

National Aeronautics and Space Administration



# Project Status Report

## High End Computing Capability Strategic Capabilities Assets Program

8 November 2011

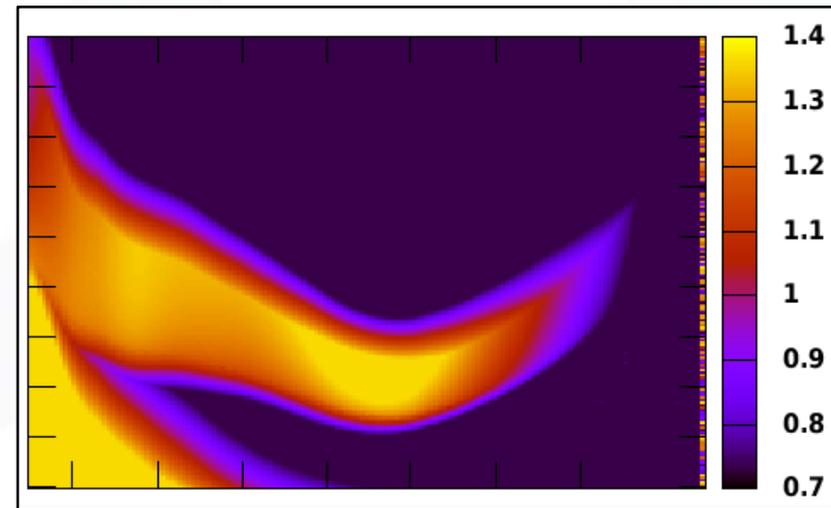
Dr. Rupak Biswas – Project Manager  
NASA Advanced Supercomputing (NAS) Division  
NASA Ames Research Center, Moffett Field, CA  
Rupak.Biswas@nasa.gov  
(650) 604-4411

# Visualization Team Completes Framework for Kepler Mission Analysis



- HECC data analysis and visualization experts completed a framework to largely automate the setup, running, and post-processing of data output by Kepler's BLENDER tool, used to validate planet candidates.
- Setup routines incorporate available orbital eccentricity constraints and accommodate "extra-dimensional" parameter scans done for difficult objects.
- Runtime processes gather and assemble results from raw output, and produce fully formatted, ordered, and labeled files for post-processing and final delivery.
- The framework systematically names all output, publishes it to an appropriate web location, and generates a message with links to the web for delivery to the Kepler team.
- Blender author Guillermo Torres said, "It's good to know that your automation efforts are paying off: you are getting ahead of us now!"

**Mission Impact:** The HECC Project's deep expertise in both visualization and data analysis technologies has enabled automation of complex, time-consuming processes to produce results for the Kepler Mission Science Team.



**Figure:** High-dimensional parameter landscape revealing viable "blends" that can mimic Kepler photometry. Chris Henze, NASA Ames.

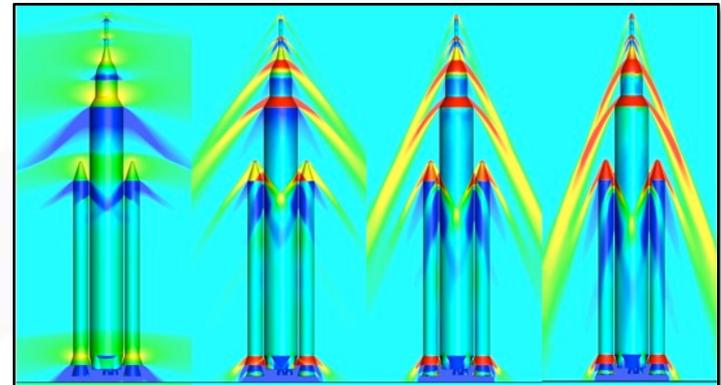
**POC:** Christopher Henze, [christopher.e.henze@nasa.gov](mailto:christopher.e.henze@nasa.gov), (650) 604-3180, NASA Advanced Supercomputing Division

# CFD Simulations Support Assessment of Seven Candidate SLS Designs



- Modeling & Simulation experts in the NASA Advanced Supercomputing Division used the Pleiades supercomputer to perform computational fluid dynamics (CFD) simulations of ascent aerodynamics for seven candidate Space Launch System (SLS) designs.
- Using NASA's Cart3D CFD code, the team performed a total of 3,334 inviscid simulations to generate a database of aerodynamic forces, moments, and line-loads for each SLS design at a range of key flight conditions throughout their ascent trajectories.
- Using around 70,000 core-hours on Pleiades, the team was able to complete the extensive databases within just two weeks.
- These analyses will help the SLS Program compare and assess the alternate configurations to select the most promising design options for further development.

**Mission Impact:** CFD analyses of candidate SLS designs enable efficient assessment and comparison their of aerodynamic performance and loads, supporting selection of the most promising design options for further development.



*Figure:* Pressure contours from CFD simulations of a candidate SLS crew launch design configuration.

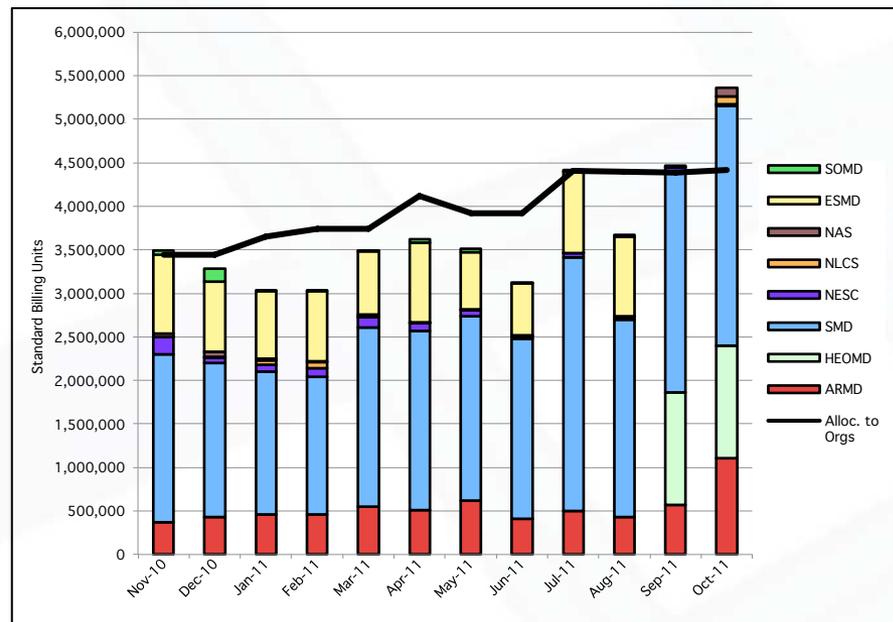
*POC:* Cetin Kiris, [cetin.c.kiris@nasa.gov](mailto:cetin.c.kiris@nasa.gov), (650) 604-4485,  
NASA Advanced Supercomputing Division

# Pleiades Sets a New Record in Utilization Exceeding Five Million SBUs in One Month



- October 2011 showed record high usage of HECC's Pleiades supercomputer.
- Over 5.4 million Standard Billing Units (SBUs) were used by NASA's Mission Directorates and mission support organizations—a 21% increase over the previous record.
- The majority of the increase is due to high user demand coupled with a stable system and efficient operations. It was further enabled by the purchase of two new racks to support the Fundamental Aeronautics Program.
- Researchers in the Aeronautics Research Mission Directorate had the biggest increase in usage, doubling their usage from previous months. The Science Missions Directorate continues to be the heaviest user.
- HECC regularly expands the computational assets to support the Mission Directorates' challenging requirements in advancing science and engineering.

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



**Figure:** Utilization of Pleiades by all Mission Directorates and support organizations. Data is shown in Standard Billing Units (SBUs), where one SBU is equivalent to one node hour of a 12-core Westmere node of Pleiades. Data is normalized to a 30-day month.

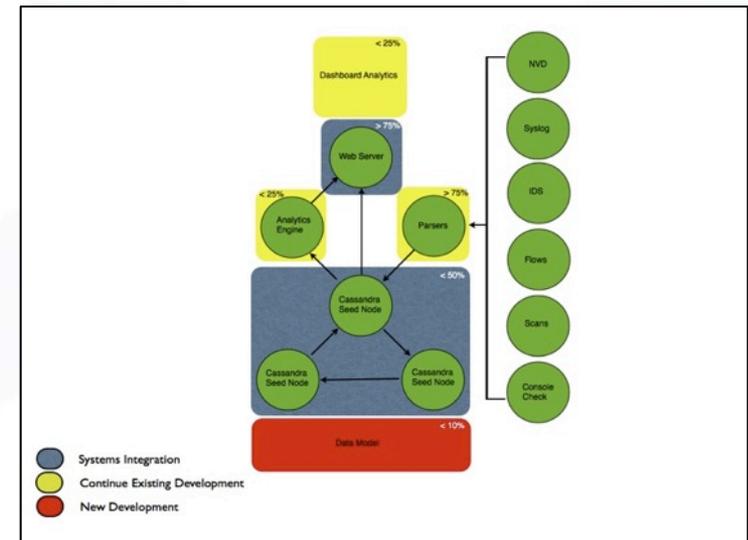
**POC:** Catherine Schulbach, [catherine.h.schulbach@nasa.gov](mailto:catherine.h.schulbach@nasa.gov)  
(650) 604-3180, NASA Advanced Supercomputing Division



# Initial Tests of New Security Storage Engine Shows 50% Performance Improvement

- The Cassandra database is designed to be a highly available, highly scalable distributed key-value store.
- Cassandra is designed for write-heavy read-light applications, which mirrors the HECC security monitoring environment perfectly.
- Developers on the HECC Security team have redesigned some of database schemas and modified parsers and analytics to interact with Cassandra instead of MySQL.
- Initial test results have shown a significant increase in performance for many HECC-developed security algorithms and parsers—by 50% or more.
- Having a distributed, scalable architecture allows HECC security analysts to use Cassandra as a cornerstone in its tool development arsenal, allowing them to analyze data in ways not possible using traditional databases, due to the scale of the datasets.

**Mission Impact:** Incorporating new technologies enables a more complete analysis of the flow data reducing the false-positive alerts, allowing the analyst to focus valuable attention on real security events.



**Figure:** The HECC Security team will store all data sources into a Cassandra cluster, using a new analytics engine to process the increased flow.

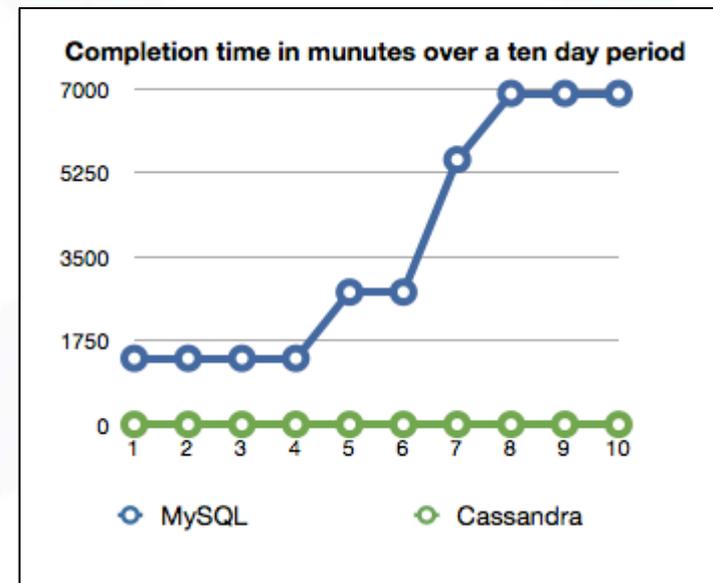
**POC:** Christopher Keller, [christopher.keller@nasa.gov](mailto:christopher.keller@nasa.gov), (408) 250-7220, NASA Advanced Supercomputing Division, Computer Sciences Corp.



# Security Team Improves Vulnerability Tracking Using 'Key Value' Philosophy

- With recent improvements to its Vulnerability Tracking System, the Security team can now correlate known vulnerabilities (called CVEs) from the National Vulnerability Database with installed packages (sets of software fixes) on HECC systems such as Pleiades and Columbia.
- Over time, as more systems were added, the existing relational MySQL approach was taking more than five days to complete.
- With more than 15,000 packages and over 120,000 CVEs being tracked by HECC security experts, performing this huge number of reads and writes in MySQL became very costly (thus the five-day completion).
- Testing revealed that doing over 15,000 lookups in Cassandra's key value database (see slide 6) to get CVE data is extremely fast—what used to take five days with MySQL now takes five minutes with Cassandra.

**Mission Impact:** Investigating next-generation database technology allows HECC security experts to re-think how to approach and solve big data problems.



**Figure:** This chart shows processing time completion in minutes over a 10-day period. The large spikes (increased system data) on days 5 and 7 cause MySQL to double its time while Cassandra stayed constant.

**POCs:** Thomas H. Hinke, [thomas.h.hinke@nasa.gov](mailto:thomas.h.hinke@nasa.gov), (650) 604-3662, NASA Advanced Supercomputing Division; Ryan Spaulding, [ryan.c.spaulding@nasa.gov](mailto:ryan.c.spaulding@nasa.gov), (650) 604-4335, NASA Advanced Supercomputing Division, Computer Sciences Corp.

# Two Pleiades Filesystems Deployed with 10x Increase in IOPS Performance



- Two additional Pleiades Lustre filesystems, named /nobackupp3 and /nobackupp4, have been deployed to replace the older nobackupp50/60 filesystems.
- The new filesystems will provide 10 times the Input/Output Operations per Second (IOPS) performance of the old filesystems.
- The 10-fold improvement in IOPS addresses a limitation in the previous RAID controller, and will provide better interactive filesystem performance to researchers.
- The remaining filesystem on Pleiades, /nobackupp30, will also be upgraded to attain a 10x improvement in IOPS.
- When the /nobackupp30 upgrade is complete, HECC's several-month project to provide larger "scratch space" for users to temporarily store files will be complete.

**Mission Impact:** Along with the increased computational capability of the Pleiades supercomputer, faster and larger "scratch space" (temporary storage) enables researchers to more fully utilize Pleiades.



*Figure:* Pleiades new Lustre filesystems provide a 10-fold improvement in the Input/Output Operations per Second performance.

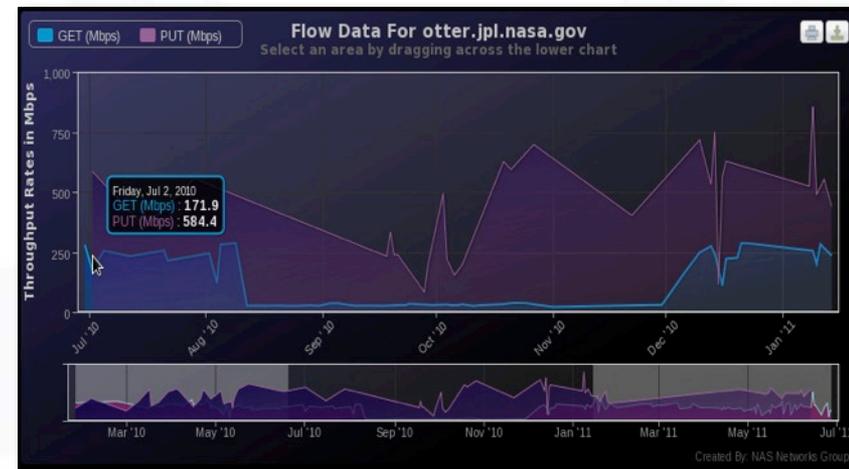
*POC:* Bob Ciotti, [bob.ciotti@nasa.gov](mailto:bob.ciotti@nasa.gov), (650) 604-4408,  
NASA Advanced Supercomputing Division

# Networks Team Enhances Visual Flow Analysis Capabilities



- The HECC Networks team completed enhancements to its FlowTools analysis software to reduce time-to-solution for identifying bottlenecks in data transfer. Enhancements include:
  - Graphs are dynamically generated based on network flow data collected showing performance history of inbound/outbound from the NASA Advanced Supercomputing facility, with drill-down views of specific timeframes.
  - Graphs can be generated for specific remote systems or sites to easily see if slow flow is affecting a single system or the entire site.
  - Visualization of both performance metrics and data volume transferred allows for simple trend analysis and to understand customer habits.
- By graphically representing network flows, engineers no longer have to sort through thousands of flow records to identify any performance changes over time.

**Mission Impact:** Software enhancements allow engineers to more quickly identify changes in network flow performance trends and determine if a problem exists—providing improved network performance to the HECC user community.



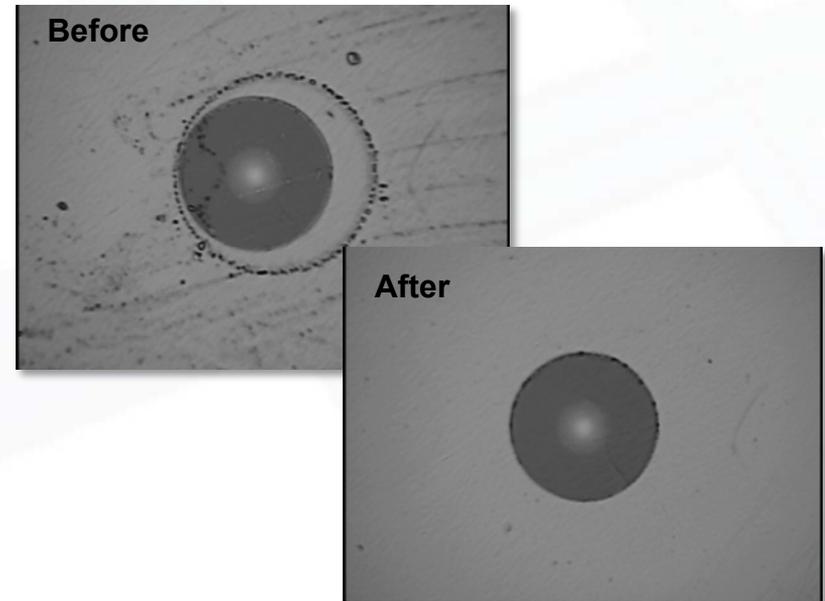
**Figure:** This example graph above shows traffic inbound (GET) and outbound (PUT) to/from the NASA Advanced Supercomputing facility at Ames to a server at NASA's Jet Propulsion Lab, over a selectable time period. The graph indicates that a firewall is negatively impacting traffic in one direction but not in the other.

