

RandCompile Removing Forensic Gadgets from the Linux Kernel to Combat its Analysis

Fabian Franzen, Andreas Chris Wilhelmer, Jens Grossklags
Technical University of Munich


December 7, 2023



1. Current State of Memory Forensics
2. How to combat Modern Forensic Tools: RandCompile
3. Evaluation
4. Future Research & Conclusion

Current State of Memory Forensics

- ▶ **Memory-Forensic:** The science of deducing **information** about an operating system **state** out of a **memory dump**
- ▶ Allows to reason about
 - ▶ Process List
 - ▶ (Cryptographic-)Secrets
 - ▶ IPs/MAC-Addresses of devices in proximity
 - ▶ ...
- ▶ Complexity depends on available information.
 - ▶ **Debugging Symbols** of operating system



Recent Developments in Linux Memory Forensics

New Challenges for analysts:

- ▶ **Structure Layout Randomization** (since 2017)
 - ▶ Binary Layout of data structures is modified at compile time.
 - ▶ Primarily a Binary Exploitation defense, but effective against forensic tools

Recent Developments in Linux Memory Forensics

New Challenges for analysts:

- ▶ **Structure Layout Randomization** (since 2017)
 - ▶ Binary Layout of data structures is modified at compile time.
 - ▶ Primarily a Binary Exploitation defense, but effective against forensic tools

Research Progress:

- ▶ Tools are capable to deal with **Structure Layout Randomization**
- ▶ OS-agnostic tools
 - ▶ Certain implementation characteristics are shared between OSES
 - ▶ Operate with minimal additional information on MacOS, Linux, Windows, and other operating systems

Forensic Gadgets

We systematized and grouped last-generation tools by the essential OS artifacts used to enable their analysis (**Forensic Gadgets**):

- ▶ **Special Comm** - "Special" strings allowing easy identification of the process list in the memory dump.

Forensic Gadgets

We systematized and grouped last-generation tools by the essential OS artifacts used to enable their analysis (**Forensic Gadgets**):

- ▶ **Special Comm** - "Special" strings allowing easy identification of the process list in the memory dump.
- ▶ **Symbol Tables** - The OS maintains a list of its own functions/global variables for dynamic loading of drivers and or additional functionality (eBPF, ftrace, ...).

Forensic Gadgets

We systematized and grouped last-generation tools by the essential OS artifacts used to enable their analysis (**Forensic Gadgets**):

- ▶ **Special Comm** - "Special" strings allowing easy identification of the process list in the memory dump.
- ▶ **Symbol Tables** - The OS maintains a list of its own functions/global variables for dynamic loading of drivers and or additional functionality (eBPF, ftrace, ...).
- ▶ **ABI Constraints** - The compiled kernel code follows predictable patterns revealing location/layout of data structures (i.e. using **offset revealing instructions**).

Forensic Gadgets

We systematized and grouped last-generation tools by the essential OS artifacts used to enable their analysis (**Forensic Gadgets**):

- ▶ **Special Comm** - "Special" strings allowing easy identification of the process list in the memory dump.
- ▶ **Symbol Tables** - The OS maintains a list of its own functions/global variables for dynamic loading of drivers and or additional functionality (eBPF, ftrace, ...).
- ▶ **ABI Constraints** - The compiled kernel code follows predictable patterns revealing location/layout of data structures (i.e. using **offset revealing instructions**).
- ▶ **Order of Fields** - The data structure layout (especially without Structure Layout Randomization) is foreseeable.

Forensic Gadgets

We systematized and grouped last-generation tools by the essential OS artifacts used to enable their analysis (**Forensic Gadgets**):

- ▶ **Special Comm** - "Special" strings allowing easy identification of the process list in the memory dump.
- ▶ **Symbol Tables** - The OS maintains a list of its own functions/global variables for dynamic loading of drivers and or additional functionality (eBPF, ftrace, ...).
- ▶ **ABI Constraints** - The compiled kernel code follows predictable patterns revealing location/layout of data structures (i.e. using **offset revealing instructions**).
- ▶ **Order of Fields** - The data structure layout (especially without Structure Layout Randomization) is foreseeable.
- ▶ **Pointer Graph** - The pointers between the kernel objects form a characteristic graph revealing e.g. the process list uniquely out of the set of objects.

