Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

10 January 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

HECC Resources Provide Critical Support for NEX Collaborative Platform

- HECC continues to provide crucial supercomputing resources that enable NASA Earth Exchange (NEX) science communities to address Earth science problems at a scale not previously achievable.
- NEX's collaboration and knowledge-sharing platform for the Earth science community combines supercomputing, Earth system modeling, workflow management, and NASA remote sensing data feeds to deliver a complete work environment for users to explore/analyze large datasets, run modeling codes, collaborate, and share results.
- In a first application of NEX, a research team from around the U.S. used the environment to adjoin and atmospherically correct a mosaic of 9,000 Landsat Thematic Mapper scenes and retrieve global vegetation density at a 30-meter resolution.
- The entire processing of the nearly 340 billion pixels in the composite took just a few hours on the Pleiades supercomputer, allowing the team to experiment with new algorithms and products within just a few days.

Mission Impact: The Pleiades supercomputing architecture, combined with a massive data storage capacity and high-speed network, enables NEX to engage large scientific communities and provide them with capabilities to perform modeling and data analysis on a grand scale not previously achievable.

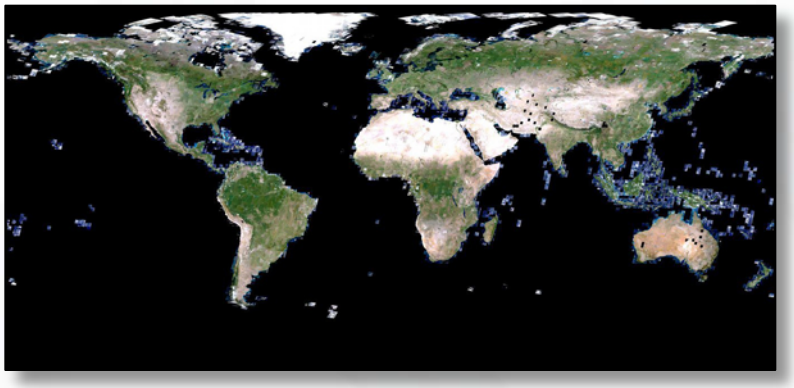

***Figure**: Image showing results from the first application of NEX: global vegetation density estimates at 30-meter resolution (nearly 340 billion pixels) from Landsat satellite data. (Andrew Michaelis, Tim Sandstrom, NASA/Ames)*

***POC:** Petr Votava, petr.votava@nasa.gov, (650) 604-4675**;** Ramakrishna Nemani, ramakrishna.nemani@nasa.gov**,** (650) 604-6185, NASA Ames Research Center*

HECC Supports High-Fidelity Simulations of the Ares I CLV Stage Separation Process

- Computational fluid dynamics (CFD) experts at NASA Ames have performed high-fidelity simulations of the Ares I Crew Launch Vehicle stage separation process.
- High-fidelity, time-accurate analyses were performed using the CFD code OVERFLOW and high-resolution computational meshes to resolve the complex flow details needed for design of Ares I stage separation systems.
- These simulations, run on Pleiades, are the only means of characterizing the complex and highly nonlinear aerodynamics of the stage separation process.
- The researchers also developed a steady-state aerodynamic database to provide key insights into this complex process.
- Access to Pleiades allowed researchers to complete these computationally demanding analyses in a timely fashion—each database simulation required 3,500 processors running for several weeks to complete.
- Over 200 simulation cases were completed to fulfill the database requirements, utilizing a tremendous amount of storage in addition to computational time.

* Note that this work was performed in FY11

Mission Impact: High-fidelity simulations, enabled by HECC resources, are critical to developing effective, reliable stage separation systems that can duplicate the flight conditions and full motor thrust levels needed to accurately characterize the stage separation flowfield.

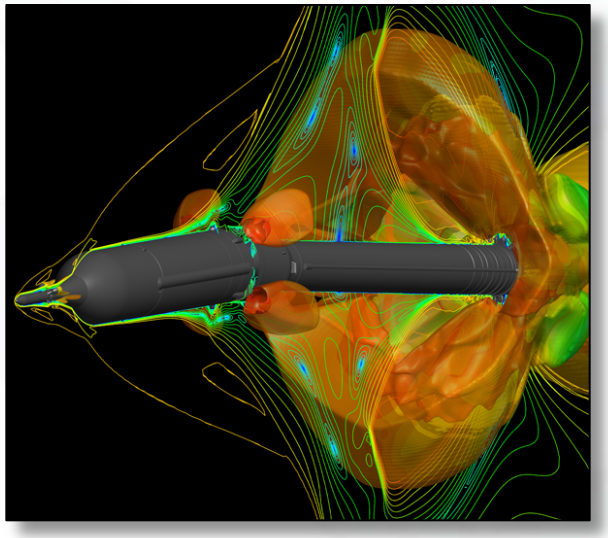

***Figure**: Cutting plane showing Mach number contours along with pressure iso-surfaces for the plumes of the Ares I launch vehicle ullage and separation motors, emphasizing the extent and complexity of stage separation aerodynamics. (Jeff Onufer, NASA/Ames)*

***POCs:** Jeff Onufer, jeffrey.t.onufer@nasa.gov, (650) 604-3982; Henry Lee, henry.c.lee@nasa.gov, (650) 604-3982, NASA Ames Research Center*

Upgrades to Pleiades Filesystems Deliver Improved Performance

- The addition of a new Pleiades Lustre filesystem, named /nobackupp6, completes a six-month process of migrating petabytes of user data to new hardware with minimal impact on user workflow.
- All HECC users have now been migrated to the newer, faster filesystems, which provide 10 times the Input/Output Operations per Second (IOPS) performance over 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.
- In addition to increasing the IOPS performance, as part of the six-month project, the Supercomputing Systems team increased the "scratch space" for users to temporarily store files.

Mission Impact: Along with the increased computational capability of the Pleiades supercomputer, improvements to the Pleiades filesystem performance including larger "scratch space" (temporary storage) enables researchers to more fully utilize this powerful resource.

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

***POCs:** Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

User Workflow Enhanced Through Secure Unattended Proxy

- The Supercomputing Systems team has added new functionality to the Secure Unattended Proxy (SUP) that allows users to remotely submit batch jobs to specific hosts within the HECC enclave.
- With the SUP, users obtain special "SUP keys" using SecurID authentication, after which they can use those keys to perform operations from unattended jobs and/or scripts.
- The new capability provides a unified approach for automated remote workflow processing—users can stage data, submit jobs, monitor job processes, and retrieve results through the same mechanism.
- This enhancement, requested by the Kepler Mission team to improve their workflow to Pleiades, increases the simplicity, robustness, and throughput of Kepler jobs; other users may be approved to use the new feature in the future.

Mission Impact: Improved workflow enables more efficient usage of HECC resources and results in higher user productivity.

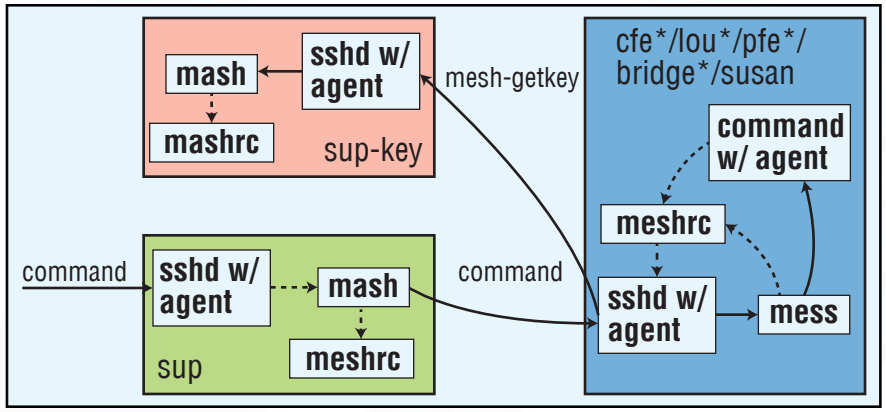

Figure: *Chart showing the architecture of the Secure Unattended Proxy, which provides a unified approach for automated remote workflow processing using HECC resources.*

POCs: *Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division;*
Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp;
Paul Kolano, paul.kolano@nasa.gov, (650) 604-4271, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Major Upgrade to hyperwall Visualization System Reduces Operating Costs

- The HECC Systems team planned, tested, and deployed a major operating system upgrade to the hyperwall.
- CentOS, a freely available, Enterprise-class Linux Distribution that is binary compatible with Red Hat Linux, reduces operating costs by eliminating the ongoing maintenance support costs of the Linux operating system.
- The hyperwall provides a good testbed to validate the feasibility of CentOS in the HECC environment prior to a large-scale deployment on the Pleiades supercomputer.
- Deploying CentOS on Pleiades will result in additional, significant savings in maintenance support costs.

Mission Impact: Planning, evaluating, and testing various available solutions reduces maintenance support costs, which enables more effective use of HECC funding.

Figure: *The in-house developed hyperwall visualization system provides a supercomputer-scale environment to visualize and explore the very large, high-dimensional datasets produced by NASA supercomputers and instruments.*

POCs: *Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.s.chan@nasa.gov, (650) 604-4613, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

Data 'Roll Forward' Doubles Capacity of Tape Media

- The "roll forward" of archive data from LTO-4 tape media and drives to LTO-5 has been completed in a process that was transparent to users; HECC staff converted 19 petabytes (PB) of data residing on 18,000 tapes.
- The conversion reduces HECC tape library slot requirements, as LTO-5 media has double the capacity of LTO-4 media.
- HECC identified significant issues with LTO-4 drives failing; HECC systems staff worked with Spectra Logic and IBM, and identified contaminated media as the root cause of the problems.
- As part of the mitigation effort, Spectra Logic provided equivalent LTO-5 media to replace the 9 PB of contaminated LTO-4 media—this was done at no cost, which provided approximately $360,000 in savings to NASA.

Mission Impact: This media conversion reduces the tape library slot requirements, enabling HECC to better scale data storage capacity within existing tape libraries.

Figure: *One of HECC's six 8-frame Spectra Logic T950 tape libraries.*

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 Participates in Annual AGU Conference

- HECC supported and participated in the annual American Geophysical Union (AGU) conference held in San Francisco, December 5–9, 2011.
- Three staff members and one user gave technical talks, including three in the NASA booth theater area:
  - "HECC Project: Passing the PetaFLOP Barrier," William Thigpen
  - "Visualization and Analysis with Adaptive Mesh Refinement Data," Patrick Moran
  - "NASA Earth Exchange: A Collaborative Supercomputing Platform," Ramakrishna Nemani
  - "Are Cloud Environments Ready for High-Performance Computing Applications," Steve Heistand (poster)
- HECC staff also supported the NASA booth, handing out printed materials generated for the SC11 conference in November, which reduced the cost of AGU participation.
- In addition, staff facilitated a tour of Ames for visitors from NASA Headquarters (see slide 12).

Mission Impact: Participation in science conferences highlights the scientific value of NASA data and provides a valuable opportunity to meet current and new HECC users to discuss computational resource and service needs.

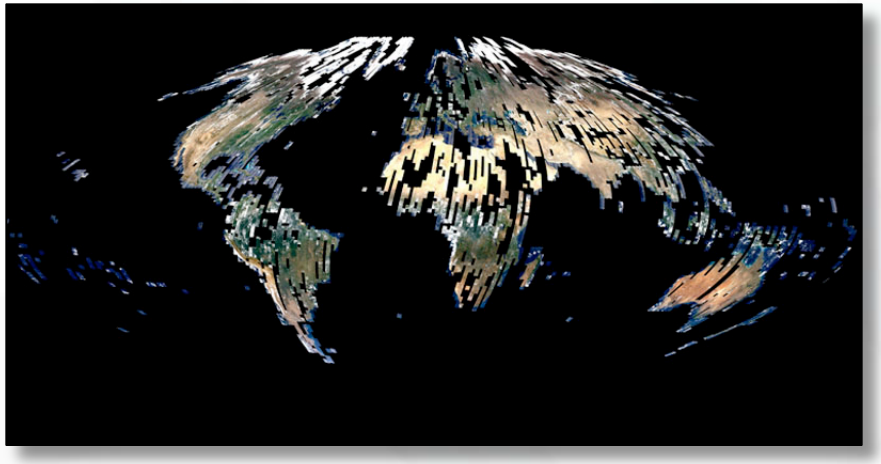

***Figure**: The American Geophysical Union (AGU) conference is the largest international conference in the geophysical sciences, attracting about 20,000 Earth and space scientists, educators, students, and policy makers.*

***POCs:** Harper Pryor, harper.pryor@nasa.gov, (650) 604-0468, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

HECC Hosts Visit by Dr. Warren Washington

- Dr. Warren Washington, pre-eminent climate scientist, delivered a presentation on "20th and 21st Century Climate Modeling, Societal Impacts, and Environmental Justice."
- A standing-room only crowd attended this Ames Director's Colloquium talk in the NAS auditorium; HECC staff sponsored a tour of Ames and the NAS facility.
- As a pioneer in climate modeling, Dr. Washington wrote the seminal book, "An Introduction to Three-Dimensional Climate Modeling;" he plays active role in the U.S. Global Change Research Program as part of the DOE Climate Change Prediction Program.
- Dr. Washington served as an advisor to five U.S. Presidents; is former chair of the National Science Board; and received the National Medal of Science from President Obama in November 2010.

Mission Impact: Forging relationships with key climate scientists and programs supports NASA's science mission; awareness of critical climate research issues enhances understanding of future computational needs; discussion of societal impacts is highly motivational for staff.

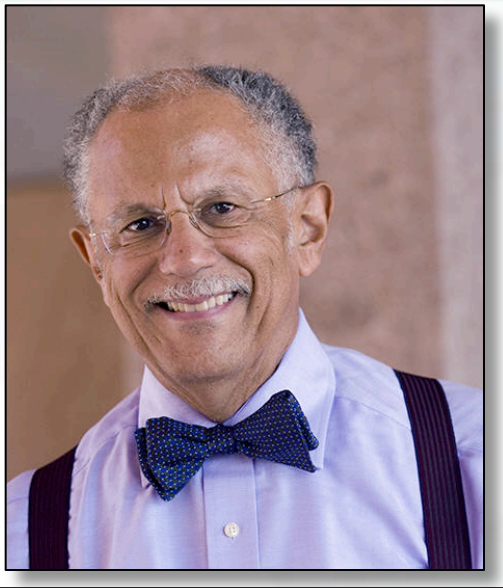

***Figure**: Above, Dr. Warren Washington, senior scientist, National Center for Atmospheric Research.*

***POCs:** Harper Pryor, harper.pryor@nasa.gov, (650) 604-0468, NASA Advanced Supercomputing Division, Computer Sciences Corp.*

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

- Requests approved: 3; New requests received: 4; Requests waiting: 8.
- Wait times are increasing again.
- Wait times are 0.7 to 2.4 months, with 6 over 1 month.
- The International Visitor Control Center has been contacted for updates on those who have been waiting more than one month.

