


BPFDoor - An Evasive Linux Backdoor Technical Analysis

 sandflysecurity.com/blog/bpfdoor-an-evasive-linux-backdoor-technical-analysis/

May 11, 2022



SANDFLY DETAILS VIEW SANDFLY

Failed: process_deleted

2022-05-10T23:54:06Z (an hour ago)

Defense Evasion Execution T1070.004

The process name 'kdmtmpflush' with PID '1343' binary has been deleted from the disk but it is still running at '/dev/shm/kdmtmpflush'. This could be a malicious program still resident in memory and trying to hide from the file system.

HOST SUMMARY VIEW HOST

HOSTNAME
bpfdoor

TARGET ADDRESS
[REDACTED]

IP (LAST SEEN)
[REDACTED]

PLATFORM
Linux bpfdoor 5.4.0-109-generic
#123-Ubuntu SMP Fri Apr 8
09:10:54 UTC 2022 x86_64

SANDFLY HUNTER DATA POINTS

process.name
kdmtmpflush SEARCH

process.cmdline
dbus-daemon --system SEARCH

Malware Linux Forensics

Date

May 11, 2022

Author

The Sandfly Security Team

Recently Kevin Beaumont revealed a new evasive backdoor targeting Linux associated with the Chinese Red Menshen threat actors. In his [article](#) he reveals that this backdoor has been operating globally for many years with potentially thousands of instances already deployed.

The backdoor has also been noted by investigators at PricewaterhouseCoopers in their latest [Cyber Threat Intelligence Retrospect Report](#) (pg. 36).

The source for this backdoor was posted by anonymously and Sandfly researchers are able to provide the following in-depth technical analysis. At a high level, it does the following:

- Goes memory resident and deploys anti-forensics and evasion to hide.
- Loads a Berkeley Packet Filter (BPF) sniffer allowing it to efficiently watch traffic and work in front of any locally running firewalls to see packets (hence BPFDoor).
- Upon receiving a special packet, it will modify the local firewall to allow the attacker IP address to access resources such as a spawned shell or connect back bindshell.
- Operations are hidden with process masquerading to avoid detection.

While the malware takes steps to evade casual observation, it is easily seen if you know where and how to look. We'll review the above and provide detection tips.

Source Build Size and Compatibility

The BPFDoor source is small, focused and well written. While the sample we reviewed was Linux specific, with some small changes it could easily be ported to other platforms (a Solaris binary reportedly exists). BPF is widely available across operating systems and the core shell functions would likely work across platforms with little modification.

The dynamically linked binary is small at about 35K on Ubuntu:

```
-rwxr-xr-x 1 root root 34952 May 11 00:03 bpfdoor
```

Statically linked it would grow to about 1MB, but the dynamically linked version would likely work on most modern Linux distributions. Cross compiling for various CPUs is also possible so this implant would likely work on embedded Linux devices as well.

Implant Operation Steps

The binary itself just needs to be downloaded onto the victim and run. It doesn't matter how or where it gets to the host as it takes care of moving itself to a suitable area once run to remain resident. However, the binary does need *root* permissions to run.

When run, the binary has an initialization sequence as follows:

- 1) Copy binary to the */dev/shm* directory (Linux ramdisk).
- 2) Rename and run itself as */dev/shm/kdmtmpflush*.

- 3) Fork itself and run fork with "--init" flag which tells itself to execute secondary clean-up operations and go resident.
- 4) The forked version timestomps the binary file `/dev/shm/kdmtmpflush` and initiates packet capture loop.
- 5) The forked version creates a dropper at `/var/run/haldrund.pid` to mark it is resident and prevent multiple starts.
- 6) The original execution process deletes the timestomped `/dev/shm/kdmtmpflush` and exits.
- 7) The forked version remains resident and monitors traffic for magic packet to initiate attacker operations such as a bindshell.

Persistence

The implant itself has no persistence mechanisms as it is highly focused on a single task. Persistence would need to be initiated by the attacker in some other way such as `rc` or `init` scripts or scheduled tasks such as with `crontab`. The initial report referenced above indicates that persistence scripts have been found.

The implant uses `/dev/shm` on Linux. This is a ramdisk and is cleared out on every reboot. For persistence reasons, the implant will need to be somewhere else on the host to survive reboots or be inserted again remotely.

Incident response teams that find this implant operating should **assume the real binary is somewhere else on the file system**. Check all system boot scripts for unusual references to binaries or paths.

Timestomping

The binary copies itself to `/dev/shm/kdmtmpflush` which is only in RAM and clears out every reboot. The interesting part of the implant is that it sets a bogus time to timestomp the binary before deletion. The relevant code is below:

```
tv[0].tv_sec = 1225394236;
tv[0].tv_usec = 0;

tv[1].tv_sec = 1225394236;
tv[1].tv_usec = 0;


utimes(file, tv);
```

The date is set to 1225394236 seconds from epoch which translates to: **Thursday, October 30, 2008 7:17:16 PM (GMT)**

We did some searches to see if this date has any significance but didn't see anything obvious. It could have some significance to the author or could be random.

The interesting part about this is the timestamp happens by the forked process before the main process tries to delete the binary. We assume this is a failsafe mechanism. If the implant should fail to load and not delete itself from `/dev/shm/kdntmpflush` then the file left behind will have an innocuous looking time on it that masks when it was created. It would also make incident response harder if you are looking for recently created files (this one looks like it was created 14 years ago).

```
root@bpfdoor:~# ls -al /dev/shm
total 36
drwxrwxrwt 2 root root 60 May 10 09:48 .
drwxr-xr-x 17 root root 3800 May 10 09:10 ..
-rwxr-xr-x 1 root root 36616 Oct 30 2008 kdntmpflush
root@bpfdoor:~#
```



PID Dropper

The implant creates a zero byte PID file at `/var/run/haldrund.pid`. This file is deleted if the implant terminates normally, but if there is a problem such as a hard shutdown or crash, the file may be left behind. The implant will not start if this file is present as it is used to mark that it may already be running.

Binary Deletion

After the binary starts, it deletes itself making recovery harder. However, recovering a deleted process binary on Linux is trivial once it is running ([see our article on how to do it](#)). But the main thing deletion does is allow the binary to avoid detection by malware scanners that rely on file scanning. The binary is simply not on the disk to be scanned and if the main binary is hidden/encrypted on the device for persistence it would be very hard to find.

