PRAETORIAN

JULY 2026 – OCTOBER 2026

The AI Attacker's Almanac 2026 Q3 Edition

What autonomous attack agents found
across 50 enterprise organizations



FOREWORD

Black October: A New Operational Tempo

In the second half of October 2026, the industry saw a preview of what we will likely live with for the next several years: the opening of a new operational tempo. In the short to medium term, AI gives attackers an operational advantage. In the long term, the same capability that is causing the problem may ultimately be the solution. Hopefully.

As the founder of an offensive security company, I have never been more excited about our market segment. Prior to AI, penetration testing had largely become a check-the-box commodity as unqualified players saturated the space. The TAM was small and for CISOs, the service was, at best, tactical.

AI has changed all of that. Attackers' adoption of AI has caused customers to rapidly shift to a continuous offensive security model to mirror the relentlessness and speed of what is happening in the wild. Proactive offensive security is now just as important as reactive post boom detection. Under the new agentic subscription model, our business is growing at a rate that we have never experienced. Agents allow us to scale in ways we never could. Unfortunately, what is good for me is also good for the attacker, and neither is a good sign for you.

This almanac series will highlight the state of play between offense and defense in the age of AI. Each edition asks one question. **Is defense keeping up?**

Nathan Sportsman

Founder & CEO, Praetorian · October 31, 2026

OCTOBER 13–29

17 DAYS

- **OCT 13 · PATCH TUESDAY**
  Microsoft fixed 1,413 vulnerabilities in one release, the largest Patch Tuesday on record.
- **OCT 15 · CONNECTWISE**
  One stolen app secret opened more than 160 customer Microsoft 365 tenants.
- **OCT 17 + 29 · AI DEVELOPER TOOLS**
  OpenCode leaked users' AI provider keys; a trojanized Claude Code extension reached up to 5,000 Cursor users in 11 hours.
- **OCT 19 · GLOBALPROTECT**
  Unit 42 tied Qilin and INC Ransom to a Palo Alto VPN flaw patched five months earlier.
- **OCT 21 · PHP2SHELL**
  Critical PHP remote code execution. An exploit attempt was logged four days before disclosure; CISA listed it as exploited within a week.
- **OCT 21 · NODE2SHELL**
  A Brotli flaw in Node.js that spread to .NET, Python and Bun. CISA listed it as exploited within a week.
- **OCT 29 · BINANCE**
  About $6.3 billion drained in roughly 20 minutes, the largest crypto theft on record.

CONTENTS

Table of contents

01	**By The Numbers** Campaign totals, exposed services, CVEs exploited, and the most common risks	PAGE 01
02	**Attack Paths** How deep paths run, where they begin, and six cases from the internet to Critical	PAGE 09
03	**Case Studies** Three customer compromises step by step and the Hyper Challenge CTF	PAGE 14
04	**Threat Actor Comparison** Eight active criminal groups and the ways in they share with our agents	PAGE 24
05	**Key Takeaways** Predictions and what to do about it	PAGE 34
06	**Praetorian Guard** Our continuous offensive security platform, a free two-week pilot and a recorded demo	PAGE 37
A	**Appendix** Definitions and MITRE ATT&CK techniques	PAGE 44



SECURITY STARTS WHERE THE ATTACK BEGINS

Q3 2026 EDITION

Two of eight attack surfaces

WHAT THIS EDITION MEASURES

What our autonomous agents found on the internet perimeter: the services, applications and attack paths an outside attacker reaches first.

WHAT IT LEAVES OUT

The crews in Section 4 also get in through phishing and the software supply chain. Praetorian Guard tests both routes, but this edition does not measure them.

WHAT COMES NEXT

A new almanac every quarter, each covering two of the eight surfaces Guard tests over a 3-month time-period. This edition pairs the external perimeter with web applications and services, which make up much of what sits on it.

01
External Perimeter
Internet-facing hosts and services
THIS EDITION

02
Web & Services
Applications and APIs on the perimeter
THIS EDITION

03
Internal Network
Hosts and services on the intranet
UPCOMING

04
Cloud Infrastructure
Cloud accounts and workloads
UPCOMING

05
LLMs & MCP Servers
Models, agents and their tools
UPCOMING

06
Source Code
CI/CD, secrets and static analysis
UPCOMING

07
Mobile
iOS and Android applications
UPCOMING

08
Users
Social engineering and phishing
UPCOMING

SECTION 1

By the numbers

01

WHAT GUARD'S AGENTS PROVED

373

CRITICAL RISK COMPROMISES

1,621

HIGH RISK COMPROMISES



50
Organizations

796
Campaigns

92
Days in the window

647	Validated attack paths
1.48	Avg. Significant Vulnerabilities Per Path
5.3	Average campaigns per organization per month

SECTION 1

Services left open to the internet

Unsurprisingly, SSH leads the way

SERVICE	ORGANIZATIONS
SSH	42
SMB / NetBIOS	36
FTP	35
SNMP	33
SIP	33
RDP	32
NFS	31
MSSQL	31



SECTION 1

CVEs used to compromise customers



7 CVEs across 6 flaws, all 7 now on CISA's KEV list

CVSS 3.1	VULNERABILITY	SCORE SOURCE AND VECTOR	CISA KEV ADDED
10.0 CRITICAL	**React / Next.js** · RCE CVE-2025-55182	CVSS 3.1 · Meta AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H	2025-12-05
9.8 CRITICAL	**Microsoft Exchange** · ProxyLogon CVE-2021-26855	NVD 3.1 (Microsoft 9.1) AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H	2021-11-03
10.0 CRITICAL	**Metabase** · SQL injection CVE-2026-72898	CVSS 3.1 AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H	2026-08-11
9.8 CRITICAL	**WordPress** · REST route CVE-2026-63030	CVSS 3.1 · WPScan AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H	2026-07-21
10.0 CRITICAL	**GitLab** · ExifTool RCE CVE-2021-22205	NVD + GitLab 3.1 AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H	2021-11-03
9.8 CRITICAL	**PHP2Shell** · out-of-memory and stale global state RCE CVE-2026-109618 · CVE-2026-109619	NVD 3.1 (9.8 / 8.1) AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H	2026-10-26

SECTION 1 AS OF OCTOBER 31, 2026

PHP2Shell drew millions of tries

Few worked, but confirmed exploited

WHAT IT IS

9.8 / 8.1 CVSS 3.1 (NVD)

Two memory-safety flaws in the core PHP runtime

CVE-2026-109618 · CVE-2026-109619

The flaw sits in PHP request processing, so fully updated WordPress and Drupal sites are exposed before app code runs.

PHP2Shell Research

https://www.praetorian.com/blog/php2shell/

CURRENT EXPLOITATION

01 **Attacked before disclosure**

AMAZON THREAT INTELLIGENCE

Honeypots logged exploitation attempts before the vulnerabilities were made public.

02 **Automated attack campaigns already underway**

SOCRADAR

A working exploit chain and target-screening scripts sat on an exposed attacker server, aimed at Europe, North America and Asia.

03 **Race to hit servers before patches land**

RELIAQUEST · RECORDED FUTURE · CLOUDFLARE

Botnet operators and financially motivated crews are racing to hit unpatched servers; few attempts have worked.

04 **Confirmed break-ins**

UNIT 42

Attackers exploited a small share of internet-facing firewall management interfaces and chained the flaw to root on servers. No victims named.

SECTION 1

The exploitable risks seen most widely

Ranked by how many organizations had the risk on a validated attack path, not by how often it shows up on the 647 paths. A risk can show up many times inside one or two organizations and still be rare. **Exposed secrets show up 191 times on attack paths, in 16 organizations.**

#	RISK CLASS	ORGANIZATIONS	FINDINGS ON PATHS
1	**Broken access control**	25	125
2	**Authentication and default credentials**	23	330
3	**Injection: command, SQL, template, code**	19	219
4	**Exposed secrets and keys**	16	191
5	**Information disclosure**	13	21
6	**Misconfiguration**	8	57

SECTION 1

Exposed secrets that presented the highest risk

Critical and High secrets on the external surface, ranked by how many organizations had each type.

#	Secret type	ORGANIZATIONS
1	**Amazon Cognito**	4
2	**AWS AppSync**	2
3	**AWS IAM Access Keys**	2
4	**Azure Storage**	2
5	**Auth0**	1
6	**Azure Service Principal**	1
7	**Azure Service Bus**	1
8	**Google Maps** (No Restriction)	1
9	**Okta**	1



SECTION 1

Common web flaws that presented the most danger

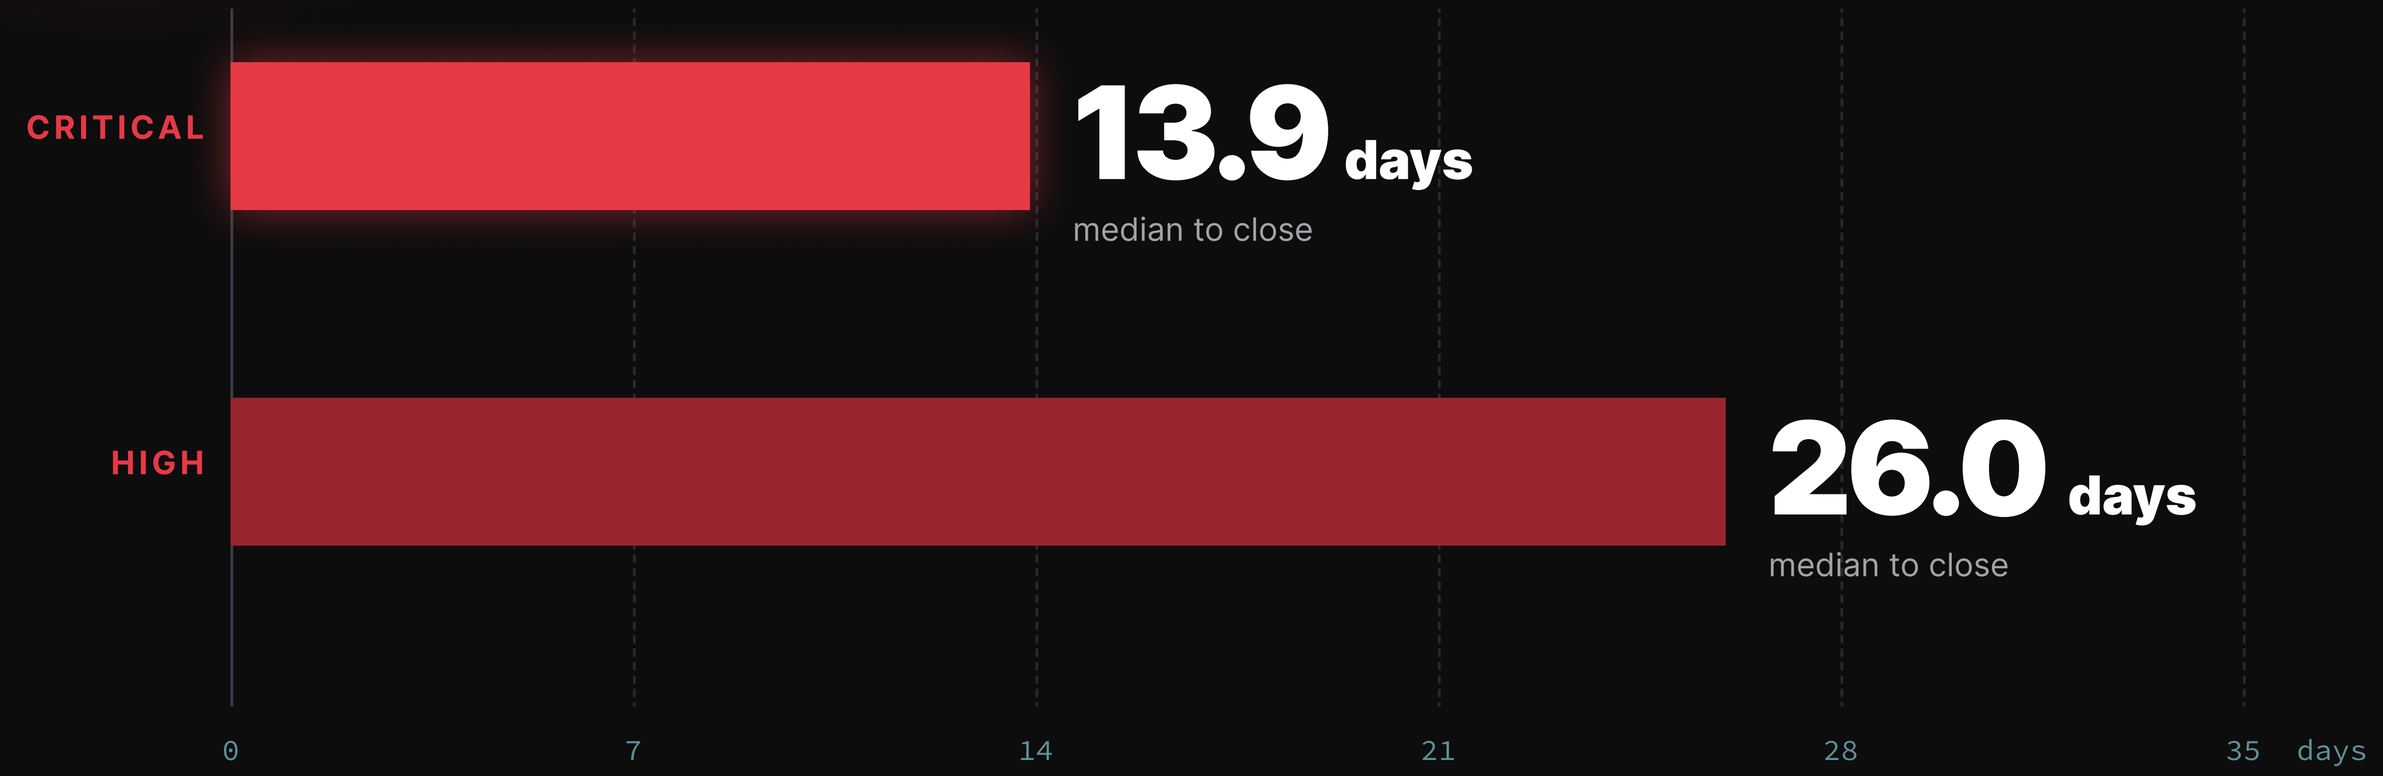
Critical and High web application flaws on the external surface, grouped by flaw class and ranked by how many organizations had each. Secrets are counted separately. **Authentication bypass or account takeover turned up at 40 organizations.**

#	Flaw class	ORGANIZATIONS
1	Authentication bypass or account takeover	40
2	Missing authentication on an API or function	38
3	Broken access control or privilege escalation	29
4	SQL or NoSQL injection	28
5	Known CVE in web software	25
6	Remote code execution or code injection	23
7	Sensitive data exposure	21
8	Object-level authorization (IDOR)	20
9	Server-side request forgery (SSRF)	20
10	Default or weak credentials	18

SECTION 1

How Long Critical & High Compromises Stay Open

How fast the closed findings closed.



