Answering the Really Important Questions

(a few of them)

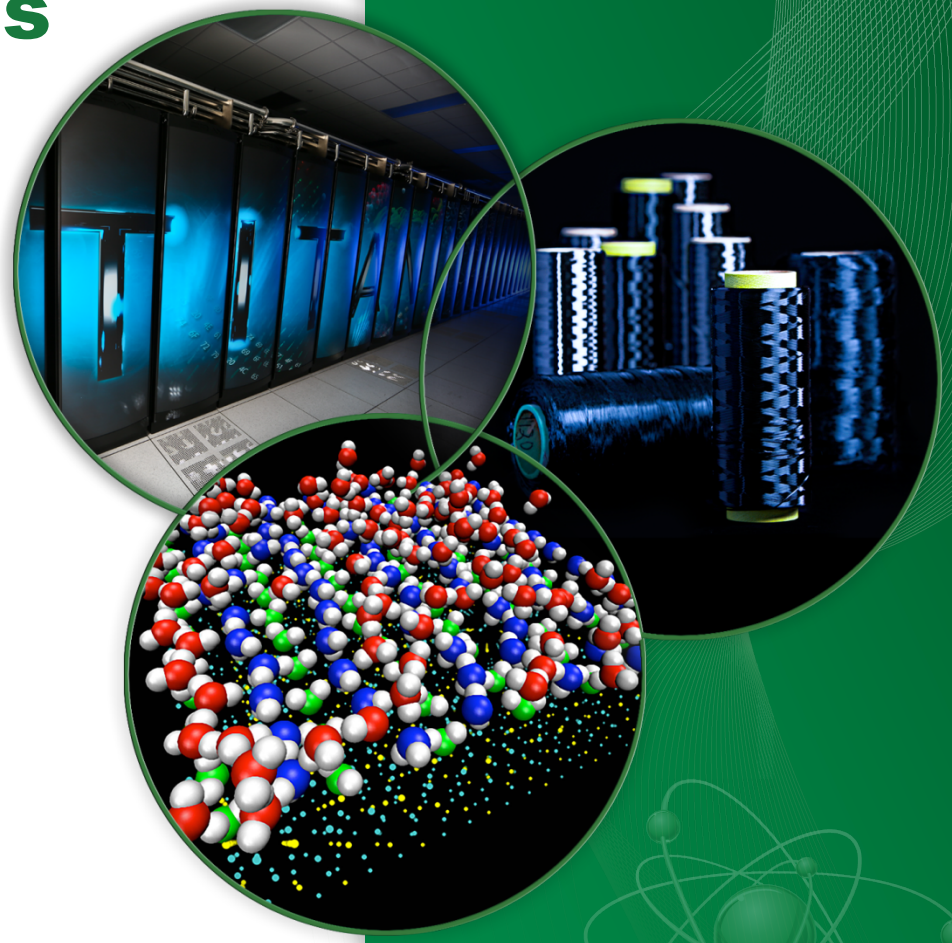

Al Geist
Corporate Fellow
Oak Ridge National Lab

CCDSC 2014
Lyon France
September 3, 2014

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory

Fault Tolerance -- Everything's fine What a/b AI's doom and gloom?

- Fundamental assumptions of applications and system software design did not anticipate **exponential growth in parallelism**
- **Fault rate proportional to number of components.** Jaguar and Titan have very different components and different number of cores, but approximately the same number of components.
- Memory is a special case. It is prone to cosmic ray errors proportional to area and circuit design. **Jaguar saw ECC bit flips at rate of 350/min (1 flip/min/TB)**
- Today's apps rely on checkpoint/restart and systems have improved RAS to handle increased fault rate. **Titan loses a node every 1.5 days but system hasn't crashed in over 7 months!**

