7AI

WHITEPAPER

COMPLIANCE · THREAT HUNTING · HEALTHCARE

SATISFYING NIST 800-53 RA-10 THREAT HUNTING REQUIREMENTS WITH AI AGENTS

A practical guide for security leaders facing a continuous threat hunting mandate, a talent market that won't deliver the headcount, and an audit deadline that doesn't care.

AUDIENCE

Security leaders in healthcare, federal contracting, and other regulated industries.

FORMAT

Capability mapping, customer case study, cost analysis, and implementation guide.

CONTENTS

WHAT'S INSIDE

01 / EXECUTIVE SUMMARY

CONTINUOUS BY MANDATE. CONTINUOUS BY DESIGN.

NIST SP 800-53 Rev. 5 control RA-10 requires organizations to "establish and maintain a cyber threat hunting capability to search for indicators of compromise in organizational systems and detect, track, and disrupt threats that evade existing controls."

For organizations in healthcare, federal contracting, and other regulated industries, that single sentence creates a significant operational challenge. Traditional approaches to threat hunting are resource-intensive, difficult to staff, and rarely achieve the continuous cadence the control demands.

This paper examines how a major regional health system satisfied RA-10 with AI-powered threat hunting and one analyst, instead of the three threat hunters it had originally budgeted for. It includes:

- A direct mapping of 7AI capabilities to RA-10 control requirements
- Real-world results from production deployment, including a multi-wave phishing campaign uncovered three weeks in
- Cost comparison between traditional and AI-augmented approaches
- Implementation considerations for regulated industries

3 → 1	24/7	7-14D	100%
HUNTER HEADCOUNT REQUIRED	HUNT CADENCE ACHIEVED	TO PRODUCTION HUNTING	AUDIT-READY DOCUMENTATION

02 / THE RA-10 COMPLIANCE CHALLENGE

WHAT THE CONTROL ACTUALLY REQUIRES

NIST SP 800-53 Rev. 5 RA-10 (Threat Hunting) states:

> NIST SP 800-53 REV. 5 · RA-10
>
> *"Employ a cyber threat hunting capability to search for indicators of compromise in organizational systems and detect, track, and disrupt threats that evade existing controls at an organization-defined frequency."*

The control's discussion further clarifies that threat hunting is an active means of cyber defense, in contrast to passive measures like firewalls, intrusion detection, sandboxing, and SIEM. The objective is to track and disrupt cyber adversaries as early as possible in the attack sequence and measurably improve the speed and accuracy of response.

BREAKING DOWN THE REQUIREMENTS

REQUIREMENT	WHAT IT MEANS	COMPLIANCE CHALLENGE
Establish and maintain	Not a one-time exercise; an ongoing capability.	Requires dedicated resources and tooling that don't lapse.
Search for IOCs in organizational systems	Active investigation across data sources, not passive monitoring.	Requires query capability across EDR, identity, cloud, and network.
Detect, track, disrupt threats that evade existing controls	Find what other tools miss.	Requires expertise beyond standard alert response.
Organization-defined frequency	A cadence the organization must document and execute.	Creates audit trail requirements and demands consistency.

02 / CONTINUED

WHY TRADITIONAL APPROACHES FALL SHORT

STAFFING CONSTRAINTS

- Experienced threat hunters command $150,000 to $180,000 base salaries.
- A functional threat hunting team requires three or more FTEs for reasonable coverage.
- The cybersecurity talent shortage makes hiring difficult, even with budget approved.
- Turnover disrupts program continuity and resets institutional knowledge.

OPERATIONAL LIMITATIONS

- Manual threat hunting is time-intensive and pulls hours that don't scale.
- Analyst bandwidth competes constantly with alert triage demands.
- Consistent documentation is difficult to maintain when hunting is opportunistic.
- Hunt frequency depends on whatever cycles happen to be available that week.

AUDIT READINESS

- Ad-hoc hunting lacks structured evidence that auditors can review.
- Demonstrating "defined frequency" requires logs, reports, and reproducible artifacts.
- Proving a continuous capability is hard when the capability itself is intermittent.

> THE CORE TENSION
>
> RA-10 demands continuous, documented hunting. Traditional staffing produces intermittent, documentation-poor hunting. The gap between what the regulation requires and what the talent market can deliver is what drives most organizations to look for a different model.