Systematization of Last Generation Tools

Tool	Year	Analysis Subject	FG 1: Special comm	FG 2: Symbol Tables	FG 3: ABI Constraints	FG 4: Order of Fields	FG 5: Pointer Graph	Recovery Scope
Linux-specific								
KATANA	2022	Offset Revealing Instructions		X	X			All structures
Trustzone Rootkit	2022	Kernel Runtime Data	X					Selected structures
LOGICMEM	2022	Kernel Runtime Data	X	X		X	X	Selected structures
AUTOPROFILE	2021	Offset Revealing Instructions		X	X	X		All structures
OS-agnostic								
FOSSIL	2023	Kernel Runtime Data	X				X	All structures
HYPERLINK	2016	Kernel Runtime Data	X				X	Selected structures

How to Combat Modern Memory Forensic Tools?

Harden **Linux** systems against **automated** forensic analysis

	Forensic Gadgets					Transformation	
	FG-1	FG-2	FG-3	FG-4	FG-5	GCC Plugin	Manual
String and Pointer Encryption	✓				✓		✓
Better Data-Order Randomization				✓			✓
Externalize <code>printf</code> Format Strings			✓			✓	
Adding Bogus Parameters with Artificial Memory Accesses			✓			✓	

- ▶ Perform **selected transformations** on the kernel to remove **four** out **five** forensic gadgets.
 - ▶ **two** are **automatically** applied (by a compiler plugin)
 - ▶ **two** applied manually in form of **kernel patch**
- ▶ **Disclaimer:** Perfect Obfuscation is in general **not possible!** This is a hardening mechanism against **automated tools**.

ABI Randomization

KATANA and AUTOPROFILE target **FG 3**

- ▶ Offset Revealing Instructions reveal layout of data structures
- ▶ ABI mandates calling convention
 - ▶ Allows a structural matching of generated machine code with the source code

Example:

```
1 do_stuff(current->mm①, current->  
  ↪ cred②, &g③);
```



```
1 mov    rdx③,0xffffffff82019c60  
2 mov    rax,QWORD PTR gs:0x16d00  
3 mov    rsi②,QWORD PTR [rax+0x10]  
4 mov    rdi①,QWORD PTR [rax+0x440]  
5 call   ffffffff811bacd0 <do_stuff>
```


ABI Randomization

Countermeasures by RandCompile

- ▶ **Shuffle** the order of the arguments at call site and implementation site
- ▶ Applied automatically to all functions through a compiler plugin.

Issues

- ▶ Functions with few parameters have few possibilities for randomization

Example:

```
1 do_stuff(current->cred①, current->mm  
↪ ②, &g③);
```



```
1 mov rdx③,0xffffffff82019c60  
2 mov rax,QWORD PTR gs:0x16d00  
3 mov rsi②,QWORD PTR [rax+0x440]  
4 mov rdi①,QWORD PTR [rax+0x10]  
5 call ffffffff811bacd0 <do_stuff>
```

ABI Randomization

We can add bogus parameters to functions with few parameters

- ▶ This can be undone by an analysis tool that has access to the source code
- ▶ Also add bogus assembly code hurting performance

Example:

```
1 int64_t bogusstuff[6];
2 do_stuff(current->cred①, current->mm
   ↪ ②, bogusstuff[0]③, &g④,
   ↪ bogusstuff[3]⑤, bogusstuff[5]
   ↪ ⑥);
```

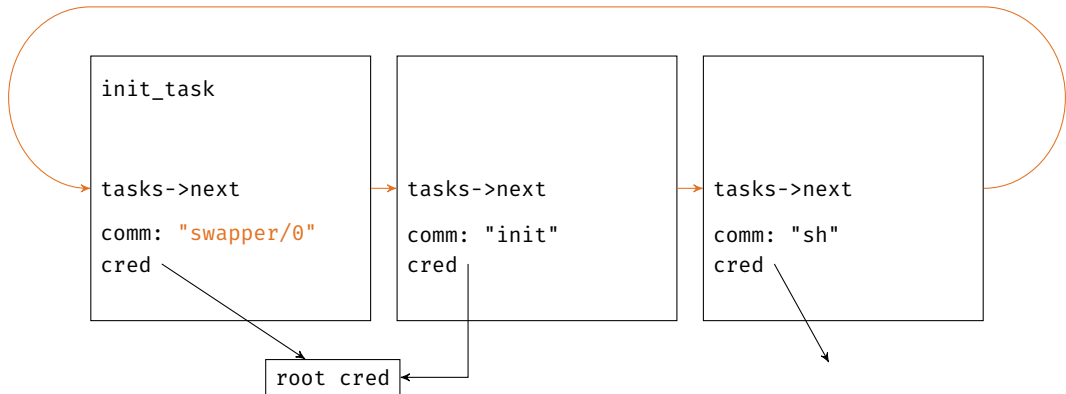


```
1 mov rcx④,0xffffffff82019c60
2 mov r8⑤,QWORD PTR [rsp+0x18]
3 mov r9⑥,QWORD PTR [rsp+0x28]
4 mov rax,QWORD PTR gs:0x16d00
5 mov rsi②,QWORD PTR [rax+0x440]
6 mov rdx③,QWORD PTR [rsp]
7 mov rdi①,QWORD PTR [rax+0x10]
8 call ffffffff811bacd0 <do_stuff>
```

