



Securing Tomorrow's Missions Today.



Owning EDR Isn't Operating It: A Fixed-Fee EDR Coverage and Efficacy Assessment for Agency CISOs

Endpoint Detection & Response: measure what your platform covers, test what it catches, and fix the gaps in the right order.

AvalonTechServices.com

contact@AvalonTechServices.com

Executive Summary	2
Owned Is Not Operating	3
Why It Concentrates in Government	3
The Mandates and the Calendar	4
What the Coverage and Efficacy Assessment Does	6
Deliverables	8
Four Numbers, Measured Not Promised	8
What Follows, Conditionally	8
Evidence Assessors Can Use	9
How an Engagement Runs	10
What the Agency Provides	11
What Is Not Included	11
How Avalon Staffs It	12
Pricing and How to Buy	13
Illustration: The Owned-Versus-Operating Gap	13
What the Assessment 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

Nearly every federal agency now owns an endpoint detection and response platform, often bought through an enterprise agreement or provided through the Continuous Diagnostics and Mitigation (CDM) program. Owning the platform is not the same as operating it. In many fleets a share of devices never got an agent, some agents stopped reporting, prevention policies still sit in audit mode, and nobody has run a controlled test to see what the console would actually catch.

Three questions expose the gap. What percentage of the authoritative inventory reports to the console? When was detection last tested, and what was missed? What is the false-positive rate the analysts live with? If the answers are a vendor dashboard number, "never" and "too high to count," the platform is paid for but not yet doing its job.

This paper is for CISOs who report a coverage number, SOC directors who live with the alert volume, ISSMs and ISSOs who need SI-3, SI-4 and IR-4 evidence, CIOs consolidating endpoint tools, and the CORs who buy the work. It describes Avalon Technology Services' Endpoint Detection & Response service, which starts with one engagement:

- **A fixed-fee EDR Coverage and Efficacy Assessment.** Inventory reconciliation, an agent-health audit, a policy and exclusion review, a controlled ATT&CK-mapped detection test, a telemetry review and a costed remediation roadmap, on the platform the agency already licenses. Three to six weeks from kickoff, inside the \$25K to \$75K market band for this work, and well under the simplified acquisition threshold.
- **Conditional follow-ons.** A tuning sprint, a small deployment or migration, or a quarterly validation retest, each scoped from the assessment and bought only if the agency chooses.

Section 9 works through an illustrative component of 6,000 endpoints. Its console reports 84% of the inventory, 92% of reporting agents are healthy, and only 67% of the fleet is protected by a healthy agent in blocking mode. The fixed fee for that assessment in the base case is \$50K.

The service implements, tests and evidences. It does not grant an authorization, certify compliance or assess its own work. The agency's assessor draws assessment conclusions and the Authorizing Official makes the risk decision.

Owned Is Not Operating

An EDR platform protects only the endpoints it is installed on, reporting from and allowed to act on. The gap between owning and operating shows up in the same ways across agencies.

What people say	What it usually means
"We're covered. We have [vendor]."	The platform is owned. Nobody has measured deployment against the inventory or tested what it detects.
"The CDM dashboard says one number; our console says another."	Inventory sources disagree. The coverage figure in the FISMA report rests on whichever source was easiest to pull.
"Our analysts ignore half the alerts."	Default policies and broad exclusions produce noise. Real detections drown in it.
"We can't put agents on those systems."	Coverage exceptions nobody decided deliberately and nobody documented, often laboratory, medical or instrument-attached devices.
"The IG asked how fast we'd catch an intrusion."	No efficacy evidence exists. The honest answer today is "we don't know."

Why It Concentrates in Government

- **The tooling arrived faster than the operations.** Post-2021 funding pushed licenses out quickly. Tuning and validation are skilled, non-recurring work that loses to daily alert triage.
- **Federal fleets are mixed.** Legacy server builds, virtual desktops, Linux and macOS pockets, cloud workloads and devices that cannot take an agent. Each is a coverage decision.
- **Every policy change rides the change board.** Moving a policy group from audit mode to blocking mode needs a change approval. A monthly change board turns a few weeks of tuning into months, so many agencies stay in audit mode indefinitely.