03 / CUSTOMER CASE STUDY

HEALTHCARE SYSTEM

BACKGROUND

A major regional healthcare system with more than 300 affiliated clinics faced RA-10 compliance requirements as part of its broader NIST 800-53 program. The organization also needed to demonstrate threat hunting capability for HIPAA Security Rule compliance under 45 CFR §164.306 and §164.308.

The VP of Security Operations had budget approved for three threat hunting positions but faced significant challenges: he could not find qualified candidates in the current market, his existing SOC team was fully utilized on alert triage, he needed audit-ready documentation for compliance reviews, and he required continuous operation, not periodic exercises.

THE SOLUTION

The customer engaged 7AI to implement AI-powered threat hunting as part of his security operations program. Three pieces of the model carried the weight:

CONTINUOUS IOC CORRELATION

AI agents ingest threat intelligence from multiple sources, including commercial feeds, ISACs, and breaking research, and immediately search across the customer environment. No manual query building. No analyst scheduling required.

STRUCTURED HUNT EXECUTION

Each threat hunt follows a documented process: hypothesis or IOC specification, data source identification and query, correlation and analysis, finding determination, and recommended response actions.

AUDIT-READY DOCUMENTATION

Every hunt produces exportable records covering the original hypothesis, data sources queried, reasoning steps and queries executed, confidence-scored conclusions, and recommended or executed response actions.

HUMAN ESCALATION

When AI agents identify potential threats, cases are escalated to human analysts with full context. Humans make decisions; AI handles search and documentation.

03 / RESULTS

UNCOVERING A PHISHING CAMPAIGN

Within weeks of deployment, the platform demonstrated its value in a real incident.

INITIAL DETECTION

A **PLAID ELITE** analyst at 7AI reviewed an alert escalated by the platform. Investigation confirmed a phishing email was malicious based on content and attachment analysis.

THREAT HUNT EXECUTION

Working with a 7AI Security Engineer, the team conducted a deeper threat hunt to determine scope. The investigation revealed the email was not isolated. It was part of a multi-wave phishing campaign actively targeting the organization.

ATTACK SOPHISTICATION

The adversary had compromised a legitimate domain and properly configured SPF, DKIM, and DMARC authentication. That allowed phishing emails to bypass traditional email security controls and land directly in user inboxes.

OUTCOME

Without proactive threat hunting, only the initial alert would have been addressed. The hunt identified the full campaign scope, enabling comprehensive response.

> "THIS IS EXACTLY WHAT RA-10 IS DESIGNED TO CATCH. THREATS THAT EVADE EXISTING CONTROLS. WE TOLD THE CISO WE WOULDN'T NEED THREE HEADCOUNT, JUST ONE ANALYST PLUS THE PLATFORM. HE APPROVED IT IN A WEEK."
>
> VP, SECURITY OPERATIONS · REGIONAL HEALTH SYSTEM

04 / CAPABILITY MAPPING

7AI TO RA-10, LINE BY LINE

The table below maps each requirement of RA-10 directly to a 7AI capability and the documentation an auditor would expect to see.

RA-10 REQUIREMENT	7AI CAPABILITY	EVIDENCE
Establish a threat hunting capability	AI agents purpose-built for IOC hunting across multiple data sources.	Platform configuration, connector documentation.
Maintain the capability	Continuous operation; no dependence on analyst availability.	Uptime records, hunt frequency logs.
Search for IOCs in organizational systems	Automated query execution across EDR, identity, cloud, and network telemetry.	Hunt execution logs, query records.
Detect threats evading existing controls	AI analysis surfaces threats missed by alerting tools.	Hunt findings, incident correlation reports.
Track threats	Case management with timeline and evidence correlation.	Investigation records, case documentation.
Disrupt threats	Response recommendations and automated actions, with human approval gates.	Response logs, action records.
Organization-defined frequency	Configurable hunt cadence; continuous operation.	Scheduling configuration, execution logs.

PLAID, IN PRACTICE

7AI's People-Led, AI-Driven model means a dedicated AI Security Engineer customizes hunts and documentation to your environment. The platform executes; named human experts shape what it does and why. **PLAID ELITE** adds 24×7 human overwatch on top of that.

05 / HIPAA SECURITY RULE ALIGNMENT

WHAT THIS MEANS FOR HEALTHCARE

For healthcare organizations, RA-10 compliance also supports HIPAA Security Rule requirements. Three citations matter most.

45 CFR §164.306 / SECURITY STANDARDS

Threat hunting supports the requirement to "protect against any reasonably anticipated threats or hazards to the security or integrity" of electronic protected health information.

