LA-UR-21-31669
Approved for public release; distribution is unlimited.

Title: An Intro to HPC @ LANL

Author(s): Bonnie, David John

Intended for: Presentation to students

Issued: 2021-11-29

Los Alamos
NATIONAL LABORATORY

An Intro to HPC @ LANL

David Bonnie – HPC-DES

11/30/21

LAUR-123456789

Who am I?

- Storage architect – making sure the bits go where they should, when they should
  - From archive to scratch, faster and safer
- Scientist at LANL since 2013
  - Student in 2009 @ LANL, 2010 @ ORNL, then again at LANL in early 2013
- Undergrad / grad at Clemson, Computer Systems Architecture
- Technical lead of Enterprise Backups, Future Archive, and Campaign Storage
- Has way too much fun building cars and old jet skis

What is HPC?

- In particular – simulation HPC

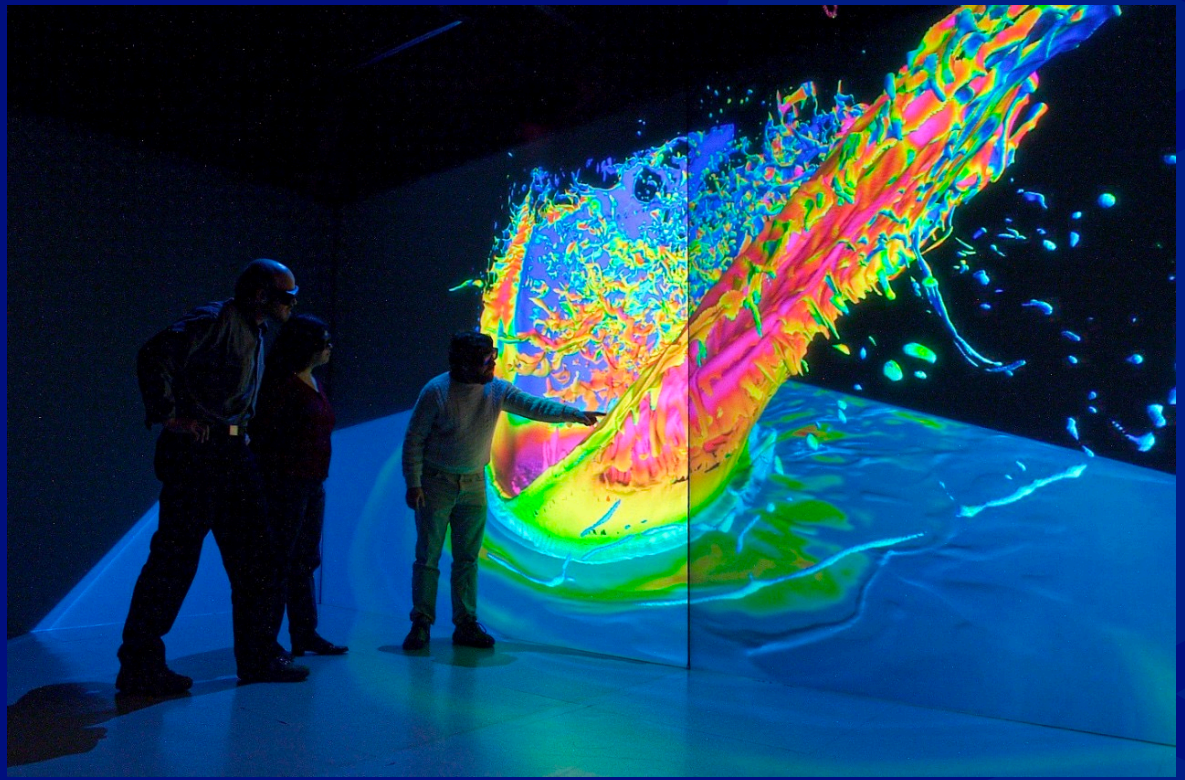

What's so difficult?