The Mandates and the Calendar

Driver	What it sets	What it means for EDR
OMB M-22-01 (Oct 8, 2021)	Issued in 2021 to direct federal civilian EDR deployment and CISA access to EDR data	The origin of most civilian deployments; cited here as history, not as a current requirement
OMB M-26-14 (May 2026)	Risk-based event logging model; replaced M-21-31	Endpoint telemetry, retention and SIEM forwarding are judged against the agency's risk-based logging approach and the CISA Logging Reference Architecture
CISA BOD 26-04 (Jun 10, 2026)	Risk-based remediation clocks from 3 days with forensic triage to the next upgrade; replaced BOD 22-01 and BOD 19-02	EDR exposure and vulnerability findings feed the same remediation queue the clocks measure
NIST SP 800-61 Rev. 3 (Apr 2025)	Incident response organized around the six NIST CSF 2.0 functions, including Govern	The frame for SOC runbooks and the detection-to-escalation handoff
CISA Federal Incident Notification Guidelines	Report incidents to CISA within one hour of identification	Detection speed decides whether the clock can be met
CISA ZTMM 2.0, DoD Zero Trust Strategy	Devices pillar; DoD target level by end of FY2027	EDR is the core of the devices pillar
DFARS 252.204-7012, NIST SP 800-171	72-hour incident reporting for defense contractors; practices 3.3, 3.6 and 3.14	Same detection and audit practices on the contractor side; CMMC Phase 2 was suspended Jul 13, 2026, Phase 1 self-assessment continues

Driver	What it sets	What it means for EDR
DoD CSRMC (Sep 24, 2025)	Cybersecurity risk management construct announced; RMF, DoDI 8510.01 and eMASS remain in use	EDR evidence still goes into the RMF record the component uses today

Three moments usually start a purchase: an inspector general or GAO finding on detection or coverage, which creates a funded POA&M item; an EDR contract expiry or consolidation that forces a migration; and fiscal year-end (July through September), which favors assessment-sized buys with one-year funds.

What the Coverage and Efficacy Assessment Does

The entry engagement answers three questions for one component on the platform it already licenses: what the platform covers, what it catches, and what to fix first. Work is vendor-agnostic and stays inside the agency's console, under a signed statement of work and rules of engagement.

Workstream	Activities
Inventory reconciliation	Reconcile the authoritative inventory (Active Directory, CDM asset data, CMDB) against what the console reports, producing a coverage percentage and a named gap list by endpoint class: workstations, member servers, domain controllers, Linux, macOS, VDI and cloud workloads
Agent-health audit	Stale check-ins, broken sensors, version drift, and agents installed but left in passive or audit mode
Policy and exclusion review	Prevention and detection policies against vendor baselines and CIS Benchmarks; every exclusion and allowlist entry listed with its justification, or the lack of one
Controlled efficacy test	Safe adversary emulation with Atomic Red Team or MITRE Caldera test cases mapped to MITRE ATT&CK techniques, on agency-provided test devices; records what fired, what was blocked, what was missed and how long detection took
Telemetry review	Endpoint telemetry, retention and SIEM forwarding against the agency's M-26-14 logging approach and the CISA Logging Reference Architecture
Remediation roadmap	Prioritized and costed, written so any qualified vendor, the incumbent or the agency's own staff can execute it

Deliverables

Deliverable	Format
Coverage and Gap Assessment Report: coverage, agent health, policy audit, efficacy results	PDF report and executive briefing
Gap register, one row per gap with endpoint class, cause and recommended fix	Excel
ATT&CK baseline: technique-by-technique results of the efficacy test	PDF and ATT&CK Navigator layer file
Costed remediation roadmap	Phased spreadsheet and briefing usable as budget justification

Four Numbers, Measured Not Promised

Every engagement reports four numbers at the start and, if follow-on work is bought, at the close: **coverage** against the authoritative inventory, **agent health**, **median time to detect** in controlled testing, and **false-positive rate**. They are measured on the agency's fleet, never promised in advance, because results depend on agency decisions and change approvals.

What Follows, Conditionally

If the agency chooses to continue, the roadmap scopes the follow-on work. A **tuning sprint** moves policy groups from audit to blocking mode with defined soak periods and rollback criteria, engineers precise exclusions with system owners, and measures alert fidelity before and after. A **small deployment or migration** extends coverage through the agency's existing management tooling (Intune, MECM, GPO, BigFix or Ansible) in pilot rings and change-board-aligned waves, with a dual-agent plan when moving between platforms. A **quarterly validation retest** reruns the same ATT&CK-mapped suite to show whether coverage and detection held. None of this is committed by buying the assessment.

