

HARD ROCK SECURITY • LIVE AT HARD ROCK CAFE BELGRADE • 08.10.2026



UNDER PRESSURE

Nemesis BAS: what happens when we put the defense to the test?

Nikos Laleas

PERSISTENT SECURITY INDUSTRIES

HOSTED BY NET++

WHO'S ON STAGE

NIKOS LALEAS

Co-founder, Persistent Security Industries

14 YEARS

in offensive security and
red teaming

RED TEAMER

I break into companies for
a living. With permission.

NEMESIS

Co-creator of the platform
you'll see today

TODAY, 12:30-13:00

THE SETLIST

#	Track	What you get
01	Under Pressure	Why owning tools is not proof
02	(It's Not) A Kind of Magic	Meet Nemesis, and how it works
03	The Show Must Go On	Testing every week, with the AI Operator
04	Don't Stop Me Now	Who gets what, and how to start
+	Encore	A quick look at the dashboard
?	After the show	Your questions

TRACK 01 • UNDER PRESSURE • AUDIENCE CHECK

HANDS UP

Who here has an EDR?

Keep it up if you've tested it against a real attack in the last 30 days.

TRACK 01 • UNDER PRESSURE

EVERY SECURITY CONTROL WORKS PERFECTLY IN THE VENDOR'S DEMO.

The real question is how it holds up in your network, with your configuration, today.

CAN YOU DETECT THIS?

01

A workstation pushing 1 GB of data over DNS to a domain nobody has seen before

02

Customer data leaving in cleartext over HTTP, to a non-standard high port

03

A workstation contacting a specific domain every 65 seconds 24/7

If the honest answer is “I think so”, that's exactly the gap.

A PENTEST IS A PHOTO. YOUR NETWORK IS A LIVE SHOW.

POINT-IN-TIME

ONCE A YEAR

- A snapshot of a single day
- Changes between tests go unseen
- Shows you were covered, once

CONTINUOUS VALIDATION

EVERY WEEK, AUTOMATICALLY

- Every change gets re-tested
- A trend over time, not one score
- Shows whether you're covered now

WHY NOW: THE PRESSURE IS RISING

27

SECONDS

Fastest breakout to a second system. Average: 29 min.

-7

DAYS TO EXPLOIT

Attacks start before the patch exists.

48%

OF BREACHES

were first spotted by an outsider, not the victim.

95%

OF SECURITY TEAMS

report at least one skills gap.

Serbia: the new NIS2-aligned Law on Information Security applies since 31 October 2025. Risk assessments are due by 2027.

Sources: CrowdStrike Global Threat Report 2026 • Mandiant M-Trends 2026 • ISC2 Cybersecurity Workforce Study 2025 • Serbian Law on Information Security

TRACK 02 OF 04

(IT'S NOT) A KIND OF MAGIC

Meet Nemesis, and how it works

TRACK 02 • MEET NEMESIS



NEMESIS

YOUR SECURITY ONLY WORKS UNTIL IT DOESN'T. TEST IT BEFORE ATTACKERS DO.

The breach and attack simulation platform by Persistent Security Industries, built in Belgium.

SIMULATE

Real attacker techniques, run safely in your network

AUTOMATE


Scheduled and repeatable, driven by AI and threat intel

VALIDATE

Prevented? Detected? Proof instead of assumptions

TRACK 02 • (IT'S NOT) A KIND OF MAGIC

HOW IT ALL CONNECTS



RIFFS, SONGS AND GIGS

THE RIFF ATOMIC

One real attacker technique, mapped to MITRE ATT&CK

THE SONG SCENARIO

Atoms chained into a full attack, the way real threat groups work

THE GIG ASSESSMENT


Scenarios played on your agents, on a schedule you choose

Assessment	
Scenario	
Atomic	T1003.002
Atomic	T1003.001
Atomic	T1055.002

Scenario	
Atomic	T1055
Atomic	T1055.002
Atomic	T1055.001


TRACK 02 • IN THE PRODUCT

READY-MADE LIBRARIES



ATOMICS

Single techniques. Write your own in TOML and upload them.



SCENARIOS

Ready-made attack chains, or generate one with AI.

