Hardening the Kernel Against Unprivileged Attacks

Claudio Canella Graz University of Technology

November 9, 2022























A Systematic Evaluation of Transient Execution Attacks and Defenses

Claudio Canella Graz University of Technology

Moritz Lipp Graz University of Technology

Frank Piessens imec-DistriNet, KU Leuven Jo Van Bulck

Benjamin von Berg Graz University of Technology

Dmitry Evtyushkin College of William and Mary Michael Schwarz Graz University of Technology

Philipp Ortner Graz University of Technology

Daniel Gruss Graz University of Technology



• out-of-order execution





- out-of-order execution
- speculative execution





- out-of-order execution
- speculative execution

What if CPU was wrong?



- out-of-order execution
- speculative execution

What if CPU was wrong?

 \rightarrow Roll back to mistake



- out-of-order execution
- speculative execution

What if CPU was wrong?

 $\rightarrow~\text{Roll back}$ to mistake

Squashed instructions are called transient

Transient-Execution Attacks





Claudio Canella





Transient

cause?







Claudio Canella



Claudio Canella

Spectre: Defense Analysis

	ing the fing											on cking Jucti, BB							
	Defense Attack	InvisiSno	SafeSne	DAMG	RSB Stud	Retpolin	Poison Iz	Index M.	Site Isol.	SLH	YSNB	IBRS	STIPB	IBP _B	Serializza	Taint T	Timer P.	Sloth	SSBD/SS
Intel	Spectre-PHT				\diamond	\diamond	•	O	0	•	0	\diamond	\diamond	\diamond	O		0		\diamond
	Spectre-BTB				\diamond	٠	\diamond	\diamond	0	\diamond	\diamond	٠	O	0	\diamond		O	\diamond	\diamond
	Spectre-RSB				O	\diamond	\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		•	\diamond	\diamond
	Spectre-STL				\diamond	\diamond	\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		O		•
ARM	Spectre-PHT				\diamond	\diamond	٠	0	0	٠	0	\diamond	\diamond	\diamond	O		0		\diamond
	Spectre-BTB				\diamond		\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		O	\diamond	\diamond
	Spectre-RSB				O	\diamond	\diamond	\diamond	O	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		\bullet	\diamond	\diamond
	Spectre-STL				\diamond	\diamond	\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		0		•
AMD	Spectre-PHT				\diamond	\diamond		0	0		0	\diamond	\diamond	\diamond	O		O		\diamond
	Spectre-BTB				\diamond		\diamond	\diamond	0	\diamond	\diamond				\diamond		0	\diamond	\diamond
	Spectre-RSB				●	\diamond	\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond		\diamond		O	\diamond	\diamond
	Spectre-STL				\diamond	\diamond	\diamond	\diamond	0	\diamond	\diamond	\diamond	\diamond	\diamond	\diamond		0		•

Symbols show if an attack is mitigated (●), partially mitigated (●), not mitigated (○), theoretically mitigated (■), theoretically impeded (■), not theoretically impeded (□), or out of scope (◇).

More Details

More details in the paper

- New Meltdown variants
- Spectre mistraining strategies
- Defense categorization and performance analysis
- Gadget classification and prevalence analysis
- ...



Claudio Canella, Jo Van Bulck, Michael Schwarz, Moritz Lipp, Benjamin von Berg, Philipp Ortner, Frank Piessens, Dmitry Evtyushkin, Daniel Gruss. A Systematic Evaluation of Transient Execution Attacks and Defenses.



KASLR: Break It, Fix It, Repeat

Claudio Canella Graz University of Technology

Martin Schwarzl Graz University of Technology Michael Schwarz Graz University of Technology

Daniel Gruss Graz University of Technology Martin Haubenwallner Graz University of Technology

Hypothesis

Load is executed, returned value is zeroed out

Hypothesis

Load is executed, returned value is zeroed out



Meltdown Attack

Hypothesis

Load is executed, returned value is zeroed out



Meltdown Attack



Performance Counter



user kernel not present



Claudio Canella










mem[*0xffffffff80000000]















	code & data	
\uparrow		$\widehat{}$
0xffff ffff 8000 0000		Oxffff ffff bfff ffff

	code & data	
\uparrow		\uparrow
0xffff ffff 8000 0000		Oxffff ffff bfff ffff





More Details

More details in the paper

- EchoLoad from and on SGX
- Meltdown in JavaScript on 32-bit Systems
- FLARE for Kernel modules, vmalloc, ...
- ...



