

National Aeronautics and Space Administration

High-End Computing Capability

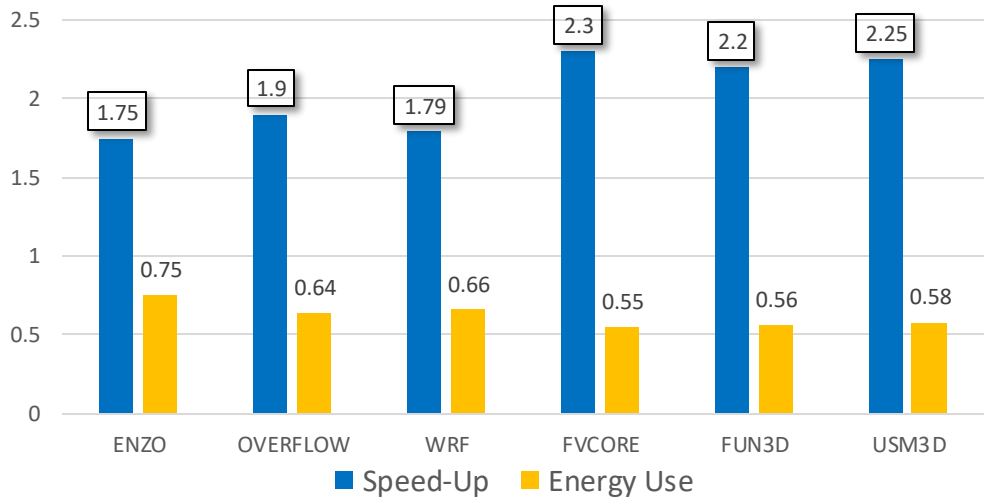
William Thigpen
HECC Component Manager
william.w.thigpen@nasa.gov

September 10, 2024

Grace CPU Superchip Evaluation Using the SBU1 Suite

- The HECC Application Performance and Productivity (APP) team ported the SBU1 Benchmark Suite to the NVIDIA Grace CPU pathfinding system installed at the NAS facility.
  - The pathfinding system consists of two NVIDIA Grace CPU Superchip nodes. Each node contains a single NVIDIA Grace superchip and 480 gigabytes of LPDDR5X memory. Each superchip consists of two NVIDIA Grace chips, with 72 ARM Neoverse V2 CPU cores per chip, connected with NVIDIA NVLink-C2C—a memory-coherent, high-bandwidth, and low-latency interconnect for superchips.
  - The SBU1 Benchmark suite is part of a set of benchmarks to determine HECC's Standard Billing Unit (SBU) rates. It consists of the following codes: OVERFLOW, FUN3D, USM3D , WRF, ENZO, and FVCORE.
- The purpose of the effort was two-fold:
  - To obtain a rough estimate of how the Grace CPU Superchip compares to AMD Rome nodes in terms of performance and energy consumption.
  - To evaluate the porting effort from Intel and AMD x86-based processors, like the AMD Rome nodes available at NAS, to the ARM architecture.
- The APP team conducted single-node experiments only, using a variety of compilers and MPI stacks. The initial observations show:
  - Out of the box, the Grace Superchip node delivered a 1.75 – 2.3x speed-up compared to a 2-processor Rome node, while consuming only 0.55 – 0.75x the amount of the energy.
  - Porting the benchmarks to run on the ARM processors was easy, requiring only few minor code modifications in two of the benchmarks. This was the first time that all six benchmarks were successfully ported on an architecture other than x86-64.

IMPACT: Porting the SBU1 Benchmarks to the NVIDIA Grace CPU system enabled an estimate of the expected performance, energy consumption, and ease of use compared to Rome nodes.



This chart displays the ratios of out-of-the-box performance and energy consumption of all SBU1 benchmarks Grace/Rome. Out of the box, Grace delivered between 1.75 – 2.3 times the performance while consuming only 0.55-0.75 times the amount of energy compared to Rome nodes.

Grace Hopper Chip Evaluated For Data Science

- The HECC Data Science team evaluated the Grace Hopper—the latest chip developed by NVIDIA. This work complements testing by the APP team (see slide 4). Designated as the GH200, this computer chip combines the latest CPU and GPU technology into one integrated circuit.
- The HECC team used their Data Science Benchmarks to evaluate the Grace Hopper chip.
  - Comparing model training times is one way to benchmark HPC systems on AI/ML tasks. However, data input and output (I/O) can have a big impact on training times.
  - Training a model using real data and synthetic data reveals the impact of data I/O.
  - The GH200 test system uses an NFS filesystem and the existing A100 GPUs at NAS are connected to the faster Lustre filesystem.
  - The STEMDL-STFC dataset was used to train a ResNet30 model with data I/O.
  - Synthetic images generated using the Image-net format were used to train a ResNet50 model training without data I/O.
  - The ReseNet50 models were trained using the new GH200 test system. The training times were compared to the training times using an A100 GPU.
- The benchmark results revealed:
  - The GH200 chip trained the ResNet50 model in 12.76 minutes without I/O.
  - The A100 GPU node trained the ResNet50 model in 23.92 minutes without I/O.
  - These results indicate that without I/O the GH200 AI chip is approximately 2 times faster than the A100 GPU in training AI models.

IMPACT: HECC evaluations show that the increased performance of the new Grace Hopper superchip will enable more scientific research and faster time to solutions.

