Los Alamos NATIONAL LABORATORY

Toward Open Object-Based Computational Storage For Analysis Query Pushdown

Qing Zheng, Jason Lee, Dominic Manno, Gary Grider, Los Alamos National Laboratory

11/12/23

LA-UR-23-32776

3 Things About Scientific Data Analytics

Data is big

Moving data is expensive

Queries often target a tiny portion of a large dataset

Query Pushdown Through Computational Storage

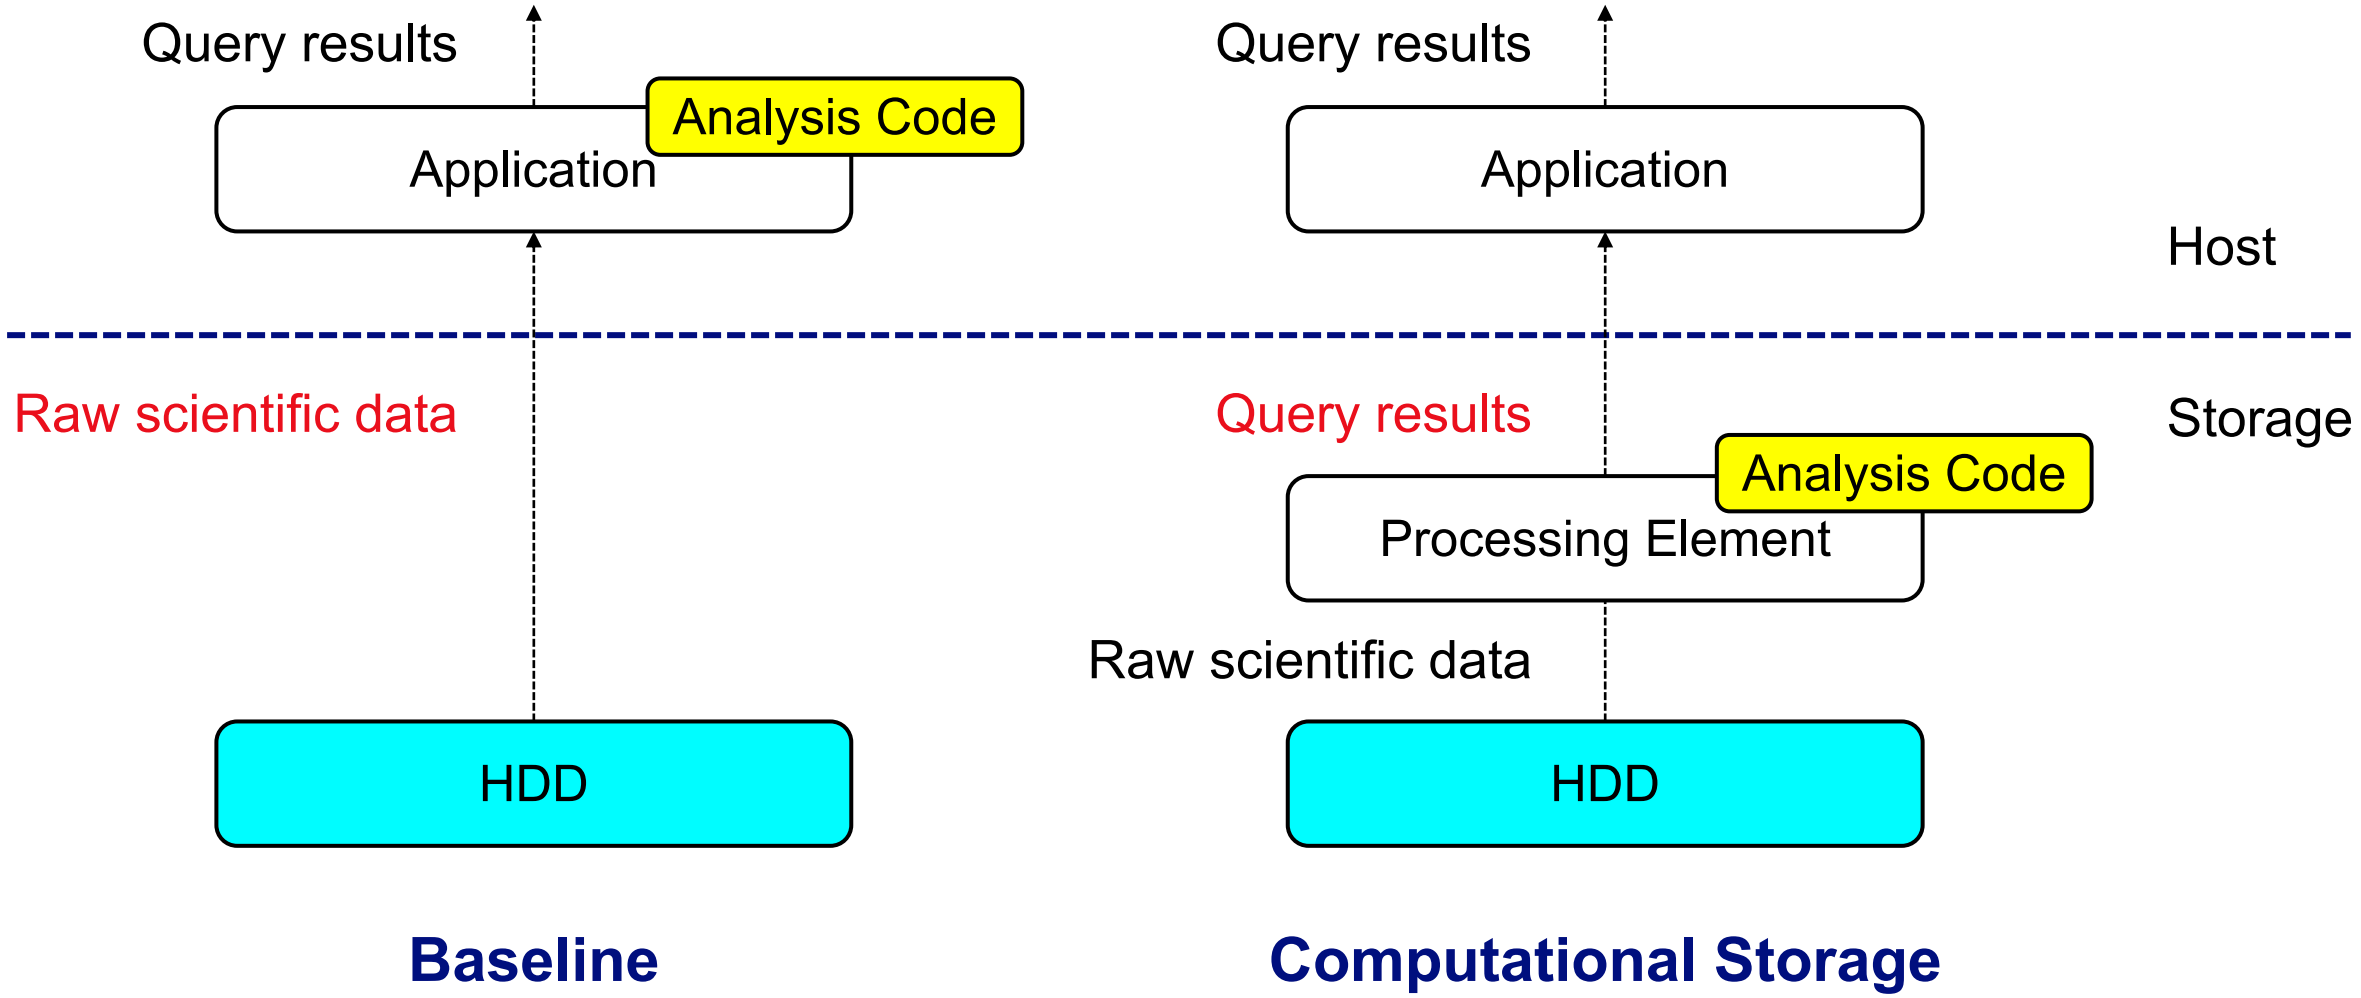

Data Agnostic vs Data Aware Offloads

Data Agnostic

- Storage does not know what's in the data (view data as byte streams)
  - Like what POSIX filesystems do today
- Example offloads: data compression, encryption, custom risc-v, eBPF functions

This effort will use the data aware approach

Data Aware

- Storage and apps agree on a data format (e.g., Apache Parquet) and a query format (e.g., Substrait)

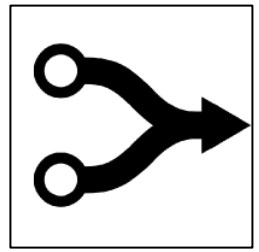

Storage Interface: Block? KV? Object?

