BLUE WATERS

SUSTAINED PETASCALE COMPUTING

NERSC's Contributions to the World's Supercomputing Community

Dr. William Kramer
Blue Waters Director, NCSA @Scale Program Director
National Center for Supercomputing Applications, University of Illinois

NCSA NSF GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION CRAY

My Random Thoughts over 100,000,000+x Performance Improvement



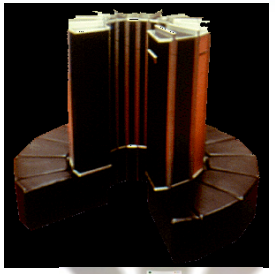




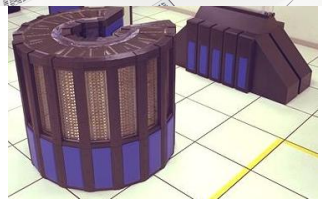




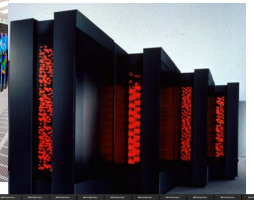
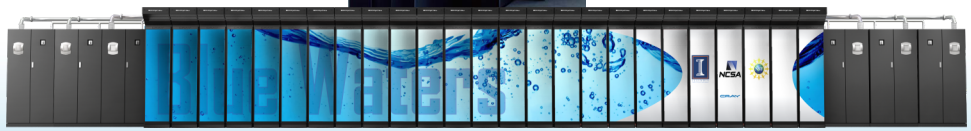

"Old Timer" Prerogatives

- These are my thoughts based on
  - Many supercomputers, storage systems, facilities, users, projects, …
  - Three world class HPC organizations
  - Working in industry, national centers, government, academia
  - Many relationships and colleagues
- I may have some details wrong – but I am sure the message is correct
  - We old timers may not remember things correctly, but we do remember the important things to pass along
  - So, let me know the correct details later if you will
- I will comment on what I am interested and what I remembered
  - If you are not interested in it, sorry
- I do not want to demean any of the many other great HPC organizations and their contributions

But this is NERSC's birthday today after all

My Definition of "NERSC"

- To me, NERSC means
- NERSC Clients and Users
- NERSC Philosophy
- NERSC Staff
- NERSC Resources
- NERSC Funders
- NERSC Leaders
- NERSC Academic Partners
- NERSC Associates, Collaborators and Friends

At SC 07, Dr. Ray Orbach, DOE's Undersecretary for Science, said in his plenary session

"NERSC was … and is today, the best run … computing facility in the world"