**POC:** Hugh La Master, [hugh.lamaster@nasa.gov](mailto:hugh.lamaster@nasa.gov), (650) 604-1056, NASA Advanced Supercomputing Division

# Network Experts Provides Fiber Troubleshooting Best Practices to Vendor

- HECC network engineers assisted Spectra Logic engineers in troubleshooting connection problems with the tape library drives that the company recently upgraded.
- They isolated the problem to the small fiber jumper cable in the drive enclosure. The fiber jumpers were not properly cleaned before being installed in the drives.
- Actions taken by the Networks team include:
  - Trained Spectra Logic engineers on proper procedures for visual inspection and cleaning of the Fibre Channel components (connectors and adapters).
  - Provided Spectra Logic engineers with photos of the fiber taken with the EXFO Inc. fiber inspection tool, to make them aware of both the problems and the tools.
  - Recommended products and tools for Spectra Logic to acquire for visual inspection and cleaning of fiber connectors and adapters.

**Mission Impact:** By providing the vendor with best practices for maintaining, troubleshooting, and testing fiber interconnects, HECC improves the quality of products received from that vendor shortening the time to incorporate the upgrade as well as having a larger community impact.



**Figure:** Photos taken with the EXFO Inc. fiber inspection tool showing a Spectra Logic tape drive fiber port before cleaning (top), and after cleaning (bottom). HECC testing proved that light loss caused by dirt shown in the “before” picture is not acceptable at 8 Gbps data rates.

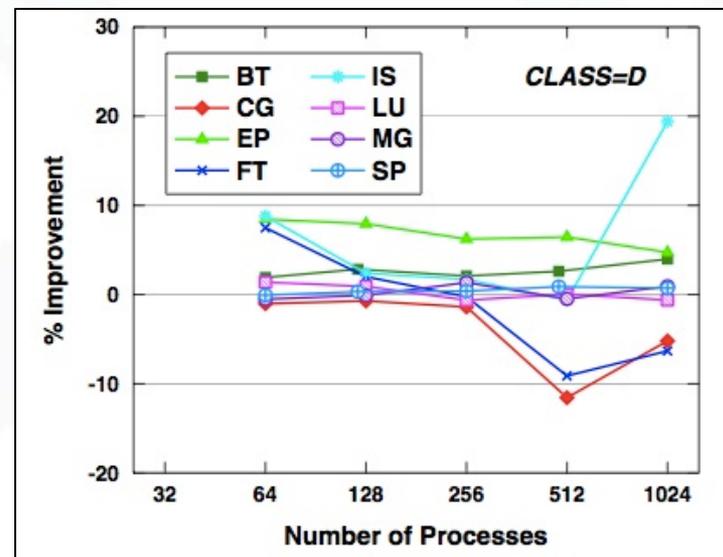
*POC: Chris Buchanan, [chris.buchanan@nasa.gov](mailto:chris.buchanan@nasa.gov), (650) 604-4308, NASA Advanced Supercomputing Division*

# HECC Detects Performance Regression Issues in New ARMD Westmere Racks



- The HECC Application Performance & Productivity (APP) team detected performance regressions during standard pre-production testing of two new Westmere racks purchased for the Aeronautics Research Mission Directorate (see *October 2011 HECC Monthly Report*), and reported the issues to the HECC Systems team.
- The APP team ran their tests using the NAS Parallel Benchmarks (NPB) along with the Standard Billing Unit (SBU) suite of six full application codes.
- Some of the codes showed worse performance despite the slightly faster clock speed of the newer Intel XEON processor (3.06 GHz vs. 2.93 GHz) in the racks; the team determined that an environmental variable crucial for pinning MPI processes to physical cores had not been set during configuration—the performance regression was resolved after the Systems team reconfigured the system.
- At large core counts (>256), the communication-intensive NPB benchmark codes (CG and FT) still ran slower than expected, most likely because of longer latency experienced by messages going through the InfiniBand switches of a nearby rack; this issue is under further investigation by the Systems team and SGI engineers.

**Mission Impact:** Routine testing of newly integrated HECC resources prior to production allows identification of possible hardware and software issues that may impact system usability and/or performance.



**Figure:** Performance of various NAS Parallel Benchmarks on the two new ARMD Westmere racks compared to other Pleiades Westmere racks (baseline). The worse performance of CG and FT at 512 and 1024 cores indicates potential communication issues in the InfiniBand network that are still being investigated.

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

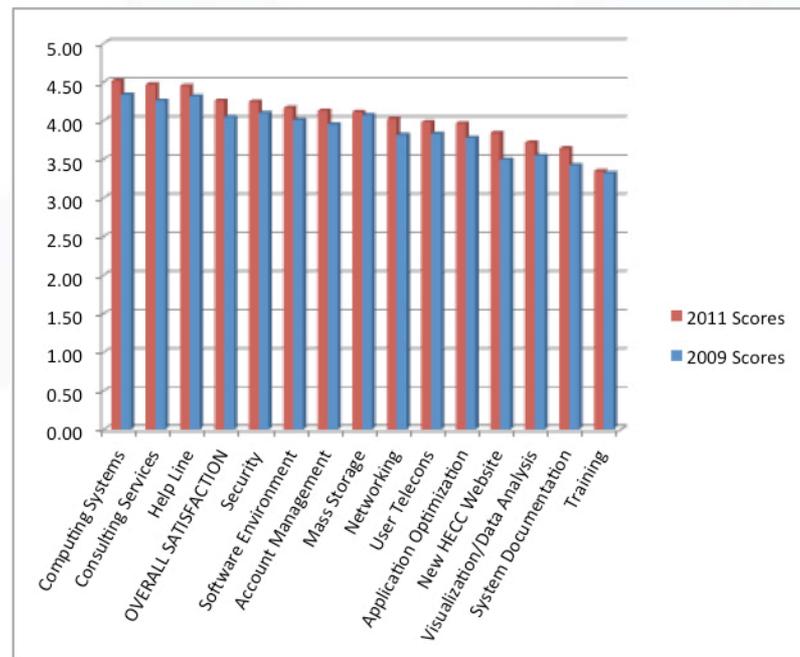
# HECC Conducts Follow-up User Survey



- The HECC project has completed a new user survey, a follow-up to its first survey in 2009.
- All NASA Mission Directorates, the NASA Engineering Safety Center, and the NASA Advanced Supercomputing Division were represented; users at NASA centers, universities, industry, and other locations responded.
- 271 users (~33% of users) assessed 14 individual aspects of service, and gave an overall evaluation.
- Scores increased slightly in all areas, with overall satisfaction scoring 4.26 out of 5.0; relative scoring is essentially unchanged (4.0 rating in 2009).
- Users provided many comments on outstanding aspects of service (computing systems, help desk, consulting) and aspects that need the most improvement (documentation, training, and web site; visualization/data analysis).
- The HECC web site was the most improved area, but further improvement is possible.
- Results are being used to guide service improvement efforts.

**Mission Impact:** User feedback provides insight into what is most important to users and what they like best and least about the services they receive. This enables service providers to better focus limited resources to provide the greatest benefit to users and assess the impact of services changes.

Scores from 2011 User Survey



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

# Tape Library Upgrade Will Triple Data Storage Capacity



- HECC has completed procurement of six 8-frame Spectra Logic T950 tape libraries, and installation of the first library is underway.
- The new libraries will hold about three times the number of tapes as the current library, providing nearly 60,000 tape slots.
- When the full upgrade is completed, the total tape storage capacity for uncompressed data will approach 110 petabytes, up from about 36 petabytes.
- The major components in the library are compatible with the next-generation library, which allows for cost-effective, in-place upgrades.

**Mission Impact:** To keep pace with the expanded computational capability of the Pleiades supercomputer, increased mass storage space is needed to enable researchers to save and re-analyze the results of their massive computations.



*Figure:* HECC's new six 8-frame Spectral Logic T950 tape libraries provide the capacity needed to support the expanded computational capabilities of Pleiades.

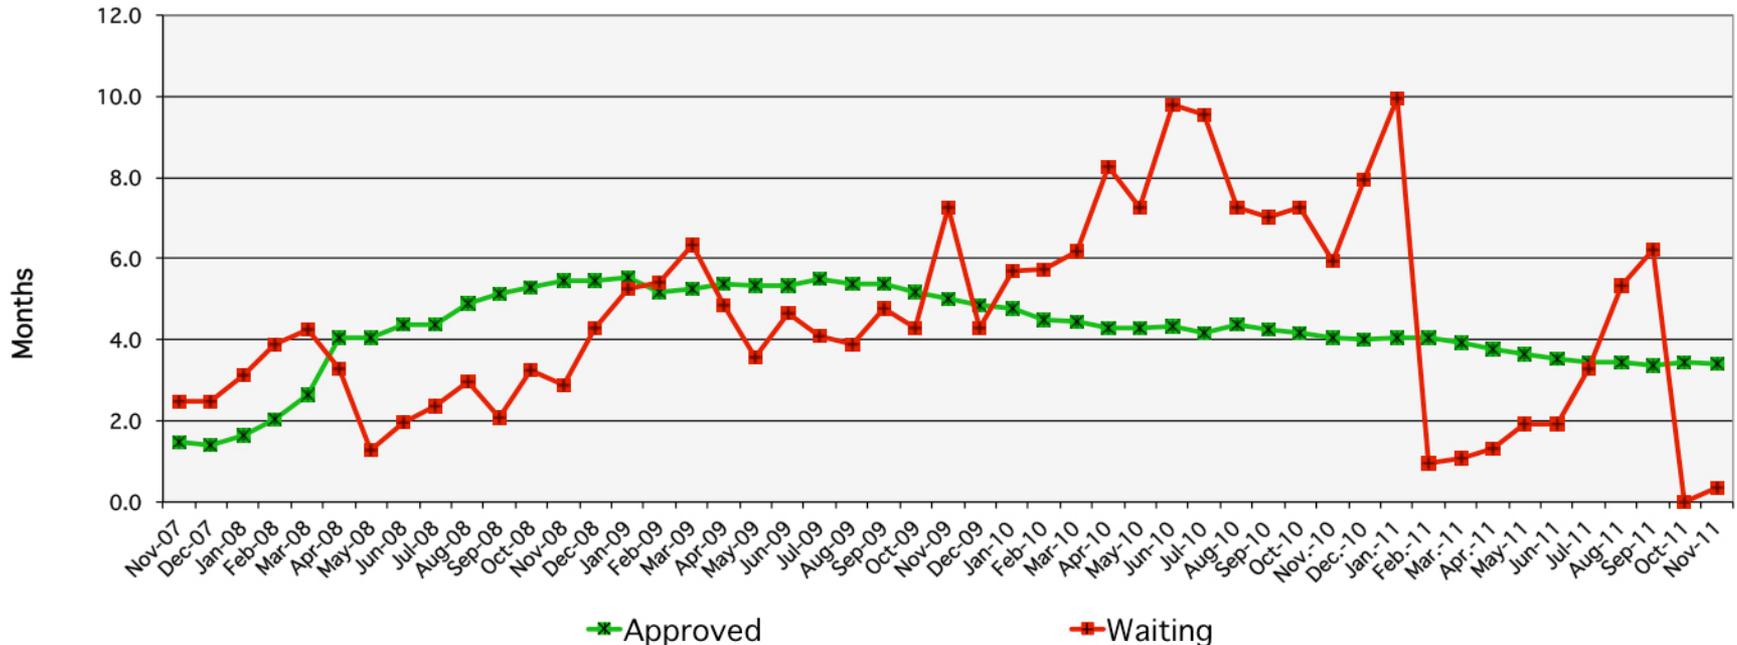
*POC:* Bob Ciotti, [bob.ciotti@nasa.gov](mailto:bob.ciotti@nasa.gov), (650) 604-4408,  
NASA Advanced Supercomputing Division

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



- Requests approved: 3; New requests received: 5; Requests waiting: 2.
- Ames Security has eliminated the requirement for Form 781, "Request for Foreign National Computer Access" because IdMAX makes it unnecessary to have this paper form. The action removes at least 4 steps from the process, and thus a big bottleneck to getting accounts for Foreign Nationals is now eliminated.

Average Wait for Requests Submitted After Aug. 1, 2007



# HECC Facility Hosts Several Visitors and Tours in October 2011



- HECC hosted 8 scheduled tour groups in October; guests received an overview of the HECC Project, demonstrations of the hyperwall-2 visualization system, and tours of the computer room floor. Among this month's visitors were:
- Elizabeth Robinson, NASA Chief Financial Officer and guests.
- Larry Gumbiner, Deputy Assistant Secretary, United Nations Conference on Sustainable Development (RIO+20); Erin Mote, USAID; Susan Povenmire and Isabel Gates, senior policy advisors, Bureau of Oceans, Environment and Science.
- 80 students from the Ames-hosted SACNAS Conference; students gathered ideas and received encouragement to pursue education in Science, Technology, Engineering, and Mathematics (STEM) fields.
- Admiral Conrad Lautenbacher, Vice President, Science Programs, Computer Sciences Corp., and past NOAA Administrator.
- 50 participants in an Ames Tweet-Up, as part of Ames' SOFIA project open house.



**Figure:** NASA Chief Financial Officer Elizabeth Robinson and guests attend a demonstration of the hyperwall-2 visualization system, as part of their tour of the NASA Advanced Supercomputing facility—images shown are from the Kepler Mission's search for Earth-size planets.

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

# Presentations and Papers



- "The Rockstar Phase-Space Temporal Halo Finder and the Velocity Offsets of Cluster Cores," Peter S. Behroozi, Risa H. Wechsler, Hao-Yi Wu, submitted to APJ, October 31, 2011.  
<http://arxiv.org/pdf/1110.4372>
- "Discretized peridynamics for brittle and ductile solids," Wenyang Liu, Jung-Wuk Hong, International Journal for Numerical Methods in Engineering, John Wiley & Sons, Ltd., October 21, 2011.\*  
<http://onlinelibrary.wiley.com/doi/10.1002/nme.3278/abstract>
- "HPC @ NASA," Rupak Biswas, presentation at the 43rd HPC User Forum, Paris, France, October 4, 2011.
- "HPC @ NASA," Rupak Biswas, presentation at the 44th HPC User Forum, Stuttgart, Germany, October 7, 2011.
- "Supercomputing @ NASA: A Brief Overview," Rupak Biswas, presentation at the National Aerospace Laboratories, Bangalore, India, October 11, 2011.
- "K&C Science Report – Phase 2, ALOS Image Mosaics for Wetland Mapping," Bruce Chapman, Kyle McDonald, Laura Hess, JAXA Kyoto & Carbon Initiative Science Team Report, October 2011.\*  
[http://www.eorc.jaxa.jp/ALOS/en/kyoto/phase\\_2/KC-Phase-2\\_report\\_chapman-small.pdf](http://www.eorc.jaxa.jp/ALOS/en/kyoto/phase_2/KC-Phase-2_report_chapman-small.pdf)

