Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

10 June 2012

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

Latest Expansion Substantially Enhances Pleiades Computational Capability

- The HECC Systems team, with SGI engineers, replaced 28 racks of Harpertown processors with 24 racks of Sandy Bridge processors, substantially enhancing the computational capability of Pleiades.
- The upgrade resulted in the same number of nodes but increased the system's peak performance by an additional 409 teraflops; the enhanced system now has a theoretical peak performance of 1.74 petaflops.
- In addition to the increase in computational capability, the new racks are equipped with fourteen data rate (FDR) InfiniBand links, which provide 56 gigabits-per-second of network bandwidth—a 40% improvement.
- The racks were integrated during Pleiades dedicated time; other activities completed during this period included upgrading the Lustre filesystem, updating software, replacing failing components, and cleaning up the InfiniBand fabric.

Mission Impact: The latest expansion of Pleiades provides increased computational capability to support the ever-growing requirements of NASA's mission directorates.

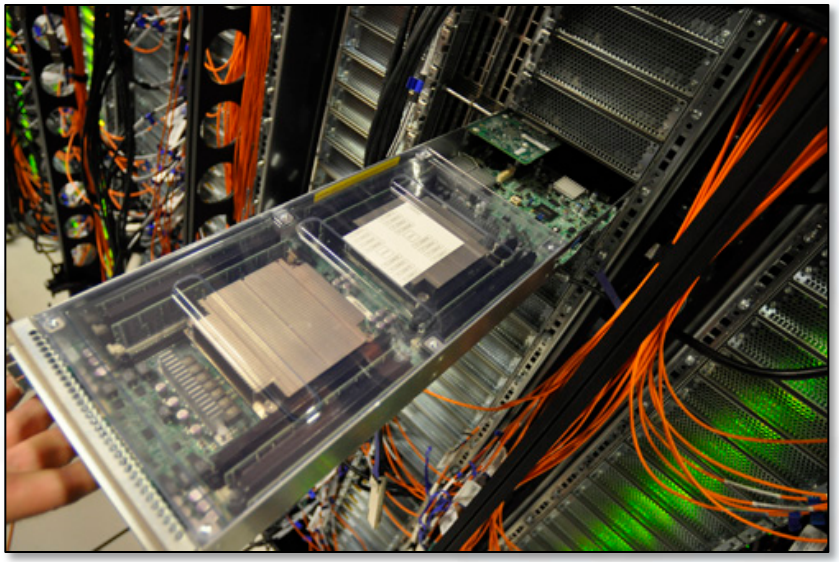

Figure: Twenty-four SGI Sandy Bridge racks added 409 teraflops of additional computational capability to Pleiades.

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

Pleiades Ranks 1st Among World's Fastest Supercomputers on the Graph500

- After the Pleiades dedicated time activities, Intel systems engineers ran the Graph500 benchmark on 1,024 Sandy Bridge nodes, attaining 262 billion traversed edges per second, which would unofficially place the system in the 1st position on the current Graph500 list.
- The Graph500 is a graph benchmark designed to measure the performance of data-intensive supercomputer applications; established in 2010, the list is updated semi-annually in November and June and complements the LINPACK benchmark to provide a more complete view of a system's performance.
- Running the Graph500 benchmark provided valuable information that will be used to help guide future software updates for better application scaling performance.

Mission Impact: Upgrades to Pleiades provide users with one of the fastest supercomputers in the world, enabling efficient computation of data-intensive applications for Agency projects across all mission directorates.

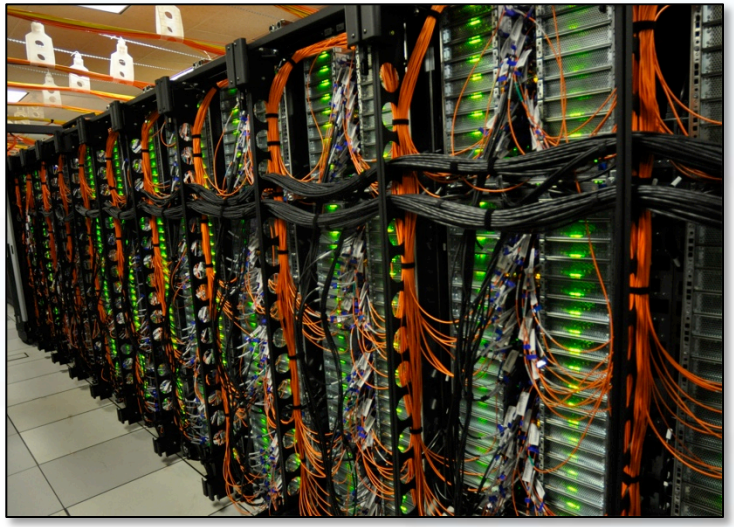

Figure: Engineers attained 262 billion traversed edges per second on the Graph500 benchmark, running on 1,024 nodes on Pleiades.

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

Archive System Relocated for Improved Data Migration

- To facilitate an upgrade from SGI's Data Migration Facility (DMF) storage management system to a new, faster, and more scalable parallel Data Migration Facility (pDMF), HECC systems engineers relocated the Lou2 archive system from the secondary facility at NASA Ames to the primary HECC computer room.
- The relocation of the archive system will simplify the process of migrating to pDMF in the near future.
- HECC staff coordinated this work to overlap with the recent Pleiades dedicated time, in order to minimize the impact on users.
- pDMF will improve access to tape-archived data by eliminating single-server bottlenecks, scaling throughput to and from tape storage, providing redundancy, and lowering cost.

Mission Impact: Reconfiguration of the archive storage system will make the upgrade to parallel DMF more efficient, resulting in increased scalability and reliability for NASA users.

Figure: The Lou2 archive system, an SGI Altix 4700, was transported from a secondary facility located about 1 kilometer away to the primary computer floor at the NASA Advanced Supercomputing facility.

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

Disk Cache Augmented for Archive Systems

- HECC systems engineers deployed new disk cache filesystems on the Lou1 and Lou2 archive systems; the increase in size and performance gives users faster access to archive storage.
- The Lou1 filesystem disk cache was expanded from 310 terabytes (230 TB available to users) to 1.8 petabytes (1.3 PB available), a 5.6x increase in capacity.
- In addition, the Lou2 filesystem disk cache was expanded to 2.4 PB (1.7 PB available), a 4.8x increase in capacity.
- The increased disk cache capacity results in faster access time for researchers by allowing data to remain on disk longer before migrating to tape storage.

Mission Impact: Increased disk cache on the HECC archive systems provide faster data access, resulting in higher user productivity.

Figure: The new filesystems provide a combined 4.2 petabytes of disk cache to the HECC archive systems.

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

HECC Network and Security Experts Improve Data Transfer Rates for Remote Users

- HECC network and security experts employed innovative techniques to significantly improve data transfer rates to and from remote sites that are required to use High Performance Network (HPN) software.
- A bug in HPN limited all HPN-requested transfers to 1.1 megabytes per second (MB/s), affecting sites such as the NASA Center for Climate Simulation (NCCS) and Langley Research Center (LaRC).
- The two teams identified and removed a faulty HPN function, and implemented new HPN interfaces on the HECC systems.
- On systems at both LaRC and NCCS, transfer rates jumped from 1.1 MB/s to over 70 MB/s with the fully HPN-enabled environment—a 63-fold improvement.

Mission Impact: Identifying and removing network bottlenecks allows HECC users to move data at optimal rates, decreasing scientific time-to-solution for the NASA user community.

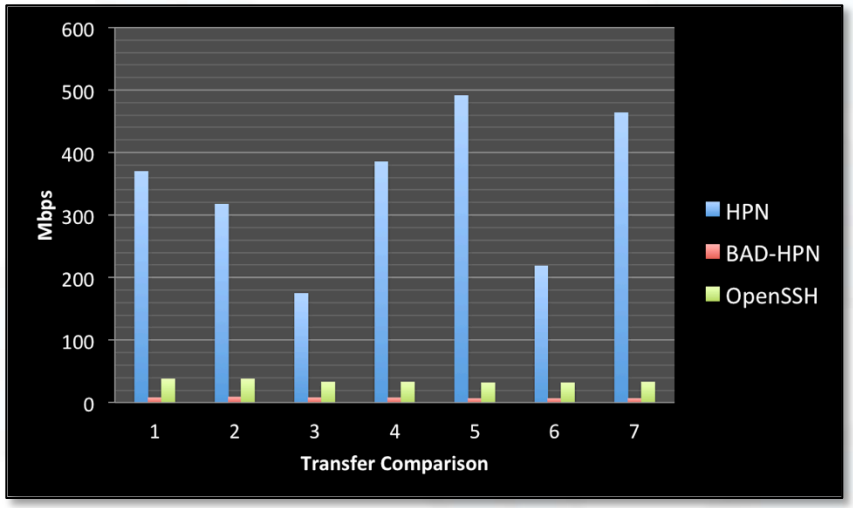

Figure: This graph shows the results of several tests from LaRC systems to NASA Ames using HPN, HPN-disabled, and OpenSSH software.

POCs: *Nichole Boscia, nichole.k.boscia@nasa.gov, (650) 604-0891*
Derek Shaw, derek.g.shaw@nasa.gov, (650) 604-4229, NASA Advanced Supercomputing Division, Computer Sciences Corp.

New Allocation Period Begins for HEOMD, NESC, and SMD

- May 1st marked the beginning of a new allocation period for the Human Exploration and Operations Mission Directorate (HEOMD), the NASA Engineering and Safety Center (NESC), and the Science Mission Directorate (SMD).
- The three groups awarded new allocations on Columbia and Pleiades to over 200 computing projects that support science and engineering activities.
- Combined awards exceeded 65 million Standard Billing Units* (SBUs)—a 20% increase over the previous year.
- The recent expansion of Pleiades makes it possible for HECC to keep pace with these continually growing demands for resources.
- The new allocation period provides an opportunity for each organization to rebalance allocations to meet computing needs for the upcoming year.

1 SBU equals 1 hour of a Pleiades Westmere 12-core node.

Mission Impact: NASA programs and projects periodically review the distribution of supercomputer time to assure consistency with mission-specific goals and objectives.

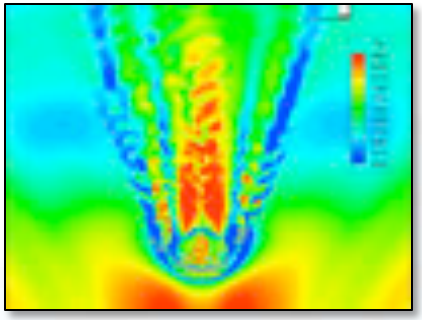

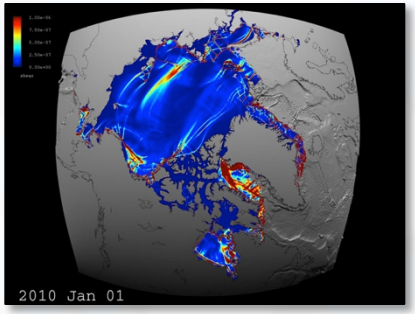

Figures: Representative images from HEOMD and SMD projects. Above left: Transition to turbulence during Mars entry. Steven Yoon, Michael Barnhardt, NASA/Ames. Right: Sea-ice shear. Gunnar Spreen, NASA/JPL, Tim Sandstrom, NASA/Ames

POC: *Catherine Schulbach, catherine.h.schulbach@nasa.gov, (650) 604-3180, NASA Advanced Supercomputing Division*

Hypersonic Turbulence and the Birth of Stars

- Knowledge of the star formation process is key to understanding how galaxies form and evolve. Researchers at the University of Barcelona are producing very large simulations to better understand the hypersonic turbulence of star formation from clouds of interstellar gas.
- The main goal of this project is to understand how stars are formed, by performing two types of numerical computations:
  - Idealized hypersonic turbulence flow experiments to capture universal statistics that will be applied in the theory of star formation.
  - Realistic astrophysical simulations to test and further guide the theory development.
- Due to the complexity of magnetized hypersonic turbulence and its interaction with gravity, these extremely large numerical computations—enabled by Pleiades—are essential to investigate the birth of stars.

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

Mission Impact: The ambitious scope of this project requires the largest hypersonic turbulence simulations in the world. These simulations could not be run without NASA's powerful and efficient high-end computing resources, particularly the Pleiades supercomputer.

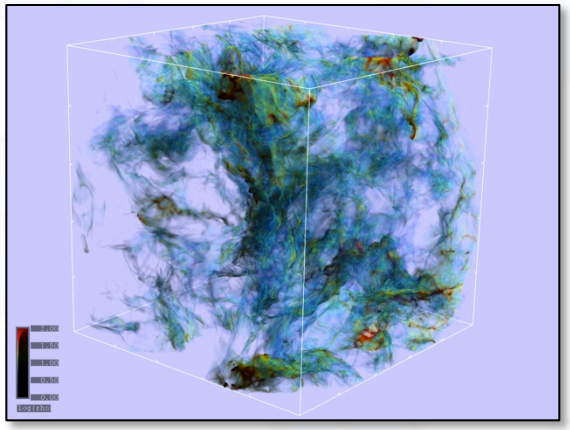

Figure: Volume rendering of the density field from an idealized experiment of magnetized hypersonic turbulence, without self-gravity. Dense filaments and cores are formed from turbulent shocks, even without the assistance of self-gravity. Tim Sandstrom, NASA/Ames

POC: *Paolo Padoan, ppadoan@icc.ub.edu, Catalan Institution for Research and Advanced Studies, Institute of Cosmos Sciences (ICREA & ICC), University of Barcelona*

High-Performance Computing for Extended Formation Flight

- Aircraft flown in formation can obtain a significant reduction in drag due to wake upwash from the lead aircraft.
- Researchers at NASA Ames are performing high-fidelity simulations of formation flight aerodynamics to provide more accurate predictions of drag reduction and to obtain energy savings.
- Formation flight involves aerodynamic phenomena spanning lengths from miles to millimeters; this wide range of scales, combined with the output's sensitivity to slight errors, makes these analyses very challenging.
- After simulation results are verified with flight-test data, they are used in models of the National Airspace System to help understand the potential of formation flight to reduce fuel consumption.

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

Mission Impact: With problem sizes that are on the order of a billion degrees of freedom, and with hundreds of cases to consider, the availability of NASA's high-end computing resources is key to this research.

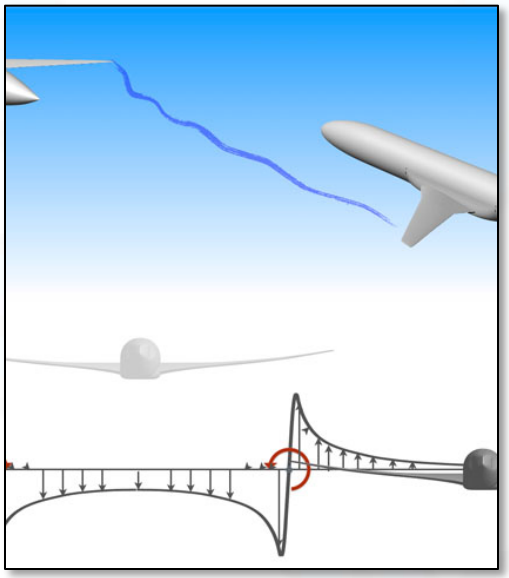

***Figure**: Formation flight can produce drag savings of around 20% due to additional lift provided by the tip-vortex of the lead aircraft. The schematic in the lower portion shows the counter-rotating vortices of the lead aircraft in red and the cross-flow velocity field induced by these vortices in black. Michael Aftosmis, NASA/Ames; S. Andrew Ning, Stanford University*

***POC:** Mike Aftosmis, michael.aftosmis@nasa.gov, (650) 604-4499, NASA Advanced Supercomputing Division*

Status of Requests for NAS Computer Accounts by non-U.S. Citizens

- Requests approved: 13; New requests received: 8; Requests waiting: 5.
- The time to approval continues to improve. The eight requests that were approved were approved in less than one month.