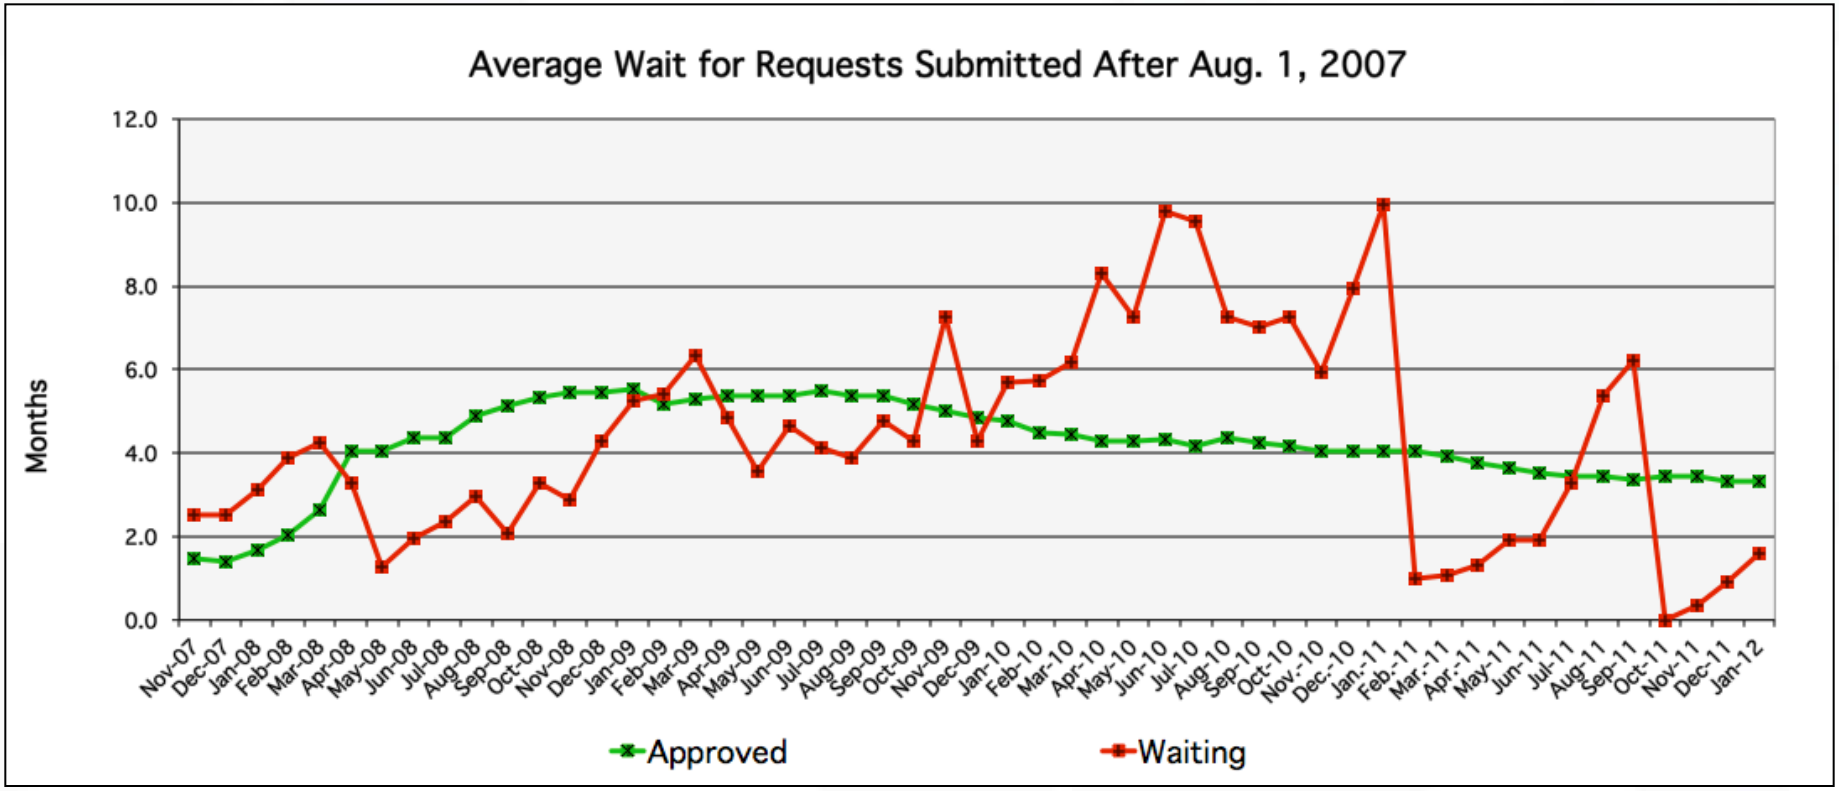

HECC Facility Hosts Several Visitors and Tours in December 2011

- HECC hosted 4 scheduled tour groups in December; guests received an overview of the HECC Project, demonstrations of the hyperwall visualization system, and tours of the computer room floor. Guests this month included:
  - Dr. Warren Washington, senior scientist and former head of the Climate Change Research Section in the Climate and Global Dynamics Division at NCAR (see slide 10).
  - As part of Ames' education outreach, a group of students from Pacific Law Academy, a charter school in the Stockton Unified School district; these students are all enrolled in calculus classes and have a demonstrated aptitude in math.
  - Visiting writers for Gizmodo, who received an overview of HECC/NAS and a computer room tour. Gizmodo is one of the five most-visited blogs on the Internet, with over 6 million unique visitors last month (*http://advertising.gawker.com/gizmodo/*).

Figure: As part of their Ames visit, students toured the NASA Advanced Supercomputing facility, which included demonstration of the hyperwall.

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

Presentations and Papers

Presentations

- **American Geophysical Union Conference, Dec. 6–9, San Francisco**
  - "HECC Project: Passing the PetaFLOP Barrier," W. Thigpen
  - "Visualization and Analysis with Adaptive Mesh Refinement Data," P. Moran
  - "NASA Earth Exchange: A Collaborative Supercomputing Platform," R. Nemani*
  - "Are Cloud Environments Ready for High- Performance Computing Applications," S. Heistand

Papers

- "The Impact of Hyper-Threading on Processor Resource Utilization in Production Applications," S. Saini, H. Jin, R. Hood, D. Barker, P. Mehrotra, R. Biswas, *18th IEEE International Conference on High Performance Computing (HiPC),* Bangalore, India, Dec 18-21, **Best Paper Award winner**.

- "Two Earth-sized Planets Orbiting Kepler 20," F. Fressin, G. Torres, J.F. Rowe, D. Charbonneau, C.E. Henze, *et al., Nature,* Published online 20 Dec. 2011.

- "A Look at the Impact of High-End Computing Technologies on NASA Missions," R. Biswas, J. Dunbar, J. Hardman, F.R. Bailey, L. Wheeler, S. Rogers, *IEEE IT Professional*, ISSN: 1520-9202, pre-print.
  http://www.computer.org/portal/web/csdl/doi/10.1109/MITP.2011.110

- "Short Wave-Length Electromagnetic Perturbations Excited Near the Solar Probe Plus Spacecraft in the Inner Heliosphere: 2.5D Hybrid Modeling," A.S. Lipatov , E.C. Sittler, R.E. Hartle, J.F. Cooper, *Planetary and Space Science,* in press, available online Dec. 2011.*
  http://www.sciencedirect.com/science/article/pii/S0032063311003527

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

News and Events

- **Powerful Pixels: Mapping the Apollo Zone on the Moon**, *article, SpaceRef*, Dec. 30, 2011 – Describes how the Apollo Zone project processes maps of the moon using Pleiades. http://www.spaceref.com/news/viewsr.html?pid=39455

- **Powerful Pixels Help Map The Apollo Zone**, *article, Moon Daily,* Dec. 29, 2011 Describes how the Apollo Zone project processes maps of the moon using Pleiades. *http://www.space-travel.com/reports/Powerful_Pixels_Help_Map_The_Apollo_Zone_999.html*

- **It's a Small World: Kepler Spacecraft Discovers First Known Earth-Size Exoplanets**, *news article, Scientific American,* Dec. 20, 2011 – Includes description of Pleiades' role in helping researchers analyze close to a billion different scenarios for the NASA's Kepler Mission. *http://www.scientificamerican.com/article.cfm?id=kepler-20-smallest*

- **Smallest planet is tinier than Earth,** *news article, NewScientist,* Dec. 20, 2011 – Describes the Kepler team's use of Pleiades to validate their finds by "modelling all other possible explanations, including brown dwarfs on NASA's fastest supercomputer," with link to Pleiades web page. *http://www.newscientist.com/article/dn21306-smallest-planet-is-tinier-than-earth.html*

- **Harvard astronomers find Earth-sized planets, the smallest yet**, *news article, Boston Globe,* Dec. 20, 2011 – Mentions how the Kepler team utilizes Pleiades to run calculations ruling out "any other astrophysical phenomenon than another planet." *http://www.boston.com/Boston/metrodesk/2011/12/harvard-astronomers-find-earth-sized-planets-the-smallest-yet/WaLRgfAvXn6O3CP7Z7jIfM/index.html*

