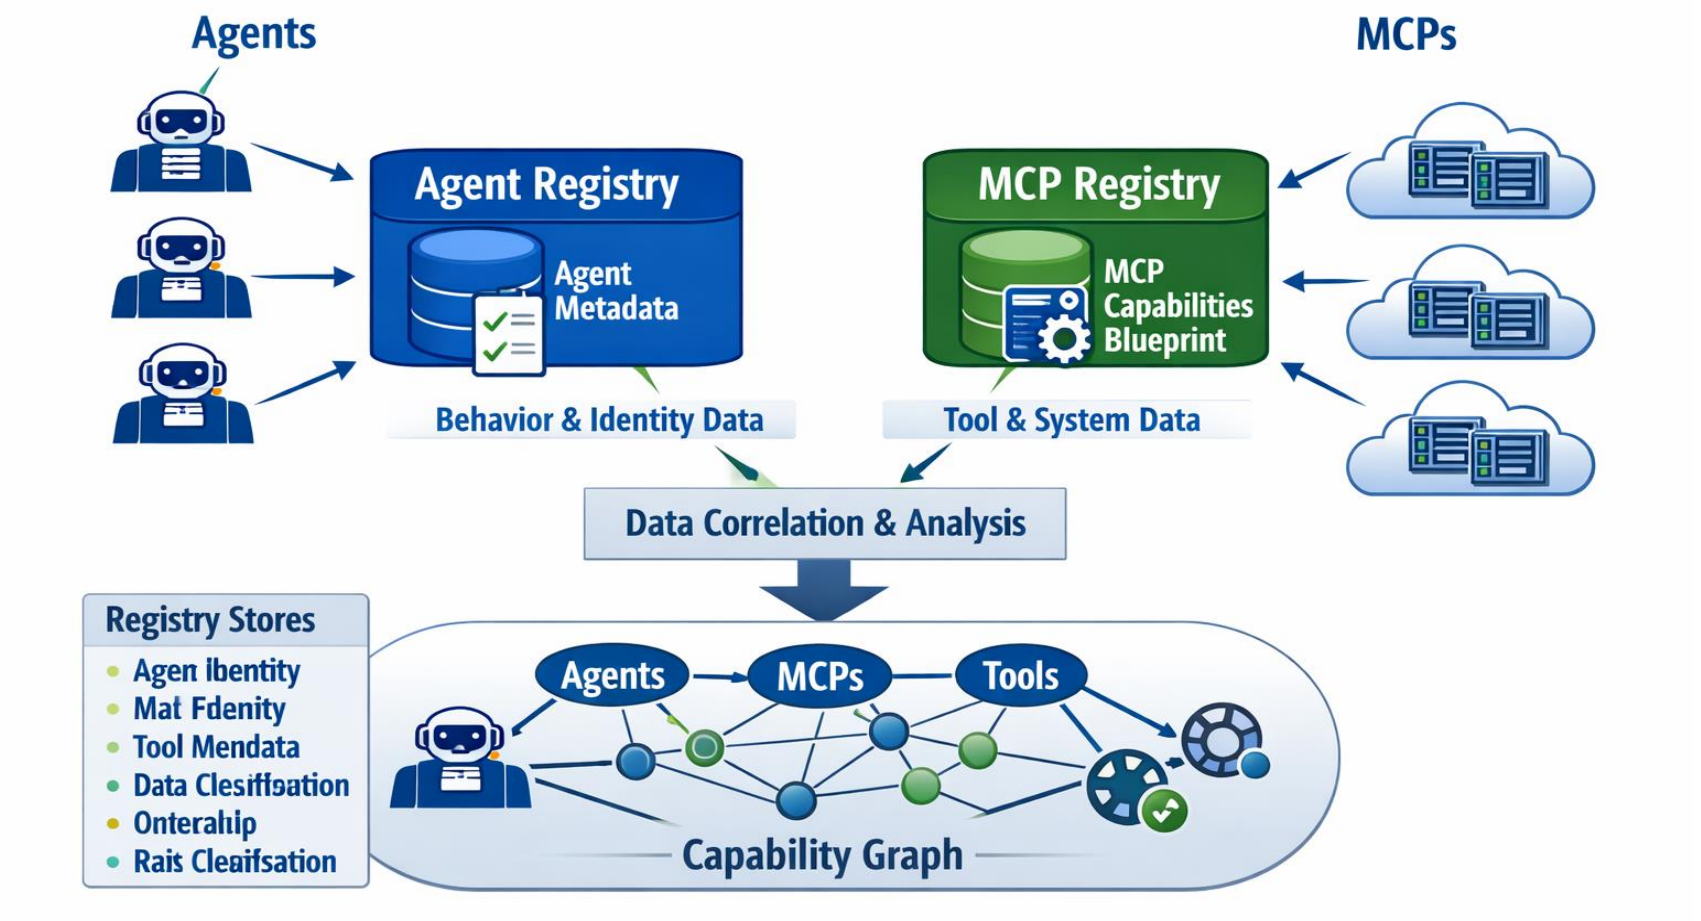

AGENT RISK TIERING AND LETHAL TRIFECTA

Agent risk is a non-binary decision, capability graphs come to rescue

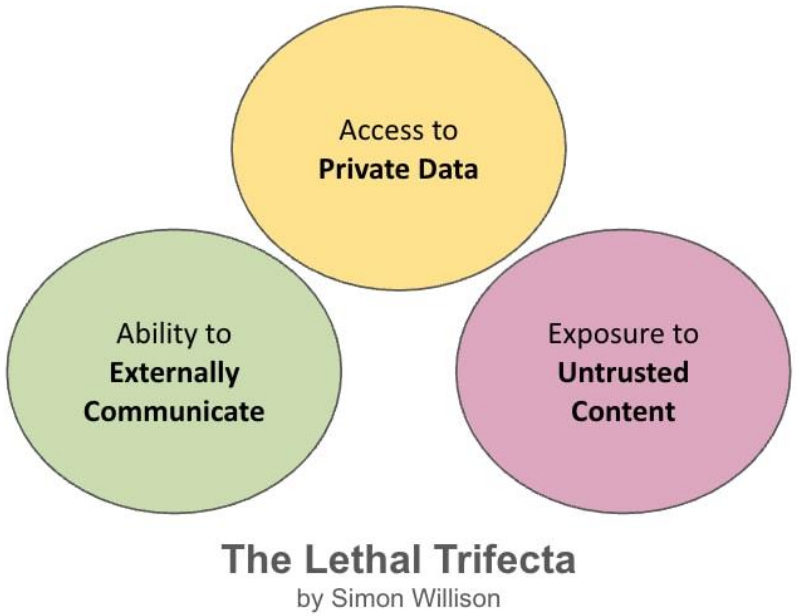

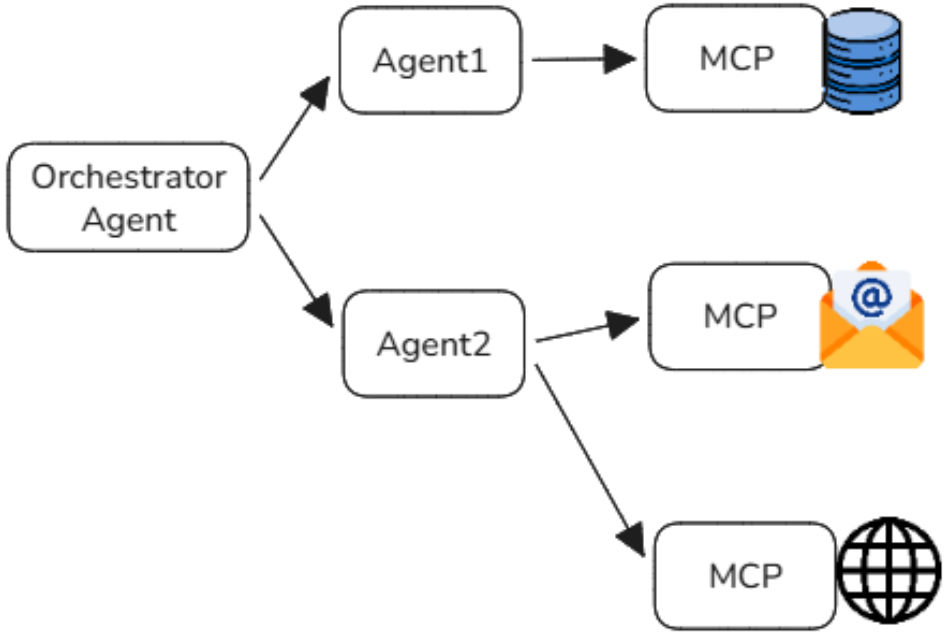

IDENTIFY AND MITIGATE THE RISK

Agent risk is a non-binary decision, capability graphs come to rescue

Tier	Examples	Risk
Tier 0	Documentation Q&A	Low
Tier 1	JIRA and confluence Agent	Medium
Tier 2	Internal Workflow agent	High
Tier 3	External facing Client Support agent	Trifecta

Architecture Controls

- Identity-aware authorization
- Capability graph enforcement
- External MCP gateway
- Human approval workflows
- Risk-based execution policies

GATEWAY FOR EXTERNAL MCP

Enterprise Registry-Controlled Access to MCPs outside the enterprise

Identity Translation
- OBO token exchange
- Credential management
- Identity mapping