Historical Perspective

National Production Supercomputing Programs

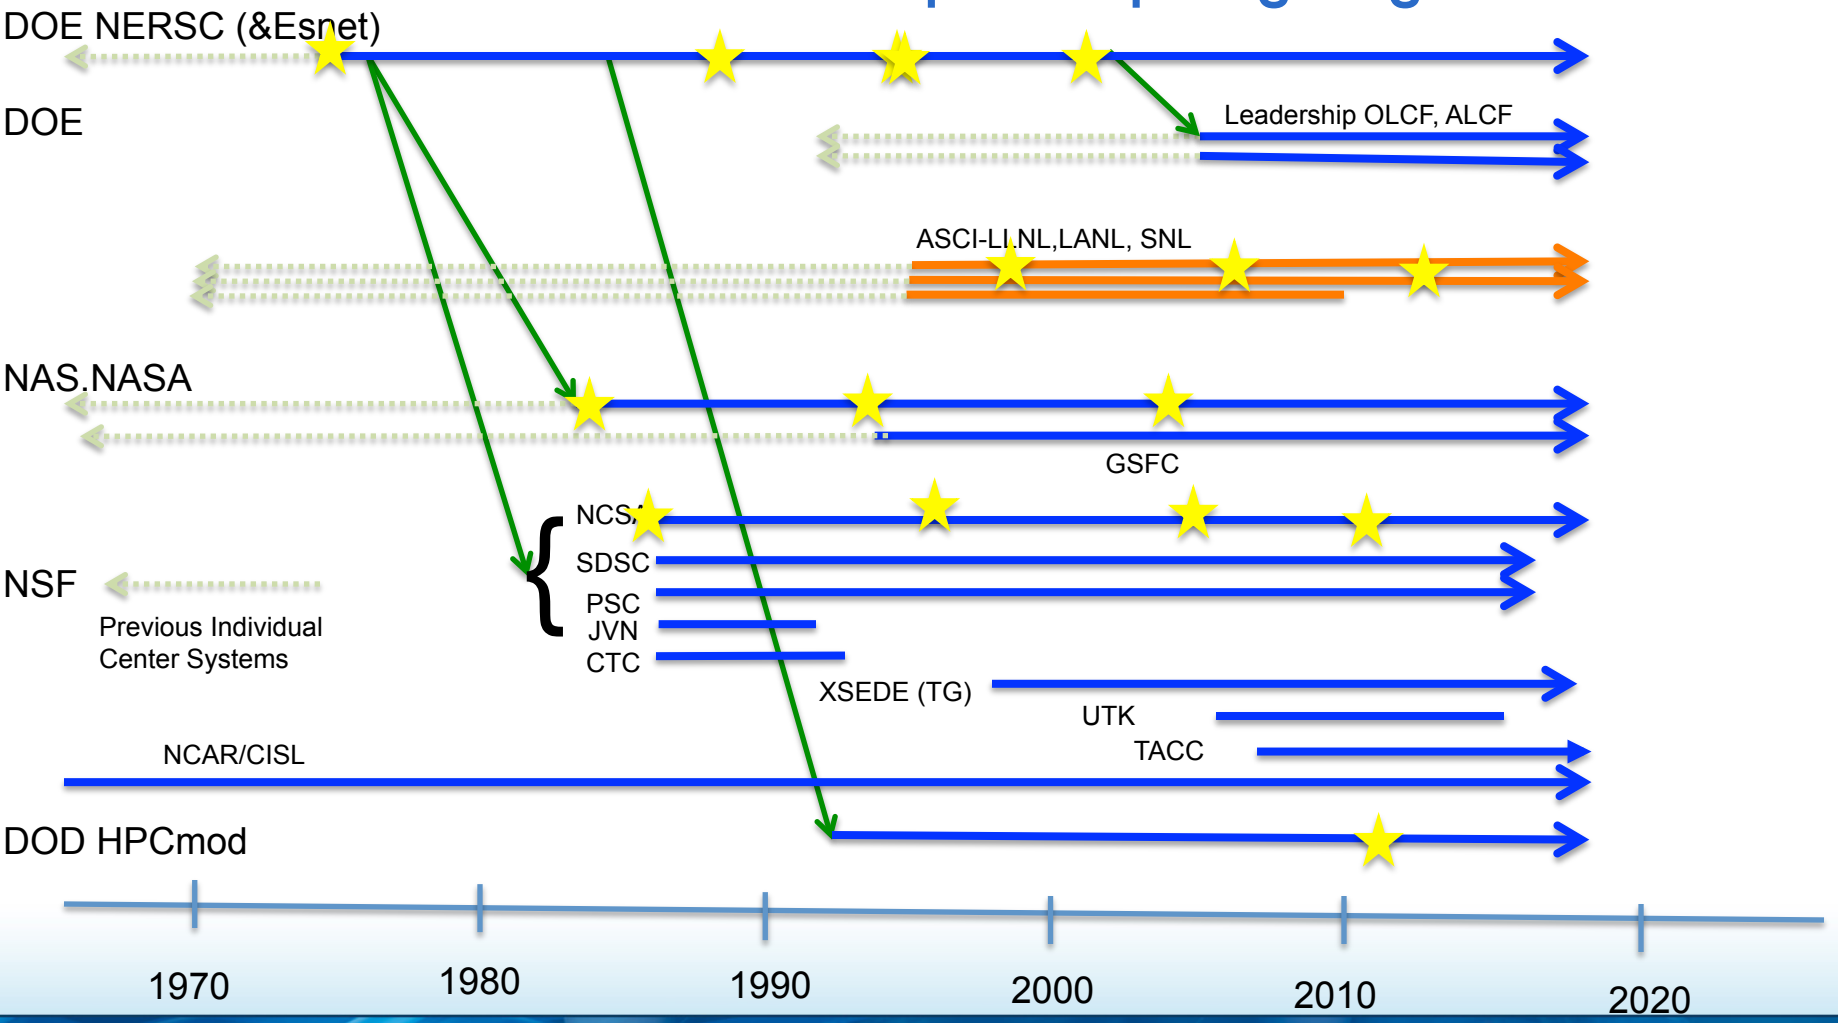

NERSC's Greatest Legacies

NERSC is the First Multi-discipline Nationally based HPC resource for Diverse, Open Science

- Only local/center based supercomputing before
  - e.g ARC, LeRC, ARL, UCB,
- Contemporary shared supercomputing was single discipline before
  - e.g NCAR, MFECC, …
- NERSC taught the rest of the HPC community how to do it
  - e.g NASA, NSF, DOD, other DOE, ….
  - International Supercomputing

NERSC's Greatest Legacies

NERSC Made Massively Parallel Processing into Production Supercomputing

- Previous MPP systems were prototyping/experimental/exploratory
- The T3E became the major production computing resource at NERSC in 1996
- NERSC, and the NERSC Users, were the first major site to commit its production and future only the distributed, highly parallel computing
  - Majority of use, effort and funding
- Within 1 year of the full T3E, all disciplines were using much more time and had larger allocations on the parallel system than on all the vector system combined.
- Now, there is only MPP computing within the Supercomputing community

NERSC's Great Legacies

NERSC's Focus on Science First

- Exceptional customer service is an integral part of NERSC
- NERSC held as the "standard candle" for high quality services
- Increases science team's productivity
- A lot of hard work and innovation goes into excellent customer service and production
- Many methods adopted by others including Leadership Facilities

NERSC's Great Legacies

NERSC and Partners Helped Create New Software Capabilities

- MPI enhancements and use
- Data Management
- Visualization
- Adaptive and load balancing techniques
- I/O and storage
- Programming Models
- Cyber Protection
- Applications
  - Astrophysics, Climate, Materials, ….

