
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Liljestrand, Hans; Gauhar, Zaheer; Nyman, Thomas; Ekberg, Jan-Erik; Asokan, N.
Protecting the stack with PACed canaries

Published in:
SysTEX '19, October 27, 2019, Huntsville, ON, Canada

DOI:
[10.1145/3342559.3365336](https://doi.org/10.1145/3342559.3365336)

Published: 27/10/2019

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Published under the following license:
Unspecified

Please cite the original version:
Liljestrand, H., Gauhar, Z., Nyman, T., Ekberg, J.-E., & Asokan, N. (2019). Protecting the stack with PACed canaries. In SysTEX '19, October 27, 2019, Huntsville, ON, Canada ACM.
<https://doi.org/10.1145/3342559.3365336>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Protecting the stack with PACed canaries

Hans Liljestrand
Aalto University, Finland
Huawei Technologies Oy, Finland
hans@liljestrand.dev

Zaheer Gauhar
Aalto University, Finland
zaheer.gauhar@aalto.fi

Thomas Nyman
Aalto University, Finland
thomas.nyman@aalto.fi

Jan-Erik Ekberg
Huawei Technologies Oy, Finland
Aalto University, Finland
jan.erik.ekberg@huawei.com

N. Asokan
University of Waterloo, Canada
asokan@acm.org

Abstract

Stack canaries remain a widely deployed defense against memory corruption attacks. Despite their practical usefulness, canaries are vulnerable to memory disclosure and brute-forcing attacks. We propose PCan, a new approach based on ARMv8.3-A pointer authentication (PA), that uses dynamically-generated canaries to mitigate these weaknesses and show that it provides more fine-grained protection with minimal performance overhead.

1 Introduction

Run-time attacks that exploit memory errors to corrupt program memory are a prevalent threat. Overflows of buffers allocated on the stack are one of the oldest known attack vectors [12, 17]. Such exploits corrupt local variables or function return addresses. Modern attacks techniques—such as return-oriented programming (ROP) [16] and data-oriented programming (DOP)[6]—can use this well-known attack vector to enable attacks which are both expressive and increasingly hard to detect. The fundamental problem is insufficient bounds checking in memory-unsafe languages such as C / C++. Approaches for hardening memory-unsafe programs have been proposed, but tend to incur high performance overheads, and are therefore impractical to deploy [19]. An exception is a technique called *stack canaries* [2], which is both efficient and can detect large classes of attacks. Consequently stack canaries are widely supported by compilers and used in all major operating systems today [2, 9].

Widely deployed stack canary implementations suffer from one or more of the following weaknesses: they 1) rely on canary values that are fixed for a given run of a program [2]; 2) store the reference canary in insecure memory, where an attacker can read or overwrite it [9]; or 3) use only a single canary per stack frame and therefore cannot detect overflows that corrupt only local variables.

The recently introduced ARMv8.3-A pointer authentication (PA) [1] hardware can be used to verify return addresses [14], effectively turning the return address itself into a stack canary. However, PA on its own is susceptible to *reuse attacks*, where an attacker substitutes one authenticated pointer with another [8]. State-of-the-art schemes harden PA return-address protection to ensure that protected return address as statistically unique to a particular control-flow path, and therefore cannot be substituted by an attacker [7].

We propose fine-grained PA-based canaries that: 1) protect individual variables from buffer overflow, 2) do not require secure storage for reference canaries, 3) can use existing return-address protection [7] as an *anchor* to produce canaries which are statistically unique to a particular function call, and 4) are efficient, since they can leverage hardware PA instructions both for canary generation and verification.

Our contributions are:

- **PCan**: A **fine-grained**, and **efficient** PA-based canary scheme (Section 5).
- A **realization of PCan** for LLVM (Section 6).
- **Evaluation** showing that PCan is more secure than existing stack canaries and has only a small performance impact (Section 7).

2 Background

A *stack canary*[2] is a value placed on the stack such that a buffer overflow will overwrite it before corrupting the return address (Figure 1). The buffer overflow can be detected by verifying the integrity of the canary before return.