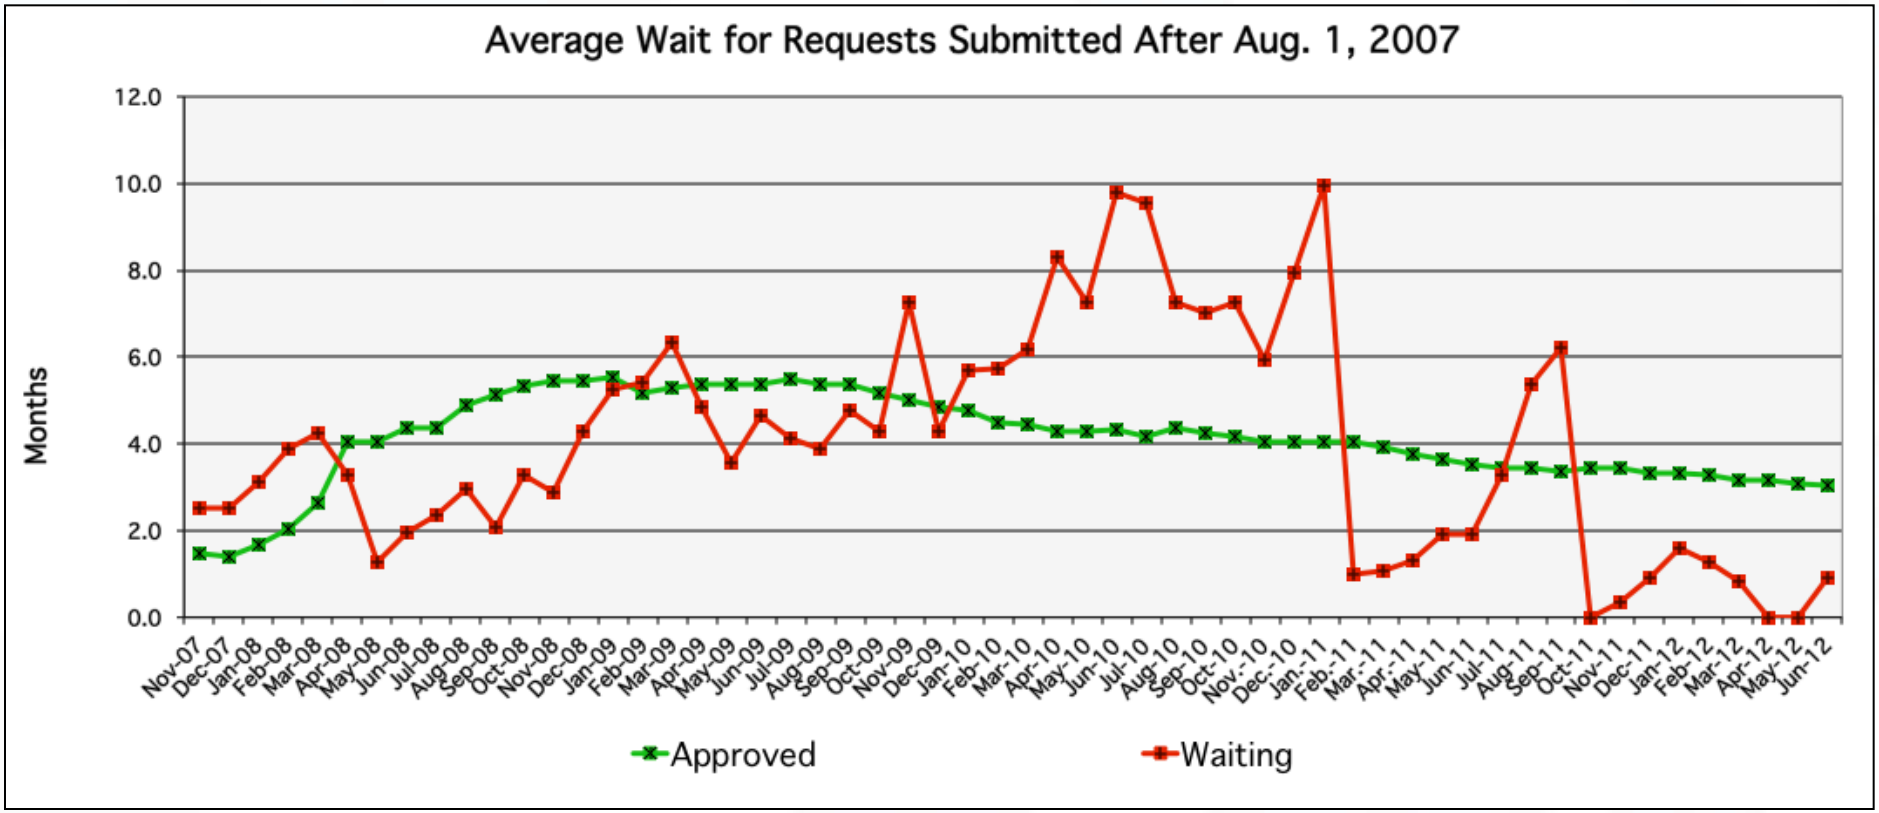

HECC Facility Hosts Several Visitors and Tours in May 2012

- HECC hosted 6 tour groups in May; guests learned about the Agency-wide missions being supported by Pleiades, and viewed scientific results on the hyperwall. Visitors included:
  - Bulgarian President, Rosen Plevneliev—along with an entourage of 2 ambassadors, 4 advisors and a large press team.
  - A group from Rutherford Appleton Laboratory (which works with the UK Space Agency that coordinates UK civil space activities), who met with HECC senior staff.
  - A technical group from ATK Aerospace Systems, who met with HECC researchers to discuss how we can support future research in computational fluid dynamics.
  - A group from the NASA's Strategic Capabilities Assets Program (SCAP).
  - A White House staffer, who visited after President Obama's recent stop-over in the San Francisco Bay Area.
  - The president of the Italian Space Agency, who was visiting Ames to discuss potential collaborations.

Figure: Bulgarian President Rosen Plevneliev and his entourage visited the NASA Advanced Supercomputing (NAS) facility, where they received an overview of HECC and NAS Division activities, along with a hyperwall demonstration, and a tour of the Pleiades supercomputer system.

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

Presentations and Papers

- **"Detached Eddy Simulation of the UH-60 Rotor Wake Using Adaptive Mesh Refinement,"** Neal Chaderjian and Jasim Ahmad, presented at the American Helicopter Society 68th Annual Forum, Fort Worth, TX, May 1–3, 2012.*
- **"Improvements to the Pegasus5 Overset CFD Software,"** Stuart Rogers, presented at the Applied Modeling & Simulation (AMS) Seminar Series, NASA Ames Research Center, May 15, 2012.*
- **"Simulation of a Heavy-Lift Slowed-Rotor Compound Helicopter,"** Jasim Ahmad, Timothy Sandstrom, Bryan Allan, presented at 24th International Conference on Parallel CFD, Atlanta, GA, May 21–25, 2012.*
- **"Magnetospheric Configuration and Dynamics of Saturn's Magnetosphere: A global MHD simulation,"** Xianzhe Jia, Kenneth C. Hansen, Tamas Gombosi, Journal of Geophysical Research, Vol. 117, A05225, 22 pp., May 18, 2012.* *http://www.agu.org/pubs/crossref/2012/2012JA017575.shtml*
- **"Aerodynamics and Debris Transport for the Space Shuttle Launch Vehicle,"** Stuart Rogers, The 2012 I.I. Glass Lecture, presented at the University of Toronto Institute for Aerospace Studies, May 22, 2012.*

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

News and Events

- **Amazon's HPC Cloud: Supercomputing for the 99 Percent,** *Wired Cloudline,* May 2, 2012 – *Feature story comparing Amazon's cloud capability for scientific applications with that of supercomputing centers; includes an image of Pleiades.*
  http://www.wired.com/cloudline/2012/05/amazon-hpc-cloud/

- **Bulgaria's President Starts Super Computer of NASA**, *Standart, English Edition,* May 19, 2012 – News blurb featuring image from President Rossen Plevneliev's tour of the NASA Advanced Supercomputing facility.
  http://paper.standartnews.com/en/article.php?d=2012-05-19&article=39160

- **Are Cloud Environments Ready for NASA HPC Applications?"** HECC user webinar, May 31, 2012 – This webinar presented the results of a study comparing the performance of two cloud environments, Amazon's AWS and SGI's Cyclone, to that of NASA's Pleiades supercomputer to determine the suitability of cloud computing for HPC applications.
  http://www.nas.nasa.gov/hecc/support/past_webinars.html

