

LAD'21

# DAOS/Lustre Integration

Bruno Faccini, Johann Lombardi (Intel)

Kevin Harms (ALCF)

The Intel logo is located in the bottom left corner of the slide. It consists of the word "intel" in a white, lowercase, sans-serif font, followed by a registered trademark symbol (®). The logo is positioned to the right of a decorative graphic of several overlapping squares in various shades of blue.

intel®

# Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

No product or component can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/benchmarks>.

Intel Advanced Vector Extensions (Intel AVX) provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.


Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

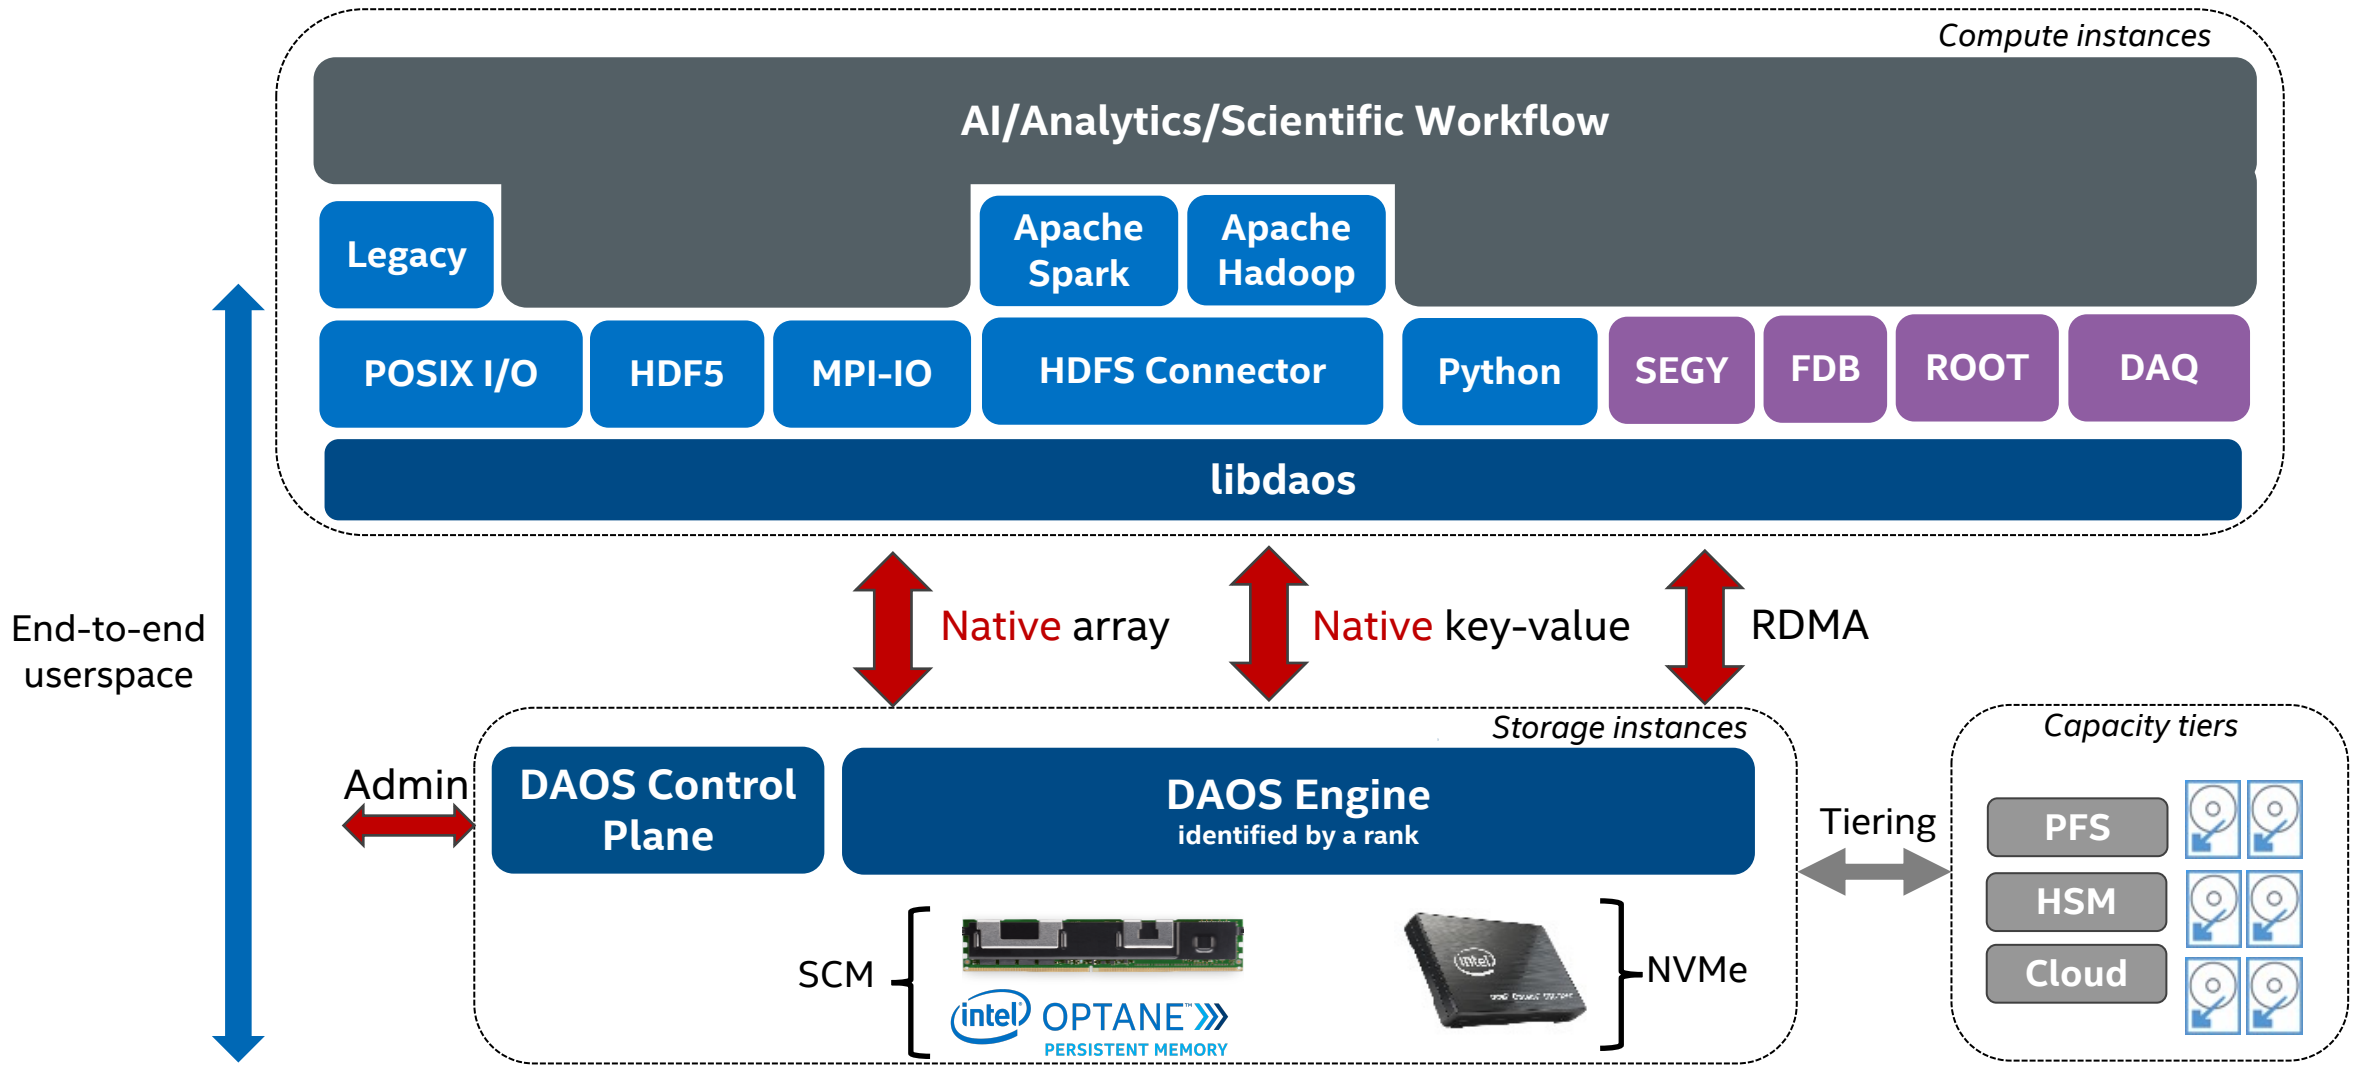
# Agenda

- DAOS/Lustre Integration Update (Johann)
- Demonstration (Bruno)
- Use Case (Kevin)