ATTACK PATHS

02

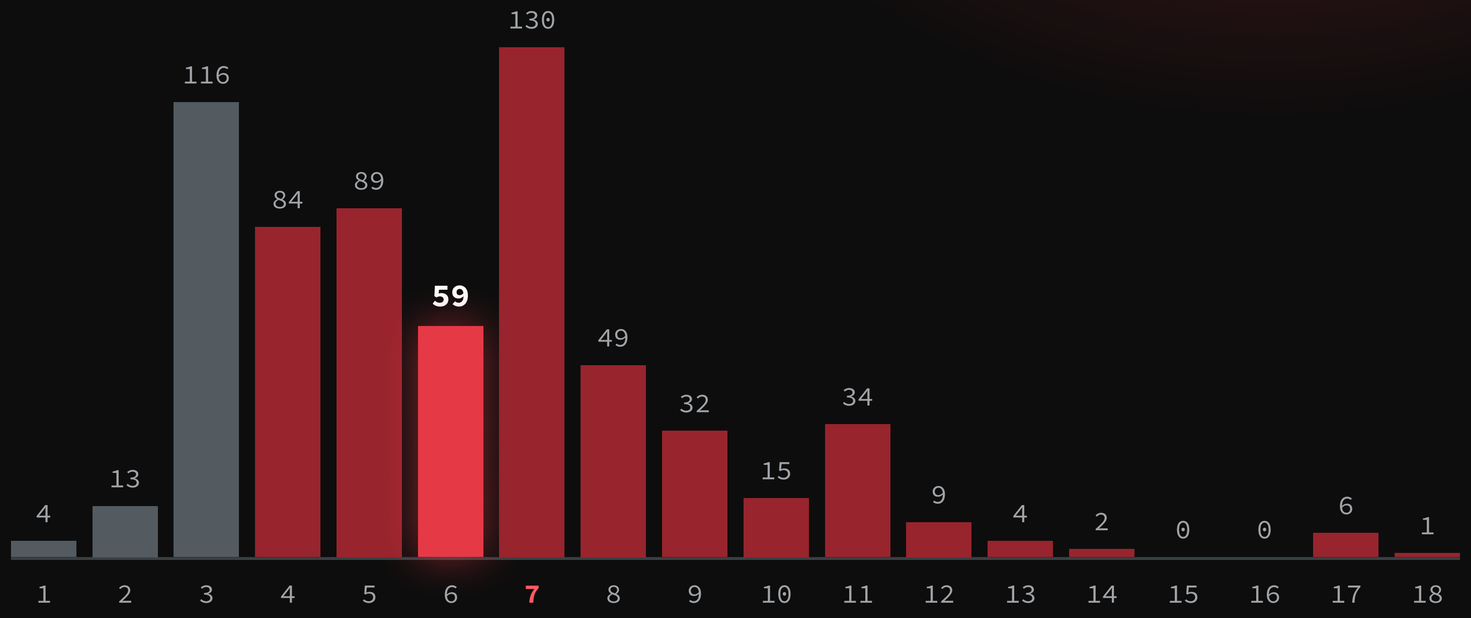
How deep the paths run

6

median steps to the last Critical or High finding

STEPS PER PATH

647 VALIDATED ATTACK PATHS



0.6%
reached it in one step

79%
took four or more steps

647
Validated attack paths

SECTION 2

Longest Attack Paths to Critical

The six longest distinct attacks that started from the internet and reached an Open or Remediated Critical finding, across four organizations. No organization, hostname or proof is shown.

CASE 1

Public-facing app → 18 steps → secrets vault

Credentials leaked in a code repository fed a password spray. Command injection, SSH lateral movement and a privilege escalation followed, ending in the secrets vault.



CASE 2

Vulnerability scan → 12 steps → host private key

An unauthenticated admin API led to command injection, then a container escape that exposed the host's SSH private key.



CASE 3

Public-facing app → 12 steps → personal data

An unauthenticated report revealed account details. Enumeration and brute force produced valid logins that reached personal data.



CASE 4

Public-facing app → 11 steps → arbitrary file read and write

Information disclosure and default credentials led to admin credentials in a configuration backup, then path traversal.



CASE 5

Public-facing app → 10 steps → host filesystem

An unauthenticated configuration endpoint allowed shell injection, then a container escape onto the host and an RSA key.



CASE 6

Public-facing app → 10 steps → customer personal data

Tokens issued without a login opened an API of bulk customer records. A document download then exposed personal data.



SECTION 2

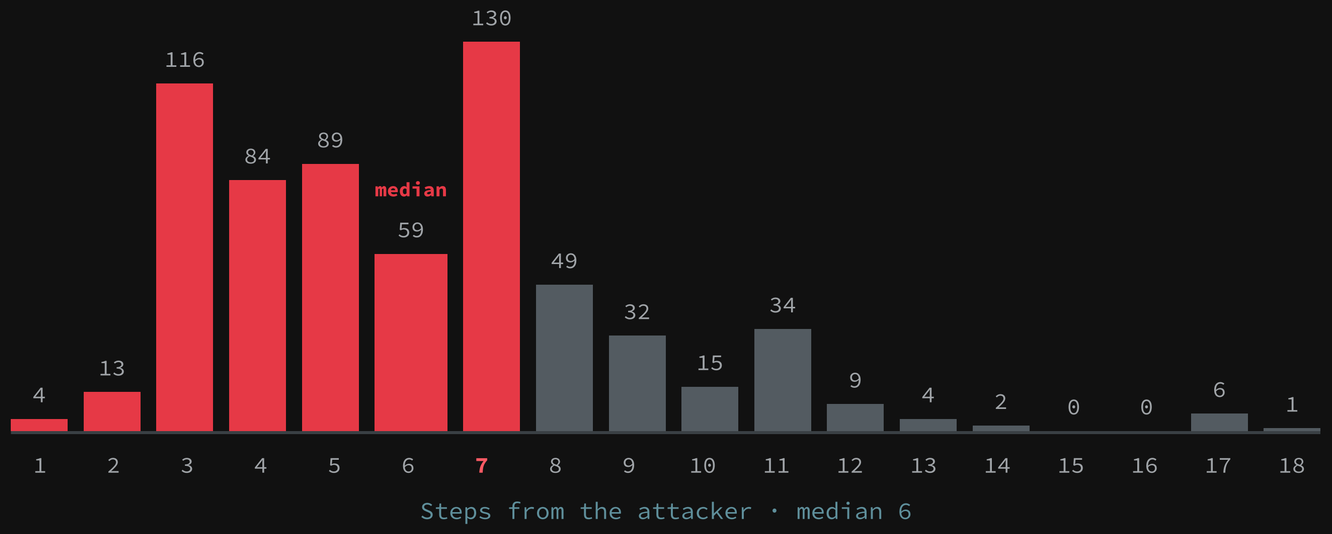
Shorter Attack Paths Do Not Mean Less Damaging

Length measures the attacker’s effort, not the damage at the end of it.



STEPS TO A CRITICAL OR HIGH FINDING

647 validated attack paths across 41 customer environments



77%

of paths reach a Critical or High risk compromise in 7 steps or fewer

THE CASE STUDIES THAT FOLLOW

6 steps	Read/write access to policy data
8 steps	Full Salesforce administrator
8 steps	Write access to the production identity directory

SECTION 2

How attack paths get in

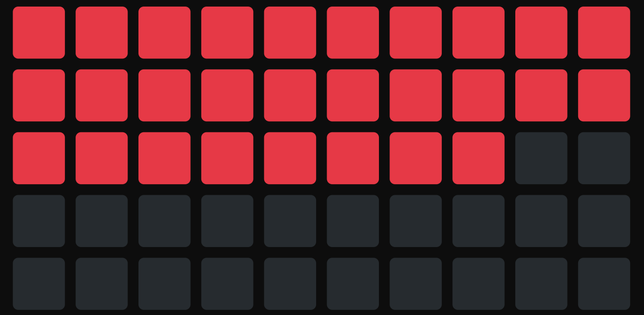
T1190	**Exploit Public-Facing Application**	28 · 427 paths
T1078	Valid Accounts	8 · 25 paths
T1078.001	Default Accounts	7 · 25 paths
T1078.004	Cloud Accounts	6 · 18 paths

These four Initial Access techniques, of 33 starting techniques, begin 495 of the 647 validated attack paths, at 32 of 41 organizations. The other 152 paths start elsewhere. A start is the first technique that gets in; reconnaissance before it is skipped. 122 paths at 21 organizations began with reconnaissance, and 74 of them went next to T1190.

Bar length is organizations; the longest bar is 28.

28

of the 50 organizations had a path that got in through an internet-facing application.



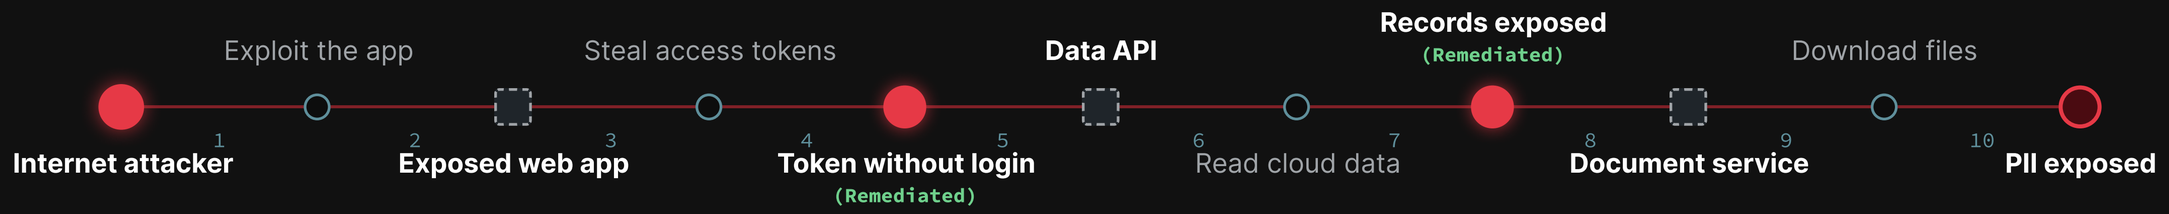
SECTION 2

Every step counts

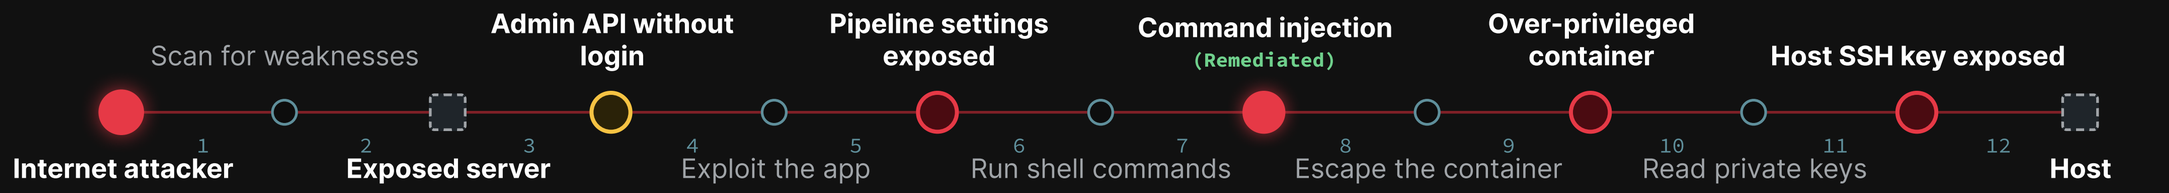
Assets, vulnerability, and technique each count as a step in our attack paths



Case 6 10 steps, exposed web app to customer personal data



Case 2 12 steps, vulnerability scan to the host's private key



Critical High Medium Technique System reached

03

Case Study 1

WHAT HAPPENED

Public JavaScript file → 6 steps → read/write access to policy data

A public insurance web app shipped its Okta client credentials and its payment gateway credentials inside the JavaScript every visitor downloads. The Okta credentials alone were enough to request an access token with read/write scopes for underwriting, policy and document data. The payment gateway credentials created a payment session. Guard's agents demonstrated both.

