SHIELD AI

AI SECURITY · DLP · GOVERNANCE

EXECUTIVE BRIEFING

Executive AI Governance Report

CinderLabs

Apr 6 — May 6, 2026

CONFIDENTIAL · BOARD / CXO

Prepared for: Executive Leadership & Audit Committee
Prepared by: SHIELD AI Security Platform
Report ID: EXEC-20260506-8QW43P
Date issued: May 6, 2026

1 · EXECUTIVE SUMMARY

During the reporting period (Apr 6 – May 6, 2026), SHIELD processed **2789** security events across **14** endpoints at CinderLabs. Event volume increased **NEW** versus the prior 30-day window, driven primarily by 970 DLP violations and 1794 shadow-AI detections. **1942** critical/high-severity events were recorded (prior period: 0). 5% of all events were resolved; 2565 remain open and require triage. Control effectiveness: **92%** of DLP-flagged uploads were blocked at the endpoint. 25 user overrides were recorded and warrant leadership review. The most-exposed destination was **us-east-1.console.aws.amazon.com**; the rule most frequently triggered was **JWT Token**. **31 AI services** observed in traffic have no governing policy — representing an immediate governance gap. Fleet coverage stands at **7%** — 13 endpoints have not reported a heartbeat in 72h, creating a monitoring blind spot.

Threat Index

100 **Critical**

Composite: severity × volume + overrides + coverage gaps. Scale 0–100.

Key Metrics (30-Day vs Prior 30-Day)

2789	1942	970	1794
TOTAL EVENTS	CRITICAL + HIGH	DLP VIOLATIONS	SHADOW AI
NEW vs prior	NEW vs prior	NEW vs prior	NEW vs prior

92%	25	5%	7%
BLOCK RATE	USER OVERRIDES	RESOLVED RATE	FLEET COVERAGE
of flagged uploads	requires review	131 of 2789	1 of 14 reporting

30-Day Event Volume

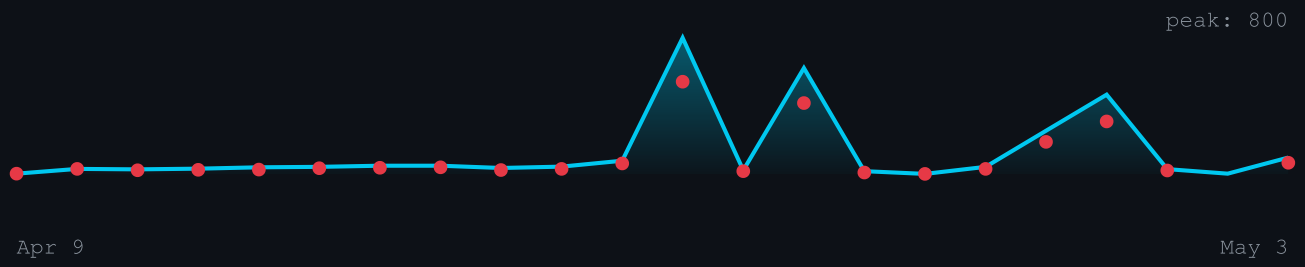

Red dots indicate days with Critical / High severity events.

2 · RISK POSTURE & CONTROL EFFECTIVENESS

Severity Breakdown

343	1599	847	0	2789
CRITICAL	HIGH	MEDIUM	LOW	TOTAL

Event Pipeline Status

2565	90	131	3
NEW (AWAITING TRIAGE)	ACKNOWLEDGED	RESOLVED	FALSE POSITIVE

Alert pipeline backlog: 2565 events remain in *new* status — SLA review recommended.

Endpoint Control Effectiveness

92%	0%	25
BLOCKED AT ENDPOINT	ALLOWED (MONITORED)	OVERRIDDEN BY USER
913 of 995	0 of 995	3% of decisions

Recent User Overrides (sample)

DESTINATION	RULE	HOST	WHEN
ChatGPT	Social Security Number	browser-147-mo1vpq5c	Apr 24
ChatGPT	Social Security Number	browser-147-mo1vpq5c	Apr 23
www.bing.com	Social Security Number	browser-147-mo1vpq5c	Apr 19
ChatGPT	Proprietary Code Marker	browser-147-mo1vpq5c	Apr 16
Copilot	Social Security Number	browser-147-mo1vpq5c	Apr 16
Claude	Google API Key	Mac	Apr 16

3 · DATA EXPOSURE ANALYSIS

This section breaks down **what data** was at risk, **which rules** triggered, and **where** it was headed. Use this to prioritize policy tuning and user awareness.

