

National Aeronautics and Space Administration

High-End Computing Capability

William Thigpen
HECC Component Manager
william.w.thigpen@nasa.gov

July 10, 2024

Selected Monthly Accomplishments

- **Improving Efficiency**
  - Worked with vendor DataDirect Networks (DDN) to develop a new architecture to increase the reliability of the Lustre filesystems. The new architecture migrates the Hot Pools processes (which move files between the small, fast-flash based storage and large, slower, spinning disk storage) from the Lustre Object Storage Service (OSS) servers to five new servers, reducing the load on the OSS servers.
- **Improving System Stability**
  - Worked with the Orbiting Carbon Observatory (OCO-3) users to move their processing off the /home5 filesystem to the nfshome45 server to avoid denial of service. The main issue: the thousands of serial processes read the same five input files, causing massive file logjam. This work restored the filesystem stability and improved accessibility for users on /home5. Also suggested reducing the number of jobs that are run concurrently to reduce the pressure on the filesystems.
  - Solved issues and made improvements to NAS processes, resulting in increased system stability, resulting in no unplanned outages in the last month.
- **Improving User Experience**
  - Released GPU Usage to the myNAS website. This allows all users (regular users, principal investigators, and managers) to see both their GPU and CPU job information in one space. Only users with GPU allocation will be shown GPU information.
- **Release of RFP for Object Store Filesystem**
  - Released the object store filesystem request for proposal (RFP) on June 25. The selected system will be used to replace the existing, soon-to-be end-of-life Lustre filesystem (export1), which is used for the data portal back-end storage.
- **Outreach and Collaboration**
  - The Visualization and Data Analysis team completed several animations for: 1) launch environments in support of SLS, combined (solid rocket booster-space shuttle main engine) exhaust and water (liquid+vapor) species using user-supplied cutplanes and scalar mapped surfaces; 2) user-supplied cutplanes and surface flow for the High Lift Prediction Workshop.

Compute Module 2 Installed at the MSF

- Compute Module 2, which will house the Aitken Milan expansion and the next-generation CPU system, was delivered and placed on June 14-15 at the Modular Supercomputing Facility (MSF) at NASA Ames.
  - The module was manufactured in Maryland and transported across the U.S. by truck in two sections.
  - At 16 feet wide and 60 thousand pounds per section, the module and its trucking rig were so large that HECC facilities experts had to set up a special coordination effort with Ames Codes JP, JCM, and Q to move the module onto the Ames site and lift it into place by crane.
- The module's supporting infrastructure was also delivered:
  - One 3,750 kilovolt-amperes transformer that provides 480 volts of power to the computers and cooling system.
  - Two 380-ton evaporative coolers that provide 90°F cooling water to the computers and two 1,200-gallons-per-minute pumps to circulate the water.
- Compute Module 2 can house 24 compute racks consuming a total of 3 megawatts of power. The initial computer installation planned for September 2024 will contain the next Aitken expansion, consisting of 512 nodes of AMD Milan processors.

IMPACT: Compute Module 2 provides three megawatts of power and cooling to allow for expansion of HECC's modular supercomputing resources.

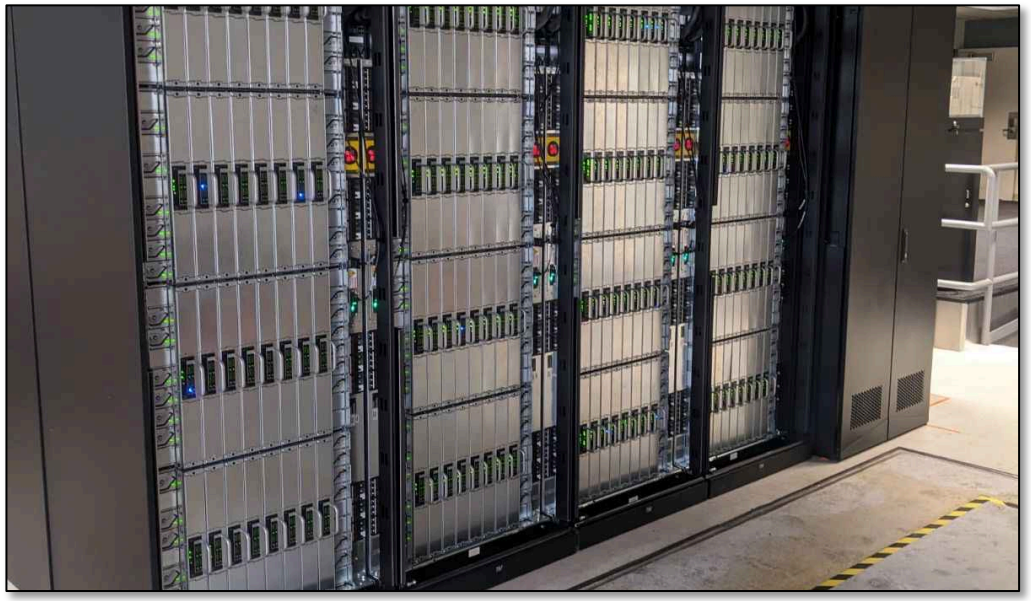


Compute Module 2 is set in place next to the existing Compute Module 1 at the Modular Supercomputing Facility. *Chris Tanner/NASA Ames*

New Aitken Expansion Nodes in Production

- HECC's Supercomputing Systems team completed the new Aitken expansion and placed the system into production—further increasing HECC's compute capability.
- The expansion increases Aitken's performance by an additional 20%, or a total of 15.7 petaflops.
- The expansion comprises four Apollo 9000 "Badger" racks with 512 AMD 7763 (Milan) compute nodes, plus one I/O rack and one coolant distribution unit. The HPE Performance Cluster Manager (HPCM) will be used to manage the cluster configuration.
  - Each node has two 64-core AMD third-generation System-on-Chip (SoC) EPYC 7763 Milan processors providing a total of 65,536 additional cores.
  - The 512 nodes have a total of 262 terabytes of memory (512 gigabytes/node).
- The Aitken expansion nodes use the PBS scheduler and use the same InfiniBand fabric as the existing Aitken nodes.
- These nodes will be moved to Compute Module 2 in September 2024.

IMPACT: This Milan expansion to the Aitken supercomputer provides more computing capacity and resources to support engineers and scientists working on NASA mission projects.

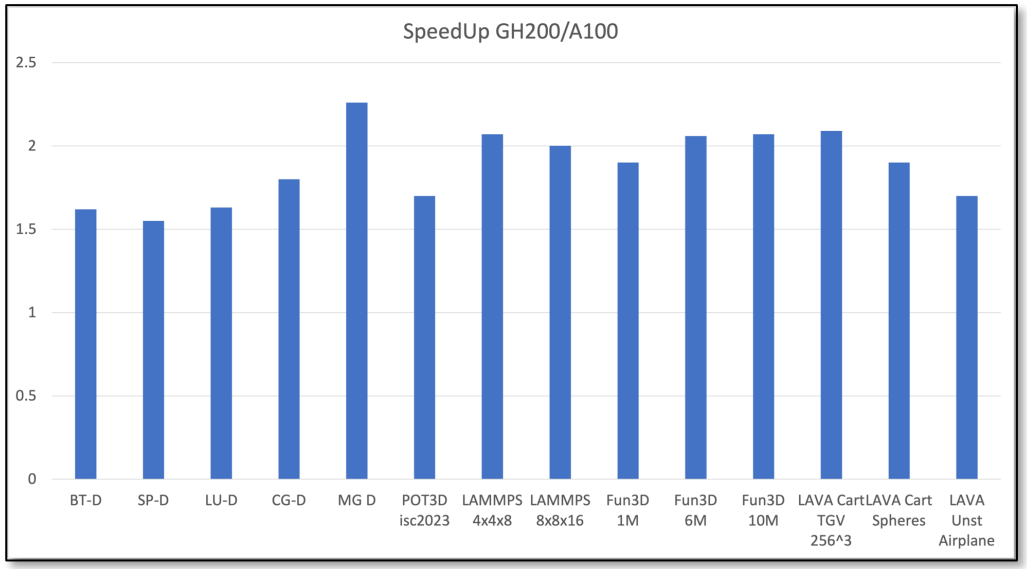


The Aikten supercomputer was expanded with four new Apollo 9000 racks to provide more computational capability for NASA users.
Alex Pelkey, NASA/Ames

GH200 Performance and Ease-of-Use Evaluation

- The Application Performance and Productivity (APP) team, in collaboration with application developers, ported a set of high-performance computing benchmarking kernels, a mini-app, and three full-scale applications to the NVIDIA GH200 pathfinding system at the NAS facility.
  - The pathfinding system consists of two NVIDIA GH200 nodes. Each node contains the NVIDIA "Grace-Hopper" superchip that combines a 72-core NVIDIA Grace CPU and an NVIDIA H100 GPU using NVLink-C2C—a memory-coherent, high-bandwidth, and low-latency interconnect for superchips—with access to 480 gigabytes of server-class LPDDR5X CPU memory and 96 GB HBM3 GPU memory.
- The purpose of the effort was two-fold:
  - To obtain a rough estimate of HPC application performance improvement and power consumption compared to the NVIDIA A100 Tensor Core GPU nodes that are currently available on HECC systems
  - To evaluate the porting effort from NVIDIA x86-hosted A100 nodes to GH200 nodes.
- The APP team conducted single-node experiments only, using a variety of test input data sets. Their initial observations show:
  - The GH200 delivered a 1.5 - 2.2x overall speed-up in performance compared to the A100 without requiring more energy. The full-scale applications fared better than the kernels and mini-apps.
  - Most of the relevant APP-developed tools and all the CPU dependencies required by the applications were ported to the GH200 without major problems.

IMPACT: The NVIDIA GH200 HPC application evaluation enabled a first glance into the expected system performance, energy usage improvements, and ease of portability for HECC users.

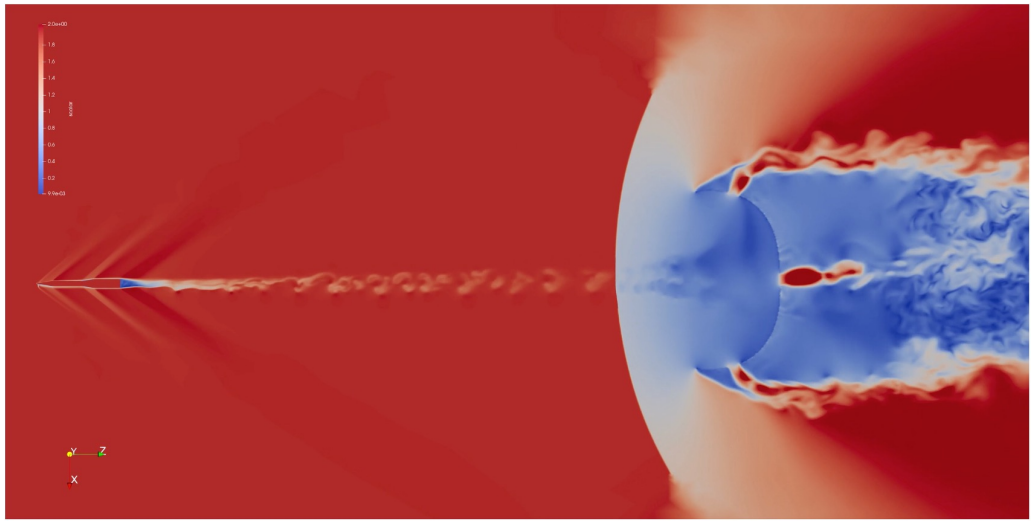


This chart shows the performance improvement of HPC benchmark kernels, mini-apps, and full-scale applications for the NVIDIA GH200 pathfinding system. A variety of relevant test input datasets was used for the evaluation effort.

Predicting Supersonic Parachute Inflation on Mars*

- Researchers at NASA's Jet Propulsion Laboratory, in collaboration with Stanford University and Stevens Institute of Technology, are using the AERO computational suite on HECC supercomputers to model the complex environment around what will be the largest supersonic parachute ever flown in the Martian atmosphere. The agency's Mars Sample Return (MSR) campaign is currently developing the largest Mars lander to date, which will be used to retrieve the planetary samples being taken by the Curiosity rover.
- The team's AERO software is giving them the ability to perform high-fidelity coupled fluid-structure interaction (FSI) simulations that incorporate both computational fluid dynamics and structural responses of the parachute model to understand how the two interact.
- Their current goal is to use flight data from the Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) to validate AERO, which will enable the tool to be used for computational evaluation of the MSR's parachute system.
- Using the Aitken and Electra supercomputers, with simulations running several weeks across more than 20 AMD Rome nodes, the team's results have identified several features of importance in the model and established many best practices for modeling the complex and varied physics of FSI simulations, which are being incorporated by groups inside and outside NASA for other computation tools.

IMPACT: Validating these computational models brings NASA one step closer to being able to accurately evaluate supersonic parachute designs without needing to perform expensive wind tunnel tests.



A cross-section simulation of airflow around the Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) model's payload and parachute structures during supersonic parachute inflation. The Mach number of the flow is show in red (high) and blue (low).
Marcus Lobbia, NASA/JPL

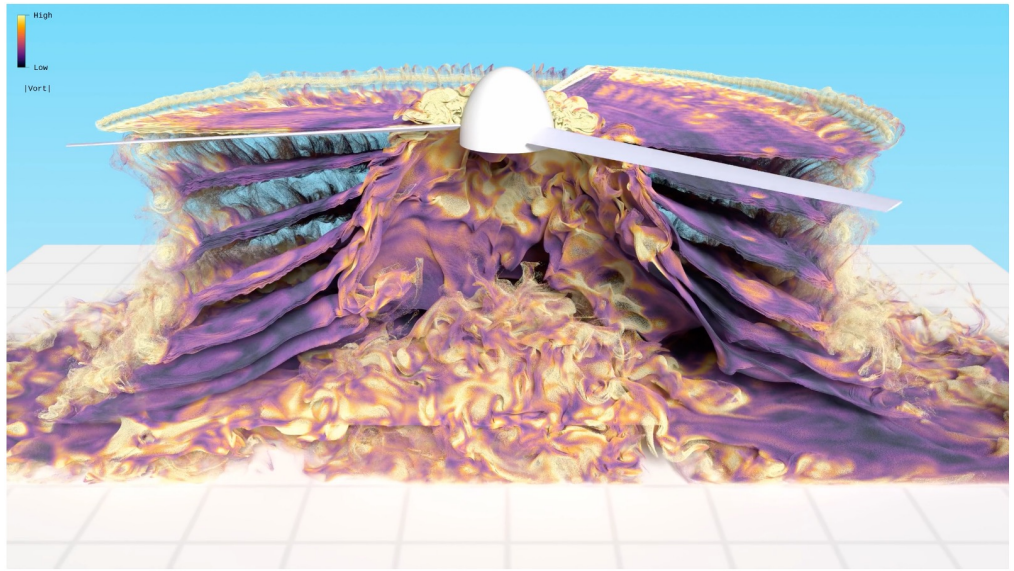
** HECC provided supercomputing resources and services in support of this work.*

Simulating NASA's Quadcopter in Ground Effect*

- NASA supports the development of urban air mobility (UAM) vehicles by studying new designs of vertical take-off and landing (VTOL) vehicles for a variety of missions, ensuring they will operate safely and efficiently.
  - Studying the complex interactions between the rotors and other components—like the fuselage or wings—while vehicles are still being developed, requires high-fidelity computational fluid dynamics (CFD).
- Aerospace engineers at NASA Ames performed CFD simulations of rotors and air taxis in "ground effect," where the addition of a ground surface makes simulations more challenging, substantially increasing the size of the computational grids vs. simulating the vehicle alone.
- The vehicles, ground, and surrounding air were modeled using hundreds of millions of grid points. NASA's OVERFLOW CFD solver was run on HECC's Pleiades, Electra, and Aitken supercomputers. The solver took about 40 days on 2,000 CPUs to complete the single-rotor case.
- These simulations enable researchers to better characterize the vehicle aerodynamics in proximity to the ground. The engineers found that interactions between the rotors lead to the emergence of dangerous areas where a person could not stand due to high winds and large gusts.
- The Ames team will compare their simulation results for different vehicle designs to highlight how ground effect impacts performance and safety.

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

IMPACT: The simulation data produced using HECC systems could be used to build maps of dangerous areas in the vehicle surroundings, which will inform safe routes for passengers and operators to embark or disembark from the vehicle.



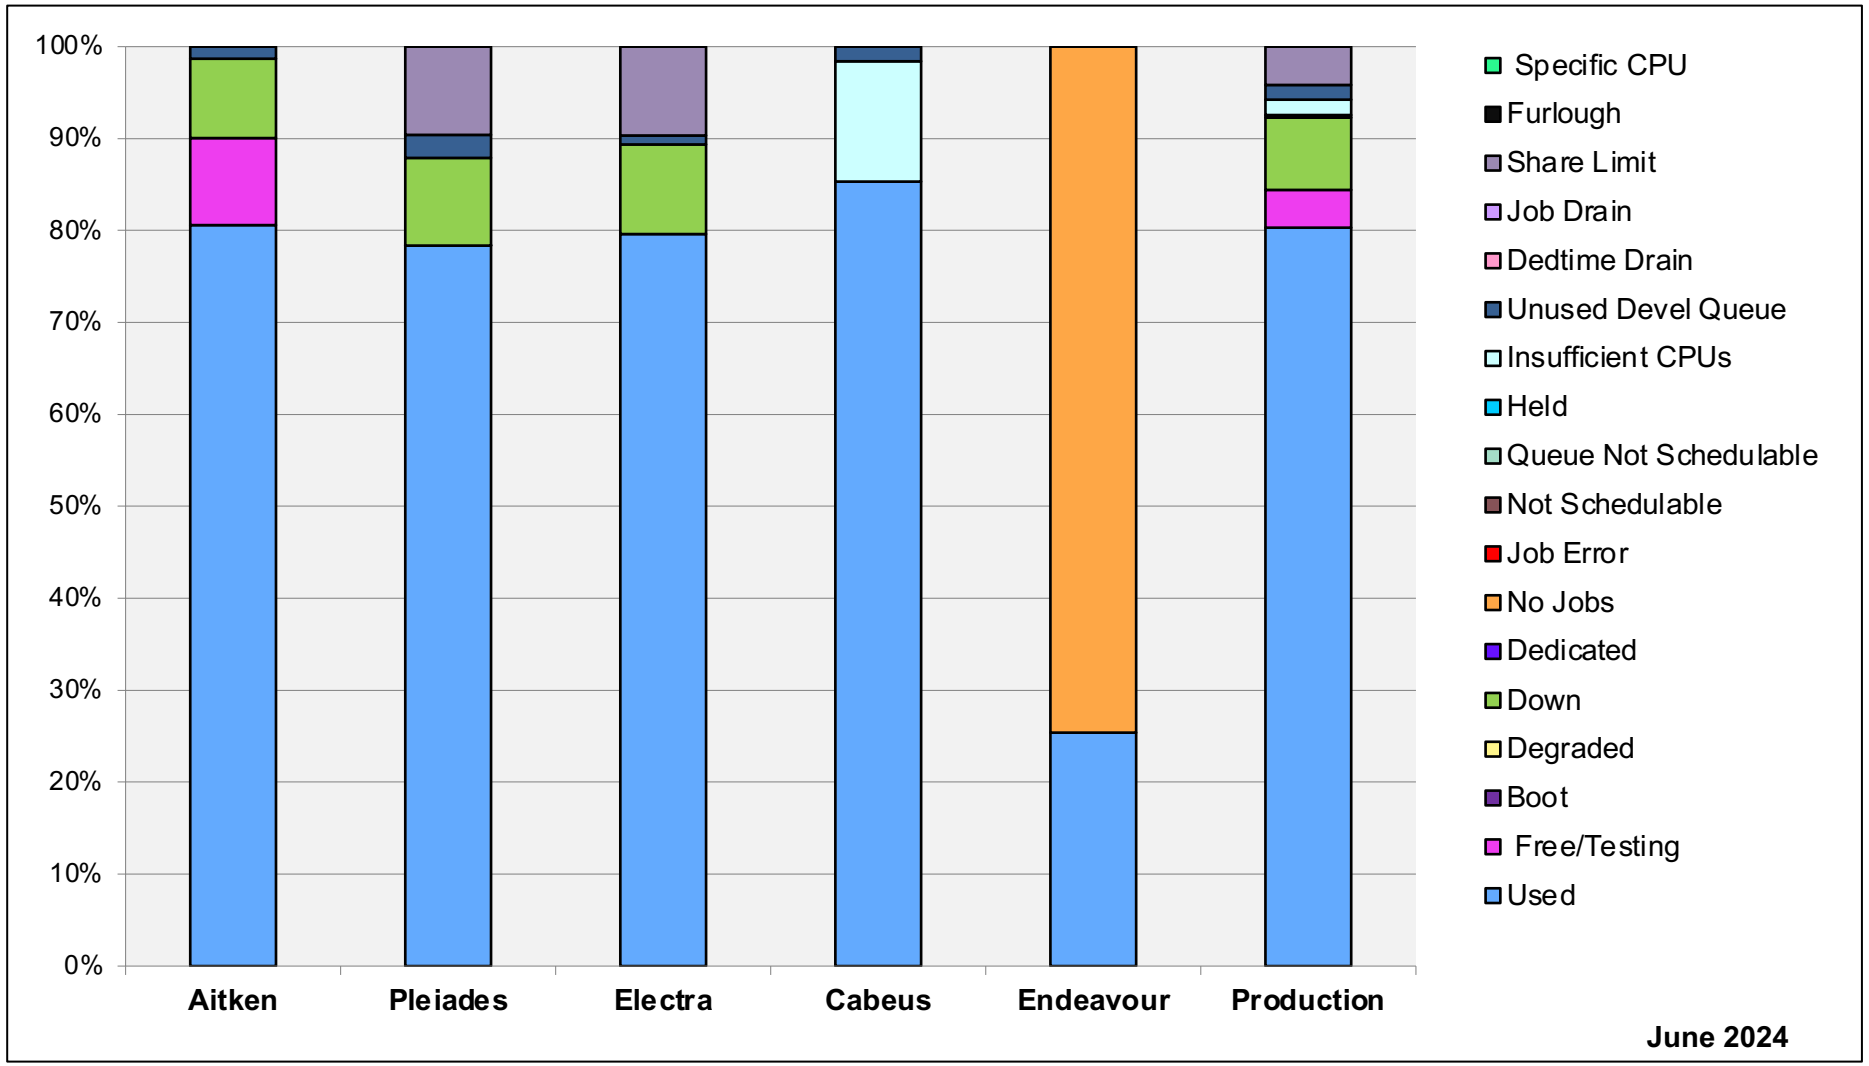
Flow visualization of NASA's quadcopter isolated rotor in hover in ground effect, with a plane located at one-radius below the rotor. The flow visualization reveals the complex wake of a single rotor in ground effect. *Tim Sandstrom, NASA/Ames*

Papers

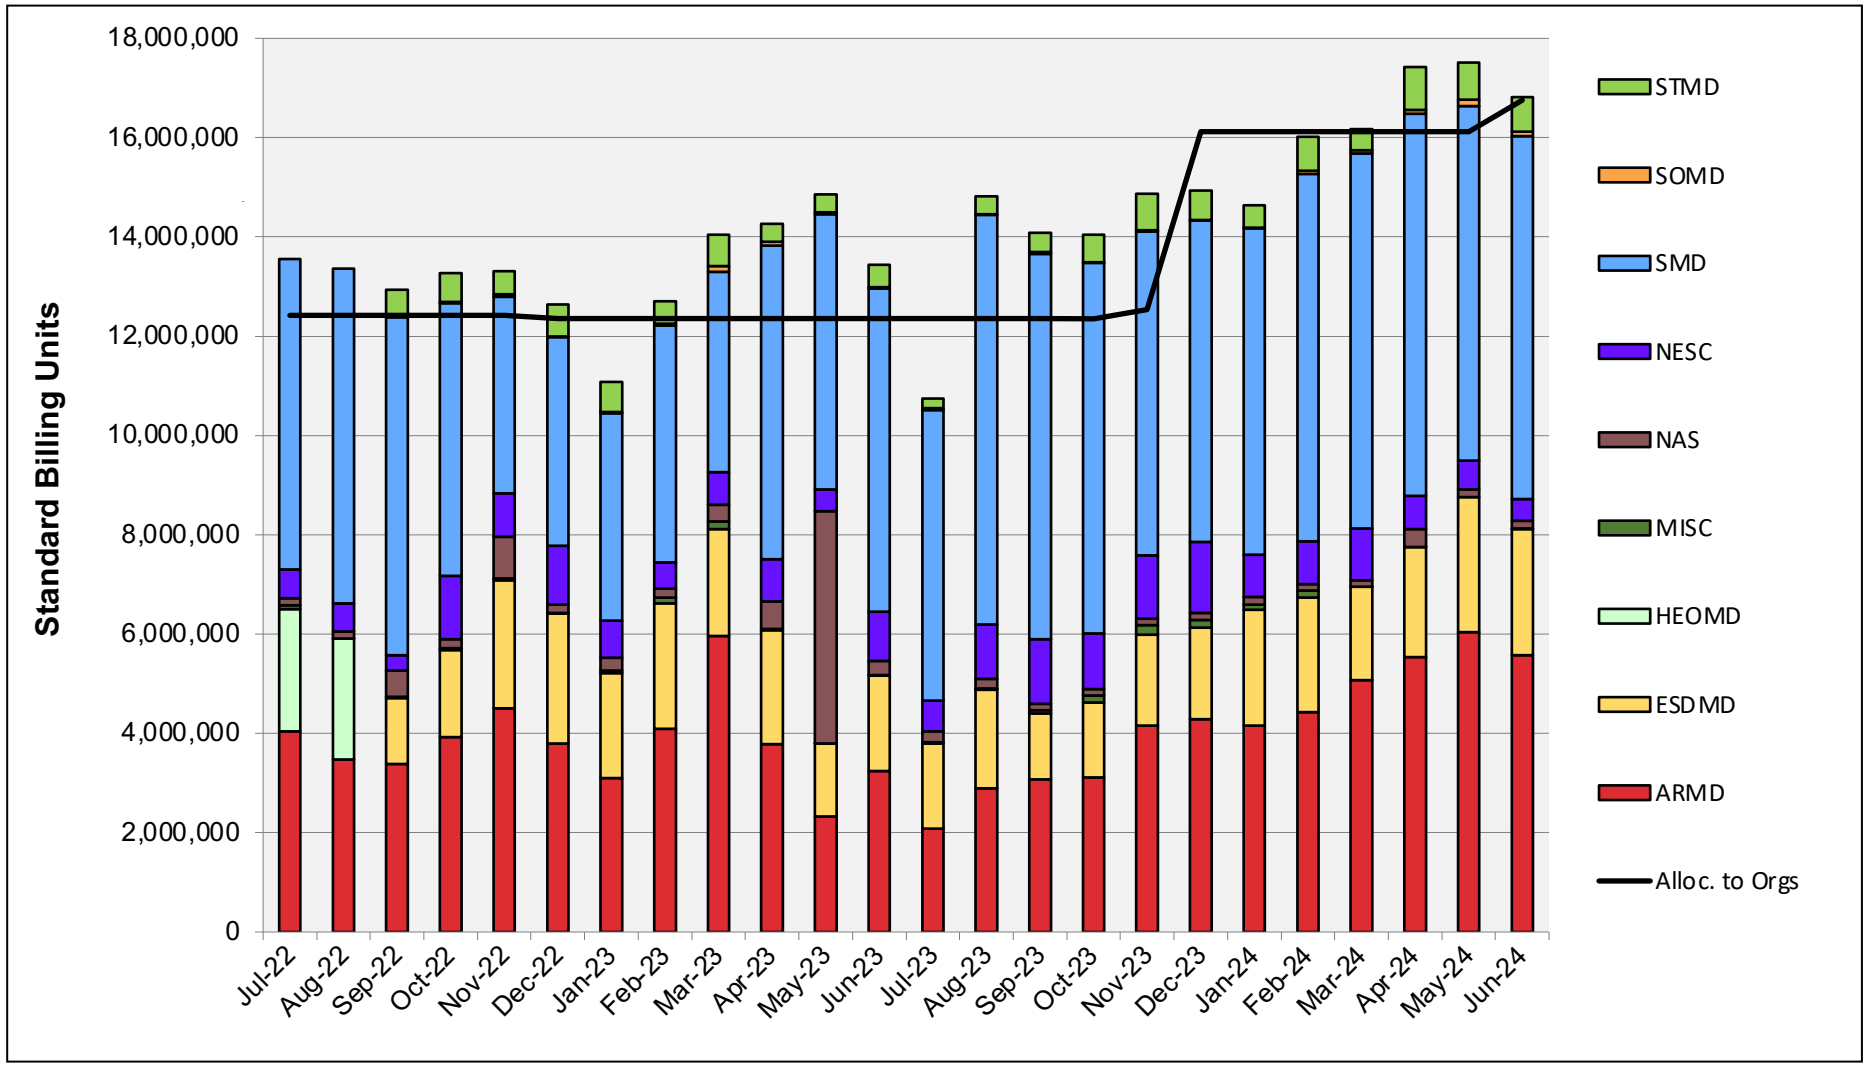
- "**TESS Giants Transiting Giants. IV. A Low-Density Hot Neptune Orbiting a Red Giant Star**," S. Grunblatt, et al., *The Astronomical Journal*, vol. 168, no. 1, published online June 5, 2024. *
  https://iopscience.iop.org/article/10.3847/1538-3881/ad4149/meta
- "**ICM-SHOX. I. Methodology Overview and Discovery of a Gas-Dark Matter Velocity Decoupling in the MACS J0018.5+1626 Merger**," E. Silich, et al., *The Astrophysical Journal*, vol. 968, no. 2, published online June 12, 2024. *
  https://iopscience.iop.org/article/10.3847/1538-4357/ad3fb5/meta
- "**2023 DZ2 Planetary Defense Campaign**," V. Reddy, et al., *The Planetary Science Journal*, vol. 5, no. 6, published online June 13, 2024.*
  https://iopscience.iop.org/article/10.3847/PSJ/ad4a6d/meta
- "**The Imprint of Dark Matter on the Galactic Acceleration Field**," A. Arora, et al., arXiv:2406.12957 [astro-ph.GA], June 18, 2024. *
  https://arxiv.org/abs/2406.12957
- "**Multistructured Accretion Flow of Sgr A* II: Signatures of a Col Accretion Disk in Hydrodynamic Simulations of Stellar Winds**," M. Balakrishnan, et al., arXiv:2406.14631 [astro-ph.HE], June 20, 2024. *
  https://arxiv.org/abs/2406.14631
- "**Understanding the Variability of Helium Abundance in the Solar Corona using Three-Fluid Modeling and UV Observations**," L. Ofman, Yogesh, S. Giordano, arXiv:2406.14897 [astro-ph.SR], June 21, 2024. *
  https://arxiv.org/abs/2406.14897
- "**Eccentricity and Inclination of Massive Planets Inside Low-Density Cavities: Results of 3D Simulations**," M. Romanova, et al., arXiv:2406.18834 [astro-ph.EP], June 27, 2024. *
  https://arxiv.org/abs/2406.18834