EVERY ATOMIC SITS ON THE MITRE ATT&CK MAP

Initial Access	Execution	Persistence	Privilege Escalation	Stealth	Defense Impairment	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration
11 techniques	20 techniques	22 techniques	13 techniques	30 techniques	18 techniques	17 techniques	34 techniques	9 techniques	17 techniques	18 techniques	9 techniques
Content Injection	BITS Jobs	Account Manipulation (7)	Abuse Elevation Control Mechanism (6)	Access Token Manipulation (5)	Disable or Modify System Firewall (3)	Adversary-in-the-Middle (4)	Account Discovery (4)	Exploitation of Remote Services	Adversary-in-the-Middle (4)	Application Layer Protocol (5)	Automated Exfiltration (1)
Drive-by Compromise	Cloud Administration Command	BITS Jobs		BITS Jobs	Brute Force (6)	Brute Force (4)	Application Window Discovery	Internal Spearphishing	Archive Collected Data (3)	Communication Through Removable Media	Data Transfer Size Limits
Exploit Public-Facing Application	Command and Scripting Interpreter (13)	Boot or Logon Autostart Execution (14)	Access Token Manipulation (5)	Build Image on Host	Domain or Tenant Policy Modification (2)	Credentials from Password Stores (8)	Browser Information Discovery	Lateral Tool Transfer	Audio Capture	Content Injection	Exfiltration Over Alternative Protocol (3)
External Remote Services	Container Administration Command	Boot or Logon Initialization Scripts (5)	Boot or Logon Autostart Execution (14)	Deobfuscate/Decode Files or Information	Downgrade Attack	Exploitation for Credential Access	Cloud Infrastructure Discovery	Remote Service Session Hijacking (2)	Automated Collection		
Hardware Additions	Deploy Container	Cloud Application Integration	Boot or Logon Initialization Scripts (5)	Direct Volume Access	Exploitation for Defense Impairment	Forced Authentication	Cloud Service Dashboard	Remote Services (8)	Browser Session Hijacking	Data Encoding (2)	Exfiltration Over C2 Channel
Phishing (4)	ESXi Administration Command	Compromise Host Software Binary	Create or Modify System Process (5)	Execution Guardrails (2)	File and Directory Permissions Modification (2)	Forge Web Credentials (2)	Cloud Service Discovery	Replication Through Removable Media	Clipboard Data	Data Obfuscation (3)	Exfiltration Over Other Network Medium (1)
Replication Through Removable Media	Exploitation for Client Execution	Create Account (3)	Domain or Tenant Policy Modification (2)	Exploitation for Stealth	Hide Artifacts (14)	Input Capture (4)	Cloud Storage Object Discovery	Software Deployment Tools	Data from Cloud Storage	Dynamic Resolution (3)	Exfiltration Over Physical Medium (1)
Supply Chain Compromise (3)	Hijack Execution Flow (12)	Create or Modify System Process (5)	Escape to Host	Hide Artifacts (14)	Hijack Execution Flow (12)	Modify Authentication Process (9)	Container and Resource Discovery	Taint Shared Content	Data from Configuration Repository (2)	Encrypted Channel (2)	Exfiltration Over Web Service (4)
Trusted Relationship	Input Injection	Event Triggered Execution (18)	Event Triggered Execution (18)	Indicator Removal (8)	Indirect Command Execution	Multi-Factor Authentication Interception	Debugger Evasion	Use Alternate Authentication Material (4)	Data from Information Repositories (6)	Fallback Channels	Scheduled Transfer
Valid Accounts (4)	Inter-Process Communication (3)	Exclusive Control	Exploitation for Privilege Escalation	Masquerading (12)	Modify Cloud Resource Hierarchy	Multi-Factor Authentication Request Generation	Device Driver Discovery		Data from Local System	Ingress Tool Transfer	Transfer Data to Cloud Account
Wi-Fi Networks	Native API	External Remote Services	Process Injection (12)	Obfuscated Files or Information (18)	Modify Registry	Network Sniffing	Domain Trust Discovery		Data from Network Shared Drive	Multi-Stage Channels	
	Poisoned Pipeline Execution	Implant Internal Image	Scheduled Task/Job (5)	Pre-OS Boot (5)	Modify System Image (2)	OS Credential Dumping (8)	File and Directory Discovery		Data from Removable Media	Non-Application Layer Protocol	
	Scheduled Task/Job (5)	Modify Authentication Process (9)	Valid Accounts (4)	Process Injection (12)	Reflective Code Loading	Steal Application Access Token	Group Policy Discovery		Data Staged (2)	Non-Standard Port	
	Serverless Execution			Rootkit	Plist File Modification		Local Storage Discovery		Email Collection (3)	Protocol Tunneling	
	Shared Modules						Log Enumeration				
	Software Development Tools						Network Service Discovery				
							Network Share				

ONE LANGUAGE

Red and blue teams talk about the same techniques

NAMED GAPS

Every miss points to the exact technique to tune

TRACK 02 • (IT'S NOT) A KIND OF MAGIC

ONE AGENT PER GOLDEN IMAGE

A small agent on one machine per distinct configuration. It talks to the platform and runs the techniques.

