



INCITE Proposal Writing Tips



James Osborn,
ALCF Computational Scientist
Argonne National Laboratory
and

Bronson Messer,
OLCF Scientific Computing Group
Oak Ridge National Laboratory

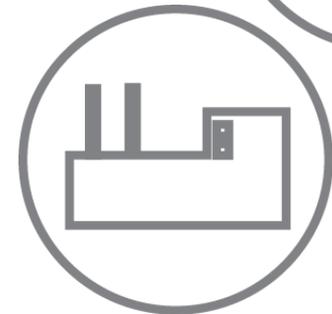
INCITE proposal writing tips

- INCITE at the DOE Leadership Computing Facilities
- What is INCITE good for? ...Not so good for?
- Walk through the Proposal Form
- User's View of Systems
- What Happens Next?
- PI Obligations
- Additional Q&A

INCITE

Innovative and Novel Computational Impact on Theory and Experiment

- Solicits large-scale, computationally intensive research projects
 - to enable high-impact scientific advances
- Open to all scientific & engineering researchers and organizations worldwide
 - including universities, laboratories and companies
- Provides large computer time and data storage allocations
 - average award in 2010 was 23 million hours



INCITE process

Call for Proposals (new and renewals) April 14 to June 30

Computational Readiness Review

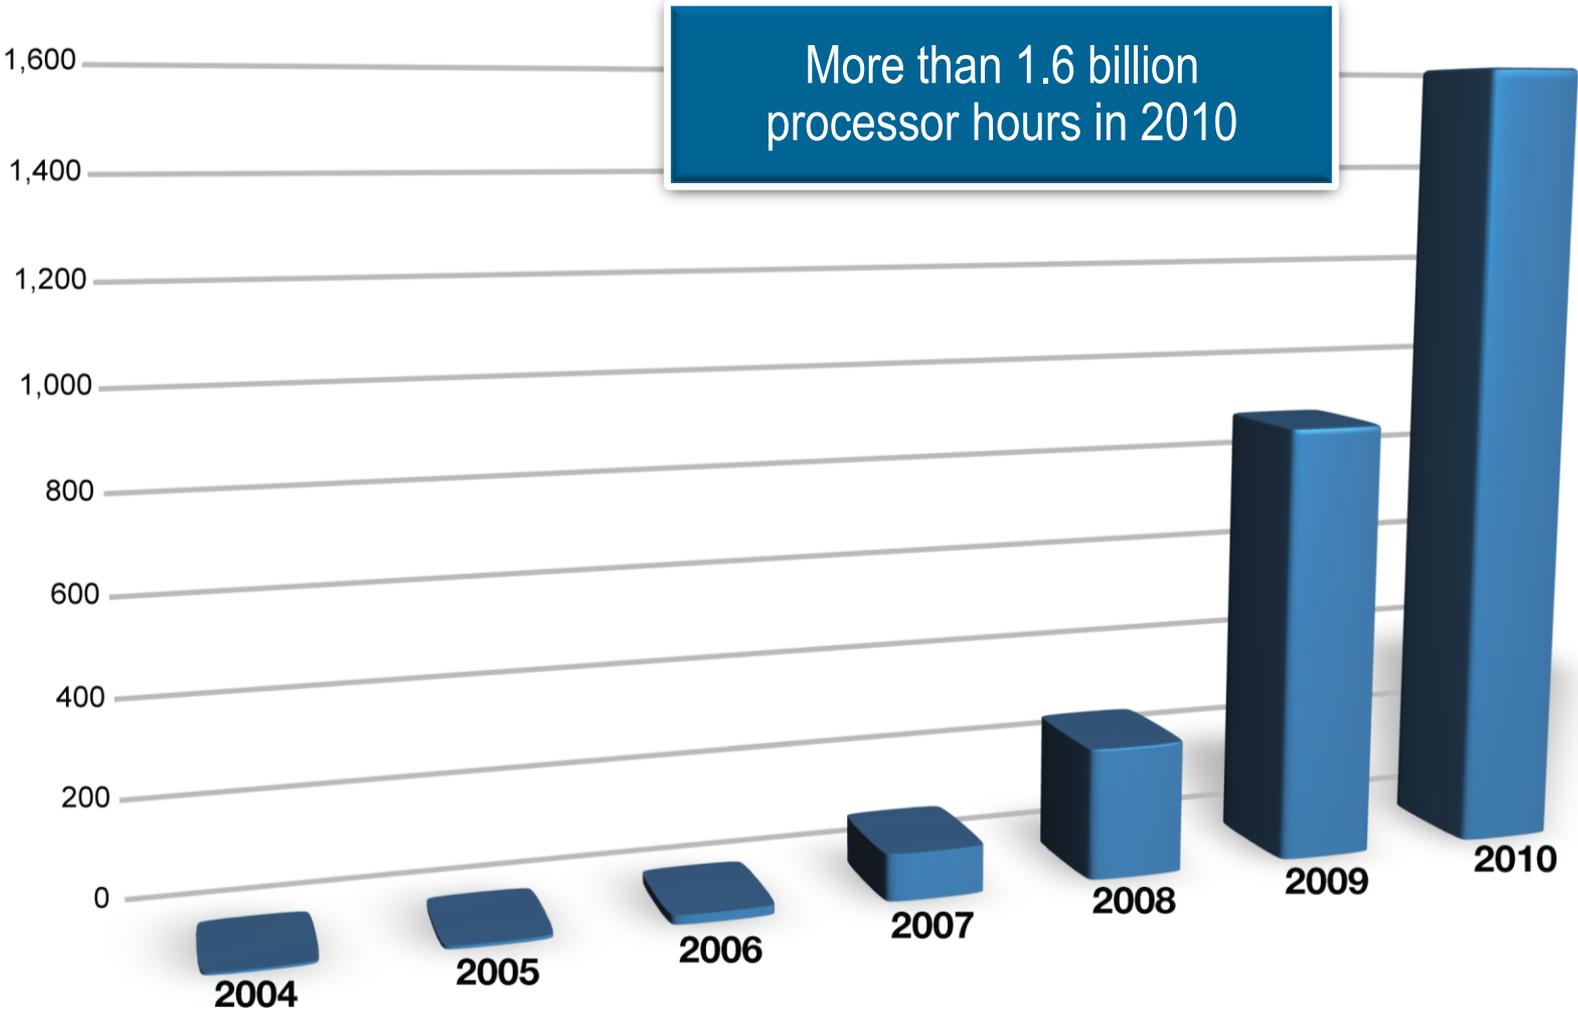
Scientific Peer Review

Announcements in early November 2010

Allocations from January to December 2011

INCITE in its 7th year

Explosive growth in program allocations



INCITE is open to researchers worldwide from academic, government, and industry organizations