However on Linux a deleted process binary is extremely suspicious. If you search for any kind of process with a deleted binary then it stands out:

```
ls -alR /proc/*/exe 2> /dev/null | grep deleted
```

```
root@bpfdoor:~# ls -alR /proc/*/exe 2> /dev/null | grep deleted
lrwxrwxrwx 1 root      root          0 May 11 00:39 /proc/1343/exe -> /dev/shm/kdmtmpflush (deleted)
root@bpfdoor:~#
```

If you are using Sandfly to protect your Linux systems you will have received multiple automated alerts about a suspicious binary running. The red arrow shows that the process is masquerading as something else. We'll talk about that in a bit.

SANDFLY DETAILS VIEW SANDFLY

Failed: process_deleted

2022-05-10T23:54:06Z (an hour ago)

Defense Evasion
Execution
T1070.004

The process name 'kdmtmpflush' with PID '1343' binary has been deleted from the disk but it is still running at '/dev/shm/kdmtmpflush'. This could be a malicious program still resident in memory and trying to hide from the file system.

HOST SUMMARY VIEW HOST

HOSTNAME
bpfdoor

TARGET ADDRESS
[REDACTED]

IP (LAST SEEN)
[REDACTED]

PLATFORM
Linux bpfdoor 5.4.0-109-generic
#123-Ubuntu SMP Fri Apr 8
09:10:54 UTC 2022 x86_64

SANDFLY HUNTER DATA POINTS

process.name kdmtmpflush	SEARCH
process.cmdline dbus-daemon --system	SEARCH

In general, any process deleting its binary after running is going to be malicious.

Masquerading

The binary masquerades its name by selecting from one of 10 names randomly:

```

/sbin/udev -d
/sbin/mingetty /dev/tty7
/usr/sbin/console-kit-daemon --no-daemon
hald-addon-acpi: listening on acpi kernel interface /proc/acpi/event
dbus-daemon --system
hald-runner
pickup -l -t fifo -u
avahi-daemon: chroot helper
/sbin/auditd -n
/usr/lib/systemd/systemd-journald

```

The names are made to look like common Linux system daemons. The implant overwrites the argv[0] value which is used by the Linux /proc filesystem to determine the command line and command name to show for each process. By doing this, when you run commands like

ps you will see the bogus name. Below you can see the process running under the masquerade name *dbus-daemon --system*.

```
root      454  0.0  0.0      0   0  7      S<   May10  0:00 [loop1]
root      455  0.0  0.0      0   0  7      S<   May10  0:00 [loop2]
root      458  0.0  0.0      0   0  7      S<   May10  0:00 [loop3]
root      460  0.0  0.0      0   0  7      S<   May10  0:00 [loop4]
systemd+  474  0.0  0.6  90188 6096  7      Ssl  May10  0:00 /lib/systemd/systemd-timesyncd
systemd+  533  0.0  0.6  10370 6952  7      Ss   May10  0:00 /lib/systemd/systemd-networkd
systemd+  552  0.0  0.9  23884 7840  7      Ss   May10  0:00 /lib/systemd/systemd-resolved
root      577  0.0  0.5  19404 5084  7      Ss   May10  0:00 /lib/systemd/systemd-udev
root      642  0.0  0.8  241028 8892  7      Ssl  May10  0:00 /usr/lib/accounts-service/accounts-daemon
root      645  0.0  0.2  8536 2896  7      Ss   May10  0:00 /usr/sbin/cron -f
message+  647  0.0  0.4  7640 4416  7      Ss   May10  0:00 /usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
do-agent  649  0.0  1.4  93768 14400  7      Ssl  May10  0:00 /opt/digitalocean/bin/do-agent --syslog
root      655  0.0  0.4  1299420 4544  7      Ssl  May10  0:00 /opt/digitalocean/bin/droplet-agent
root      656  0.0  1.4  29864 14400  7      Ss   May10  0:00 /usr/bin/python3 /usr/bin/networkd-dspatcher --run-startup-triggers
root      657  0.0  0.0  236410 8804  7      Ssl  May10  0:00 /usr/lib/polkit/polkitd --no-debug
syslog    658  0.0  0.4  224344 4568  7      Ssl  May10  0:00 /usr/sbin/rsyslogd -n -tNONE
root      660  0.0  3.6  726544 36108  7      Ssl  May10  0:01 /usr/lib/snmpd/snmpd
root      662  0.0  0.7  16664 7612  7      Ss   May10  0:00 /lib/systemd/systemd-logind
root      664  0.0  1.3  394784 13340  7      Ssl  May10  0:00 /usr/lib/udisks2/udisksd
daemon    665  0.0  0.2  3792 2248  7      Ss   May10  0:00 /usr/sbin/atd -f
root      670  0.0  0.0      0   0  7      I    May10  0:00 [kworker/0:0-mm_percpu_wq]
root      677  0.0  0.2  7352 2108  tty50  Ss+  May10  0:00 /sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 tty50 vt220
root      680  0.0  0.1  3020 3000  tty1   Ss+  May10  0:00 /sbin/agetty -o -p -- \u --noclear tty1 linux
root      724  0.0  0.6  12172 6560  7      Ss   May10  0:00 sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startups
root      727  0.0  1.0  814448 10012  7      Ssl  May10  0:00 /usr/sbin/ModemManager
root      735  0.0  1.6  108184 16096  7      Ssl  May10  0:00 /usr/bin/python3 /usr/share/unattended-upgrades/unattended-upgrade-shutdown --wait-for-signal
root      919  0.0  0.9  18392 9552  7      Ss   May10  0:00 /lib/systemd/systemd --user
root      920  0.0  0.3  103340 3388  7      S    May10  0:00 (sd-pam)
root      1079  0.0  0.7  13804 7444  7      Ss   May10  0:00 sshd: root@pts/1
root      1150  0.0  0.4  10040 4668  pts/1  Ss+  May10  0:00 -bash
root      1165  0.0  0.7  13808 7404  7      Ss   May10  0:00 sshd: root@pts/2
root      1242  0.0  0.4  10040 4120  pts/2  Ss+  May10  0:00 -bash
root      1252  0.0  0.7  13804 7408  7      Ss   May10  0:00 sshd: root@pts/3
root      1318  0.0  0.4  10040 4200  pts/3  Ss+  May10  0:00 -bash
root      1343  0.0  0.0  2588 104  7      Ss   May10  0:00 dbus-daemon --system
root      1446  0.0  0.0      0   0  7      I    May10  0:00 [kworker/u2:2-events_power_efficient]
root      1468  0.0  0.4  7208 4216  7      Ss   May10  0:00 /usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
root      1494  0.0  0.7  13800 7876  7      Ss   May10  0:00 sshd: root@pts/4
root      1570  0.0  0.5  10040 5396  pts/4  Ss+  May10  0:00 -bash
root      1740  0.0  0.0      0   0  7      I    00:00 0:00 [kworker/B:1-cgroup_destroy]
root      1832  0.0  0.0      0   0  7      I    00:14 0:00 [kworker/u2:0-events_unbound]
root      1835  0.0  0.9  14240 8632  7      Ss   00:19 0:00 sshd: root@pts/5
root      1934  0.0  0.5  10040 5320  pts/5  Ss   00:39 0:00 -bash
root      1984  0.0  0.3  10768 3840  pts/5  R+   00:55 0:00 ps -auxw
```