CONFIGURATION 1

Workstation • Windows 11

Agent 1

CONFIGURATION 2


Kali VM

Agent 2

CONFIGURATION 3

Member server • Server 2022

Agent 3



TRACK 02 • (IT'S NOT) A KIND OF MAGIC

SAFE TO PLAY IN PRODUCTION

AGENTS CALL OUT

Nothing calls in. No inbound ports.

ONE AGENT PER GOLDEN IMAGE


Not on every endpoint: a light rollout.

CLEAN UP, EVERY TIME

Reverted after each technique.

YOUR DATA, YOUR CHOICE

EU SaaS, on-prem or air-gapped.



EVERY TECHNIQUE PLAYS IN THREE BEATS

PLAY IT 1 • RUN

The agent executes the technique, like a real attacker

CHECK IT 2 • VERIFY

Did your controls block it or raise an alert?

PACK IT UP 3 • CLEAN UP

Changes reverted, files removed. Nothing stays on stage.

```
description = "Executes a basic Nmap scan using the target IP and optional ports with a specified timing template."
name = "Simple Nmap Scan (Linux)"
recommendation = "Ensure proper network permissions before scanning and be aware of scan impact on target systems."
supported_platforms = [ "linux" ]
technique_id = "T1046"
technique_name = "Network Service Scanning"
type = "actionrunner"

[[input_arguments]]
default = "127.0.0.1"
description = "Target IP address or hostname."
name = "target"
placeholder = "target"
type = "string"

[executor]
command = 'Cmd="nmap -T#{timing}"; if [ -n "#{ports}" ]; then Cmd+=" -p #{ports}"; fi; Cmd+=" #{target}"; echo "Executing: $Cmd"; eval "$Cmd"'
elevation_required = false
name = "bash"
success_command = "echo 'Scan completed'"
```


PREVENTED

DETECTED

MISSED

RUNNING AN ASSESSMENT

- 01 Pick scenarios from the library, or build your own
- 02 Link the agents you deployed
- 03 Run it on demand, or put it on a schedule
- 04 Read the verdict per atomic, with outputs and IOCs
- 05 Export the report, or ask the AI to explain the run



TRACK 02 • IN THE PRODUCT

EVERY TECHNIQUE GETS A VERDICT

Scenarios

Q

Get sigma rules

Enable edit

SCENARIO

#ATOMICS

GTI Gap - Credential Access & Evasion - 2026-07-13

5

Q

ATOMIC	TECHNIQUE	STATUS	AGENT	PREVENTED	DETECTED	
Credential Dumping: SAM, SYSTEM, and SE...	T1003.002	Completed	Windows11			
Dump LSASS.exe Memory using comsvcs.dll	T1003.001	Completed	Windows11			
Portable Executable Injection	T1055.002	Completed	Windows11			
Powershell Mimikatz (Local)	T1003.001	Completed	Windows11			
Process Injection via mavinject.exe	T1055.001	Completed	Windows11			

THE SCENARIO

Five atomics, each mapped to a MITRE ATT&CK technique

THE VERDICT

Green shield: blocked or seen. Red shield: it got through.

THE NEXT FIX

Last row: not prevented, not detected. Start there.

TRACK 03 OF 04

THE SHOW MUST GO ON

Testing every week, with the AI Operator

03

PROVE WHAT HOLDS. FIX WHAT DOESN'T. REPEAT.



Every new EDR policy, GPO, firewall rule or admin account is a gig your controls haven't rehearsed.

A NEW THREAT HITS THE NEWS. NEMESIS BUILDS THE TEST.


01 Threat intel comes in

02 The AI Operator proposes a scenario

03 **Your team approves it**

04 Agents run it on your hosts

05 Results and the next step come back



TRACK 03 • THE AI OPERATOR, IN THE PRODUCT

FROM A QUESTION TO A FIX

1 • ASK

ASK OR EXPLORE

Daily security briefing

Attack surface discovery

DORA compliance analysis

Create custom scenario

Start from a prompt or today's threat intel

2 • CHOOSE

PICK YOUR BATTLE

CHOOSE ONE

1 • Readiness assessment

2 • AI-attack emulation

3 • MITRE threat emulation

The operator proposes tests. You choose.