Advancing the state of the art across a range of disciplines

Accelerator physics

Astrophysics

Bioenergy

Chemical sciences

Computer Science

Engineering

Climate research

Environmental science

Life sciences

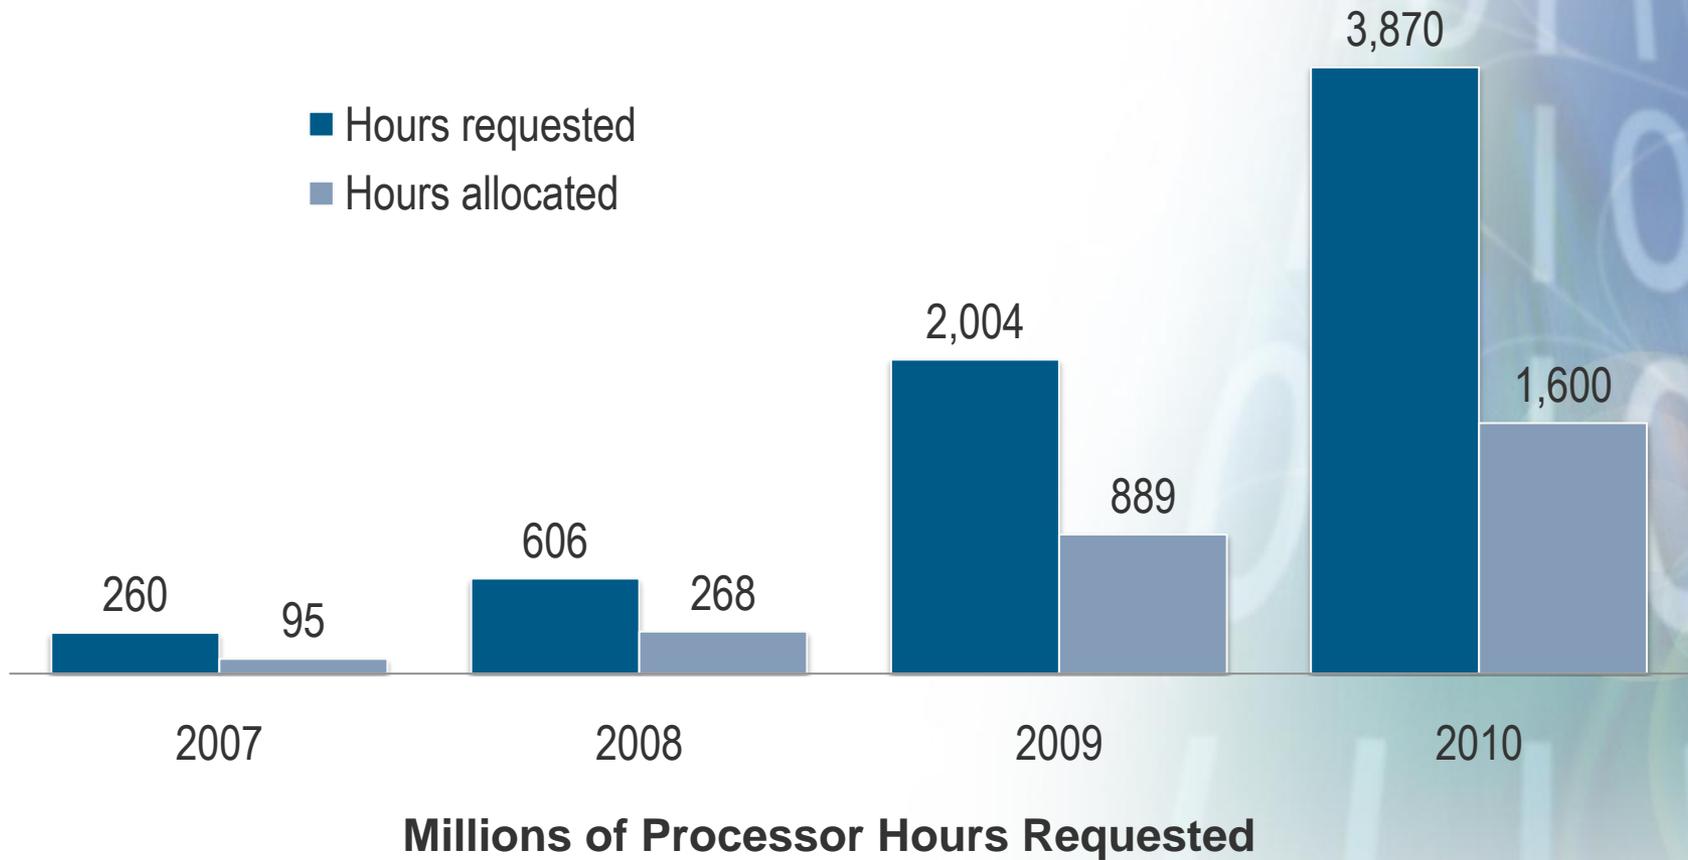
Materials science

Fusion energy

Nuclear physics

- No designated number of hours for a particular science area
- INCITE looks at all recommendations, focusing on potential for scientific or technological impact, then readiness

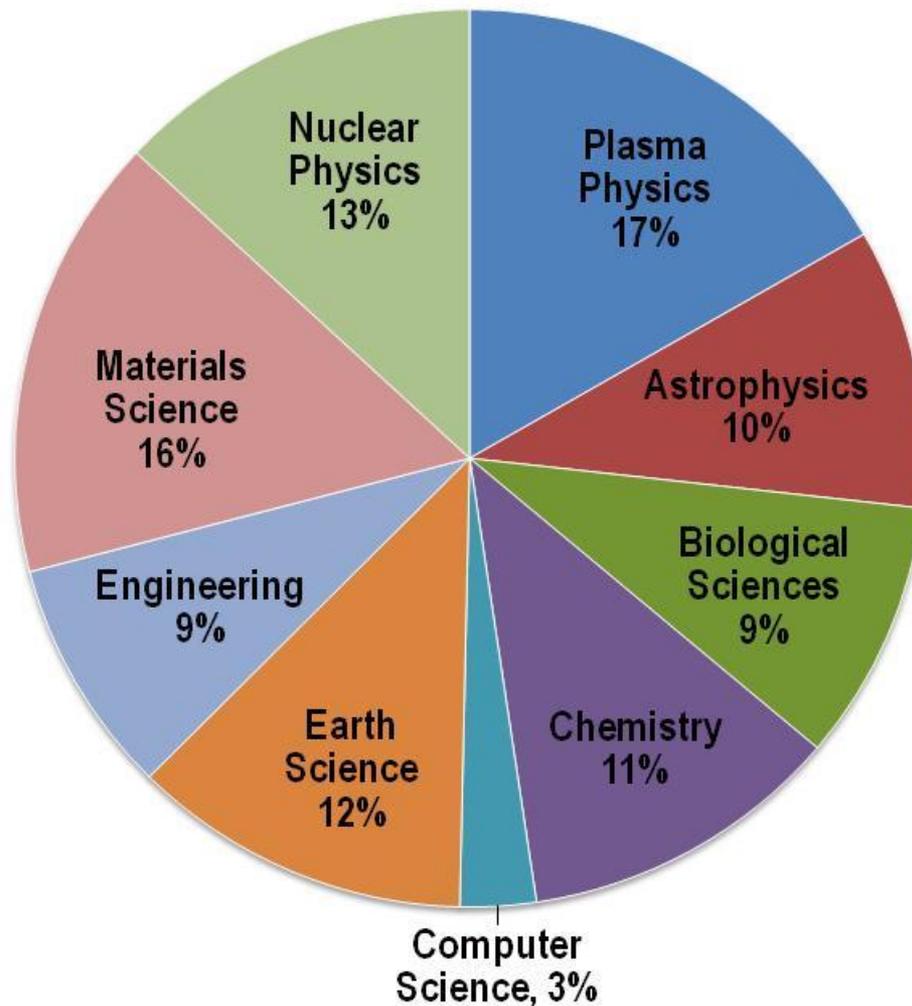
The INCITE program is typically 2.5× oversubscribed