- Breaking up the problem
- Synchronization
- Crashes
  - by proxy, restart
- Precision
- Visualization
- Scale
- Scale
- Scale

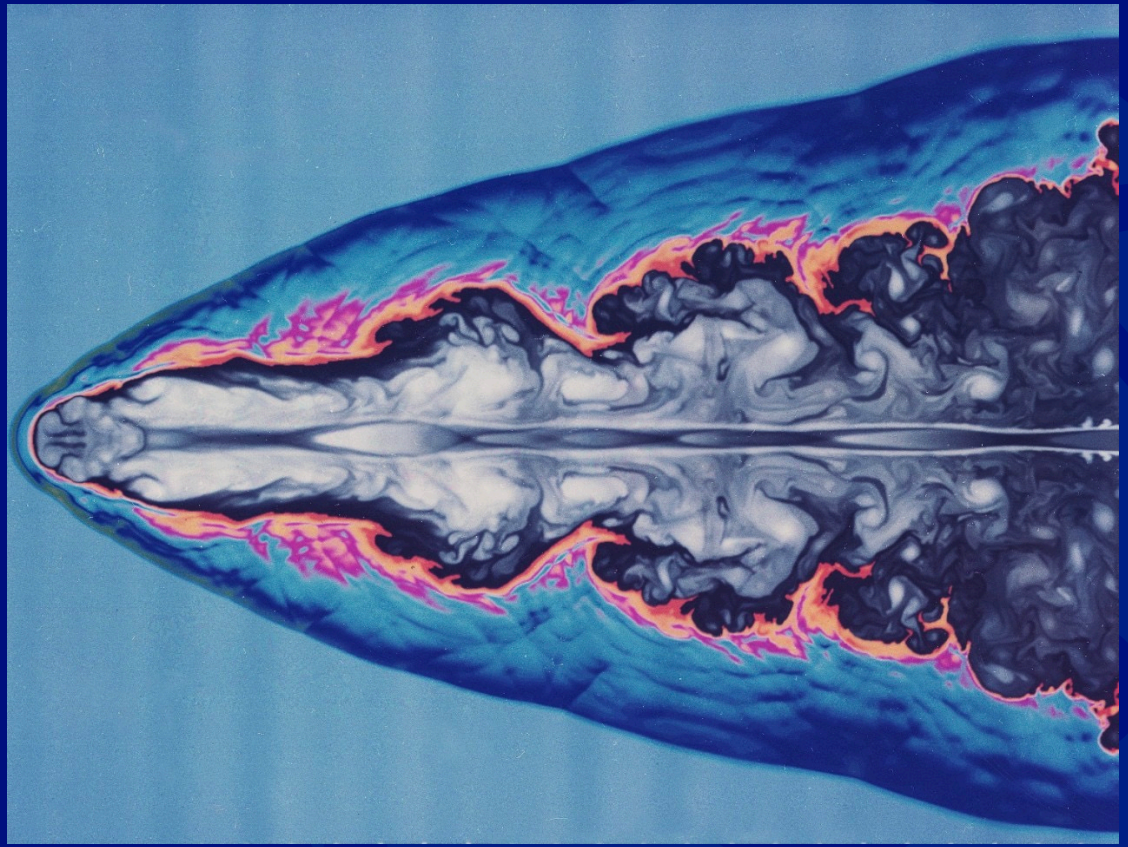

Speed of light problems

- Sheer size!
- 2 PiB RAM
- 20,000 nodes
- ~1m cores
- 112 cabinets
- ~4 PiB Flash
- ~80 PiB HDD
  - (another 40 racks)


CRAY
Alliance for Computing at Extreme Scale
Los Alamos
NATIONAL LABORATORY
Sandia National Laboratories
TRINITY
DANGER
UNAUTHORIZED PERSONNEL KEEP OUT

Not a small number of wires…

- Miles and miles of cables!


