Hive Pro

Threat Level

Red

HiveForce Labs

THREAT ADVISORY

ATTACK REPORT

ShadowRay Strikes Back: Inside the Multi-Purpose Ray Cluster Takeover

Date of Publication	Admiralty Code	TA Number
November 25, 2025	A1	TA2025358

Summary

Attack Discovered: September 2024
Targeted Countries: Worldwide
Malware: XMRig
Campaign: ShadowRay 2.0
Attack: The ShadowRay 2.0 campaign marks a striking reminder of how quickly threat actors are evolving in the AI era, turning exposed Ray clusters into a global playground for cryptojacking, data theft, and multi-purpose botnet activity. By exploiting unresolved flaws in the Ray ecosystem, attackers leaned on AI-generated payloads, DevOps-style malware pipelines, and Ray's own orchestration features to spread silently across thousands of nodes. Their operation blended stealth and scale, masquerading miners as system processes, siphoning sensitive data from AI workloads, hijacking GPUs for profit, and even battling rival cryptojackers for dominance. With no official patch and hundreds of thousands of Ray servers still exposed, ShadowRay 2.0 underscores the urgent need for secure configurations, strong access controls, and better visibility across fast-moving AI infrastructure.

Attack Regions

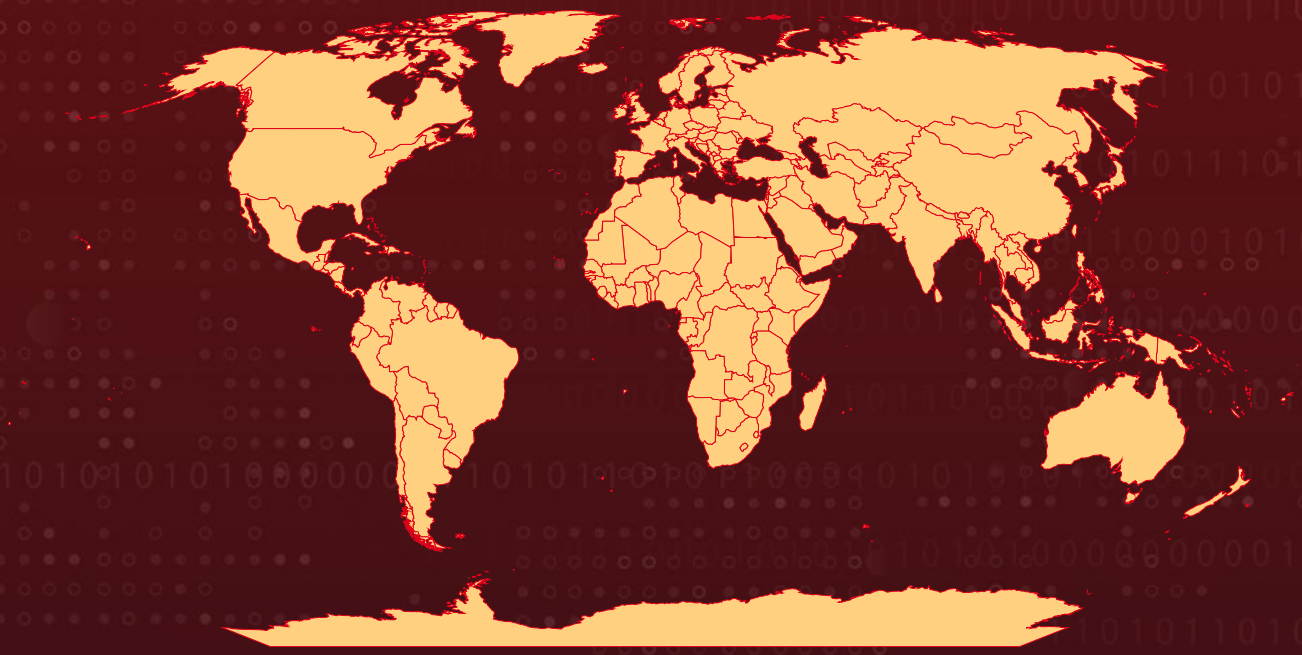

CVE

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2023-48022	Anyscale Ray Remote Code Execution Vulnerability	Anyscale Ray	✗	✗	✗

Attack Details

#1 In November 2025, a renewed and far more sophisticated attack wave exploited the vulnerability (CVE-2023-48022) in the Ray AI framework. This campaign, attributed to an actor known as IronErn440. The attackers relied on DevOps-style infrastructure, using GitLab to host and evolve their malware, and abused Ray's built-in orchestration features to create a self-propagating cryptojacking operation. Their payloads, many of which appeared to be AI-generated, were engineered to operate quietly by throttling CPU usage and disguising harmful processes. These activities may have started as early as September 2024, fueled by more than 230,000 misconfigured Ray servers exposed online.