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

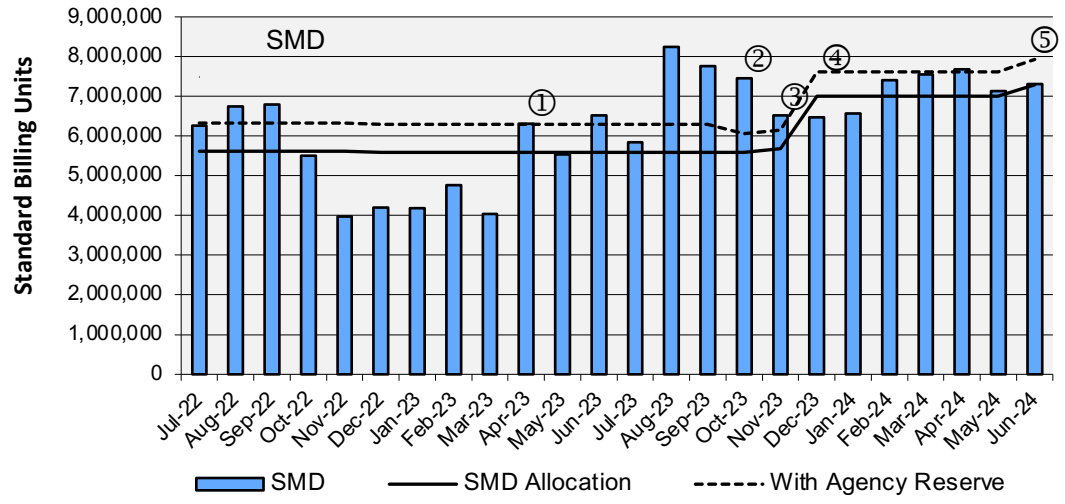
HECC Utilization

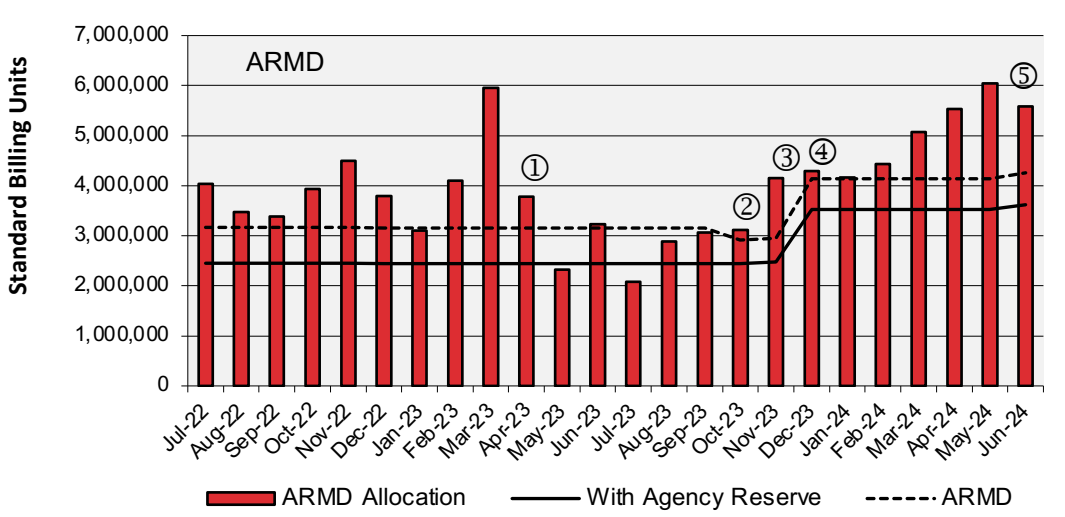


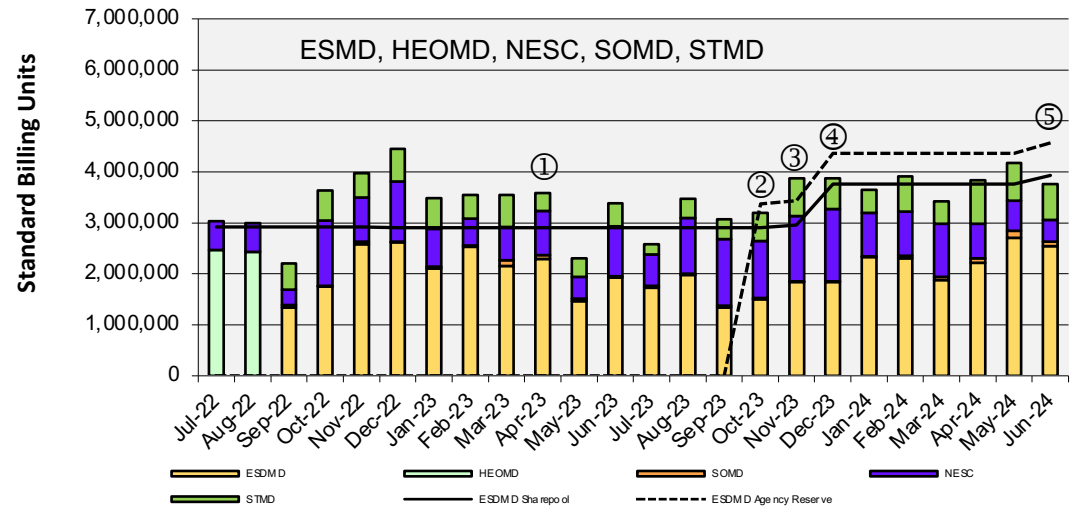
HECC Utilization Normalized to 30-Day Month



