

## Linux Meltdown patches lead to high performance overheads - Operating Systems - Security

by Admin - Tuesday, February 13, 2018

<http://itsupportnetwork.com/linux-meltdown-patches-lead-to-high-performance-overheads-operating-systems-security/>

(adsbygoogle = window.adsbygoogle || []).push({});

### Netflix engineer estimates 2% slowdown from Linux patches for Meltdown flaw

The latest Meltdown patches for Linux machines could impact performance by up to 800% in some circumstances, an assessment by a famous software engineer has revealed.

Netflix engineer and computing performance expert, Brendan Gregg - also known as dTrace Daddy - developed a "microbenchmark" to assess the Linux kernel page table isolation (KPTI) patch for the Meltdown CPU design flaw. The benchmark revealed that the patch will significantly increase overheads.

"The patches that workaround Meltdown introduce the largest kernel performance regressions I've ever seen," Gregg said in a blog post.

To understand the KPTI overhead, Gregg said there are at least five factors at play, which start with the syscall rate. When the rate of calls to the system reaches 50,000 per second per processor, the performance overhead hits 2%, and climbs as the syscall rate increases.

Next are context switches - a 'save point' for the current progress state of a process to restart it from that point later. These add similar overheads to the syscall rate, while page fault rates add "a little more" of a performance overhead as well.

Large amounts of 'hot' data - more than 10MB of data that's accessed frequently - create further performance issues due to the amount of time it takes to flush the memory cache, turning a 1% overhead created by syscall cycles into a 7% overhead, Gregg said. "This overhead can be reduced by A) pcid, available in Linux 4.14, and B) Huge pages," he added.

Lastly, certain cache access patterns can add another 10% overhead, he explained.

To explore these variants, Gregg wrote a simple microbenchmark where he could vary the syscall rate and the working set size.

He then analysed performance during the benchmark, and used other benchmarks to confirm findings. He discovered that the variables to watch out for are the syscall numbers, the Linux kernel being used and if

it supports process-context identifiers, and the size of the pages being used - using larger pages so you have fewer pages to track.

Gregg also looked into the impact of the patches on the AWS infrastructure used by his employer, Netflix. He said: "Practically, I'm expecting the cloud systems at my employer (Netflix) to experience between 0.1% and 6% overhead with KPTI due to our syscall rates, and I'm expecting we'll take that down to less than 2% with tuning".

The worst case Gregg found was an 800% overhead, saying that people can expect an impact of anywhere between 1% and 800%.

"Where you are on that spectrum depends on your syscall and page fault rates, due to the extra CPU cycle overheads, and your memory working set size, due to TLB flushing on syscalls and context switches," he said, but concluded that through tuning some of the systems, it's possible to "significantly reduce the overheads".

Linux creator Linus Torvalds has previously lambasted Intel's approach to fixing the Meltdown and Spectre flaws.

He criticised Intel's plan to continue shipping affected chips with an optional patch to apply, as well as the performance issues that resulted from some initial Intel patches. Intel last week released fresh patches that it says do not impact performance.

```
(adsbygoogle = window.adsbygoogle || []).push({});
```

[Source link](#)

---

[IT Support Network](#)