

The SNIA logo consists of a small square icon with a dot inside, followed by the letters 'SNIA' in a bold, sans-serif font. The background of the banner features a blue and teal color scheme with binary code (0s and 1s) and circuit-like patterns.

PERSISTENT MEMORY  
**PM** SUMMIT


JANUARY 18, 2017 | SAN JOSE, CA

# The SNIA NVM Programming Model: Latest Developments and Challenges



Andy Rudoff, Intel Corporation

## “Programming Model” – Four meanings (at least)

# Programming Model: SW Interface to HW

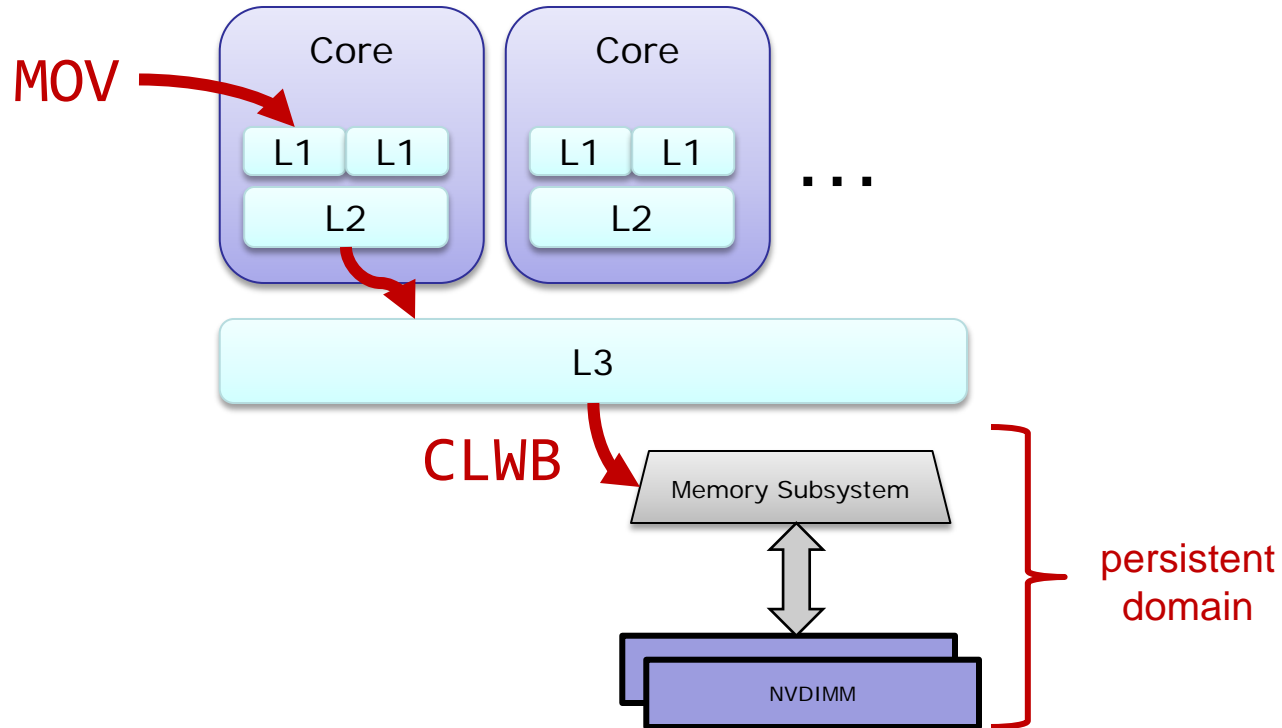


# Programming Model: SW Interface to HW

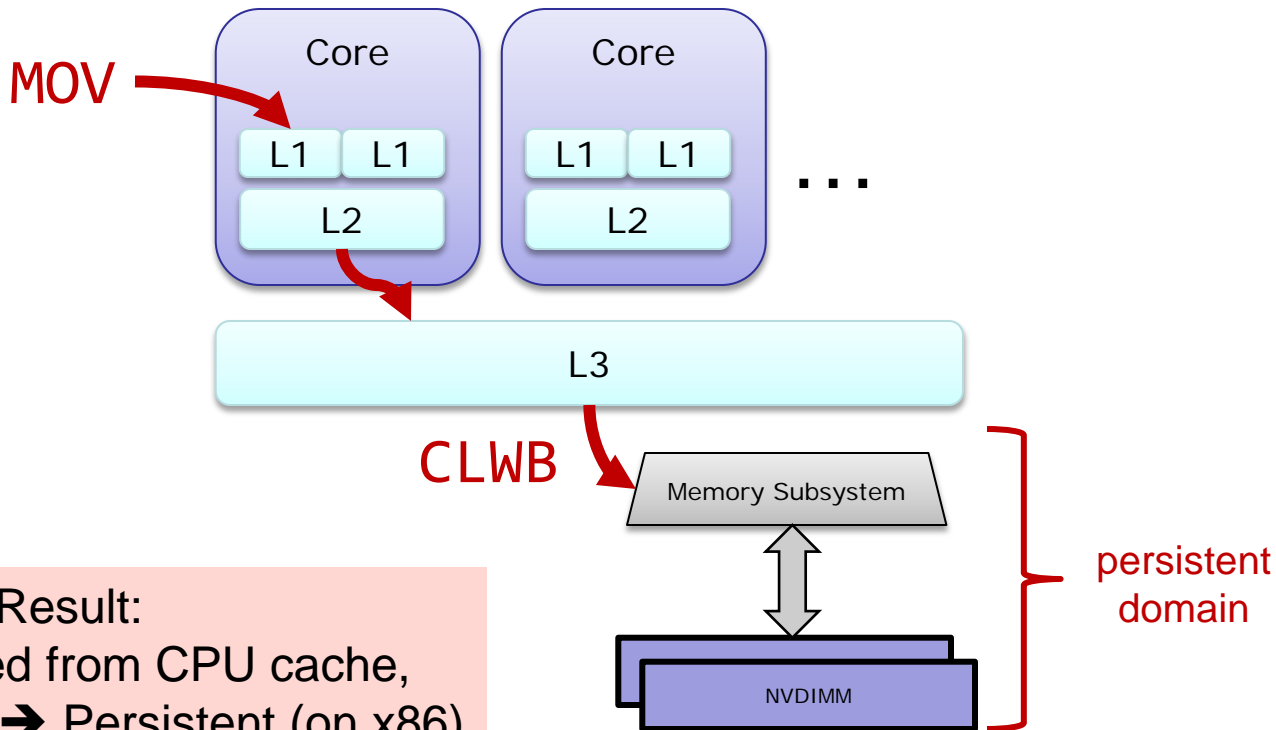


