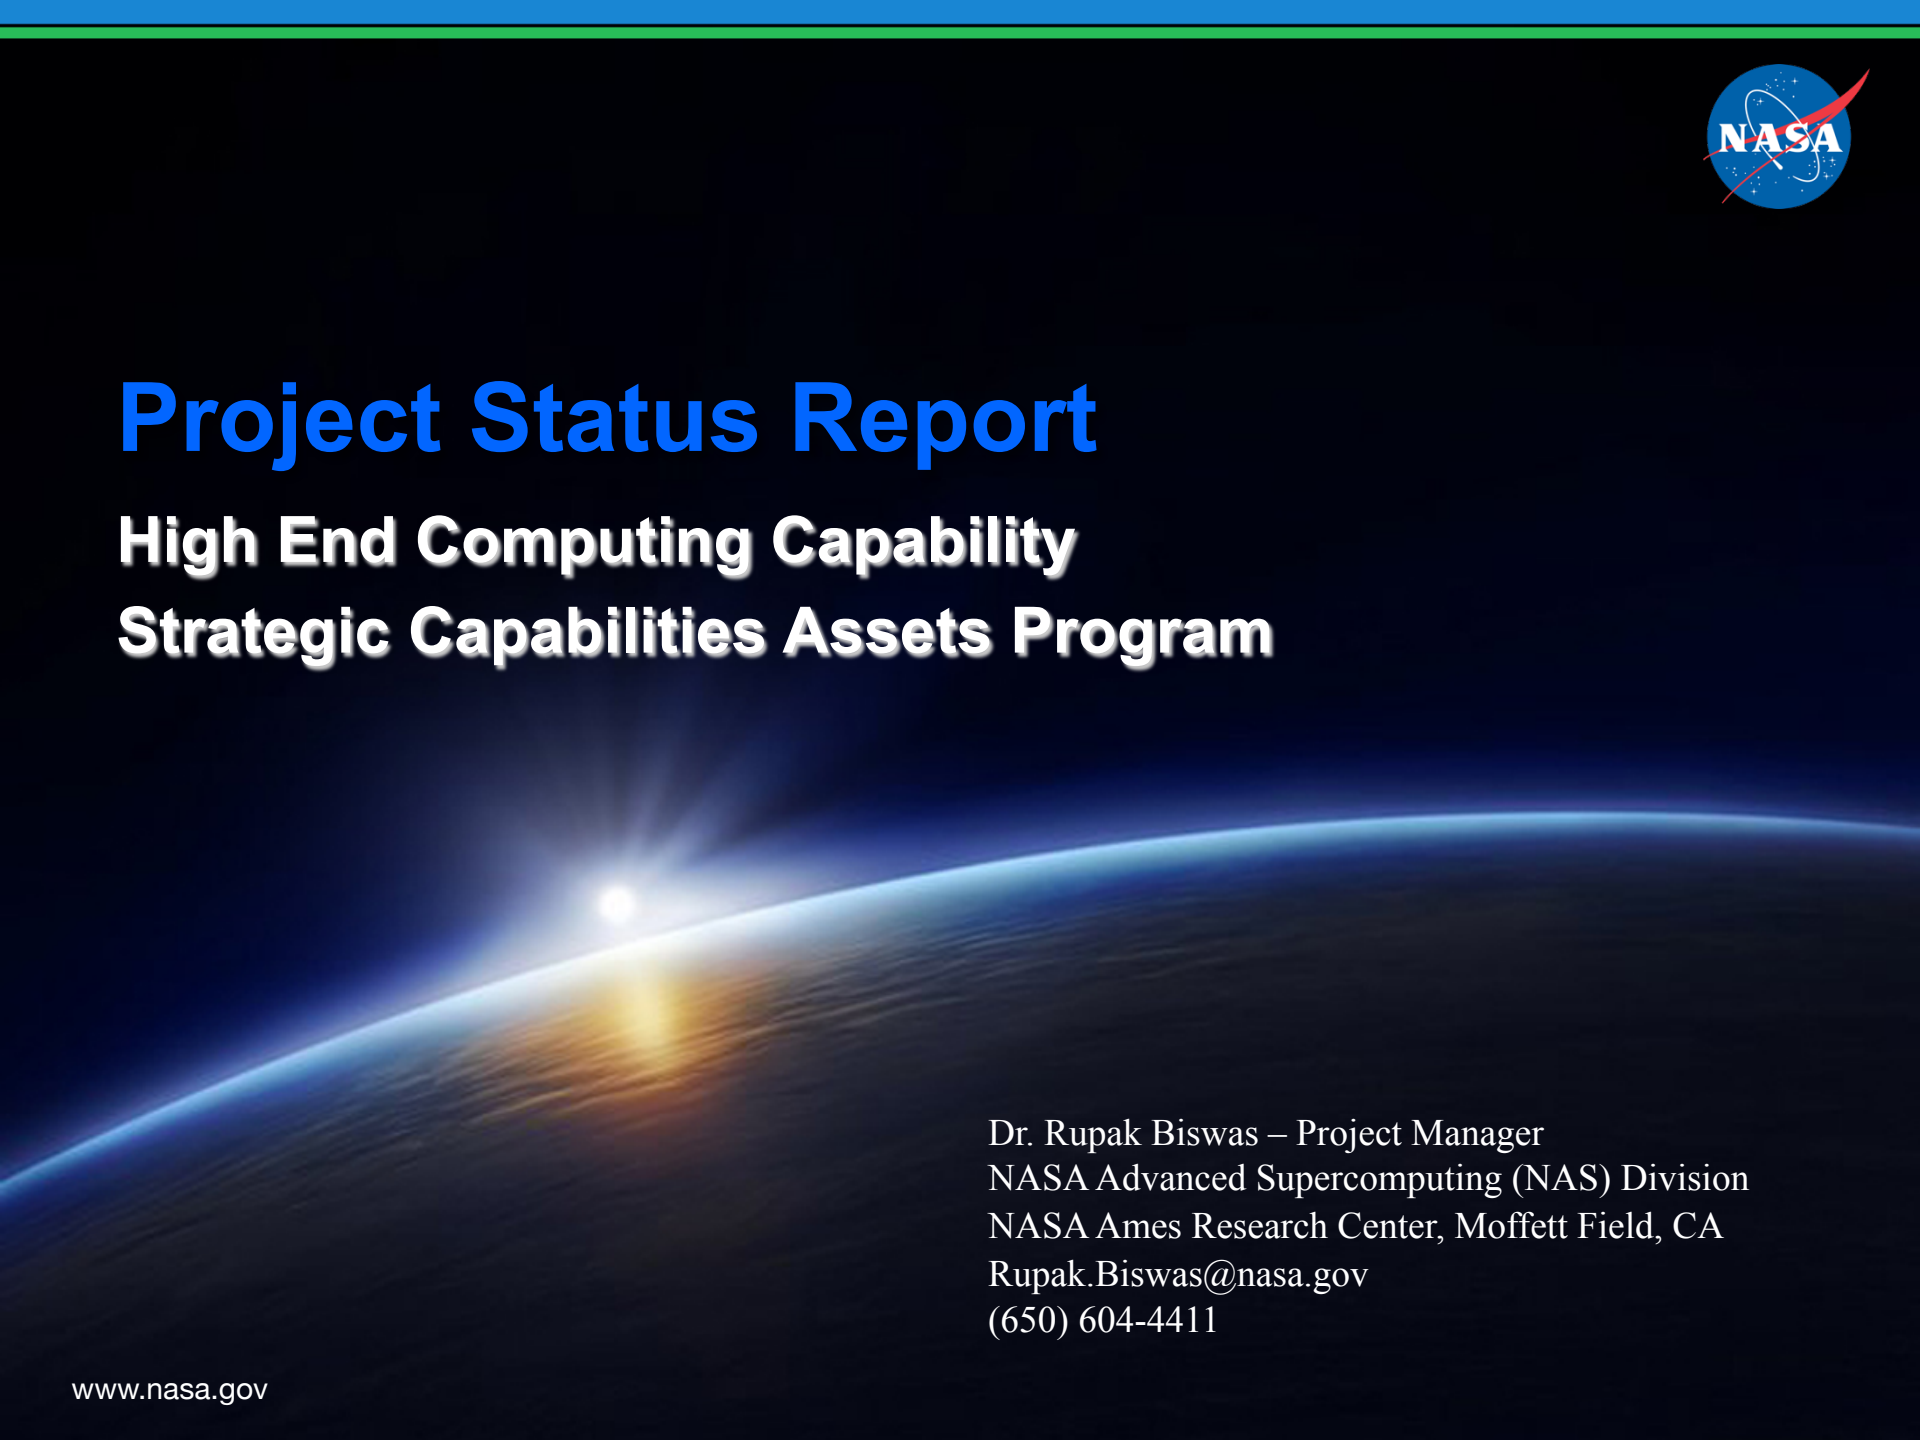

Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

Dr. Rupak Biswas – Project Manager
NASA Advanced Supercomputing (NAS) Division
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

New Testbed System, 'Maia,' Deployed for Performance Testing

- The HECC Supercomputing Systems team has deployed a new SGI testbed system named Maia, which contains Intel Xeon Phi coprocessors.
- Using the testbed system, the team will evaluate the Intel Xeon Phi technology for potential use in the next-generation HECC supercomputer.
- Maia is composed of 128 nodes, each containing 2 Sandy Bridge processors, 2 Xeon Phi coprocessors, and 64 gigabytes of memory, for a peak performance of 301.4 teraflops.
- HECC staff is working closely with SGI and Intel to mature the Xeon Phi coprocessor technology into a productive platform for users. To this end, the HECC staff has developed enhancements to the Phi computing environment to improve usability.

Mission Impact: Evaluating leading-edge technologies reduces risks and increases the probability of selecting cost-effective systems that provide the best solution to meet the Agency's supercomputing requirements.

Maia is a 128-node testbed system with a computing environment that combines dual Intel Xeon Phi coprocessors with Intel Sandy Bridge processors. It is named after the fourth brightest star in the Pleiades open star cluster.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

Dynamic PBSPro Queue Management Improves Responsiveness and Utilization

- The HECC Systems team enhanced Pleiades to dynamically manage queues to meet the needs of development jobs. Users running jobs in the Pleiades development queue are experiencing an 8% better turnaround time.
- In the past, a static set of queue resources was reserved for users' development work; these resources were at times either insufficient or oversized for the demand.
- Dynamic management resizes queue resources to meet demand at all times, with a particular emphasis on having free resources available during working hours when users are actively developing software.
- As a result, not only are more resources available to the developers at peak times, but over 100,000 additional SBUs are available for production jobs per month.

Mission Impact: HECC's dynamic queue management approach increases user productivity as software is developed and improved, and increases overall system utilization for both development and production uses.

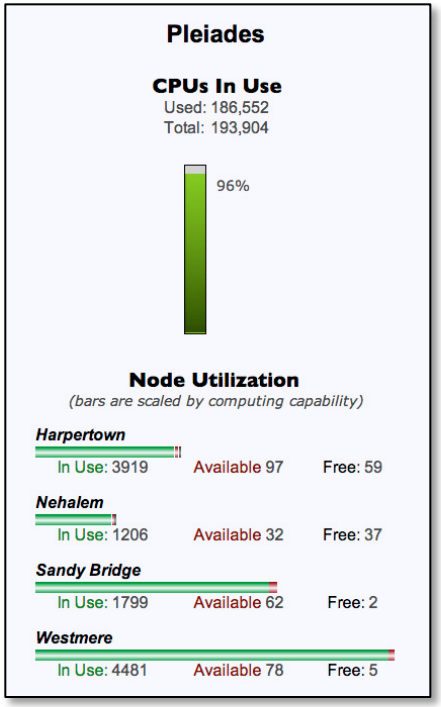

Dynamic resizing of the development queue on Pleiades helped achieve system utilization averages of over 90 percent, while meeting the diverse resource requirements for both quick turnaround software development jobs and large production jobs.

POC: Greg Matthews, gregory.matthews@nasa.gov, (650) 604-1321, NASA Advanced Supercomputing Division, Computer Sciences Corporation

Critical Maintenance Helps Secure Reliability and Availability of HECC Resources

- The HECC Facilities team coordinated critical maintenance work on electrical and mechanical infrastructure during a scheduled operational shutdown.
- Work was successfully completed on schedule, minimizing interruption to HECC users. The work included:
  - Cleaning of the chilled water plant (consisting of 4 chillers and the building's cooling tower) to mitigate the risk of extensive unscheduled downtime due to silt-clogged tubes and piping
  - Fire alarm system maintenance, in accordance with National Fire Protection Association regulations
  - Installation of Rotary Uninterruptable Power Supply (RUPS) power feeders into the building's high-voltage substation and switch gear
  - Various other building maintenance activities, including electrical infrastructure repair

Mission Impact: Periodic facilities maintenance helps ensure that critical HECC resources operate with optimal efficiency and reliability. In addition, a significant step toward completing the installation of the RUPS system was accomplished during the scheduled shutdown.

Workers clean part of the chilled water infrastructure outside of the NASA Advanced Supercomputing (NAS) Division building.

POC: Scott Prevost, scott.prevost@nasa.gov, (650) 604-4350, NASA Advanced Supercomputing Division, Computer Sciences Corp.

HECC High-Speed InfiniBand Network Extended Using New 'ColorChip' Technology

- The HECC team extended the high-speed InfiniBand network from the primary HECC facility to a secondary facility (>1 km away), using the new ColorChip technology.
  - ColorChip uses existing single-mode fiber between the two facilities to establish a Quad Data Rate (QDR) 40 Gbps connection;
  - This high-speed connection enabled the team to mount a Lustre filesystem located in the primary facility (N258) onto a Pleiades Harpertown cluster located in a secondary facility (N233A);
  - Establishing the connection required extensive testing and collaboration to identify and eliminate errors.
- Implementing this new technology increases network performance and decreases Ames Research Center (ARC) campus infrastructure requirements.

Mission Impact: The new ColorChip technology enables long-distance InfiniBand interconnects at double the speed and 1/8 the required fiber infrastructure of previous technology.

The ColorChip DragonFly Pluggable Quad Small Form Factor (QSFP+) long-reach optics connector is designed to work up to 10 km using 40Gbps Ethernet and InfiniBand networking equipment.

POC: Harjot Sidhu, harjot.s.sidhu@nasa.gov, (650) 604-4935, NASA Advanced Supercomputing Division, Computer Sciences Corp.

HECC System Maintenance Activities Performed Successfully

- During April's scheduled facility downtime, the HECC Supercomputing System team performed maintenance activities on Pleiades and the archive storage systems. Systems testing on the Endeavour cluster was also performed.
- Pleiades activities included:
  - Replacing 245 InfiniBand cables;
  - Updating RAID firmware and NFS and Lustre software;
  - Enabling Quality of Service (QoS) on InfiniBand;
  - Reducing the number of multicast groups by a factor of 20x.
- Archive storage system activities included:
  - Reconfiguration of the fibre channel tape infrastructure to support the planned merger of the Lou archival systems into a single system;
  - Validating high availability functionality and filesystem performance of the Lou archival system;
  - Updating firmware on the RAID subsystem.

Mission Impact: Periodic maintenance performed on HECC systems ensures a more stable environment for users by identifying and addressing potential issues before they can unexpectedly impact system availability.

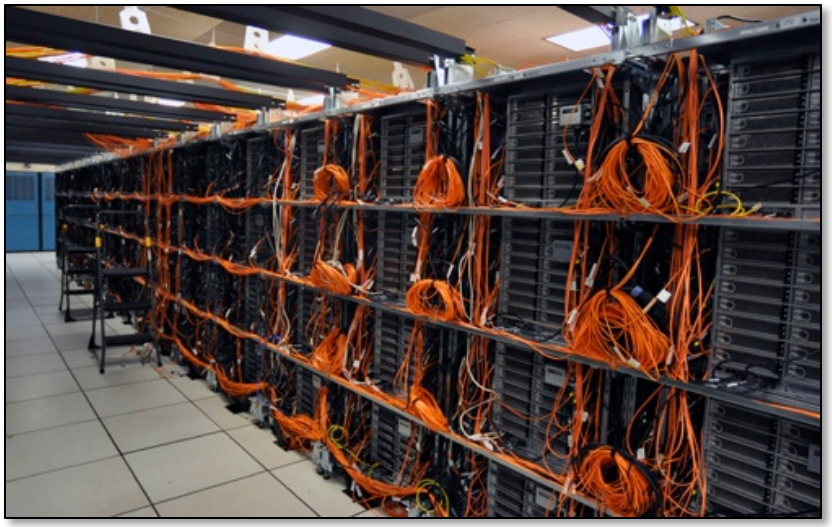

System maintenance activities performed during the recent scheduled downtime at the NAS facility included replacing 245 InfiniBand cables.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Science Corp.

Operating System Mapping Tool Helps to Improve User File Transfers

- The HECC Networks team implemented a tool that uses open-source (p0f) software to perform TCP fingerprinting of a remote user's operating systems (OS) in real-time. Network flows have unique characteristics that allow the source OS to be identified based on a database of TCP fingerprints.
- Data collected is filtered to show only the OS software that is being run by users connecting into the HECC enclave. Information about the OS is used when analyzing traffic flows to HECC systems.
- Specific tuning parameters are needed to optimize performance for different OS. This tool allows the Networks team to immediately inform users how to tune their system for faster file transfers.
- OS distribution mapping provides the Networks team with critical information that enables them to more easily troubleshoot and resolve users' connection problems.