HECC Utilization Normalized to 30-Day Month







① 4 Rome Apollo Racks added to Aitken

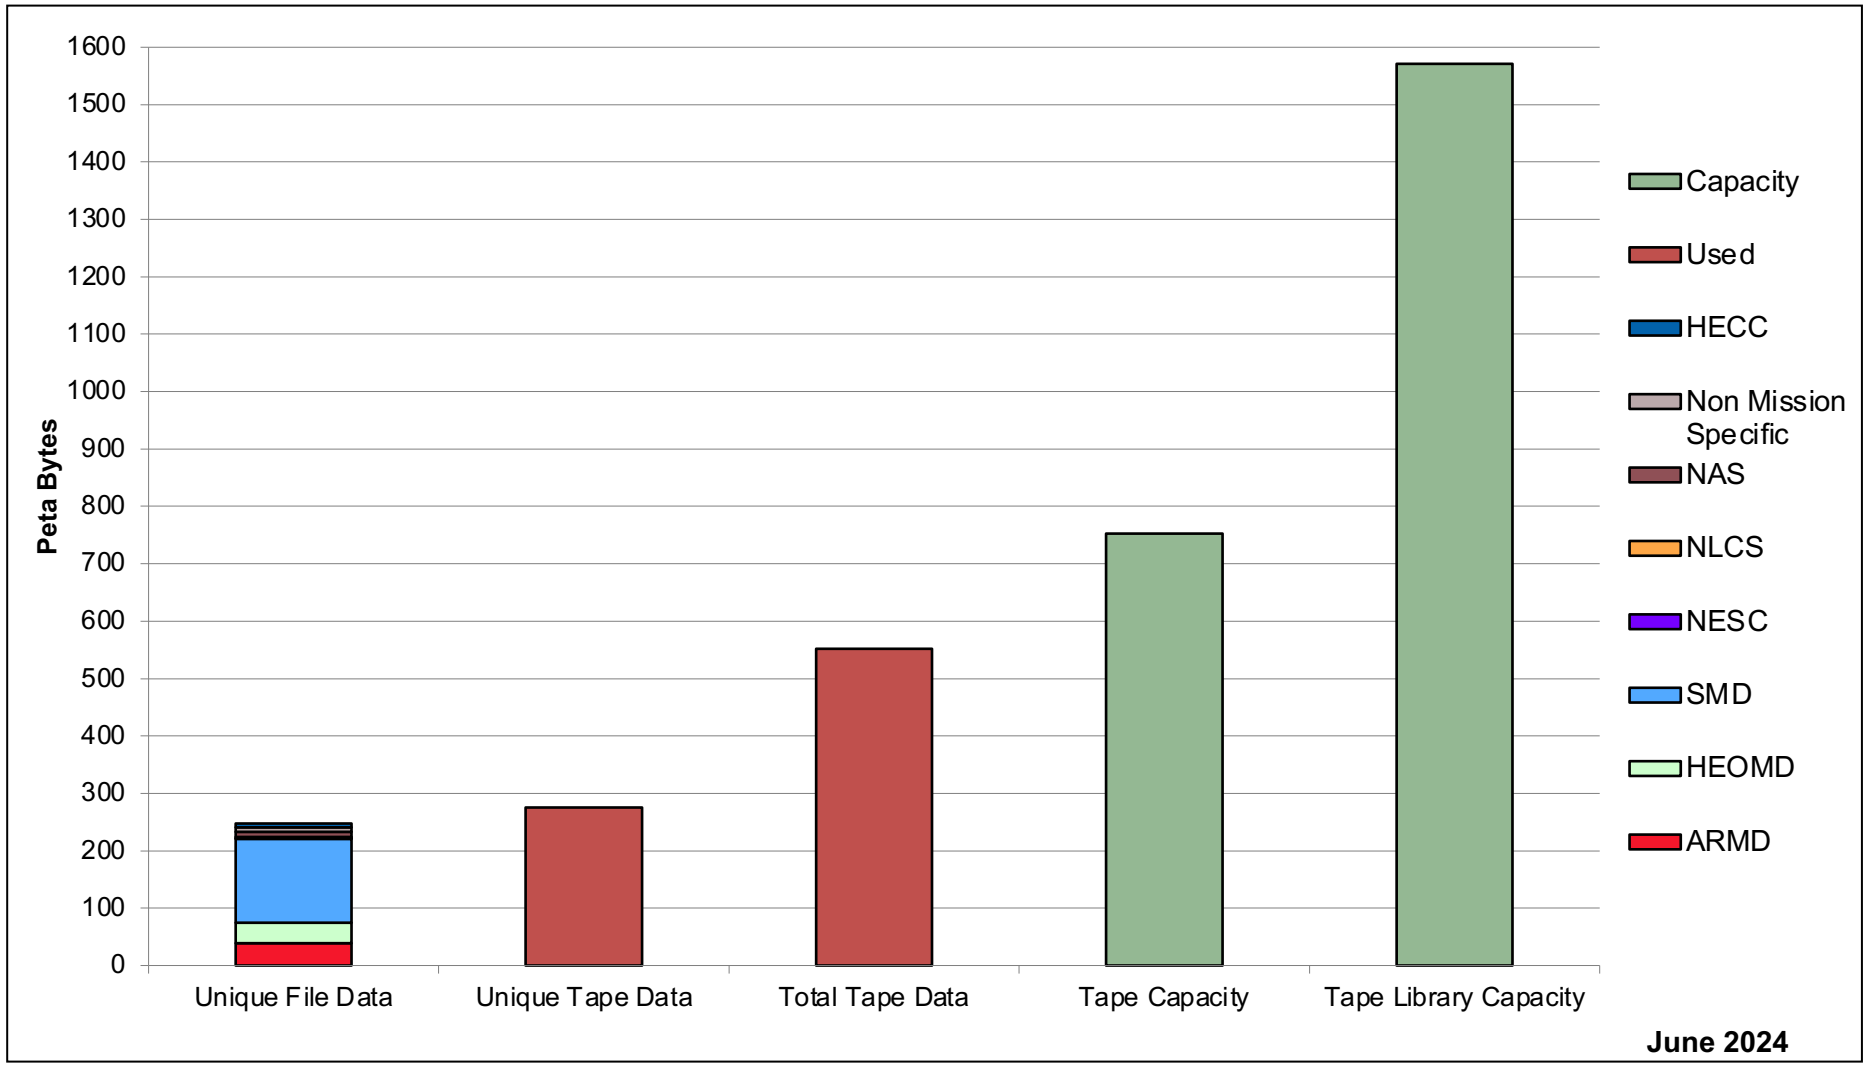
② 4 Rome Apollo Racks added to Aitken

③ Cabeus v100 nodes added

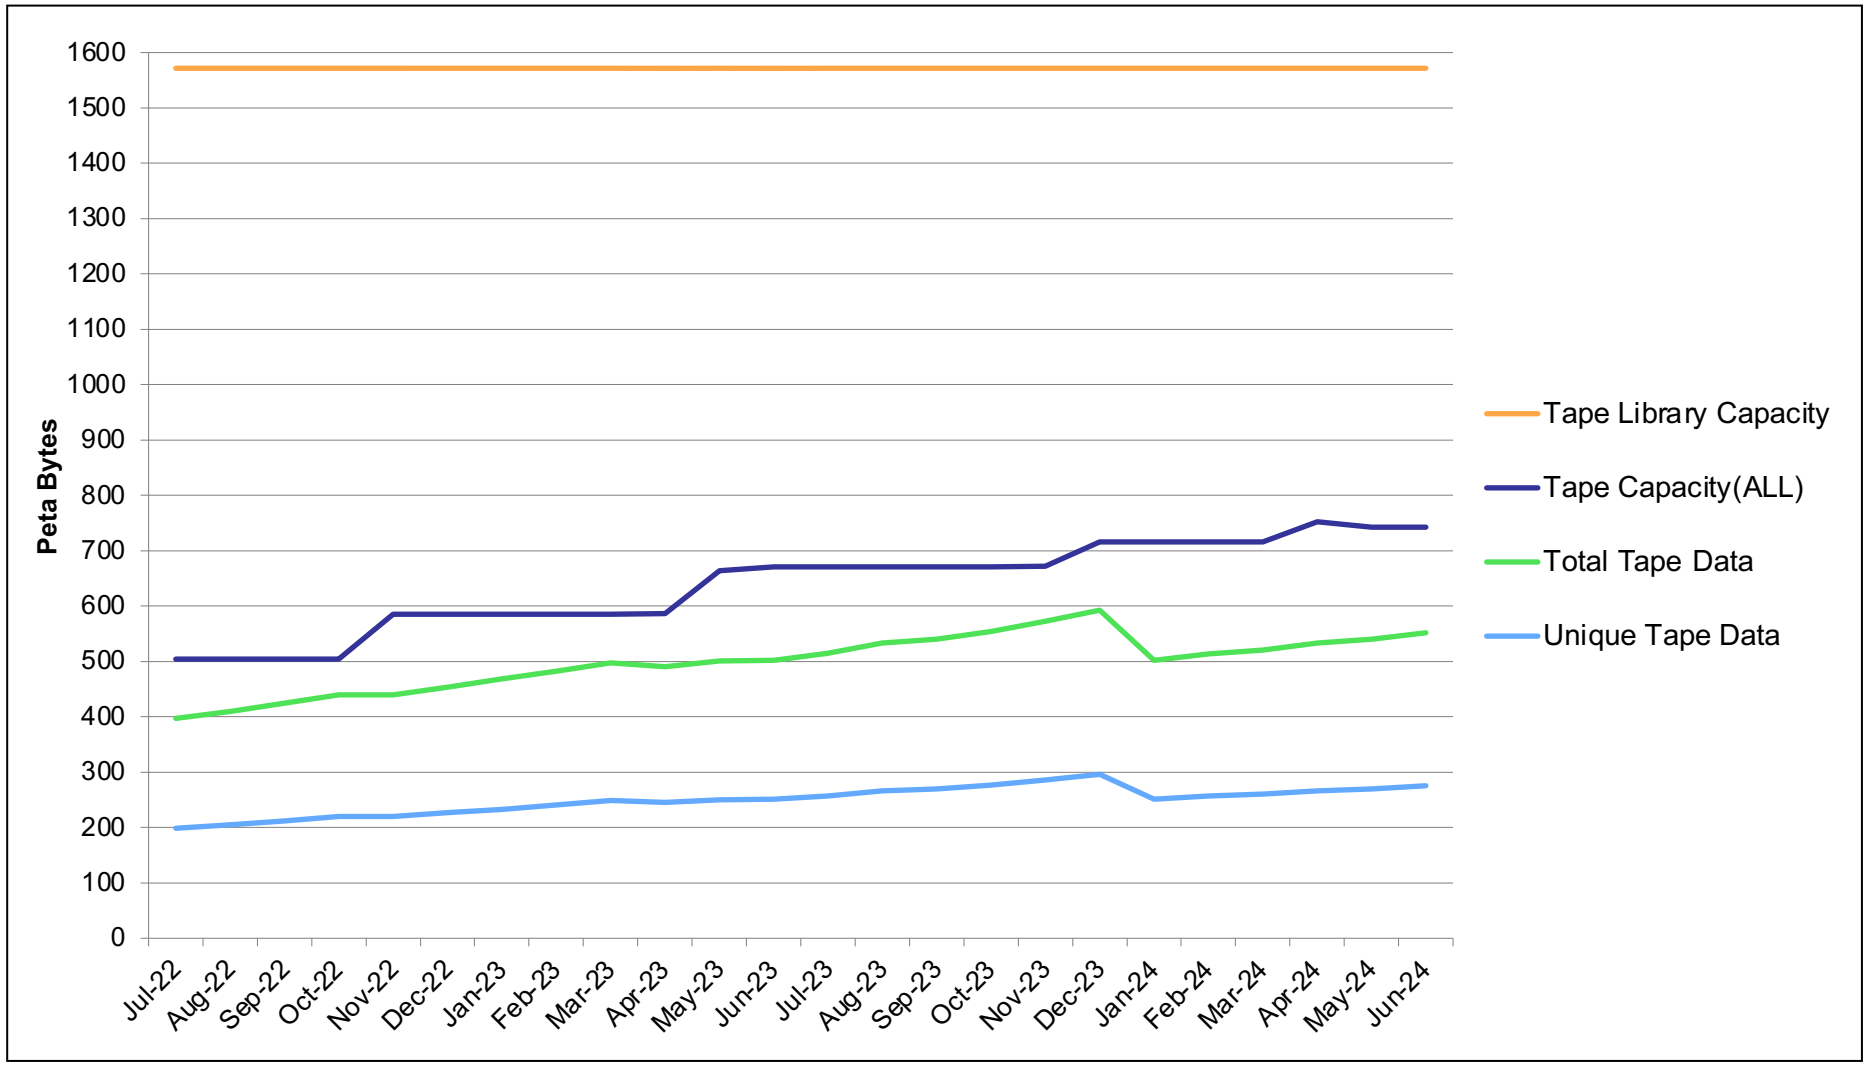
④ Cabeus Milan a100 nodes added

⑤ Aitken Milan Nodes added

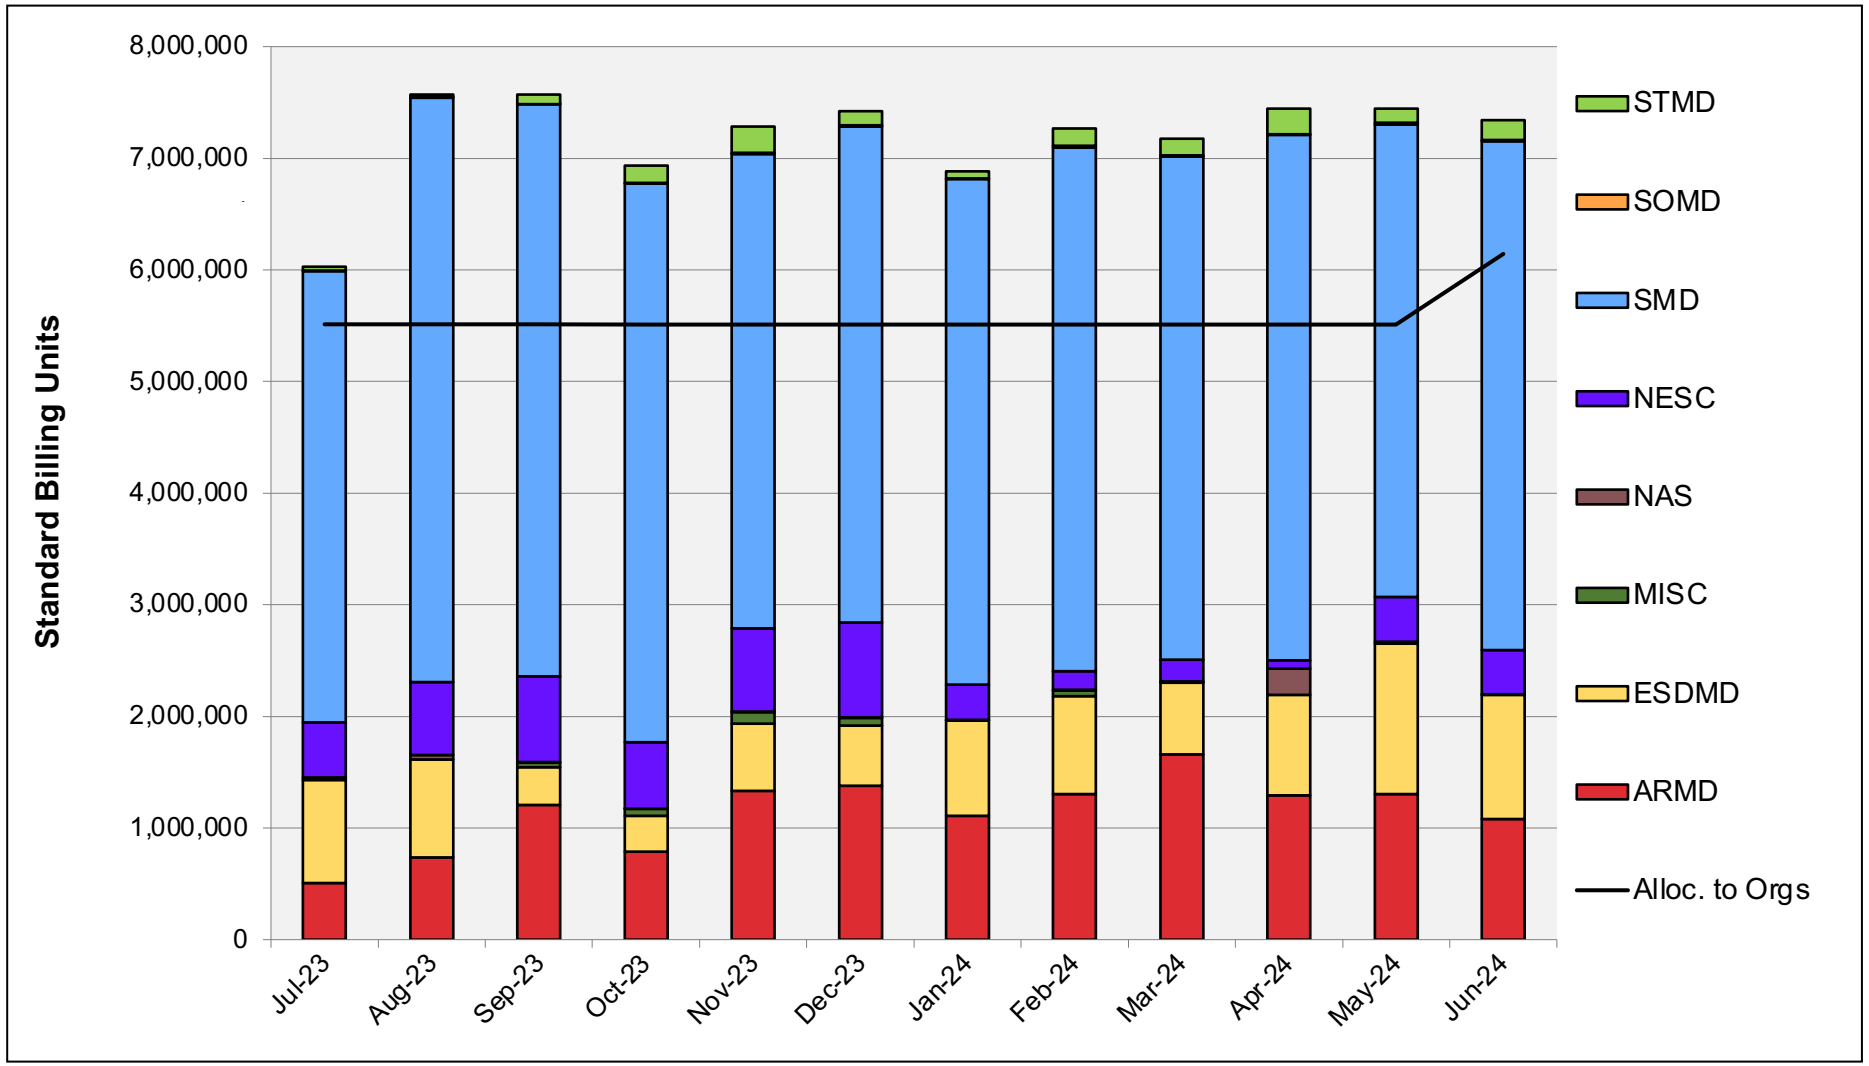
Tape Archive Status



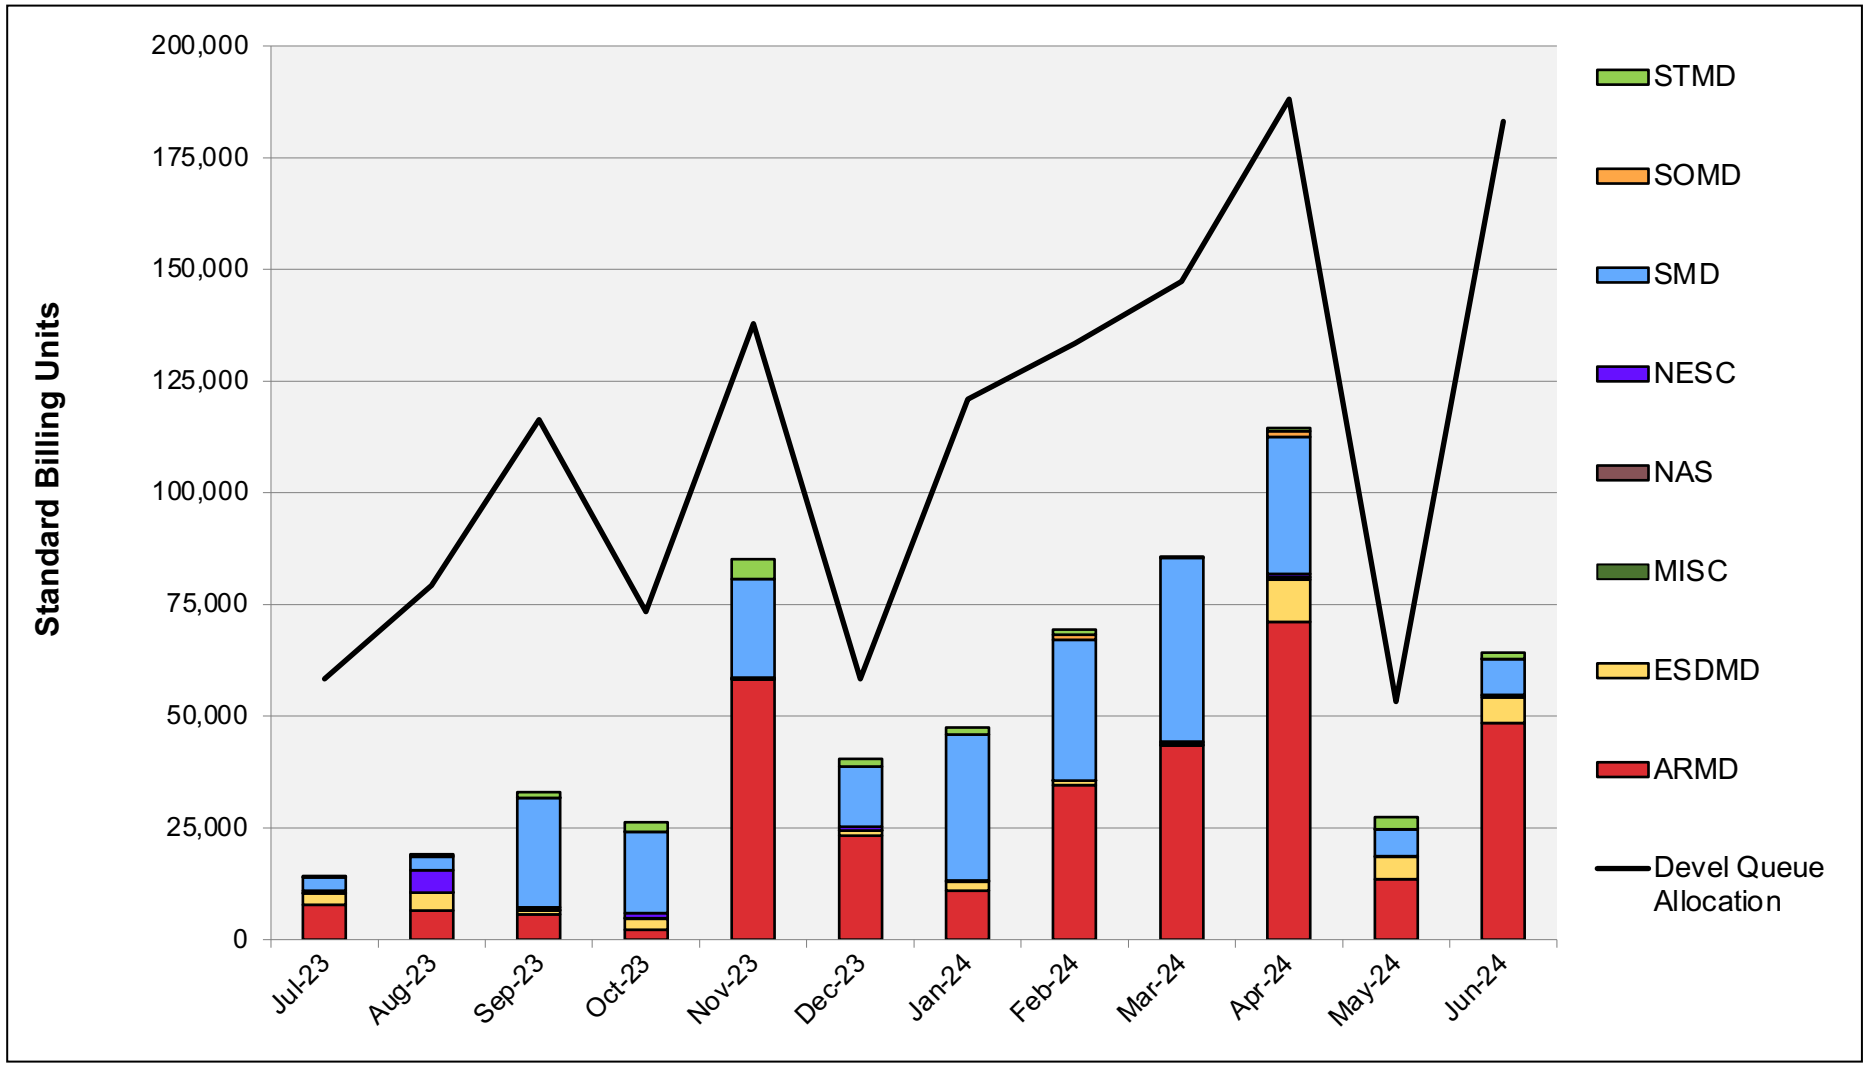
Tape Archive Status



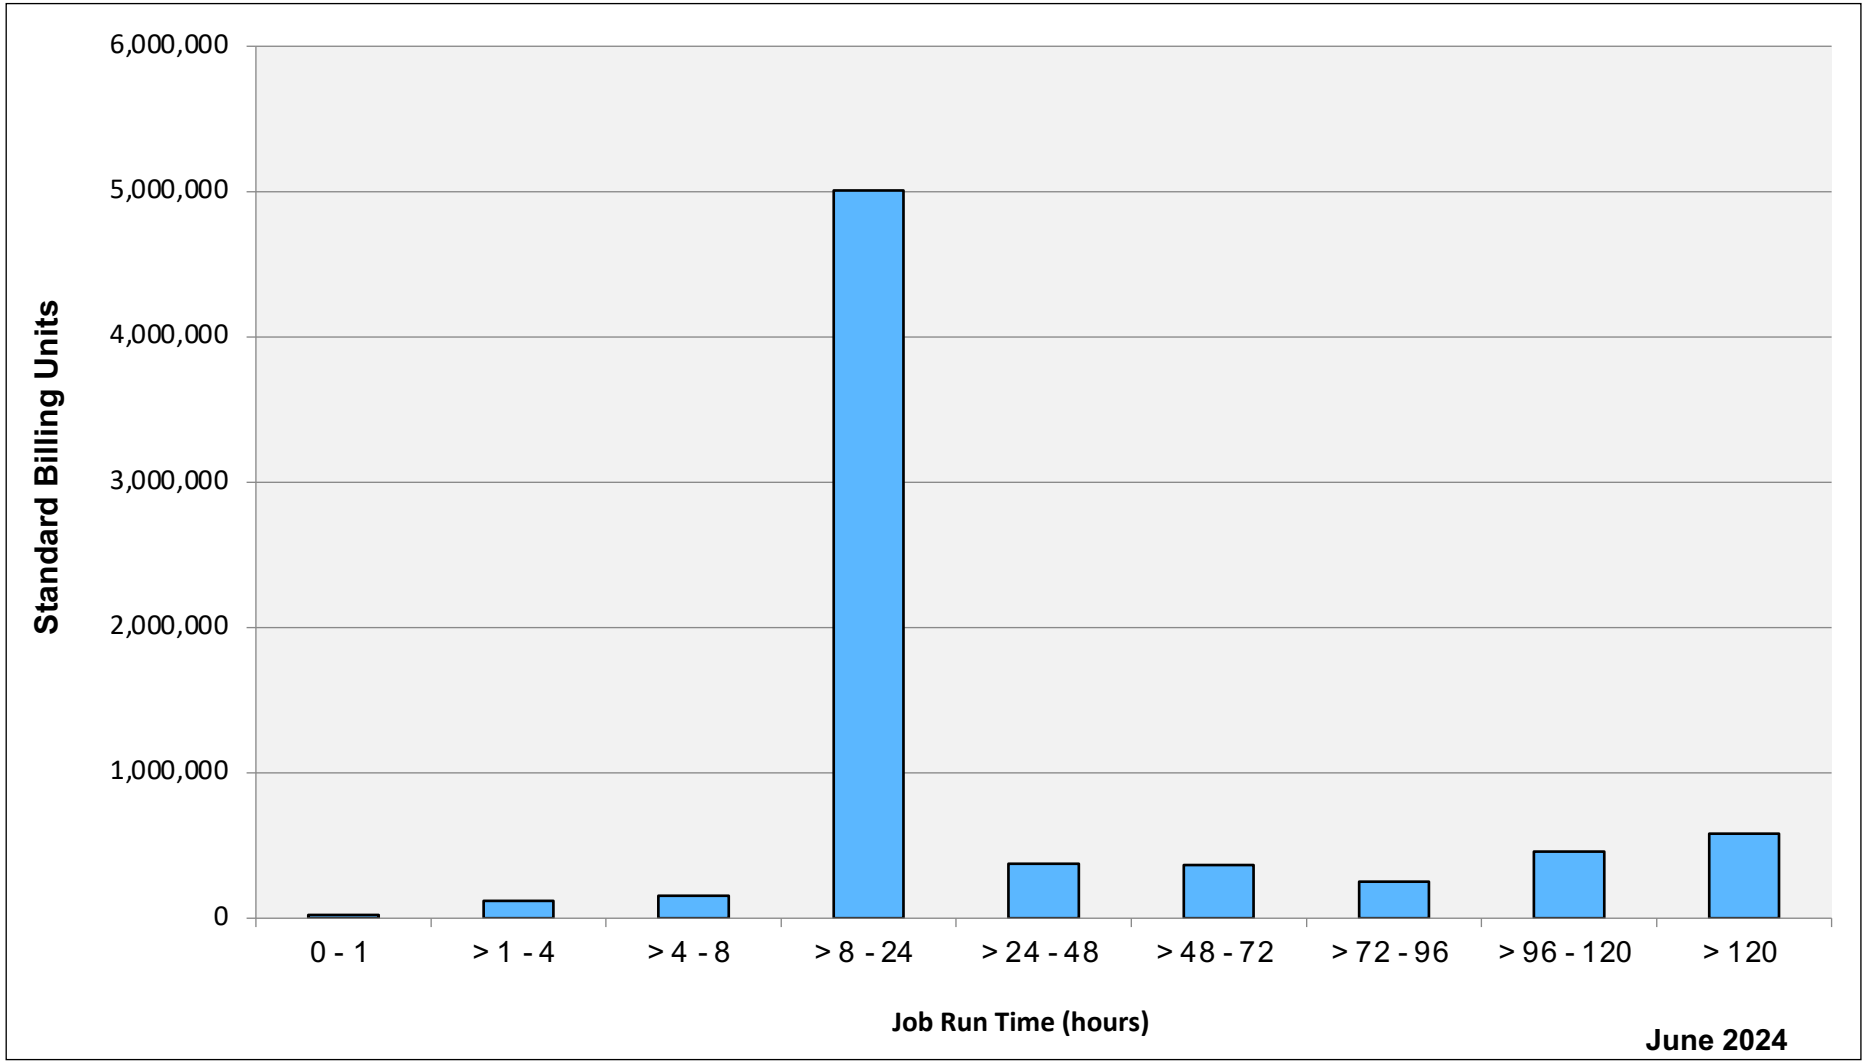
Aitken: SBUs Reported, Normalized to 30-Day Month



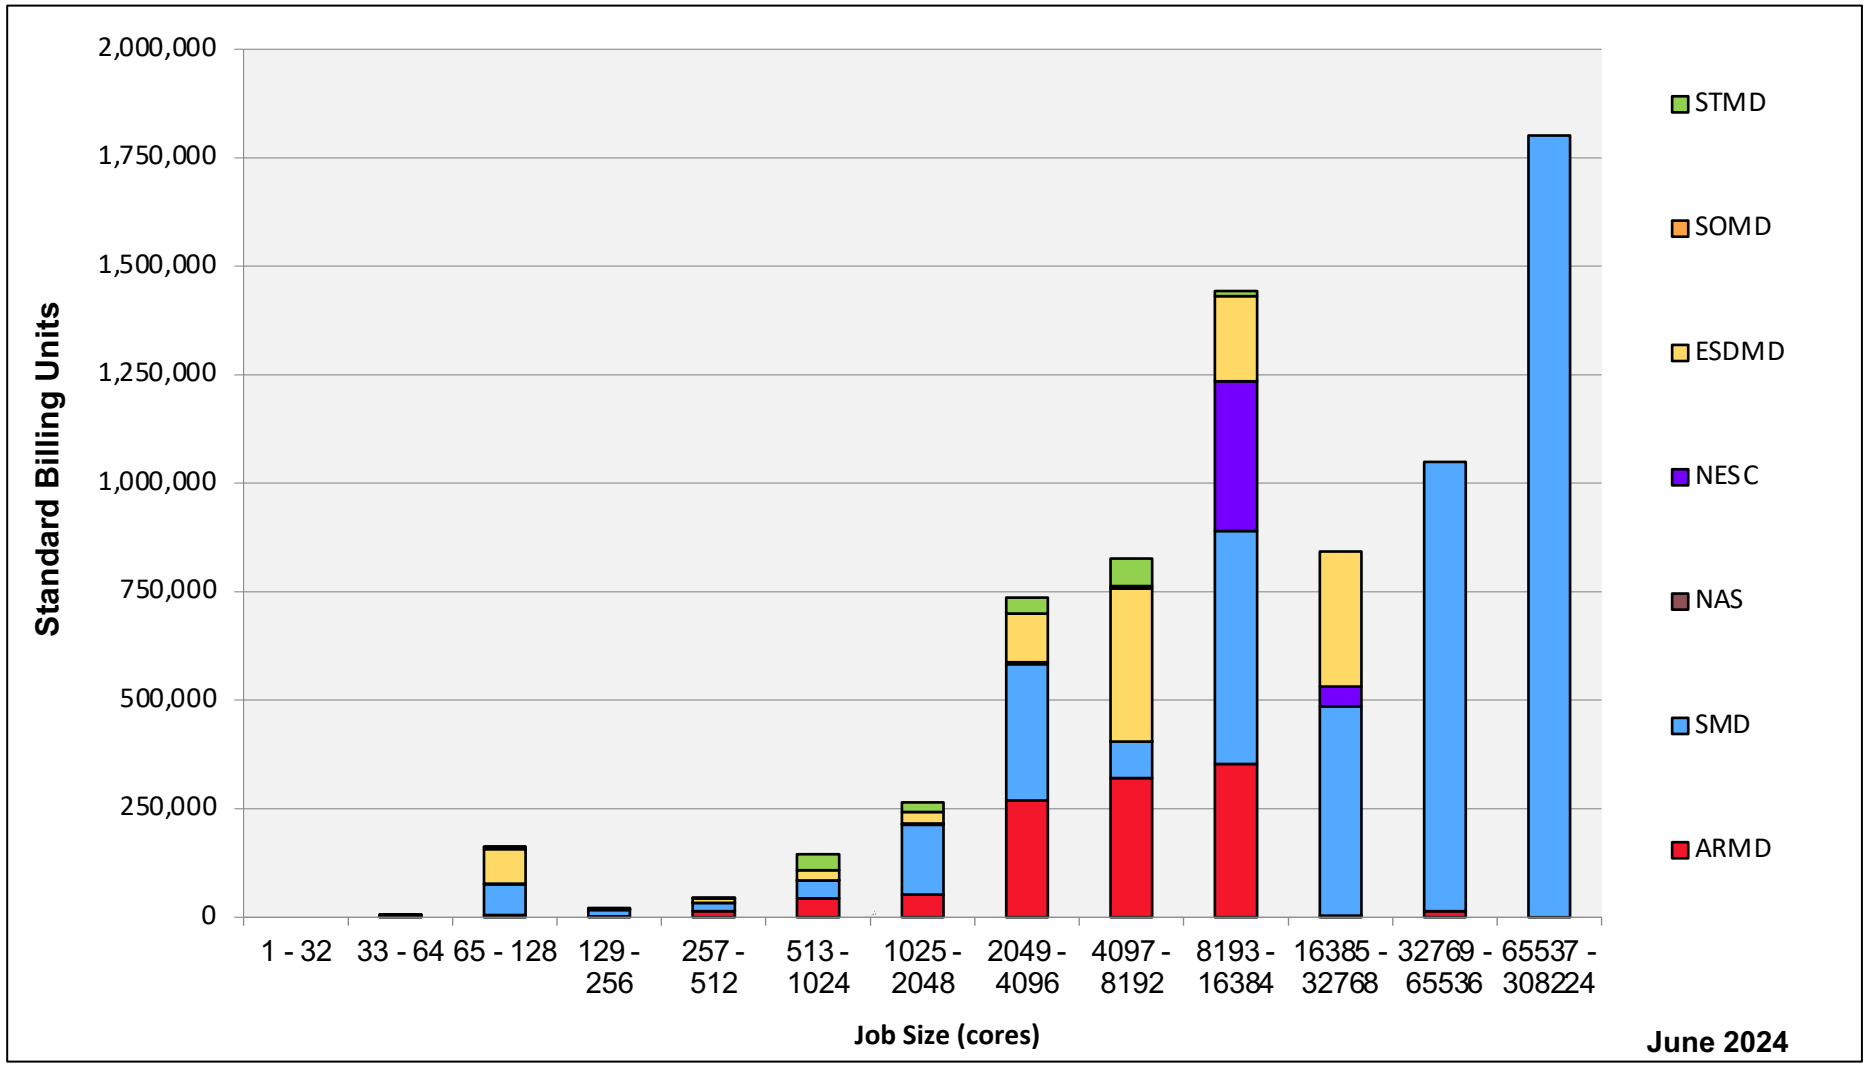
Aitken: Devel Queue Utilization



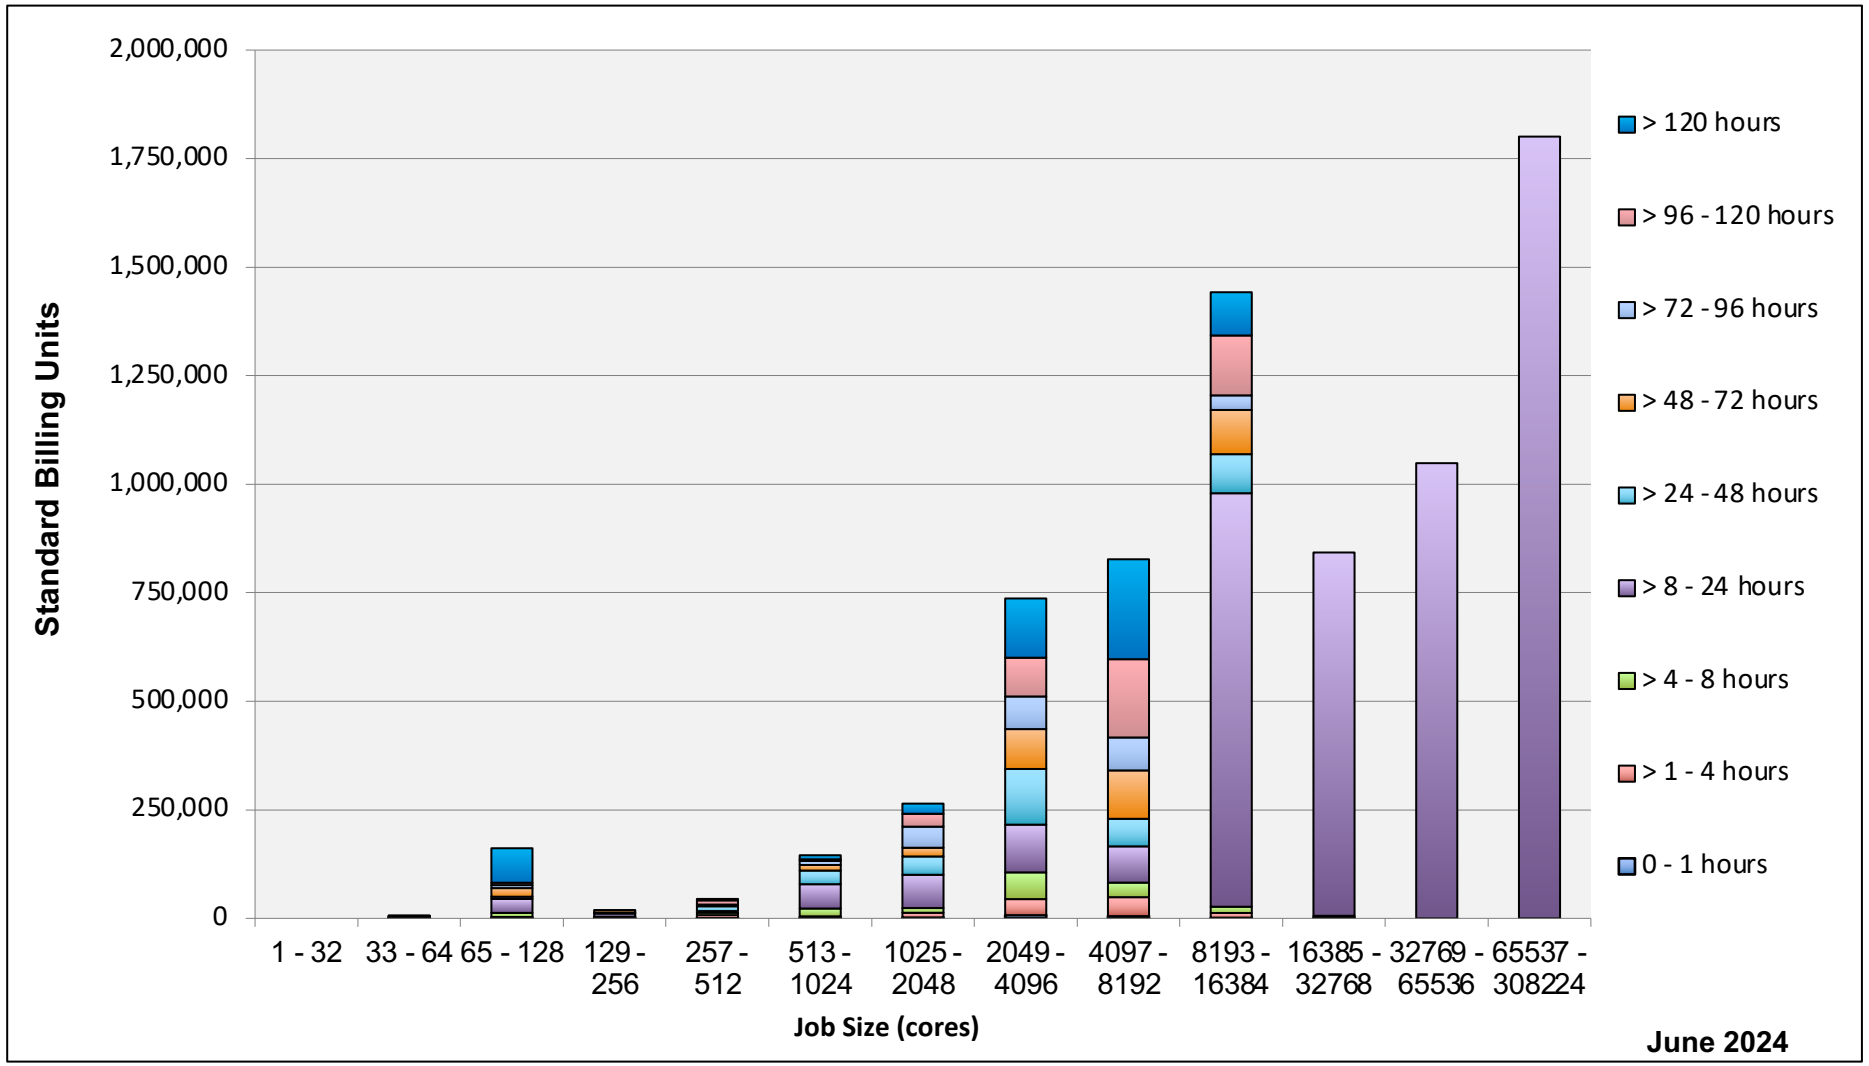
Aitken: Monthly Utilization by Job Length



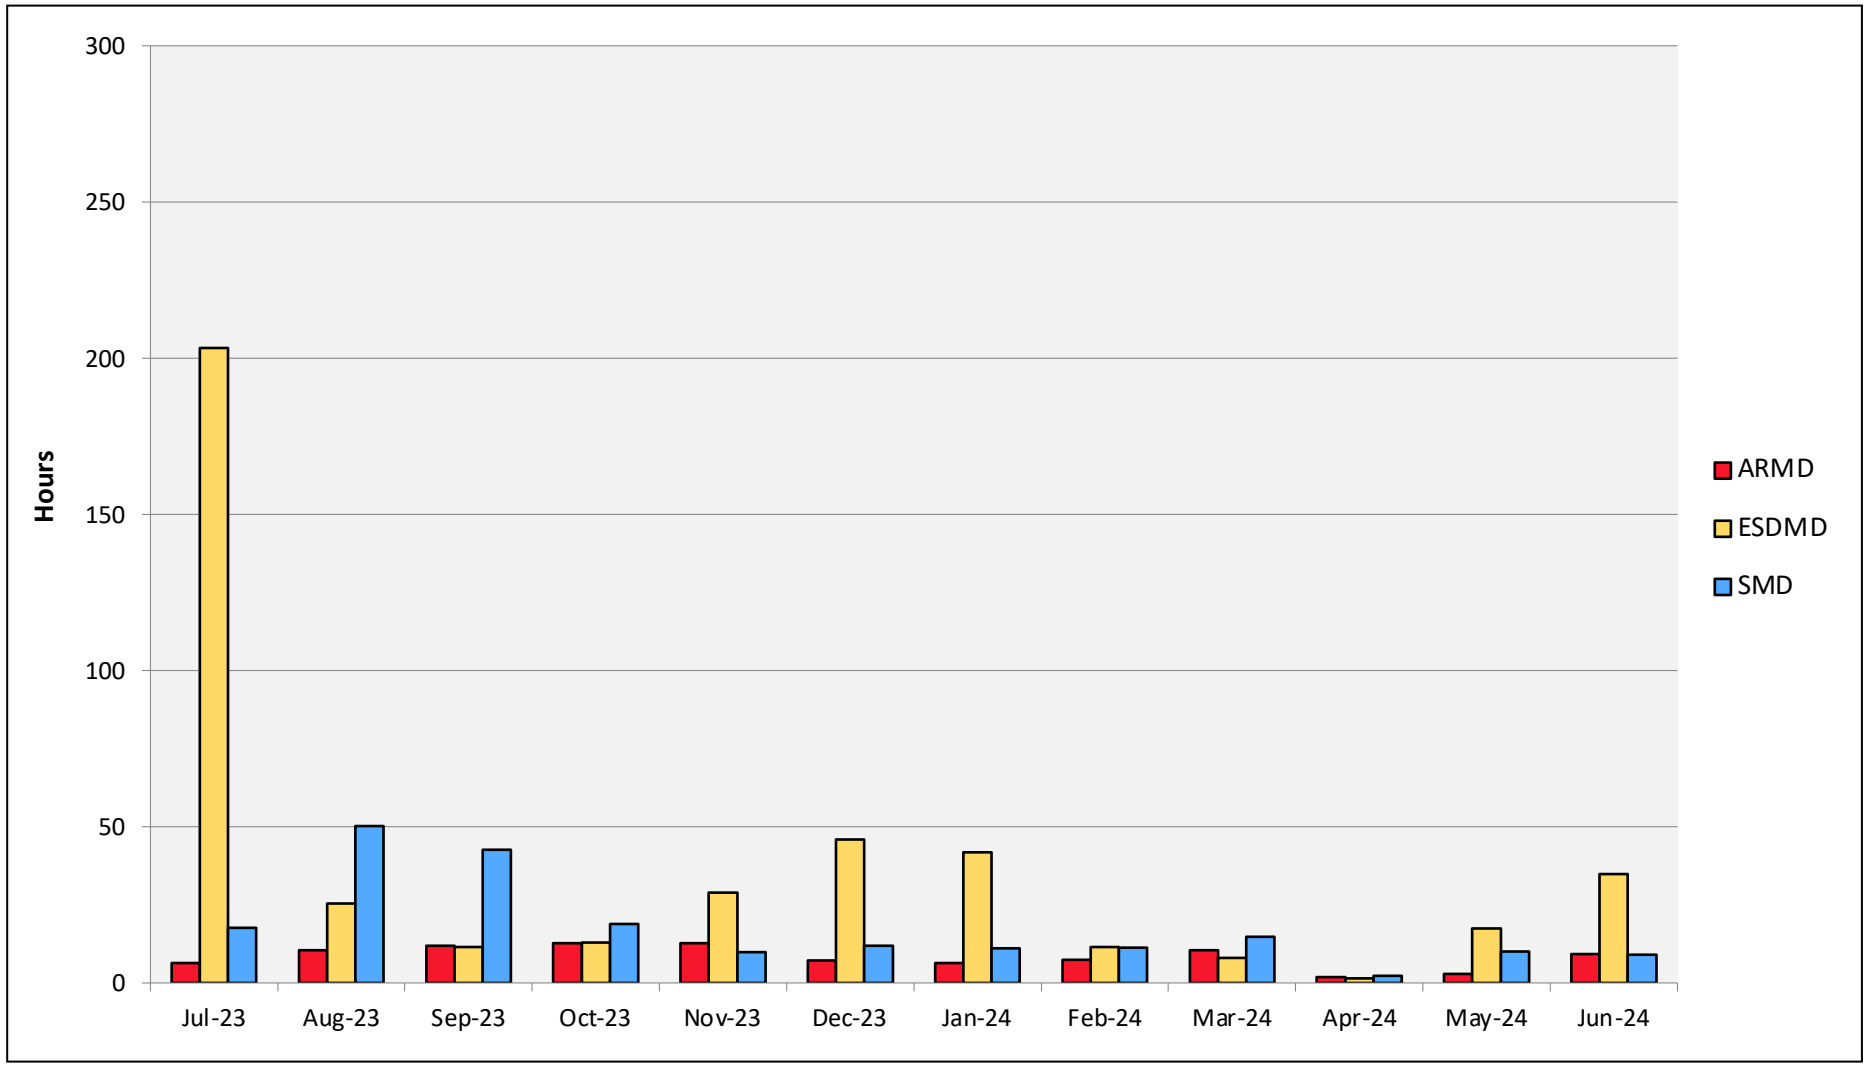
Aitken: Monthly Utilization by Job Size



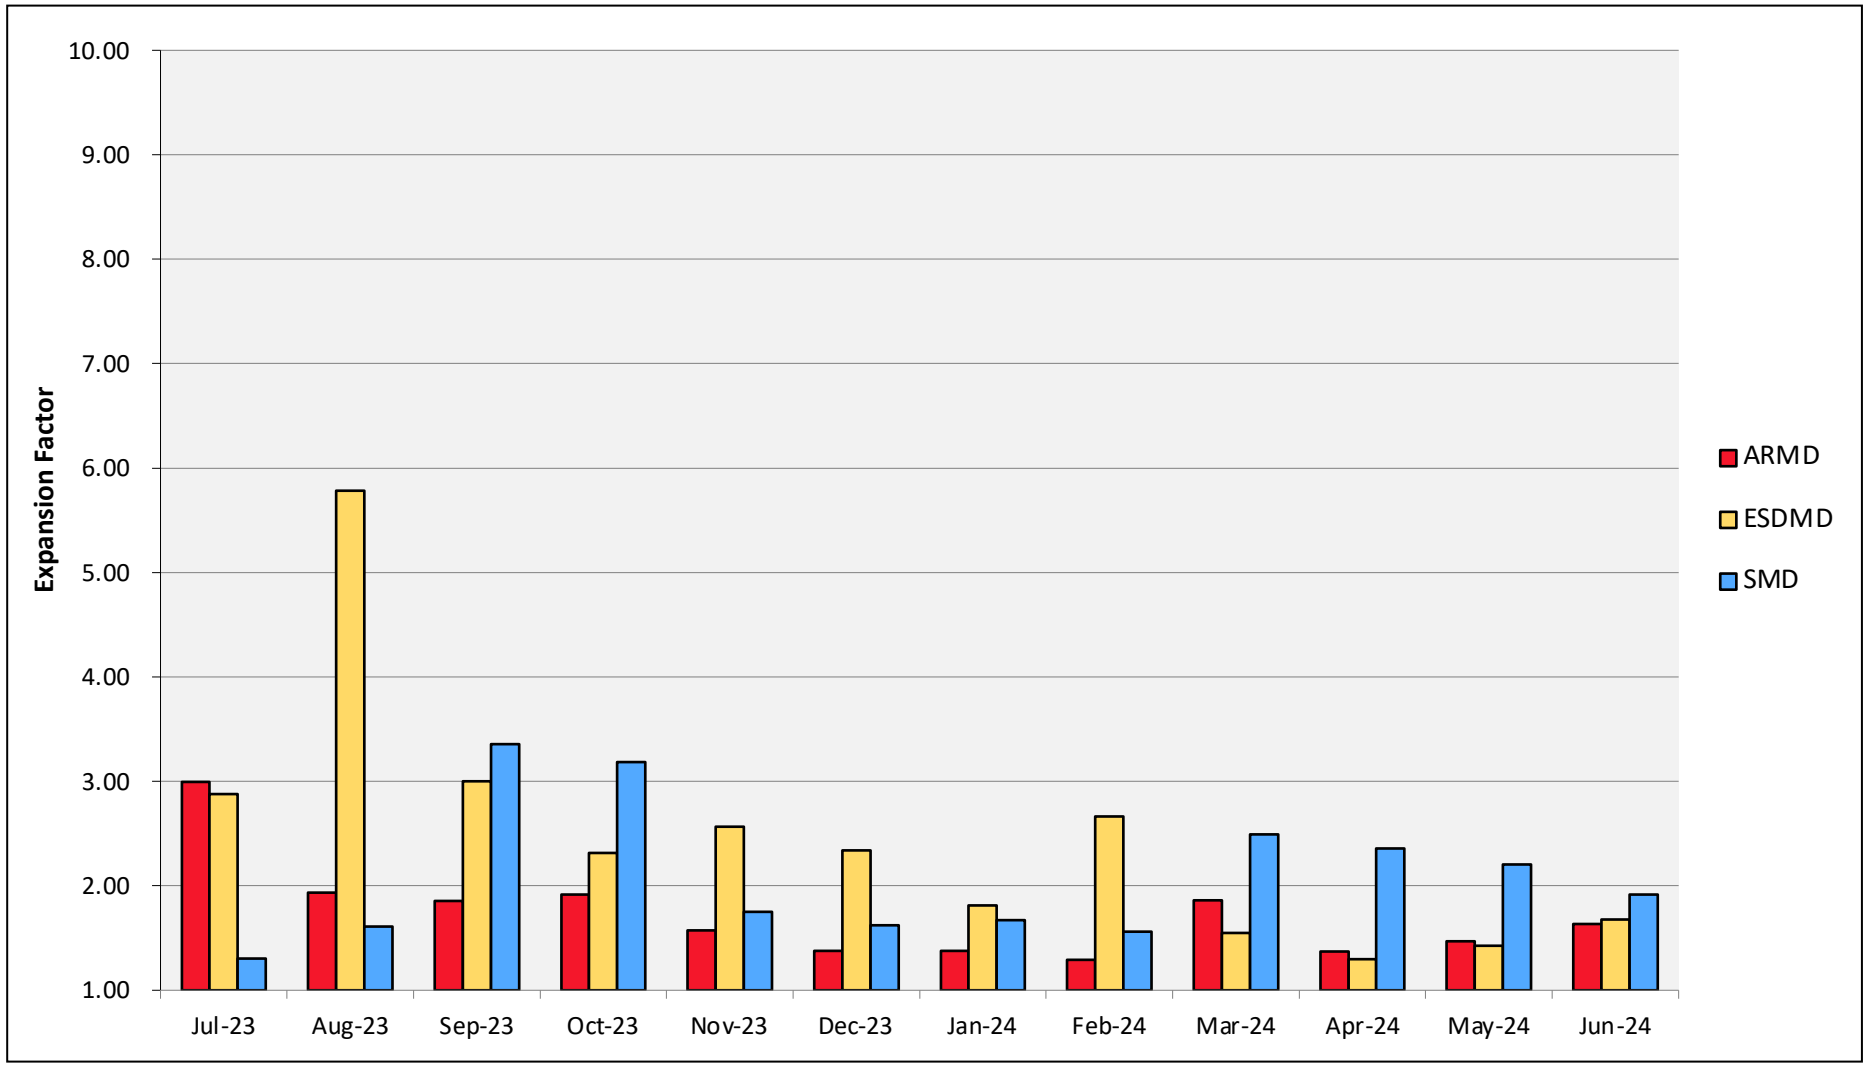
Aitken: Monthly Utilization by Size and Length