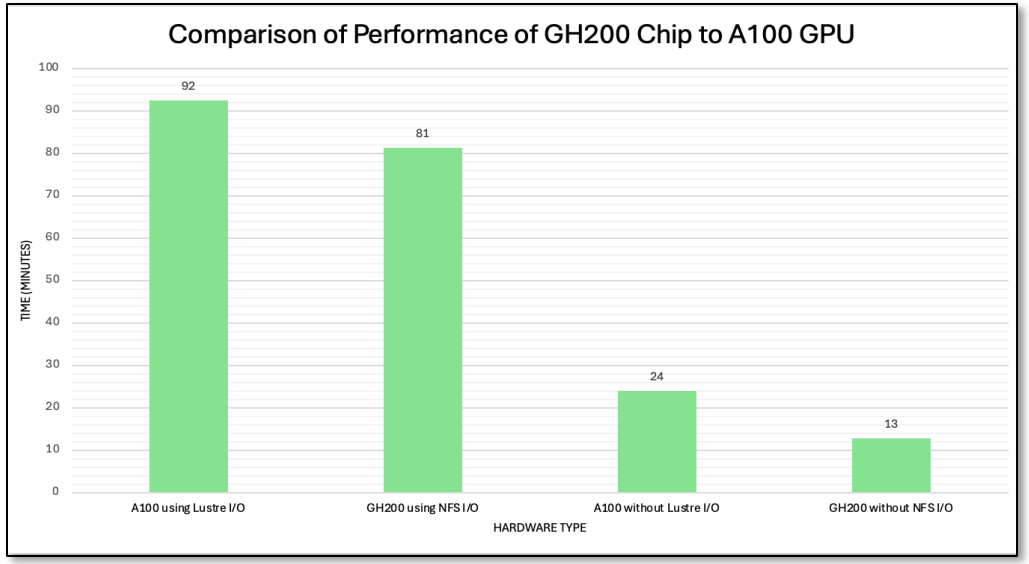
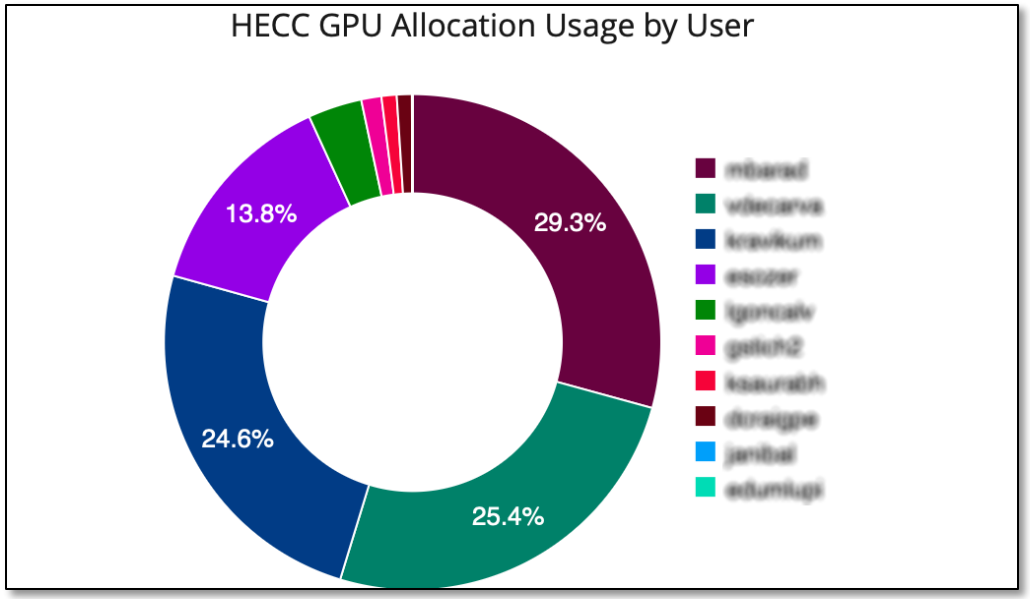


Chart comparing the performance of the NVIDIA GH200 system with the NVIDIA A100 system, with and without I/O. The GH200 performs about 2 times faster than the A100 system when training models. However, data I/O plays a key role in training time and can significantly increase it. Note that the filesystem used by the A100 (Lustre) is faster than that used by GH200 (NFS): if GH200 used Lustre, it would be faster.

Tools Team Updates myNAS for GPU Allocations

- The HECC Tools team released a new version of the myNAS web portal (portal.nas.nasa.gov) for users and principal investigators (PIs) to incorporate GPU allocation and usage information. This version added:
  - Tables showing year-to-date GPU allocations/usage and daily estimated usage by GID for those users and PIs who have GPU allocations. Users without GPU allocations won't see the additional tables.
  - Both GPU and regular HECC allocation/usage data are now downloadable in Excel, PDF, and CSV formats.
- New versions of the myNAS Allocator and Manager Editions were also released, incorporating the GPU allocation/usage information, detailed jobs data for Cabeus, and other improvements.
- The front-end code base for all myNAS editions was consolidated and streamlined to make updates and feature additions more efficient in the future.
- Extensive improvements were also made to the myNAS back-end software, which was updated to work with newer versions of the Portable Batch System job scheduler to provide more detailed job information.
- Most of the back-end scripts and tools were optimized and refactored to be more dynamic, which will reduce the time required to add new GPU types to myNAS.

IMPACT: The myNAS portal provides users, principal investigators, allocators, and managers with a powerful set of tools to analyze and optimize use of HECC resources



This screenshot from the myNAS Principal Investigator Edition shows the percentage of GPU allocation used by each user for a particular GID, names blurred for privacy.

Internet2 Membership Agreement Completed

- The NASA Advanced Supercomputing Division (NAS) forged a membership agreement with the non-profit technology consortium, Internet2, for the agency.
  - NAS staff facilitated an agency-wide agreement with Internet2, allowing NASA centers throughout the agency to take advantage of the advanced capabilities at a lower cost than through individual center membership agreements.
  - The agreement allows for a networking peering arrangement between NAS and Internet2, so users can rely on resilient, low-latency, high-throughput network connectivity between many key academic institutions, research labs, and selected cloud service providers, such as Amazon Web Services (AWS).
  - The agreement, reached in July 2024, allows additional Centers to become members at no additional cost due to Agency-level pricing specified in the agreement: Centers may still seek need to their own Internet2 network peering arrangements.
- The Internet2 advanced capabilities:
  - Enable research, science, and academic enterprise using the fastest, most advanced network for research and education in the U.S.
  - Support many science applications in Earth sciences, high-energy physics, astronomy and astrophysics, life sciences, and engineering by providing trust and identity services, low-cost cloud services connectivity, and advanced networking capabilities.

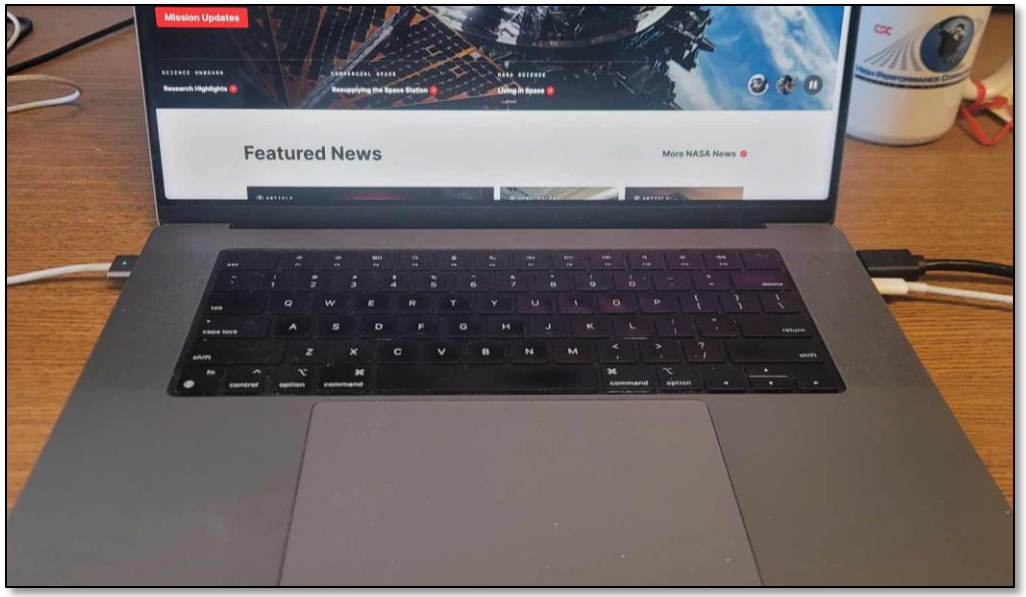
IMPACT: HECC users from academic institutions, research labs, and cloud service providers will benefit from resilient, low-latency, high-throughput network peering and connectivity provided by Internet2. Additionally, users at other NASA centers may benefit from the agreement.

