

RAD Security

The Agentic Core for Modern Security Operations



Product Deployment Quickstart

Deployment Guide Overview

This deployment guide outlines how RAD Security leverages an agentic core to transform modern security operations. We will explore our innovative approach to integrating AI, contextual intelligence, and automated response capabilities to enhance your organization's defense posture.

01

Welcome to RAD Security

02

The Four Pillars of Agentic SecOps

03

Introducing RAD FusionAI

04

The Knowledge Fabric

05

Integrations as Intelligence

06

Runtime/Cloud/Identity Context

07

RADBot: The Agentic Analyst

08

Operationalize RAD with Agentic Workflows

09

Use Cases (Vulnerability Intelligence & AI-Powered Investigation)

10

Proof of Concept Flow

11

Quick Onboarding Steps

12

Next Steps & Contact

Welcome to RAD Security

RAD Security redefines how security operations think. It's not another scanner or dashboard. RAD is an agentic reasoning layer that turns every signal from your stack into explainable action. Built for teams that want reasoning, not rules. Deploy fast. Learn faster. Evolve continuously.



Agentic Correlation

RAD Security redefines how security operations think, acting as an agentic reasoning layer.



Explainable Action

Turns every signal from your stack into clear, explainable action.



Reasoning, Not Rules

Built for teams that want intelligent reasoning, rather than rigid, predefined rules.



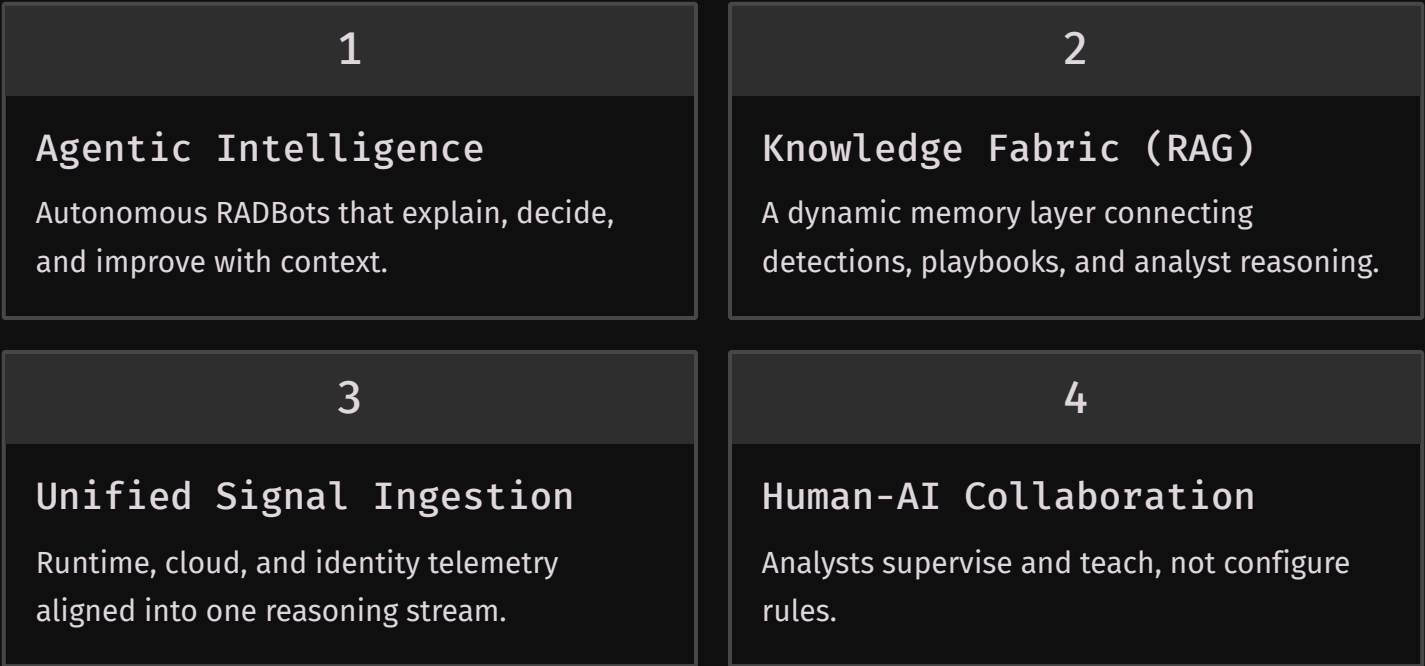
Fast Evolution

Deploy fast, learn faster, and evolve continuously to meet changing threats.




