BlueRock | Adspirer

CUSTOMER CASE STUDY: ADSPIRER

How Adspirer Secures the Full Lifecycle of AI Advertising Operations with BlueRock

BlueRock helped Adspirer secure and scale AI-powered advertising operations by adding runtime visibility, policy enforcement, and operational control across its MCP infrastructure.

Executive Synopsis

Adspirer is building an AI-native advertising platform that enables marketing teams to execute campaigns, analytics, and optimization workflows across Google Ads, Meta, LinkedIn, TikTok, Klaviyo, and other systems directly from conversational AI interfaces such as Claude and ChatGPT.

The core offering relies on operating MCP servers at scale to process tens of millions of tool calls for several thousand enterprise customers. As adoption accelerated, Adspirer recognized that security, runtime visibility, and operational control would become foundational requirements for scaling their advertising infrastructure while supporting enterprise environments.

To ensure an enterprise-ready security model, Adspirer integrated BlueRock across the full lifecycle of its MCP infrastructure — from CI/CD security validation during the build process through runtime observability and operational control once deployed into production environments.

Using BlueRock's MCP Trust Registry, Adspirer continuously scans its server builds prior to deployment to identify MCP-specific vulnerabilities that traditional static application security testing (SAST) scans do not pick up.

> **"As an agentic AI company, Adspirer powers AI performance marketing agents that plan, launch, and optimize campaigns across Meta, Google, and more. With thousands of decisions flowing through our MCP infrastructure, our customers expect reliability and security. BlueRock gives us the visibility and control we need to run agentic teams with confidence.**
>
> **— Abhi Mekala, CEO of Adspirer**

Once deployed using BlueRock, Adspirer gets full runtime observability, and agentic execution layer guardrails to protect against undesirable tool execution or malicious runtime actions.

By integrating BlueRock across both development and runtime operations, Adspirer strengthened enterprise trust while gaining continuous visibility, protection, and operational control throughout its agentic infrastructure lifecycle.

How Adspirer Is Building the Future of Agentic Marketing

Adspirer represents a new category of AI-native platforms designed around autonomous execution rather than traditional software workflows.

Instead of forcing users into entirely new interfaces, Adspirer allows marketing teams to work directly from AI systems they already use while AI agents coordinate actions across advertising and analytics platforms behind the scenes.

This creates significant operational leverage for marketing teams, but it also introduces a new class of infrastructure and security concerns over MCP.
As a service provider helping coordinate large-scale ad campaigns, Adspirer needed greater visibility into how tools were being invoked, how workflows executed at runtime, and how to apply operational controls without slowing down innovation.