Pointer & String Encryption

HYPERLINK and FOSSIL analyse the pointer graph of kernel objects (FG 5).

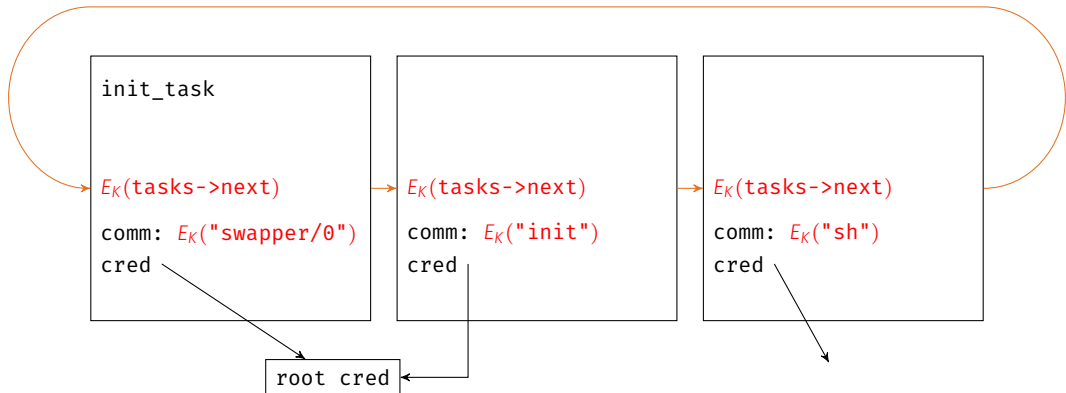
- ▶ e.g. the process information objects are connected by a **linked list**.
- ▶ first process in list contains well-known string (FG 1).



Pointer & String Encryption

HYPERLINK and FOSSIL analyse the pointer graph of kernel objects (FG 5).

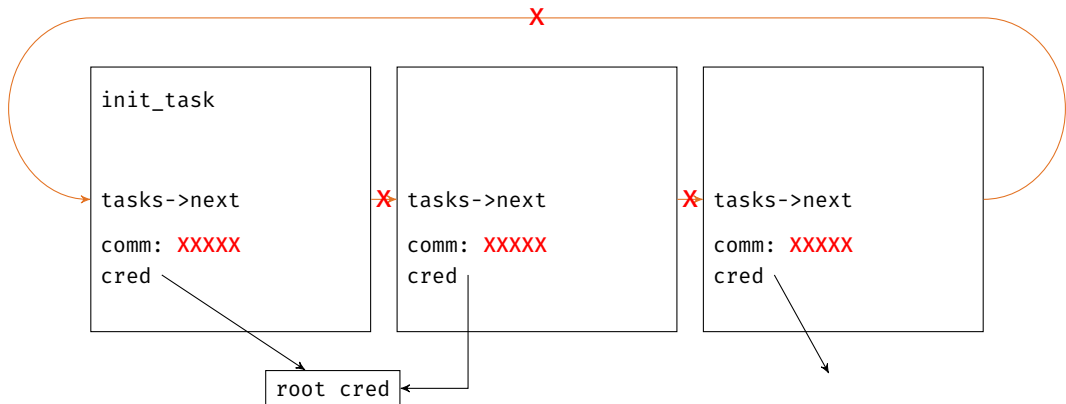
- ▶ e.g. the process information objects are connected by a **linked list**.
- ▶ first process in list contains well-known string (FG 1).
- ▶ **Encrypt** Pointers and Strings in process information objects



Pointer & String Encryption

HYPERLINK and FOSSIL analyse the pointer graph of kernel objects (FG 5).

