


Update on NASA High End Computing (HEC) & Survey of Federal HEC Capabilities

March 6, 2017

Tsengdar Lee
HEC Portfolio Manager
Science Mission Directorate




Modular Supercomputing Facility (MSF)

- Status
 - MSF is operational
 - Site is monitored by security 24x7 over cctv
 - Site is secured with 9' fence with card controlled gate access
 - Module status is monitored by control room 24x7 and inspected each shift
 - Electra is operational as of 1 January 2017
 - Congress has just approved the 2nd module be installed in 2017.
 - Electra
 - 16 SGI ICE-X Racks with Intel Xeon processors E5-2680v4 (Broadwell processors) (1.24 PF**;
147 TB+; 4,654 SBU⁺⁺/hr)
 - 1,152 nodes (dual-socket blades)
 - 2,304 processors (32,256 cores)
 - Dual-plane fully populated 7D hypercube InfiniBand FDR[^] Network with Metro-X IB extenders

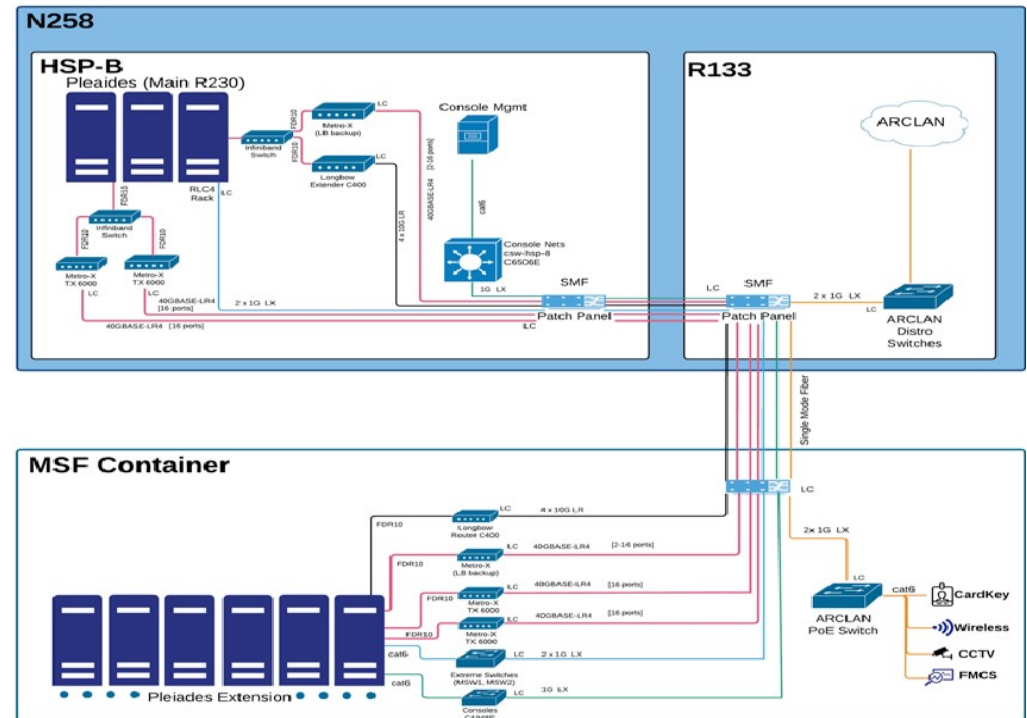


* TF = Teraflops = one million million (10^{12}) floating-point operations per second; ** PF = 10^3 TF;
+ TB = 10^{12} Bytes; ++ SBU = Standard Billing Unit for tracking computing usage
(<https://hec.nasa.gov/user/policies/sbus.html>); ^ FDR = 56Gb/s serial high bandwidth links



MSF Initial Test Results

- Benchmark Tests
 - LINPACK+ – 1.096 petaflops (PF) – 96th on November 2016 TOP500 List
 - HPCG++ – 25.188 teraflops (TF) – 46th on November 2016 HPCG List
- Application Tests
 - Extensive testing was conducted in November and December identifying performance issues focused on connectivity with the remote I/O infrastructure.
 - After system modifications, there is minimal to no impact on job run times.
- Operational Facility Early Results
 - Numerous facility tests were conducted in November and December with modifications made to improve the efficiency of the module.
 - PUE numbers are consistently below 1.03, greatly exceeding our 1.06 target.