This masquerading tactic has been around for a while. While it works, the real process name is still visible inside Sandfly with the masqueraded versions also showing. This kind of difference in real process name vs. the command line values indicates also there is a problem.

SANDFLY HUNTER DATA POINTS	
process.name kdmtmpflush	SEARCH
process.cmdline dbus-daemon --system	SEARCH
process.command dbus-daemon --s	SEARCH
process.path /dev/shm/kdmtmpflush	SEARCH
process.uid_name root	SEARCH

If you find a suspicious process ID (PID), you can quickly investigate what the real name may be by going to `/proc/<PID>` and doing a simple `ls` command:

```
cd /proc/<PID>
ls -al
```

You will see the `exe` link which will be pointing to the real binary location which can confirm at least what the binary was called when started. Also you'll see the timestamp on the file is when the binary was started which can help bracket the time of intrusion. Linux also helpfully labels the binary as "(deleted)".


```
root@bpfdoor:/# cd /proc/1343
root@bpfdoor:/proc/1343# ls -al
total 0
dr-xr-xr-x  9 root root 0 May 11 00:39 .
dr-xr-xr-x 163 root root 0 May 10 23:35 ..
-r--r--r--  1 root root 0 May 11 01:01 arch_status
dr-xr-xr-x  2 root root 0 May 11 01:01 attr
-rw-r--r--  1 root root 0 May 11 01:01 autogroup
-r-----  1 root root 0 May 11 01:01 auxv
-r--r--r--  1 root root 0 May 11 01:01 cgroup
--W-----  1 root root 0 May 11 01:01 clear_refs
-r--r--r--  1 root root 0 May 11 00:39 cmdline
-rw-r--r--  1 root root 0 May 11 01:01 comm
-rw-r--r--  1 root root 0 May 11 01:01 coredump_filter
-r--r--r--  1 root root 0 May 11 01:01 cpuset
lrwxrwxrwx  1 root root 0 May 11 00:39 cwd -> /root
-r-----  1 root root 0 May 11 01:01 environ
lrwxrwxrwx  1 root root 0 May 11 00:39 exe -> '/dev/shm/kdntmpflush (deleted)!'
dr-x-----  2 root root 0 May 11 00:39 fd
dr-x-----  2 root root 0 May 11 00:39 fdinfo
-rw-r--r--  1 root root 0 May 11 01:01 gid_map
-r-----  1 root root 0 May 11 01:01 io
-r--r--r--  1 root root 0 May 11 01:01 limits
-rw-r--r--  1 root root 0 May 11 01:01 loginuid
dr-x-----  2 root root 0 May 11 01:01 map_files
-r--r--r--  1 root root 0 May 11 00:39 maps
-rw-----  1 root root 0 May 11 01:01 mem
-r--r--r--  1 root root 0 May 11 01:01 mountinfo
-r--r--r--  1 root root 0 May 11 01:01 mounts
-r-----  1 root root 0 May 11 01:01 mountstats
dr-xr-xr-x 54 root root 0 May 11 01:01 net
dr-x--x--x  2 root root 0 May 11 01:01 ns
-r--r--r--  1 root root 0 May 11 01:01 numa_maps
-rw-r--r--  1 root root 0 May 11 01:01 oom_adj
-r--r--r--  1 root root 0 May 11 01:01 oom_score
```

Environment Wipe

The last thing the implant does before going fully resident is wipe out the process environment. When you start a process on Linux it stores a lot of useful forensic information in `/proc/<PID>/environ`. This area can often reveal useful information such as SSH connections that started the process, usernames, etc.

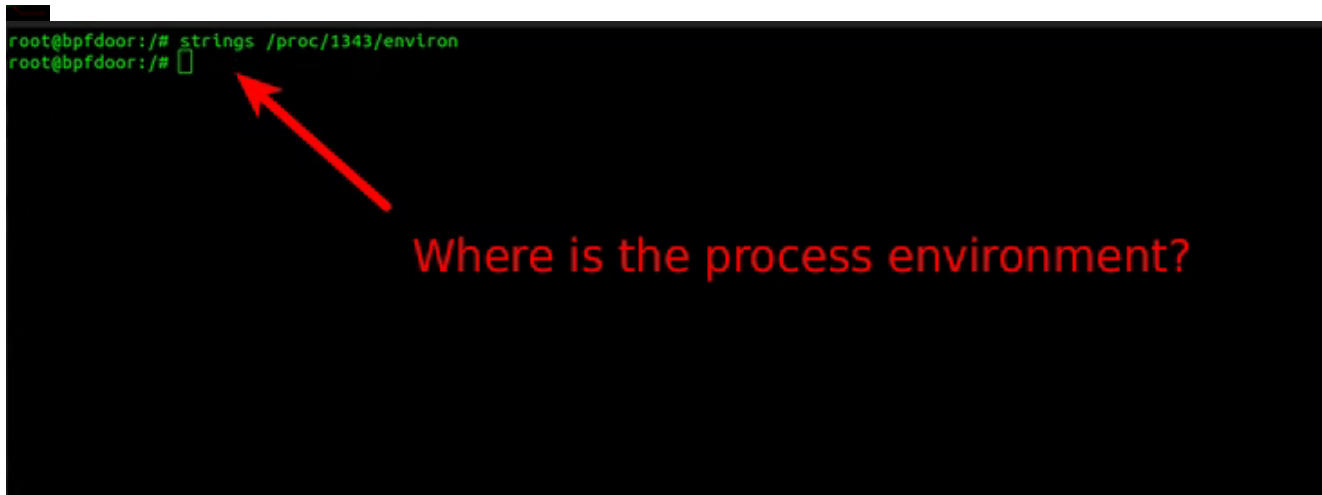
The environment wipe the implant uses is interesting because it wipes out `envp[]` (third argument to `main()` which is where environment variables are passed in as array in Linux). See below for explanation of how to use `argv` iteration to get environment variables:

[argv prints out environment variables](#)

This is a pretty good mechanism and ensures there are no environment variables left on the running process. Although it won't work in this case, we have an article on how to do live process environment forensics here:

[Using Linux Process Environment Variables for Live Forensics](#)

The end result is that the implant leaves the environment completely blank which can happen under some normal circumstances, but usually not. Below we see the fake `dbus` implant and the real `dbus` on the same box. The real `dbus` environment has some data with it. A completely empty environment is unusual for normal processes.



```
root@bpfdoor:/proc/1343# strings /proc/1468/envIRON
HOME=/root
LANG=C.UTF-8
LOGNAME=root
PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/games:/usr/local/games:/snap/bin:/snap/bin
SHELL=/bin/sh
USER=root
XDG_RUNTIME_DIR=/run/user/0
XDG_DATA_DIRS=/usr/local/share:/usr/share:/var/lib/containers/desktop
DBUS_SESSION_BUS_ADDRESS=unix:path=/run/user/0/bus
MANAGERPID=919
LISTEN_PID=1468
LISTEN_FDS=1
LISTEN_FDNAMES=dbus.socket
INVOCATION_ID=2fd0b1fe61c049f297222c29e4fa8942
JOURNAL_STREAM=9:30115
root@bpfdoor:/proc/1343#
```

Real dbus process environment.

BPF Filter Activation and Analysis

Once the implant has done its anti-forensics and evasion housekeeping, it goes into the packet capture loop. The packet capture function loads a BPF filter (hence the name). BPF is a highly efficient way to filter packets coming into a system which massively reduces CPU load by preventing all packets from needing to be analyzed by the receiver. It effectively operates as a very efficient pre-filter and only passes likely valid packets for review to the implant.

With BPFDoor they have taken a BPF filter and converted it to BPF bytecode. This keeps the implant small and less reliant on system libraries that may need to parse a human readable BPF filter, and allows for filters that the normal expression syntax cannot represent.

We have reverse engineered the bytecode below to show you what it is doing. It does basically two things:

- Grabs either an ICMP (ping), UDP or TCP packet.
- Sees if it has magic data value. If not then reject.

The commented assembly and pseudocode is here:

```

l0:    ldh [12]                // A = halfword from offset 12 [Ethernet: EtherType]
l1:    jeq #0x800, l2, l29    // if EtherType==IPv4 goto l2; else goto l29
l2:    ldb [23]              // A = byte from packet offset 23 [IPv4: Protocol]
(data begins at offset 14 of ethernet packet; so this is offset 9 in the IPv4 packet)
l3:    jeq #0x11, l4, l9     // if Protocol==UDP goto l4, else goto l9
l4:    ldh [20]              // A = IPv4 flags+fragment offset
l5:    jset #0x1fff, l29, l6 // ...if fragmentation offset != 0, goto l29
l6:    ldx 4*([14]&0xf)      // X = IPv4 Header Length
l7:    ldh [x+22]           // A = halfword from offset X+22... first halfword of
UDP datagram data
l8:    jeq #0x7255, l28, l29 // if A==0x7255 goto l28, else goto l29
l9:    jeq #0x1, l10, l17    // if Protocol==ICMP goto l10, else goto l17 (jumped
to from l3; register contains IP protocol)
l10:   ldh [20]              // A = IPv4 flags+fragment offset
l11:   jset #0x1fff, l29, l12 // ...if fragmentation offset != 0, goto l29
l12:   ldx 4*([14]&0xf)      // X = IPv4 Header Length
l13:   ldh [x+22]           // A = halfword from offset X+22... first halfword of
ICMP data
l14:   jeq #0x7255, l15, l29 // if A==0x7255 goto l15, else goto l29
l15:   ldb [x+14]           // A = byte from offset X+14... ICMP Type
l16:   jeq #0x8, l28, l29   // if ICMP Type == Echo Request (ping) goto l28, else
goto l29
l17:   jeq #0x6, l18, l29   // if Protocol==TCP goto l18, else goto l29 (jumped
to from l3; register contains IP protocol)
l18:   ldh [20]              // A = IPv4 flags+fragment offset
l19:   jset #0x1fff, l29, l20 // ...if fragmentation offset != 0, goto l29
l20:   ldx 4*([14]&0xf)      // X = IPv4 Header Length
l21:   ldb [x+26]           // A = byte from offset X+26... Assume no IPv4
options so X=20; packet offset 46 = TCP segment offset 12
l22:   and #0xf0            // A = A & 0xF0 (high nibble of TCP offset 12 = Data
offset = TCP header size in 32-bit words)
l23:   rsh #2               // A = A >> 2 (this has the effect of multiplying the
high nibble by four, e.g A>>4 then A<<2). A now contains number of bytes in the TCP
header
l24:   add x                 // A = A + X. Adding IPv4 header length + TCP header
length.
l25:   tax                  // X = A
l26:   ldh [x+14]           // A = halfword from packet offset X+14 (14 is
ethernet header, x is IPv4+TCP header, so this offset is first TCP payload byte)
l27:   jeq #0x5293, l28, l29 // if A==0x5293 goto l28, else goto l29
l28:   ret #0xffff          // return match
l29:   ret #0               // return doesn't-match

```

Pseudocode. "return false" means packet doesn't match; "return true" means packet matches.

```

if (EtherType == IPv4) {
    if (Packet is an additional piece of a fragmented packet)
    {
        return false;
    }
    else
    {
        if (Protocol == UDP && data[0:2] == 0x7255)
        {

```

```

        return true;
    }
    else if (Protocol == ICMP && data[0:2] == 0x7255 && ICMP Type == Echo
Request)
    {
        return true;
    }
    else if (Protocol == TCP && data[0:2] == 0x5293)
    {
        return true;
    }
    else {
        return false;
    }
}
}
else
{
    return false;
}
}

```

To get past the filter you will need to send the right data in the packet as shown above.