ENTRY

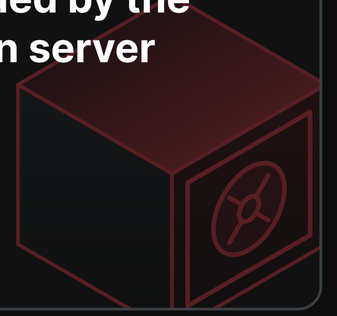
Public insurance web app, no login

TURN

Okta client credentials in its JavaScript

RESULT

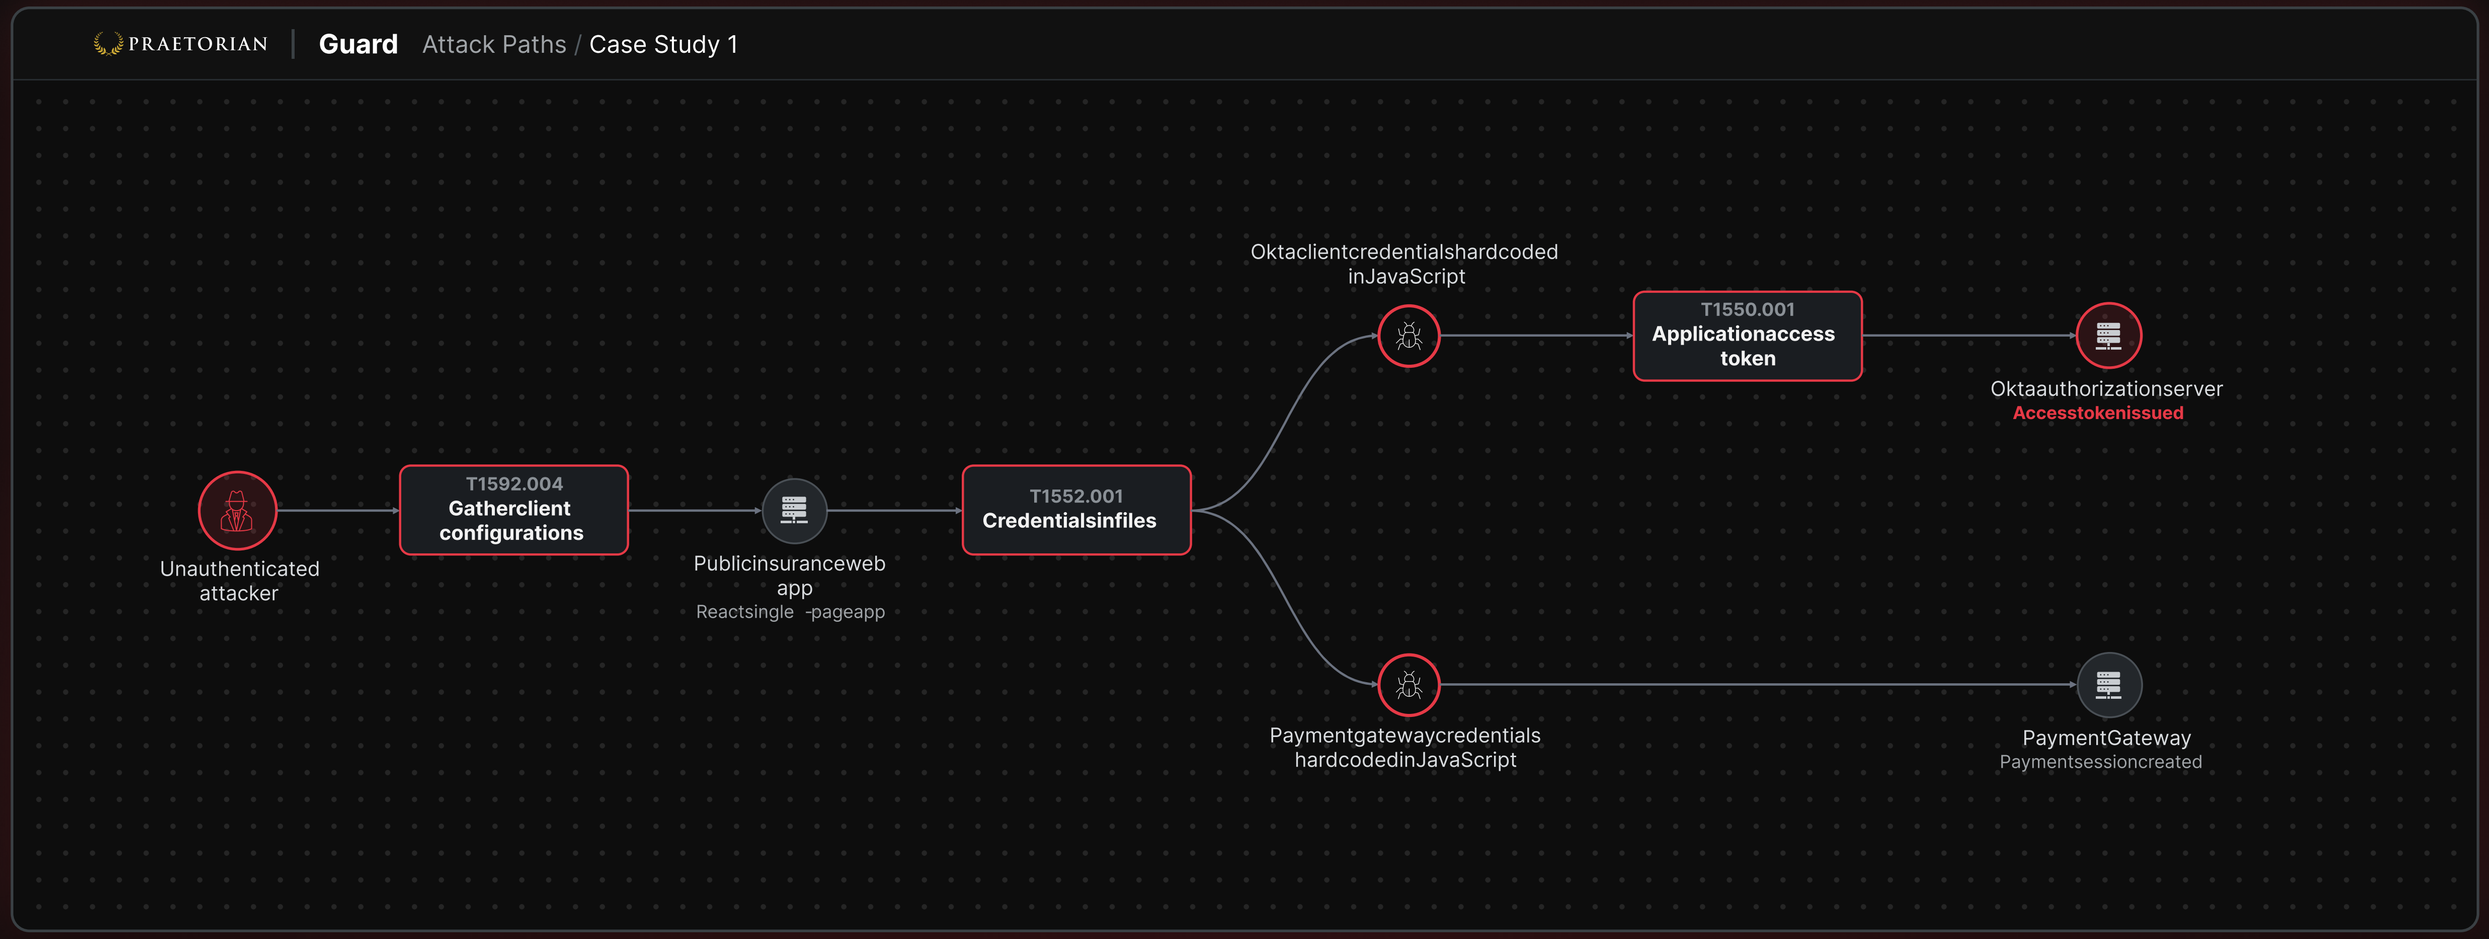
Access token issued by the Okta authorization server



SECTION 3

Case Study 1

Credentials in a public JavaScript file gave read/write access to policy and underwriting data



SECTION 3

Case Study 2

WHAT HAPPENED

Unauthenticated portal → 8 steps → full CRM admin

A guest API that needed no login read any record. Changing a profile ID made the attacker a system administrator who could sign in as any user through single sign-on. The same entry exported payment card data and deleted records.

ENTRY

Salesforce guest API, no login

TURN

A changed profile ID granted system admin

RESULT

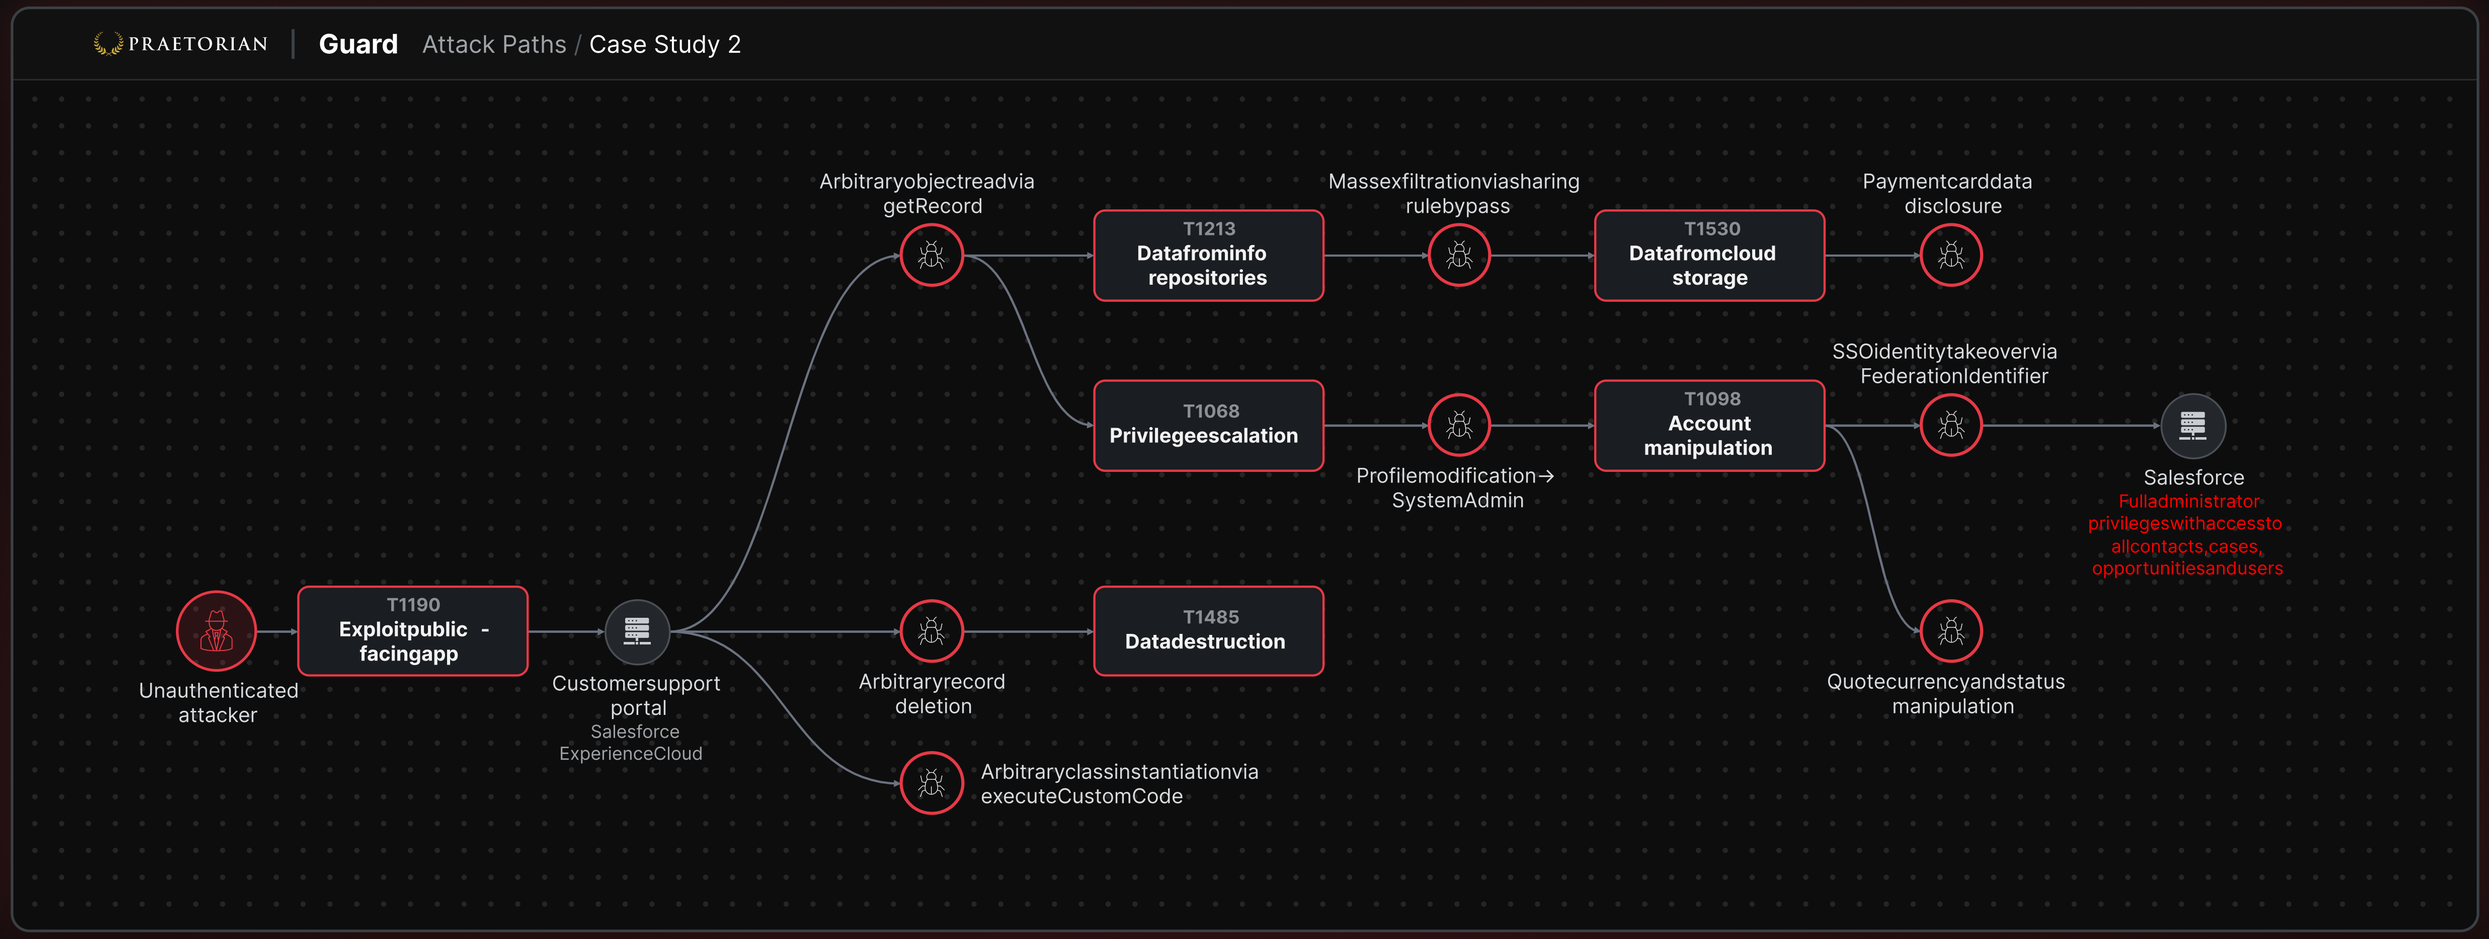
Sign in as any user, export and delete records