+ LINPACK = linear algebra benchmark suite to measure computational speed; ++ HPCG = high performance conjugate gradient benchmark to measure computational speed; * PUE = power usage effectiveness




HECC Growth - Last Calendar Year

| Systems | Racks | | Cores | | TF* | | SBU ⁺ /hr | |
|--------------|------------|------------|----------------|----------------|--------------|--------------|----------------------|---------------|
| | Jan-16 | Feb-17 | Jan-16 | Feb-17 | Jan-16 | Feb-17 | Jan-16 | Feb-17 |
| Pleiades | 163 | 160 | 211,872 | 245,400 | 5,340 | 7,232 | 26,020 | 32,230 |
| Electra | | 16 | | 32,256 | | 1,239 | | 4,654 |
| Merope | 36 | 56 | 13,824 | 21,504 | 162 | 252 | 1,152 | 1,792 |
| Total | 199 | 232 | 225,696 | 299,160 | 5,502 | 8,723 | 27,172 | 38,676 |
| Growth | | 16.58% | | 32.55% | | 58.54% | | 42.34% |





* TF = Teraflops = one million million (10^{12}) floating-point operations per second
+ SBU = Standard Billing Unit for tracking computing usage (<https://hec.nasa.gov/user/policies/sbus.html>)

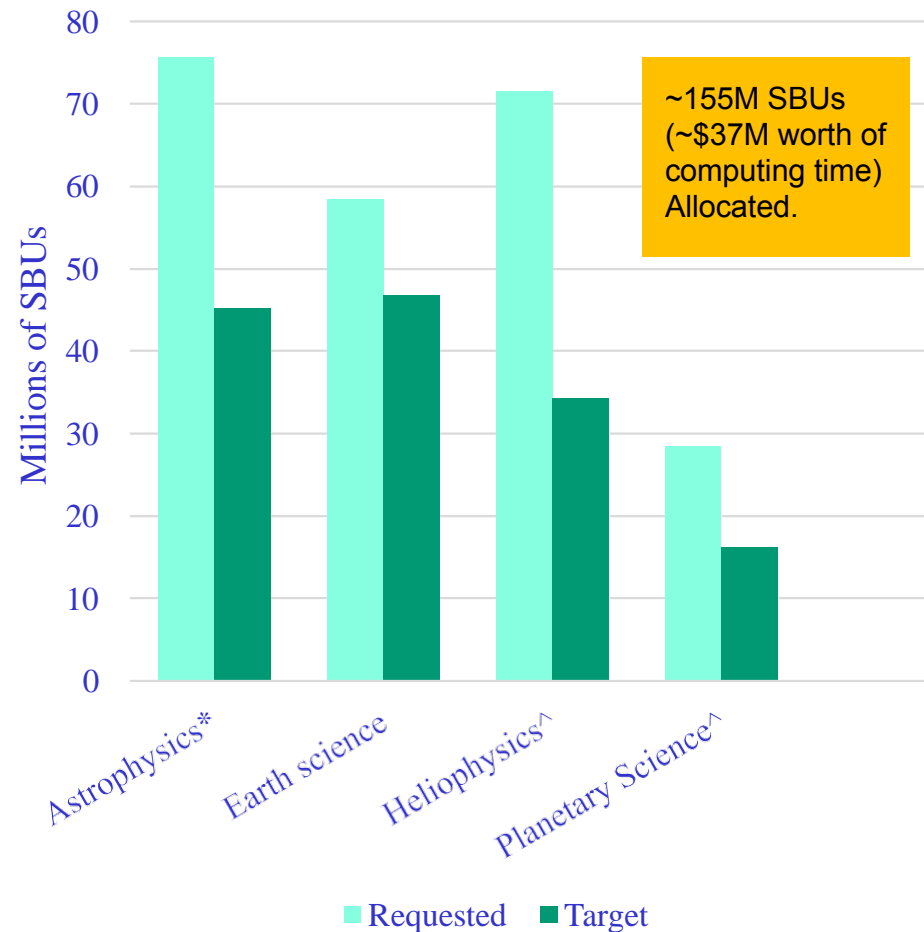


2017 Capacity Oversubscription

NAS Oversubscription




NAS Oversubscription



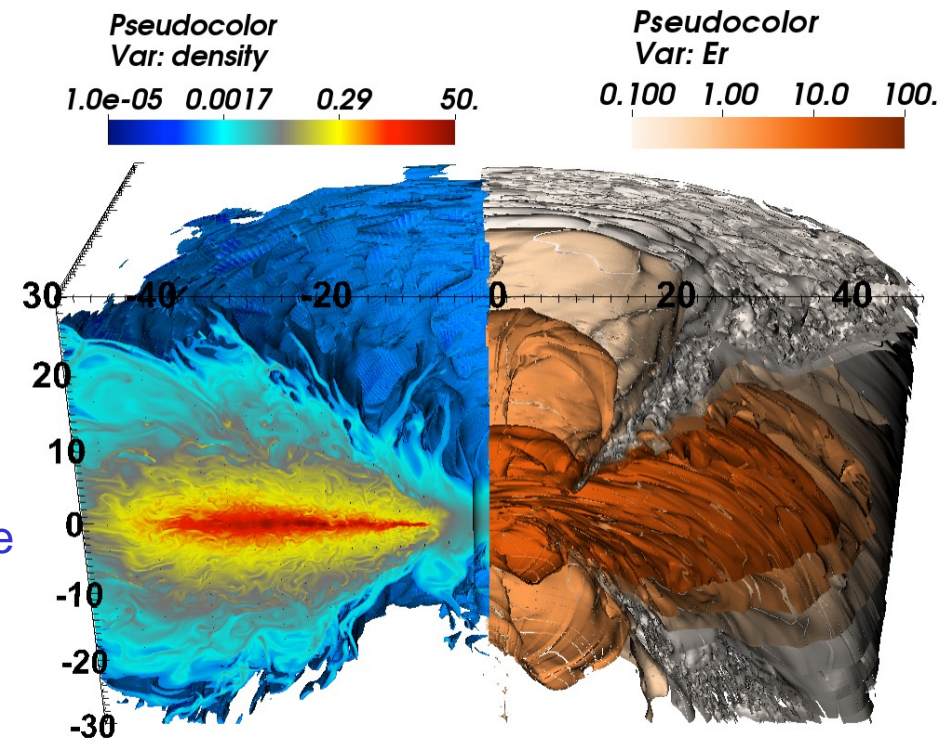
* Available capacity numbers are based on baseline resource distributions.

