redhat.

Meltdown / Spectre

Patrick Ladd

Technical Account Manager

January 10th 2018

For your manager...

https://www.youtube.com/watch?v=syAdX44pokE

```
meltdown@meltdown: ./meltdown
e01d8110: 61 78 20 6f 72 20 73 74  61 74 65 20 6d 61 63 68 |ax or state mach|
e01d8120: 69 6e 65 2c 20 69 74 20  69 73 20 62 65 69 6e 67 |ine, it is being|
e01d8130: 20 75 73 65 64 20 77 69  74 68 20 61 75 74 68 6f | used with autho|
```

https://youtu.be/bReA1dvGJ6Y

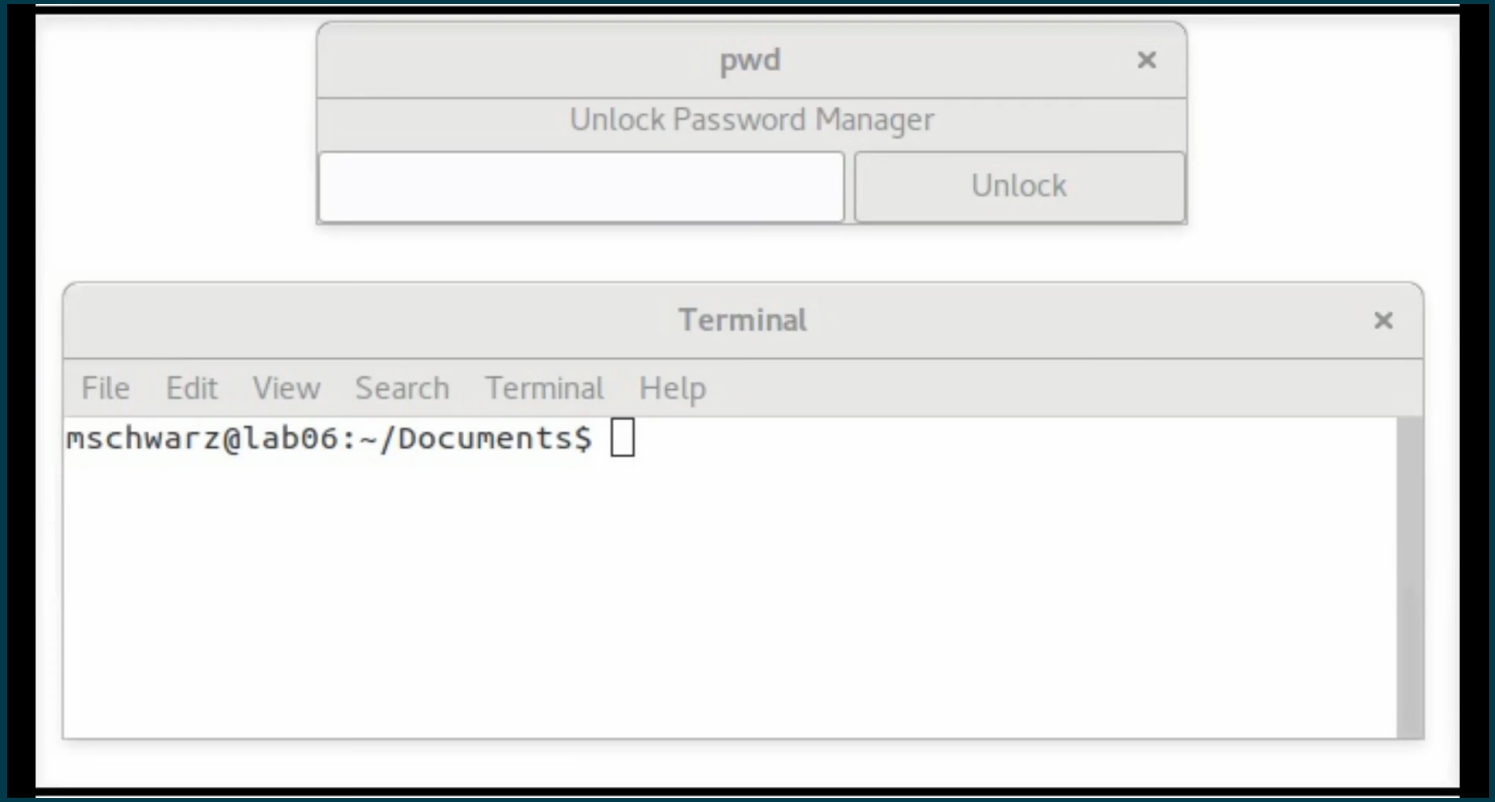

https://youtu.be/RbHbFkh6eeE

History / Timeline

Technical Details

Speculative Execution

Side-Channel Attacks

Meltdown

Exploit Code

```
raise_exception();
// the line below is never reached
access(probe_array[data * 4096]);
```

