BLACK CELL
Protecting critical infrastructures

WHITEPAPER

RED TEAM SERVICE WHITEPAPER

From Adversary Simulation to Measurable Resilience

Created by

Lajos Illich | Offensive Security

January 2026

TABLE OF CONTENTS

01

ABOUT BLACK CELL

Black Cell is a European cybersecurity company focused on protecting critical infrastructures and the organizations that support them. Our business units cover SOC, Integration, Offensive Security, Cloud Security, Compliance, and ESM (Enterprise Security Monitoring). We take a customer first approach that starts with listening, then shaping solutions to fit the way our clients operate.

Our teams are adaptable and draw on deep knowledge across industries and technologies, from IT and Cloud to ICS/OT. We engage for the long term, providing continual support, service improvement, and measurable outcomes over the lifecycle of the relationship. Clients rely on us to connect regulatory requirements with technology choices and to guide organisational transformation that sticks.

We combine architecture, implementation, and managed operations to close gaps quickly and build sustainable capability. Local presence in Central Europe matters to us, with teams in Budapest and Frankfurt that understand the regional context. This proximity helps us respond faster, coordinate with partners, and keep stakeholders aligned. Above all, we aim to be a trusted partner who strengthens resilience today and prepares you for what comes next.

DISCLAIMER

The information provided in this document is for general guidance only and is used at your own risk. No contractual or advisory relationship is created between Black Cell and any person accessing or using this document or any part of it. Black Cell accepts no liability for any actions, decisions, or consequences arising from the use of this material.

References to third-party sources are included where appropriate. Black Cell is not responsible for the content, accuracy, or availability of external sources, including websites mentioned in this publication.

CONTACT

Béla Droppa
CEO
bela.droppa@blackcell.io

COPYRIGHT NOTICE

02

WHY US?

Our Red Team engagements are designed to be realistic, measurable, and safe to execute. Our red team consists of experts with extensive experience, with knowledge spanning over various fields of both defensive and offensive cyber security.

We deliver red team engagements using a capability-driven model rather than relying on a single individual's skillset. This supports continuity, peer review, and strong operational governance.

We pay outmost care for our client's business continuity, implementing several safeguards which are specified in the engagement's "Rules of Engagement" document.

As the engagement concludes, we produce a comprehensive report of the engagement, which includes tailored sections for the different levels of executives with actionable strategic recommendations.

As part of our service, we provide a closure workshop in which our team presents their findings to the client's stakeholders.

2.1. SENIOR-LED EXECUTION AND QUALITY GATES

Engagements are led by experienced operators and a designated Red Team Lead who is accountable for operational decisions, stakeholder alignment, and risk management throughout the exercise. We apply internal quality gates to ensure that findings are evidence-based, reproducible, and tied to agreed mission objectives and measurable outcomes.

Quality gates typically include:

- Safety and scope compliance checks (rules-of-engagement adherence, stop/pause readiness, and change-control discipline).
- Evidence standards (timestamped logs, screenshots/artifacts, and reproduction notes) so findings are not “one-off observations” but defensible and actionable.
- Peer review of the attacker narrative and the remediation logic to ensure recommendations address systemic patterns, not only isolated weaknesses.

2.2. OPERATIONAL GOVERNANCE AND FAST DECONFLICTION

We establish clear operational governance from day one. This typically includes a dedicated, logged collaboration channel for day-to-day execution and status updates, a defined escalation ladder, and at least two reachable operator contacts for 24/7 deconfliction when defenders observe suspicious activity. If an incident is raised during the exercise, we follow a structured deconfliction process to rapidly confirm whether the activity is part of the engagement or a real-world event, and we act accordingly. This reduces uncertainty for SOC/IT teams and keeps the exercise controlled and aligned with business constraints.

2.3. SAFE-BY-DESIGN EXECUTION

We apply safety controls throughout the engagement to reduce operational risk while preserving realism. These controls are explicitly documented in the “Rules of Engagement” document and agreed upon with stakeholders before execution.

- **Time and change control:** we define permitted testing windows, support “freeze windows” where operations pause on request, and manage scope/time changes through documented change requests (CR).
- **Stop/Pause (“kill switch”):** we immediately suspend activity upon business-impact signals, unexpected data exposure, physical safety risks, or an explicit client stop request. We then place offensive infrastructure into a safe state and resume only on written approval.
- **Data handling safeguards:** evidence extraction is minimized and controlled. If sensitive or personal data is encountered, it is handled with strict access control, protected in transit and at rest, and retained only for a limited period for reporting purposes.
- **Controlled C2 operations:** communications are encrypted and mutually authenticated. Outbound destinations are restricted to what is required for mission objectives, and persistence mechanisms are removed at the end of the engagement. During the engagement, all operator actions are logged.