- NAS: Astro - 21.7%, Earth Science - 53.8%, Helio - 13.9%, Planetary - 10.6%
- Astrophysics allocation includes ~20M SBUs from the Agency reserve




Application Support: ATHENA

- Magnetohydrodynamics (MHD) code used to simulate black holes
- Previous I/O characterization
 - Identified bottleneck with IOT tool
 - Sped up I/O
 - 10x in 3264-rank case
 - 36x in 5968-rank case
- Current Performance Characterization
 - Developed set of kernels from full code
 - Using tools: Vector Advisor, Paraver, op_scope MPIprof
 - Transformations to reduce memory pressure yielded 12% improvement in a kernel
 - Investigating performance on Intel Xeon Phi cluster
- Next Steps
 - Conduct formal interview with PI
 - Investigate issues identified in interview
 - Write summary report



Snapshot of the global structure of a radiation-dominated black hole accretion disk with super-Eddington accretion rate from a global 3D radiation MHD simulation done in Pleiades. Left image shows the density, and the right shows the radiation energy density (both in units of the initial maximum value). Coordinate axes are labeled in units of Schwarzschild⁶ radius. Pleiades has been used for Athena/Athena++ code to perform simulations to study various astrophysical systems, including properties of black hole accretion disks, stability and structures of massive star envelopes as well as stream-stream collision in tidal disruption events.




Survey of Federal Agency HEC Capabilities

DoD, DoE, NSF

**Response to Heliophysics Subcommittee, NASA
Advisory Council**


**Survey Conducted by NASA HEC
NCCS, NASA Goddard Space Flight Center**

February 27, 2017
L. Parnell, NCCS Consultant



Federal Agency HEC Sites Surveyed

- DoD HPCMP
 - AFRL, ARL, ERDC, Navy/SSC, MHPCC
- DoE
 - NNSA (LLNL)
 - SC (ALCF, OLCF, NERSC)
- NSF
 - Centers: NCSA, PSC, SDSC, TACC
 - XSEDE




DoD HPCMP HEC Systems Summary

| Agency | Compute Nodes | | Aggregated TFLOPS |
|-------------|---------------|-----------|-------------------|
| | # Nodes | # Cores | |
| DoD HPCMP | 30,406 | 843,318 | 20,076 |
| DoE ASCR | 87,232 | 6,644,623 | 78,650 |
| DoE NNSA | 28,932 | 516,370 | 8,817 |
| NSF Centers | 38,288 | 1,042,360 | 26,338 |
| NSF XSEDE | ? | 220,548 | 16,556 |
| NASA | 11,472 | 246,048 | 7,250 |

Caveats:

- Not all systems are reported.
- Inconsistent reporting especially when accelerators are used.



Agency FY 2017 Budget Requests and Prior-Year HEC Funding (\$M)

High-Capability Computing Systems Infrastructure & Applications

| Agency \ PCA | 2017 HCSIA | 2016 HCSIA Estimates | 2015 HCSIA Actuals | Enabling R&D (EHCS) | Software & Systems (HCSS) |
|-----------------------|------------|----------------------|--------------------|---------------------|---------------------------|
| NSF | 183.2 | 180.4 | 189.0 | 131.0 | 86.5 |
| DoD ^b | 81.9 | 80.8 | 80.8 | 216.4 | 12.9 |
| DOE ^c | 393.6 | 374.4 | 378.8 | 208.3 | 17.5 |
| NIH ^d | 194.6 | 194.6 | 194.6 | 23.1 | 30.0 |
| NIST | 8.1 | 8.1 | 8.1 | 18.0 | 15.7 |
| NASA | 60.9 | 62.8 | 67.3 | 11.0 | 4.9 |
| NOAA | 36.0 | 29.7 | 23.0 | | |
| NNSA | | | | 30.0 | |
| Total ^{a, d} | 958.3 | 930.8 | 941.6 | 647.5 | 167.5 |


Headings Key: PCA – Program Component Area

EHCS – Enabling R&D for High-Capability Computing Systems

HCSIA – High-Capability Computing Systems Infrastructure & Applications


HCSS – High Confidence Software & Systems

Source: [FY2017 NITRD Supplement to the President's Budget](#)