The Four Pillars of Agentic SecOps

RAD Security is built upon fundamental principles that ensure comprehensive, intelligent, and adaptive security operations. These four pillars form the core of our agentic approach:




These pillars work in concert to create a robust and continuously learning security environment, transforming raw signals into actionable intelligence through a central Agentic Core.



Introducing RAD FusionAI

The challenge isn't visibility, it's understanding. Security teams drown in alerts that lack history and context. RAD adds the missing layer of reasoning. It listens to the tools you have in place today, and RAD proprietary cloud collectors and makes them think together. Instead of dashboards, you get dialogues. Instead of triage queues, you get decisions.




The Knowledge Fabric

RAD's Knowledge Fabric encodes the institutional memory of your security operations. It's not static documentation - it's a living vector database of detections, resolutions, and insights. It learns from:


- Runtime and workload fingerprints
- Cloud posture and configuration history
- Analyst feedback and decisions
- External intelligence and advisories

Every new RADBot explanation references this shared fabric to stay accurate and aligned with your organization.



Integrations as Intelligence


RAD thrives in ecosystems. It doesn't replace your tools — it learns from them. By seamlessly connecting with your existing security infrastructure, RAD gathers critical context and operational insights. This deep integration allows RAD to build a comprehensive understanding of your environment, leading to more accurate detections and informed responses.



Each integration is a knowledge contributor, expanding RAD's collective reasoning.

Runtime, Cloud, and Identity – Context Providers

RAD treats infrastructure as signal, not surface area. Runtime — Shows what's actually executing. Cloud — Reveals where it runs and who owns it. Identity — Explains who's acting and why. These streams give RADBots the ground truth for contextual reasoning.



RADBot: The Agentic Analyst

RADBot is your autonomous analyst — a thinking system trained on your environment's memory. It doesn't send alerts. It starts conversations. Examples:

Reasoning Summary

'This deviation matches finance-cluster drift pattern from March.'

High Confidence

'Reachable CVE confirmed via live OpenSSL call chain.'

Approved Action

'No action: matches analyst-approved baseline.'

Every output is explainable, traceable, and stored for continuous improvement.

The screenshot displays the RADBot chat interface. On the left is a sidebar with a 'history' section containing three items: 'GitHub Security Check: Investigating Po...' dated 10/7/2025, 'CVE Triage Support: Navigating Securit...' dated 10/8/2025, and 'Mapping CVEs to Cloud Security Postur...' dated 10/3/2025. The main chat area has a title bar 'GitHub Security Check: Investigating Potential Accidental Public Repository Commits'. A user message asks: 'Have we had any incidents where we accidentally committed to a public GitHub repository'. The RAD Assistant responds with a detailed analysis. It states: 'Based on the incident reports in your system, yes, you have had incidents where AWS access keys were accidentally committed to public GitHub repositories. Here are the two documented cases:'. It then lists two incidents: 'Incident 1 - March 15, 2024 (Successful Response)' and 'Incident 2 - August 1-3, 2024 (Security Breach)'. Each incident includes details about the developer, repository, exposed keys, detection time, outcome, and response.

