

National Aeronautics and Space Administration



# 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

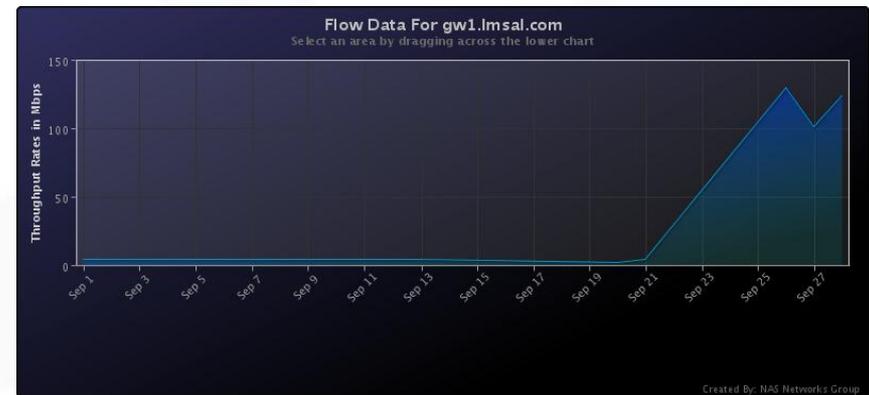
10 October 2011

# Network Team Increases User's Data Transfer Performance by 40x



- User Mark Cheung, Lockheed Martin, was getting a transfer rate of only 3-megabits per second (Mbps) when transferring data to HECC resources via the Secure Copy Protocol over the Internet—so slow that he resorted to driving to NASA Ames to bring data on disk whenever he needed to move it.
- HECC network engineers isolated the problem to an older version of Secure Shell running on the Lockheed system and small Transmission Control Protocol buffers on the remote host.
- HECC sent suggestions for system tuning to Lockheed engineers, and the user is now getting up to 130 Mbps transfer, a 40-fold improvement.
- Cheung sent an email stating, “We can't thank you enough for your counsel and advice. It has been invaluable. Still some things to do but this is a huge breakthrough for us. You obviously know your job very well.”

**Mission Impact: By evaluating all components of file transfers, HECC Network Engineers are able to identify and resolve issues affecting user productivity.**



**Figure:** Graph shows data transfer performance before and after optimizing the path between Lockheed Martin and the NAS.

**POC:** Chris Buchanan, [chris.buchanan@nasa.gov](mailto:chris.buchanan@nasa.gov), (650) 604-4308, NASA Advanced Supercomputing Division



# Two Westmere Racks Added to Pleiades for ARMD

- Two new SGI Altix ICE systems arrived on 9/26; these additional Pleiades racks will be used by the Fundamental Aeronautics Program (FAP) in the Aeronautics Research Mission Directorate (ARMD) (See slide 6).
- Each system is comprised of 64 nodes of 12-core Westmere processors, and together will annually provide ~850,000 node hours or standard billing units (SBUs) of processing.
- The addition provides a 12% increase to the ARMD/FAP computing capacity.
- Over the last year, FAP used approximately 97% of the ~6 million SBUs used by ARMD; usage is expected to continue increasing.
- The two new racks will be available to ARMD/FAP users the first week of October 2011.
- By doing this augmentation through HECC, FAP was able to take advantage of the pricing the HECC project gets and add a significant compute capability without having to pay for all of the ancillary support infrastructure.

**Mission Impact:** Increasing Pleiades' system capacity provides NASA mission directorates more resources for the accomplishment of their goals and objectives.



*Figure:* Two SGI Altix ICE 8400 systems, purchased to augment ARMD modeling and simulation throughput, were recently added to Pleiades.

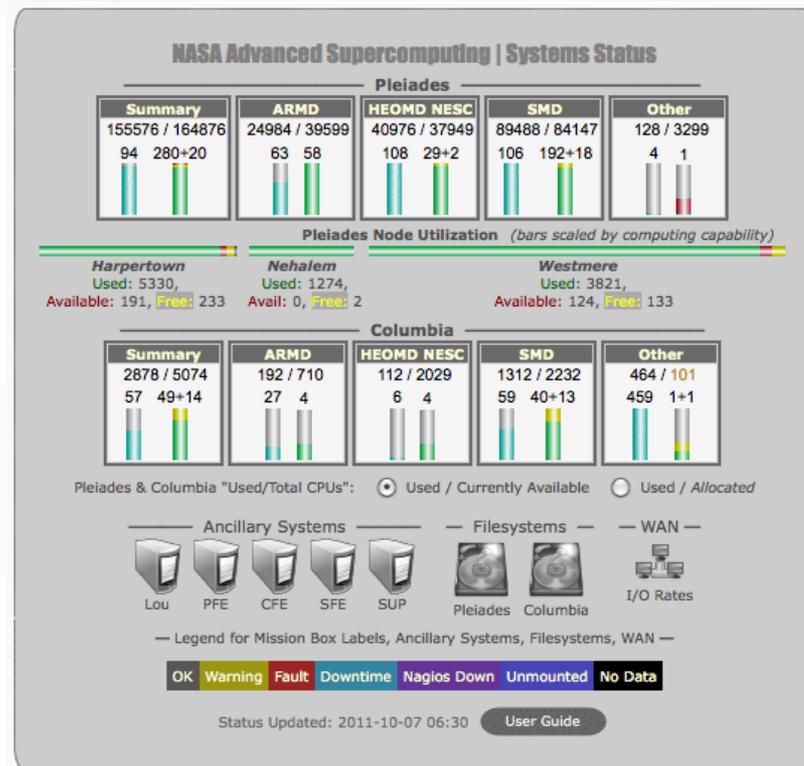
*POC:* Catherine Schulbach, [catherine.h.schullbach@nasa.gov](mailto:catherine.h.schullbach@nasa.gov),  
(650) 604-3180, NASA Advanced Supercomputing Division

# HECC Operations Modified to Support New HEO Mission Directorate



- With the creation of the Human Exploration and Operations Mission Directorate (HEOMD), the HECC program made several operational modifications to keep pace with the agency's organizational changes.
- Extending the concept of system "shares" to make sure that each mission directorate has access to the computational resources it is allocated, a new "share" was created for HEOMD.
- This new share replaced those previously allocated to the Exploration Systems Mission Directorate (ESMD) and the Space Operations Mission Directorate (SOMD).
- The very small NASA Engineering and Safety Center share that was combined with SOMD is now combined with HEOMD.
- System status displays for users and Control Room analysts, system monitoring scripts, and monthly usage reports were also updated to show the change from ESMD and SOMD to HEOMD.

**Mission Impact:** Modifying HECC operations to reflect NASA's organizational changes keeps status and usage information relevant for users and staff.



**Figure:** Example of the System Status page showing each mission directorate's share of HECC systems and the utilization of its resources.

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



# Installation of New Racks Brings Pleiades to 113,408 Cores

- HECC Systems experts and SGI engineers installed 2 new Westmere racks on Pleiades, with the following configuration enhancements:
- The new racks add another 1,536 Westmere cores (113,408 cores total), increasing the system's peak performance by an additional 18 teraflops – the system now has a 1.342 petaflops theoretical peak performance.
- The Systems team successfully added the racks using live integration techniques, which enabled the system to remain in use while the expansion was in progress; this saved over 7 million hours of computing time that would have been lost if the system had been brought down for the integration.

**Mission Impact:** This expansion of Pleiades provides increased computational capability to keep up with the requirements of the Aeronautics Research Mission Directorate.



*Figure:* Two additional SGI Westmere racks were installed in September 2011, adding another 18 teraflops of computational capability to Pleiades.

*POC: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408,  
NASA Advanced Supercomputing Division*

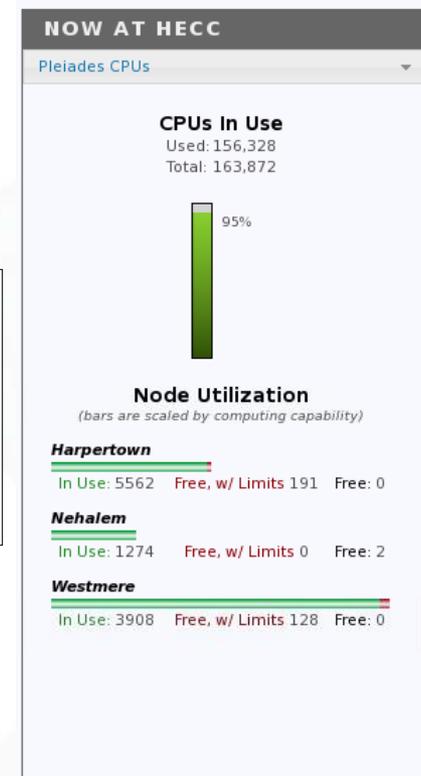


# PBSPro Configuration Change Improves Scheduling Performance

- The HECC Systems team changed the configuration of the job scheduler, PBSPro, to limit any given user to 300 total jobs at one time.
- This change, based on the observation that the vast majority of users run just a few dozen jobs at a time, helps to maintain scheduler performance and responsiveness by removing scenarios where a single user submits many hundreds of jobs at once (often by accident).
- For those users who have a legitimate need for many simultaneous jobs the HECC APPs team provides assistance with packing many subjobs into a smaller number of batch jobs, often improving the turnaround time and efficiency.

**Mission Impact:** Increases in scheduling performance assist HECC in keeping the computational capability as fully utilized as possible, both for batch jobs and interactive use.

*Figure:* Pleiades utilization in the range of 90+% is partly the result of a well-performing and efficient scheduler. Job requests are matched with free nodes in accordance with scheduling policy that ensures fairness for missions, groups, and users.



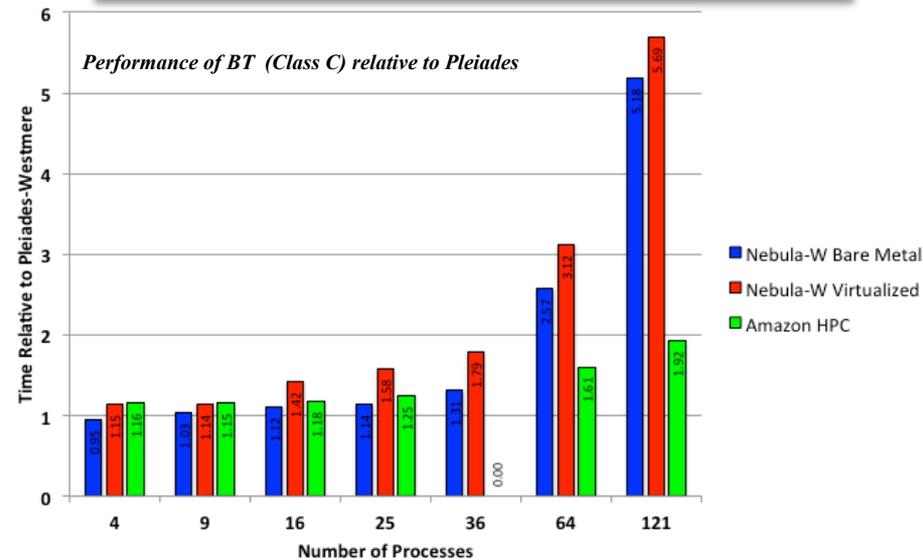
*POC: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408,  
NASA Advanced Supercomputing Division*

# APP Team Completes Evaluation of Nebula – NASA’s Cloud Computing Offering



- The HECC Application Performance & Productivity (APP) group completed an extensive evaluation of Nebula – NASA’s offering in cloud-computing – for HPC applications.
- The benchmark set consisted of the HPC Challenge Benchmarks, the NAS Parallel Benchmarks (NPB), and the IO Benchmark along with six full applications regularly used by NASA scientists and engineers.
- Initial testing lead to the discovery of some deficiencies in the environment. Turning on Virtio and Jumbo Frames by the Nebula team improved the application performance by 5-10x.
- The virtualization software layer plus 10GigE network on Nebula hurt application performance relative to Pleiades, which has a faster InfiniBand interconnect.
- Along with performance, the team members also assessed the ease of use of the Nebula environment.

**Mission Impact:** Performance assessment of Nebula relative to the HECC system Pleiades and the commercial Cloud offering from Amazon will allow SMD project managers to determine the suitability of Nebula for SMD applications.



**Figure:** Graph shows timing of BT benchmark one of the NPB suite on configurations of Nebula (bare metal and with virtualization software) and Amazon EC2 relative to Pleiades.

**POC:** Piyush Mehrotra, Piyush.Mehrotra@nasa.gov, (650) 604-5126,  
NASA Advanced Supercomputing Division

# Older Columbia Systems Made Available to External Organizations



- Recently, older systems that were part of the Columbia supercomputer were removed from service and made available to other government organizations.
- A total of 12 older Altix 3700 and Altix BX-2 systems (162 racks) were deinstalled in September.
- Current technology now in use within the HECC environment makes the older racks expendable but they can be useful at other sites.
- NASA property “excessing” procedures were followed to assure proper transfer from the HECC program, and included notification of availability to other agencies.
- With the removal of this hardware, HECC is now able to install current-generation computing and storage systems to meet increasing requirements for users.

**Mission Impact:** The removal of older computer systems makes room for new computing and storage enhancements to support NASA’s science and engineering users.



**Figure:** This picture, taken in September 2004, shows the HECC computer floor and Columbia. Most of the Columbia-era racks have now been replaced with the current technology available on the Pleiades supercomputer.

**POC:** John Parks, [john.w.parks@nasa.gov](mailto:john.w.parks@nasa.gov), (650) 604-4225,  
NASA Advanced Supercomputing Division

# HECC Facility Hosts Several Visitors and Tours in September 2011



- With the end of summer, HECC tours began to slow down. During September, staff hosted 7 scheduled tour groups, most of which received an overview of the HECC program, demonstrations of the hyperwall-2 visualization system, and tours of the computer room floor. This month's groups included:
  - Dr. Pedro Espina, Office of Science and Technology Policy, Executive Director of the National Science and Technology Council.
  - 2011 Software Design and Productivity (SDP) Summit attendees.
  - A group of technical members from the Department of Defense Joint Strike Fighter team.
  - Participants in the Agency Modeling and Simulation Forum.
  - Sasi Pillay, the new Agency CTO for IT.
  - 20 new interns and docents from the Ames Educations Associates Program.



*Figure:* Visualization team lead Chris Henze uses the hyperwall-2 to show results of a global ocean simulation to a group from the Agency Modeling and Simulation Forum.

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

# Presentations and Papers



- “A vortex in Saturn’s upper atmosphere as the driver of electromagnetic periodicities at Saturn: Magnetospheric and ionospheric responses,” M. G. Kivelson, X. Jia (2), and T. I. Gombosi, EPSC-DPS Joint Meeting 2011, October 2–7, Nantes, France\*  
<http://meetingorganizer.copernicus.org/EPSC-DPS2011/EPSC-DPS2011-58.pdf>
- “The Magnetosphere of Saturn,” T.I. Gombosi, K.C. Hansen, X. Jia and M.G. Kivelson, EPSC-DPS Joint Meeting 2011, October 2–7, Nantes, France\*  
<http://meetingorganizer.copernicus.org/EPSC-DPS2011/EPSC-DPS2011-25.pdf>
- "An Overview of NITRD Activities Related to High Performance Computing", Bryan A. Biegel, 42nd HPC User Forum, Sep. 6-8, 2011, San Diego, CA.
- "Software Challenges in High Performance Computing", Bryan A. Biegel, Software Design and Productivity Cross Agency and National Needs Summit, Sep. 20-23, 2011, NASA Ames Research Center.