Agency HEC Site Funding, Users & Allocations

| Agency | Site | FY 16 Funding (\$M) | Number of Projects | Number of Users | Allocations | Allocation Policy, Process (Methodology is very similar across all agencies and sites) |
|--------|----------|---------------------|--------------------|-----------------|--|--|
| HPCMP | Various* | 230 | | ~1,800 | Users: 95% MPO: 5% | Policy: Allocate 90-95% of user-available machine resources. Retain 5-10% for discretionary use. Process: 1.Projects submit proposals (quarterly to annually) 2.Requirements are validated 3.Proposals are evaluated & ranked 4.Allocations are made to approved proposals (quarterly to annually) 5.Progress is reviewed, assessed |
| DoE | LLNL | | | 3,000+ | | |
| | ALCF | 77 | 350 | 990 | INCITE: 60% Challenge:30% % Sites: 10% | |
| | OLCF | 104 | 316 | 1176 | | |
| | NERSC | 86 | 750 | 5000+ | | |
| NSF | NCSA | | | 3,500+ | Users: >90% Sites: <= 10% | |
| | PSC | | | | | |
| | SDSC | | | | | |
| | TACC | | | | | |
| | XSEDE | | | | | |




Backup Information



List of Acronyms

- ALCF = Argonne Leadership Computing Facility
- DoD = Department of Defence
- DoE = Department of Energy
- FDR = 56Gb/s serial high bandwidth links
- HPCMP = High Performance Computing Modernization Program
- LLNL = Los Alamos National Laboratory
- NERSC = National Energy Research Scientific Computing Center
- NCSA = National Center for Supercomputing Applications
- NNSA = National Nuclear Security Administration
- NSF = National Science Foundation
- OLCF = Oak Ridge Leadership Computing Facility
- PF = Petaflops = 10^{15} floating-point operations per second
- PSC = Pittsburgh Supercomputing Center
- SBU = Standard Billing Unit for tracking computing usage (<https://hec.nasa.gov/user/policies/sbus.html>)
- SDSC = San Diego Supercomputing Center
- TACC = Texas Advanced Computing Center
- TB = 10^{12} Bytes
- TF = Teraflops = one million million (10^{12}) floating-point operations per second
- XSEDE = Extreme Science and Engineering Discovery Environment




Survey of Federal Agency HEC Capabilities

DoD, DoE, NSF

**Response to Heliophysics Subcommittee, NASA
Advisory Council**


**Survey Conducted by NASA HEC
NCCS, NASA Goddard Space Flight Center**

February 27, 2017
L. Parnell, NCCS Consultant



Federal Agency HEC Sites Surveyed


- DoD HPCMP
 - AFRL, ARL, ERDC, Navy/SSC, MHPCC
- DoE
 - NNSA (LLNL)
 - SC (ALCF, OLCF, NERSC)
- NSF
 - Centers: NCSA, PSC, SDSC, TACC
 - XSEDE



DoD HPCMP HEC Systems Summary

Number of users: ~1,800


| Vendor / Model | Compute Nodes (Standard / Large + Phi, GPU) | | | Peak TFLOPS |
|----------------|--|---------|------------------|-------------|
| | # Nodes | # Cores | GB Memory / Node | |
| Cray XC30 | 1,215 / 8 + 124 | 30,592 | 64 128 | 786 |
| Cray XC30 | 2,370/ 0 + 32 | 57,200 | 64 | 1280 |
| Cray XC30 | 1,183 / 8 + 156 | 30,144 | 64 128 | 822 |
| Cray XC40 | 1,523 / 8 + 168 | 51,008 | 128 512 | 2000 |
| Cray XC40 | 3,098/32+ 32 | 101,184 | 128 512 | 3770 |
| Cray XC40 | 1,523 / 8 + 168 | 51,008 | 128 512 | 2000 |
| Cray XE6 | 4,716 | 150,912 | 64 | 1500 |
| Cray XE6m | 460 | 14,720 | 64 | 135 |
| IBM IDPx | 1,220 / 4 + 12 | 19,776 | 32 256 | 411 |
| IBM IDPx | 756 | 12,096 | 32 | 252 |
| SGI ICE X | 4,590 | 73,440 | 32 | 1500 |
| SGI ICE X | 3,216 / 4 + 356 | 125,888 | 128 768 | 5620 |
| SGI ICE X | 3,456 / 4 + 32 | 125,440 | 128 1 | 4660 |



DoE ASCR HEC Systems Summary

Number of users: 7,166+

| Site / System | Vendor / Model | Compute Nodes | | | Peak TFLOPS | Notes |
|-------------------------------|----------------|--------------------------|------------------------------|---------------------------------------|-------------|-----------------------------------|
| | | # Nodes | # Cores | GB Memory / Node | | |
| ORNL - LCF | | | | | | |
| Titan - Cray XK7 | | 18,688 | AMD: 4,784,128 | 32GB + 6GB | 27,000 | +GPUs: 18,688 |
| Eos - Cray XC30 | | | 11,776 | 64 | | Titan Support |
| ANL LCF | | | | | | |
| Mira - IBM / Blue Gene/Q | | 49,152 | 786,432 | 16 | 10,000 | |
| Theta - Intel (Xeon Phi)/Cray | | >2,500 | 231,935 | 208 | 9,650 | |
| NERSC | | | | | | |
| CORI - Cray XC40 (Xeon, Phi) | | Xeon: 1,988 / Phi: 9,304 | Haswell: 63,616 Phi: 632,672 | Haswell: 203 TB Phi: 1 PB (aggregate) | 30,000 | + PDSF & Genepool – dedicated sys |
| Edison - Cray XC30 | | 5.600 | 134,064 | 64 | 2,000 | |




