

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Code sandboxing

Alpha Abdoulaye - Pierre Marsais

EPITA Systems/Security Laboratory (LSE)

July 14, 2017

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Introduction

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Limit usage of some resources such as system calls and shared object functions
- But not from the whole program (we trust our `libc.so`, `ld.so`, ...)

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Needed when executing untrusted code on your machine.
- Allow or deny use of some “resources”
- Usually these “resources” are accessed through syscalls
- We already have namespaces(7) and seccomp(2)

Are all needs fulfilled ?

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

There is no “ready-to-use” solution for:

- Function usage
- Library usage

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Aim for:

- Speed
- Reliability
- Security

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Solutions

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Trap at each function call
- Check if the call is righteous
- Continue as if nothing happened

CODE:

```
...
call my_func@plt
...
```

PLT:

```
PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
```

```
PLT(n): // my_func@plt
  jmp *(my_func@GOT)
  push n
  jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
  resolver address
...
```

```
my_func@GOT:
  PLT(n) + 6
...
```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0) ←
...

```

GOT:

```
...
GOT(2):
  resolver address
...
my_func@GOT:
  PLT(n) + 6
...

```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
PLT(0):
push GOT(1)
jmp *(GOT(2)) // resolver
...
```

```
PLT(n): // my_func@plt
jmp *(my_func@GOT)
push n
jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
resolver address
...
my_func@GOT:
PLT(n) + 6
...
```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
  ...
```

```
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0)
  ...
```

GOT:

```
...
GOT(2):
  resolver address
  ...
my_func@GOT:
  PLT(n) + 6
  ...
```

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
  resolver address
...
my_func@GOT:
  PLT(n) + 6
...
```

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
  resolver address
...
my_func@GOT:
  PLT(n) + 6
...
```

How a library function is called

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
PLT(0):
push GOT(1)
jmp *(GOT(2)) // resolver
...
```

```
PLT(n): // my_func@plt
jmp *(my_func@GOT)
push n
jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
resolver address
...
my_func@GOT:
PLT(n) + 6
...
```

How a library function is called

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0) ←
...
```

GOT:

```
...
GOT(2):
  resolver address
...
my_func@GOT:
  PLT(n) + 6
...
```


How a library function is called

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0) ←
...
```

GOT:

```
...
GOT(2):
  resolver address
...
my_func@GOT:
  PLT(n) + 6
...
```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
  ...
```

```
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0)
  ...
```

GOT:

```
...
GOT(2):
  resolver address
  ...
my_func@GOT:
  my_func
  ...
```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
→ PLT(0):
  push GOT(1)
  jmp *(GOT(2)) // resolver
...
```

```
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
  push n ←
  jmp PLT(0)
...
```

GOT:

```
...
GOT(2):
  resolver address
...
```

```
my_func@GOT:
  my_func
...
```

- Code sandboxing
- Alpha Abdoulaye - Pierre Marsais
- Introduction
- Solutions
- Elf trickery
- Virtualization
- Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
...
PLT(n): // my_func@plt ←
        jmp *(my_func@GOT) →
...

```

GOT:

```
...
my_func@GOT:
    my_func
...

```

- Code sandboxing
- Alpha Abdoulaye - Pierre Marsais
- Introduction
- Solutions
- Elf trickery
- Virtualization
- Conclusion

CODE:

```
...
call my_func@plt
...
```

PLT:

```
...
PLT(n): // my_func@plt ←
        jmp *(my_func@GOT) →
...

```

GOT:

```
...
my_func@GOT:
    my_func
...

```

CODE:

```
...
call my_func@plt
...
```

PLT:

```
...
PLT(n): // my_func@plt ←
  jmp *(my_func@GOT) →
...
```

GOT:

```
...
my_func@GOT:
  my_func
...
```

- Code sandboxing
- Alpha Abdoulaye - Pierre Marsais
- Introduction
- Solutions
- Elf trickery
- Virtualization
- Conclusion

CODE:

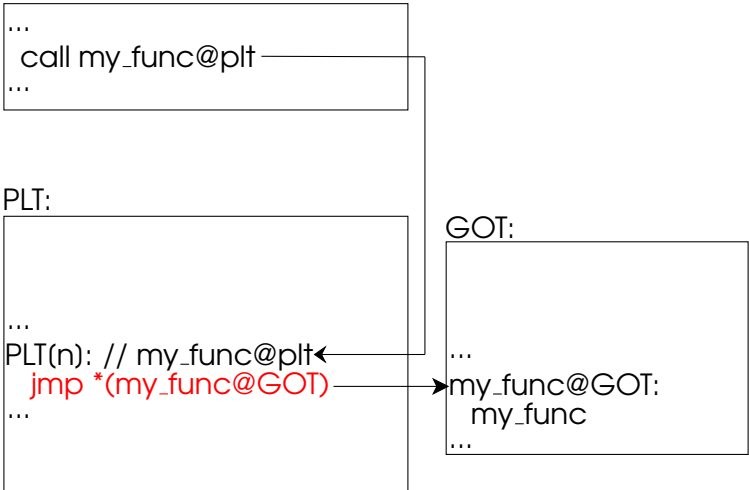
```
...
call my_func@plt
...
```

PLT:

```
...
PLT(n): // my_func@plt
        jmp *(my_func@GOT)
...
```

GOT:

```
...
my_func@GOT:
    my_func
...
```



- Code sandboxing
- Alpha Abdoulaye - Pierre Marsais
- Introduction
- Solutions
- Elf trickery
- Virtualization
- Conclusion

CODE:

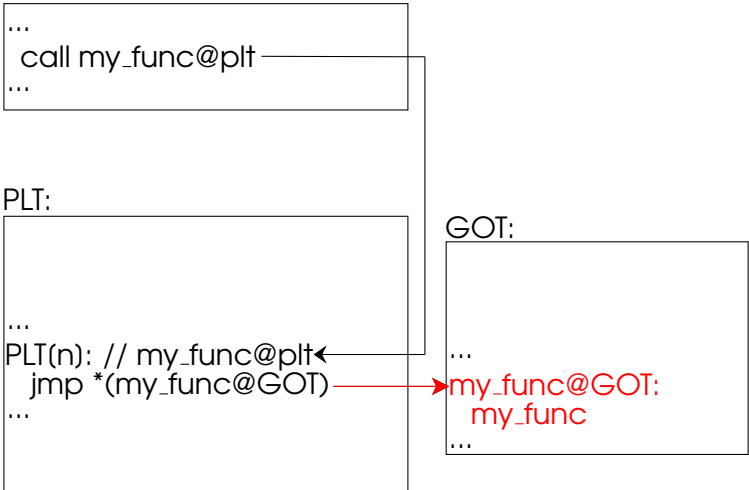
```
...
call my_func@plt
...
```

PLT:

```
...
PLT(n): // my_func@plt ←
        jmp *(my_func@GOT)
...
```

GOT:

```
...
my_func@GOT:
    my_func
...
```



Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

We have two solutions:

- Disallow GOT reads of the sandboxed ELF
- Disallow code execution of executable mapping

Then handle the rights violation and check if the resource access is allowed or not.

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

GOT protection can be bypassed.
The correct solution would be unallowing execution of
executable mappings.

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Elf trickery

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Can we solve our problem without privileged code ?

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Change mapping rights
- Handle mapping violation

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- `ptrace(2)`
- `procfs(5)`

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

No:

- How to change mapping permissions from the tracer ?
- What about non-GOT data on GOT pages ?
- What about multithreaded programs ?

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

How to change mappings from the tracer ?

- We can link an ELF to the sandboxed binary.
- We can use signal handlers in order to protect and unprotect the GOT.
- Use ELF constructors to setup everything.

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

How to handle non-GOT data on GOT pages ?

- GOT doesn't necessarily start and end at pages boundaries
- We can force this, with a custom linker script
- All we need is to customize the default linker script to align the GOT and export its size

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- LD_BIND_NOW=1
- Cache authorized GOT access

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

We can't allow a lot of stuff for the sandboxed application:

- We currently need to link an object to the sandboxed application
- `mprotect` can't be used to `PROT_READ` the GOT
- `SIGSEGV` can't be handled
- Libraries addresses can be leaked

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Address space leaks
 - `/proc/self/*`
 - `auxv`
 - some syscalls
 - addresses on stack and structures
- Functions pointers in structures
- `dlopen(3)`, `dlsym(3)`...

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Idea taken from OpenBSD
- If the user gets a libc address, and knows what libc is used, it can easily call any function
- The problem arise for any libs, but the libc is the more annoying for us
- We currently have a script to randomize the glibc
- Additional work needed for other libraries

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Virtualization

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Extended Page Table
- Additional translation level
- Hardware assisted
- Solve multi-threading problem

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Lightweight
- Extendable
- “C++ in Kernel”
- Multi-platform

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- `ptrace(2)`
- `/proc/[pid]/maps`
- `/proc/[pid]/pagemap`
- `linkmap`, `symbols`, etc.

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- vmcall to report to hypervisor
- Virtual Machine Control Structure
 - VM State
 - Global Configuration
- VM Exits
- Enable EPT violation
- Convert to Virtualization Exception

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

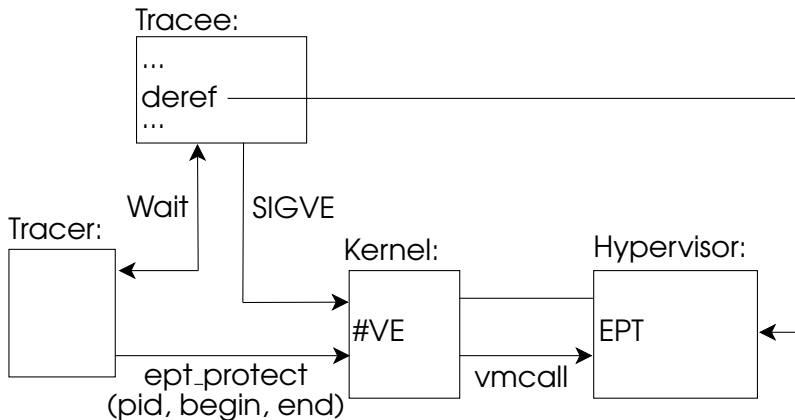
Elf trickery

Virtualization

Conclusion

- Handle #VE
- Trap on protected code
- Protect executable and check
- Decide!!!!
- And so on...

- Code sandboxing
- Alpha Abdoulaye - Pierre Marsais
- Introduction
- Solutions
- Elf trickery
- Virtualization
- Conclusion



Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Conclusion

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Be sure that our solution is foolproof
- handle multithreaded programs
- Work on performance
- What about statically linked ELF's ?
- ROP ?

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

Questions ?

Code
sandboxing

Alpha
Abdoulaye -
Pierre
Marsais

Introduction

Solutions

Elf trickery

Virtualization

Conclusion

- Alpha Abdoulaye - alpha@lse.epita.fr
- Pierre Marsais - pierre.marsais@lse.epita.fr