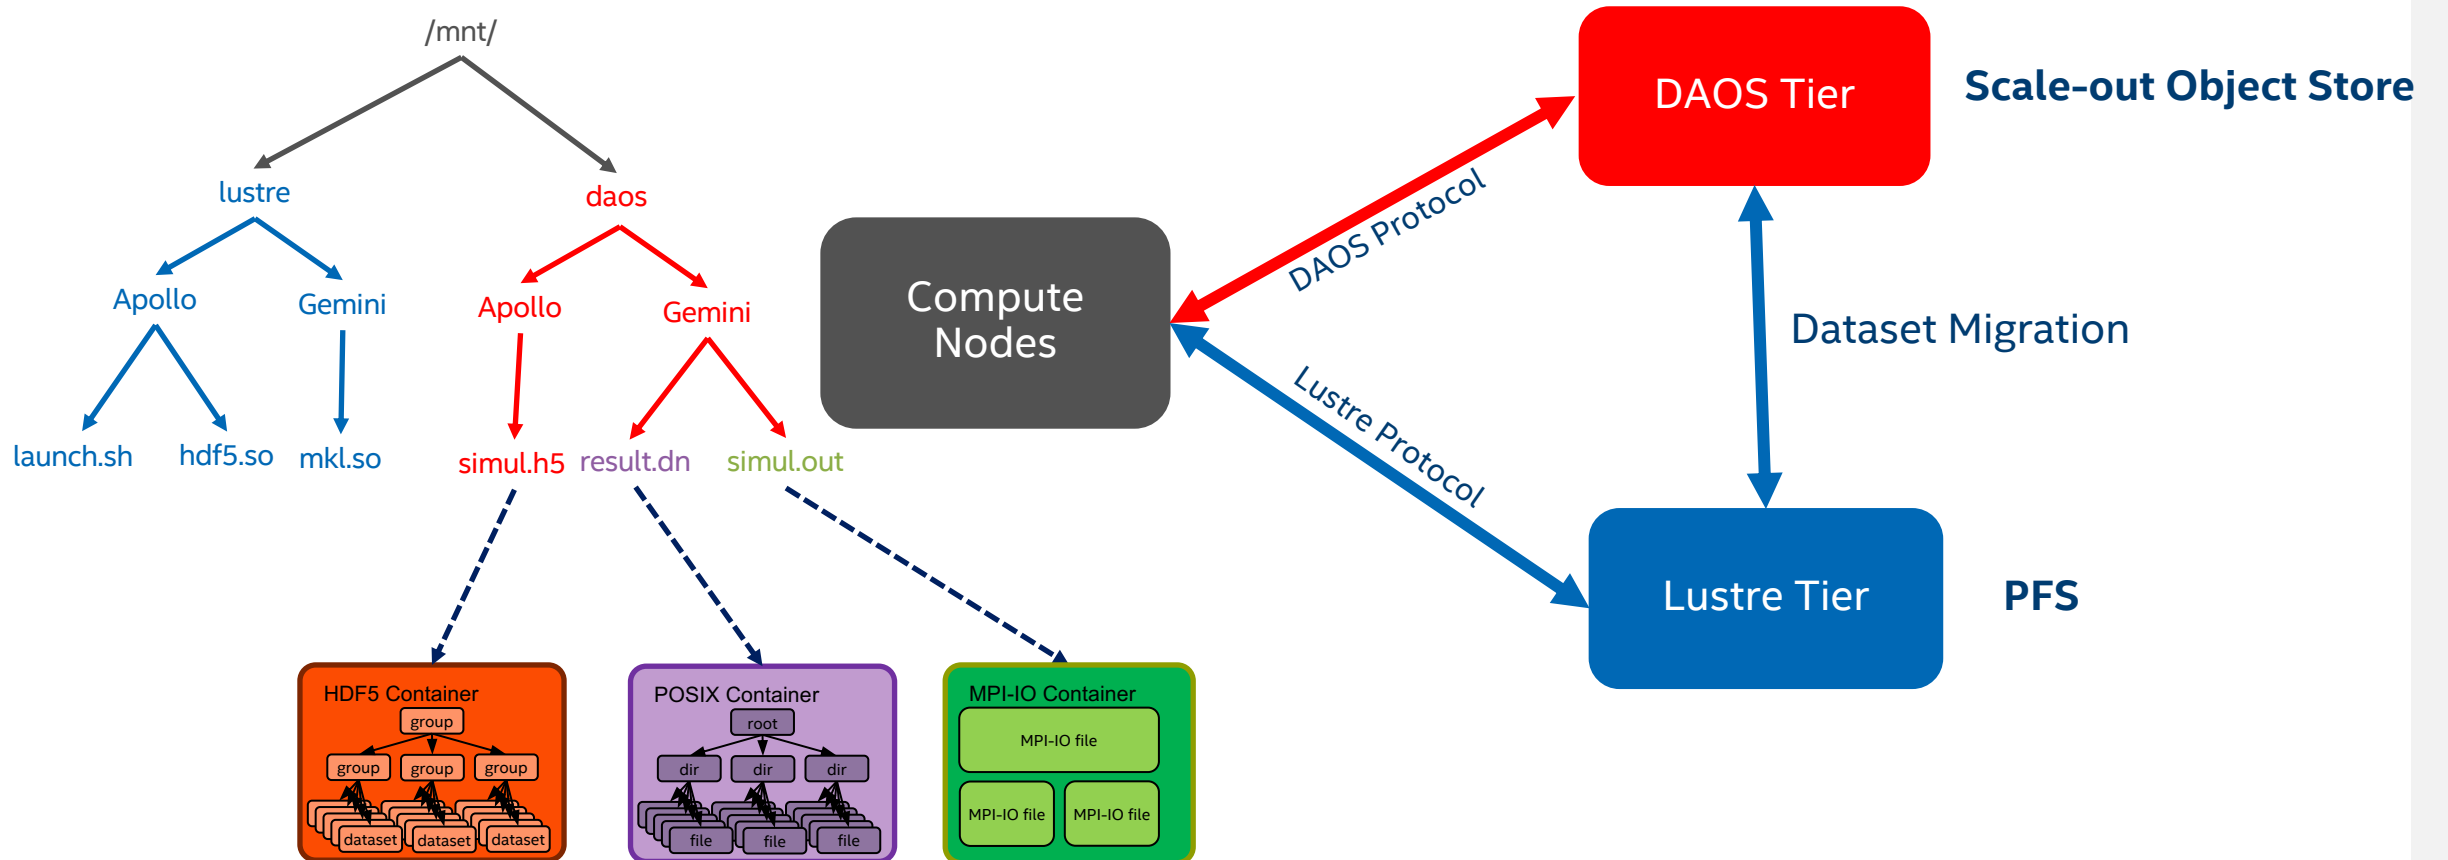


# DAOS: Nextgen Object Store