The initially proposed canaries were randomly generated 32-bit values assigned at process startup and stored within the process memory [2]. The canaries must remain confidential to prevent an attacker \mathcal{A} from avoiding detection by writing back the correct canary when triggering the buffer overflow. *Terminator canaries* [3], consisting of string terminator values (e.g., $0x0$, EOF, and $0xFF$) can prevent \mathcal{A} from using string operations to read or write whole canaries, thereby thwarting run-time canary harvesting. Another approach is to re-generate canaries at run-time, for instance by masking them with the return address [4]. However, such techniques rely on the secrecy of the masking value.

© 2019. Copyright held by the owner/author(s). This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in SysTEX '19, October 27, 2019, Huntsville, ON, Canada <https://doi.org/10.1145/3342559.3365336>.




Figure 1. A stack canary is a value placed on the stack so that it will be overwritten by a stack buffer that overflows to the return address. It allows detection of overflows by verifying the integrity of the canary before function return.

Multi-threaded and forked processes can be vulnerable to guessing attacks if the canaries are shared. If the child process or thread is restarted after a crash, \mathcal{A} can execute a large number of guesses without resetting the canary. Moreover, such attacks can utilize incremental—e.g., byte-for-byte—guessing to efficiently find the canary value [10].

Strong adversaries with arbitrary memory read or write access can trivially circumvent any canary based solution; using reads alone allows \mathcal{A} to first read the correct canaries from memory and then perform a sequential overwrite that writes back the correct canaries while corrupting other data.

2.1 Stack canaries in modern compilers

Modern compilers such as LLVM/Clang and GCC provide the `-fstack-protector` feature that can detect stack-buffer overflows¹. It is primarily designed to detect stack overflows that occur in string manipulation. The default `-fstack-protector` option includes a canary only when a function defines a character array that is larger than a particular threshold. The default threshold value in GCC and LLVM is 8 characters, but in practice the threshold is often lowered to 4 to provide better coverage. However, a stack overflow can occur on other types of variables. The `-fstack-protector-all` option adds a canary to **all** functions. However, it can incur a substantial use of stack space and run-time overhead in complex programs.

The `-fstack-protector-strong` option provides a better trade-off between function coverage, run-time performance, and memory cost of stack canaries. It adds a canary to any function that 1) uses a local variable’s address as part of the right-hand side of an assignment or function argument, 2) includes a local variable that is an array, regardless of the array type or length, and 3) uses register-local variables.

Today `-fstack-protector-strong` is enabled by default for user-space applications in major Linux distributions, such as Debian and its derivatives². The `-fstack-protector` protects non-overflowing variables by rearranging the stack

¹<https://lists.lvm.org/pipermail/cfe-dev/2017-April/053662.html>

²<https://wiki.debian.org/Hardening>




Figure 2. PA verifies pointers using an embedded PAC generated from a pointer’s address, a 64-bit modifier and a hardware-protected key. (Figure from [7])




Figure 3. The signed return address generated by `-msign-return-address` effectively serves as a canary by allowing detection of stack-buffer overflows.

such that an overflow cannot corrupt them; but this protection cannot protect other buffers. On AArch64 the LLVM/Clang implementation of `-fstack-protector` uses a single reference canary value for the whole program. This in-memory reference canary is used to both store and verify the stack canary on function entry and return, respectively.

2.2 ARMv8-A Pointer Authentication

ARMv8.3-A PA is a instruction set architecture (ISA) extension that allows efficient generation and verification of pointer authentication codes (PACs); i.e., keyed message authentication codes (MACs) based on a pointer’s address and a 64-bit modifier [1]. The PAC is embedded in the unused bits of a pointer (Figure 2). On 64-bit ARM, the default Linux configuration uses 16-bit PACs. GNU/Linux has since 5.0 provided support for using PA in user-space. PA provides new instructions for generating and verifying PACs in pointers, and a generic `pacga` instruction for constructing a 32-bit MAC based on two 64-bit input registers. After a PAC is added to a pointer, e.g., using the `pacia` instruction, it can be verified later using the corresponding authentication instruction, in this case `autia`. A failed verification does not immediately cause an exception. Instead, PA corrupts the pointer so that any subsequent instruction fetch or dereference based on it causes a memory translation fault. The `pacga` instruction is an exception as it outputs the produced PAC to a given destination register; verification in this case must be performed manually by comparing register values.