The Internet2 network is designed for the requirements of data-intensive research and supports the computing needs of the research and education community. *Logo courtesy of Internet2*

ESS Team Completes Nasmac Retirement

- The NAS Engineering Systems and Services (ESS) team finished the migration of all NAS users from Apple Mac laptops owned by the NASA Advanced Supercomputing Division (commonly designated as “Nasmacs”) to laptops owned, procured, and supported by the NASA Enterprise Service Desk (ESD) group.
- The new arrangement removes Mac laptops from the NAS Division security plan. This eliminates the need to patch machines based on user interaction. The result is a better security posture and fewer overdue vulnerabilities.
- The change also allows for more flexibility in NAS user choices: users can choose to move to Windows-based laptops if those systems meet their needs better than the standard Mac laptops previously maintained by the ESS team.

IMPACT: The new Mac support arrangement allows for more timely security patching and fewer overdue patches moving forward.

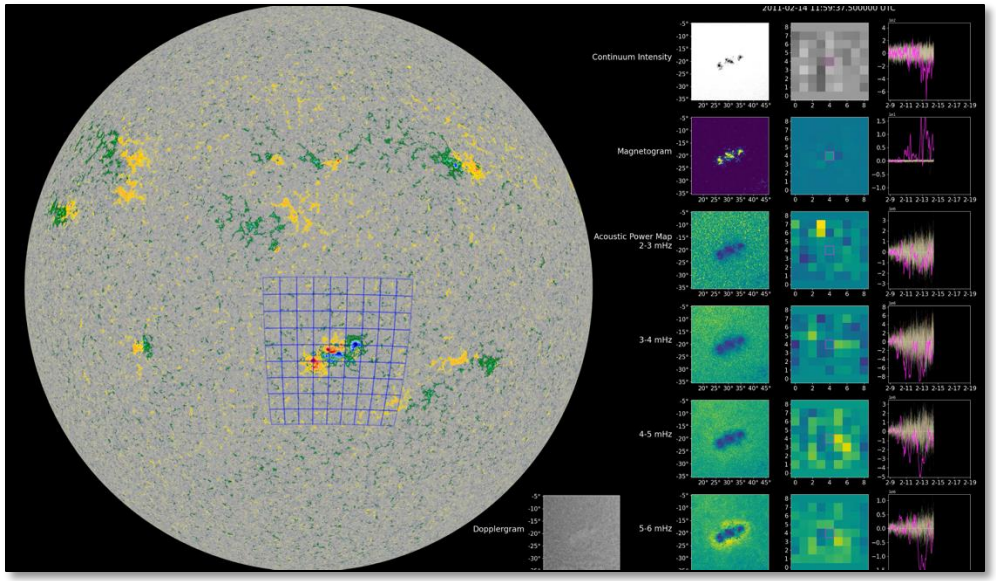


A typical Apple Mac laptop that is now supported by the NASA Enterprise Service Desk group.

Predicting the Appearance of Sunspots Using ML *

- Scientists at NASA Ames are developing machine learning (ML) technologies to enable early warning systems for space storms by analyzing the time evolution of turbulent convection on the Sun's surface. Predicting solar storms is increasingly important to mitigate the impact on space exploration, near-Earth activities, and human health.
- The team employed ML methods to analyze the dynamics of the turbulent motions on the Sun's surface and automatically recognize changes in oscillations associated with an emerging flux from the solar interior, which drives the formation of sunspots.
  - They used observations of magnetic fields, continuum intensity, and dopplergrams obtained from the Helioseismic and Magnetic Imager (HMI) onboard the Solar Dynamics Observatory (SDO)
  - The ML model demonstrated the ability to predict the appearance of an active region 18 hours before it became visible on the solar surface, without false positives.
- HECC's Pleiades supercomputer was used to process the observations from SDO/HMI and run the machine learning models to predict the appearance of sunspots on the solar surface.

IMPACT: This project, made possible by HECC resources, will enable development of new robust machine learning-driven models to predict the formation of sunspots before they become visible on the solar surface.



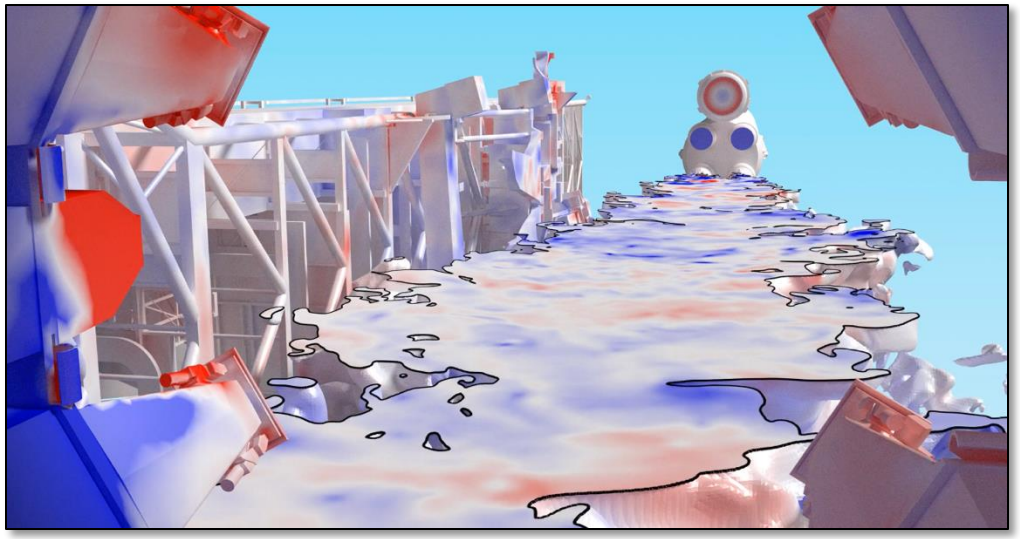
Example of data flow to forecast the emergence of active region AR11158 on the solar surface. On the left, the Solar Dynamics Observatory's Helioseismic and Magnetic Imager (HMI) magnetogram with an overplotted 9 by 9 grid shows area targeted for analysis of different observables (right panels).
Spiridon Kasapis, Nina McCurdy, Timothy Sandstrom, NASA/Ames