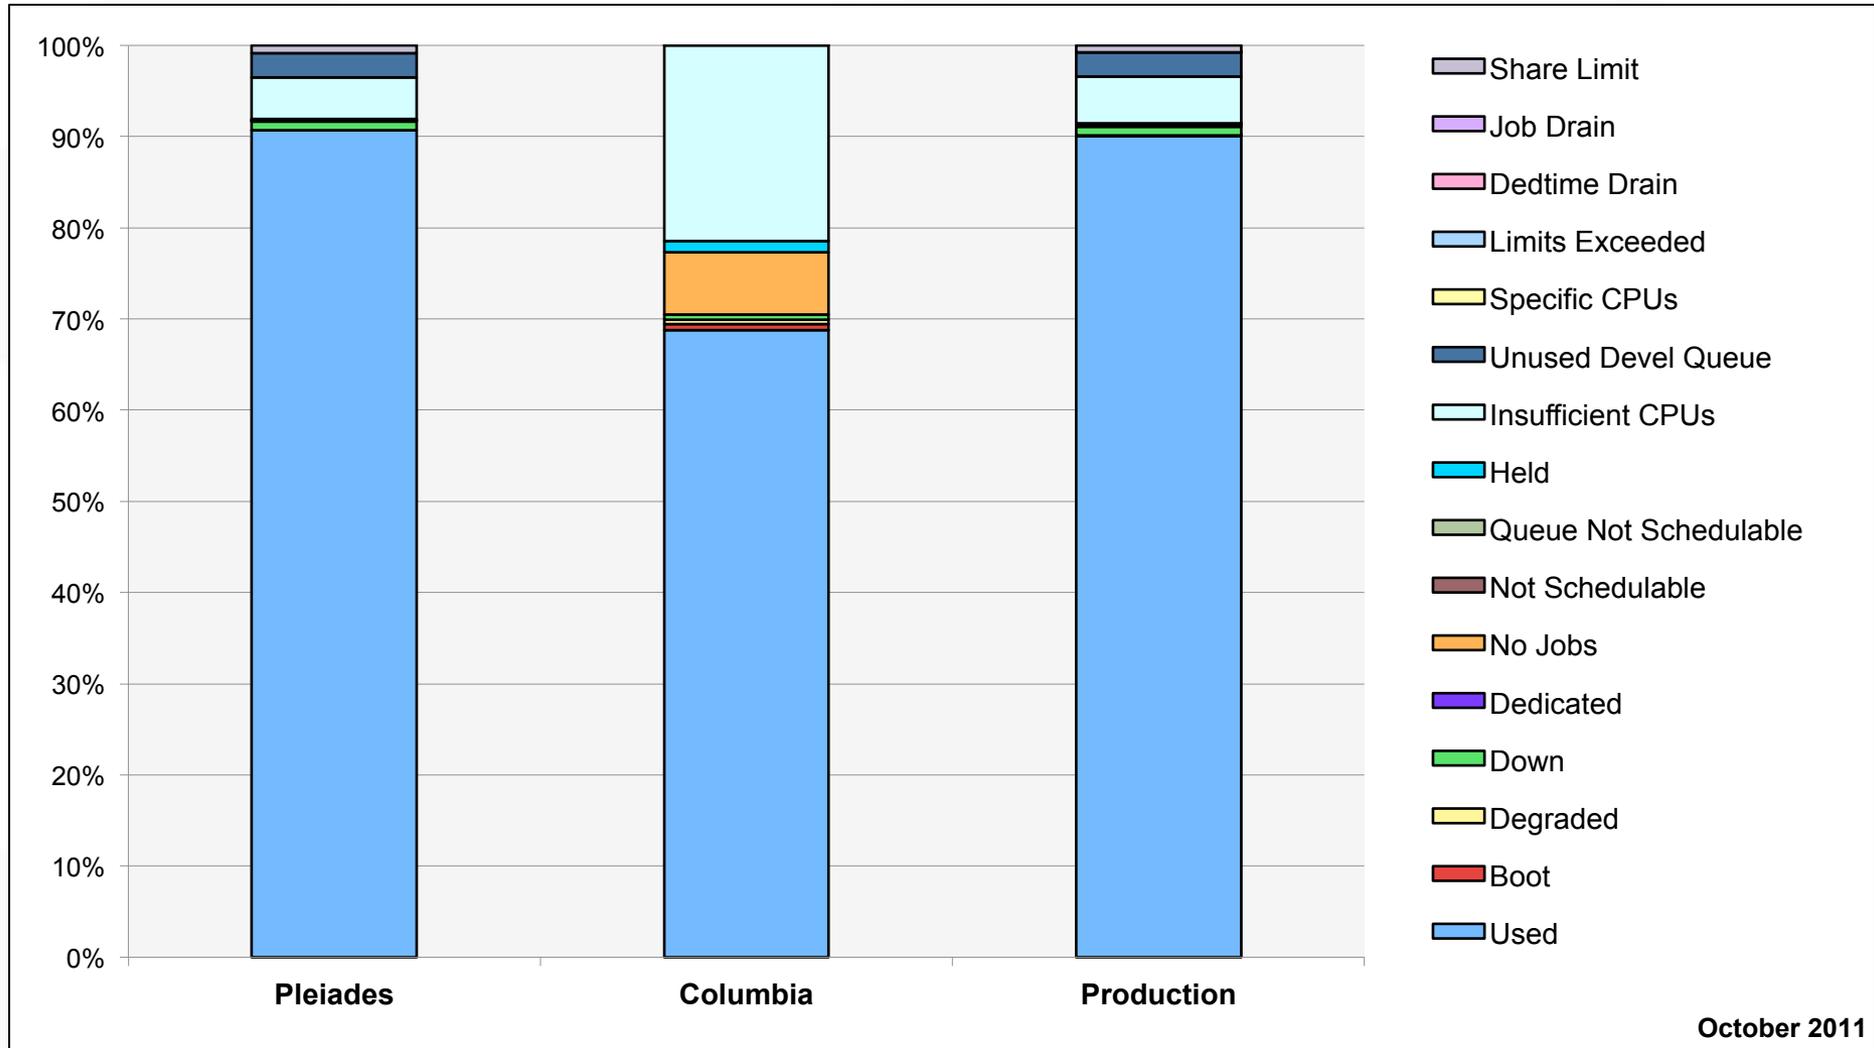
\* HECC provided supercomputing resources and services in support of this work

# News and Events



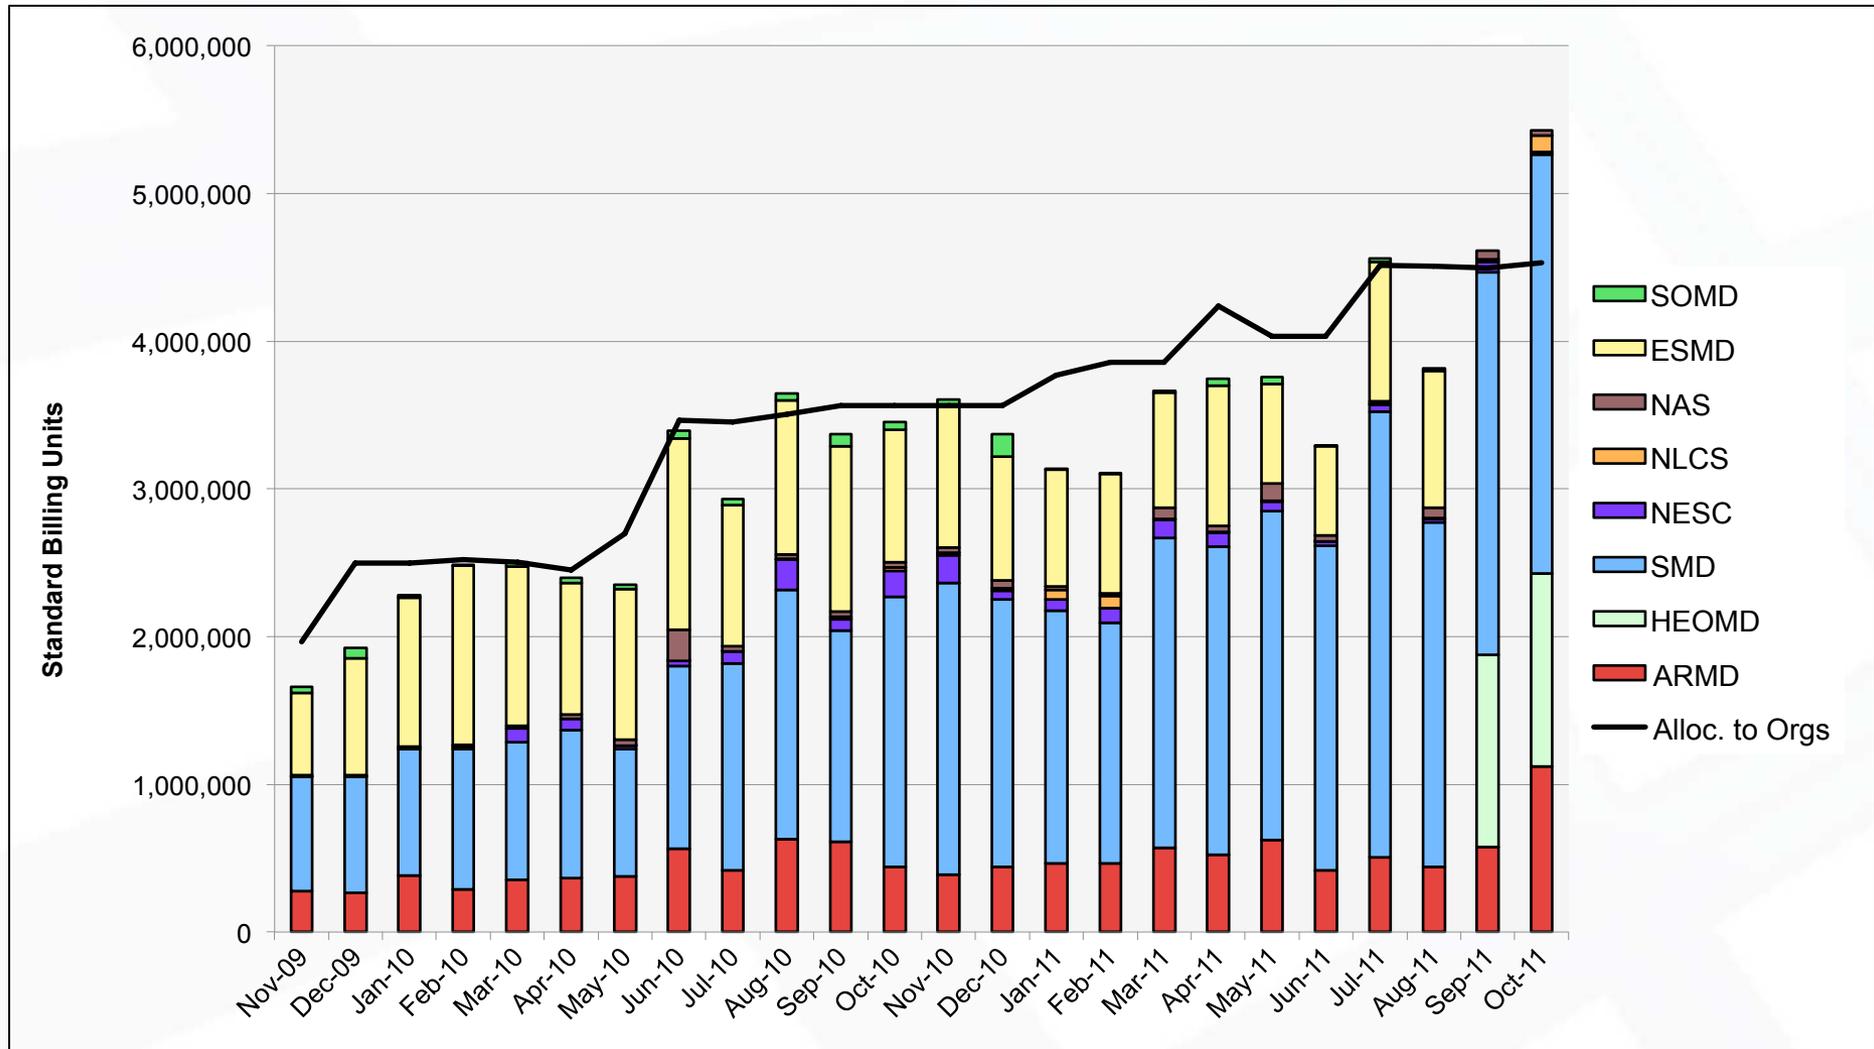
- **NASA Ames Inks Lustre Support Deal with Whamcloud**, *press release, Whamcloud*, October 24, 2011 – Whamcloud, Inc., has announced a Lustre support contract with the NASA Advanced Supercomputing (NAS) Division at Ames Research Center, to provide better operational continuity for NAS's high-performance computing environment. Picked up by several media outlets, including SpaceRef, MarketWatch, and local newspapers.  
<http://www.whamcloud.com/news-and-events/nasa-ames-inks-lustre-support-deal-with-whamcloud/>
- **Oak Ridge Labs Builds Fastest Supercomputer**, *article, Information Week*, October 11, 2011 – Mentions Pleiades as one of five top systems operated by the Federal Government.  
<http://www.informationweek.com/news/government/enterprise-architecture/231900554>

# HECC Utilization

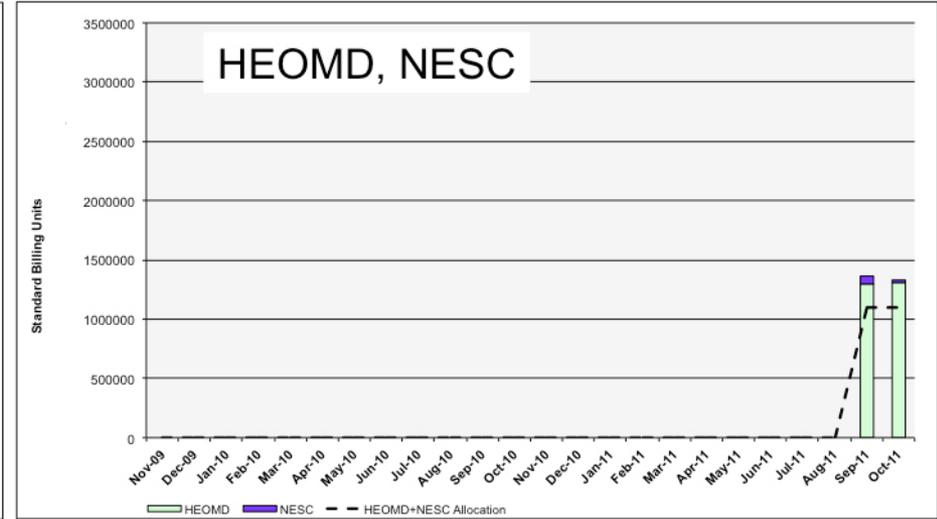
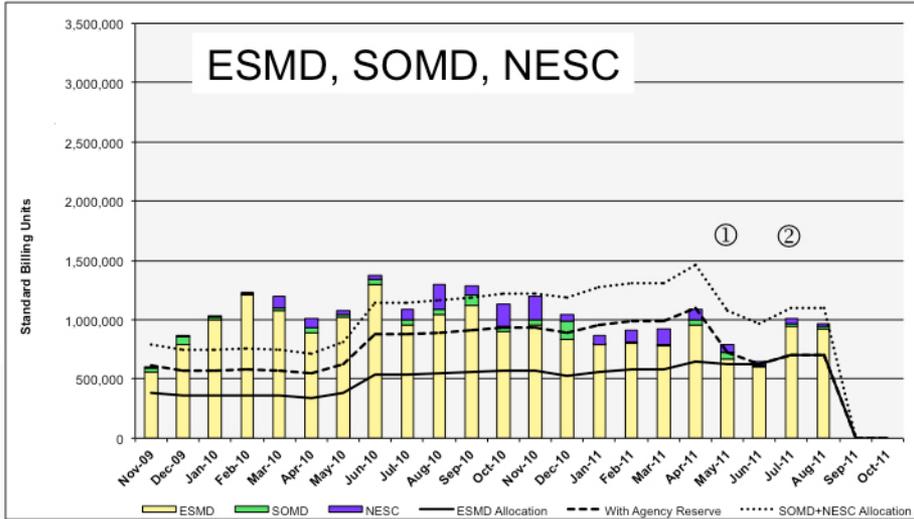
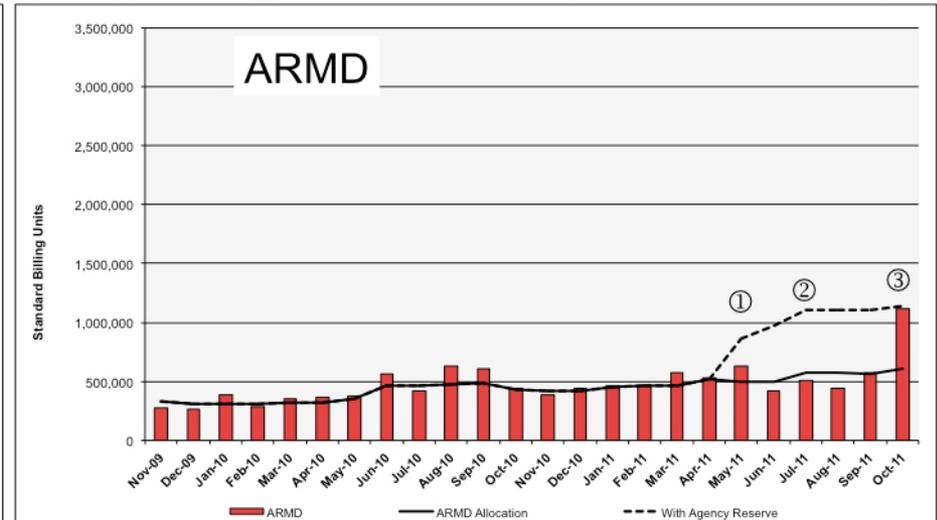
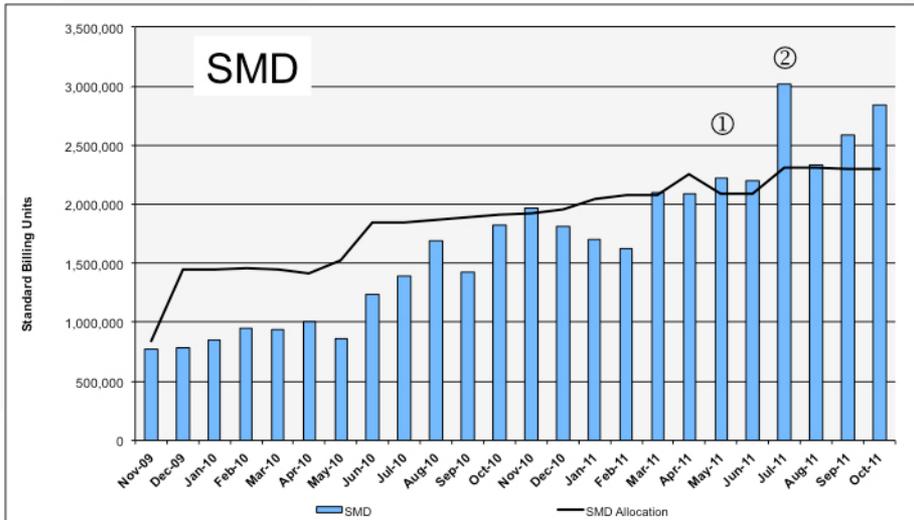