Block

- Best for **data agnostic operations** (compression, encryption)

KV

- Best for **row-based** applications such as various PIC (particle-in-cell) codes

Object (each can be a Parquet fragment)

- Enable **columnar** analytics often seen in grid-based codes in addition to PIC codes

Prior work at Los Alamos looked at these (ZIA, KV-CSD, C2) in collaboration with Aeon, Eideticom, Nvidia, SK hynix, Seagate

This effort will use objects

Standardization

Query Pushdown

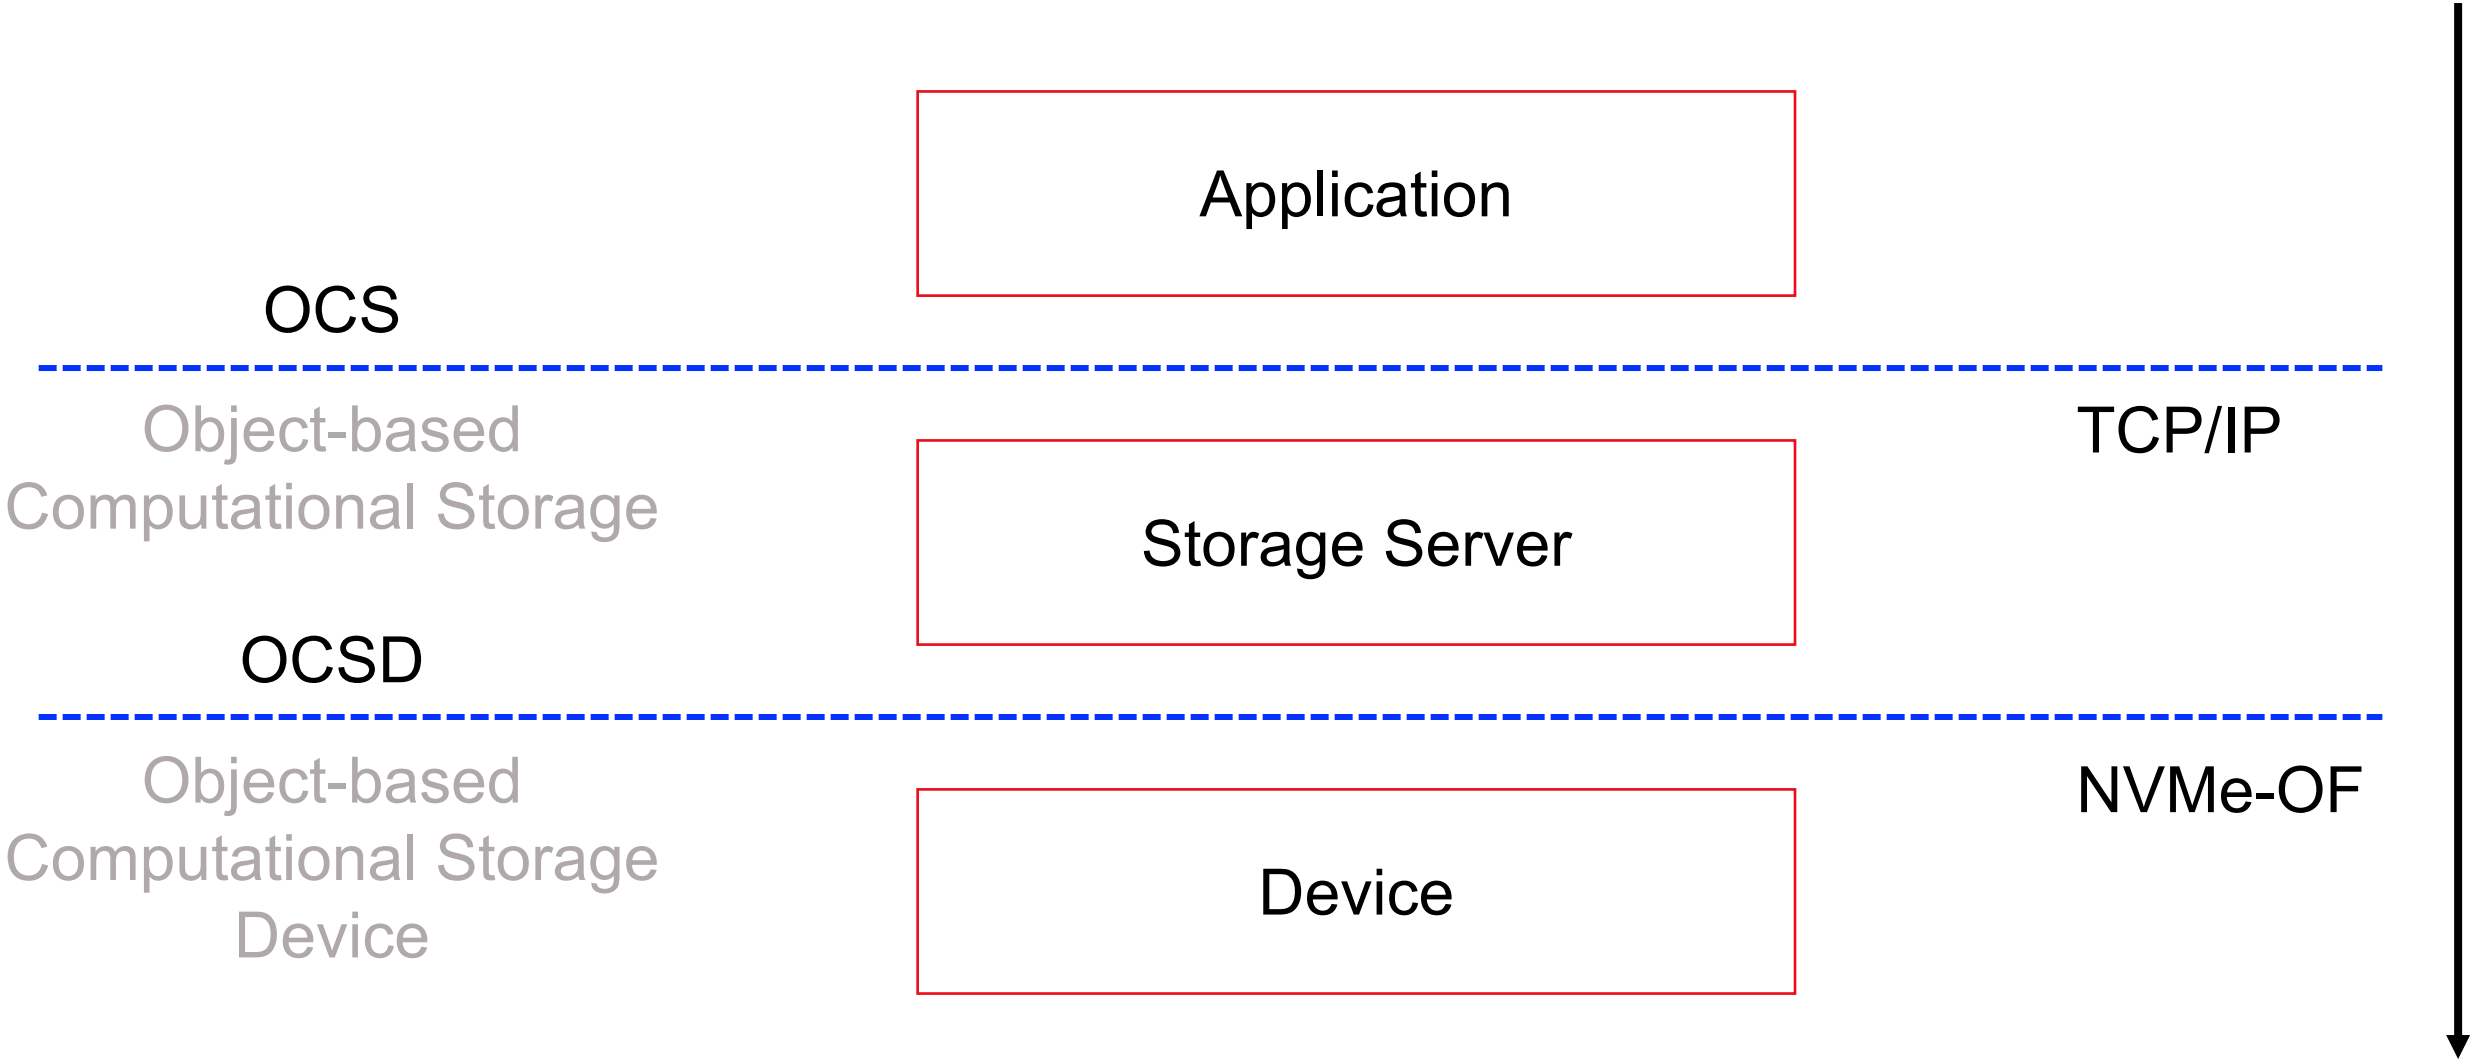

Industry Ecosystem