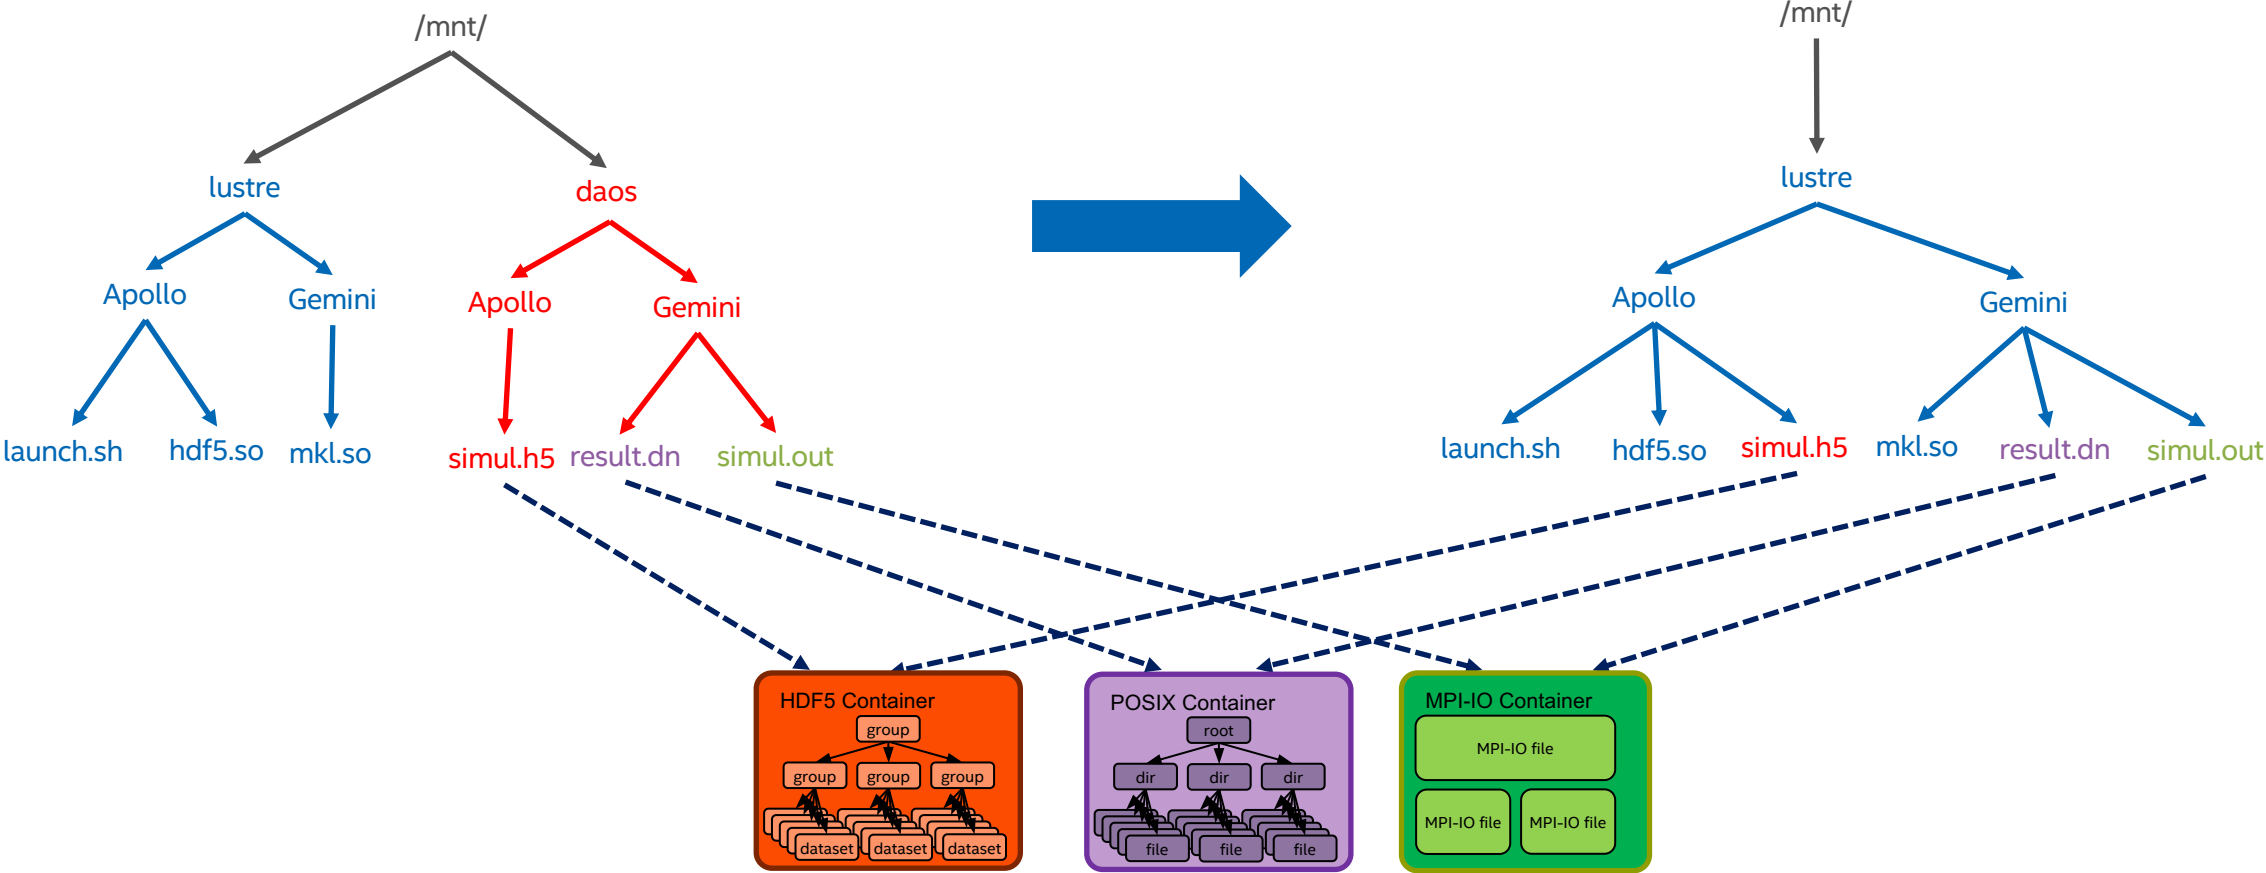
- Generic I/O middleware supported today
- Domain-specific data models under development in co-design with partners