October 2011

# HECC Utilization Normalized to 30-Day Month

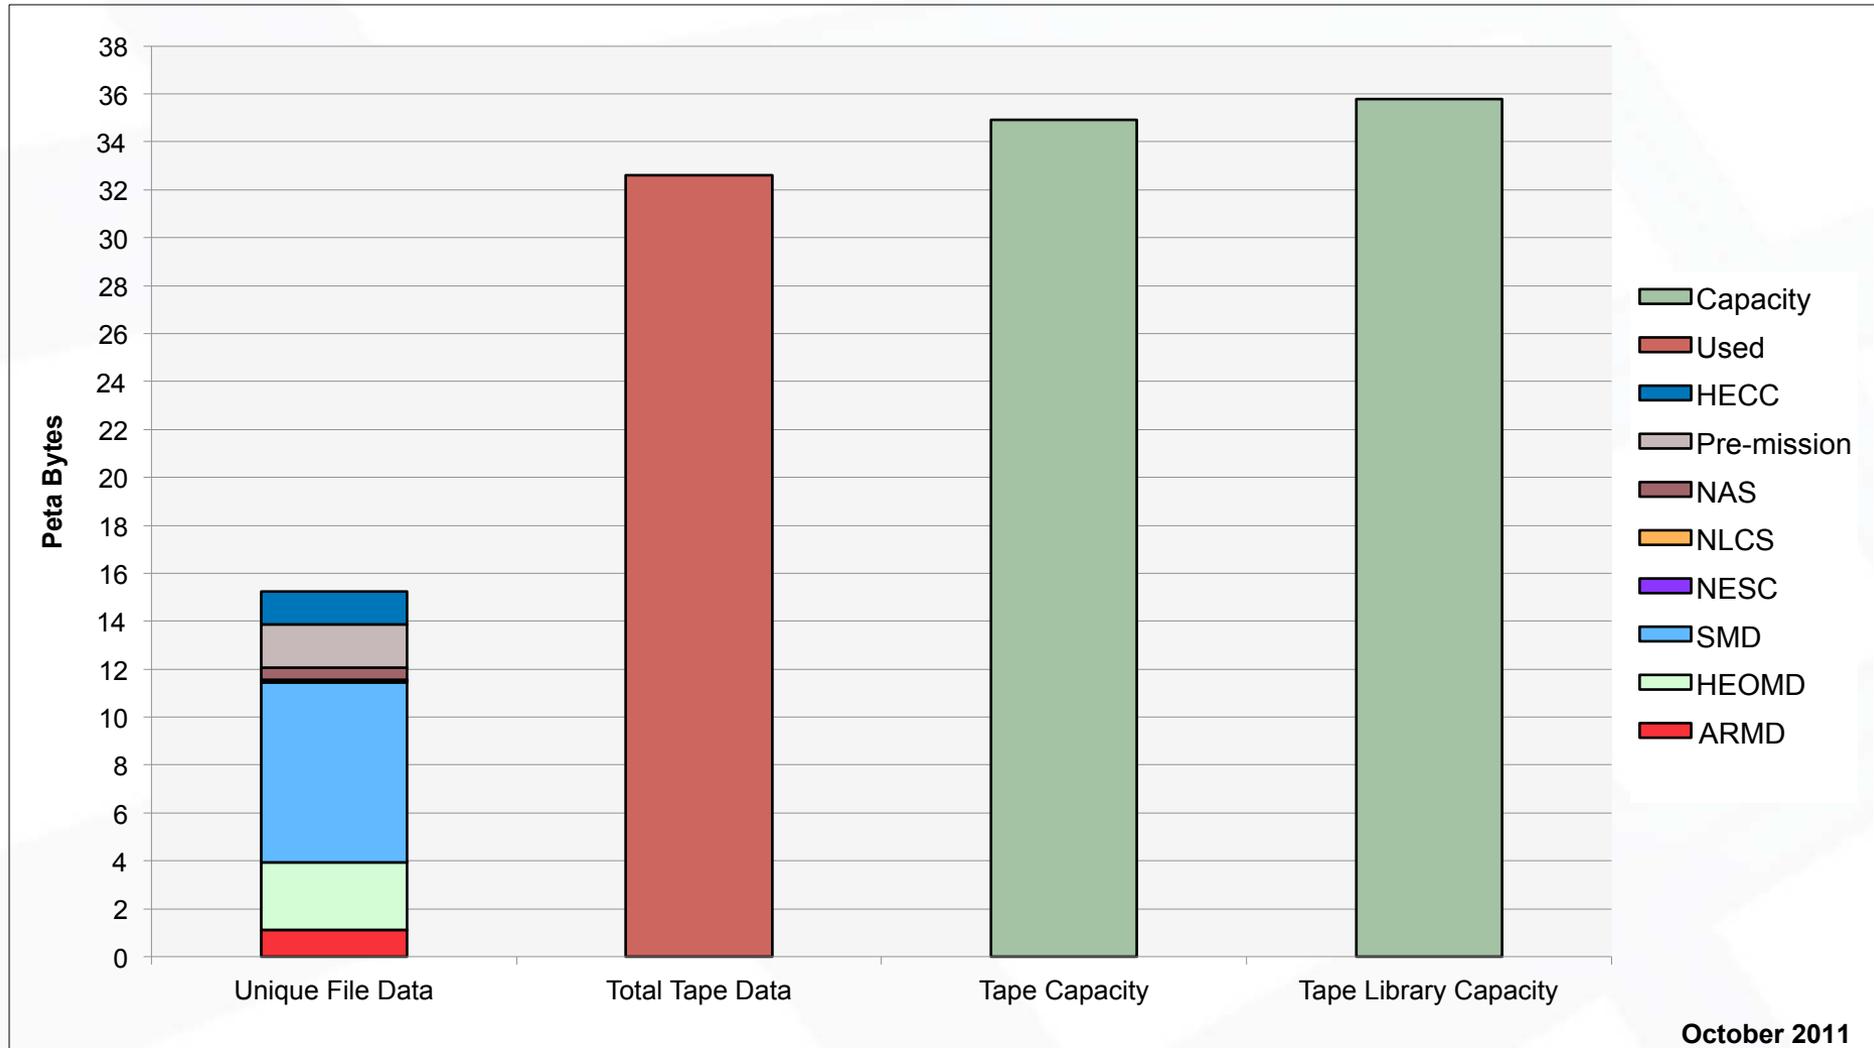


# HECC Utilization Normalized to 30-Day Month



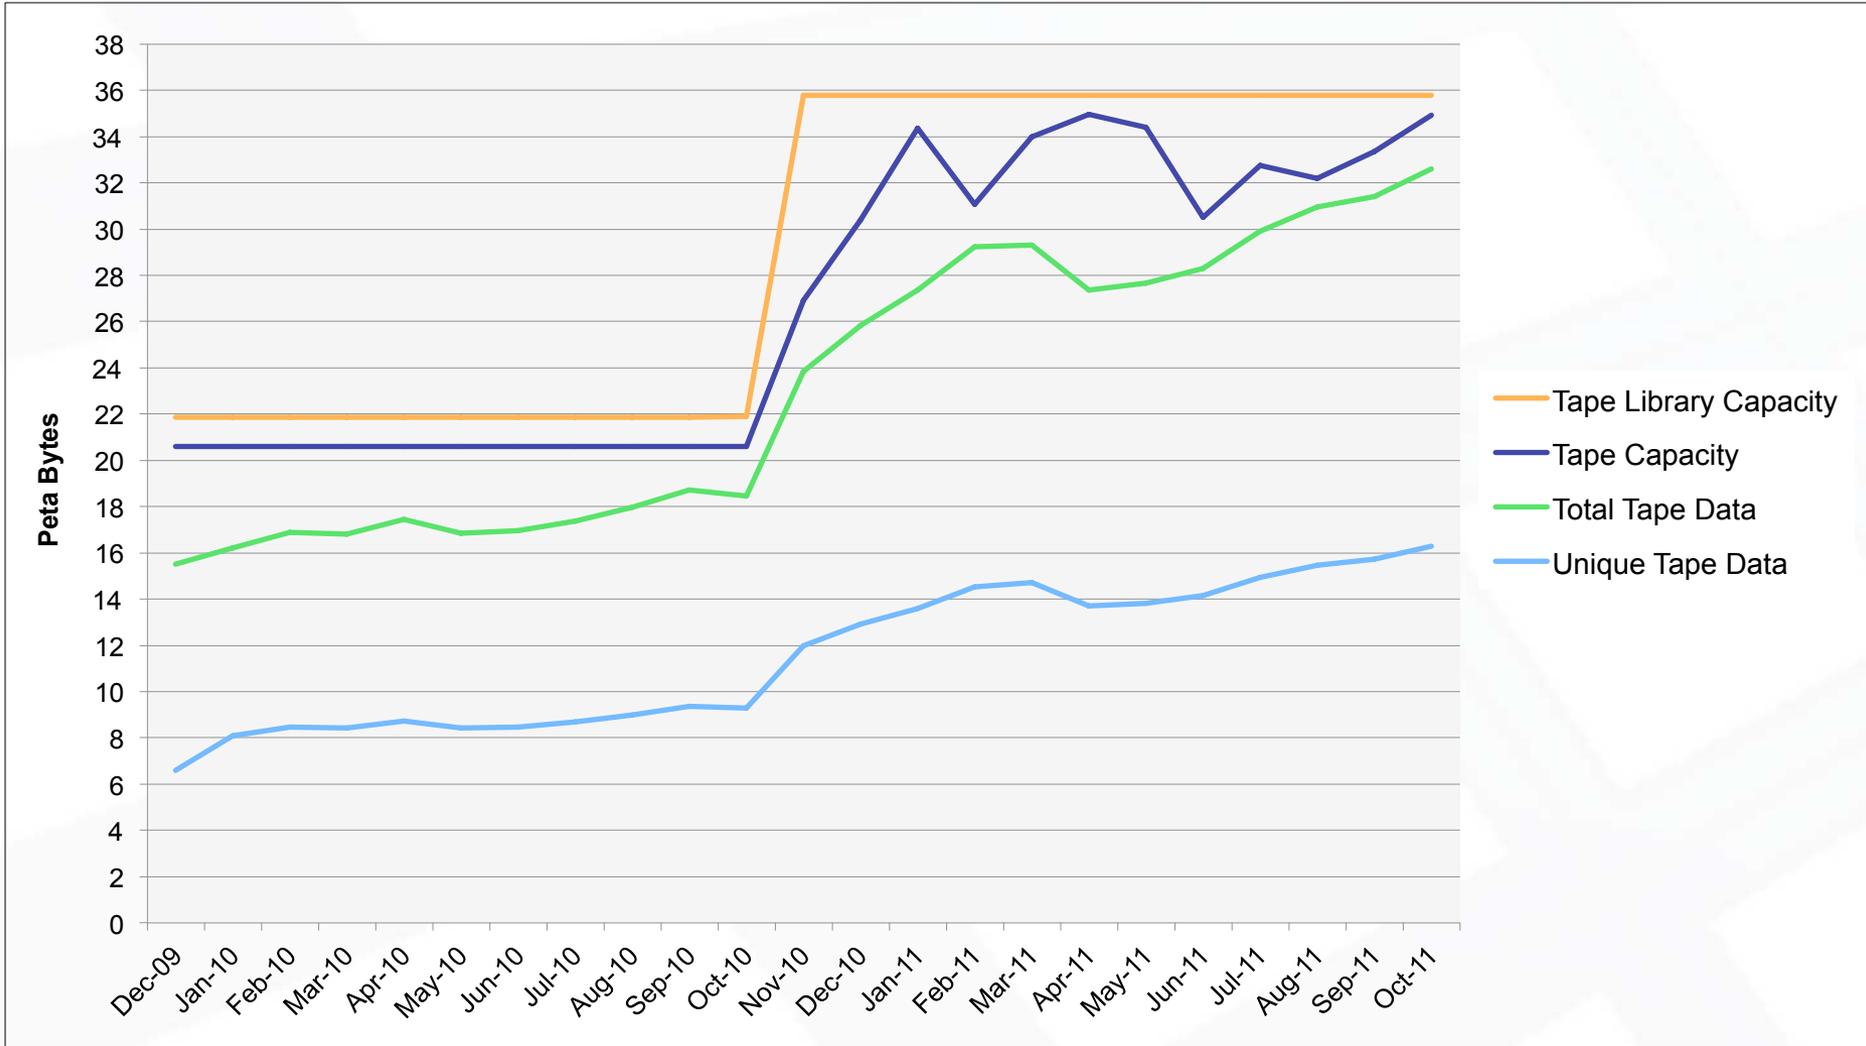
① Allocation to orgs. decreased to 75%, Agency reserve shifted to ARMD ② 14 Westmere racks added ③ 2 ARMD Westmere racks added

