



Contribution ID: 96

Type: **not specified**

DAMON-based Pages Migration for {C,G,X}PU [un]attached NUMA nodes

Modern systems feature increasingly complex NUMA (Non-Uniform Memory Access) topologies, often with multiple nodes that may or may not be equipped with CPUs, GPUs, or other accelerators. This complexity makes it crucial to migrate memory pages efficiently based on access patterns.

DAMON, a Linux kernel subsystem, offers effective monitoring of system and workload data access patterns. It also includes DAMOS, a feature that enables system operations to be performed based on these monitoring results. Several companies and researchers, including SK Hynix, Micron, and Meta, have developed memory tiering and dynamic memory interleaving solutions for CXL memory utilizing DAMON. There are plans and potential to further extend DAMON to support a broader range of NUMA topologies.

This session will provide a concise overview of existing DAMON solutions for CXL memory. It will also outline plans for expanding DAMON's capabilities to support general NUMA topologies. Furthermore, the session aims to facilitate discussions with other core and device memory management system developers regarding any concerns or unaddressed requirements.

Primary author: PARK, SeongJae

Presenter: PARK, SeongJae

Session Classification: Device and Specific Purpose Memory MC

Track Classification: Device and Specific Purpose Memory MC