Mission Impact: Identifying user operating systems in real-time as they connect to the HECC enclave provides critical information that can help optimize user file transfers and resolve any performance issues.

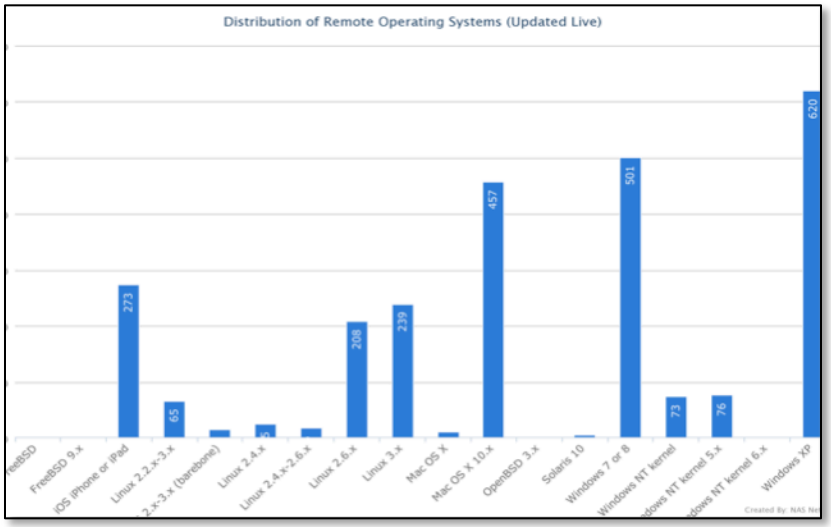

The chart shows an example of how output data from the OS Mapping tool can be used to determine how many users are running each operating system when they access HECC systems.

POC: Nichole Boscia, nichole.boscia@nasa.gov, (650) 604-0891, NASA Advanced Supercomputing Division, Computer Science Corporation

Code Optimization Saves Computing Time for Space Launch System Analyses

•Optimization experts in the HECC Application Performance and Productivity (APP) team solved a puzzling issue for a Pleiades user running 3D simulations of the Space Launch System (SLS) fuel tank.

•The user's job unexpectedly ran out of memory when using the "mpirun" command to start 128 or more processes on Pleiades.

•The APP investigation involved:

- Discovery of a configuration issue on an obscure Message Passing Interface (MPI) module used in the job script, which caused all MPI processes to run on a single node, rather than on 16 nodes of the 128-process job;
- This issue led to an out-of-memory message when the user tried to scale up the number of processors used for the job runs;
- APP staff reworked the PBS script to eliminate the problem, and made several other changes to enhance job performance and reliability.

•These improvements resulted in a savings of more than 90% of the computing used each time the job was run.

Mission Impact: Debugging and optimizing user code enables users to reduce the total runtime for simulations. With a faster turnaround time for their applications, HECC users are able to meet their milestones on time, and at reduced computational costs.

Optimization of user code run on Pleiades, pictured above, resulted in more than 90% savings of computing time for 3D SLS fuel tank simulations.

POCs: Johnny Chang, johnny.chang@nasa.gov, (650) 604-4356; Daniel Kokron, daniel.s.kokron@nasa.gov, (310) 286-3959, NASA Advanced Supercomputing Division, Computer Sciences Corporation

HECC Security Team Completes Successful Annual Testing of Security Controls

- The HECC Security team recently completed annual testing of security controls, as required by the Agency, to maintain authorization to operate information technology systems.
- Security plan certification lasts three years. During that time, a subset of the NIST SP 800-53 security controls are required to be tested each year–either through staff interviews or by examination of the actual control.
- This year, over the course of three months, 92 controls were tested against values identified in the NAS Security Plan. No issues were found and the results were uploaded to the Agency.
- The team will next conduct contingency plan training and testing, as required by the Agency.

Mission Impact: Annual security control testing is required by NASA to maintain authorization to operate information technology systems.

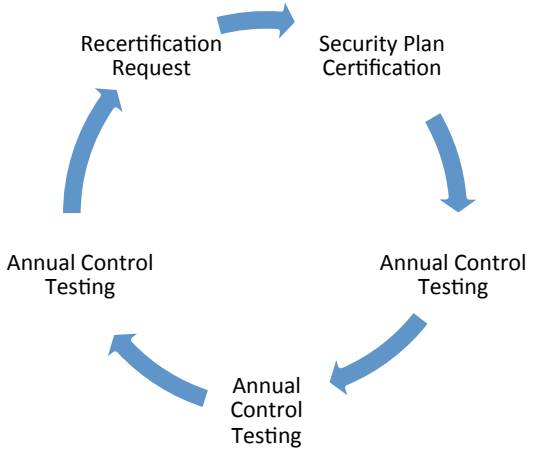

The three-year security plan and control testing cycle, as performed by the HECC Security team, in order to meet Agency information systems security requirements.

POC: Derek Shaw, derek.g.shaw@nasa.gov, (650) 604-4229, NASA Advanced Supercomputing Division, Computer Sciences Corp.

New Queue Improves Turnaround Time of Jobs for SMD Operational Projects

•A new queue, smd_ops, was created on Pleiades to improve turnaround time for jobs supporting the Science Mission Directorate's (SMD) operational projects, including:

- Kepler
- Interface Region Imaging Spectrograph (IRIS)
- Carbon in Arctic Reservoirs Vulnerability Experiment (CARVE)
- Arctic Boreal Vulnerability Experiment (ABoVE)

- The queue changes size dynamically based on demand and can provide priority access for up to approximately half of SMD's Pleiades resources.
- Dynamic resizing is done by creating and scheduling high priority jobs that move resources into the smd_ops queue when a user from one of the operational projects submits a job.
- If two or more operational projects submit jobs at the same time, the scheduler uses a "fair share" approach to distribute the resources.
- By restricting the queue to SMD resources, there is no impact on other Mission Directorates, while turnaround time for high-priority SMD projects is improved.

Mission Impact: Providing priority access to required resources allows NASA SMD researchers working on operational projects to meet their processing deadlines and accomplish their mission objectives.

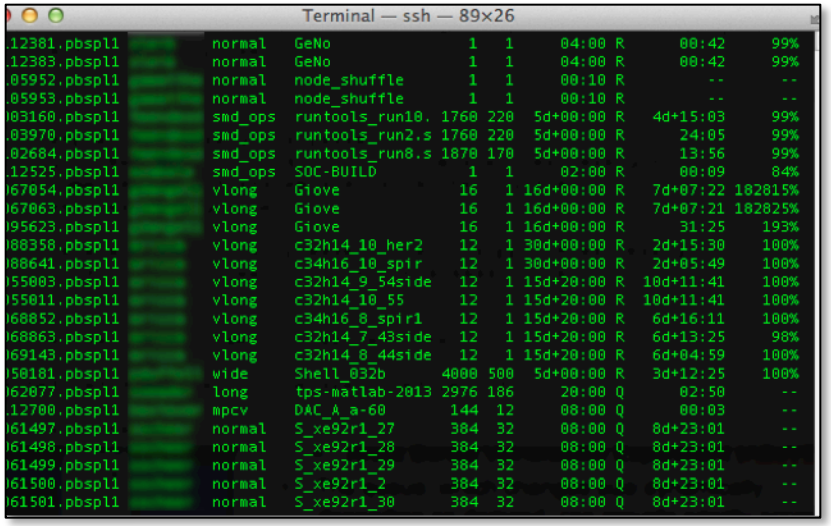

Status of a portion of the nearly 900 jobs running on Pleiades on May 3, 2013. Near the top are 4 jobs running in the new smd_ops queue. No jobs in the smd_ops queue were waiting for computing resources.

POC: Greg Matthews, gregory.matthews@nasa.gov, (650) 604-1321, NASA Advanced Supercomputing Division, Computer Sciences Corporation

New Features in Automated Transfer Tool Greatly Increase Performance, Convenience

- The HECC Supercomputing Systems team improved their existing Shift (Self-Healing Independent File Transfer) tool with additional features to speed up and simplify user transfer tasks.
  - Local and remote tar file creation/extraction with 8-host 64 gigabyte (GB) creation/extraction measured at 15x–30x faster than the standard tar utility.
  - Optimized remote transport makes parallelization of single-file remote transfers practical, with 8-host remote transfer of a single 128-GB file measured at 434 megabytes per second.
  - Native synchronization mechanism with 8-host sync of a single 64-GB file measured at 6.5x faster than the rsync utility.
  - Automatic Lustre striping transparently optimizes later I/O access by parallel jobs.
- These new features allow HECC users to spend less time managing file transfer details and waiting for transfers to complete.

Mission Impact: Enhancements to HECC's automated file transfer capabilities enable more efficient usage of scientific and engineering users' time, and result in higher productivity.

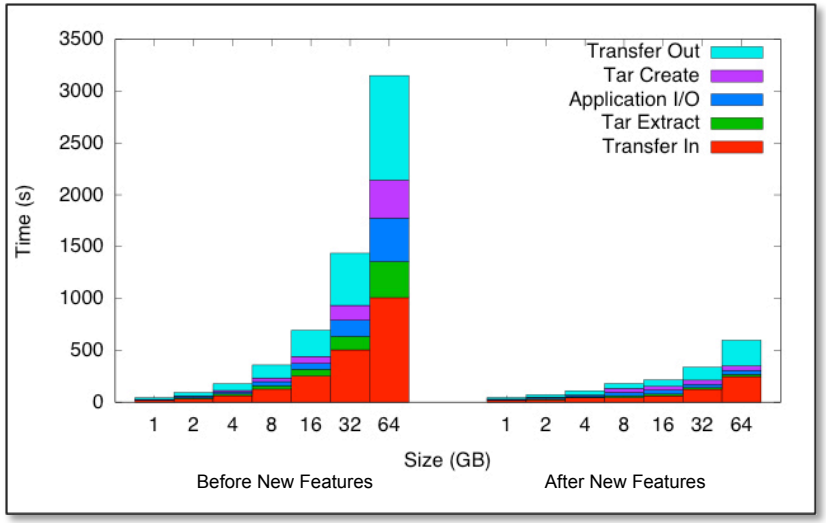

Through file transfer parallelization, embedded high-speed tar creation/extraction, and automatic Lustre striping, Shift can achieve significant reductions in turnaround time for the most common workflow activities.

POCs: Paul Kolano, paul.kolano@nasa.gov, (650) 604-4271; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Sciences Corporation

HECC Facility Hosts Several Visitors and Tours in April 2013

•HECC hosted 8 tour groups in April; guests learned about the agency-wide missions being supported by Pleiades, and viewed scientific results on the hyperwall system. Visitors this month included:

- A diverse group of researchers who were attending the AIAA Modeling and Simulation Technical Committee Spring Meeting, which was held at Ames;
- "CBS Sunday Morning" correspondent Barry Petersen, who interviewed members of the Kepler team while viewing their findings on the NAS hyperwall;
- Kevin H. Knuth of the Departments of Physics and Informatics, University at Albany, New York;
- A delegation of senior water officials from the government of Morocco, who were in the U.S. on a study tour that is being organized by USAID;
- A delegation from the Technical University of Denmark.

Aerospace Engineer Stu Rogers shows members of the "Solar Impulse" solar-powered aircraft team the main computer room of the NASA Advanced Supercomputing (NAS) facility during a tour this month.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputer Division

Papers and Presentations

- **"Multi-Species Turbulent Mixing Under Supercritical-Pressure Conditions: Modeling, Direct Numerical Simulation and Analysis Revealing Species Spinodal Decomposition,"** E. Masi, J. Bellan, K. G. Harstad, N. A. Okong'o, Journal of Fluid Mechanics, vol. 721, pp. 578-626, April 2013. *
  http://journals.cambridge.org/action/displayAbstract?fromPage=online&aid=8869499
- **"Saturation of the Magneto-Rotational Instability in Strongly Radiation-Dominated Accretion Disks,"** Y.-F. Jiang, J. M. Stone, S. W. Davis, The Astrophysical Journal, vol. 767, no. 2, April 5, 2013. *
  http://iopscience.iop.org/0004-637X/767/2/148
- **"Composition of Low Redshift Halo Gas,"** R. Cen, arXiv:1304.3466 [astro-ph.CO], April 11, 2013. *
  http://arxiv.org/abs/1304.3466
- **"The Clustering of Galaxies at z ≈ 0.5 in the SDSS-III Data Release 9 BOSS-CMASS Sample: a Test for the ΛCDM Cosmology,"** S. E. Nuza, et al., Monthly Notices of the Royal Astronomical Society, April 19, 2013. *
  http://mnras.oxfordjournals.org/content/early/2013/04/18/mnras.stt513.full
- **"Stability Enhanced High-Order Hyperviscosity-Based Shock Capturing Algorithm,"** U. K. Kaul, AIAA Journal (Early Edition), April 2013. *
  http://arc.aiaa.org/doi/pdf/10.2514/1.J051704

** HECC provided supercomputing resources and services in support of this work*

Papers and Presentations (continued)

- **"The Observable Thermal and Kinetic Sunyaev-Zel'dovich Effect in Merging Galazy Clusters,"** J. J. Ruan, et al., arXiv:1304.6088 [astro-ph.CO], April 22, 2013. * *http://arxiv.org/abs/1304.6088*
- **"Pleiades: NASA's First Petascale Supercomputer,"** R. Biswas, W. Thigpen, R. Hood, R. Ciotti, P. Mehrotra, C. Henze, J. Parks, B. Biegel, in Contemporary High Performance Computing: From Petascale Toward Exascale, April 23, 2013. *http://www.crcpress.com/product/isbn/9781466568341*
- **54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference,** April 8-11, 2013, Boston, Massachusetts
  - "Navier-Stokes Based Unsteady Aerodynamic Computations of Launch Vehicles Undergoing Forced Coupled Oscillations," G. Guruswamy * *http://arc.aiaa.org/doi/pdf/10.2514/6.2013-1887*
- **2013 OpenFabrics Alliance User Day Workshop,** April 19, 2013, Monterey, California
  - "NFS Over RDMA," J. Becker
- **HPC User Forum,** April 29-May 1, 2013, Tucson, Arizona
  - "Visualization Infrastructure at the NASA Advanced Supercomputing Facility," C. Henze
- **SGIUG Conference 2013,** April 30-May 2, 2013, San Diego, California
  - "Resource Matching for Combining Quick-turnaround and Batch Workloads in PBSPro," G. Matthews
  - "Setting Up the UV2000 at NASA Ames," H. Yeung
  - "pDMF at the NASA Advanced Supercomputing Facility," M. Cary

** HECC provided supercomputing resources and services in support of this work*

News and Events

•**NASA's Kepler Seeks to Answer: Is Anybody Out There?,** *CBS Sunday Morning,* April 28, 2013—CBS Sunday correspondent Barry Petersen visits Pleiades and the hyperwall-2 at NAS, and talks to Kepler scientists about the spacecraft's search for planets in the "Goldilocks" zone of stars which might contain life.
http://www.cbsnews.com/8301-3445_162-57581737/nasas-kepler-seeks-to-answer-is-anybody-out-there/