# Tape Archive Status

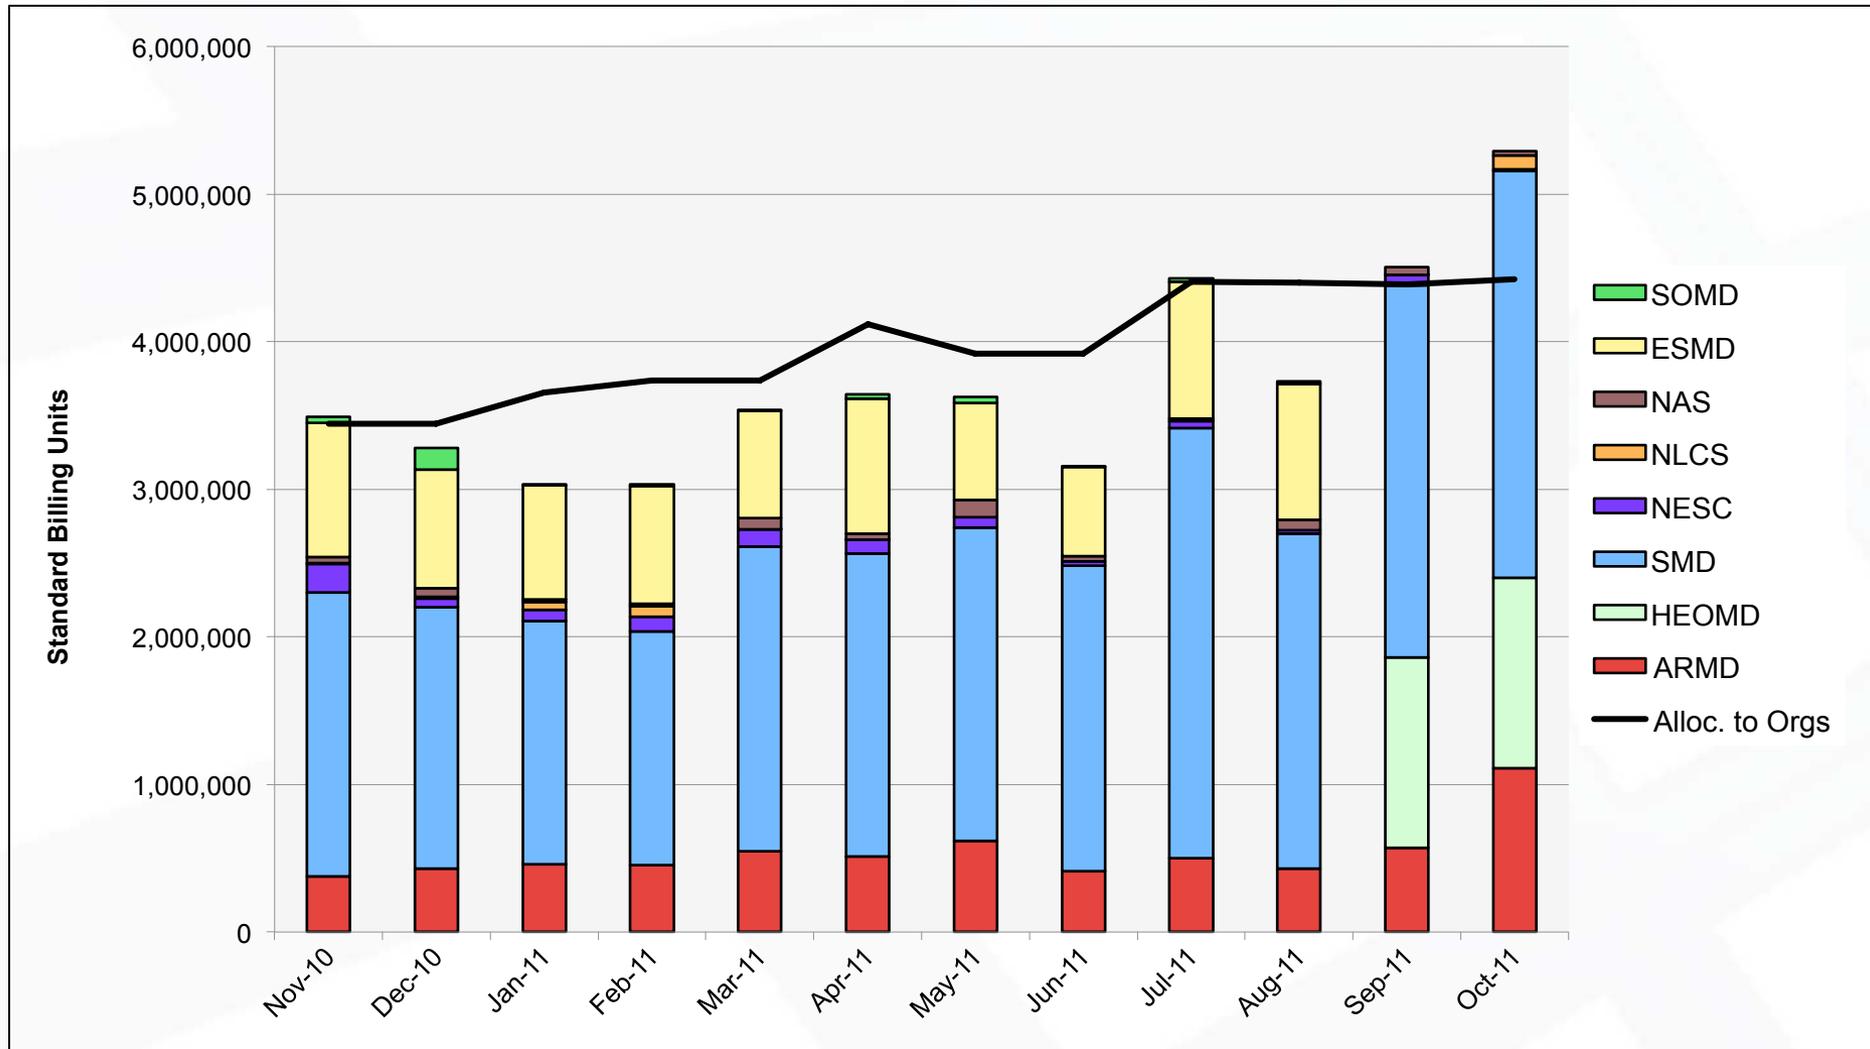


October 2011

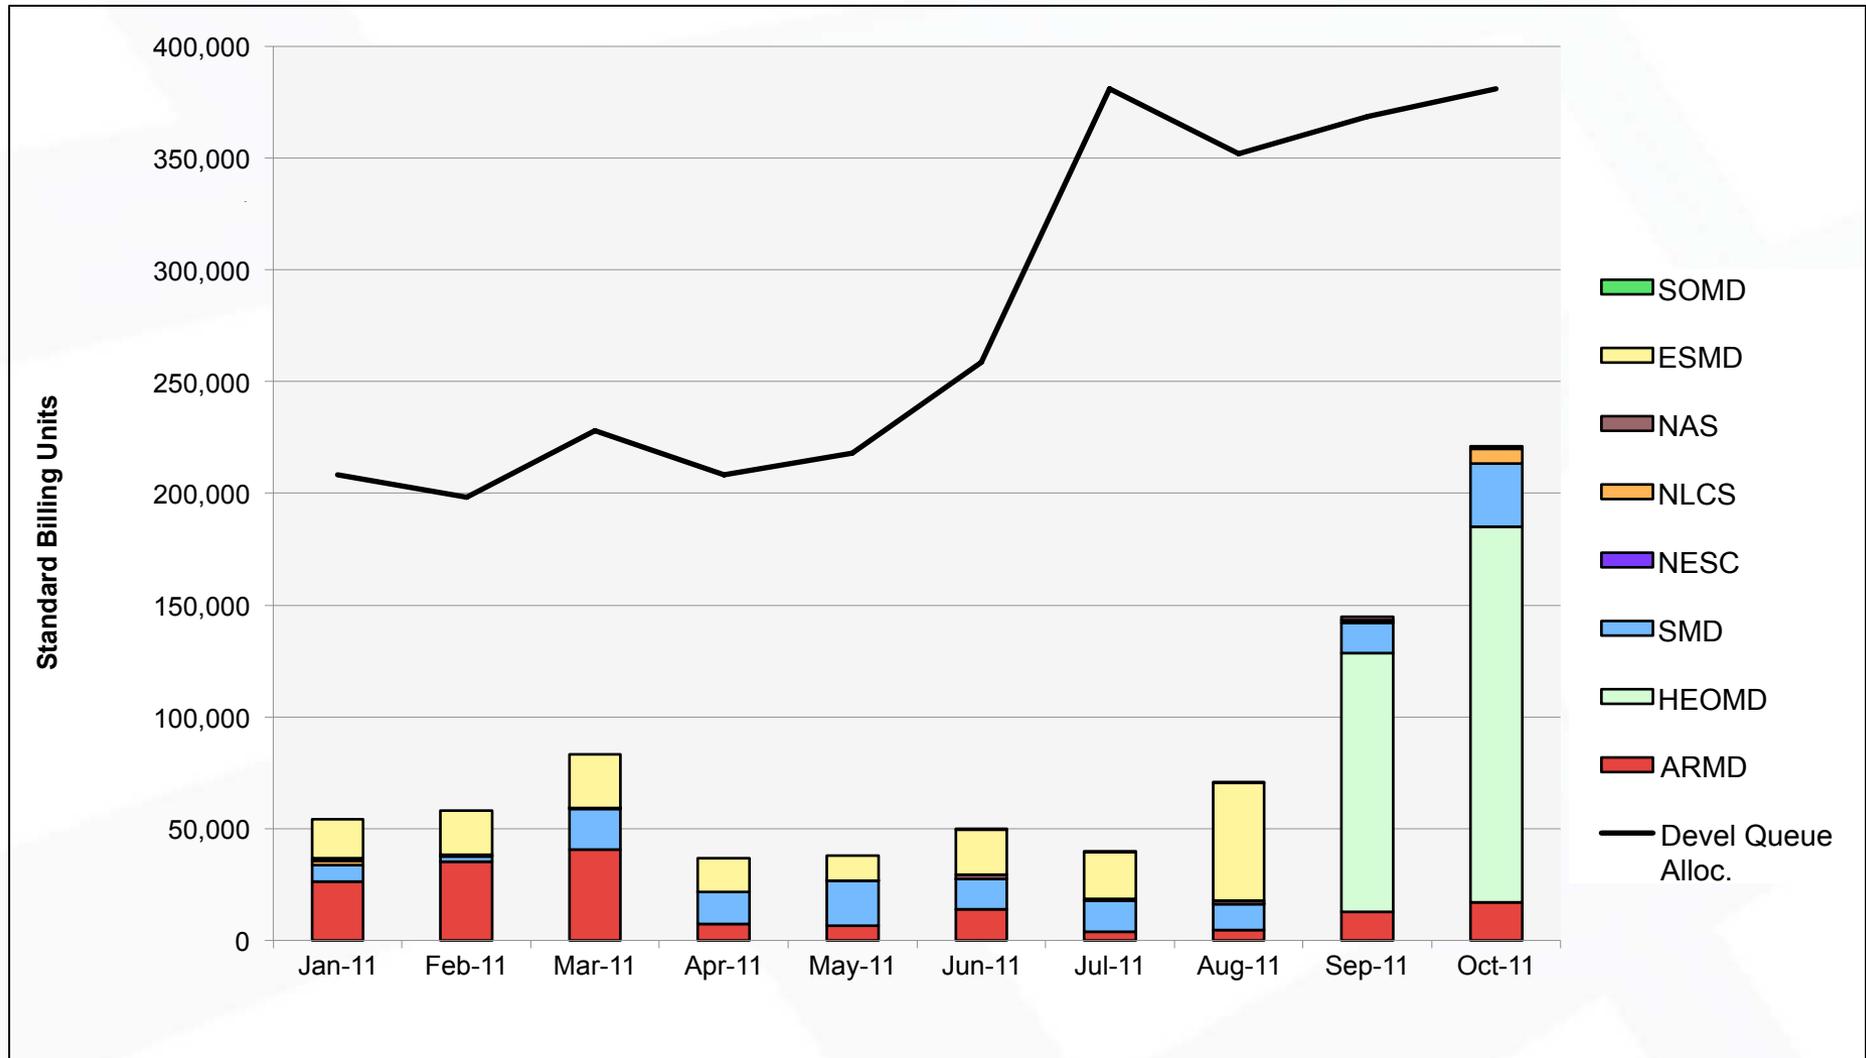
# Tape Archive Status



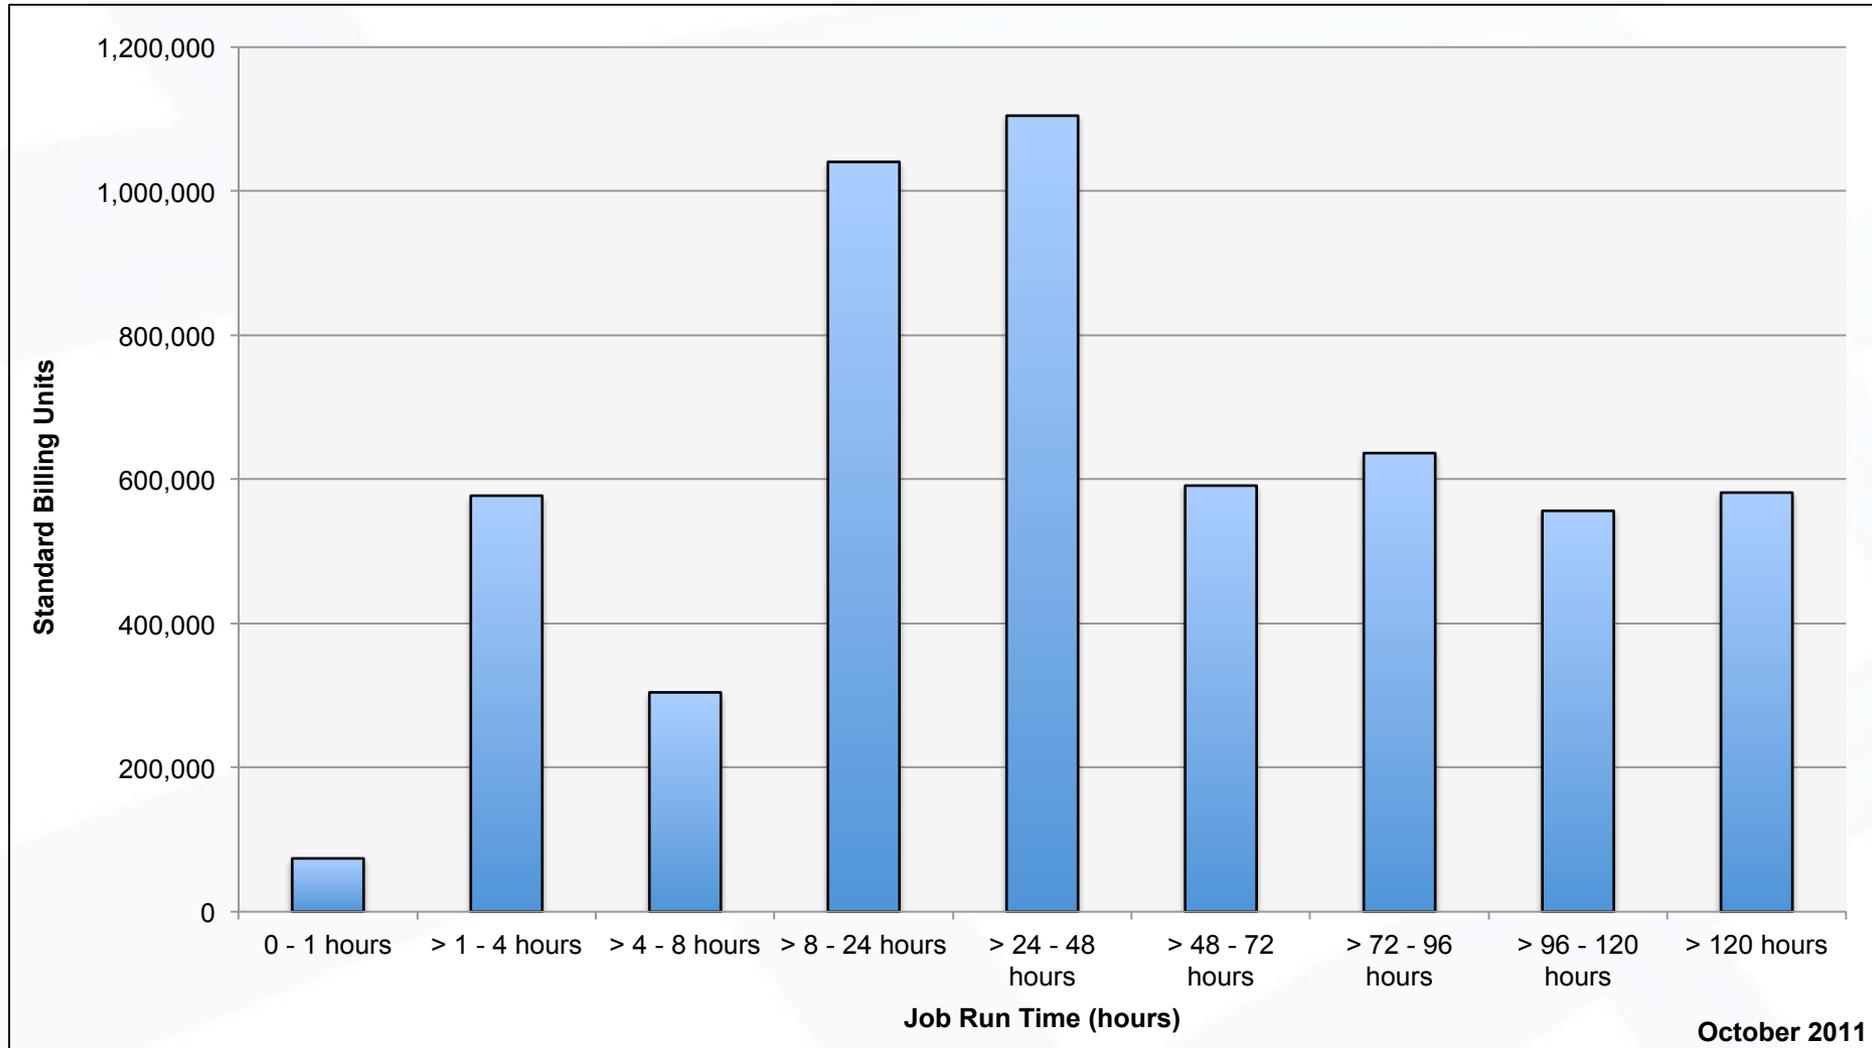
# Pleiades: SBUs Reported, Normalized to 30-Day Month