**Result:**  
Persistent Memory hardware  
accessed like memory  
(cache coherent).  
Described by ACPI on x86.

# Programming Model: Instruction Set Architecture

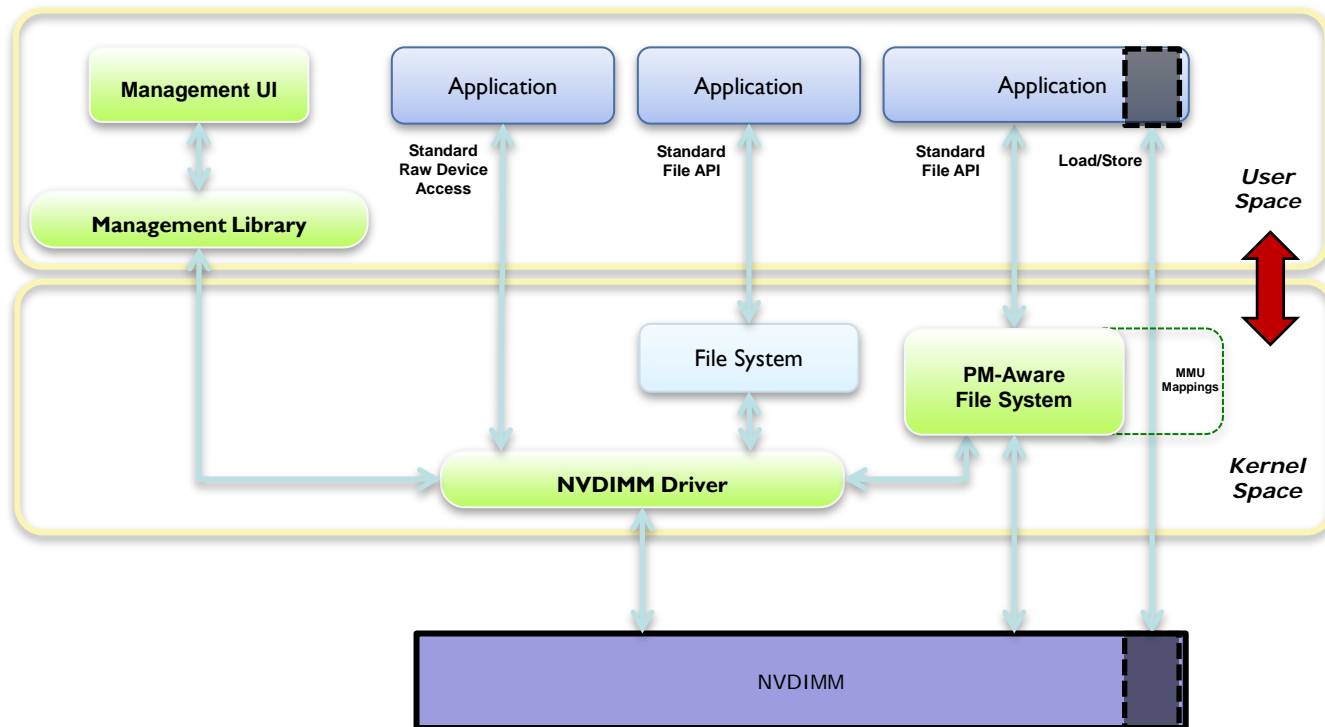


# Programming Model: Instruction Set Architecture

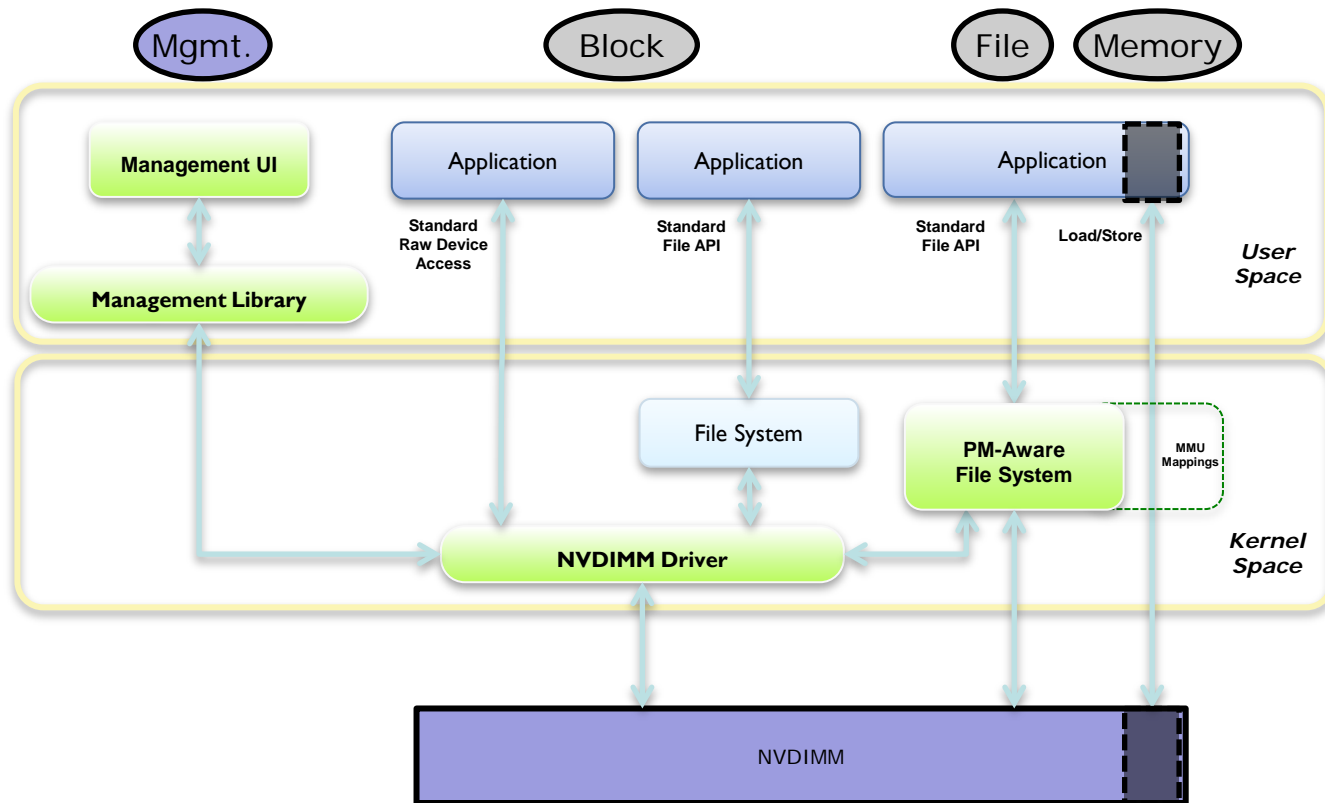