The filter rejects most traffic from entering the main packet decoding loop so the implant will run with very little CPU signature. Packets that make it through the BPF check are quickly checked for a valid password before any further operations take place.

Packet Capture and Firewall Bypass

The relevance of the BPF filter and packet capture is that it is sniffing traffic at a lower level than the local firewall. **This means that even if you run a firewall the implant will see, and act upon, any magic packet sent to the system.** The firewall running on the local host will not block the implant from having this visibility. This is an important point to understand.

Further, if you have a network perimeter firewall, but it allows traffic to a host on ICMP, UDP or TCP to **any** port, then the implant can see it. Any inbound traffic to the sever can activate the magic packet check and initiate a backdoor.


For instance, if you run a webserver and lock it down so only port TCP 443 can be accessed, the attacker can send a magic packet to TCP port 443 and activate a backdoor. Same for any UDP packet or even a simple ICMP ping. We'll talk about how it does this below.

Locating Packet Sniffing Processes

Finding a process sniffing packets on Linux by hand is not always obvious. It's just not normal for most people to check for such things and as a result something like BPFDoor can remain around for a long time unnoticed.

However, with this malware in a wait state loop searching for packets you can look for traces left under the process stack by viewing `/proc/<PID>/stack`. With BPFDoor we can see references to function calls in the Linux kernel that indicate the process is likely grabbing packets.

```
"stack": [  
  "__skb_wait_for_more_packets+0x103/0x170",  
  "__skb_recv_datagram+0x67/0xc0",  
  "skb_recv_datagram+0x3b/0x60",  
  "packet_recvmsg+0x75/0x4b0",  
  "sock_recvmsg+0x70/0x80",  
  "__sys_recvfrom+0x19e/0x1d0",  
  "__x64_sys_recvfrom+0x29/0x30",  
  "do_syscall_64+0x57/0x190",  
  "entry_SYSCALL_64_after_hwframe+0x44/0xa9"  
]
```




You can search the entire `/proc/*/stack` area for any process that is showing packet capture functions like the above:

```
grep packet_recvmsg /proc/*/stack  
grep wait_for_more_packets /proc/*/stack
```

False alarms with this search are possible, but you can narrow down possible candidates like below. The red arrow is the PID in question running BPFDoor.

```
root@bpfdoor:/# grep packet_recvmsg /proc/*/stack
/proc/6412/stack:[<0>] packet_recvmsg+0x75/0x4b0
root@bpfdoor:/#
```



The above is time consuming though and not practical in many cases. Instead, we recommend an automated sweep from Sandfly to quickly identify all processes with packet sockets in operation. With this, BPFDoor is immediately found. The packet capture socket in operation shows up easily and there is no mistaking that this process is reading network traffic.

SANDFLY DETAILS

[VIEW SANDFLY](#)

Failed: process_running_sniffer_operating_tuned

2022-05-09T01:01:15Z (2 minutes ago)

Credential Access

Discovery

T1040

The process name 'kdmtmpflush' with PID '7383' shows evidence it may be operating as a network sniffer. Network sniffing can be a valid activity for some system processes (e.g. systemd), but you should investigate this process to be sure it is legitimate if you don't know why it is being run on this system and grabbing network traffic.

HOST SUMMARY

[VIEW HOST](#)

HOSTNAME

bpfdoor

TARGET ADDRESS

IP (LAST SEEN)

PLATFORM

Linux bpfdoor 5.4.0-107-generic
#121-Ubuntu SMP Thu Mar 24
16:04:27 UTC 2022 x86_64

SANDFLY HUNTER DATA POINTS

process.name

kdmtmpflush

[SEARCH](#)

process.cmdline

avahi-daemon: chroot helper

[SEARCH](#)

process.command

[SEARCH](#)


```
"file_descriptors": [  
  {  
    "path": "/dev/pts/0",  
    "type": "block",  
    "class": "device",  
    "number": 0  
  },  
  {  
    "path": "/dev/pts/0",  
    "type": "block",  
    "class": "device",  
    "number": 1  
  },  
  {  
    "path": "/dev/pts/0",  
    "type": "block",  
    "class": "device",  
    "number": 2  
  },  
  {  
    "path": "socket:[65652]",  
    "type": "socket",  
    "class": "packet",  
    "number": 3  
  }  
],  
},
```

RC4 Encryption and Passwords

To access the implant you need not just a magic packet, but also the correct password. The leaked source has some passwords set, but there is no reason to believe these are used in actual deployment.

The implant uses RC4 as the encryption layer. RC4 is a very robust cipher for this application and is the only truly secure cipher you can write on a napkin. It's a great choice for small implant code like this.

In the case of the implant we will deduct a few points because they did not throw out the first few kilobytes from the cipher stream which can weaken it, but overall this cipher is a good choice to keep things small and fast.

The implant can take an optional password. The password is compared against two internal values. If it matches one value it will setup a local bindshell backdoor. If it matches another it will do a reverse bindshell connect-back to the specified host and port.