HECC Utilization

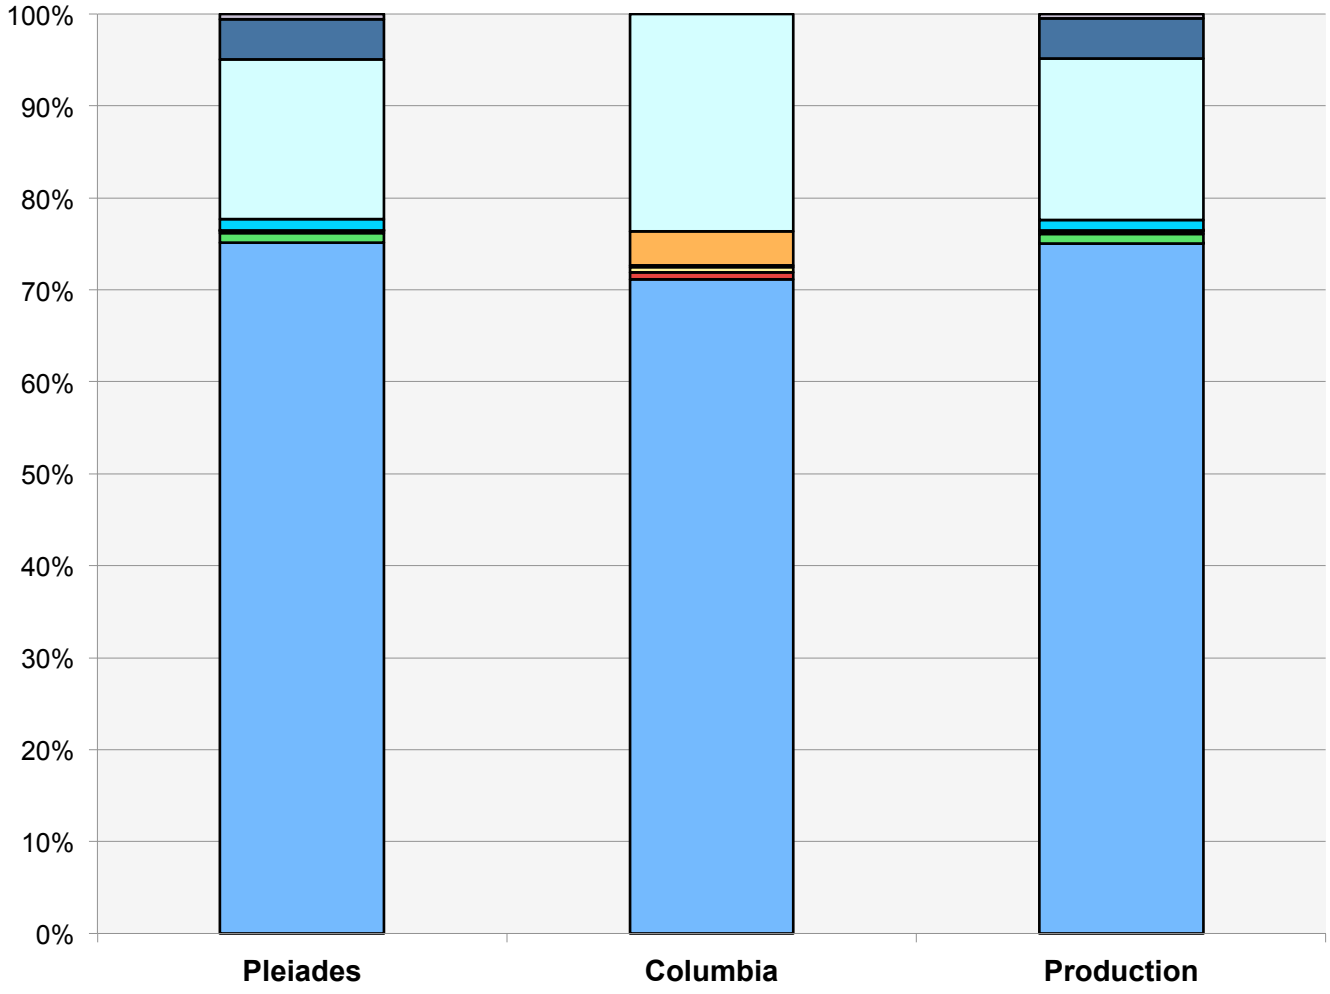

December 2011

HECC Utilization Normalized to 30-Day Month

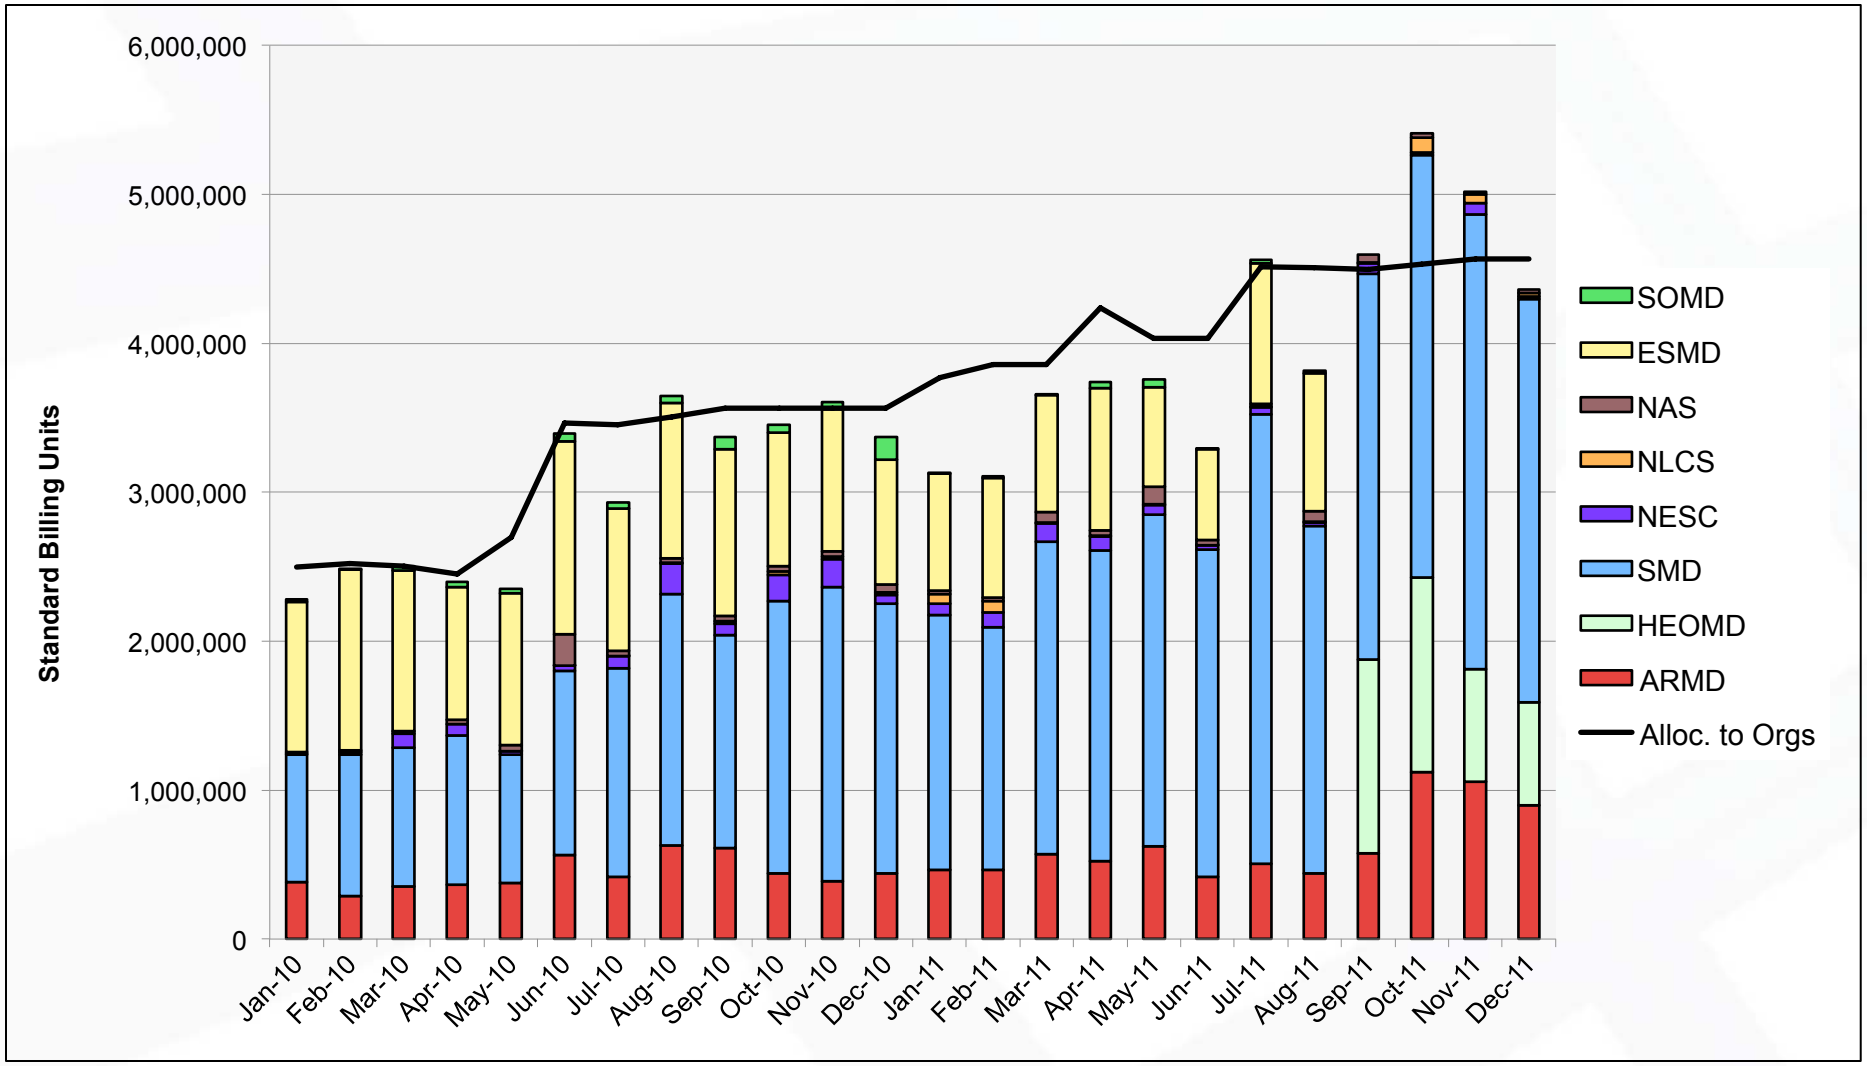

HECC Utilization Normalized to 30-Day Month

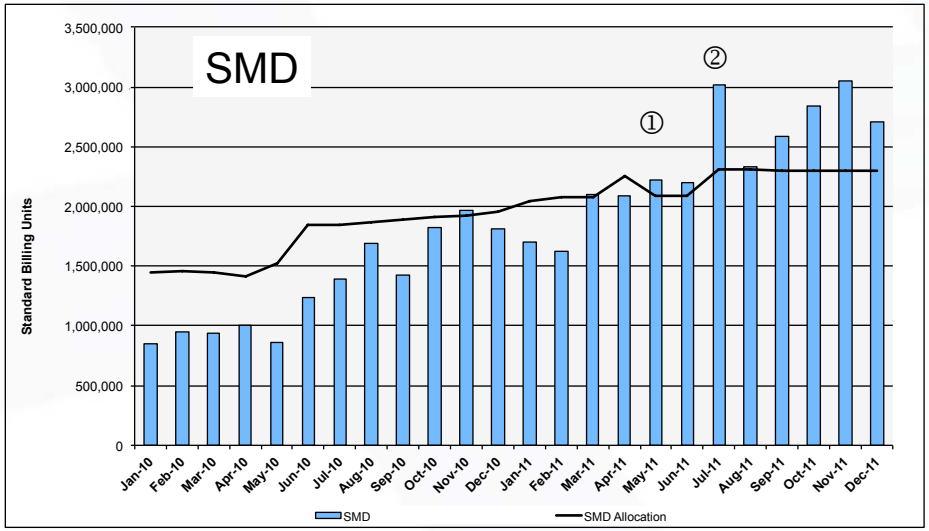

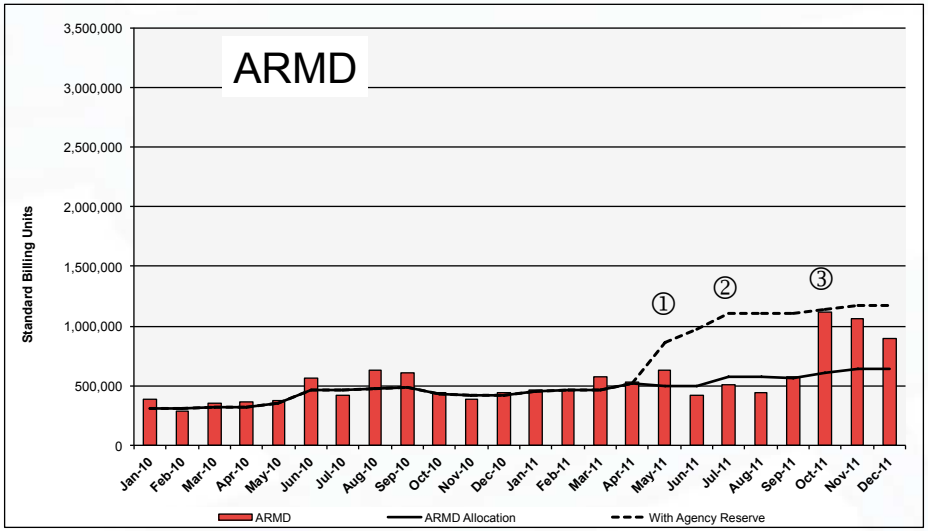

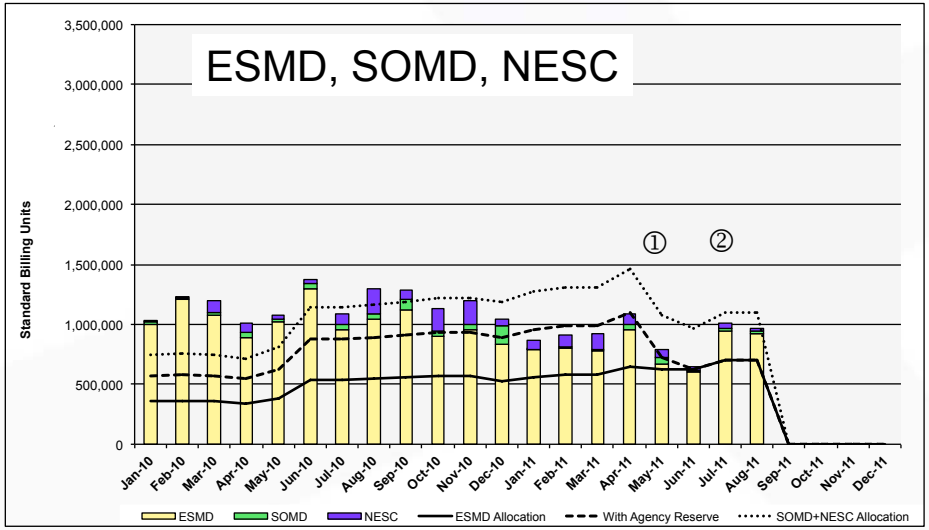

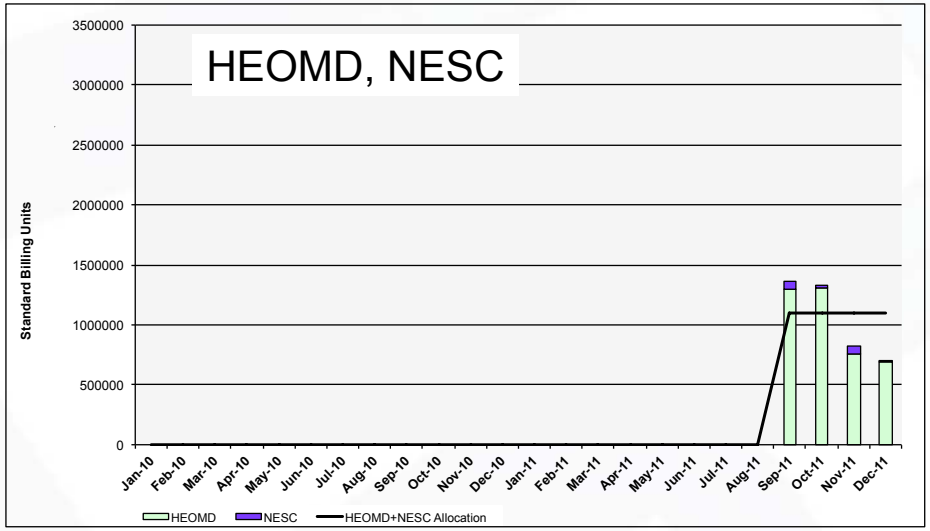

① Allocation to orgs. decreased to 75%, Agency reserve shifted to ARMD ② 14 Westmere racks added ③ 2 ARMD Westmere racks added

