



Securing Tomorrow's Missions Today.



Is Someone Already Inside? A Fixed-Fee Baseline Threat Hunt on Your Own Telemetry

Threat Hunting & Advanced Persistent Threat Detection: test specific hypotheses against the logs you already hold, and leave with findings, a gap map and detections your SOC keeps.

AvalonTechServices.com

contact@AvalonTechServices.com

Executive Summary	2
Why Alerts Are Not Answers	3
Why It Concentrates in Government	3
The Mandates and the Calendar	4
What the Baseline Hunt Does	6
Deliverables	8
Where the Hunt Stops	8
Evidence Assessors Can Use	9
How an Engagement Runs	10
What the Agency Provides	10
What Is Not Included	11
How Avalon Staffs It	11
Pricing and How to Buy	12
Illustration: How Much of the Hunt Plan Your Telemetry Can Test	13
The Visibility Gap Map	14
What the Hunt Costs	15
Common Pitfalls	16
Conclusion and Next Step	16
Appendices	17
Appendix A: Glossary	17
Appendix B: Illustration Assumptions	17
Appendix C: References	18

Executive Summary

Most federal security operations centers are built to answer one question: what did the tools flag? The adversaries that matter most are designed to avoid that question. They sign in with valid credentials, use the same administrative tools the agency's own staff uses, and leave no malware signature behind. CISA's Volt Typhoon advisory (AA24-038A) and its living-off-the-land guidance describe exactly this tradecraft. An empty alert queue is not a clean environment.

Threat hunting asks the other question: what did the tools miss? A hunter starts from how real adversaries operate, writes specific hypotheses (if an intruder were using forged Kerberos tickets here, these events would appear in these logs), and tests each one against the agency's own telemetry. NIST SP 800-53 Rev. 5 names this as its own control, RA-10 (Threat Hunting), but does not allocate it to any 800-53B baseline. That is a large part of why agency hunting is uneven, and why many CISOs cannot tell an inspector general whether the network is compromised today.

This paper is for the people who carry that question: CISOs and deputy CISOs who fund the answer, SOC directors whose teams and tools the hunt runs on, ISSMs and owners of high value assets who need dated evidence for the authorization record, and the CORs who buy the work. It describes Avalon Technology Services' Threat Hunting & Advanced Persistent Threat Detection service, which starts with one engagement:

- **A fixed-fee Baseline Hunt.** One unclassified enclave of up to about 2,500 endpoints and one identity plane, hunted against a signed, ATT&CK-mapped hunt plan on the agency's existing EDR and SIEM, with agency-approved open-source collectors where coverage is thin. Six to seven weeks of work after one to three weeks of agency-controlled access setup, priced inside a \$60K to \$120K band.
- **Conditional follow-ons.** A Standard Hunt for larger or multi-enclave scopes, and scheduled business-hours hunt sprints with refreshed hypotheses, bought only if the agency chooses.

Section 9 works through an illustrative enclave with a 12-hypothesis hunt plan. The agency's telemetry can fully test 6 hypotheses (50%), partly test 3 and cannot test 3, because cloud audit logging was never turned on. The untestable rows become the logging roadmap. The fixed fee for that hunt is \$96K.

Every hunt ends with one of two results: validated findings with preserved evidence, or a scoped, dated, evidence-backed negative for the systems and time window examined. Either way the agency keeps the visibility gap map and the detection rules. The service

does not certify an environment clean, grant an authorization or respond to an active intrusion; those belong to the agency, its Authorizing Official and its incident response provider.

Why Alerts Are Not Answers

Detection tooling and hunting look for different things. Tools match known patterns and raise alerts. Hunting looks for behavior that fits an adversary's goals but trips no rule. The gap between them shows up in the same phrases across agencies.