# Pleiades: Devel Queue Utilization

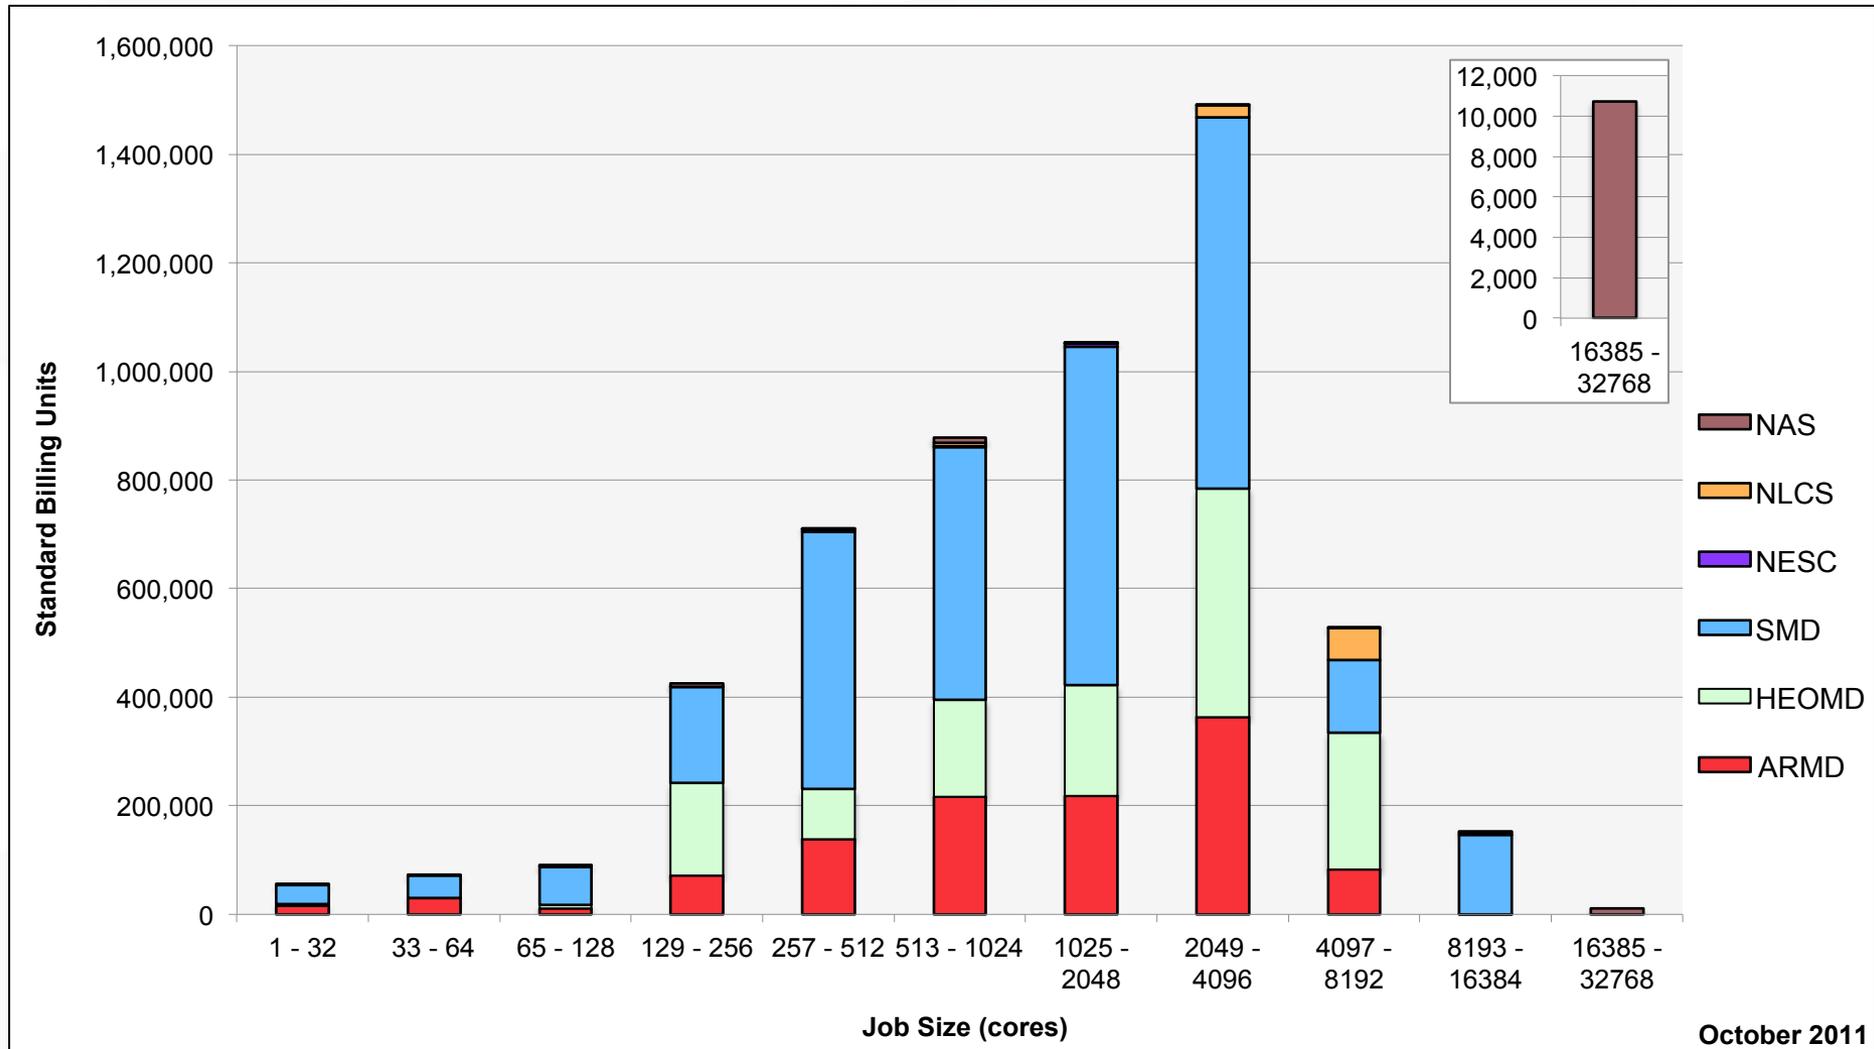


# Pleiades: Monthly SBUs by Run Time



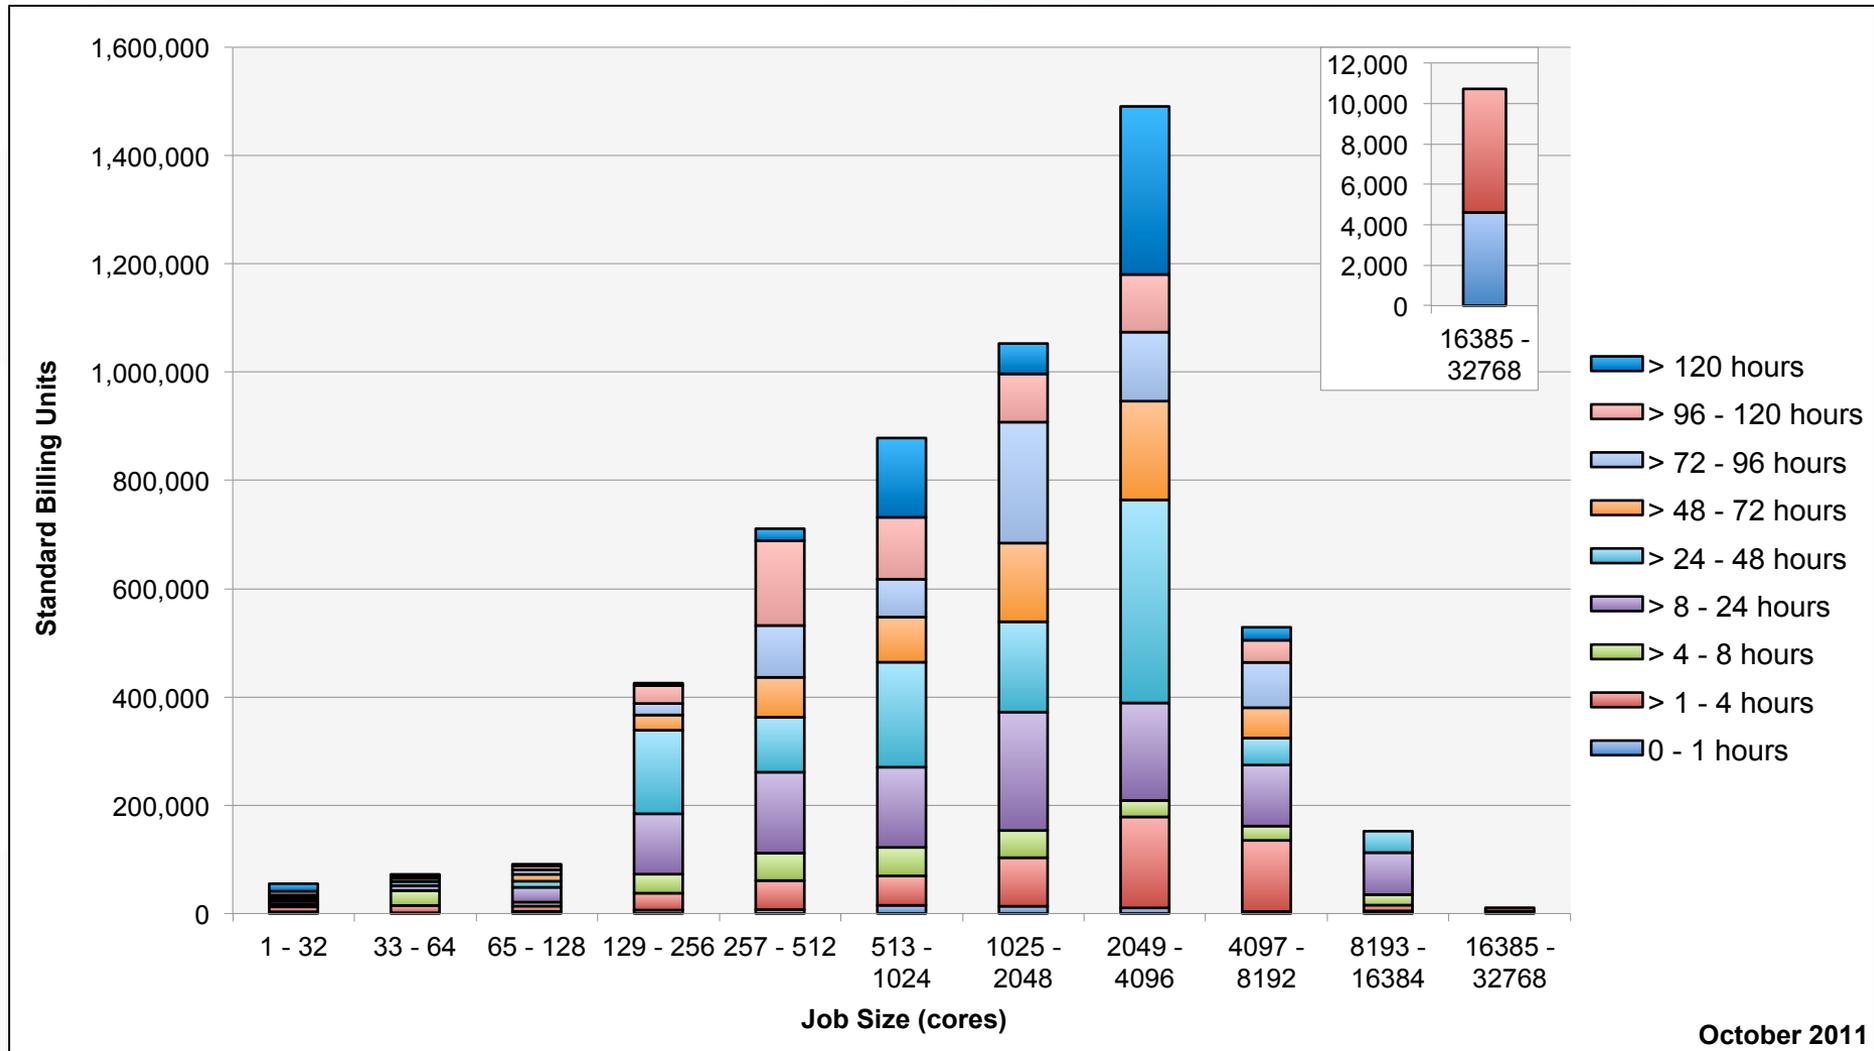
October 2011

# Pleiades: Monthly Utilization by Size and Mission



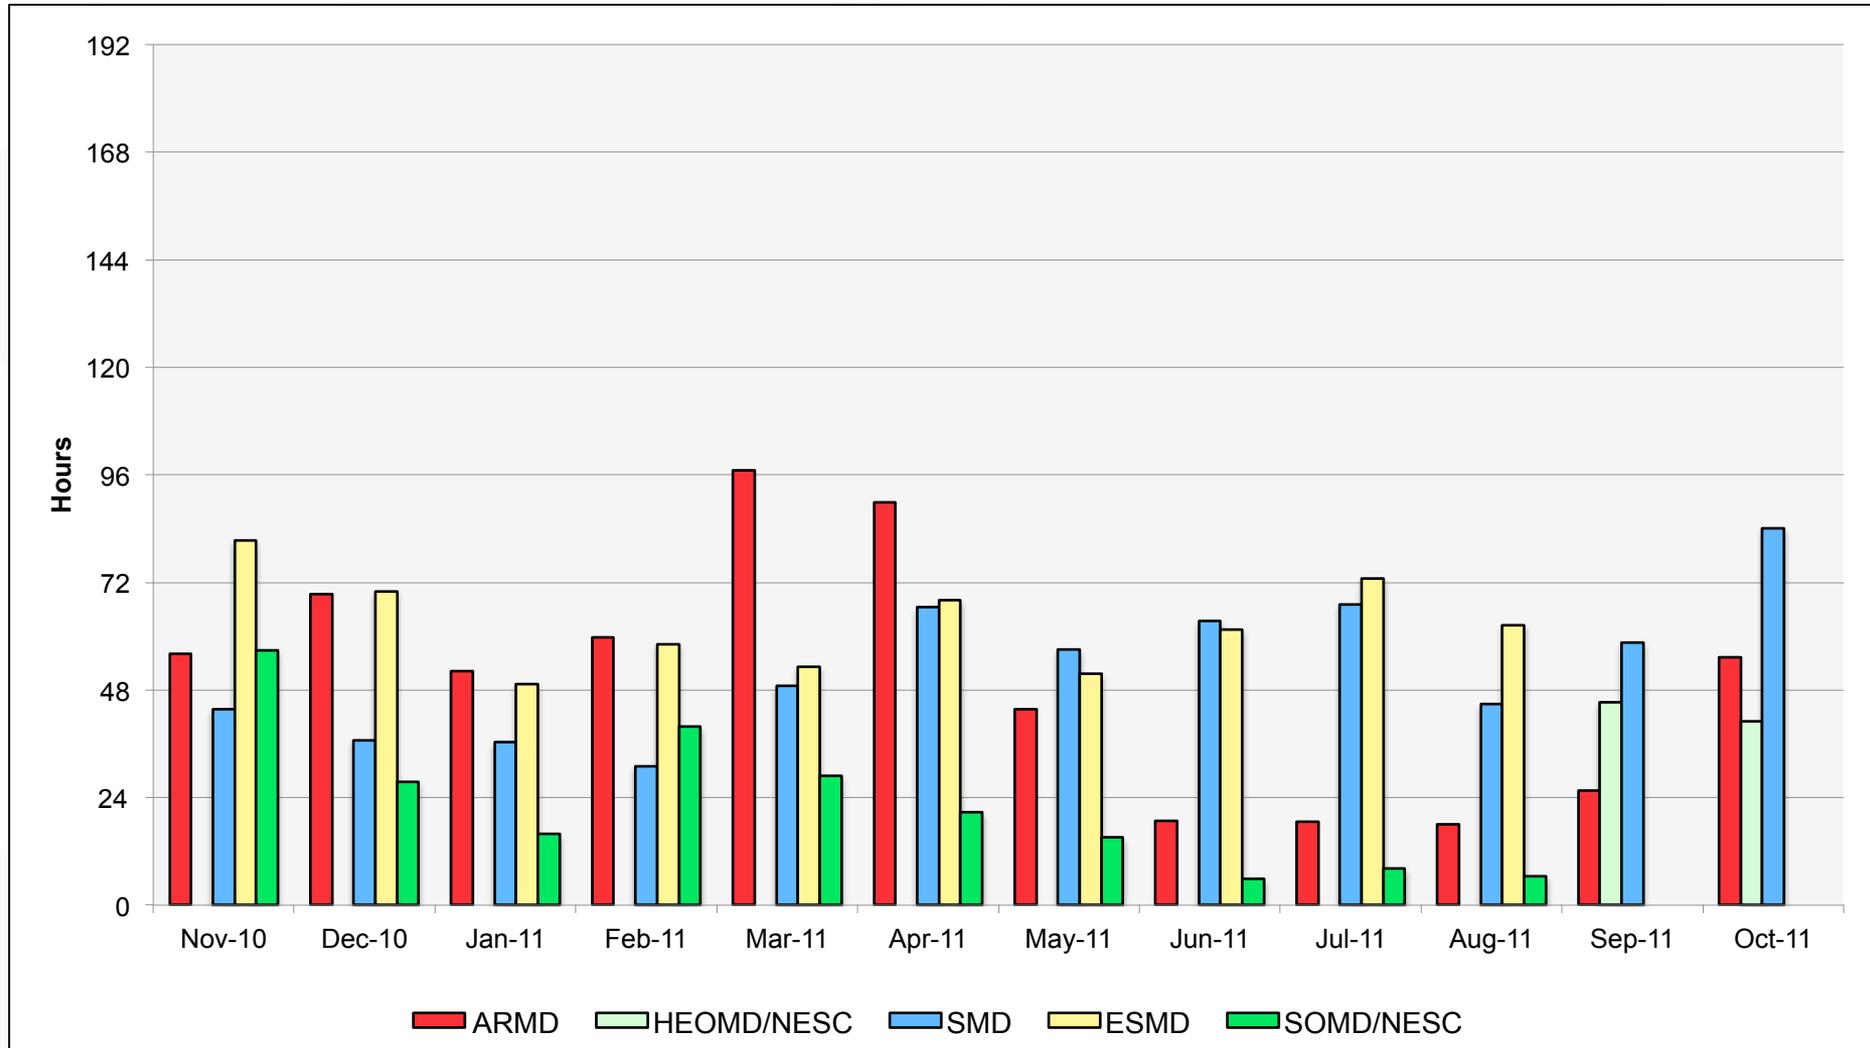