Policy Enforcement
- Tool allowlists
- Risk-aware authorization

Audit and Control
- E2E telemetry
- Usage analytics

Cost Governance
- Token consumption
- API usage

MCP PROXY GATEWAY

FASTMCP · Gravitee · MCPProxy · truefoundry

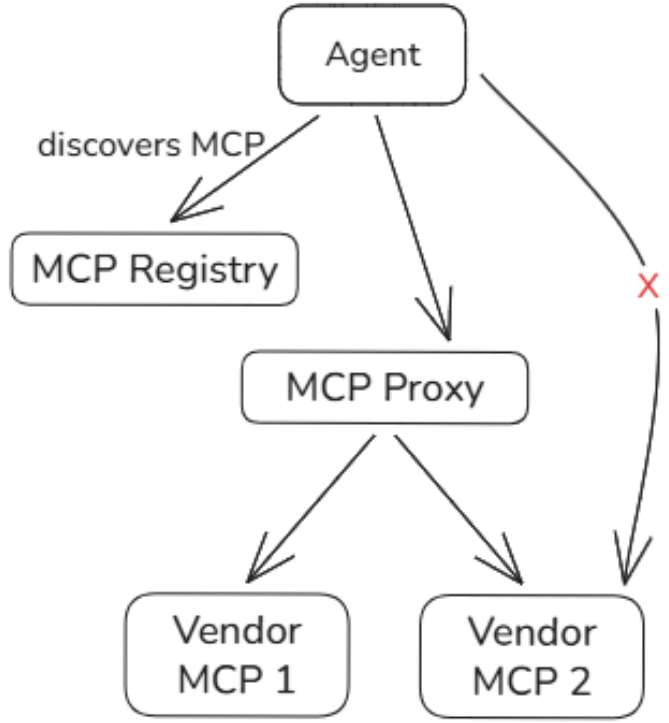

MCP RISK → ENTERPRISE ARCHITECTURE MAPPING

MCP Risk	Primary Architecture Control
MCP01: Token Mismanagement & Secret Exposure	Identity Plane, Workload Identity, Secretless Authentication
MCP02: Privilege Escalation via Scope Creep	Policy Enforcement, Capability Graph Analysis
MCP03: Tool Poisoning	MCP Registry, Approval Workflow, Tool Certification
MCP04: Supply Chain Attacks & Dependency Tampering	Runtime Governance, Artifact Validation, Provenance Controls
MCP05: Command Injection & Execution	Execution Controls, Sandboxing, Policy Enforcement
MCP06: Intent Flow Subversion	Intent Governance, Policy Engine, Approval Gates
MCP07: Insufficient Authentication & Authorization	Identity-Aware Control Plane, AuthN/AuthZ Layer
MCP08: Lack of Audit & Telemetry	Audit Pipeline, Observability Plane, Trace Correlation
MCP09: Shadow MCP Servers	Discovery Registry, Registration Controls, Ownership Tracking
MCP10: Context Injection & Over-Sharing	Context Firewall, External MCP Proxy, Data Governance

THE DREAM ...