3 • RUN

ACT ON RESULTS

✓ Completed

AI-attack emulation assessment

Agents run it. Results come back scored.

4 • ITERATE

FIX AND REPEAT

CHOOSE ONE

1 Generate missing atomics

2 Analyze gaps

3 Generate interactive visualization

4 Re-run detection

Detection rules, missing atomics, re-run

TRACK 03 • THE AI OPERATOR, IN THE PRODUCT

ONE-CLICK INTELLIGENCE, EVERYWHERE



Explain this finding

Any result, explained in plain language



Suggest remediation

Concrete steps to close the gap



Generate Detection Rules

Detection rules for whatever got through



Executive summary

A summary your board can read




AI suggestions ▼

What to test next, based on your results

Anything you can click in the platform, you can also ask the AI Operator to do in a chat.

GIVE IT AN OBJECTIVE. IT PLANS THE ATTACK.



AN OBJECTIVE

Not a fixed script: a goal, like reaching domain admin

PHASES ON THE FLY

The AI builds each phase like a scenario, as it goes

ADAPTS TO RESULTS

Picks the next atomics based on what was blocked or seen

TRACK 04 OF 04

DON'T STOP ME NOW

Who gets what, and how to start

ONE RUN, THREE ANSWERS

CISO

EVIDENCE

One page for the board and auditors: covered, at risk, trending

IT MANAGER

PRIORITIES

A short, ranked list of what to fix first, and how

ADMINS & SOC

DETECTIONS

Exactly which techniques your EDR and SIEM missed, so you can tune them

Mapped to MITRE ATT&CK, with NIS2 and DORA alignment built in.

WHAT TEAMS USE NEMESIS FOR

BREACH AND ATTACK SIMULATION

SOC validation	Detection rule testing	Ransomware readiness
Baseline hardening	Security tool comparison	Compliance and cyber drills
Data loss prevention	Threat group simulation	Purple teaming

PLUS: ATTACK SURFACE

Domain monitoring

Look-alike domains and cloned websites

Asset monitoring

Your public IPs, domains and certificates

Port scanning

Newly opened or closed ports, as they happen

DETECTION RULES FROM REAL RUNS

```
title: BCDEdit Boot Recovery Tampering
id: a]ransomware-t1490-bcdedit-recovery
status: experimental
description: >
  Detects bcdedit.exe being used to disable Windows recovery features,
  a hallmark of ransomware pre-encryption preparation.
  Observed in Nemesis BAS ransomware readiness assessment.
author: Nemesis BAS - Persistent Security Industries
date: 2026/09/03
references:
  - https://attack.mitre.org/techniques/T1490/
tags:
  - attack.impact
  - attack.t1490
logsource:
  category: process_creation
  product: windows
detection:
  selection_binary:
    ImageEndswith: '\bcdedit.exe'
```

FOUND A GAP

A technique ran and nobody saw it.

GET THE RULE

Nemesis gives you a Sigma rule for it, mapped to MITRE ATT&CK.

PROVE THE FIX

Re-run the scenario and watch the shield turn green.

THREE THINGS TO TAKE HOME

RELIEVING THE PRESSURE

01 **TOOLS ARE NOT PROOF**
Test them the way an attacker would.

02 **YOUR DEFENSES DRIFT**
Validate continuously, not once a year.

03 **START SMALL**
One agent, one scenario, one report.

TRACK 04 • DON'T STOP ME NOW

START WITH ONE GIG

1

AGENT

on one
golden image

1


SCENARIO

a real-world
attack scenario

1

REPORT

prevention and
detection scores,
with fixes



Ask the Net++ team about a proof of concept.

WHAT THE CISO TAKES TO THE BOARD

SCORES


Prevented, detected or missed, per technique

COVERAGE

Mapped to MITRE ATT&CK, so every gap has a name

TREND

Prevention and detection over time, run after run



AFTER THE SHOW



QUESTIONS?

Next on stage: Smells Like a Breach, from attack to response, live.

Nikos Laleas

nikos.laleas@persistent-security.net

NEMESIS BY PERSISTENT SECURITY INDUSTRIES • IN SERBIA VIA NET++