SECTION 3

Case Study 2

A guest API that needed no login led to org wide compromise of Salesforce



SECTION 3

Case Study 3

WHAT HAPPENED

Leaked cloud credentials → 8 steps → the identity directory

Service credentials left in files on a web app gave read and write access to every account in the production identity directory. The same credentials listed the privileged admins and opened the reporting workspaces that held client data.

ENTRY

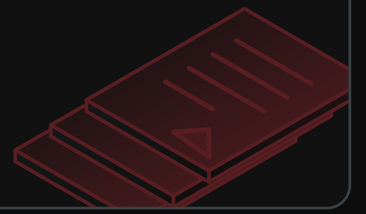
Files on an internet-facing web app

TURN

Service credentials for the cloud identity tenant

RESULT

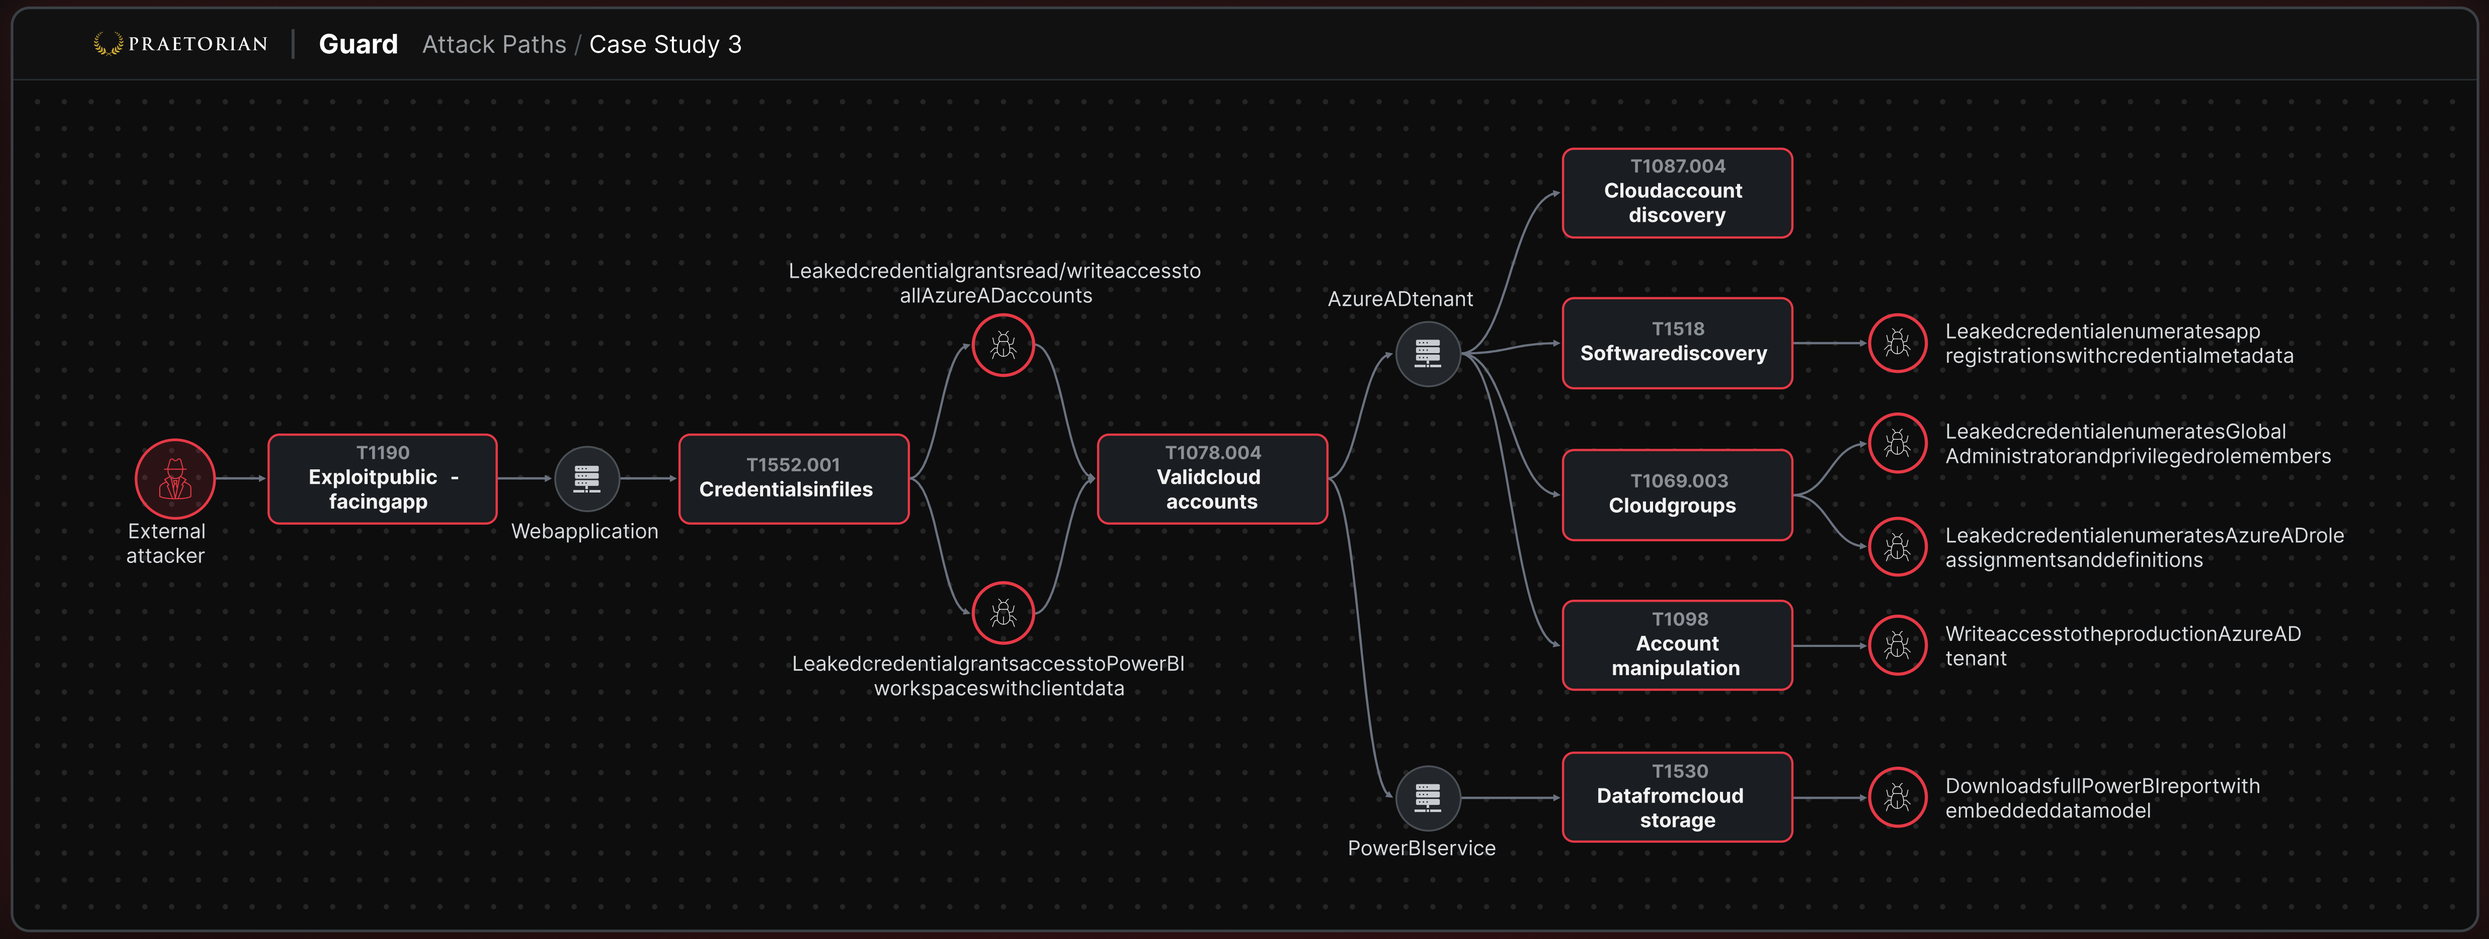
Read and write on every account in the identity directory



SECTION 3

Case Study 3

Credential in Internet facing web file accessible by anyone led to org wide compromise of cloud provider and all customer data



SECTION 3

Case Study 4

Case Study 4 is the Hyper Challenge CTF, an agentic AI hackathon held by a Global 500 company. We nearly skipped it, entered late, and Guard finished it in an afternoon.

WE ALMOST SKIPPED IT

It looked like a mentoring event for students. It was an agentic AI hackathon whose previous edition was billed as the largest yet. The red team track: reach a sensitive customer database from the outside, through a lab planted with decoys.

WE STARTED LATE

We entered 11 days after the competition began, with one person supervising the agent part time and one IP address to start from. We kept a human in the loop by choice, to see how the challenge was built. In full autonomy, the platform would have completed the task even faster.

GUARD FINISHED IN AN AFTERNOON

It chained web flaws and leaked credentials across network boundaries, recognized every decoy as misdirection, and read the production master key.

KEY TAKEAWAY

AI can now go from the internet to full CTF compromise including internal secrets in an afternoon.

That is why it is in this edition: the paths in this report are where attacks like this begin.



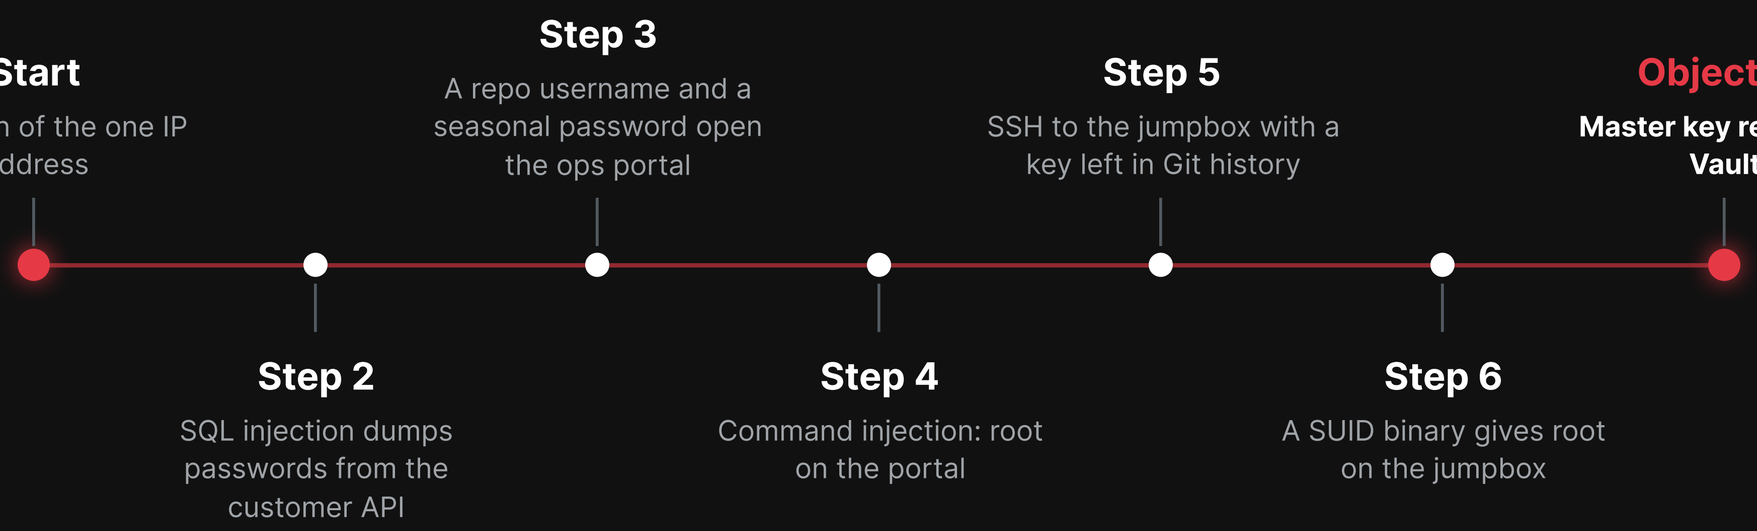
$875	4	5	51
platform cost	network boundaries	decoy flags	logged actions

SECTION 3

One IP Address to Production Secrets

Guard’s path through the lab, in the order the platform logged it.

