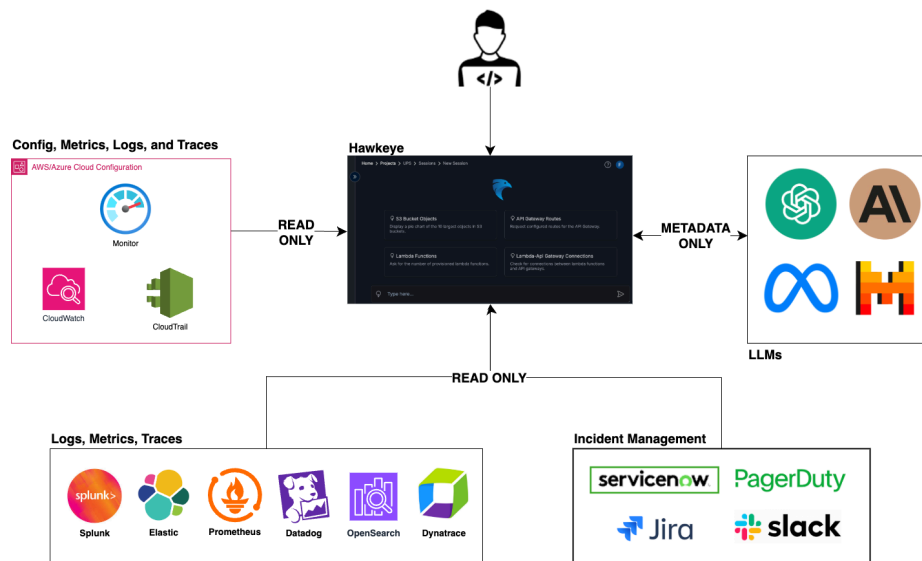


How Hawkeye Works- Deep Dive: Secure GenAI-Powered IT Operations

Modern IT operations generate an overwhelming amount of telemetry data across dozens of tools and platforms. While traditional approaches struggle to process this complexity, Hawkeye takes a fundamentally different approach - using GenAI to transform how we analyze and respond to IT incidents. Let's look under the hood to understand how Hawkeye works and why our security-first architecture sets us apart.



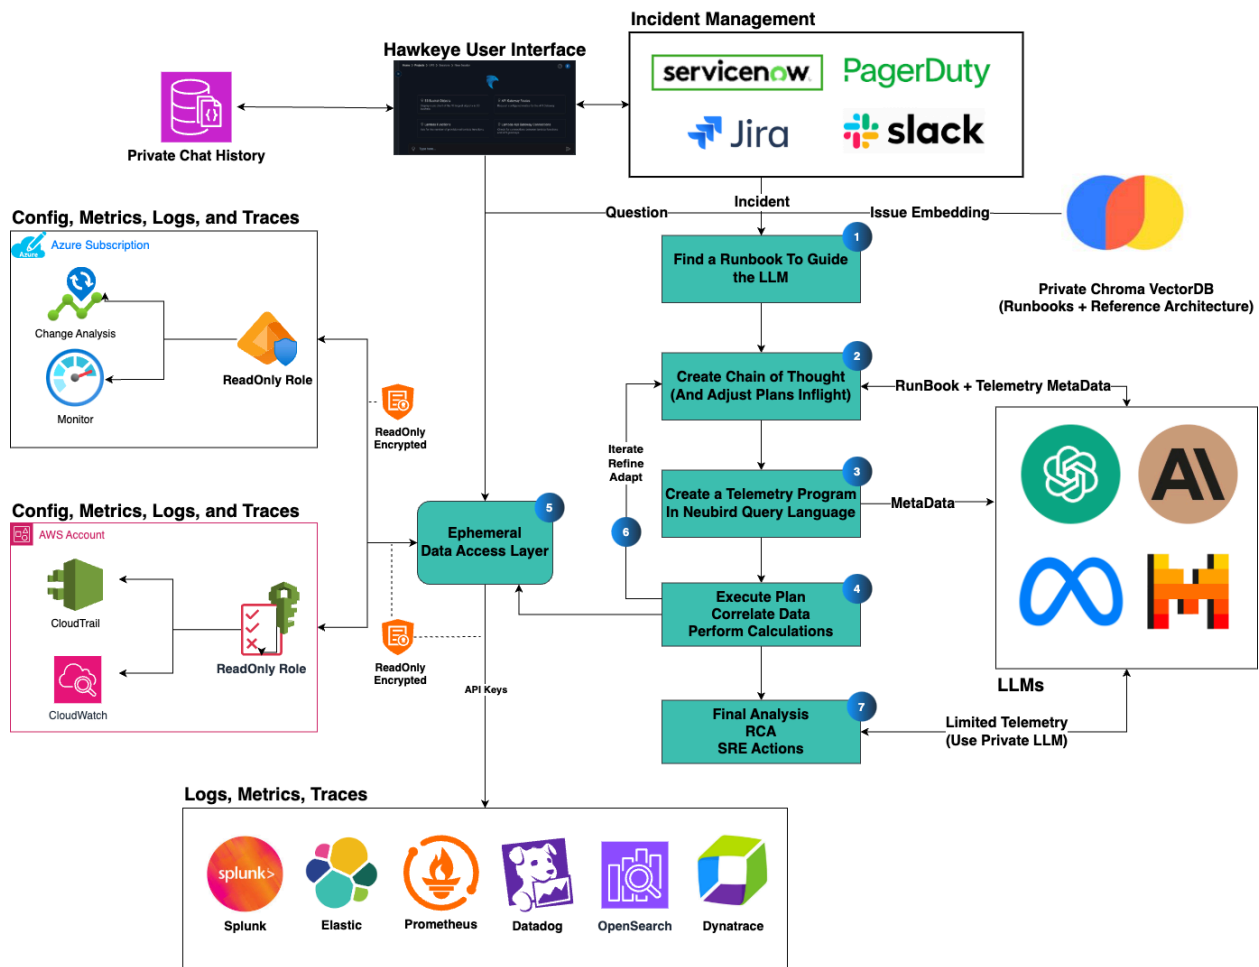
The Foundation: Security and Privacy Built in