There is a third option though and that is if no password is specified, then a function is called that sends a single UDP packet with the value "1". This might be some kind of operation check status to show the implant is still running. Relevant code below:

```
if ((s_len = sendto(sock, "1", 1, 0, (struct sockaddr *)&remote, sizeof(struct
sockaddr))) < 0) {
    close(sock);
    return -1;
}
```

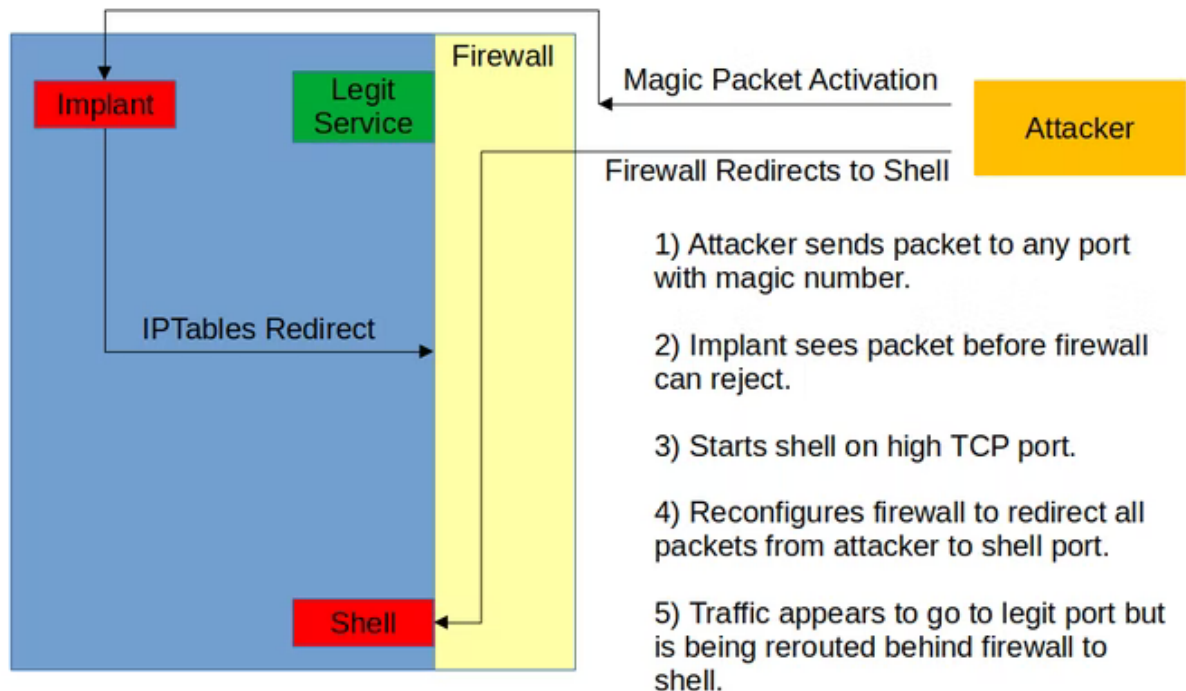
Note the above is not passed through the RC4 encryption. If you are investigating this on your network, it may be worthwhile looking for single UDP packets with just the data "1" in them and nothing else from many hosts over time to a single host IP (controller).

Firewall Bypass for Bindshell Backdoor

The implant has a unique feature to bypass the host firewall and make the traffic look legitimate. When the magic packet is received by the host, the implant will spawn a new instance and change the local *iptables* rules to do a redirect from the requesting host to the shell port.

For instance, the implant can redirect all traffic from the attacker using TCP port 443 (encrypted web) to the shell. Externally, the traffic will look like TLS/SSL traffic but in fact the attacker is interacting with a remote *root* shell on the system.

Let's review what this means with a diagram:



The steps are as follows if the actor requests the system open a local shell:

- 1) Implant is listening to all traffic coming to the host regardless of firewall.
- 2) Implant sees magic packet.
- 3) Magic packet can contain IP address, port and password for attacker. Otherwise it uses the origin IP address.
- 4) Depending on password, implant will open local or connect-back backdoor.
- 5) Implant selects a TCP port sequentially starting at 42391 up to 43391.
- 6) Implant binds to first unused port in this range.
- 7) For local shell, *iptables* is called to redirect all traffic from attacker host IP from the requested port to the bound port range from the above steps.
- 8) Attacker connects with TCP to the defined port they requested (e.g. TCP port 443).
- 9) Traffic from that host is redirected from the port to the shell routing behind the firewall.
- 10) Observed traffic still appears to be going to host on a legitimate port, but in fact is being routed to the shell on the host.

If this is confusing, then let's look at the shell in action using SSH as our target port.

NOTE: We disabled the RC4 encryption in the implant for example purposes to use *netcat*. The real implant would require the correct password and RC4 encryption layer to see these results.

Below we connect to a host on SSH port 22. We get back a normal OpenSSH banner. Then on another window we send the magic packet on UDP to the host (or TCP or ICMP). The implant sees this packet and that we want to use TCP port 22 as our shell access port. The implant starts a shell on a high TCP port and then redirects the traffic. When we connect again to port 22, instead of SSH we now get a shell with *root* access.

Here again is the important point: **All other users still get SSH. Only the attacker's traffic is redirected to the shell even though it goes to the same SSH port!**

The image shows a terminal window with the following content:

```
root@sandflysecurity:/ # nc 147.182.254.128 22
SSH-2.0-OpenSSH_8.2p1 Ubuntu-4ubuntu0.4
Invalid SSH identification string.
root@sandflysecurity:/ # nc 147.182.254.128 22
3438[ \u0127 \W]\$ id
uid=0(root) gid=0(root) groups=0(root)
[\u0127 \W]\$ whoami
root
[\u0127 \W]\$ ls -al
ls -al
total 72
drwxr-xr-x 19 root root 4096 May 10 23:35 .
drwxr-xr-x 19 root root 4096 May 10 23:35 ..
lrwxrwxrwx 1 root root 7 Apr 11 21:39 bin -> usr/bin
drwxr-xr-x 4 root root 4096 May 10 09:13 boot
drwxr-xr-x 17 root root 3800 May 10 23:35 dev
drwxr-xr-x 95 root root 4096 May 10 09:15 etc
drwxr-xr-x 2 root root 4096 Apr 15 2020 home
lrwxrwxrwx 1 root root 7 Apr 11 21:39 lib -> usr/lib
lrwxrwxrwx 1 root root 9 Apr 11 21:39 lib32 -> usr/lib32
lrwxrwxrwx 1 root root 9 Apr 11 21:39 lib64 -> usr/lib64
lrwxrwxrwx 1 root root 10 Apr 11 21:39 libx32 -> usr/libx32
drwx----- 2 root root 16384 Apr 11 21:42 lost+found
drwxr-xr-x 2 root root 4096 Apr 11 21:39 media
drwxr-xr-x 2 root root 4096 Apr 11 21:39 mnt
drwxr-xr-x 3 root root 4096 May 10 09:11 opt
dr-xr-xr-x 164 root root 0 May 10 23:35 proc
drwx----- 6 root root 4096 May 11 02:25 root
drwxr-xr-x 27 root root 880 May 11 02:25 run
lrwxrwxrwx 1 root root 8 Apr 11 21:39/sbin -> usr/sbin
drwxr-xr-x 6 root root 4096 Apr 11 21:42 snap
drwxr-xr-x 2 root root 4096 Apr 11 21:39 srv
dr-xr-xr-x 13 root root 0 May 10 23:35 sys
drwxrwxrwt 13 root root 4096 May 11 02:25 tmp
drwxr-xr-x 14 root root 4096 Apr 11 21:40 usr
drwxr-xr-x 13 root root 4096 Apr 11 21:41 var
[\u0127 \W]\$
```

