UMAP: One Approach to Exposing Emerging NVRAM To Applications

Adam Manzanares
Western Digital Research

Emerging NVRAM technologies are generally classified as either lower latency block storage devices or as higher latency and persistent replacements for DRAM. In this session, I would like to discuss the properties of current OS mechanisms capable of exposing NVRAM devices and how they give us a clue as to what device properties are desirable for the given mechanism. It is my belief that properties of existing software mechanisms together with device specifications give an indicator of how the NVRAM should be exposed to the application.

NVRAM can be exposed through the traditional block layer interface and presented to applications as a low latency block device. This is a well understood interface and there are limitations to the lowest latency achievable by the software mechanism itself. A second approach DAX, considers the device to be memory but, this memory is accessed like a device and not managed as memory by the OS. A third approach considers the NVRAM to be higher latency memory and is managed by the OS. The first and third approaches leverage existing mechanisms and will likely require further optimizations to leverage NVRAM properties. DAX is an interesting approach aimed at dealing with NVRAM, but it assumes that direct access is the optimal access path to the NVRAM.

I would like to introduce UMAP, which is a user directed, kernel managed virtual memory management framework. UMAP exposes NVRAM as memory to applications, while allowing the OS to treat the NVRAM as a device. In addition, UMAP has a pluggable caching and device layer, which can be tuned for their respective targets. UMAP can also leverage VMA based hinting to make decisions on how to expose the NVRAM that is backed by UMAP.