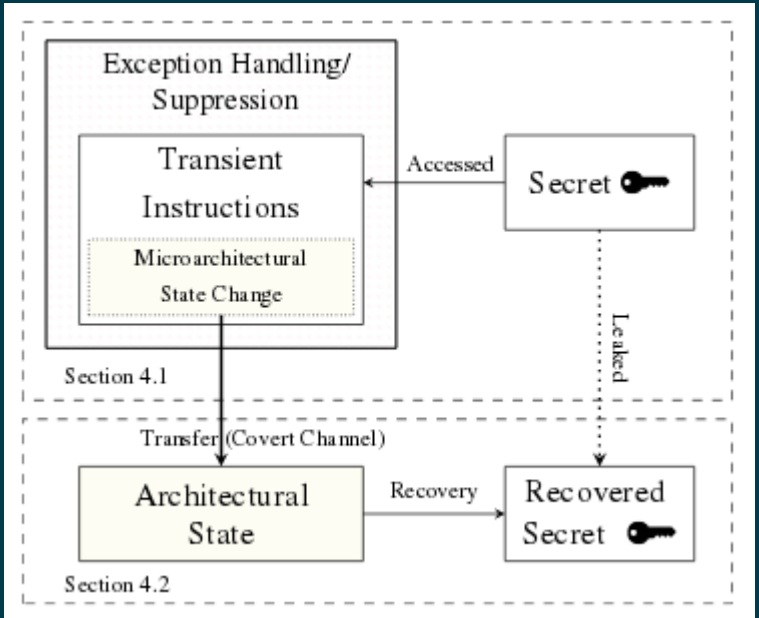

Spectre

Exploit Code

Conditional Branch Exploit

```
if (x < array1_size)
    y = array2[array1[x] * 256];
```

Indirect Branch Exploit

Other Variants

Patching

Patching Meltdown

Kernel Patch

Patching Spectre

Variant #1 – Bounds Check Violations

- Kernel Patch

Variant #2 – Branch Predictor

- Kernel Patch & CPU Microcode Updates

Impact

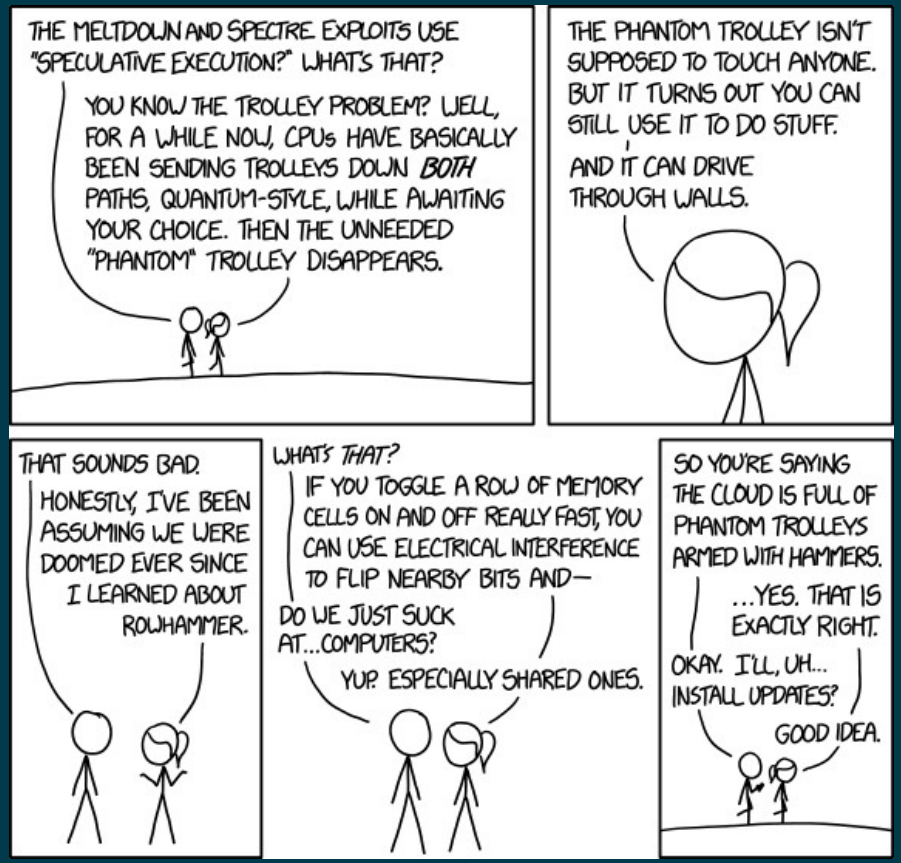
THE MELTDOWN AND SPECTRE EXPLOITS USE "SPECULATIVE EXECUTION?" WHAT'S THAT?
YOU KNOW THE TROLLEY PROBLEM? WELL, FOR A WHILE NOW, CPUs HAVE BASICALLY BEEN SENDING TROLLEYS DOWN BOTH PATHS, QUANTUM-STYLE, WHILE AWAITING YOUR CHOICE. THEN THE UNNEEDED "PHANTOM" TROLLEY DISAPPEARS.
THE PHANTOM TROLLEY ISN'T SUPPOSED TO TOUCH ANYONE. BUT IT TURNS OUT YOU CAN STILL USE IT TO DO STUFF.
AND IT CAN DRIVE THROUGH WALLS.
THAT SOUNDS BAD.
HONESTLY, I'VE BEEN ASSUMING WE WERE DOOMED EVER SINCE I LEARNED ABOUT ROWHAMMER.
WHAT'S THAT?
IF YOU TOGGLE A ROW OF MEMORY CELLS ON AND OFF REALLY FAST, YOU CAN USE ELECTRICAL INTERFERENCE TO FLIP NEARBY BITS AND—
DO WE JUST SUCK AT...COMPUTERS?
YUP. ESPECIALLY SHARED ONES.
SO YOU'RE SAYING THE CLOUD IS FULL OF PHANTOM TROLLEYS ARMED WITH HAMMERS.
...YES. THAT IS EXACTLY RIGHT.
OKAY. I'LL, UH... INSTALL UPDATES?
GOOD IDEA.