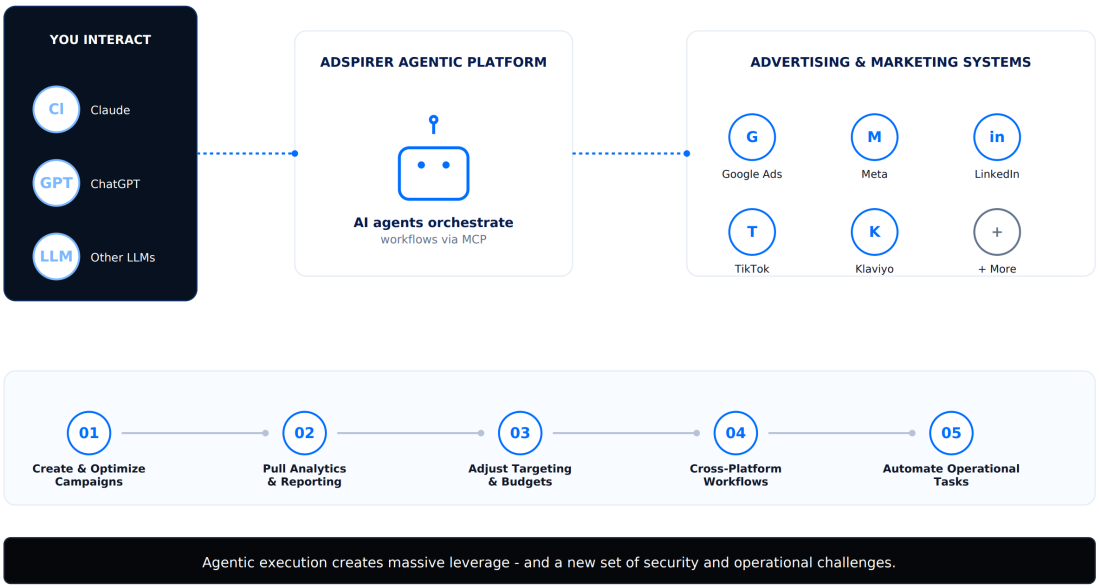

Through MCP infrastructure, users can:

- Create and optimize campaigns
- Pull analytics and reporting
- Adjust targeting and budgets
- Coordinate cross-platform advertising workflows
- Automate operational marketing tasks through natural language

Security also became increasingly important to the company's broader growth strategy. As Adspirer expanded enterprise engagements, customers and partners expected stronger assurances around MCP security, runtime protection, and operational governance before connecting sensitive advertising environments and production marketing systems.

"The future of marketing is agentic. We believe it has to be secure by default."
— Adspirer

The Challenge

Adspirer's platform was built around public-facing MCP infrastructure connecting AI agents directly to downstream advertising and analytics systems.

As the company scaled, traditional security and logging approaches no longer provided sufficient visibility or operational confidence for autonomous execution environments. The company also recognized that securing agentic systems required controls that extended beyond production runtime alone and into the software delivery lifecycle itself.

Key Challenges

- Continuously evaluate MCP security posture before deployment
- Understand how agents interacted with tools and downstream systems
- Monitor runtime behavior across agentic workflows
- Reduce exposure to prompt injection, tool abuse, and unsafe execution paths
- Improve operational investigations during enterprise security evaluations

The team needed a way to:

- Ensure the code base for the MCP servers was robust and adhered to best practices for securing MCP
- Strengthen enterprise trust during security reviews and customer evaluations
- Better understand how AI agents interacted with tools and downstream systems
- Monitor runtime behavior across agentic workflows
- Reduce exposure to prompt injection, tool abuse, and unsafe execution paths
- Protect against both known and unknown exploits given accelerated time to exploit in the AI era
- Improve operational investigation capabilities
- Scale infrastructure confidently as adoption accelerated

This became especially important as Adspirer underwent enterprise security evaluations and worked to demonstrate that its agentic infrastructure could safely operate across highly sensitive advertising systems and enterprise marketing environments.