2.4. EVIDENCE-DRIVEN OUTCOMES YOU CAN ACT ON

We document actions with timestamps and supporting artifacts so that defenders can correlate what happened with telemetry, measure time-to-detect/time-to-respond where feasible, and reproduce key steps in a controlled manner. Reporting is structured for multiple audiences (board/executive/technical) and includes an attacker narrative that connects tactics and techniques to concrete impact and prioritized remediation.

2.4.1. EXAMPLE: “CONTROLLED REALISM AT SCALE”

In a recent multi-location exercise pattern, one site prevented a physical access attempt while another allowed a controlled internal foothold. That contrast helped leadership validate that strong controls do work, but also revealed that where “internal foothold” is achieved, the decisive risk concentrates in identity and privilege layers. This is the kind of real-world, actionable learning outcome our approach is built to deliver.

03

RED TEAMING SERVICE

The Red Teaming service provides a structured, intelligence-led simulation of real-world cyber adversaries targeting an organization's people, processes, and technology. The primary goal is to evaluate how effectively an organization can resist, detect, and respond to a determined attacker operating under realistic conditions. Rather than focusing on isolated weaknesses, the service examines how individual gaps can be combined into meaningful attack paths that result in tangible impact.

Red team engagements are designed to reflect the behavior of modern threat actors, including their use of stealth, persistence, and adaptive decision-making. Techniques and tooling are selected to emulate credible adversaries relevant to the organization's industry and risk profile. This approach enables a more accurate assessment of defensive capability than traditional testing methods that rely on narrowly scoped techniques.

All activities are conducted within a clearly defined and formally approved rules-of-engagement framework. This framework establishes legal authorization, outlines acceptable techniques, and defines operational boundaries. Depending on the engagement model, testing may be conducted covertly to preserve realism, or with limited awareness among defensive stakeholders to support specific testing objectives.

3.1. REALISTIC ATTACK SCENARIOS

We model credible adversary behavior across people, process, and technology. Scenarios are selected based on the threat model and tailored objectives, and may combine multiple pathways into a single, end-to-end attacker narrative.

Scenario patterns we commonly emulate include:

- **Multi-site physical entry simulation (social engineering)**
  - Role-based entry attempts (e.g., maintenance/audit/IT-service pretexts) to test reception and staff verification practices.
  - Outcome comparison across sites to identify inconsistency in control enforcement.
- **Controlled internal foothold via approved implant device**
  - A managed device placed on-site (approved beforehand) provides a realistic “inside-the-perimeter” starting point.
  - Post-placement activities focus on low-noise discovery and safe internal operations under “Rules of Engagement” constraints.
- **Identity-led compromise chains**
  - We test how quickly an attacker can progress once any authenticated identity is obtained (user or machine), emphasizing directory services, certificate-based authentication, and privilege escalation paths.
  - The aim is to evaluate whether the organization can prevent escalation, detect it early, and contain it effectively.
- **Enumeration of client-specific services**
  - We validate how services employed by our client (e.g., service desk portals) or legacy components can expose pivot opportunities (e.g., injection flaws or outdated components), and we observe whether exploitation phases are detected and triaged.

3.1.1. EXAMPLE “WHY IDENTITY BECOMES THE DECISIVE LAYER”

A repeated observation in real-world-style engagements is that controls can perform well against unauthenticated probing, yet the attacker’s success probability increases significantly after internal presence is achieved and at least one authenticated identity is obtained. In such cases, the organization’s resilience depends on identity hardening, certificate service configuration, and the speed and consistency of response to high-signal identity events.

3.2. TAILORED OBJECTIVES AND “CROWN JEWEL” FOCUS

Each engagement is tailored to the client’s environment and business priorities. Working together with our client’s approved stakeholders, we define the primary objectives of the engagement. These objectives can be specific privileges to achieve or systems/infrastructure objects to compromise.

- **Crown jewels and mission outcomes:** objectives aligned to business-critical systems (e.g., sensitive file repositories, privileged administrative tiers, identity bridge components, or ERP/production-supporting systems).
- **Realistic constraints and dependencies:** permitted test windows, change-control for scope/time modifications, and explicit stop/pause conditions to protect business continuity.
- **Meaningful success criteria:** where engagement conditions allow, we include measurable indicators such as time-to-detect and time-to-respond expectations for key attacker actions.

3.3. HOW YOUR ORGANIZATION BENEFITS

Red teaming provides decision-grade evidence of resilience under realistic attacker behavior. Instead of producing only a list of vulnerabilities, the engagement demonstrates how weaknesses can be chained into meaningful attack paths that create business impact and which controls actually stop or slow the attacker.