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

Launch Pad Rocket Plume Simulations for Artemis*

- To support NASA's Artemis missions launching aboard the Space Launch System, the Launch, Ascent, and Vehicle Aerodynamics (LAVA) team at NASA Ames pushed the state of the art for predicting exhaust plume impingement in launch environments using HECC resources.
- Exhaust plumes pose a safety hazard by creating unwanted forces and vibrations that may damage the launch pad and support structures. During the Artemis I launch in 2022, the doors to the launch tower's elevators were damaged by pressure generated by these plumes.
- The LAVA team performed twelve simulations, focusing on six different instants in time after ignition along two different trajectories. The first trajectory followed the path taken by the Artemis launch in November 2022; the second used a simulated trajectory known as Artemis II Max Drift—which is a potential worst-case scenario where the vehicle drifts very close to the launch tower.
- Each high-order accurate LAVA Cartesian plume simulation typically used 150 AMD Rome nodes for 1-3 days on the Aitken supercomputer to simulate four seconds of real-world time, accumulating roughly 100 terabytes of volume and surface data.
- Launch pad engineers at Kennedy Space Center are actively hardening the launcher tower in preparation for future launches. These simulations are helping to identify and quantify plume impingement loads, which will improve understanding and safety for future launches.

IMPACT: Fluid dynamics simulations of plume impingement on the launch pad and tower at various stages of liftoff are useful for understanding the complex flows that may damage the launch environment and endanger astronauts during future missions.



This image shows a slice through the exhaust plumes colored by pressure (red is high, blue is low) of the Artemis I launch, using LAVA. Black line indicates the 90% plume exhaust/air interface. The view is from the exhaust hole on the mobile launcher, looking straight up at the vehicle. *Timothy Sandstrom, Michael Barad, NASA/Ames*

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

Papers

- "**TESS Discovery of Two Super-Earths Orbiting the M-Dwarf Stars TOI-6002 and TOI-5713 Near the Radius Valley**," M. Ghachoui, et al., arXiv:2408.00709 [astro-ph.EP], August 1, 2024. *
  https://arxiv.org/abs/2408.00709
- "**The Orbiting Carbon Observatory-2 (OCO-2) and In-Situ CO2 Data Suggest a Larger Seasonal Amplitude of the Terrestrial Carbon Cycle Compared to Many Dynamic Global Vegetation Models**," R. Lei, et al., *Remote Sensing of Environment*, vol. 312, published online August 3, 2024. *
  https://www.sciencedirect.com/science/article/abs/pii/S0034425724003444
- "**Electrostatic Waves and Electron Holes in Simulations of Low-Mach Quasi-Perpendicular Shocks**," A. Bohdan, et al., arXiv:2408.01699 [physics.space-ph], August 3, 2024. *
  https://arxiv.org/abs/2408.01699
- "**Web-Based Visualization and Analytics of Petascale Data: Equity as a Tide that Lifts all Boats**," A. Panta, et al., arXiv:2408.11831 [cs.HC], August 6, 2024. *
  https://arxiv.org/abs/2408.11831
- "**Evolution of BD-14 3065b (TOI-4987b) from Giant Planet to Brown Dwarf as Possible Evidence of Deuterium Burning at Old Stellar Ages**," J. Subjak, et al., *Astronomy & Astrophysics*, vol. 688, published online August 9, 2024. *
  https://www.aanda.org/articles/aa/abs/2024/08/aa49028-23/aa49028-23.html
- "**Experimental Determination of the Unusual CH Stretch Frequency of Protonated Fullerenes**," L. Finazzi, et al., *The Astrophysical Journal*, vol. 971, no. 2, published online August 16, 2024. *
  https://iopscience.iop.org/article/10.3847/1538-4357/ad6302/meta

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

Papers (cont.)

- "**Single-Star Warm-Jupiter Systems Tend to be Aligned, Even Around Hot Stellar Hosts: No T_{eff}–λ Dependency**," X.-Y. Wang, et al., arXiv:2408.10038 [astro-ph.EP], August 19, 2024. *
  https://arxiv.org/abs/2408.10038
- "**YSO Jets Magnetocentrifugally Driven by Reconnecting Atmospheric Avalanche Accretion Streams Above Inner Circumstellar Disks**," Y. Tu, et al., arXiv:2408.10938 [astro-ph.SR], August 20, 2024. *
  https://arxiv.org/abs/2408.10938
- "**TOI-2490b – The Most Eccentric Brown Dwarf Transiting in the Brown Dwarf Desert**," B. Henderson, et al., *Monthly Notices of the Royal Astronomical Society*, vol. 533, issue 3, published online August 20, 2024. *
  https://academic.oup.com/mnras/article/533/3/2823/7737668
- "**Effects of Alloying Elements on Twinning in Ni-Based Superalloys**," V. Borovikov, et al., *Superalloys 2024* (conference paper), published online August 21, 2024. *
  https://link.springer.com/chapter/10.1007/978-3-031-63937-1_97
- "**Generation of Top-Boundary Conditions for 3D Ionospheric Models Constrained by Auroral Imagery and Plasma Flow Data**," J. van Irsel, et al., *Journal of Space Physics*, vol. 129, issue 8, published online August 21, 2024. *
  https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2024JA032722
- "**On the Formation and Detectability of H_2CNCN and Its Progenitors**," R. Fortenberry, V. Esposito, *The Astrophysical Journal*, vol. 972, no. 1, published online August 23, 2024. *
  https://iopscience.iop.org/article/10.3847/1538-4357/ad685c/meta

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

Papers (cont.)