What the CISO hears	What it usually means
"Our SIEM is a data lake nobody swims in."	Logs are collected and retained but never queried for anything the rules do not already catch
"We passed the pen test, but I can't tell the IG we're not compromised."	A penetration test finds exploitable paths; it does not look for an attacker already using them
"CISA hunted us once. Nothing since."	A point-in-time federal hunt, scheduled to national priorities, left no detections the agency owns
"If someone used valid credentials and PowerShell, would we see it?"	The living-off-the-land question; the honest answer depends on logging no one has checked
"I can't prove a negative, and I keep being asked to."	What the agency needs is a scoped, dated negative backed by evidence, not a certification

Why It Concentrates in Government

- **SOCs are built around tickets.** Most federal SOC's are staffed to work an alert queue. Hypothesis-driven searching has no ticket, so it rarely gets hours.
- **Logging is uneven.** Mixed on-premises and cloud estates mean no single console sees everything. Cloud audit logs are often off, short-retained or owned by another team.
- **The threat is patient and well resourced.** Federal networks are targeted by state actors who pre-position quietly and avoid anything a signature can catch.
- **Edge devices are a common way in.** End-of-support routers, firewalls and VPN appliances are a recurring entry point, which is why CISA BOD 26-02 now sets deadlines to inventory and retire them.

The Mandates and the Calendar

Driver	What it requires	What it means for a hunt
NIST SP 800-53 Rev. 5, RA-10	Threat hunting capability; not allocated to the 800-53B baselines	The direct control; arrives through agency policy or overlays, so evidence of it is often missing
SI-4, AU-6, CA-7, RA-3(3), IR-4(4), PM-16	System monitoring, audit review, continuous monitoring, dynamic threat awareness, incident correlation, threat awareness program	The controls a hunt exercises and evidences alongside RA-10
OMB M-26-14 (May 2026)	A risk-based logging model; replaced M-21-31	The hunt consumes logging and tests it; the gap map is written against the agency's M-26-14 priorities
CISA BOD 26-04 (Jun 10, 2026)	Risk-based remediation clocks, from 3 days with forensic triage to the next upgrade	Findings that touch exploited vulnerabilities feed the same clocks
CISA BOD 26-02 (Feb 6, 2026)	End-of-support edge devices: inventory in 3 months, decommission in 12, replace in 18	Edge devices still in service are a priority hunting surface
NIST SP 800-61 Rev. 3 (Apr 2025)	Incident response organized around CSF 2.0's six functions, including Govern	Defines the handoff when a hunt confirms an intrusion
Federal Incident Notification Guidelines	Report to CISA within one hour of identifying an incident	The agency reports; the hunt's stop, notify and preserve protocol is written to that clock
DoD CSRMC (Sep 24, 2025)	Cybersecurity risk management construct announced; RMF, DoDI 8510.01 and eMASS remain in use	For unclassified DoD corporate IT, hunt evidence goes into the existing RMF record

Hunts get bought at predictable moments: an advisory naming the agency's sector or installed technology, an intrusion at a peer agency, an inspector general finding on detection, the question "are they actually gone?" after an incident, an authorization decision on a high value asset, or a new CISO's baseline. Fiscal year-end (July through September) suits a fixed-fee purchase of this size.

What the Baseline Hunt Does

The Baseline Hunt is hypothesis-driven and technique-based, not an indicator sweep. Running a list of known-bad hashes and addresses through a SIEM is a scan. A hunt asks whether specific adversary behavior is present, and documents what could and could not be tested.

Workstream	Activities
Threat profile and hypotheses	An agency-specific threat profile from CISA advisories, sector targeting and the agency's own incident history, turned into testable hypotheses mapped to MITRE ATT&CK techniques
Visibility assessment (first)	What the agency can actually see: EDR coverage, log sources, retention and cloud audit configuration, documented before hunting so gaps are known rather than discovered mid-hunt
Endpoint hunting	Persistence, anomalous execution, credential-access artifacts and living-off-the-land use of administrative tools, through the agency's EDR console or Velociraptor where approved
Identity and directory hunting	Kerberoasting artifacts, forged ticket indicators, anomalous privileged logons, dormant-account use, suspicious federation and token activity
Network hunting	Zeek, flow and DNS metadata for beaconing, unusual east-west movement and covert channels
Cloud control-plane hunting	Microsoft 365 and Entra, AWS GovCloud or Azure Government audit logs: consent grants, mailbox rules, role escalation and service-principal behavior
Validation and triage	Every candidate finding validated with preserved evidence and classified: confirmed malicious, suspicious requiring an agency decision, or benign and explained
Detection engineering handoff	Hunt logic that produced value converted to Sigma rules, SIEM queries and EDR watchlists the SOC keeps

