

Accelerating Resilient AI Agent Development **Using Hyper-Synthetic Data**

Organization:

Leading Global AI & Cloud Provider

Target Audience:

AI Engineering, Security Product Management, and SecOps Teams

Core Objective:

Train, evaluate, and stress-test both deterministic security tools and autonomous AI agents without relying on real customer production data.

The Challenge: Sourcing High-Fidelity Data to Build & Train Next-Gen Security AI & SIEM Tools

Building reliable security AI agents requires testing them against complex attack scenarios. However, testing in sterile labs creates fragile AI that fails in the real world, while using real customer networks introduces catastrophic privacy risks and regulatory compliance blocks.

As a global leader in cloud computing and artificial intelligence, the engineering teams faced a major bottleneck in their product development lifecycle: sourcing clean, diverse, and realistic security data to build and evaluate next-generation security solutions.

The provider required vast volumes of telemetry to train two fundamentally distinct technical targets:

01. Deterministic SIEM & SecOps Tools

Rule-based central platforms require standardized data centralization, automated threat ingestion, and predictable SOAR execution.

Without centralized, predictable telemetry, security automation breaks down. SOC teams are forced back into manual alert processing, driving up Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR). Unpredictable SOAR execution leads to missed detections, broken remediation playbooks, and catastrophic alert fatigue that lets active breaches slip through unnoticed.

02. Non-Deterministic Agentic AI

Reasoning AI agents designed for autonomous triage, threat hunting, and investigation. These agents require comprehensive evaluation datasets ("evals") to test reasoning loops, tool invocation via APIs, and Model Context Protocol (MCP) servers, and verdict accuracy.

Insufficient evaluation data causes non-deterministic AI agents to hallucinate, misunderstand operational context, or take improper containment actions. In a live enterprise SOC, a rogue or untrusted AI agent risks shutting down critical production infrastructure, dropping legitimate traffic, or completely misidentifying active threat vectors.

The Dual Friction Points

Strict Customer Data Restrictions:

Due to rigorous privacy compliance and strict corporate data governance, product teams could not tap into real customer production networks or log feeds containing Personally Identifiable Information (PII) to train or benchmark their models.

The “Needle in a Pile of Needles” Problem:

Traditional lab environments only generated isolated attack traffic. Lacking realistic background business activity, these static labs forced AI models to search for a “needle in a pile of needles.” To reliably function in real enterprise environments, the models needed to learn how to distinguish subtle adversary behavior from normal user noise—finding a threat needle in a true enterprise haystack.

The Solution: Generating Hyper-Synthetic Data in SimSpace AI Proving Grounds


To overcome the data bottleneck, the provider integrated advanced cyber simulations into their continuous evaluation and training pipeline through the SimSpace AI Proving Grounds. SimSpace's cyber range provided a set of accurate, labeled, representative enterprise environments capable of generating high-quality streams of hyper-synthetic telemetry.

SimSpace Generates Hyper-Synthetic in the AI Proving Grounds

Core Capability	Technical Description
Programmable User Emulation (UE)	Generates realistic user activity (email, web browsing, cloud access, application interactions) to create patterns of life in the environment.
Programmable Attack Emulation (AE)	Creates realistic, full attack chain scenarios that recreate the most notable threat actors compromising the environment.
Multi-Cloud Enterprise Terrain	Production-grade cloud topologies with modern security toolstacks in isolated environments.
Range-as-Code & MCP Automation	API-driven VM cloning, MCP server integration, and instant range resets enable rapid, automated iteration loops.

Key Technical Capabilities Deployed:

- ▶ **Programmable User Emulation (UE):** SimSpace injected authentic background enterprise traffic—simulating day-to-day employee email activity, web browsing, application interactions, and cloud access. This created a noisy, realistic baseline environment alongside attack scenarios.
- ▶ **Programmable Attack Emulation (AE):** SimSpace deployed end-to-end, multi-stage adversary campaigns that emulated the sophisticated tactics, techniques, and procedures (TTPs) of high-profile threat actors. This gives the provider's engineering teams the realistic compromise scenarios needed to rigorously benchmark agentic response pathways and stress-test SIEM detection rules.
- ▶ **Multi-Cloud Topology:** SimSpace constructed realistic network topologies to mirror real-world customer architecture.
- ▶ **Automated Range Orchestration:** Utilizing Range-as-Code, API integrations, and MCP servers, the provider automated environment provisioning, VM cloning, and instant range resets. This allowed continuous, programmatic execution of complex attack scripts alongside automated user emulation.



The Results: The "4Ss" Framework

By replacing static lab setups with SimSpace’s hyper-synthetic data engine, the global technology provider established a high-velocity development loop defined by the 4 Ss:

Pillar	Operational Impact
01. Speed	API-driven VM cloning and automated resets reduced environment staging from days to minutes, dramatically accelerating iteration cycles.
02. Scale	High-velocity training loops generated petabyte-scale synthetic telemetry across diverse attack surface variations without risking customer PII.
03. Sophistication	SimSpace enabled testing against cutting-edge, complex threats—including AI-Empowered Identity Attacks, Deepfakes, and exploits targeting Vulnerable LLMs/Applications—embedded inside noisy user environments.
04. Stamina	24/7 continuous simulation loops enabled round-the-clock model evaluation, shifting human engineers from manual testing to strategic architecture and policy tuning.

Detailed Impact Summary

01. Time (Accelerated Model Evals & Iteration Velocity)

- ▶ Engineering teams can now test and benchmark new AI capabilities (such as tool calling, context retention, and reasoning pathways) against synthetic scenarios.
- ▶ Environment provisioning and scenario resets dropped from days to minutes, allowing product teams to shorten model evaluation cycles and ship agentic capabilities to production significantly faster.

02. Capability (Dual-Target Optimization & Multi-Layered Fidelity)

- ▶ The exact same range infrastructure successfully fed structured telemetry into deterministic SIEM rule engines while simultaneously providing interactive sandbox environments to evaluate non-deterministic AI agents.
- ▶ Advanced user emulation injected realistic enterprise activity into the range, giving AI agents the sophisticated training required to recognize subtle lateral movement without triggering false positives.

03. Cost Savings (Elimination of In-House Range R&D & Infrastructure Overhead)

- ▶ Building, maintaining, and continually updating a proprietary, hyper-realistic cyber simulation environment in-house would have required millions of dollars in dedicated engineering headcounts, infrastructure maintenance, and ongoing adversary research.
- ▶ Leveraging SimSpace's turn-key platform delivered an immediate, production-grade synthetic data engine—saving years of internal tool development costs and significantly lowering total cost of ownership (TCO).

04. Trust (Zero Privacy Risk & Responsible AI Governance)

- ▶ By building on clean-room hyper-synthetic data, product teams eliminated compliance hurdles, removing the need to navigate customer data agreements while maintaining production-grade fidelity.
- ▶ Rigorous, reproducible evaluation datasets ("evals") established a trusted benchmark to validate AI agent reasoning pathways, minimize hallucinations, and ensure safe deployment into production enterprise environments.

“SimSpace allows us to create dynamic datasets over and over again, different types of attacks, changing the environment to really replicate our different customers, so we can train our agents without risking customer data.”

— Senior Product Manager, Leading Global AI & Cloud Provider

Prove Resilience with Preemptive Cyber Simulation in the AI Proving Grounds

Talk to an AI security
expert at SimSpace.com