Incident 1 - March 15, 2024 (Successful Response)


- **Developer:** Marcus Thompson
- **Repository:** `cloudstore-frontend-v2`
- **Keys Exposed:** AKIAIOSFODNN7EXAMPLE
- **Detection Time:** 4 minutes (09:36:38 UTC after 09:32:15 UTC commit)
- **Outcome:** Successful mitigation - No unauthorized access occurred
- **Response:** Automated monitoring detected the exposure quickly, keys were revoked within 5 minutes, and the incident was fully resolved

Incident 2 - August 1-3, 2024 (Security Breach)

- **Developer:** Amy Foster
- **Repository:** `cloudstore-mobile-sdk`
- **Keys Exposed:** AKIAI44QH8DHBEXAMPLE
- **Detection Time:** 2 days (undetected due to monitoring configuration error)
- **Outcome:** Cryptocurrency mining attack - External actors exploited the keys
- **Impact:** Unauthorized EC2 instances launched across 4 AWS regions for mining operations

Use Case: Vulnerability & Reachability Intelligence

RAD focuses on contextual risk, not volume. It maps each CVE to runtime evidence, identity exposure, and business relevance — then ranks what truly matters. Example: 'CVE-2025-0421 reachable via active OpenSSL call chain in production. High confidence drift detected.'



Use Case: AI-Powered Incident Investigation

Traditional incident investigations are slow, manual, and costly. RAD Security's agentic AI dramatically transforms this, ****reducing investigation time from 4-6 hours to just 15-20 minutes****.

RADBot acts as an intelligent analyst, automating evidence collection across your security ecosystem. This automation not only speeds up investigations but also leads to ****savings of \$50,000+ annually per analyst****. Engage with RADBot via a conversational chat interface for clear explanations and actionable insights.



Alert Ingestion

Real-time feeds from security tools.



Automated Evidence

RADBot pulls context from all sources.



AI Analysis

Intelligent correlation for insights.



Conversational Guidance

Chat for explanations and next steps.



Rapid Resolution

Expedited, explainable actions.

This streamlined approach boosts efficiency, cuts exposure, and enables faster, more accurate threat response.

Proof of Concept Flow

Goal: show agentic reasoning in action – not just detections.

Step 1: Connect sources (runtime, cloud, identity)	Unified data stream
Step 2: RADBot builds context	Knowledge fabric forms
Step 3: Trigger behavioral drift	Agentic explanation generated
Step 4: Analyst validates	Feedback stored & retrievable
Step 5: Integrate third-party data	Full reasoning cycle visible

Within 30 minutes:

-  RAD Core reasoning live
-  Context flowing from multiple systems
-  First deviation explained in natural language
-  Analyst feedback loop active
-  Integrations enriching agentic reasoning

Experience rapid deployment and immediate insights with these initial milestones:

Quick Onboarding – Start Your Agentic SecOps Journey

Get RAD thinking in minutes. This setup connects your data sources, enriches the Knowledge Fabric, and deploys lightweight sensors for live reasoning.



Connect Data Sources

Seamlessly integrate existing security tools and platforms to feed RAD with crucial data.



Enrich Knowledge Fabric

Populate the living vector database with your historical data and organizational context for deeper reasoning.



Deploy Lightweight Sensors

Install agents for real-time runtime telemetry across your environment with minimal overhead.

Welcome to RAD Security! I'm here to help you get started. First, let's understand what you're looking to achieve. What's your primary security goal?

Choose your primary security focus:

☐ Vulnerability Management

☒ Compliance & Governance

☐ Cloud Security Posture

☐ Incident Response

vulnerability

Great choice! Now let's set up your organization. What would you like to name your tenant and account?

Enter your organization details:

e.g., Acme Corp - Production

test-org

Perfect! Now let's connect your infrastructure. Which cloud provider or platform would you like to integrate?

Select your cloud provider:

☒ Amazon Web Services (AWS)

☐ Google Cloud Platform

☐ Microsoft Azure

☐ Kubernetes Cluster

aws

Excellent! To better customize your security experience, could you tell me a bit about your organization?

What best describes your company?