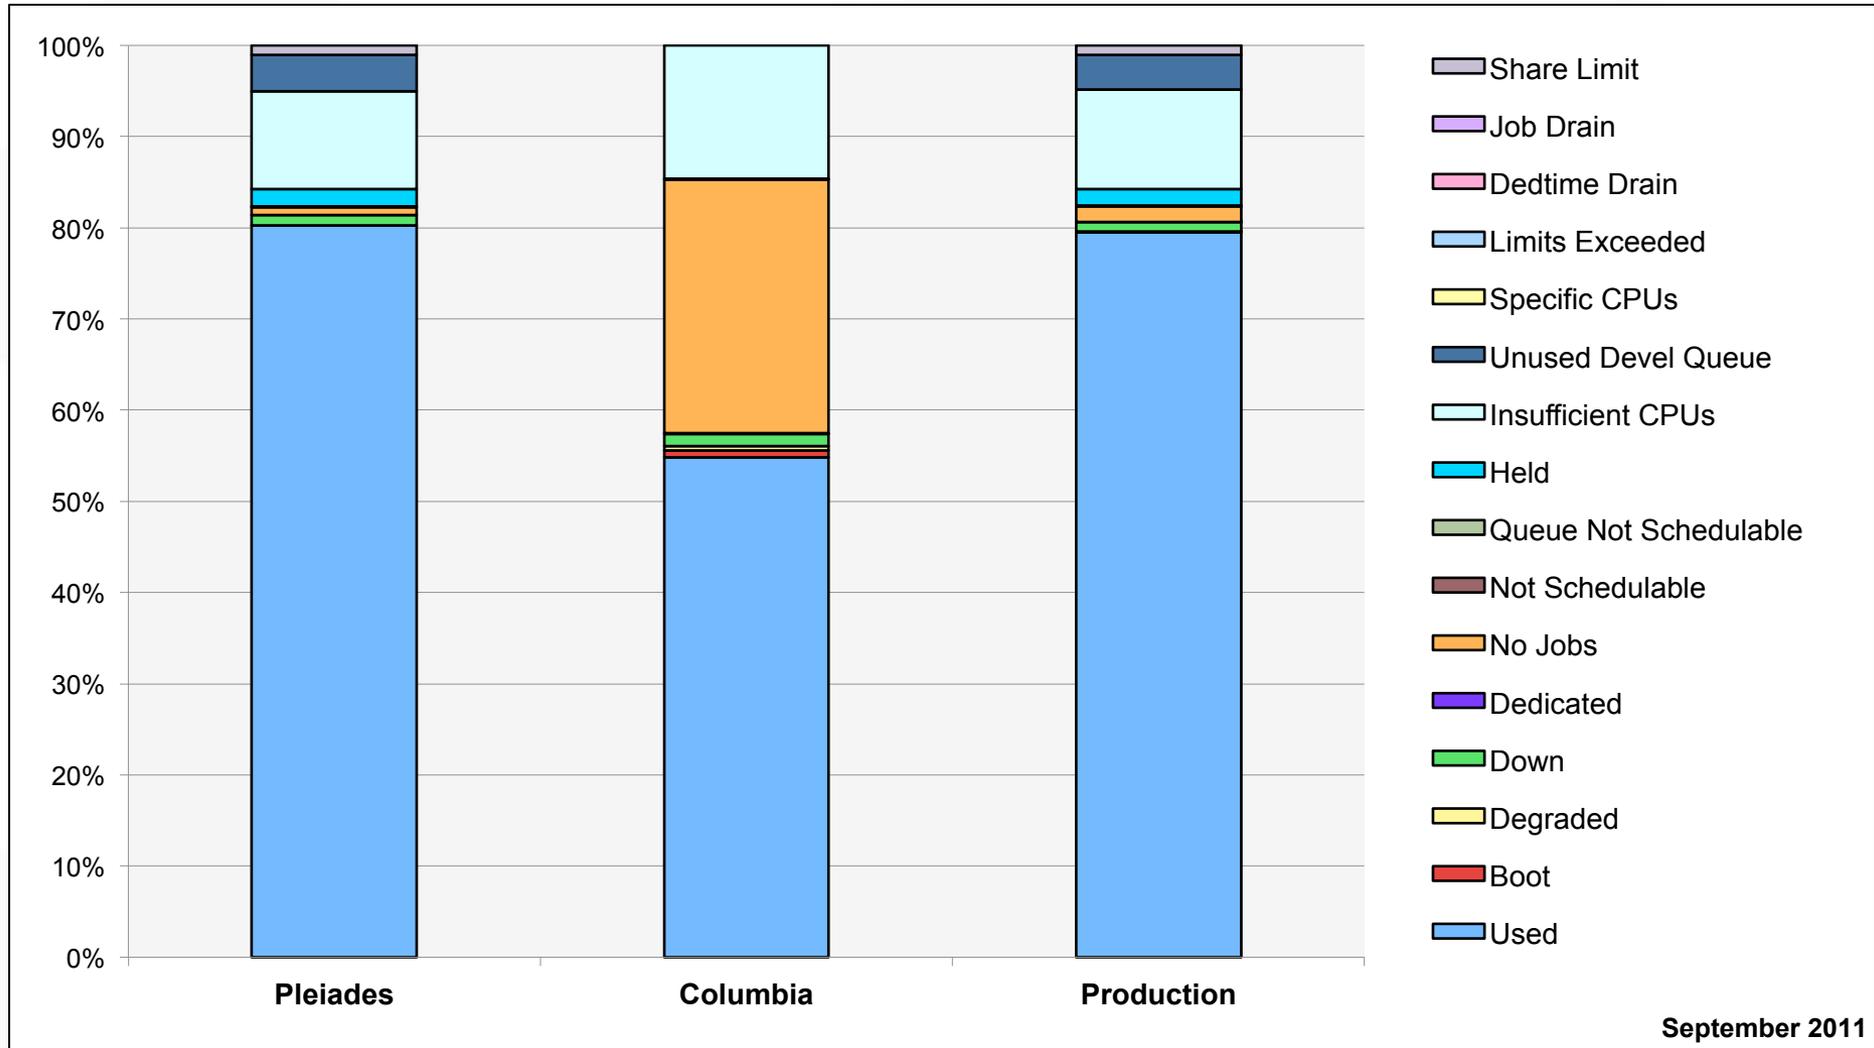
\* HECC provided supercomputing resources and services in support of this work

# News and Events



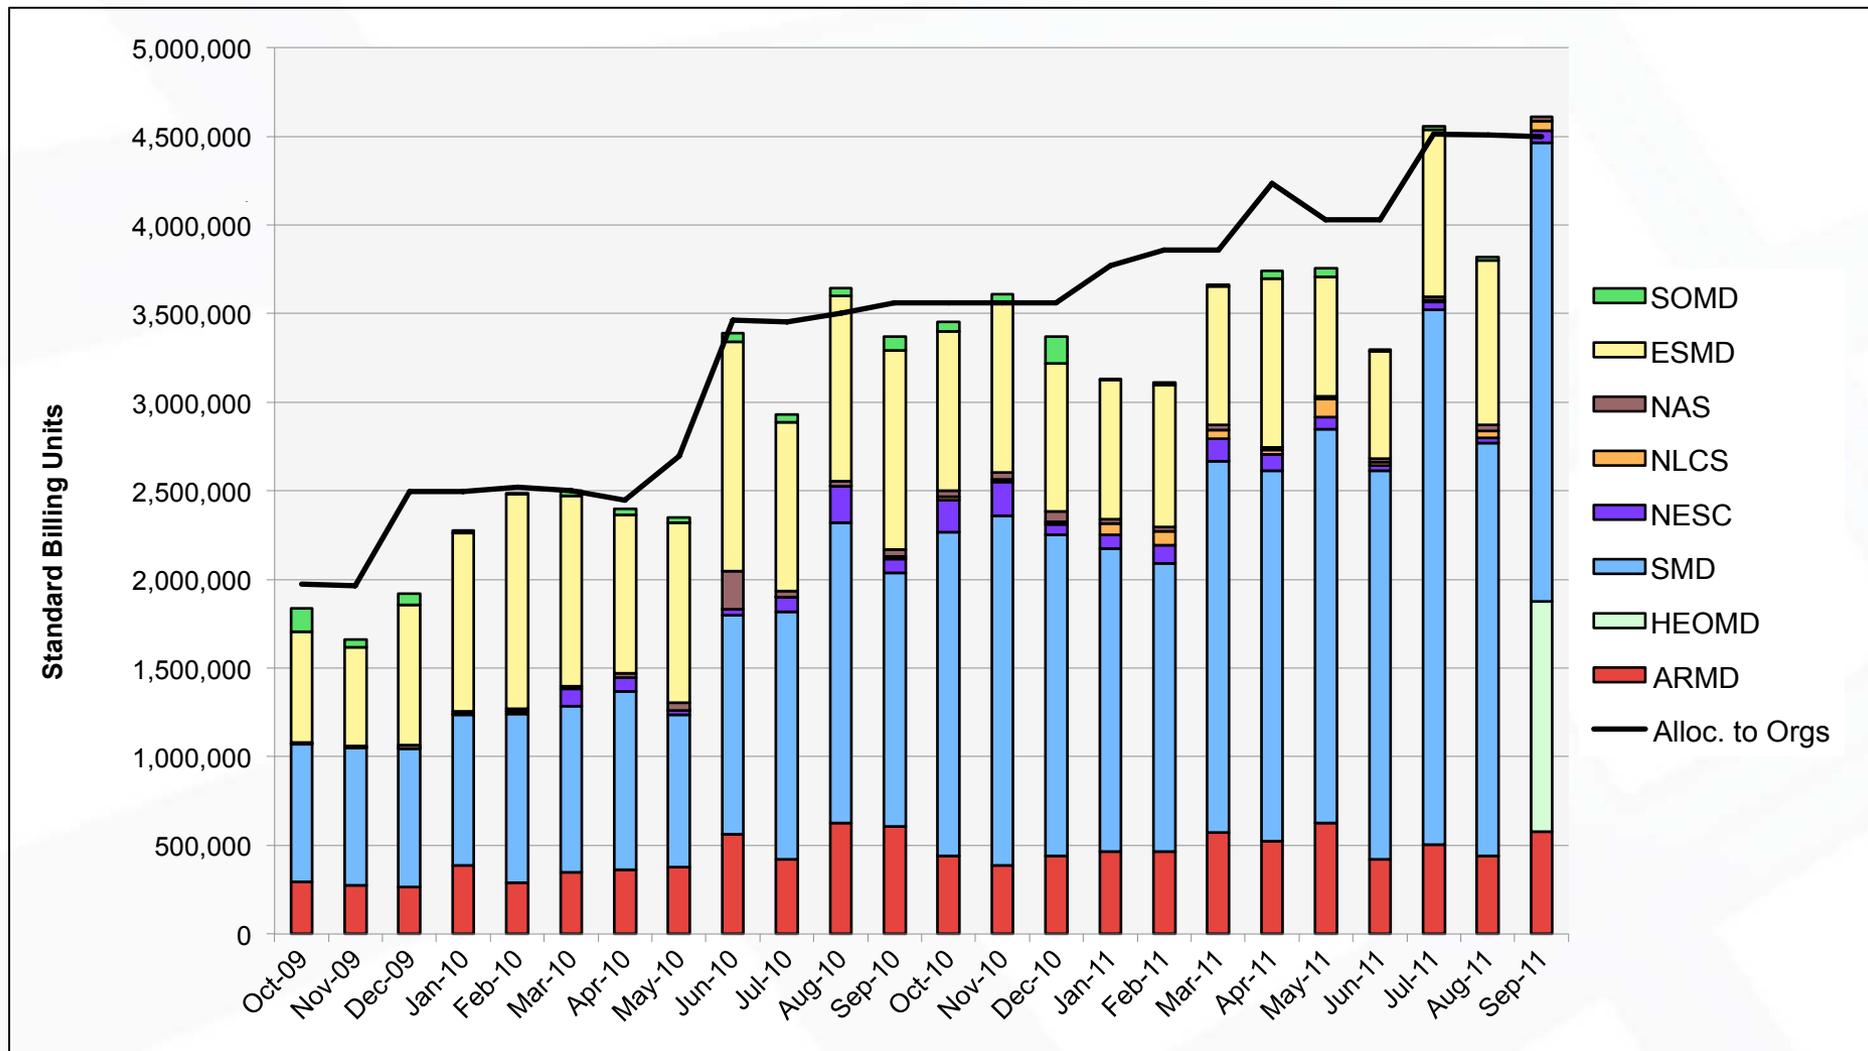
- **NASA Supercomputer Enables Largest Cosmological Simulations** , *press release, Ames Research Center*, September 29, 2011 – Scientists have generated the largest and most realistic cosmological simulations of the evolving universe to-date, thanks to NASA's powerful Pleiades supercomputer. Story received wide coverage by many media outlets, including HPCwire, SupercomputingOnline, SpaceRef, Information Week, Network World, International Business Times, San Francisco Chronicle, San Jose Mercury News.  
<http://www.nasa.gov/centers/ames/news/releases/2011/11-77AR.html>
- **42nd HPC User Forum:** NAS management and staff provided broad support for the HPC User Forum, the most recent meeting of which was held in San Diego, CA, Sep. 6-8, 2011. Rupak Biswas, HECC Project Manager, serves on the Steering Committee, helped to develop the agenda, engaged a full slate of leading speakers, and chaired a session. HECC Systems Lead Bob Ciotti served on a panel on the future of Lustre. Other NAS staff attended the meeting.  
<http://www.hpcuserforum.com/registration/sandiego/sandiegoagenda.pdf>
- **National Workshop on Software Challenges:** On Sep. 20-23, Ames hosted the "Software Design and Productivity Cross Agency and National Needs Summit", which focused on challenges and potential Federal government-led solutions in software supporting National missions and priorities. NAS provided substantial support to this interagency and multi-sector (government, academia, industry) workshop. NAS Deputy Chief Bryan Biegel organized a half-day session on "Software Challenges in HPC", and served on a plenary panel providing agency viewpoints on software challenges. HECC APP Group Lead Piyush Mehrotra co-chaired one of the three breakout groups focusing on the greatest software challenges facing Federal agencies in the short, medium and long term.  
<http://www.nitrd.gov/subcommittee/sdp/materials/September2011SDPSummit.pdf>