Current versions of GCC and LLVM/Clang provide the `-msign-return-address` feature that uses PA to protect return addresses [14]. It signs the return address with the stack




Figure 4. When using only one canary per stack frame, an attacker could overflow a stack buffer to corrupt local variables without overwriting the canary.

pointer (SP) as modifier using `pacia lr, sp`. The integrity of the return address is verified before return by issuing the corresponding authentication instruction `autia lr, sp`. Signed return addresses provides similar protection to stack canaries, i.e., if a stack-buffer overflow corrupts the return address, this is detected when the return address is verified before returning from a function (Figure 3). However, PA is vulnerable to reuse attacks where previously encountered signed pointers can be used to replace latter signed pointers using the same key and modifier [8]. For instance, `-msign-return-address` can be circumvented by reusing a prior return address signed using the same SP value.

3 Adversary Model

In this work we consider an adversary \mathcal{A} that attempts to compromise the memory safety of a user-space process by exploiting a stack-buffer overflow. We do not consider adversaries at kernel or higher privilege levels. However, \mathcal{A} can: 1) trigger any existing stack-buffer overflow, 2) use stack-buffer over-reads to read memory and 3) repeatedly restart the process and any child processes or threads in an attempt to brute force canaries. Adversaries with arbitrary memory read or write access cannot be thwarted with canary-based approaches and are beyond the scope of this work.

We assume that \mathcal{A} can analyze the target binary and therefore knows the exact stack layout of functions (barring variable-length buffers). This enables \mathcal{A} to target individual local variables reachable from a particular buffer overflow without overflowing canaries past the local variables (Figure 4). If \mathcal{A} further manages to exploit a buffer over-read or other memory disclosure vulnerability, they could also overflow past the canary by simply replacing the correct canary value during the overflow.

4 Requirements

To detect linear buffer overflows we require a design that fulfills the following requirements:

- R1** Each canary value should be statistically unique.
- R2** Reference canaries must not be modifiable by \mathcal{A} .
- R3** A stack-buffer overflow must always corrupt a canary.

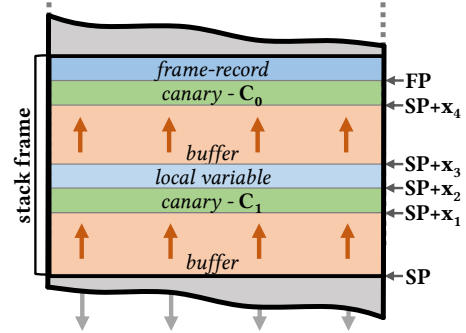


Figure 5. To detect all overflows we inject canaries after any vulnerable stack buffer.

5 Design

We propose PCan, a PA-based canary design that employs multiple function-specific canaries. By placing canaries after any array that could overflow (Figure 5), PCan can detect overflows that only corrupt local variables. This prevents \mathcal{A} from performing precise overflows that corrupt only local variables without overwriting the canary (**R3**). To exploit an overflow without detection \mathcal{A} is instead forced to learn the correct canary values and write them back into place.

In contrast to traditional approaches, PCan avoids exposing reference canaries in memory (**R2**). Instead, canaries are re-generated or verified directly using PA. \mathcal{A} thus cannot manipulate the reference canaries, and must instead leak the specific on-stack canary or attempt a brute-force attack.