The hunt runs on the agency's existing stack (for example Microsoft Defender and Sentinel, CrowdStrike, SentinelOne, Splunk or Elastic), supplemented only with agency-approved open-source tools: Velociraptor for endpoint collection, Sigma for portable detections, YARA for artifact scanning and Zeek for network metadata. Analysis happens on agency consoles. No appliance is installed, no license is required, and no data leaves the agency boundary.

Deliverables

Deliverable	Format
Hunt plan: hypotheses, ATT&CK map, data sources, rules of engagement; signed before hunting starts	PDF or Markdown
Findings report, executive volume, with a defensible-negative statement where it applies	PDF, briefing-ready
Findings report, technical volume: per-finding evidence, timeline, assets, ATT&CK mapping, remediation priority	PDF plus a CSV or JSON finding index
Evidence package per confirmed finding: artifacts, hashes and a handling log	Retained on agency systems under agency records direction
Detection content pack	Native SIEM format plus Sigma source
Telemetry and visibility gap assessment, written against M-26-14 priorities and SI-4 and AU-6	PDF table and prioritized source list
Executive outbrief and SOC handoff session	Live briefing and slide deck

Where the Hunt Stops

If the hunt confirms an active intrusion, it stops, notifies and preserves. The hunt lead tells the agency's designated contact at once, preserves the evidence already gathered, and hands the finding to the agency's incident response provider. The agency makes the report to CISA within one hour under the Federal Incident Notification Guidelines. Containment, eviction and recovery are separate work under NIST SP 800-61 Rev. 3. Validation hours per finding are capped in the statement of work so a hunt does not turn into unpriced incident response.

Evidence Assessors Can Use

Control	Evidence the hunt produces
RA-10	Signed hunt plan, hypotheses tested and not tested, and findings, all dated
RA-3(3)	Threat profile drawn from current advisories and sector targeting
SI-4	Visibility assessment of monitoring coverage, and new detections deployed to the SIEM
AU-6	Documented audit-log review and correlation across endpoint, identity, network and cloud
IR-4(4)	Correlated findings handed to incident response with preserved evidence
CA-7	A dated point-in-time input to the continuous monitoring record
PM-16	Threat profile and hypotheses reusable by the agency's threat awareness program

In CSF 2.0 terms the work sits in Detect (DE.CM, DE.AE) with inputs to Identify (ID.RA). In RMF terms it supports the Monitor step and gives the Authorizing Official current, dated evidence about detection capability. These are implementer artifacts. The agency's assessor draws assessment conclusions, and the Authorizing Official decides.

Accurate	Not accurate
Documented, evidence-based implementation of RA-10, SI-4, CA-7 and AU-6	Makes the agency FISMA compliant
A point-in-time assessment of compromise indicators across in-scope systems	Certifies the environment clean
A scoped, dated, evidence-backed negative	"You are not compromised"
Evidence for the continuous monitoring record	Results in or speeds an ATO

How an Engagement Runs

Phase	Timing	What happens
Pre-engagement	1–3 weeks, agency-controlled	Statement of work and rules of engagement, account provisioning, console access, data-handling agreement
Kickoff and visibility assessment	Week 1	Confirm scope and telemetry reality, document coverage gaps, finalize the collection plan
Hypotheses and instrumentation	Weeks 1–2	Threat profile, hunt plan sign-off, approved collectors deployed
Active hunt sprints	Weeks 2–5	Iterative hunting across the four lanes; a weekly findings brief so nothing waits for the final report
Validation and reporting	Week 6	Evidence validation, report drafting, detection content packaging
Outbrief and closeout	Weeks 6–7	Executive and technical debriefs, SOC handoff, artifact turnover

Pre-engagement is the schedule's long pole, and the agency controls it. The statement of work ties delivery dates to an access service level, with day-for-day relief when accounts or console access arrive late.