2 TON

Beyond just physical size -

- Scaling is difficult and creates new issues to overcome
  - How do you boot tens of thousands of nodes?
  - How do you start up a job?
  - How do you monitor status?
  - How do you find what is broken?
  - How do you even wire it up?
  - SCALE

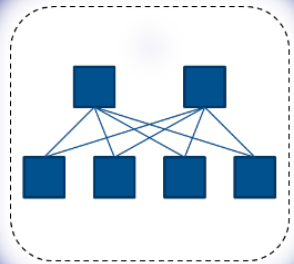
Fat Tree

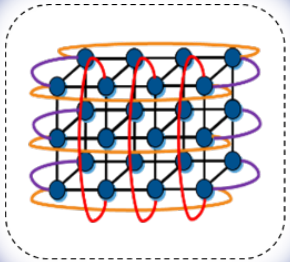
Torus

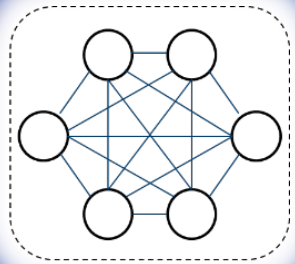
Dragonfly

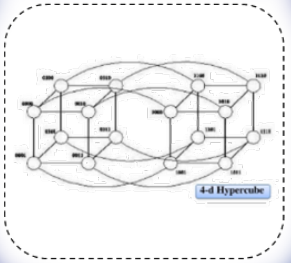

Hypercube

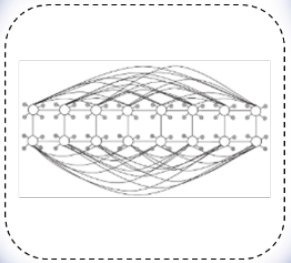
HyperX

Checkpoint/Restart

- Cosmic rays, failing hardware, power blip, bug in application, etc all cause jobs to crash and require restarting
- 2 PiB of RAM state to dump and reload
- 1 TB/s at this scale still requires almost 35 minutes to save!