THE RUN, STEP BY STEP



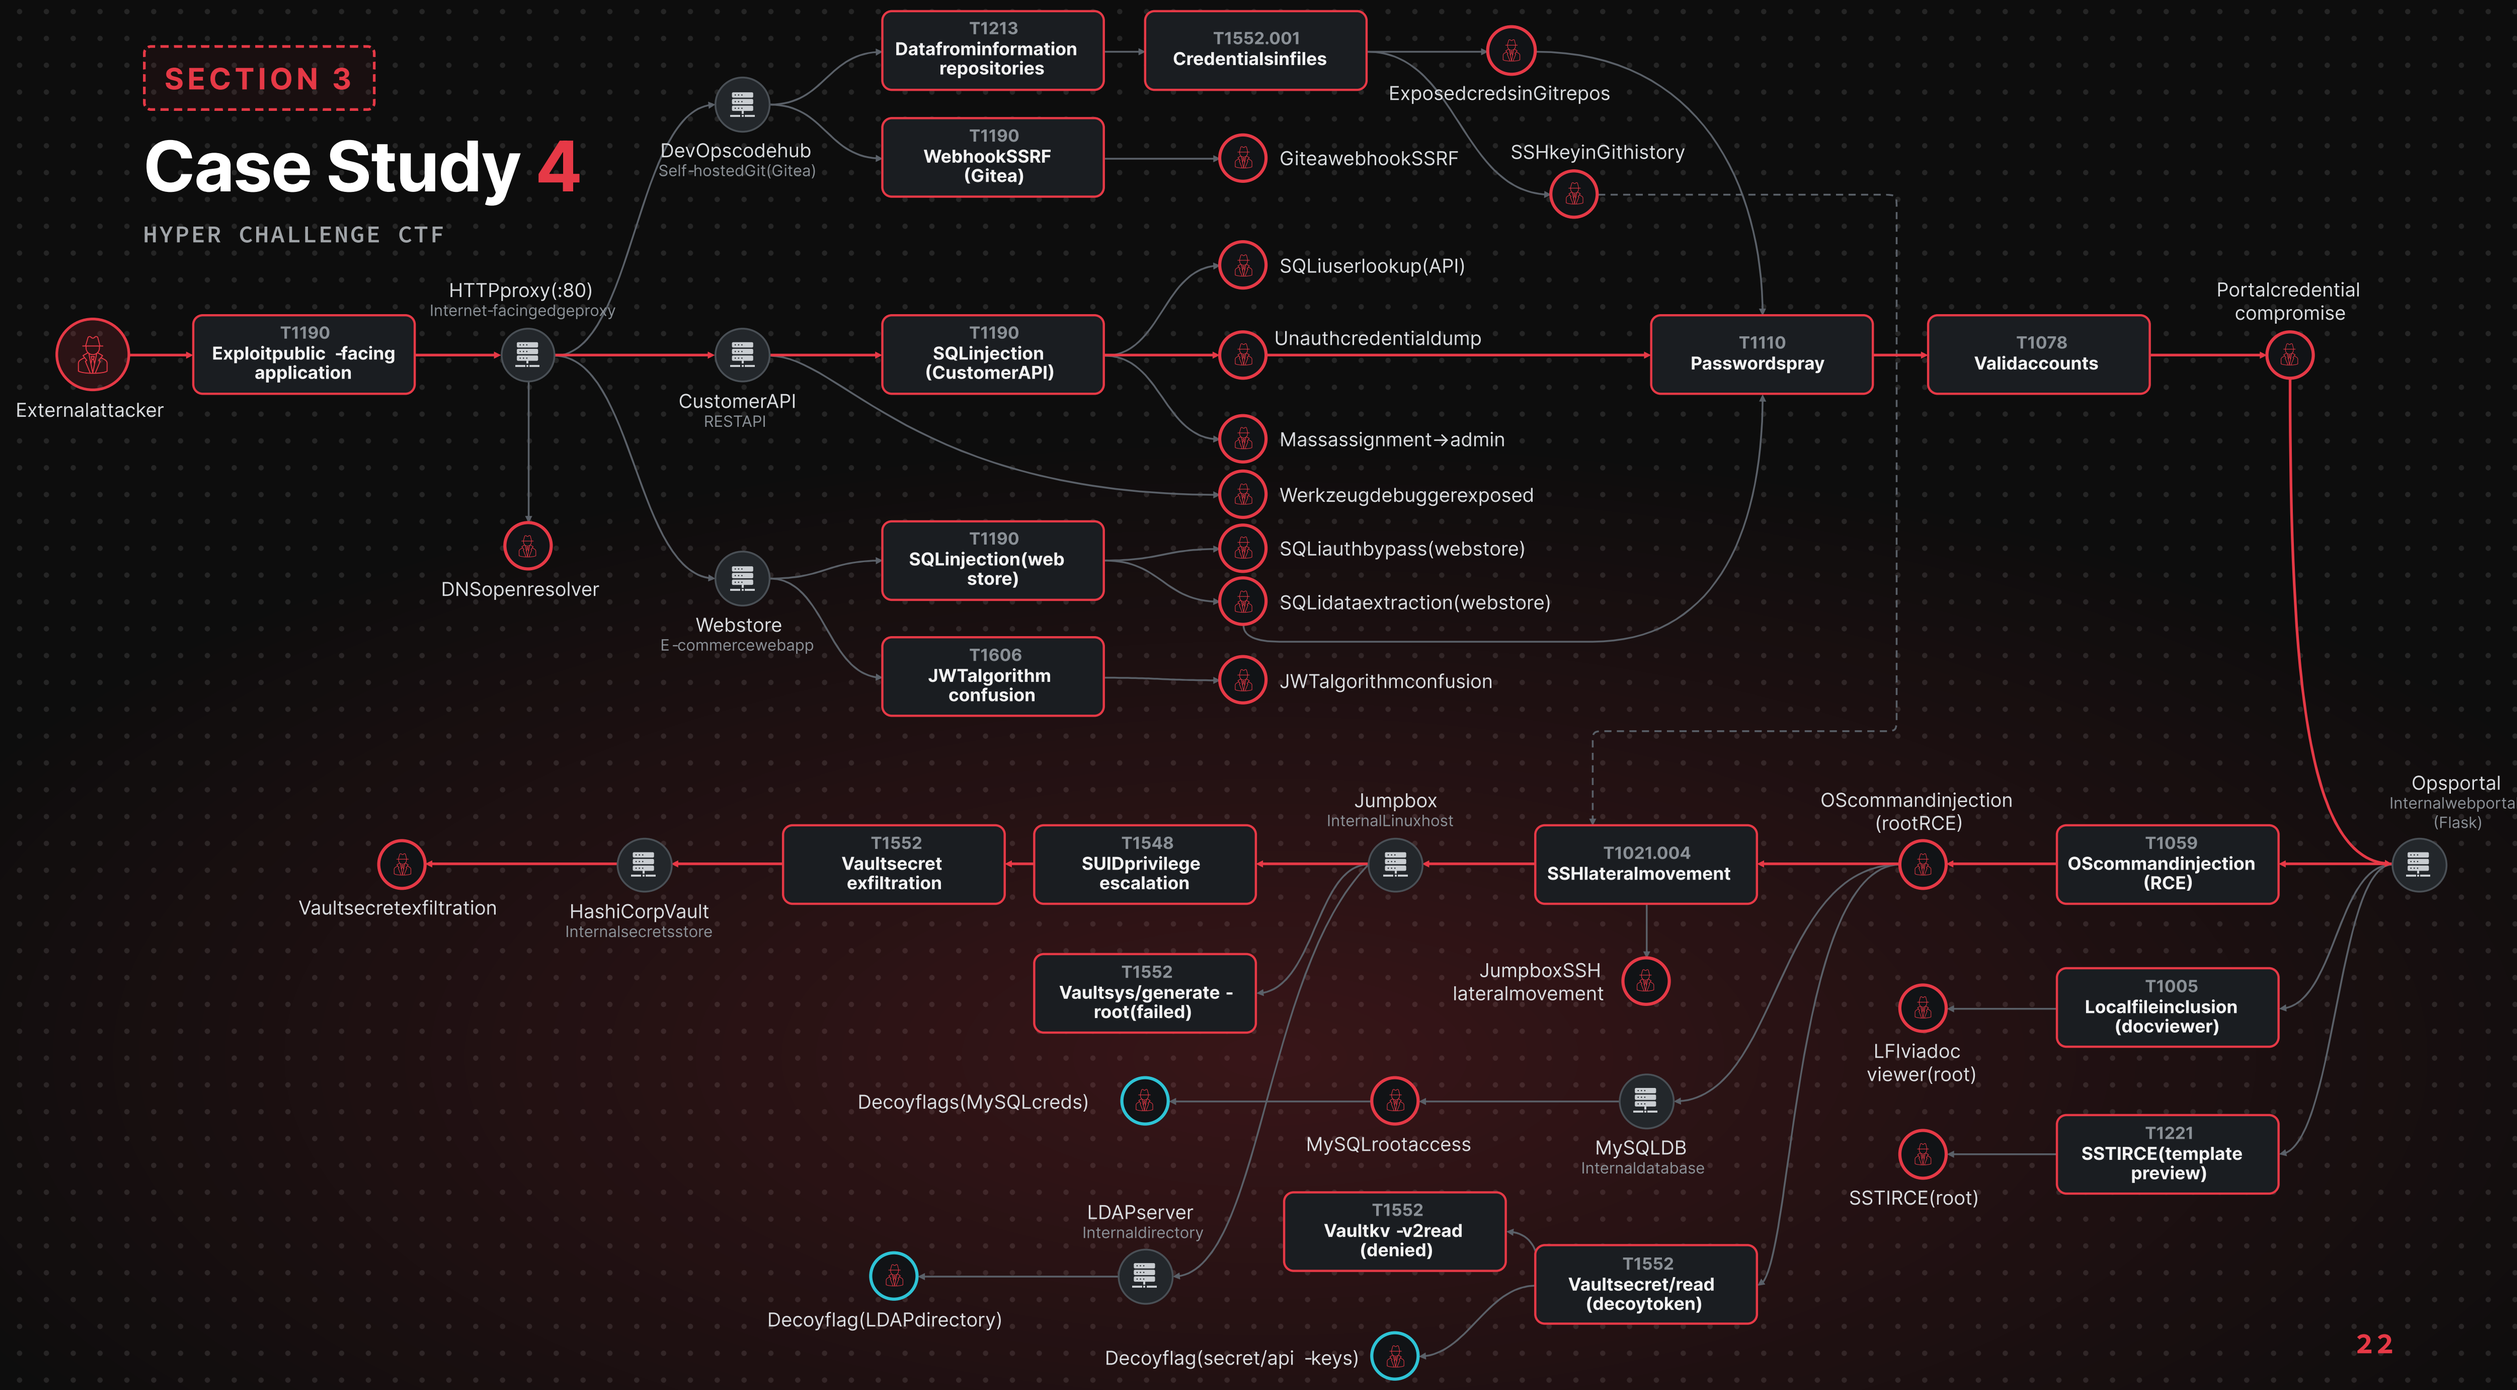
Decoy flags at Vault, the database and the directory: the agent reads each one and keeps going

Run with a human in the loop. Fully autonomous Hannibal mode would not have waited on the operator.

