Offensive validation

Prove your detections fire against real adversary techniques

Rhombus runs adversary techniques against an isolated range on the kit. Defenders confirm each attack produces the telemetry and detection it should. It turns "we think we would see it" into a repeatable, evidence-backed answer.

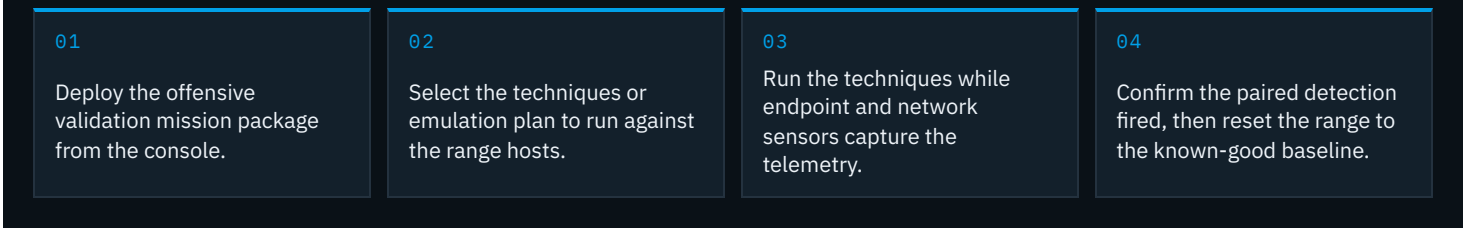
01 · WHY IT MATTERS

- A detection you have never triggered is an assumption, not a control.
- Repeatable attack playbooks let you prove coverage and catch silent regressions.
- Every technique maps to MITRE ATT&CK, so gaps are named in a shared language.

02 · HOW IT WORKS ON RHOMBUS

- Deploy the range from the console as a mission package — an isolated Active Directory forest plus the attack tooling.
- Three attack engines drive the range — single-technique atomic tests, deterministic playbook simulations, and chained agent-based emulation.
- A full-featured penetration-testing workstation provides manual tradecraft through a browser-based desktop.
- The console can reset the range to its delivered baseline between cohorts; Kali content refreshes through signed bundles.

FIG · 01 FROM DEPLOYMENT TO RESULTS



03 · DEFENSE & GOVERNMENT

- Cyber protection teams validate detection coverage before an operation or inspection.
- Defensive cyber operations generate positive controls to prove sensors and rules work.
- Instructors detonate techniques live so trainees watch telemetry appear and hunt it.

04 · COMMERCIAL & ENTERPRISE

- Banks and financial services can validate security operations detections against current attacker tradecraft.
- Enterprise incident response teams can rehearse the run, detect and validate loop in a safe range.
- Managed security service providers can demonstrate detection efficacy to their clients.

05 · OUTCOMES

- Each technique is paired to the telemetry and detection it should produce.
- Runs are repeatable, so detection gaps and regressions surface consistently.
- Results map to MITRE ATT&CK for briefing coverage in a common framework.

06 · CONSIDERATIONS

- The range is deliberately vulnerable and isolated; it is not a hardened production asset.
- Agent-based emulation runs fully offline once staged; some atomic content is staged to the offline mirror as a prerequisite.
- Sizing the range and its sensors is a prerequisite; scope every run to range assets only.

INTEGRATED TOOLS

Kali Linux | Atomic Red Team | PurpleSharp | MITRE Caldera

Commercial products licensed separately · Trademarks of their owners; no endorsement implied