- "**High-Fidelity Numerical Simulation of Longitudinal Thermoacoustic Instability in a High-Pressure Subscale Rocket Combustor**," V. R. Hasti, R. Ranjan, *Aerospace Science and Technology*, vol. 154, published online August 26, 2024. * https://www.sciencedirect.com/science/article/abs/pii/S1270963824006187
- "**Carbon Emissions from the 2023 Canadian Wildfires**," B. Byrne, et al., *Nature*, published online August 28, 2024. * https://www.nature.com/articles/s41586-024-07878-z
- "**TOI-2379 b & TOI-2384 b: Two Super-Jupiter Mass Planets Transiting Low-Mass Host Stars**," E. Bryant, et al., *Monthly Notices of the Royal Astronomical Society*, published online August 29, 2024. * https://academic.oup.com/mnras/advance-article/doi/10.1093/mnras/stae2034/7745505

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

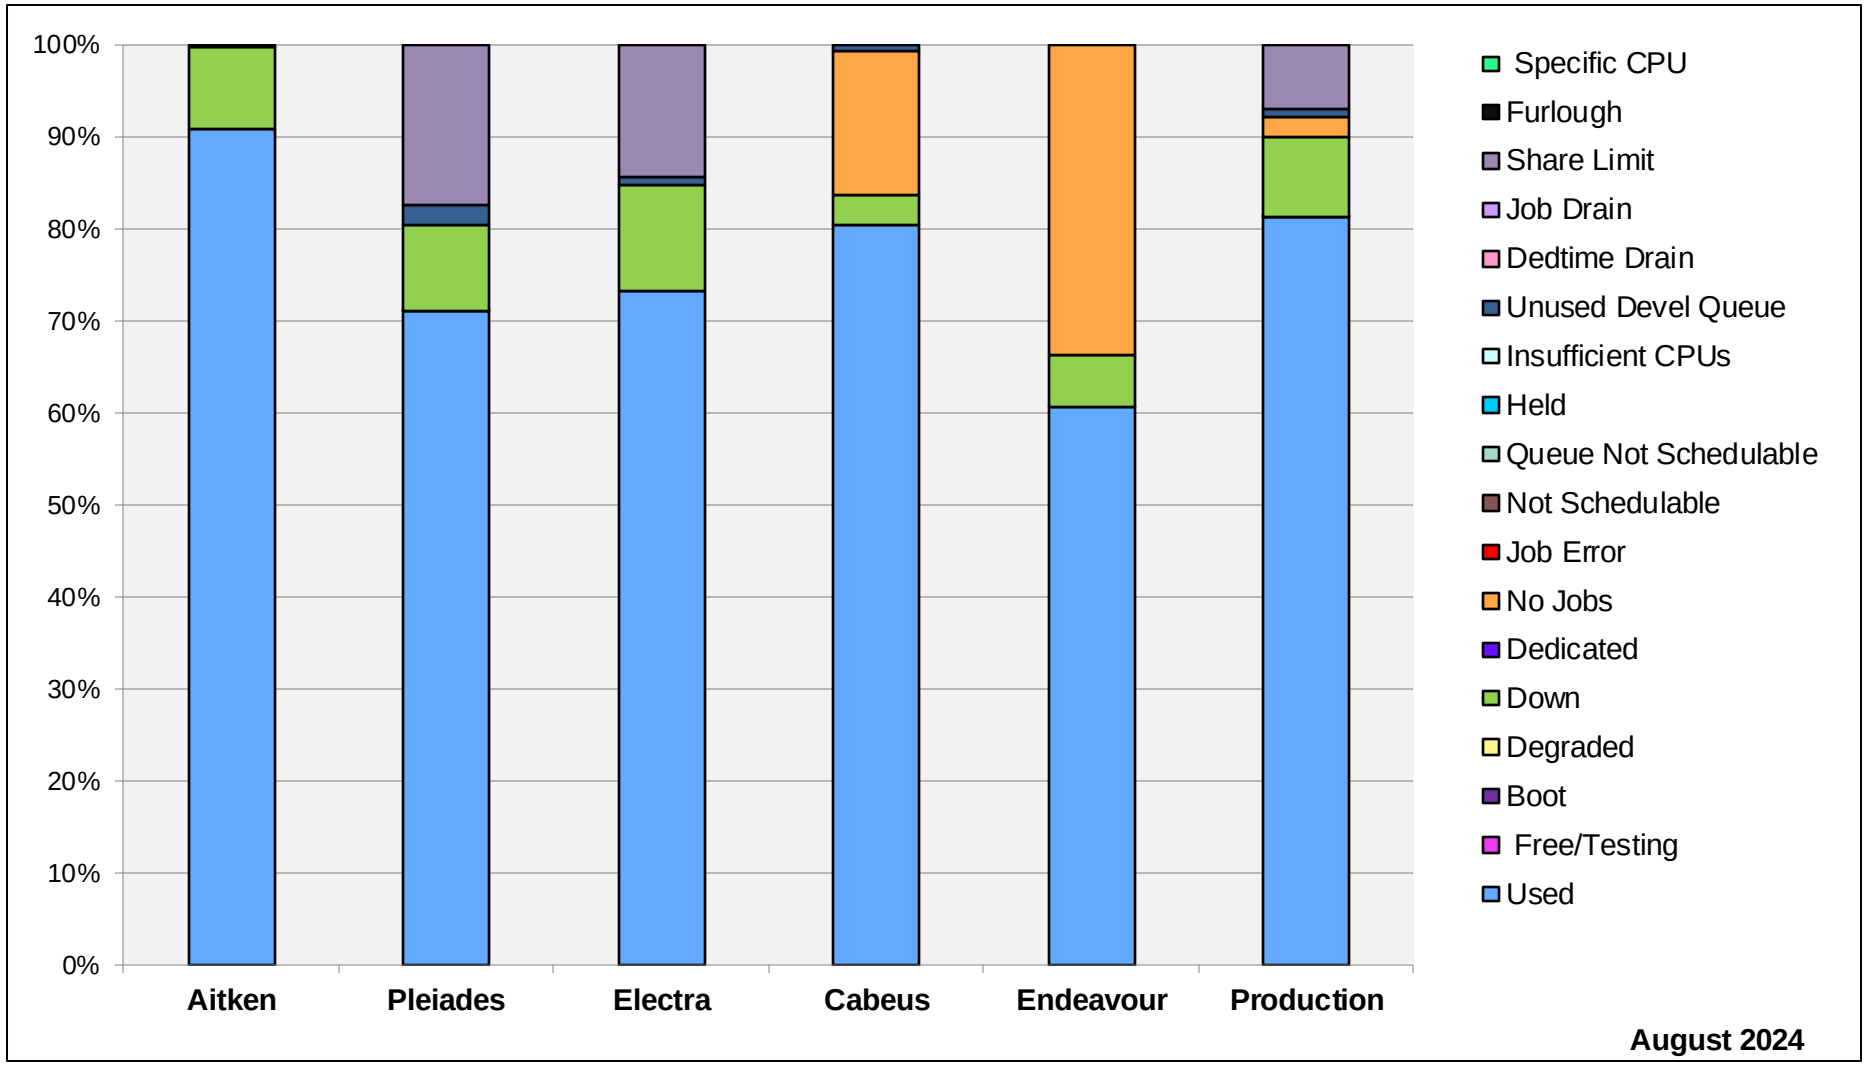
News & Events

- **NASA Ames to Host Supercomputing Resources for UC Berkeley Researchers**, *NASA*, August 2, 2024—Under a new agreement, NASA will host supercomputing resources for the University of California, Berkeley at the NASA Advanced Supercomputing facility – the agency's premiere supercomputing center. The agreement is part of an expanding partnership between Ames and UC Berkeley and will support the development of novel computing algorithms and software for a wide variety of scientific and technology a reas. *https://www.nasa.gov/general/nasa-ames-to-host-supercomputing-resources-for-uc-berkeley-researchers/*
  - **NASA Ames to host supercomputing resources for UC Berkeley**, *Scientific Computing World* *https://www.scientific-computing.com/article/nasa-ames-host-supercomputing-resources-uc-berkeley*