# Targeted Storage Architecture



# Unified Namespace Concept



# Unified Namespace Concept

```
jlombard@jprisaca-MOBL:~$ cd /mnt/lustre
jlombard@jprisaca-MOBL:/mnt/lustre$ ls -R
Apollo Gemini

./Apollo:
hdf5.so  launch.sh simul.h5

./Gemini:
mkl.so   result.dn simul.out




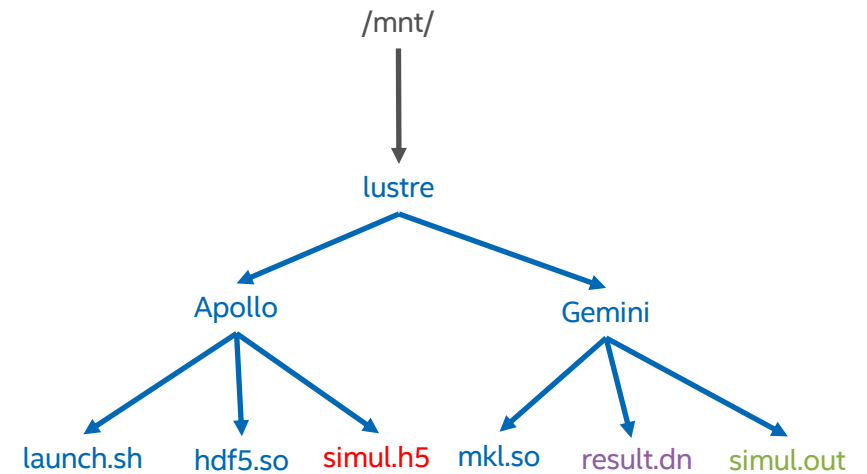