October 2011

# Pleiades: Monthly Utilization by Size and Length

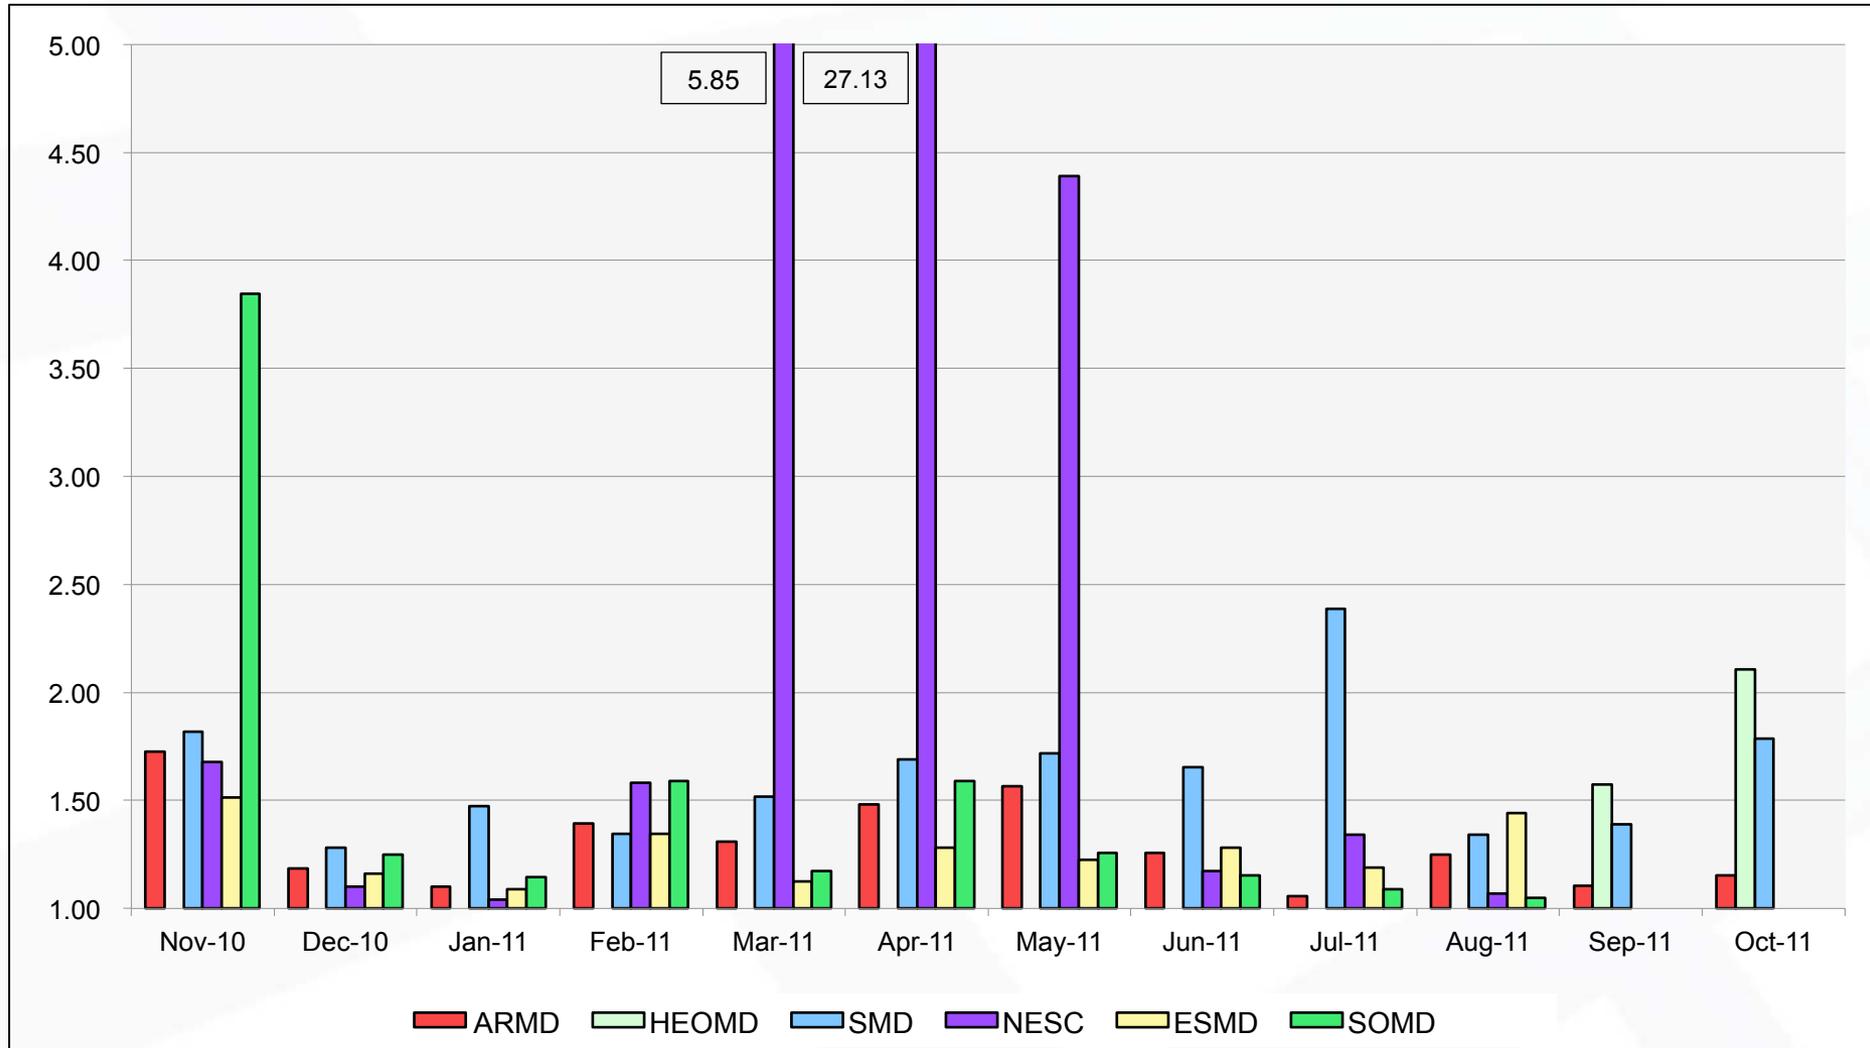


October 2011

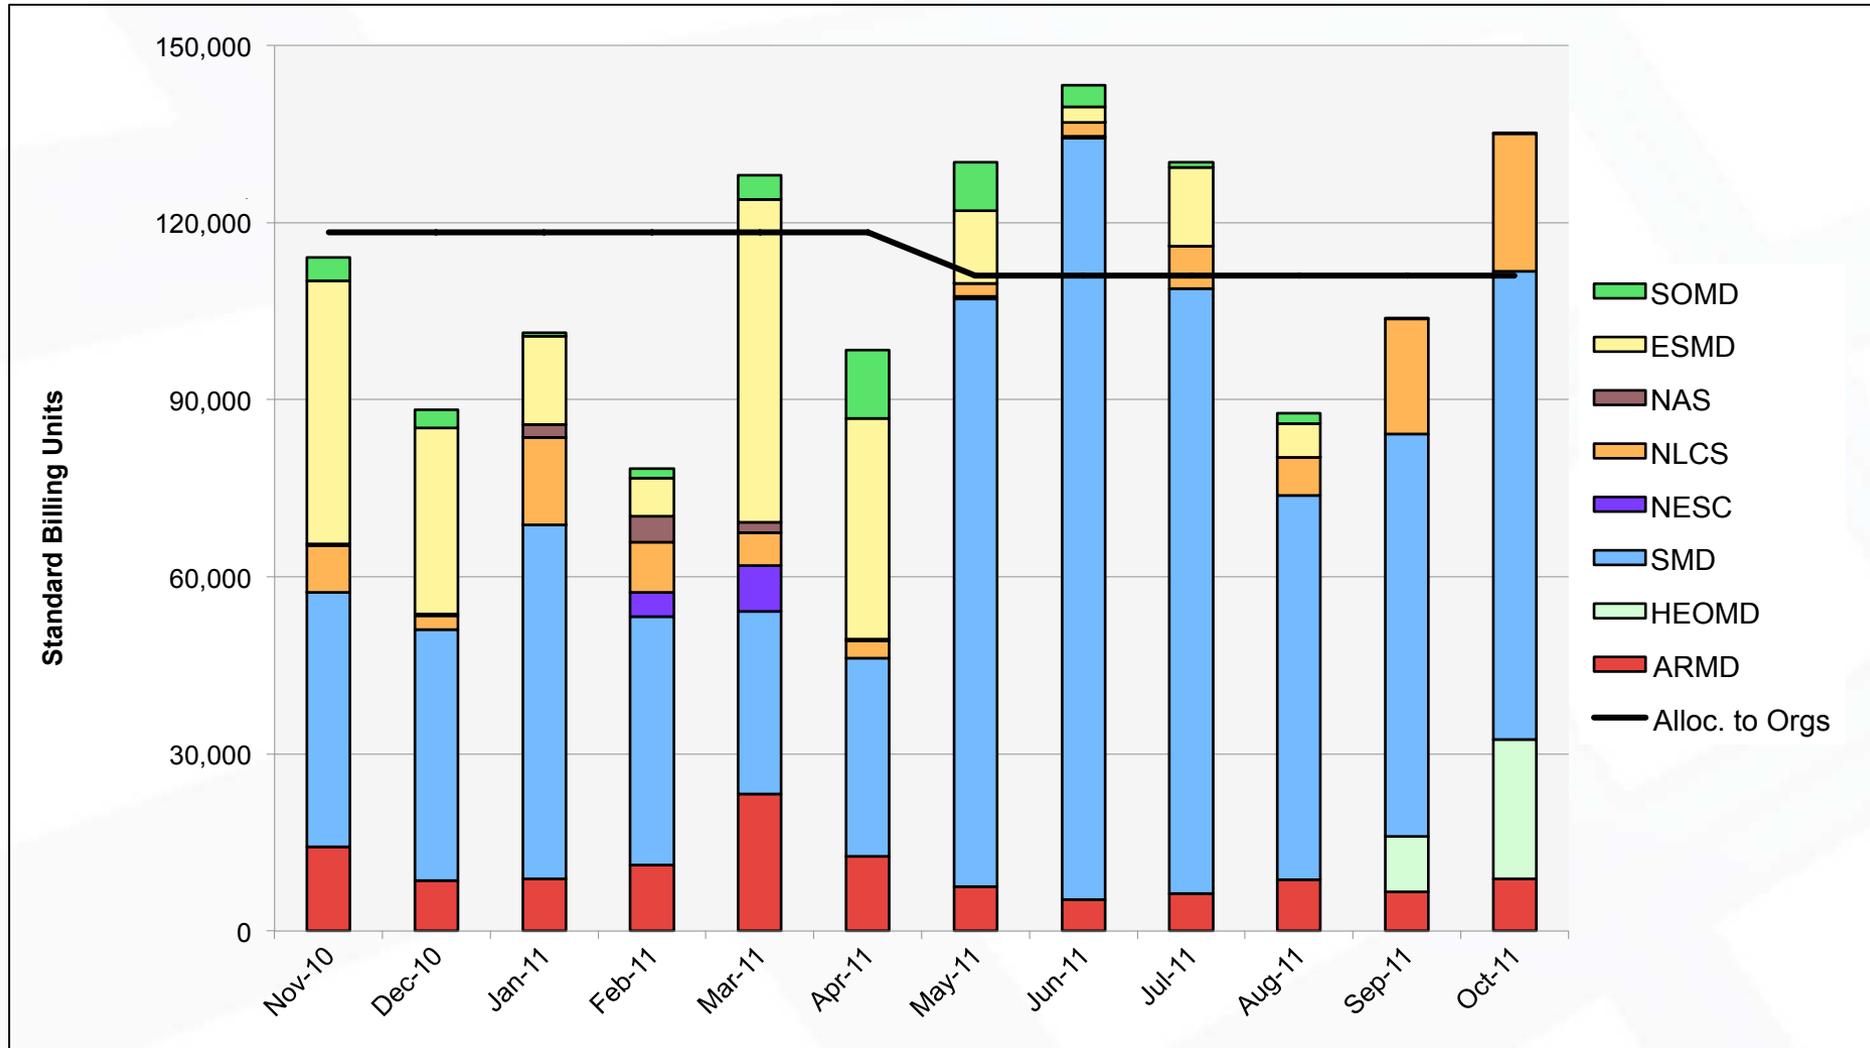
# Pleiades: Average Time to Clear All Jobs