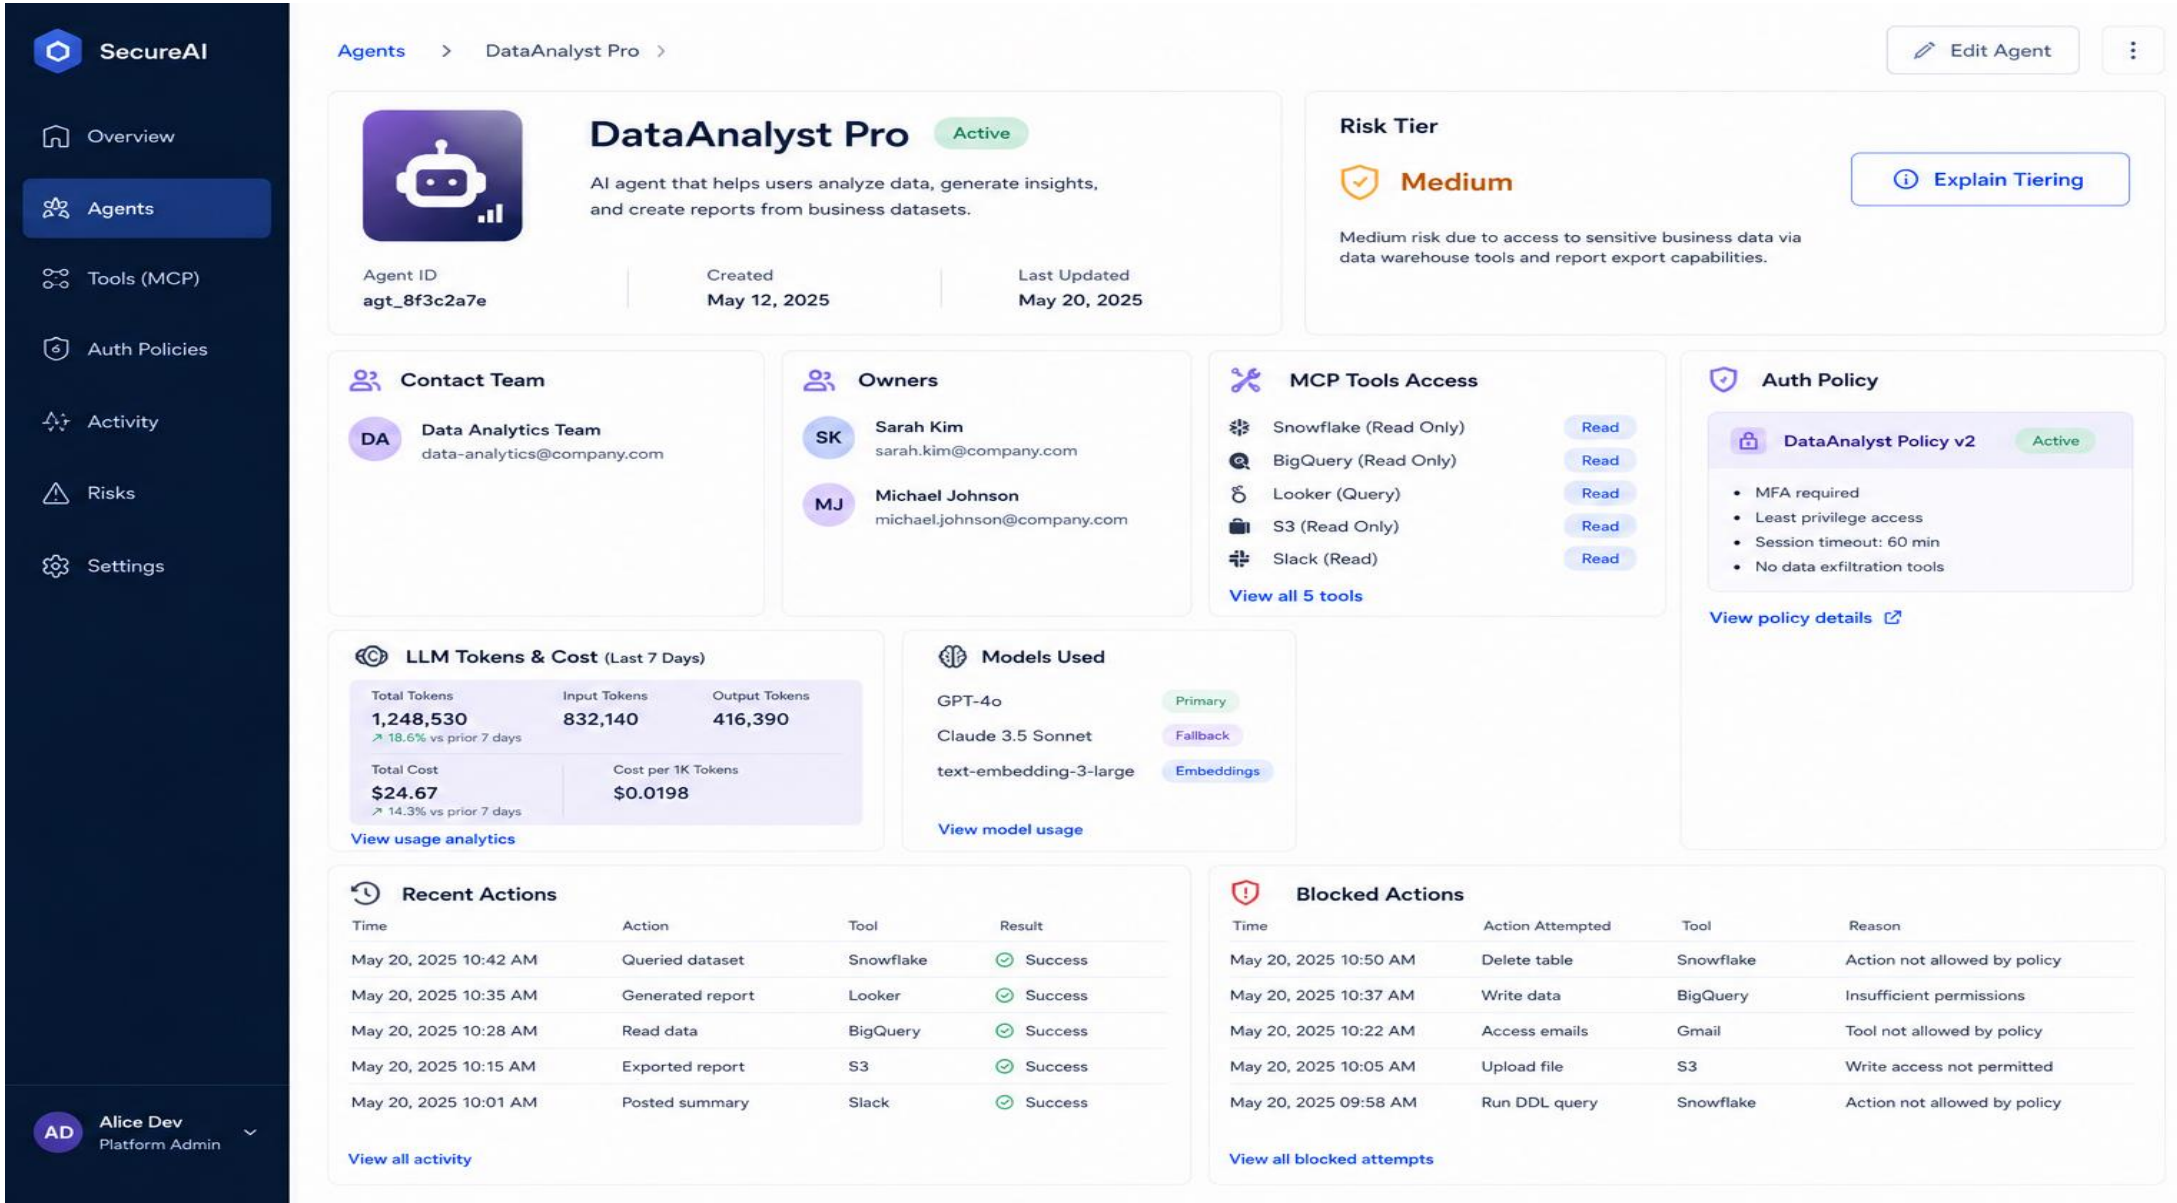
SecureAI
Overview
Agents
Tools (MCP)
Auth Policies
Activity
Risks
Settings
Alice Dev
Platform Admin
Agents > DataAnalyst Pro >
Edit Agent
DataAnalyst Pro
Active
AI agent that helps users analyze data, generate insights, and create reports from business datasets.
Agent ID
agt_8f3c2a7e
Created
May 12, 2025
Last Updated
May 20, 2025
Risk Tier
Medium
Explain Tiering
Medium risk due to access to sensitive business data via data warehouse tools and report export capabilities.
Contact Team
DA
Data Analytics Team
data-analytics@company.com
Owners
SK
Sarah Kim
sarah.kim@company.com
MJ
Michael Johnson
michael.johnson@company.com