./Gemini/result.dn:
dir1 dir2 dir3

./Gemini/result.dn/dir1:
file1 file2 file3

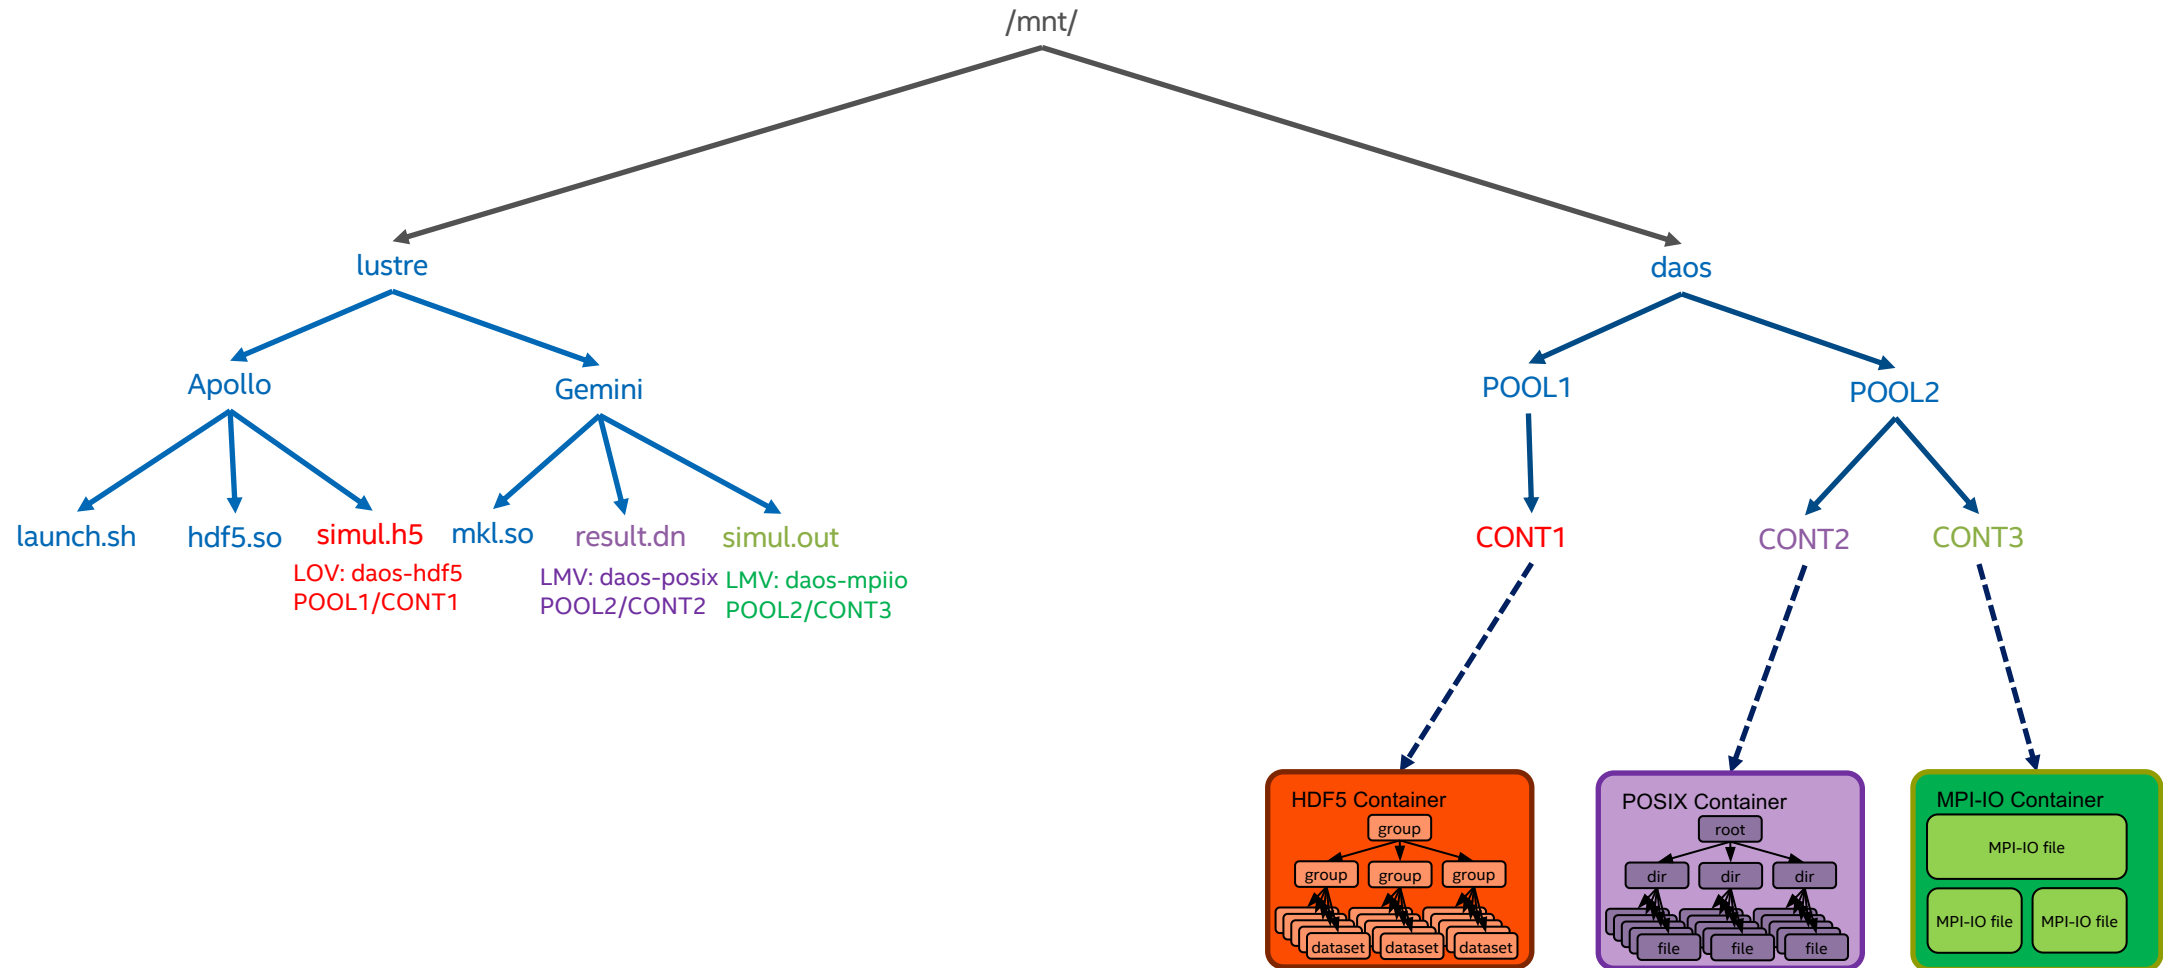
./Gemini/result.dn/dir2:
file1 file2 file3

./Gemini/result.dn/dir3:
file1 file2 file3

./Gemini/simul.out:
file1 file2 file3
jlombard@jprisaca-MOBL:/mnt/lustre$
```