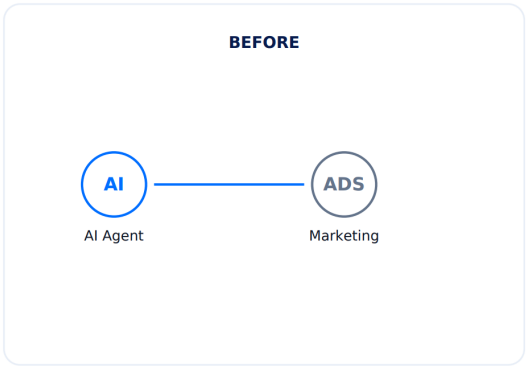

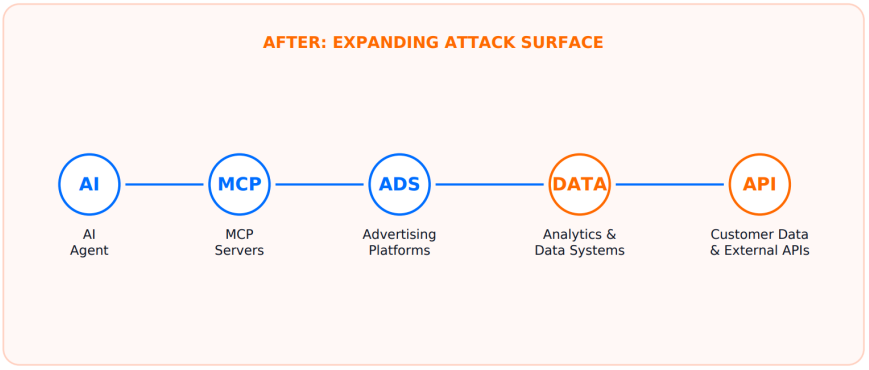

Why Adspirer Chose BlueRock

Adspirer needed more than traditional API monitoring or perimeter security. The company needed protection across the full lifecycle of its MCP infrastructure, from build-time validation to runtime visibility and execution-time controls. BlueRock addressed this through three connected capabilities.

Understanding the Risk Profile of MCP Infrastructure

Challenge

Adspirer needed a way to continuously evaluate the security posture of its MCP environment as agents interacted with tools connected to sensitive advertising systems and enterprise workflows.

Solution

BlueRock's MCP Trust Registry provided continuous MCP server scanning and risk analysis mapped against frameworks including:

- OWASP MCP Top 10
- MITRE CWE
- MAESTRO

The platform gave Adspirer visibility into exposed tools, configuration risks, security posture, and remediation opportunities while supporting broader enterprise trust and customer assurance initiatives.

Gaining Visibility Into Agentic Runtime Behavior

Challenge

Traditional logging and monitoring systems were not designed to explain how autonomous AI workflows behaved after execution began. Adspirer needed visibility into how agents interacted with tools, what actions were being executed downstream, and how to investigate abnormal behavior quickly.

Solution

BlueRock's runtime observability platform provided visibility across the full agentic action path, including:

- Agent decisions
- MCP tool calls
- Downstream execution behavior
- Runtime anomalies and behavioral drift

This gave Adspirer significantly improved operational telemetry across autonomous workflows and simplified investigation processes when abnormal behavior occurred.

Applying Security Controls at the Moment of Execution

Challenge

Adspirer needed a way to apply operational protections directly where autonomous actions occurred rather than relying solely on static perimeter controls or post-event analysis.

Solution

BlueRock introduced runtime guardrails that enforced policy controls across:

- Agent-to-tool interactions
- Agent-to-data movement
- Execution boundaries

This allowed Adspirer to apply security and operational controls directly at execution time while preserving the speed and flexibility required for AI-powered marketing workflows.

Technical Integration & Operational Scaling

BlueRock integrated directly alongside Adspirer's MCP workloads running on GCP infrastructure, providing runtime visibility, observe-only deployment, and execution-time controls without disrupting existing conversational AI workflows.