Performance Impact

Measurable: 8-19%

Highly cached random memory, with buffered I/O, OLTP database workloads, and benchmarks with high kernel-to-user space transitions are impacted between 8-19%.

- OLTP Workloads (tpc)
- sysbench
- pgbench
- netperf (< 256 byte)
- fio (random I/O to NvME).

Performance Impact

Modest: 3-7%

Database analytics, Decision Support System (DSS), and Java VMs are impacted less than the “Measurable” category. These applications may have significant sequential disk or network traffic, but kernel/device drivers are able to aggregate requests to moderate level of kernel-to-user transitions

- SPECjbb2005
- Queries/Hour
- Overall analytic timing (sec)

Performance Impact

Small: 2-5%

HPC (High Performance Computing) CPU-intensive workloads are affected the least with only 2-5% performance impact because jobs run mostly in user space and are scheduled using cpu-pinning or numa-control.

- Linpack NxN on x86
- SPECcpu2006

Performance Impact

Minimal: <2%

Linux accelerator technologies that generally bypass the kernel in favor of user direct access are the least affected

- DPDK (VsPERF at 64 byte)
- OpenOnload (STAC-N)
- Userspace accesses to VDSO like get-time-of-day are not impacted

Tuning

Kernel Tuning Parameters

noibrs / ibrs_enabled

Indirect Branch Restricted Speculation (ibrs)

CVE-2017-5715
Variant #2 / Spectre

noibpb / ibpb_enabled

Indirect Branch Prediction Barriers (ibpb)

CVE-2017-5715
Variant #2 / Spectre

nopti / pti_enabled

Kernel Page Table Isolation (pti)

CVE-2017-5754
Variant #3 / Meltdown

Architectural Defaults

Automatically Depending on Detected Architecture

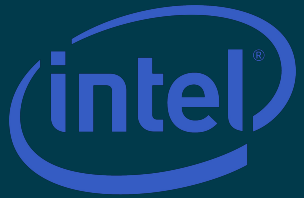

Variants 1,2,3 Enabled

`pti 1 ibrs 1 ibpb 1` ->fix variant#1,2,3

`pti 1 ibrs 0 ibpb 0` -> fix variant#1 #3 (for older Intel systems with no microcode update available)

Variants 1,2 Enabled
Not Vulnerable to #3

`pti 0 ibrs 0 ibpb 2` -> fix variant #1 #2 if the microcode update is applied

`pti 0 ibrs 2 ibpb 1` -> fix variant #1 #2 on older processors that can disable indirect branch prediction without microcode updates

What Do I Do Now?

Resources

Resources

Red Hat Customer Portal

Master Vulnerability Page
https://access.redhat.com/security/vulnerabilities/speculativeexecution

Performance Impact
https://access.redhat.com/articles/3307751

Tunables
https://access.redhat.com/articles/3311301

Resources

Red Hat

What you need to know – Jon Masters

https://www.redhat.com/en/blog/what-are-meltdown-and-spectre-here%E2%80%99s-what-you-need-know

Q&A Webinar - *Thursday, January 11 at 11:00AM EST*

https://onlinexperiences.com/Launch/QReg/ShowKey=47447&AffiliateData=701f2000000tsoPAAQ&

Resources

External

Meltdown Paper
https://meltdownattack.com/meltdown.pdf

Spectre Paper
https://spectreattack.com/spectre.pdf

Google Project Zero
https://googleprojectzero.blogspot.ca/2018/01/reading-privileged-memory-with-side.html

Resources

External

Meltdown Exploit POC Code
https://github.com/IAIK/meltdown

THANK YOU

plus.google.com/+RedHat

facebook.com/redhatinc

linkedin.com/company/red-hat

twitter.com/RedHatNews

youtube.com/user/RedHatVideos