The canaries are generated with PA, using a modifier m consisting of a 16-bit function identifier and the least-significant 48 bits from SP: $m = SP * 2^{16} + \text{function-id}$. This modifier makes canaries function-dependent and, when SP differs, different for each call to the same function (**R1**). Because the canaries are generated at run-time and the PA keys are randomly set on each execution, the generated canaries are also statistically unique for each execution. To avoid detection \mathcal{A} must acquire the exact stack-canary belonging to the specific function call and cannot rely on pre-calculated canaries or those belonging to other function calls.

5.1 PA-based canaries

PA-based return-address protection [7, 8, 14] already effectively serves as a canary by detecting return-address corruption. We therefore propose a design that can be efficiently and easily integrated with existing return-address protection schemes, but also provide a stand-alone setup. The first canary in a function's stack frame, protecting the return-address, is either a `pacga`-generated stand-alone canary or the signed return address:

$$C_0 = \begin{cases} \text{pacga}(\text{SP}, m) & \text{if stand-alone} \\ \text{signed_return_address} & \text{if combined} \end{cases}$$

We denote a canary loaded from the stack with C' to indicate that it might have been corrupted by \mathcal{A} . Verification of C'_0 is done either by re-generating the stand-alone pacga canary or by relying on the return-address protection to verify it. To verify using pacga we re-generate C_0 and check that $C_0 = C'_0$.

Subsequent canaries, $C_i, i > 0$, consist of signed pointers to the previous canary:

$$C_i = \text{pacda}(C_{ptr_{i-1}}, m), i > 0$$

where C_{ptr_i} is a pointer to C_i . Verification of $C'_i, i > 0$ is done by authenticating and loading the canary to retrieve C'_{i-1} . If any C'_i is corrupted, authentication fails, causing the subsequent load to fault (Section 2.2) A successful chain of loads will yield C_0 , which is then verified as detailed above.

Our stand-alone scheme is more powerful than `-msign-return-address` in that it does not rely solely on the SP value. However, other schemes might provide better protection for the return address. For instance, PACStack [7] proposes a scheme that uses statistically unique modifiers to protect return addresses by maintaining the head of a chain of PACs in a single reserved register. We propose that PCan could be combined with such a mechanism by defining m as the PACStack authentication token $auth_i$ and C_0 as the PACStack protected return address. Because the m in this case would be statistically unique to a specific call-flow this would also harden the canaries C_i for $i > 0$.

6 Implementation

We implement PCan as an extension to LLVM 8.0 and using the stand-alone pacga approach (Section 5.1). To instrument the LLVM Intermediate Representation (IR) we added new LLVM intrinsics for generating and verifying PCan canaries. These intrinsics, along with instructions for storing and loading the canaries, are added through IR transformations before entering the target-specific compiler backend. We define corresponding target-specific intrinsics to leverage built-in register allocation before converting the intrinsics to hardware instructions in the pre-emit stage.

6.1 Canary creation

To instrument the function prologue PCan locates LLVM `alloca` instructions that allocate buffers in the entry basic block of each function. A new 64-bit allocation for the canaries is added after each existing `alloca`. Intrinsics for generating the canaries and storing them are then added. The instrumented code will generate a larger stack-frame to accommodate the canaries and include code to generate and store the canary values (Listing 1).

6.2 Canary verification

To verify canaries in the function epilogue, PCan loads them in reverse order, starting from the last C_n (Listing 2). Each canary C'_i is authenticated using `autda` and then dereferenced to acquire the next canary C'_{i-1} . To verify the final

```

1 canary-creation:
2  mov  x8, sp
3  movk x8, #3, lsl #48 ; x8 ← mod
4  pacga x10, sp, x8 ; x10 ← C0
5  sub  x9, x29, #0x8 ; x9 ← Cptr0
6  pacda x9, x10 ; x9 ← C1
7  str  x9, [sp, #40] ; store C1
8  stur x10, [x29, #-8] ; store C0