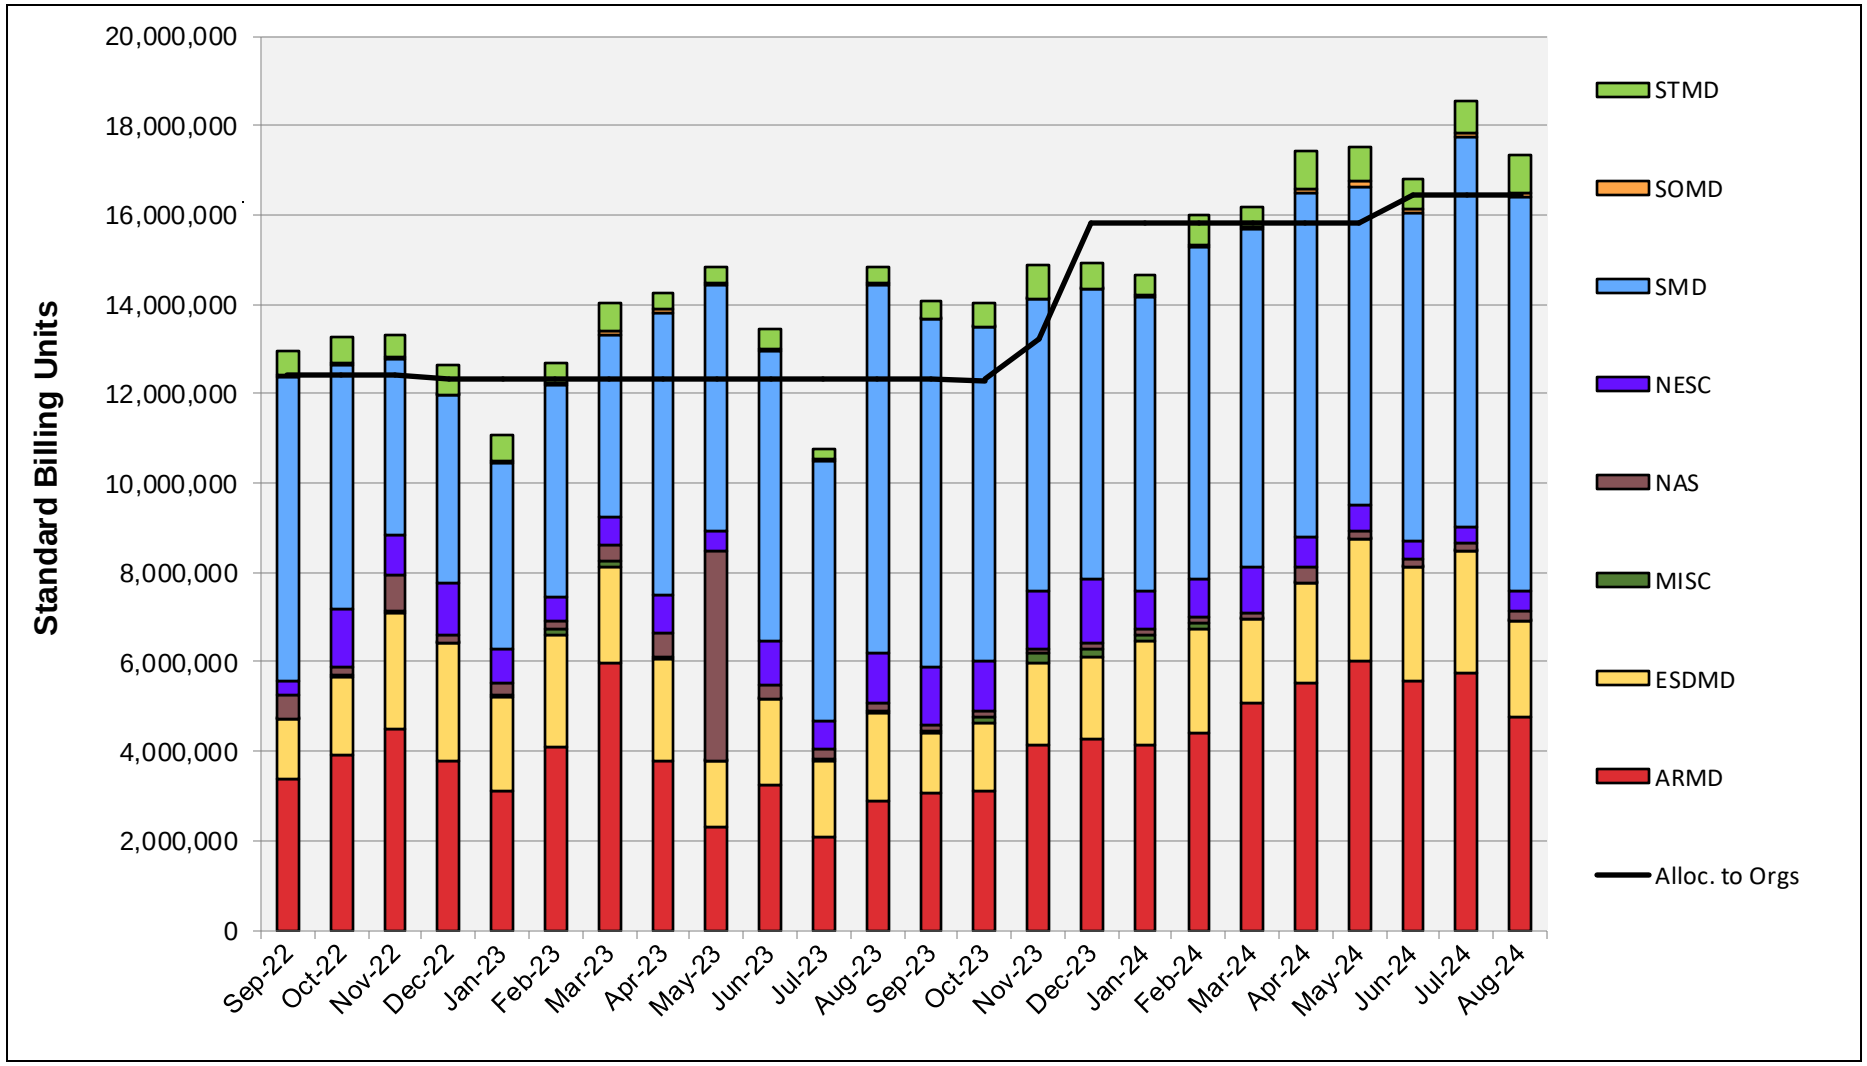
Social Media Coverage

- UC Berkeley – NASA Partnership
  - NASA Ames: X 47 reposts, 202 likes; Facebook 151 likes, 26 shares.
- Solar Magnetic Field Research
  - NASA Advanced Supercomputing Division: X 2 repost, 1 like.