DoE NNSA HEC Systems Summary

Livermore Computing General Availability Compute Platforms

Number of Users: 3,000+

| System | Nodes | Total Cores | Total Memory (GB) | TFLOPS Peak | Vendor | Program(s) |
|--------|--------|-------------|-------------------|-------------|---------|------------------|
| Ansel | 324 | 3,888 | 7,776 | 43.5 | Dell | M&IC |
| Cab | 1,296 | 20,736 | 41,472 | 431.3 | Appro | ASC, M&IC |
| Vulcan | 24,576 | 393,216 | 393,216 | 5,033.2 | IBM | ASC, M&IC, HPCIC |
| Quartz | 2,688 | 96,768 | 344,064 | 3,251.4 | Penguin | ASC, M&IC |
| Borax | 48 | 1,762 | 6,144 | 58.1 | Penguin | ASC |



NSF Centers HEC Resources

NCSA, PSC, SDSC, TACC

| Site | Vendor / Model | Compute Nodes (Regular / Large / Extreme Memory) | | | Peak TFLOPS | Notes |
|-----------------------|-----------------------|---|---------|--------------------------|----------------|------------------------------------|
| | | # Nodes | # Cores | GB Memory/ Node | | |
| NCSA: Blue Waters | Cray XE6 / XK7 hybrid | 22,640 / 4,228 | 396,032 | 32 / 64 | 13,340 | XK: +4,228 GPUs |
| PSC: Bridges* | Intel Xeon | 800 / 42 / 4 | 27,040 | 128 / 3,000 / 12,000 | 895 / 895 | |
| SDSC: | | | | | | |
| Comet* | Dell Intel Xeon | 1,944 / 4 | 47,776 | 128 (+320 Flash) / 1,500 | 2,000 | + 144 GPUs |
| Gordon* | Appro/ Intel | 1,024 | 16,384 | 64 | 341 | 300 TB flash |
| Triton shared cluster | Vars. | | | ~4,000 | 100+ | UCSD/UC "condo" |
| TACC: | | | | | | |
| Stampede* | Dell Intel | 6,400 | 522,080 | 40 | 9,600 | Xeon + PHI |
| Lonestar 5 | Cray XC40 | 1,252 / 2 / 8 | 30,048 | 64 / 1,000 / 512 | | Academic |
| Wrangler* | Dell/ EMC (10PB DSSD) | | 3,000+ | | 62 | Data Analysis, IU, UC - Replicated |

*Access and allocations provided through XSEDE




NSF XSEDE Compute Resources & Usage*

Number of Users: 3,500+

| Name | Site | CPUs | Peak TFlops | Utilization | Running Jobs | Queued Jobs | Other Jobs |
|-------------------|------------|--------|-------------|-------------|--------------|-------------|------------|
| Stampede | UT Austin | 102400 | 9600 | 86% | 814 | 281 | 108 |
| Comet | SDSC | 47616 | 2000 | 91% | 1734 | 2334 | 137 |
| XStream | Stanford | 1300 | 1012 | 52% | 68 | 9 | 86 |
| SuperMIC | LSU CCT | 7200 | 925 | | | | |
| Bridges Reg. Mem | PSC | 21056 | 895 | | | | |
| Bridges Large Mem | PSC | 160 | 895 | | | | |
| Jetstream | Indiana | 15360 | 516 | | | | |
| Gordon Cluster | SDSC | 16384 | 341 | 20% | 123 | 32 | 19 |
| Beacon | NICS | 768 | 210 | | | | |
| Red | UN Lincoln | 6000 | 100 | | | | |
| Wrangler | UT Austin | 2304 | 62 | | | | |

XSEDE: a single, virtual cyberinfrastructure that scientists use to interactively share computing resources, data, and expertise.

*Usage provided real time. Sample shown obtained 2/13/17 at 3:00 EST.



Agency FY 2017 Budget Requests and Prior-Year HEC Funding (\$M)

High-Capability Computing Systems Infrastructure & Applications

| Agency \ PCA | 2017 HCSIA | 2016 HCSIA Estimates | 2015 HCSIA Actuals | Enabling R&D (EHCS) | Software & Systems (HCSS) |
|-----------------------|------------|----------------------|--------------------|---------------------|---------------------------|
| NSF | 183.2 | 180.4 | 189.0 | 131.0 | 86.5 |
| DoD ^b | 81.9 | 80.8 | 80.8 | 216.4 | 12.9 |
| DOE ^c | 393.6 | 374.4 | 378.8 | 208.3 | 17.5 |
| NIH ^d | 194.6 | 194.6 | 194.6 | 23.1 | 30.0 |
| NIST | 8.1 | 8.1 | 8.1 | 18.0 | 15.7 |
| NASA | 60.9 | 62.8 | 67.3 | 11.0 | 4.9 |
| NOAA | 36.0 | 29.7 | 23.0 | | |
| NNSA | | | | 30.0 | |
| Total ^{a, d} | 958.3 | 930.8 | 941.6 | 647.5 | 167.5 |