MCP Tools Access
Snowflake (Read Only) Read
BigQuery (Read Only) Read
Looker (Query) Read
S3 (Read Only) Read
Slack (Read) Read
View all 5 tools
Auth Policy
DataAnalyst Policy v2
Active
MFA required
Least privilege access
Session timeout: 60 min
No data exfiltration tools
View policy details
LLM Tokens & Cost (Last 7 Days)
Total Tokens
1,248,530
↗ 18.6% vs prior 7 days
Input Tokens
832,140
Output Tokens
416,390
Total Cost
$24.67
↗ 14.3% vs prior 7 days
Cost per 1K Tokens
$0.0198
View usage analytics
Models Used
GPT-4o Primary
Claude 3.5 Sonnet Fallback
text-embedding-3-large Embeddings
View model usage
Recent Actions
Time Action Tool Result
May 20, 2025 10:42 AM Queried dataset Snowflake Success
May 20, 2025 10:35 AM Generated report Looker Success
May 20, 2025 10:28 AM Read data BigQuery Success
May 20, 2025 10:15 AM Exported report S3 Success
May 20, 2025 10:01 AM Posted summary Slack Success
View all activity
Blocked Actions
Time Action Attempted Tool Reason
May 20, 2025 10:50 AM Delete table Snowflake Action not allowed by policy
May 20, 2025 10:37 AM Write data BigQuery Insufficient permissions
May 20, 2025 10:22 AM Access emails Gmail Tool not allowed by policy
May 20, 2025 10:05 AM Upload file S3 Write access not permitted
May 20, 2025 09:58 AM Run DDL query Snowflake Action not allowed by policy
View all blocked attempts

THANK YOU