HECC Utilization

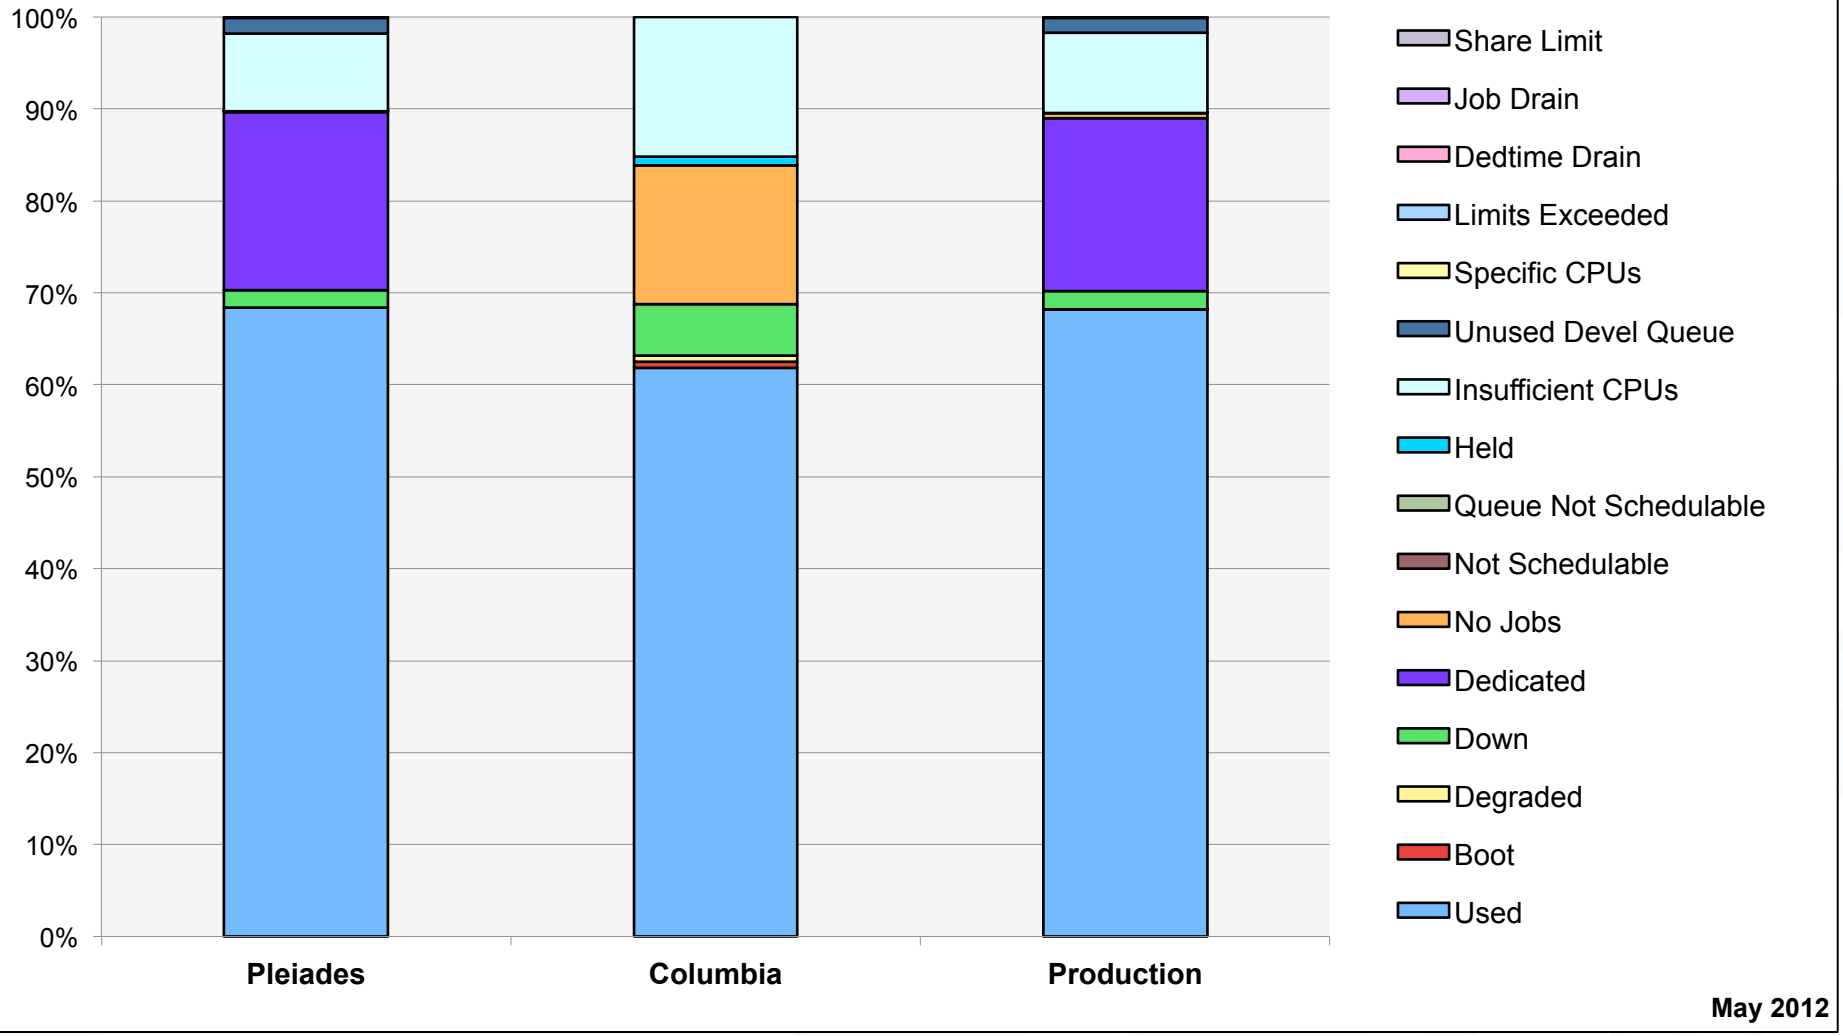

HECC Utilization Normalized to 30-Day Month

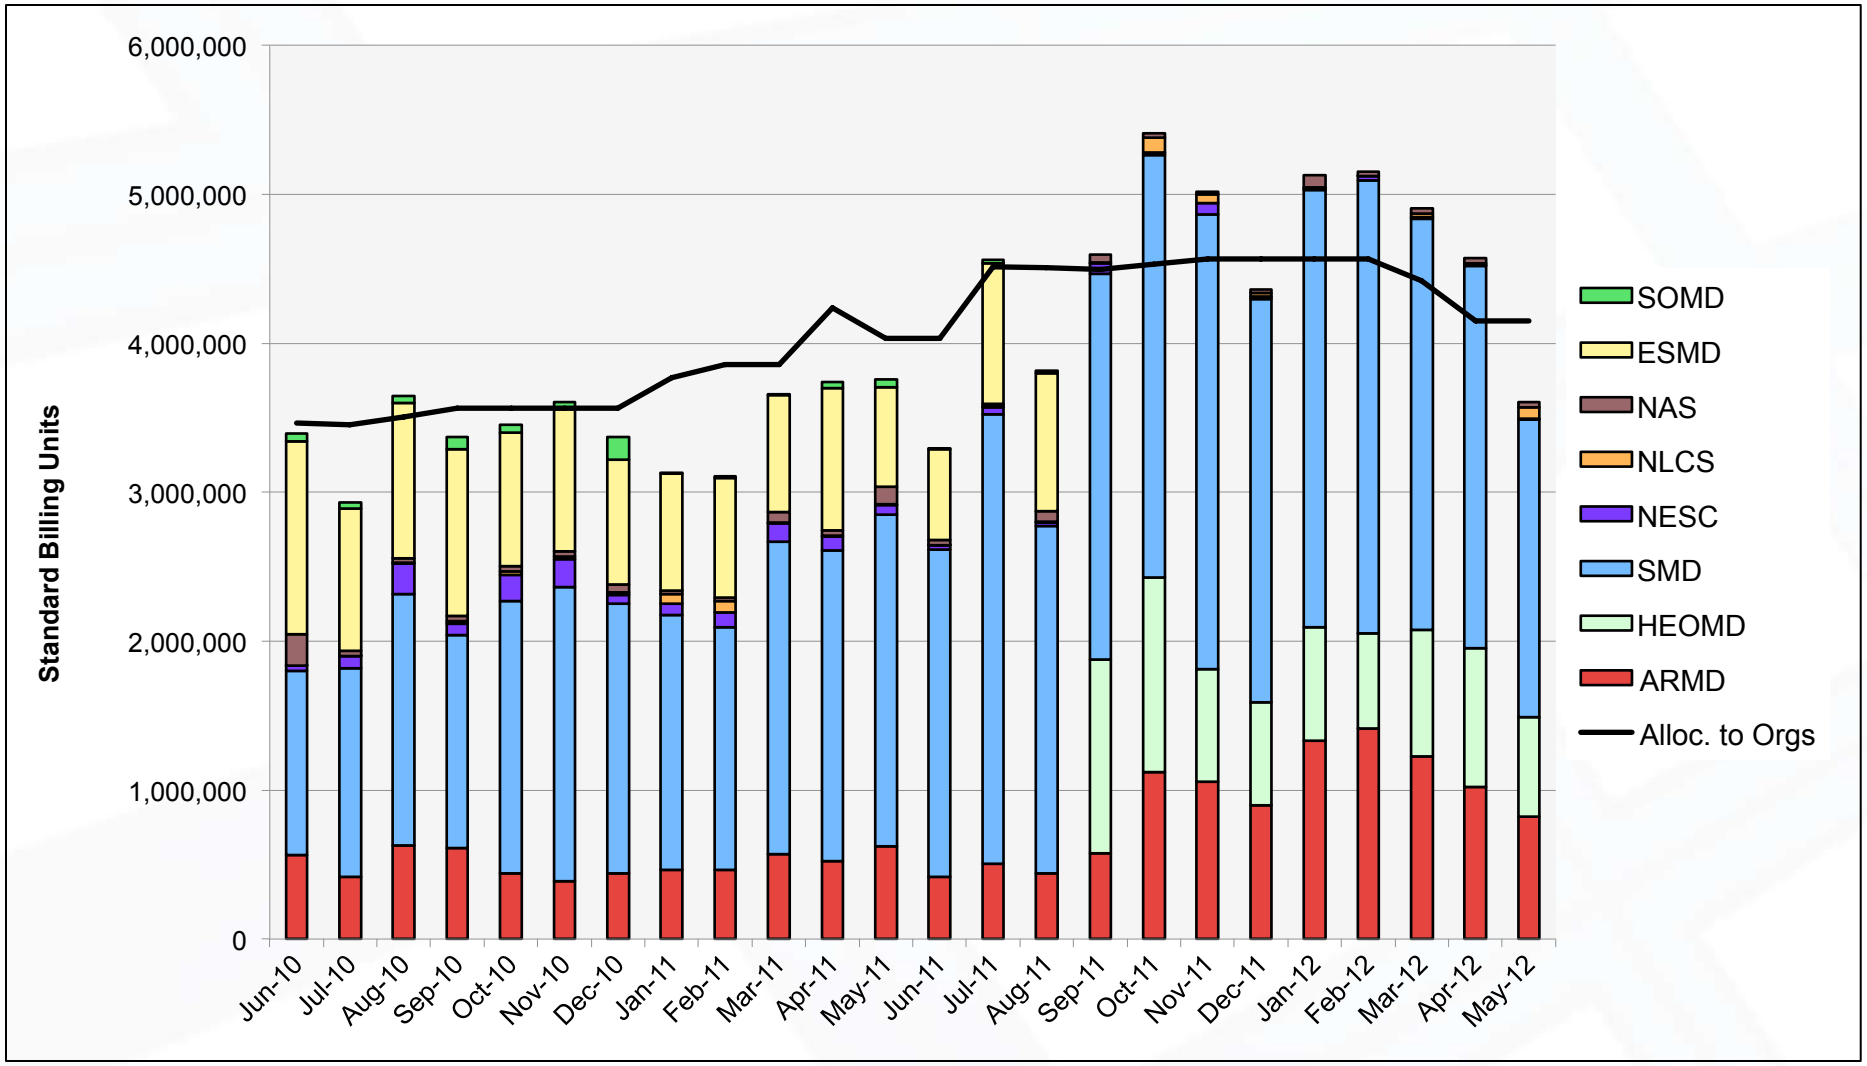

HECC Utilization Normalized to 30-Day Month

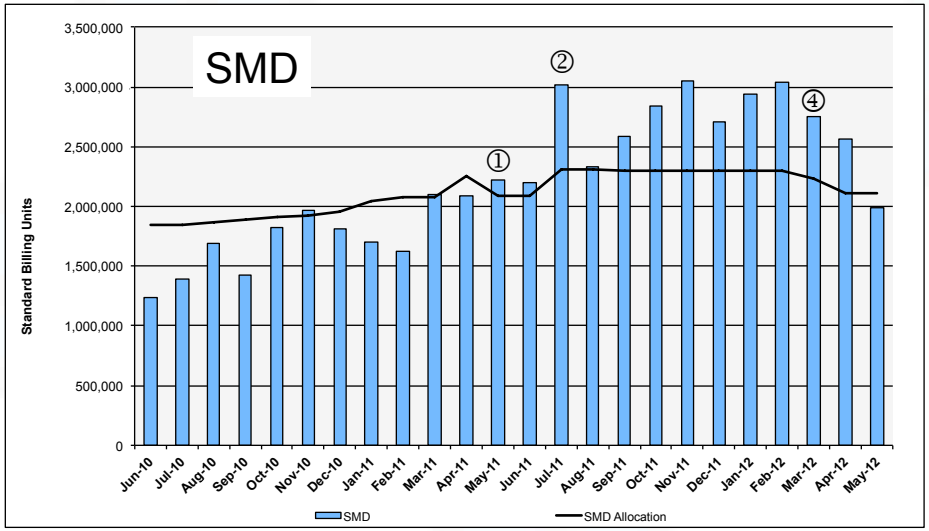

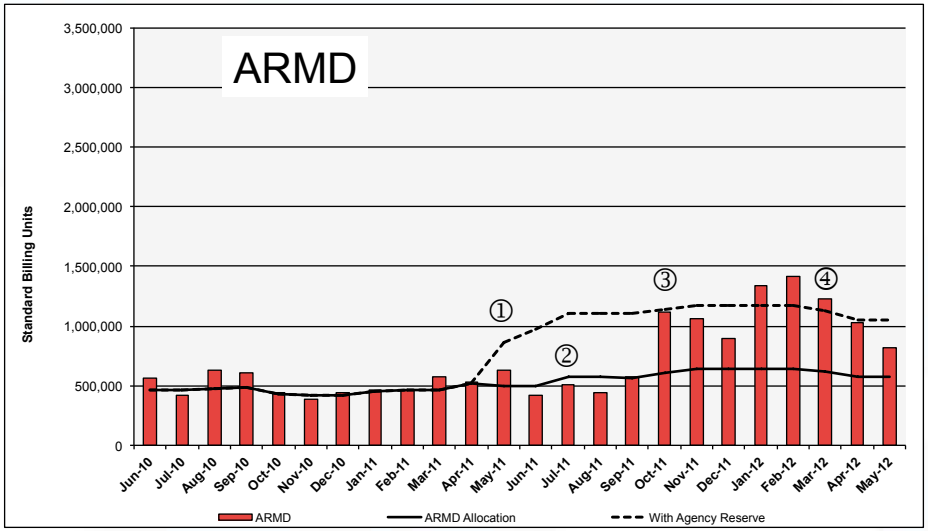

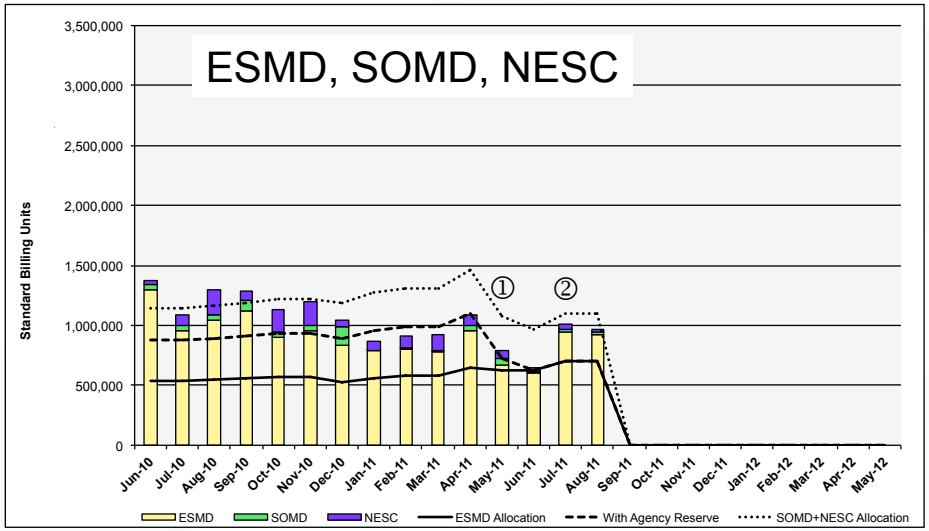

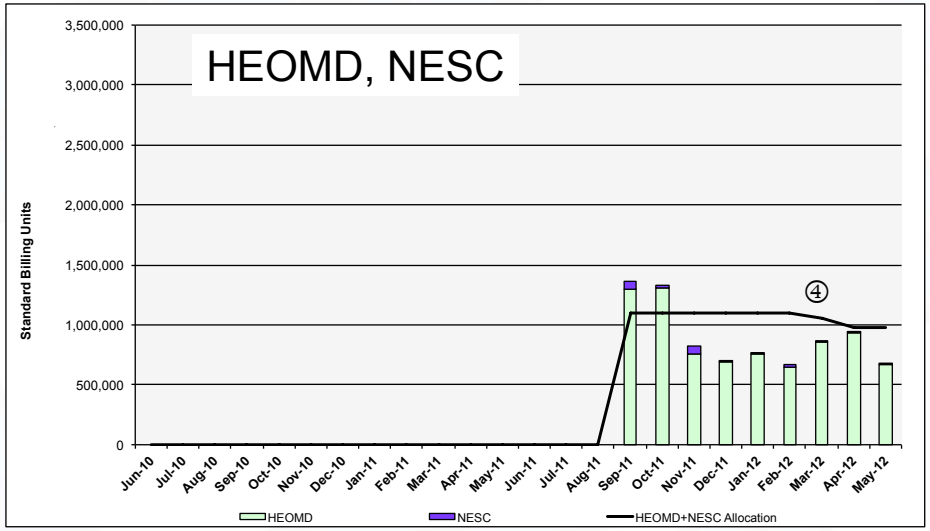

① Allocation to orgs. decreased to 75%, Agency reserve shifted to ARMD ② 14 Westmere racks added
③ 2 ARMD Westmere racks added ④ 28 Harpertown racks removed

