Spectre-v2 and retpoline on Skylake

Red Hat, Inc.

Andrea Arcangeli <aarcange at redhat.com>

Microarchitecture Workshop
Boston University

20 Feb 2019

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

call

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

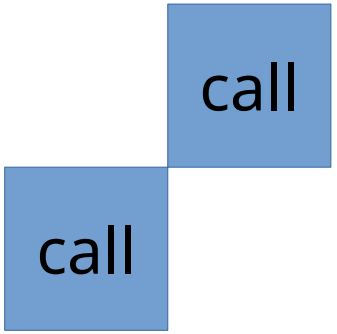

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

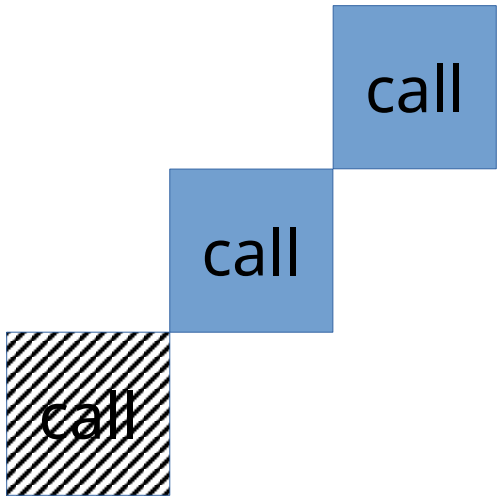

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

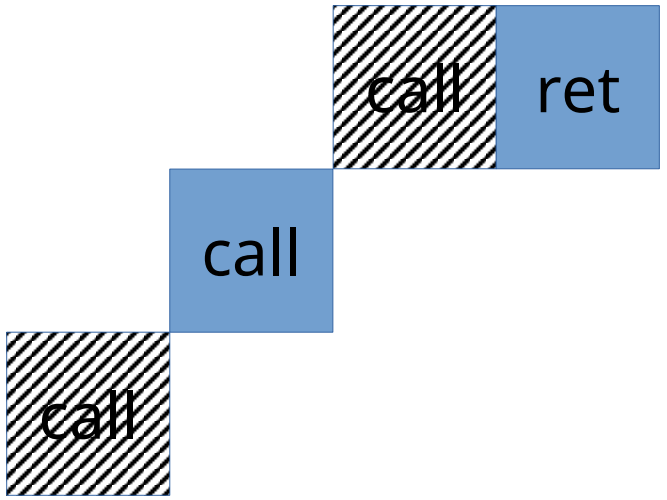

redhat.

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

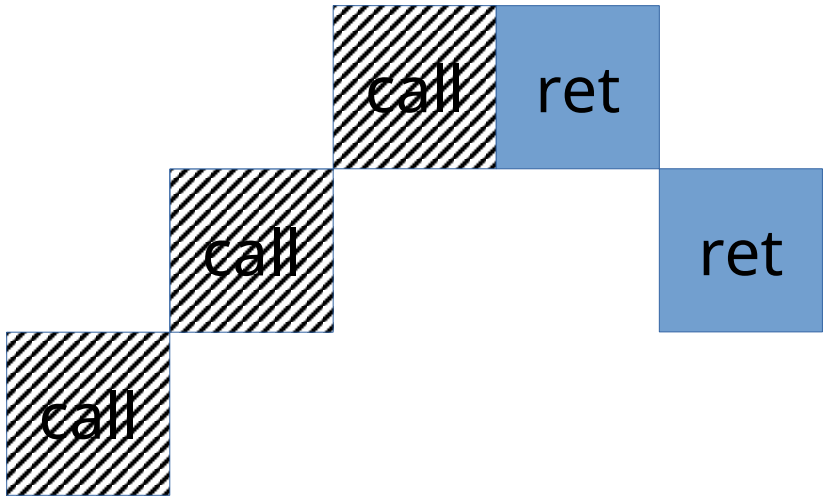

“ret” pulling from BTB

- In some recent CPU the “ret” instruction will pull the *“poison”* in the BTB if the RSB stack is empty (i.e. RSB underflow). Specs chatper 5.2
- Assume a stack of 2 entries:

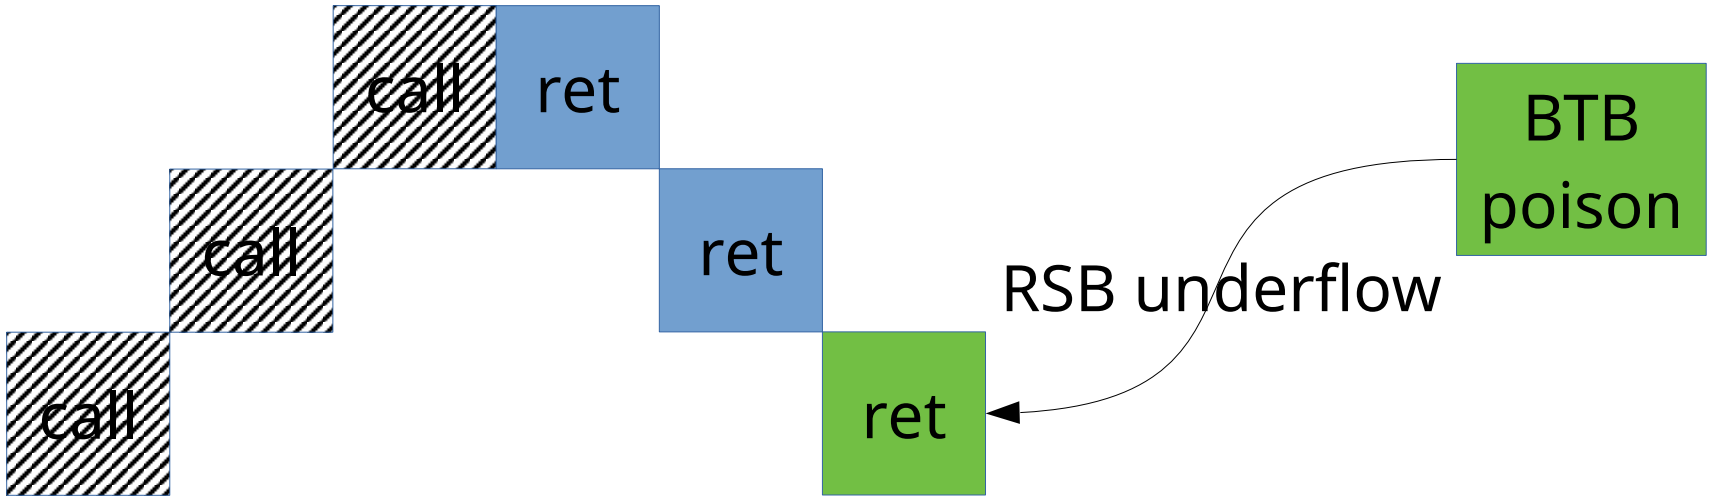

Other way the RSB may be drained

- IBPB
- Various SGX and microcode update WRMSR
- Imbalance between CALL instructions and RET instructions
  - Context switch
  - Longjmp
- MWAIT C6 sleep

FILL_RETURN_BUFFER mitigation

- Coverage:
  - Context switch
  - VMEXIT
- Lack of coverage:
  - IBPB
    - Not an issue because the BTB is flushed too
  - interrupts/nmi
  - MWAIT
  - **Natural underflow after >16-32 rets**

Question:
is the following possible at all?

1) Find a syscall or another kernel entry point that causes a RSB underflow in a “ret” instruction at a fixed kernel address where userland can still control some register content

2) Train the BTB at an alias of the possible randomized fixed kernel addresses to execute a ROP spectre-v1 gadget through the RSB-underflowing “ret”

3) Use the gadget to first derandomize the fixed kernel address of the RSB-underflowing “ret” instruction and then to circumvent the spectre-v2 retpoline mitigation on Skylake

Q/A

redhat.