**Result:**  
Stores flushed from CPU cache,  
globally-visible → Persistent (on x86)

# Programming Model: Exposing to Applications

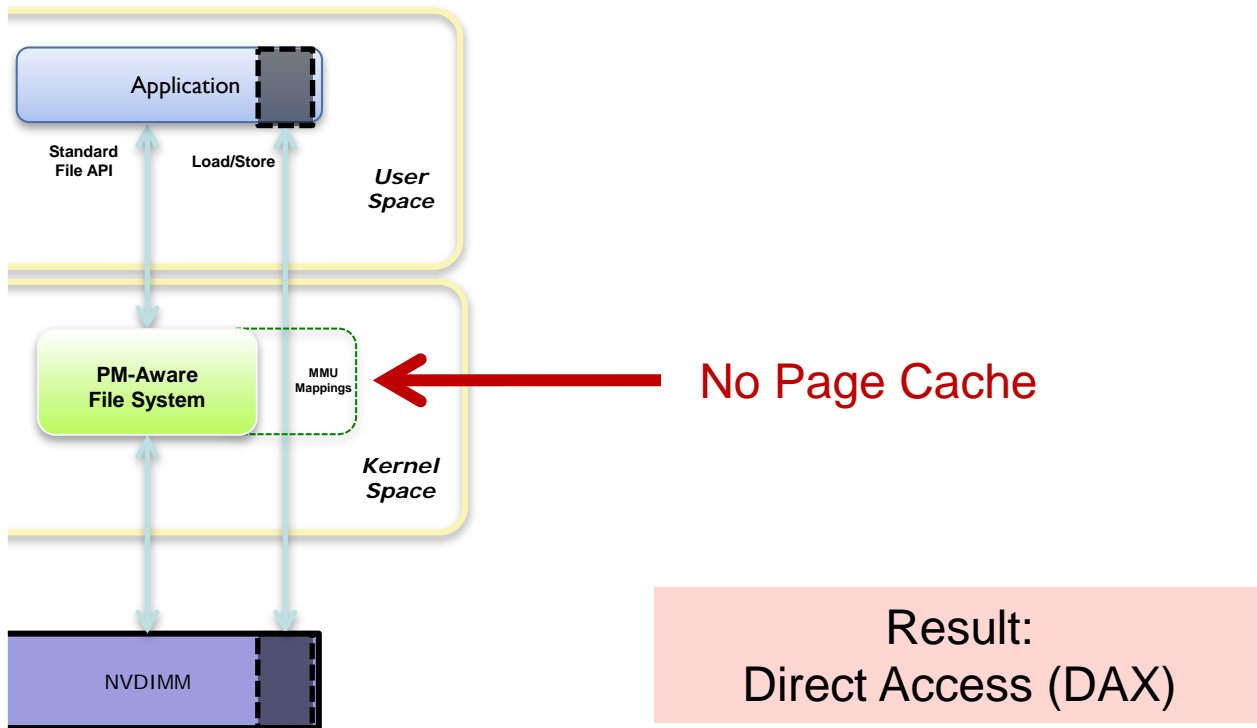


# SNIA NVM Programming Model

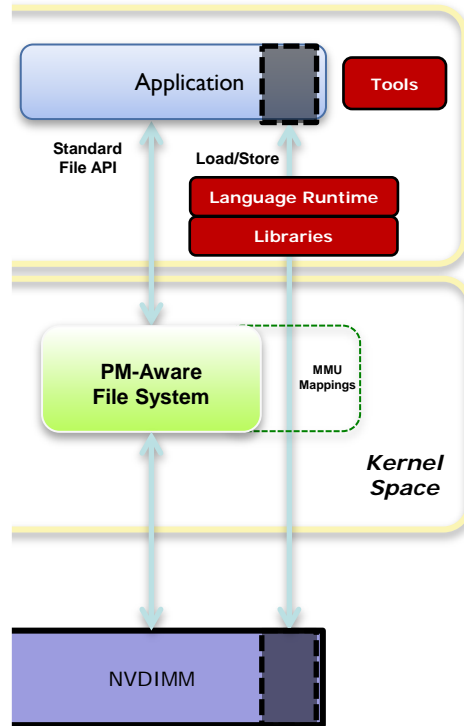




# Memory-Mapped Files



# Programming Model: The Programmer Experience



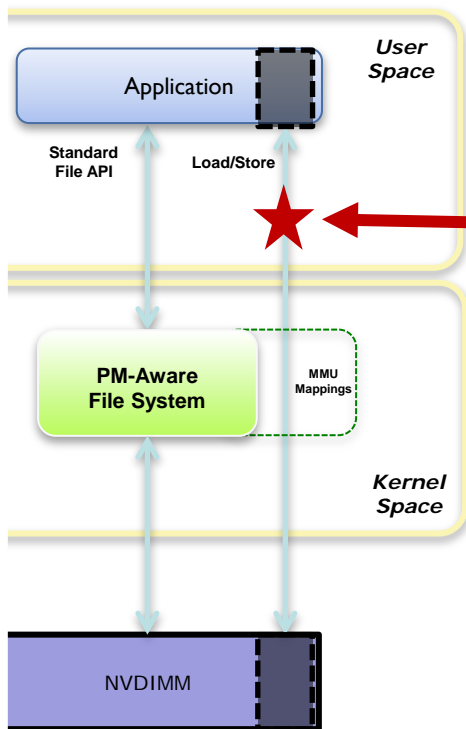
**Result:**  
Safer, less error-prone,  
idiomatic in common languages

## What's Next for the SNIA NVM Programming Model?

- More details on flushing to persistence
  - ◆ Includes flushing to remote persistence
- Continue to refine the error model
- Transactions
- APIs?