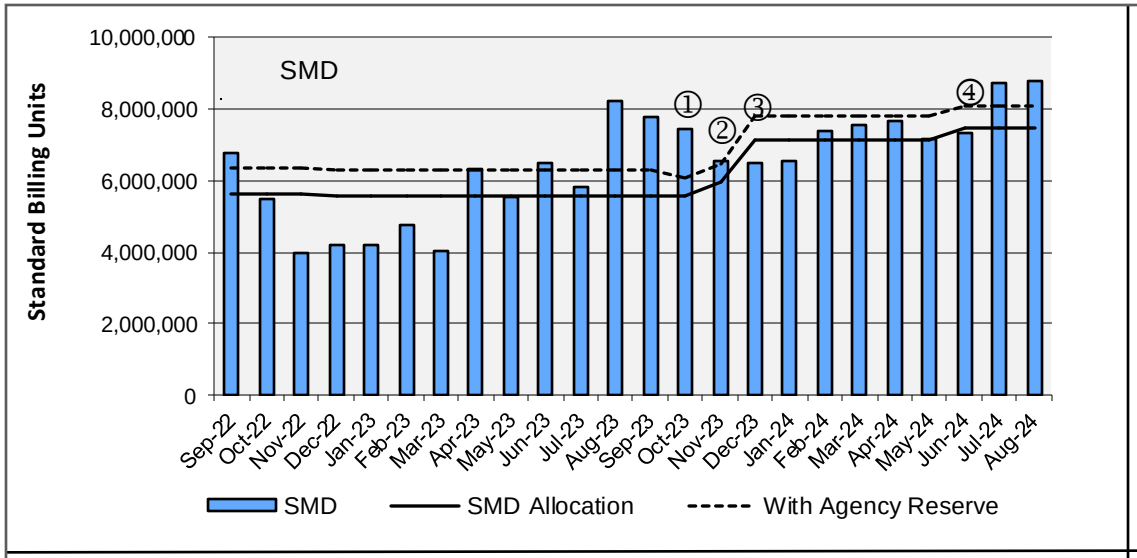
HECC Utilization

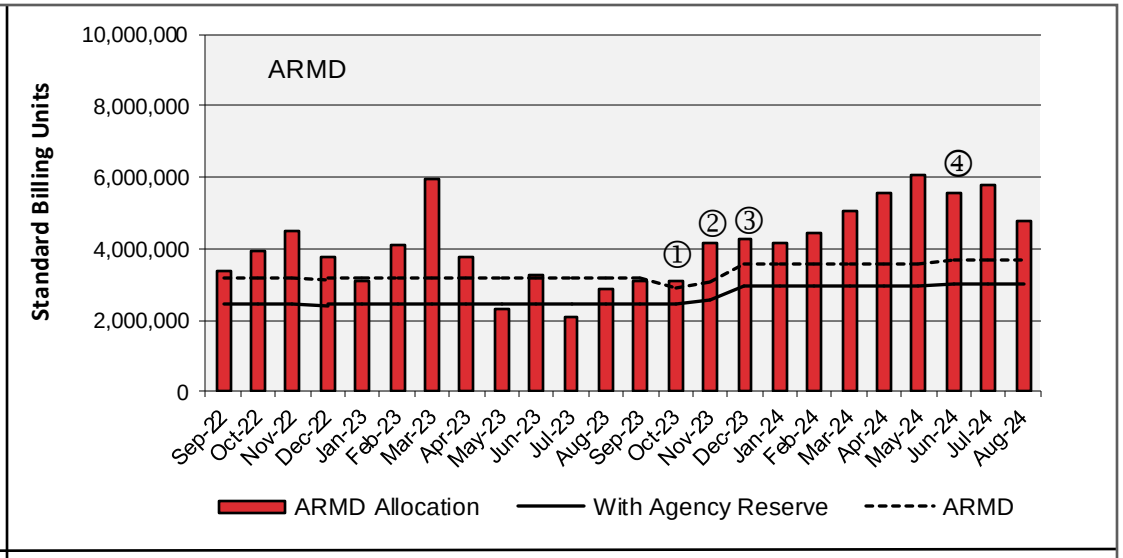


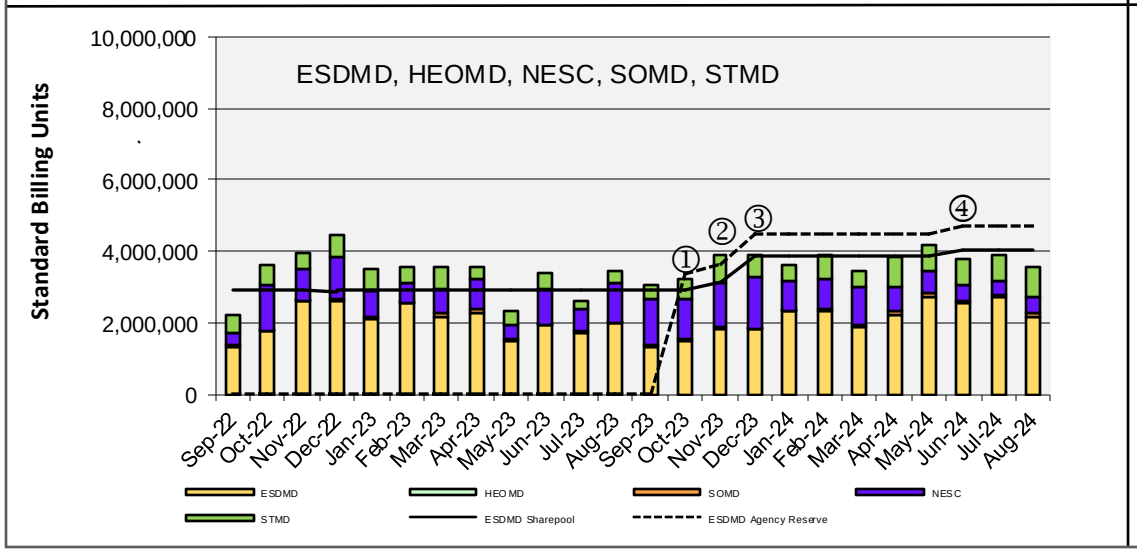
HECC Utilization Normalized to 30-Day Month