45 CFR §164.308 / ADMINISTRATIVE SAFEGUARDS

Threat hunting fulfills elements of the risk analysis and risk management requirements by proactively identifying threats that could impact ePHI before they cause harm.

AUDIT CONSIDERATIONS

HHS OCR auditors increasingly expect evidence of proactive security measures beyond reactive controls. Documented threat hunting demonstrates security program maturity in a way that point-in-time risk assessments simply cannot.

> FOR THE AUDIT BINDER
>
> An auditor asking for evidence of "established and maintained" threat hunting wants to see hunt cadence, IOC sources, query records, findings, and response actions. With AI-driven hunting, that evidence is generated as a byproduct of the work, not a separate exercise after the fact.

06 / COST ANALYSIS

THE MATH, IN PLAIN NUMBERS

TRADITIONAL APPROACH: DEDICATED THREAT HUNTING TEAM

COST ELEMENT	ANNUAL COST
3 FTE Threat Hunters (fully loaded)	$450,000 to $540,000
Tooling and platform costs	$50,000 to $100,000
Training and certification	$15,000 to $30,000
Management overhead	$30,000 to $50,000
Total	**$545,000 to $720,000**

Challenges: hiring difficulty, turnover risk, coverage gaps, documentation burden.

AI-AUGMENTED APPROACH: 7AI THREAT HUNTING

COST ELEMENT	ANNUAL COST
7AI Platform (threat hunting module)	Customer-specific
1 FTE analyst for escalation handling	$150,000 to $180,000
Total	**Significant reduction**

Benefits: continuous operation, built-in documentation, scales with threat intel volume.

WHAT THE MATH MISSES

Direct cost is only part of the story. The four points below are where security leaders tell us the difference shows up first.

- **Audit readiness.** Built-in documentation cuts compliance preparation time by an order of magnitude.
- **Consistent coverage.** No gaps from vacation, turnover, or competing priorities.
- **Scale.** Hunt volume can increase without a proportional cost increase.
- **Analyst elevation.** The existing team focuses on decisions, not searches. That's the work they were hired for.

07 / IMPLEMENTATION

HOW TO GET TO PRODUCTION

PREREQUISITES

1. **Data source connectivity.** Threat hunting requires access to EDR, identity, cloud, and network telemetry.
2. **Threat intelligence feeds.** IOC sources for hunt input, including commercial feeds and ISACs.
3. **Escalation workflow.** A defined process for human review of AI findings.
4. **Documentation requirements.** A clear understanding of what your auditors need to see.

DEPLOYMENT TIMELINE

Typical implementation: 7 to 14 days from kickoff to operational threat hunting.

PHASE	DURATION	ACTIVITIES
Discovery	Days 1 to 2	Data source inventory, compliance requirements review.
Configuration	Days 3 to 7	Connector deployment, hunt parameter setup.
Validation	Days 8 to 10	Test hunts, documentation review.
Production	Days 11+	Operational threat hunting, ongoing refinement.

SUCCESS METRICS

- Hunt execution frequency (daily, continuous).
- IOCs searched per period.
- True positive rate on findings.
- Mean time from IOC publication to hunt completion.
- Audit documentation completeness at the end of every quarter.

08 / CONCLUSION

A SUSTAINABLE PATH TO RA-10

NIST SP 800-53 Rev. 5 RA-10 requires organizations to maintain a continuous threat hunting capability. Traditional staffing approaches struggle to satisfy that requirement at the cadence and consistency the control demands.

AI-powered threat hunting offers a sustainable alternative that:

- Meets the control's requirement for an established, maintained capability.
- Executes hunts at organization-defined frequency without staffing constraints.
- Produces audit-ready documentation automatically as part of the work itself.
- Frees security analysts to focus on decisions, investigations, and the strategic work that requires human judgment.

For organizations facing RA-10, the question is no longer whether to implement threat hunting. The control is clear. The question is how to operationalize it in a way that's sustainable, auditable, and effective. For one major healthcare system, the answer was AI agents and one analyst, instead of three threat hunters they couldn't hire.

ABOUT 7AI

7AI's AI agents autonomously investigate security alerts, proven in production at Fortune 500 scale. The 7AI Platform includes AI-powered threat hunting as part of its agentic security capabilities, available both as a standalone platform and through **PLAID ELITE** fully managed security operations.

WEBSITE	TALK TO OUR TEAM	SEE IT IN ACTION
7ai.com	7ai.com/contact	7ai.com/demo