Annotations on the image:

- Red arrow pointing to the SSH banner: "Before backdoor activation. SSH port answers normally for attacker host."
- Red arrow pointing to the root shell prompt: "After backdoor activation. SSH port for attacker host is redirected to shell running on high TCP port."
- Red arrow pointing to the file listing: "External traffic observation will think this is normal SSH TCP port 22 traffic."

You'll see also when you connect the value "3458" is sent above. This may be a version identifier for the implant operator.

The redirect rules for the shell access are seen below. Here you see that the TCP port 42392 was found available and the shell bound to it. All TCP port 22 traffic from our attacker host (192.168.1.1) is now routed to this shell on this port. The traffic looks like encrypted SSH communications going to TCP port 22, but in reality is being directed to the shell port once it hits the *iptables* rule for the attacker host *only*.

```
/sbin/iptables -t nat -D PREROUTING -p tcp -s 192.168.1.1 --dport 22 -j REDIRECT --to-ports 42392
/sbin/iptables -D INPUT -p tcp -s 192.168.1.1 -j ACCEPT
```

Connect-Back Bindshell

The implant also has the ability to connect back to a host as defined in the magic packet. Operation here is largely the same, but this may not be as stealthy having a system connect outbound (and many orgs may block servers talking outbound). The first method of packet redirect is far more dangerous and harder to find as it will look like legitimate traffic going to the server and not originating outbound.

Status Check

As discussed above, if you do not supply any password or an incorrect password in a magic packet the implant will simply send out "1" on UDP. We believe this is some kind of status check to allow keeping tabs on many systems.

Organizations may want to consider running a network scan against their hosts sending a magic packet on ICMP, UDP or TCP with a known UDP port you monitor to see if any systems respond. Any host responding has an active implant.

Shell Anti-Forensics and Evasion

The shell is spawned by forking a controller and finally the shell itself. The controlling PID will be masquerading running under the name:

```
/usr/libexec/postfix/master
```

The shell itself will be running under the name:

```
qmgr -l -t fifo -u
```

In the *ps* listing you'll see the following. Here the implant is restarted under a new bogus name. Then you see the two processes masking the shell operation when it is operating.

```

root      642  0.0  0.5 1299420 5632 ?        Ssl  20:51  0:00 /opt/digitalocean/bin/droplet-agent
root      648  0.0  1.8 29864 18340 ?        Ss   20:51  0:00 /usr/bin/python3 /usr/bin/networkd-dispatcher --run-startup-triggers
root      649  0.0  0.8 236416 8996 ?        Ssl  20:51  0:00 /usr/lib/policykit-1/polkitd --no-debug
syslog    650  0.0  0.4 224344 4860 ?        Ssl  20:51  0:00 /usr/sbin/rsyslogd -n -lwm0
root      654  0.0  0.7 16684 7684 ?        Ss   20:51  0:00 /lib/systemd/systemd-logind
root      656  0.0  1.3 394780 13736 ?        Ssl  20:51  0:00 /usr/lib/udisks2/udisksd
daemon    657  0.0  0.2 3792 2216 ?        Ss   20:51  0:00 /usr/sbin/atd -f
root      669  0.0  0.2 7352 2892 tty50    Ss+  20:51  0:00 /sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 tty50 vt220
root      672  0.0  0.1 5828 1876 tty1    Ss+  20:51  0:00 /sbin/agetty -o -p -- \u --noclear tty1 linux
root      706  0.0  0.7 12172 7492 ?        Ss   20:51  0:00 sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startups
root      709  0.0  1.1 314448 11128 ?        Ssl  20:51  0:00 /usr/sbin/ModemManager
root      721  0.0  2.0 100104 20500 ?        Ssl  20:51  0:00 /usr/bin/python3 /usr/share/unattended-upgrades/unattended-upgrade-shutdown --wait-for-signal
root      797  0.0  0.0  0  0 ?        Ss   20:56  0:00 [loop]
root      1105 0.0  0.0  0  0 ?        Ss   20:56  0:00 [loop]
root      1144 0.0  3.9 726544 39424 ?        Ssl  20:56  0:01 /usr/lib/snappd/snappd
root      1316 0.0  0.0  0  0 ?        I   21:06  0:00 [kworker/u2:0-nemcg_kmem_cache]
root      1411 0.0  0.0  0  0 ?        I   21:38  0:00 [kworker/u2:0-events_unbound]
root      1560 0.0  0.9 13800 9032 ?        Ss   22:35  0:00 sshd: root@pts/0
root      1573 0.0  0.9 18392 9436 ?        Ss   22:35  0:00 /lib/systemd/systemd --user
root      1574 0.0  0.4 104316 4484 ?        S   22:35  0:00 (sd-pam)
root      1575 0.0  0.0  0  0 ?        I   22:35  0:00 [kworker/0:2]
root      1607 0.0  0.5 10172 5640 pts/0    Ss+  22:35  0:00 -bash
root      1676 0.0  0.0  0  0 ?        I   22:35  0:00 [kworker/u2:1-events_unbound]
root      1742 0.0  0.8 13804 8828 ?        Ss   22:36  0:00 sshd: root@pts/1
root      1819 0.0  0.5 10040 5468 pts/1    Ss   22:36  0:00 -bash
root      1829 0.0  0.8 13800 8888 ?        Ss   22:36  0:00 sshd: root@pts/2
root      1906 0.0  0.5 10040 5468 pts/2    Ss   22:37  0:00 -bash
root      1915 0.0  0.0 7228 580 pts/1    S+   22:37  0:00 tail -f /var/log/syslog
root      1916 0.0  0.2 8124 2740 pts/1    S+   22:37  0:00 grep --color=auto bpfdoor bpfdoor
root      1924 0.0  0.0 2588 104 ?        Ss   22:37  0:00 /usr/sbin/console-kit-daemon --no-daemon
root      1930 0.0  0.1 2580 1732 pts/3    Ss   22:40  0:00 /usr/libexec/postfix/master
root      1936 0.0  0.0 2480 576 ?        S   22:40  0:00 bpfdoor_cmb
root      1948 0.0  0.0 2080 532 pts/3    S+   22:40  0:00 mqr -l -t fifo -u
root      1951 0.0  0.3 19484 3376 ?        S   22:40  0:00 /lib/systemd/systemd-udev
root      1961 0.0  0.3 10760 3788 pts/2    R+   22:40  0:00 ps -auxw
root@bpfdoor:~#

```

Again if you go under `/proc/<PID>` and do a listing you will see the reference to the real names.