HECC Utilization
April 2013

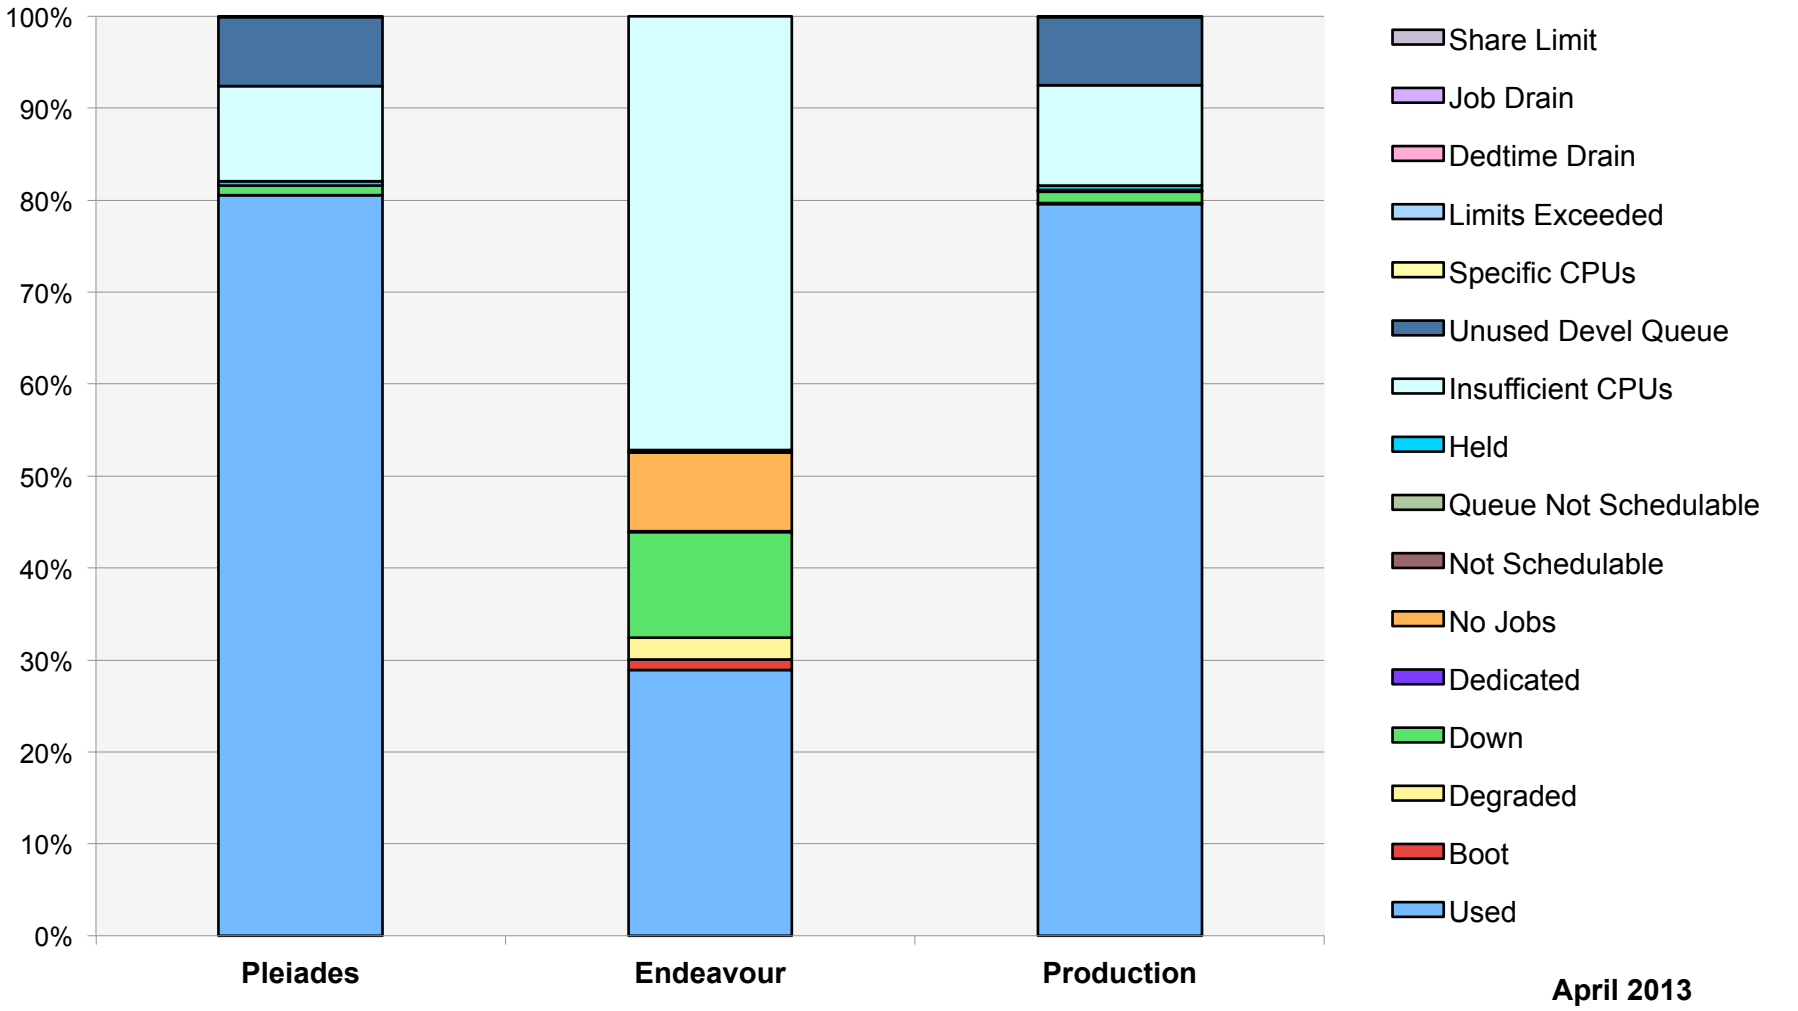

HECC Utilization
Normalized to 30-Day Month

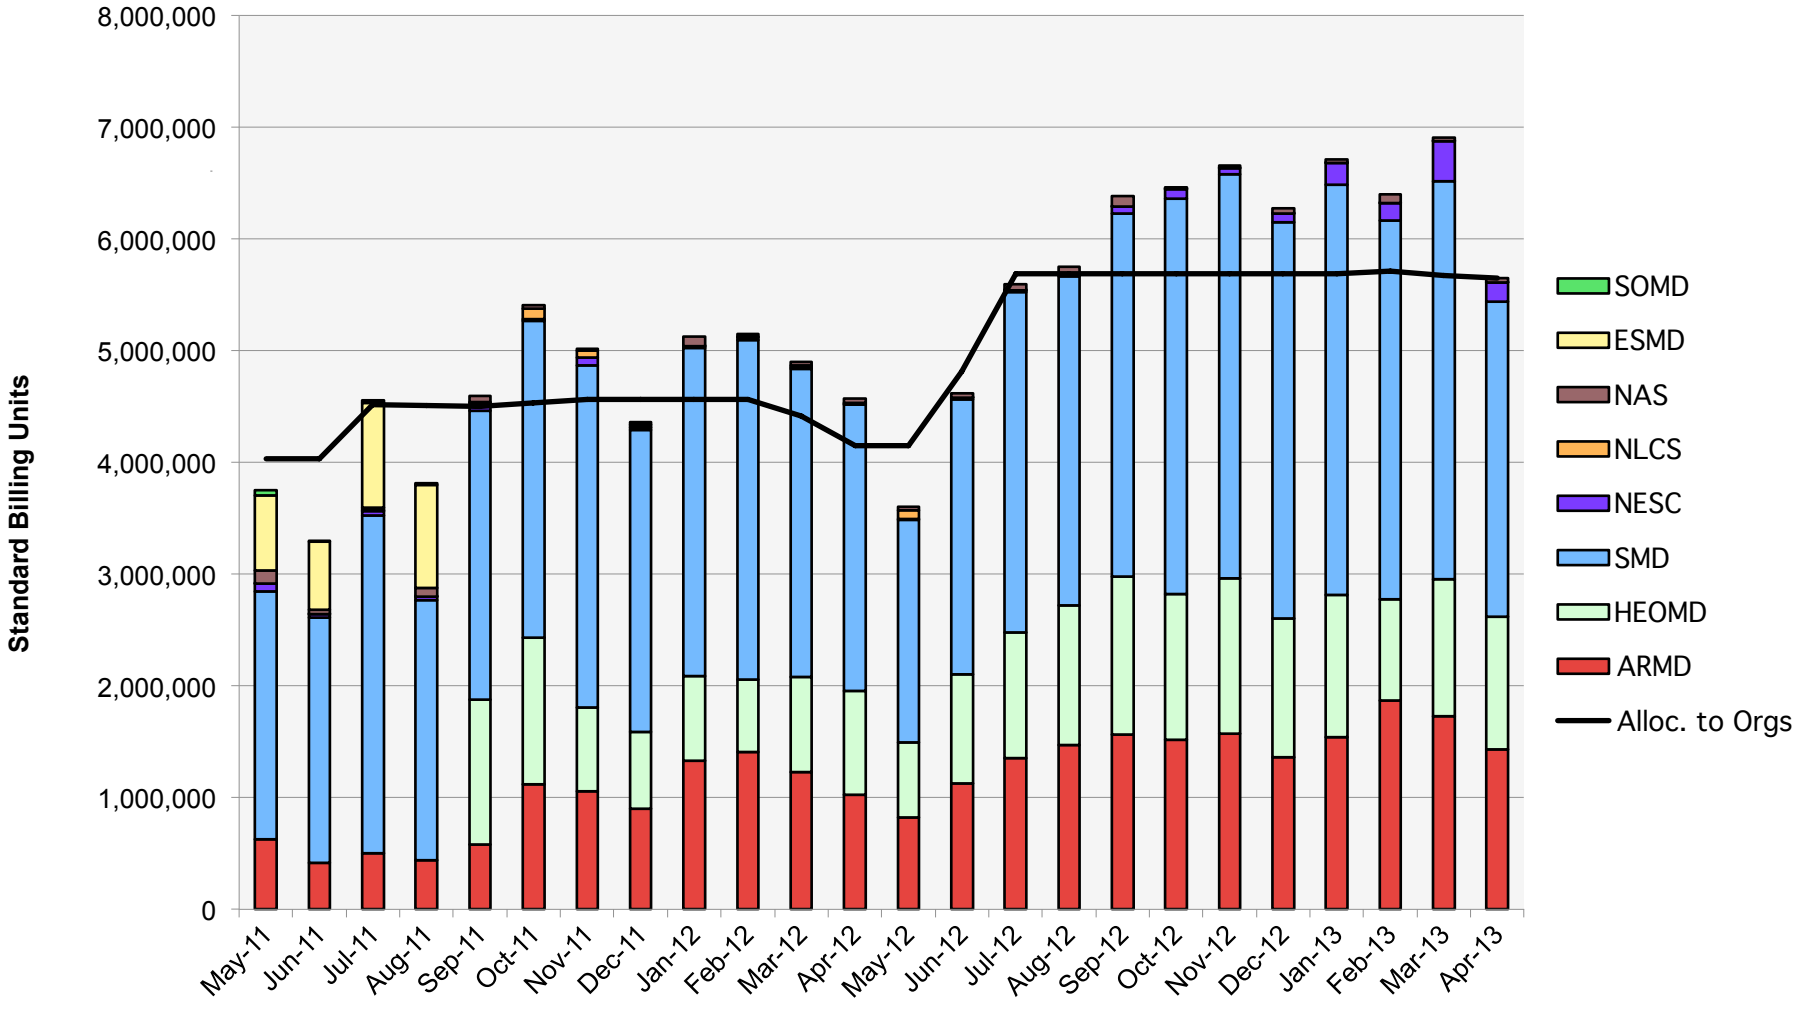

HECC Utilization
By Mission Directorate

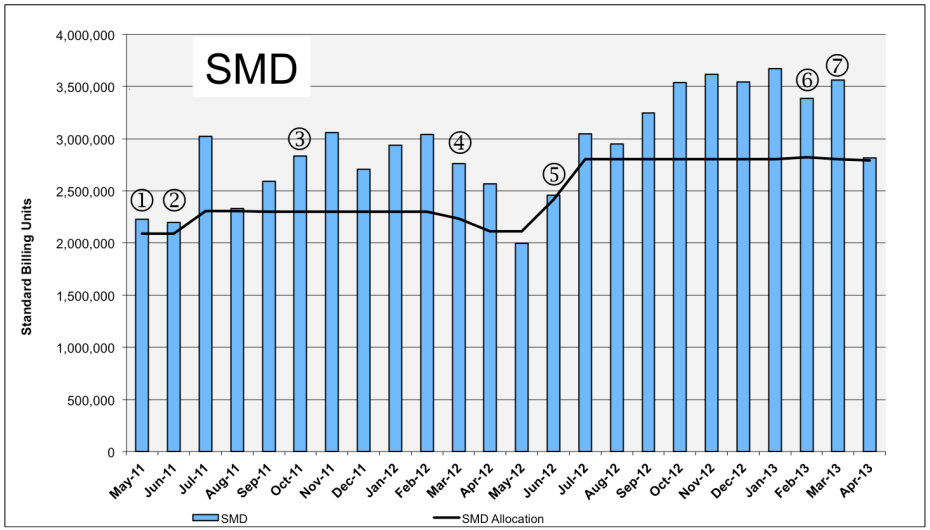

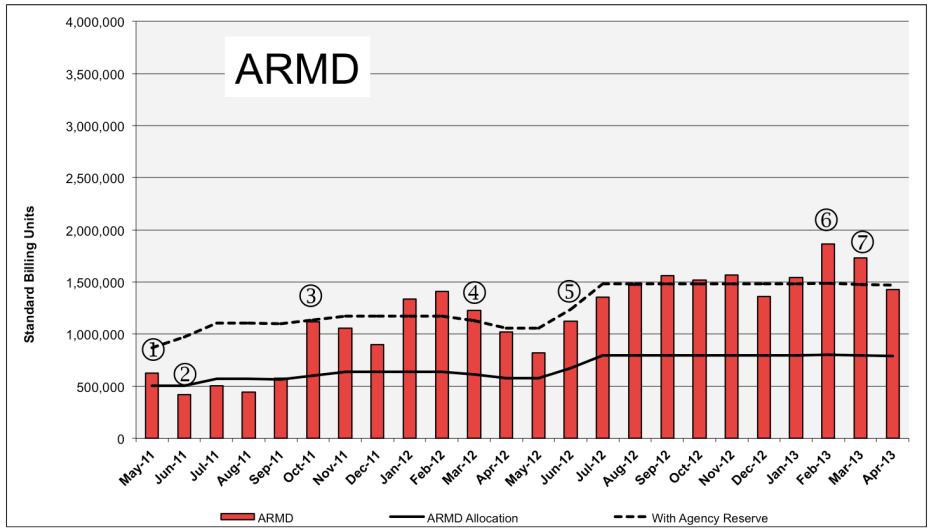

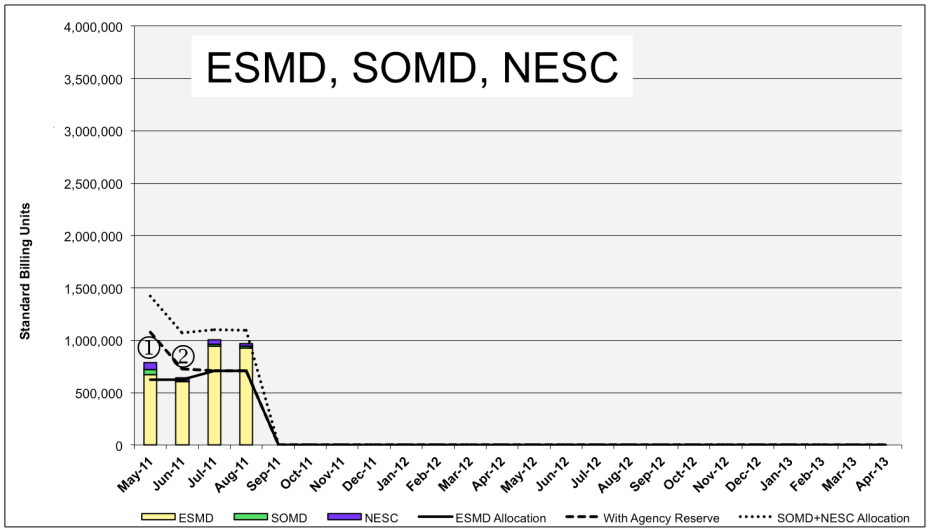

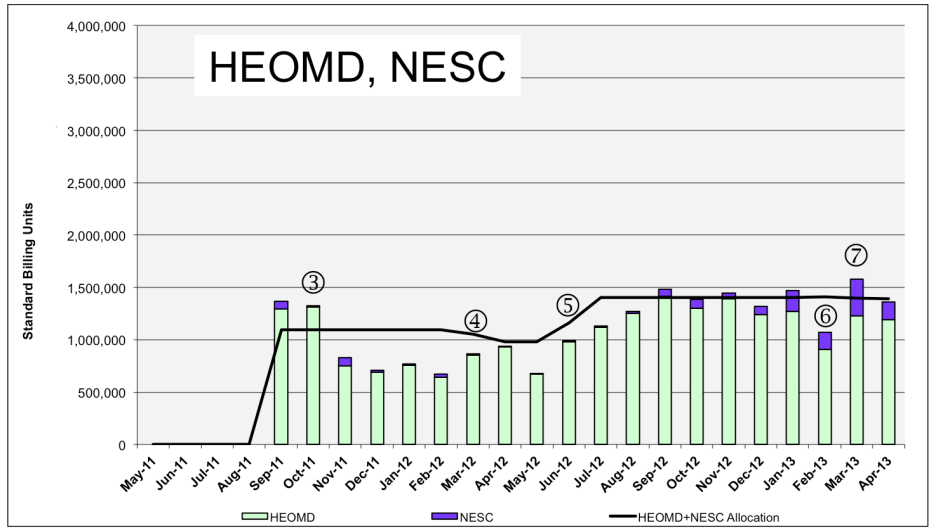

① Allocation to orgs. decreased to 75%, Agency reserve shifted to ARMD ② 14 Westmere racks added ③ 2 ARMD Westmere racks added
④ 28 Harpertown racks removed ⑤ 24 Sandy Bridge racks added ⑥ Columbia 21, 23,24 Endeavour 2 added ⑦ Columbia 22 removed, Endeavour 1 added