Before diving into Hawkeye's technical architecture, it's crucial to understand our foundational security principles:

- **Zero Data Storage:** Hawkeye operates as a completely ephemeral platform. We process your telemetry data in real-time and never store historical information. Once an analysis session ends, all data is automatically purged from memory.
- **Read-Only by Default:** Every connection to your infrastructure uses strictly read-only permissions. This isn't just a policy - it's architecturally enforced, making it technically impossible for Hawkeye to modify your systems or data.
- **Customer-Controlled Access:** You maintain complete control through customer-specific external IDs and custom trust policies. Access can be revoked instantly at any time.

The Architecture: Step by Step

Let's walk through how Hawkeye processes an incident or investigation, following our architectural diagram below (**Diag. Hawkeye from Neubird architecture step by step**):



Diag. Hawkeye from Neubird architecture step by step

Step 1. Finding the Right Approach Using Runbooks

When an incident occurs Hawkeye first step is selecting the appropriate analysis strategy. Using your private ChromaDB vector database, Hawkeye identifies similar historical patterns and successful investigation approaches. It uses an embedding of your issue and ChromaDB's fast similarity search - without ever storing any of your telemetry data.

As it learns more about your investigations this vector database can build up knowledge about investigation plans that work best for your systems.

2. Creating the Investigation Plan Using LLMs Reasoning Capabilities

At this step, the LLM's reasoning capabilities are leveraged to formulate a dynamic investigation plan, one that may be inspired by the retrieved information from step one but leverages the generative power of the LLM to adapt it. Hawkeye constructs a detailed chain of thought for the investigation, adapting its approach based on:

- The type of incident or investigation
- Available telemetry sources described through metadata
- Description of your architecture based on available information
- Historical patterns of similar issues

No configuration or telemetry data is included in the prompts to the LLMs. The choice of LLM changes based on the up to date benchmarks to achieve the best results.

Step 3. Telemetry Program Generation

Here's where Hawkeye's innovation shines. Instead of sending your sensitive telemetry data to an LLM, Hawkeye:

- Creates a specialized telemetry retrieval program
- Uses our fine-tuned LLM only for program logic, never for data processing
- Ensures all actual data handling happens in isolated memory space

The fine-tuning of the LLM (currently based on Llama 3.2 70B) is done by Neubird using only synthetic data programs and leveraging [LARK](#) files to control and validate the syntax of the generated telemetry program is valid and will produce results.

Step 4. Secure Data Processing

Hawkeye executes the telemetry program in a secure, ephemeral runtime environment:

- Correlates data across multiple sources
- Performs necessary calculations and mathematical analysis in python
- Maintains strict memory isolation for each customer's telemetry data
- Automatically purges all data after processing an investigation

Step 5. Real-Time Data Access Layer

Hawkeye's second secret weapon is its secure data access layer. Queries to access data are all written in a common syntax which the fine-tuned LLM can generate with 100% accuracy, resulting in reliable and precise data access no matter what the data source is. Our secure data access layer:

- Uses temporary credentials with minimal scope
- Implements read-only access across all integrations
- Supports major cloud providers (AWS, Azure, GCP) and observability tools
- Never stores your telemetry data on disk
- Leverages schema on read technology, avoiding issues with schema drift

Step 6. Continuous Refinement

As the investigation progresses, Hawkeye:

- Iteratively refines its analysis approach based assertion on the data performed by the secure data processor, allowing to adapt to new information without sending telemetry data to the LLM
- Maintains audit trails of its reasoning and investigation steps
- Never uses your data to train or improve its models

Step 7. Final Analysis and Actions

Once all facts are available and the investigation has converged Hawkeye will produce:

- Detailed root cause analysis
- Clear evidence for all findings
- Specific recommended actions

In order to protect your telemetry and configuration data, Hawkeye leverages a privately hosted open source LLM to produce this final analysis.

Getting Started

Ready to transform your IT operations with Hawkeye? [Here's how to begin.](#)

Conclusion

Hawkeye represents a fundamental shift in IT operations - combining the power of GenAI with unwavering commitment to security and privacy. By processing complex telemetry data in real-time while maintaining zero data persistence, we're transforming how teams handle incidents and investigations.

Ready to see Hawkeye in action? Contact us to schedule a demo and learn how we can help transform your IT operations while maintaining the highest security standards.