

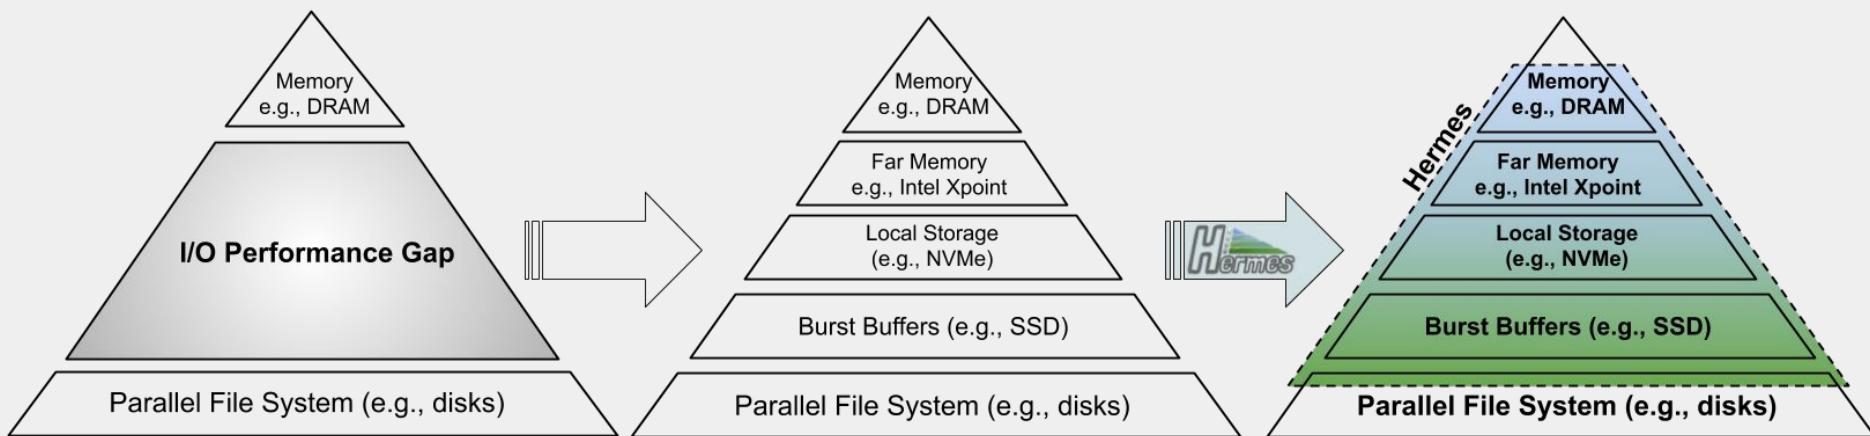


Anthony Kougkas
akougkas@iit.edu

Hermes Overview

2

- Modern storage system designs include multiple tiers of storage organized in a **deep storage hierarchy**. The goal is to mask the I/O gap.
- Hermes is a new distributed I/O buffering software library that abstracts the complexity of managing these multiple tiers and aims to boost performance.



Hermes Project

The team

Collaborative project
funded by NSF



ILLINOIS INSTITUTE
OF TECHNOLOGY



The **UDF** Group



ILLINOIS

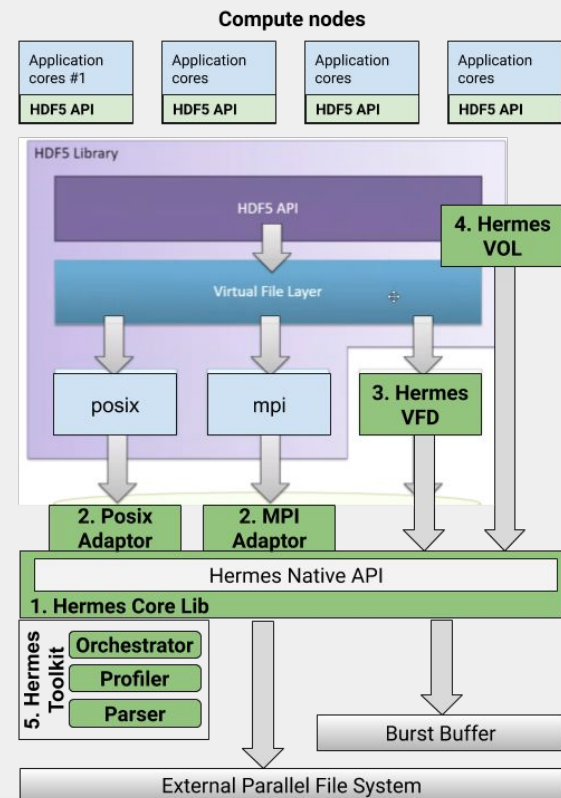
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN



Hermes ecosystem

4

1. Hermes core library
 - a. Manages tiers transparently
 - b. Facilitates data movement in the hierarchy
 - c. Provides native buffering API
2. Hermes Adapters
 - a. POSIX, MPI-IO, Object Store, etc
 - i. Intercept I/O calls to Hermes
 - ii. Boosts legacy app support
3. Hermes VFD
4. Hermes VOL
 - a. Captures application's behavior and provides hints to Hermes core lib
5. Hermes Toolkit






Hermes Container Library (HCL)

5

- Deep Storage Hierarchy
 - Spans multiple tiers within a node...
 - ...but also multiple nodes in the cluster
- Applications need to distribute data structures across multiple nodes.
 - Hermes Container Library
 - <https://github.com/HDFGroup/hcl>
 - We invite the community to try it out
 - And please give us feedback.





HDF5 VOL experience

6

- We have built a couple of VOL plugins within the Hermes project
 - A global data prefetcher (**HFetch**)
 - An intelligent compression framework (**HCompress**)
- Main VOL functionality we appreciate
 - The ability to intercept HDF5 calls with minimal code changes
- VOLs can now be dynamically loaded using environment variables
 - Makes it easier for us developers
 - No need to recompile
 - Allowed us to experiment with different configurations (On or Off the plugin)



Thank you