SECTION 3

Case Study 4

HYPER CHALLENGE CTF

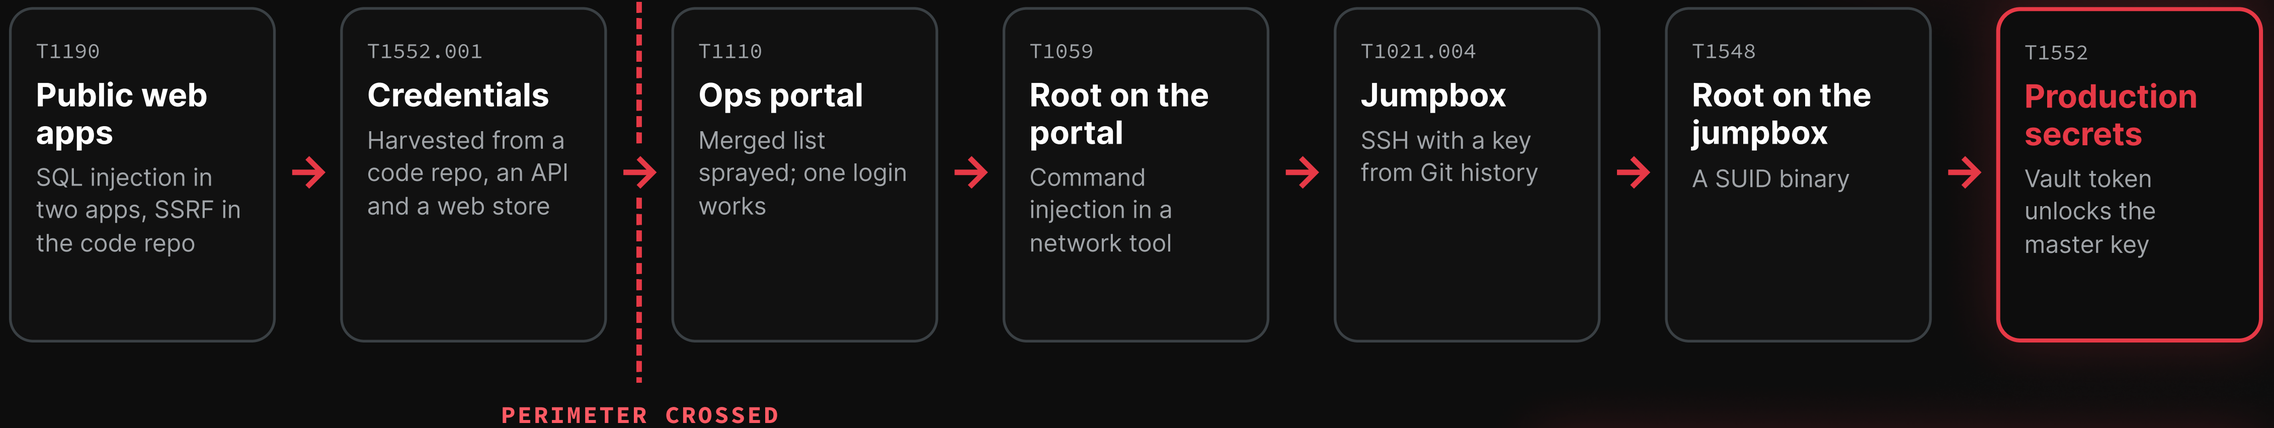


SECTION 3

From the Perimeter to the Inside

The paths in this edition reach a Critical or High finding. In the CTF, the agent kept going: past the perimeter, across network boundaries, to production secrets.

THE CTF CHAIN, FROM ONE IP ADDRESS TO PRODUCTION SECRETS



THIS EDITION

How attackers get in. Median: 6 steps from the attacker to a Critical or High finding on the external attack surface.

THE CTF

What comes after. 7 stages, 4 network boundaries and 5 decoy flags, in one afternoon for $875.

WHAT IT SHOWS

A perimeter finding is a starting point. From one IP address, the agent reached the data an attacker wants.

SECTION 4

Threat Actor Comparison

04

We mapped the techniques behind our compromises to eight active groups: the six busiest ransomware crews of Q2 2026, plus two that target SaaS and developer pipelines.

01
Qilin
G1050 · Water Galura

02
The Gentlemen
G1058

03
DragonForce
No MITRE entry yet

04
Akira
G1024

05
LockBit 5.0
S1202 · software only

06
INC Ransom
G1032

07
ShinyHunters
G1057

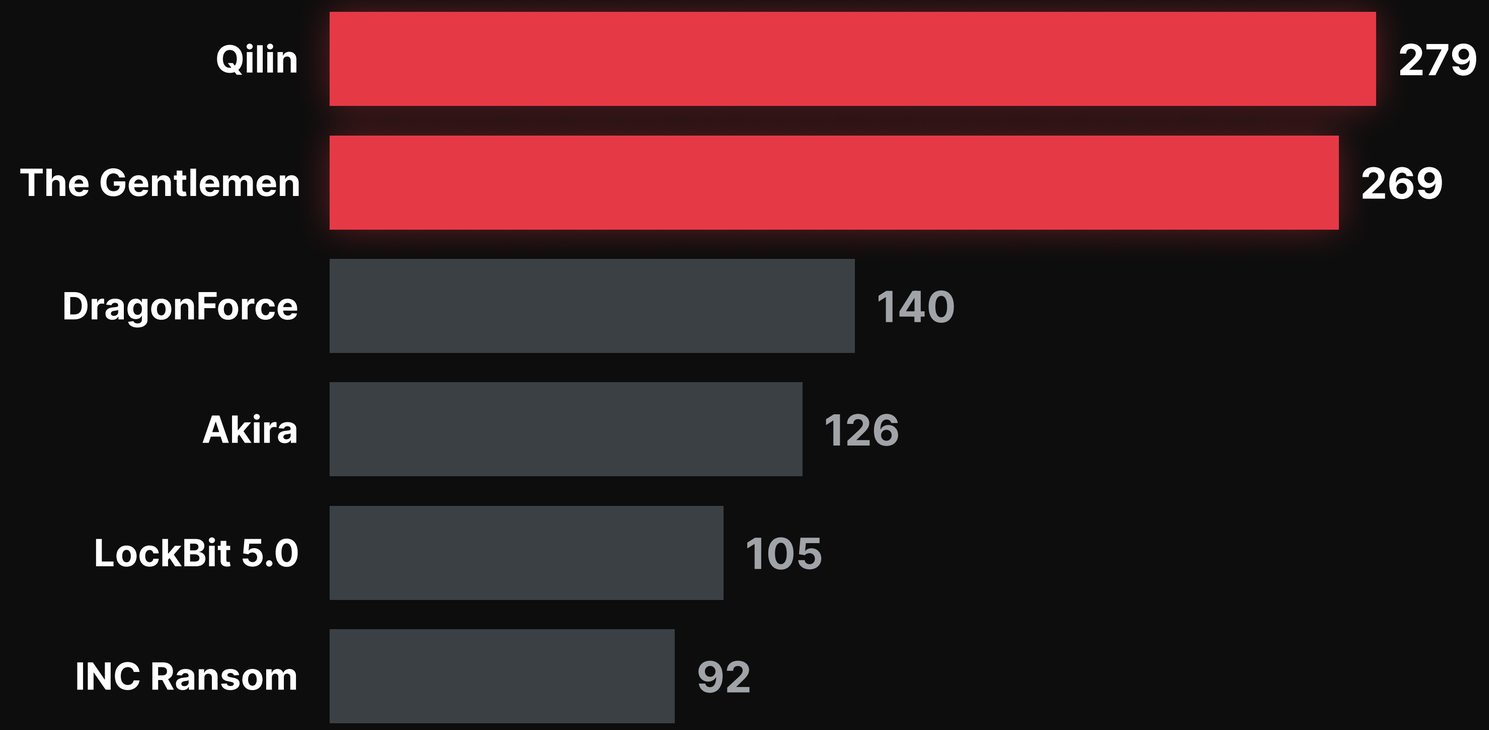
08
TeamPCP
G1056



SECTION 4

Two crews nearly doubled the next group

Victims named on ransomware leak sites, top six groups, Q2 2026



The Gentlemen grew 62% from Q1 and out-posted Qilin in June, though Qilin led the quarter.

93 groups posted at least one victim in Q2, up from 71 in Q1.

All six break in through internet-facing flaws and stolen logins, compared on the next slide.

SECTION 4

All eight use our top two Initial Access techniques

Initial Access techniques only (MITRE ATT&CK TA0001): the six most active ransomware groups of Q2 2026, plus ShinyHunters and TeamPCP, against the ways our agents got in.

GROUP	Q2 2026 VICTIMS	T1190 exploit app · 28/50	T1078 valid accounts · 8/50	T1566 phishing · not in this edition	T1133 remote access · 3/50	MAIN WAY IN
Qilin	**279**	**Yes**	**Yes**	**Yes**	**Yes**	Fortinet flaws, VPN logins
The Gentlemen	**269**	**Yes**	**Yes**	—	**Yes**	FortiGate bypass, brute-forced VPN
DragonForce	**140**	**Yes**	**Yes**	**Yes**	**Yes**	SQL and edge servers, MSPs
Akira	**126**	**Yes**	**Yes**	**Yes**	**Yes**	SonicWall and Cisco VPNs
LockBit 5.0	**105**	**Yes**	**Yes**	**Yes**	**Yes**	Citrix Bleed, Fortinet, RDP
INC Ransom	**92**	**Yes**	**Yes**	**Yes**	**Yes**	Citrix NetScaler, Fortinet EMS
ShinyHunters	Not ransomware	**Yes**	**Yes**	**Yes**	—	Vishing for SSO logins
TeamPCP	Sells access	**Yes**	**Yes**	—	—	React2Shell, poisoned dev tools

T1133: organizations with a Critical or High finding that exposed SSH or SFTP login access, open or resolved.

Group techniques: MITRE ATT&CK, CISA and FBI advisories, vendor threat reports

SECTION 4

Attackers hold three of our CVEs

Of the seven CVEs our agents used on page 03, these three are also in named groups' hands.

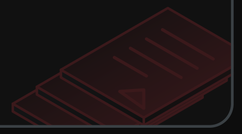
React / Next.js RCE

CVE-2025-55182

WHO ELSE USES IT

TeamPCP (Operation PCPcat); Weaxor and BQTLock ransomware; China-nexus groups Earth Lamia and Jackpot Panda.

EXPLOITED SINCE DEC 2025



Exchange ProxyLogon

CVE-2021-26855

WHO ELSE USES IT

HAFNIUM and other state groups. DearCry ransomware was deployed through it in March 2021.

EXPLOITED SINCE 2021



WordPress REST route

CVE-2026-63030

WHO ELSE USES IT

INC Ransom and The Gentlemen affiliates built an exploit. Exploited in the wild per CISA KEV (Jul 2026); not yet tied to a named group’s victims.

EXPLOITED SINCE JUL 2026



No named-group activity found for the other four: Metabase, GitLab and the two PHP2Shell CVEs.

Sources: AWS; Halcyon; SOCRadar (Oct 2026); CISA AA21-062A; CISA KEV; Sophos

SECTION 4

ShinyHunters: steal the data, then extort

MITRE G1057 · aka ShinyCorp, UNC6240, Bling Libra · active since 2019

KNOWN FOR

Pay-or-leak extortion of data stolen from cloud and SaaS platforms. No encryption.

HOW IT GETS IN

Voice phishing for SSO logins and MFA codes; stolen OAuth tokens from SaaS integrations.

WHO IT HAS HIT

140+ organizations since 2025 per the FBI, including Salesforce tenants reached through Salesloft Drift and Gainsight tokens. At least $70M paid in extortion.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8
T1552.001	**Credentials In Files** Finds passwords, keys and tokens left in files, configs and repos.	7

Showing the 3 most common of 8 starts ShinyHunters shares with our agents.

SECTION 4

TeamPCP: steal the keys, then sell access

MITRE G1056 · aka PCPCat, ShellForce, DeadCatx3, UNC6780 · active since 2025

KNOWN FOR

Worm-driven credential theft and supply-chain attacks on CI/CD pipelines.

HOW IT GETS IN

Exposed Docker APIs, Kubernetes, Ray and Redis; then trojanized open-source tools and packages.

WHO IT HAS HIT

Trivy, KICS, LiteLLM and the Telnyx SDK in early 2026, 66+ npm packages and the European Commission's AWS environment. Sells access to ransomware partners Vect and CipherForce.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8
T1552.001	**Credentials In Files** Finds passwords, keys and tokens left in files, configs and repos.	7

Showing the 3 most common of 9 starts TeamPCP shares with our agents.

SECTION 4

Qilin: encrypt, then leak

MITRE G1050 (Water Galura) · aka Agenda, Gold Feather · active since 2022

KNOWN FOR

Ransomware-as-a-service: encrypts Windows, Linux and ESXi, then extorts with a leak site.

HOW IT GETS IN

FortiGate flaws CVE-2024-21762 and CVE-2024-55591, leaked VPN credentials on portals without MFA, phishing.

WHO IT HAS HIT

1,358 claimed victims from April 2025 to March 2026, up 443%. Its 2024 attack on Synnovis cancelled 6,000+ London hospital appointments.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8
T1083	**File and Directory Discovery** Lists files and folders to find what to steal or encrypt.	1

SECTION 4

The Gentlemen: in via the firewall

MITRE G1058 · run by the admin “hastalamuerte” · active since Sep 2025

KNOWN FOR

Fast-growing ransomware-as-a-service: steals data for weeks, then encrypts and leaks.

HOW IT GETS IN

FortiGate auth bypass CVE-2024-55591, brute-forced Fortinet VPN, OWA and Microsoft 365 logins, bought access.

WHO IT HAS HIT

352 claimed victims by May 10, 2026, and 269 in Q2 alone, up 62%. Out-posted Qilin in June 2026.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8

From MITRE G1058 and public reporting.

Sources: S2W; Cybereason; CUHK ITSC; Group-IB; DeepStrike Q2 2026

SECTION 4

Akira: in through the VPN

MITRE G1024 · code overlaps with Conti · active since 2023

KNOWN FOR

Double extortion across Windows, Linux, ESXi and Nutanix. About $244M in ransoms by late 2025.

HOW IT GETS IN

SonicWall SSL VPN flaw CVE-2024-40766 and Cisco ASA flaws, VPN accounts without MFA, brute force, bought access.

WHO IT HAS HIT

About 980 leak-site victims in 2025 and 126 in Q2 2026. Arete saw it as the most active group in September 2026.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8
T1566.002	**Spearphishing Link** Sends a targeted link that steals logins. Not measured in this edition.	—

SECTION 4

DragonForce: ransomware for rent

Cartel-style RaaS · works with Scattered Spider affiliates · active since 2023

KNOWN FOR

A cartel model: affiliates run their own brands on DragonForce's ransomware and leak site.

HOW IT GETS IN

Unpatched edge and SQL servers, stolen VPN and RDP logins, help-desk impersonation, compromised MSPs.

WHO IT HAS HIT

253 claimed victims in H1 2026 and 590+ in total. Its 140 in Q2 2026 were the third-most of any group.