What the Agency Provides

- A sponsor, usually the CISO or SOC director, who can unblock access decisions within days, not weeks.
- Read and query access to EDR and SIEM consoles, or approved data exports, and approval to deploy lightweight collectors where EDR coverage is missing.
- Network diagrams and a current asset inventory.
- A SOC point of contact for four to eight hours a week for context and finding triage.

- A named incident response contact and provider for the stop, notify and preserve handoff.

What Is Not Included

Included	Separate work
Stop, notify and preserve on a confirmed intrusion	Incident response, containment, eviction and recovery
Looking for evidence of actual attackers	Penetration testing, which simulates one
A bounded hunt with a start and an end	Continuous monitoring, a managed SOC or 24x7 coverage
Triage-level analysis of discovered binaries	Deep malware reverse engineering
Prioritized remediation guidance	Performing the fixes
Detection rules plus one working session with the SOC	Ongoing detection engineering
Unclassified enclaves	Classified enclaves and cross-domain work

How Avalon Staffs It

Avalon takes one engagement at a time, works business hours, and works in unclassified environments only. The Baseline Hunt is led by a hunt lead full time through validation, with a hunt analyst during the active sprints, part-time threat intelligence support up front, and Avalon's founder as project manager and compliance liaison. Avalon has no past federal engagements under this service, and it cites no credential, clearance or past engagement that is not current and verifiable.

Requirement	How it is handled
Hunt lead	Named personnel whose certifications and investigation level match the agency's requirements, listed in the proposal, hired or subcontracted
Hunt analyst	Named in the proposal with hunt experience on the agency's EDR and SIEM platforms, as a condition of the engagement
Civilian access	Public Trust investigation at the tier the agency requires, as a condition of the engagement
DoD privileged access	Personnel qualified under DoD 8140 for the relevant work role, as a condition of the engagement
Response coverage	None offered. No on-call bench or response-time commitment; a confirmed intrusion goes to the agency's IR provider
Insurance and contract vehicles	Any coverage the solicitation requires is a condition of award; Avalon does not hold a GSA Schedule or a governmentwide contract today

Pricing and How to Buy

Engagement	Market-norm range	Pricing basis
Baseline Hunt: one enclave, up to ~2,500 endpoints, one identity plane	\$60K–\$120K	Firm fixed price
Standard Hunt: up to ~15,000 endpoints or 2–3 enclaves, one cloud tenant (conditional)	\$120K–\$300K	Firm fixed price
Scheduled hunt sprints after a baseline (conditional)	Scoped after the baseline	Firm fixed price per sprint

Both hunt tiers fit under the \$350,000 simplified acquisition threshold, so they can be bought as open-market purchases. Purchases between the \$15,000 micro-purchase threshold and the SAT are generally set aside for small business when two or more

capable small businesses are expected to compete. Hunting inside a larger security operations task order can be subcontracted to Avalon by the prime. Avalon does not hold a GSA Schedule or a governmentwide contract today.

Adding the hunt to the incumbent SOC contract has one structural weakness: the team that runs the detections would be testing its own work. Contingency pricing is declined, because a hunt recovers no dollars to share; a fixed price against a named scope keeps incentives clean.

Illustration: How Much of the Hunt Plan Your Telemetry Can Test

The figures below are a planning illustration for one representative unclassified enclave, not a quote and not a result from any engagement. Every hypothesis, data-source status and hour is an assumption.

The hunt plan holds 12 hypotheses, three per lane, each mapped to an ATT&CK technique and tagged with the data sources it needs. The visibility assessment marks each source **available**, **partial** (part of the enclave, or retained too briefly) or **absent**. A hypothesis is only as testable as its weakest required source.