```

Listing 1. For a function with two vulnerable stack buffers PCan generates and stores two canaries.

```

1 canary-verification:
2  ldr  x8, [sp, #40] ; x8 ← C'1
3  mov  x8, sp
4  movk x8, #3, lsl #48 ; x8 ← mod
5  autda x8, x29 ; authenticate C'1
6  ldr  x8, [x8] ; x8 ← C'0
7  pacga x9, sp, x9 ; x9 ← C0
8  cmp  x8, x9 ; check C0 = C'0

```

Listing 2. To verify the integrity of canaries PCan first loads C'_1 , then authenticates it before using it to load C'_0 , which in turn is compared to the re-generated C_0 .

canary, C'_0 , PCan first re-generates C_0 and then performs a value comparison. Upon failure, an error handler is invoked, otherwise the function is allowed to return normally. As suggested in Section 5.1, the final canary can be replaced with a return-address protection scheme. The return address then serves as a canary that is verified using the corresponding protection scheme (e.g., `-msign-return-address`).

7 Evaluation

Due to lack of publicly available PA-capable hardware we have used an evaluation approach similar to prior work [7, 8]. We used the ARMv8-A *Base Platform Fixed Virtual Platform (FVP)*, based on *Fast Models 11.5*, which supports ARMv8.3-A for functional evaluation. For performance evaluation we used the PA-analogue from prior work [7] and performed measurements on a 96board Kirin 620 HiKey (LeMaker version) with an ARMv8-A Cortex A53 Octa-core CPU (1.2GHz) / 2GB LPDDR3 SDRAM (800MHz) / 8GB eMMC, running the Linux kernel v4.18.0 and BusyBox v1.29.2.

7.1 Performance

We evaluated the performance of PCan using the SPEC CPU 2017³ benchmark, and running it on the HiKey board. We cross-compiled the benchmarks on an x86 system using whole program LLVM⁴, and timed the execution of the individual benchmark programs using the time utility. Results are reported normalized

³<https://www.spec.org/cpu2017/>

⁴<https://github.com/travitch/whole-program-llvm>

benchmark	stack-protector		PCan	
505.mcf_r	-4.78%	(4.55)	-0.05%	(0.13)
519.lbm_r	-0.01%	(0.01)	0.04%	(0.02)
525.x264_r	-0.01%	(0.01)	1.80%	(0.01)
538.imagick_r	-0.01%	(0.01)	0.19%	(0.01)
544.nab_r	0.05%	(0.24)	-0.18%	(0.16)
557.xz_r	0.00%	(0.03)	0.04%	(0.06)
geo.mean.	-0.08%		0.03%	

Table 1. SPEC CPU 2017 performance overhead of PCan and `-fstack-protector-strong` compared to an uninstrumented baseline (standard error in parenthesis). Results indicate that both schemes incur a negligible overhead (geometric mean of 0.3% and < 0%, respectively).

to a baseline measured without PCan instrumentation and compiled with `-fno-stack-protector` (Table 1). We compare this baseline to two different setups; one using only `-fstack-protector-strong` and another using `-fno-stack-protector` and PCan instrumentation. Our results indicate that PCan incurs a very low overhead with a geometric mean of 0.30%. In some cases `-fstack-protector-strong` caused the benchmarks, we suspect this is caused by it rearranging the stack. Measurements were repeated 20 times and all binaries were compiled with `-O2` optimizations enabled.

7.2 Security

The initial `pacga` canaries used by PCan provide similar security to traditional canaries. To perform an overflow while avoiding detection \mathcal{A} must achieve the following goals: 1) find the location of canaries in relation to the overflowed buffer, 2) leak the specific canary values on the stack, and 3) write back the correct canaries when performing the buffer overflow. In our adversary model step 1) is trivial; \mathcal{A} can inspect the binary to analyze the stack layout. Step 2) could be achieved by leaking or modifying the in-memory reference values, but because PCan generates canaries on-demand, \mathcal{A} is forced to leak the values from the stack (**R2**). Moreover, because the canaries are statistically unique to a function and SP value \mathcal{A} cannot rely on finding just any canary and substitute it with one in the overflowed stack frame (**R1**). This limits the scope of attacks, as both the memory leak and overflow must happen within the lifetime of the attacked stack frame. By using multiple canaries—one after each buffer—PCan can detect overflows that only touch local variables (**R3**). Based on our evaluation PCan thus provides comprehensive protection with an overhead similar to currently deployed defenses.