☒ Startup (1-50 employees)

☐ Growing Company (51-200 employees)

☐ Enterprise (200+ employees)

☐ Prefer not to say

Step 1 – Connect Third-Party Sources to the Knowledge Fabric

RAD learns from your environment. Add your most important data sources first so the Agentic Core has rich context to reason with.

Recommended sources:

- Container registries & CI pipelines
- Ticketing and Communication Platforms (Slack, Jira)
- Identity providers (Okta, Azure AD)
- Security telemetry (Qualys, Rapid7 EDR, Crowdstrike, etc.)
- External intel feeds (Vuln DB, CVE, MITRE)

In the UI: Data Sources → Integrations → Choose and Configure Integration

The screenshot displays the 'Integrations' section of the RAD interface. At the top, there's a navigation bar with 'Data Sources' and 'Integrations' tabs, along with 'Refresh' and '+ Add New' buttons. Below this, a sub-navigation bar shows 'Connections', 'Integrations' (selected), 'Asset Inventory', and 'Knowledge Base'. The main content area is divided into two sections: 'Connected Integrations' (5 items) and 'Marketplace' (12 items). The 'Connected Integrations' section lists five integrations: Jira Cloud (Ticketing), CrowdStrike Falcon Insight (EDR), CrowdStrike Falcon Spotlight (Vulnerabilities), Okta Identity (IAM), and Google Workspace (IAM). Each integration card shows a status of 'Connected' and an 'Add Integration' button. The 'Marketplace' section lists three integrations: Splunk Enterprise (SIEM), Qualys (Vulnerabilities), and Microsoft Defender (EDR). These cards show a status of 'Disconnected' and an 'Add Integration' button. A left sidebar contains a 'CATEGORIES' list with buttons for 'All', 'Ticketing', 'SIEM', 'Vulnerabilities', 'EDR', and 'IAM'.

Step 2 – Deploy RAD Sensors

Deploy RAD's proprietary sensors to gain unparalleled visibility. These sensors are designed to collect maximum runtime telemetry, critical SBOM data, Kubernetes monitoring, and comprehensive cloud telemetry, feeding rich context directly to the Agentic Core.

Runtime Telemetry

Capture detailed execution data from applications and systems for deep behavioral analysis and threat detection.

SBOM Data Collection

Automatically generate and maintain Software Bill of Materials (SBOM) to track components and dependencies across your estate.

Kubernetes Monitoring

Gain granular insights into your Kubernetes clusters, including pod activity, network flows, and container security posture.

Cloud Telemetry

Collect extensive data from your cloud environments (AWS, GCP, Azure) to understand resource usage, configurations, and security events.

In the UI: Data Sources → Connections → Select Deployment Method → Verify Data Stream

Data Sources > Connections

Refresh

Add New

Connections

Integrations

Asset Inventory

Knowledge Base

Total Connections

14 ~ 5 cloud 9 k8s

Active Providers

5 ~ All active

Connection Health

100% ~ Healthy

Kubernetes Clusters

9 ~ 2 healthy

Cloud Providers (5)

Clusters (9)

Add Cluster

Kubernetes Clusters

Provider	Cluster Name	K8s Version	Resources	Health Status	Actions
Linode	cluster-reassign-3	1.27.5	1 Nodes 361 Workloads	Disconnected	...
Linode	k3d with plugins v1.0.25	1.27.4	1 Nodes 307 Workloads	Disconnected	...
Linode	marketplace-testing	1.29.6	2 Nodes 383 Workloads	Disconnected	...
Linode	Max Marketplace Test	1.29.0	2 Nodes 397 Workloads	Disconnected	...
Linode	max-marketplace-test	1.29.6	2 Nodes 584 Workloads	Disconnected	...
Linode	sri	1.27.15	5 Nodes 463 Workloads	Disconnected	...
Linode	TEST_PRD_KSOC_nonEKS	1.27.4	3 Nodes 312 Workloads	Disconnected	...
Linode	gke-cluster-standard	1.33.4-gke.1245000	9 Nodes 697 Workloads	Healthy	...
AWS	pe-prd-us-west-2	1.31	36 Nodes 2902 Workloads	Healthy	...