Claudio Canella, Michael Schwarz, Martin Haubenwallner, Martin Schwarzl, Daniel Gruss. KASLR: Break It, Fix It, Repeat.





The "seccomp" system provides an opt-in feature made available to userspace

The "seccomp" system provides an opt-in feature made available to userspace, which provides a way to reduce the number of kernel entry points available to a running process.

The "seccomp" system provides an opt-in feature made available to userspace, which provides a way to reduce the number of kernel entry points available to a running process. This limits the breadth of kernel code that can be reached

The "seccomp" system provides an opt-in feature made available to userspace, which provides a way to reduce the number of kernel entry points available to a running process. This limits the breadth of kernel code that can be reached, possibly reducing the availability of a given bug to an attack.



Effort





Effort

тостои







Effort

тостои

Stateless







Effort

ТОСТОИ

Stateless

Claudio Canella

16

Automating Seccomp Filter Generation for Linux Applications

Claudio Canella Graz University of Technology

Daniel Gruss Graz University of Technology Mario Werner Graz University of Technology

Michael Schwarz CISPA Helmholtz Center for Information Security

P1: Static Analysis



















More Details

More details in the paper

- Implementation details
- Information on overapproximation
- Detailed evaluation
- ...



Claudio Canella, Mario Werner, Daniel Gruss, Michael Schwarz. Automating Seccomp Filter Generation for Linux Applications.



SFIP: Coarse-Grained Syscall-Flow-Integrity Protection in Modern Systems

Claudio Canella Graz University of Technology

Daniel Gruss Graz University of Technology Sebastian Dorn Graz University of Technology

Michael Schwarz CISPA Helmholtz Center for Information Security

Syscall-Flow-Integrity Protection


Syscall-Flow-Integrity Protection



Syscall-Flow-Integrity Protection



Syscall-Flow-Integrity Protection



Claudio Canella

State Machine Analysis

Application	#S tates	Average Transitions
busybox	23.52	15.99
coreutils	26.64	16.66
pwgen	18	13.56
muraster	29	18.89
nginx	107	74.05
ffmpeg	55	49.07
memcached	86	43.16
mutool	53	32.26

State Machine Analysis

Application	#States	Average Transitions
busybox	23.52	15.99
coreutils	26.64	16.66
pwgen	18	13.56
muraster	29	18.89
nginx	107	74.05
ffmpeg	55	49.07
memcached	86	43.16
mutool	53	32.26

Origin Analysis

Application	Total #Offsets	Avg #Offsets
busybox	102.64	3.75
coreutils	116.71	4.42
pwgen	84	4.42
muraster	193	4.6
nginx	318	3.0
ffmpeg	279	4.98
memcached	317	3.69
mutool	278	4.15

Origin Analysis

Application	Total #Offsets	Avg #Offsets
busybox	102.64	3.75
coreutils	116.71	4.42
pwgen	84	4.42
muraster	193	4.6
nginx	318	3.0
ffmpeg	279	4.98
memcached	317	3.69
mutool	278	4.15



Return-Oriented Programming

• uses exisiting code to exploit a program



- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)



- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)
- chains gadgets together for an exploit



- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)
- chains gadgets together for an exploit
- SFIP restricts ROP chains via



- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)
- chains gadgets together for an exploit
- SFIP restricts ROP chains via
 - \bullet syscall transitions \rightarrow not every sequence is possible



Return-Oriented Programming

- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)
- chains gadgets together for an exploit

SFIP restricts ROP chains via

- \bullet syscall transitions \rightarrow not every sequence is possible
- \bullet syscall origins \rightarrow not every location is valid



Return-Oriented Programming

- uses exisiting code to exploit a program
- jumps to parts of functions (gadgets)
- chains gadgets together for an exploit

SFIP restricts ROP chains via

- \bullet syscall transitions \rightarrow not every sequence is possible
- \bullet syscall origins \rightarrow not every location is valid

Conclusion

SFIP imposes more significant constraints on control-flow-hijacking attacks than seccomp



More Details

More details in the paper

- Implementation details
- Extensive security discussion
- Mimicry attacks
- ...