8 Related Work

After the seminal article “Smashing the Stack for fun and profit” [12], the notion of *canaries* as a protection against buffer overflow was first introduced in StackGuard [2], and initial GCC compiler support appeared at the same time. StackGuard proposes to use a random canary, stored at the top of the stack (or in the thread local storage memory area), during program launch to thwart canary harvesting from the compiled code. The threat of canary harvesting and the added protection (especially for C) provided by terminator canaries was identified shortly thereafter [3]. The problem of canary copy and re-use was already identified by Etoh and Yoda in 2000 [4], where the stack-frame based canary protection was augmented by masking the canary value with the function return address. Later, Strackx et al. [18] argue against the futility of storing secrets in program memory, which supports using PA to generate canaries dynamically.

Another shortcoming of canary integrity are cases when the canary mechanism is subject to brute-force attacks, e.g., in the context of process forking. \mathcal{A} could use the canaries in forked child processes as oracles to perform brute-force guessing of canary values. Published solutions against this form of attack includes DynaGuard [13] and DCR [5]. Both solutions keep track of canary positions in the code, and re-initialize all canaries in a child process, at considerable performance overhead. DCR optimizes the canary location tracking by chaining canaries using embedded offsets - we inherit this notion of chaining canaries from their work, although we deploy these for canary validation whereas DCR uses the mechanism for canary rewriting. By combining the SP in the canaries PCan mitigates such attacks, but full protection would require a similar approach of re-initializing canaries on fork. Finally, the polymorphic canaries by Wang et al. [20] optimize away the need to rewrite canaries during fork, by adding a function-specific random mask to the stack canary, which effectively removes the opportunity for systematic canary brute-forcing.

Heap protection with canaries has received much less attention than stack protection, possibly because the optimal balance between validation and performance overhead when canaries are applied to the heap remains an open problem. The first paper on the subject was Robertson et al. in 2003 [15], but a more recent mechanism — HeapSentry by Nikiforakis et al. [11] puts effort on the unpredictability (randomness) of the heap canaries. HeapSentry consists of a wrapper for the allocator and a kernel module, and exhibits overheads at around 12%. Pointer bounds checking schemes offer protections stronger than canaries alone, but in comparison incur significant performance overheads [19].

9 Future Work

Our current approach only protects stack-based variables with a static size. Canaries for dynamic allocations cannot

be verified in the prologue because they might be either out of scope or overwritten by later dynamic allocations, and are currently not used by PCan. To prevent attacks that corrupt dynamic allocations, we propose to add instrumentation that protects dynamic allocations based on their life-time, i.e., which verifies the associated canaries immediately when the allocation goes out of scope. The existing LLVM allocation life-time tracking could be leveraged to implement this addition without significant changes to the compiler.

We plan to refine and expand our canary approach by using compile-time analysis—i.e., the StackSafetyAnalysis⁵ of LLVM—to omit instrumentation of buffers that can be statically shown to be safe. In some cases \mathcal{A} could achieve their goal before function return, i.e., before the canary corruption is detected. Such attacks could be detected earlier by utilizing the StackSafetyAnalysis to add checks after vulnerable steps during function execution, before the function epilogue.

We also plan to extend PCan instrumentation to cover heap allocations, similar to HeapSentry [15]. Because the PA-keys are managed by the kernel, PCan could be used for HeapSentry-like consistency checks from within the kernel, e.g., before executing system-calls.