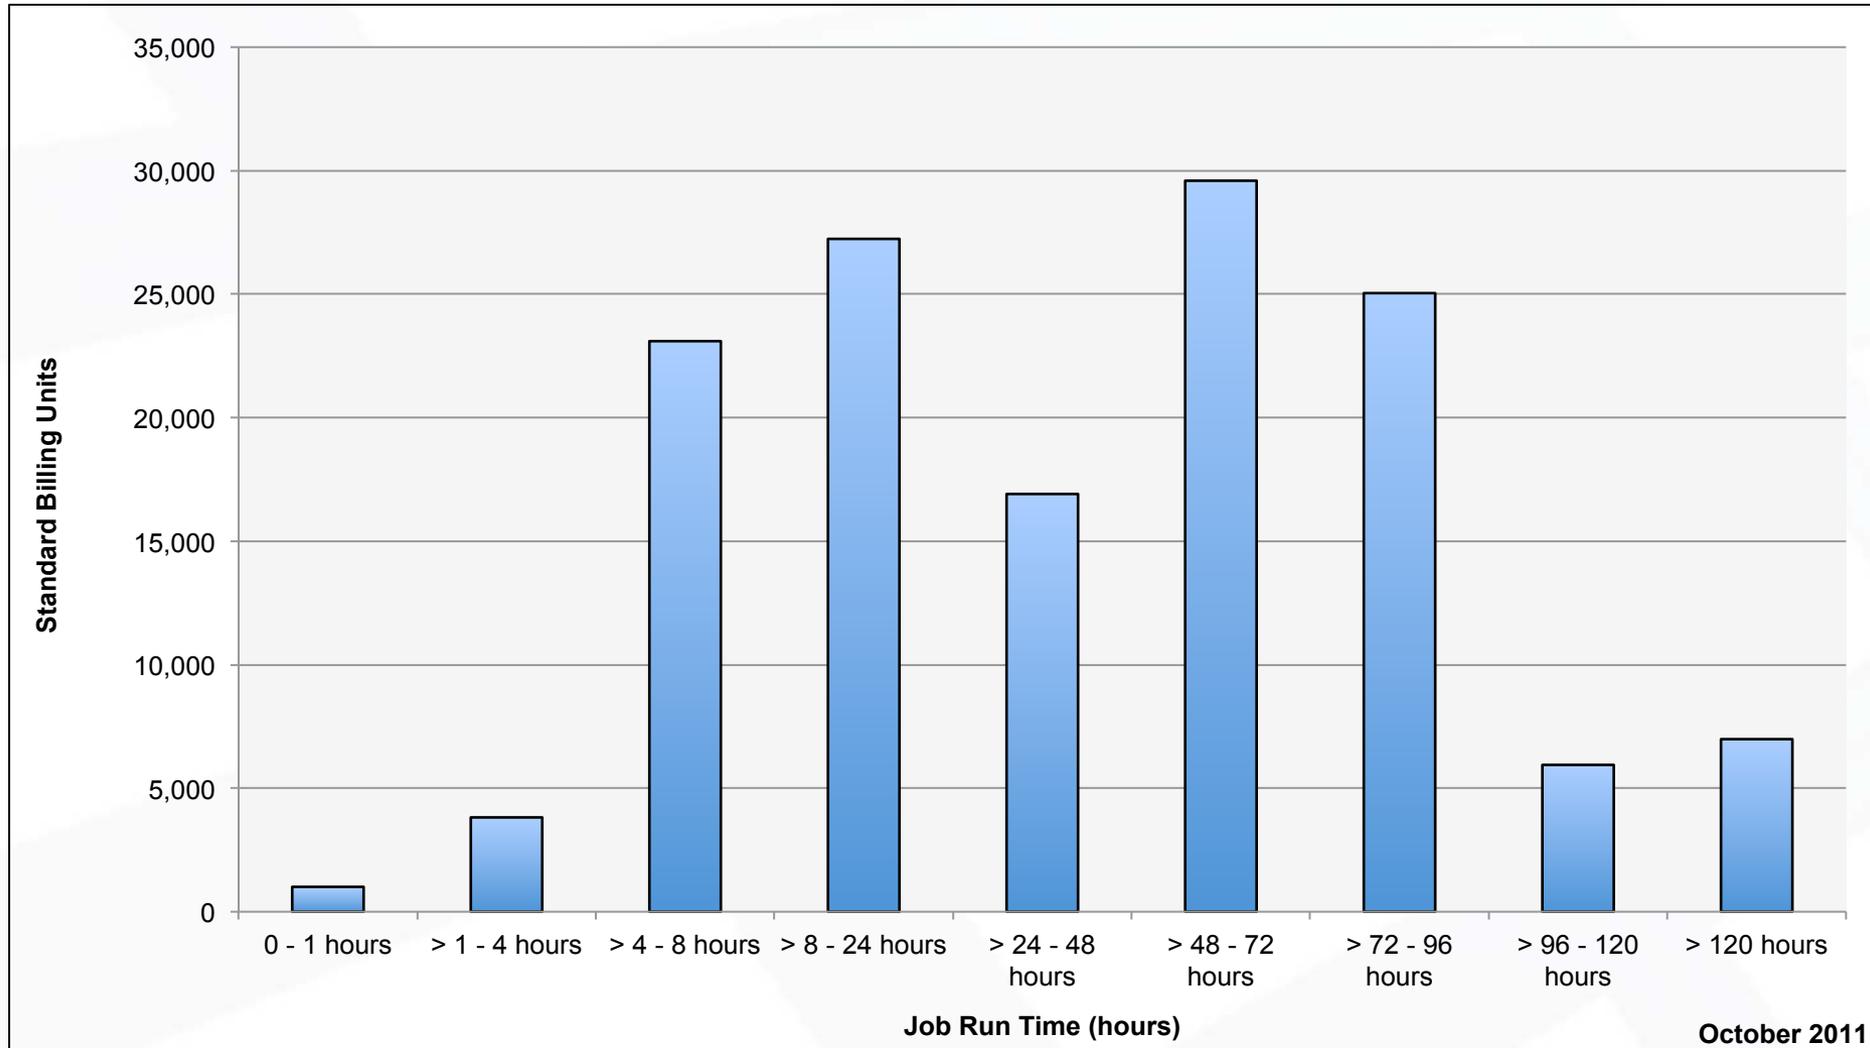
# Pleiades: Average Expansion Factor



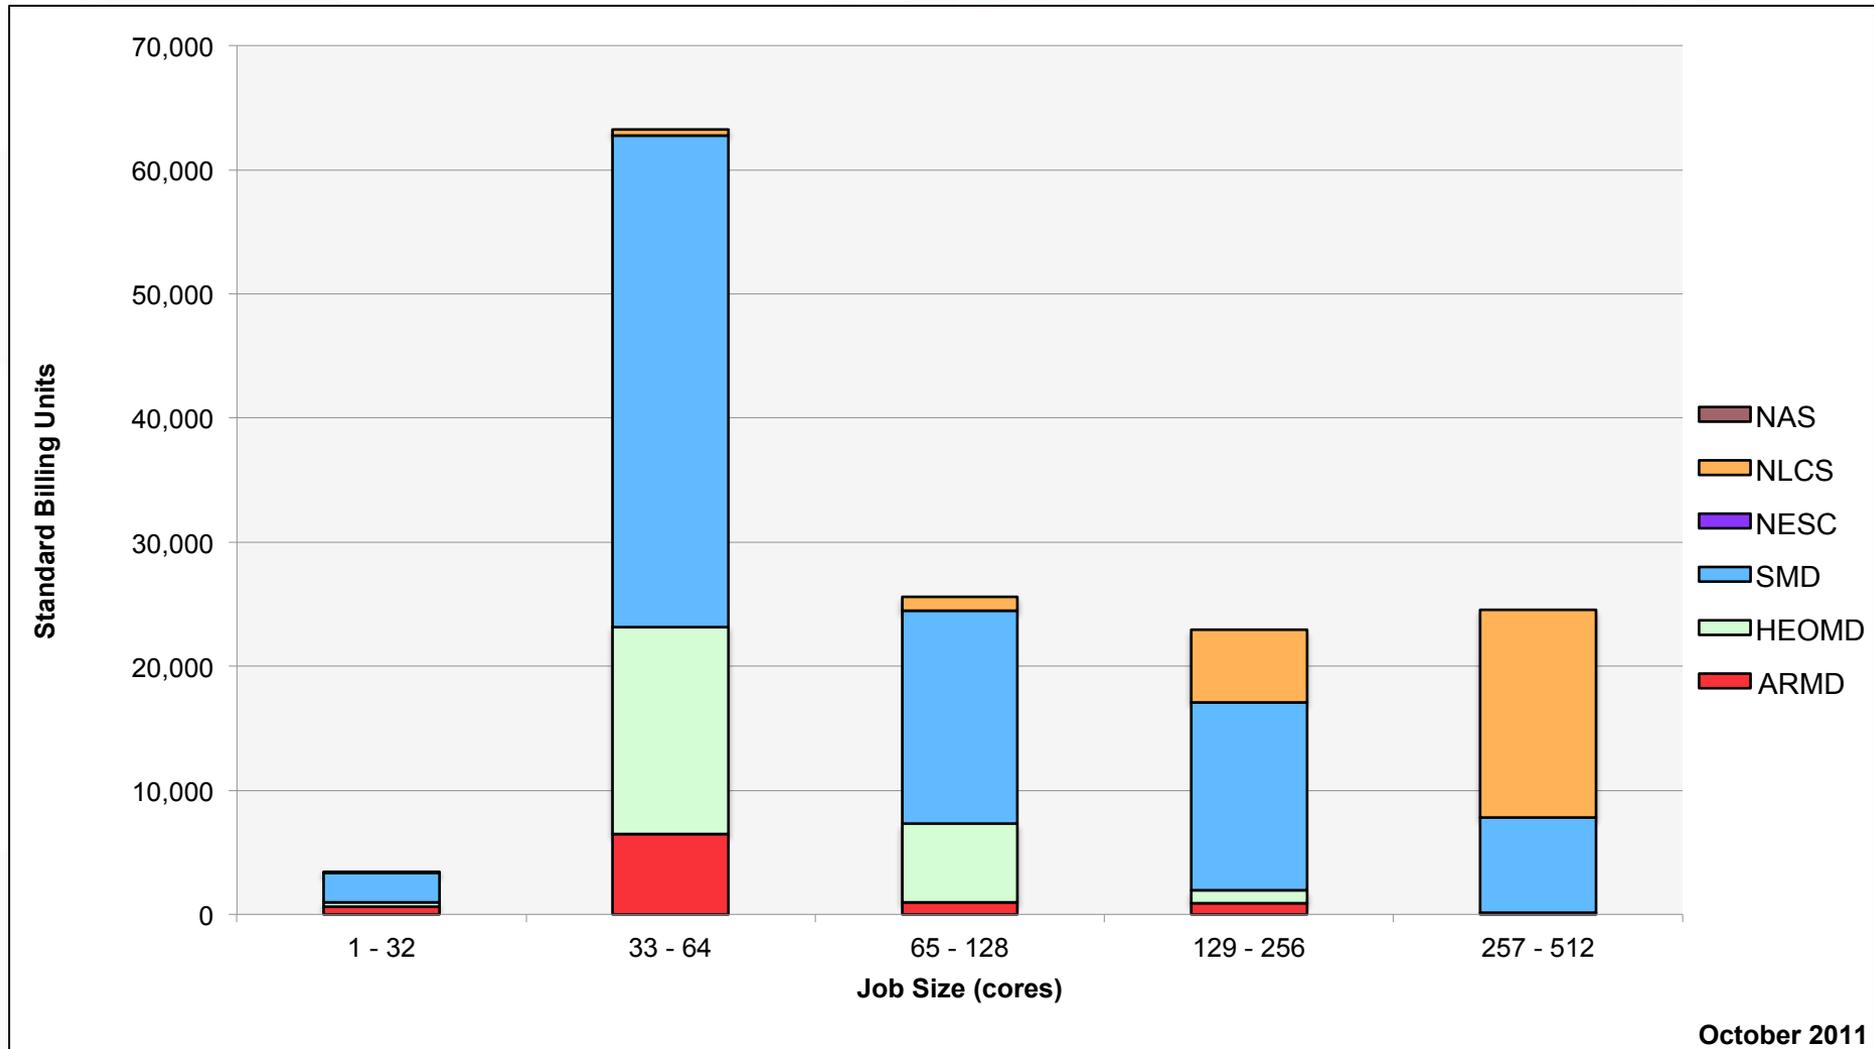
# Columbia: SBUs Reported, Normalized to 30-Day Month



# Columbia: Monthly SBUs by Run Time

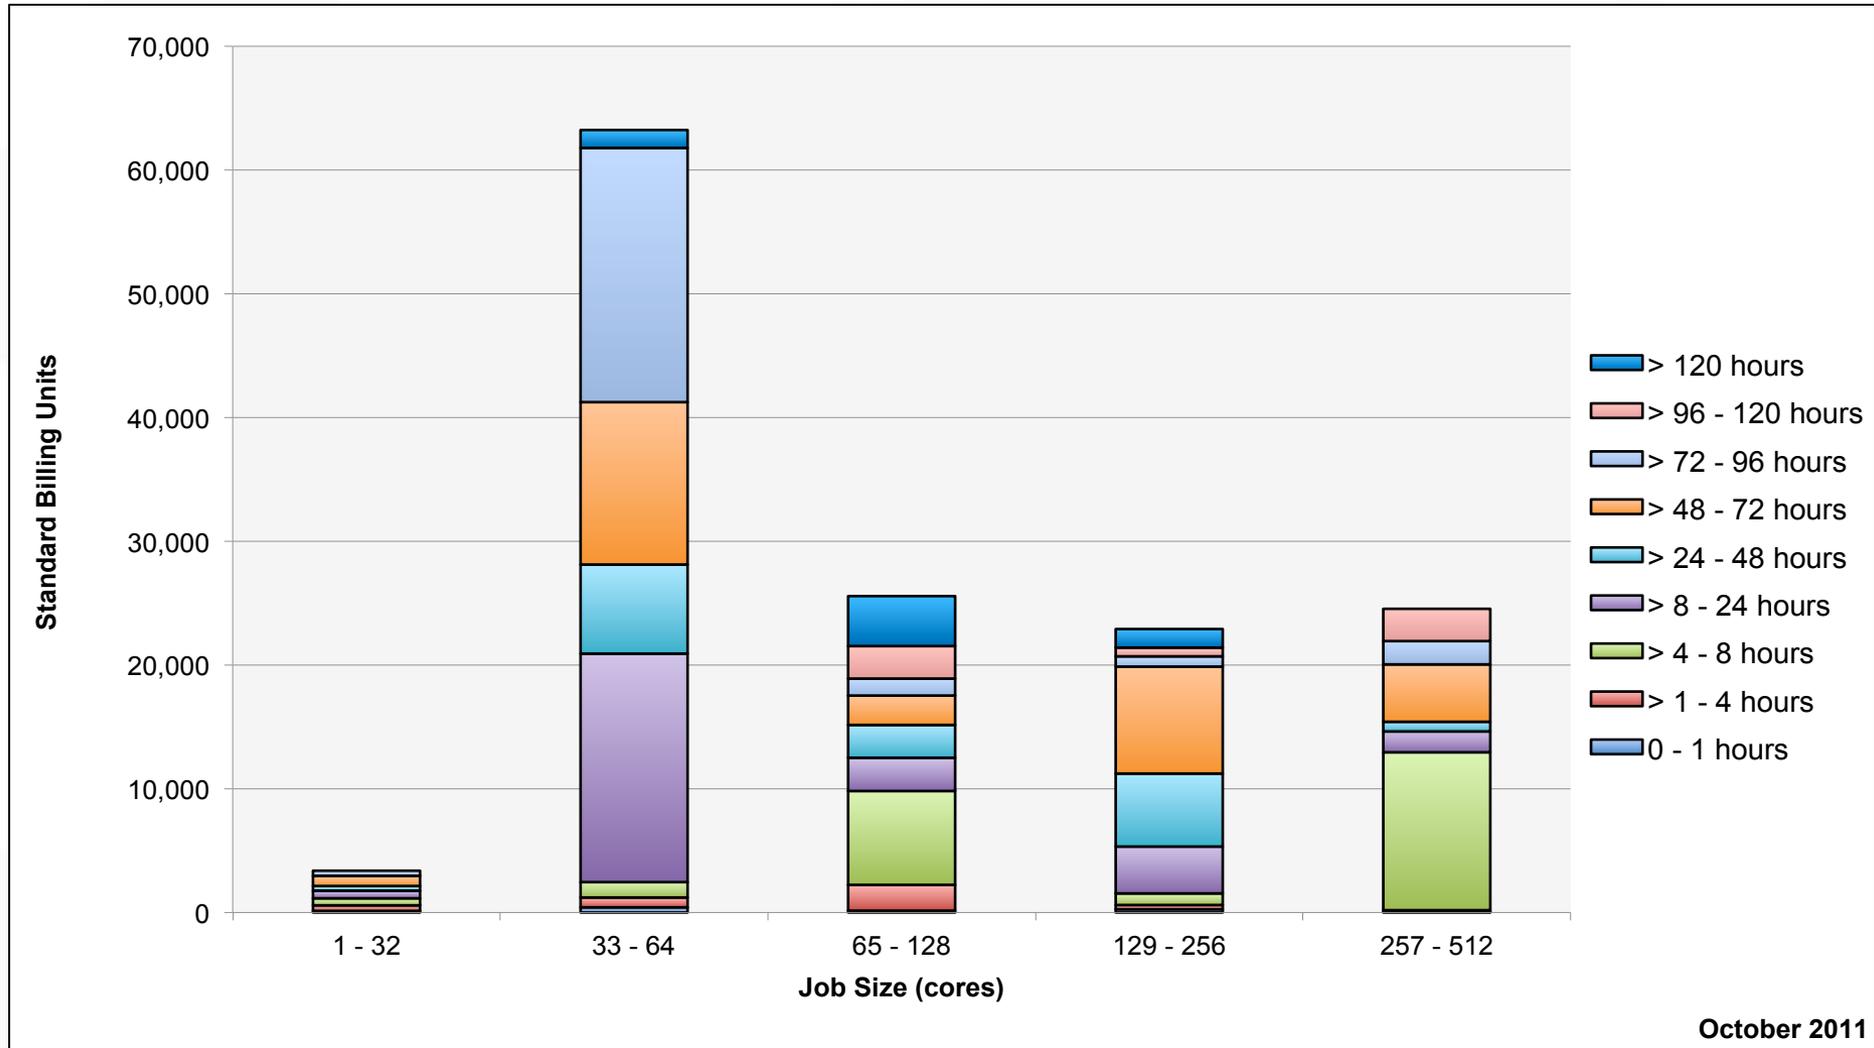


# Columbia: Monthly Utilization by Size and Mission



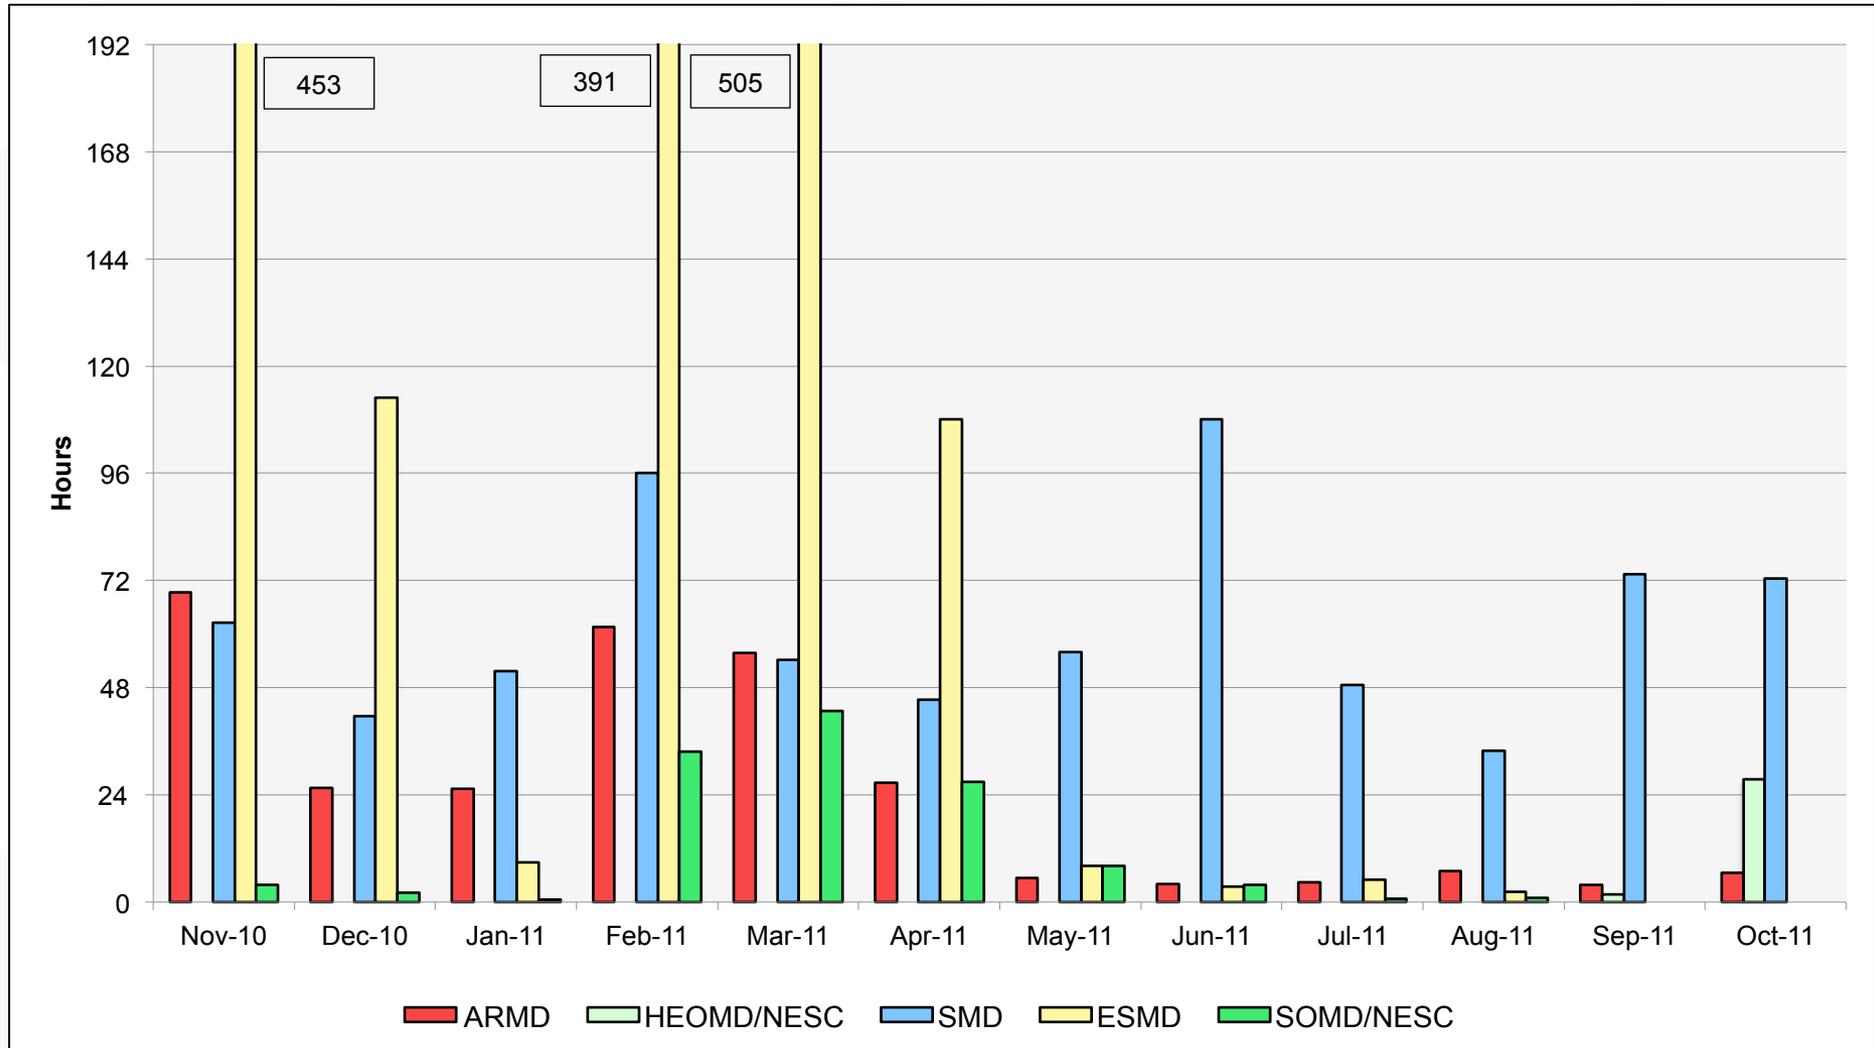
October 2011

# Columbia: Monthly Utilization by Size and Length



October 2011

# Columbia: Average Time to Clear All Jobs



# Columbia: Average Expansion Factor

