SC17 BIRDS OF A FEATHER SESSION

Argonne NATIONAL LABORATORY

ANALYZING PARALLEL I/O

PHIL CARNS
Argonne National Laboratory

carns@mcs.anl.gov

JULIAN KUNKEL
German Climate Computing Center (DKRZ)

kunkel@dkrz.de

November 16
Denver, Colorado

MOTIVATION FOR ANALYZING PARALLEL I/O

Or… why is this still such an important topic?

I/O performance is critical to application performance (and variability), but it is getting even harder to understand.

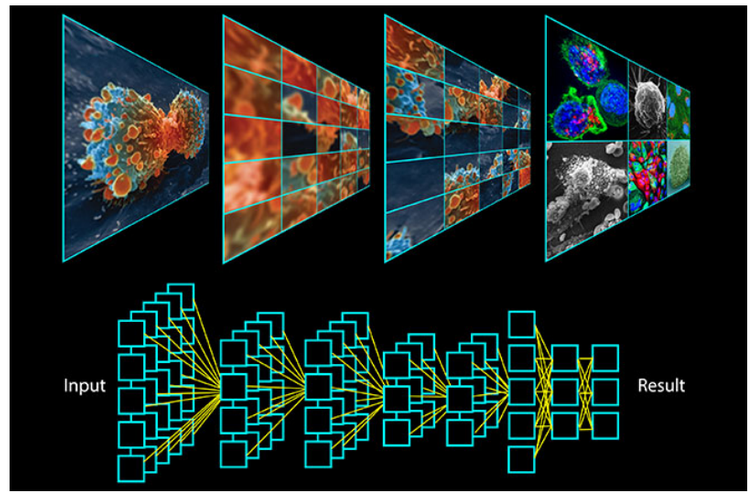

https://www.exascaleproject.org/candle-accelerates-solutions-toward-three-top-cancer-challenges/

What do HPC applications look like?

- Conventional MPI, workflows, deep learning, data analytics, and more.

New existential question: what *is* I/O and where does it happen?

- File system? External store? On-node? Mmap/NVRAM? Wide area?

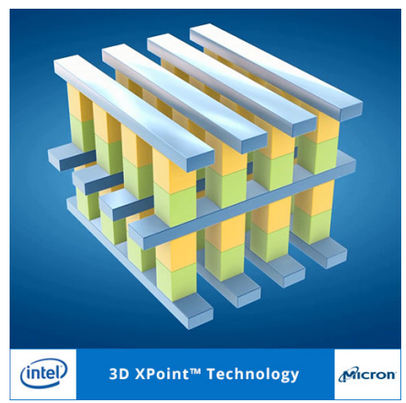

Argonne NATIONAL LABORATORY

KEY CHALLENGES

- **Instrumentation**:
  - What do we measure and how at the app level?
  - What do we measure and how at the system level?
- **Analysis**:
  - How do we correlate instrumentation data?
  - How do we extract knowledge from data?
- **Application**:
  - Develop best practices and tune applications
  - Improve system software
  - Design and procure better systems

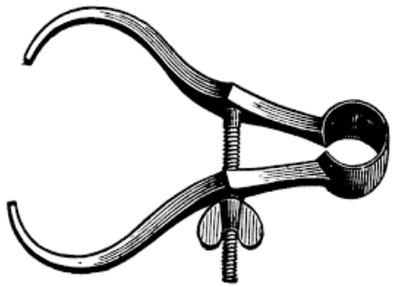

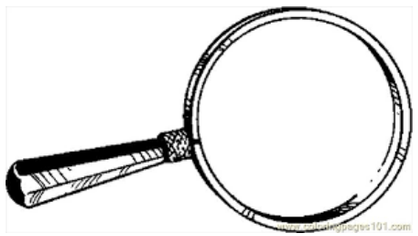

The SC17 and PDSW-DISCS programs have several examples of work in this area (too many to list this year)!

WHAT WILL WE ACCOMPLISH TODAY

- Guest speakers will highlight examples of the start of the art in art in tools, techniques, and applied results:
  - Glenn Lockwood (LBL)
  - Ross Miller (ORNL)
  - Eugen Betke (DKRZ)
  - Wolfgang Frings (JSC)

- Community discussion
  - Questions, suggestions, and comments
  - Ask the speakers about their work
  - Identify gaps and areas for improvement in the field
- Your questions and feedback can have a real impact on the community
- We will summarize discussion in a public report

DISCUSSION

Help the community to solve ongoing challenges

- What should we do about byte-addressable NVM?
- How do we motivate vendors to make I/O monitoring an integral part of their products?
- Should storage monitoring include archive (tape) activity?
- Is there a need to record data lineage?
- What kinds of applications are not well served by current tools?
- What kind of information about I/O accesses is of interest for users?
  - Do we need different tools for users and facility operators?

https://wr.informatik.uni-hamburg.de/events/2017/bof-analyzing