Tape Archive Status

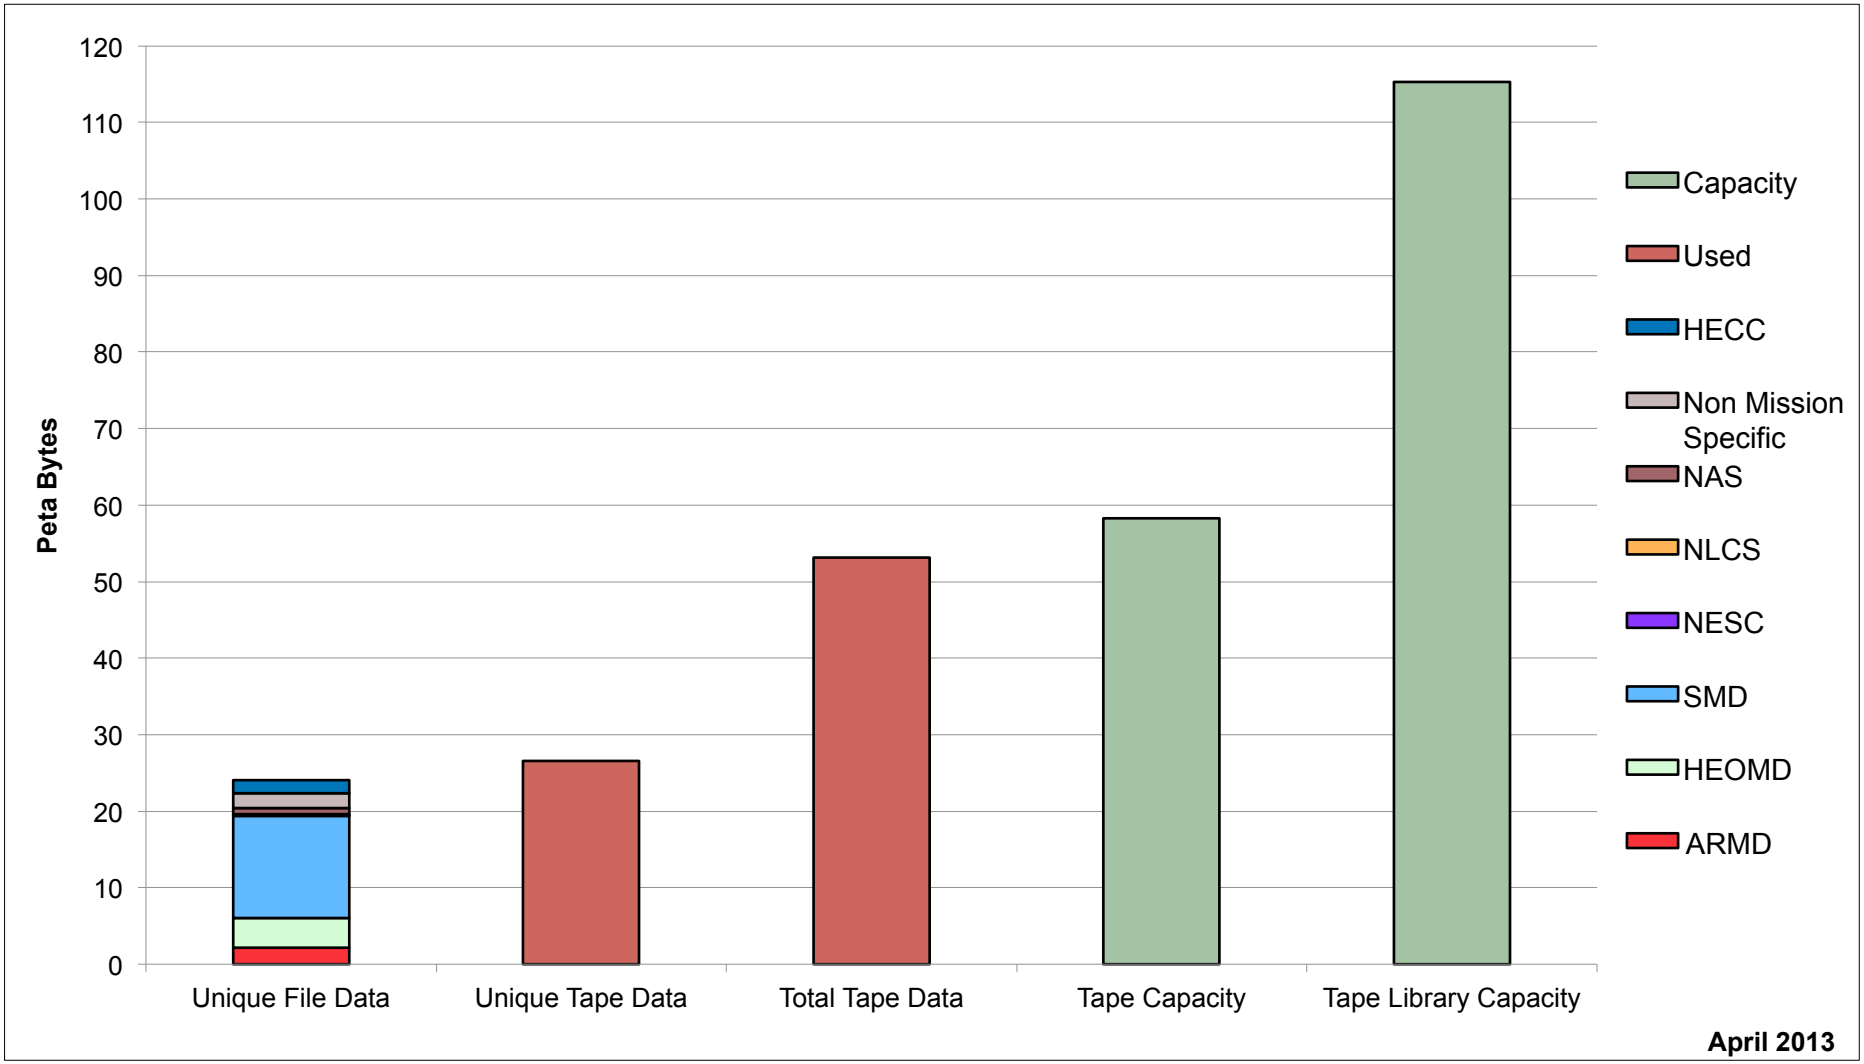

Tape Archive Status

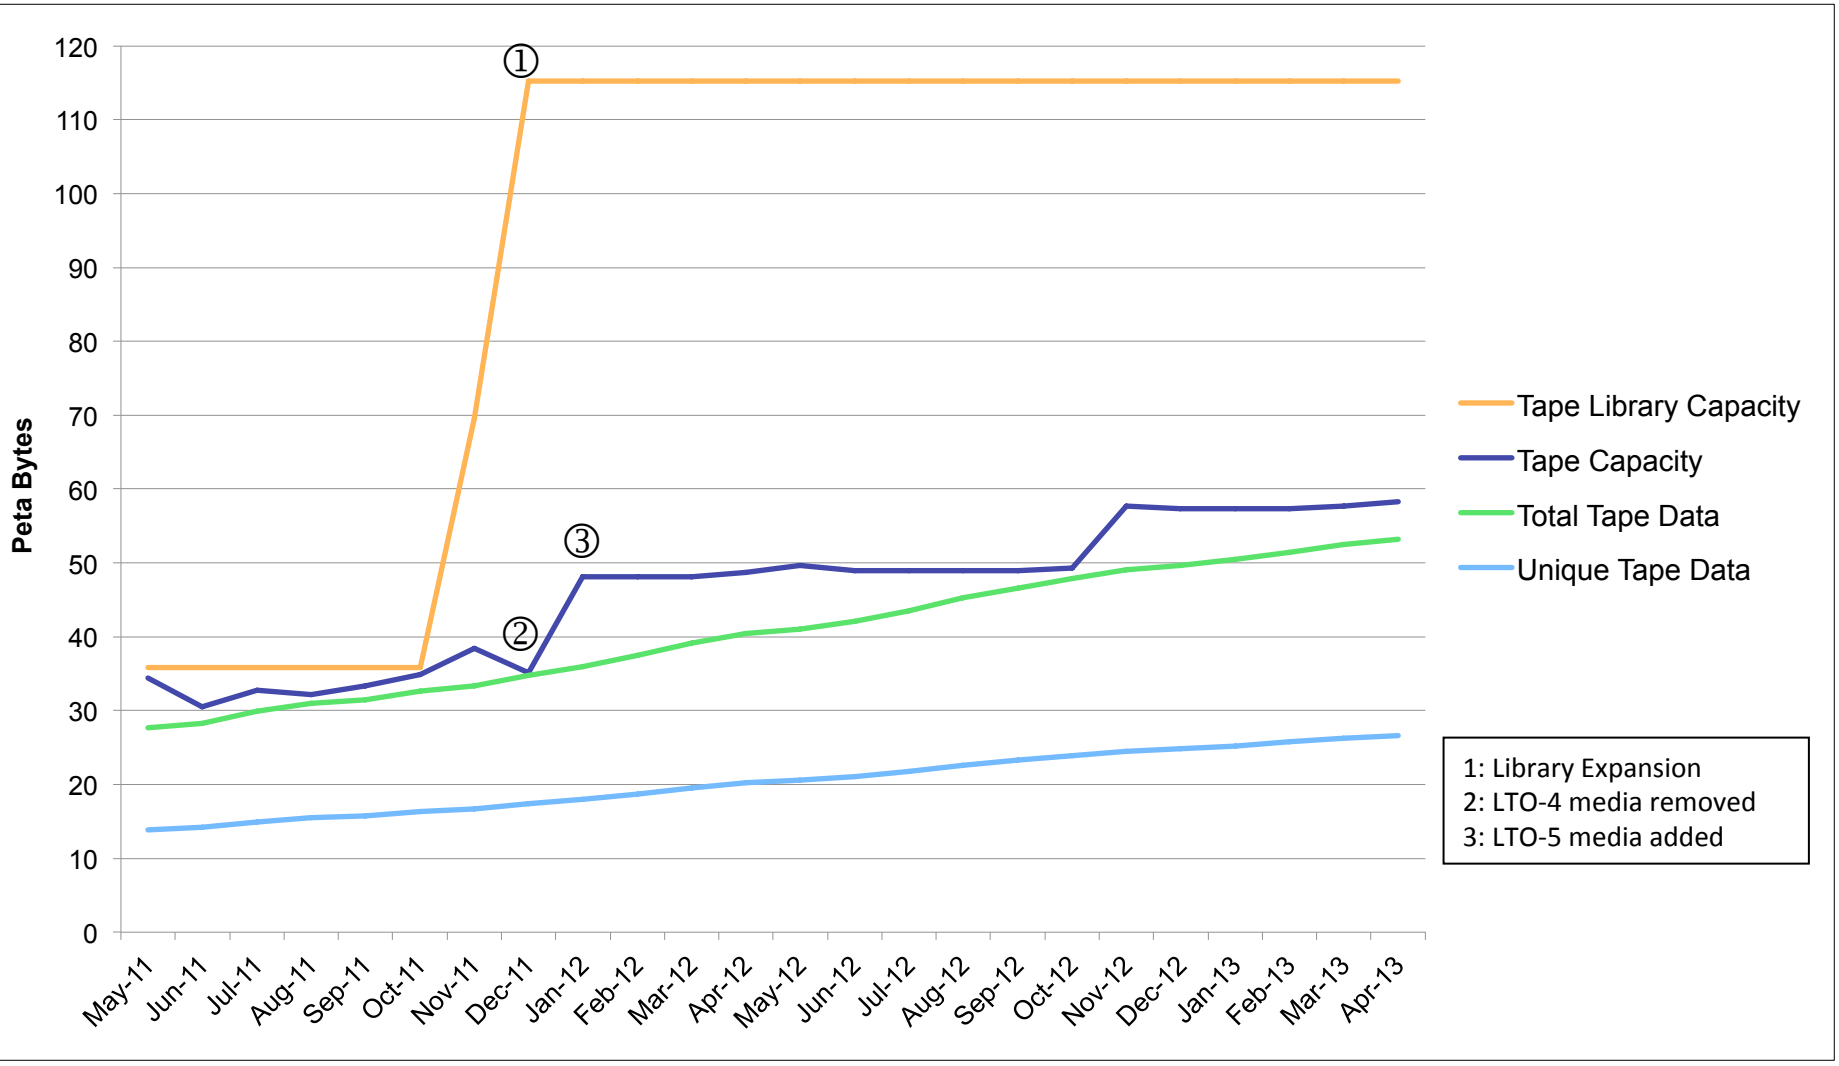

Pleiades:
SBUs Reported, Normalized to 30-Day Month

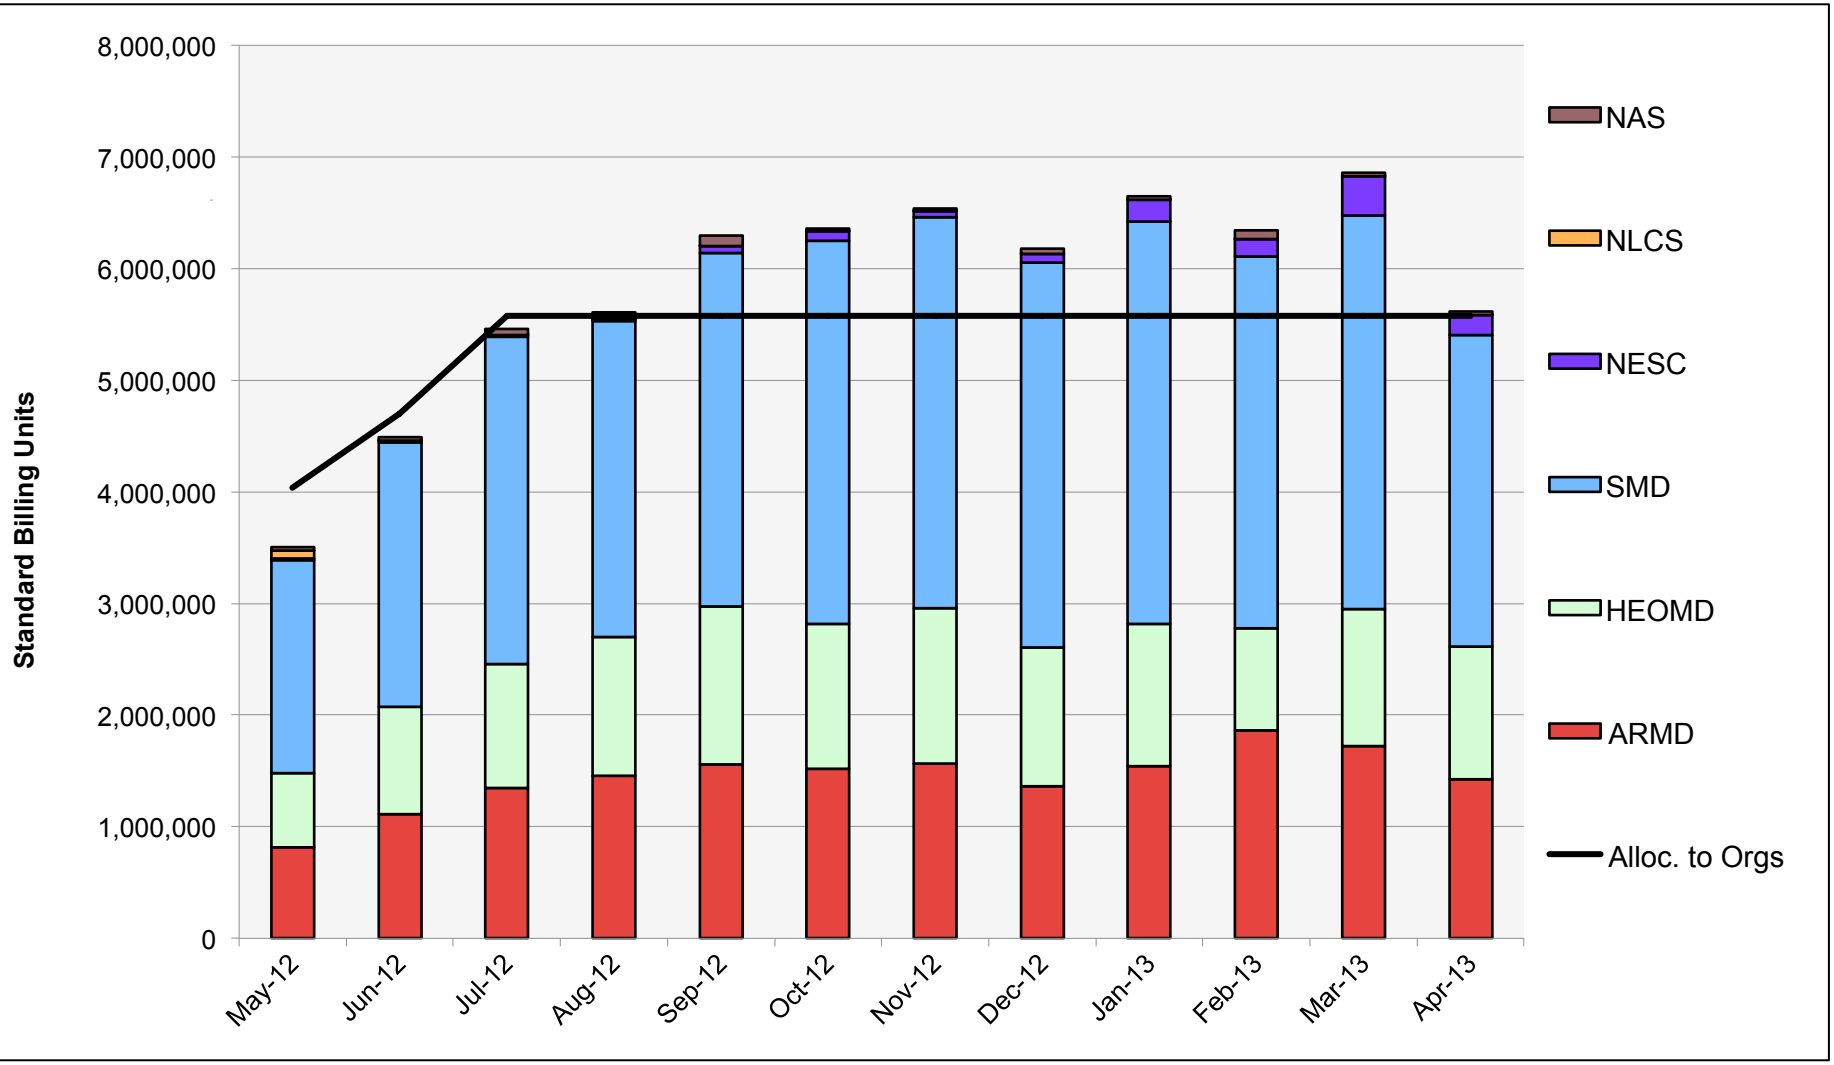

Pleiades:
Devel Queue Utilization

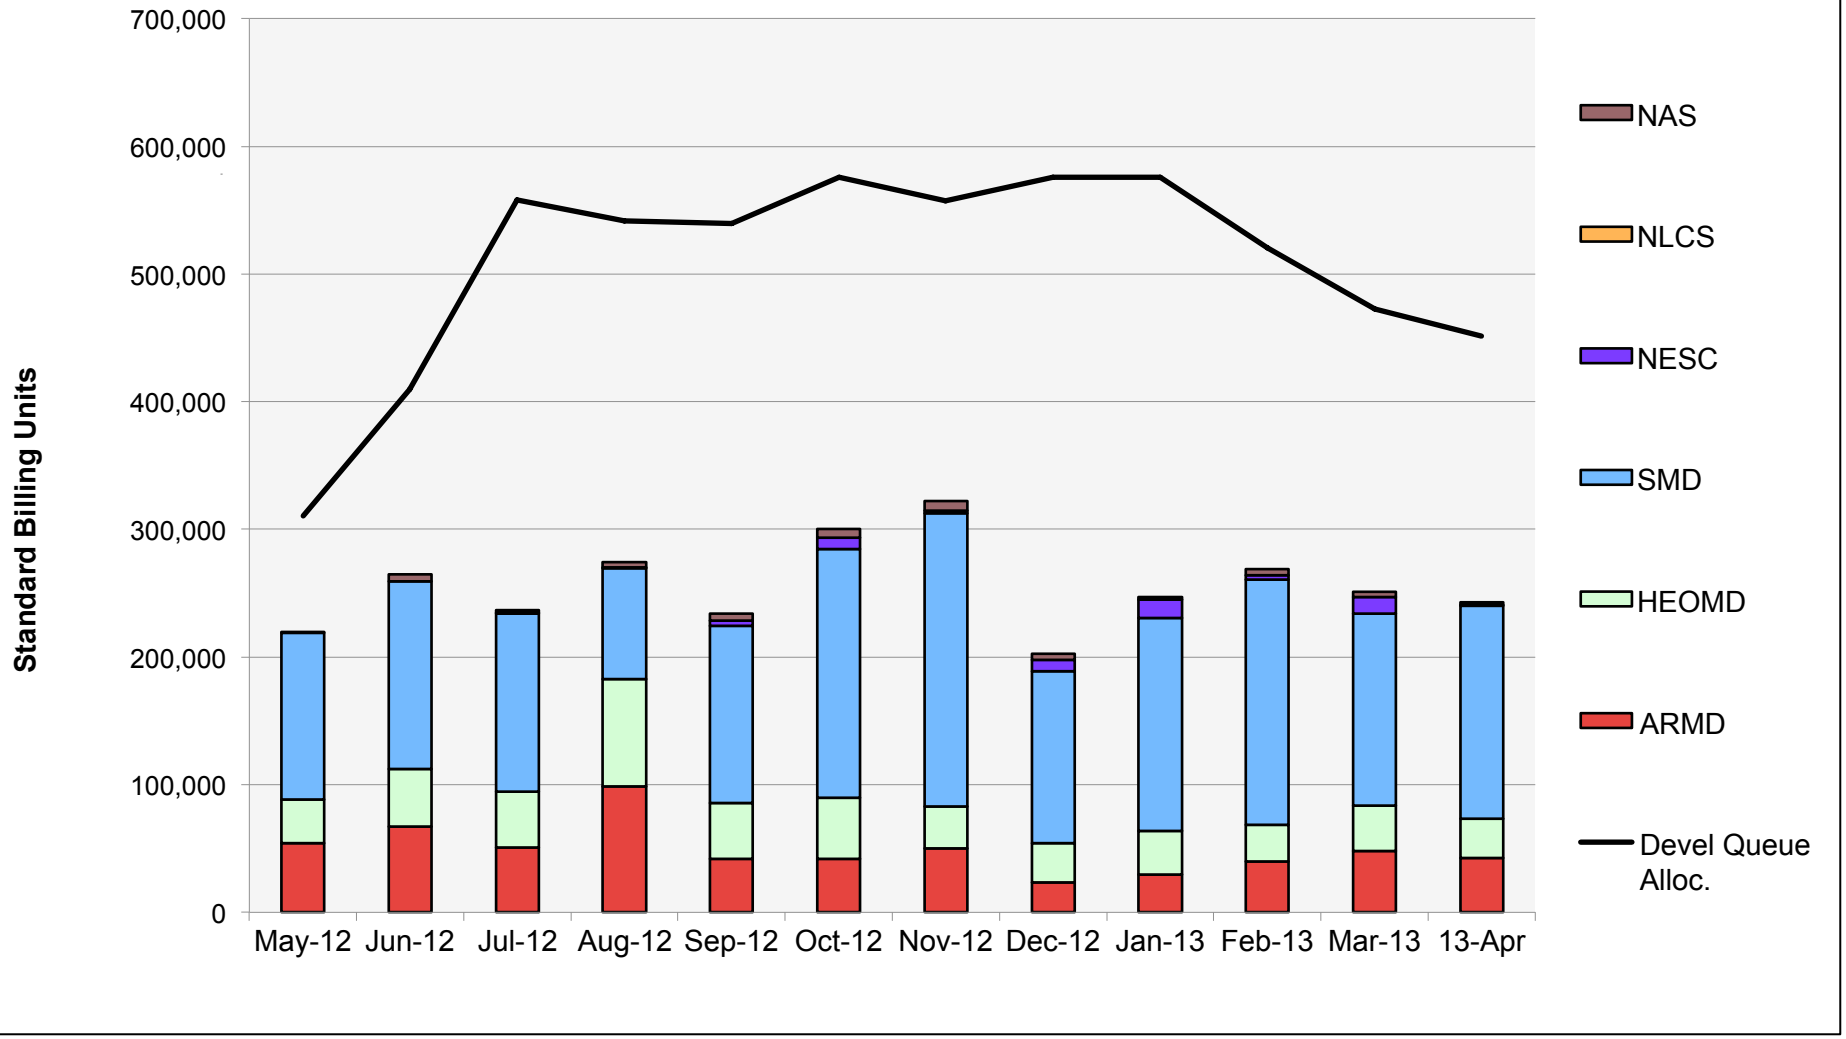

Pleiades:
Monthly Utilization by Job Length

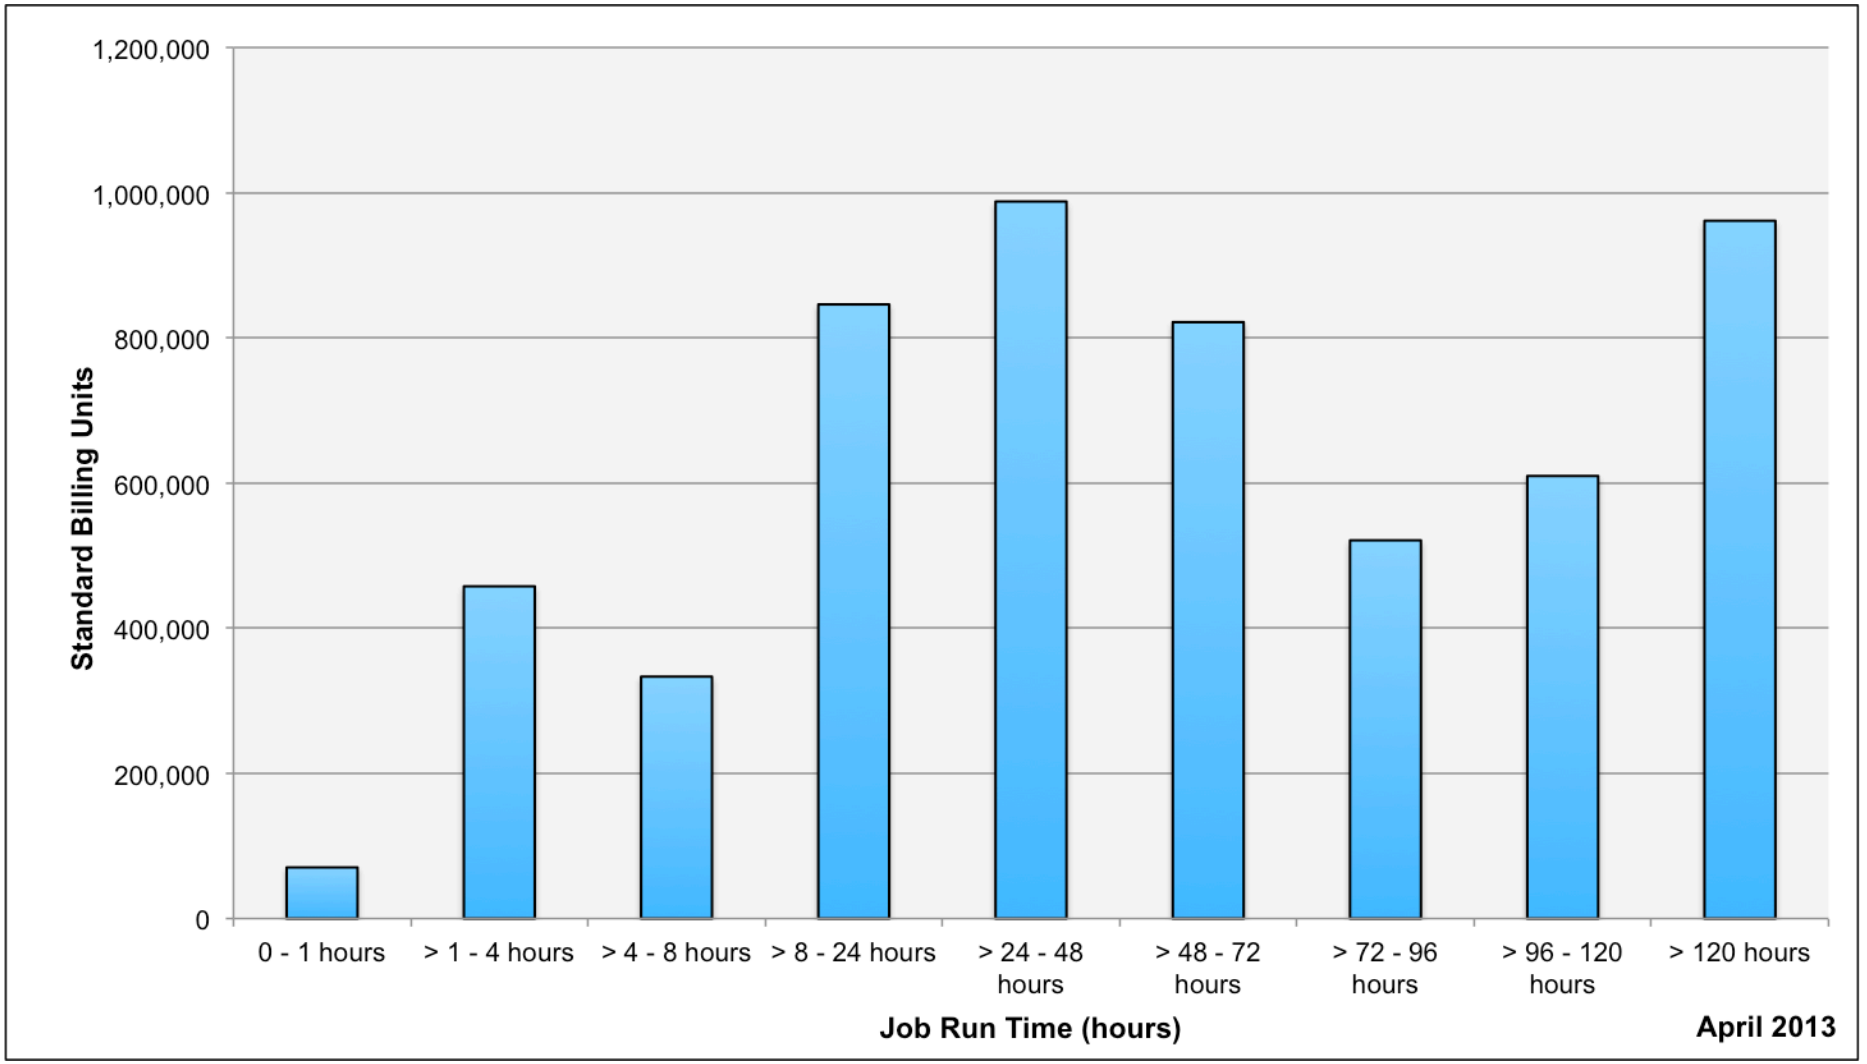

Pleiades:
Monthly Utilization by Size and Mission

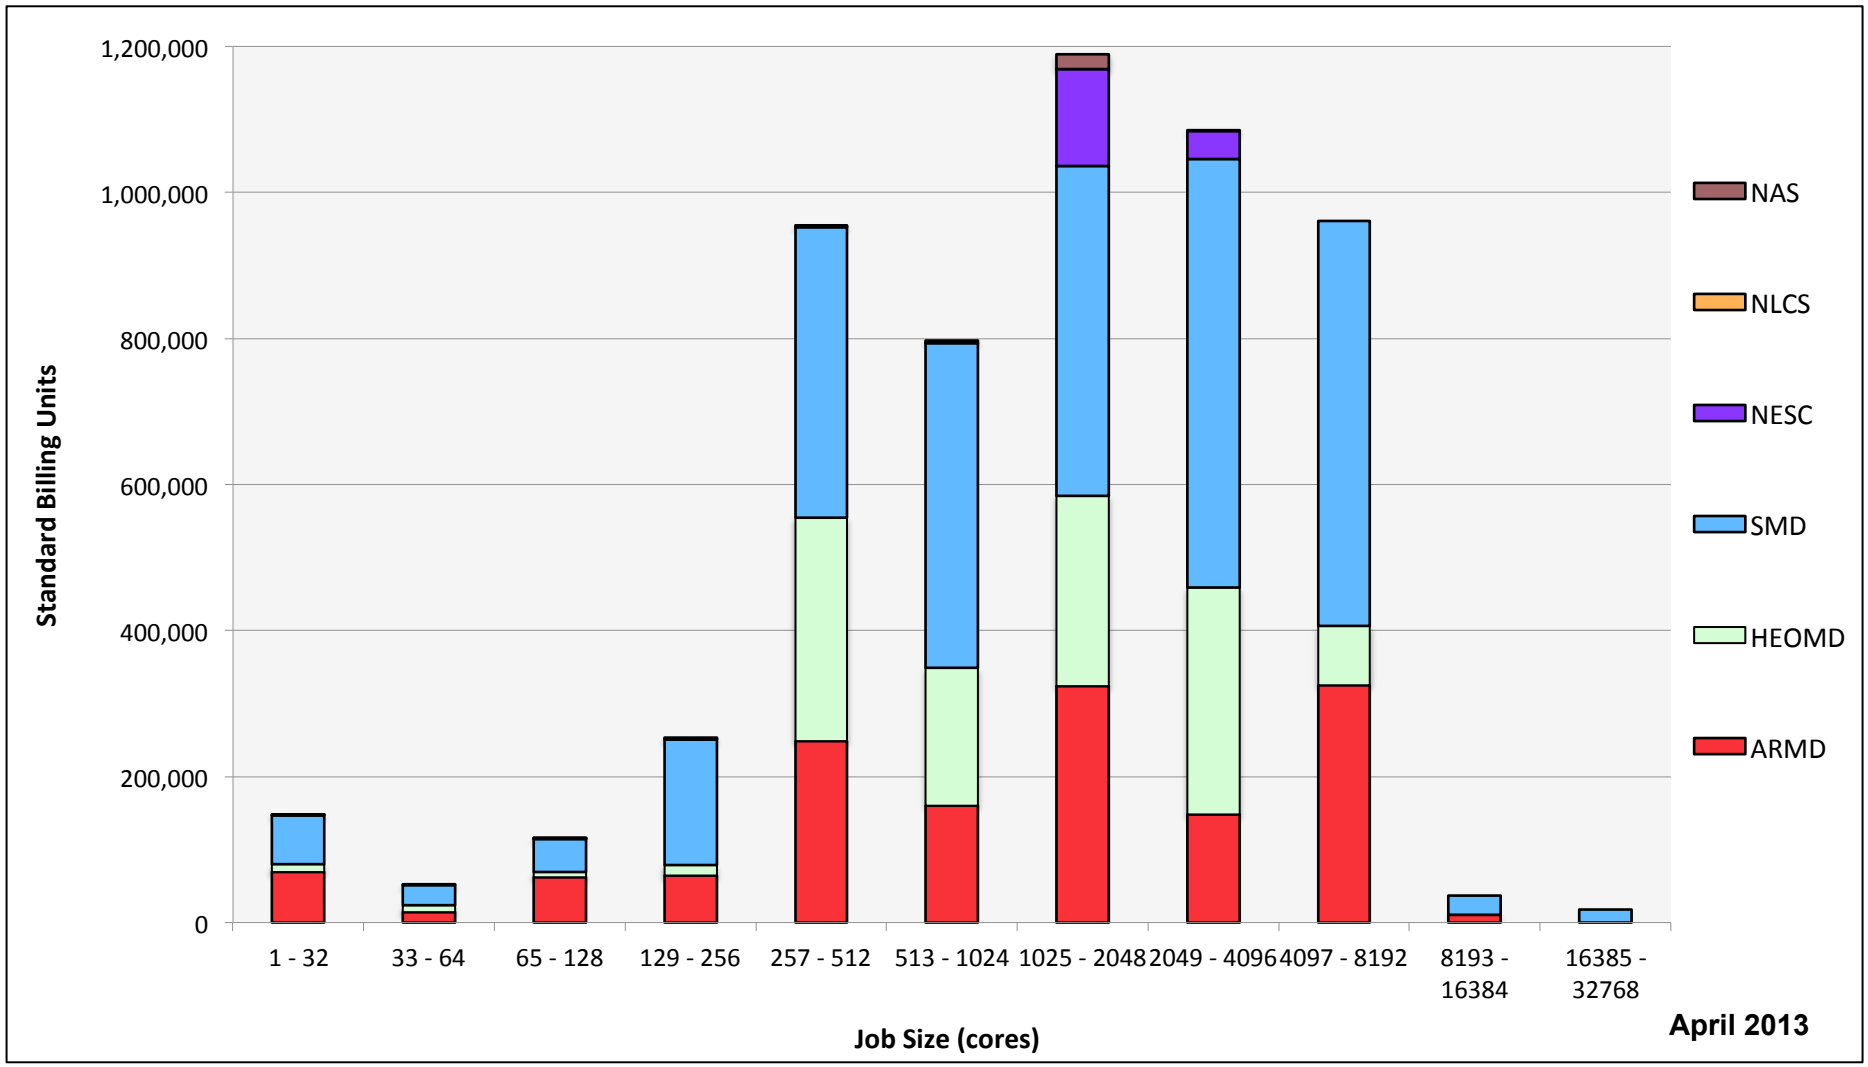