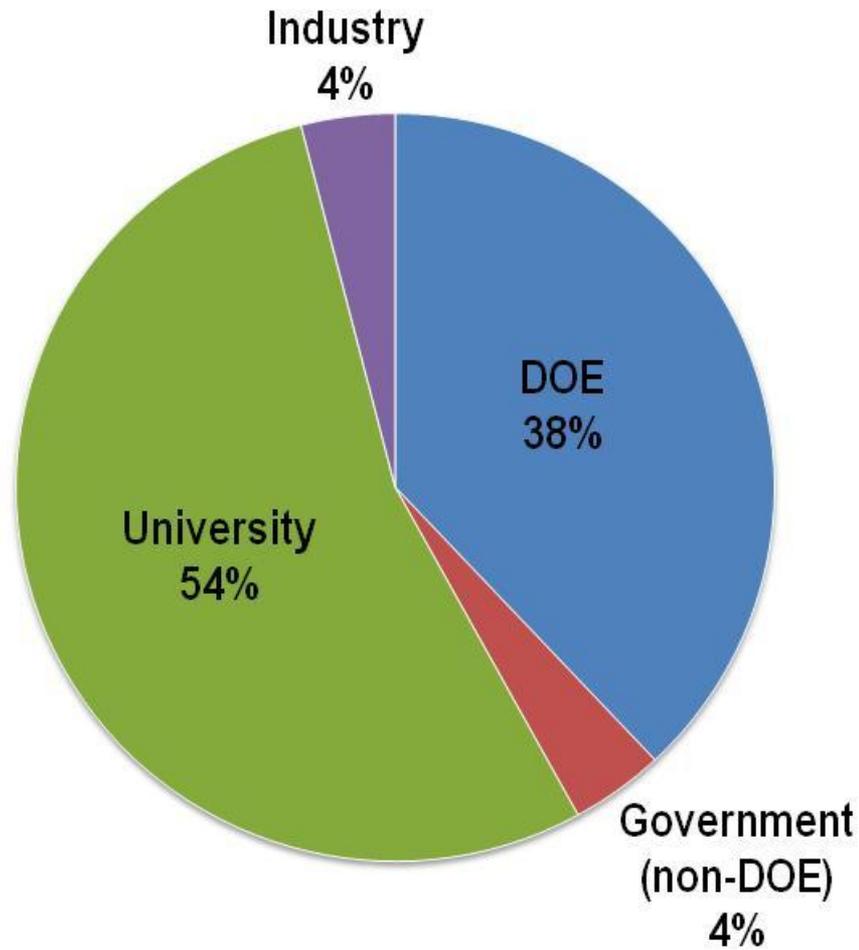
INCITE 2010 awards:

35 new projects and 34 renewal projects

41% of new submittals and 83% of renewals



2010 project demographics by PI affiliation



Review process: Computational readiness

- Reviewers: Center staff who are expert in these systems
- Both centers review each proposal

Criteria for new proposals	Criteria for renewals
<ul style="list-style-type: none">• Appropriateness for requested resources• Appropriateness of computational approach• Technical readiness	<ul style="list-style-type: none">• Met technical and computational milestones• On track to meet future milestones

Review process: Science Panel review

- Reviewers: Domain experts drawn from institutions worldwide

Criteria for new proposals	Criteria for renewals
<ul style="list-style-type: none">• Scientific and/or technical merit• Appropriateness of proposal method• Team qualifications• Reasonableness of requested resources	<ul style="list-style-type: none">• Change in scope• Met technical/scientific milestones• On track to meet future milestones• Impact relative to other proposals under consideration

Outline of the proposal form

1. Principal Investigator and Co-Principal Investigators
2. Project Title
3. Research Category
4. Project Summary (50 words)
5. Computational Resources Requested
6. Funding Sources
7. Other High Performance Computing Support for this Project
8. Project Narrative, other Materials
 - 8.a Executive Summary (1 page)
 - 8.b Project Narrative, including computational readiness, job characterization (15 pages)
 - 8.c Personnel Justification (1 page)
9. Application Packages
10. Proprietary and Sensitive Information
11. Export Control
12. Monitor Information
13. Outreach

Let's get started

- The story behind each science campaign is different. Use your best judgment when writing your proposal.
- Help and advice is available from many sources, e.g., LCF staff, DOE Program Managers, colleagues with INCITE projects.

What is INCITE good for?

- High impact science and engineering with specific objectives
- Computationally intensive runs that you cannot do anywhere else
- Jobs that can use at least 20% of the system for hours to weeks
- Campaigns requiring millions of CPU hours
- Computations that are efficient on LCF systems

A Few Critical Questions to ask

- Is another resource more appropriate?
- Is both the scale of the runs *and* the time demands of the problem of LCF scale?
- Do you need specific LCF hardware?
- Is your production run size larger than the minimum size allowed on the system?
- Do you have the people ready to do this work?
- Do you have a post-processing strategy?
- Do you need a workflow?

Some of the above characteristics are negotiable, so make sure to discuss atypical requirements with the centers

Some projects may be appropriate for other centers

- Laws regulate what can be done on these systems
 - LCF systems have cyber security plans that bound the types of data that can be used and stored on them
- Some kinds of information we cannot have
 - Personally Identifiable Information (PII)
 - Information requiring an export control license
 - Classified Information or National Security Information
 - Unclassified Controlled Nuclear Information (UCNI)
 - Naval Nuclear Propulsion Information (NNPI)
 - Information about development of nuclear, biological or chemical weapons, or weapons of mass destruction
- Inquire if you are unsure or have questions

Proprietary data and export control

Back to LCF center requirements: it isn't only how MUCH data, but what KIND of data (input, output, software)

- Are you using proprietary input or software? Are you producing proprietary output?
 - Proprietary materials are things that you plan to reserve rights to
- Does your project have export control classification?
 - Export control can be related to the topic under study (ex. fuel rods)
 - Export control can be relevant to the software application you're using

This information helps the centers determine which user agreement is appropriate and the levels of data protection needed.

It is a small world...