# NAS Utilization

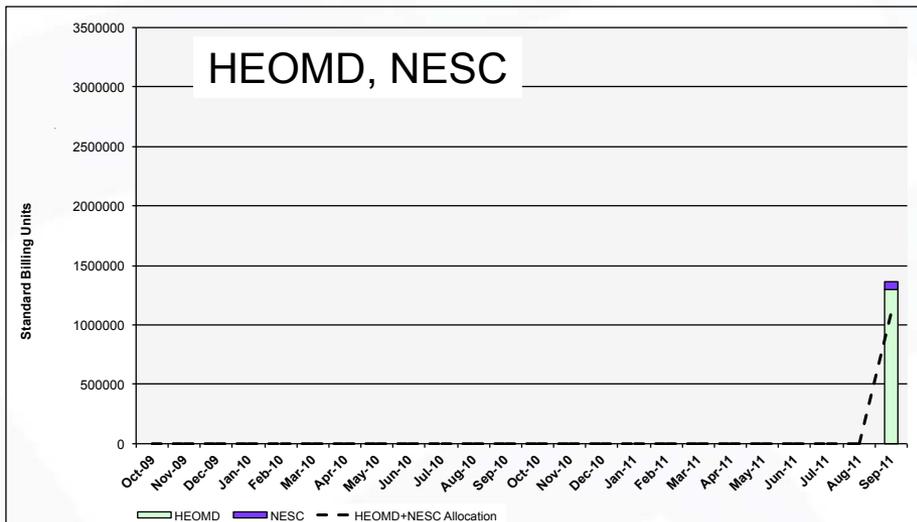
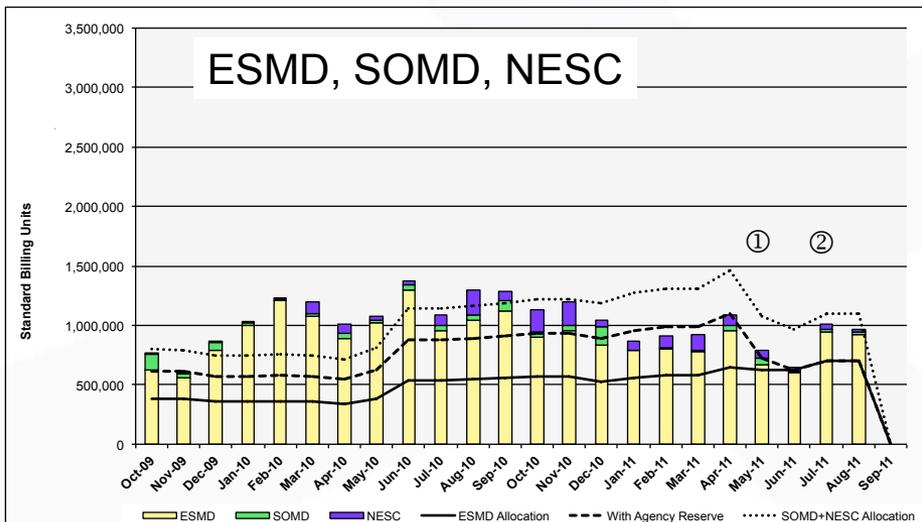
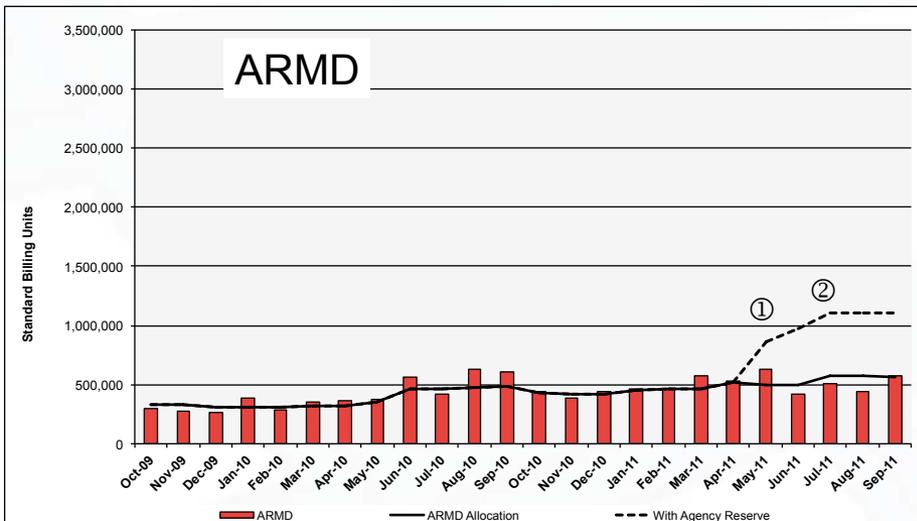
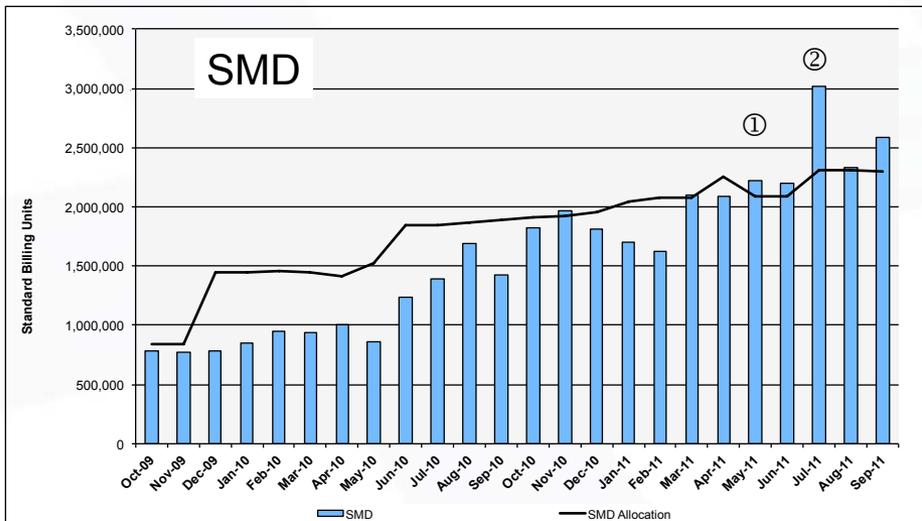


September 2011

# NAS Utilization Normalized to 30-Day Month

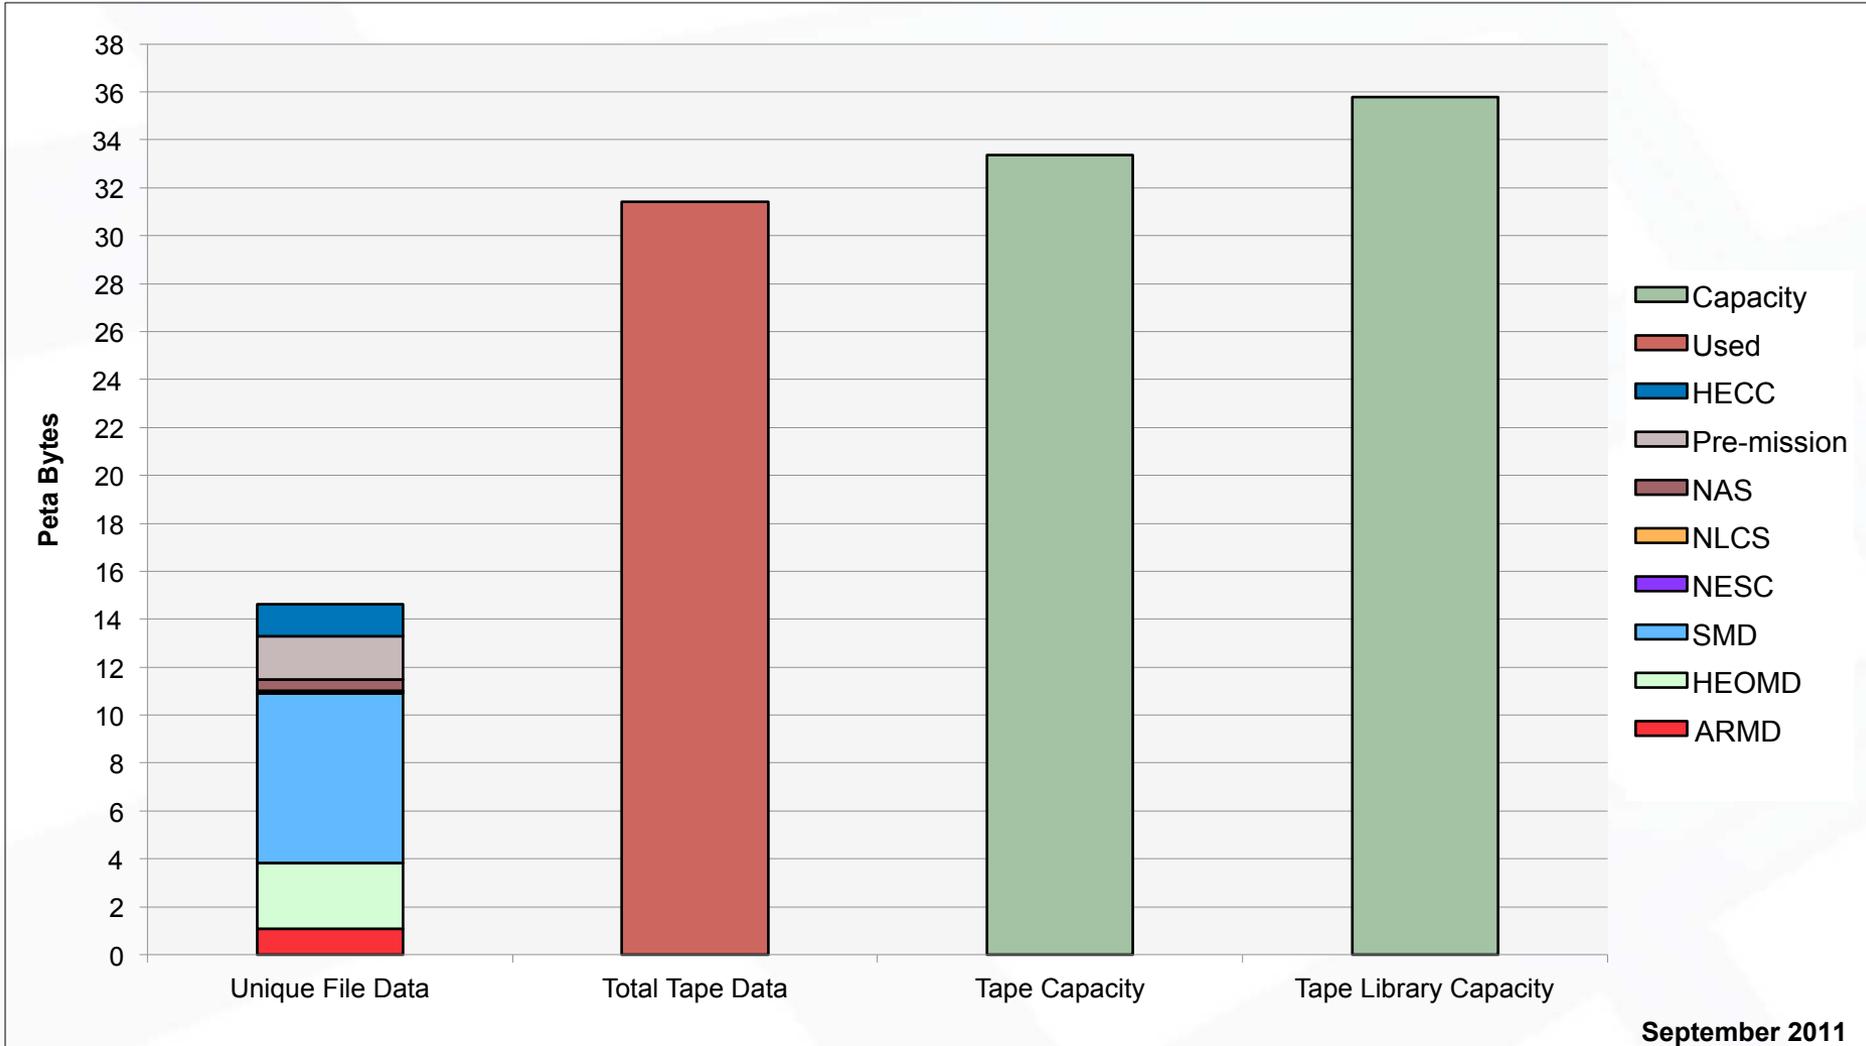


# NAS Utilization Normalized to 30-Day Month



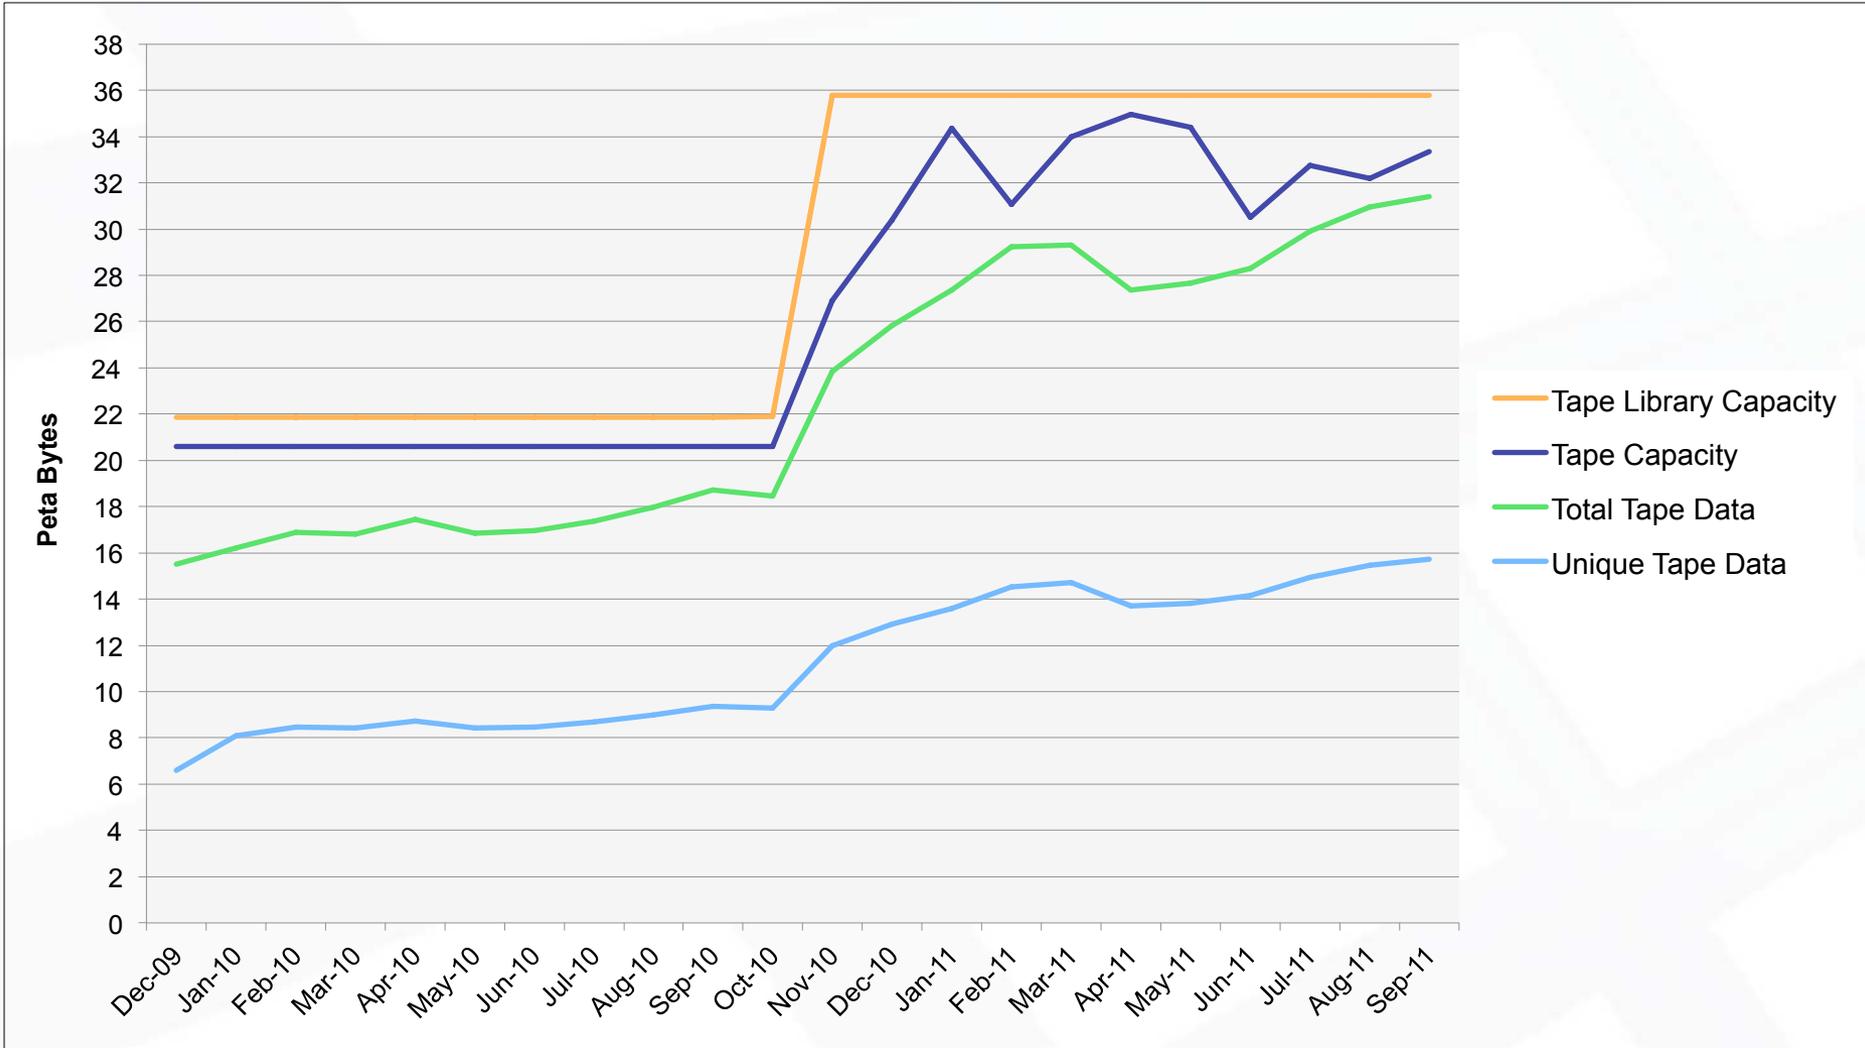
① Allocation to orgs. decreased from 80% to 75%, Agency reserve shifted to ARMD    ② 14 Westmere racks added