Pleiades:
Monthly Utilization by Size and Length

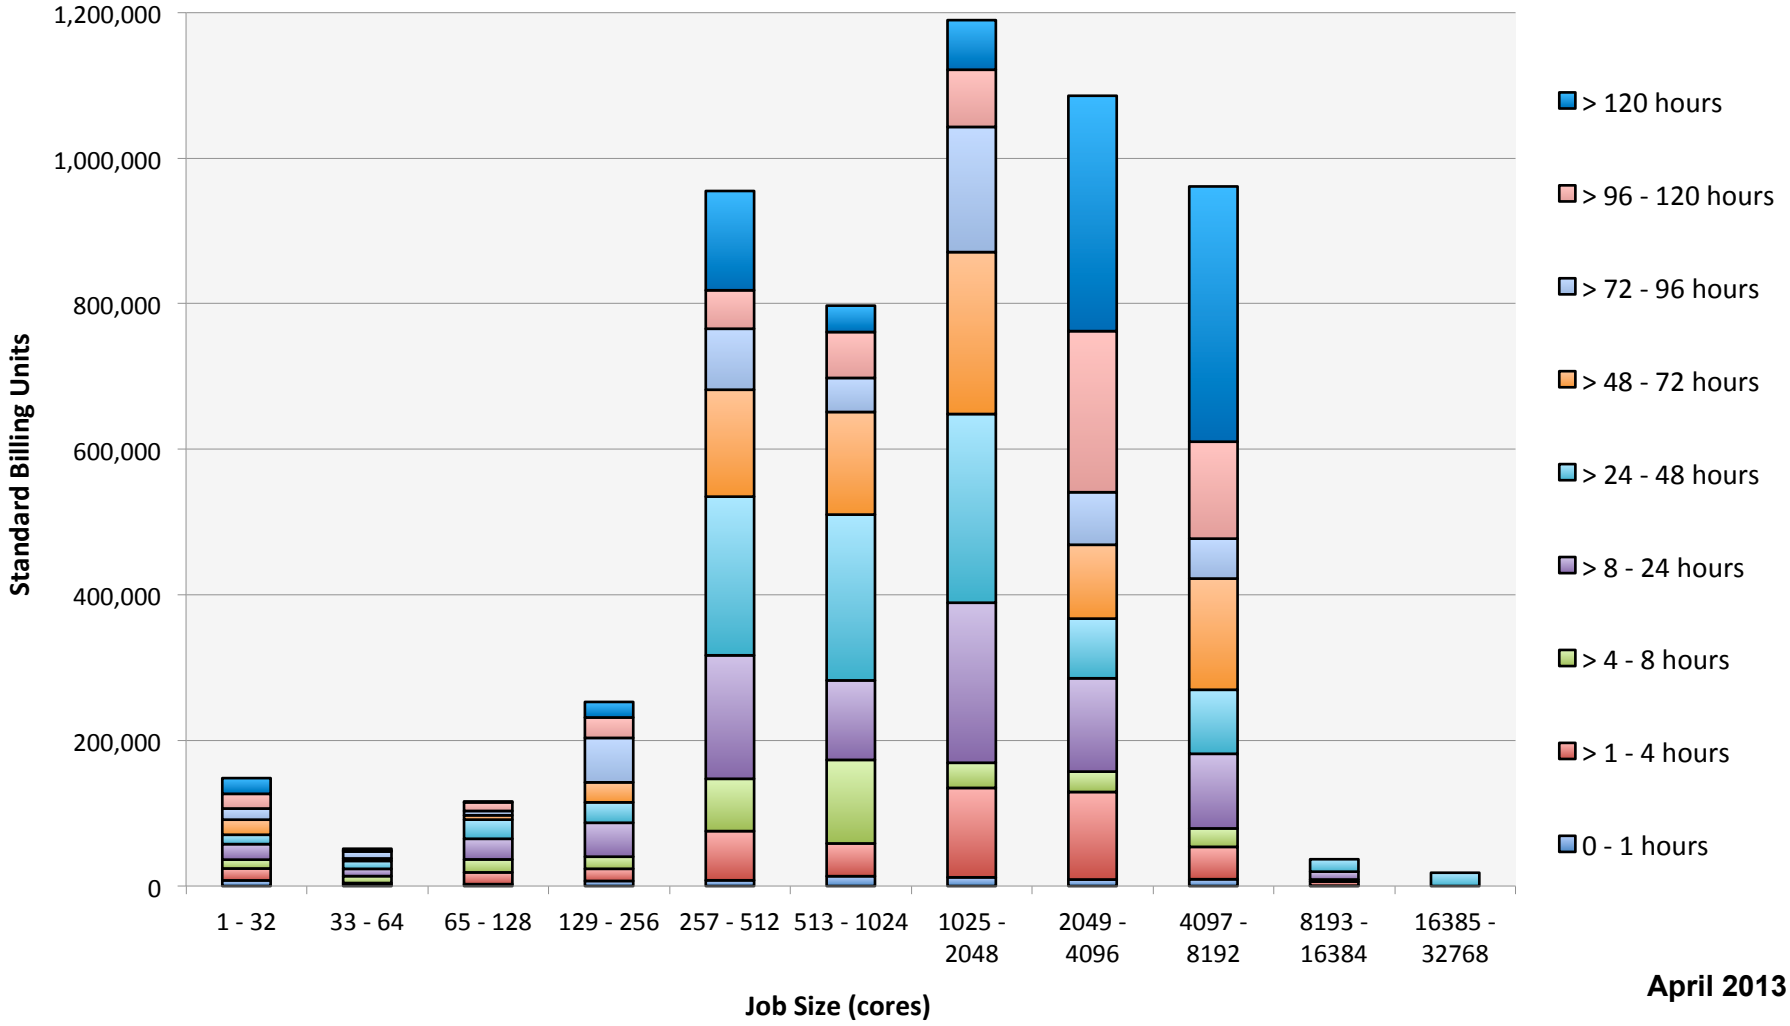

Pleiades:
Average Time to Clear All Jobs

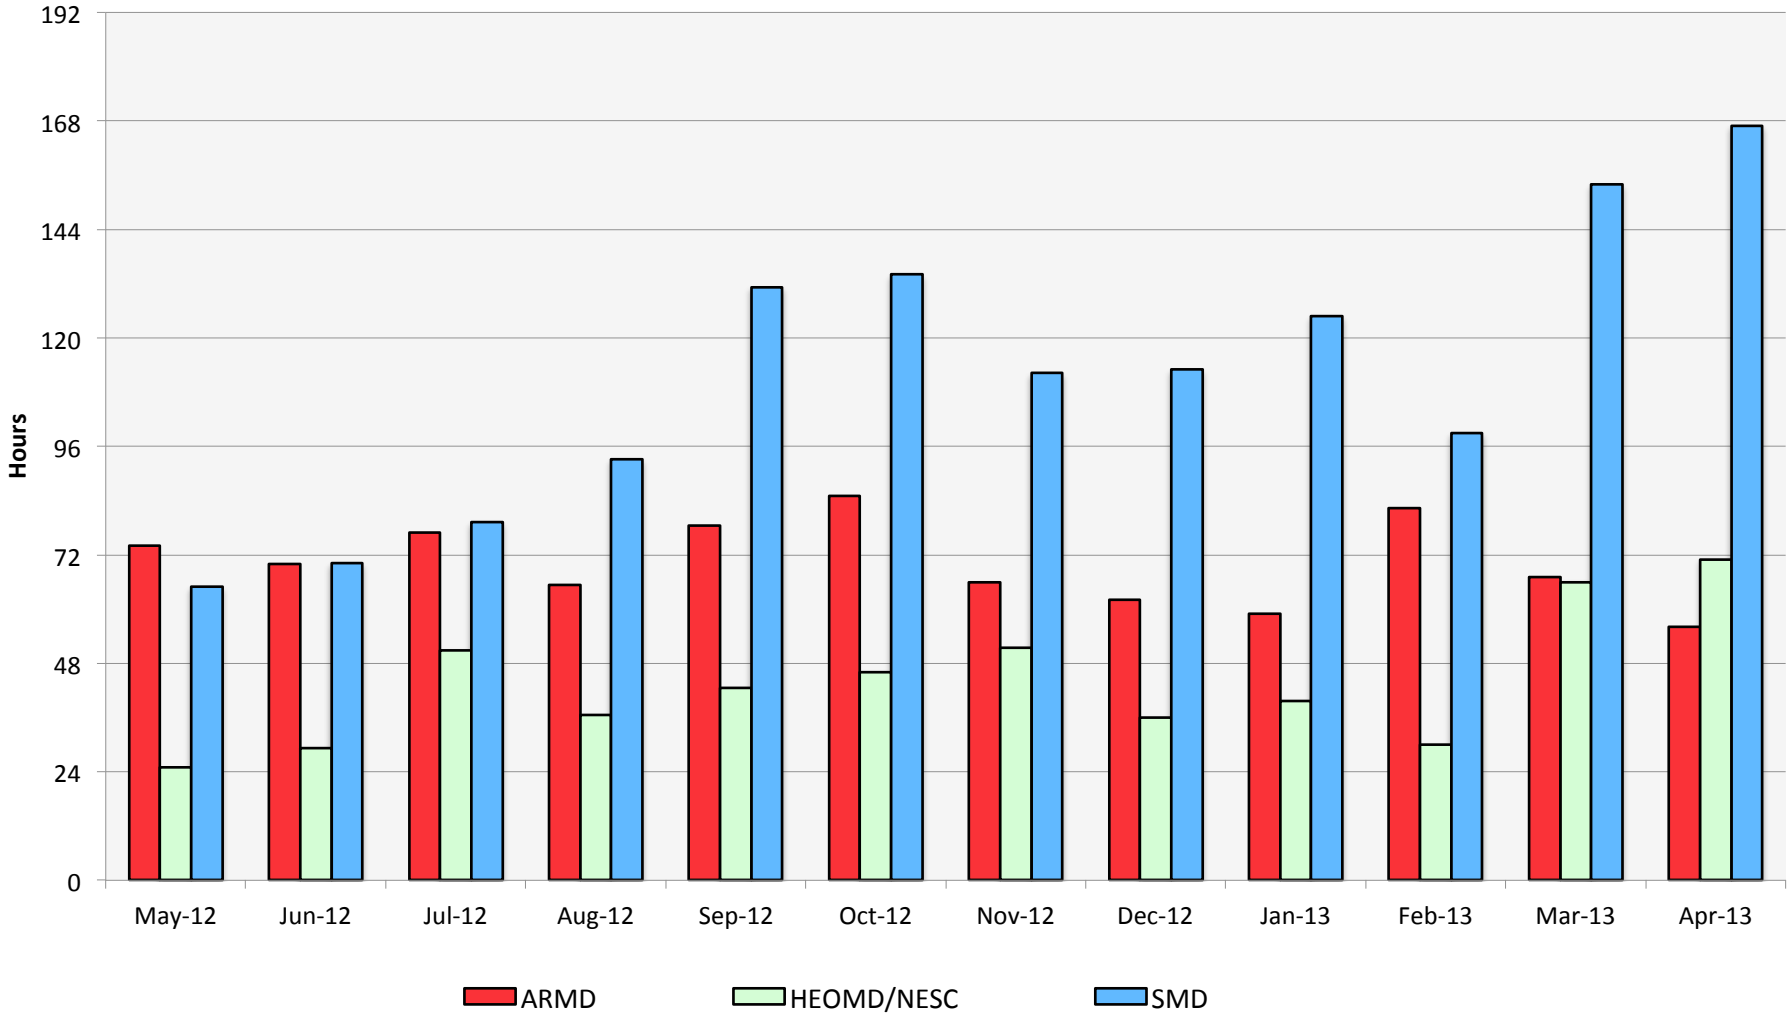

Pleiades:
Average Expansion Factor (with Anomalies)