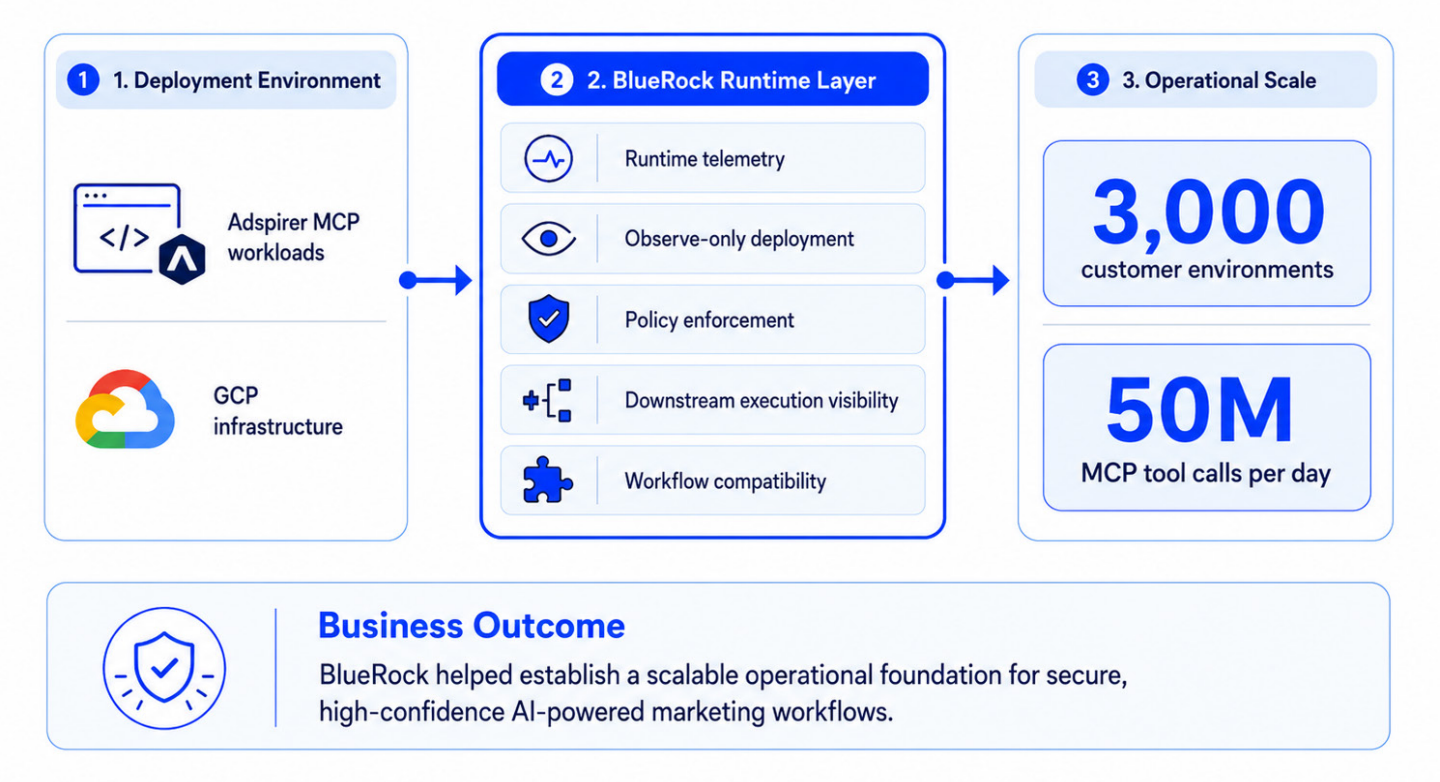

As adoption accelerated, the partnership also established a scalable operational foundation capable of supporting large-scale agentic environments and growing enterprise demand.

Conclusion & Results

BlueRock helped Adspirer strengthen operational confidence and enterprise readiness during a period of rapid platform growth. By extending security and operational controls across development and runtime environments, Adspirer established a more complete lifecycle approach for safely scaling agentic infrastructure into enterprise environments.

Key Results

Improved Visibility

Into MCP activity and execution behavior.

Reduced Complexity

During investigation workflows across agentic systems.

Strengthened Security

Across development and runtime environments.

Supported Assurance

For enterprise customer security requirements.

Scaled with Confidence

Across agentic marketing workflows.

Adspirer and BlueRock share a common belief that the future of enterprise software will increasingly be driven by autonomous systems operating across tools, workflows, and production infrastructure.

As organizations adopt agentic workflows at larger scale, runtime visibility, execution-layer controls, and operational trust will become foundational requirements for enterprise AI adoption.

Together, the companies are helping define what secure, scalable lifecycle operations for agentic infrastructure look like across development, deployment, and runtime environments.

> "The speed of agentic marketing, with guardrails that keep it safe at scale." — Adspirer

About the Companies

About Adspirer

Adspirer is an AI-powered advertising platform that enables marketing teams to operate campaigns through conversational AI systems using MCP infrastructure. The platform connects autonomous workflows across major advertising and analytics systems including Google Ads, Meta, LinkedIn, TikTok, and Klaviyo.

About BlueRock

BlueRock.io provides runtime observability, guardrails, and trust infrastructure for agentic systems and MCP environments. Its platform helps organizations securely operate, understand, and control autonomous execution across AI-native infrastructure.

bluerock.io