# Tape Archive Status

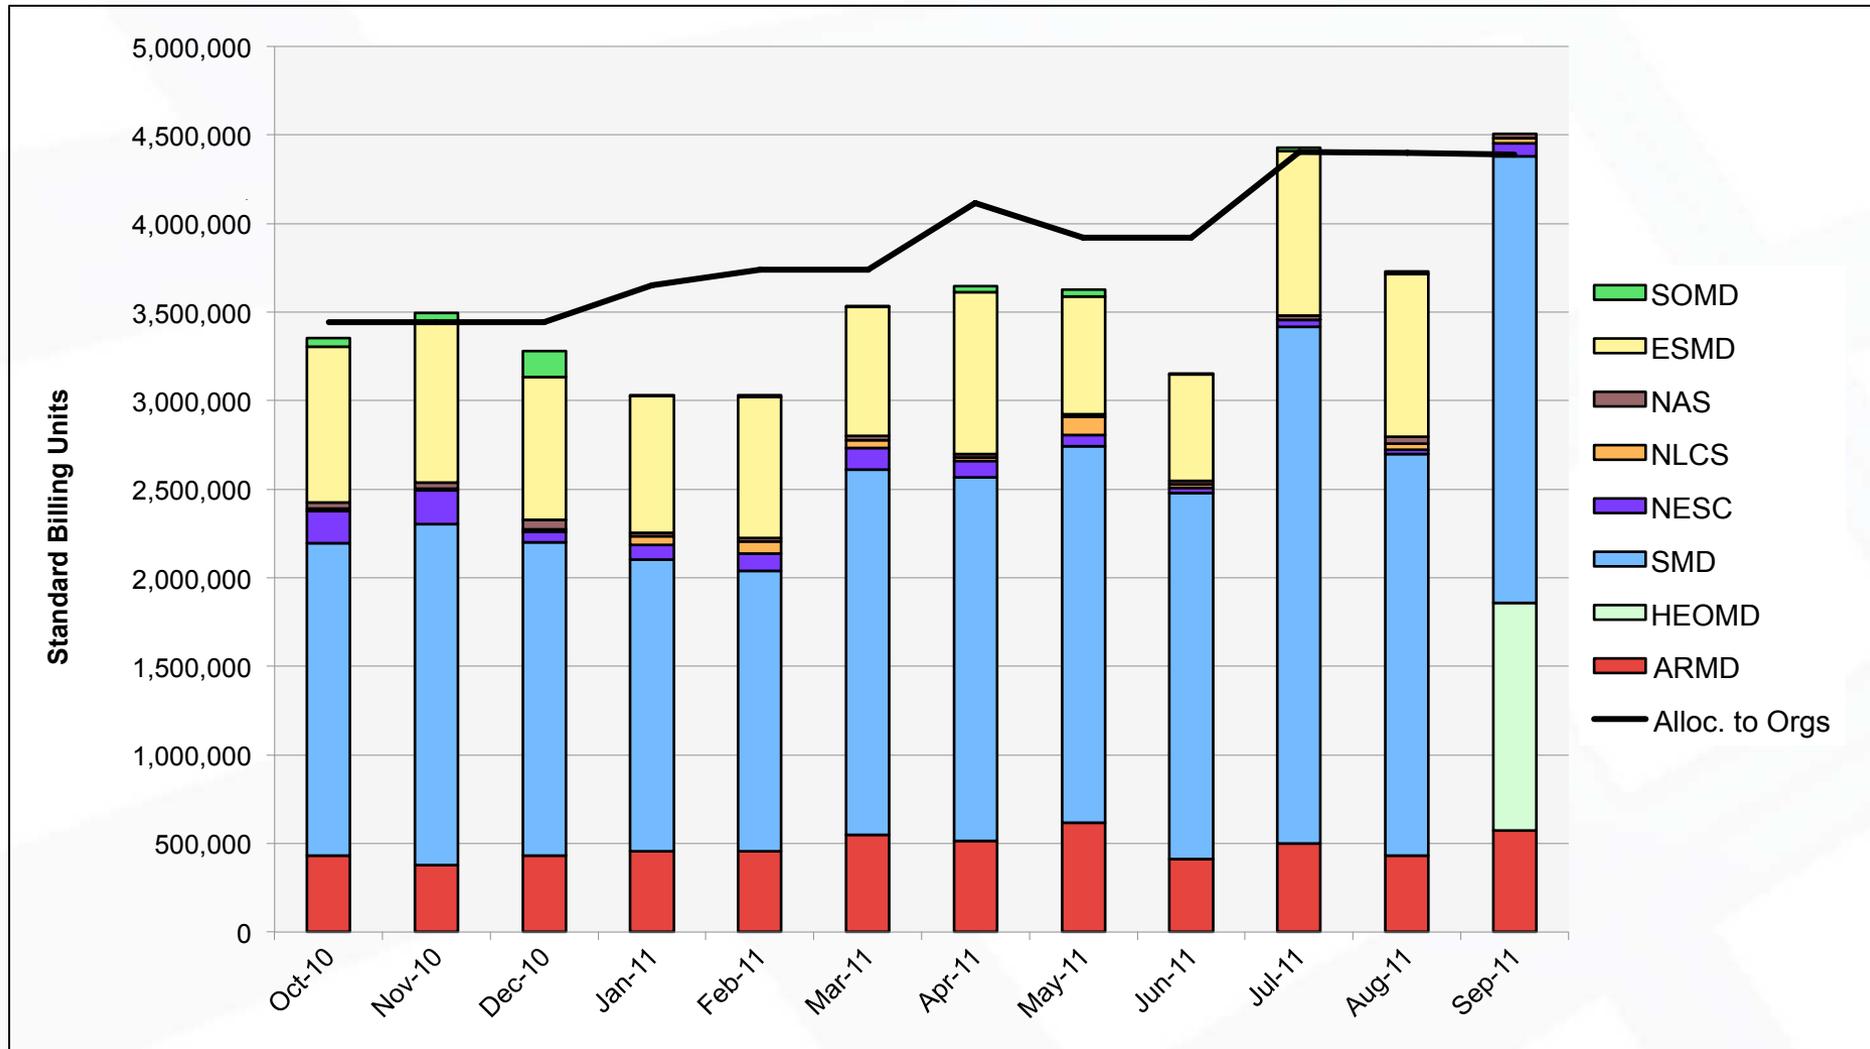


September 2011

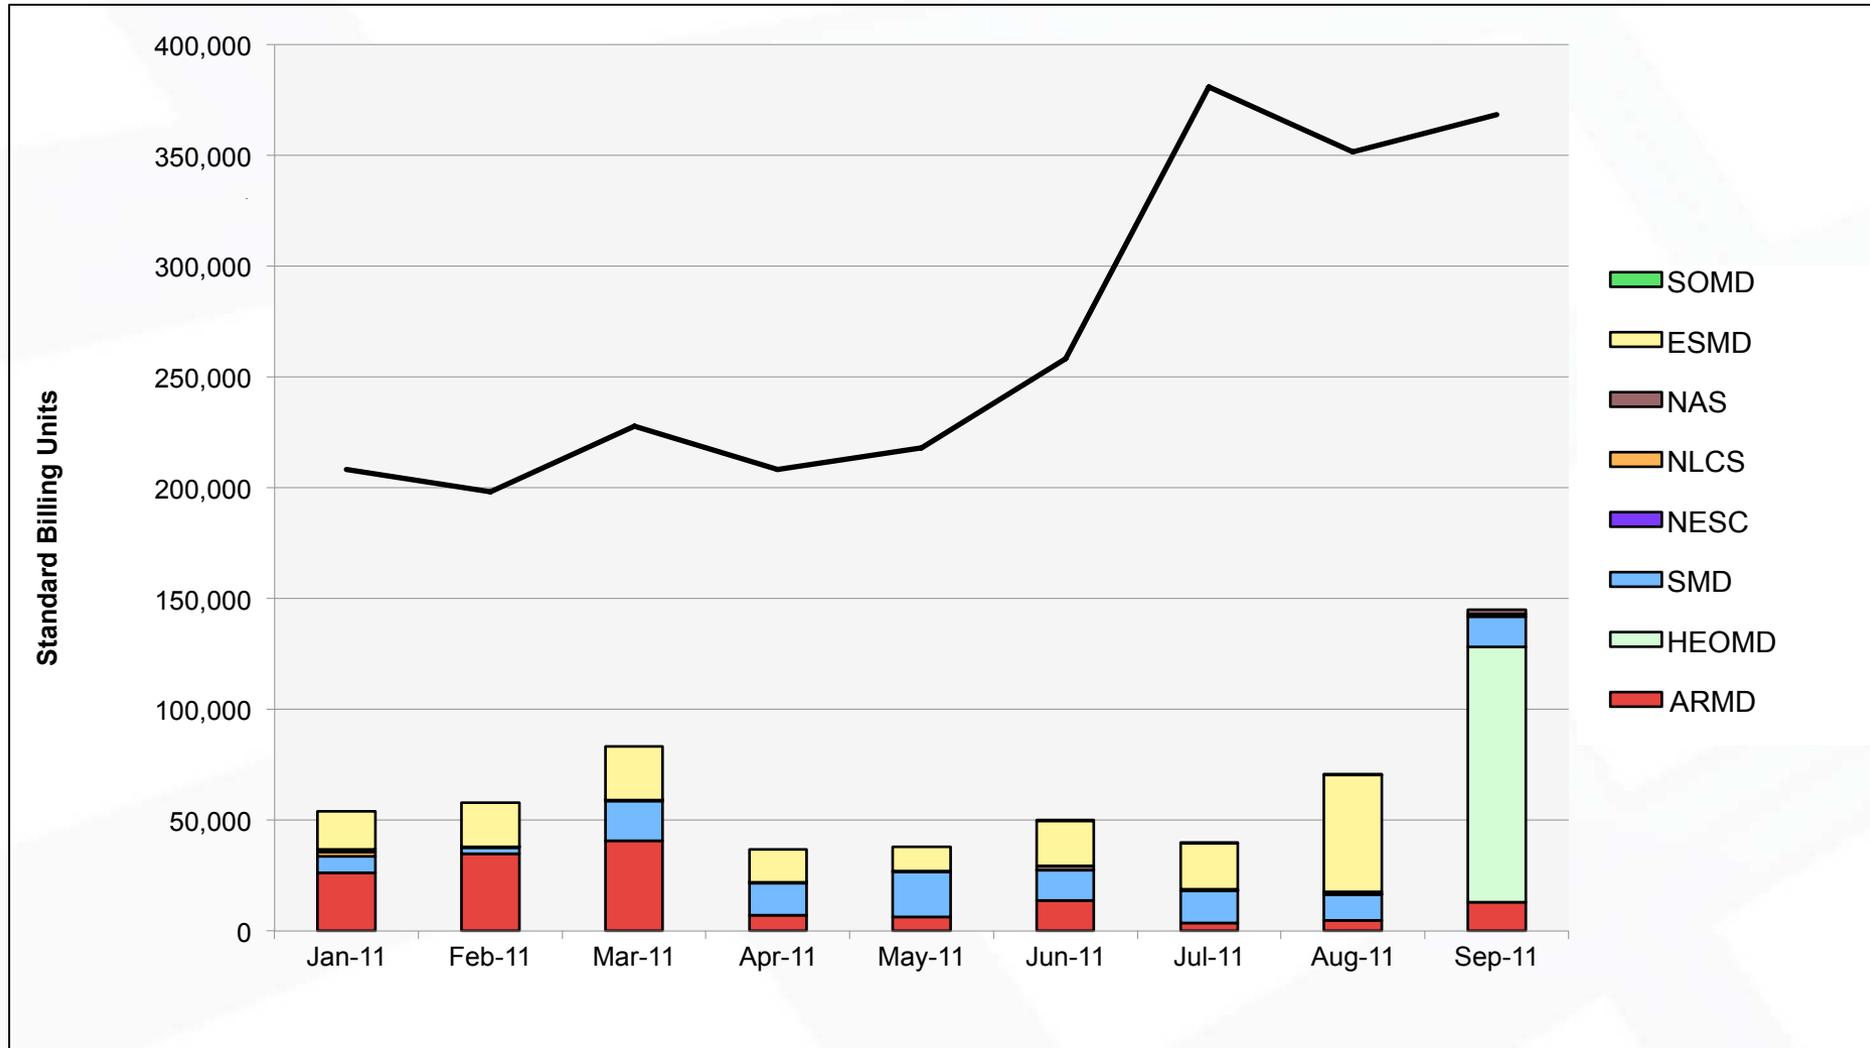
# Tape Archive Status



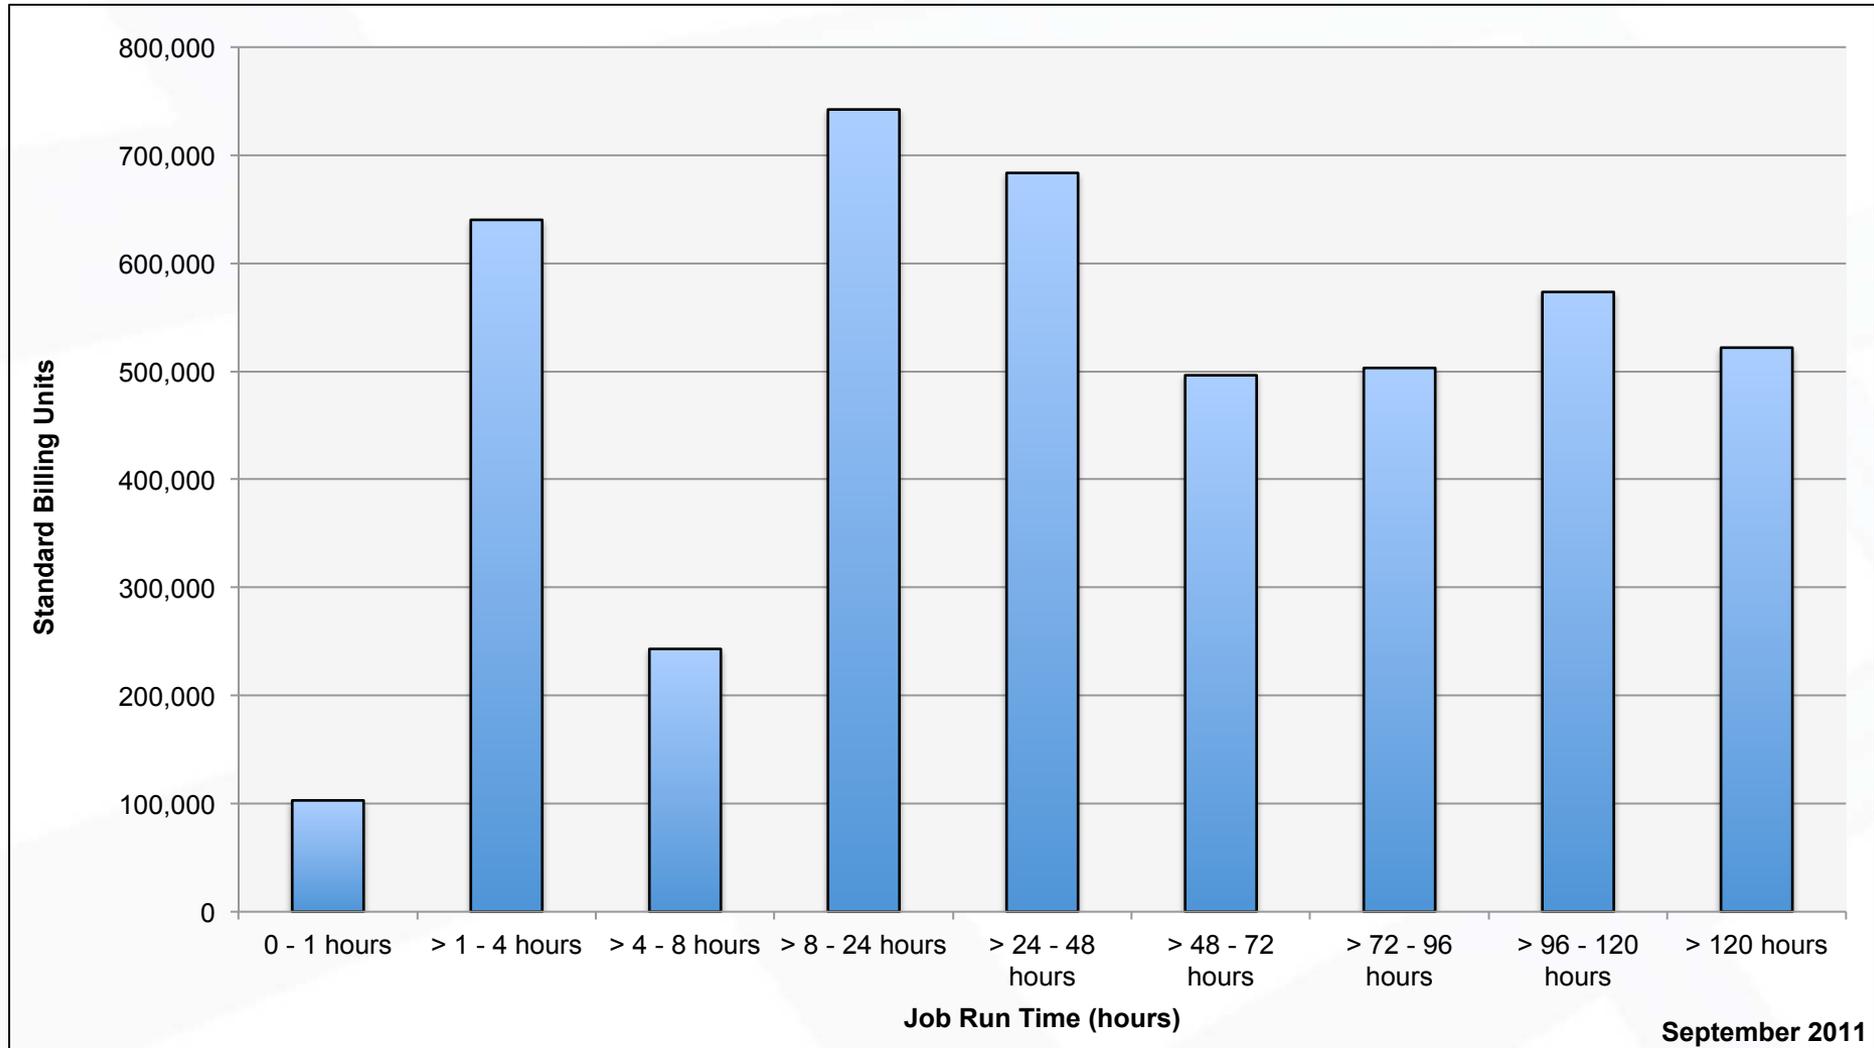
# Pleiades: SBUs Reported, Normalized to 30-Day Month