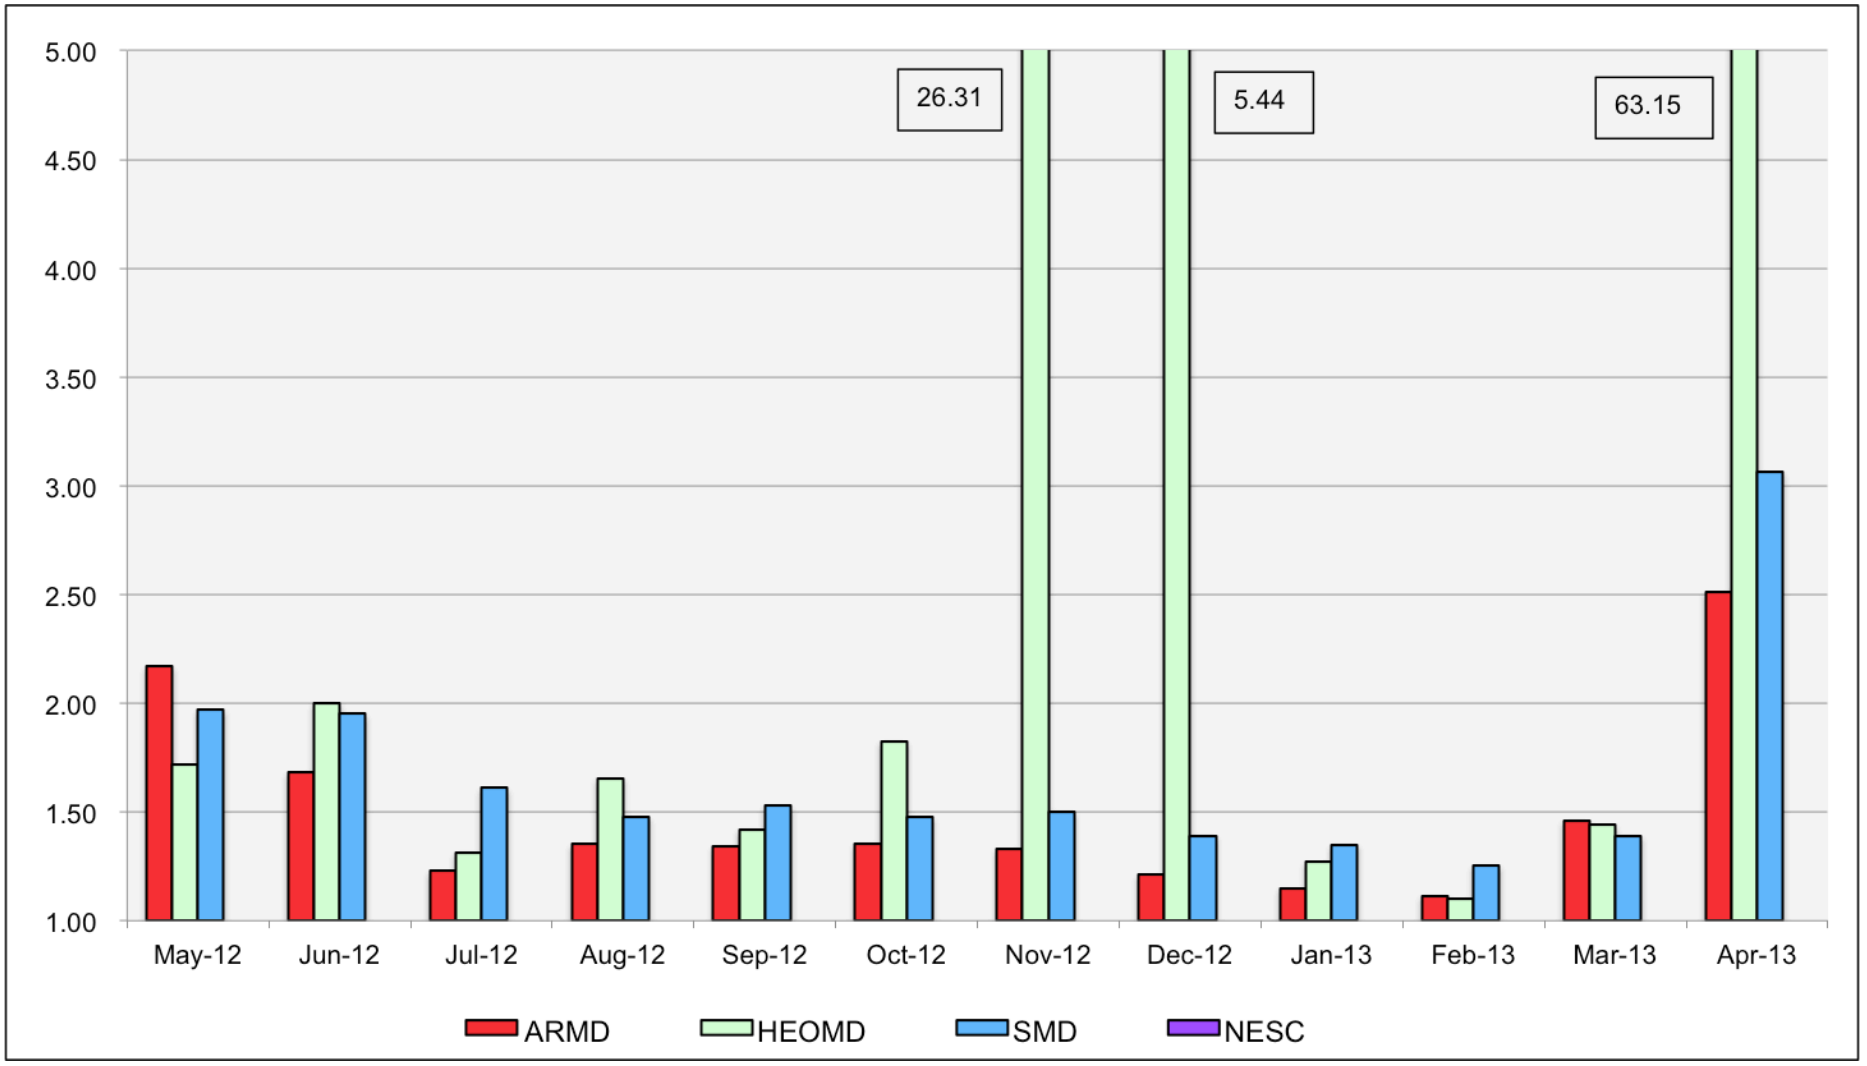

Pleiades:
Average Expansion Factor (Anomalies Removed)