Tape Archive Status

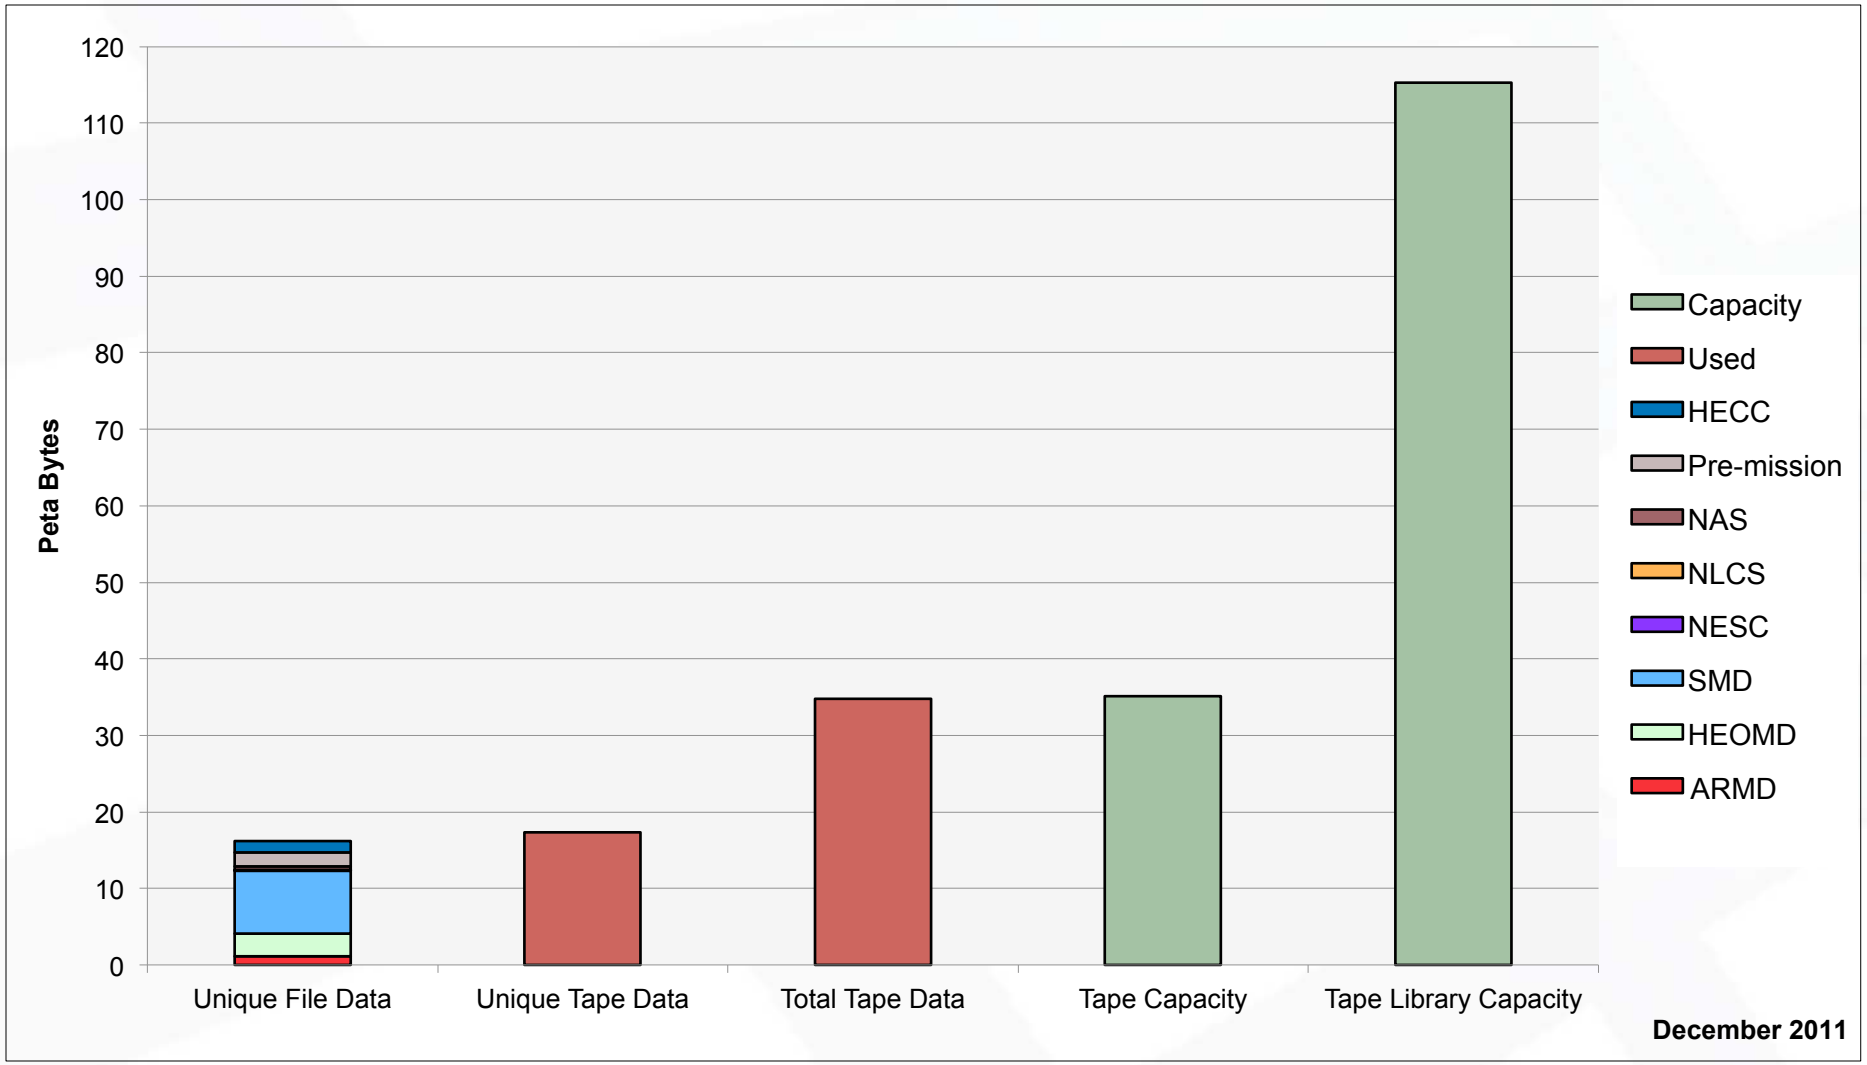

Tape Archive Status

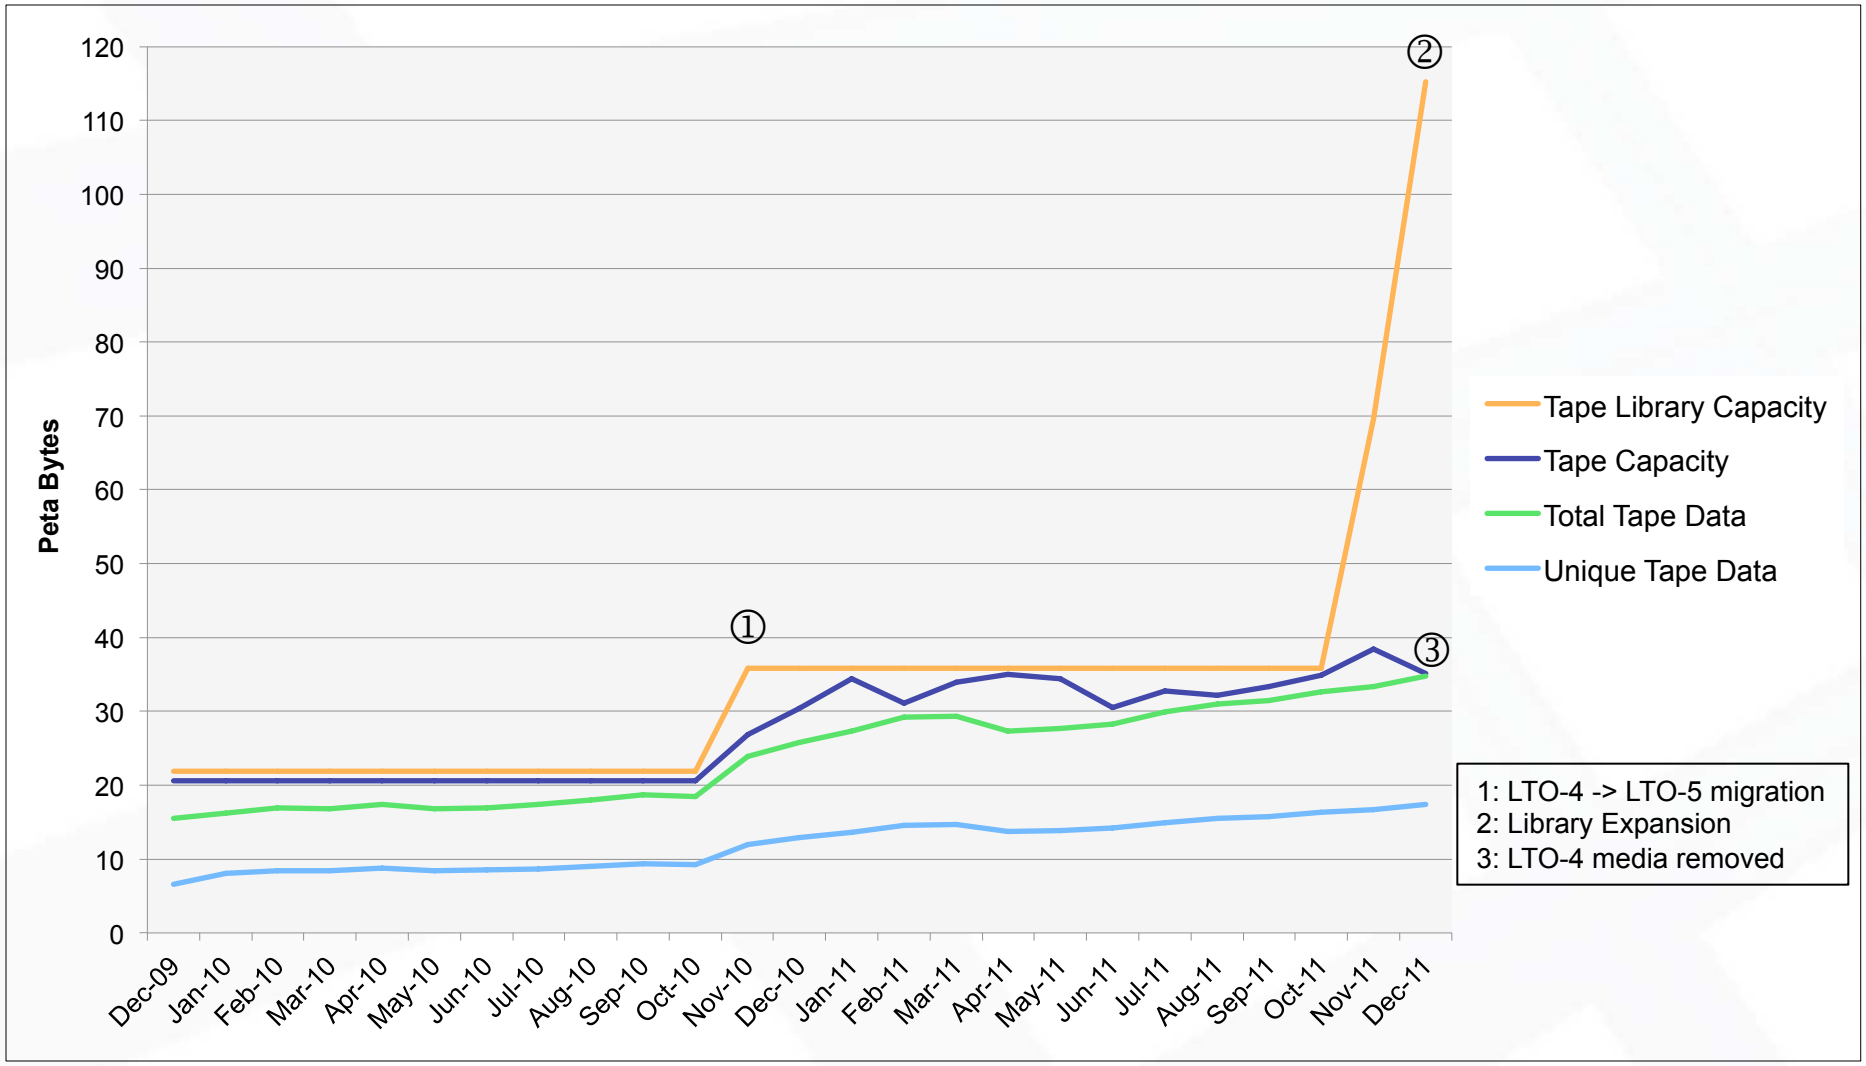

Pleiades:
SBUs Reported, Normalized to 30-Day Month

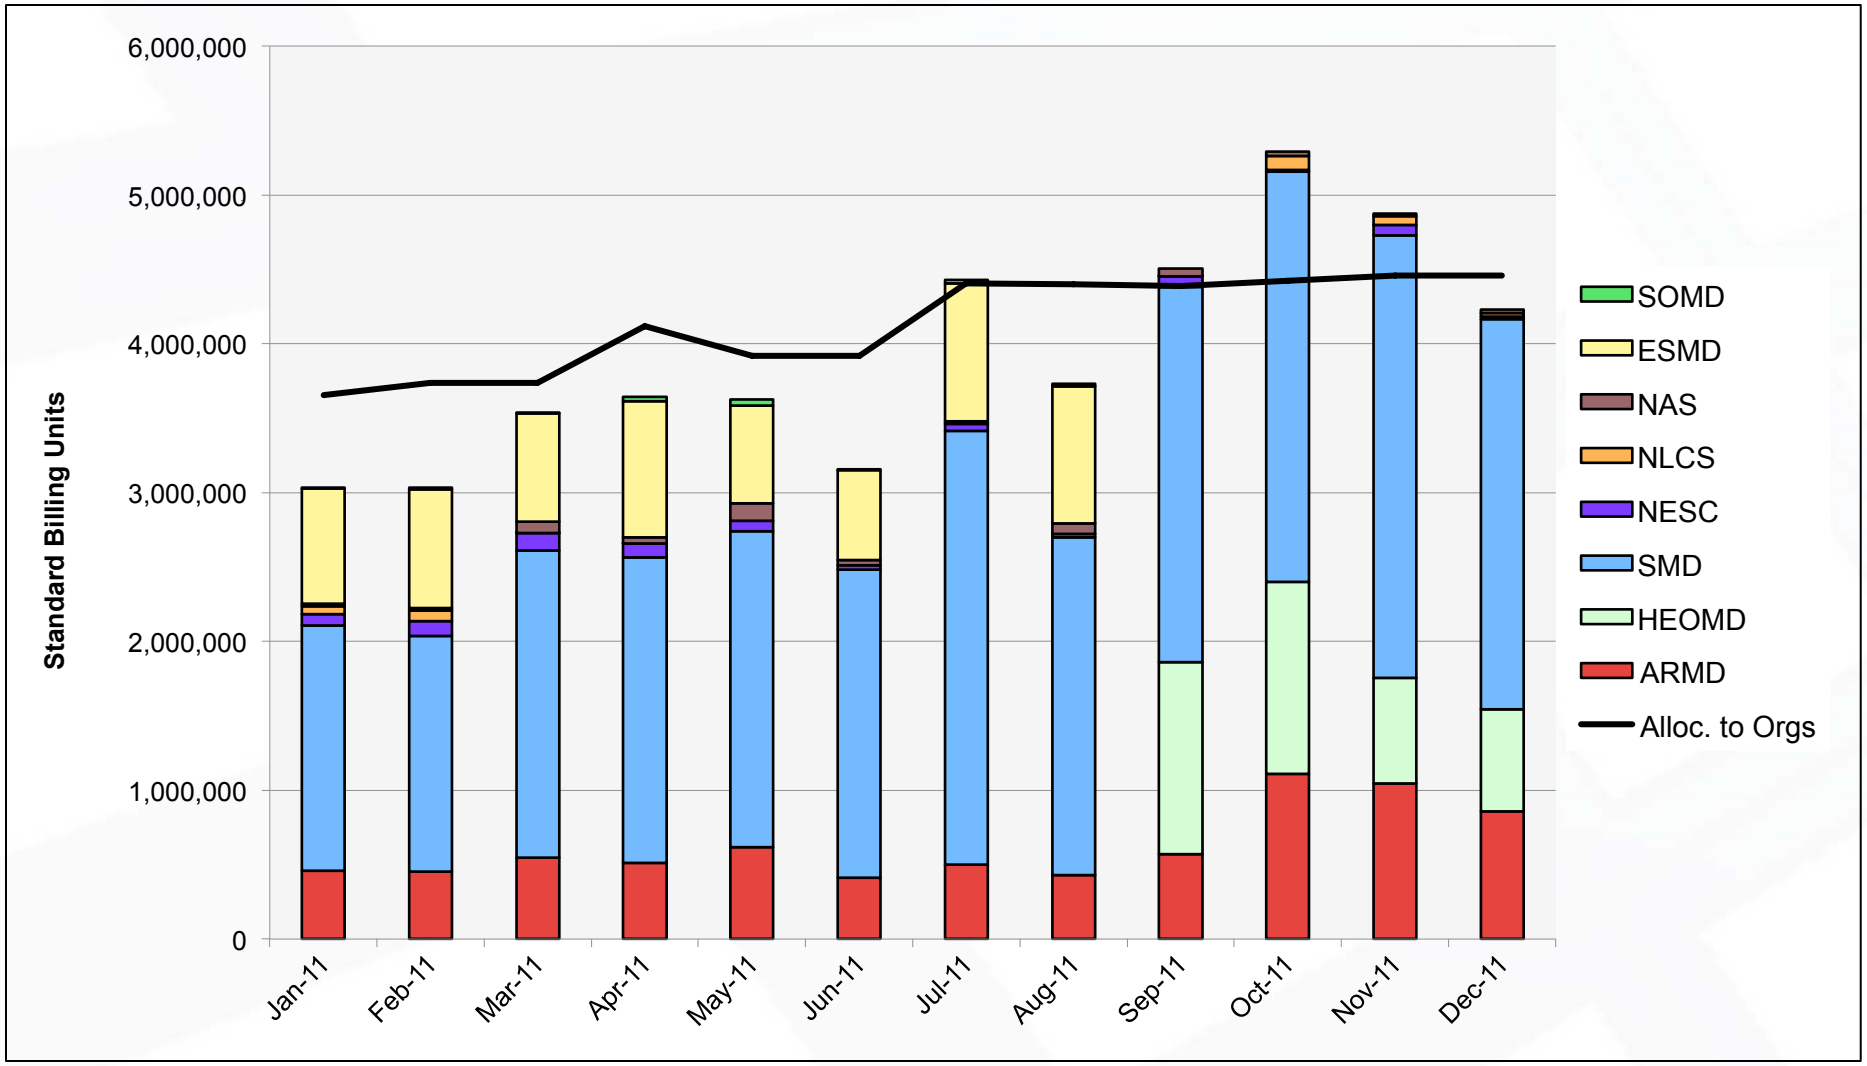

Pleiades:
Devel Queue Utilization

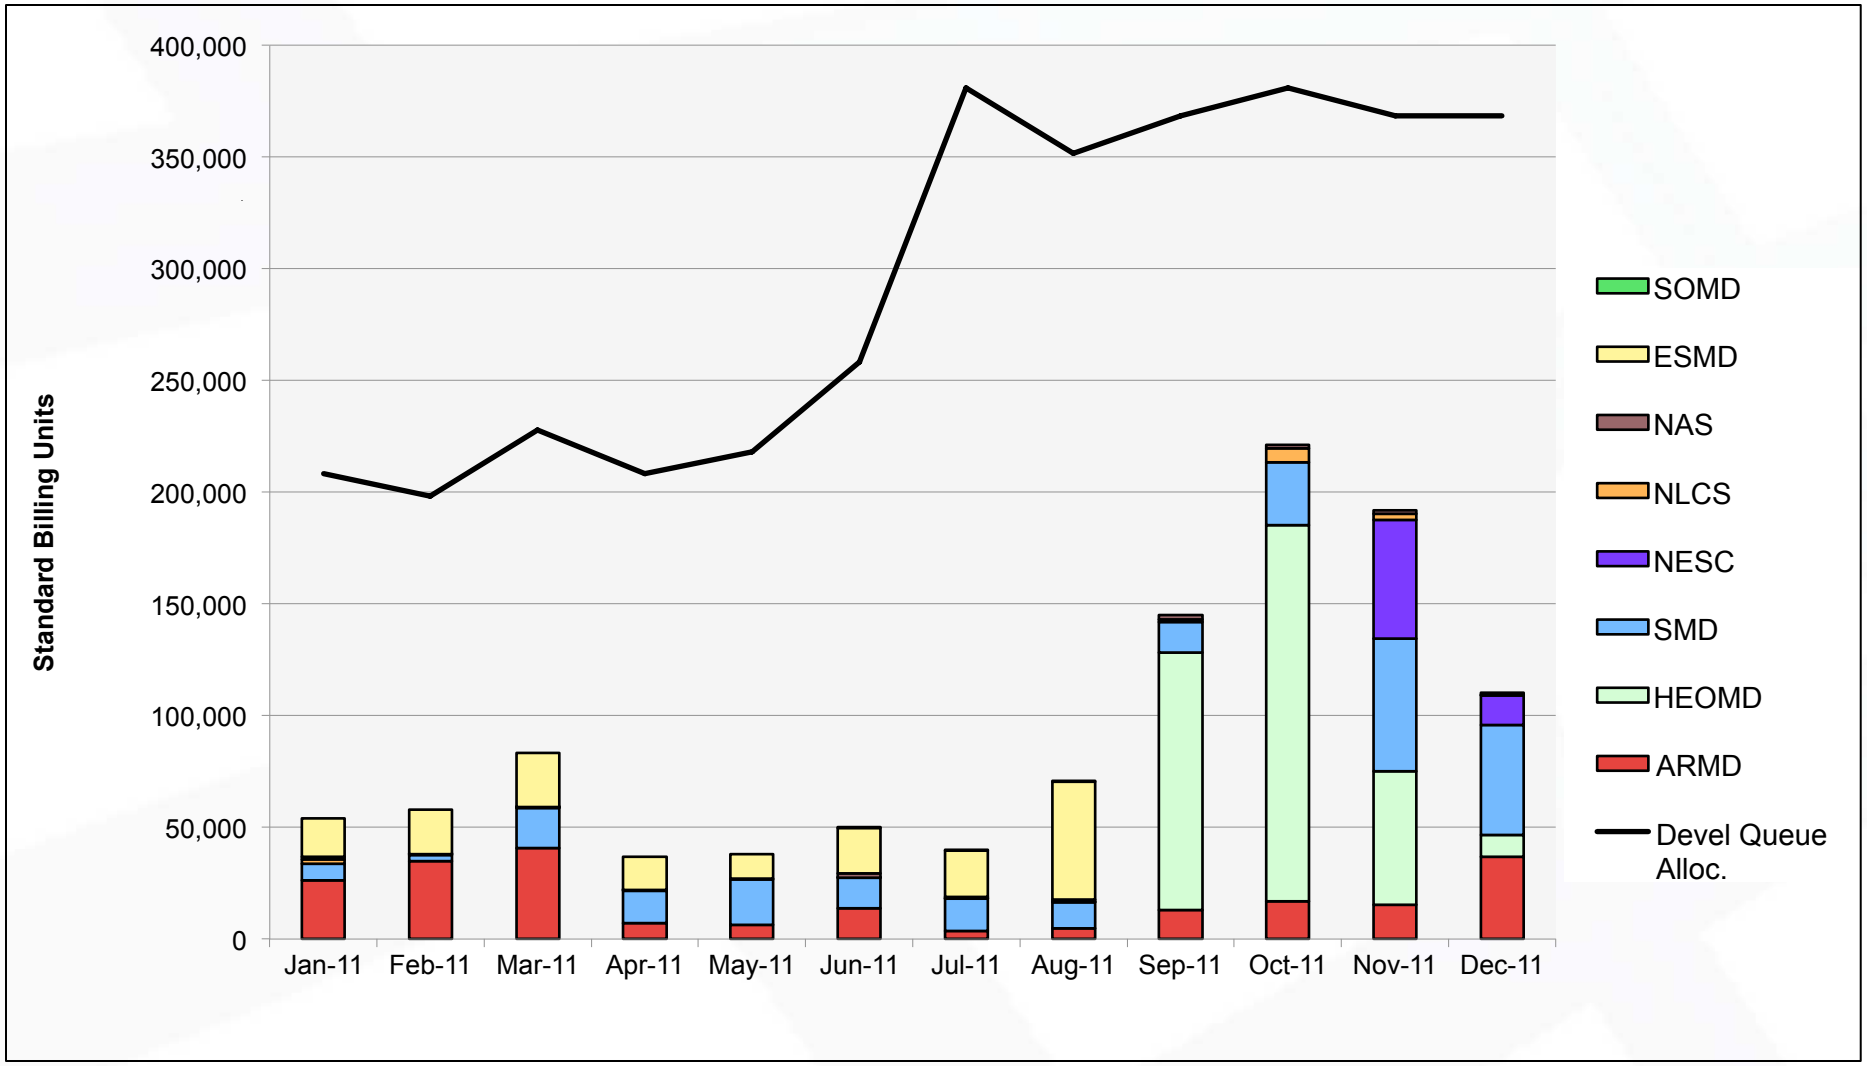

Pleiades:
Monthly SBUs by Run Time

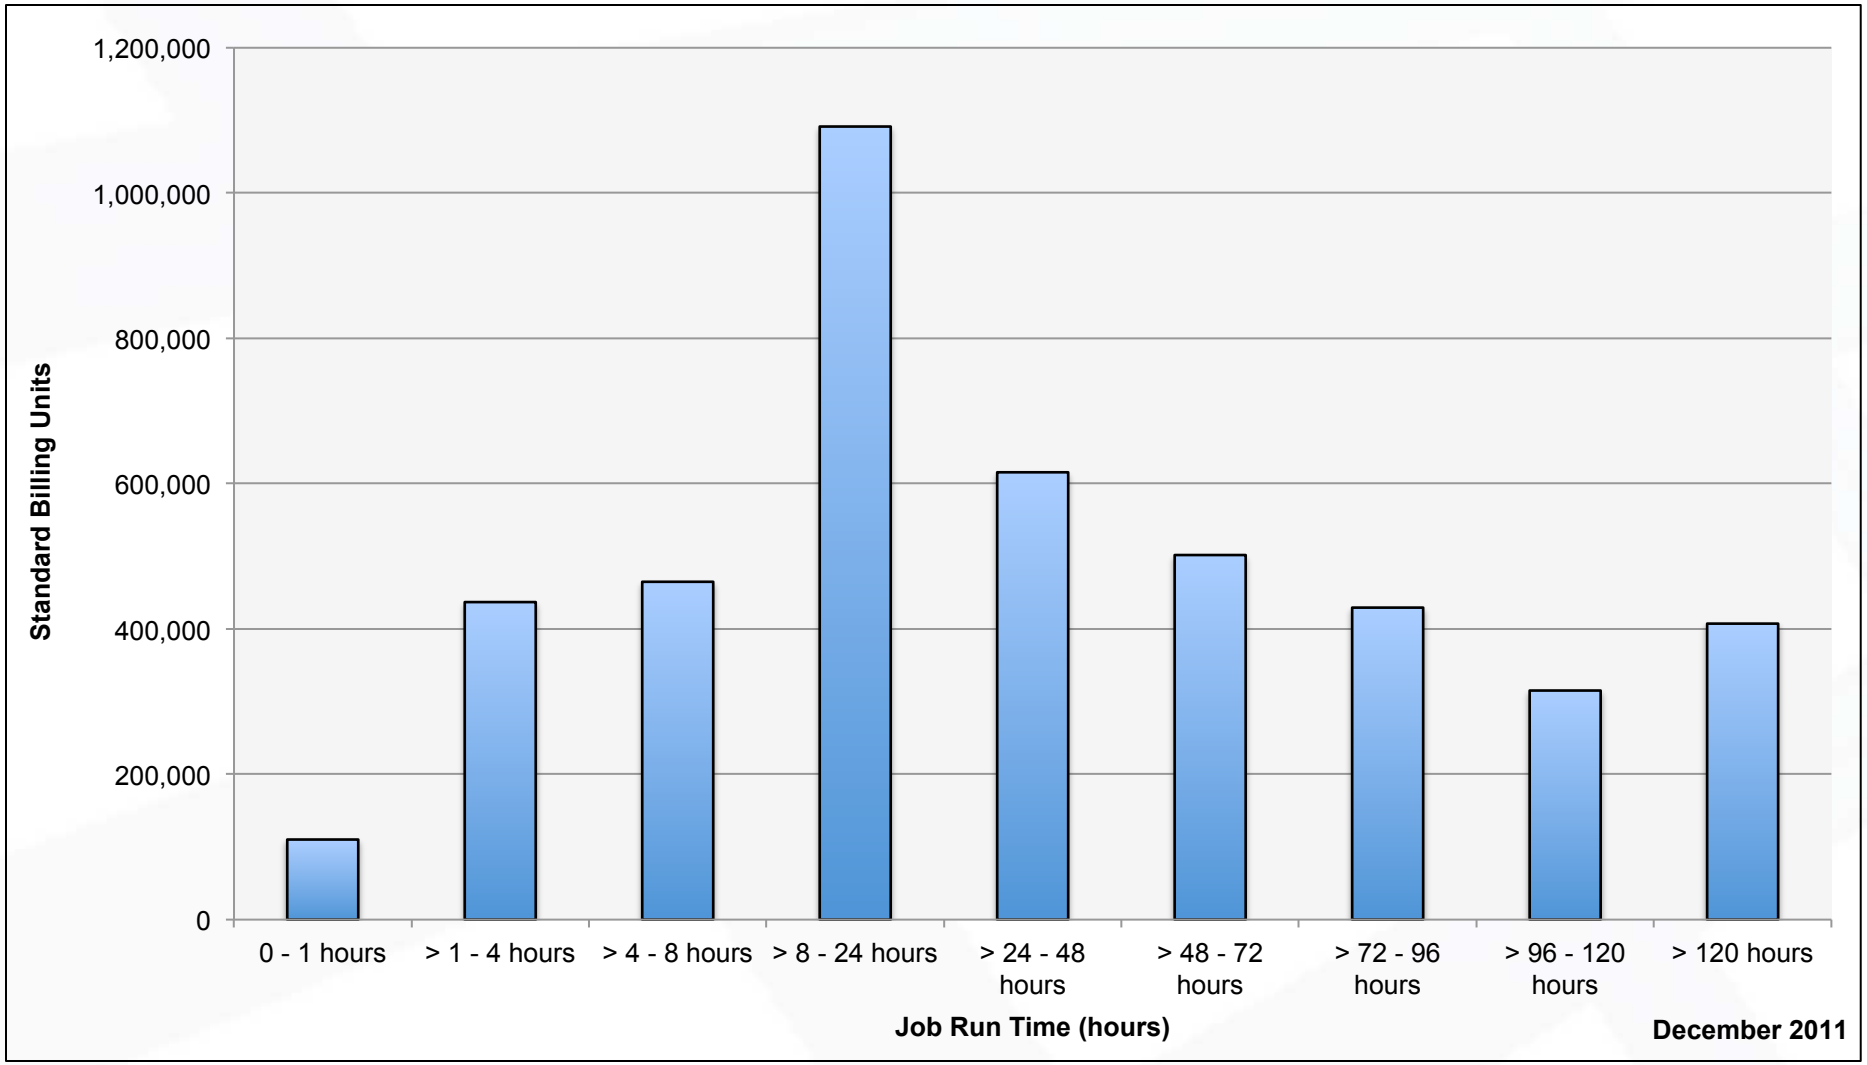

Pleiades:
Monthly Utilization by Size and Mission

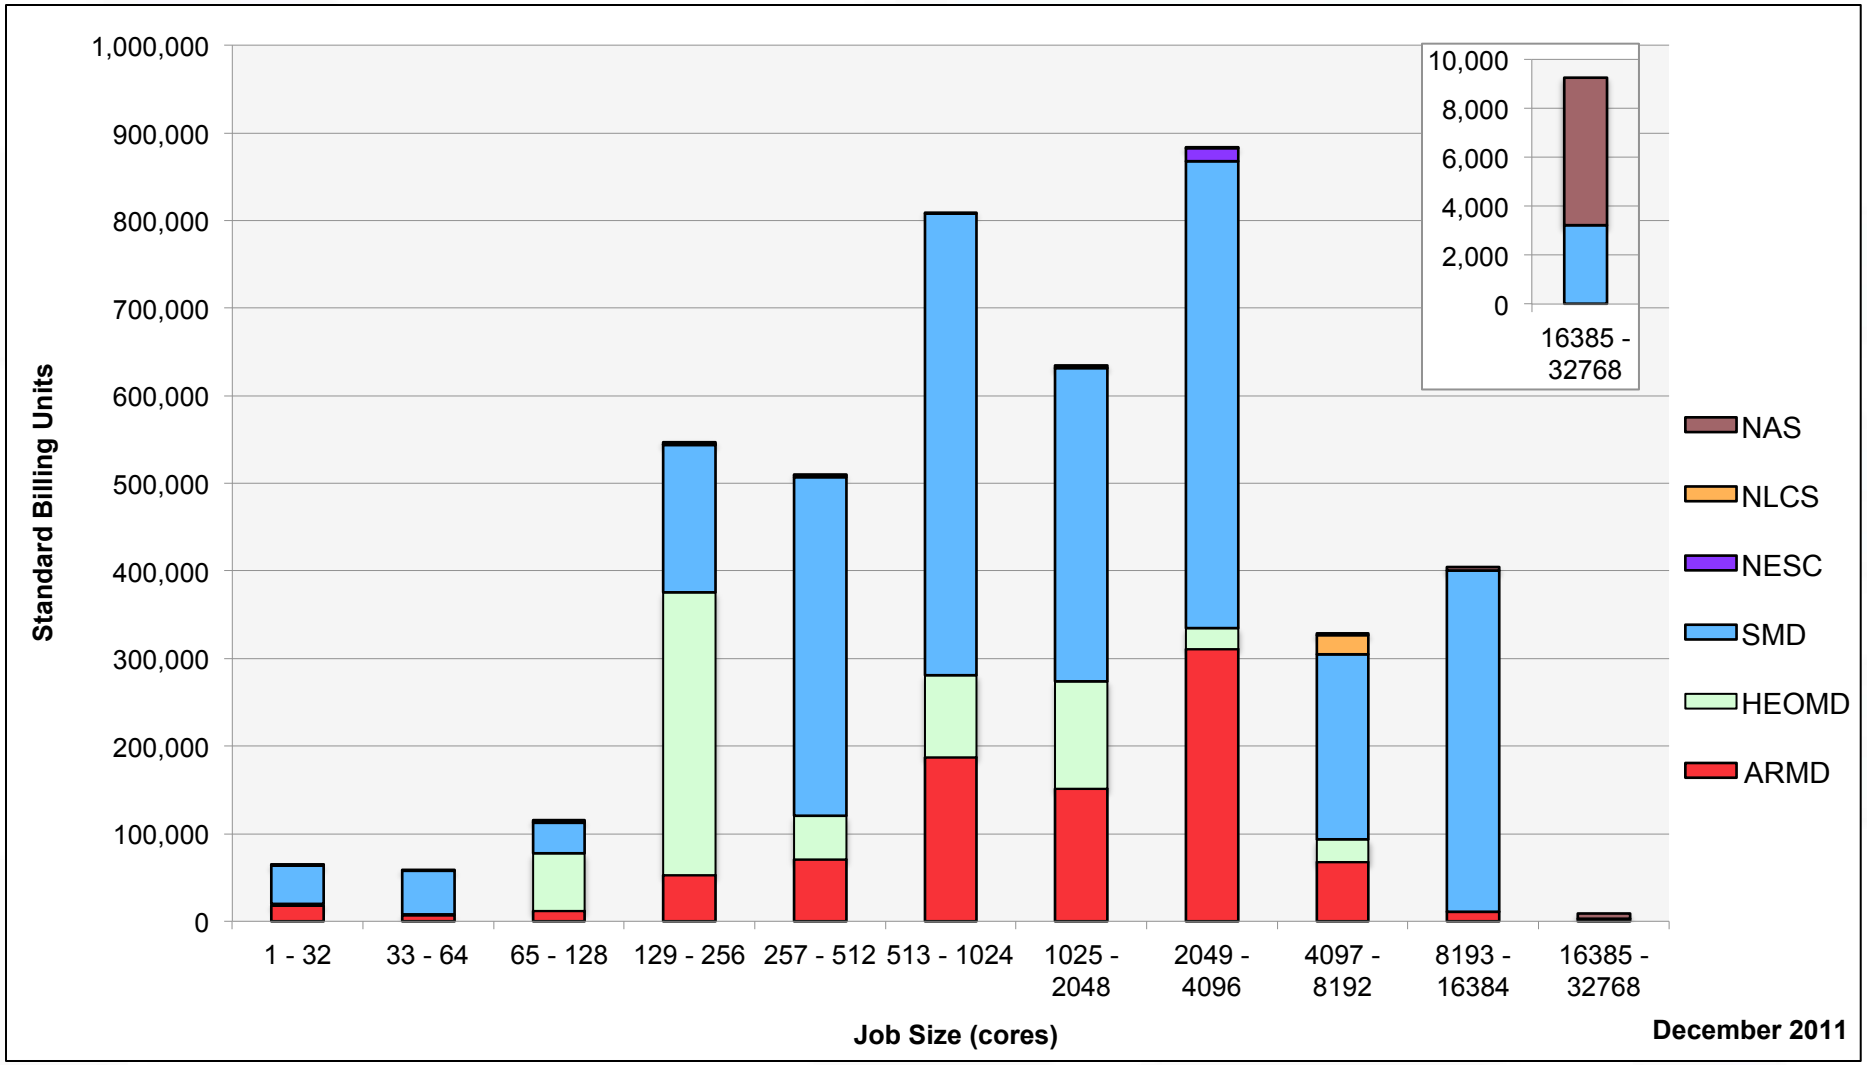

Pleiades:
Monthly Utilization by Size and Length

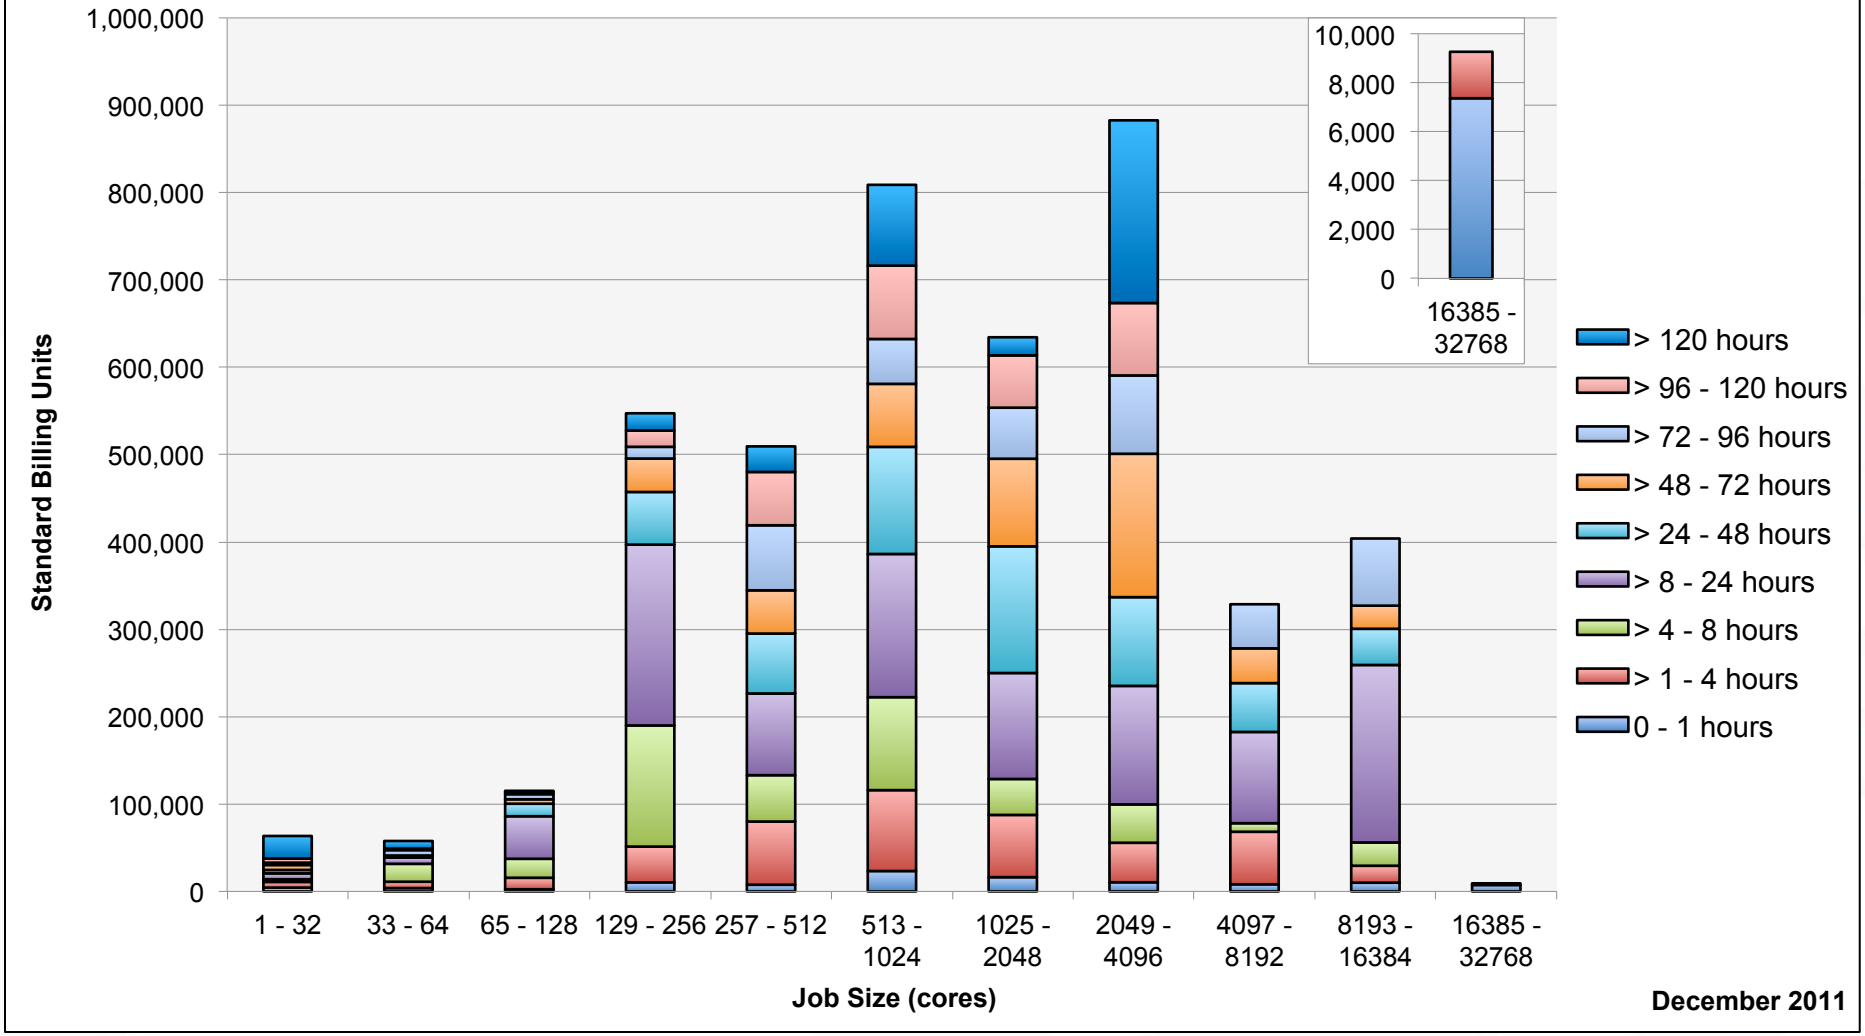

Pleiades:
Average Time to Clear All Jobs

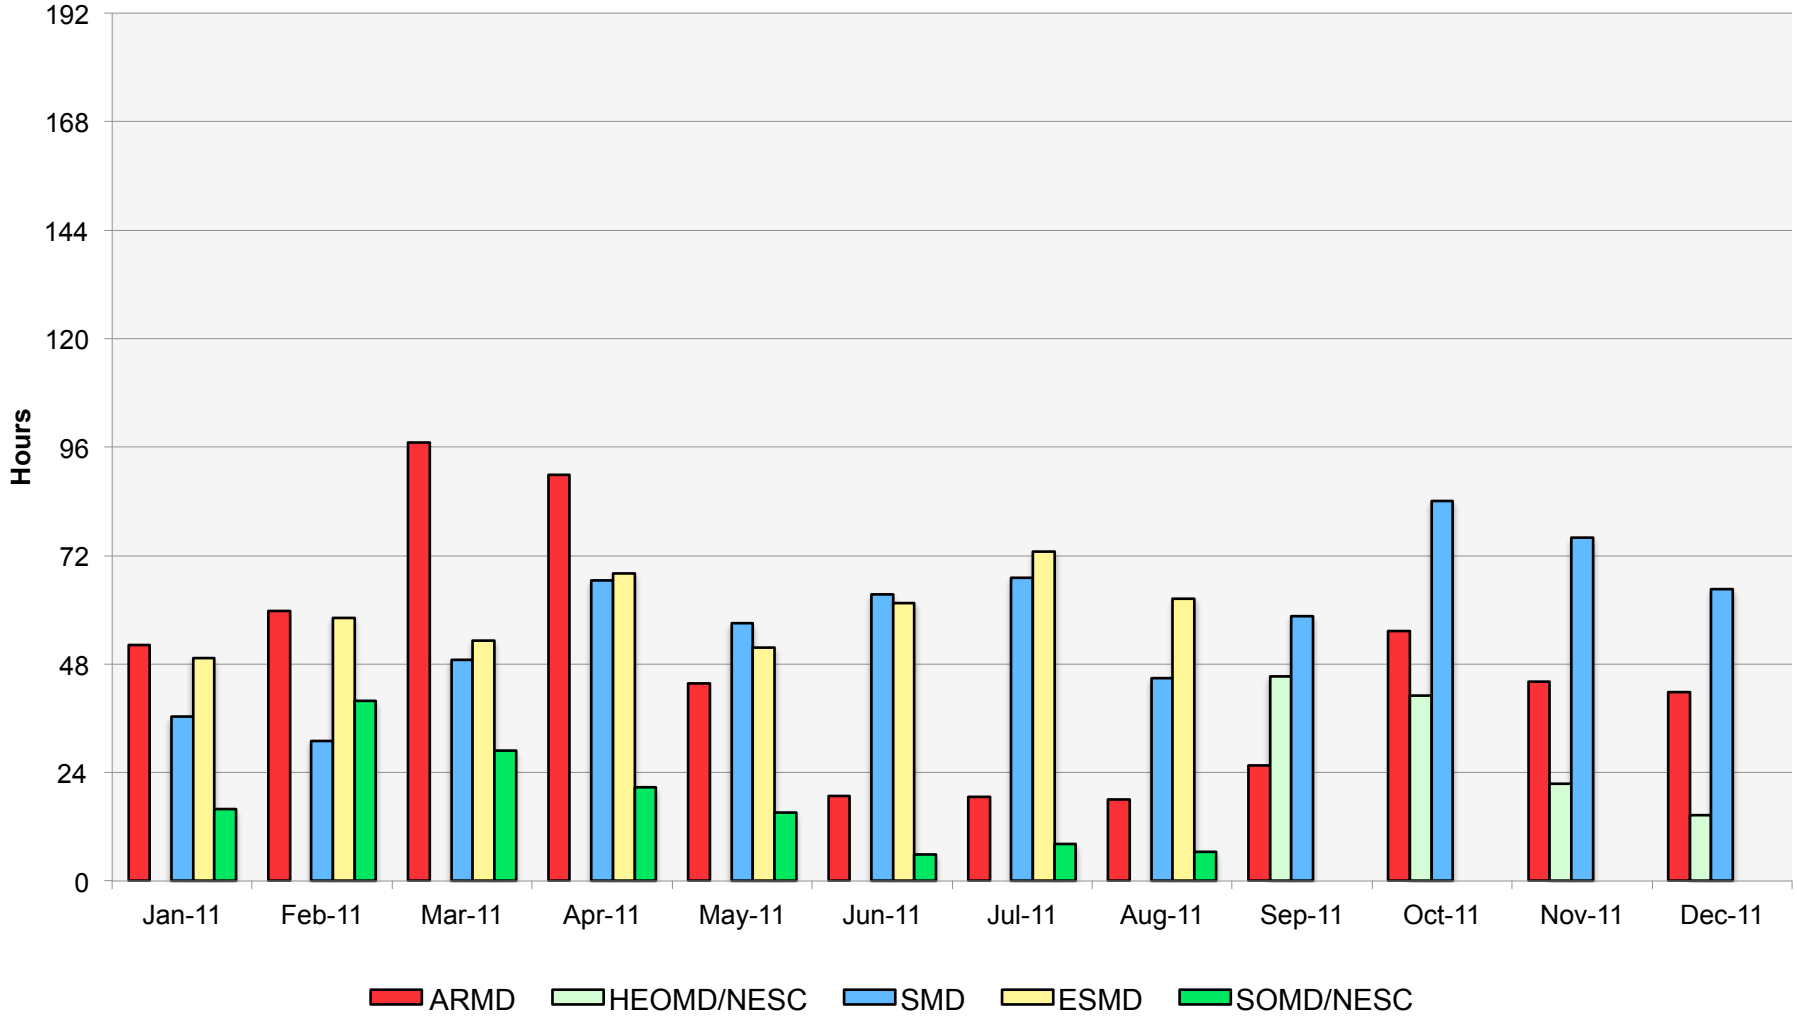

Pleiades:
Average Expansion Factor

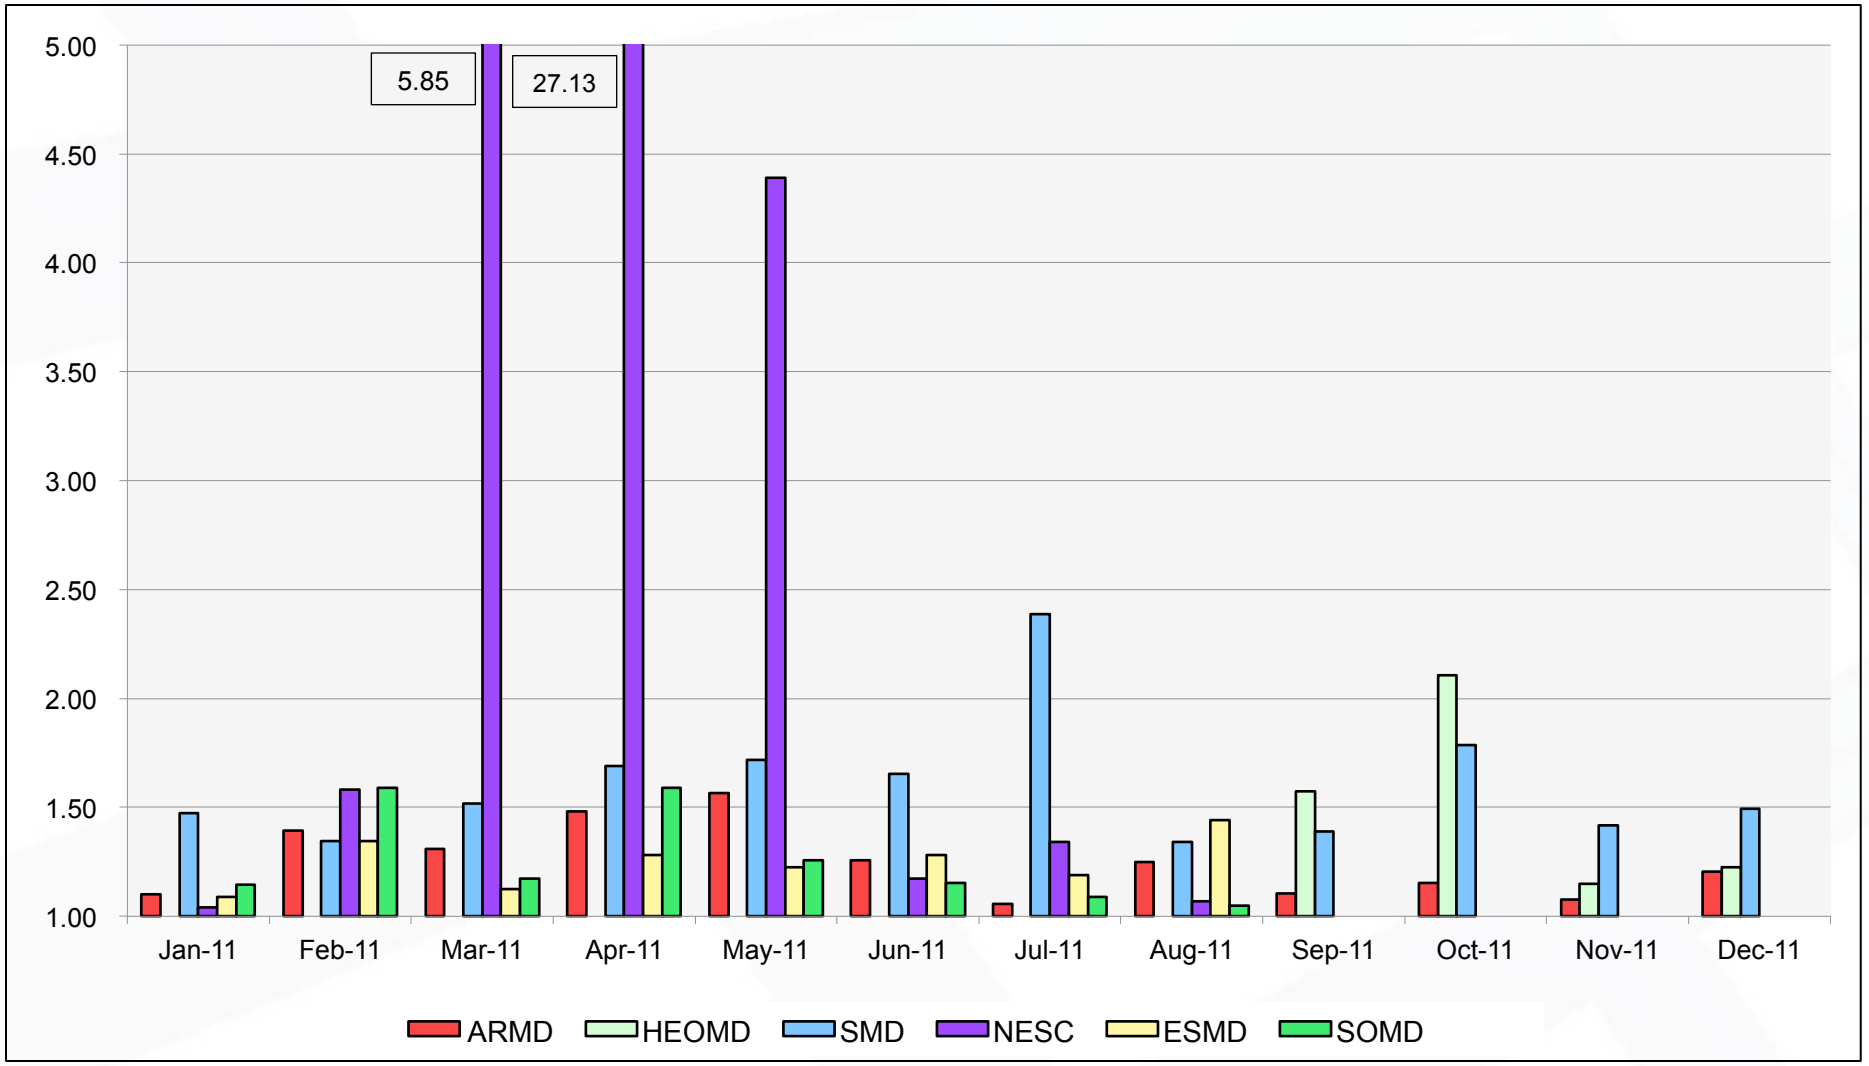