Claudio Canella, Sebastian Dorn, Daniel Gruss, Michael Schwarz. SFIP: Coarse-Grained Syscall-Flow-Integrity Protection in Modern Systems.



• Deepened understanding of transient-execution attacks





- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow\,$ Proposed widely-adopted naming scheme



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow\,$ Proposed widely-adopted naming scheme
 - $\rightarrow~$ Highlighted insufficient defenses



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow\,$ Proposed widely-adopted naming scheme
 - $\rightarrow~$ Highlighted insufficient defenses
- Hardened kernel against microarchitectural KASLR breaks



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow\,$ Proposed widely-adopted naming scheme
 - \rightarrow Highlighted insufficient defenses
- Hardened kernel against microarchitectural KASLR breaks
- Reduced the exposed attack surface of the kernel



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow~$ Proposed widely-adopted naming scheme
 - $\rightarrow~$ Highlighted insufficient defenses
- Hardened kernel against microarchitectural KASLR breaks
- Reduced the exposed attack surface of the kernel
 - $\rightarrow~$ Automated Seccomp



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow~$ Proposed widely-adopted naming scheme
 - $\rightarrow~$ Highlighted insufficient defenses
- Hardened kernel against microarchitectural KASLR breaks
- Reduced the exposed attack surface of the kernel
 - $\rightarrow~$ Automated Seccomp
 - $\rightarrow\,$ Syscall-Flow-Integrity Protection



- Deepened understanding of transient-execution attacks
 - $\rightarrow~$ Found new attack variants
 - $\rightarrow~$ Proposed widely-adopted naming scheme
 - \rightarrow Highlighted insufficient defenses
- Hardened kernel against microarchitectural KASLR breaks
- Reduced the exposed attack surface of the kernel
 - $\rightarrow~$ Automated Seccomp
 - $\rightarrow\,$ Syscall-Flow-Integrity Protection
 - $\rightarrow~\mbox{Enabled}$ complex argument checks



My PhD in Numbers































Hardening the Kernel Against Unprivileged Attacks

Claudio Canella Graz University of Technology

November 9, 2022

References

- [Can+19a] C. Canella, D. Genkin, L. Giner, D. Gruss, M. Lipp, M. Minkin, D. Moghimi,
 F. Piessens, M. Schwarz, B. Sunar, J. Van Bulck, and Y. Yarom. Fallout: Leaking Data on Meltdown-resistant CPUs. In: CCS. 2019.
- [Can+19b] C. Canella, J. Van Bulck, M. Schwarz, M. Lipp, B. von Berg, P. Ortner, F. Piessens, D. Evtyushkin, and D. Gruss. A Systematic Evaluation of Transient Execution Attacks and Defenses. In: USENIX Security Symposium. Extended classification tree and PoCs at https://transient.fail/. 2019.
- [Can+20] C. Canella, M. Schwarz, M. Haubenwallner, M. Schwarzl, and D. Gruss. KASLR: Break It, Fix It, Repeat. In: AsiaCCS. 2020.
- [Can+21a] C. Canella, A. Kogler, L. Giner, D. Gruss, and M. Schwarz. Domain Page-Table Isolation. In: arXiv:2111.10876 (2021).

- [Can+21b] C. Canella, M. Werner, D. Gruss, and M. Schwarz. Automating Seccomp Filter Generation for Linux Applications. In: CCSW. 2021.
- [Can+22] C. Canella, S. Dorn, D. Gruss, and M. Schwarz. SFIP: Coarse-Grained Syscall-Flow-Integrity Protection in Modern Systems. In: arXiv:2202.13716 (2022).
- [Gru+16] D. Gruss, C. Maurice, A. Fogh, M. Lipp, and S. Mangard. Prefetch Side-Channel Attacks: Bypassing SMAP and Kernel ASLR. In: CCS. 2016.
- [HWH13] R. Hund, C. Willems, and T. Holz. Practical Timing Side Channel Attacks against Kernel Space ASLR. In: S&P. 2013.
- [JLK16] Y. Jang, S. Lee, and T. Kim. Breaking Kernel Address Space Layout Randomization with Intel TSX. In: CCS. 2016.
- [Sch+19] M. Schwarz, C. Canella, L. Giner, and D. Gruss. Store-to-Leak Forwarding: Leaking Data on Meltdown-resistant CPUs. In: arXiv:1905.05725 (2019).