Lane	Hypothesis	ATT&CK	Result
Endpoint	Persistence through scheduled tasks or services on servers	T1053	Testable
Endpoint	Administrative tools used outside approved admin hosts (living off the land)	T1059.001	Testable
Endpoint	Credential dumping from LSASS memory	T1003.001	Testable
Identity	Kerberoasting of service accounts	T1558.003	Testable
Identity	Forged Kerberos tickets (golden or silver)	T1558.001	Testable
Identity	Dormant or departed-staff accounts signing in	T1078.002	Testable
Network	Periodic beaconing to rare external domains	T1071.004	Partially testable
Network	Unusual east-west remote administration between workstations	T1021.002	Partially testable

Lane	Hypothesis	ATT&CK	Result
Network	Tunneling or covert channels over permitted protocols	T1572	Partially testable
Cloud	Illicit OAuth consent grants to unfamiliar applications	T1528	Untestable
Cloud	Inbox forwarding rules created to external addresses	T1114.003	Untestable
Cloud	Privileged role or service-principal escalation	T1098.003	Untestable

Endpoint and identity telemetry are healthy, so all 6 of those hypotheses are testable. Network metadata covers only part of the enclave, so the 3 network hypotheses can be tested only where sensors exist. Cloud tenant and mailbox audit logs are not collected, so the 3 cloud hypotheses cannot be tested at all. Overall, 6 of 12 (50%) are fully testable.

That result is itself a finding. The report states which hypotheses were and were not tested, so the negative is honest about its boundaries.

The Visibility Gap Map

Data source	Status	Hypotheses blocked	Priority
Cloud tenant audit logs (M365/Entra or IaaS control plane)	Absent	T1528, T1114.003, T1098.003	1
Mailbox audit logs	Absent	T1114.003	2
DNS query logs	Partial	T1071.004, T1572	3
East-west flow or Zeek metadata	Partial	T1021.002, T1572	4

Turning on cloud tenant audit logging is first because it unblocks a whole lane, and in most tenants it is configuration, not a purchase. Extending network sensors to the rest of the enclave comes next. Each row maps to SI-4 and AU-6 and is written against the

agency's M-26-14 logging priorities, so it can go straight into a plan of action or a budget request.

What the Hunt Costs

Scenario	Hours	Rate-band range	Fixed fee	Tier band
Lean (4 weeks active)	339	\$54K–\$88K	\$68K	Baseline \$60K–\$120K
Base (Baseline Hunt, one enclave)	479	\$77K–\$125K	\$96K	Baseline \$60K–\$120K
Standard (8 weeks, two analysts)	923	\$148K–\$240K	\$185K	Standard \$120K–\$300K

The base case is 240 hunt-lead hours (six weeks), 160 analyst hours (four weeks), 20 threat intelligence hours, 49 project management hours at about 18% over seven weeks, and 10 compliance hours: 479 hours in all. At a \$200 planning rate that is \$96K, inside the Baseline band. Across the full \$160 to \$260 rate band it runs \$77K to \$125K, so the Baseline band holds at blended rates up to about \$250 an hour. A lean four-week hunt comes in near \$68K. A Standard scope with two analysts over eight weeks prices near \$185K, inside the Standard band. Every scenario sits under the simplified acquisition threshold.

Common Pitfalls

Pitfall	How to avoid it
A hunt that turns into free incident response	Stop, notify and preserve; cap validation hours per finding; hand off to the IR provider
Endpoint and enclave creep	Name the enclaves and the endpoint band; price additions by modification
Unpaid detection engineering	Cap the handoff at the content pack plus one working session
"Keep watching" drift after the outbrief	A bounded end date; ongoing hunting only as scheduled sprints
Overselling the compliance value	Use the accurate phrasing; the AO and the assessor decide
Access friction eating the schedule	Visibility assessment first and an access service level with day-for-day relief

Conclusion and Next Step

An agency cannot answer "are we compromised?" from an alert queue. It can answer a narrower and more useful question: for these systems, in this window, did we find evidence of these specific adversary behaviors, and what could we not see? That answer, backed by evidence and dated, is what an inspector general, an Authorizing Official and a CISO can actually use.

