Enabling Scientific Application I/O on Cloud Filesystems

Milo Polte, Esteban Molina-Estolano,
John Bent, Garth Gibson, Carlos Maltzahn, Maya B. Gokhale, Scott Brandt

"Super" clusters aren't just supercomputers any more

milo polte - milop@cs.cmu.edu

"Super" clusters aren't just supercomputers any more

≈1.75 Petaflops
(Top500.org)

milo polte - milop@cs.cmu.edu

"Super" clusters aren't just supercomputers any more

≈1.75 Petaflops
(Top500.org)

milo polte - milop@cs.cmu.edu

"Super" clusters aren't just supercomputers any more

≈1.75 Petaflops
(Top500.org)

450,000 servers
worldwide
(2006, NY Times)
≈20-100 Petaflops

milo polte - milop@cs.cmu.edu

"Sacred" Scientific Apps

- Climate simulations, comp biology, astro, cyber security, etc.
- "Untouchables"
- Often require a POSIX filesystem or MPI-IO interface, many writers concurrently writing a single file
- Mismatch with Cloud Filesystem semantics
  - HDFS as a Cloud Filesystem example: one-writer/file, sequential writes only, no rewrites

milo polte - milop@cs.cmu.edu

Goal and Approach

milo polte - milop@cs.cmu.edu

Goal and Approach

- **Our goal:** Running unmodified scientific apps on a cloud filesystem

milo polte - milop@cs.cmu.edu

Goal and Approach

- **Our goal:** Running unmodified scientific apps on a cloud filesystem
- **Our approach:** A FUSE-based interposition layer