Headings Key: PCA – Program Component Area

EHCS – Enabling R&D for High-Capability Computing Systems

HCSIA – High-Capability Computing Systems Infrastructure & Applications


HCSS – High Confidence Software & Systems

Source: [FY2017 NITRD Supplement to the President's Budget](#)



Agency HEC Site Funding, Users & Allocations

| Agency | Site | FY 16 Funding (\$M) | Number of Projects | Number of Users | Allocations | Allocation Policy, Process (Methodology is very similar across all agencies and sites) |
|--------|----------|---------------------|--------------------|-----------------|---|--|
| HPCMP | Various* | 230 | | ~1,800 | Users: 95% MPO: 5% | Policy: Allocate 90-95% of user-available machine resources. Retain 5-10% for discretionary use. Process: 1.Projects submit proposals (quarterly to annually) 2.Requirements are validated 3.Proposals are evaluated & ranked 4.Allocations are made to approved proposals (quarterly to annually) 5.Progress is reviewed, assessed |
| DoE | LLNL | | | 3,000+ | | |
| | ALCF | 77 | 350 | 990 | INCITE: 60% Challenge:30 % Sites: 10% | |
| | OLCF | 104 | 316 | 1176 | | |
| | NERSC | 86 | 750 | 5000+ | | |
| NSF | NCSA | | | 3,500+ | Users: >90% Sites: <= 10% | |
| | PSC | | | | | |
| | SDSC | | | | | |
| | TACC | | | | | |
| | XSEDE | | | | | |




Detailed Agency HEC Resource Allocation Process Descriptions

A. Department of Defense HPCMP

1. Requirements Determination

- Each DoD Service/Agency defines a set of computational projects supporting RDT&E program
 - DoD Services: Army, Navy, Air Force
 - DoD Agencies: DTRA, DARPA, MDA
- Project requirements are captured in an annual survey of computational project leaders
 - Survey captures requirements for upcoming fiscal year and four years into the future
 - Extensive data collected on each project
 - Project description
 - Hardware requirements – core-hours per year per system
 - Applications software requirement
 - Individual job requirements
 - Training requirements
- Each Service/Agency validates its overall requirements through HPC Advisory Panel Principals of each Service and Agency
Requirements tend to exceed available HPC cycles by 3X



Detailed Agency HEC Resource Allocation Process

Descriptions

A. DoD HPCMP, continued

2. Allocation Management Process

- 100% of all available computational resources at DOD Supercomputing Resource Centers are allocated
 - Based on theoretical maximum of 8,760 hours per year
- 5% reserved for HPCMPO allocations
- 25% allocated to Frontier Projects
 - These are high-impact, computationally-intensive projects
 - Selections based on technical merit review as well as Service/Agency prioritization
- 70% allocated to Services/Agencies for distribution
 - 30/30/30/10 (Army/Navy/Air Force/Agencies) split
 - Mechanisms in place support Service/Agency allocations for high priority, advance reservations, and dedicated workloads
- Overall limit on high-priority work of 35%
 - Frontier and Service/Agency high priority projects



Detailed Agency HEC Resource Allocation Process

Descriptions

B. Department of Energy's National Laboratories


- The Office of Science computers are dedicated to tackling the modeling, simulation, and big datasets behind ambitious research.
 - The projects that use these facilities range from needing 10 million to more than 100 million hours of computing time.
- The DoE-SC user facilities' dedication to the larger research community makes them particularly unique.
 - They allocate time to researchers based on the strength of competitive scientific proposals.
 - As with all DOE user facilities, they are free to use as long as the scientist makes their data available to the larger scientific community. Discretionary allocations consume 10% of resources.
- The DOE's INCITE program provides 60% of allocations to computationally intensive, large-scale research projects that aim to address "grand challenges" in science and engineering.
 - The program conducts a two-part review of all proposals: a peer review by an international panel of experts and a computational-readiness review.
 - The annual call for proposals is issued in April, and the allocations are awarded in millions of core-hours for one to three years.



Detailed Agency HEC Resource Allocation Process Descriptions

B. DoE National Laboratories, continued

- The DOE's ALCC program allocates 30% of resources to projects directly related to the DOE's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources.
 - The DOE conducts a peer review of all proposals based on scientific and technical merit of the project; appropriateness of the proposed method or approach; competency and adequacy of personnel and proposed resources; and the reasonableness and appropriateness of the proposed allocation request.
 - The yearlong allocation cycle runs from July 1 to June 30.
- The NNSA unclassified Livermore Computing facilities are allocated in a manner similar to that of the DoD HPCMP.
 - Project leaders submit proposals for resources based on their requirements. Proposals are evaluated and available appropriate facility resources are allocated based on the requirements




Detailed Agency HEC Resource Allocation Process

Descriptions

C. National Science Foundation

- Allocations are made through the Extreme Science and Engineering Discovery Environment (XSEDE), which is composed of multiple partner institutions known as Service Providers, each of which contributes one or more allocatable services.
 - Resources include High Performance Computing (HPC) machines, High Throughput Computing (HTC) machines, visualization, data storage, testbeds, and services.
- An XSEDE allocation is a grant that provides approved research projects with 12 months of access to the XSEDE cyberinfrastructure.
 - Allocations may include computing and visualization time, data storage space, and extended collaborative support services at one or more XSEDE service provider sites.
 - Allocations are awarded to researchers who serve as principal investigators (PIs) of approved projects.
 - After receiving an allocation, each project's PI can create XSEDE accounts for other researchers on the project team, giving them access to the project's allocated digital services.