3.3.1. WHAT EXECUTIVES GAIN

- **Risk clarity:** a concrete "attacker story" that shows what could happen, how quickly, and under what conditions.
- **Prioritization:** a small set of high-impact improvements, typically centered on identity and privilege layers, that reduce the likelihood of worst-case outcomes.
- **Evidence for investment:** objective validation of which security controls and processes work well, and where gaps remain, supporting smart spending and targeted maturity improvements.

3.3.2. WHAT TECHNICAL TEAMS GAIN

- **Actionable remediation:** findings with evidence and ATT&CK mapping to support detection engineering and control tuning.
- **Operational readiness insights:** measurement-oriented feedback (e.g., time-to-detect/time-to-respond for key steps), enabling teams to track improvement between exercises.
- **Differentiation between "noisy vs. quiet" phases:** clear identification of where telemetry is strong (alerts triggered) and where stealthier enumeration and identity-layer activity may require additional monitoring.

3.3.3. EXAMPLE: “HIGH ROI WITHOUT BUYING NEW TOOLS”

In one engagement pattern, many improvements were identified as targeted configuration and operational changes — especially around identity and authentication hardening — rather than requiring additional tooling purchases. This kind of outcome typically yields fast ROI because it reduces systemic risk at the layer where high-impact compromise paths converge.

4

EXECUTION AND DELIVERABLES

4.1.EXECUTION

Execution follows an adaptive, hypothesis-driven approach rather than a fixed sequence of actions. Initial assumptions about access paths, defensive coverage, and detection capability are continuously tested and refined as the engagement progresses. This allows the red team to pivot based on observed conditions, closely mirroring the operational behavior of real attackers.

Typical phases of execution may include:

- Targeted reconnaissance to identify exposed services, trust relationships, and potential entry points
- Gaining initial access using techniques consistent with the defined threat model
- Establishing persistence and escalating privileges while minimizing detection
- Conducting internal reconnaissance and lateral movement to reach high-value assets
- Simulating command-and-control activity and executing defined objectives
- Observing and documenting defensive detection, alerting, and response actions

Throughout the engagement, care is taken to limit operational risk, avoid unnecessary disruption, and maintain the integrity of client systems and data.

4.2. DETAILED REPORTING, TIMELINE, AND VISUALIZATIONS

Our reporting is built around a complete attacker narrative that reconstructs the engagement end-to-end. The objective is to make the results consumable by leadership and actionable for defenders.

The report typically includes:

- **Board Summary:** business impact, systemic risk exposure, and strategic meaning.
- **Executive Summary:** scope, objectives, key themes, and what worked vs. what failed.
- **Technical Summary and Findings:** evidence-backed findings with MITRE ATT&CK mapping and clear remediation guidance.
- **Engagement Timeline:** a time-sequenced reconstruction of key attacker actions, decision points, and observed detection/response events.
- **Visualizations:** attack-path illustrations and diagrams that show how the attacker moved from foothold to objective, including pivot points and control gaps.

We include timing and readiness indicators for key events (e.g., time-to-detect and time-to-respond), enabling teams to benchmark improvement between engagements and align remediation with operational readiness goals.

4.2.1. EXAMPLE: “TIMELINE DRIVES LEARNING”

A detailed timeline can reveal that some high-signal actions were detected quickly, while other phases (such as stealthy identity-layer enumeration) produced fewer signals. This directly informs what telemetry to improve and which playbooks to tune.

4.3. ADVERSARY & DEFENDER PERSPECTIVE

We explicitly document not only what the adversary did, but what the defenders could see, what was detected, and where visibility gaps exist. This includes distinguishing “noisy” phases that trigger alerts from “quiet” phases that may require improved telemetry, analytics, and playbooks.

4.3.1. EXAMPLE: “KNOWN SIGNATURES VS. BLIND SPOTS”

In some environments, alerts may trigger primarily when known attack signatures appear, while other reconnaissance or enumeration behavior can be less visible. We document these differences so teams can extend detection logic from signature-based triggers to behavior and identity-event monitoring. This approach better aligns with modern attacker tradecraft and provides a more robust defense against adversaries.

4.4. Measured outcomes (resilience perspective)

Where engagement conditions allow, we evaluate outcomes using measurable indicators across prevention, detection, and response. We use event-level timing to support operational learning (e.g., time-to-detect and time-to-respond for critical attacker actions) and to help teams measure improvement between engagements.

4.4.1. EXAMPLE: “EVENT-BASED METRICS BEAT AVERAGES”

Rather than relying only on overall averages, organizations often gain more value by measuring detection and containment time for specific high-impact events (e.g., privileged group membership changes or new principal creation). This makes improvement programs concrete, testable, and repeatable.