Addendum 1 to
RIDL: Rogue In-flight Data Load

Stephan van Schaik*, Alyssa Milburn*, Sebastian Österlund*, Pietro Frigo*, Giorgi Maisuradze†‡,
Kaveh Razavi*, Herbert Bos*, and Cristiano Giuffrida*

*Department of Computer Science
Vrije Universiteit Amsterdam, The Netherlands
{s.j.r.van.schaik, a.a.milburn, s.osterlund, p.frigo}@vu.nl,
{kaveh, herbertb, giuffrida}@cs.vu.nl

†CISPA Helmholtz Center for Information Security
Saarland Informatics Campus
giorgi.maisuradze@cispa.saarland

***Abstract*—On Nov 12, 2019, we disclose *TSX Asynchronous Abort* (TAA), a "new" speculation-based vulnerability in Intel CPUs as well as other MDS-related issues. In reality, this is no new vulnerability. We disclosed TAA (and other issues) as part of our original RIDL submission to Intel in Sep 2018. Unfortunately, the Intel PSIRT team missed our submitted proof-of-concept exploits (PoCs), and as a result, the original MDS mitigations released in May 2019 only partially addressed RIDL.**

At the request of Intel, and to protect their users, we redacted parts of the original RIDL paper and did not release the RIDL test suite with our PoCs on the MDS disclosure date (May 14, 2019). This addendum provides an analysis of Intel's original (flawed) MDS mitigation and an explanation for the "Misaligned Read" and the "TSX" columns in Table I, which we redacted from the original RIDL paper. Additional updated information on RIDL, TAA, the disclosure process, our now public test suite and TAA-optimized exploits can be found at https://mdsattacks.com.

A. *Flawed MDS mitigation*

Intel's original microcode update, which modifies the VERW instruction to clear CPU buffers and mitigate MDS, is flawed in that it clears the buffers using stale (potentially sensitive) data on several of the CPUs we used for testing (e.g., i7-7700K). Intel states this bug is only present in Skylake client CPUs. This means that data can be leaked across privilege boundaries using RIDL even if SMT has been disabled and the recommended VERW mitigation has been applied.

The non-microcode versions of the mitigation provided in the MDS whitepaper[1] appear to correctly clear the CPU buffers, but at a much higher performance cost. Our RIDL paper originally reported the intended behavior of Intel's mitigation. Unfortunately, at Intel's request, we had to withhold any comment on the flawed mitigation from the paper, in order to comply with the second embargo. The new microcode recently released by Intel still does *not* fix the issue, as we still see leaks with RIDL PoCs shared with Intel in May.

B. *TSX Asynchronous Abort*

TSX transactions can be aborted by for instance flushing a cache line before the transaction, then loading from the same cache line inside the transaction. This causes the processor to abort the transaction despite execution of instructions in the pipeline continuing until retirement, allowing information to be leaked via the load from various internal CPU buffers—including store buffers—using RIDL. Intel refers to this RIDL variant as the *TSX Asynchronous Abort* (TAA) vulnerability.

This vulnerability is present even on CPUs which Intel claims are not vulnerable to MDS, such as recent Cascade Lake CPUs. Although no microcode MDS mitigations were available on these CPUs when MDS was disclosed in May, Intel recently (September 2019) provided microcode updates. We believe this vulnerability can be mitigated by disabling TSX. Our original RIDL paper reported results for TAA for a variety of CPUs in the "TSX" column in Table 1; we withheld the explanation at Intel's request to comply with the second embargo.

C. *Alignment faults*

Alignment faults (e.g., due to the AC flag or aligned vector instructions) can be used to cause exceptions and perform RIDL attacks. Although this vulnerability is not mitigated by the silicon fixes for Meltdown/MFBDS (page faults), it appears to be mitigated on Intel's latest CPUs. Alignment faults and split loads across cache lines can be used to leak data from a variety of sources, including load ports (as originally reported by Intel) and (indirectly) also store and fill buffers.

Our original RIDL paper reported results for both split loads and alignment faults for a variety of CPUs in the "Misaligned read" column in Table 1. Such results showcased leaks that were not explained by Intel's original MDS whitepaper (e.g., store-to-load leaks). Since then, Intel's whitepaper has undergone a number of updates. Again, at Intel's request, we withheld a full explanation of our results from the paper.

D. *Conclusion*

This research—whose details were withheld from the public version of the RIDL paper due to responsible disclosure considerations—further supports the arguments presented in our original paper. As demonstrated by the TAA vulnerability (still present in recent Intel CPUs) and the flawed MDS mitigation, RIDL-class vulnerabilities are non-trivial to fix or mitigate, and current "spot" mitigation strategies for resolving these issues are questionable. Moreover, we question the effectiveness of year-long disclosure processes and also raise concerns on their disruptive impact on the academic process. We continue to work with Intel to improve their coordinated disclosure process and collaboration with academia.

[1] https://software.intel.com/security-software-guidance/insights/deep-dive-intel-analysis-microarchitectural-data-sampling