Detailed Agency HEC Resource Allocation Process Descriptions

C. National Science Foundation, continued

XSEDE offers four types of allocations to accommodate various use cases:

- **Startup allocations** -intended primarily for researchers who are new to the XSEDE environment, enabling them to quickly gain access to XSEDE digital services.
 - They are well suited for researchers who need to develop their codes, test their applications, or experiment with XSEDE systems.
 - Faculty members and postdoctoral researchers at US academic research institutions, as well as K-12 teachers, are welcome submit requests for Startup allocations at any time
 - Startup allocations provide each project with a maximum of 200,000 service units to be used within the 12-month time limit, on either a single computational system or a combination of systems.
- **Education allocations** - specifically designed for academic instruction or training activities, and are intended for faculty who want to incorporate XSEDE digital services into to their courses about the use of advanced cyberinfrastructure technologies.
 - PIs on Education allocations must be faculty at a US academic institution. You are welcome to submit a request for an Education allocation at any time through the XSEDE User Portal.
 - Education allocations provide each project with a maximum of 200,000 service units to be used within the 12-month time limit.
 - You may request time on either a single computational system or a combination of systems.




Detailed Agency HEC Resource Allocation Process


Descriptions

C. National Science Foundation, continued

- **Research allocations** - reserved for projects with requirements that Startup allocations cannot accommodate.
 - With a Research allocation, an approved project can make extensive use of XSEDE's high-performance and high-throughput compute systems, scientific visualization systems, and data storage systems, and receive expert consulting and support from Extended Collaborative Support Service (ECSS) staff.
 - Requests for Research allocations are reviewed four times a year by the XSEDE Resource Allocations Committee (XRAC).
 - Research allocation requests may be submitted using the XSEDE User Portal, but each request must include a detailed document that describes the work the project plans to accomplish and justifies the need for the requested resources.
- **Campus Champions allocations** - reserved for the Campus Champions program which establishes representatives at US educational institutions who distribute knowledge of XSEDE's advanced cyberinfrastructure technologies to the researchers, educators, and students on their campuses.
 - Campus Champions help users on their campuses set up Startup allocations and provide them with the XSEDE contacts they need to resolve any problems.
 - Campus Champions are also responsible for hosting training and information sharing sessions on their campuses to make researchers, faculty, and students aware of XSEDE's resources and services, and how XSEDE can help them achieve their research goals.



Additional/Backup Information



DoE ASCR Computing Upgrades At a Glance

| System attributes | NERSC Now | OLCF Now | ALCF Now | | NERSC Upgrade | OLCF Upgrade | ALCF Upgrades |
|------------------------------|--------------------------------|-----------------------------|----------------------------|---|---|--|--|
| Name Planned Installation | Edison | TITAN | MIRA | Theta 2016 | Cori 2016 | Summit 2017-2018 | Aurora 2018-2019 |
| System peak (PF) | 2.6 | 27 | 10 | >8.5 | > 30 | 200 | 180 |
| Peak Power (MW) | 2 | 9 | 4.8 | 1.7 | < 3.7 | 13.3 | 13 |
| Total system memory | 357 TB | 710TB | 768TB | >480 TB DDR4 + High Bandwidth Memory (HBM) | ~1 PB DDR4 + High Bandwidth Memory (HBM)+1.5PB persistent memory | > 2.4 PB DDR4+ HBM + 3.7 PB persistent memory | > 7 PB High Bandwidth On-Package Memory Local Memory and Persistent Memory |
| Node performance (TF) | 0.460 | 1.452 | 0.204 | > 3 | > 3 | > 40 | > 17 times Mira |
| Node processors | Intel Ivy Bridge | AMD Opteron Nvidia Kepler | 64-bit PowerPC A2 | Intel Knights Landing Xeon Phi many core CPUs | Intel Knights Landing many core CPUs Intel Haswell CPU in data partition | Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS | Knights Hill Xeon Phi many core CPUs |
| System size (nodes) | 5,600 nodes | 18,688 nodes | 49,152 | >2,500 nodes | 9,300 nodes 1,900 nodes in data partition | ~4,600 nodes | >50,000 nodes |
| System Interconnect | Aries | Gemini | 5D Torus | Aries | Aries | Dual Rail EDR-IB | 2 nd Generation Intel Omni-Path Architecture |
| File System | 7.6 PB 168 GB/s, Lustre® | 32 PB 1 TB/s, Lustre® | 26 PB 300 GB/s GPFS™ | 10PB, 210 GB/s Lustre initial | 28 PB 744 GB/s Lustre® | 120 PB 1 TB/s GPFS™ | 150 PB 1 TB/s Lustre® |