Tape Archive Status

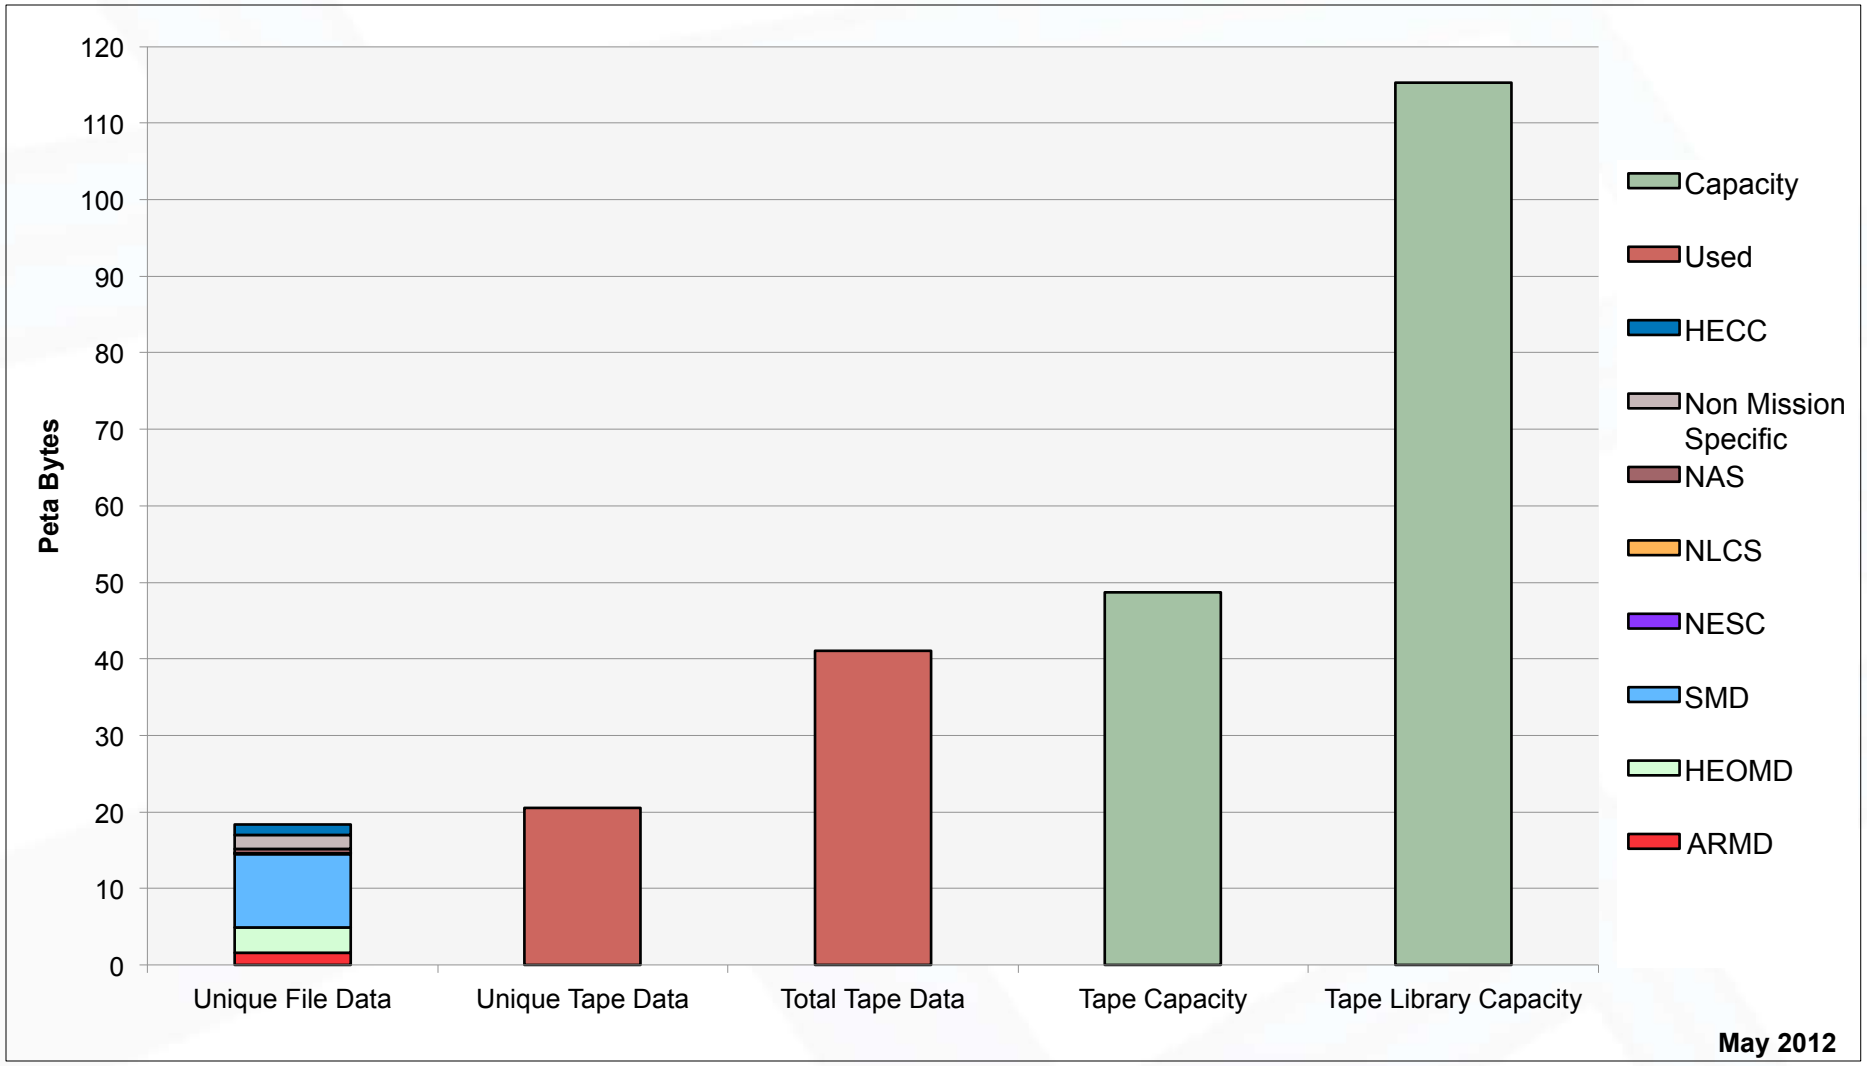

Tape Archive Status

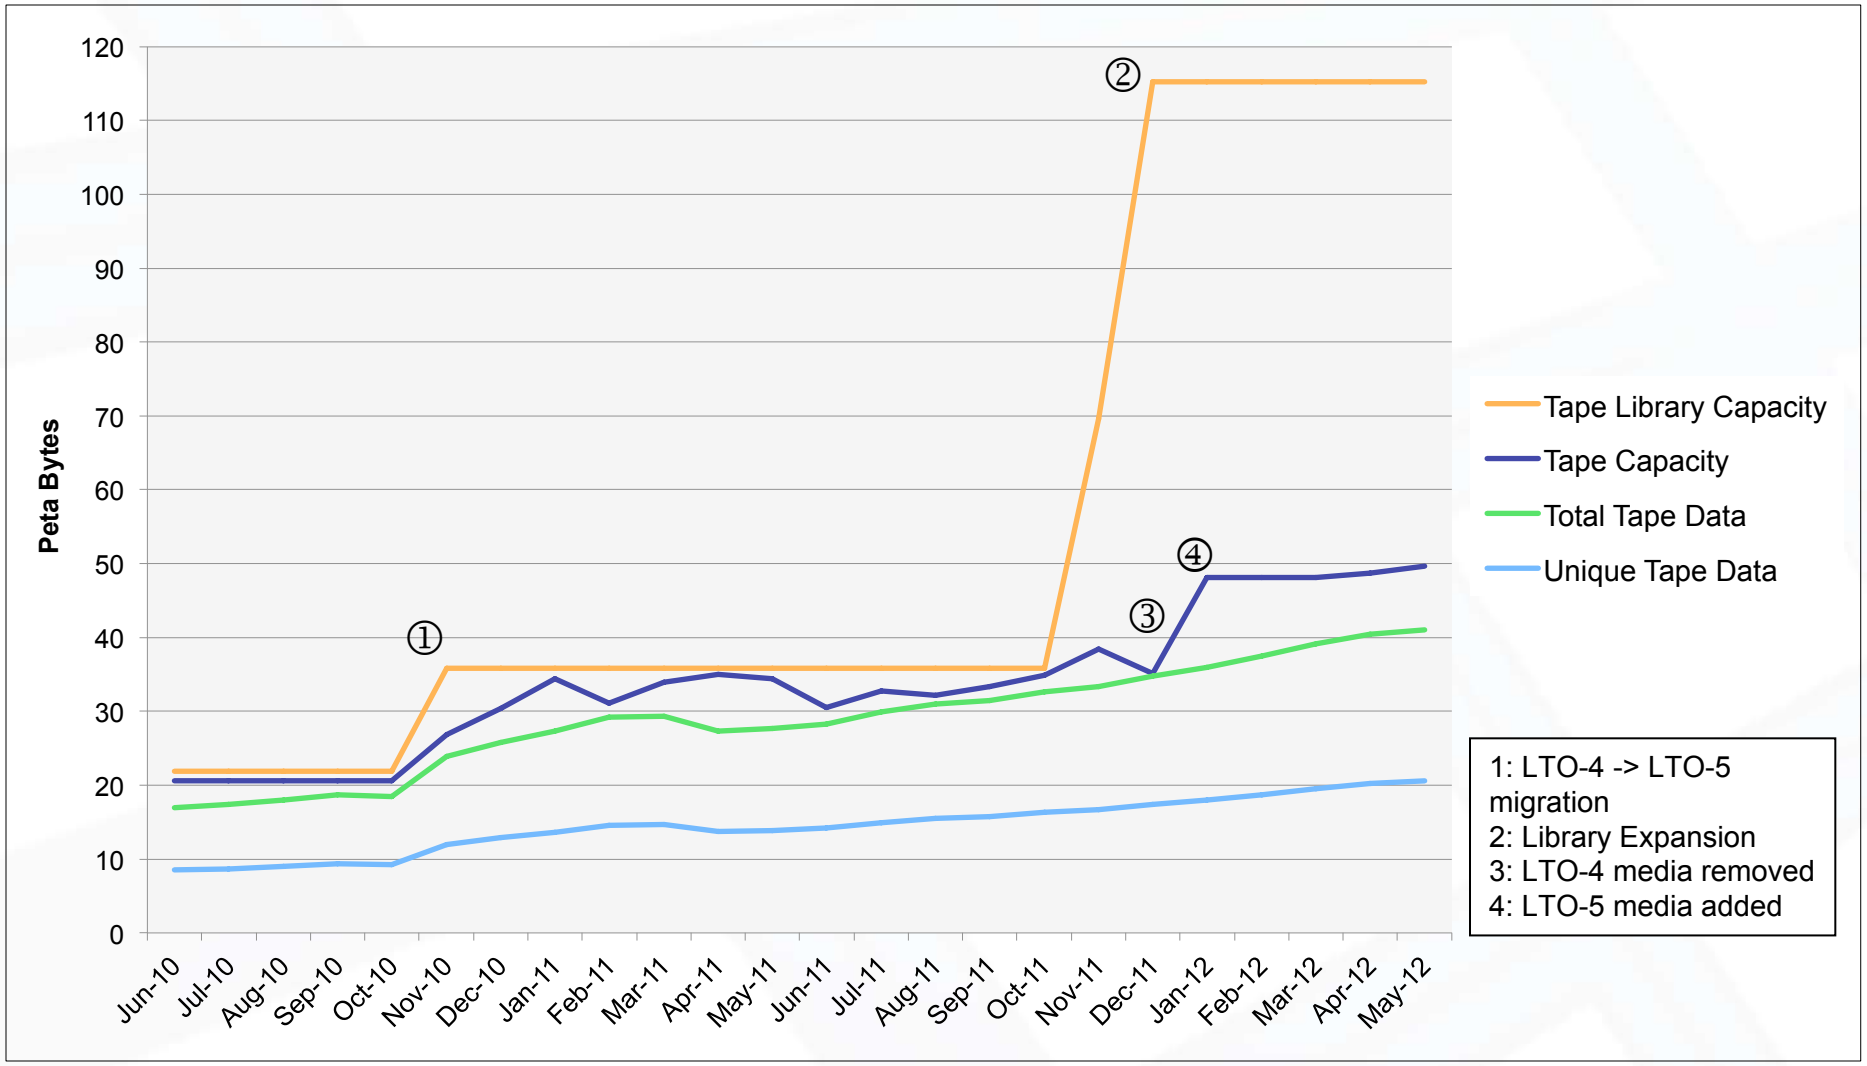

Pleiades:
SBUs Reported, Normalized to 30-Day Month

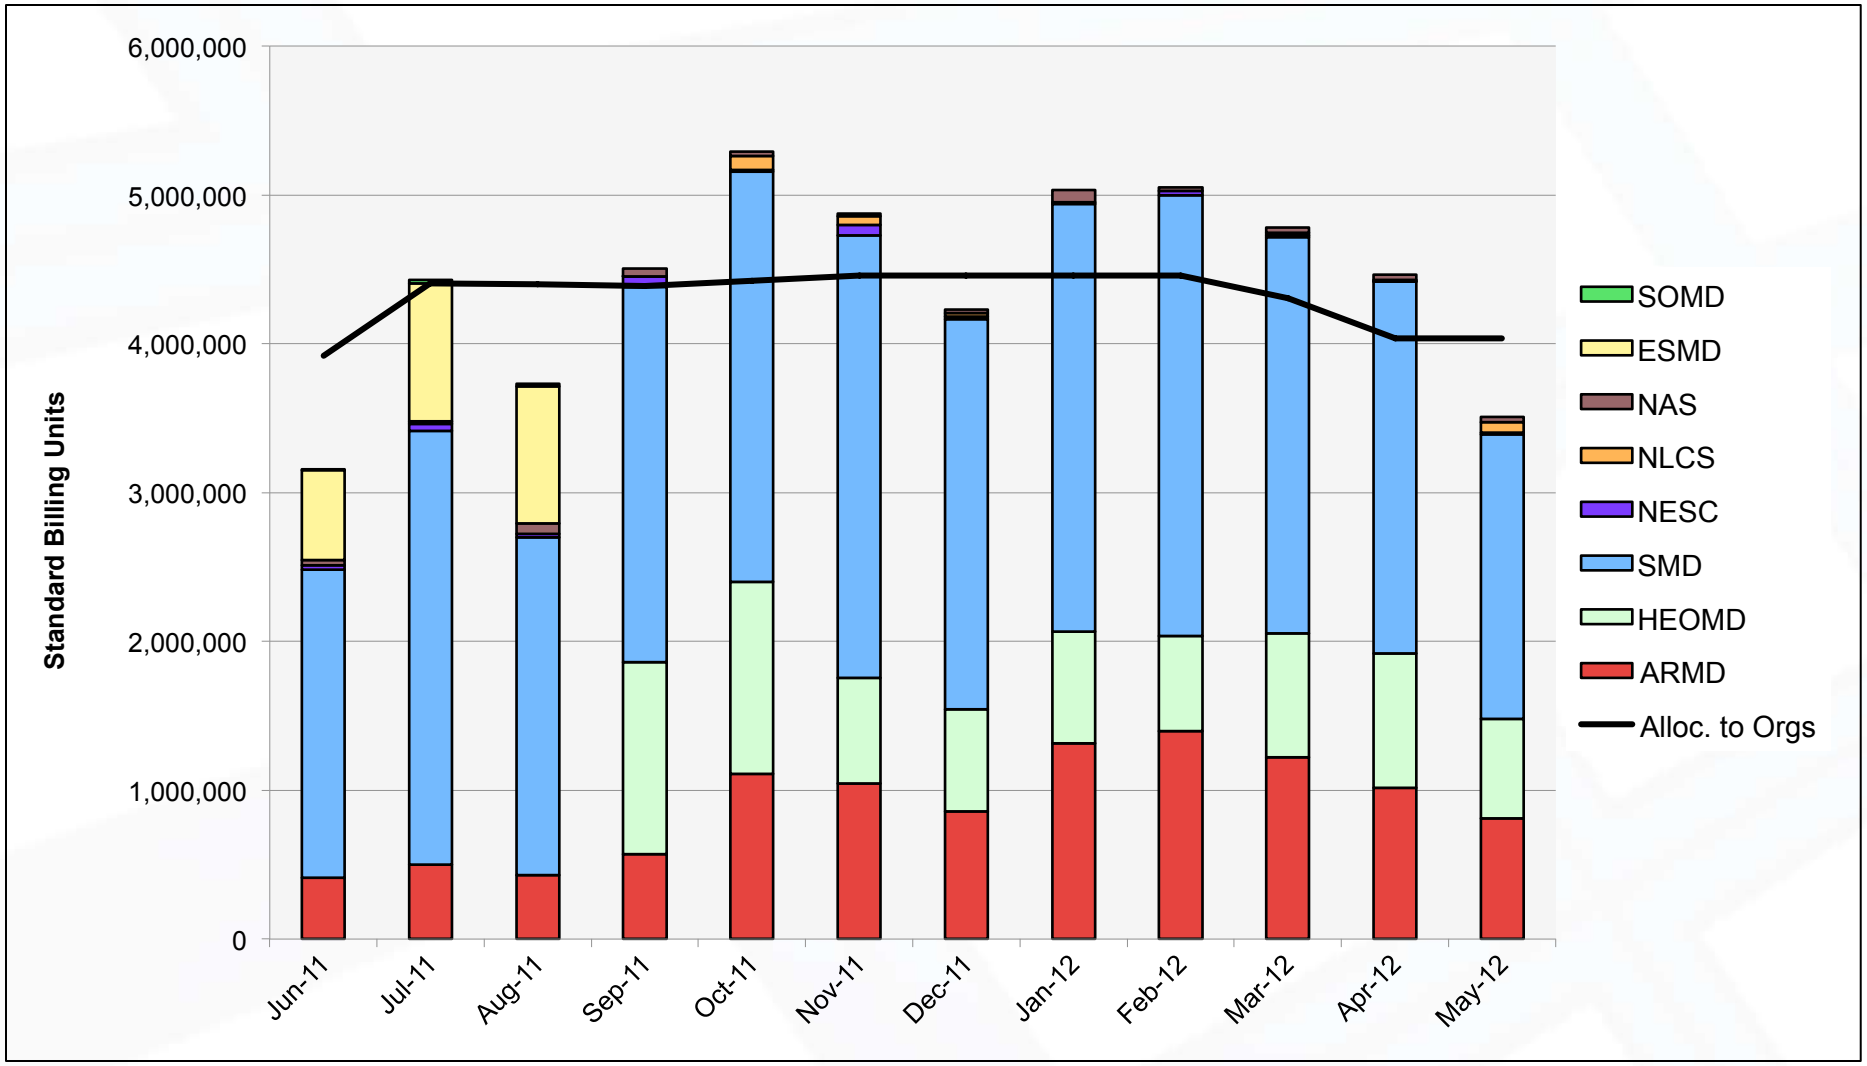

Pleiades:
Devel Queue Utilization

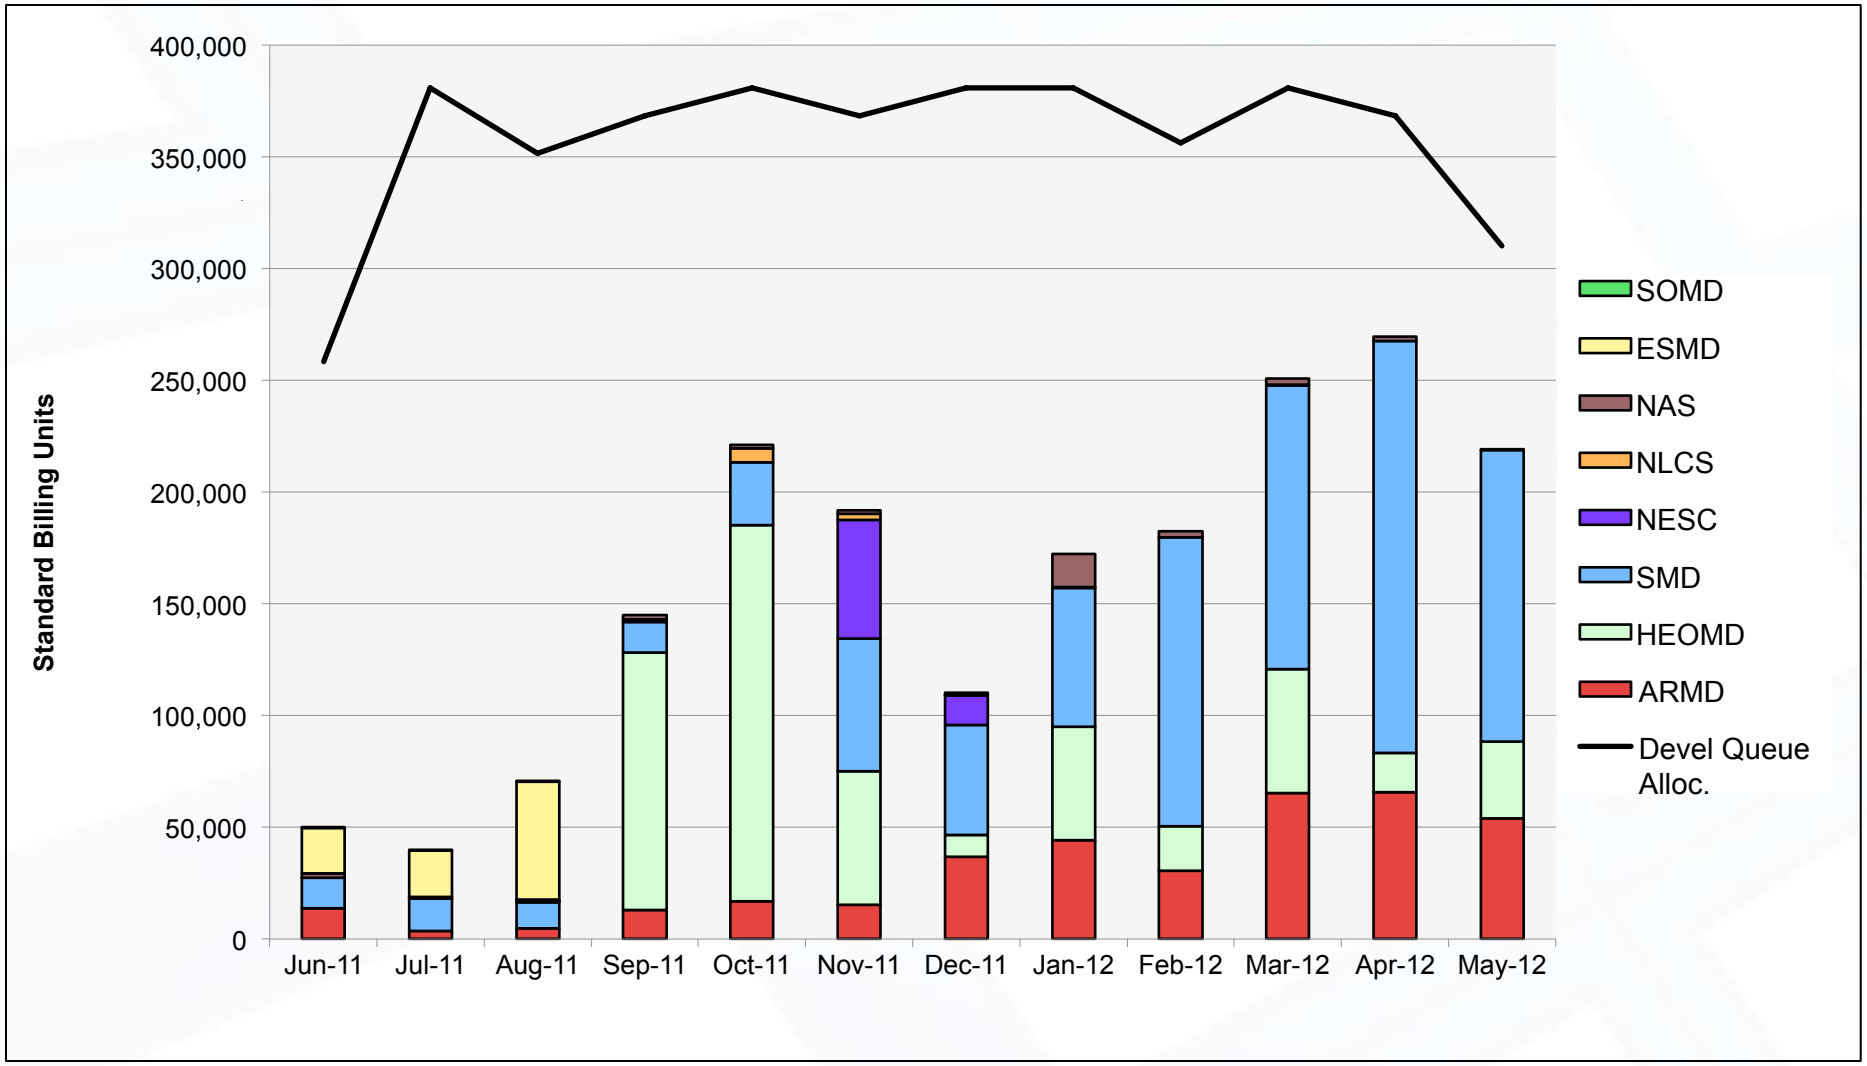

Pleiades:
Monthly SBUs by Run Time

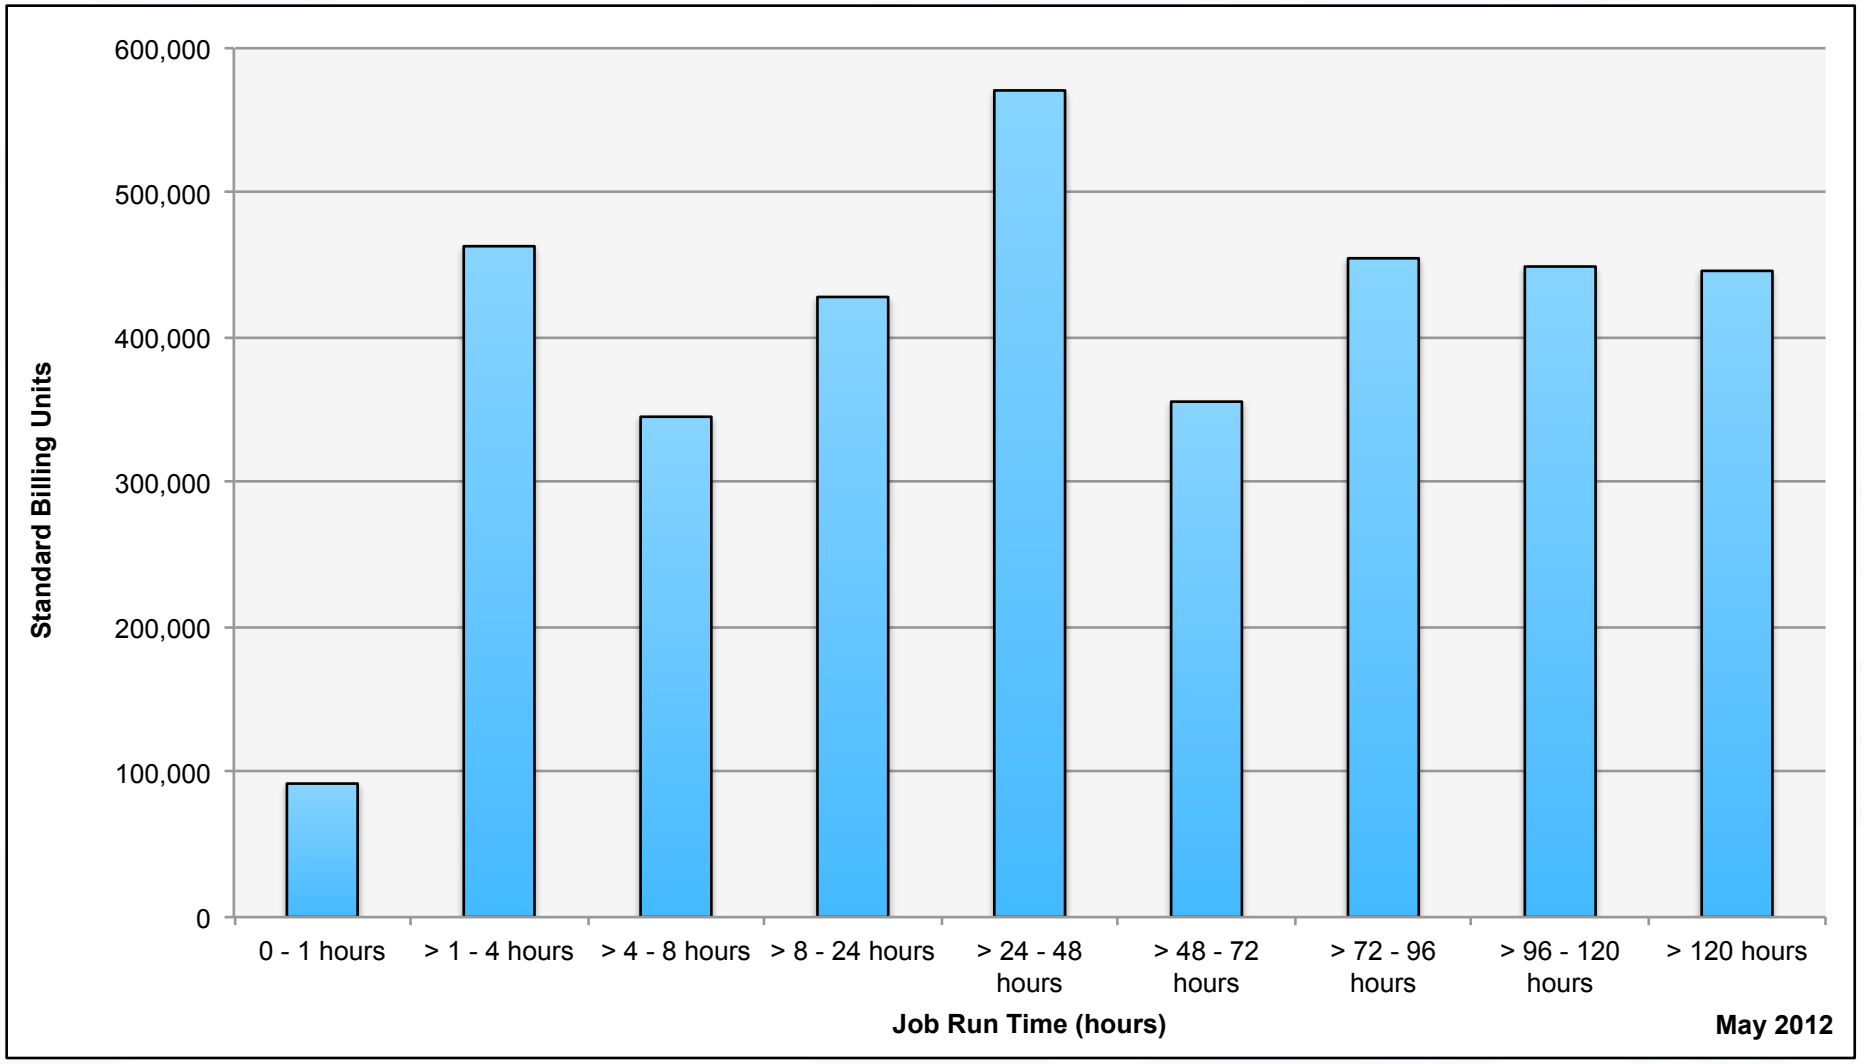

Pleiades:
Monthly Utilization by Size and Mission

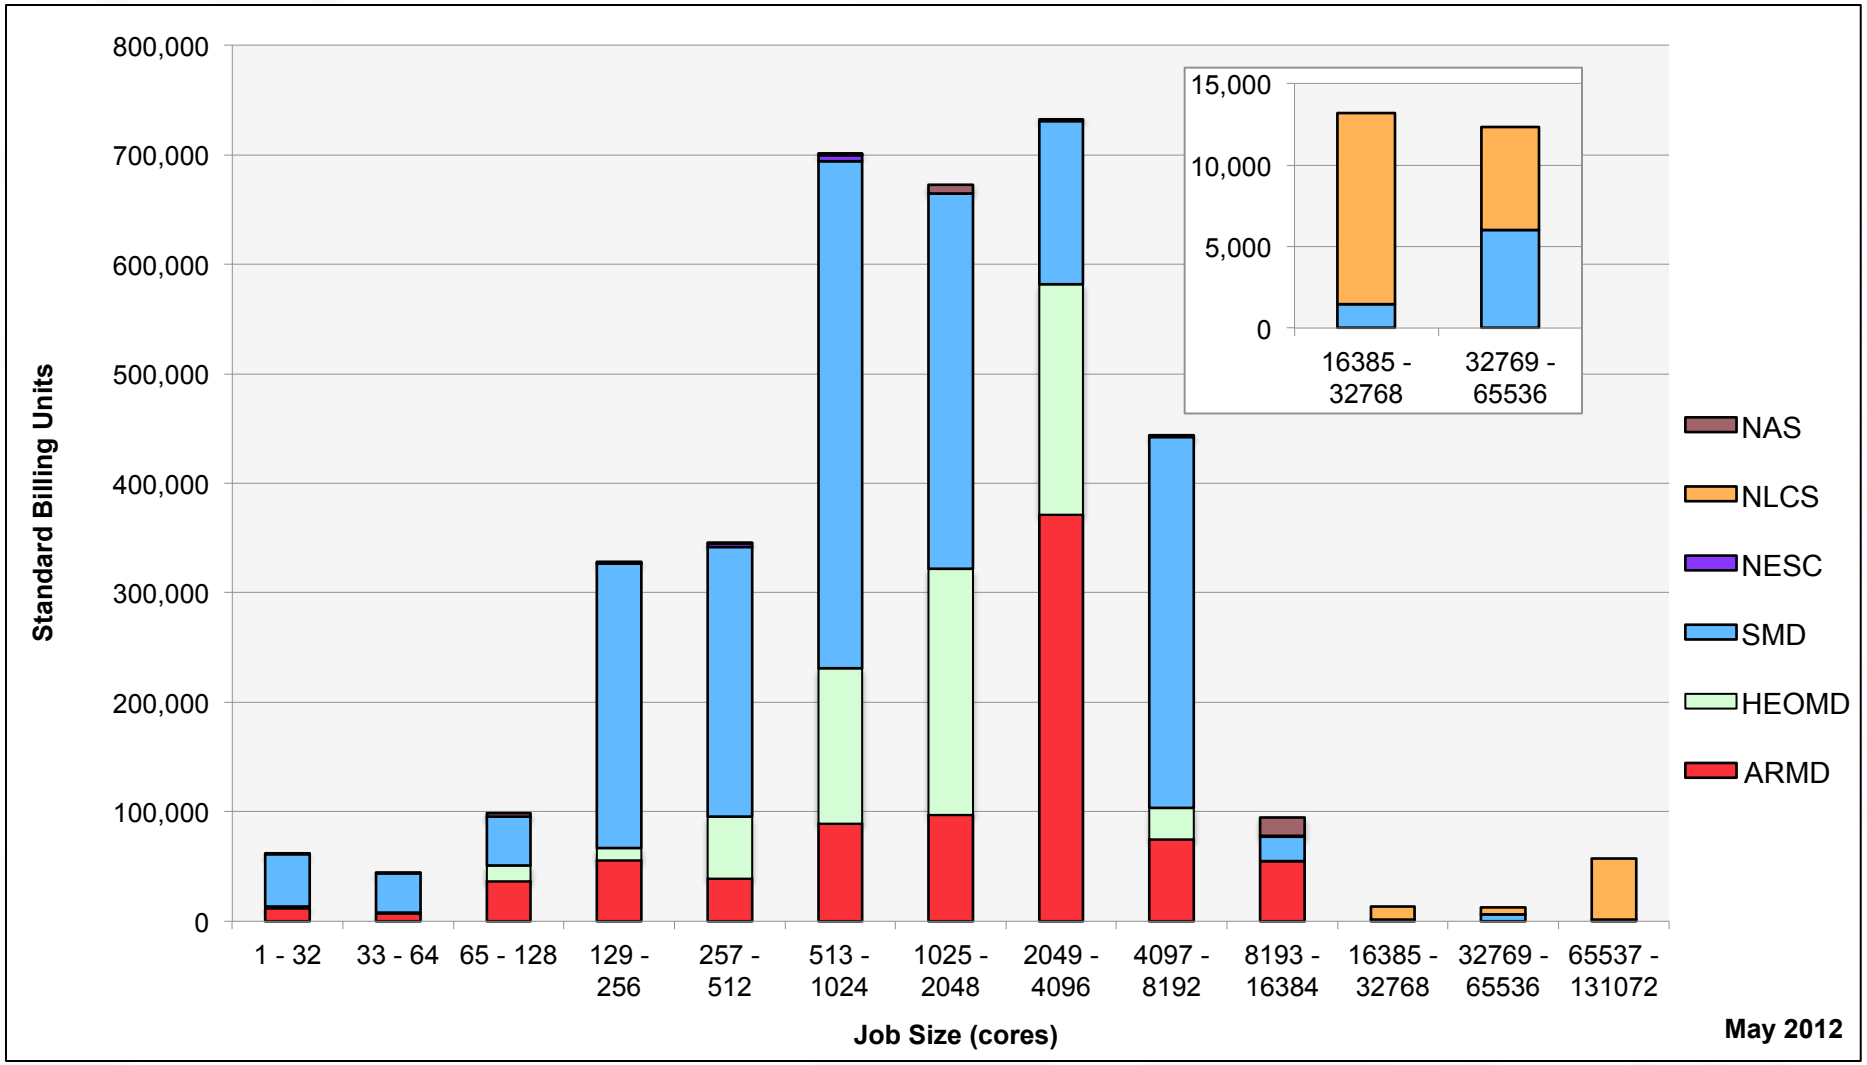

Pleiades:
Monthly Utilization by Size and Length

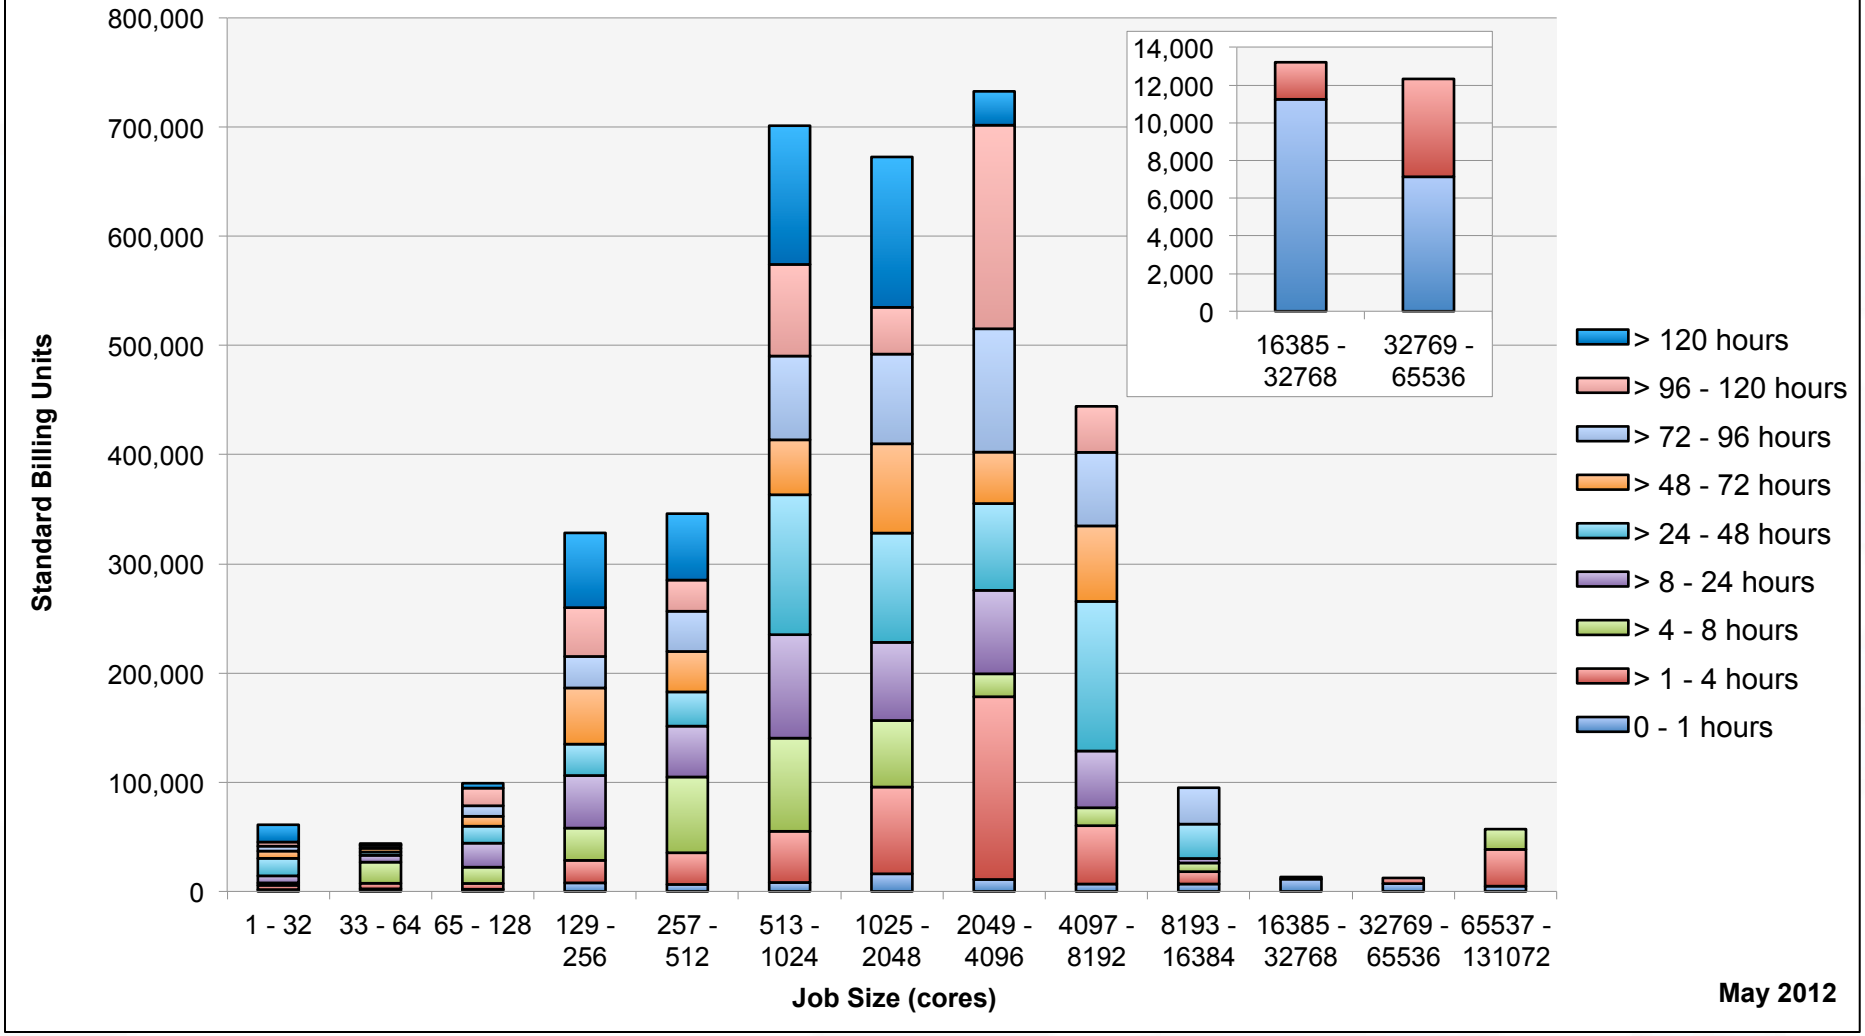

Pleiades:
Average Time to Clear All Jobs

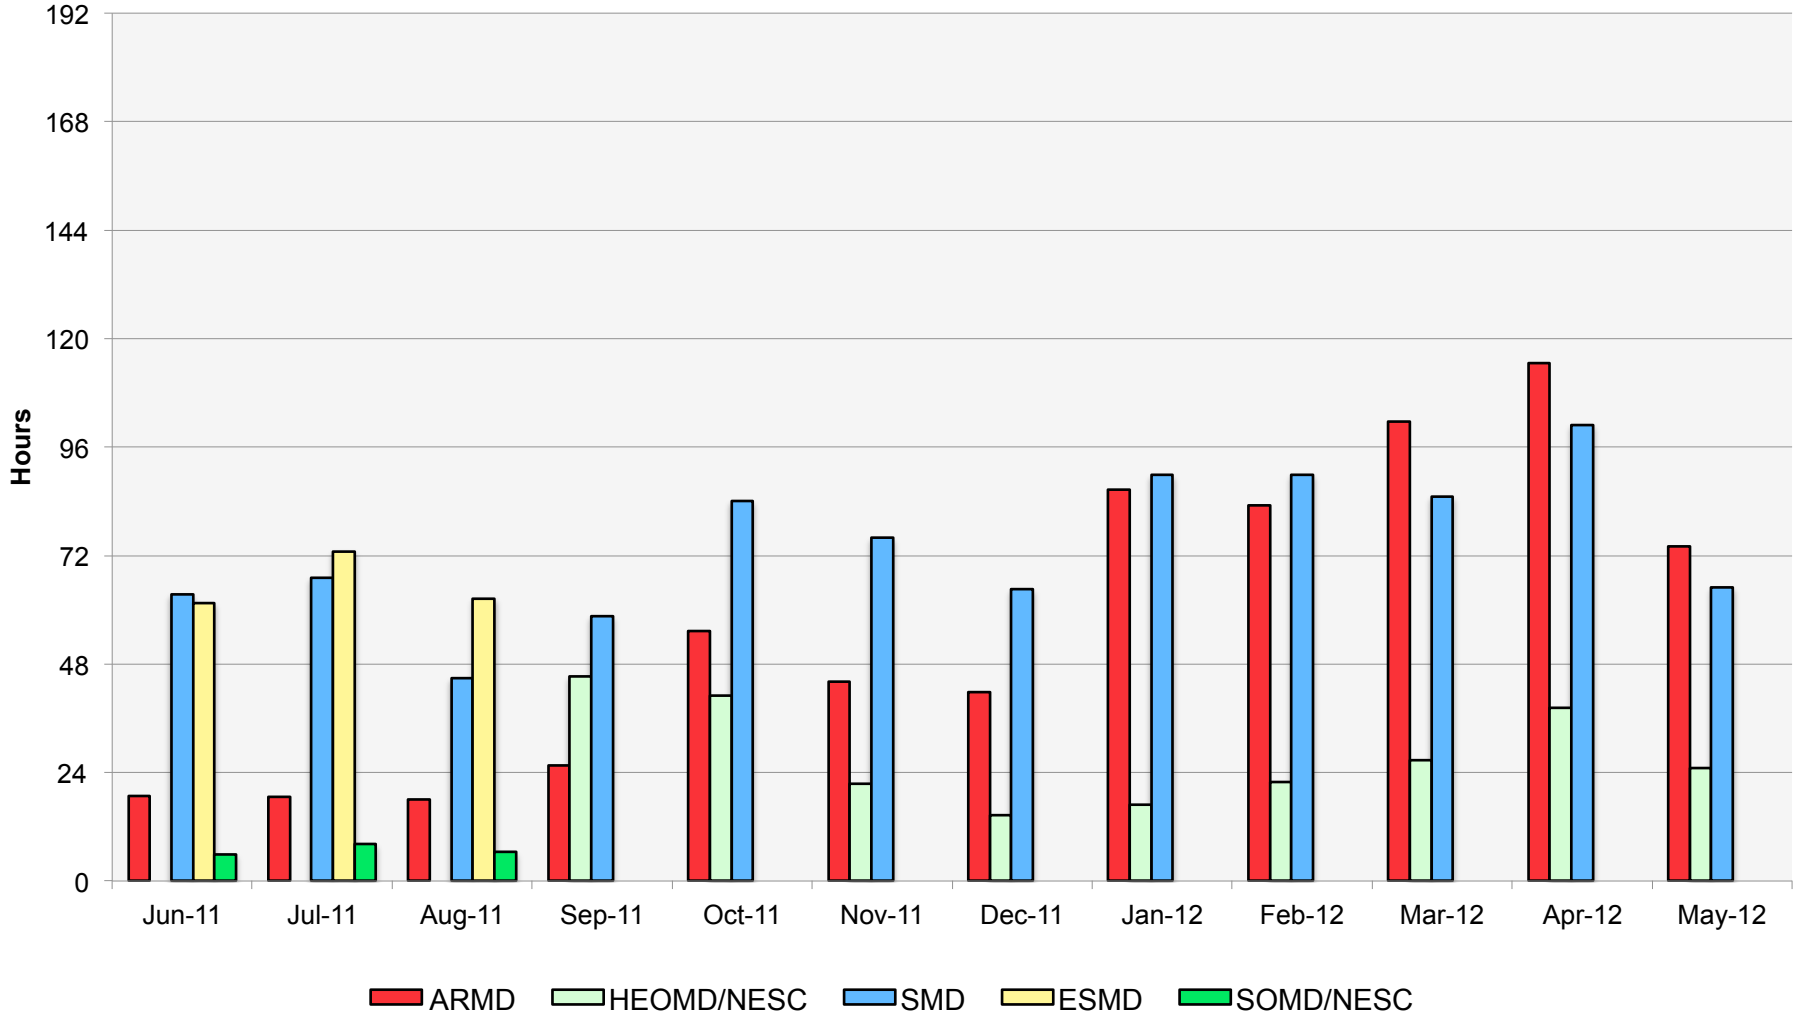

Pleiades:
Average Expansion Factor

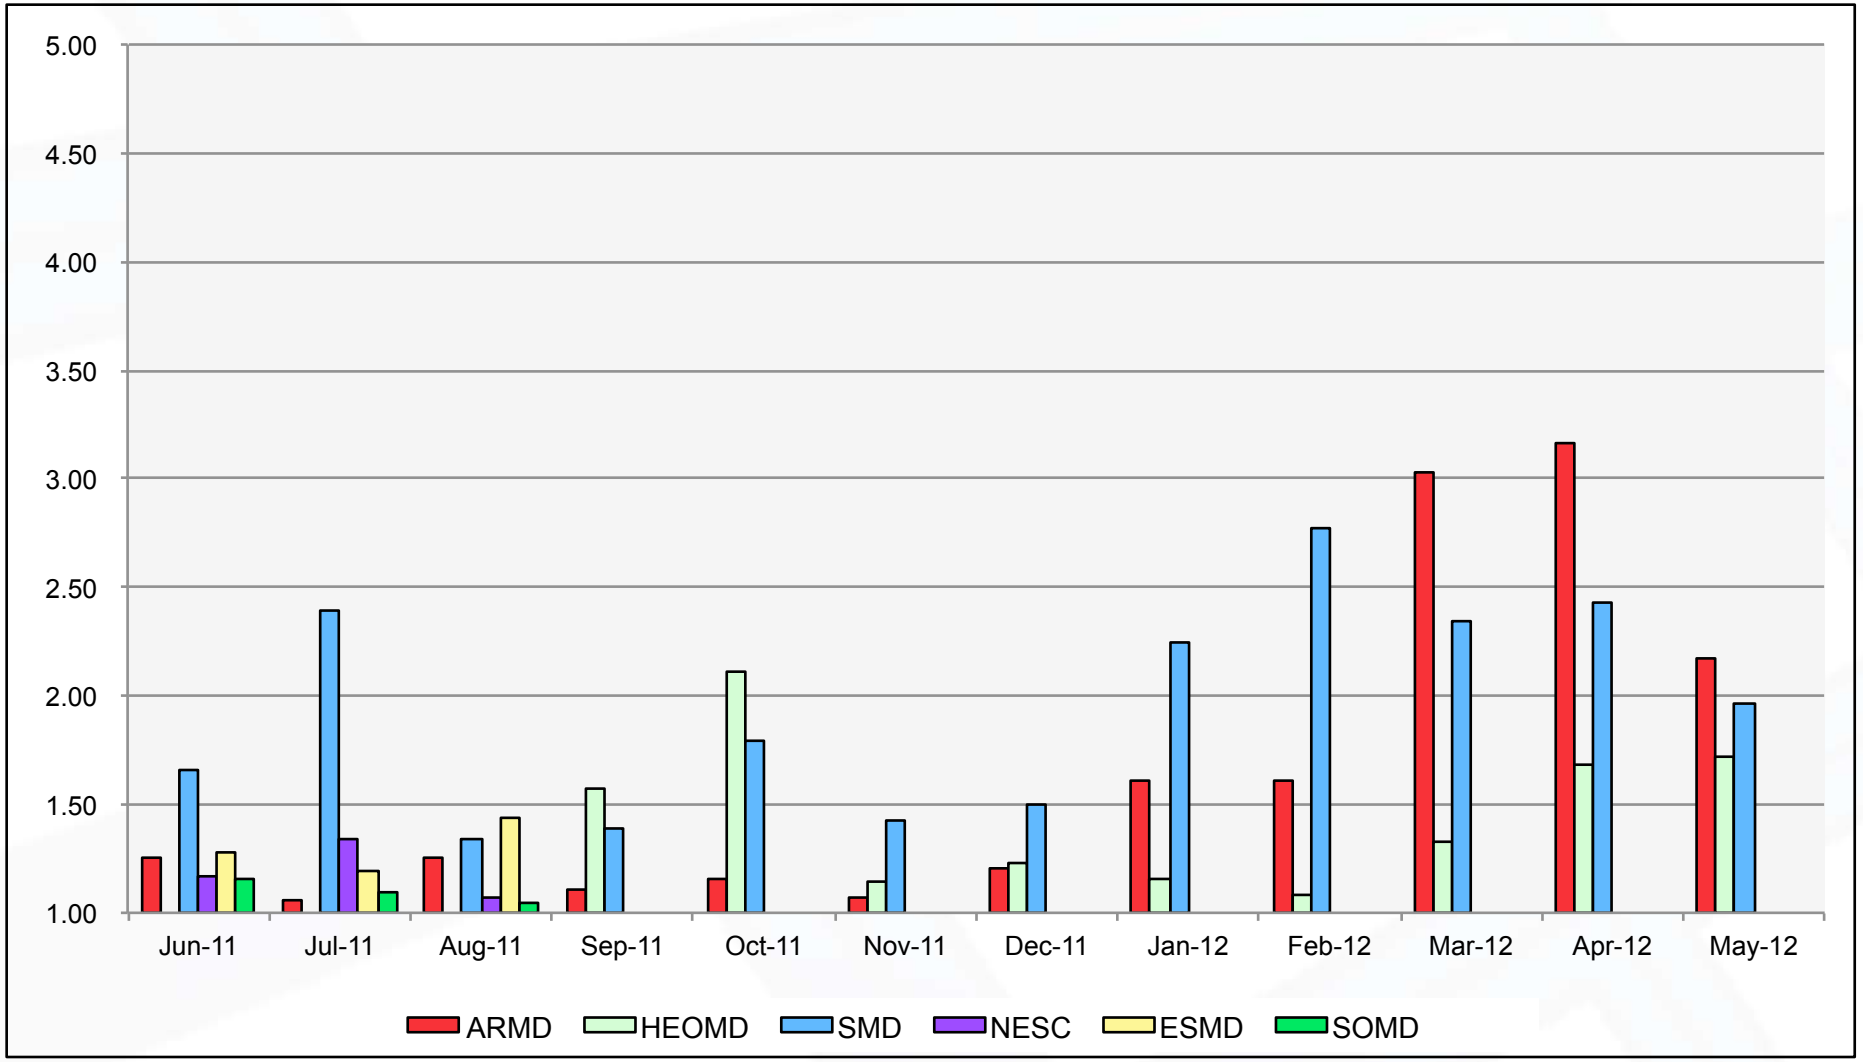

