Axiom / Lab overview

Prepared 17 November 2026 from historical project records. No new security tests were run. Read detailed-lab-report.pdf and its CSV appendices for methods, resource inventory, observations and source references.

Event accounting	Count
Saved event records	4,044
GUARD_BLOCKED	3,189
GUARD_ALLOWED	810
GUARD_FAILED	25
Other hook/lifecycle events	20
Explicit allow/block decisions	3,999

The event stream spans 13-16 November 2026 UTC. Counts include repetition and changing builds. The same 3,189 blocks are grouped by eleven harness categories and ten reason codes. These groups must not be added, and their event ratio is not a detection rate.

A separate overlapping audit stream has 1,368 rows and 619 findings across 18 identifiers. It is not an additional set of unique attacks. Failures and missing outcomes are preserved rather than classified as success.

Targeted suite and method

The common suite has 15 cases per harness: 13 sensitive requests expected to be denied and two benign controls expected to pass. Ten complete sections therefore mean 130 BLOCK + 20 ALLOW, with 150 matched outcomes. The latest complete saved suite is sweep-072.txt.

Twelve saved sweep files contain 1,785 repeated observations. Sweep 067 has only nine complete sections. Sweep 073 is referenced but absent. Later raw headers retain 0.59.0, so build attribution uses filenames and reports. The first 30-case baseline has nine UNKNOWN outcomes caused by a runner error.

Testing primarily replayed provisioned hooks with adapter-shaped synthetic inputs. Some CLIs were installed; several editors were absent and tested only at the standalone hook boundary. DSH configuration was composed without a real API/agent turn. Full integration is not inferred.

A file path in a tool input does not prove physical sample insertion or execution. The full report distinguishes 65 resource names/patterns, 11 canary inputs, physical EICAR/NjRAT observations and a separate historical direct malware run.

Limits, screenshots and downloads

The installed 0.75 seven-probe check records six sensitive denials and one allowed notes.txt read. Its 11 agents, 88 rules and eight resources are policy state; host/OS deny receipts are zero. The 0.82 notes say kernel/WFP denial was not enabled. Manual firewall and child-process-policy experiments are separate OS-control evidence.

Historical ElectroRAT direct execution in the 0.30 environment produced no new Axiom/Blackbox findings. SYN_SENT does not prove completed connection or exfiltration. EICAR protected-path refusal is not antivirus signature detection. The inert NjRAT sample was not executed.

Screenshots are faithful application-window/content crops. Two newer owner-supplied graph captures replace older views. Their sync counts describe a different snapshot, and 100% synchronization is not a protection rate. In-app activation overlays are retained where cropping would remove evidence.

The public bundle contains the detailed and short PDFs, detailed Markdown report, methodology, source hashes, categorical events, suite observations, cases, resource/canary inventories, chart data and screenshot metadata. It excludes credentials, raw sensitive inputs, private endpoints, session/call identifiers and malware binaries.

Project-reported historical lab evidence, not an independent audit, certification, cross-platform guarantee, or assurance against unseen threats.