- DOE program managers interact on many occasions, within the Office of Science, within DOE, and with other science agencies
- Advice is often sought and given
- So it could help to let the science program that funds your work know how significant your planned INCITE allocation will be to your work and their program

Have your elevator speech ready

Science & Engineering Story Tips

- Audience
 - Computational-science-savvy senior scientists / engineers, and faculty
 - Not everyone will be well versed in your approach
- Story Elements
 - What the problem is, and its significance
 - Key objectives, key simulations/computations, project milestones
 - Approach to solving the problem, its challenging aspects, preliminary results
 - Impact of a successful computational campaign - the big picture
 - Reasons why it is important to carry out this work now

INCITE has become quite competitive

Administrative Tips

- **Principal Investigator, other contacts**
 - Author/PI needs to ensure that all contact information is current
 - Ask around your organization for the Institutional contact (i.e. the person at your institution who can sign a user agreement)
- **Project Title**
 - Pick a title you will be proud of seeing in many, many places
 - Be succinct - no kidding!
- **Project Summary**
 - 2-sentences suitable for the public (e.g., *Science News*, *Scientific American*, web, PowerPoint)

Note: the PI is unlikely to also be the Institutional Contact

Computational Resource Request

- How many years will your project need (1-3)?
- Things that will slow you down the first year
 - Porting and code development
 - Learning to use a new center
 - User agreements for all the institutions and people involved
 - Paperwork for proprietary use
- Mind the units!
 - **Processor (Core) Hours** - for the system you will be running on
 - Units are core-hours, however you are charged for all cores on a node
 - On Blue Gene you are charged for all cores in your partition
 - Large partitions are in increments of a rack (1,024 nodes, 4,096 cores)
 - On jaguar, hours are core-hours, but you must request cores in units of nodes (i.e. 12 cores)
 - **Disk Storage** - in gigabytes for both Home and Scratch space
 - **Mass (Tape) Storage** - in gigabytes or terabytes (specify)

Multiple Facilities

- If your project needs a single primary resource
 - Identify the primary resource in (5a)
- If your project needs multiple primary resources
 - Identify each of the primary resources in (5a)
 - By adding resources for the same year
 - Justify your need for and ability to effectively use multiple resources in section (8)

Computational Approach

- Essential to show experience and credibility in section 8.b
 - Successful proposal teams demonstrate a very clear understanding of petascale computing and can optimally use these resources to accomplish the stated scientific/technical goals
- Programming Languages, Libraries and Tools used (8.b)
 - Check that what you need is available on the system
- Description of underlying formulation (8.b)
 - Don't assume reviewers know all the codes
 - Make it clear that the code you plan to employ is the correct tool for your research plan
 - If you plan to use a private version of a well-known code, delineate the differences

Computational Plan

- Application Packages (9)

- List of all the software application packages/suites to be used

Note: Long lists may reduce credibility

- What will be used to set up computations (and where)?
- What are the codes for the main simulation/modeling?
- What will be used to analyze results (and where)?

Application Credentials

- Port your code before submitting the proposal
 - Check to see if someone else has already ported it
 - Request a start up account if needed
- It is very hard to embarrass a >150,000-processor system
 - Prove application scalability in your proposal
 - Run example cases at full scale
 - Make examples as close to production runs as possible
 - If you cannot show proof of runs at full scale, then provide a very tight story about how you will succeed

A sample of codes with local expertise available at ALCF and OLCF

Application	Field	Argonne	Oak Ridge
FLASH	Astrophysics	✓	✓
MILC,CPS	LQCD	✓	✓
Nek5000	Nuclear energy	✓	
Rosetta	Protein Structure	✓	
DCA++	Materials Science		✓
ANGFMC	Nuclear Structure	✓	
REDSTICK	Nuclear Structure		✓
Qbox	Chemistry	✓	✓
LAMMPS	Molecular Dynamics	✓	✓
NWChem	Chemistry	✓	✓
GAMESS	Chemistry	✓	✓
MADNESS	Chemistry		✓
CHARMM	Molecular Dynamics	✓	✓
NAMD	Molecular Dynamics	✓	✓
AVBP	Combustion	✓	
GTC	Fusion	✓	✓
Allstar	Life Science	✓	
CPMD, CP2K	Molecular Dynamics	✓	
CCSM3	Climate	✓	✓
HOMME	Climate	✓	✓
WRF	Climate	✓	✓
Amber	Molecular Dynamics	✓	✓
enzo	Astrophysics	✓	✓
Falkon	Computer Science/HTC	✓	✓
s3d	Combustion		✓
DENOVO	Nuclear Energy		✓
LSMS	Materials Science		✓
GPAW	Materials Science	✓	

Programming Models

- Parallel Programming System
 - MPI (MPICH2) is the work horse on both platforms
 - Cray MPI based on MPICH2
 - ARMCI/Global Arrays is available
 - OpenMP on nodes
 - Some groups have “rolled their own” at lower level, e.g., QCD
- Special needs? Inquire
 - e.g., Python, custom kernel

Tools/Libraries/Packages supported at OLCF

- Analysis
 - Matlab
 - NCO
 - UDUNITS
 - Ferret
 - gnuplot
 - grace
 - IDL
 - ncl
 - ncview
 - PGPLOT
- Libraries
 - Communication
 - BLACS
 - Global Arrays
 - I/O
 - HDF5
 - netcdf
 - pnetcdf
 - szip
- Math
 - acml
 - ARPACK
 - ATLAS
 - Aztec
 - BLAS
 - fftpack
 - fftw
 - gsl
 - hypre
 - LAPACK
 - Cray libsci
 - METIS
 - MUMPS
 - ParMETIS
 - PETSc
 - pspline
 - ScaLAPACK
 - sprng
 - Sundials
 - SuperLU
 - Trilinos
- UMFPACK
- Program Dev
 - cmake
 - doxygen
 - git
 - mercurial
 - subversion
 - Debugging
 - Totalview
 - valgrind
 - Performance
 - Apprentice2
 - craypat
 - fmpi
 - gptl
 - mpe2
 - mpip
 - PAPI
 - TAU

Tools/Libraries/Packages supported at ALCF

- Analysis
 - Gnuplot
 - Paraview
 - Visit
- Libraries
 - Communication
 - BLACS
 - DCMF
 - Global Arrays
 - I/O
 - HDF5
 - netcdf
 - pnetcdf
 - szip
 - Math
 - ATLAS
 - BLAS/GotoBLAS
 - essl
 - Fftpack
 - Fftw
- p3dfft
- hypre
- LAPACK
- Mass/massv
- METIS
- MUMPS
- ParMETIS
- PETSc
- SCALAPACK
- Spooles
- SuperLU
- Trilinos
- Valgrind
- coreprocessor
- Performance
 - Darshan
 - HPC Toolkit (IBM)
 - HPCToolkit (Rice)
 - Fpmpi
 - mpe2
 - mpip
 - PAPI
 - TAU
 - UPC (universal performance counter)
- Program Dev
 - cmake
 - Doxygen
 - git
 - mercurial
 - subversion
 - Debugging
 - gdb
 - Totalview
- Other
 - ZeptoOS

If a you have a requirement that is not listed, ask!