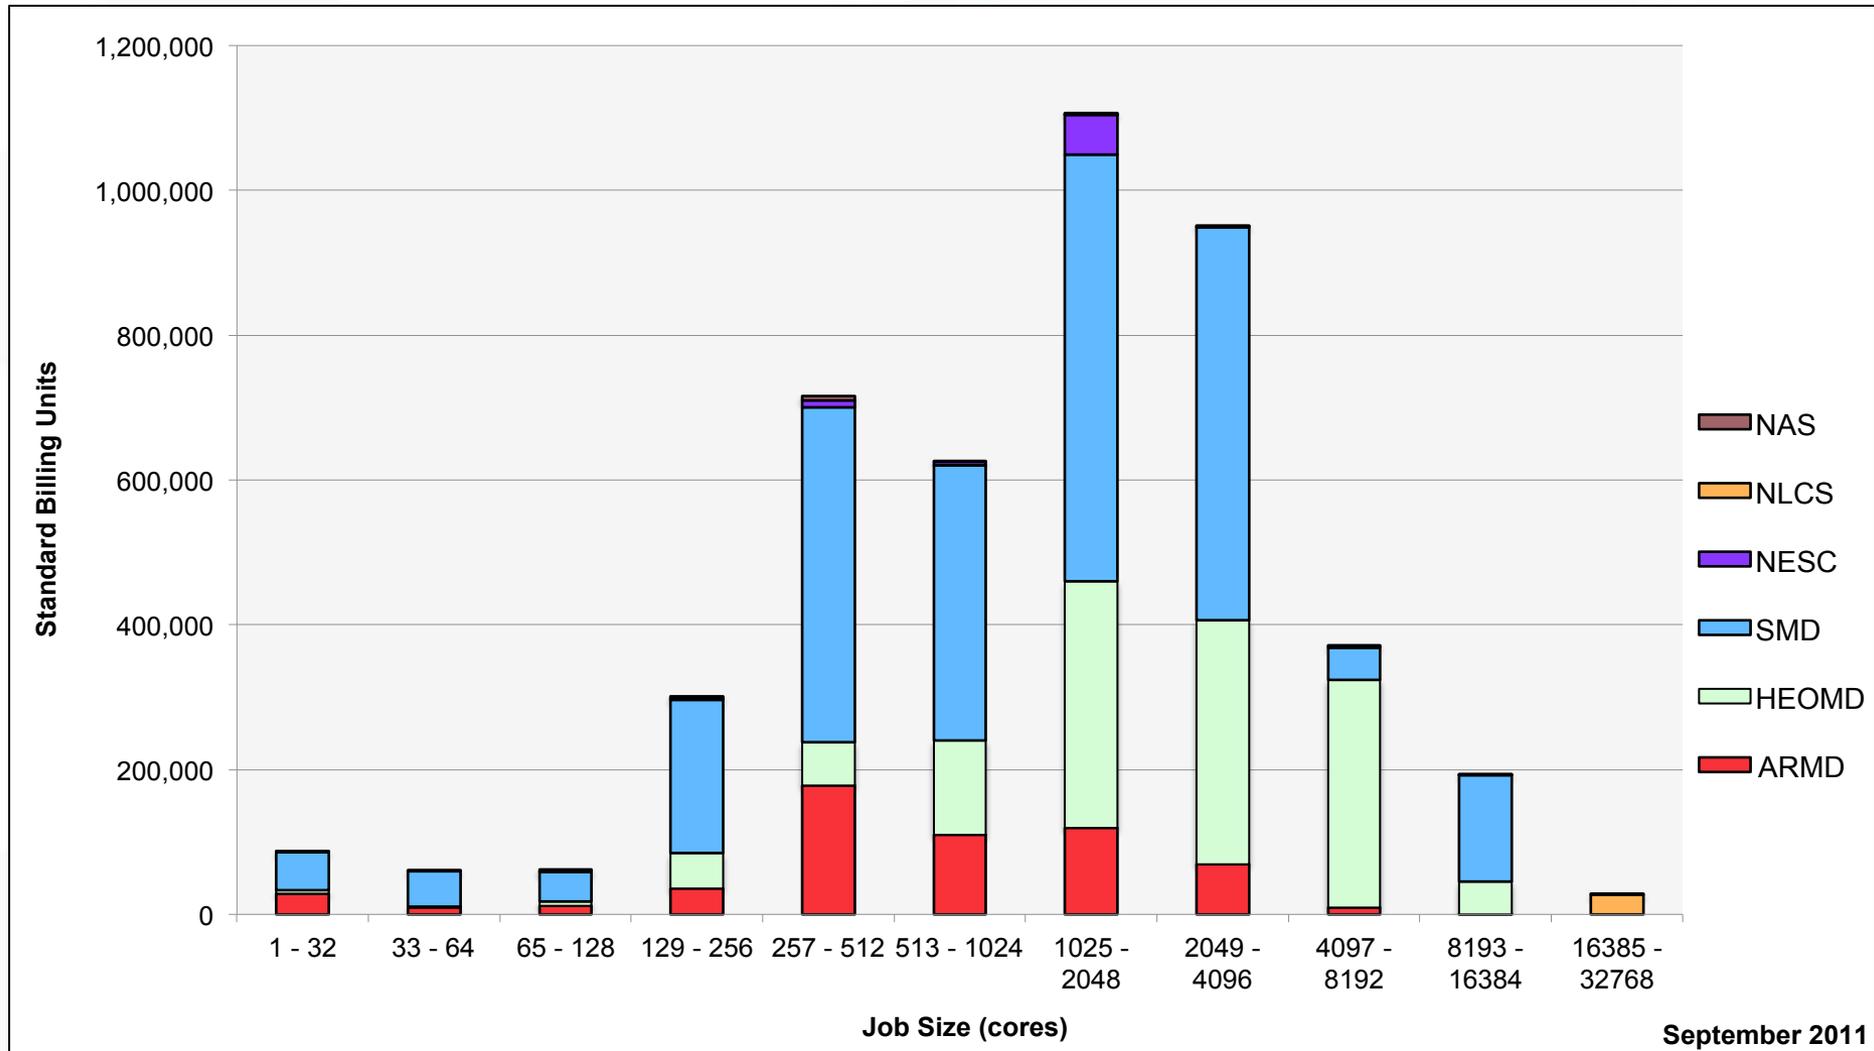
# Pleiades: Devel Queue Utilization



# Pleiades: Monthly SBUs by Run Time

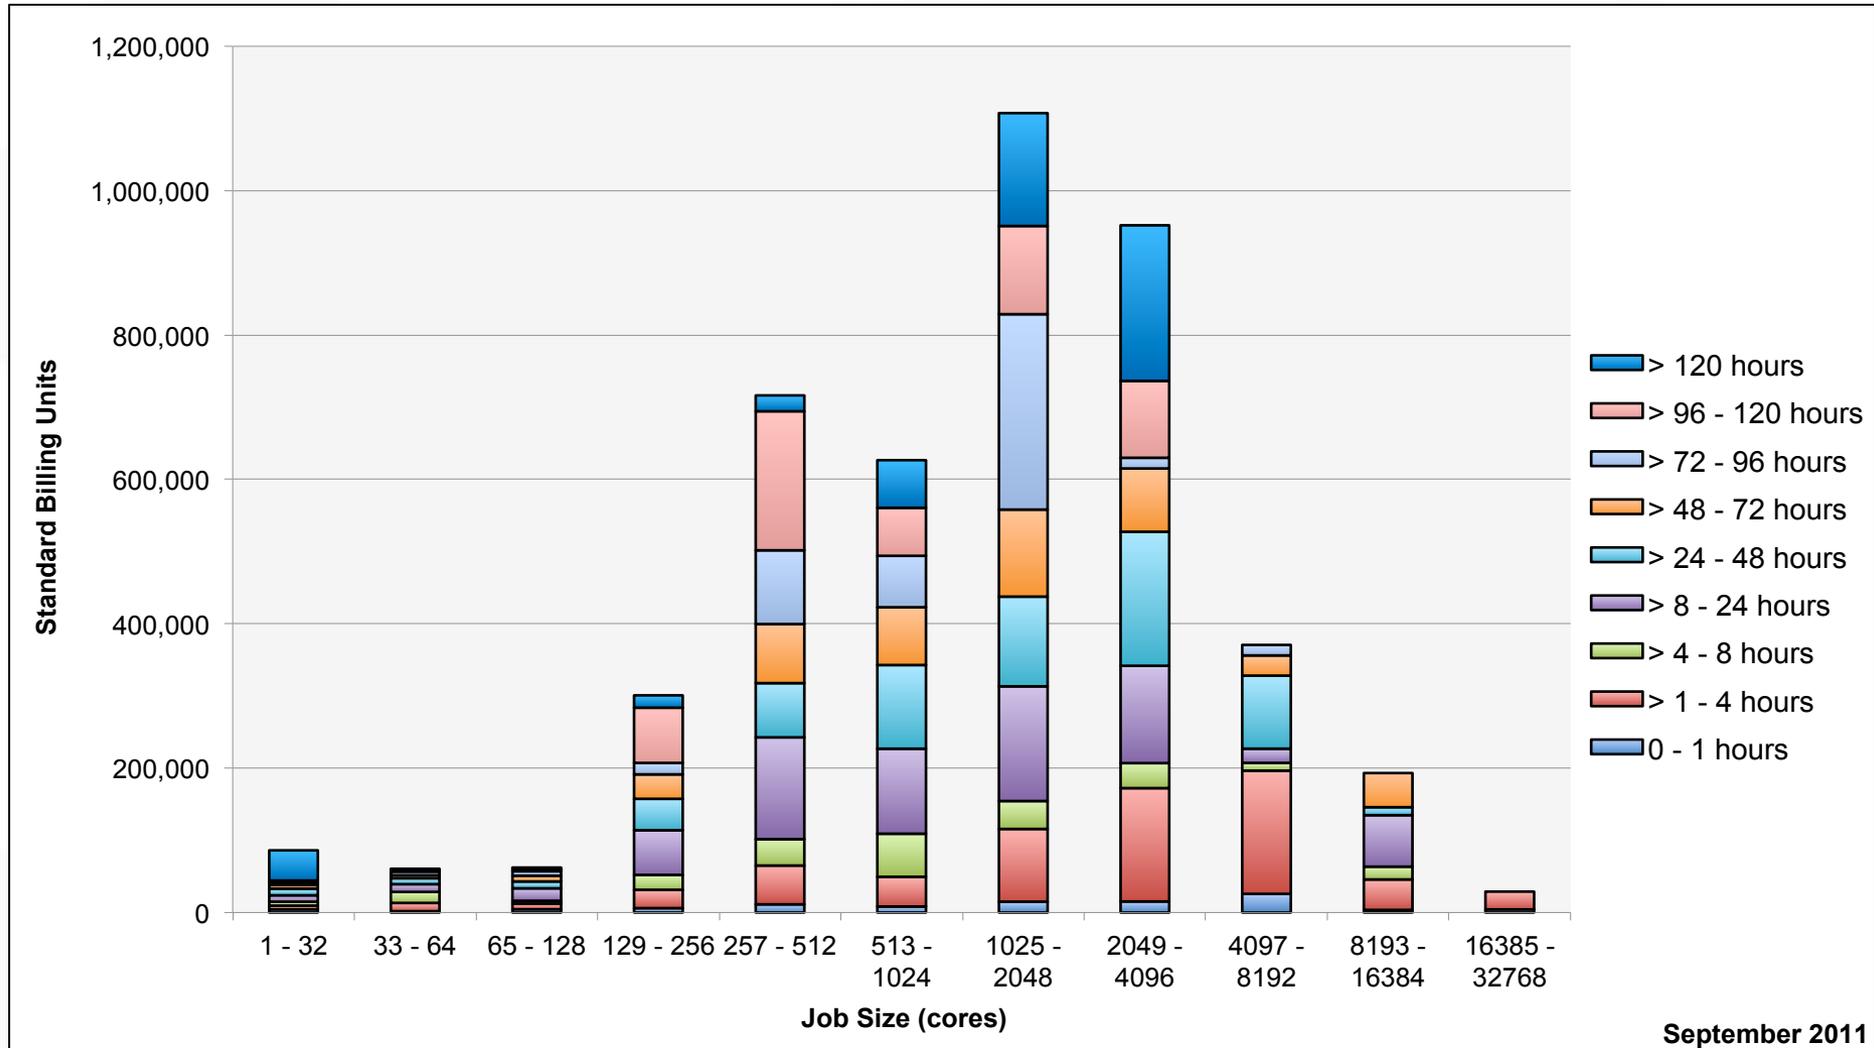


# Pleiades: Monthly Utilization by Size and Mission