#2 ShadowRay originally gained attention in March 2024 when it was linked to large-scale attacks targeting AI workloads. The vulnerability, CVE-2023-48022, enabled unauthenticated remote code execution through Ray's Jobs API, allowing adversaries to compromise thousands of servers across industries. These systems were abused for cryptomining, theft of sensitive data, and extraction of information from ongoing AI tasks. Although some related issues were addressed over time, the core flaw was never directly patched because it stemmed from Ray's design assumptions that clusters would operate within secured and isolated environments.

#3 By late 2025, a dramatic resurgence in malicious activity targeting Ray clusters, marking the rise of ShadowRay 2.0. Over 200,000 Ray servers were found exposed to the internet, with many already compromised. The lack of a definitive vendor patch and the heavy dependence on users to secure their own deployments created ideal conditions for attackers. The new campaign mirrored earlier attack patterns, remote code execution, payload delivery, stealthy persistence, cryptomining, and data theft, while introducing fresh techniques and unique indicators, pointing to different threat actors evolving the exploitation chain.

#4 This second-generation campaign unfolded in two rapid waves. The first wave emerged in early November when attackers used GitLab to host evolving payloads, until the account was taken down on November 5. Just days later, the operation resurfaced through a new GitHub repository launched on November 10, which continued to push updates at high velocity. Even after the repository was removed on November 17, yet another one appeared the same day, highlighting the actors' persistence and agility. Across both waves, the attackers demonstrated an advanced operational approach, AI-generated reconnaissance scripts, multi-stage Python payloads, hidden persistence mechanisms, GPU-optimized cryptojacking, large-scale victim scanning, and even DDoS tooling, effectively transforming compromised Ray clusters into multi-purpose botnets capable of mining, spreading autonomously, stealing sensitive data, and launching offensive operations across the internet.

Recommendations

- **Lock Down Your Ray Cluster Ports:** Make sure your Ray dashboard and Jobs API are not exposed to the public internet. Place them behind a firewall, VPN, or private network segment. If attackers can't reach your cluster, they can't exploit it.
- **Turn On Authentication Everywhere:** Ray is often deployed without authentication by mistake. Enable access controls, use strong API keys, and integrate identity controls wherever possible. Treat Ray like any other production-critical system.
- **Restrict Who Can Submit Jobs:** Limit job submission permissions to a small group of trusted users or service accounts. This stops attackers from abusing Ray's job features to run malicious code.
- **Monitor for Sudden CPU or GPU Spikes:** Unusual compute usage, especially near 100%, is often the first real-world clue of cryptomining. Set up alerts for unexpected CPU/GPU activity, new processes, or unknown Python scripts running at night.
- **Keep Your Infrastructure Updated:** Regularly update your Ray version, your OS, and cloud images. Even though CVE-2023-48022 isn't directly patched, newer releases have improvements around security guidance and deployment hardening.

Potential MITRE ATT&CK TTPs

TA0042 Resource Development	**TA0001** Initial Access	**TA0002** Execution	**TA0003** Persistence
TA0005 Defense Evasion	**TA0007** Discovery	**TA0010** Exfiltration	**TA0011** Command and Control
TA0040 Impact	**T1190** Exploit Public-Facing Application	**T1588** Obtain Capabilities	**T1588.006** Vulnerabilities

T1059 Command and Scripting Interpreter	**T1059.006** Python	**T1106** Native API	**T1053** Scheduled Task/Job
T1053.003 Cron	**T1543** Create or Modify System Process	**T1543.002** Systemd Service	**T1087** Account Discovery
T1082 System Information Discovery	**T1041** Exfiltration Over C2 Channel	**T1071** Application Layer Protocol	**T1071.001** Web Protocols
T1105 Ingress Tool Transfer	**T1036** Masquerading	**T1027** Obfuscated Files or Information	**T1562** Impair Defenses
T1562.004 Disable or Modify System Firewall	**T1498** Network Denial of Service	**T1496** Resource Hijacking	**T1588.005** Exploits

Indicators of Compromise (IOCs)

TYPE	VALUE
IPv4	18[.]228[.]3[.]224, 45[.]95[.]168[.]100, 185[.]215[.]180[.]70, 104[.]194[.]151[.]181, 121[.]160[.]102[.]68, 54[.]154[.]170[.]233, 158[.]160[.]123[.]117, 193[.]29[.]224[.]83, 162[.]248[.]53[.]119, 103[.]127[.]134[.]124, 18[.]230[.]118[.]147, 67[.]217[.]57[.]240, 45[.]61[.]150[.]83
Domain	*.oast.fun, pool.supportxmr.com, gulf.moneroocean.stream