STARTS THIS GROUP SHARES WITH OUR AGENTS		ORGS OF 50
T1190	**Exploit Public-Facing Application** Uses a flaw in an internet-facing app or service to get in.	28
T1078	**Valid Accounts** Signs in with real stolen, reused or default credentials.	8
T1566.002	**Spearphishing Link** Sends a targeted link that steals logins. Not measured in this edition.	—

TAKE AWAYS

05

Key takeaways

01 —

Applications remain the most common way in

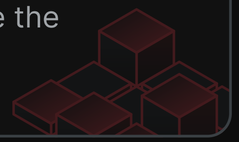
Exploiting a public-facing app was how our agents got in at 28 of 50 orgs, more than any other technique.



02 —

Organizations don’t know their internet footprint

SSH was open to the internet at 42 of 50 orgs, SMB at 36, FTP at 35 and RDP at 32. SMB and RDP rarely need to face the internet at all.



03 —

Credentials and secrets are left exposed

Authentication flaws and default credentials sat on attack paths at 23 of 50 orgs, and exposed secrets at 16.



04 —

Organizations can’t patch Critical CVEs fast enough

All 7 CVEs Guard exploited are now on CISA’s KEV list, two of them since 2021. Closed Criticals took a median 13.9 days.



05 —

Proven paths matter more than vulnerability counts

79% of the 647 paths Guard proved took four or more steps. Scanners score each finding alone and miss the chain.



06 —

Threat actors use the same ways in our agents do

All eight crews we track rely on exploiting public-facing apps and logging in with valid accounts. Guard looks for the same ways in they do.



TAKE AWAYS

Predictions

01 Web flaws will remain one of the most common ways in

Zero days are not needed. Common web flaws in internet-facing applications were the way in at 28 of 50 organizations and will remain the most common way in.

02 That said, zero days will drop at record speed

Black October brought Microsoft's largest Patch Tuesday ever, with an exploited Windows zero day, while PHP2Shell and Node2Shell were exploited within days. CISA has listed more exploited flaws in 2026 than in all of 2025.

03 Exploitation before disclosure becomes routine

Honeypots logged PHP2Shell exploit attempts before publication. As the circle in the know increases leading up to disclosure, leak likelihood increases. We knew PHP was coming before it went public.

04 AI driven attacks will outrun human-speed defense

The operational tempo is increasing. Teams that triage, determine asset ownership, and patch by hand will fall further behind and their organizations will become increasingly susceptible to AI attack.

05 Reachability and exploitability dictate priority

Attackers now move at machine speed and teams cannot patch everything. The only vulnerabilities that deserve focus are the ones an attacker can reach and exploit.

06 Securing LLMs and MCP servers will be a major headache

LLMs are an asset, a tool, and a target. Many LLMs are deployed insecurely and are easy to compromise. Like with every technology wave, adoption will out-run security.



TAKE AWAYS

Three recommendations for leadership



PATCH CYCLES HAVE COLLAPSED • THREE MOVES TO MAKE NOW

01 Reexamine capital allocation towards cybersecurity

Traditional cyber spend no longer covers the threat. Bain estimates many organizations will need up to 2× current levels. The amount of additional capital is unique to each organization, but in all cases the answer is “More.”

02 Shift to continuous, AI-driven penetration testing

Deploy agentic AI in security operations to continuously test critical assets from the attacker’s perspective. With continuous coverage, focus on reachability and exploitability vs. theoretical risk and best practices. Prioritize high-value, forward-facing assets.

03 The patch window has vanished. Automate remediation

Remediation is now the bottleneck. AI-driven discovery surfaces issues faster than human teams can triage and patch. Invest in automated remediation: AI-proposed fixes routed straight into source control and ticketing, with retest-on-deploy confirming closure. The board metric shifts from “number of closures” to “time from discovery to verified fix.”

Source: Bain (21 Apr 2026). https://www.bain.com/insights/claude-mythos-and-ai-cybersecurity-wake-up-call/

06

SECTION 6

Praetorian Guard

Our Continuous Offensive Security Platform



PRAETORIAN

You'd never run your EDR 2 weeks a year. Why are you still running offensive security that way?



Your attack surface changes every week. A test that runs once a year only sees the week it ran.

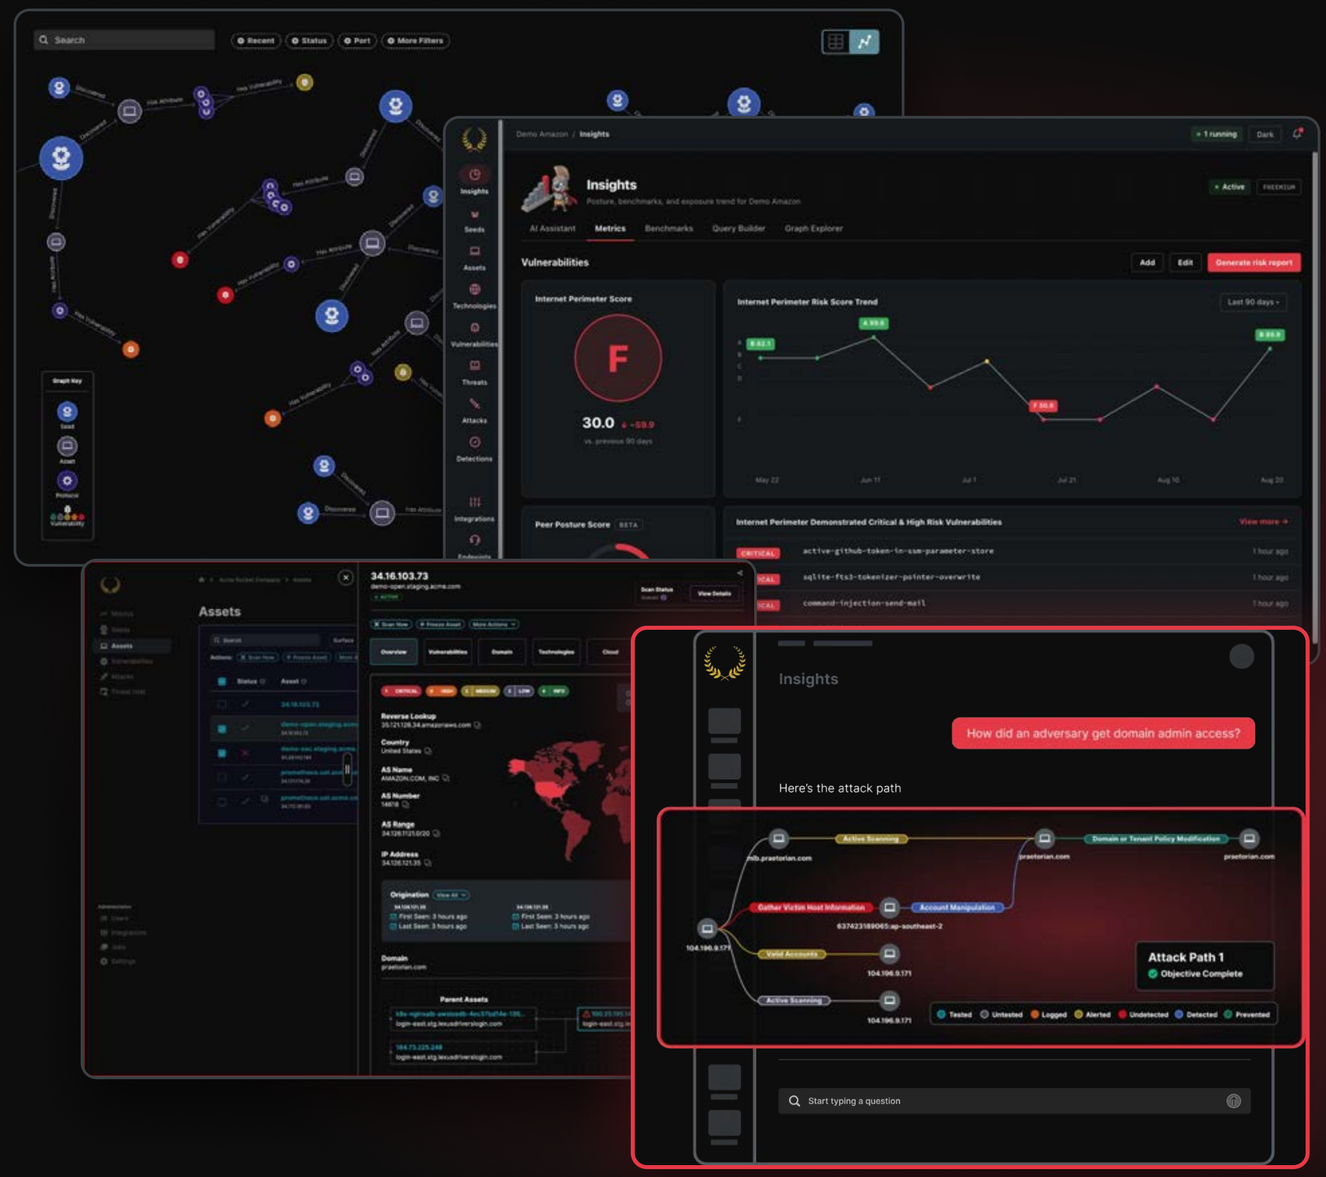
SECTION 6

Meet Praetorian Guard.

An AI First Offensive Security Platform

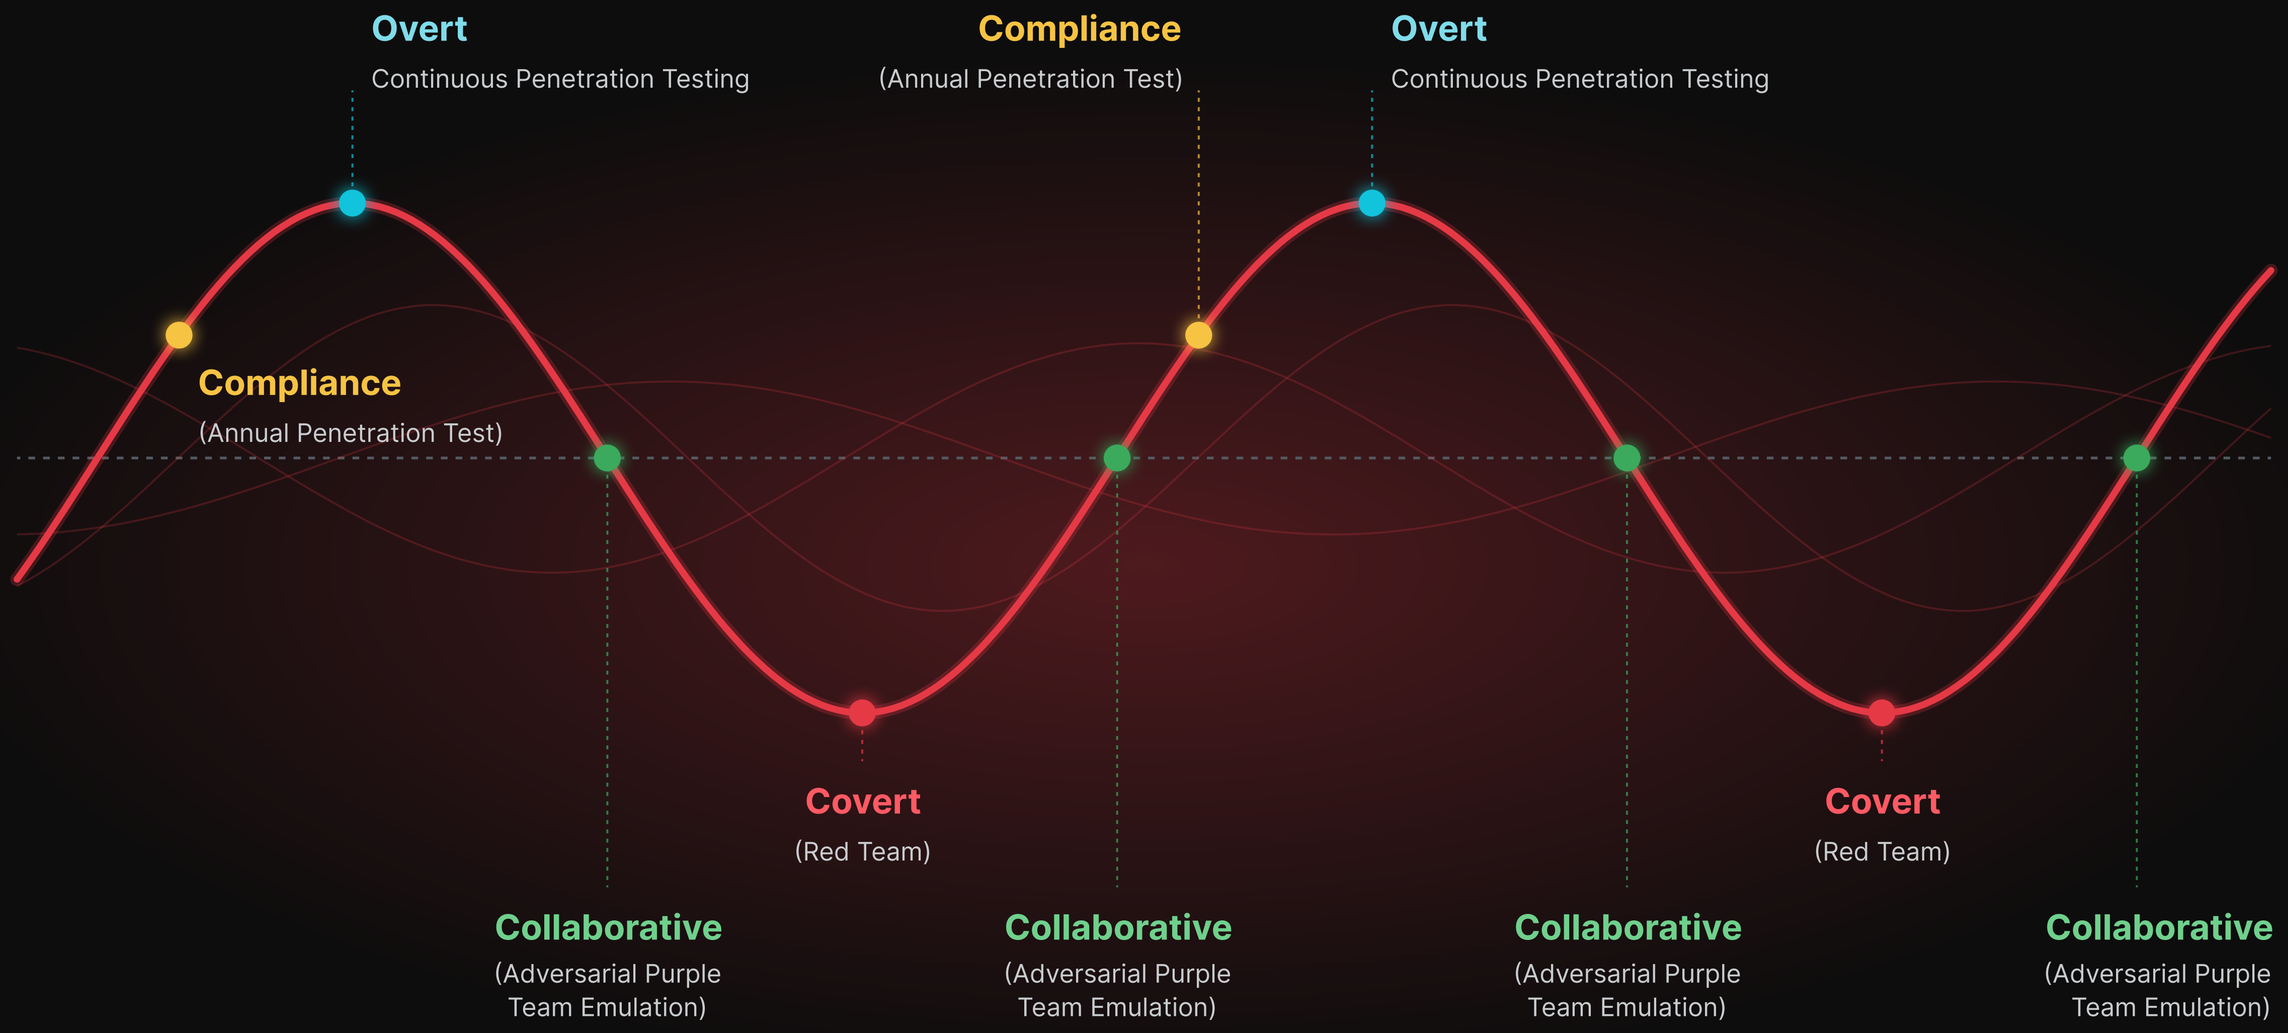
Continuously penetrate your environment to find the breach points