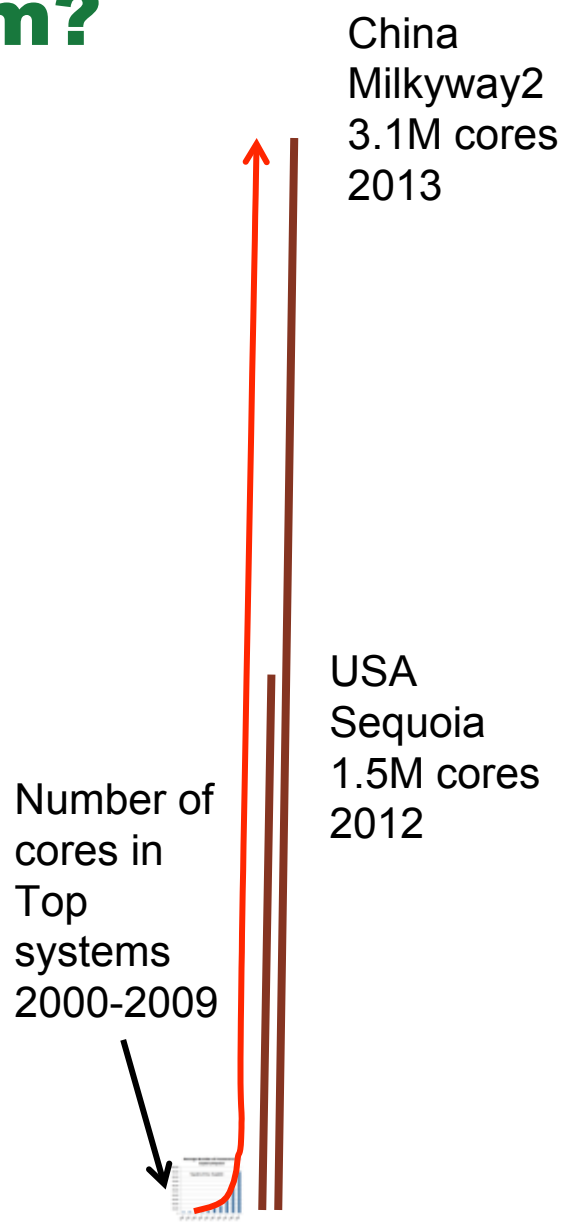

Fault Tolerance – Don't Worry, Be Happy! . . . Sort of

- **Chekpoint will get a huge boost as NVRAM on node** becomes the norm. Time to chekpoint drops to only a few seconds to a few minutes.
  - This provides the ability to have a higher chekpoint rate. High enough even for exascale… BUT
- **Increased danger of wrong answers** as Undetected error rates increase which don't trigger a restart and can get written out into the chkpt file.
- **Need to reduce undetected error rate** through improved detection in HW & SW

I didn't want the (grapes) right answer anyway

Ideal Exascale Research Program given a budget of 100 Million Euro

- **Step 0 Don't do what the USA has done**.
  - Exascale Plan delayed to FY16 due to two restarts (long story)
  - $100M sent to vendors to do "research", which vendors admit will not be used in their exascale systems.
- **Step 1 Set up long-term partnerships** between your major Computer Centers and vendors
- **Step 2 Have vendors (or vendor consortia) develop an viable roadmap** to exascale based on the needs and constraints of the Computer Centers and their users
- **Step 3 Fund vendors and research community** to do the long-range research needed to address power, resilience, productivity of systems on the roadmap
- **Step 4 Centers negotiate with their partner vendors** to procure a series of ever bigger systems on that vendor's roadmap. Providing users a long-term common environment.

When will we have a Linpack-exaflop machine?