Computational Campaign

- Describe the kind of runs you plan with your allocation
 - L exploratory runs using M nodes for N hours
 - X big runs using Y nodes for Z hours
 - P analysis runs using Q nodes for R hours
- Big runs often have big output
 - Show you can deal with it and understand the bottlenecks
 - Understand the size of results, where you will analyze them, and how you will get the data there

Parallel Performance

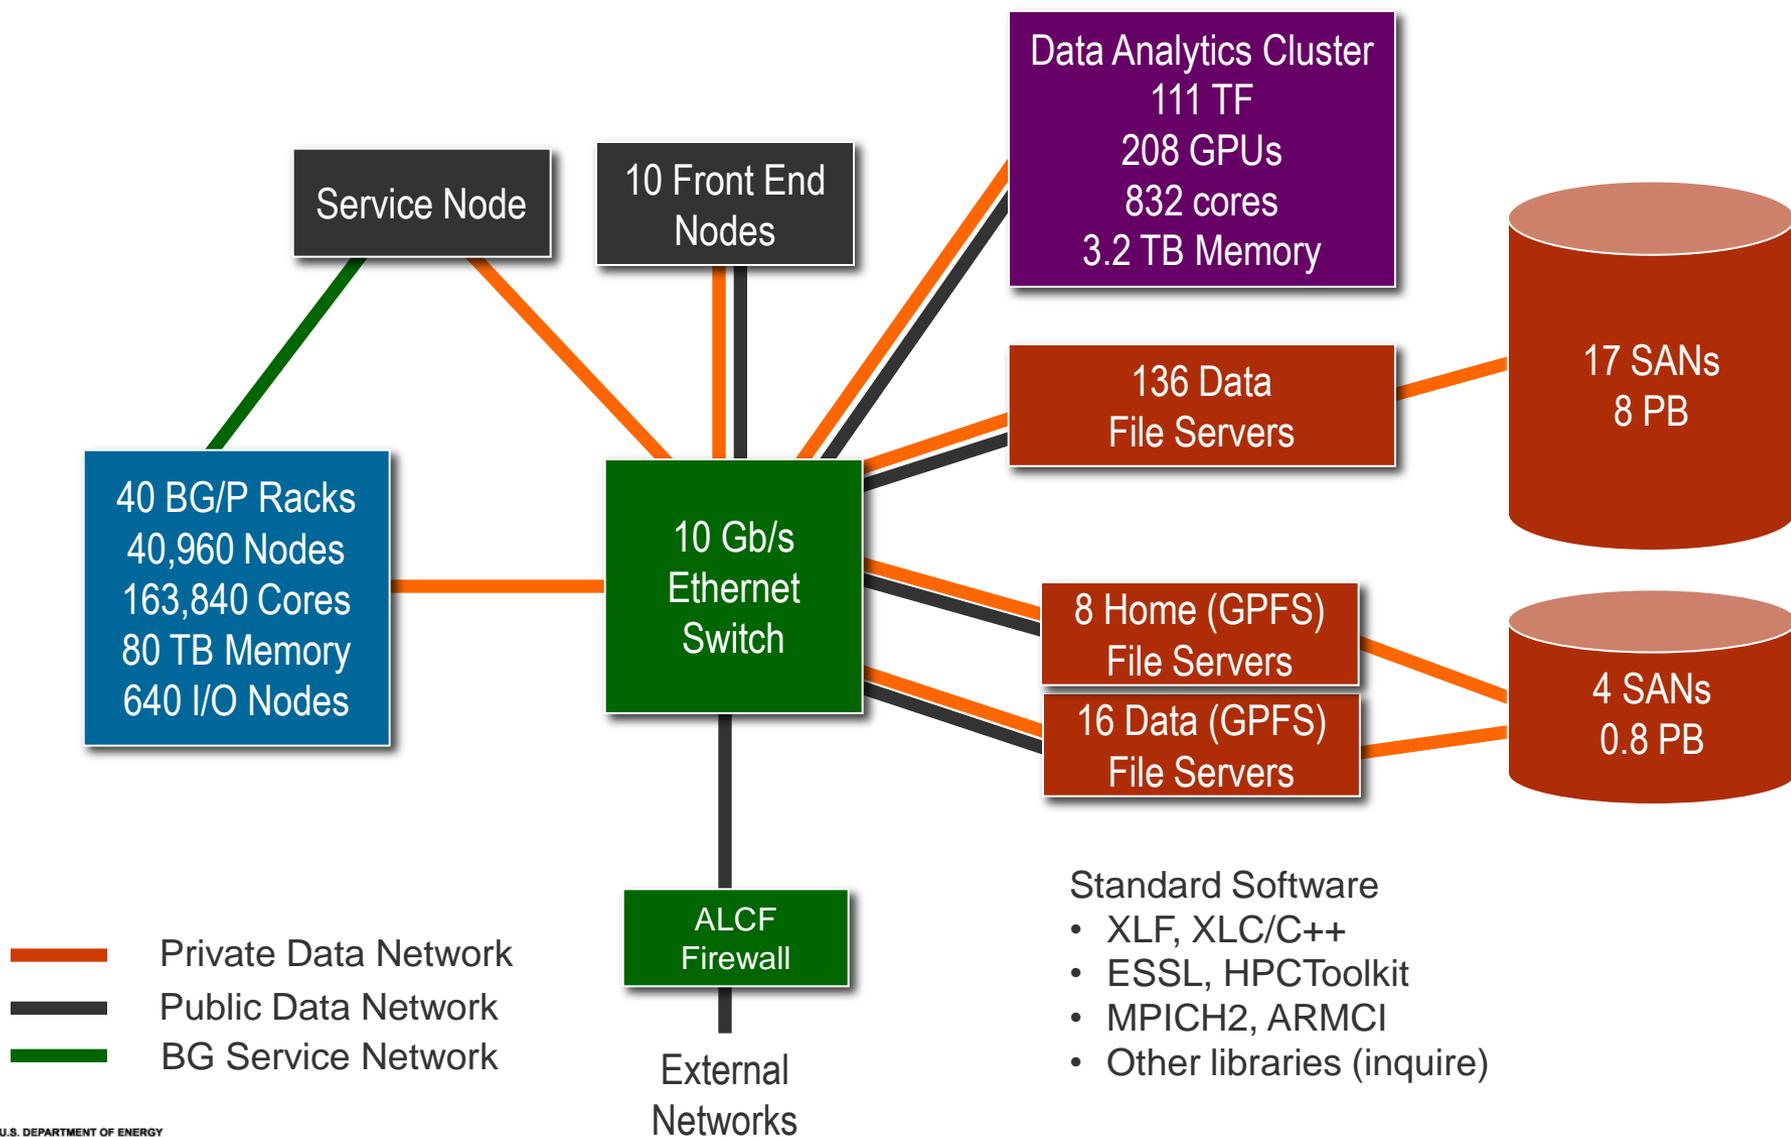
direct evidence is required

- Pick the approach relevant to your work and show results
 - **Strong Scaling Data**
 - Increase resources (nodes) while doing the same computation
 - **Weak Scaling Data**
 - Increase problem size as resources are increased
- Performance data should support the required scale
 - Use similar problems to what you will be running
 - Show that you can get to the range of processors required
 - Best to run on the same machine, but similar size runs on other machines can be useful
 - Describe how you will address any scaling deficiencies
- Be aware of scaling data from other groups and literature