Columbia:
SBUs Reported, Normalized to 30-Day Month

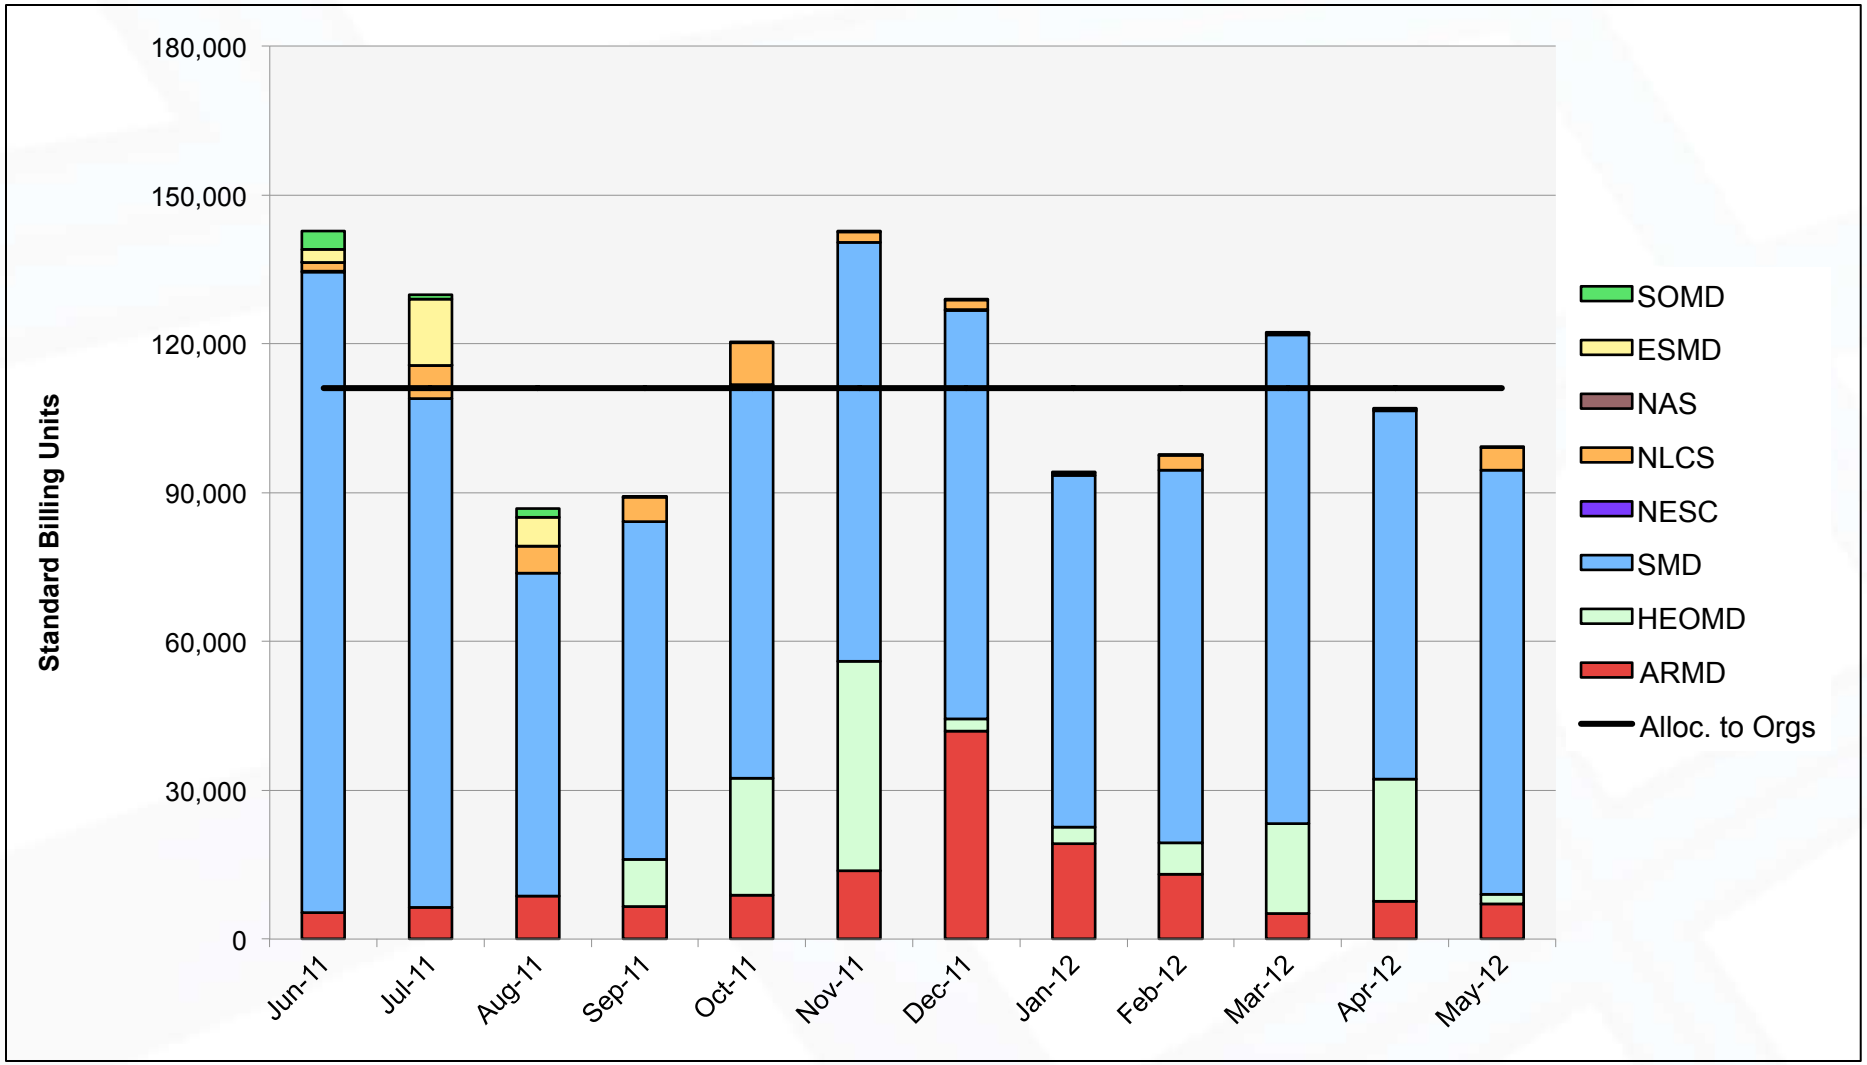

Columbia:
Monthly SBUs by Run Time

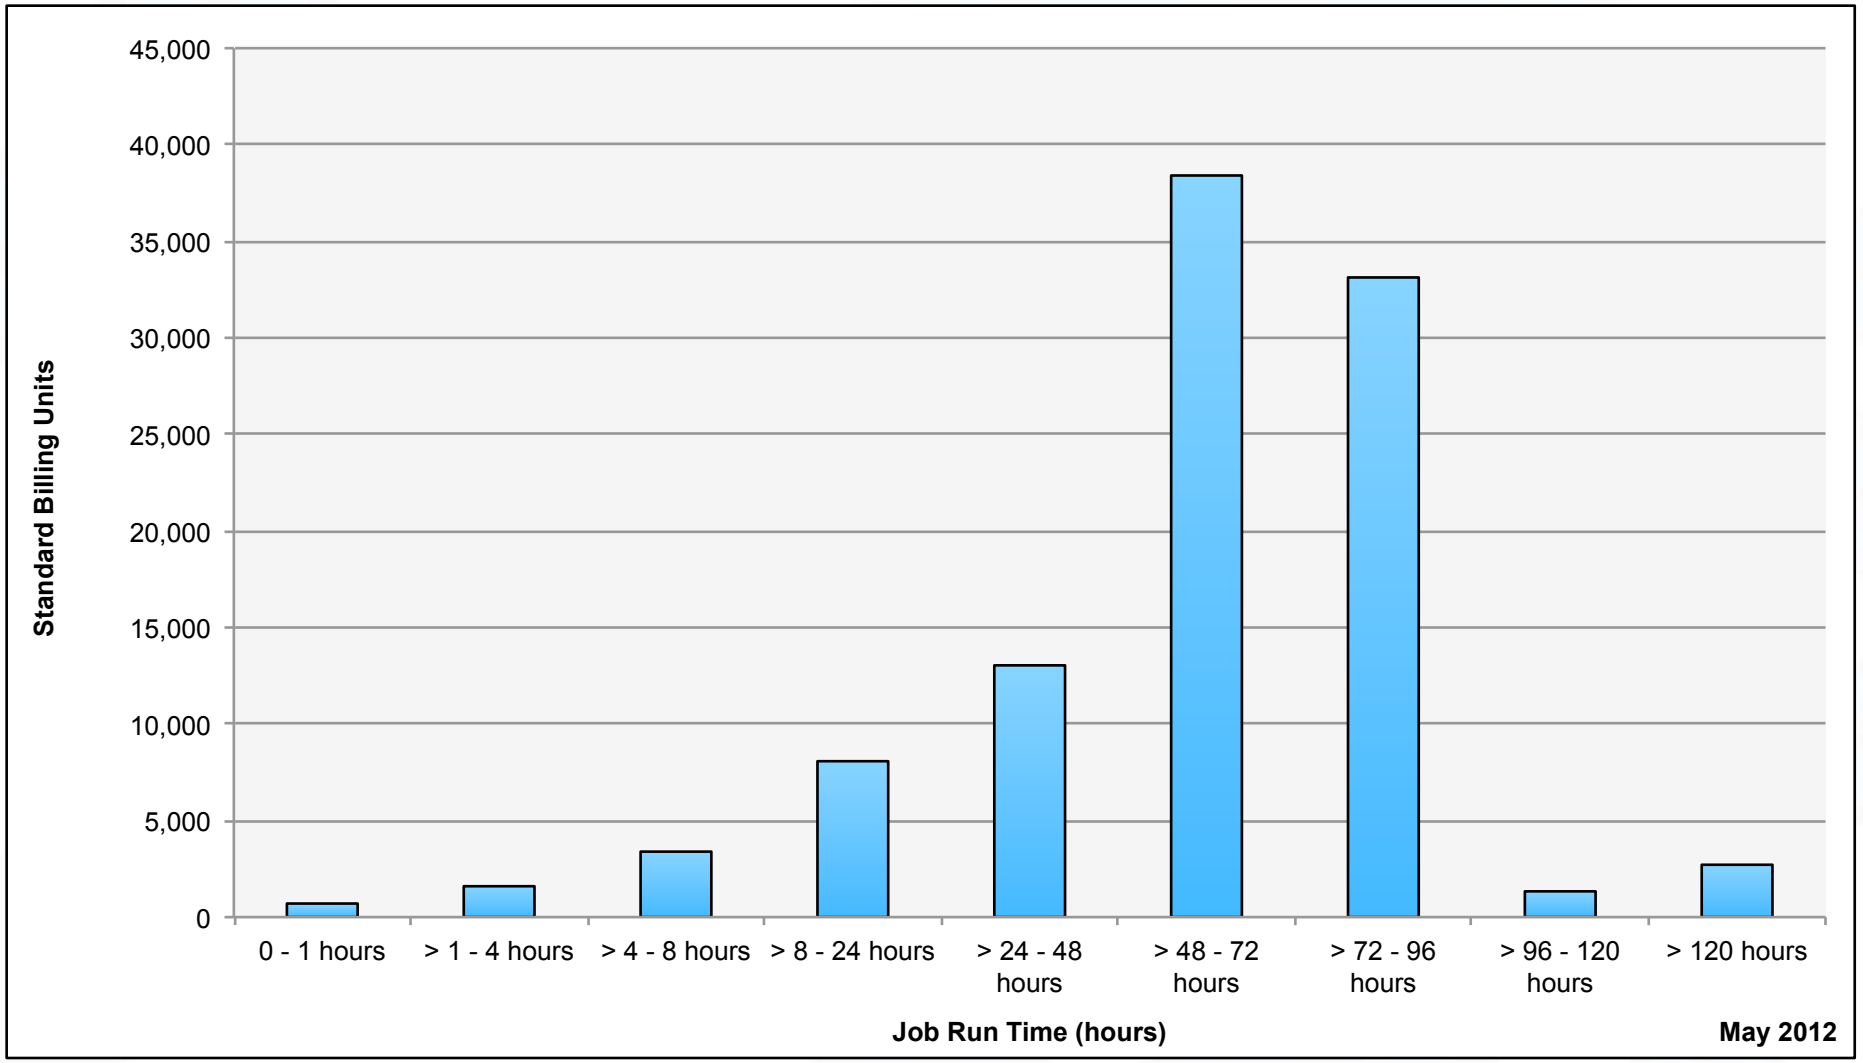

Columbia:
Monthly Utilization by Size and Mission