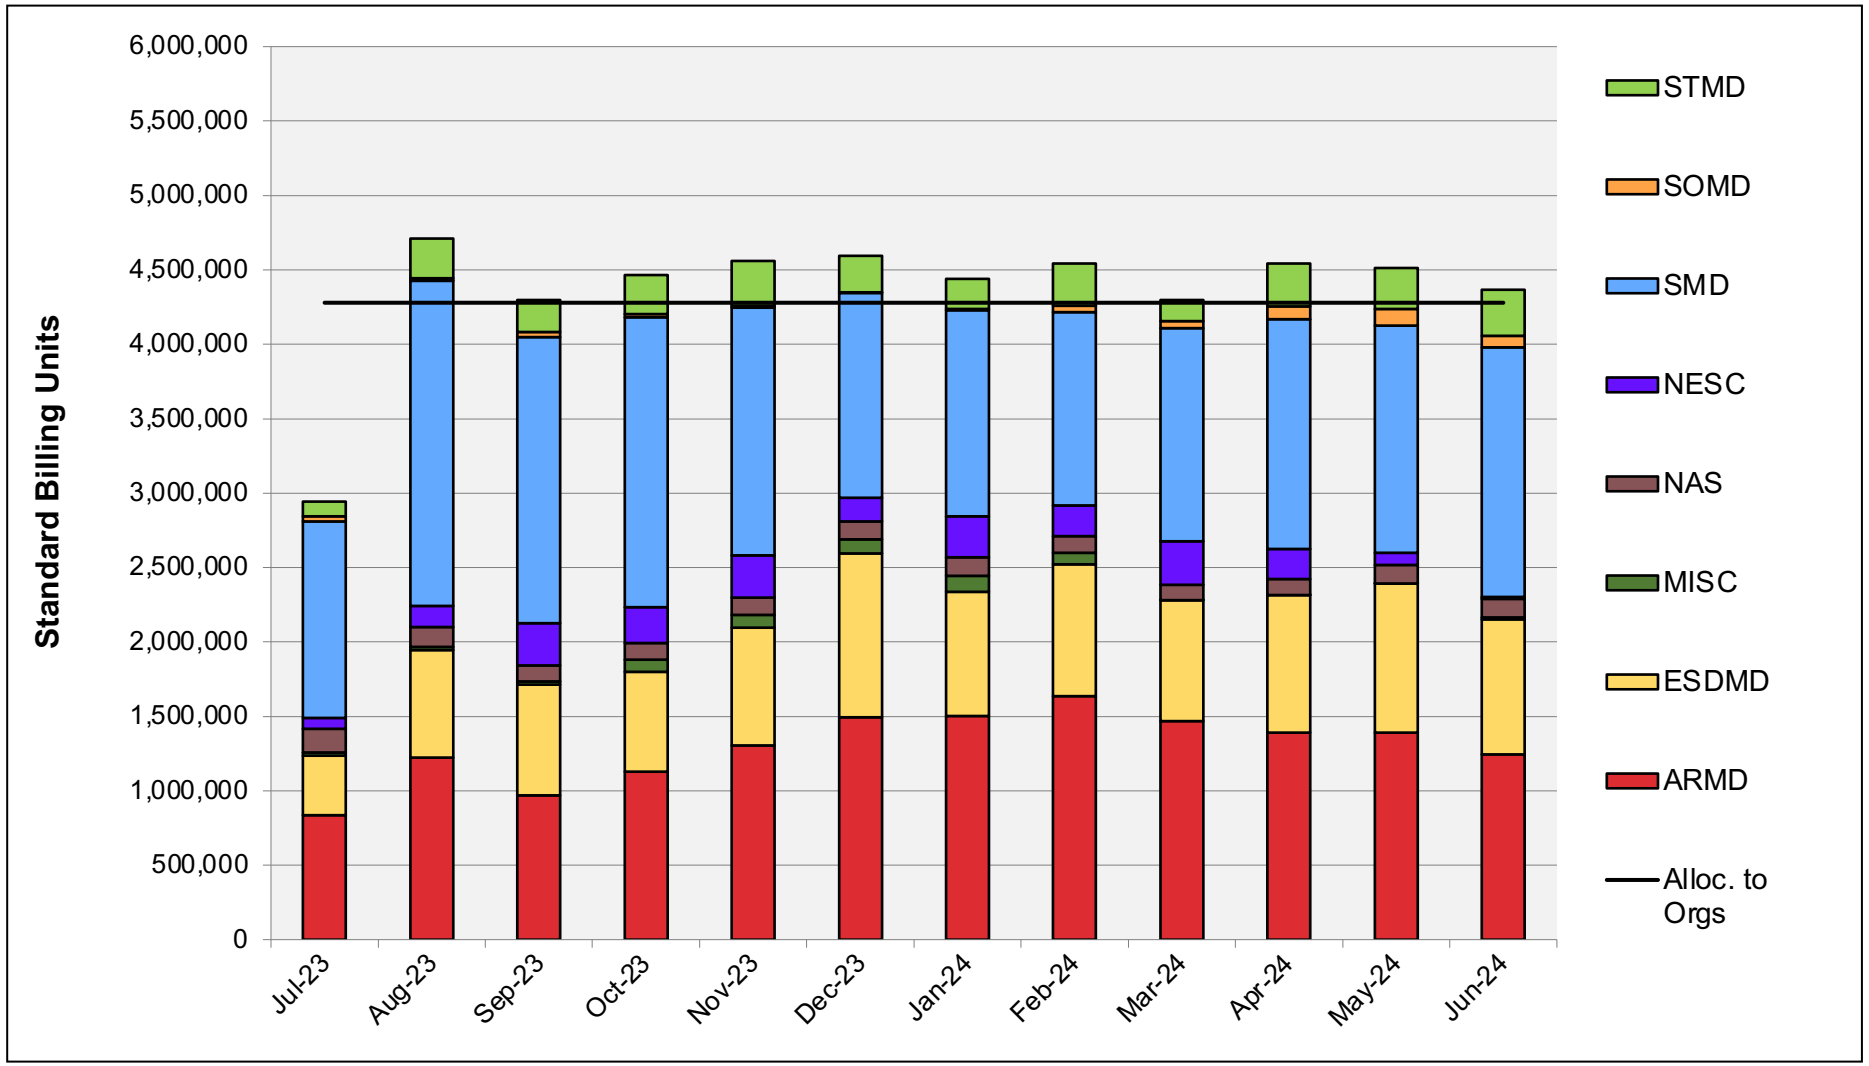
Aitken: Average Time to Clear All Jobs



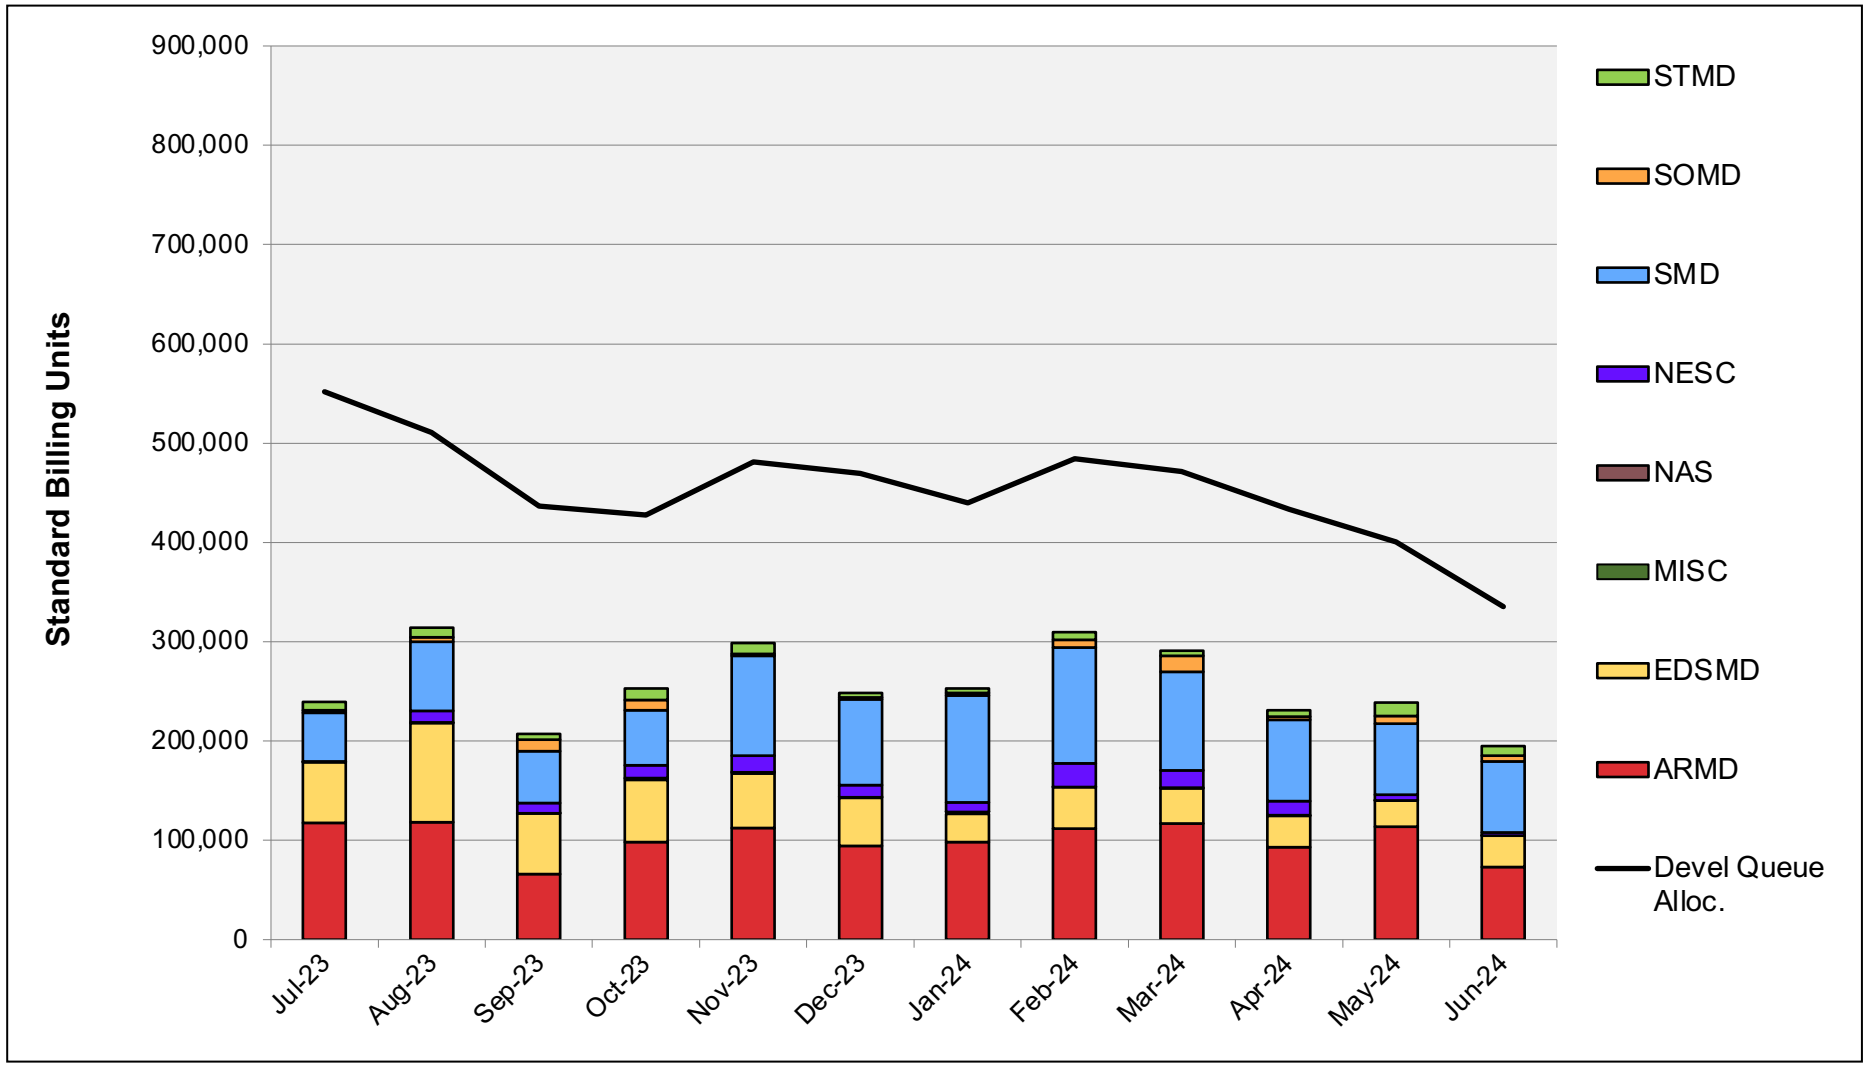
Aitken: Average Expansion Factor



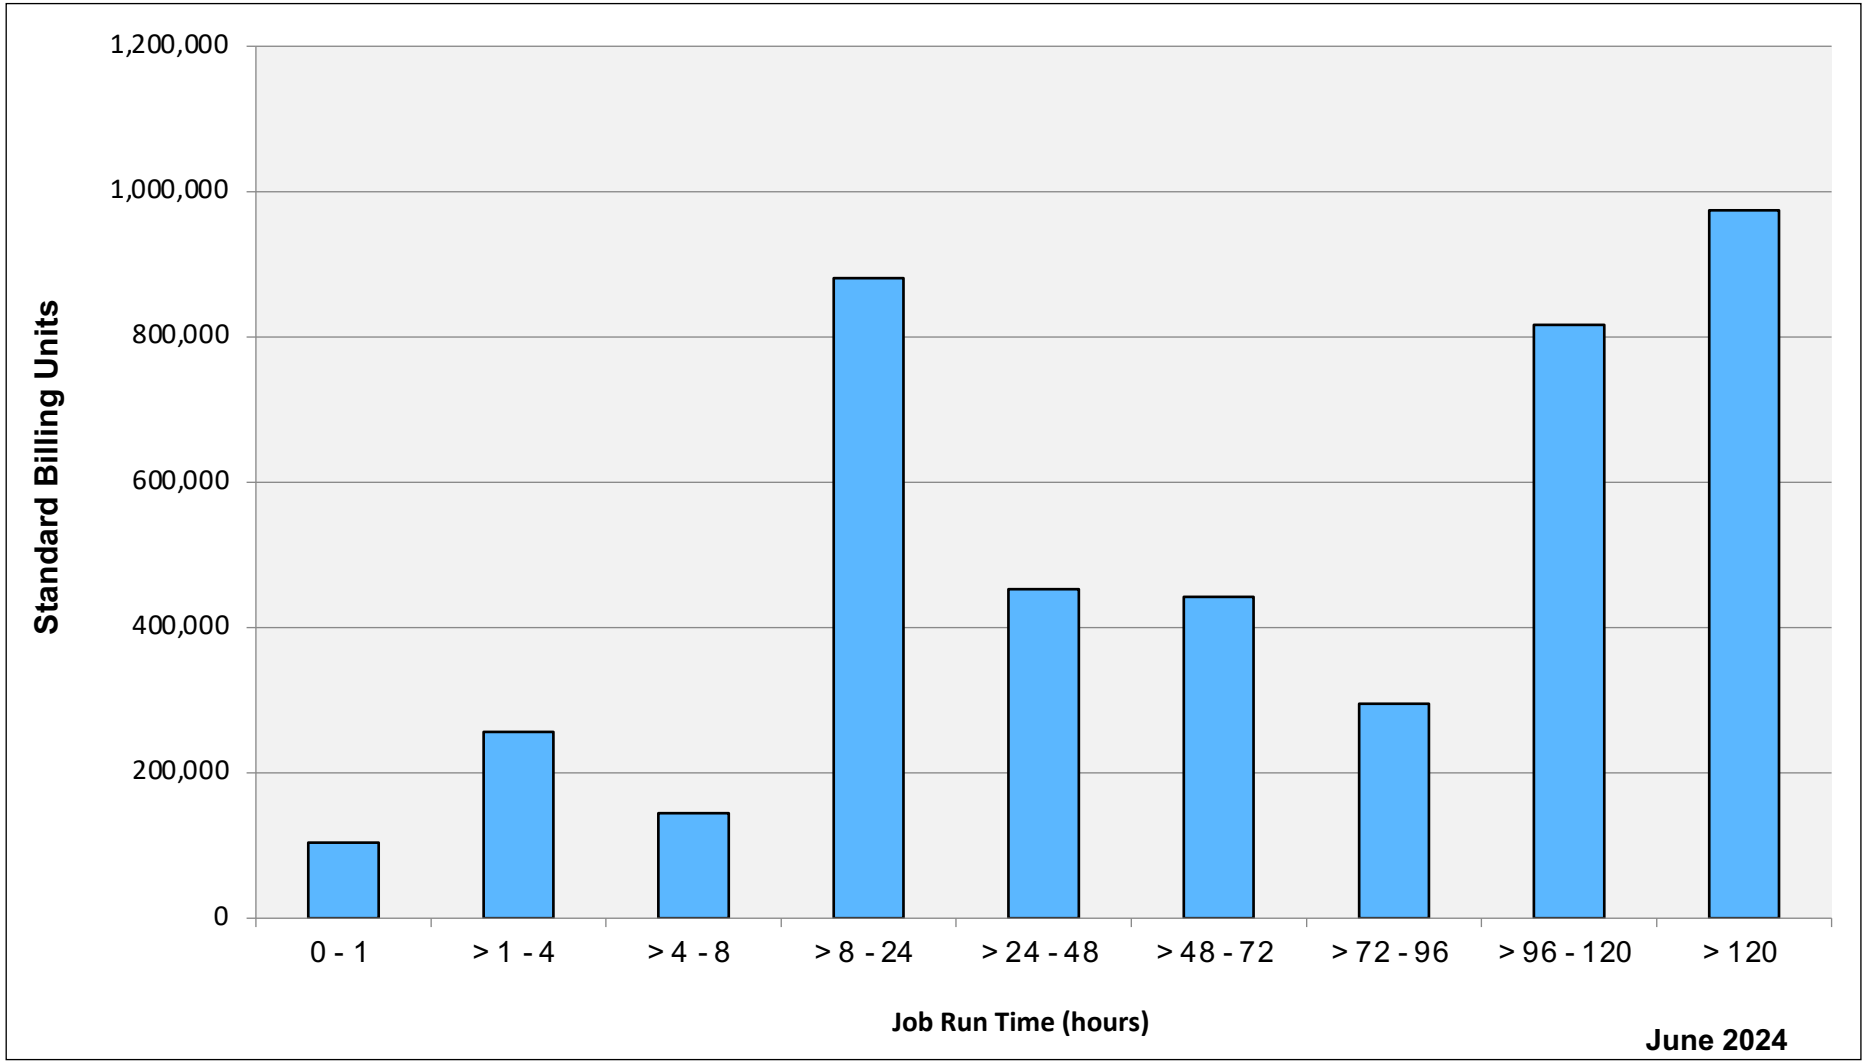
Pleiades: SBUs Reported, Normalized to 30-Day Month



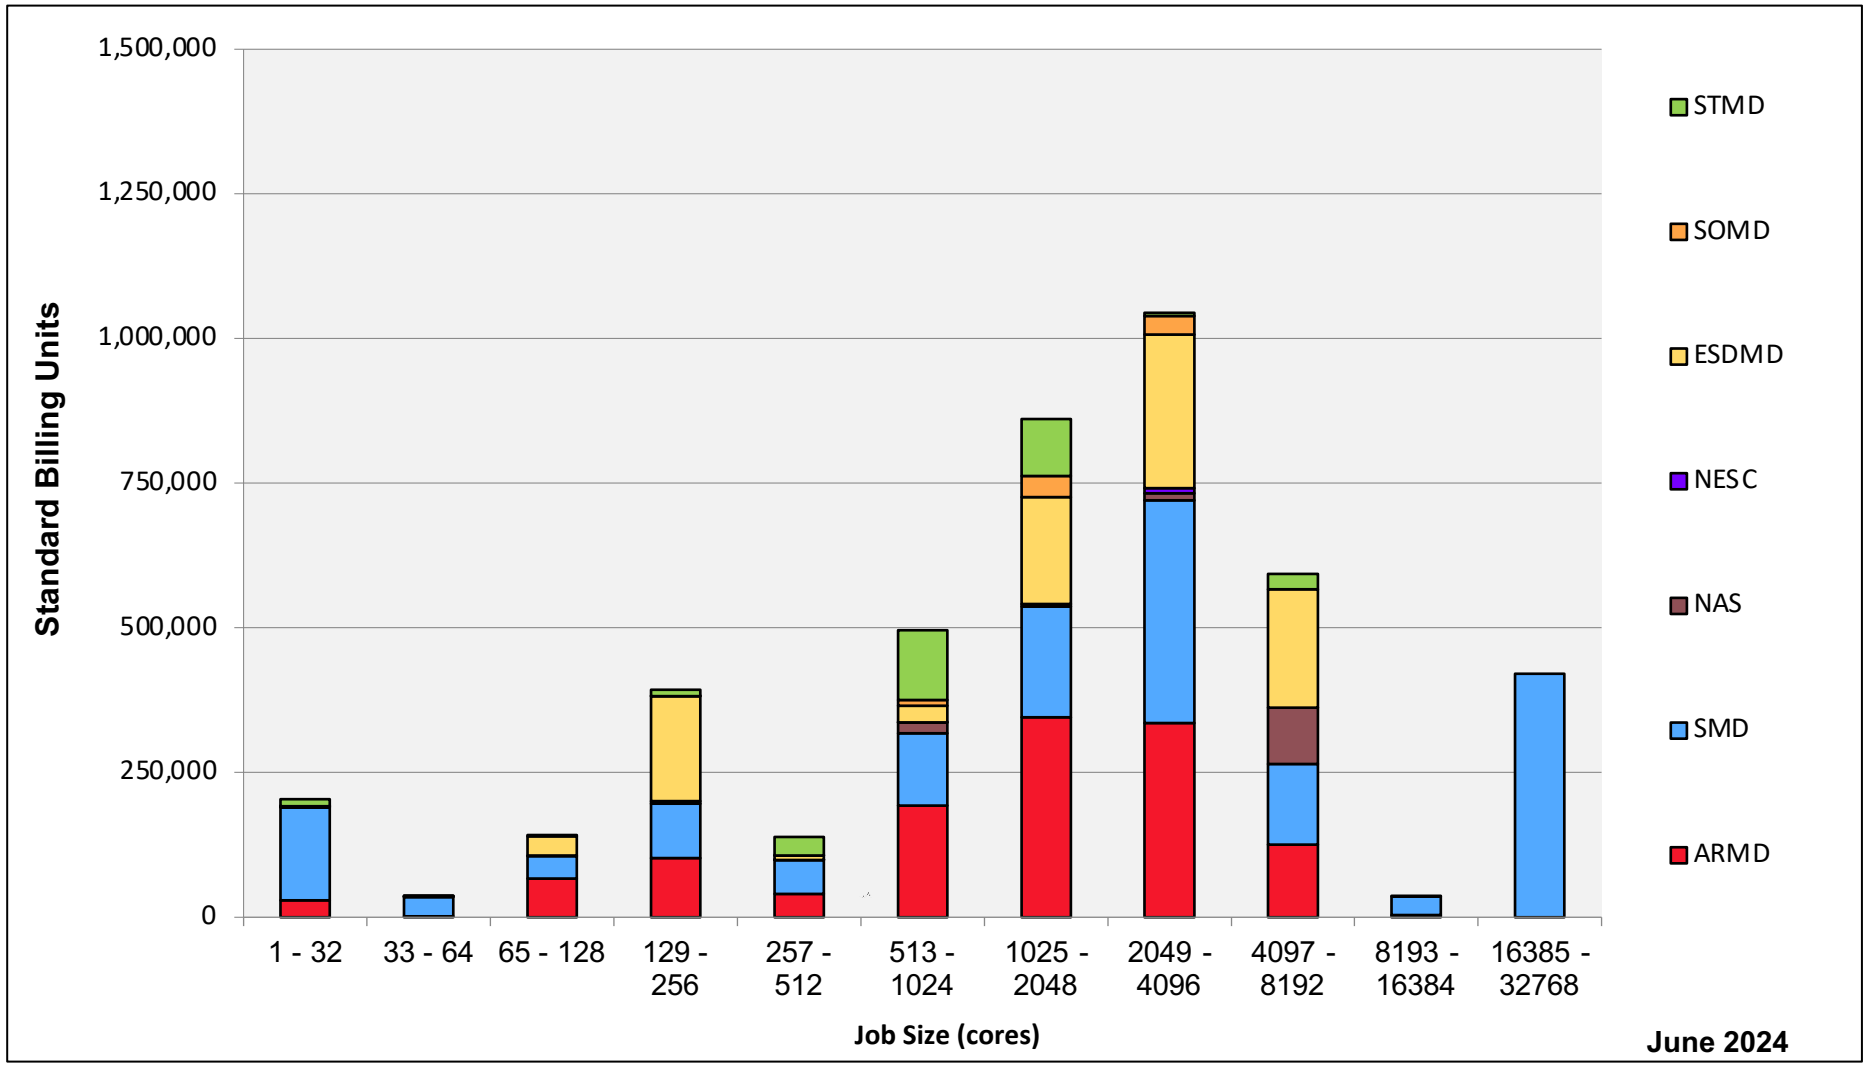
Pleiades: Devel Queue Utilization



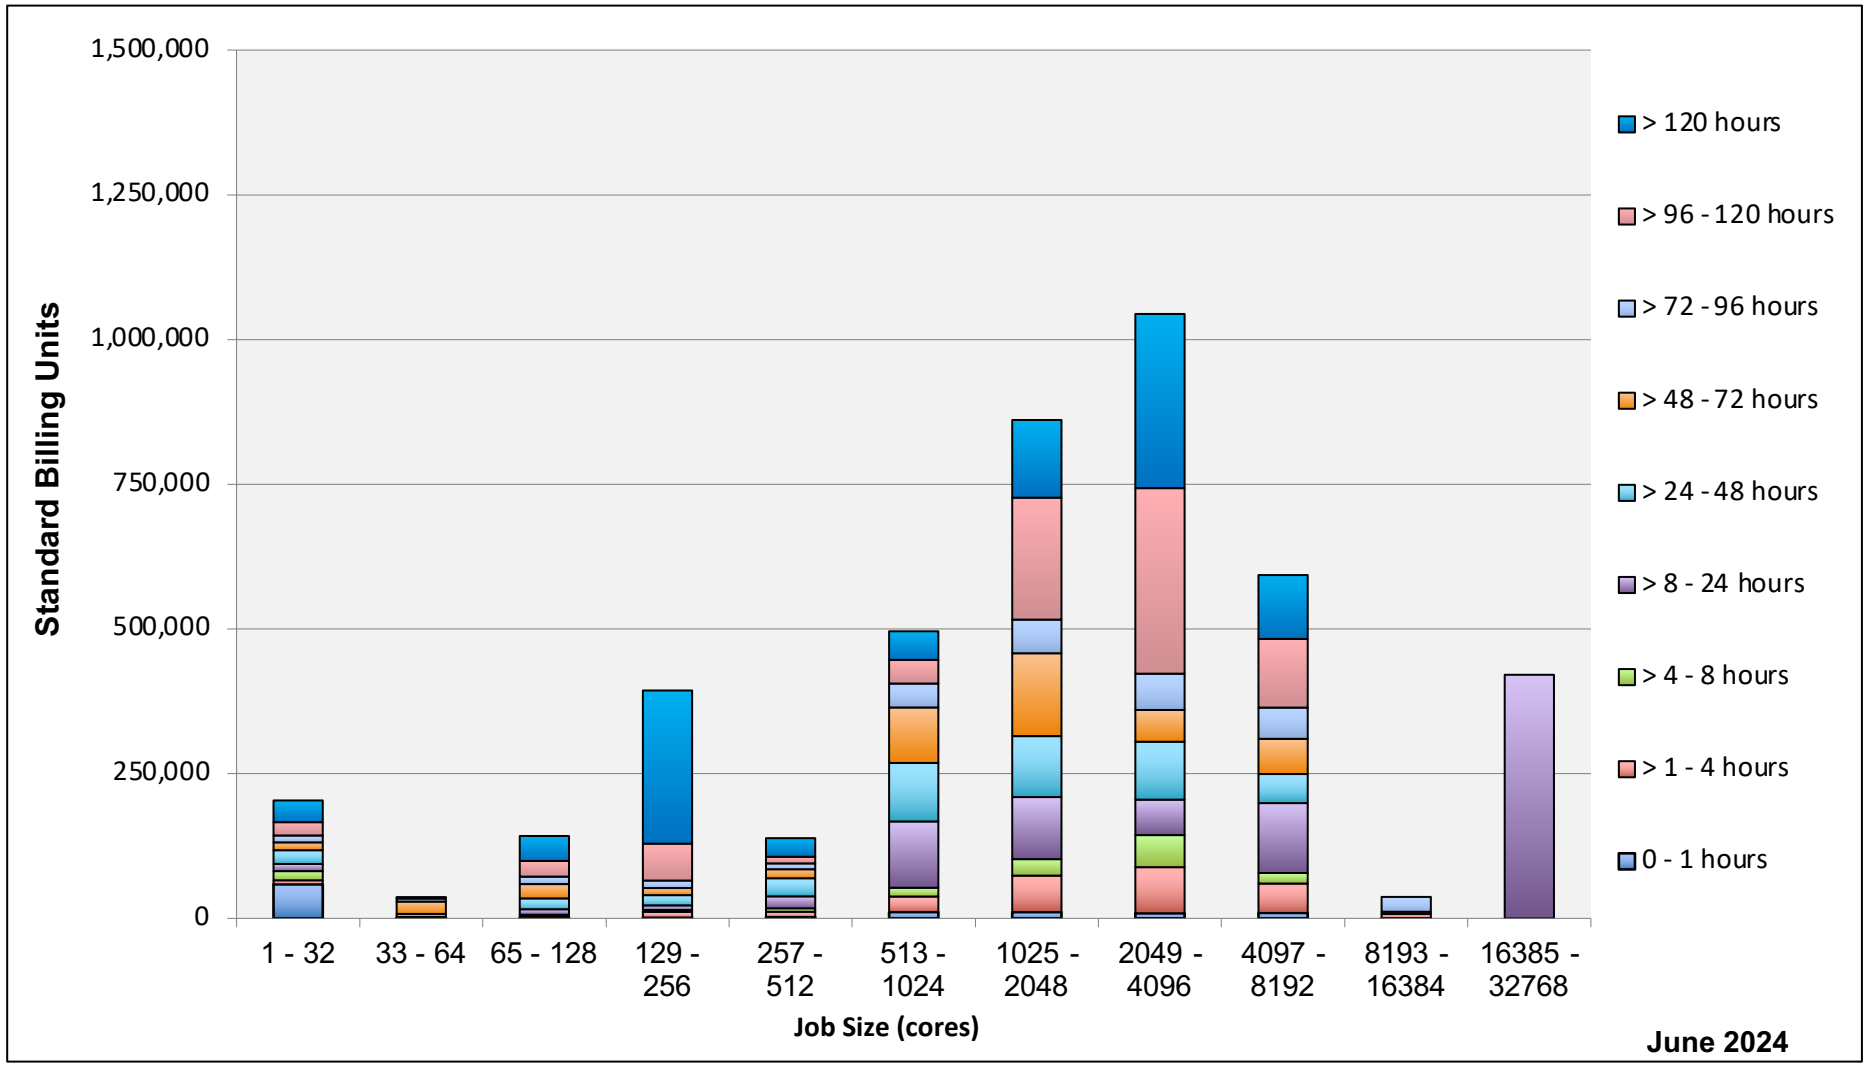
Pleiades: Monthly Utilization by Job Length