TYPE	VALUE
Sub Domain	bwqqvqfgsseplyoltois92rdukv0mm5th.oast.fun
Monero Wallet Address	45MinZ6ECgTgxn8gbm5gAsK9ATrEN6N95hbH3g4r5N4bKwH8QxuFygw3G7VwHwAusR9L35E4YjWYdTJaWDjbMGDCKYNz5X1
ZANO Wallet Address	KrQtbtsrPTqSTzQwZZisiyJxgtcDMwrdVrQ
Mining Pool Address	eu.zano.k1pool.com
SSH Public Key	ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIHy6WMgqslpdUCaumLmlUcBjBjuAk4KspADxbcAKrzYd root@archtop
GitLab Repository	gitlab.com/ironern440-group/ironern440-project
GitLab User	ironern440-group, least3654, thisisforwork440-ops
URLs	hxxps[:]//gitlab[.]com/ironern440-group/ironern440-project/-/raw/main/mon[.]sh, hxxps[:]//gitlab[.]com/ironern440-group/ironern440-project/-/raw/main/aa[.]sh, hxxps[:]//gitlab[.]com/ironern440-group/ironern440-project/-/raw/main/run[.]sh, hxxps[:]//gitlab[.]com/ironern440-group/ironern440-project/-/raw/main/run-CN[.]sh, hxxps[:]//github[.]com/xmrig/xmrig/releases/download/v6[.]16[.]4/xmrig-6[.]16[.]4-linux-static-x64[.]tar[.]gz, hxxps[:]//github[.]com/rigelminer/rigel/releases/download/1[.]22[.]3/rigel-1[.]22[.]3-linux[.]tar[.]gz, hxxp[:]//45[.]61[.]150[.]83/1mmy/xd[.]sh, hxxp[:]//45[.]61[.]150[.]83/1mmy/cloud, hxxp[:]//67[.]217[.]57[.]240[:]666/files/netsh

TYPE	VALUE
SHA256	6f445252494a0908ab51d526e09134cebc33a199384771acd58c4a87f1ffc063, 1f6c69403678646a60925dcffe8509d22bb570c611324b93bec9aea72024ef6b
MD5	1f63fa7921c2f5fb8f8ffa430d02ac4a
SHA1	779a8af3b9838a33d1e199da3fc2f02a49e7c13e
Filename	dns-filter, .python3.6, rigel, python3.7.3, netsh, sockstress, mon.sh, aa.sh, aa_clean.sh, run.sh, run-CN.sh, .ddns.sh, xd.sh, cloud.txt
File Path	/usr/lib/dev/systemdev/dns-filter, /tmp/dns, /var/tmp/.ddns.sh, /etc/init.d/dns-filter, ~/.bashrc

Patch Details

Formal Patch not yet released for the flaw CVE-2023-48022.

References

https://www.oligo.security/blog/shadowray-2-0-attackers-turn-ai-against-itself-in-global-campaign-that-hijacks-ai-into-self-propagating-botnet

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a free demo with **HivePro Uni5**: Threat Exposure Management Platform.

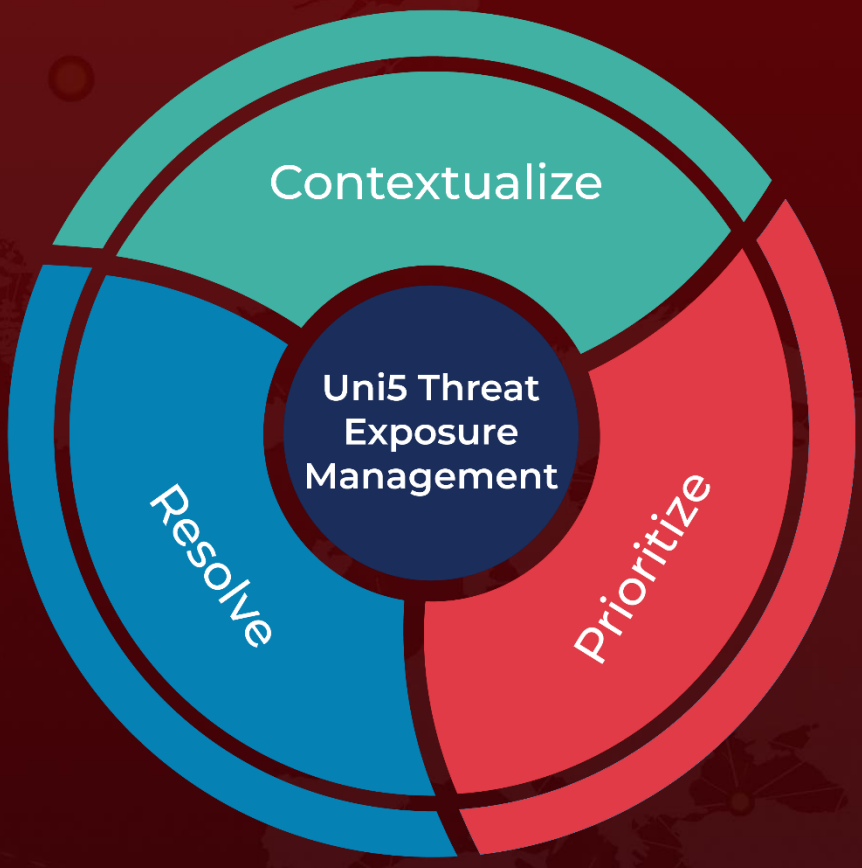

REPORT GENERATED ON

November 25, 2025 • 6:30 AM

More at www.hivepro.com