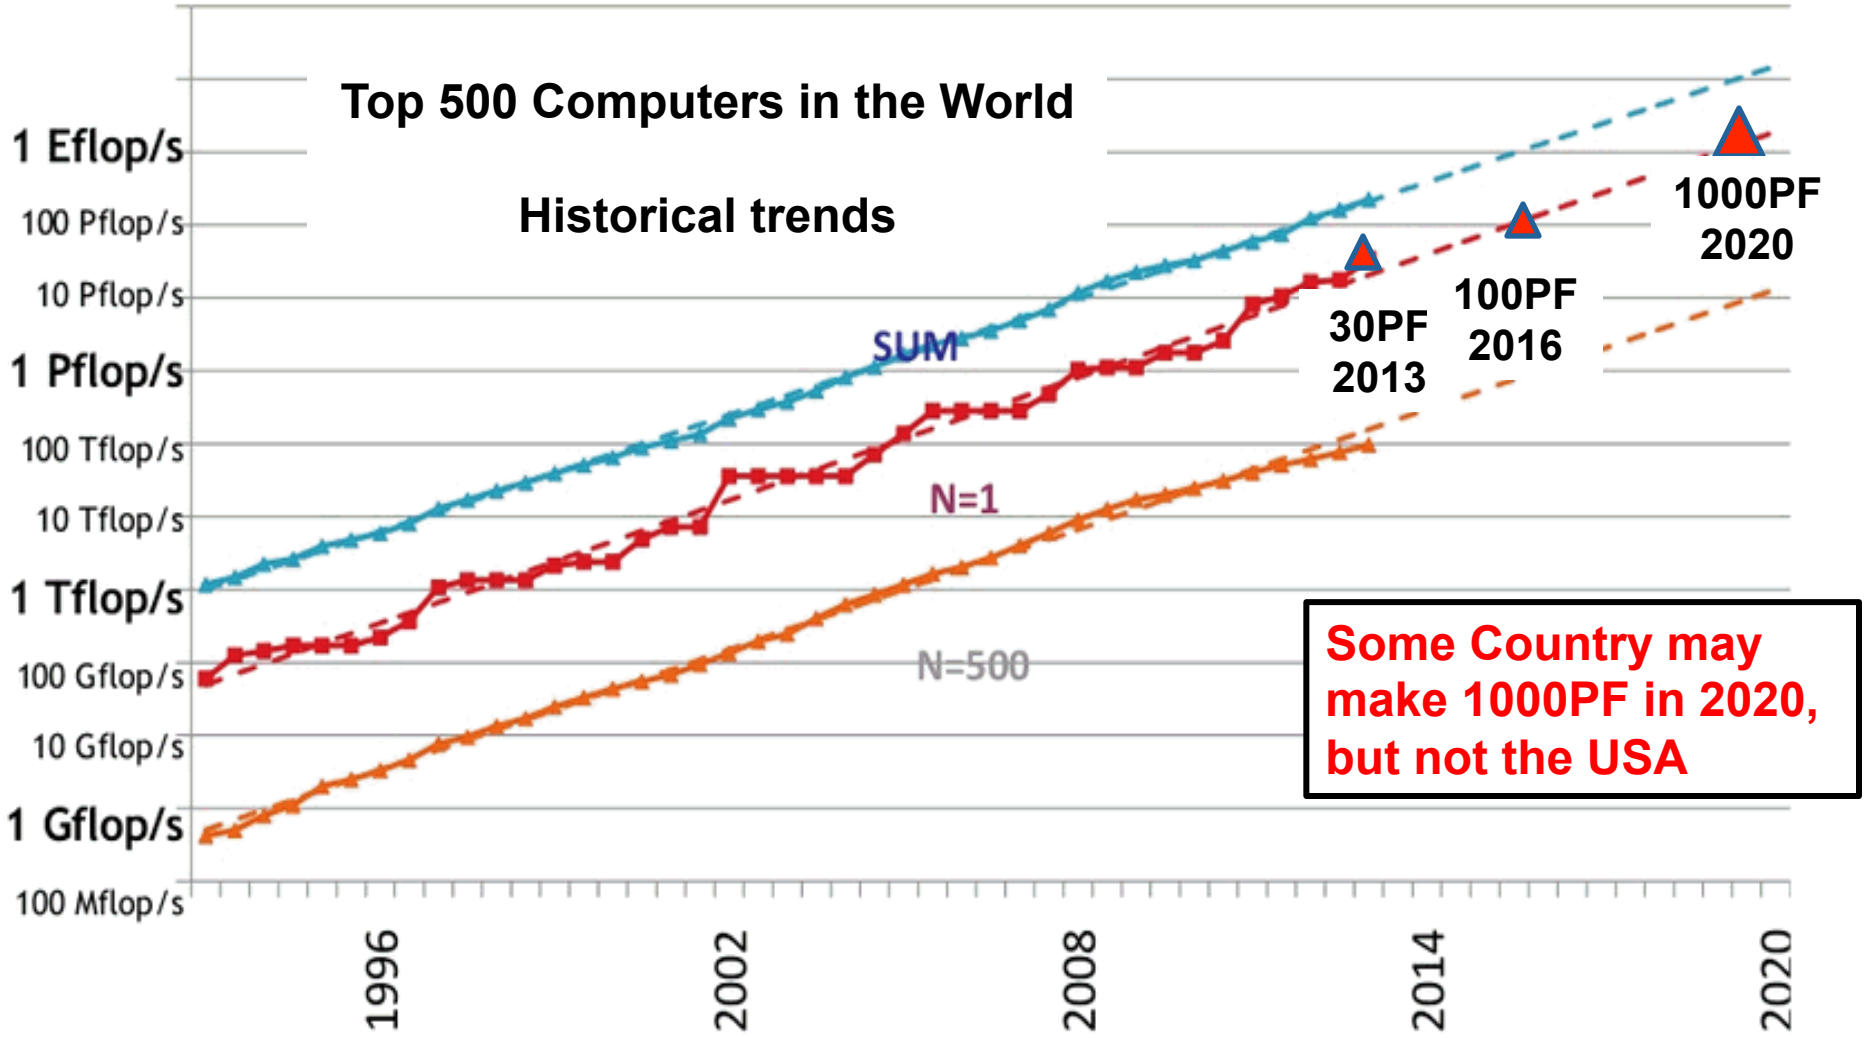

Exascale in the USA not until 2022

DOE Facilities have a fixed 4-5 year cadence
Present Roadmap for Largest US supercomputers 2012 - 2022

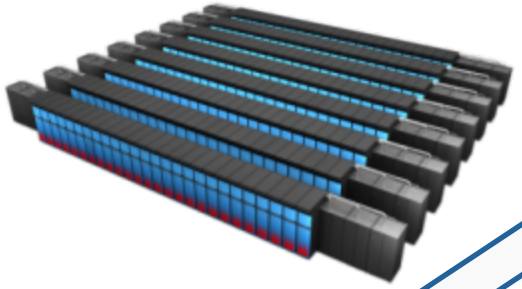

2022 CORAL-2 1000 PF

2019 Trinity-2 250-300 PF

2017 CORAL 100-200 PF

2015 Trinity 60 PF