1. Risk discovery at speed and scale of AI
2. Offered as a standalone or managed solution
3. Risk prioritization through demonstrated impact
4. Seamless integration with your security tech stack
5. Cost efficiencies through AI and automation

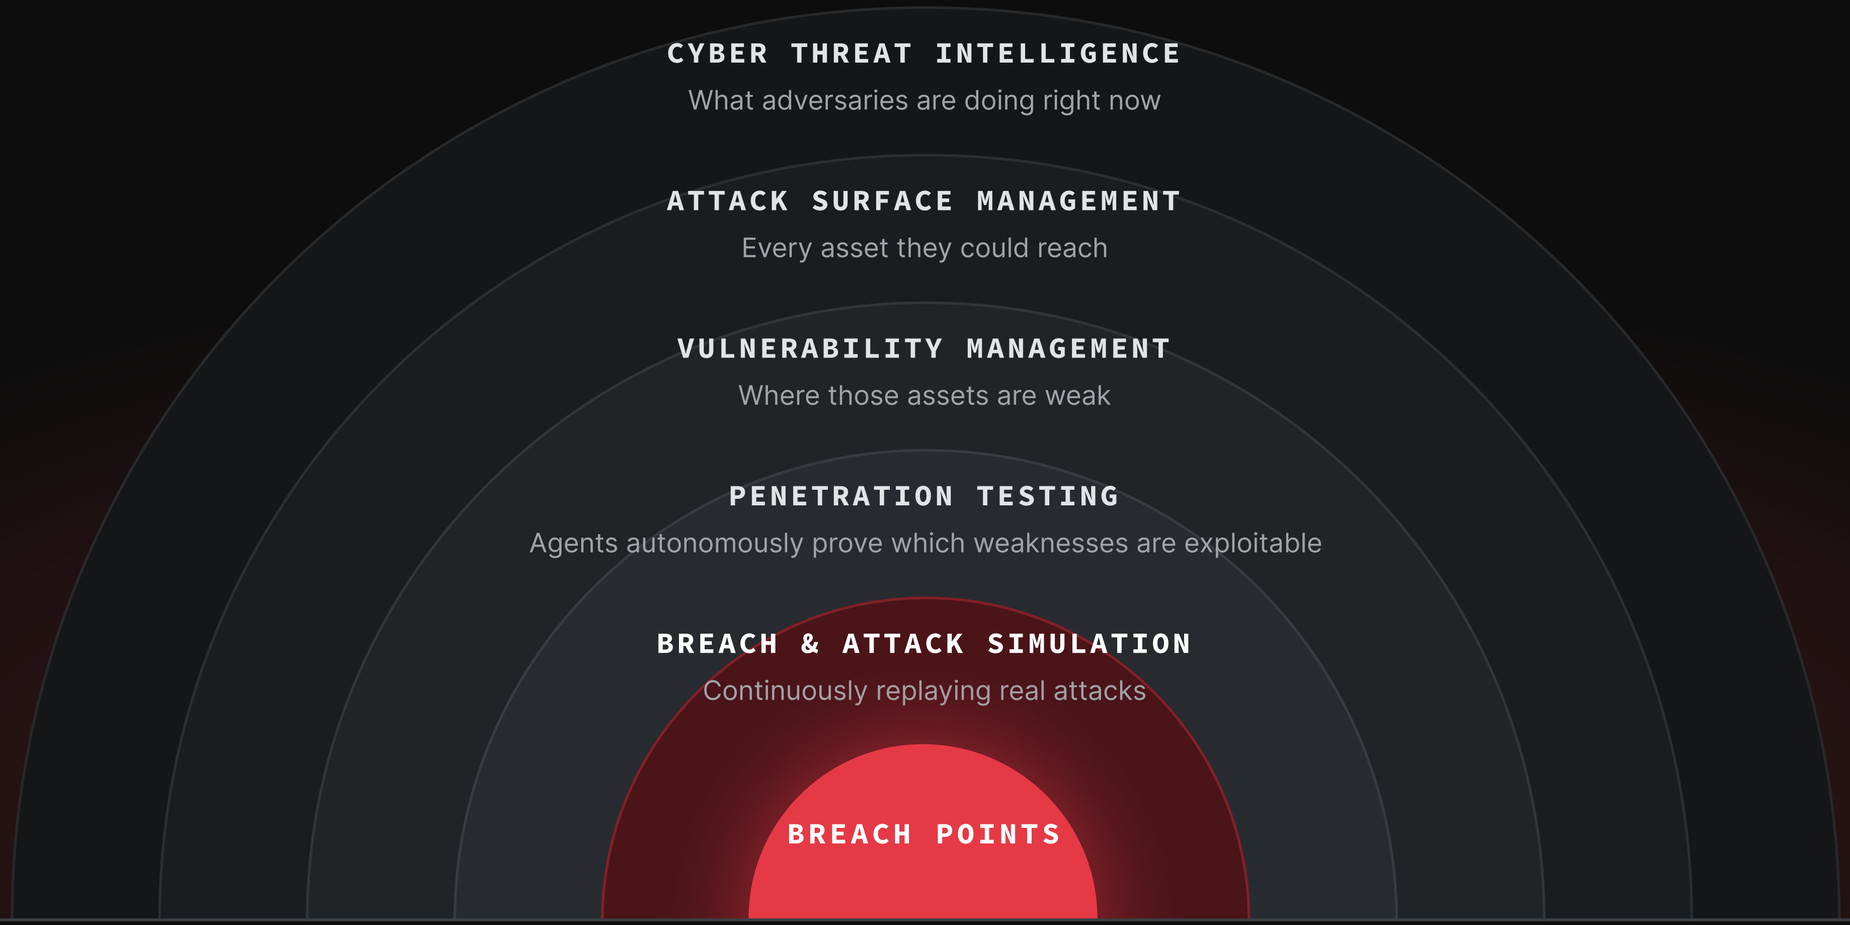


SECTION 6

The Sine Wave of Continuous Offensive Security



Multiple point products as modules under one platform

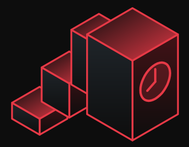


Leverage our platform modules, integrate your existing tech stack, or mix a combination of the two. Every ring is an input; we're focused on the outcome

SECTION 6

Free Pilot

Your free, two-week proof of value runs all three phases on your own environment — **as easy as 1 · 2 · 3.**



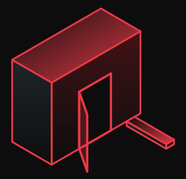
1 PHASE 01 · ASSETS

Enumerate Assets

Map your external attack surface — domains, IPs, cloud, and exposed services.

Vulnerability enumeration follows automatically.

Enumerate Your Assets



2 PHASE 02 · VULNERABILITIES

Locate Vulnerabilities

Every discovered asset is scanned for potential weaknesses that could lead to a compromise.

This runs on its own.

Enumerate Your Risks



3 PHASE 03 · COMPROMISE

Find the Compromise

An autonomous fleet of agents finds the breach points and demonstrates where compromise paths exist

Find The Compromise

SECTION 6

See Praetorian Guard in Action

Watch Demo Now

Open the link above to watch the recorded product demo any time.



APPENDIX

Definitions

Step

One move on an attack path, whether it uses a vulnerability, a technique, or an asset. The steps slide counts from the attacker to the last Critical or High finding; the case studies count the longest route drawn on each graph.

Attack path

A chain of steps an agent demonstrated, from where the attacker started to what they reached.

Compromise

Unauthorized access that results in the breach of a system's confidentiality or integrity. Availability is excluded from the triage because platform guardrails prohibit denial-of-service and load testing without explicit pre-approval and environmental controls.

Open and Remediated

Open is not yet fixed; Remediated is fixed. Risks closed any other way are left out. By the numbers and remediation use status at the freeze; paths, steps, exposures and case studies use status when the data was pulled.

Organization

A legal entity whose organization is considered "Enterprise" by size and revenue including the Global 500 and Fortune 500

Campaign

One autonomous attack campaign against a specific customer. A campaign can be against a specific target or targets, a specific attack surface, or all attack surfaces.

Agent run and subagent run

An agent run is one top-level agent conversation. A subagent run is one an agent started to take on part of its task.

Starting technique

The MITRE ATT&CK technique a path used to get in, such as T1190; reconnaissance before it is skipped.

Risk class

A group of similar findings, matched by a shared CVE or a saved grouping decision.

Significant vulnerability

A Critical or High-risk vulnerability on the external attack surface that Praetorian Guard found and confirmed as real, either still open or since remediated. Findings in triage, rejected as false positives, closed any other way, imported from third-party tools, or credentials no longer working are not counted.

APPENDIX

MITRE ATT&CK techniques

T1190 Initial Access

Exploit Public-Facing Application

An attacker uses a bug, flaw or misconfiguration in an internet-facing system, such as a web app, VPN, API or database, to get a first foothold.

Defend: Patch exposed services quickly, put web apps behind a WAF and keep off the internet what does not need to be on it.

attack.mitre.org/techniques/T1190

T1078 Initial Access, Persistence

Valid Accounts

An attacker signs in with legitimate credentials that were stolen, reused, bought or left at default, so the activity looks like a normal user.

Defend: Require phishing-resistant MFA, remove default accounts and alert on unusual sign-ins.

attack.mitre.org/techniques/T1078

T1552.001 Credential Access

Credentials In Files

Part of Unsecured Credentials. An attacker searches local files, config files, scripts and code repositories for stored passwords, API keys and tokens.

Defend: Keep secrets in a vault, scan repositories for secrets and rotate any key found in code.

attack.mitre.org/techniques/T1552/001

T1566.002 Initial Access

Spearphishing Link

Part of Phishing. An attacker sends a targeted email or message with a link to a fake login page or a malicious download.

Defend: Filter and rewrite links, train staff to report lures and require MFA so a stolen password alone fails.

attack.mitre.org/techniques/T1566/002

T1083 Discovery

File and Directory Discovery

An attacker lists files and folders on a host or share to find valuable data and decide what to steal or encrypt.

Defend: Hard to block outright; alert on bulk enumeration and unusual access to file shares.

attack.mitre.org/techniques/T1083

T1133 Initial Access, Persistence

External Remote Services

An attacker signs in through an internet-facing remote-access service, such as a VPN, SSH, RDP or Citrix gateway, often with stolen or default credentials.

Defend: Put remote access behind phishing-resistant MFA, take services off the internet that do not need to be there and alert on unusual remote sign-ins.

attack.mitre.org/techniques/T1133