NERSC's Great Legacies

NERSC merged Big Compute and Big Data Decades Before We Knew to Call it Big Data

- T3E and PDSF arrived at NERSC on the same day in 1996 and started production the same year
- HPSS collaboration
- Global Unified Parallel File System
  - Envisioned at NERSC – others said it could not be done
    - Need shown when T3E was in service
  - Emulated elsewhere – i.e. Spider, NCSA, DOD,….
  - Many others tried but failed before NERSC showed it was possible
- More data comes to NERSC than leaves
- A high percentage of stored data is actually used

NERSC's Great Legacies

NERSC Focus on Real Performance

- Real time to solution for real problems = real sustained performance
- System choices (back to MFECC) was always driven by science needs
- SSP method is adapted in many places
  - DOD HPC Mod, NCSA, Australian Met, ….

NERSC's Great Legacies

NERSC is the best at showing its real impact on scientific progress

- It is hard to show science impact
- NERSC has been showing its impact for a long time
  - Volume
  - Uniqueness
  - Unexpected
- Many other organizations leverage it and benefit from it

NERSC's Great Legacies

NERSC and LBNL Helped Re-invigorated Federal Supercomputing Funding – Particularly for Open Science

- Circa 2004 - Earth Simulator just announced
- Informal meeting at LBNL with a new Office of Science appointee Ray Orbach (not yet confirmed)
- We discussed the issues, the importance, the motivations and the benefits of increased HPC investment
- Since that time, the Federal budget for High End Computing has doubled since that discussion
- The open computational science share of funding more than doubled

NERSC's Great Legacies

NERSC Systems Are Entirely Open and on the Internet – NERSC created New Cyber-protection Methods to Enable This

- Implemented and enhanced *Bro* (Vern Paxson)
- NERSC showed high performance, unencumbered networking with the philosophy of "Monitor and React rather than Risk"
  - Major systems are not behind slow, disruptive firewalls
- Science DMZ implemented and proven by ESnet and NERSC
- NERSC Influenced policy to remain open and efficient while others restricted
- Monitor and React is only viable if it works. NERSC remained in service when many other sites were down for compromises.
  - Technology and expertise

NERSC's Great Legacies

Serving the Grand Challenges and the overall Computational Community

- Persistent support for all types of computing
- DOE INCITE program started at NERSC to serve at scale
  - Colliding Black Holes and others
  - Showed great steps forward can be done by giving large resources to well qualified teams
- Resource management for really big, large work and lots of smaller jobs was solved
- Missing middle has no better place to work than NERSC

NERSC's Great Legacies

NERSC Focus on Science First

- Exceptional Customer Service is an integral part of NERSC
- NERSC held as the "standard candle" for high quality services
  - A required chapter in any "Best Practices" document

NERSC's Great Legacies

NERSC Re-invigorated HPC at Berkeley

- When NERSC arrived, no one at Berkeley was using NERSC
- Within 5 years, Berkeley had the largest share (by institution) of NERSC allocations
- Built strong programs, with Berkeley, to attract new computationally focused staff
- Synergies enabled many new discoveries and insights and at least one Nobel Prize

NERSC and Vendors

- Clearly a give and take relationship
- LLNL/NERSC BVSS is the model acquisition method for many places
  - Holistic evaluation
- NERSC helped many vendors be better that then benefits all other customers
  - Cray (in all its variations) – vector, MPP, data
  - IBM – SP at scale, Power IH architecture, GPFS, HPSS, …
  - Networking Vendors – Force 10, Juniper, ….
  - Storage Vendors -
- NERSC made some hard decisions, but in retrospect, they were the right decisions
  - Cray Computer, IBM, SGI, …
- SSP method assures real, meaningful performance for applications
- NERSC provides honest, fair and meaningful opinion

NERSC Diaspora

- A person recently told me that a measure of impact of an Organization is the diaspora of its alumni
- NERSC benefited from the diaspora of others and developed "best practices" by importing people who knew the best practices
  - Being one of the best helps get the people with the best practices
- NERSC staff have made many other organizations richer
  - NERSC alumni have gone on to work at other centers, industry and commercial, academia and government
  - Some even return to NERSC and LBNL after their "sabbaticals"

My Summary of NERSC

- NERSC was there first
- NERSC is there the longest
- NERSC is open science at unprecedented scales and impact
- NERSC is a perennial leader of the Supercomputing Community
- The Nation, DOE and the Supercomputing Community should be proud of NERSC and thankful for NERSC

I have some other interests now. So, I am happy to paraphrase Ray today

NERSC is today, and will be in the future, one of the two best run HPC computing facilities in the world.

I will let you can guess which is the other facility.

Happy 40th Anniversary NERSC!!!

COMMENTS?
WHAT DID I MISS?
WHAT DID I GET WRONG?
WHAT ARE YOUR LEGACIES?