2012 Titan 26 PF and Sequoia 20PF

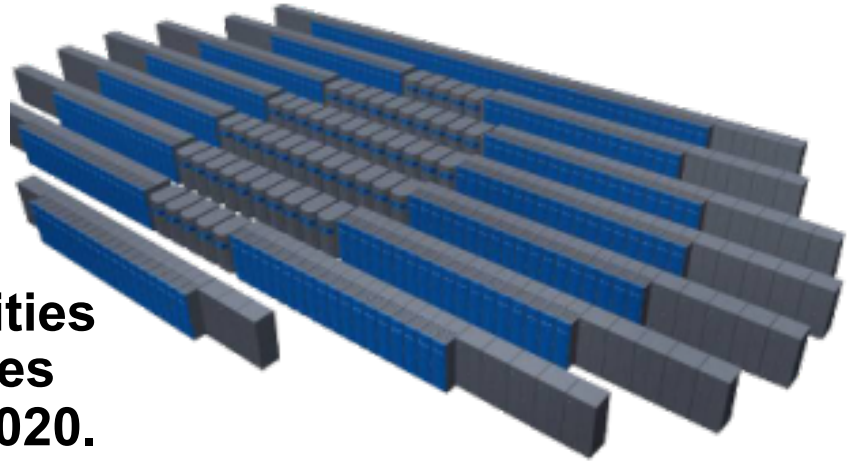

Power constraints of 20-30 MW facilities and pay-off schedules of 4 year leases limit accelerating this Roadmap to 2020.

What will it look like Architecturally? Physically?

When the first application reaches 1/4 exaflop on this?