10 Conclusion

Canaries are a well-established protection mechanism against errors in memory-unsafe programs. We present PCan, which provides hardware-assisted integrity-protection for canaries, inhibiting the most prevalent canary-circumvention techniques. Furthermore, we propose the notion of fine-grained canaries, where canaries not only protect the return address, but also detect overflows in individual data objects. We make available our compiler prototype at <https://github.com/pointer-authentication/PCan-llvm>, and measure its performance impact. Finally we point out further optimizations for fine-grained canaries, as well a solution path for protecting dynamic allocations.

11 Acknowledgments

This work was supported in part by the Intel Collaborative Research Institute for Collaborative Autonomous & Resilient Systems (ICRI-CARS), the Academy of Finland under grant nr. 309994 (SELIoT), and a Google ASPIRE award.

References

- [1] ARM Ltd. 2017. ARMv8 Architecture Reference Manual, for ARMv8-A architecture profile (ARM DDI 0487C.a). https://static.docs.arm.com/ddi0487/ca/DDI0487C_a_armv8_arm.pdf.
- [2] Crispin Cowan et al. 1998. StackGuard: Automatic adaptive detection and prevention of buffer-overflow attacks. In *Proc. USENIX Security '98*. 63–78.
- [3] Crispin Cowan et al. 1999. Protecting systems from stack smashing attacks with StackGuard. *Linux Expo* (1999).
- [4] Hiroaki Etoh and Kunikazu Yoda. 2000. *Protecting from stack-smashing attacks*. Technical Report. IBM Research Division, Tokyo Research Laboratory.
- [5] William H Hawkins et al. 2016. Dynamic canary randomization for improved software security. In *Proc. ACM CISR '16*. 9:1–9:7.
- [6] Hong Hu et al. 2016. Data-Oriented Programming: On the Expressiveness of Non-control Data Attacks. In *Proc. IEEE S&P '16*. 969–986.
- [7] Hans Liljestrand et al. 2019. Authenticated Call Stack. In *Proc. ACM/EDAC/IEEE DAC'19*. Article 223.
- [8] Hans Liljestrand et al. 2019. PAC it up: towards pointer integrity using ARM pointer authentication. In *Proc. USENIX Security '19*. 177–194.
- [9] David Litchfield. 2003. Defeating the stack based buffer overflow prevention mechanism of Microsoft Windows 2003 Server. In *Black Hat Asia '03*.
- [10] Hector Marco-Gisbert and Ismael Ripoll. 2013. Preventing brute force attacks against stack canary protection on networking servers. In *Proc. IEEE NCA '13*. 243–250.
- [11] Nick Nikiforakis et al. 2013. HeapSentry: kernel-assisted protection against heap overflows. In *Proc. DIMVA 13'*. 177–196.
- [12] Elias Levy (Aleph One). 1996. Smashing the stack for fun and profit. *Phrack* 7, 19 (1996). <http://phrack.org/issues/49/14.html>
- [13] Theofilos Petsios et al. 2015. Dynaguard: Armoring canary-based protections against brute-force attacks. In *Proc. ACM ACSAC '15*. 351–360.
- [14] Qualcomm. 2017. Pointer Authentication on ARMv8.3. <https://www.qualcomm.com/media/documents/files/whitepaper-pointer-authentication-on-armv8-3.pdf>.
- [15] William K Robertson et al. 2003. Run-time detection of heap-based overflows. In *Proc. USENIX LISA '03*. 51–60.
- [16] Hovav Shacham. 2007. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proc. ACM CCS '07*. 552–561.
- [17] Solar Designer. 1997. lpr LIBC RETURN exploit. <http://insecure.org/spl0its/linux.libc.return.lpr.sploit.html>
- [18] Raoul Strackx et al. 2009. Breaking the memory secrecy assumption. In *Proc. ACM EuroSec '09*. 1–8.
- [19] László Szekeres et al. 2013. SoK: Eternal war in memory. In *Proc. IEEE S&P '13*. 48–62.
- [20] Zhilong Wang et al. 2018. To detect stack buffer overflow with polymorphic canaries. In *Proc. IEEE/IFIP DSN '18*. IEEE, 243–254.

⁵<https://llvm.org/docs/StackSafetyAnalysis.html>