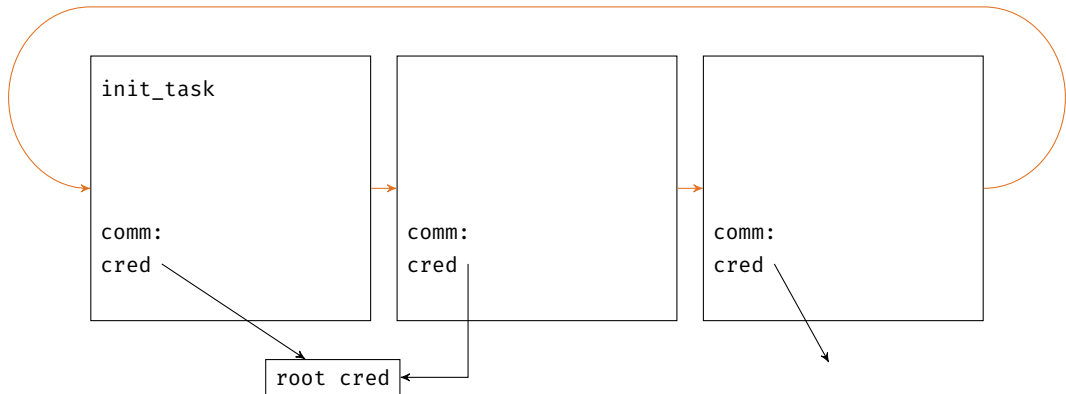
- ▶ e.g. the process information objects are connected by a **linked list**.
- ▶ first process in list contains well-known string (FG 1).
- ▶ **Encrypt** Pointers and Strings in process information objects



Pointer & String Encryption

HYPERLINK and FOSSIL analyse the pointer graph of kernel objects (FG 5).

- ▶ e.g. the process information objects are connected by a **linked list**.
- ▶ first process in list contains well-known string (FG 1).
- ▶ **Encrypt** Pointers and Strings in process information objects
 - ▶ Store **Encryption Key** as **immediate** value in the **compiled machine code**.



Evaluation

Effectiveness against Offset Revealing Instruction based Analysis

	Base	RandCompile (no bogus)	RandCompile (-printk, -memref)	RandCompile (-printk)	RandCompile (full)
List Modules	✓	✓	✓	✓	✓
Members reconstructed	2	2	2	2	2
Task Listing	✓	✗	✗	✗	✗
Members reconstructed	6	5	5	4	4
List Files	✓	✗	✗	✗	✗
Members reconstructed	16	15	8	7	7
Dmesg Log	✓	✓	✓	✓	✗


- ▶ We perform the core analysis of **KATANA** with and without RandCompile.
 - ▶ Already **a single fault** during reconstruction causes a fault!

Effectiveness against Kernel Runtime Data Analysis

- ▶ Encryption of the string "**swapper/0**" (FG-1) is most effective.
 - ▶ Stops LOGICMEM, Trustzone Rootkit, and HYPERLINK from operating
 - ▶ FOSSIL analysis performance is degraded. It depends on the analysts queries.
- ▶ Pointer Encryption
 - ▶ Degrades analysis opportunities of LOGICMEM, Trustzone Rootkit, and HYPERLINK further
 - ▶ Further degrades attack possibilities of FOSSIL
 - ▶ Future Work: Encrypt also other kernel pointers

Performance


Results using the `lmbench` Microbenchmark (runtimes are normalized to 1):



Performance

Results using the **lmbench** Microbenchmark (runtimes are normalized to 1):

- ▶ Less than 1-3 percent overhead on average




- ▶ Are you applying only sound transformations?
 - ▶ **Yes**. RandCompile does not change the semantic/core functionality of the Linux kernel.
- ▶ Does not **confidential computing** (CC) (like AMD-SEV) mitigate this problem?
 - ▶ RandCompile **complements** protection of CC approaches. I.e. AMD-SEV expects a Linux kernel to not trust his drivers.
- ▶ Can this be used as a **binary exploitation** defense?
 - ▶ Yes. In combination with Control Flow Integrity protections, it makes abusing existing kernel functions in ROP chains harder.
- ▶ Is it a problem that the defenses are applied at **compile** time?
 - ▶ **Partially**. Applying them during runtime would allow for more widespread use. Applying them at compile time adds diversity to the binary layout.

Conclusion

Conclusion

- ▶ RandCompile is an obfuscation tool for the Linux Kernel to harden it various memory forensic tools.
- ▶ It is **effective** against **modern** forensic analysis tools.
- ▶ It completes and extends the Structure Layout Randomization, a mainlined Linux kernel feature.



We have source code!