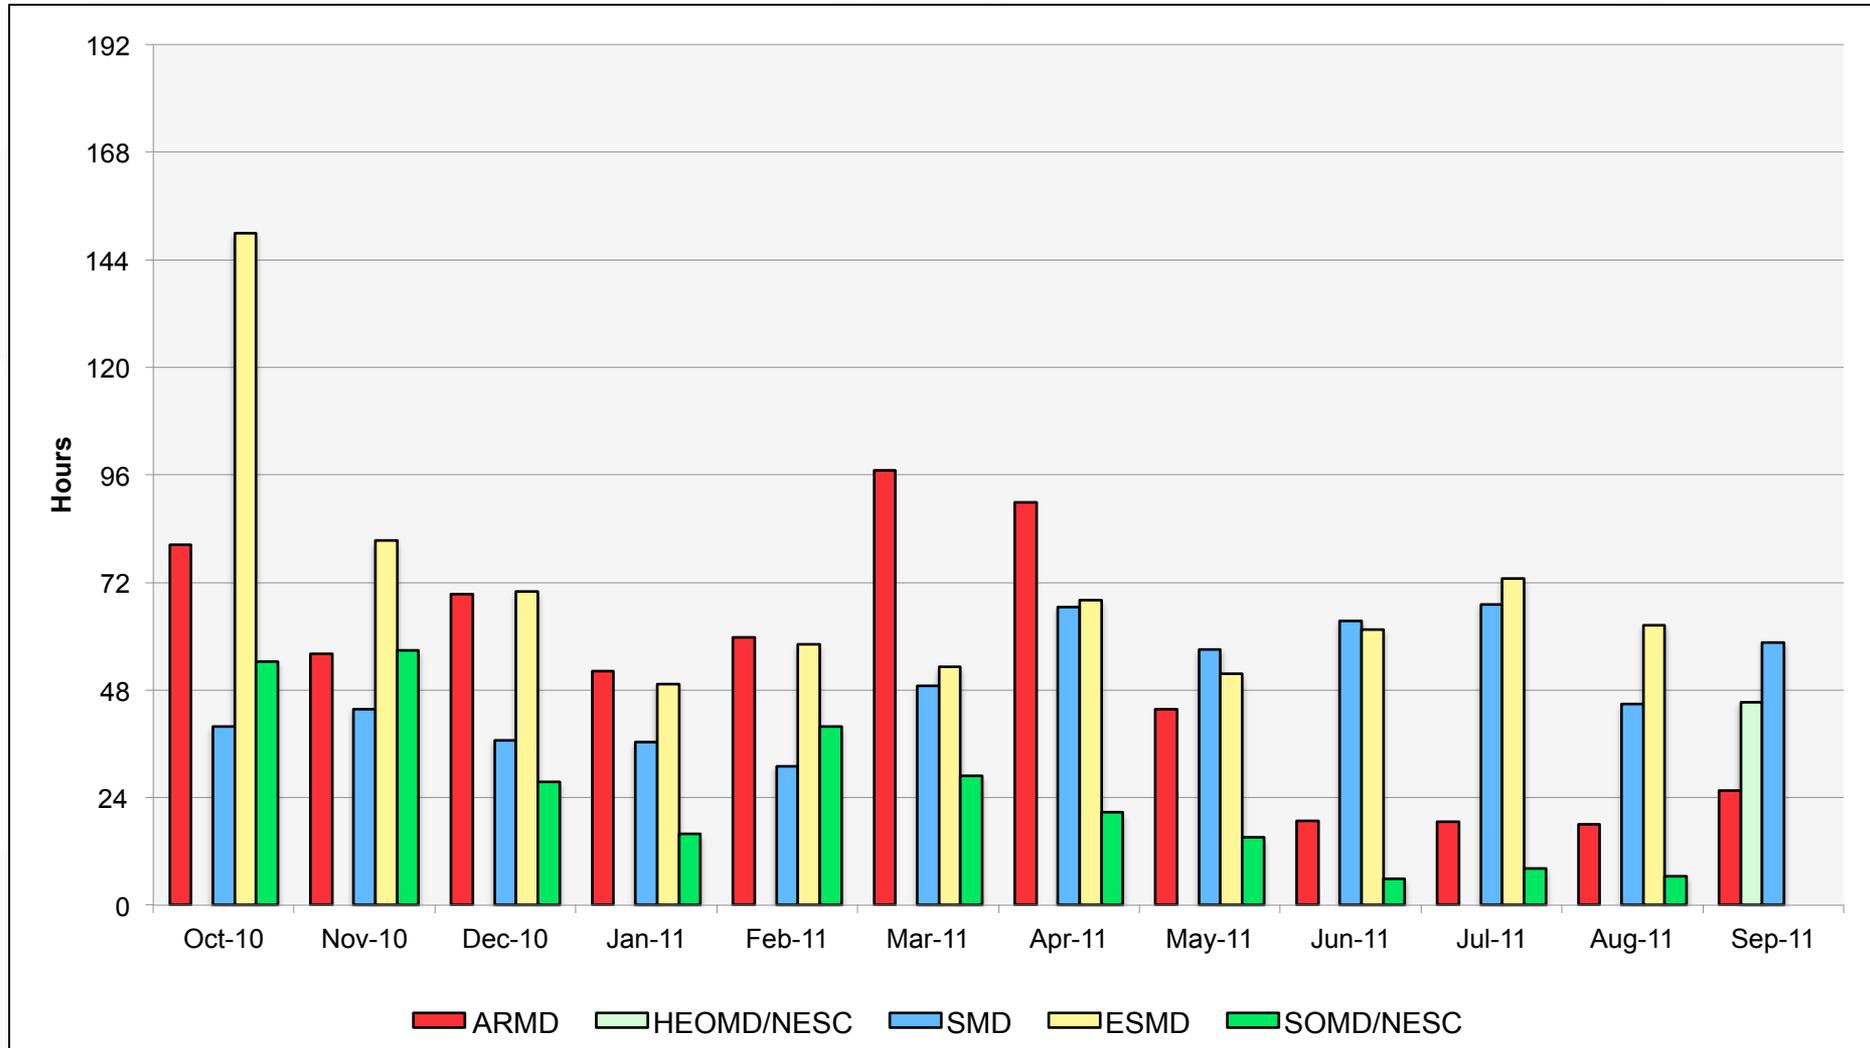


# Pleiades: Monthly Utilization by Size and Length

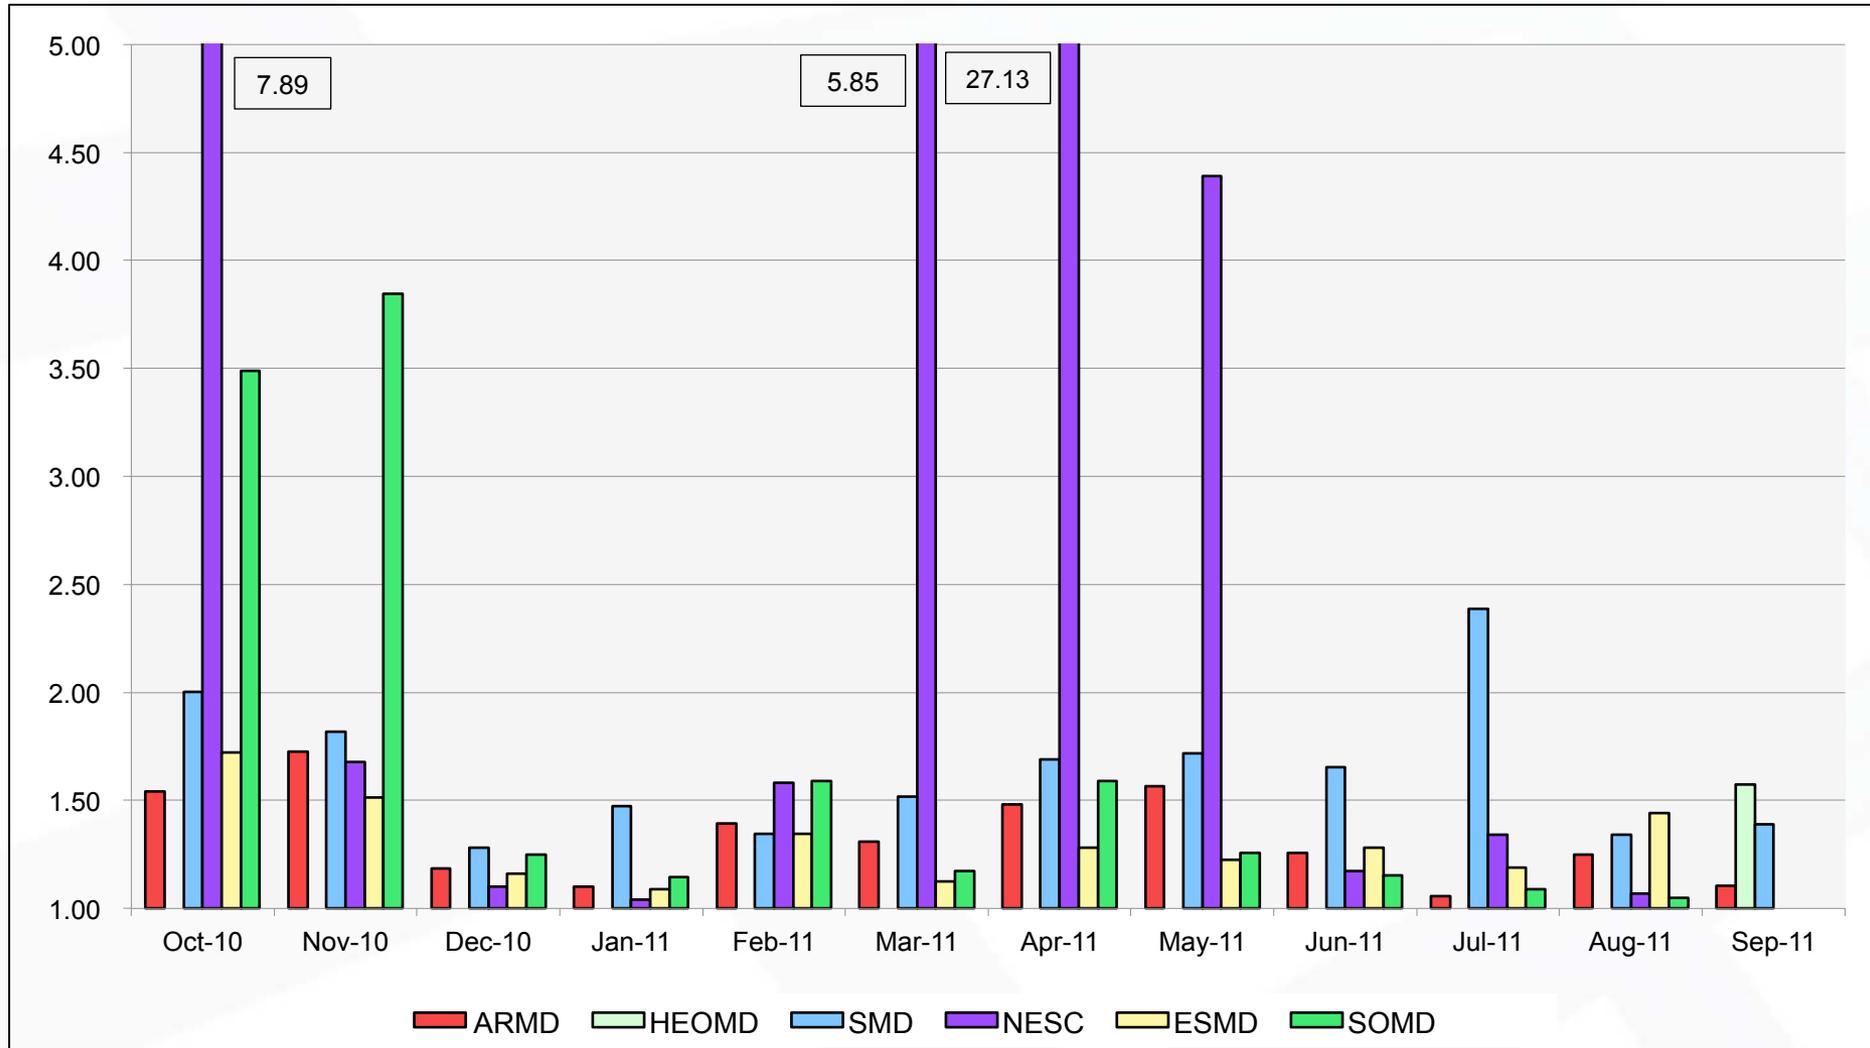


September 2011

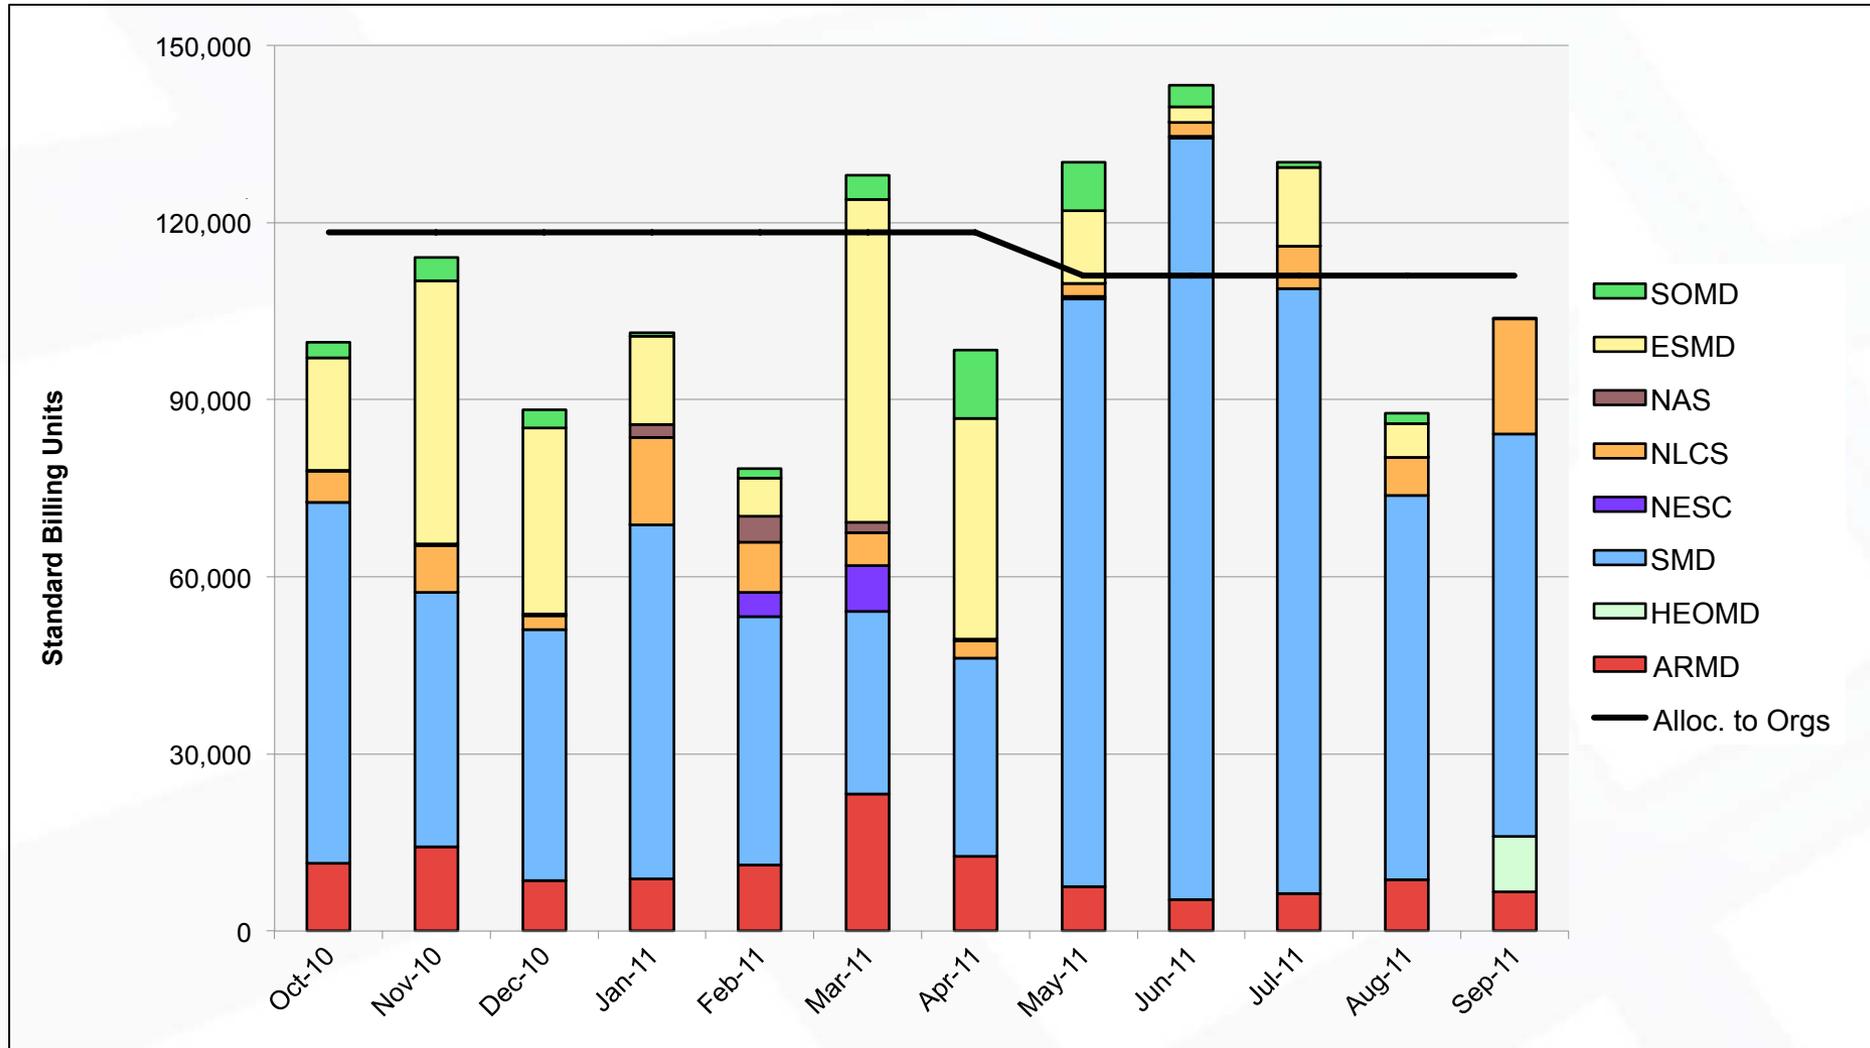