milo polte - milop@cs.cmu.edu

Decoupling Concurrency

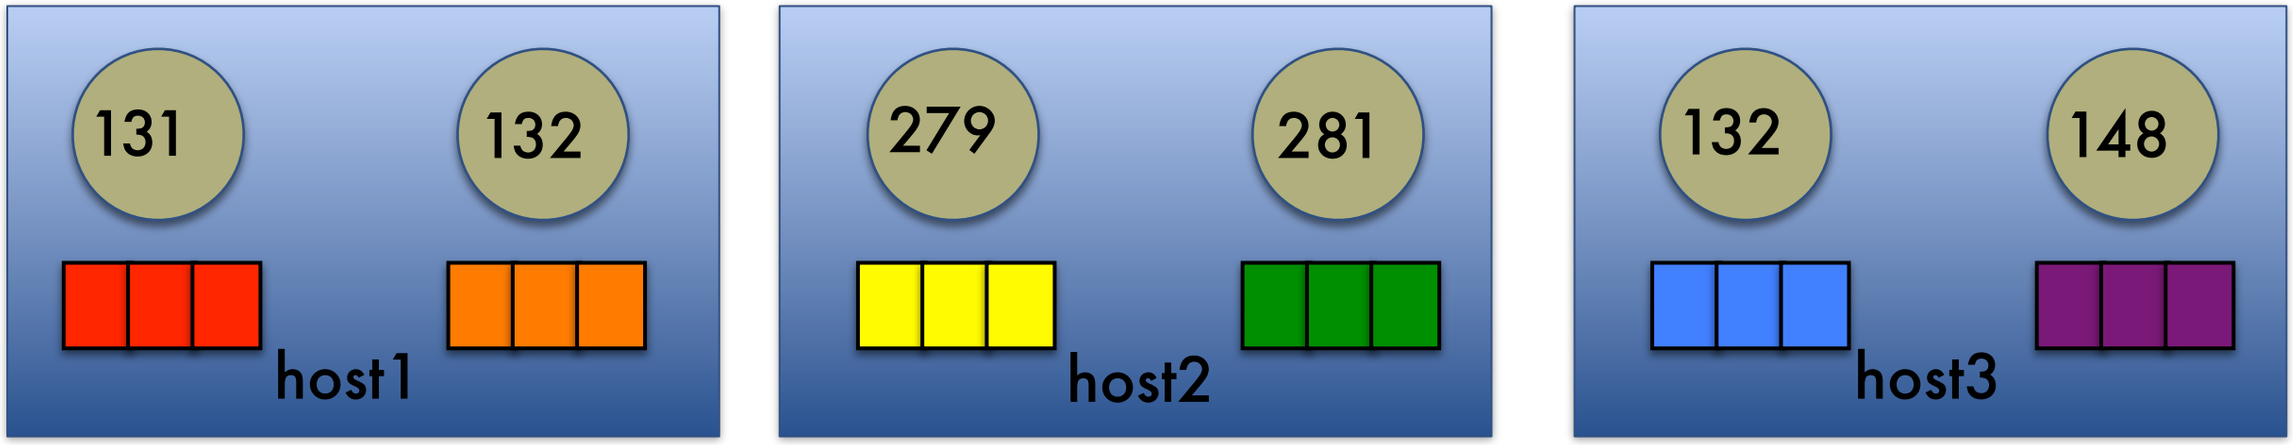

Underlying Cloud File System