I/O Requirements

- **Restart I/O** - Application initiated program restart data
 - I/O technique used, e.g., MPI I/O, HDF5, raw
 - Number of processors doing I/O, number of files
 - Sizes of files and overall dump
 - Periodicity of the checkpoint process
- **Analysis I/O** - Application written files for later analysis
 - I/O technique used, e.g., pNetCDF, pHDF5
 - Number of processors doing I/O, number of files
 - Sizes of files and overall dump
- **Archival I/O** - Data archived for later use/reference
 - Number and sizes of files
 - Retention length
 - If archived remotely, the transport tool used, e.g., GridFTP

ALCF Blue Gene/P “Intrepid” System



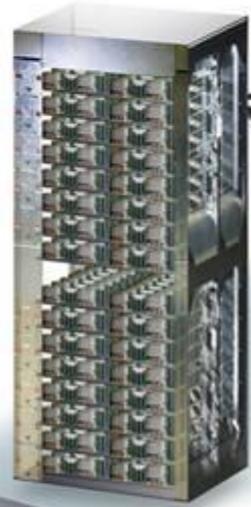
Blue Gene/P Hardware Summary

System



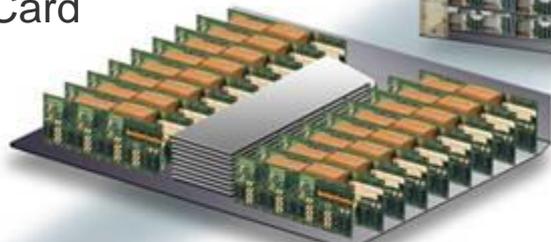
40 Racks
40960 nodes
163,840 cores
0.556 PF
80 TB DDR2

Rack



32 node-cards
1024 nodes
4096 cores
13.9 TF
2 TB DDR2

Node Card



32 nodes
128 cores
435 GF
64GB DDR2

Compute Card



1 node
4 cores
13.6 GF
2GB DDR2

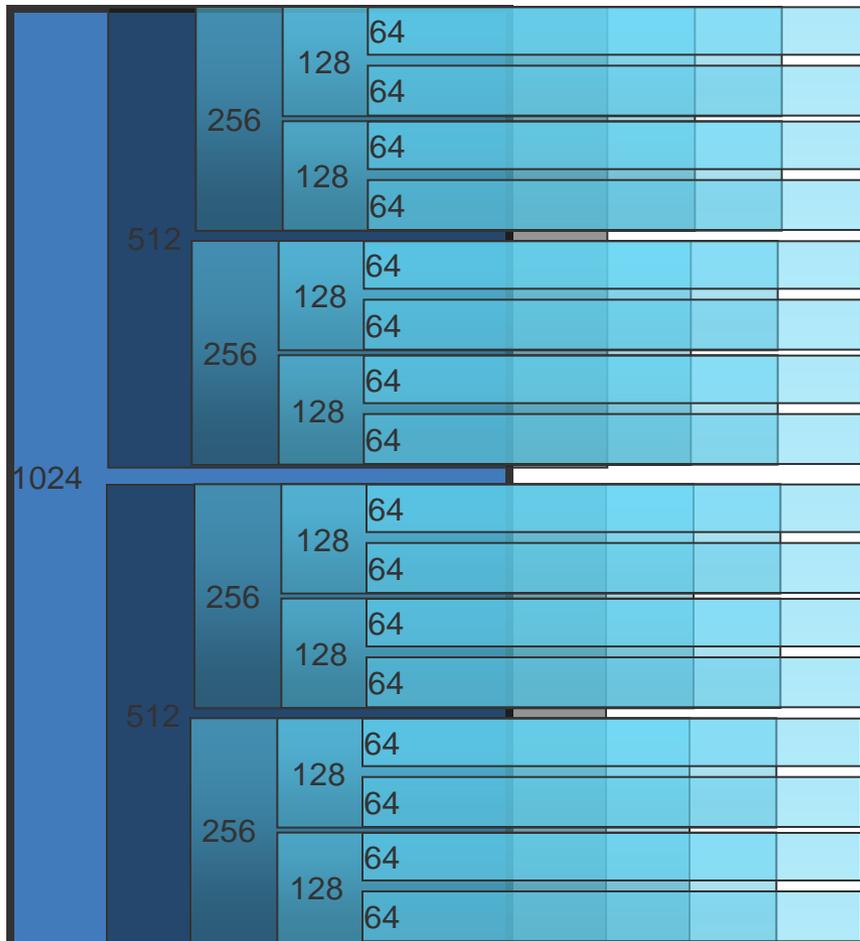
Chip



4 cores
13.6 GF

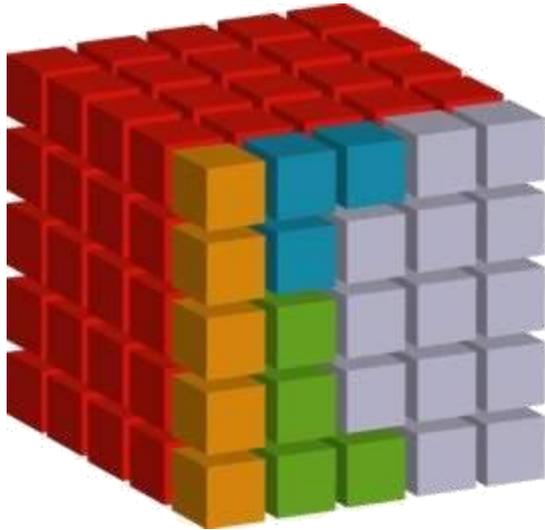
BG Partitions (“blocks”)

Single Rack



- 1 I/O node for each 64 compute nodes, hardwired to specific set of 64
 - *Minimum partition size of 64 nodes*
- Partition sizes: 64, 128, 256, 512, 1024, ...
 - *Any partition < 512 nodes will get a mesh network layout and not a torus.*
 - *Any partition <512 nodes will get a non-optimal I/O tree network.*
 - *Do not do performance testing on <512 nodes*
- Smaller partitions are enclosed inside of larger ones
 - *Not all partitions are available at all times*
 - *Once a job is running on one of the smaller partitions, no jobs can run on the enclosing larger partitions*
- Configuration changes frequently
 - *partlist shows partition state*
- Processes are spread out in a pre-defined mapping, alternate and sophisticated mappings are possible

Jaguar



Compute PE

Login PE

Network PE

System PE

I/O PE



- **Jaguar consists of an XT4 and an XT5 partition**
- **Ability to submit to either partition from a single set of login nodes will be available**

XT4 partition



Quad-core Processors	7,832
Memory / Core	2 GB
System Memory	62 TB
Disk Bandwidth	44 GB/s
Disk Space	900 TB
Node Size	4-core, 34 GF