## Why is “Flushing” a Focus?

# Memory-Mapped Files



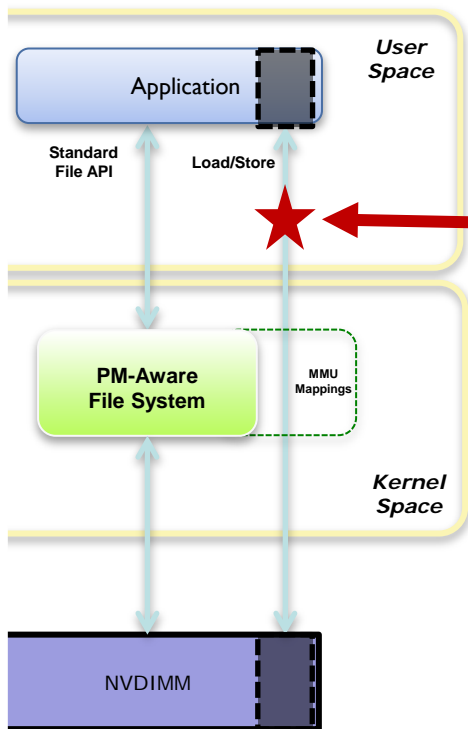
## Standard Flush

`msync()`

`FlushViewOfFile()`

`FlushFileBuffers()`

# Memory-Mapped Files



## Standard Flush

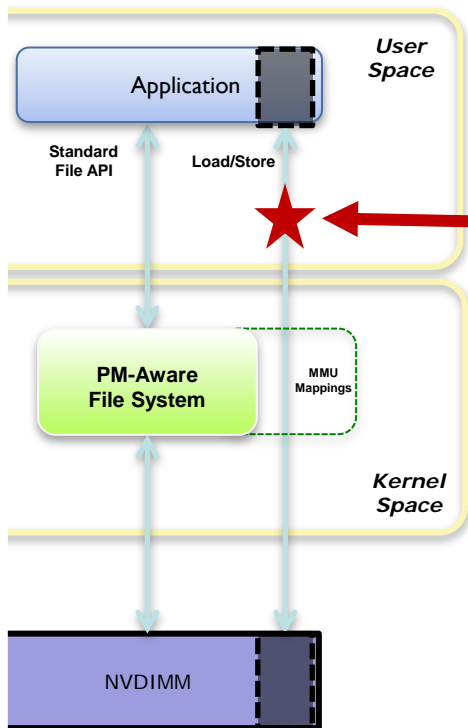
`msync()`

`FlushViewOfFile()`

`FlushFileBuffers()`

30-year-old mechanism:  
Just as error-prone today  
as it was 30 years ago

# Memory-Mapped Files



## Standard Flush

`msync()`

`FlushViewOfFile()`

`FlushFileBuffers()`

30-year-old mechanism:  
Just as error-prone today  
as it was 30 years ago

**Not Atomic**



# Java PersistentSortedMap

```
PersistentSortedMap employees = new PersistentSortedMap();
```

```
...
```

```
employees.put(id, data);
```

No flush calls.  
Transactional.  
Java library handles it all.

See “pilot” project at: <https://github.com/pmem/pcj>

# What Lies Between:

## Memory-Mapped Files

### Standard Flush

```
msync()
```

```
FlushViewOfFile()
```

```
FlushFileBuffers()
```

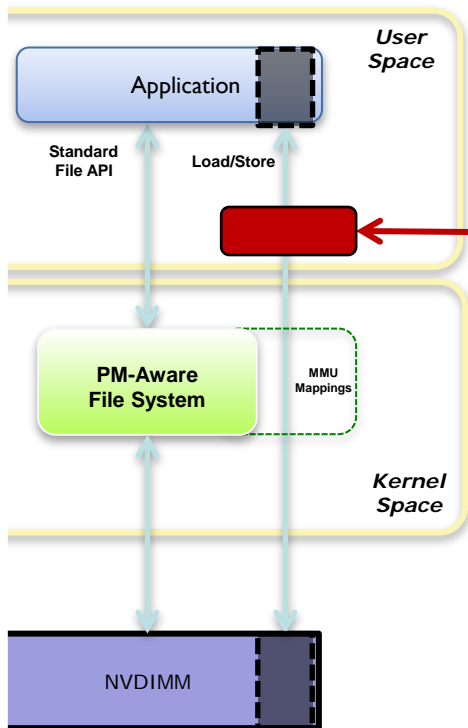
## High-Level Language Support

```
PersistentSortedMap
```

```
employees.put(id, data);
```

# NVM Libraries: pmem.io

C/C++ on Linux and Windows



NVM Libraries

- Open Source
    - <http://pmem.io>
  - libpmem
  - libpmemobj
  - libpmemblk
  - libpmemlog
  - libvmem
- Transactional

# Types of Store Barriers

Barrier Type	Current Status
Standard API	Fully specified Fully supported: <ul style="list-style-type: none"><li>• Linux (ext4, XFS)</li><li>• Windows (NTFS)</li></ul>
Optimized Flush	Specified, but evolving (ask when safe) <ul style="list-style-type: none"><li>• Linux: <b>unsafe</b> except Device DAX<ul style="list-style-type: none"><li>• (and new file systems)</li></ul></li><li>• Windows: safe</li></ul>
Remote Flush	Proposals under discussion (works today with extra round trip)
Deep Flush	Upcoming Specification
Transactions	Built on above via libraries and languages <b>Much more language support to do</b>

# Summary

- ◆ SNIA NVM Programming Model
  - ◆ Both Windows and Linux support basic model
  - ◆ Evolving to meet ongoing needs of community
  - ◆ <http://snia.org/nvmp>
- ◆ Interesting work ongoing on remote PM
  - ◆ Expect specs this year in SNIA & other forums
- ◆ Continued refinement of the SNIA model in TWG
  - ◆ Flush semantics
  - ◆ Error semantics
  - ◆ Transactions, and more...
- ◆ Ongoing community work on libraries & languages
  - ◆ <https://github.com/pmem/nvml>
  - ◆ <http://pmem.io>