ORNL's Jaguar was first computer to run application at sustained 1 PF
- It took less than 2 months after delivery for this to happen

Physically

- 300-400 cabinets
- Consume 25-30 MW power
- Likely the last generation of HPC in CMOS (5-7 nm feature size)

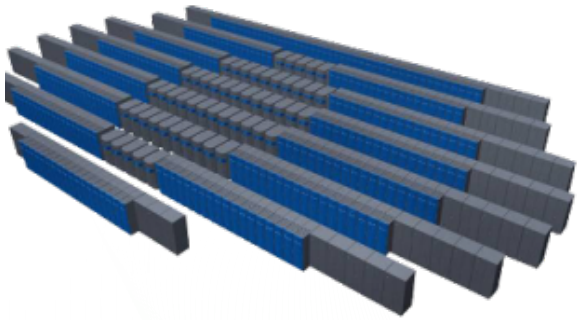

Architecturally

Two diverse trends:
- Sea of many-core CPUs with millons of nodes and a billion cores
- Sea of GPUs controlled by few CPUs with 100,000 large nodes

Big change in memory architecture
- High Bandwidth stacked memory and NVRAM on all these nodes

Another Really Important Question

- What if we do “Nothing”?

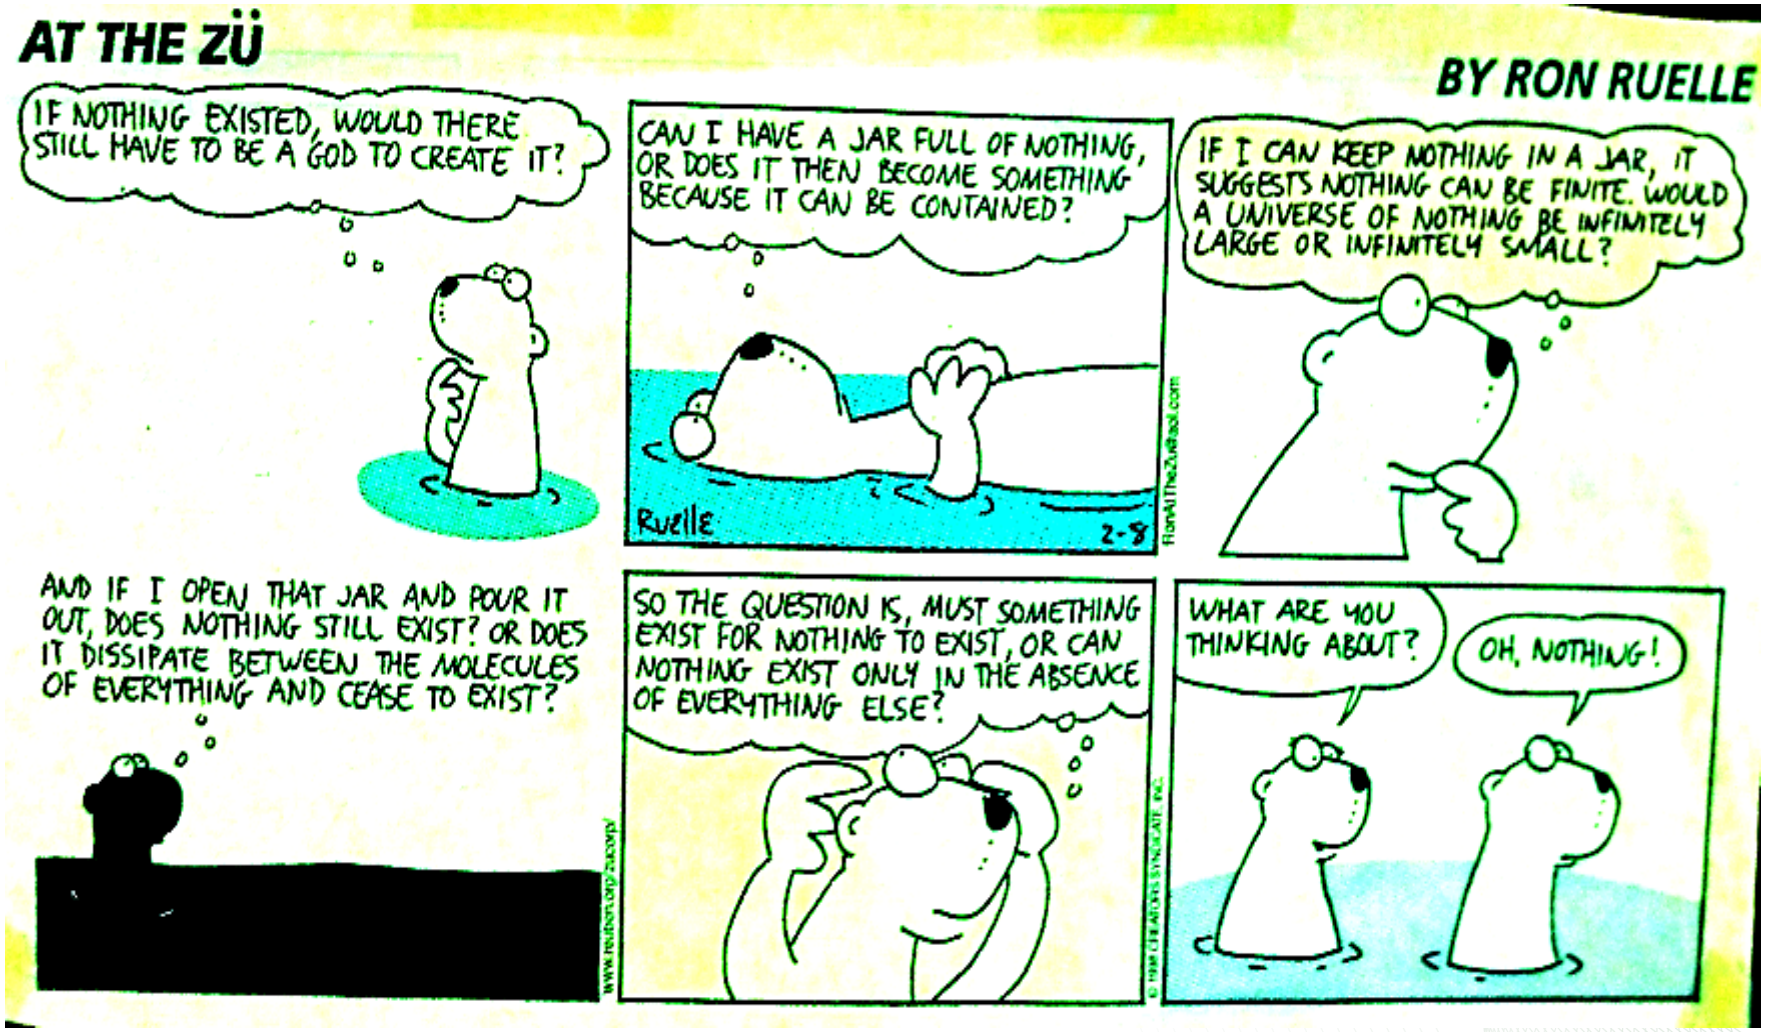

What has AI been up to? Not so important
CORAL Collaboration ORNL, ANL, LLNL)

Current DOE Leadership Computers

Titan (ORNL)
2012 - 2017

Sequoia (LLNL)
2012 - 2017

Mira (ANL)
2012 - 2017

Objective - Procure 3 leadership computers to be sited at ANL, ORNL and LLNL in CY17

Leadership Computers RFP requests >100 PF, 2 GB/core main memory, local NVRAM, and science performance 4x-8x the max(Titan, Sequoia)

Approach

Competitive process - one RFP (issued by LLNL) leading to 2 R&D contracts and 3 computer procurement contracts

For risk reduction and to meet a broad set of requirements, 2 architectural paths will be selected

Once Selected, Multi-year Lab-Awardee relationship to co-design computers

Both R&D contracts jointly managed by the 3 Labs

Each lab manages and negotiates its own computer procurement contract, and may exercise options to meet their specific needs

Understanding that long procurement lead-time may impact architectural characteristics and designs of procured computers