Step 3 – Add Knowledge Sources

Beyond live telemetry and system configurations, RAD can enrich its agentic reasoning with your organization's accumulated knowledge. By integrating internal documents such as penetration test reports, security policies, compliance documentation, and contracts, RAD gains invaluable context to understand risks, validate actions, and provide highly relevant insights.

Penetration Test Reports

Integrate findings from past security assessments, detailing vulnerabilities, exploit scenarios, and recommended remediations.

Security Policies & Procedures

Provide access to internal guidelines, acceptable use policies, and operational security procedures for compliance checks.

Compliance Documentation

Upload regulatory frameworks (e.g., GDPR, HIPAA), audit reports, and control mappings to inform compliance-driven analysis.

Contracts & Agreements

Incorporate vendor agreements, SLAs, and data sharing policies to contextualize third-party risks and responsibilities.

Contracts & Agreements

Incorporate vendor agreements, SLAs, and data sharing policies to contextualize third-party risks and responsibilities.

Institutional Knowledge

Include post-incident review documents, architectural decisions, and other tribal knowledge to build a comprehensive historical record.

In the UI: Data Sources → Knowledge Base → Add Source → Upload / Modify Documents

Data Sources

Knowledge-base

Refresh

Add New

Connections

Integrations

Asset Inventory

Knowledge Base

Knowledge Base

Add File

File Name	Status	Collections	Uploaded	Last Updated	Actions
SOC2_Compliance_Checklist.pdf	Ready	grc	Sep 9, 2025	Sep 9, 2025	...
openapi.txt	Ready	eol	Aug 21, 2025	Aug 21, 2025	...
crowdstrike-global-threat-report-2024.pdf	Ready	threats	Jul 24, 2025	Jul 24, 2025	...
incident_report_006_missing_security_patch_noncritical.md	Ready	incidents	Jul 24, 2025	Jul 24, 2025	...
incident_report_005_information_disclosure_stack_trace.md	Ready	incidents	Jul 24, 2025	Jul 24, 2025	...

Step 4 – Operationalize RAD with Agentic Workflows

RADBot is your conversational AI assistant, transforming SecOps with intelligent, agentic workflows:

- **Chat with RADBot** using natural language for complex security operations.
- **Automate workflows** by asking RADBot to analyze, propose, and execute tasks with analyst approval.
- **Execute complex tasks** like 'Investigate this suspicious login' or 'Multi-step incident containment' for immediate action and reduced response times.



Automate Incident Response

RADBot investigates alerts, correlates events, and executes pre-approved remediation playbooks, from isolating compromised hosts to patching vulnerabilities.



Proactive Threat Hunting

Direct RADBot to hunt for specific threats or anomalous behaviors across your environment, identifying stealthy attacks before they escalate.



Streamlined Compliance Reporting


Generate detailed compliance reports by querying RADBot for relevant security posture data, audit trails, and policy adherence metrics.



Contextual Data Enrichment

Automatically pull and present all relevant context for any security event, including user identity, asset criticality, and historical threat intelligence.

In the UI: RADBot Chat Interface → Initiate Conversation / Review Automated Workflow Logs



Ready to Revolutionize Your Security Operations?

Harness the power of agentic AI to transform your SecOps. RAD Security integrates unparalleled visibility, accelerated response, and intelligent automation to dramatically enhance your security posture, allowing your team to focus on strategic initiatives rather than reactive firefighting.

Experience the future of cybersecurity where threats are anticipated, incidents are contained with unprecedented speed, and your security team operates with peak efficiency.

Take the Next Step

Discover how RAD Security can integrate seamlessly into your environment and empower your SecOps team. Don't just respond to threats—anticipate and neutralize them with intelligent automation.

[Schedule a Demo](#)

[Get Started](#)

Connect With Us

- Website: radsecurity.ai
- Email: hello@rad.security