Livermore Computing Systems Summary

| System | Top500 Rank TFLOP/s | Program | Manufacture/ Model | Processor Architecture | OS | Inter- connect | Avg Power Demand (KW) | Nodes | Cores | Memory (GB) | Peak |
|-----------------------------------|------------------------|----------------|-----------------------|--|----------|-------------------|--------------------------------|--------|-----------|----------------|----------|
| Unclassified Network (OCF) | | | | | | | | | | | |
| Vulcan | 21 | ASC+M&IC+HPCIC | IBM BGQ | IBM PowerPCA2 | RHEL/CNK | SD Torus | TBD | 24,576 | 393,216 | 393,216 | 5,033.2 |
| Cab (TLCC2) | | ASC+M&IC+HPCIC | Appro | Intel Xeon E5-2670 | TOSS | IB QDR | 564 | 1,296 | 20,736 | 41,472 | 426.0 |
| Quartz (CTS-1) | 41 | ASC+M&IC | Penguin | Intel Xeon E5-2695v4 | TOSS | Omni-Path | TBD | 2,688 | 96,768 | 344,064 | 3251.4 |
| Ansel | | M&IC | Dell | Intel Xeon EP X5660 | TOSS | IB QDR | TBD | 324 | 3,888 | 7,776 | 43.5 |
| RZMerl (TLCC2) | | ASC+ICF | Appro | Intel Xeon E5-2670 | TOSS | IB QDR | TBD | 162 | 2,592 | 5,184 | 53.9 |
| RZZeus | | M&IC | Appro | Intel Xeon E5-2670 | TOSS | IB QDR | TBD | 162 | 2,592 | 5,184 | 53.9 |
| RZManta | | ASC | IBM | IBM Power8+ | RHEL | IB EDR | TBD | 36 | 720 | 11,520 | 597.6 |
| Ray | | ASC+M&IC | IBM | IBM Power8+ | RHEL | IB EDR | TBD | 54 | 1,080 | 17,280 | 896.4 |
| Catalyst | | ASC+M&IC | Cray | Intel Xeon E5-2695v2 | TOSS | IB QDR | TBD | 324 | 7,776 | 41,472 | 149.3 |
| Syrah | | ASC+M&IC | Cray | Intel Xeon E5-2670 | TOSS | IB QDR | TBD | 324 | 5,184 | 20,736 | 107.8 |
| Surface | | ASC+M&IC | Cray | Intel Xeon E5-2670 | TOSS | IB FDR | TBD | 162 | 2,592 | 41,500 | 451.9 |
| Borax (CTS-1) | | ASC+M&IC | Penguin | Intel Xeon E5-2695v4 | TOSS | N/A | TBD | 48 | 1,728 | 6,144 | 58.1 |
| RZTrona (CTS-1) | | ASC | Penguin | Intel Xeon E5-2695v4 | TOSS | N/A | TBD | 48 | 1,728 | 6,144 | 58.1 |
| Herd | | M&IC | Appro | AMD Opteron 8356, 6128 Intel XE7-4850 | TOSS | IB DDR | 7 | 9 | 256 | 1,088 | 1.6 |
| OCF Totals | Systems | 14 | | | | | | | | | 11,149.4 |
| Classified Network (SCF) | | | | | | | | | | | |
| Pinot (TLCC2, SNSI) | | M&IC | Appro | Intel Xeon E5-2670 | TOSS | IB QDR | TBD | 162 | 2,592 | 10,368 | 53.9 |
| Sequoia | 4 | ASC | IBM BGQ | IBM PowerPCA2 | RHEL/CNK | SD Torus | TBD | 98,304 | 1,572,864 | 1,572,864 | 20132.7 |
| Zin (TLCC2) | 164 | ASC | Appro | Intel Xeon E5-2670 | TOSS | IB QDR | TBD | 2,916 | 46,656 | 93,312 | 961.1 |
| Jade (CTS-1) | 42 | ASC | Penguin | Intel Xeon E5-2695v4 | TOSS | Omni-Path | TBD | 2,688 | 96,768 | 344,064 | 3251.4 |
| Shark | | ASC | IBM | IBM Power8+ | RHEL | IB EDR | TBD | 36 | 720 | 11,520 | 597.6 |
| Max | | ASC | Appro | Intel Xeon E5-2670 | TOSS | IB FDR | TBD | 324 | 5,184 | 82,944 | 107.8 |
| Agate (CTS-1) | | ASC | Penguin | Intel Xeon E5-2695v4 | TOSS | N/A | TBD | 48 | 1,728 | 6,144 | 58.1 |
| SCF Totals | Systems | 7 | | | | | | | | | 25,162.6 |
| Combined Totals | | 21 | | | | | | | | | 36,312.0 |

| System Category | TFLOP/s | % of Total |
|--------------------|---------|---------------|
|--------------------|---------|---------------|

| System Category | TFLOP/s | % of Total |
|--------------------|---------|---------------|
|--------------------|---------|---------------|