Interposition Layer

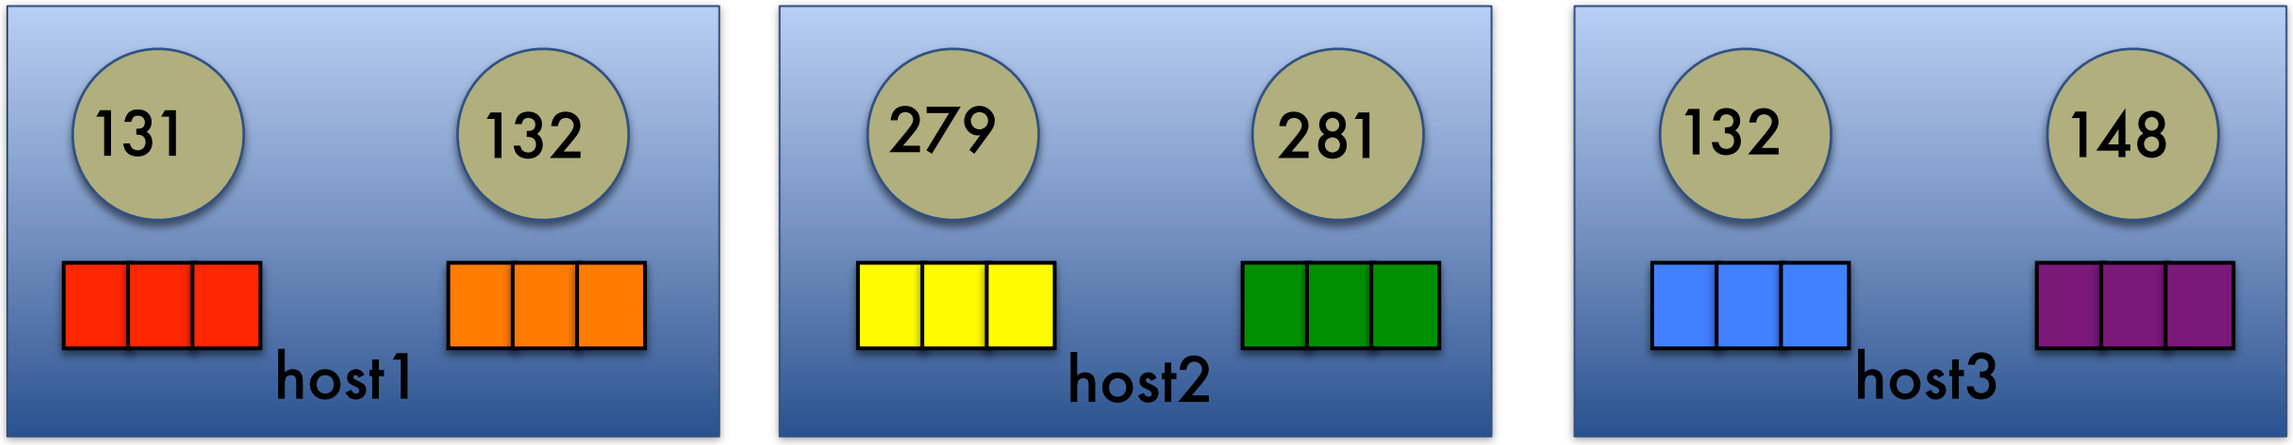

Virtual Interposition Layer

Underlying Cloud File System

Apps write normally