# Pleiades: Average Time to Clear All Jobs



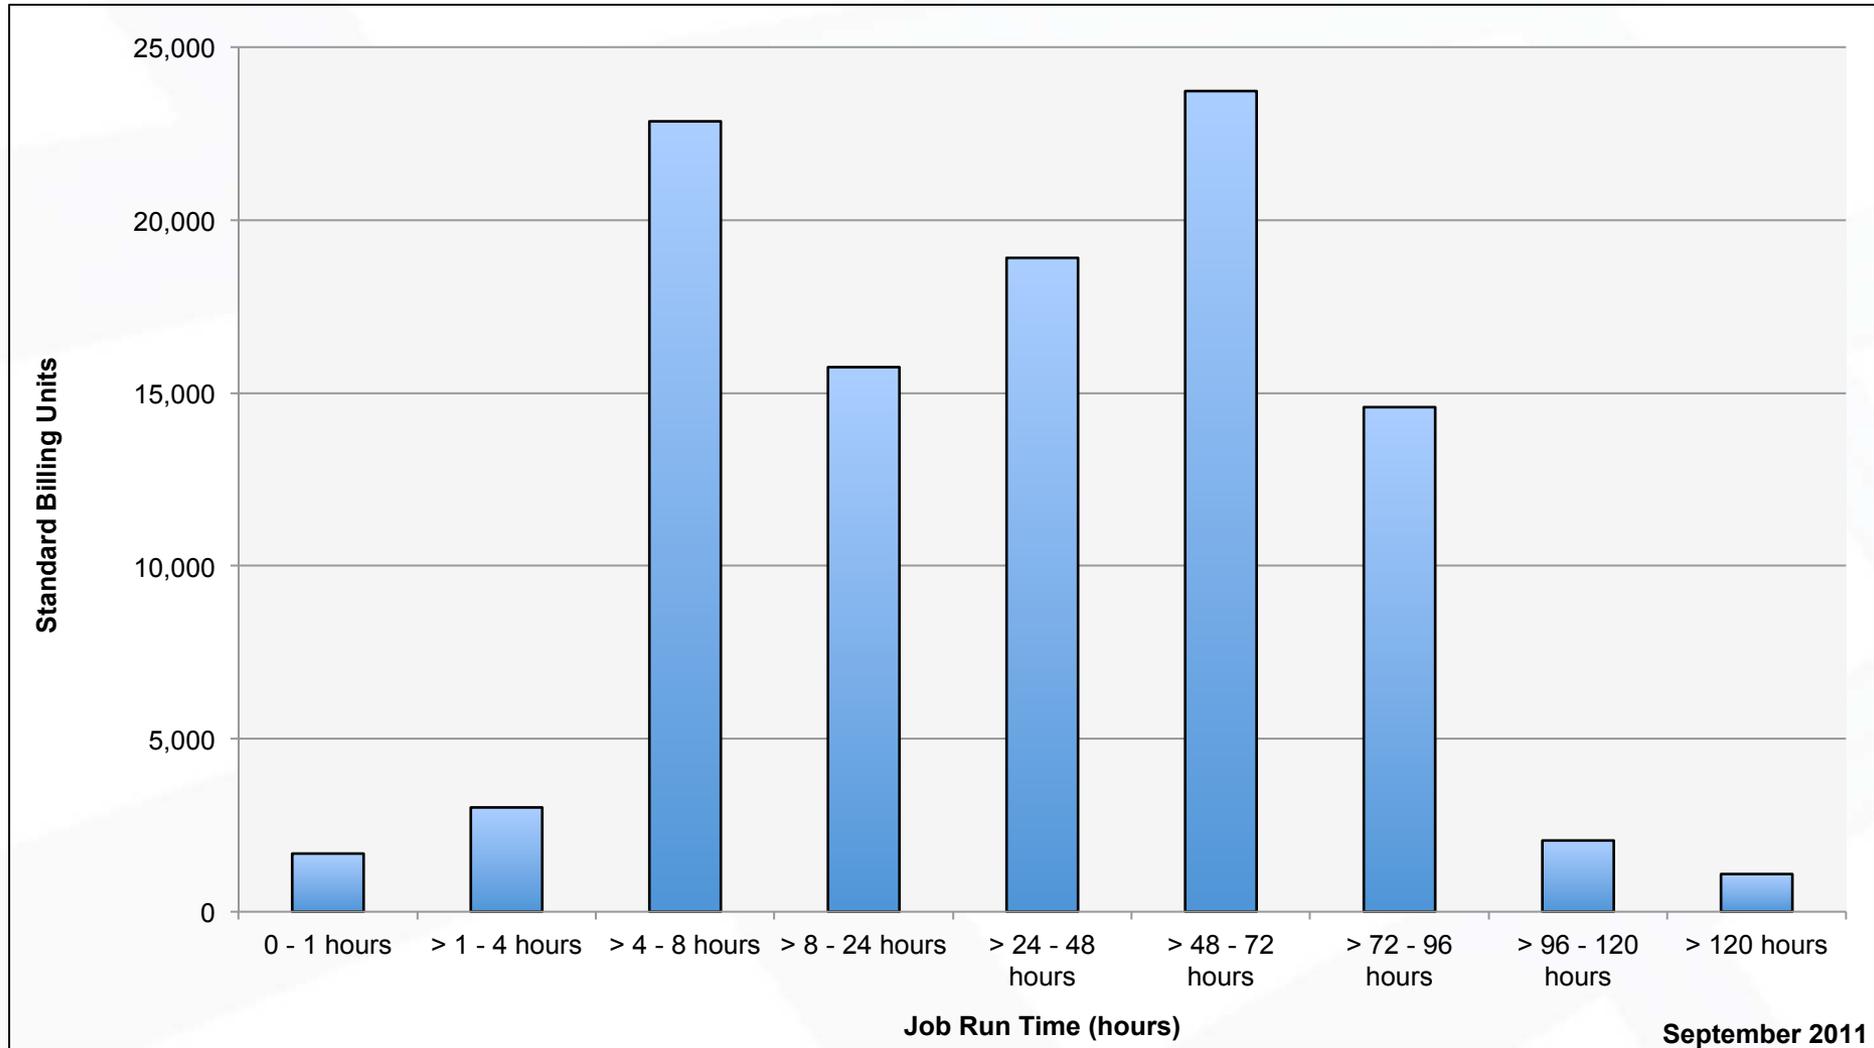
# Pleiades: Average Expansion Factor



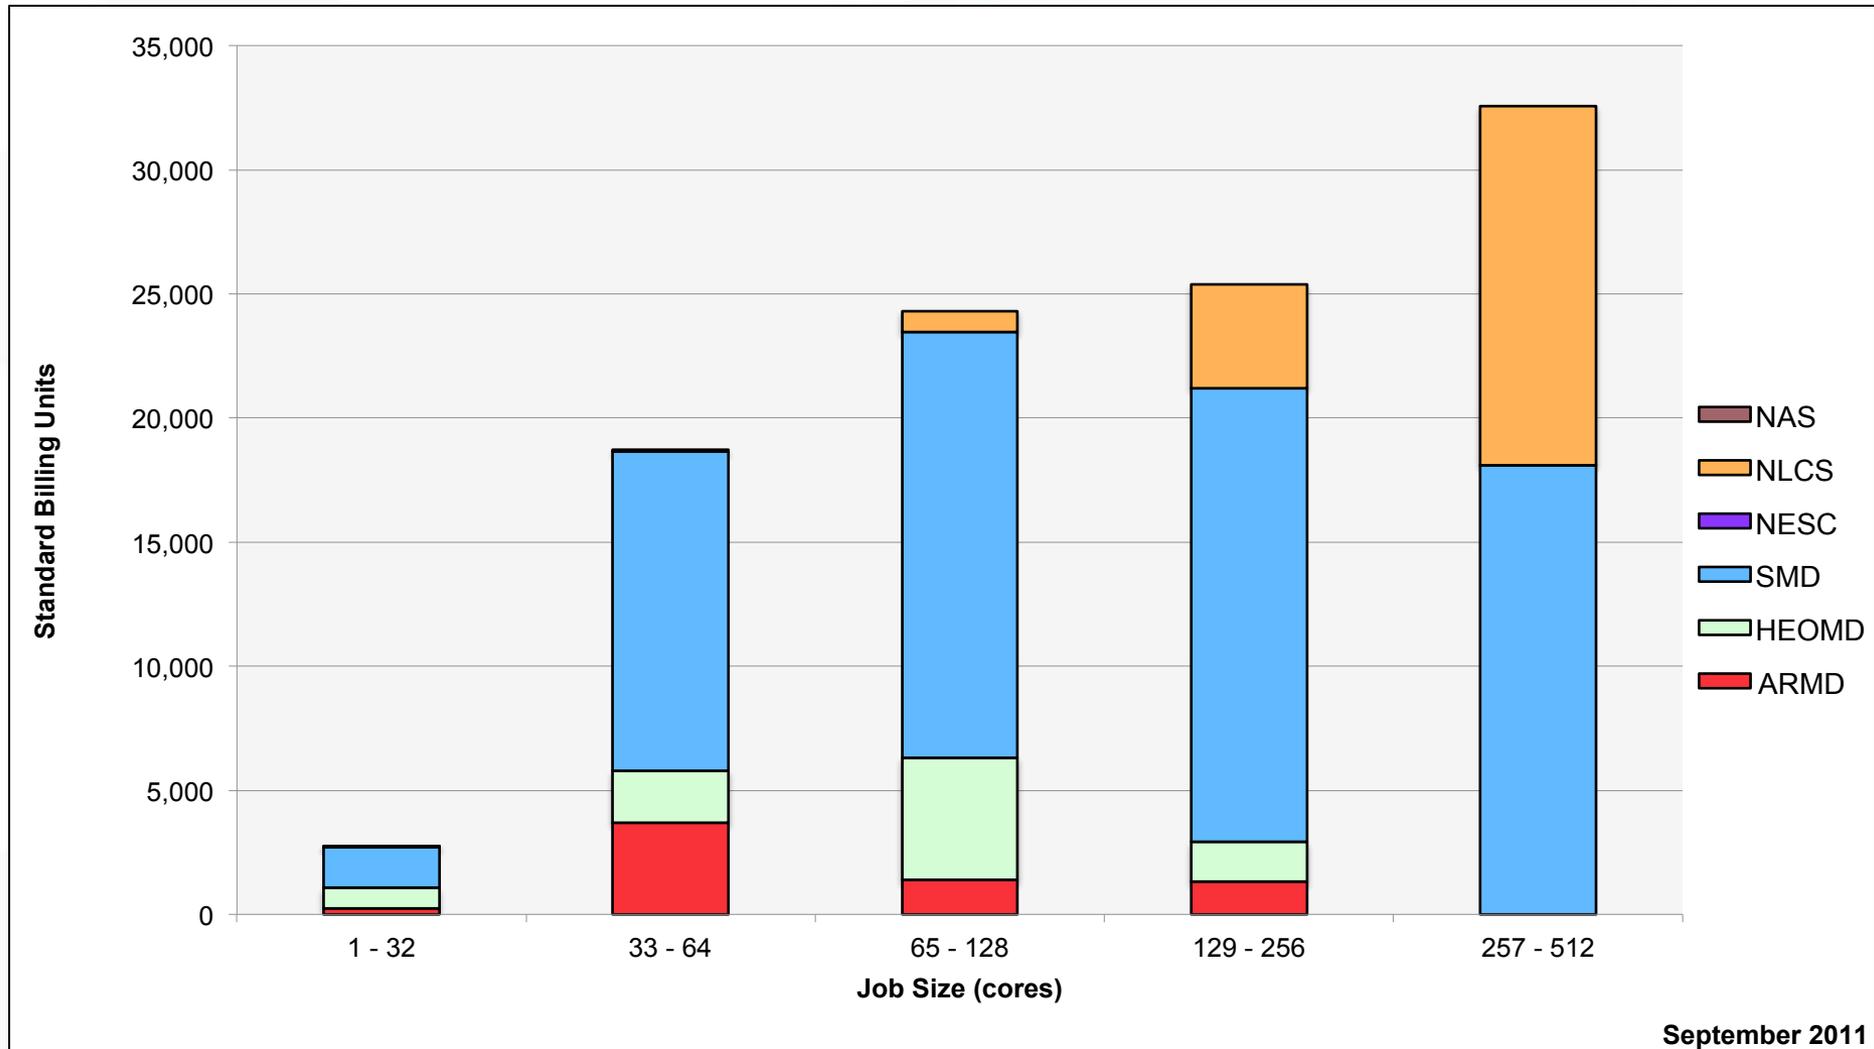
# Columbia: SBUs Reported, Normalized to 30-Day Month