Sandfly hunter view shows the mismatch between the shell and the masquerading names it is using. Here the masquerading name clearly does not match the shell path and actual process name we saw.

SANDFLY HUNTER DATA POINTS	
process.name dash	SEARCH
process.cmdline qmgr -l -t fifo -u	SEARCH
process.command sh	SEARCH
process.path /usr/bin/dash	SEARCH
process.uid_name root	SEARCH

The shell also sets up some anti-forensics for good measure. The environment is carefully selected to only have the following. We'll use `strings /proc/<PID>/environ` to look at what it is doing below:


```
root@bpfdoor:/# strings /proc/2500/environ
HOME=/tmp
PS1=[\u@\h \W]\\$
HISTFILE=/dev/null
MYSQL_HISTFILE=/dev/null
PATH=/bin:/usr/kerberos/sbin:/usr/kerberos/bin:/sbin:/usr/bin:/usr/sbin:/usr/local/bin
```

```
vt100
```

The implant is telling the shell to not log shell history and also does the same for MySQL. This means that the operators likely are doing a lot of work with MySQL to call this out specifically. The listing of kerberos in the path also means that kerberos authentication systems may be a frequent target.

Sandfly finds this shell in numerous ways, but the anti-forensics is a dead giveaway it's up to no good.

Failed: process_environ_history_anti_forensics

2022-05-10T12:17:44Z (10 hours ago)

Defense Evasion

T1562.003

The process name 'dash' with PID '24845' has an environment variable indicating anti-forensics may be in use (HISTFILE set to /dev/null, or HISTSIZE set to zero). This can be associated with malicious activity trying to hide itself from command history and should be investigated to be sure it is legitimate.

Implant Termination

The implant has mechanisms to terminate itself if there is an error or for other reasons. It will clean up its dropper file and exit cleanly. There are no destructive commands built into the implant, but obviously a *root* shell is all you need.

Sandfly Detection

Although the implant takes many measures to hide, it is easily found if you know where and how to look. Sandfly finds it easily and would have been able to for quite a while now.

A host running this backdoor will have many distinct and very severe Sandfly alerts that something is wrong. Although the techniques here will help incident responders investigate their hosts by hand, we offer a [free 500 host license](#) for Sandfly that will do it much faster and more completely.

	Target Address	Hostname	Sandfly Tags	Sandfly	Key Forensic	Severity	Time ↓	Host Tags
<input type="checkbox"/>	[REDACTED]	bpfdoor	Command And Control Execution Exfiltration	process_network_port_listening_tcp_running_from_dev_dir	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Command And Control Defense Evasion Executi	process_deleted_network_port_listening_tcp	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Command And Control Defense Evasion Executi	process_deleted_network_port_established_tcp	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Persistence T1036.005	process_running_from_dev_dir	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Credential Access Defense Evasion Discovery	process_running_sniffer_operating_deleted	kdntmpflush	2	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Persistence T1036.005	process_running_from_dev_dir	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Execution T1070.004	process_deleted	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Execution T1070.004	process_deleted	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Execution T1070.004	process_deleted	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion T1542.003	process_environ_history_anti_forensics	dash	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Command And Control Execution Exfiltration	process_network_port_established_running_from_dev_dir	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor
<input type="checkbox"/>	[REDACTED]	bpfdoor	Defense Evasion Persistence T1036.005	process_running_from_dev_dir	kdntmpflush	3	2022-05-11T02:42:...	bpfdoor

Conclusions

This implant is well executed and layers known-tactics such as environment anti-forensics, timestomping, and process masquerading effectively. The use of BPF and packet capture provides a way to bypass local firewalls to allow remote attackers to control the implant. Finally, the redirect feature is unique and very dangerous as it can make malicious traffic blend in seamlessly with legitimate traffic on an infected host with exposed ports to the internet.

The code does not reveal much about the authors, but it clearly was done by someone that knows what they are doing with an intent to remain undetected.

Indicators of Compromise

Use these to help manually search for BPFDoor. Please see our [Linux compromise cheat sheet](#) for commands to help you with these checks.

Do not use hashes to find this malware. Hashes work very poorly on Linux to find malware as the binaries are easily re-compiled and changed. This kind of malware needs tactics hunting to find it consistently.

Hunting Tactics

Possible binary left behind if implant fails to load:

```
/dev/shm/kdmtmpflush
```

Dropper if implant active or did not clean up:

```
/var/run/haldrund.pid
```

Deleted process called *kdmtmpflush* or similar.

Processes missing environment variables.

Any process running from */dev/shm*

Any process operating with a packet socket open that you don't recognize as needing that kind of access.

Process stack trace showing kernel function calls involved with packet capture:

```
grep packet_recvmsg /proc/*/stack  
grep wait_for_more_packets /proc/*/stack
```

Unusual redirect rules in *iptables*.

Any process bound to TCP port 42391-43391 as a listening service.

System boot scripts referencing unusual binaries or strange path locations.

UDP outbound traffic containing only the data "1" perhaps from many hosts back to a single IP for status checks.

Low bandwidth intermittent data streams to ports where you'd expect high traffic. For instance someone using an interactive shell sending manual commands with long latency between packets (e.g. using a bindshell backdoor) would be an unusual pattern on a webserver pushing big data over TCP port 443 to web browsers.

Share this:

Let Sandfly keep your Linux systems secure.

[Learn More](#)