Evidence Assessors Can Use

Control	Evidence the work produces
SI-3	Malicious code protection configuration, policy exports and coverage by endpoint class
SI-4, SI-4(2), (4), (5)	Monitoring coverage, automated analysis settings, inbound and outbound traffic visibility on hosts, and alert routing
IR-4, IR-5, IR-6	Detection-to-escalation runbooks, incident tracking hooks and the reporting path to the agency's incident team
AU-2, AU-6, AU-12	Endpoint event selection, review workflow and SIEM forwarding aligned to the M-26-14 approach
CM-8	Agent-derived inventory reconciliation against the authoritative source
CA-7	Repeatable ATT&CK-mapped test results that can feed continuous monitoring

For systems that handle CUI, the same work maps to NIST SP 800-171 requirements 3.3.1, 3.3.2 and 3.3.5, 3.6.1 and 3.6.2, and 3.14.2 through 3.14.7. These are implementer artifacts. If the EDR console is a cloud service, its FedRAMP status covers the vendor's offering, not the agency's systems; under the FedRAMP Consolidated Rules (June 25, 2026) offerings are "FedRAMP Certified" in classes A through D, and each vendor's status should be checked on the FedRAMP Marketplace. The agency still authorizes its own systems.

Accurate	Not accurate
Implements and validates EDR capability addressing SI-3, SI-4 and IR-4, with SSP-ready implementation evidence	Achieves compliance
Closes endpoint coverage and telemetry gaps	Satisfies M-22-01
Produces evidence for the assessor and the authorization record	Delivers your ATO
Supports CMMC practice implementation	Makes you CMMC-certified

How an Engagement Runs

Phase	Typical duration	What happens
Kickoff and access	1–2 weeks	Console access, points of contact, change-board calendar, inventory sources, rules of engagement signed
Assess and baseline	2–4 weeks	Reconciliation, agent health, policy review, efficacy test, telemetry review, roadmap
Tuning sprint (optional follow-on)	4–8 weeks	Audit-to-block progression with soak and rollback, exclusion engineering, alert-fidelity measurement
Small deployment or migration (optional)	8–12 weeks	Pilot ring of about 5% across every hardware and OS class, then waves aligned to the change board
Closeout (with follow-ons)	1–2 weeks	Closing ATT&CK retest, runbooks, one tabletop, recorded knowledge transfer

The standalone assessment runs three to six weeks from kickoff. Follow-on calendars are set more by the change board than by endpoint count, and each disconnected enclave adds two to four weeks.

What the Agency Provides

- A sponsor, usually the CISO or SOC director, with authority to decide audit-to-blocking questions.
- Console access, or standing supervised sessions, provisioned in the first two weeks. The schedule starts when access is granted, not at award.
- Endpoint-management and directory administrators for four to eight hours a week, and representative test devices for every hardware and OS class.
- The SSP, an inventory export, network diagrams, and the current EDR contract and license terms.
- Change-board slots for any follow-on work. The statement of work lists these as dependency milestones with schedule relief attached.

What Is Not Included

Included	Separate work
Assessment, tuning and deployment in the agency's licensed platform	License procurement or resale
Runbooks and knowledge transfer for the agency's SOC	24x7 monitoring or managed detection and response
Stopping and handing off on evidence of an active intrusion	Active incident response and forensic acquisition
One-time ATT&CK-mapped efficacy testing	Ongoing threat hunting
EDR integration with the existing SIEM	SIEM build-out or vulnerability scanning programs
Compensating-control recommendations for agent-ineligible devices	Network detection, OT or medical-device monitoring
Unclassified environments	Classified enclaves and cross-domain work

No console access, data collection or active testing happens before a signed statement of work and rules of engagement. If the assessment turns up evidence of an intrusion, testing stops and the agency's incident response process takes over, as the statement of work states.