Top AI Services & Destinations (events)

SERVICE / DESTINATION	EVENTS
us-east-1.console.aws.amazon.com	1300
www.loom.com	573
ChatGPT	172
riverside.com	165
Browser Navigation	144
Claude	66
www.facebook.com	48
Gemini	42
Copilot	31
app-na2.hubspot.com	30

Data Classes Exposed

CLASS	#
Credentials & Secrets	765
Other / Uncategorized	72
Financial (PAN, bank, tax)	61
PII (personal data)	41
healthcare	25
Source Code / IP	14
PII (personal data)	8
jailbreak	6
exfiltration	2
injection	1

Top Triggered Rules

RULE	#
JWT Token	526
Password in Text	196
Social Security Number (no dashes)	67
Social Security Number	41
Credit Card - Visa	30
ICD-10 Diagnosis Code	17
IBAN Number	15
AWS Access Key ID	10

Top Endpoints by Event Volume

HOST	EVENTS	DLP	SHADOW
browser-147-mo1vpq5c	2558	869	1683
Jeromies-MacBook-Air.local	45	23	18
browser-146-mnc280w1	44	25	15
browser-146-mnowd5ut	42	21	19
Mac	37	15	20
UKn0	34	9	21
browser-146-mnc068pc	29	8	18

4 · FLEET HEALTH & POLICY COVERAGE

Endpoint Agent Fleet

14	0	14	13
TOTAL AGENTS	ONLINE (24H)	ACTIVE STATUS	STALE >72H

Monitoring gap: 13 endpoints have not reported in 72h. DLP protection cannot be verified on these devices.

Policy Inventory

28	3	25	0
TOTAL POLICIES	ALLOW	MONITOR	BLOCK

Policy Coverage Gap

31 AI services observed in traffic during the period have **no governing policy**. These represent ungoverned Shadow AI paths and must be reviewed.

Compliance & Regulatory Alignment

Events captured in this period are relevant to the following control frameworks:

FRAMEWORK	CONTROL(S)	RELEVANCE
SOC 2 (TSC)	CC6.1 / CC7.2	Logical access & system monitoring — endpoint DLP + event detection.
NIST 800-53	SI-4, AC-4, AU-12	Information-system monitoring, information-flow control, audit generation.
NIST AI RMF	MEASURE 2.7	AI system security risks identified, tracked, and managed.
ISO 27001:2022	A.8.16, A.8.12	Monitoring activities & data-leakage prevention.
GDPR	Art. 32	Security of processing — technical measures against unauthorized disclosure.
CMMC 2.0	SI.L2-3.14.6 / AU.L2-3.3.1	Monitor system for attacks; create & retain audit records.
EU AI Act	Art. 15	Accuracy, robustness & cybersecurity for AI systems in use.

5 · PRIORITIZED RECOMMENDATIONS

The following actions are prioritized by risk reduction and ease of implementation. Each is mapped to a suggested owner.

PRIORITY	OWNER	ACTION	EXPECTED IMPACT
Critical	SecOps	Triage and close all **343** critical-severity events within 48 hours. Escalate any unresolved to CISO daily standup.	Reduces highest-likelihood breach vectors.
High	IT Operations	Restore connectivity or redeploy SHIELD agent on **13** stale endpoints (>72h since last heartbeat).	Closes monitoring blind spots; restores DLP enforcement at endpoint.
High	AI Governance	Classify and assign a policy (Allow / Monitor / Block) to **31** AI services observed in traffic but lacking governance.	Eliminates unreviewed Shadow AI paths; satisfies approval-workflow control.
High	CISO / HR	Review the **25** user overrides of DLP blocks. Determine whether disciplinary action, retraining, or rule tuning is required.	Reduces insider-risk exposure; demonstrates enforcement follow-through.
Medium	CISO / People Ops	**1794** Shadow AI detections indicate employees are routing around the approved catalog. Publish the approved AI tool list and run a 15-minute awareness refresher.	Drives usage toward sanctioned tools; improves compliance posture.
Medium	SecOps	**2565** events remain in *new* status. Add an SLA (e.g. 5 business days to acknowledge) and report variance weekly.	Reduces alert fatigue and unseen-risk exposure.
Medium	AI Governance	**us-east-1.console.aws.amazon.com** is the top-exposed destination (1300 events). Reassess whether it should be Approved, Sanctioned with DLP, or Blocked.	Targeted control at the single largest data-flow risk.

6 · METHODOLOGY & DEFINITIONS

Reporting Period

All metrics cover a rolling **30-day window** ending May 6, 2026. Trend comparisons use the immediately preceding 30-day window (Mar 7 – Apr 6).

Key Definitions

DLP Violation	An upload or paste blocked or flagged by an endpoint DLP rule (financial, PII, credentials, etc.).
DLP Override	A user explicitly overrode a DLP block. Treated as high-signal insider-risk indicator.
Shadow AI	Use of an AI service not listed in the organization's approved catalog.
Threat Index	Composite 0–100 score: 8×Critical + 4×High + 2×Medium + 1×Low + 5×Overrides + 3×StaleAgents, capped at 100.
Block Rate	% of DLP-flagged events where the upload was prevented at the endpoint.
Fleet Coverage	% of deployed agents reporting a heartbeat within the last 72 hours.
Policy Coverage Gap	Count of distinct AI services observed in traffic that have no matching Allow/Monitor/Block policy.

Data Sources

SHIELD endpoint agents, browser extension telemetry, and the managed AI service catalog. All data is confined to the tenant's own environment; no information is shared across organizations.