Query Pushdown

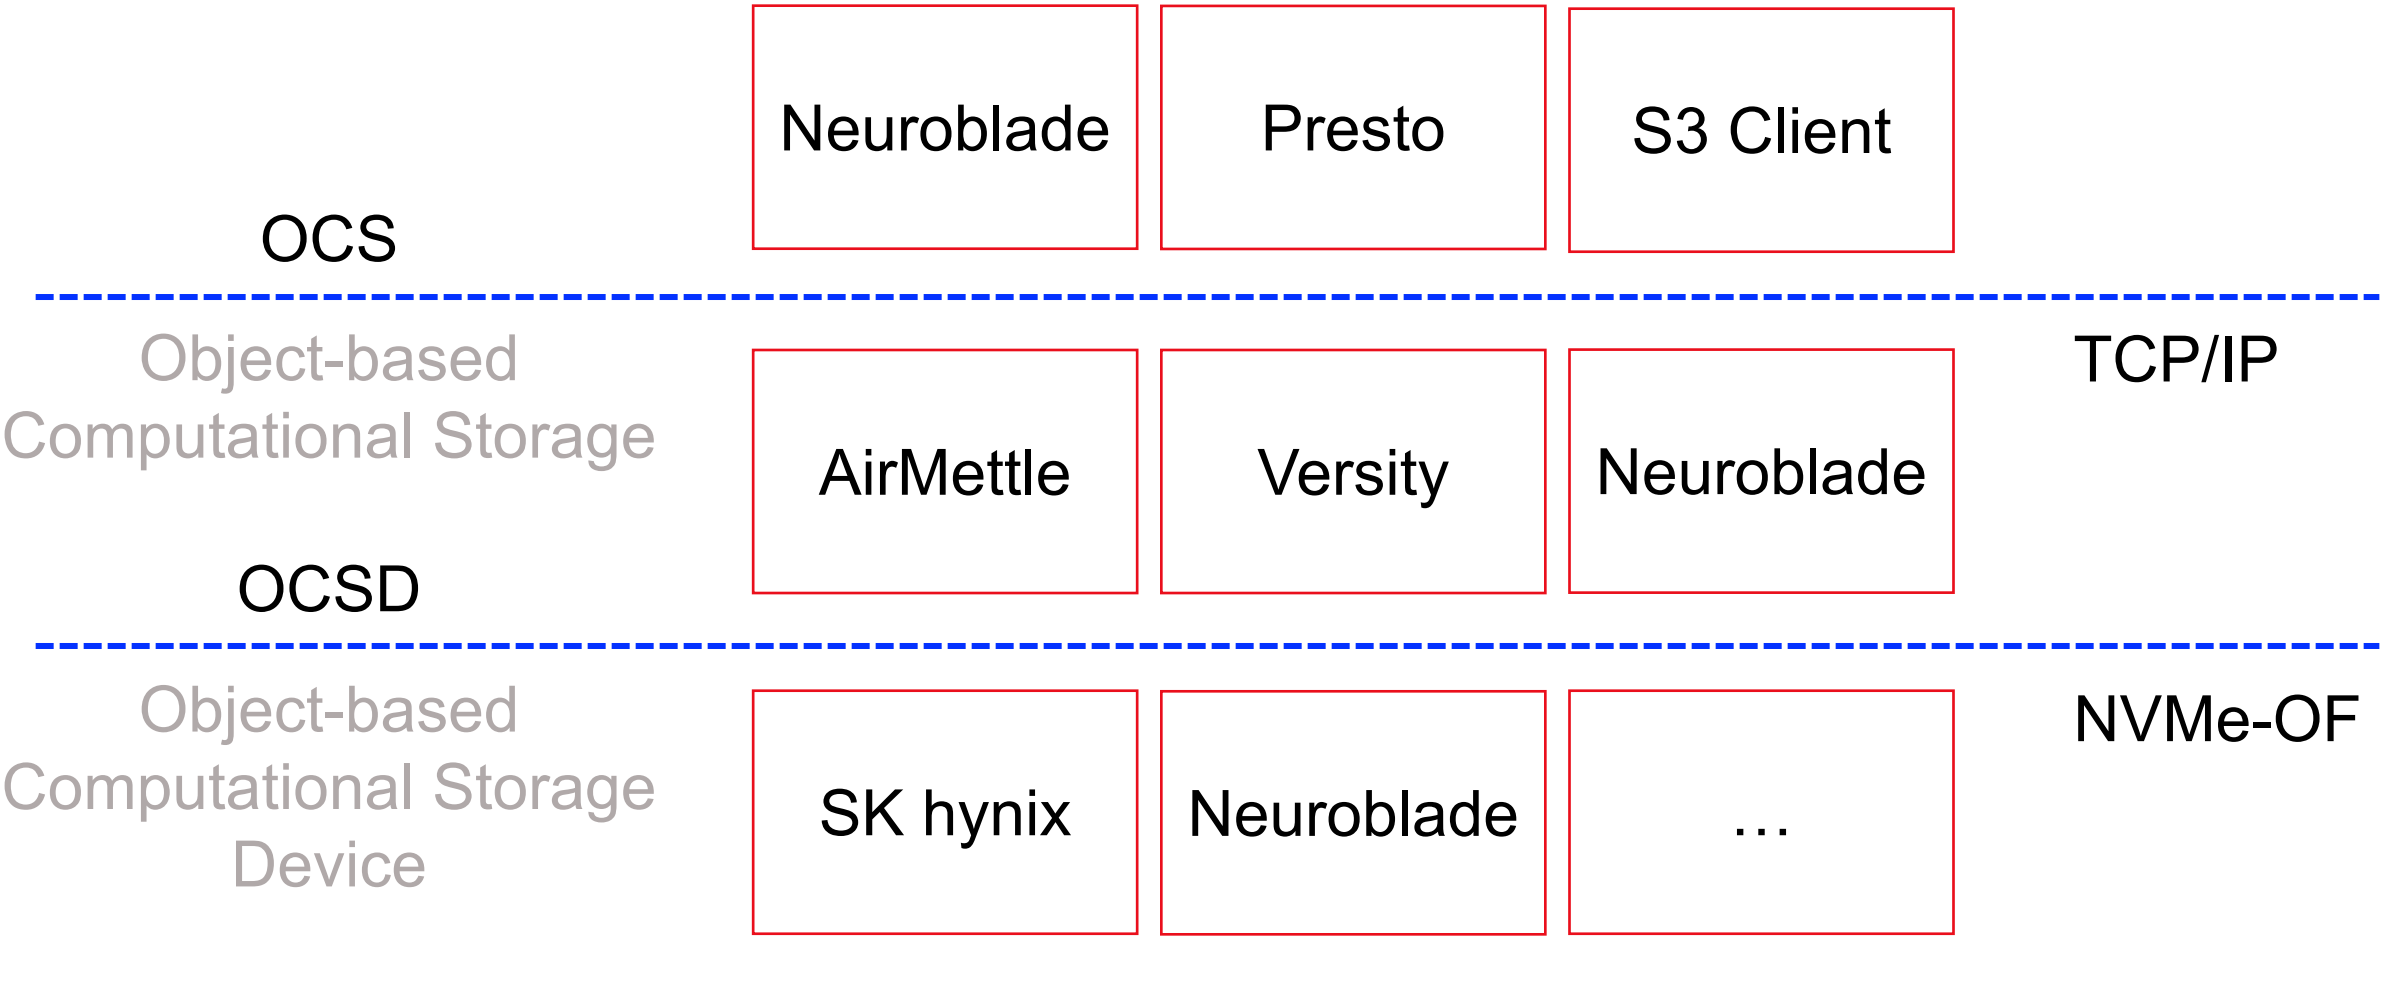


LANL/SK hynix Demo
at Exhibition Hall
Booth #2101

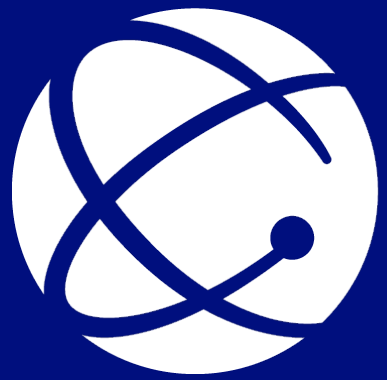

Los Alamos

NATIONAL LABORATORY