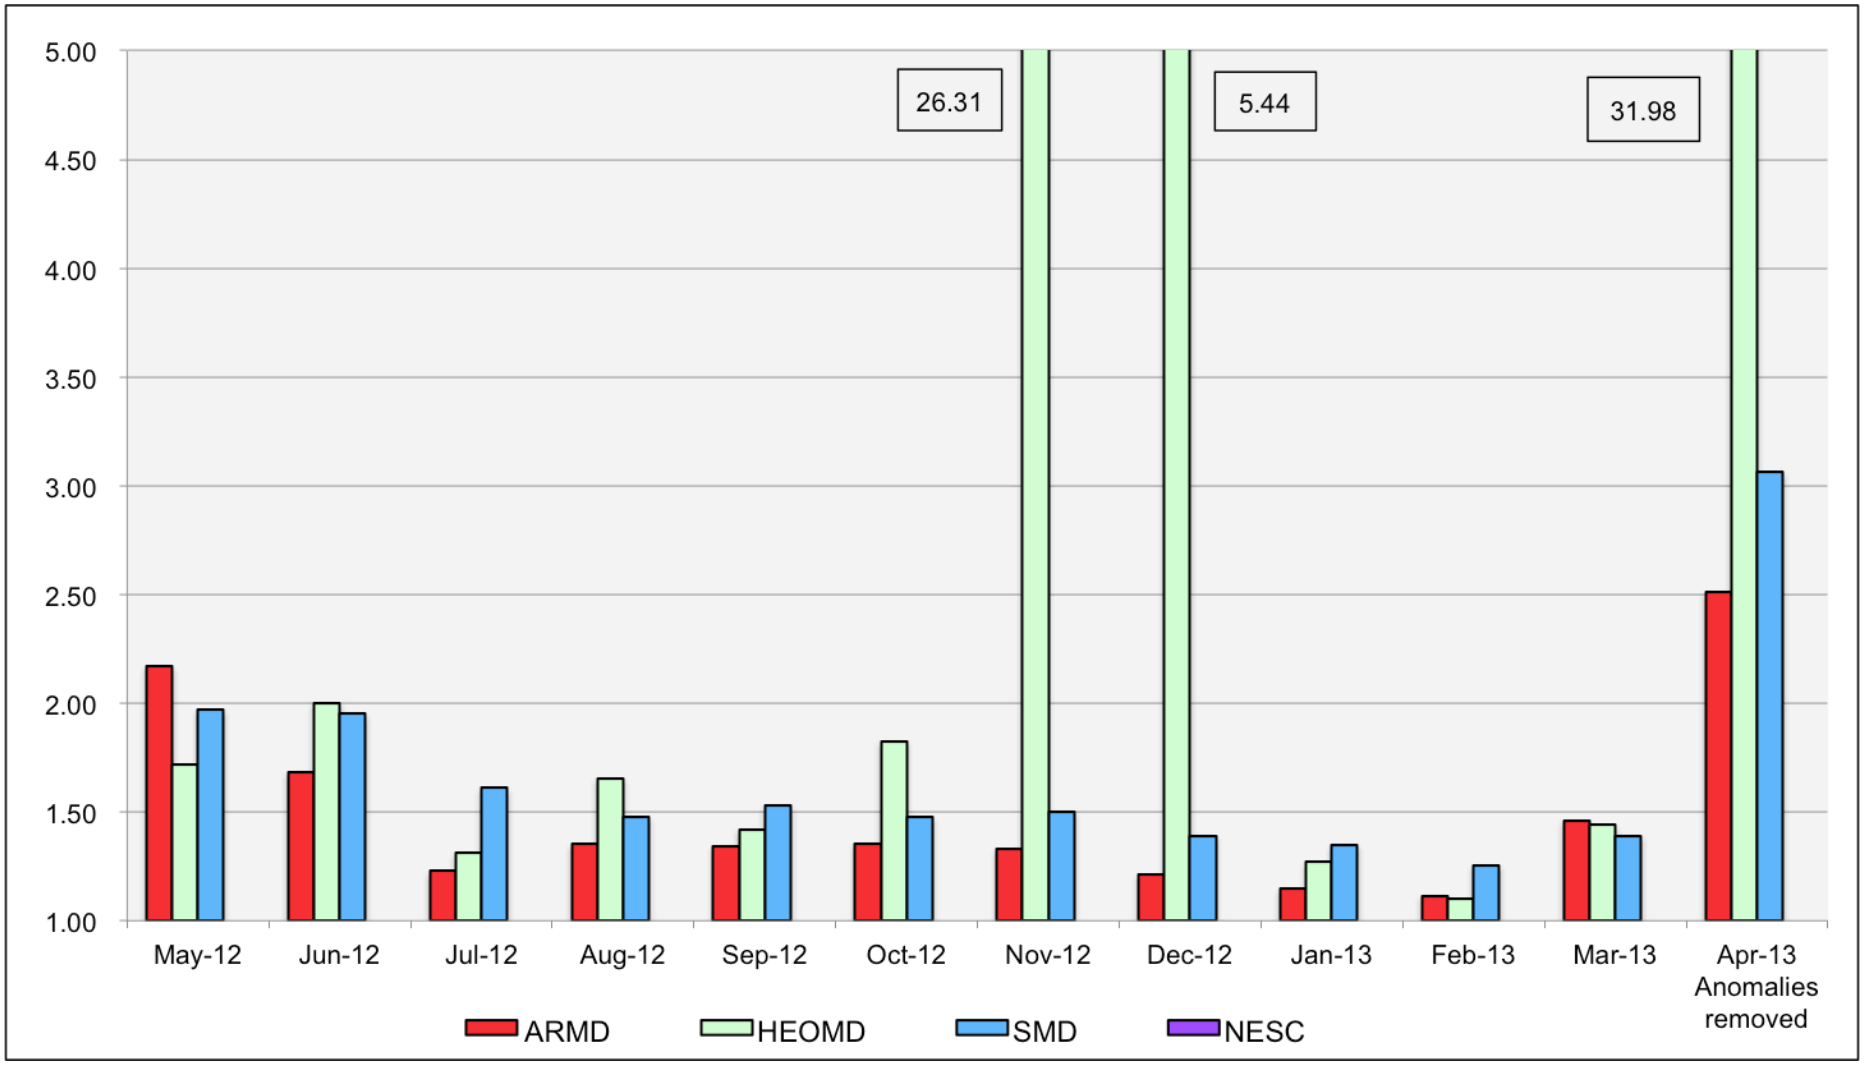

Endeavour:
SBUs Reported, Normalized to 30-Day Month

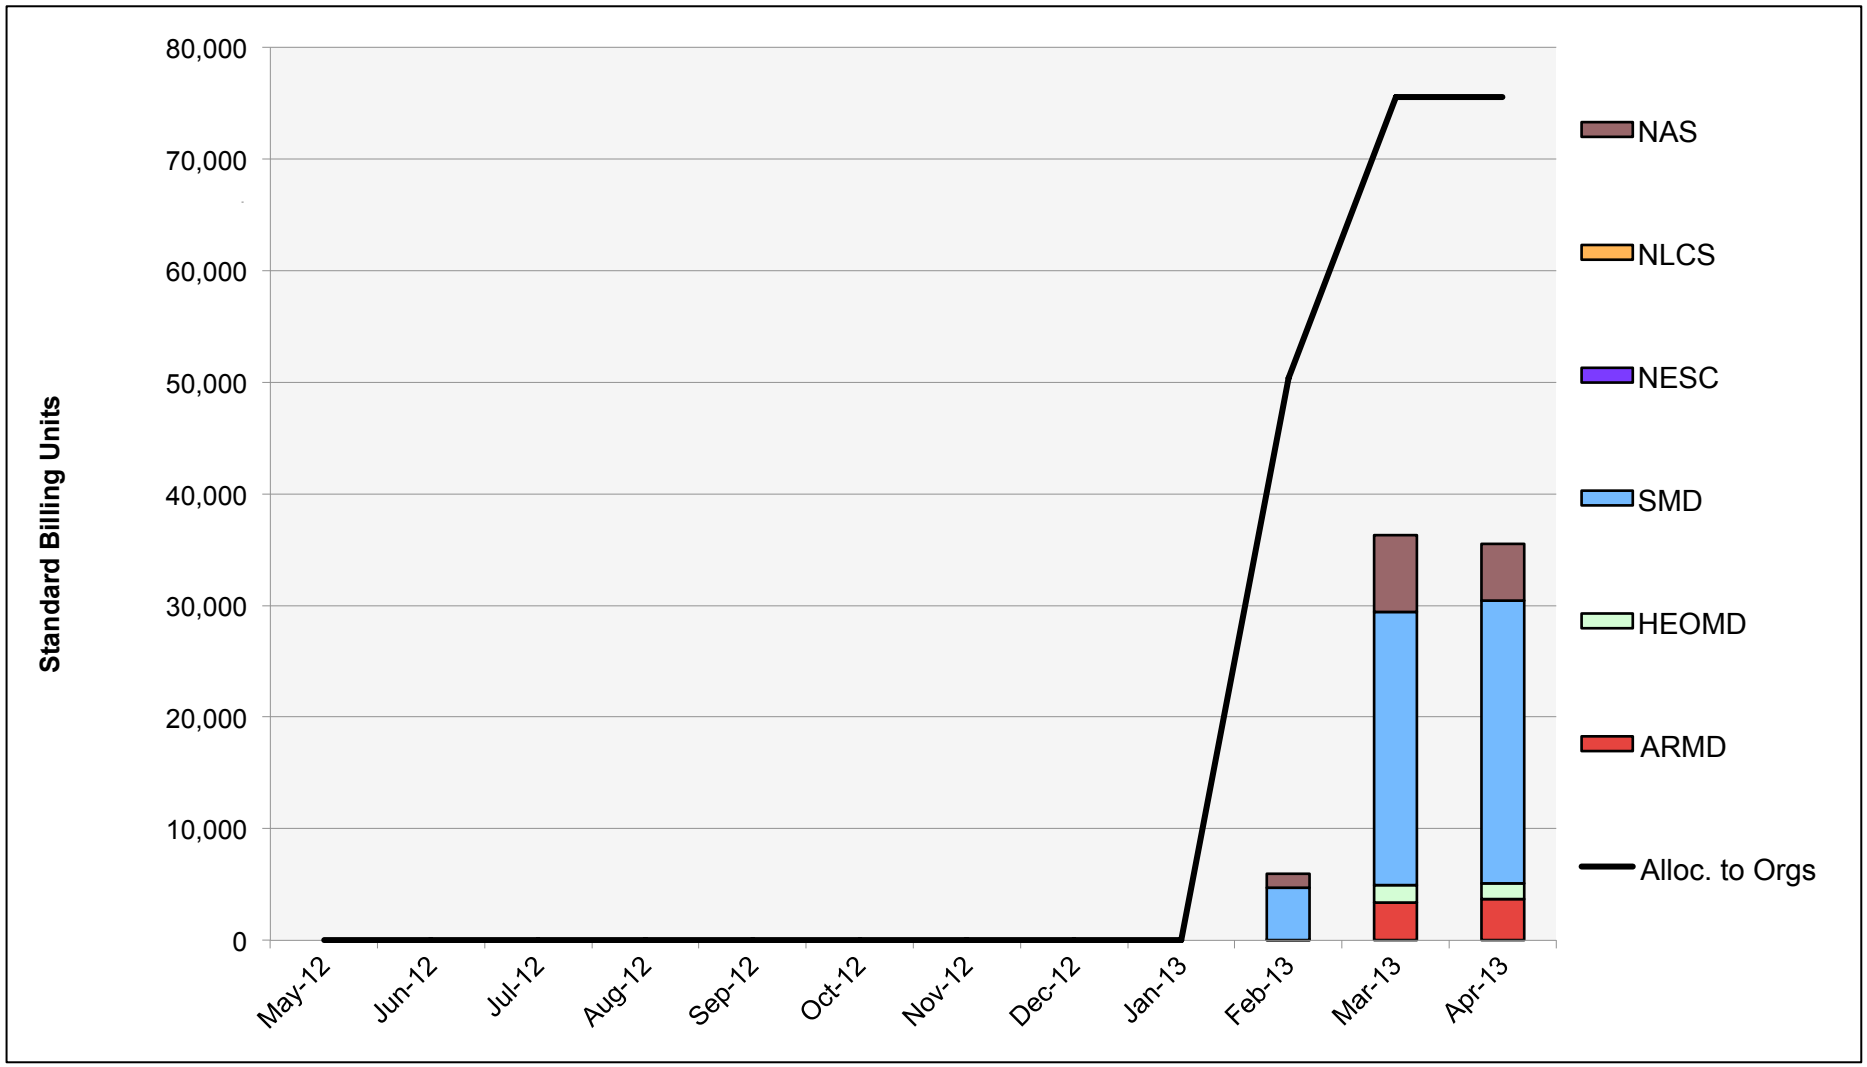

Endeavour:
Monthly Utilization by Job Length

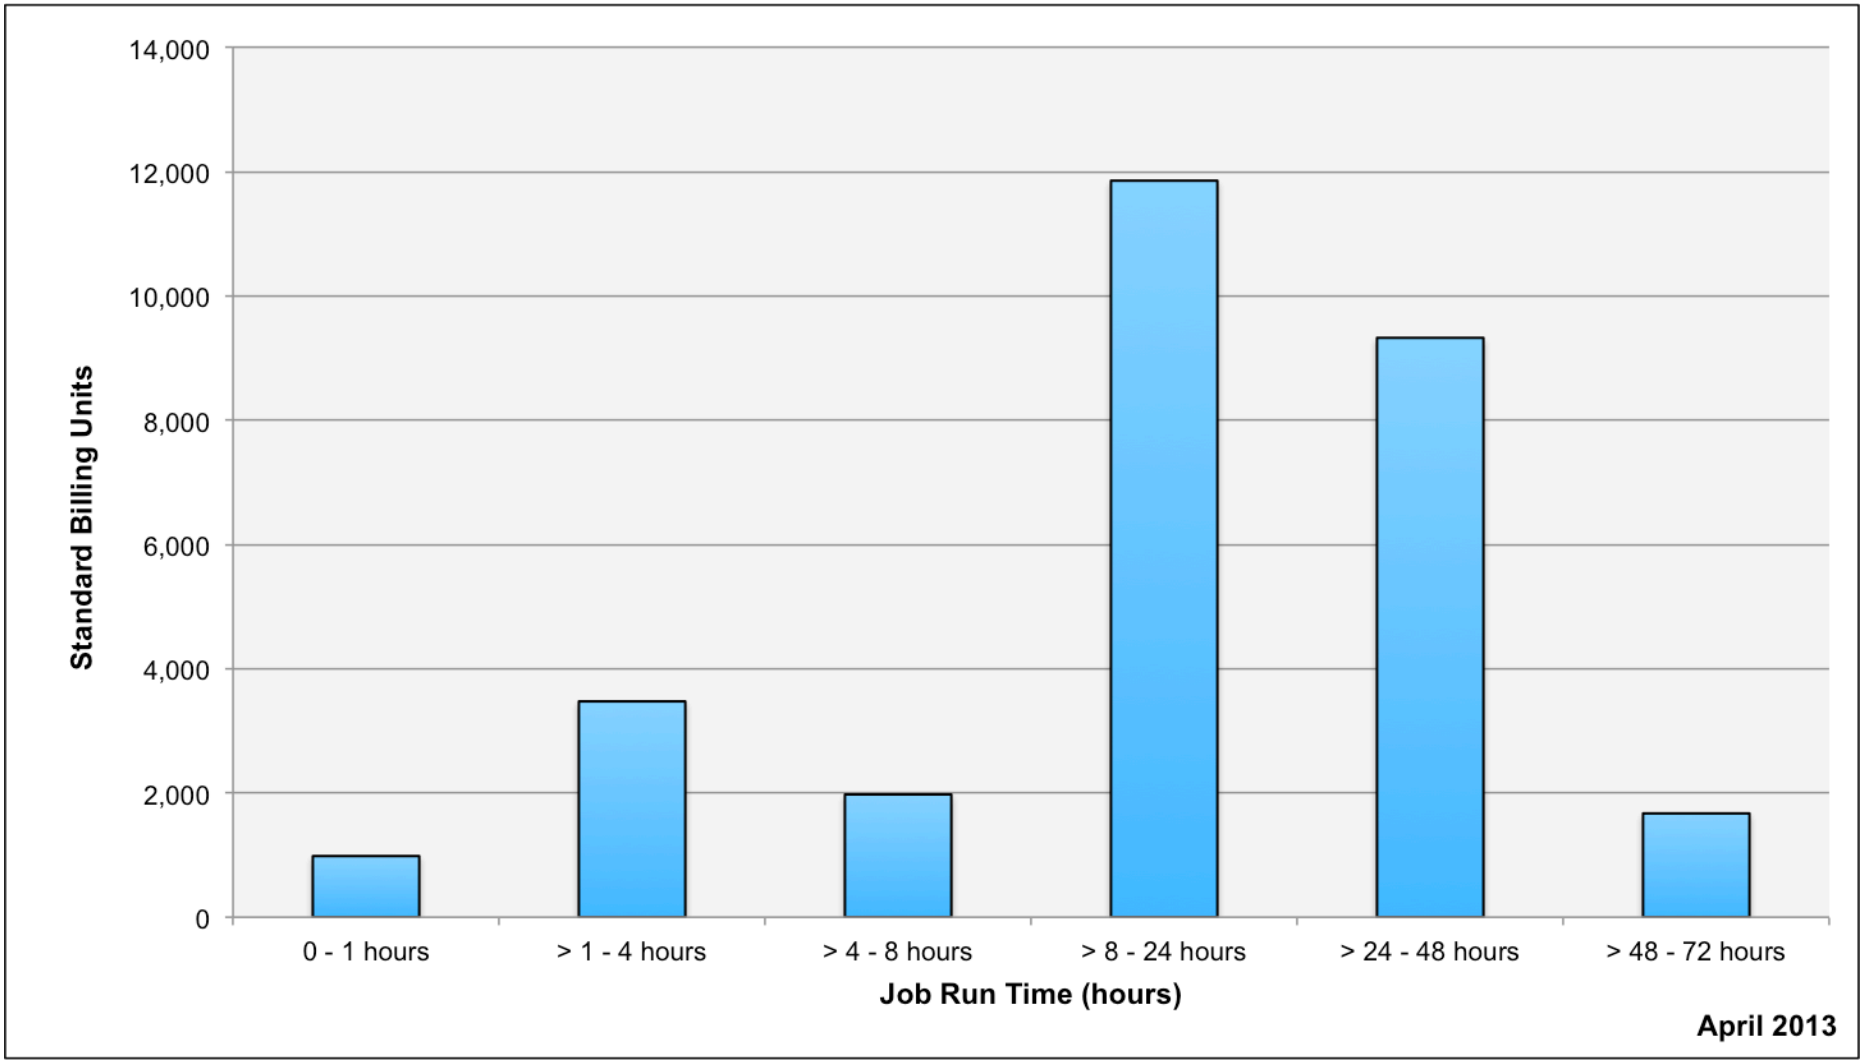

Endeavour:
Monthly Utilization by Size and Length

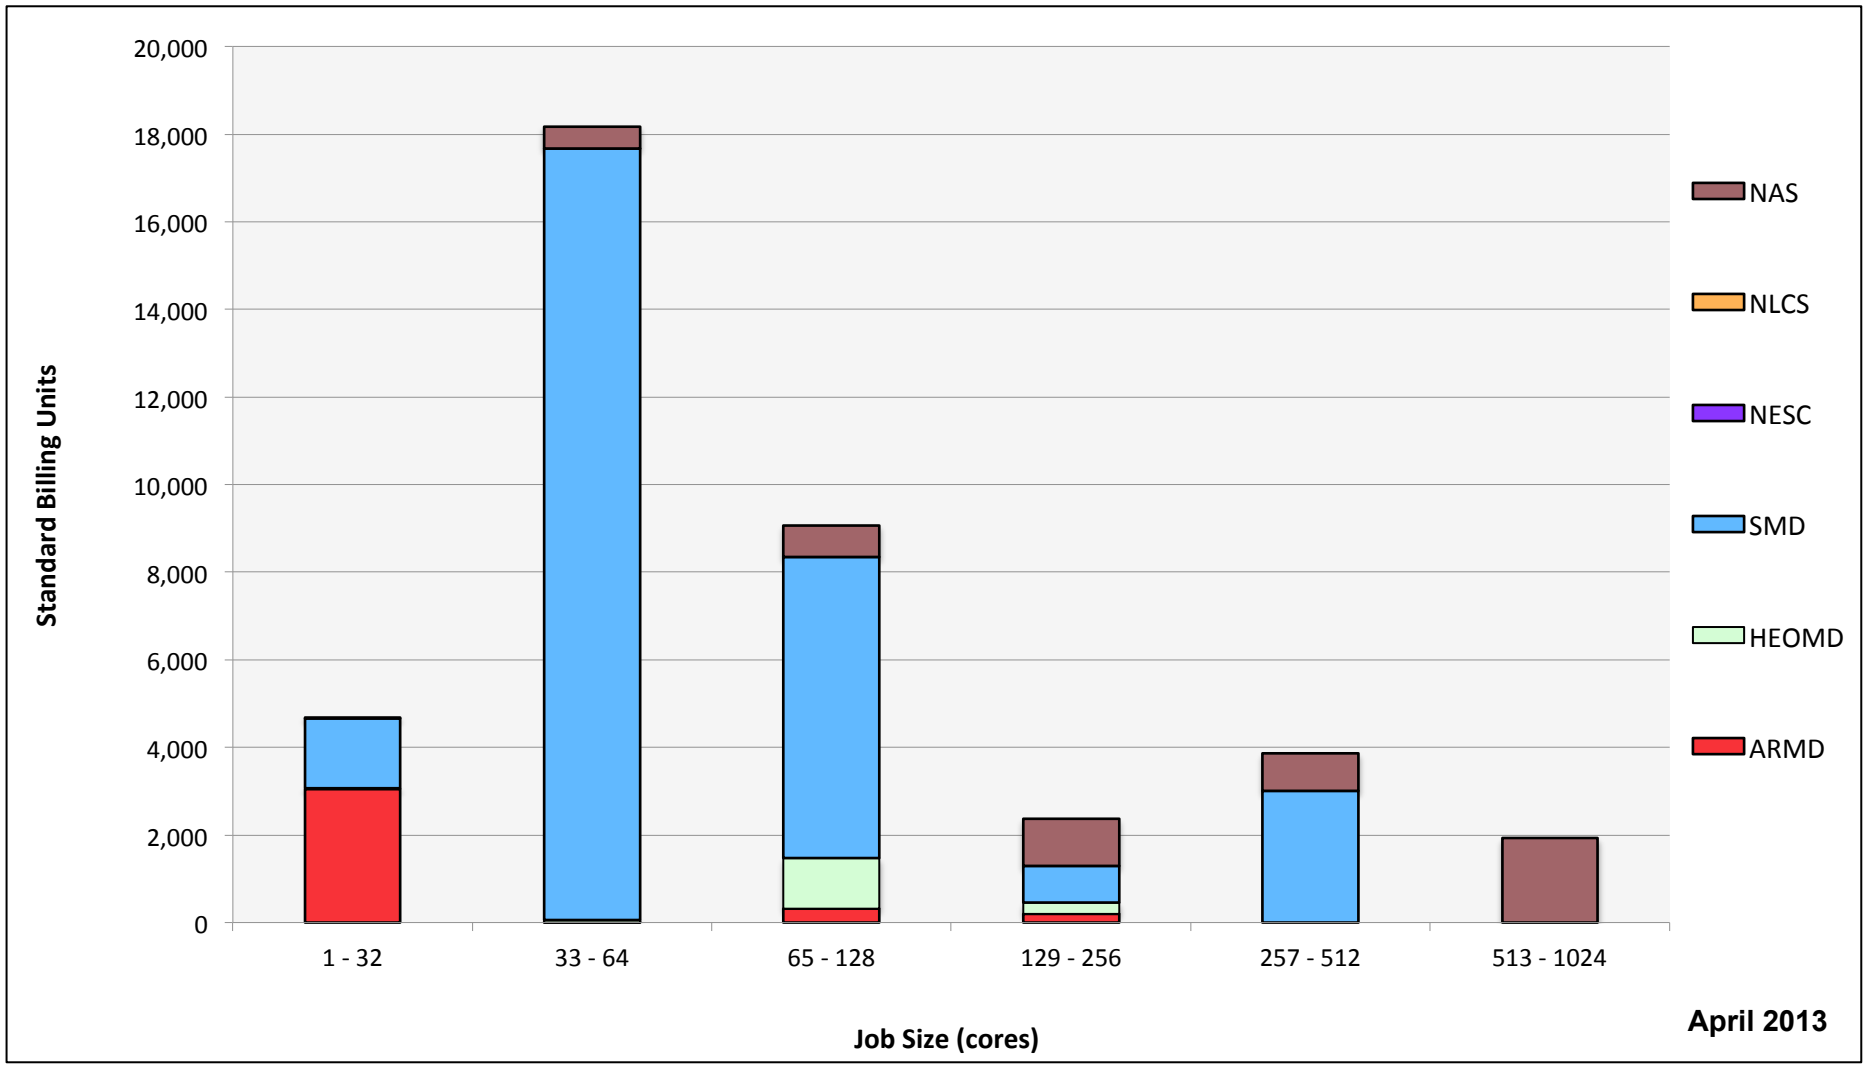

Endeavour:
Monthly Utilization by Size and Mission

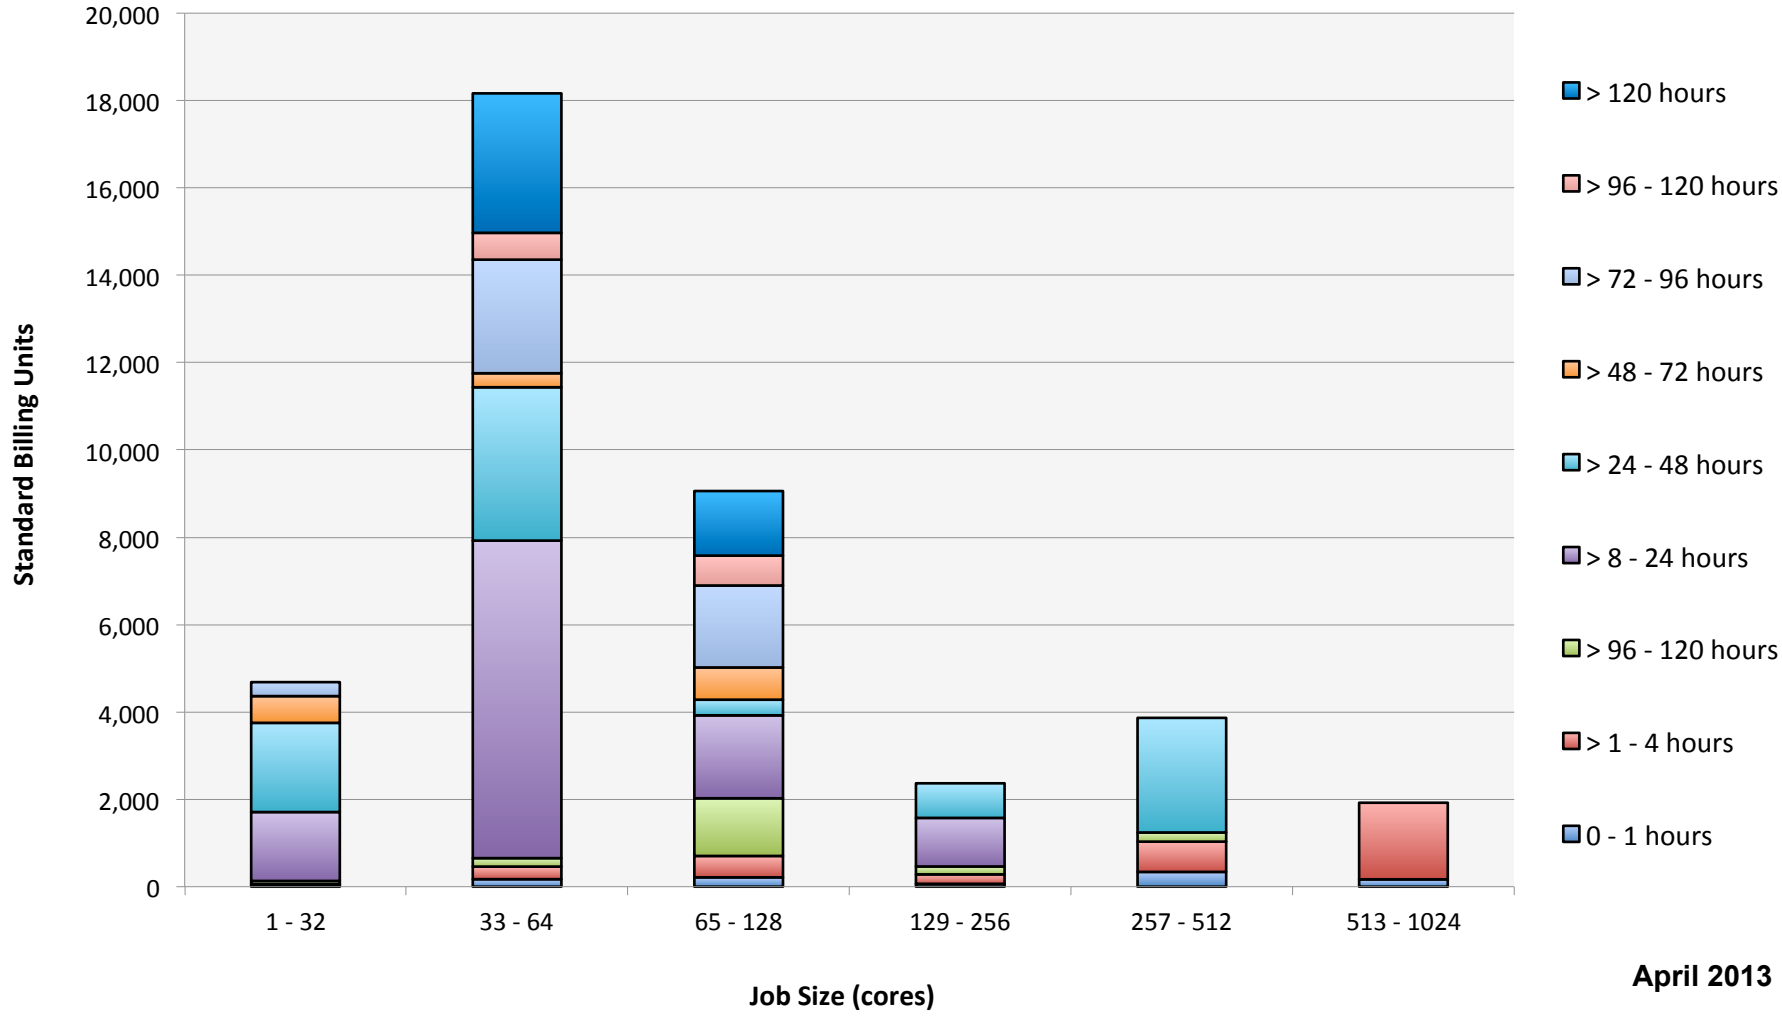

Endeavour:
Average Expansion Factor