XT5 partition

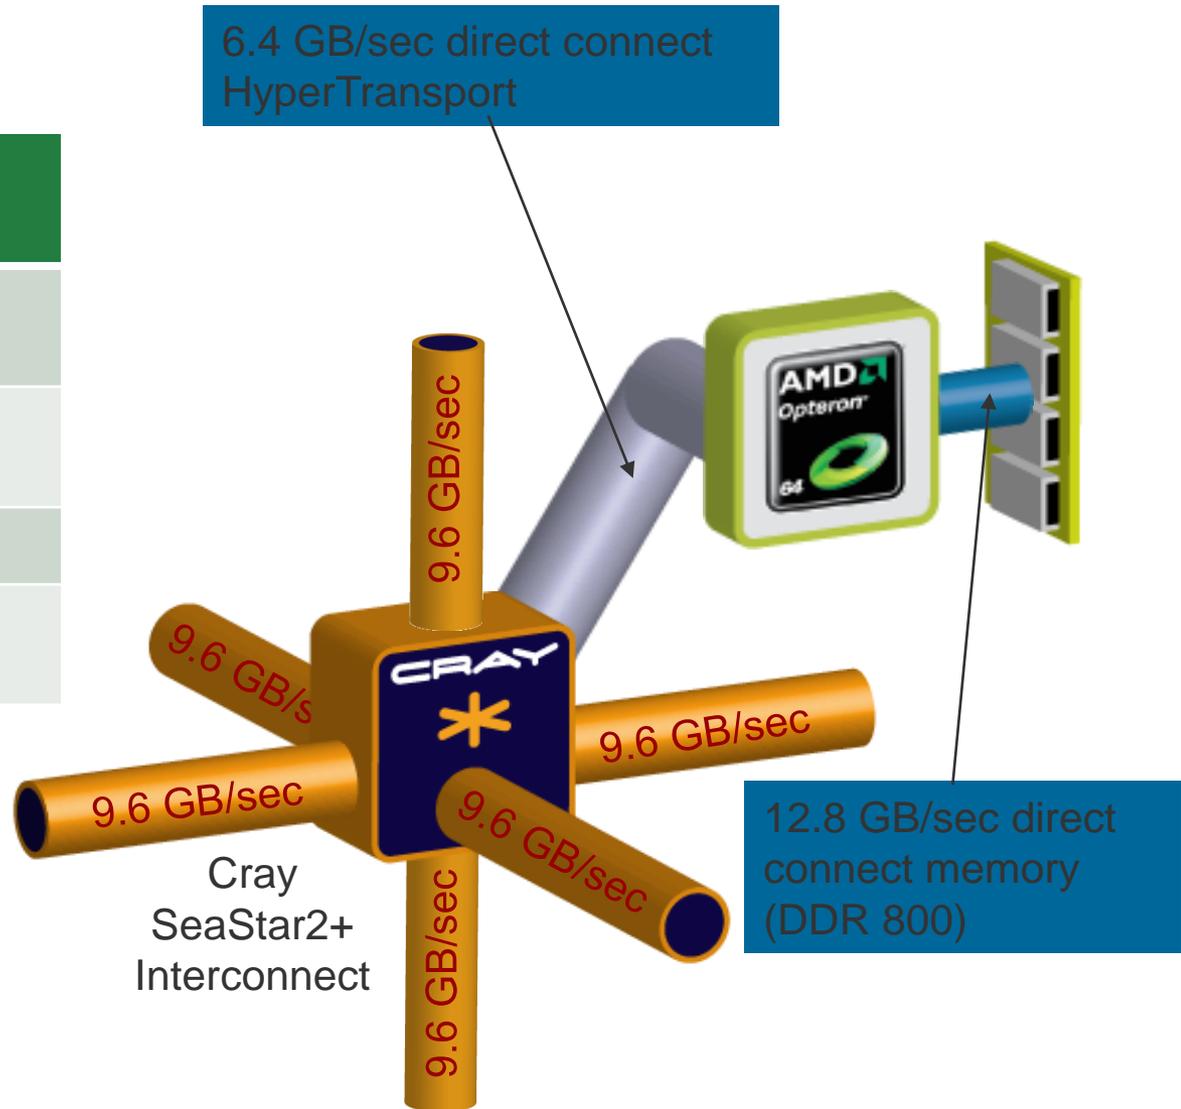


Peak performance	2.332 PF
System memory	300 TB
Disk space	10 PB
Disk bandwidth	240+ GB/s
Compute Nodes	18,688
AMD "Istanbul" Sockets	37,376
Size	5,000 feet ²
Cabinets	200 (8 rows of 25 cabinets)

XT4 node

Cray XT4 Node Characteristics

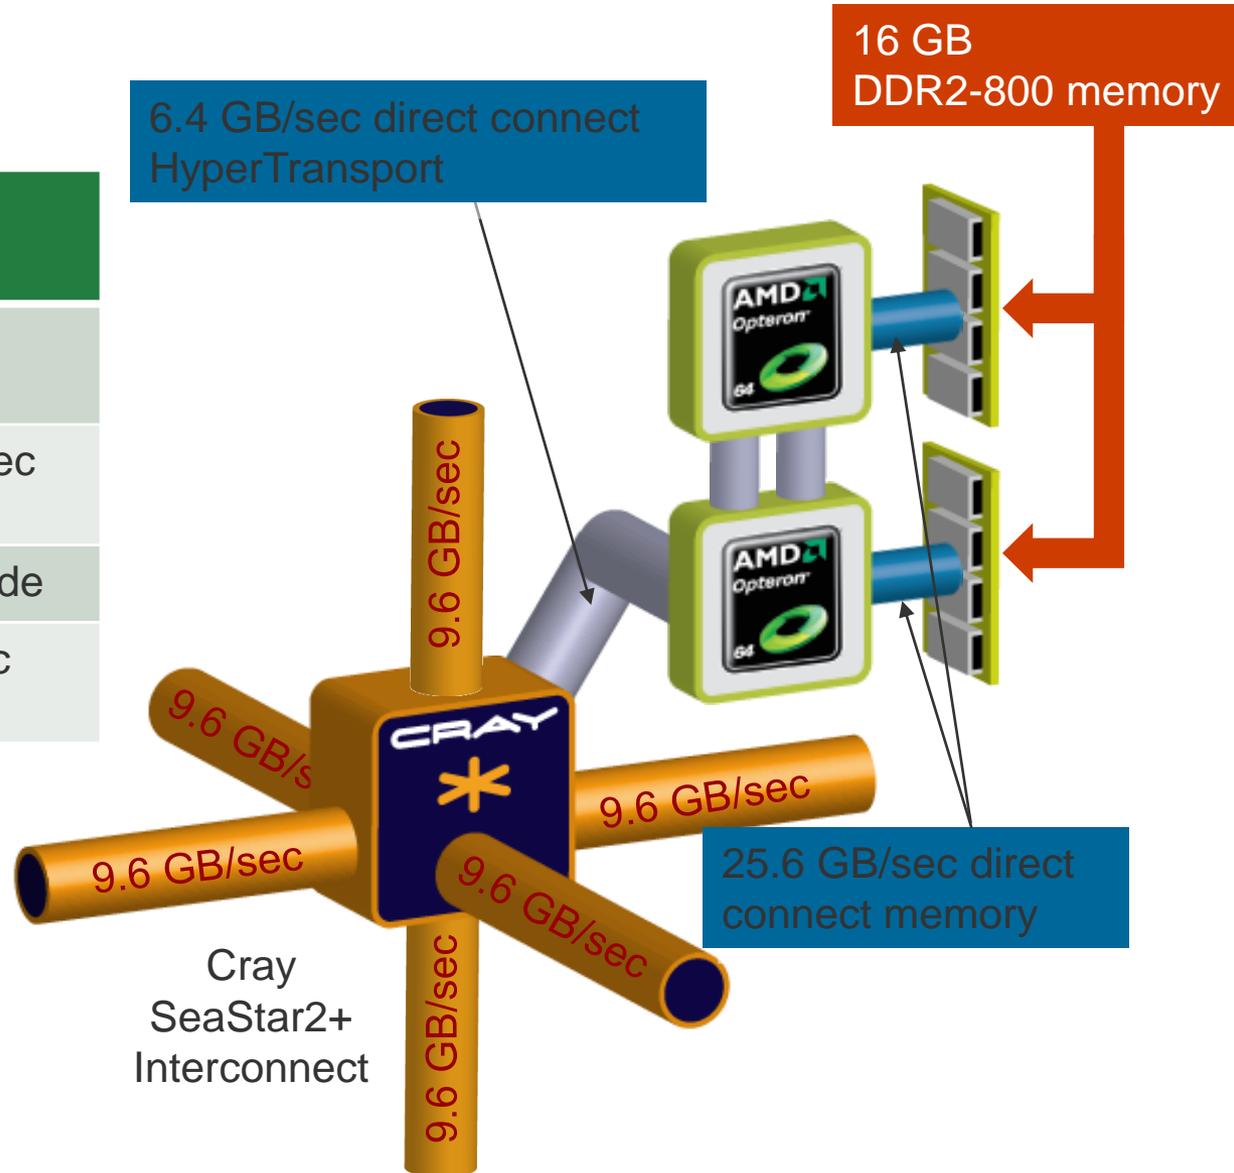
Number of Cores	4
Peak Performance	> 35 Gflops/s
Memory Size	8GB per node
Memory Bandwidth	12.8 GB/sec