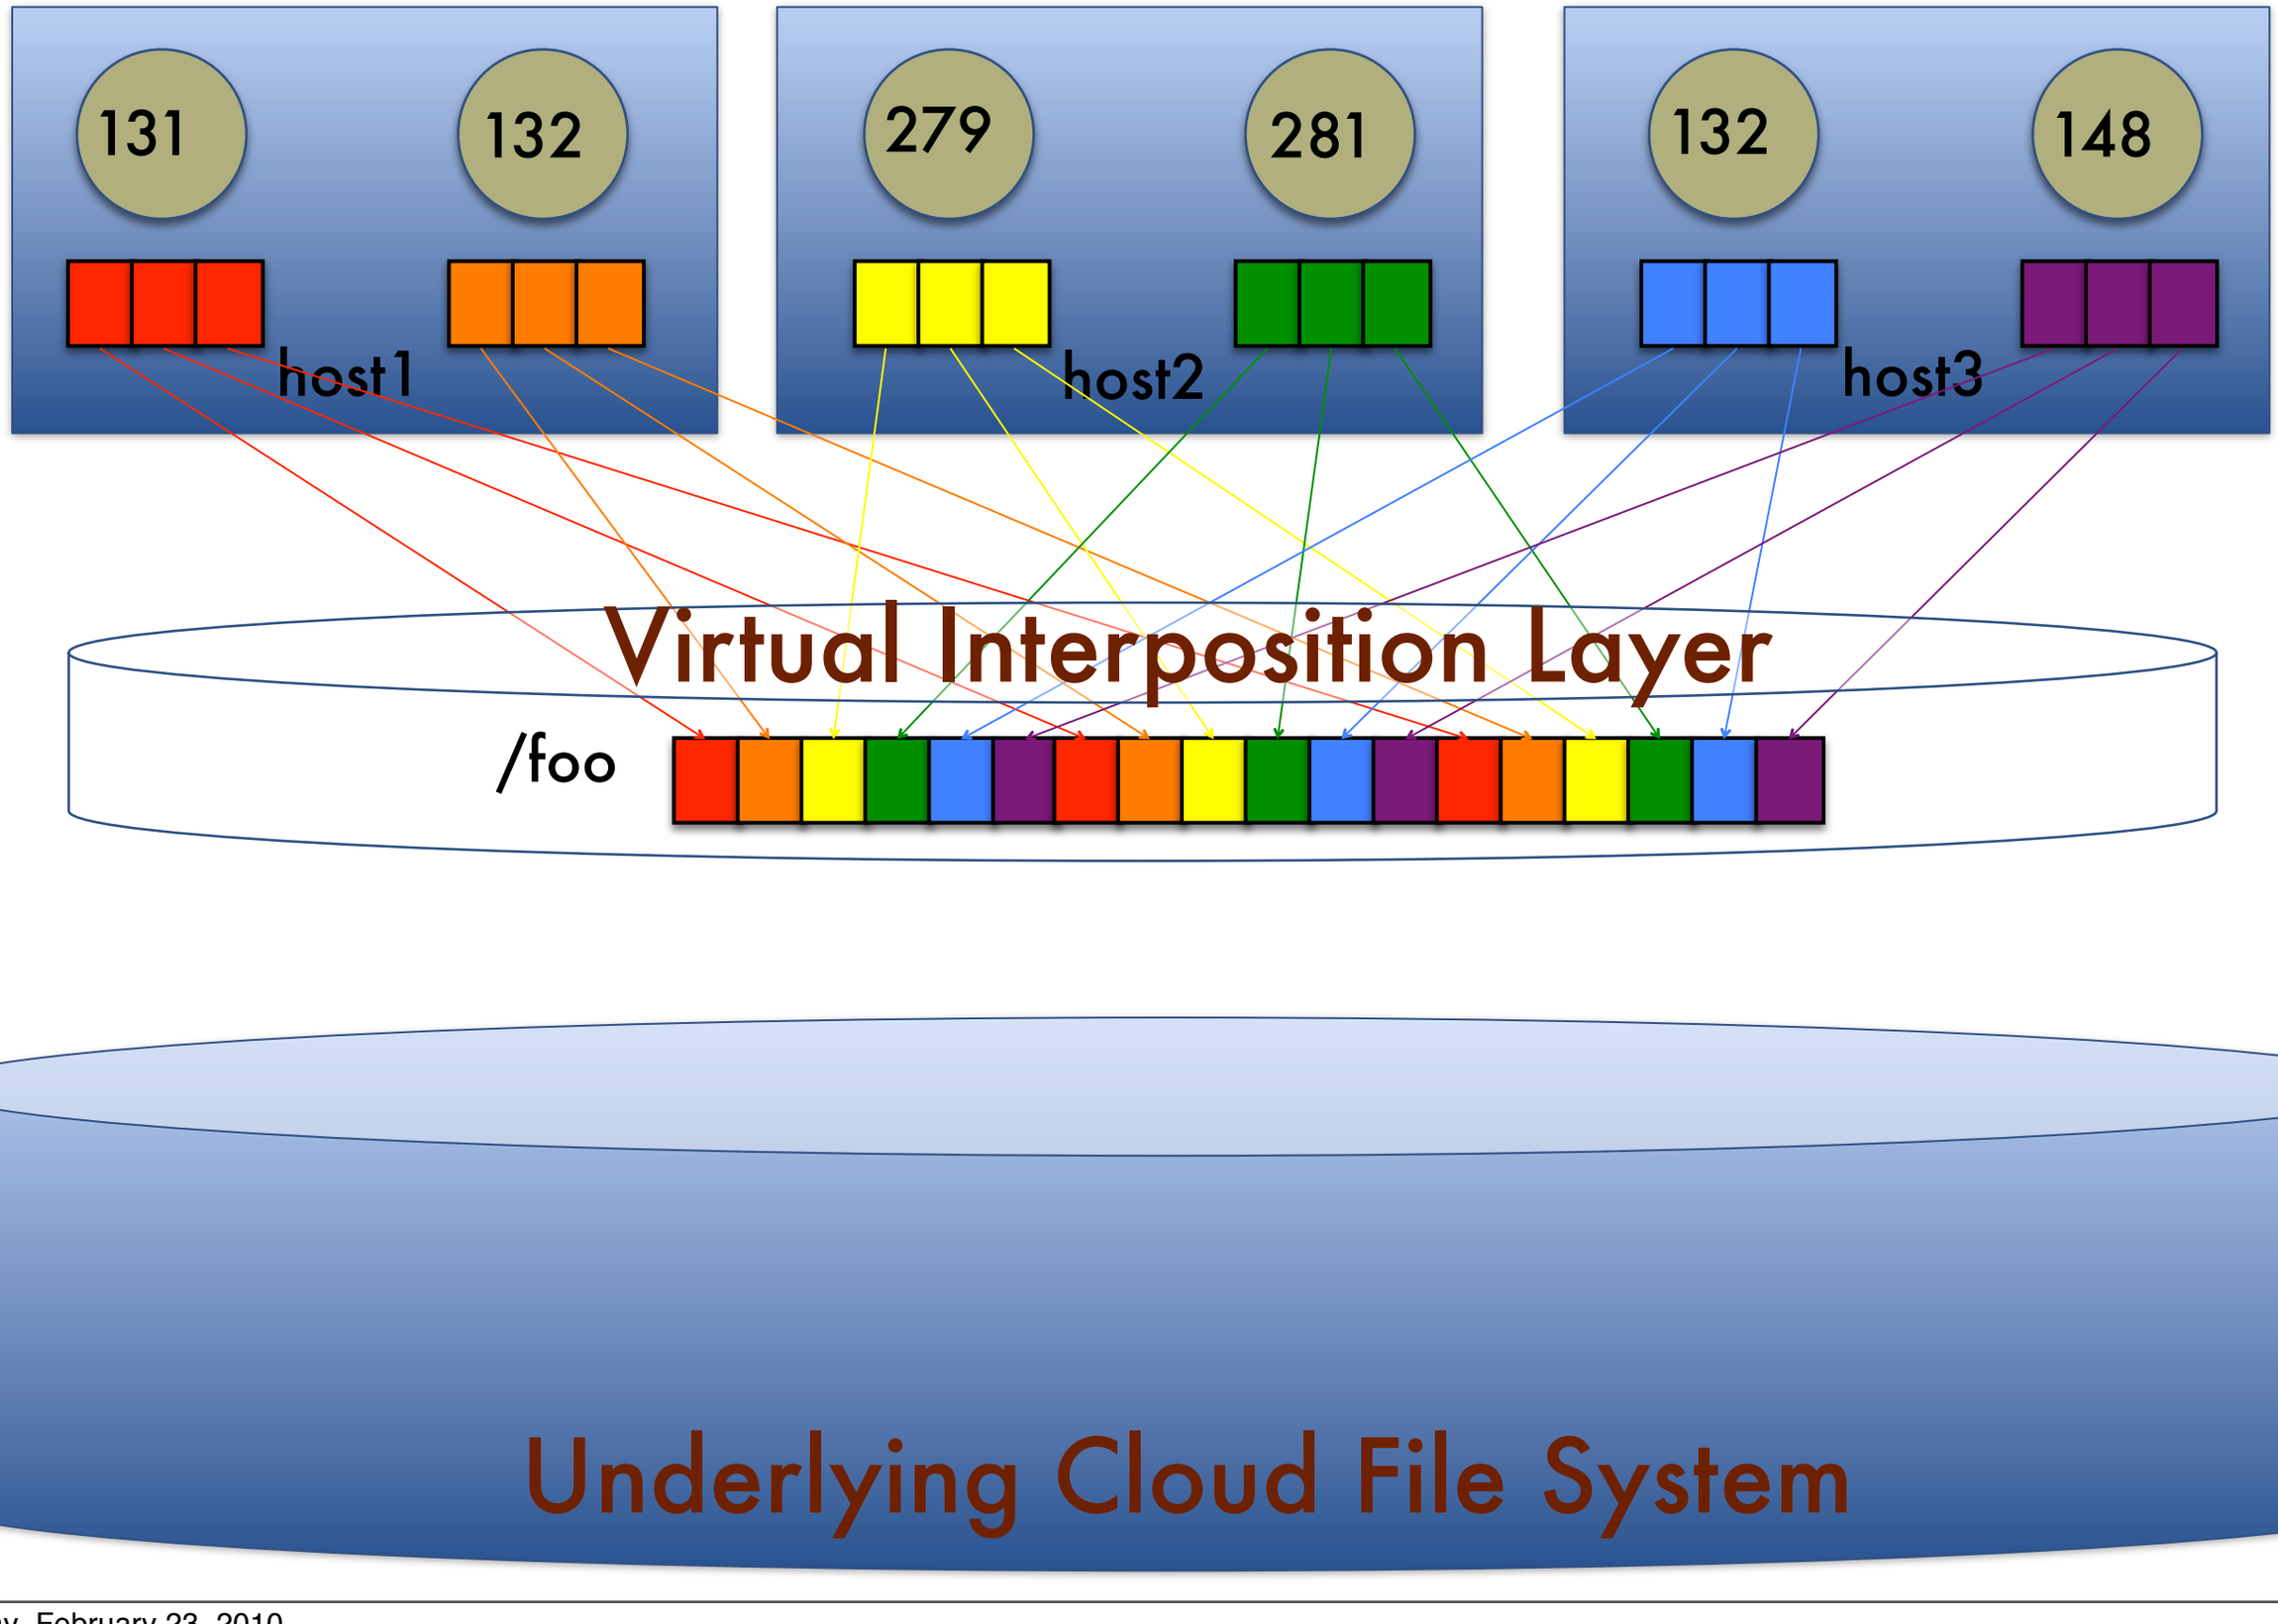

"Container" in Underlying Cloud Filesystem

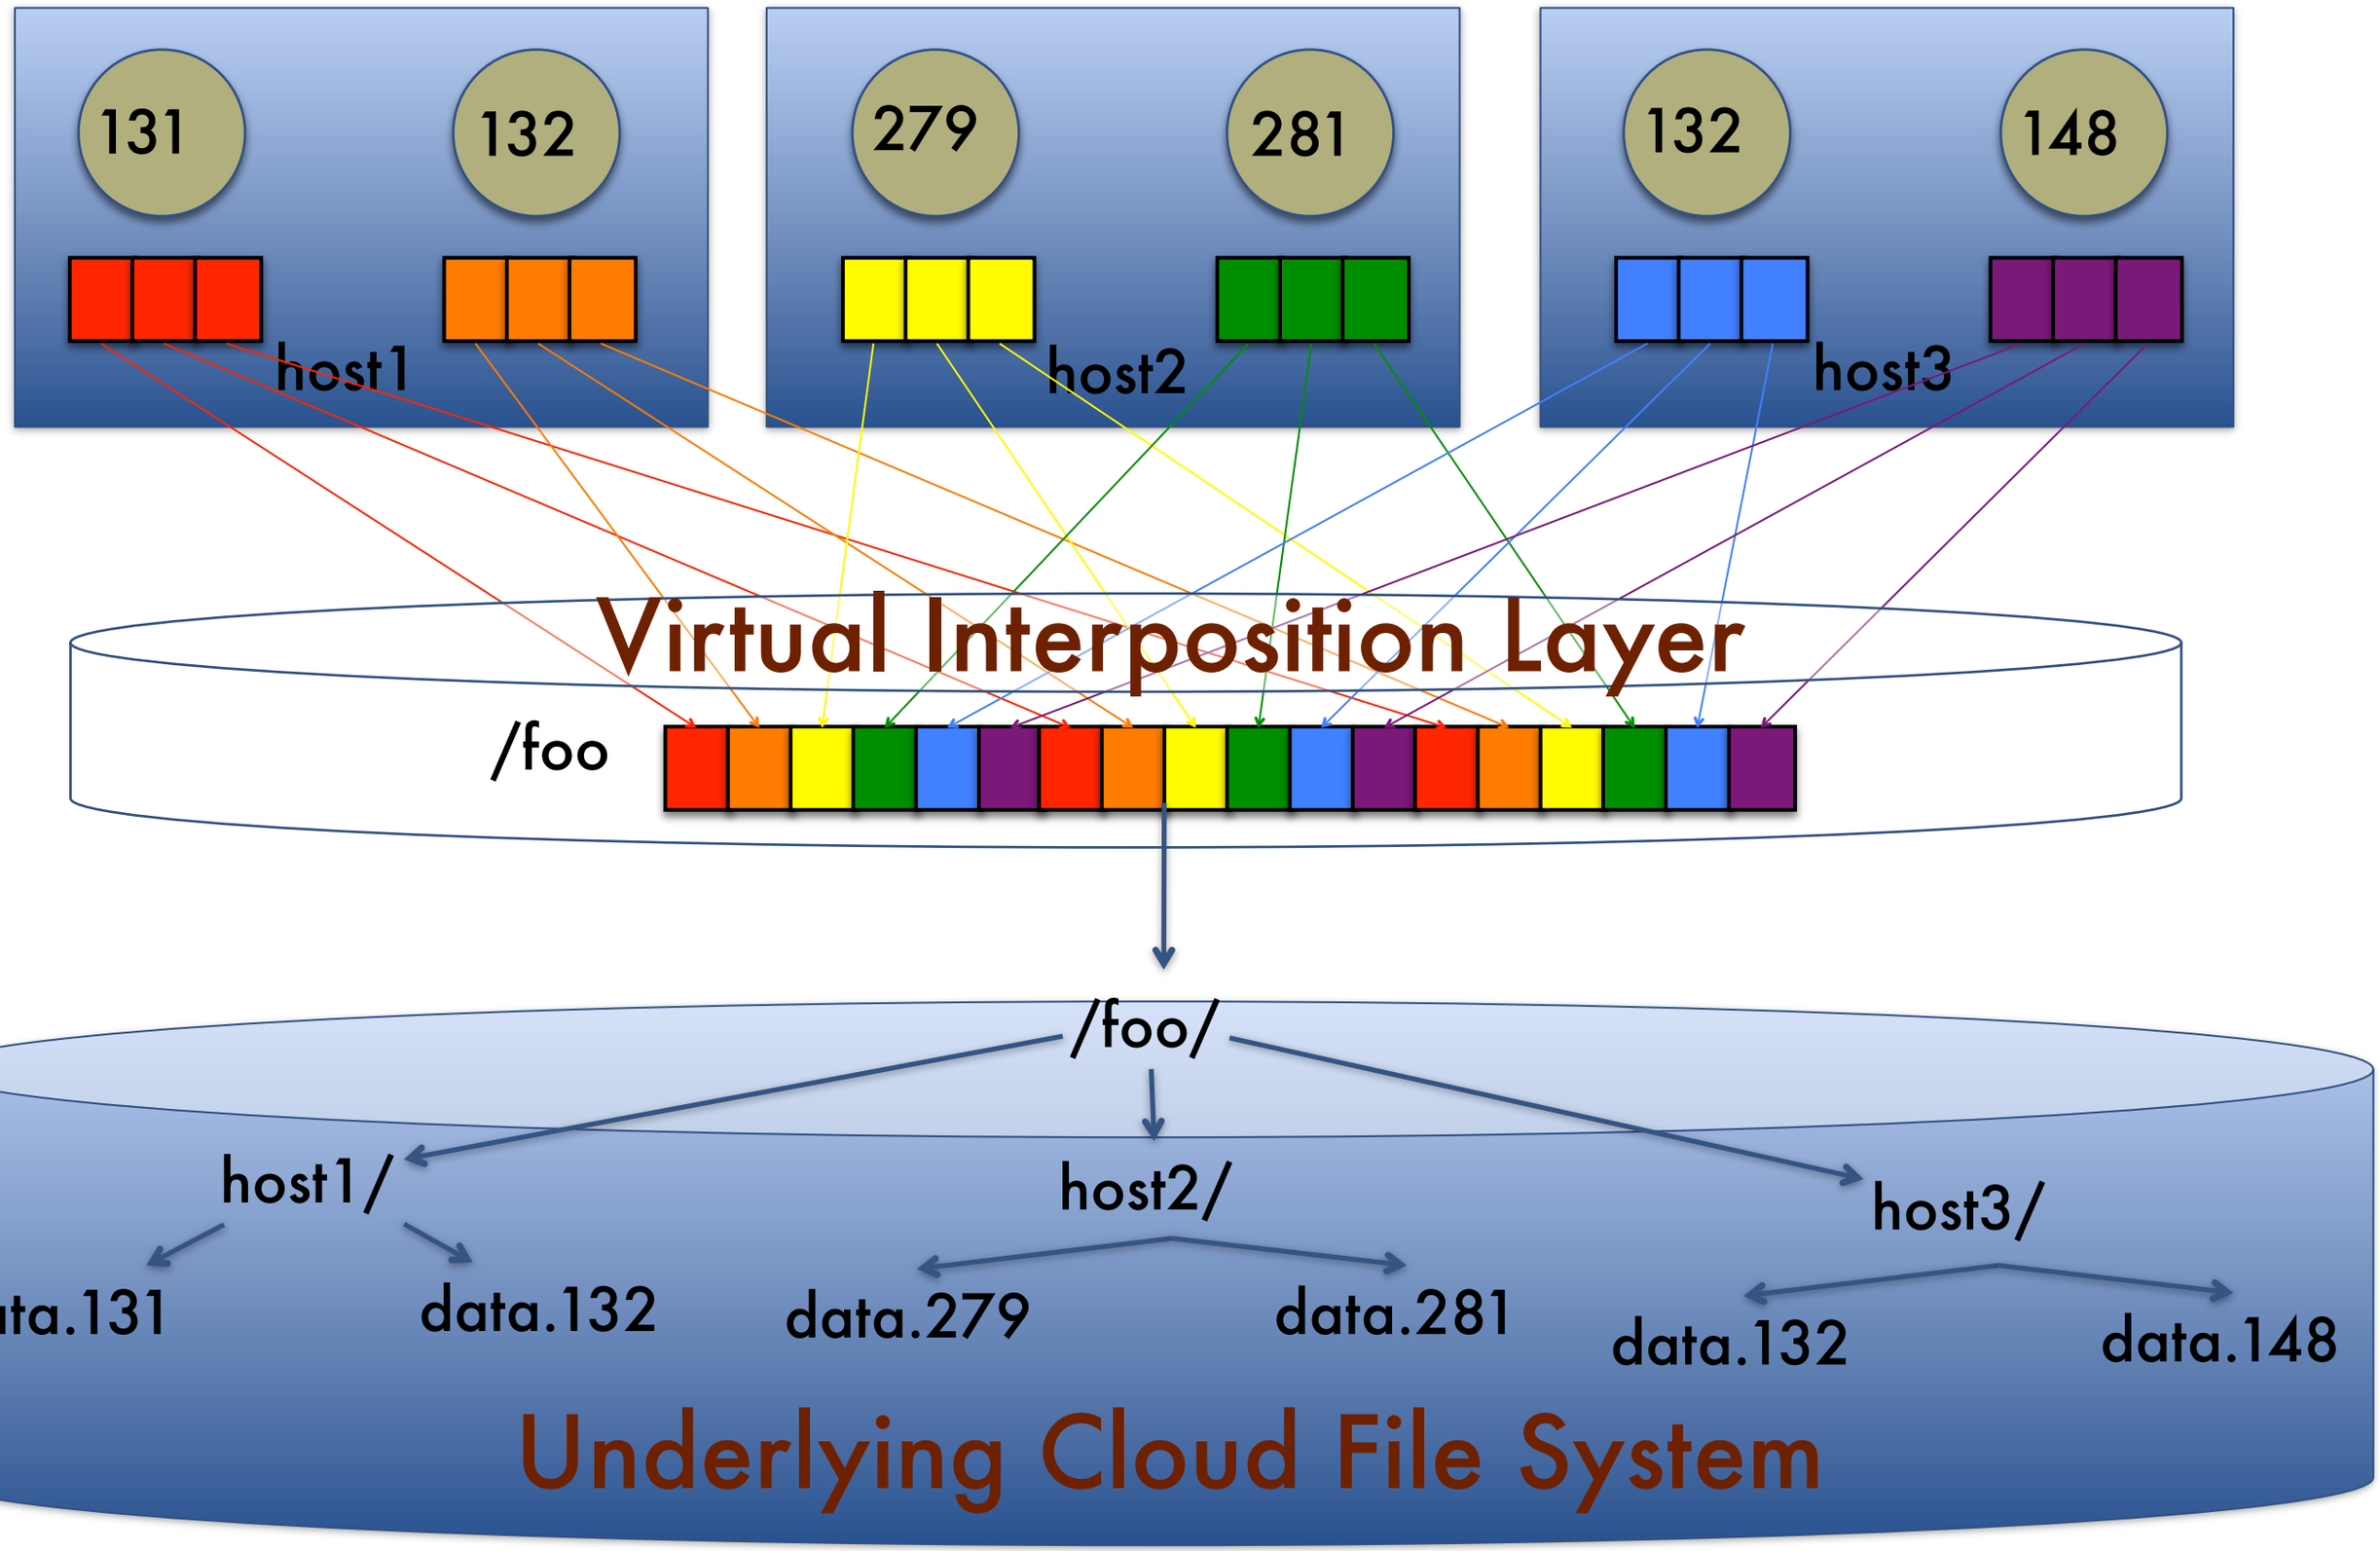