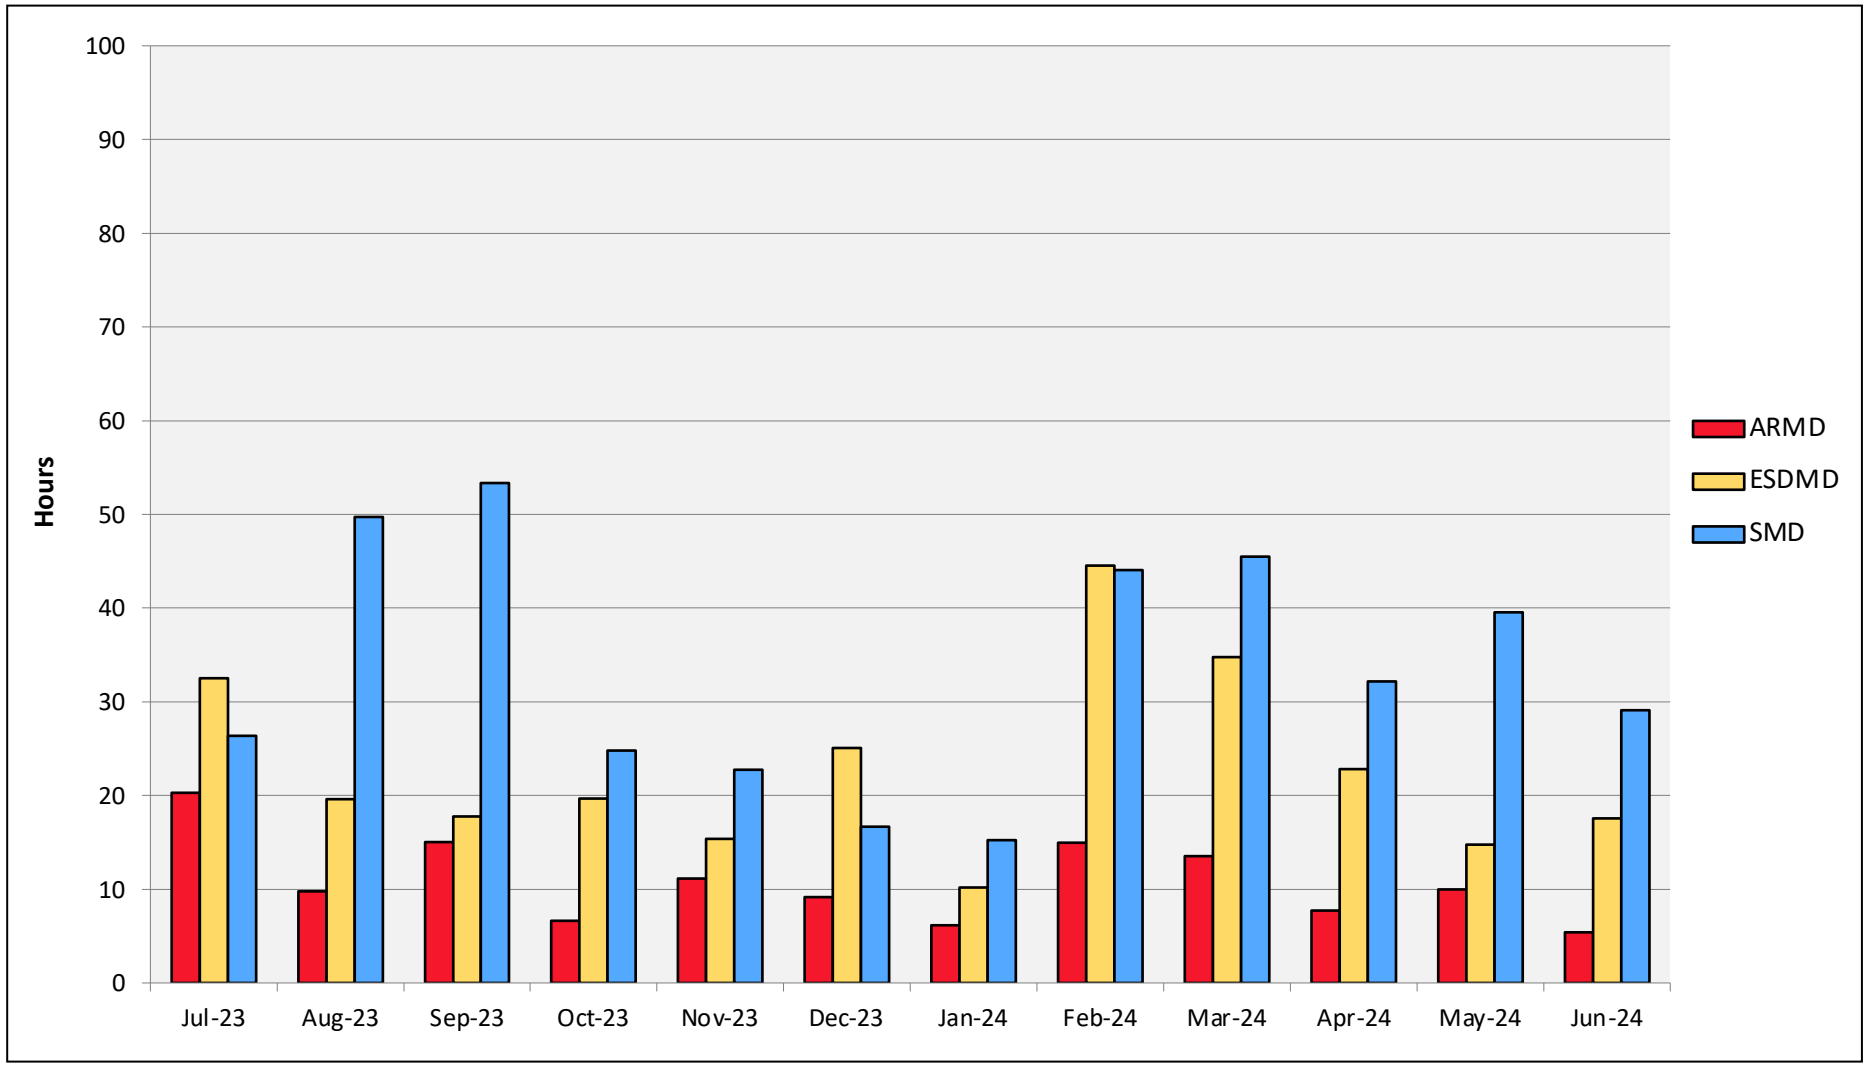
Pleiades: Monthly Utilization by Job Size



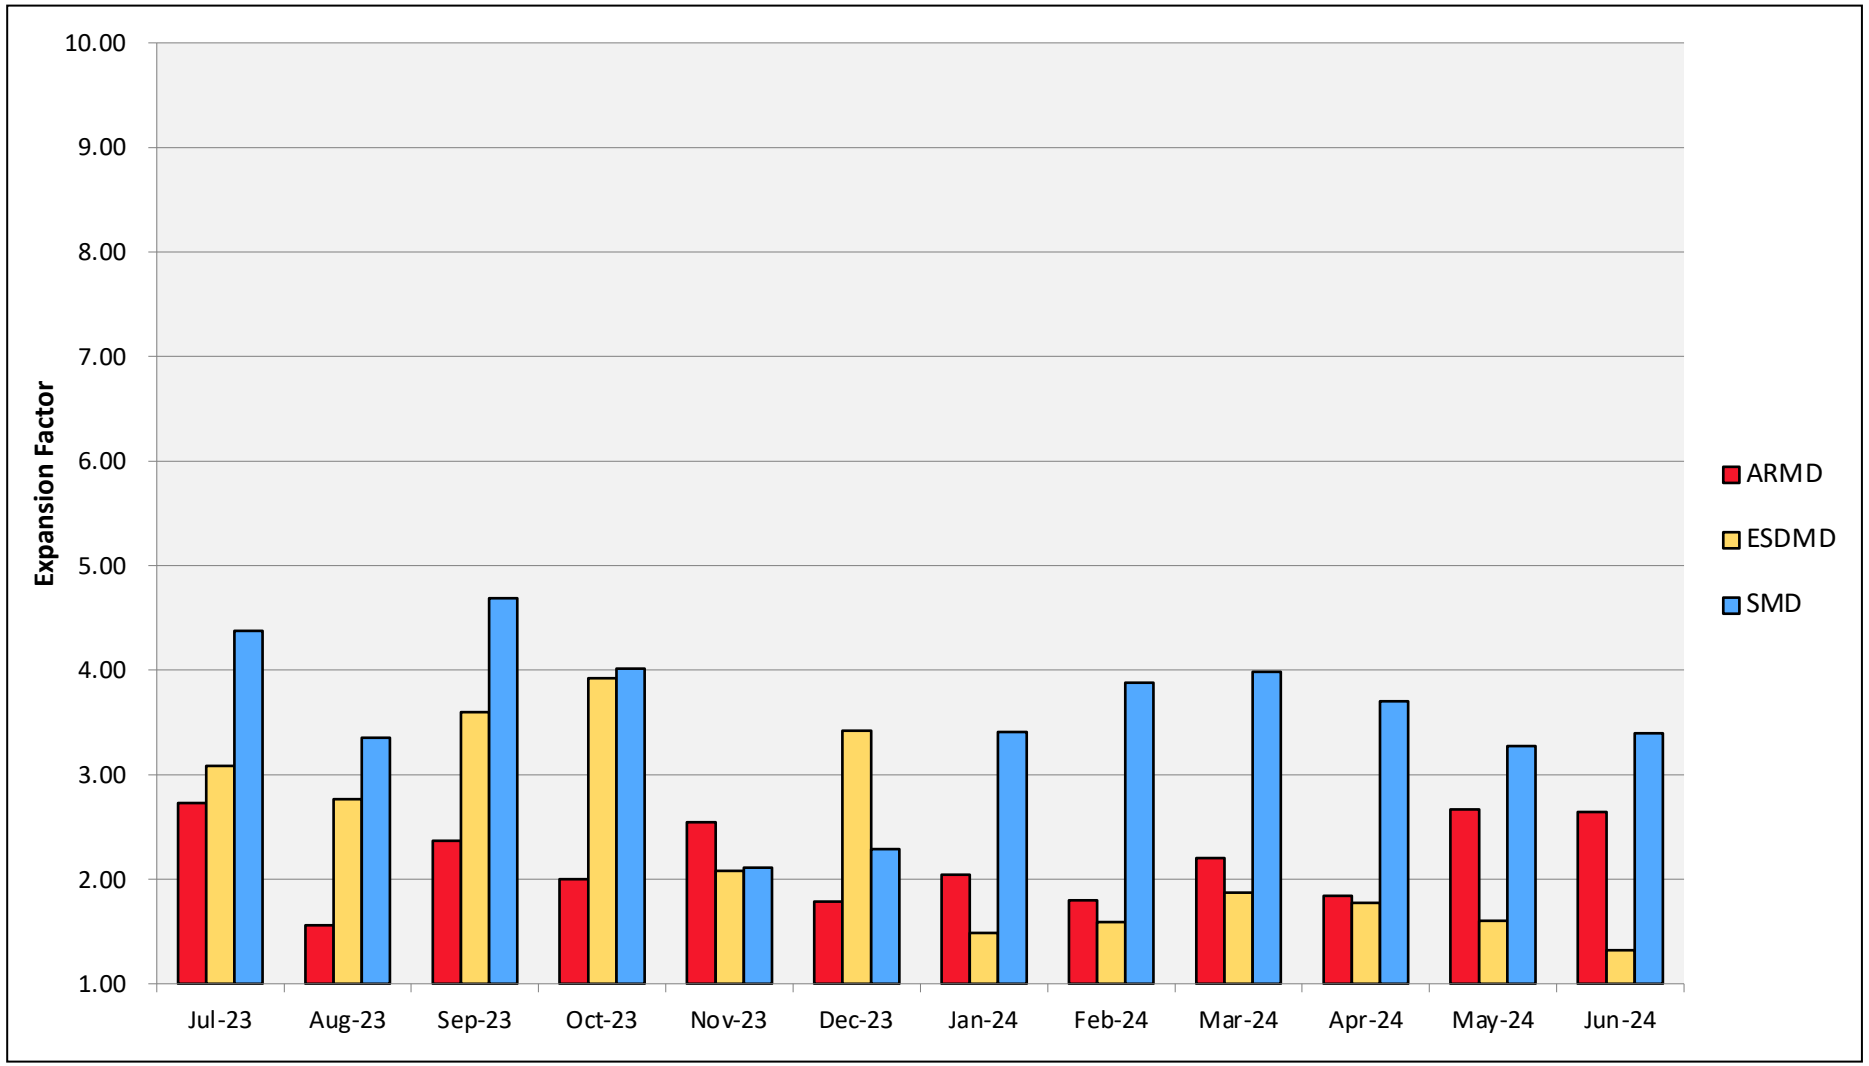
Pleiades: Monthly Utilization by Size and Length



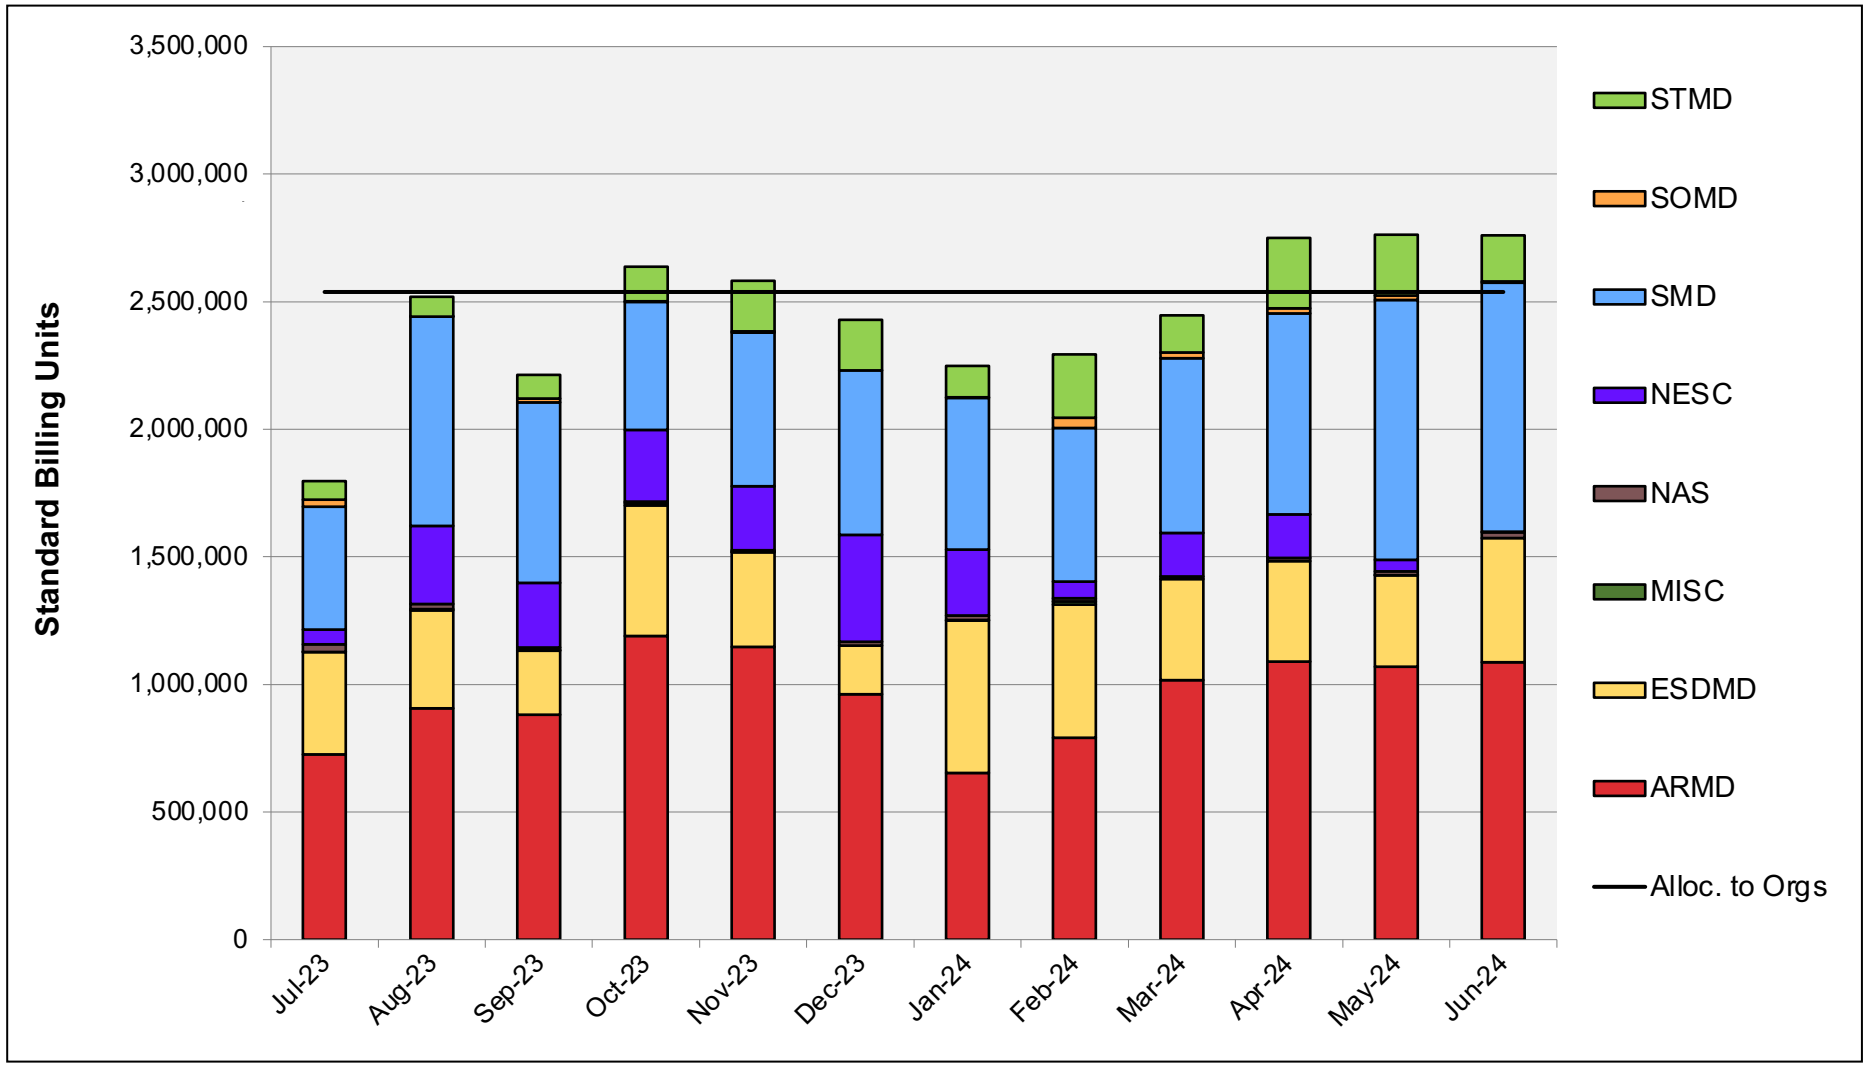
Pleiades: Average Time to Clear All Jobs



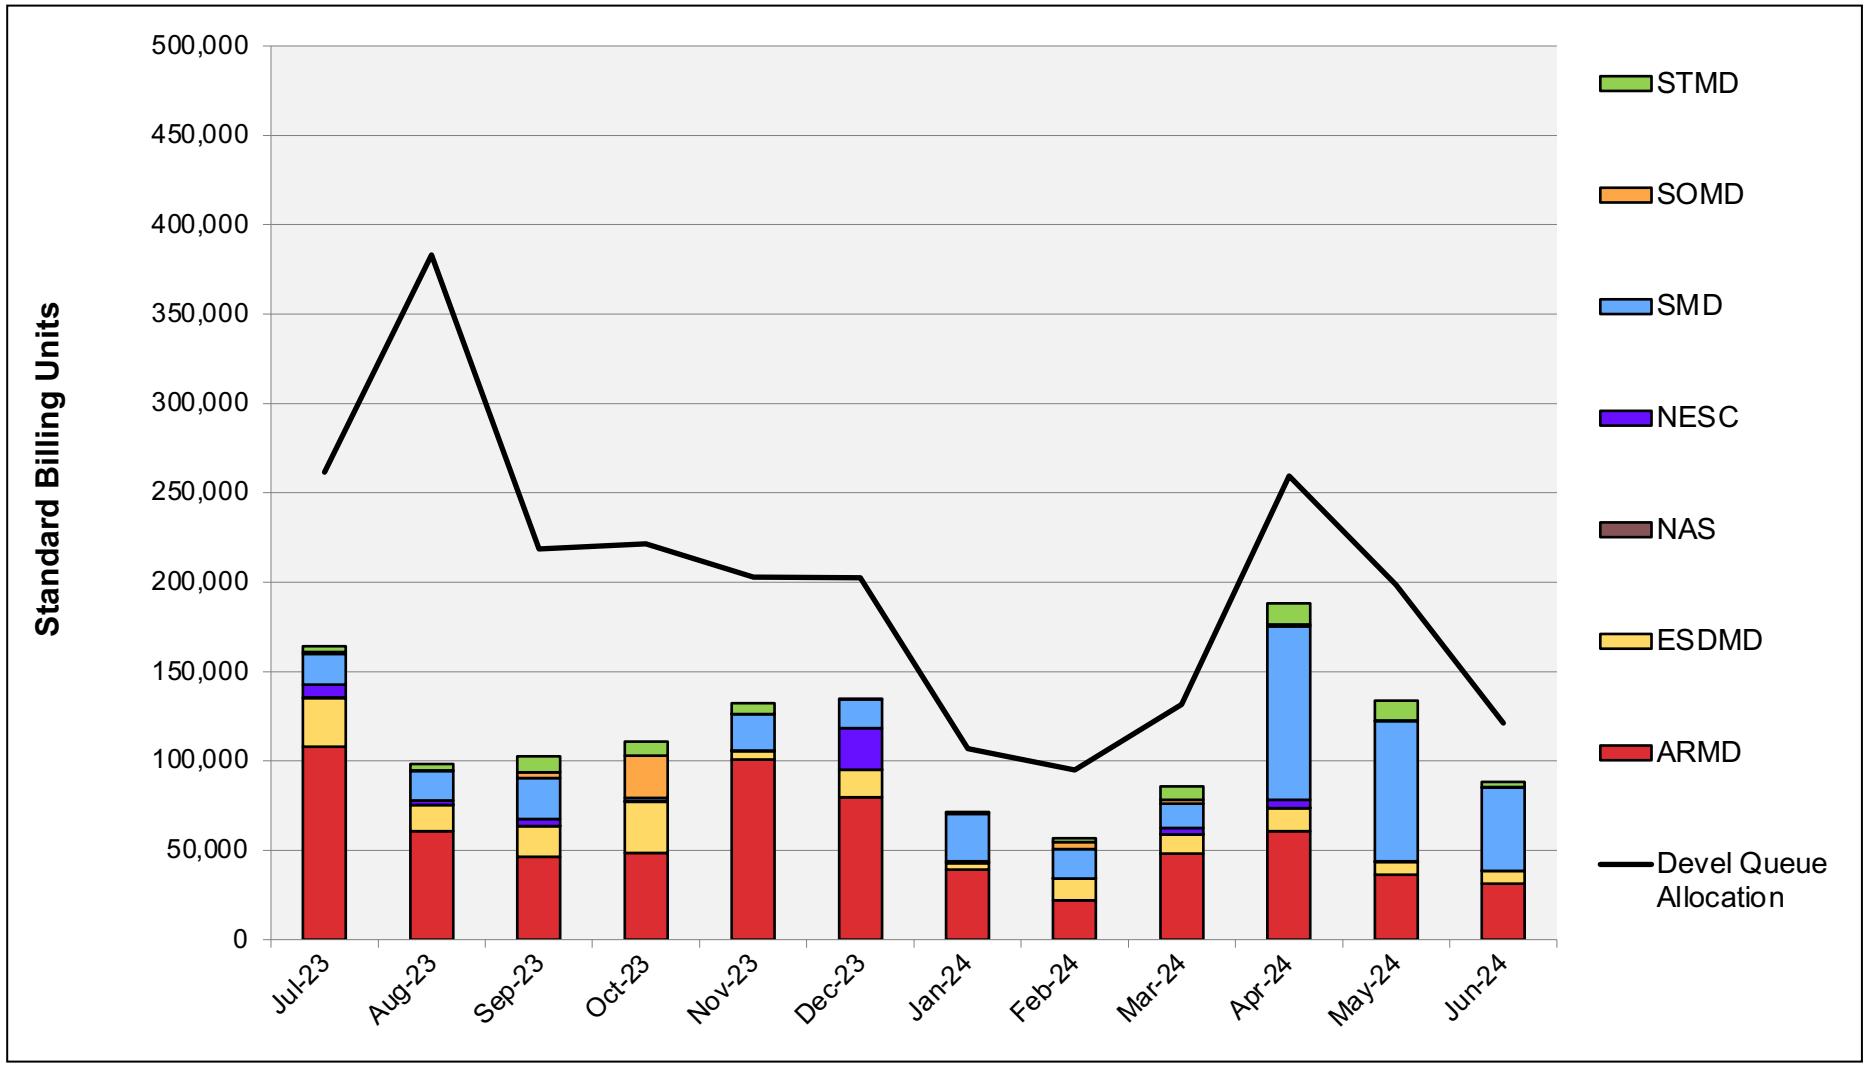
Pleiades: Average Expansion Factor



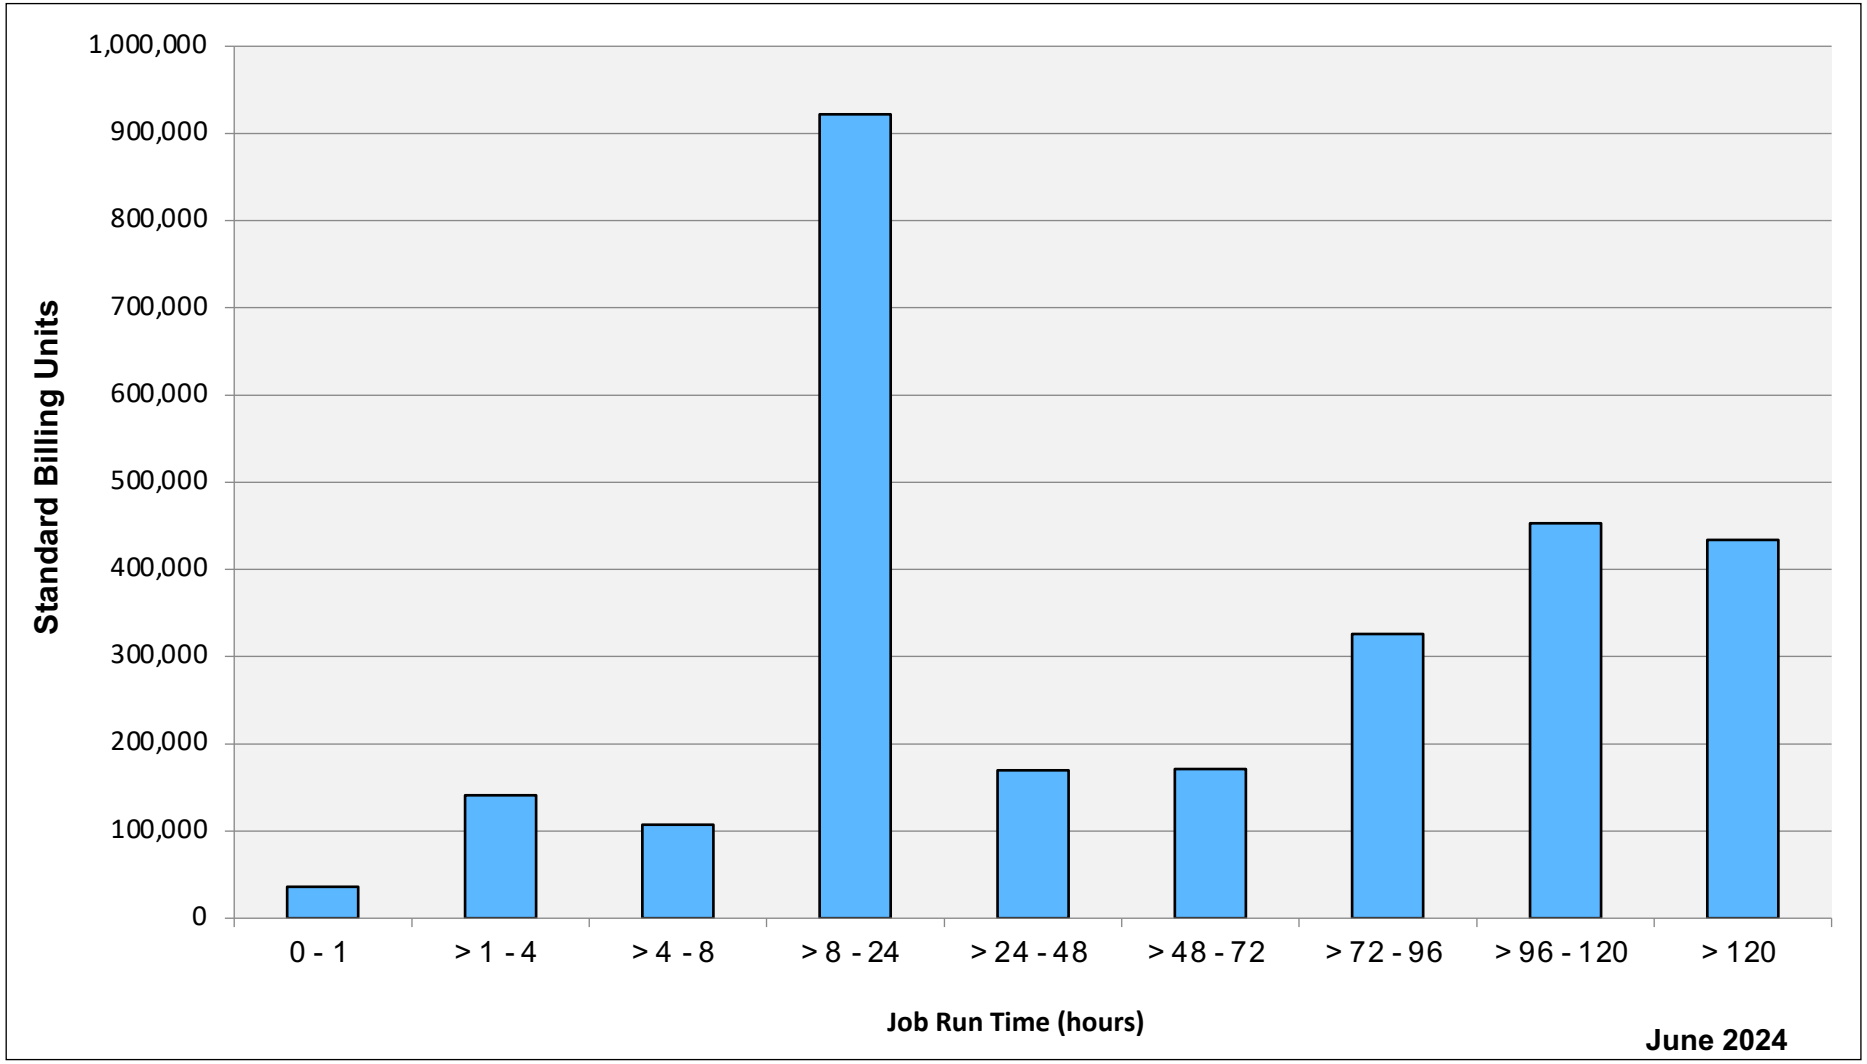
Electra: SBUs Reported, Normalized to 30-Day Month