XT5 node

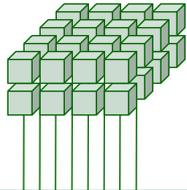
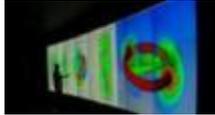
Cray XT5 Node Characteristics

Number of Cores	12
Peak Performance	125 Gflops/sec
Memory Size	16 GB per node
Memory Bandwidth	25.6 GB/sec

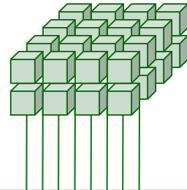


Center-wide filesystem: Spider

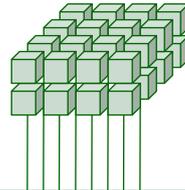
Everest
Powerwall



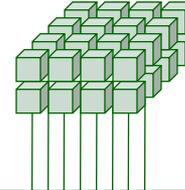
Remote
Visualization
Cluster



End-to-End
Cluster



Application
Development
Cluster

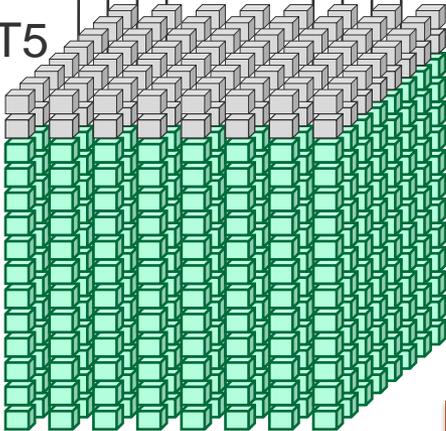


Data Archive
25 PB
HPSS

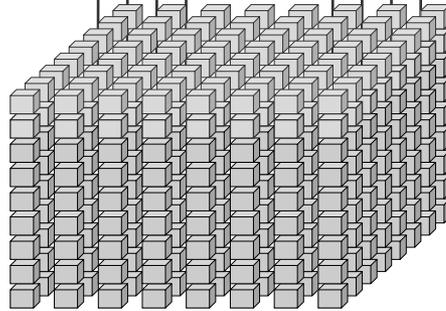


Scalable I/O Network (SION)
4x DDR Infiniband Backplane Network

XT5

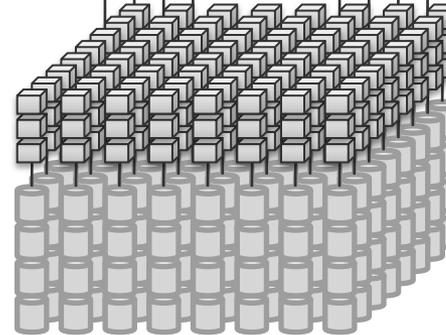


XT4



Login

Spider



ORNL's External login nodes and shared storage provide a single entry for users into a cluster of supercomputers

Ready to submit?

- You may save your proposal at any time without having the entire form complete.
- Your Co-PI's may also log in and edit your proposal.
- Required fields must be completed for the form to be successfully submitted; an incomplete form may be saved for later revisions.
- After submitting your proposal, you will not be able to edit it.

SUBMIT

Review processes revisited

Computational Readiness

- Appropriateness for the requested resources
- Appropriateness of the computational approach
- Technical readiness
- Progress (for renewals)

Reviewers can send questions to the PI's to clarify points in the proposal: Respond promptly!

Scientific Peer Review

- Scientific and/or technical merit
- Appropriateness of the proposal method
- Team qualifications
- Reasonableness of requested resources for the proposal

Science panel members have copies of the readiness reports; they may have additional questions for the PI's.

You've received an award, what's next?

- Notice comes from INCITE Manager, Julia White, in November 2010
- Welcome and startup information comes from centers
 - **Agreements to sign: start this process as soon as possible!**
 - **How to get accounts**
- User Support is geared to help you succeed
- Centers provide expert-to-expert assistance to help you get the most from your allocation

User agreements must be signed

- Nonproprietary Research
 - Nonproprietary user agreement – universities, national laboratories
 - Proprietary user agreement – for proprietary research
- Proprietary Research is permitted
 - Results retained by researcher or their organization
 - Full cost recovery. Data protection considerations
 - Carefully read DOE guidelines, and begin discussions early
- Process
 - Each institution must sign a User Agreement - by one who can commit the institution, e.g., attorneys, contract office, etc
 - Every user must have their own account, and must individually complete a facility user agreement

PI Obligations

- Provide quarterly status updates (on supplied template)
 - Milestone reports
 - Publications, awards, journal covers, presentations, etc. related to the work
- Provide highlights on significant science/engineering accomplishments as they occur
- Submit annual renewal request
- Complete annual surveys
- Encourage your team to be good citizens on the computers
- Use the resources for the proposed work
- Let us know about problems and issues

It is a small world...

For the continued success of the INCITE Program and availability of its resources that enable cutting edge research, let the program that funds your work know how significant your planned INCITE allocation will be to your work and their program

Have your elevator speech ready

Contact information



- www.alcf.anl.gov
- support@alcf.anl.org
- 1-866-508-9181



- www.olcf.ornl.gov
- help@nccs.gov
- 865-241-6536



- INCITE@DOEleadershipcomputing.org
- 865-235-8216

Additional Questions?