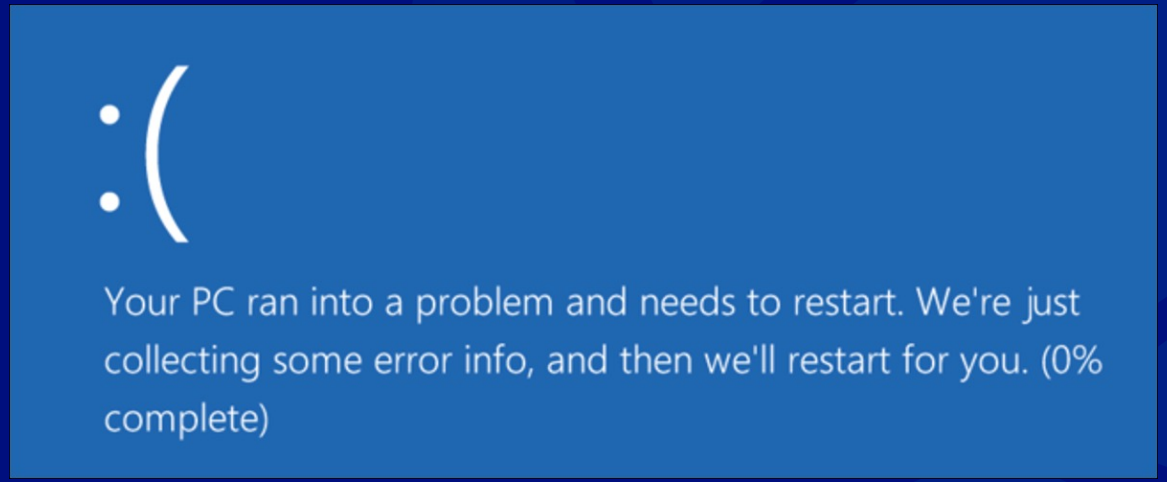

Specialized…everything

- Storage, in particular, is weird at this scale

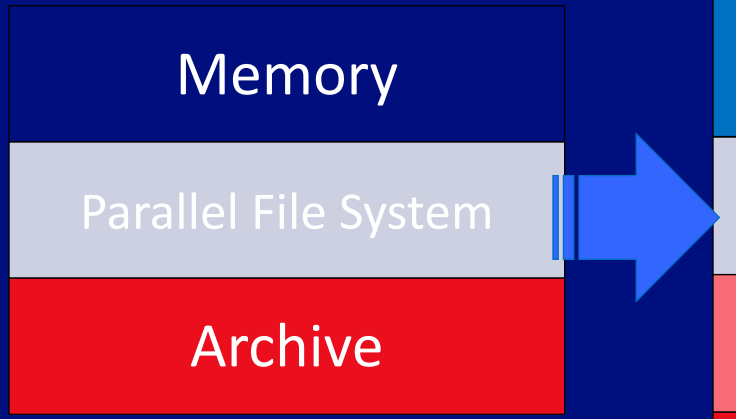

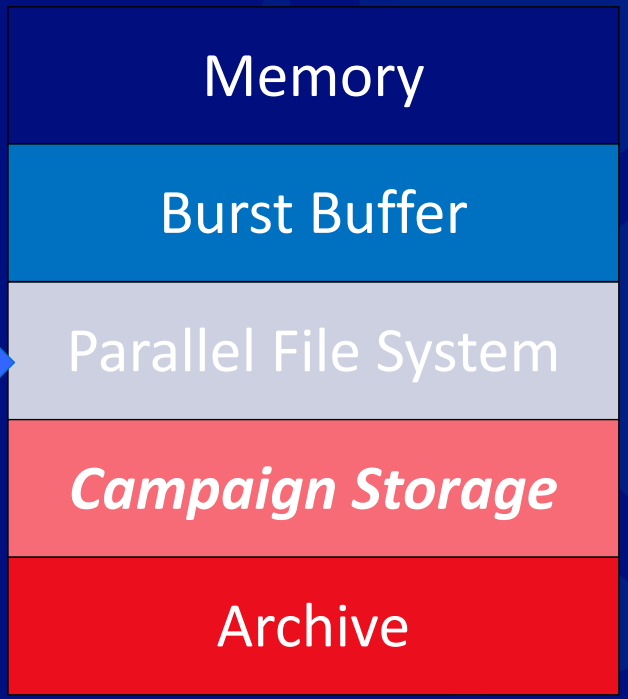

1-2 PB/sec
Residence – hours
Overwritten – continuous

4-6 TB/sec
Residence – hours
Overwritten – hours

1-2 TB/sec
Residence – days/weeks
Flushed – weeks

100-300 GB/sec
Residence – months-years
Flushed – months-years

10s GB/sec (parallel tape)
Residence – forever

How do you store a 1 PiB file?

- Around 100 HDDs at minimum
- …but 100 HDDs aren't fast, so we need lots of HDDs!
- …but now we have to coordinate and break the file into smaller pieces
- What happens when a disk fails?
  - So they need protection of some sort
  - How do you protect across thousands of HDDs?
- What about long-term storage?

How do you move a 1 PiB file?

- …now we need a cluster, just to move the data around at a reasonable speed
- …which uses an MPI job, just like the simulations, to coordinate data movement
- …which breaks up the big files into work chunks for bandwidth
- …which also allows restarts if we keep track of state!
- (SCALE)