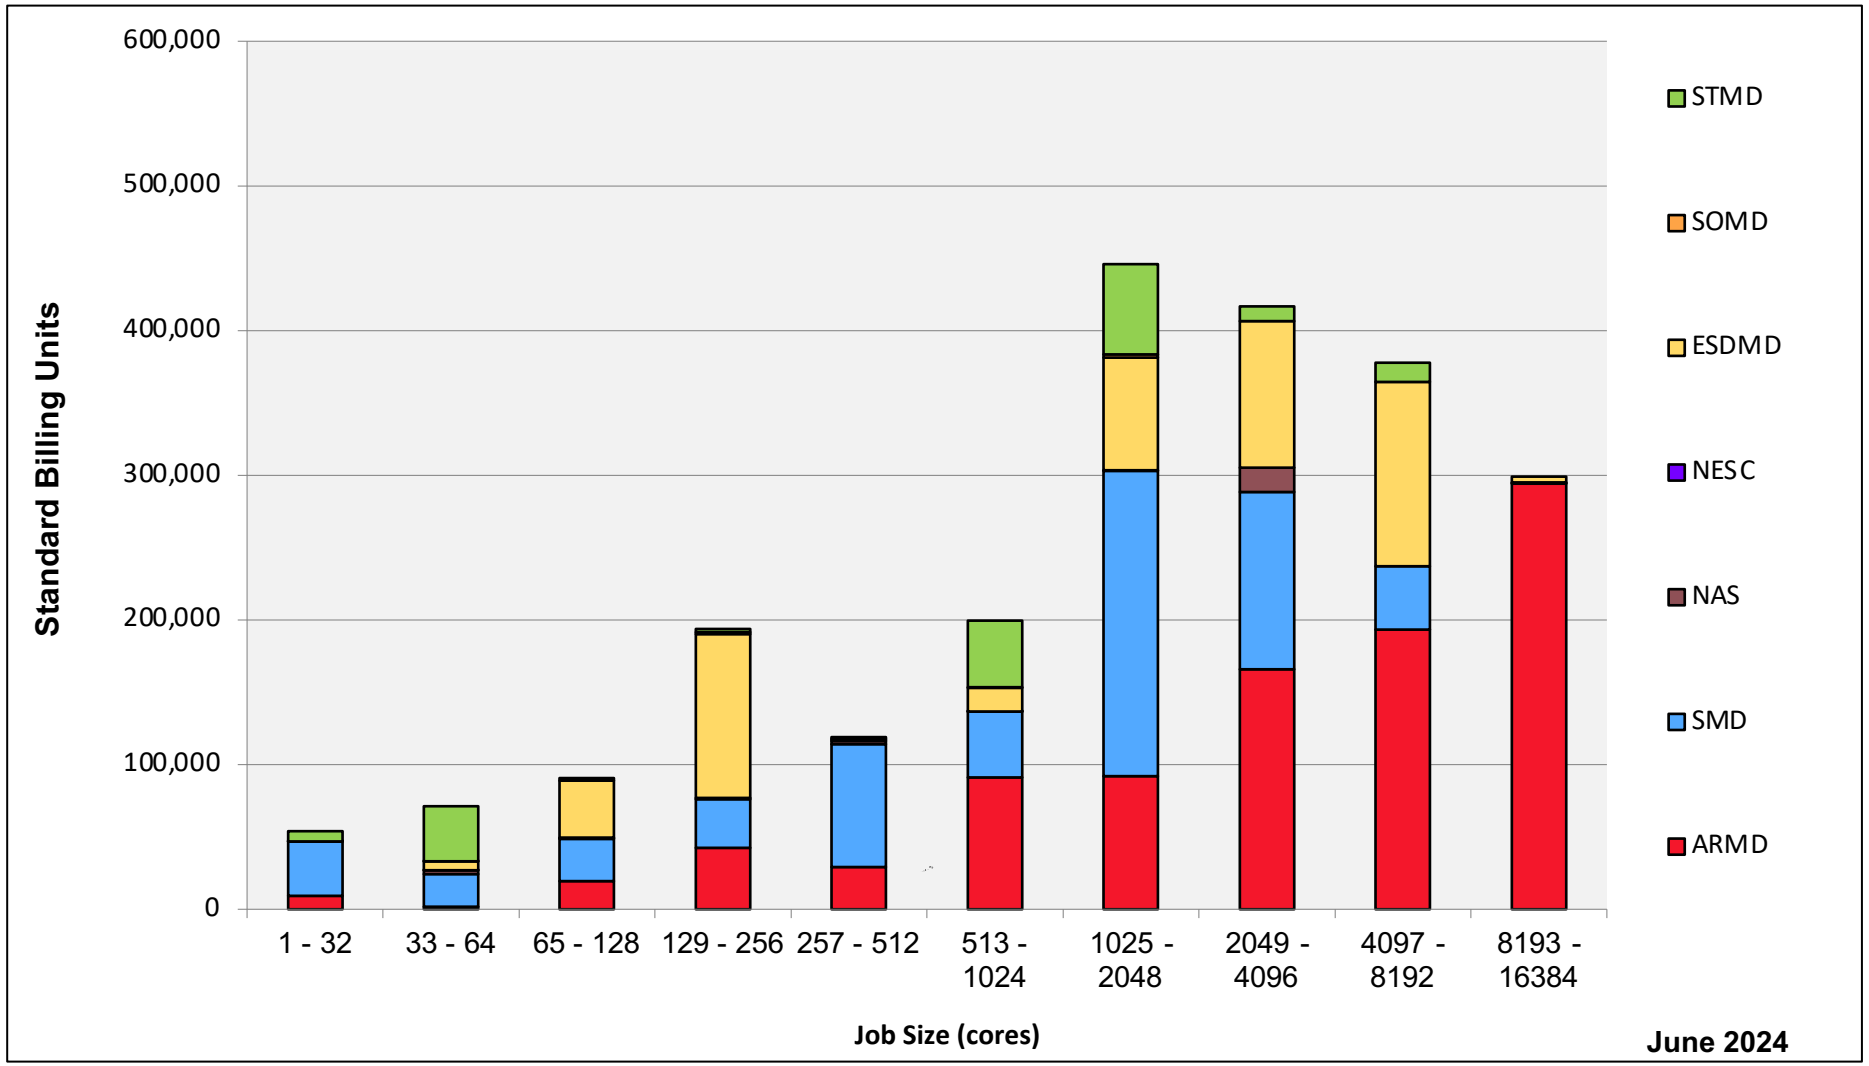
Electra: Devel Queue Utilization



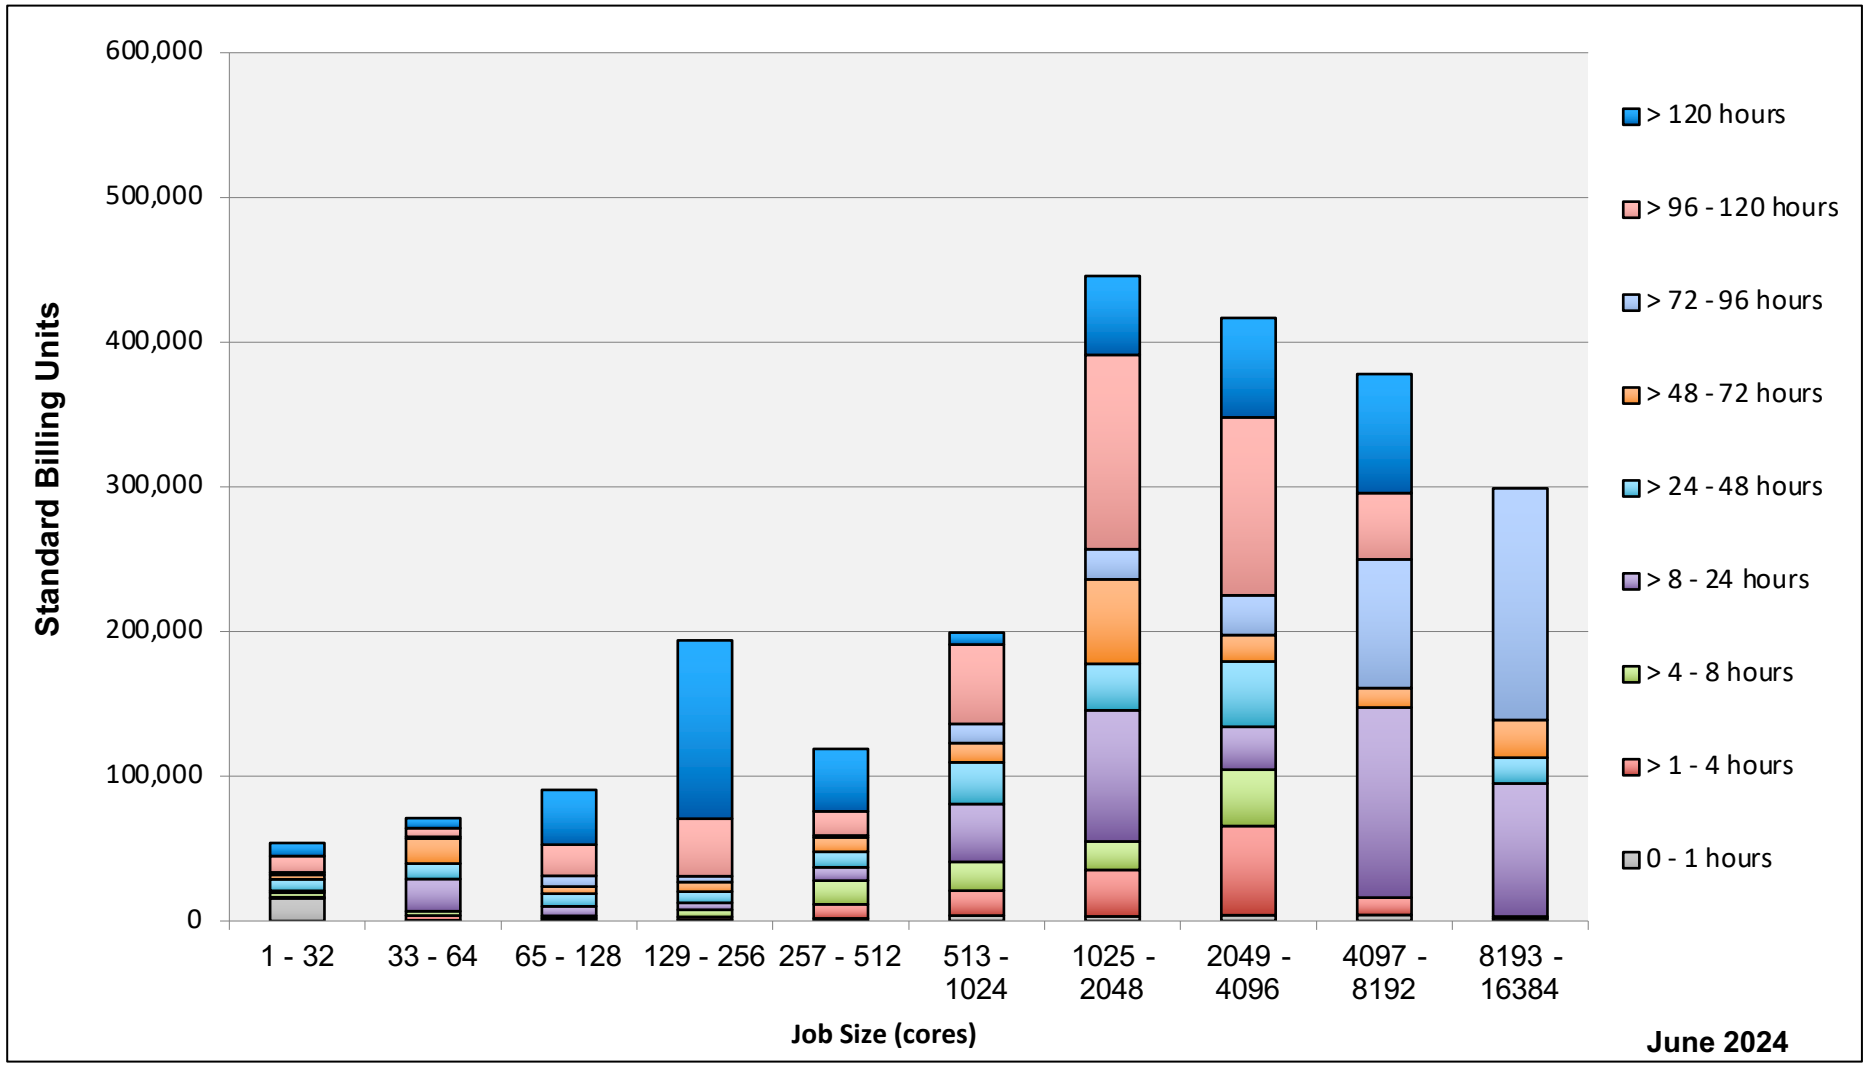
Electra: Monthly Utilization by Job Length



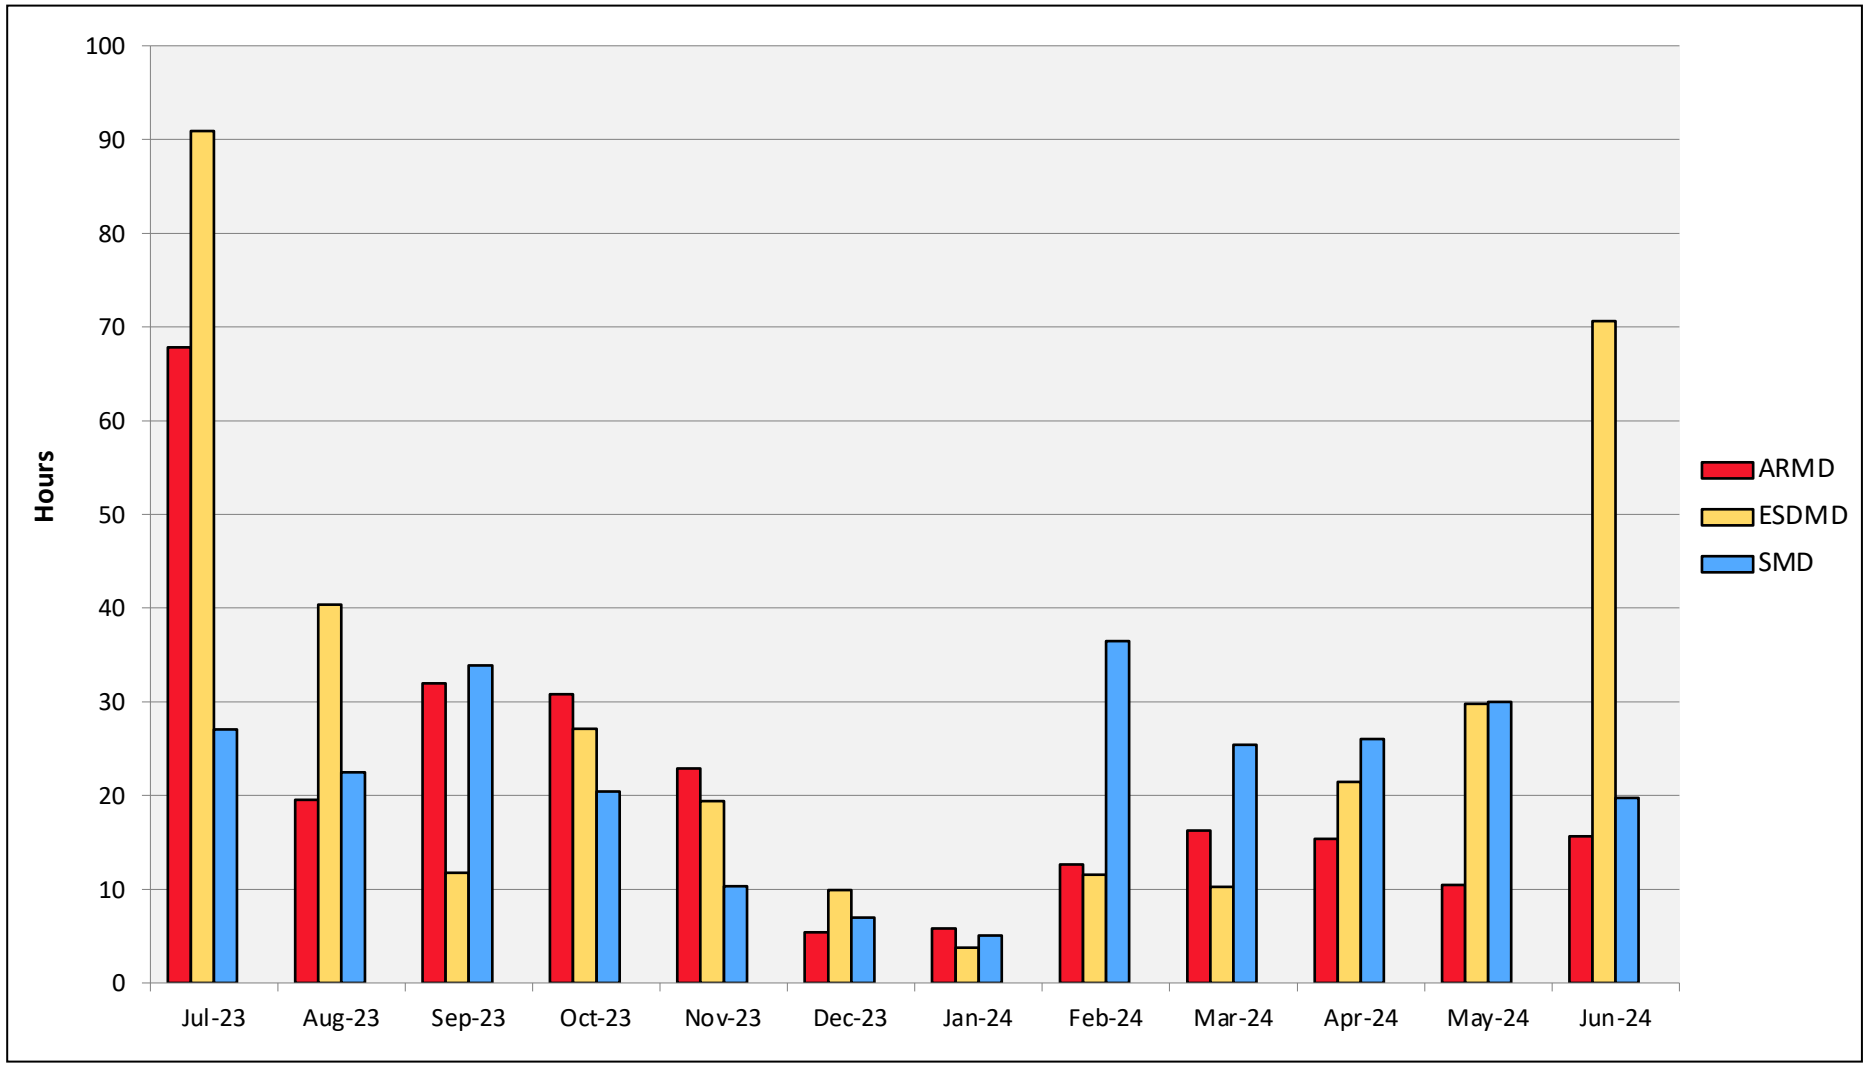
Electra: Monthly Utilization by Job Size



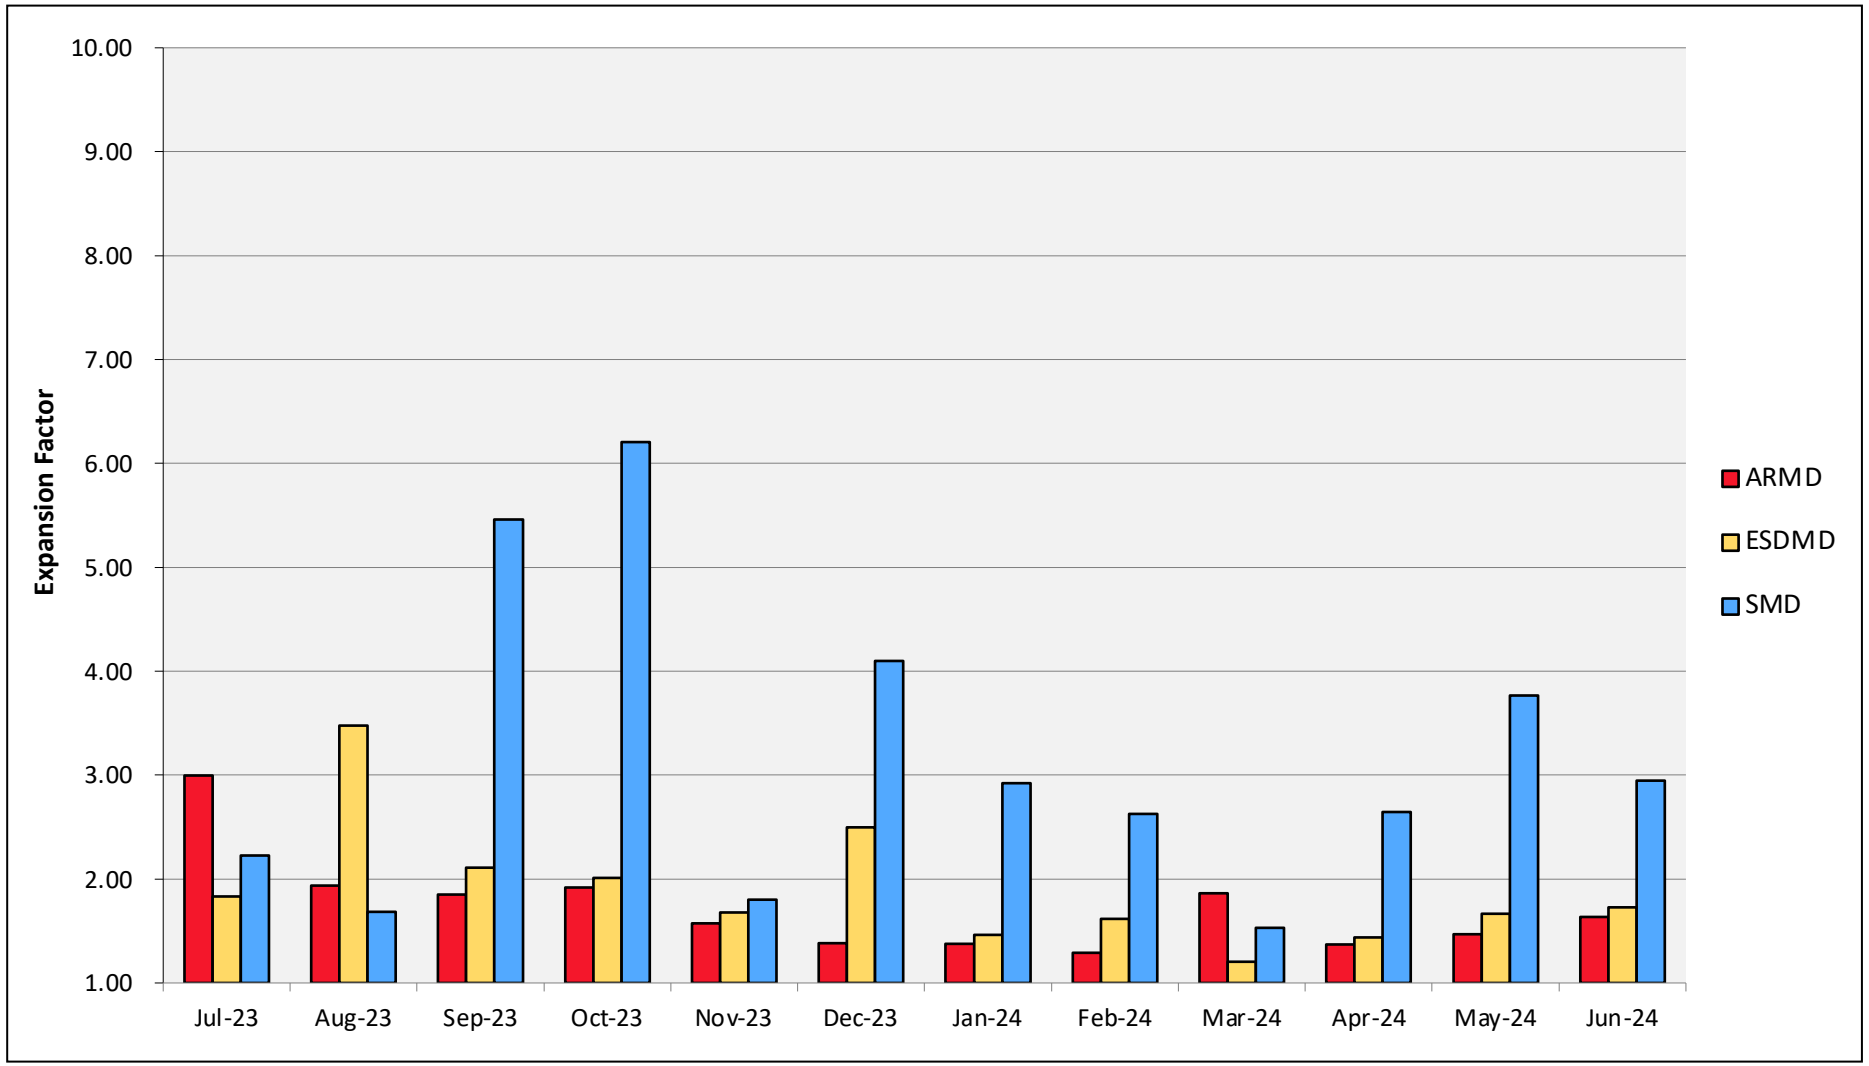
Electra: Monthly Utilization by Size and Length