Columbia:
SBUs Reported, Normalized to 30-Day Month

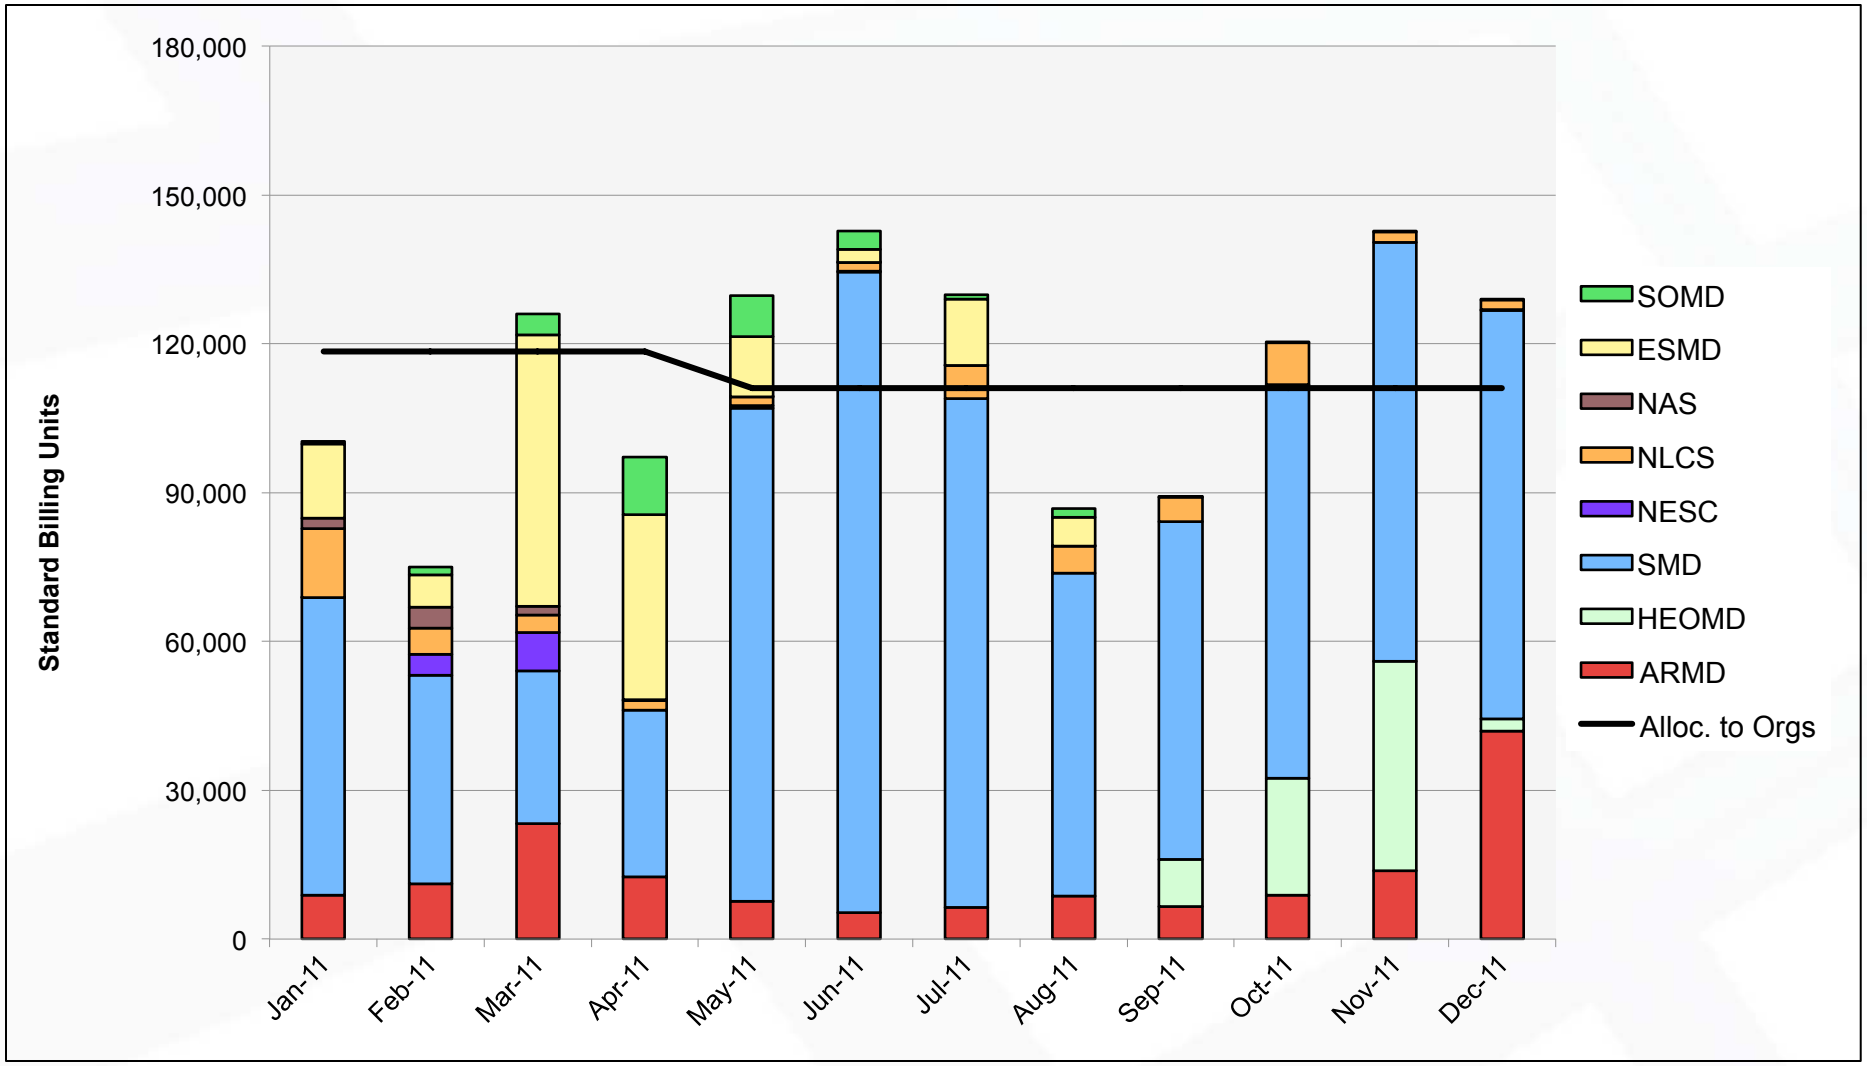

Columbia:
Monthly SBUs by Run Time

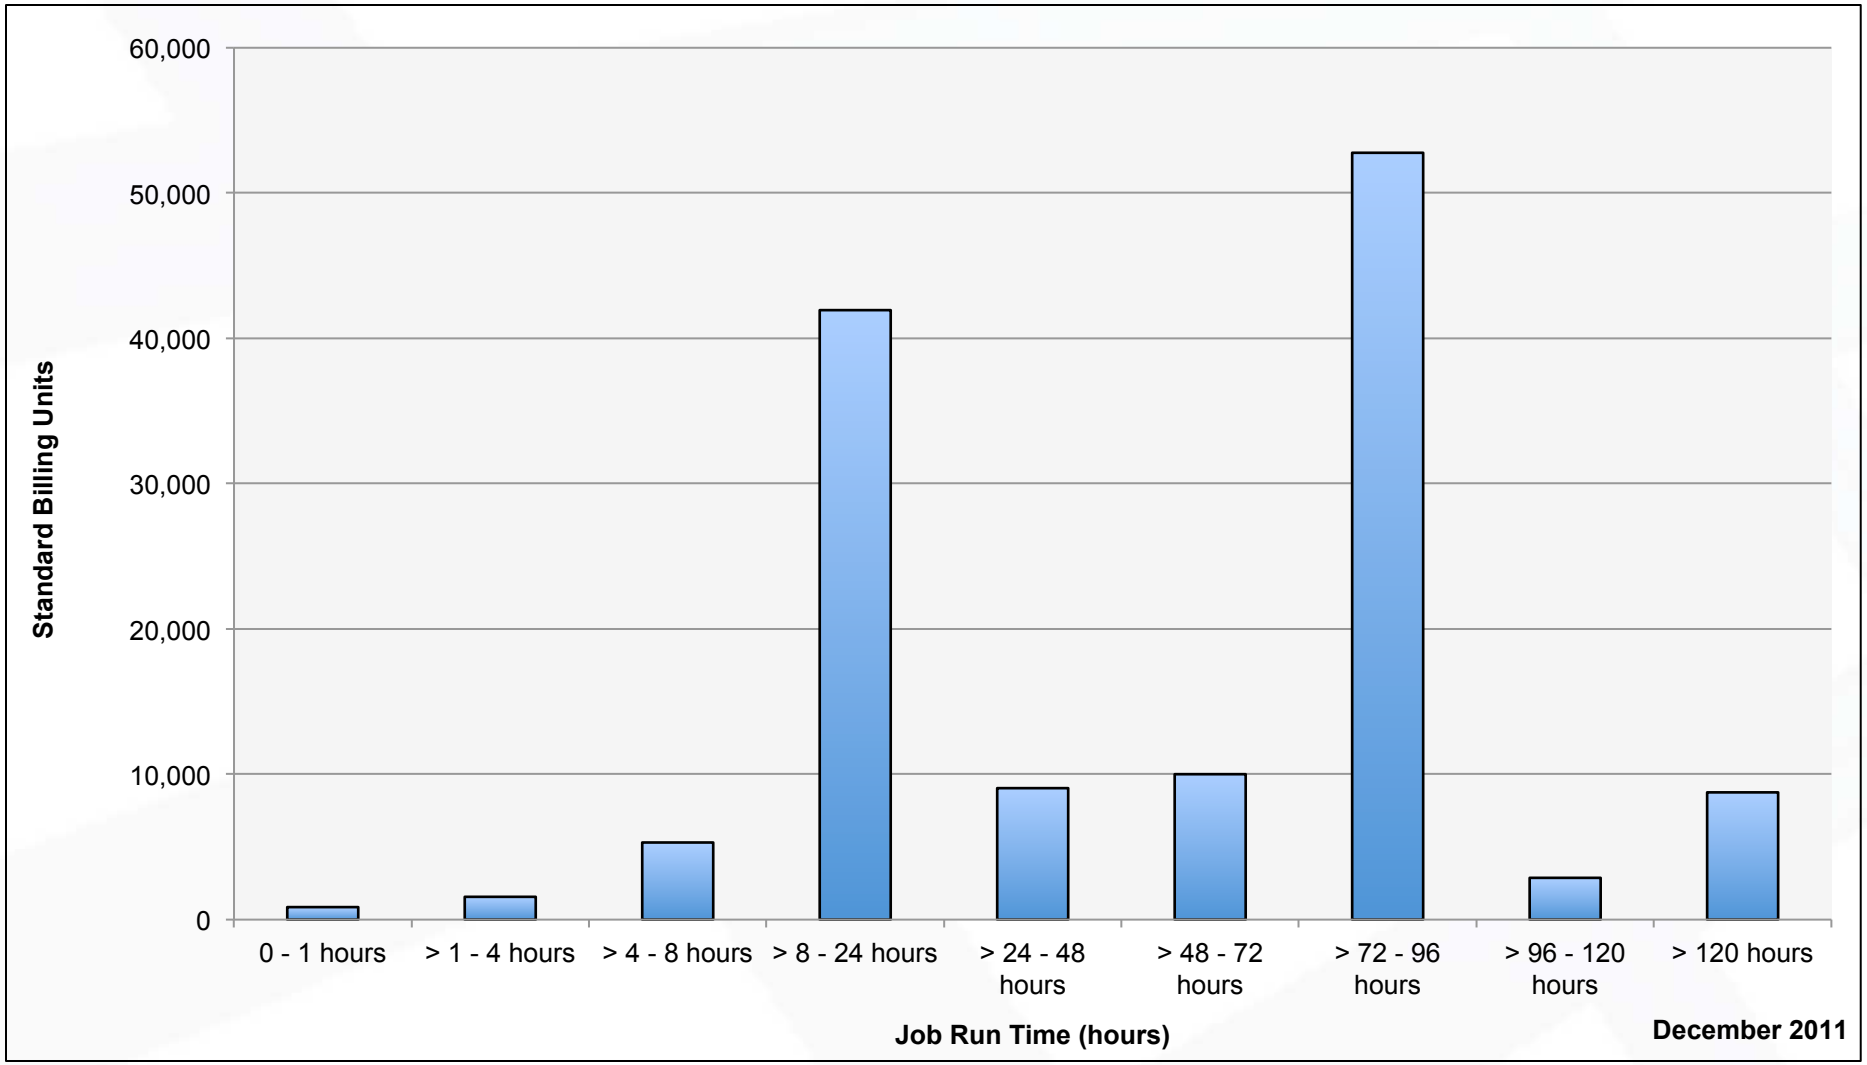

Columbia:
Monthly Utilization by Size and Mission

