

How to leverage multi-tiered storage to accelerate I/O

Sept 26th 2022
CS492

Anthony Kougkas - akougkas@iit.edu





Agenda


2

- Modern systems and the I/O bottleneck
- Hermes project overview
- ChronoLog project overview



I/O Performance Gap

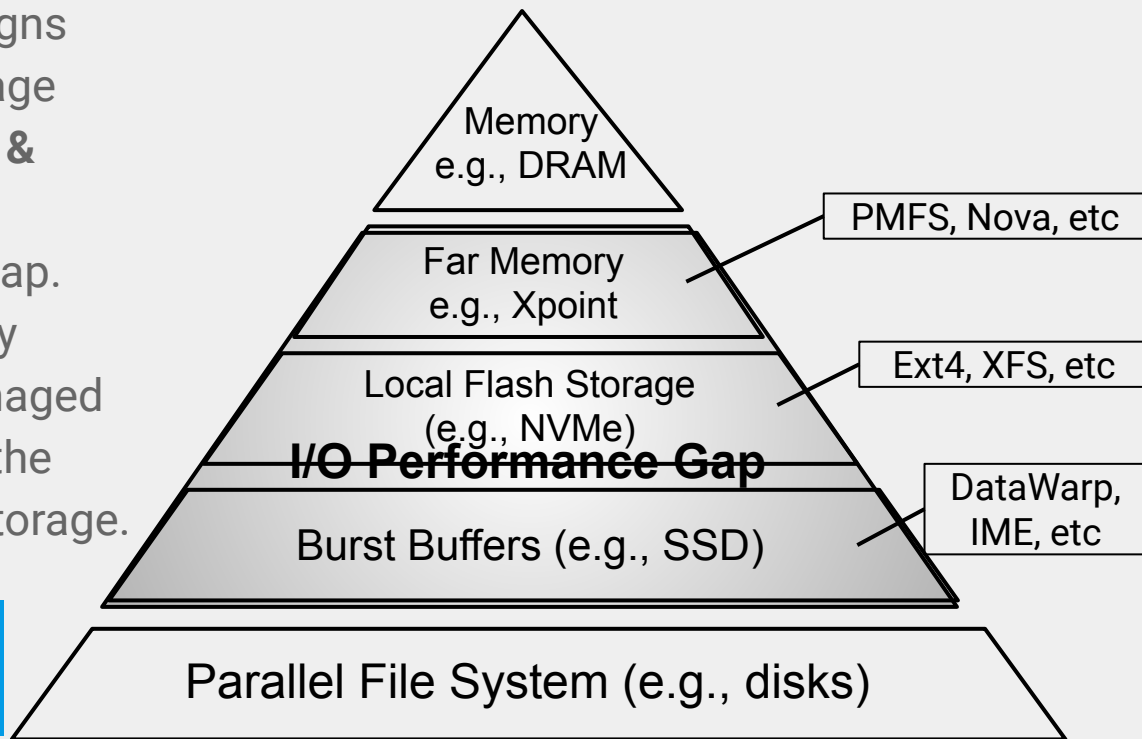
- Traditionally memory systems and storage demonstrate wildly different performance.
 - Access latency
 - Bandwidth
 - Data representation
- Applications experience performance degradation due to slow remote access to storage.





Hardware Answer

- Modern storage system designs include multiple tiers of storage organized in a **deep memory & storage hierarchy (DMSH)**.
The goal is to mask the I/O gap.
- Each system is independently designed, deployed, and managed making very difficult to reap the benefits of the hierarchical storage.

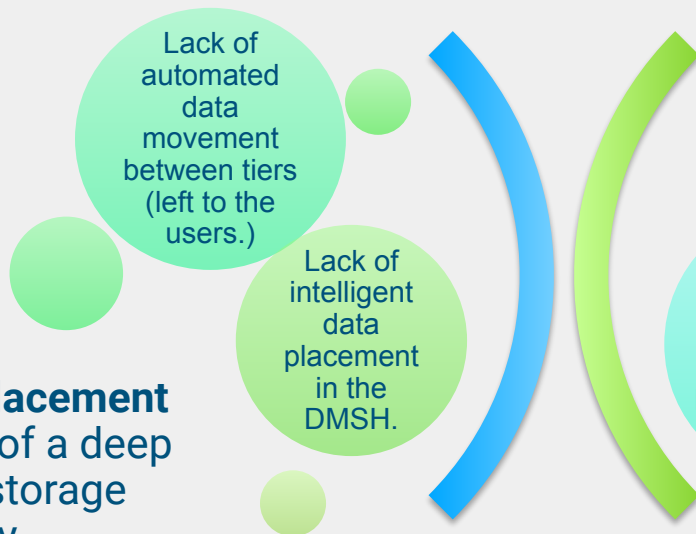


Ideally, the presence of multiple tiers of storage should be **transparent** to applications without having to sacrifice **I/O performance**.



Multi-tiered systems today

Complex data placement
among the tiers of a deep
memory and storage
hierarchy



Independent
management of each tier
of the DMSH



DMSH systems require

6



efficient and transparent **data movement** through the hierarchy



new data placement algorithms



effective metadata management



an efficient communication fabric



Leveraging multi-tiered storage

7

NSF OCI-1835764

Hermes:

A Multi-Tiered Distributed I/O Buffering System



NSF CSSI 2104013

ChronoLog:

A Distributed Tiered Shared Log Store



Learn more

- <http://www.cs.iit.edu/~scs/assets/projects/Hermes/Hermes.html>
- <http://www.cs.iit.edu/~scs/assets/projects/ChronoLog/ChronoLog.html>
- <https://github.com/HDFGroup/hermes>
- <https://github.com/HDFGroup/hcl>
- <https://github.com/scs-lab/ChronoLog>

Contact us

akougkas@iit.edu

Thank you



SCALABLE COMPUTING
SOFTWARE LABORATORY

ILLINOIS INSTITUTE
OF TECHNOLOGY