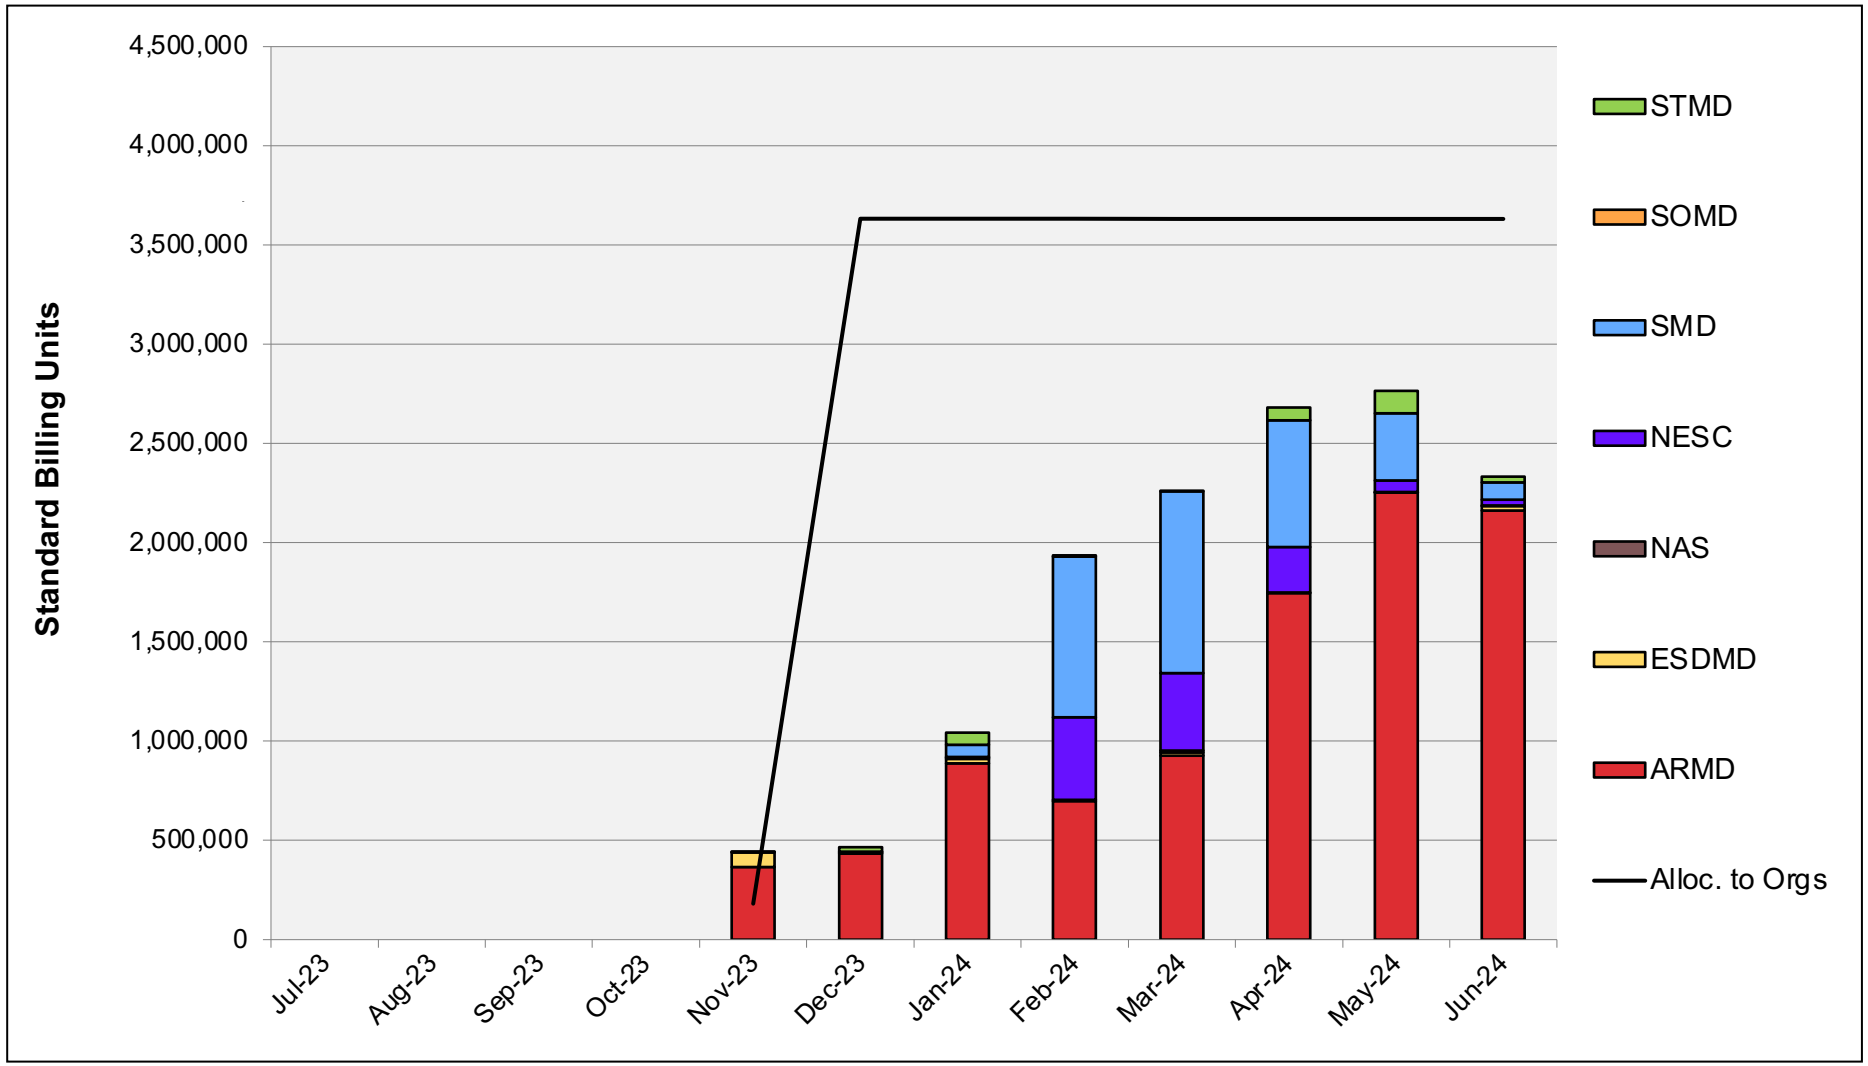
Electra: Average Time to Clear All Jobs



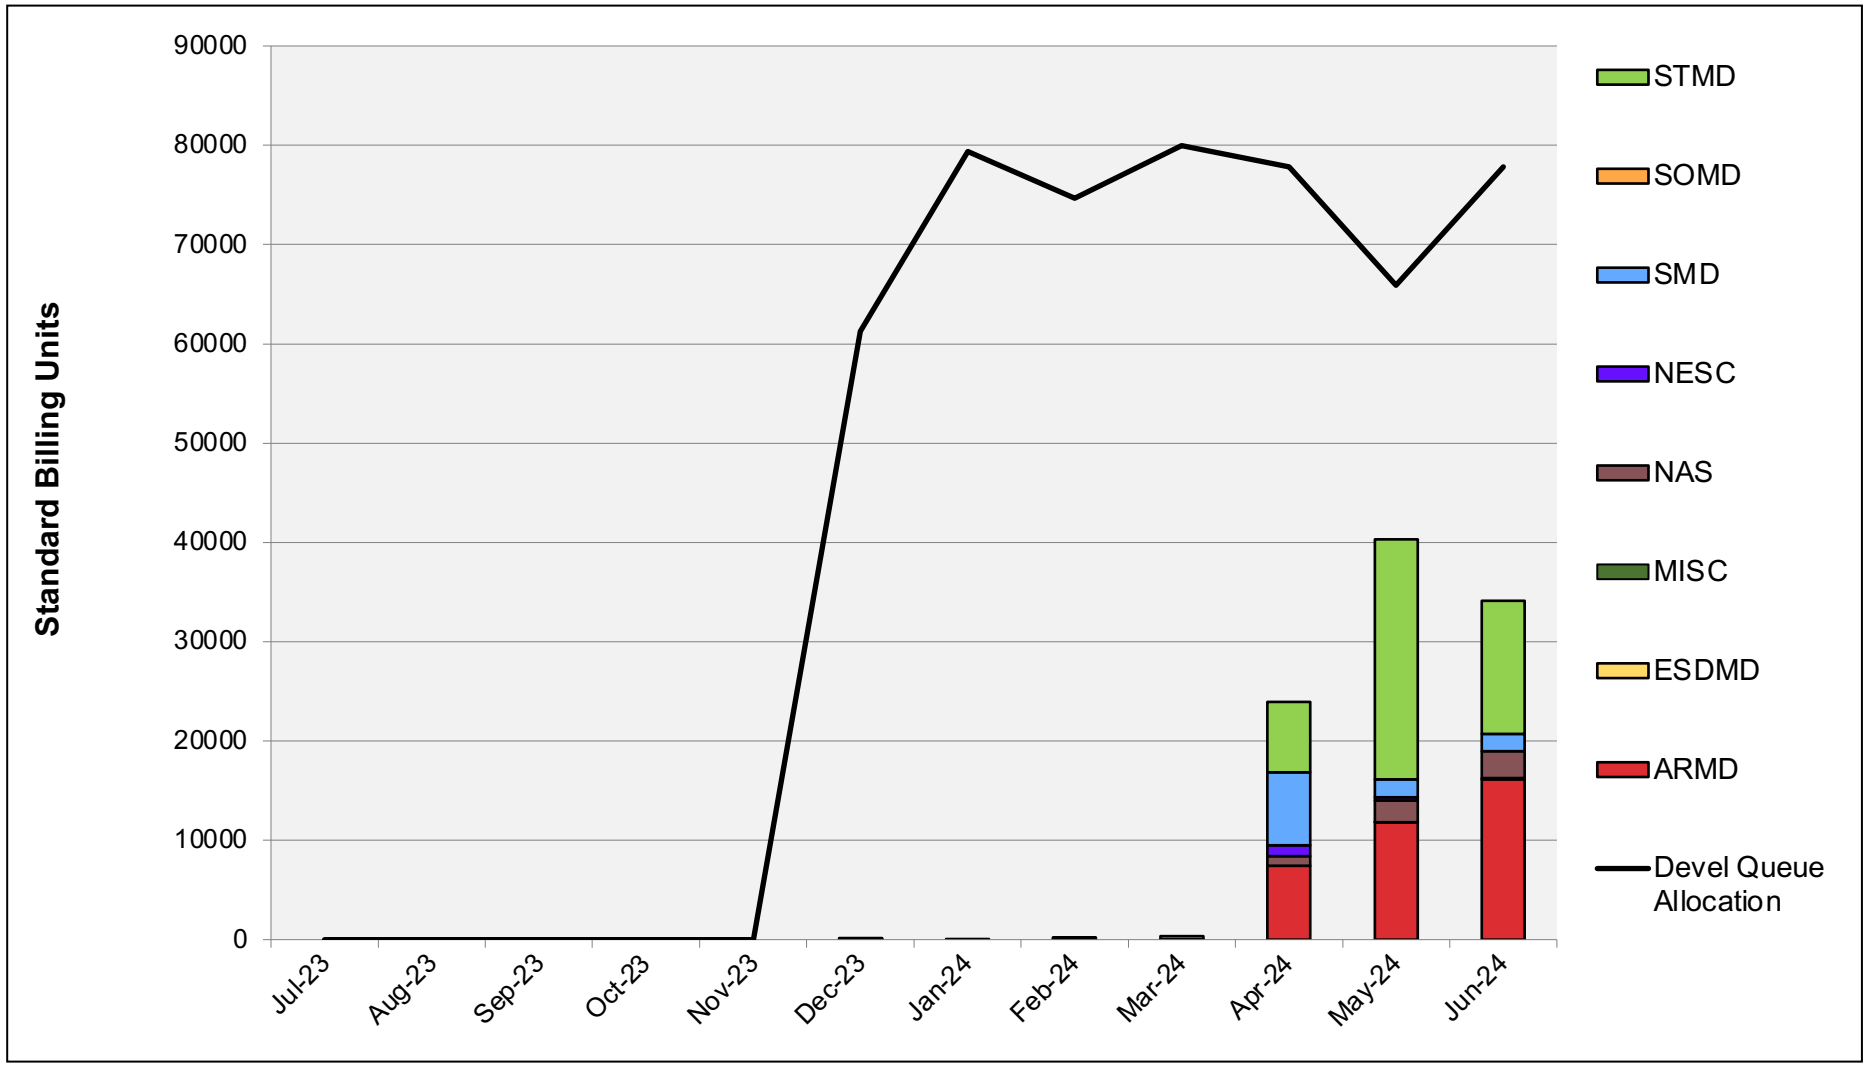
Electra: Average Expansion Factor



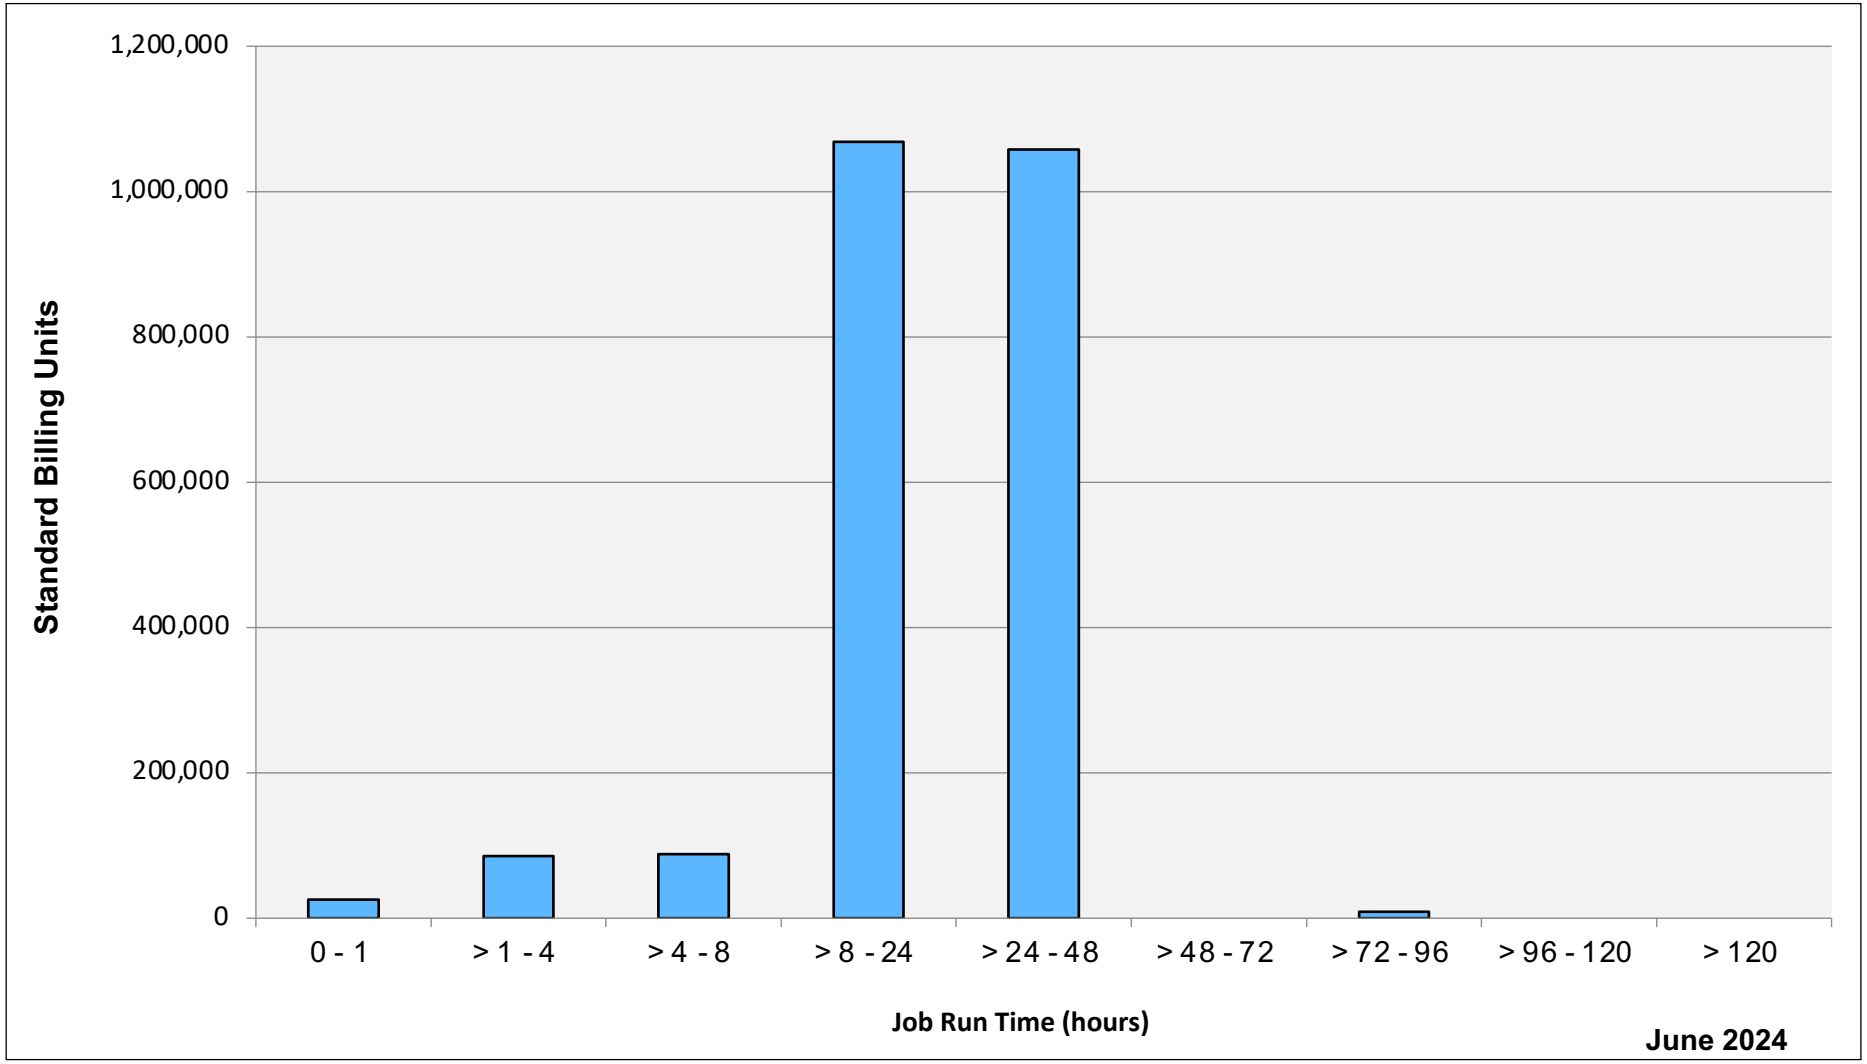
Cabeus: SBUs Reported, Normalized to 30-Day Month



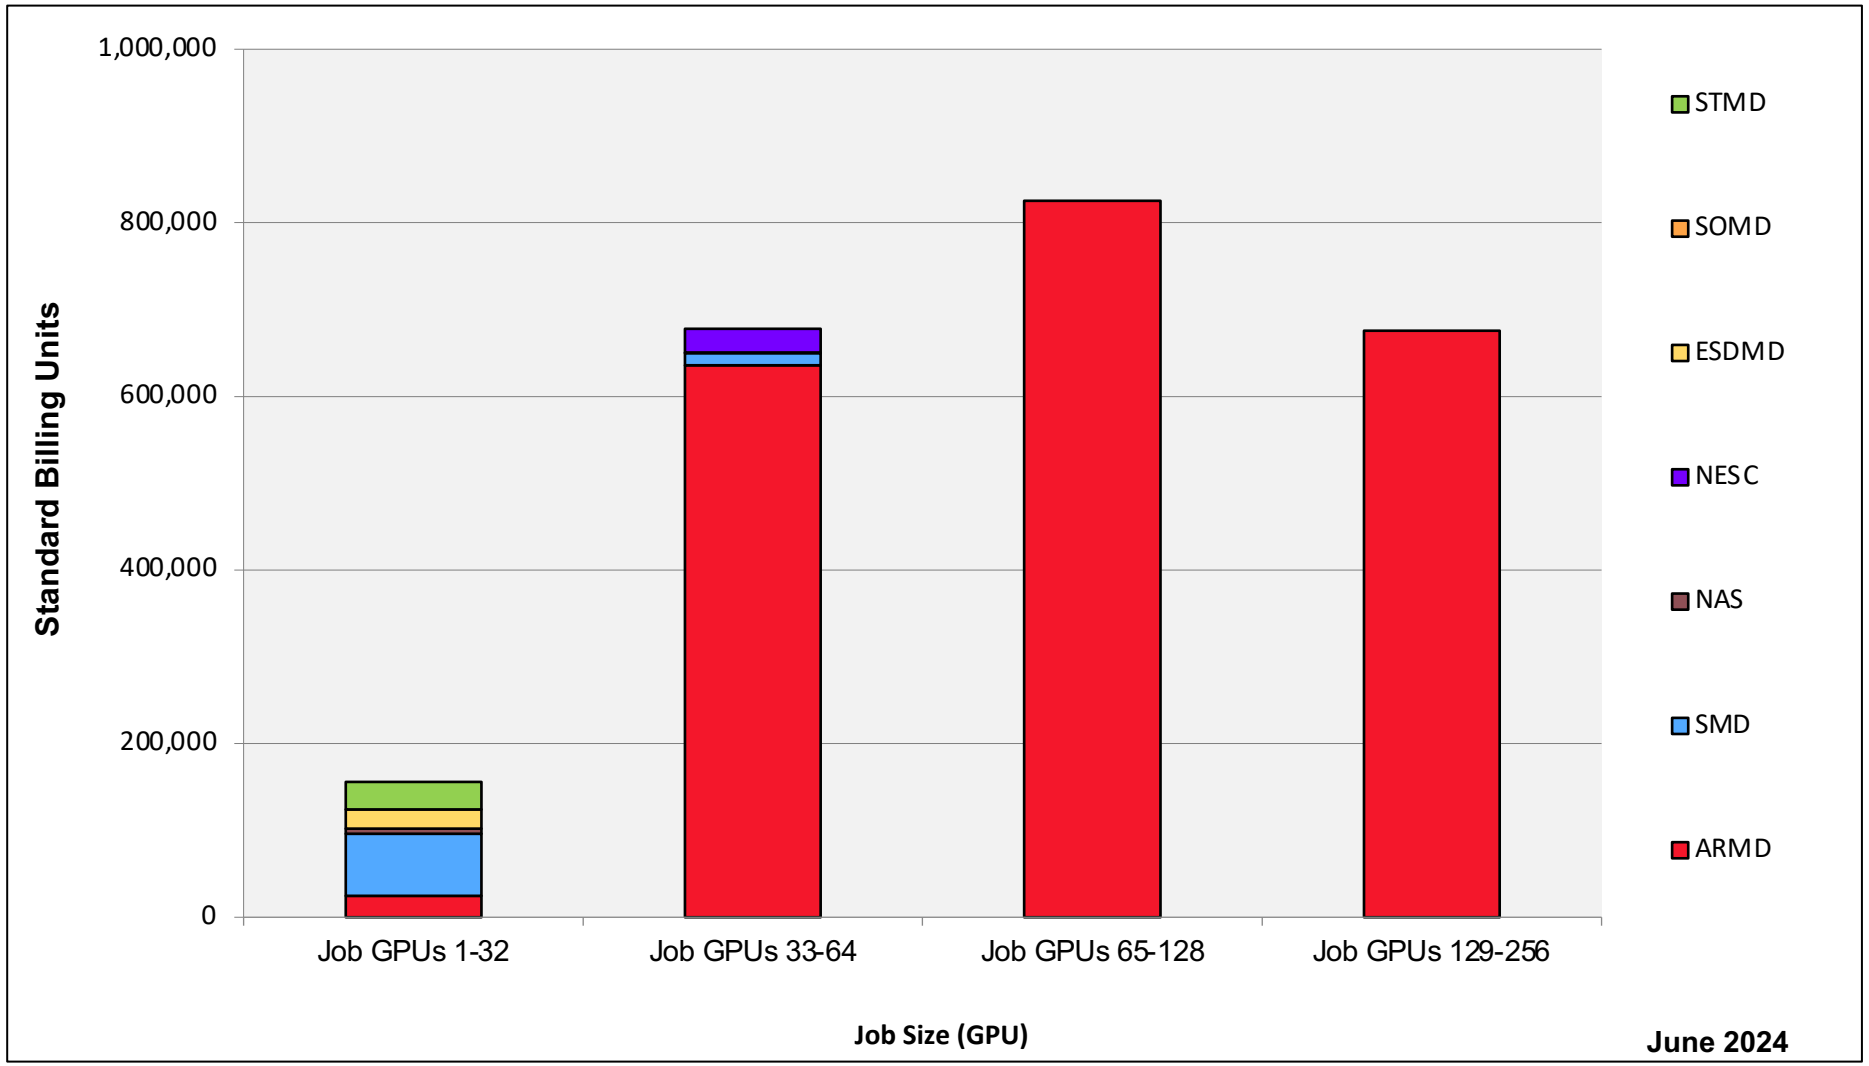
Cabeus: Devel Queue Utilization



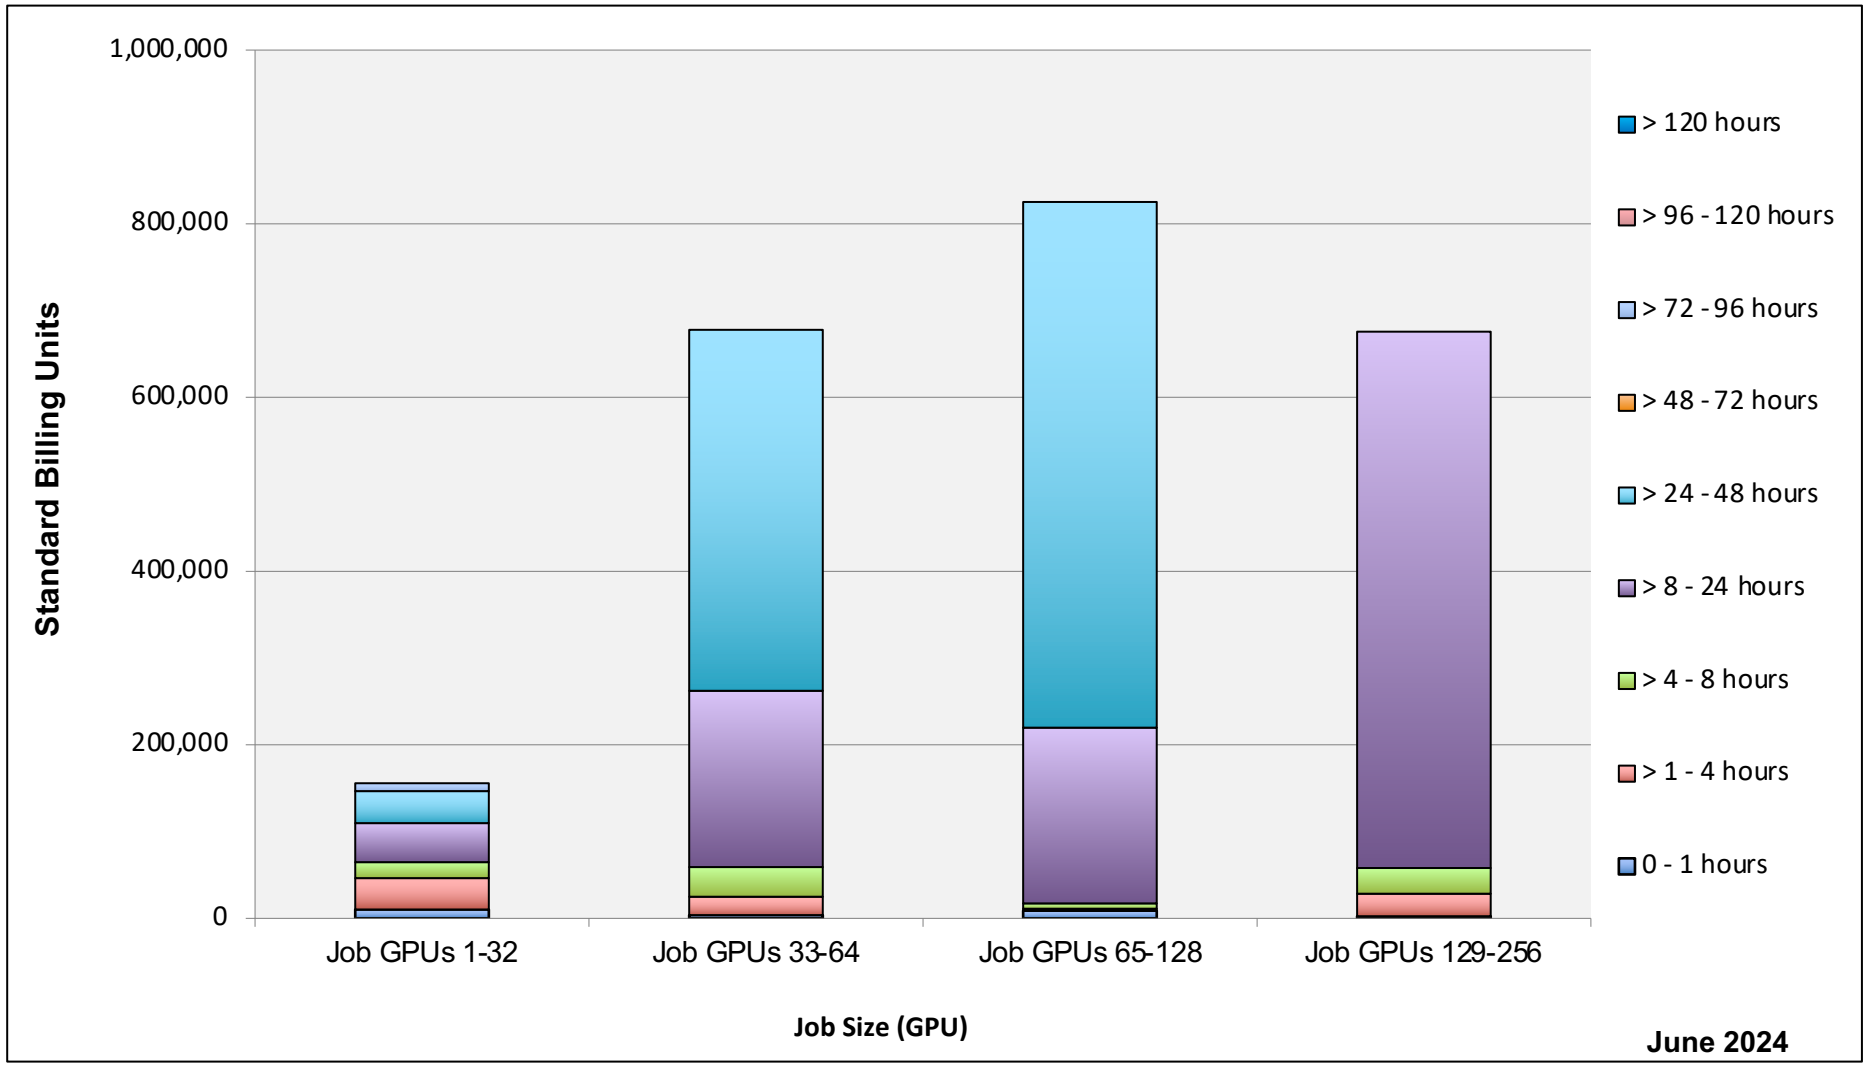
Cabeus: Monthly Utilization by Job Length