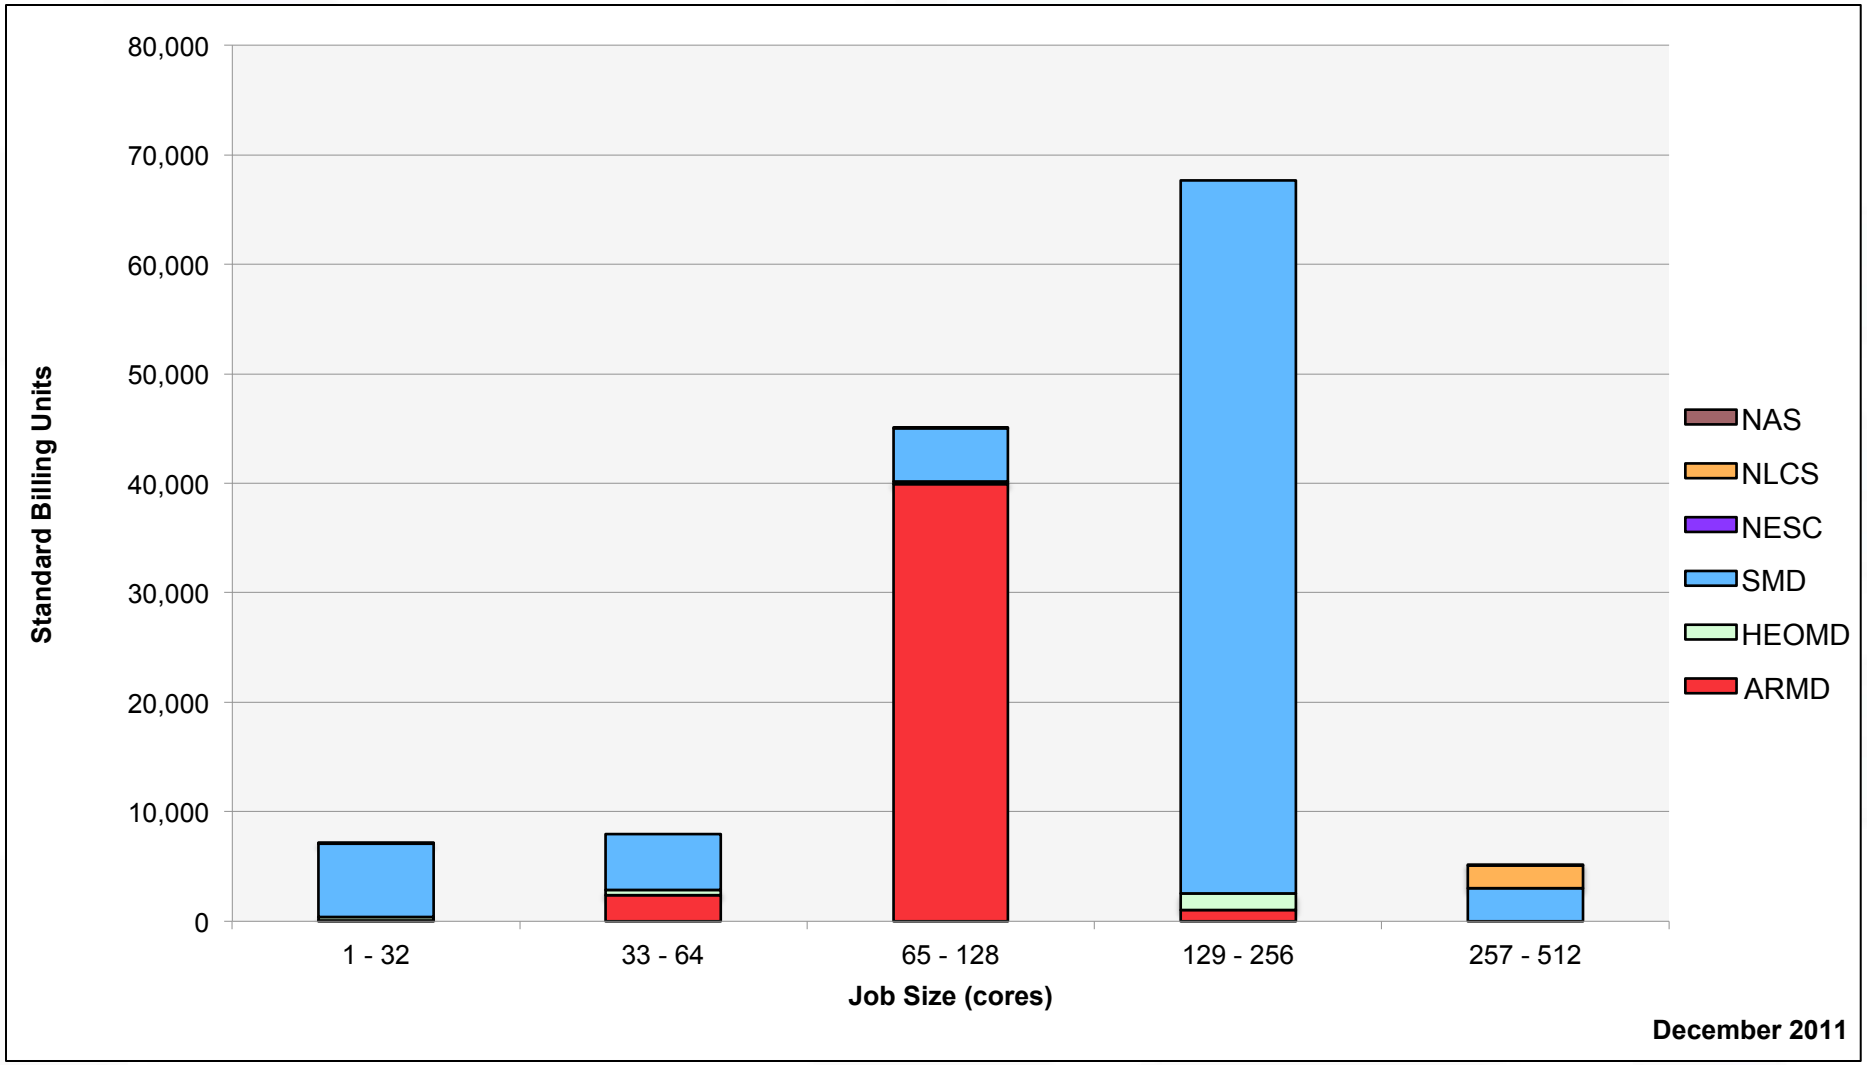

Columbia:
Monthly Utilization by Size and Length

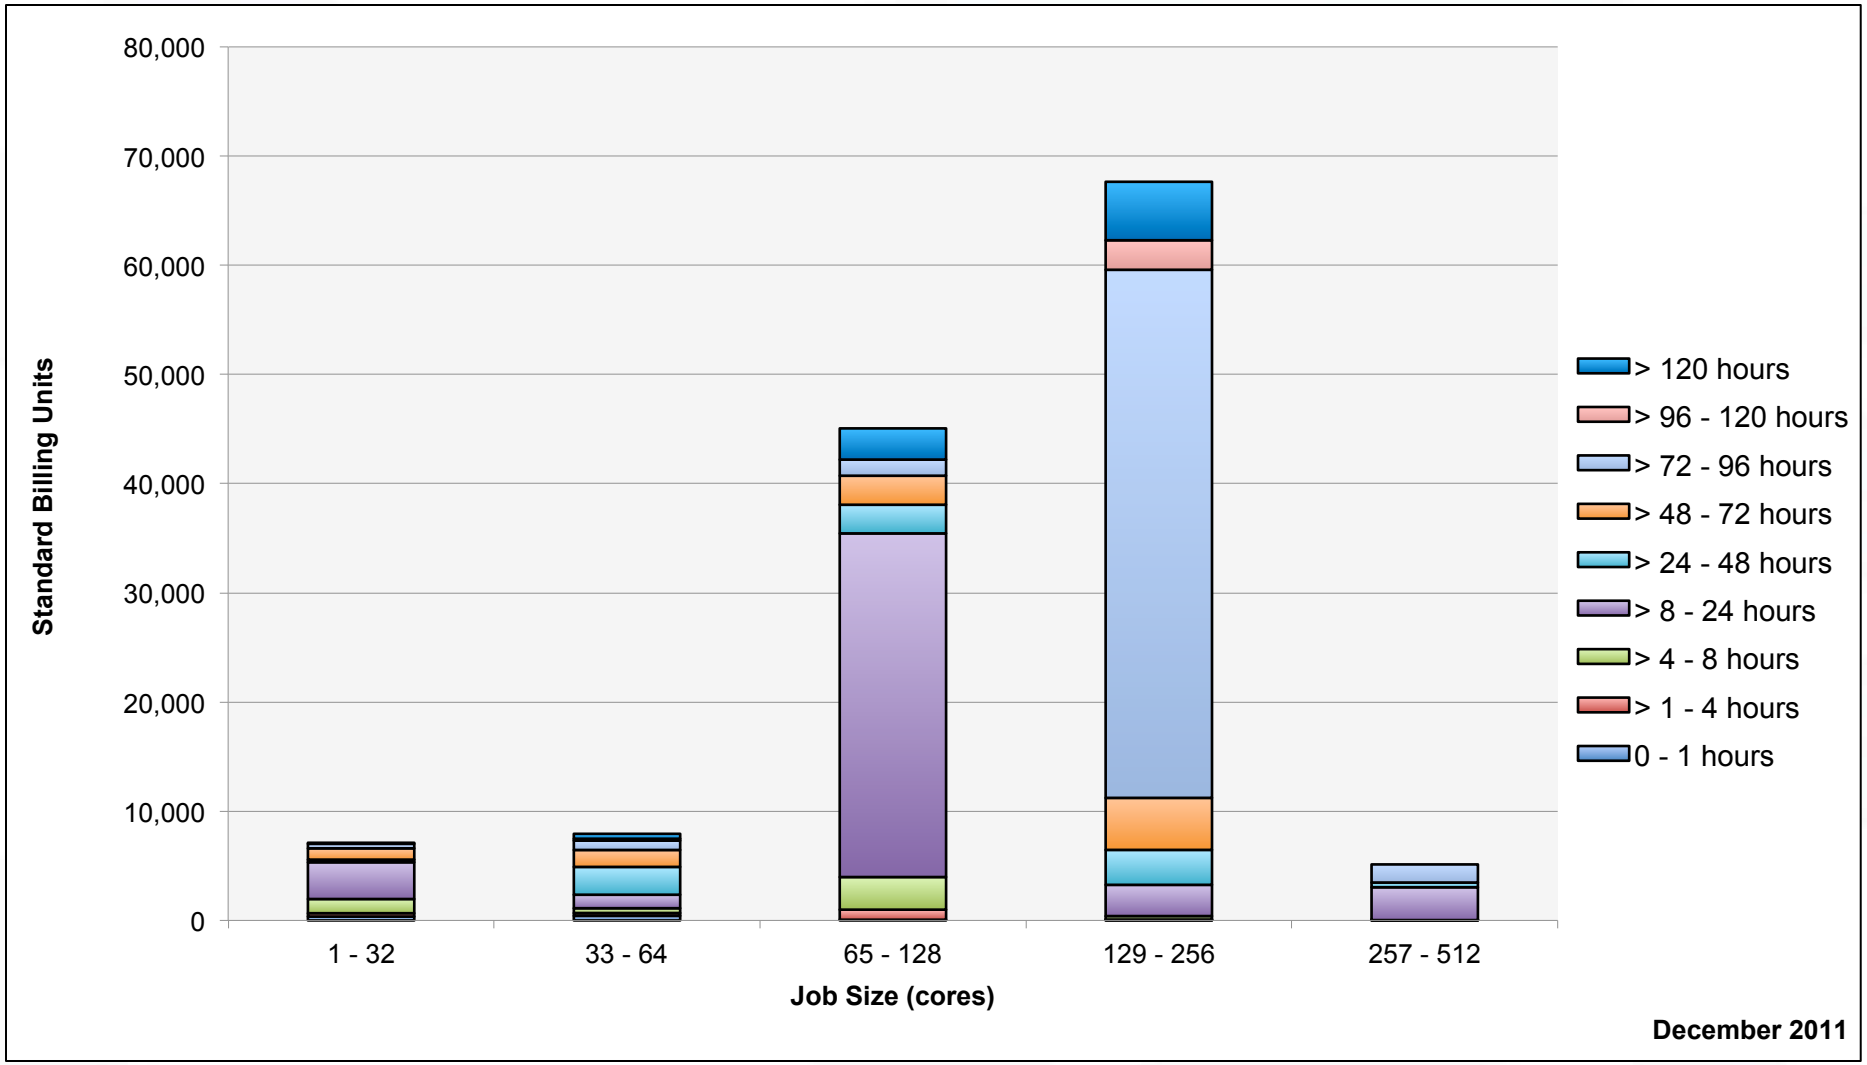

Columbia:
Average Time to Clear All Jobs

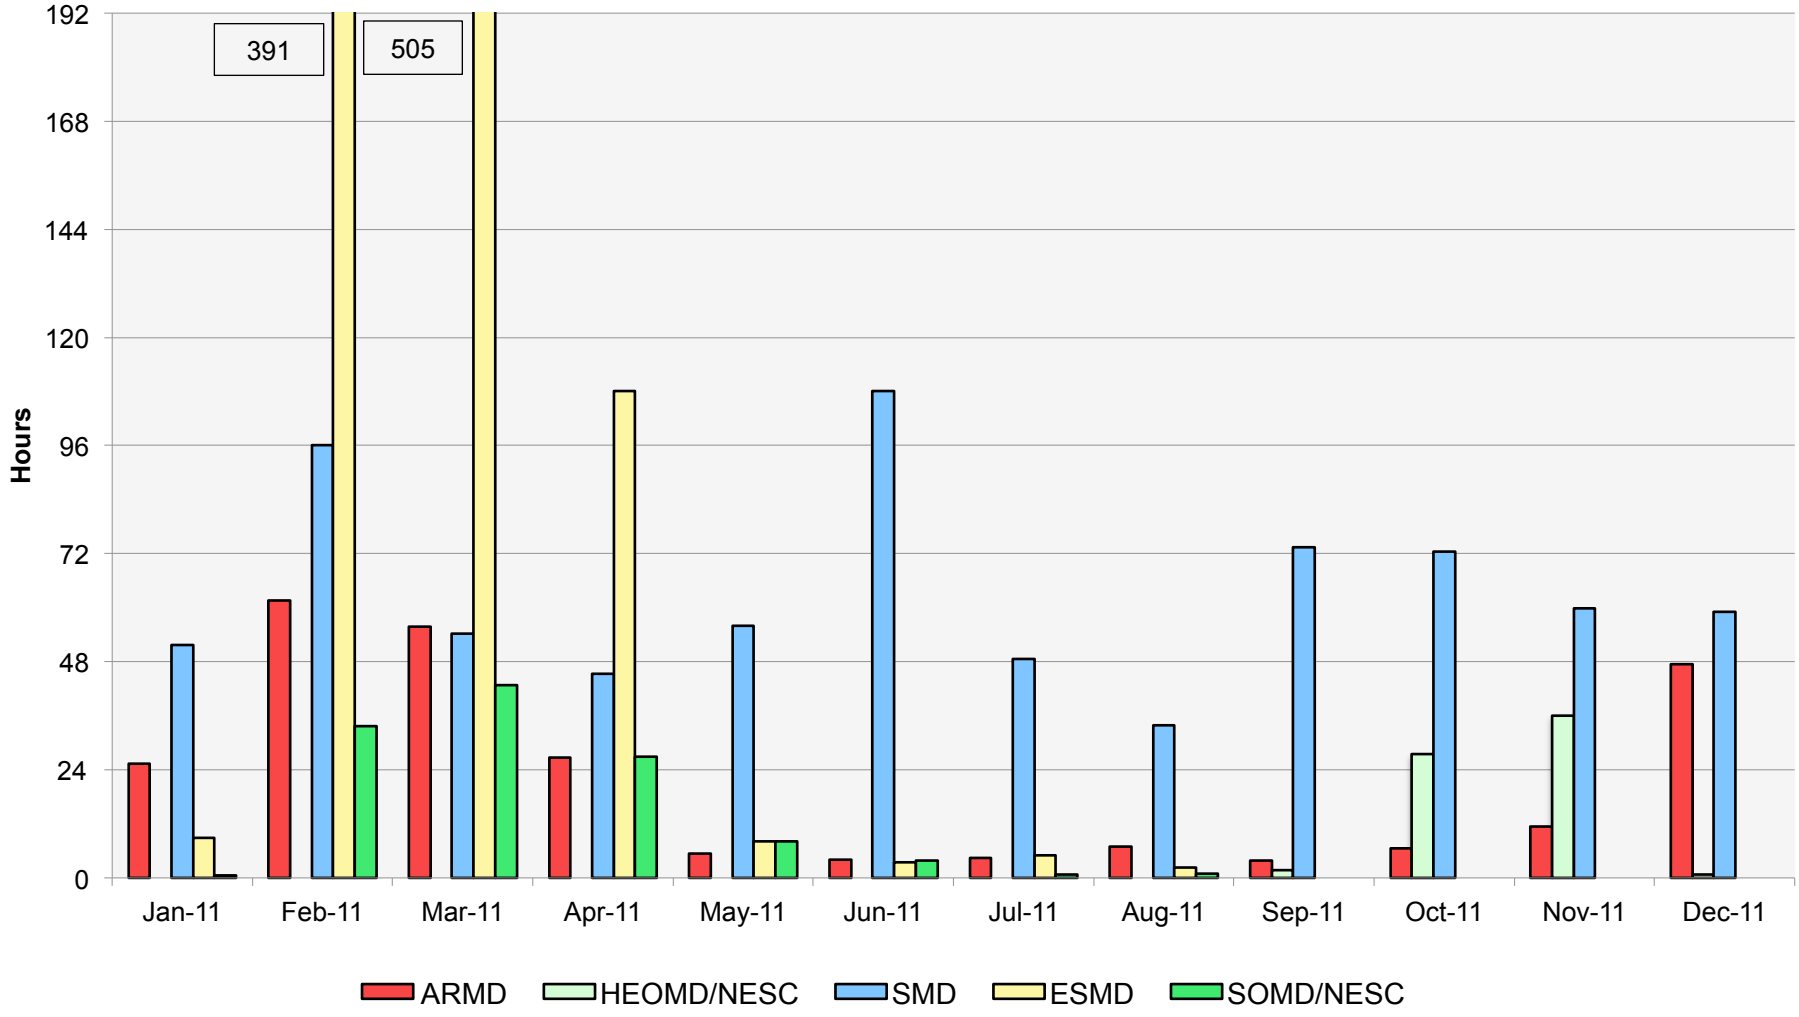

Columbia:
Average Expansion Factor