The natural first step is the fixed-fee Baseline Hunt: one unclassified enclave, under the simplified acquisition threshold, run on the agency's own telemetry, ending in validated findings or a defensible negative, a visibility gap map and detections the SOC keeps. To discuss a hunt, contact Avalon Technology Services at contact@AvalonTechServices.com.

Appendices

Appendix A: Glossary

ATT&CK

MITRE's knowledge base of adversary tactics and techniques, used to frame hypotheses and report findings.

Defensible negative

A scoped, dated statement that no evidence of the tested behaviors was found in the systems and window examined, with the untested areas named.

Hypothesis

A specific, testable statement about adversary behavior and the evidence it would leave in named data sources.

Living off the land (LOTL)

Using legitimate administrative tools and valid credentials to avoid detection.

Sigma

An open, vendor-neutral format for writing detection rules that convert to SIEM queries.

Velociraptor

An open-source endpoint collection and triage tool, deployed only with agency approval.

Appendix B: Illustration Assumptions

All values are planning assumptions for one representative unclassified enclave. None is drawn from a specific agency or a completed engagement.

Assumption	Value	Basis
Hunt plan	12 hypotheses, 3 per lane, each mapped to one ATT&CK technique	Planning assumption
Telemetry status	Endpoint and identity available; network partial; cloud tenant and mailbox audit absent	Planning assumption
Testability rule	A hypothesis takes the weakest status of its required sources	Model rule
Team loading	Hunt lead full time; analyst during sprints; TI front-loaded; PM about 17.5%; compliance in reporting	KB delivery model
Rates	\$160–\$260 an hour fully burdened; fee at a \$200 planning rate	KB market-norm band
Durations	Pre-engagement 1–3 weeks; work weeks 1–7	KB phase durations
Price check	Base fee \$96K	Inside the \$60K–\$120K Baseline band

Appendix C: References

1. **Office of Management and Budget**, *M-26-14, Ensuring Effective and Efficient Agency Logging and Network Visibility to Defend Against Evolving Cyber Threats*, May 2026
<https://www.whitehouse.gov/wp-content/uploads/2026/05/M-26-14-Ensuring-Effective-and-Efficient-Agency-Logging-and-Network-Visibility-to-Defend-Against-Evolving-Cyber-Threats.pdf>
2. **National Institute of Standards and Technology**, *SP 800-53 Rev. 5, Security and Privacy Controls for Information Systems and Organizations*
<https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final>
3. **National Institute of Standards and Technology**, *SP 800-61 Rev. 3, Incident Response Recommendations and Considerations for Cybersecurity Risk*

Management, April 2025

<https://csrc.nist.gov/pubs/sp/800/61/r3/final>

4. **National Institute of Standards and Technology**, *Cybersecurity Framework 2.0*
<https://www.nist.gov/cyberframework>
5. **Cybersecurity and Infrastructure Security Agency**, *AA24-038A, PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure*
<https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>
6. **Cybersecurity and Infrastructure Security Agency**, *Federal Incident Notification Guidelines*
<https://www.cisa.gov/federal-incident-notification-guidelines>
7. **Tenable**, *CISA BOD 26-04 FAQ: Vulnerability Remediation Impact, June 11, 2026*
<https://www.tenable.com/blog/cisa-bod-26-04-FAQ-vulnerability-remediation-impact>
8. **MITRE**, *ATT&CK*
<https://attack.mitre.org>
9. **MITRE**, *D3FEND*
<https://d3fend.mitre.org>
10. **SigmaHQ**, *Sigma detection rule format*
<https://sigmahq.io>
11. **Centre Law & Consulting**, *Inflation Updates to FAR Thresholds: What Changed on October 1, 2025*
<https://centrelawgroup.com/blog/inflation-updates-to-far-thresholds-what-changed-on-october-1-2025/>