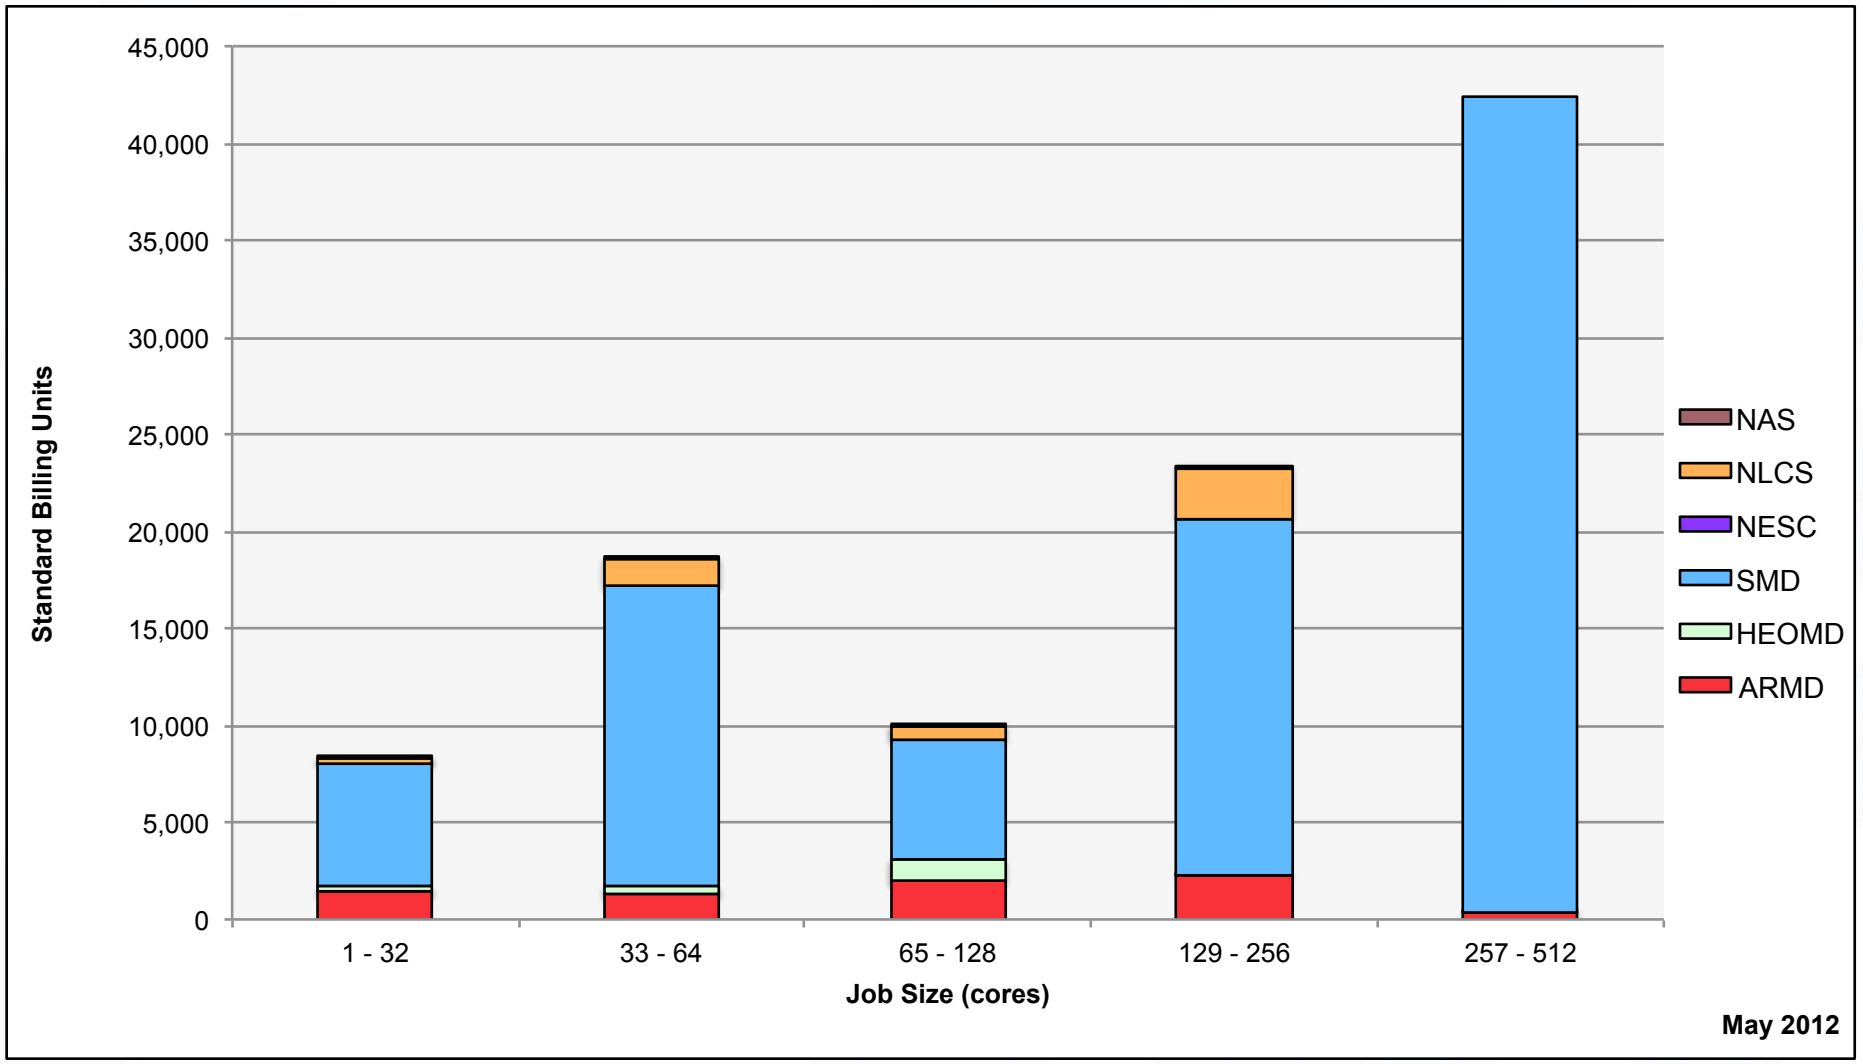

Columbia:
Monthly Utilization by Size and Length

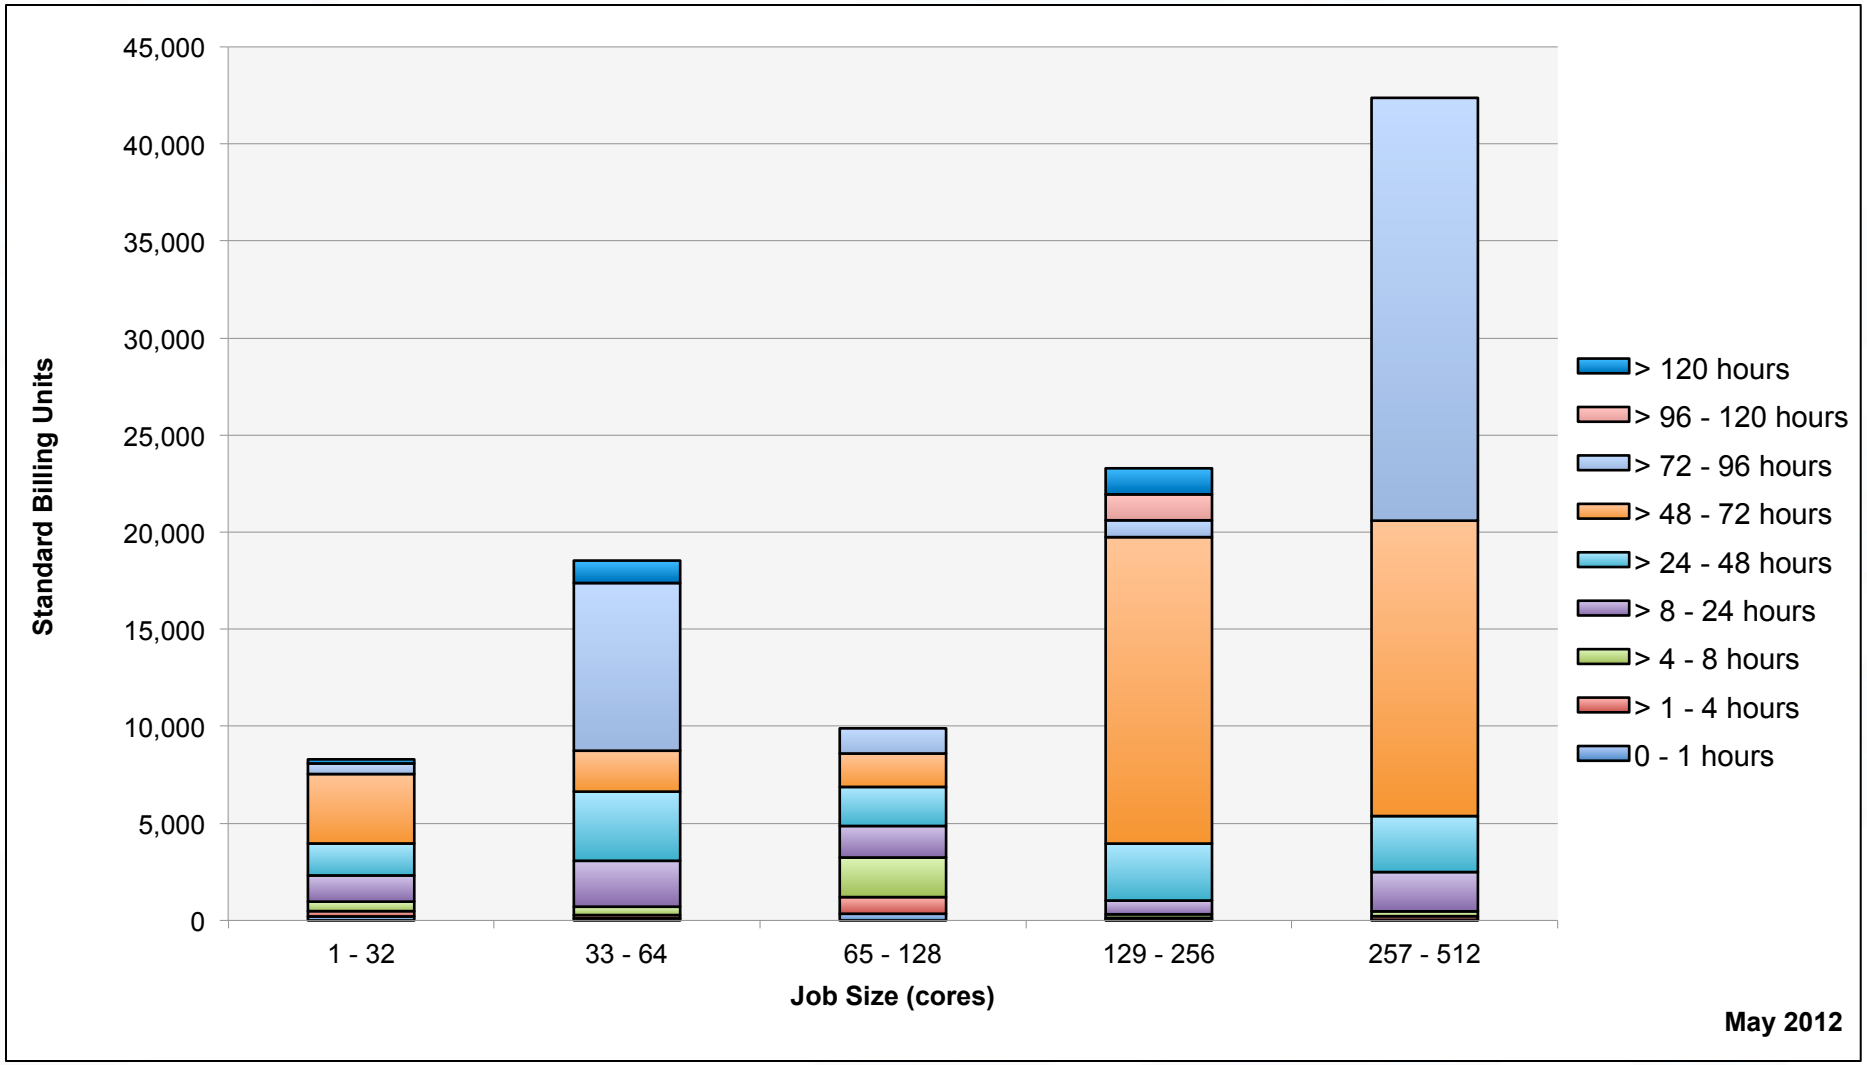

Columbia:
Average Time to Clear All Jobs

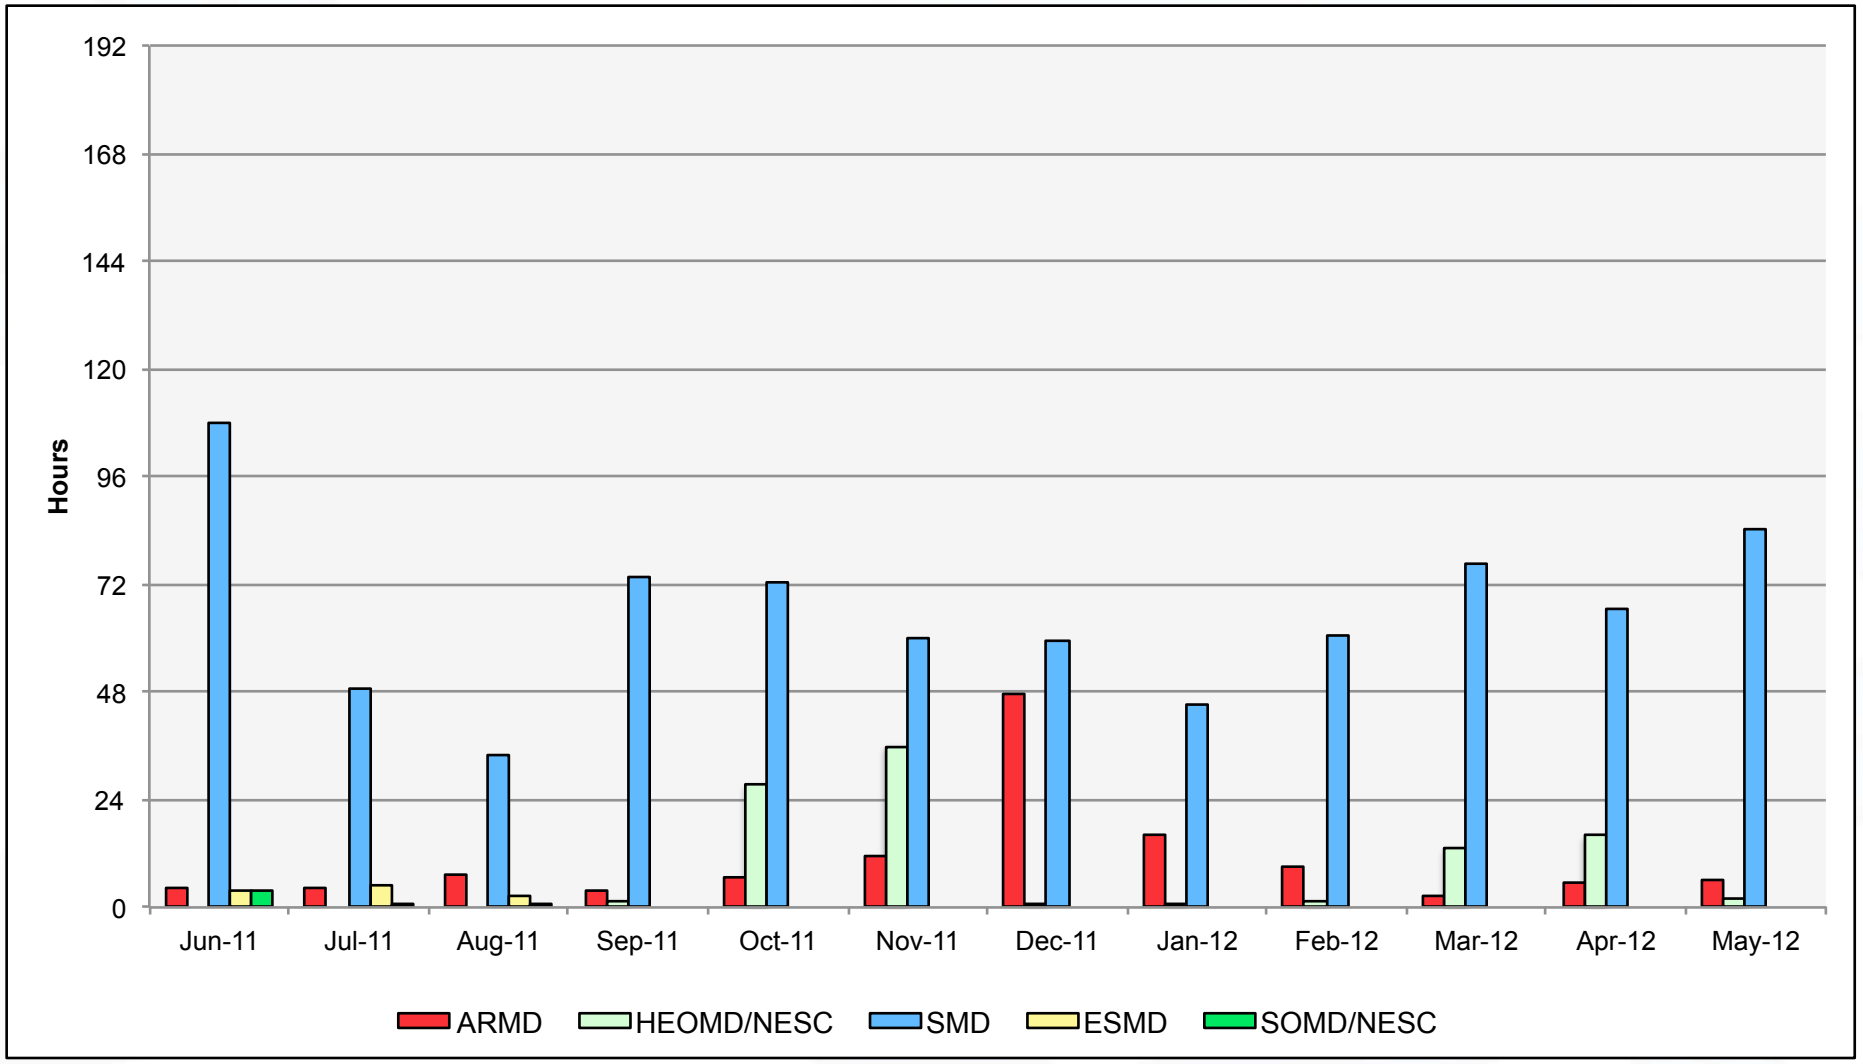

Columbia:
Average Expansion Factor