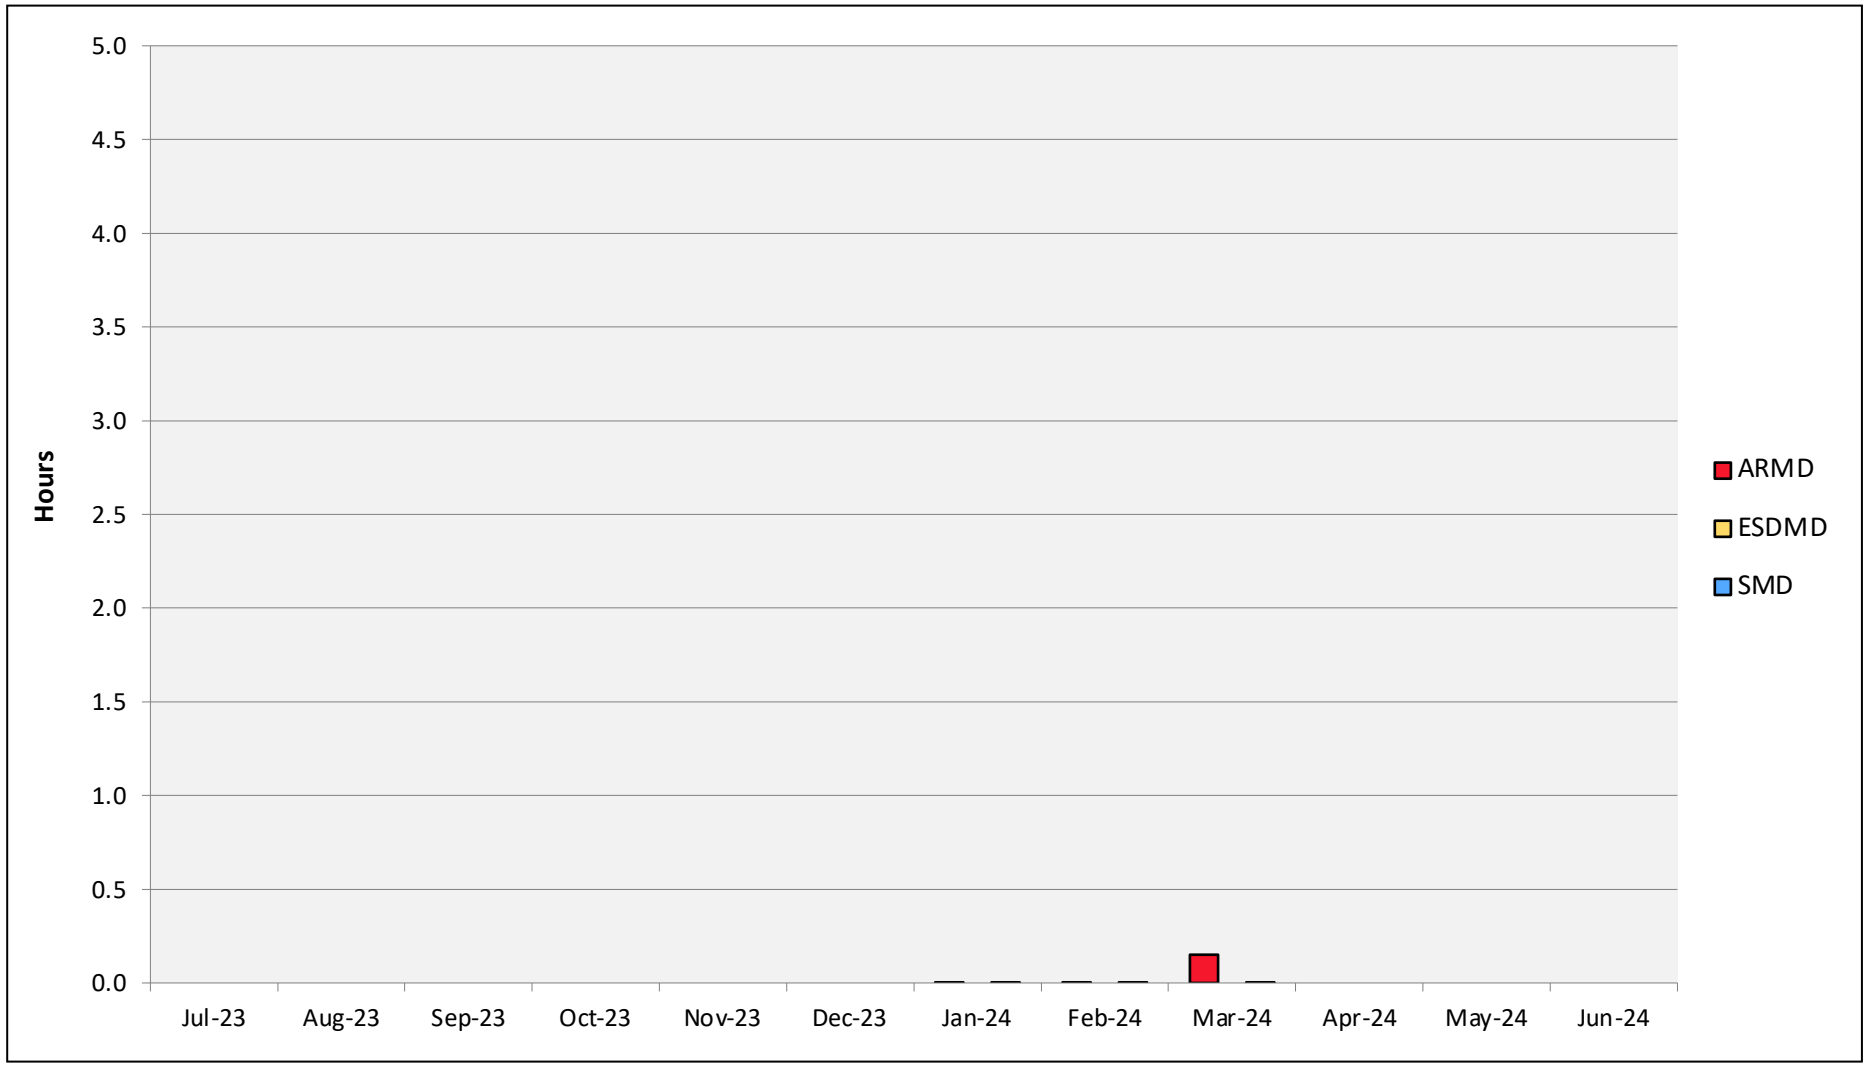
Cabeus: Monthly Utilization by Job Size



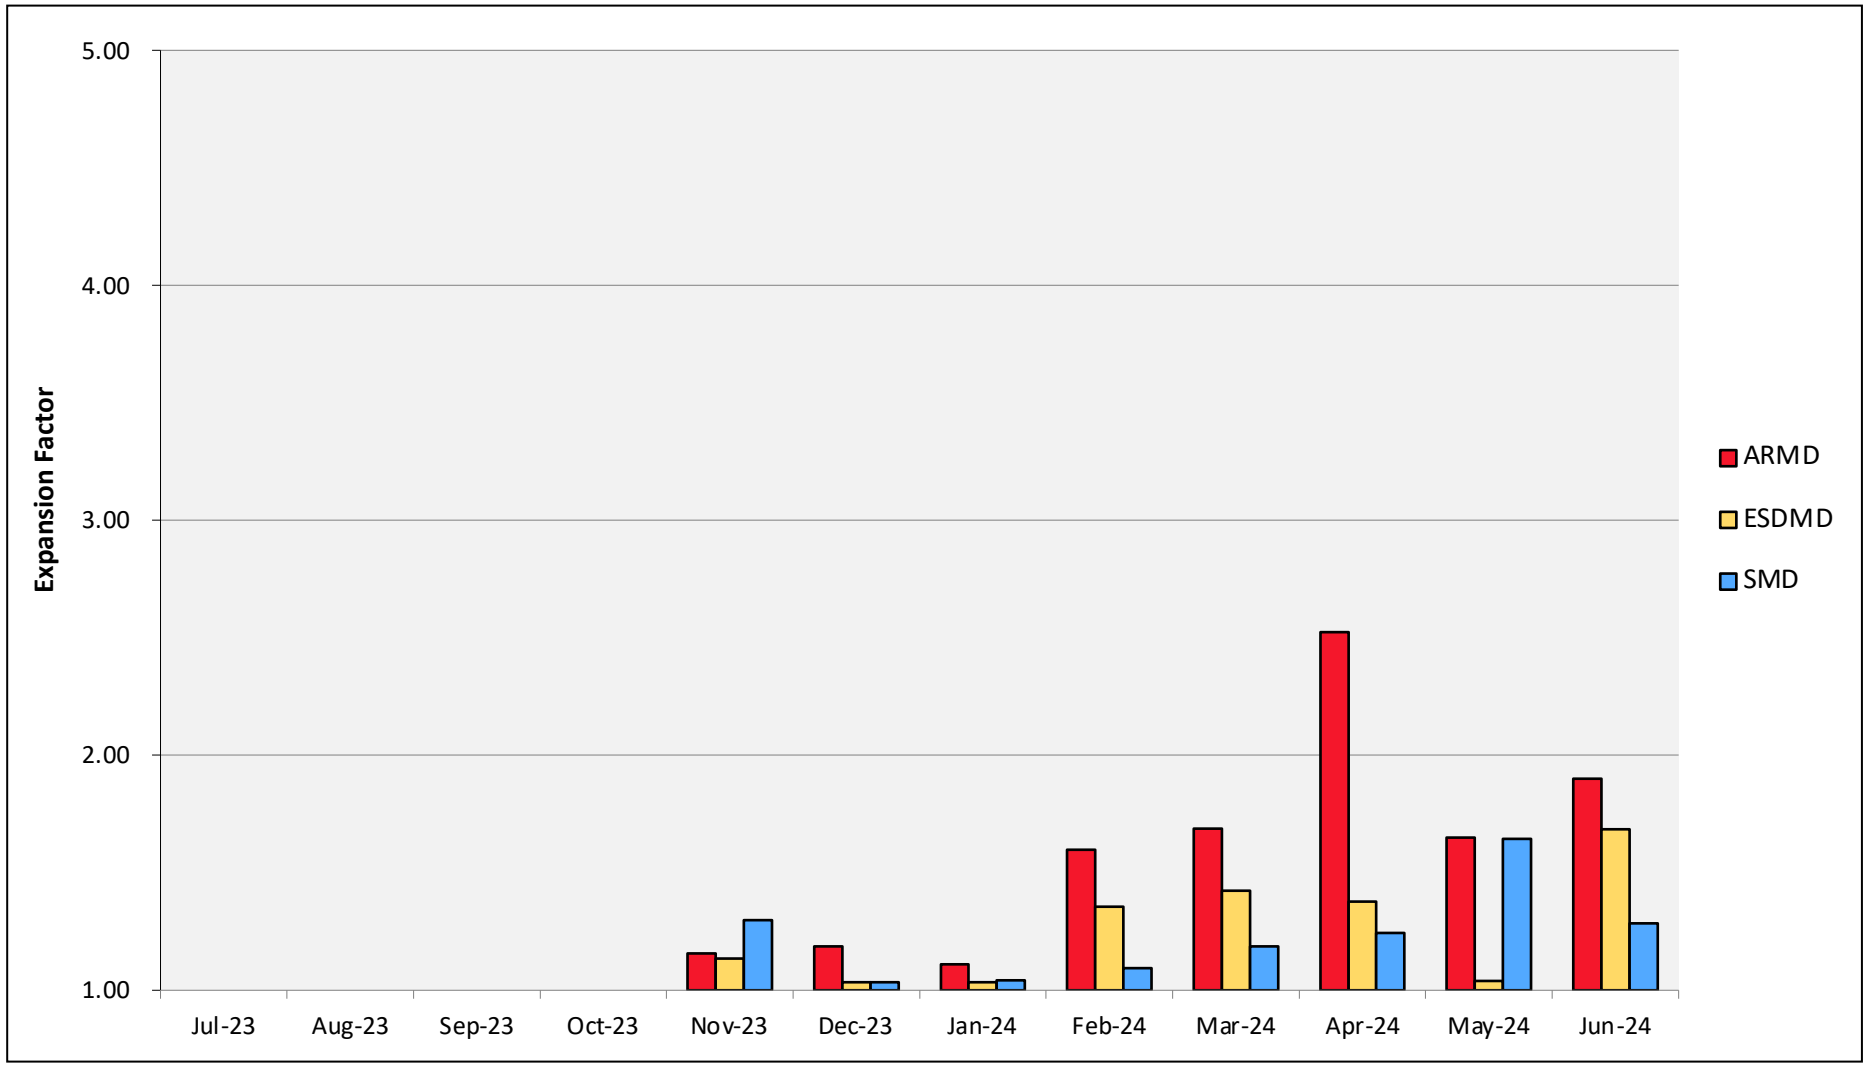
Cabeus: Monthly Utilization by Size and Length



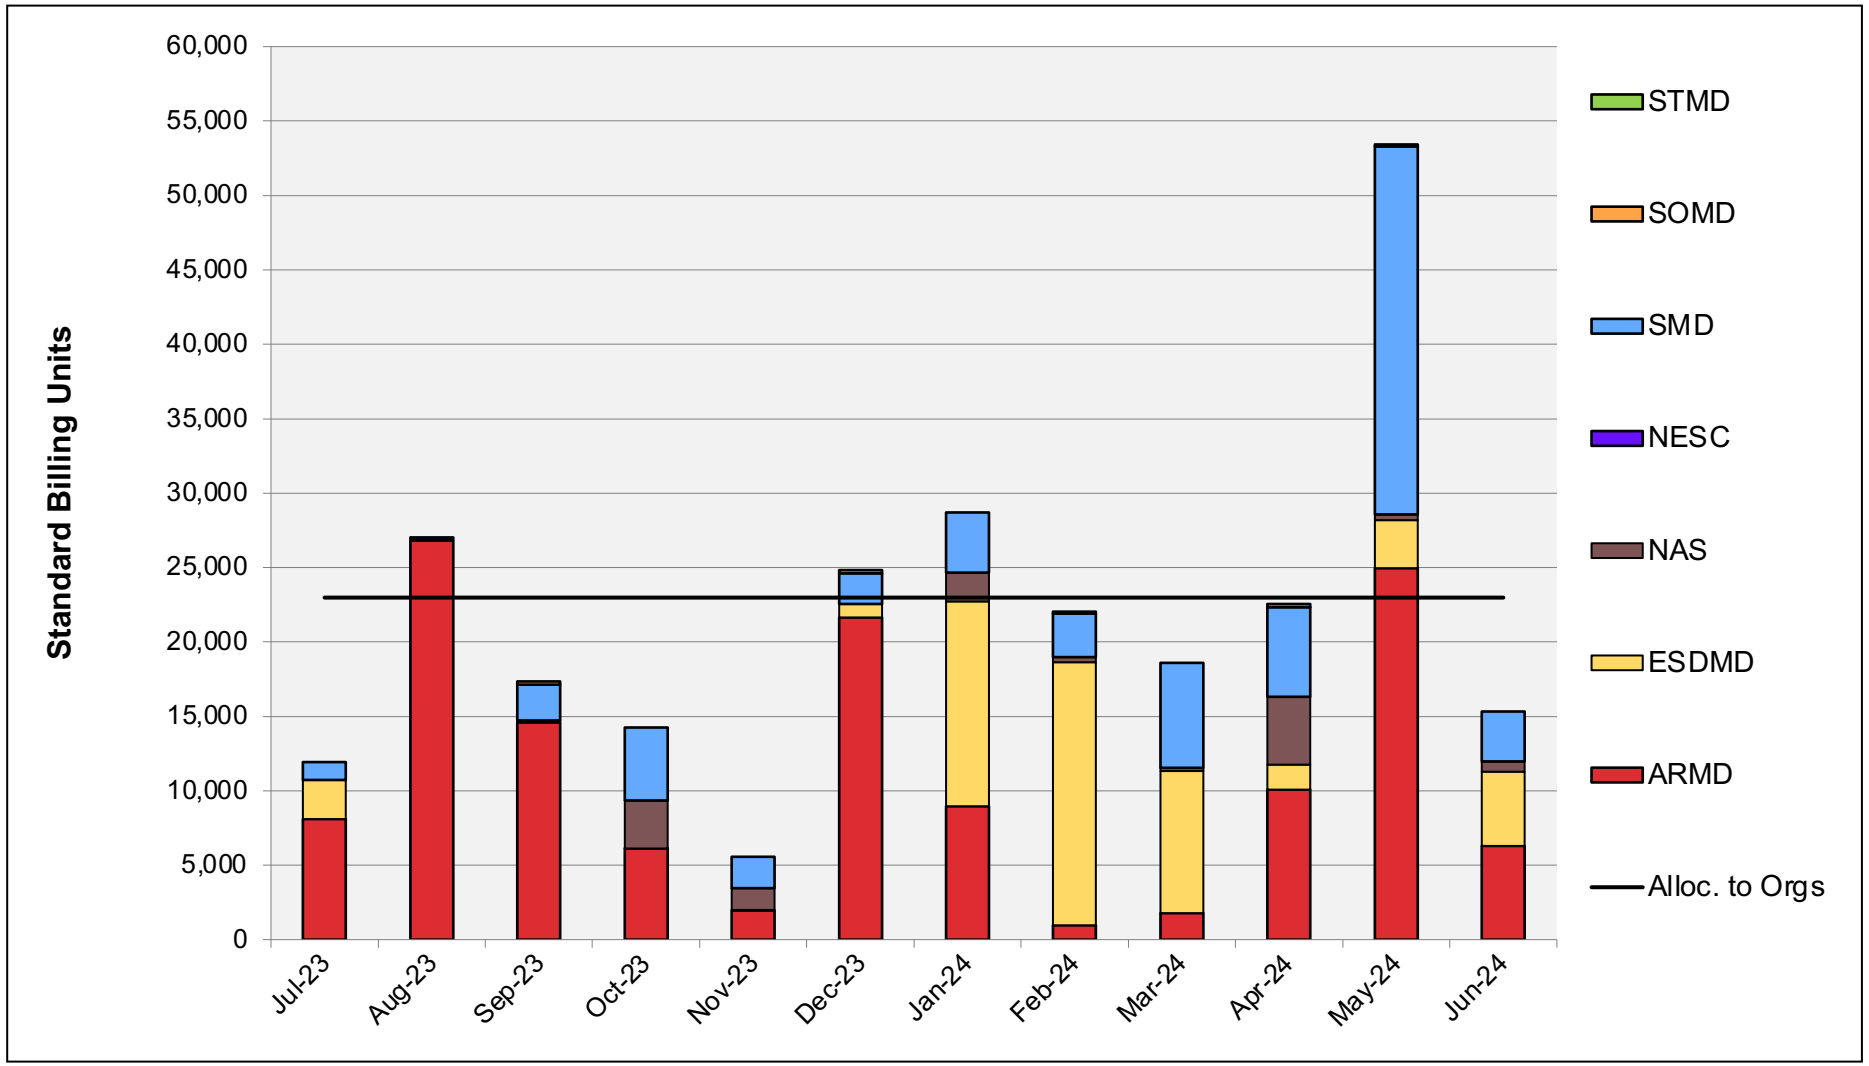
Cabeus: Average Time to Clear All Jobs



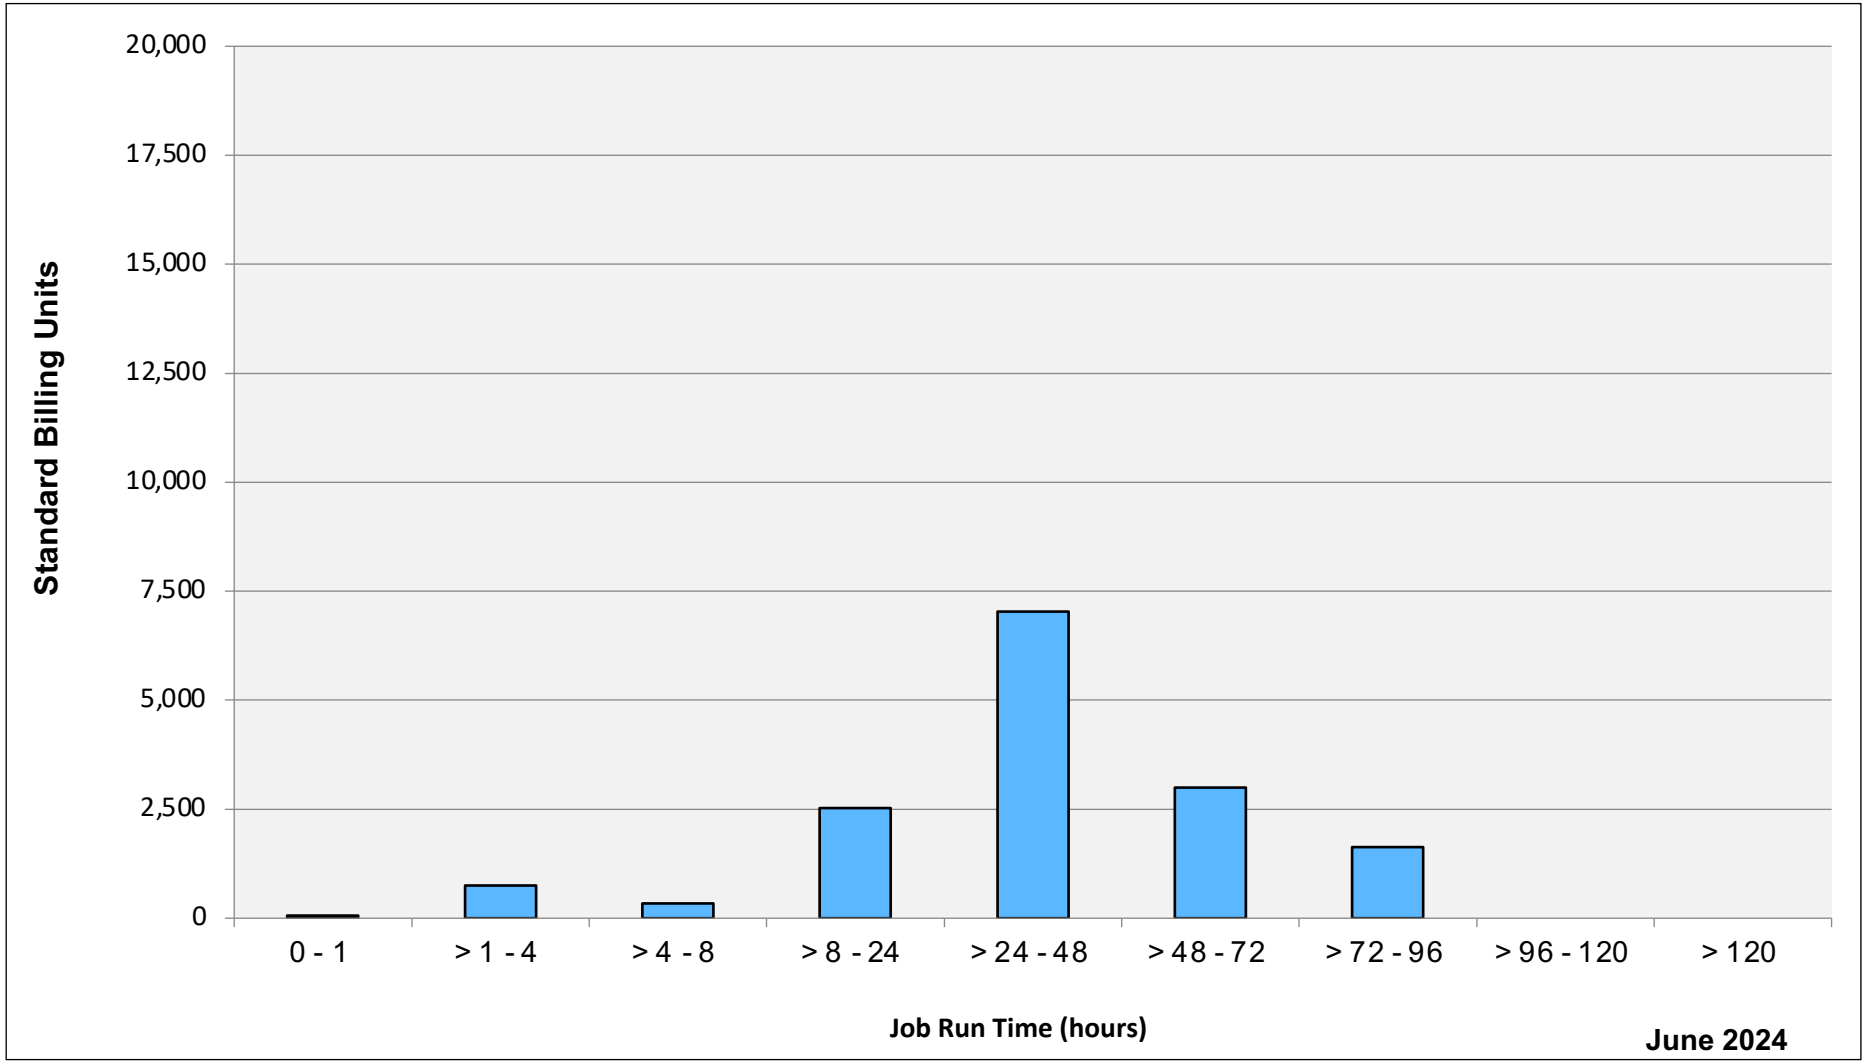
Cabeus: Average Expansion Factor



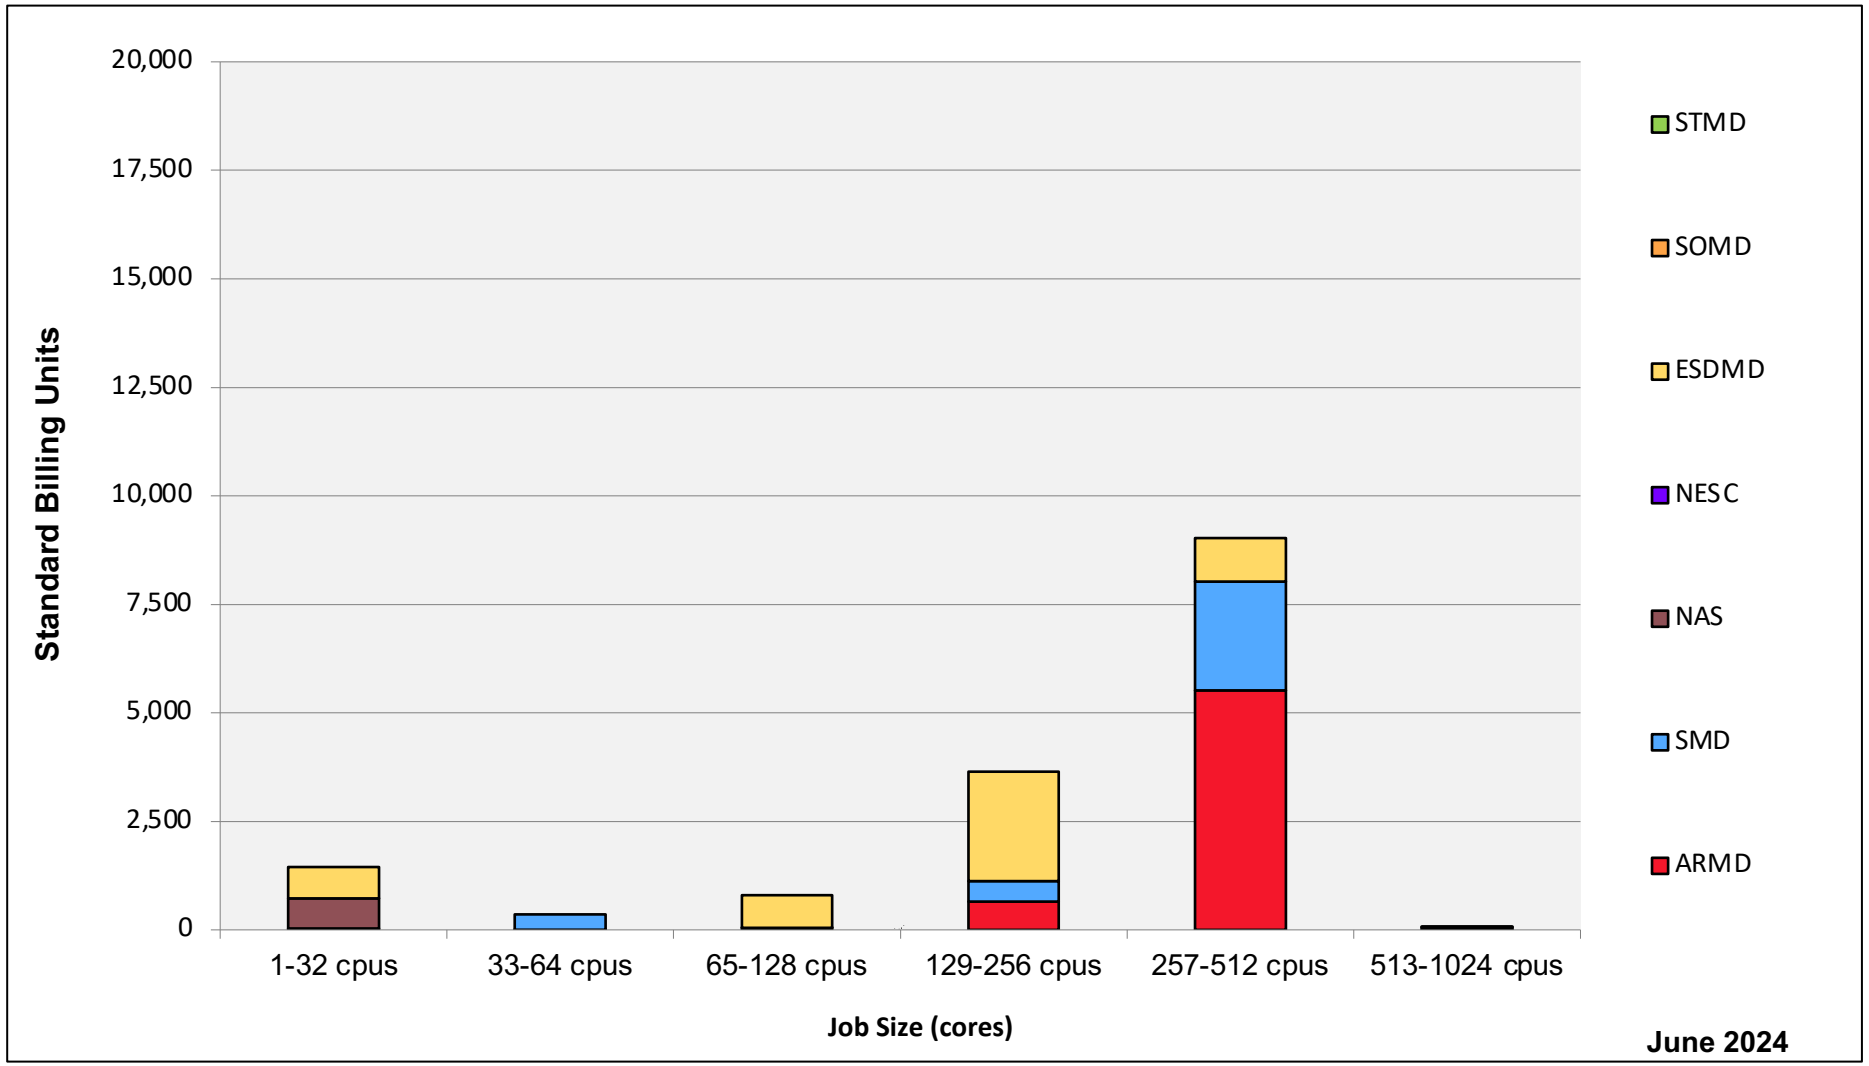
Endeavour: SBUs Reported, Normalized to 30-Day Month



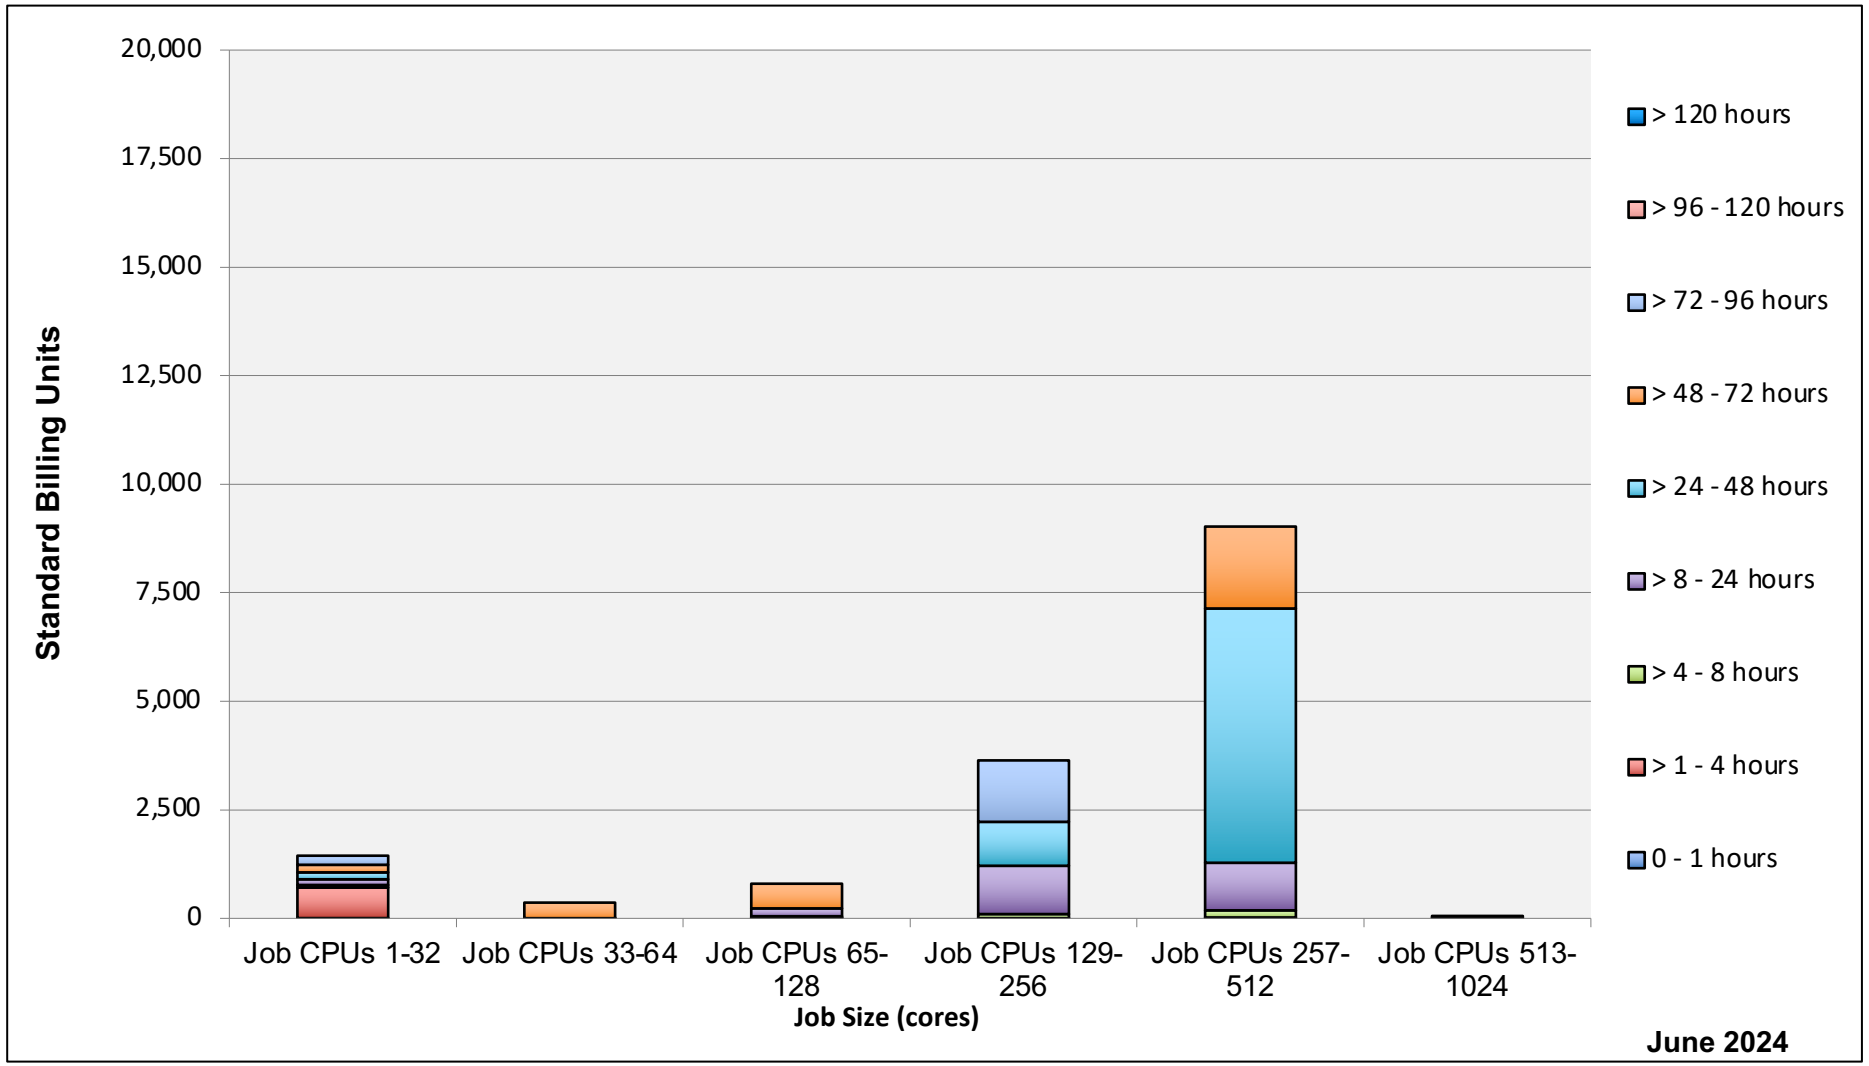
Endeavour: Monthly Utilization by Job Length



Endeavour: Monthly Utilization by Job Size



Endeavour: Monthly Utilization by Size and Length



Endeavour: Average Expansion Factor

