

# Writing a fine-grained access pattern oriented lightweight kernel module using DAMON/DAMOS in 10 minutes

Thursday, September 23, 2021 8:00 AM (45 minutes)

## DAMON and DAMOS

DAMON[1] is a framework for general data access monitoring of kernel subsystems. It provides best-effort high quality monitoring results while incurring only minimal and upper-bounded overhead, due to its practical overhead-accuracy tradeoff mechanism. On a production machine utilizing 70 GB memory, it can repeatedly scan accesses to the whole memory for every 5ms, while consuming only 1% single CPU time.

On top of it, a data access pattern-oriented memory management engine called DAMON-based Operation Schemes (DAMOS) is implemented. It allows clients to implement their access pattern oriented memory management logic with very simple scheme descriptions. We implemented fine-grained access-aware THP and proactive reclamation using this engine in three lines of scheme and achieved remarkable improvements[2].

As of this writing (2021-05-28), the code is not in the mainline but available at its development tree[3], and regularly posted to LKML as patchsets[4,5,6]. Nevertheless, the code has already merged in the public Amazon Linux kernel trees[7,8], and all Amazon Linux users can use DAMON/DAMOS off the box. We are also supporting the two latest upstream LTS stable kernels[9,10].

## Agenda

In this talk, I will briefly introduce DAMON/DAMOS and present how you can write a fine-grained data access pattern oriented lightweight kernel module on top of DAMON/DAMOS. With the talk, I will write an example module and evaluate its performance on live. A data access-aware proactive reclamation kernel module for production use will also introduced as a use case. After that, I will discuss my future plans for improving DAMON and improving other kernel subsystems using DAMON/DAMOS.

[1] <https://damonitor.github.io> (<https://damonitor.github.io/>)

[2] <https://damonitor.github.io/doc/html/latest/vm/damon/eval.html>

[3] <https://github.com/sjp38/linux/tree/damon/master> (<https://github.com/sjp38/linux/tree/damon/master>)

[4] <https://lore.kernel.org/linux-mm/20210520075629.4332-1-sj38.park@gmail.com/>

[5] <https://lore.kernel.org/linux-mm/20201216084404.23183-1-sjpark@amazon.com/>

[6] <https://lore.kernel.org/linux-mm/20201216094221.11898-1-sjpark@amazon.com/>

[7] <https://github.com/amazonlinux/linux/tree/amazon-5.4.y/master/mm/damon>

[8] <https://github.com/amazonlinux/linux/tree/amazon-5.10.y/master/mm/damon>

[9] <https://github.com/sjp38/linux/tree/damon/for-v5.4.y>

[10] <https://github.com/sjp38/linux/tree/damon/for-v5.10.y>

## I agree to abide by the anti-harassment policy

I agree

**Primary author:** PARK, SeongJae

**Presenter:** PARK, SeongJae

**Session Classification:** Kernel Summit

**Track Classification:** Kernel Summit