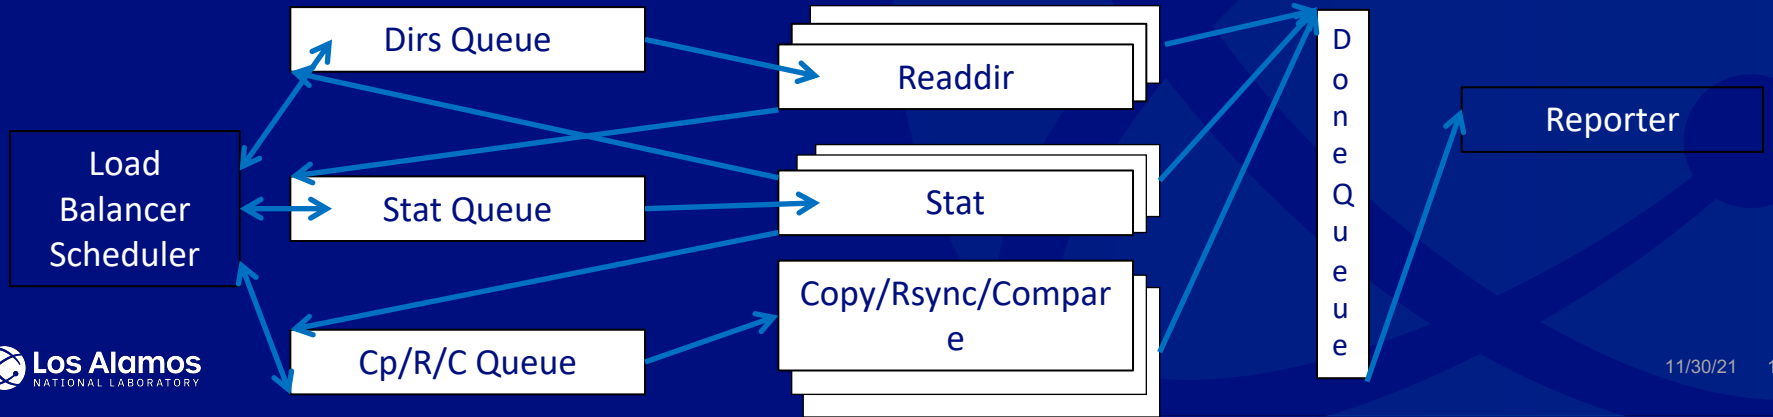

HPC is difficult!

- You have to get in the habit of thinking about everything at ludicrous scale
- Simple things become very difficult engineering problems
- Simple failures can cause cascade effects that make it hard to troubleshoot

- Even something simple like our data movement tool!
  - A singly linked list became the bottleneck for data movement

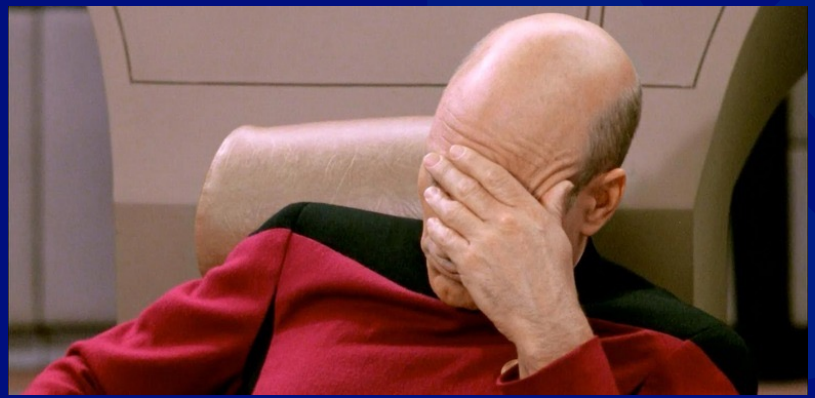

- Information Science and Technology Institute (ISTI) Summer Schools:
  - Applied Machine Learning Research Internship
  - Cyber Security School
  - Quantum Computing School
  - Supercomputer Institute
  - Codesign School
  - Data Science at Scale School
  - Parallel Computing Research Internship
  - Computational Physics Workshop

Be part of something extraordinary

Los Alamos
NATIONAL LABORATORY

Feel free to reach out!

- My email: dbonnie@lanl.gov
- HPC Student Liason: Julie Ann Wiens - jwiens@lanl.gov