How Avalon Staffs It

Avalon takes one engagement at a time, works in unclassified environments only, and works business hours. It offers no response-time commitment it cannot staff. The assessment is delivered by a named EDR technical lead at full time, who also runs the efficacy test, with Avalon's founder as project manager and compliance liaison at a quarter-time. 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
Technical lead	Named personnel whose platform certifications and investigation level match the agency's requirements, hired or subcontracted and listed in the proposal
Efficacy testing	A test lab and an ATT&CK-mapped test library in place before the first test, 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
Live response, forensics, retainers	Not offered today. Any future offer requires named responders, tooling, evidence storage and insurance in place first
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
EDR Coverage and Efficacy Assessment	\$25K–\$75K	Firm fixed price
Tuning sprint (conditional follow-on)	\$40K–\$120K	Firm fixed price
Deployment or migration, under 2,500 endpoints (conditional)	\$75K–\$150K	Firm fixed price from a labor build-up
Quarterly validation retest (conditional)	\$15K–\$40K per quarter	Fixed price per retest; no response commitment

The assessment, a tuning sprint and a small deployment all fit under the \$350,000 simplified acquisition threshold, so they can be bought as open-market purchases with one-year funds. 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. Larger deployments, or EDR work inside a security operations support contract, can be subcontracted to Avalon by the prime. Avalon does not hold a GSA Schedule or a governmentwide contract today.

Pricing is a labor build-up with endpoint classes and enclave count named in the statement of work. Performance-based pricing tied to coverage or detection results is declined, because those results depend on agency decisions the contractor does not control.

Illustration: The Owned-Versus-Operating Gap

The figures below are a planning illustration for one representative component, not a quote and not a result from any engagement. Every count is an assumption.

For each endpoint class the assessment counts inventory, reporting, healthy and blocking endpoints. Three ratios follow: **coverage** (reporting divided by inventory), **agent health** (healthy divided by reporting) and **effective protection** (blocking divided by inventory).

Endpoint class	Inventory	Reporting	Healthy	Blocking	Gap
Workstations	4,000	3,700	3,450	3,100	300
Member servers	800	680	600	420	120
Domain controllers	20	20	20	20	0
Linux servers	400	220	180	120	180
macOS	150	90	80	80	60
VDI	500	300	280	280	200
Cloud workloads	130	40	35	20	90
Total	6,000	5,050	4,645	4,040	950

Endpoint class	Coverage	Agent health	Effective protection
Workstations	92%	93%	78%
Member servers	85%	88%	52%
Domain controllers	100%	100%	100%
Linux servers	55%	82%	30%
macOS	60%	89%	53%
VDI	60%	93%	56%
Cloud workloads	31%	88%	15%
Total	84%	92%	67%

Gap: endpoints in the inventory that do not report to the console. Effective: healthy agents in blocking mode as a share of inventory.

Read across the totals. Against the authoritative inventory, coverage is 84%: 950 endpoints never report. Of the agents that do report, 405 are unhealthy, and another

605 are healthy but still in audit mode. Only 4,040 endpoints, 67% of the fleet, have a healthy agent that will act on what it sees. In all, 1,960 endpoints sit outside effective protection.

The classes tell different stories. Domain controllers are fully covered. Workstations look healthy at 92% coverage but still account for the largest raw gap. Linux servers report at 55% and cloud workloads at 31%, usually images built without the agent. Member servers have reasonable coverage but only 52% effective protection, the typical pattern where mission-system owners asked for audit mode and nobody revisited it.

Median time to detect and false-positive rate are not modeled; the assessment measures them.

What the Assessment Costs

The cost check uses one technical lead at full time and the founder at a quarter-time, priced across the catalog's blended rate band of \$185 to \$320 an hour, with the fixed fee shown at a planning rate of \$250 an hour.

Scenario	Weeks	Hours	Rate-band range	Fixed fee
(a) Clean inventory	4	200	\$37K–\$64K	\$50K
(b) No authoritative inventory	5	250	\$46.25K–\$80K	\$62.5K
(c) Many agent-ineligible devices	4.5	225	\$41.62K–\$72K	\$56.25K

In the base case, four weeks (one of kickoff, three of assessment) is 200 hours, \$37K to \$64K across the rate band, with a fixed fee of \$50K: inside the \$25K to \$75K market band and far under the simplified acquisition threshold. Avalon's price-to-win target of \$25K to \$40K for small and mid environments is reachable at this scope only at about \$200 an hour or less, or with a narrower three-week scope (150 hours, \$27.75K to \$48K). The fee is never set below cost to hit the target.