Writes redirected to log files

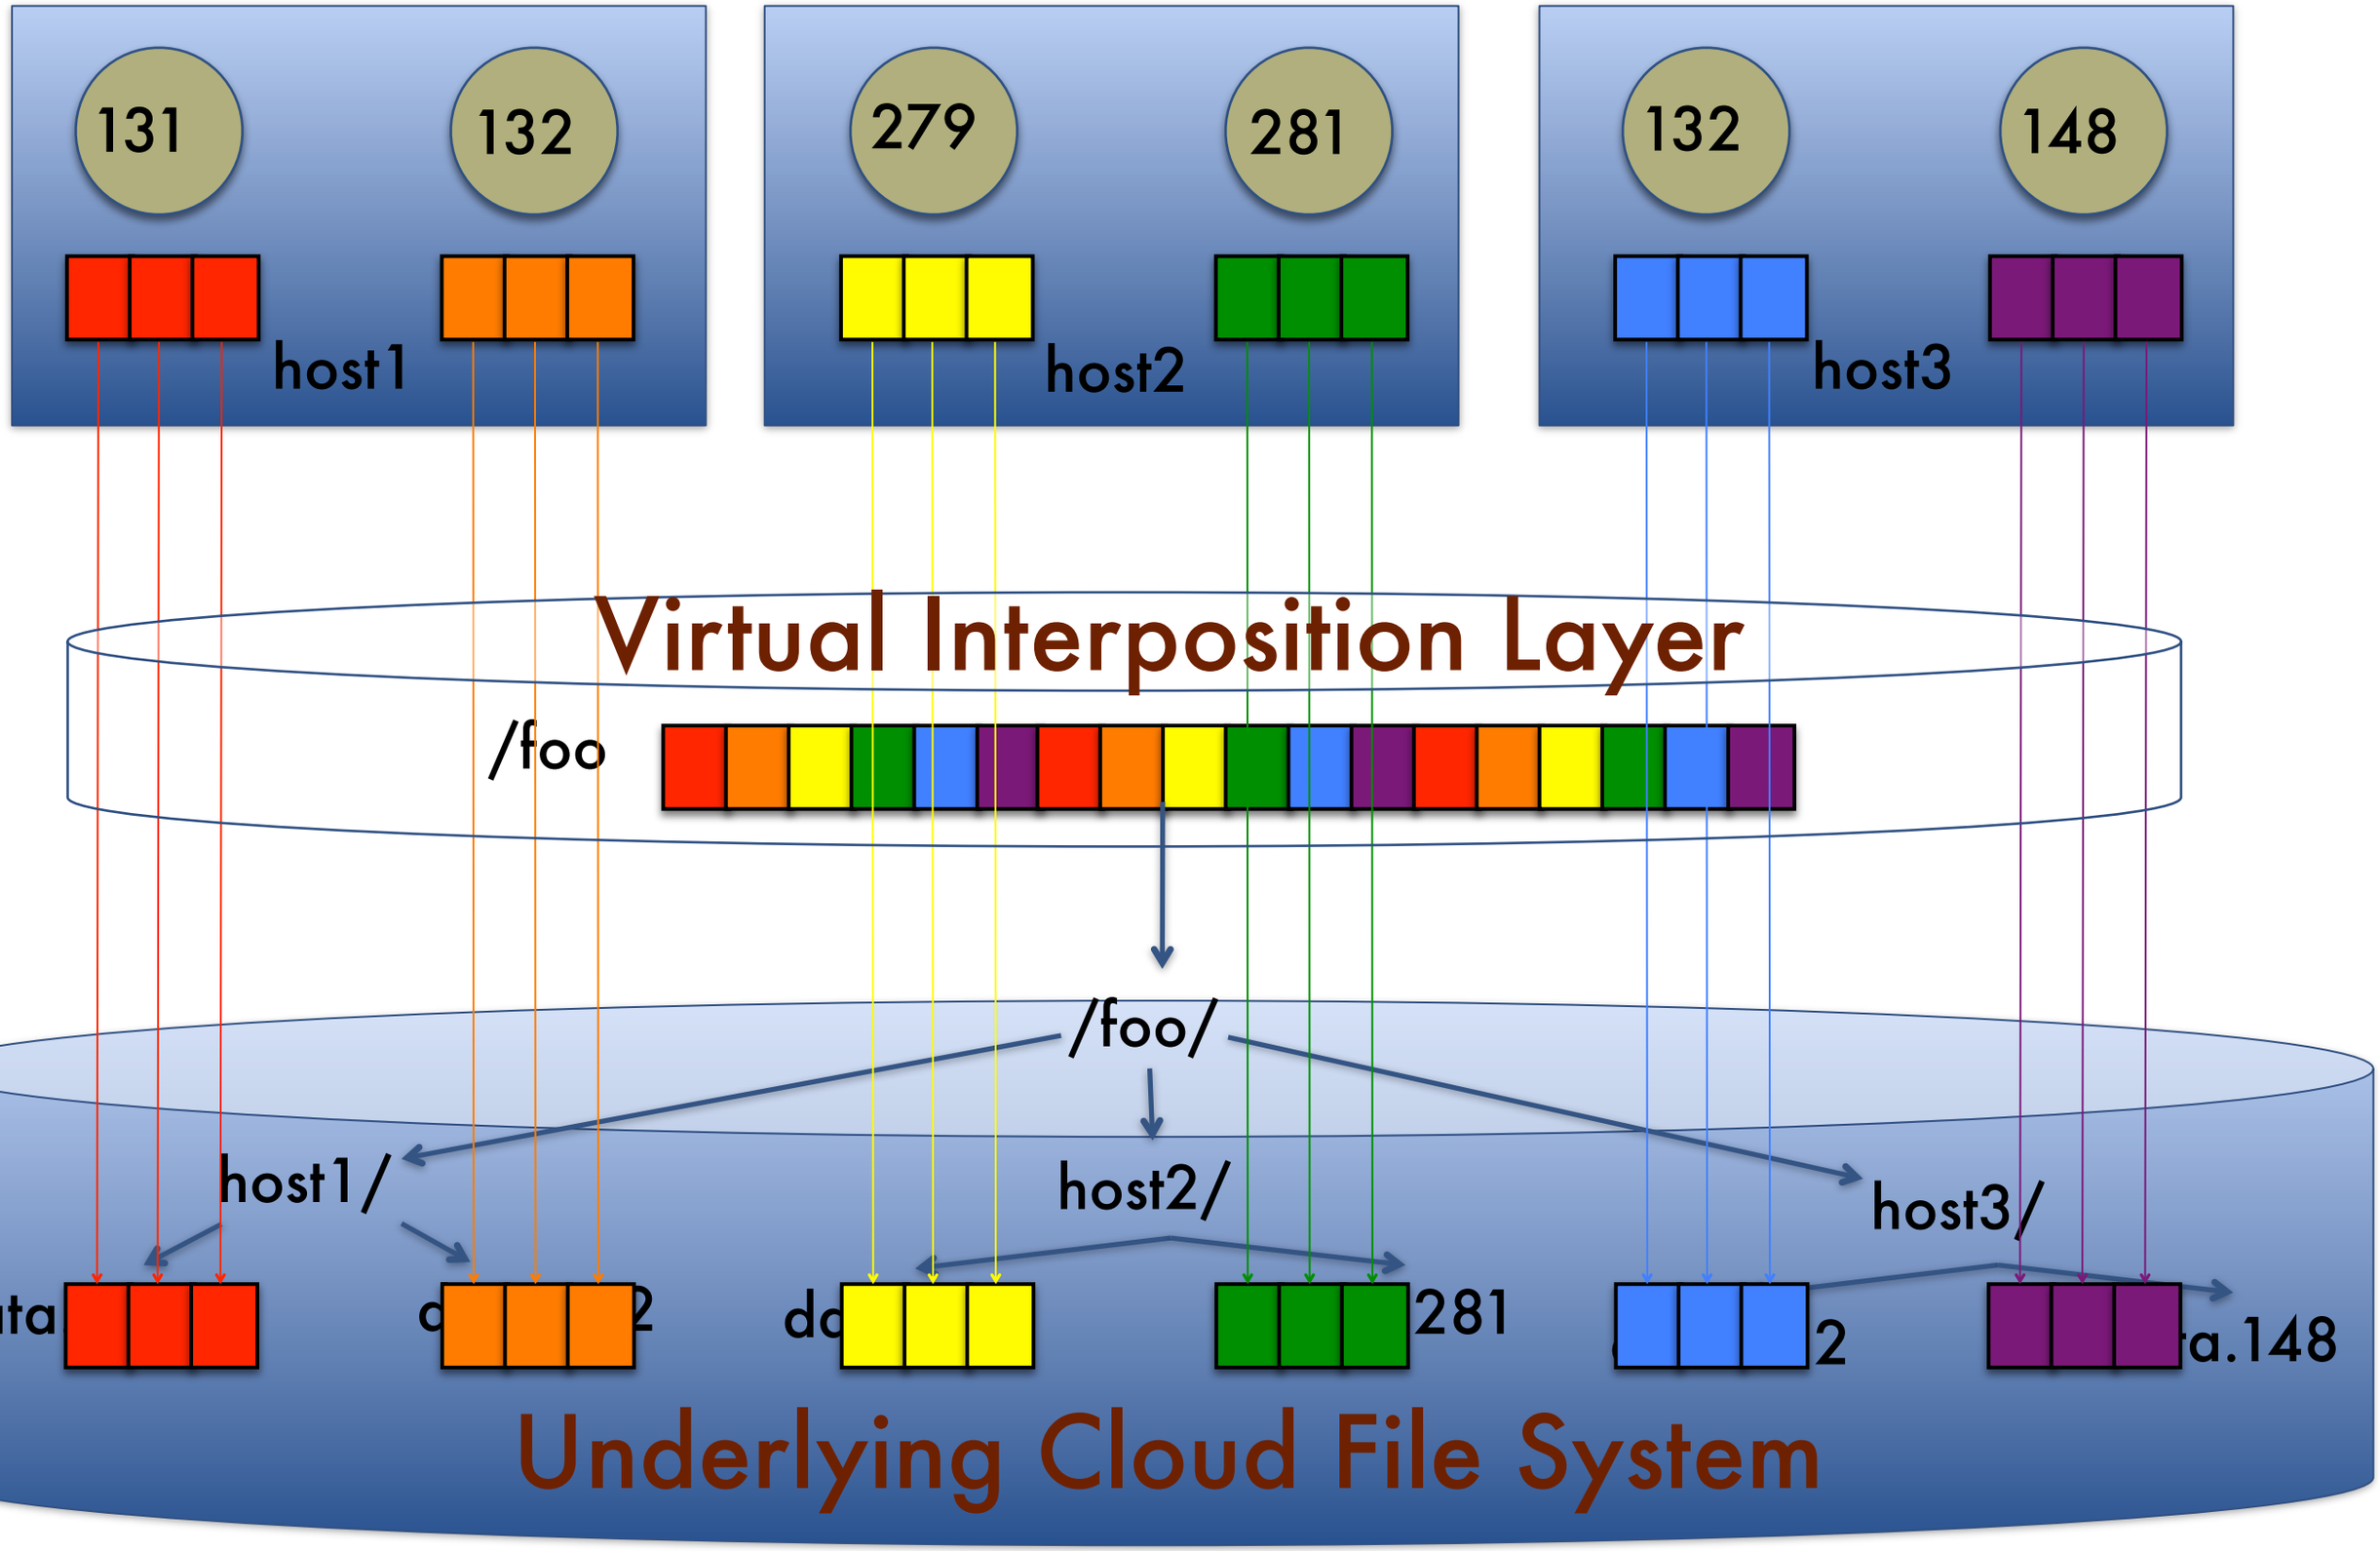

Summary

- An interposition layer lets us use a cloud filesystem storage for Scientific Apps
- Permits re-open, concurrent writes, etc.
- Writes decoupled, reads aggregated
- No changes to filesystem or apps

milo polte - milop@cs.cmu.edu

Current Status

- Preliminary interposition layer written
- Plans to test a variety of scientific apps on HDFS and PVFS
- Question: How to integrate with MapReduce?
- Come talk to us at our poster

milo polte - milop@cs.cmu.edu