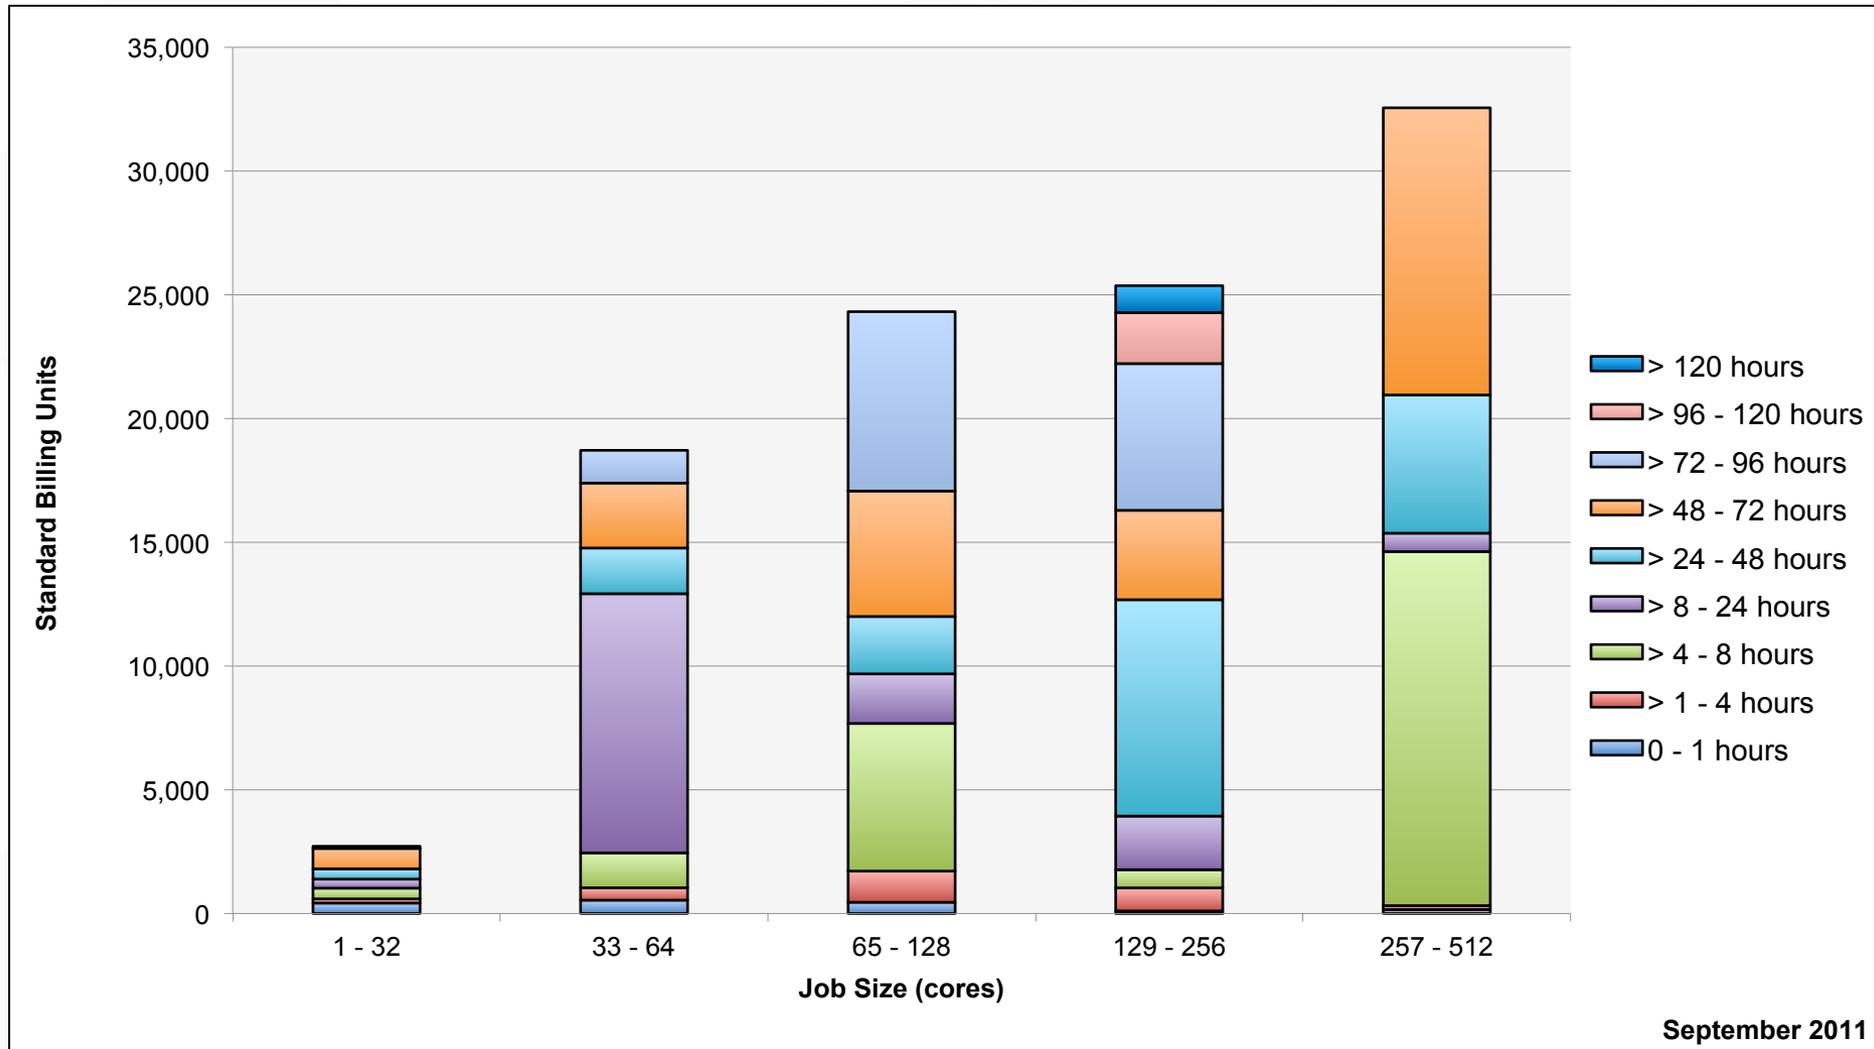
# Columbia: Monthly SBUs by Run Time



# Columbia: Monthly Utilization by Size and Mission

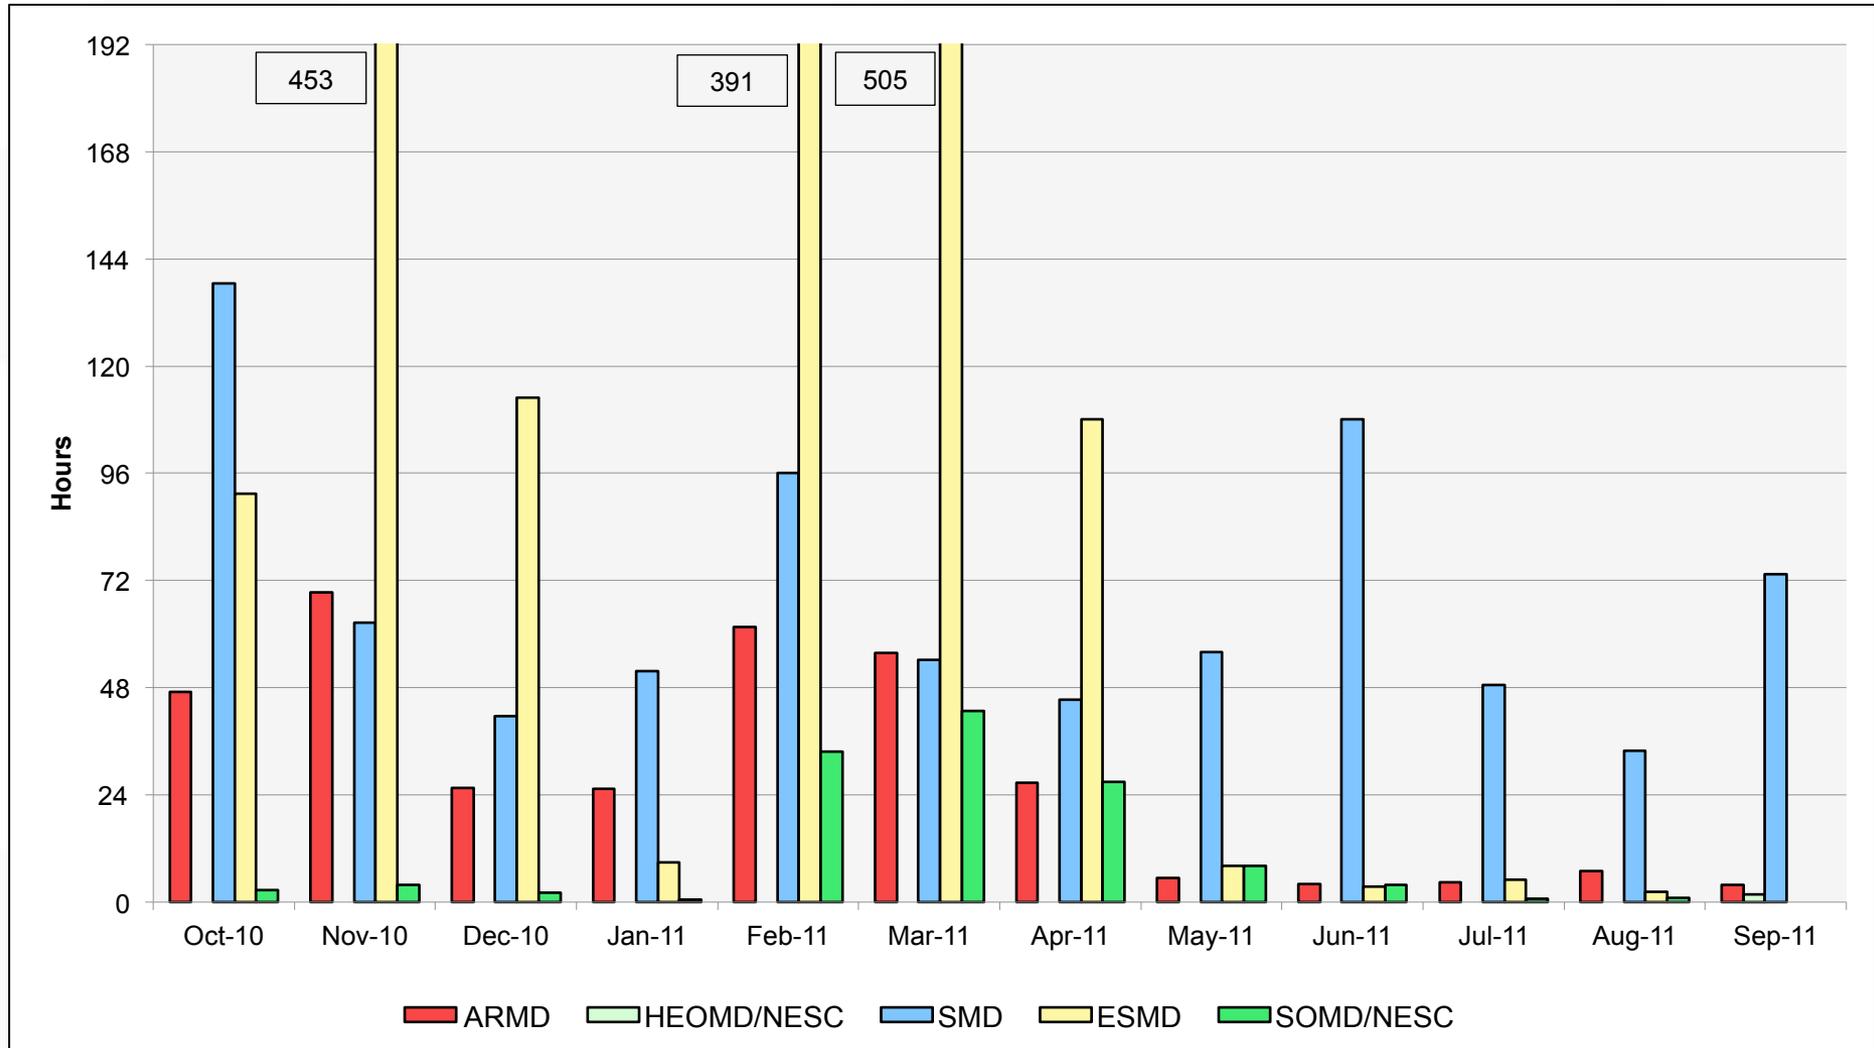


# Columbia: Monthly Utilization by Size and Length



September 2011

# Columbia: Average Time to Clear All Jobs



# Columbia: Average Expansion Factor