Scenario (b) shows why the inventory question gets priced explicitly. Without an authoritative inventory, reconciling directory, CDM and CMDB sources becomes the first work item and adds a week; at the top of the rate band the range reaches \$80K, above the \$75K market ceiling, so the statement of work should name it as a separate line. Scenario (c), a fleet with many agent-ineligible devices, adds a compensating-control register for each excluded class and stays inside the band at \$56.25K. The

recommended structure is a four-week firm fixed price with endpoint classes and enclave count named in the statement of work.

Common Pitfalls

Pitfall	How to avoid it
Scope creep: cloud workloads, network-team dashboards, one more enclave, patching	Name endpoint classes and enclaves in the statement of work; price additions from a unit-priced change menu
Blocking-mode haste	Audit-mode soak periods and rollback criteria per policy group, written into the contract
Overselling compliance	Describe the work as implemented and evidenced controls; the assessor and the AO draw the conclusions
Agency dependencies slipping silently	Dependency milestones with schedule relief for access, test devices and change-board slots
An incident discovered mid-test	Stop testing and hand off to the agency's incident response process, as the statement of work states

Conclusion and Next Step

An EDR platform is a capability only on the endpoints it reports from, stays healthy on, and is allowed to act on. Most agencies have never measured those three things against an authoritative inventory, and fewer have tested what the console catches. A short, fixed-fee assessment turns "we have EDR" into four measured numbers, a named gap list and an order of work the agency can fund and compete.

The natural first step is the EDR Coverage and Efficacy Assessment: one component, three to six weeks, a defined scope under the simplified acquisition threshold, and a roadmap any qualified vendor can execute. To discuss an assessment, contact Avalon Technology Services at contact@AvalonTechServices.com.

Appendices

Appendix A: Glossary

EDR

Endpoint detection and response: an agent and console that record endpoint activity, detect malicious behavior and can block or isolate.

Audit mode

A prevention policy that records what it would have blocked without blocking it.

Agent health

Whether a reporting agent is current, checking in and functioning, rather than stale, broken or out of date.

Coverage

In this paper, endpoints reporting to the console divided by endpoints in the authoritative inventory.

Effective protection

In this paper, endpoints with a healthy agent in blocking mode divided by the authoritative inventory.

MITRE ATT&CK

A public knowledge base of adversary tactics and techniques used as the coverage language for efficacy testing.

Authorizing Official (AO)

The senior official who accepts risk and authorizes a system to operate.

Appendix B: Illustration Assumptions

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

Assumption	Value	Basis
Fleet	6,000 endpoints in seven classes	Planning assumption
Class counts	Inventory, reporting, healthy and blocking counts as shown in Section 9	Planning assumption
Team loading	Technical lead 1.0 FTE; founder as PM and compliance liaison 0.25 FTE; 40 hours a week	KB delivery model, adapted
Rates	\$185–\$320 an hour blended; fee at \$250	Catalog rate card
Durations	Kickoff 1–2 weeks; assess 2–4 weeks; standalone 3–6 weeks	KB phase durations
Price check	Base fee \$50K	Inside the \$25K–\$75K market band
Not modeled	Median time to detect; false-positive rate	Measured in the engagement

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. **Office of Management and Budget**, *M-22-01, Improving Detection of Cybersecurity Vulnerabilities and Incidents on Federal Government Systems through Endpoint Detection and Response*, October 8, 2021
<https://www.whitehouse.gov/wp-content/uploads/2021/10/M-22-01.pdf>
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**, *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>
5. **National Institute of Standards and Technology**, *SP 800-171 Rev. 2, Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations*
<https://csrc.nist.gov/pubs/sp/800/171/r2/upd1/final>
6. **Cybersecurity and Infrastructure Security Agency**, *Federal Incident Notification Guidelines*
<https://www.cisa.gov/federal-incident-notification-guidelines>
7. **Cybersecurity and Infrastructure Security Agency**, *Zero Trust Maturity Model, Version 2.0*
<https://www.cisa.gov/zero-trust-maturity-model>
8. **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>
9. **MITRE**, *ATT&CK*
<https://attack.mitre.org/>
10. **MITRE**, *Caldera*
<https://caldera.mitre.org/>
11. **FedRAMP**, *FedRAMP Marketplace*
<https://marketplace.fedramp.gov/>
12. **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/>