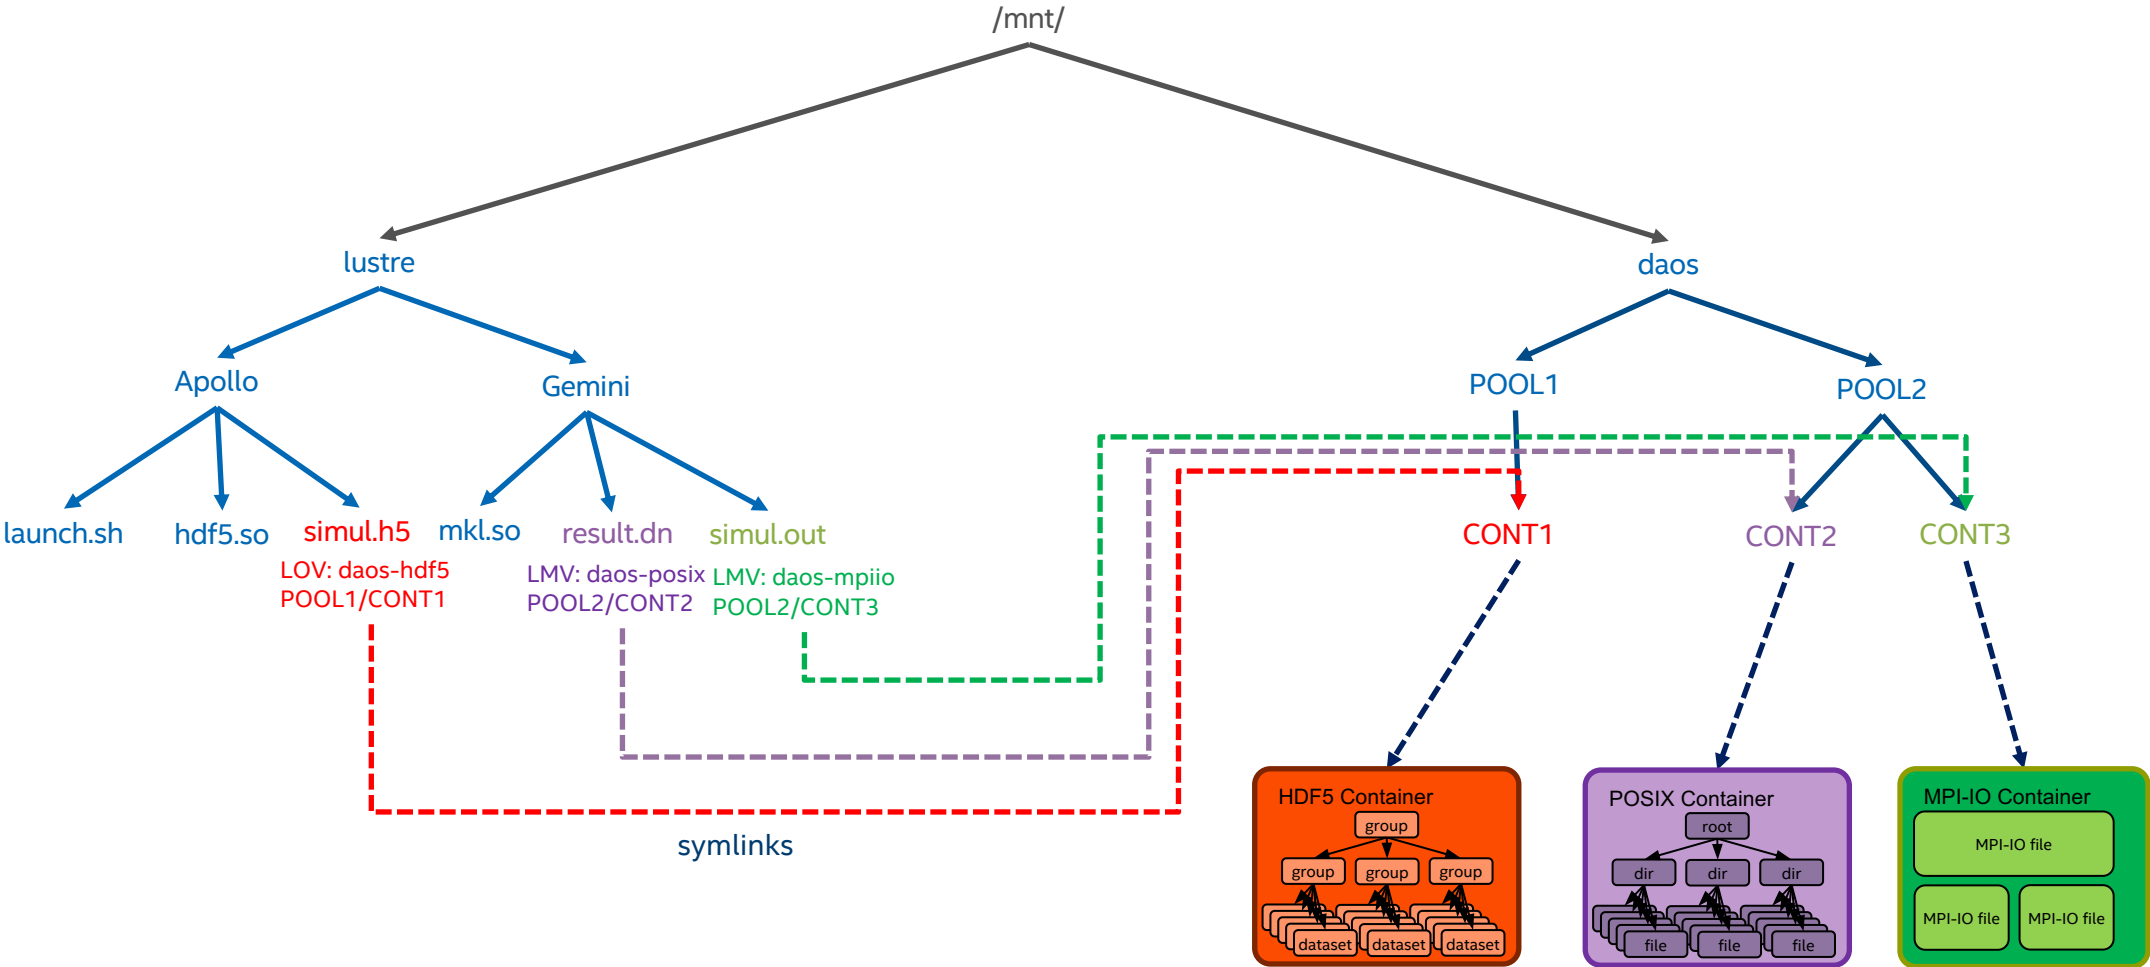


# Unified Namespace Implementation





# Unified Namespace Implementation



# Lustre Changes

- Foreign LOV/LMV
  - On-disk format
  - Lustre API/ioctl() changes
  - Lustre tools changes
  - lfsck compatibility changes
- Expose special files/directories as symlinks
  - Explored auto-bind-mount
  - Client-only changes



Available in 2.13



Available in 2.15

The Intel logo is centered on a solid blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter 'i'. To the right of the word "intel" is a registered trademark symbol (®) enclosed in a white circle.

intel®