HECC Utilization Normalized to 30-Day Month







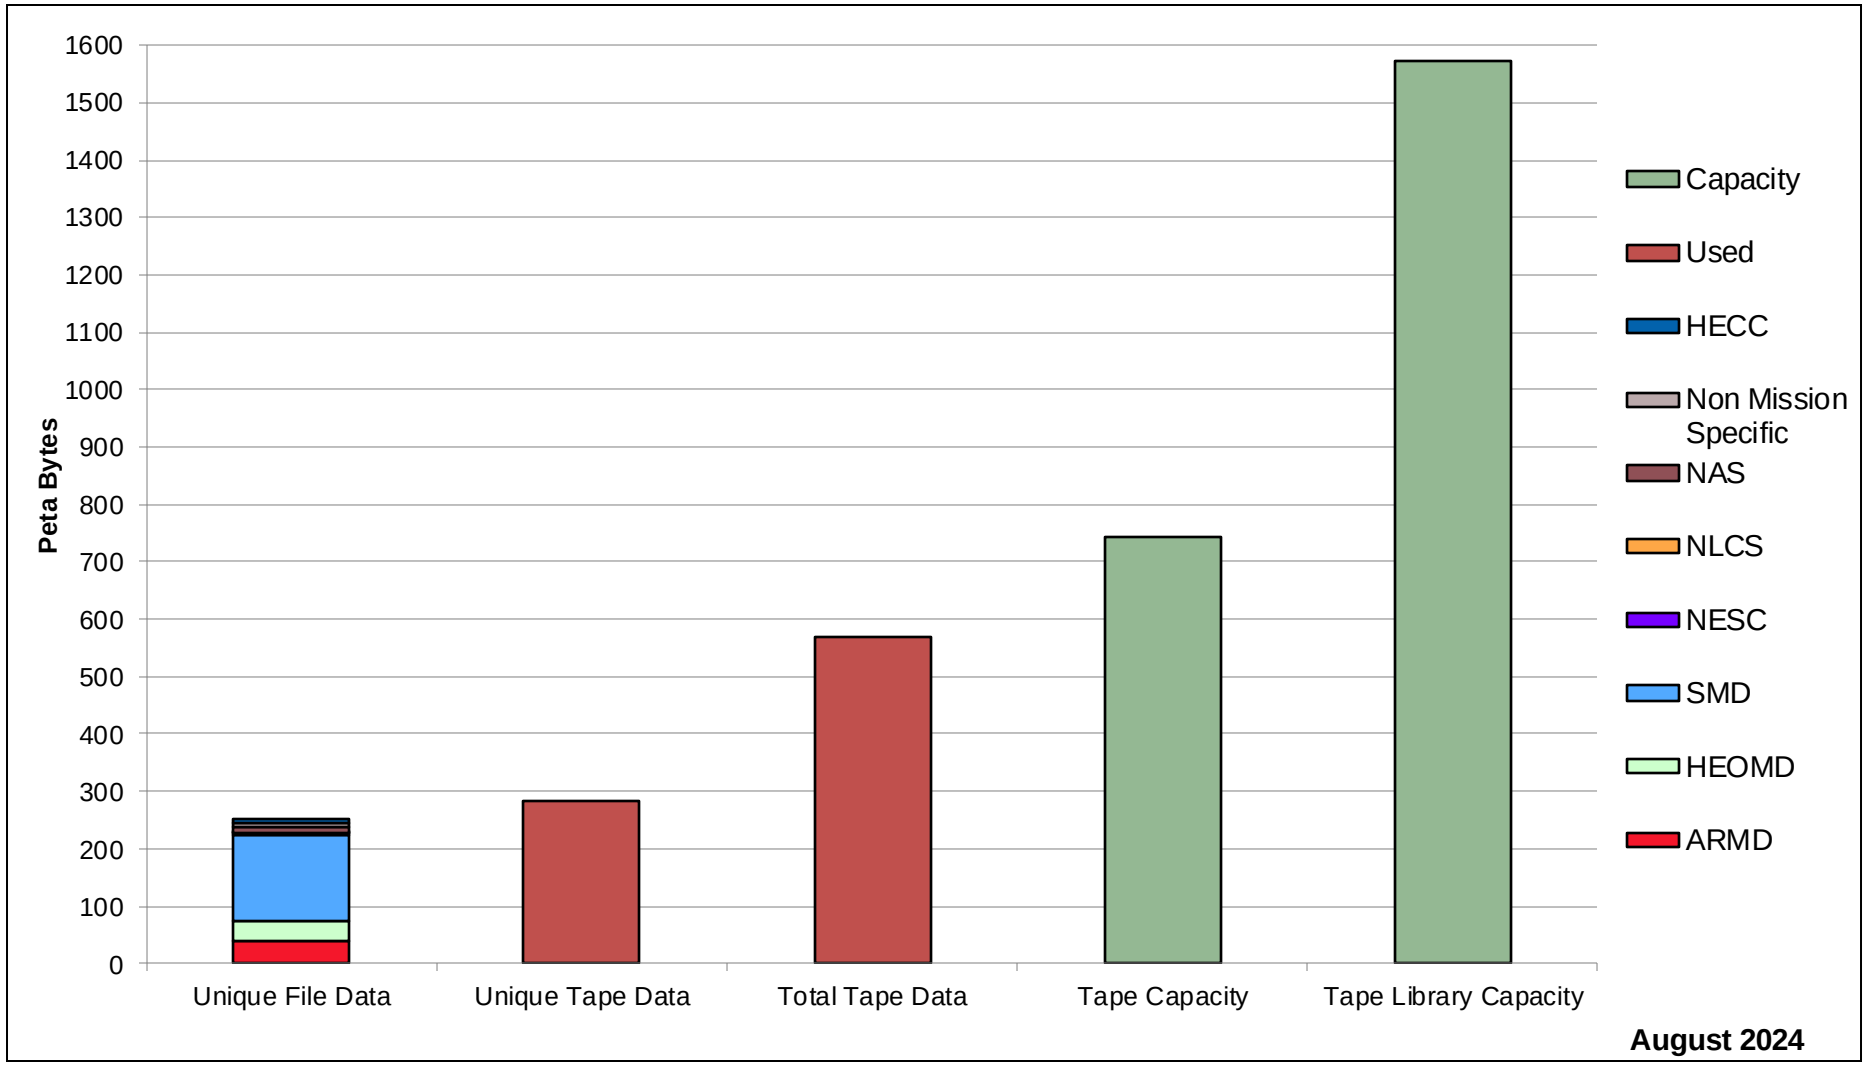
① 4 Rome Apollo Racks added to Aitken

② Cabeus v100 nodes added

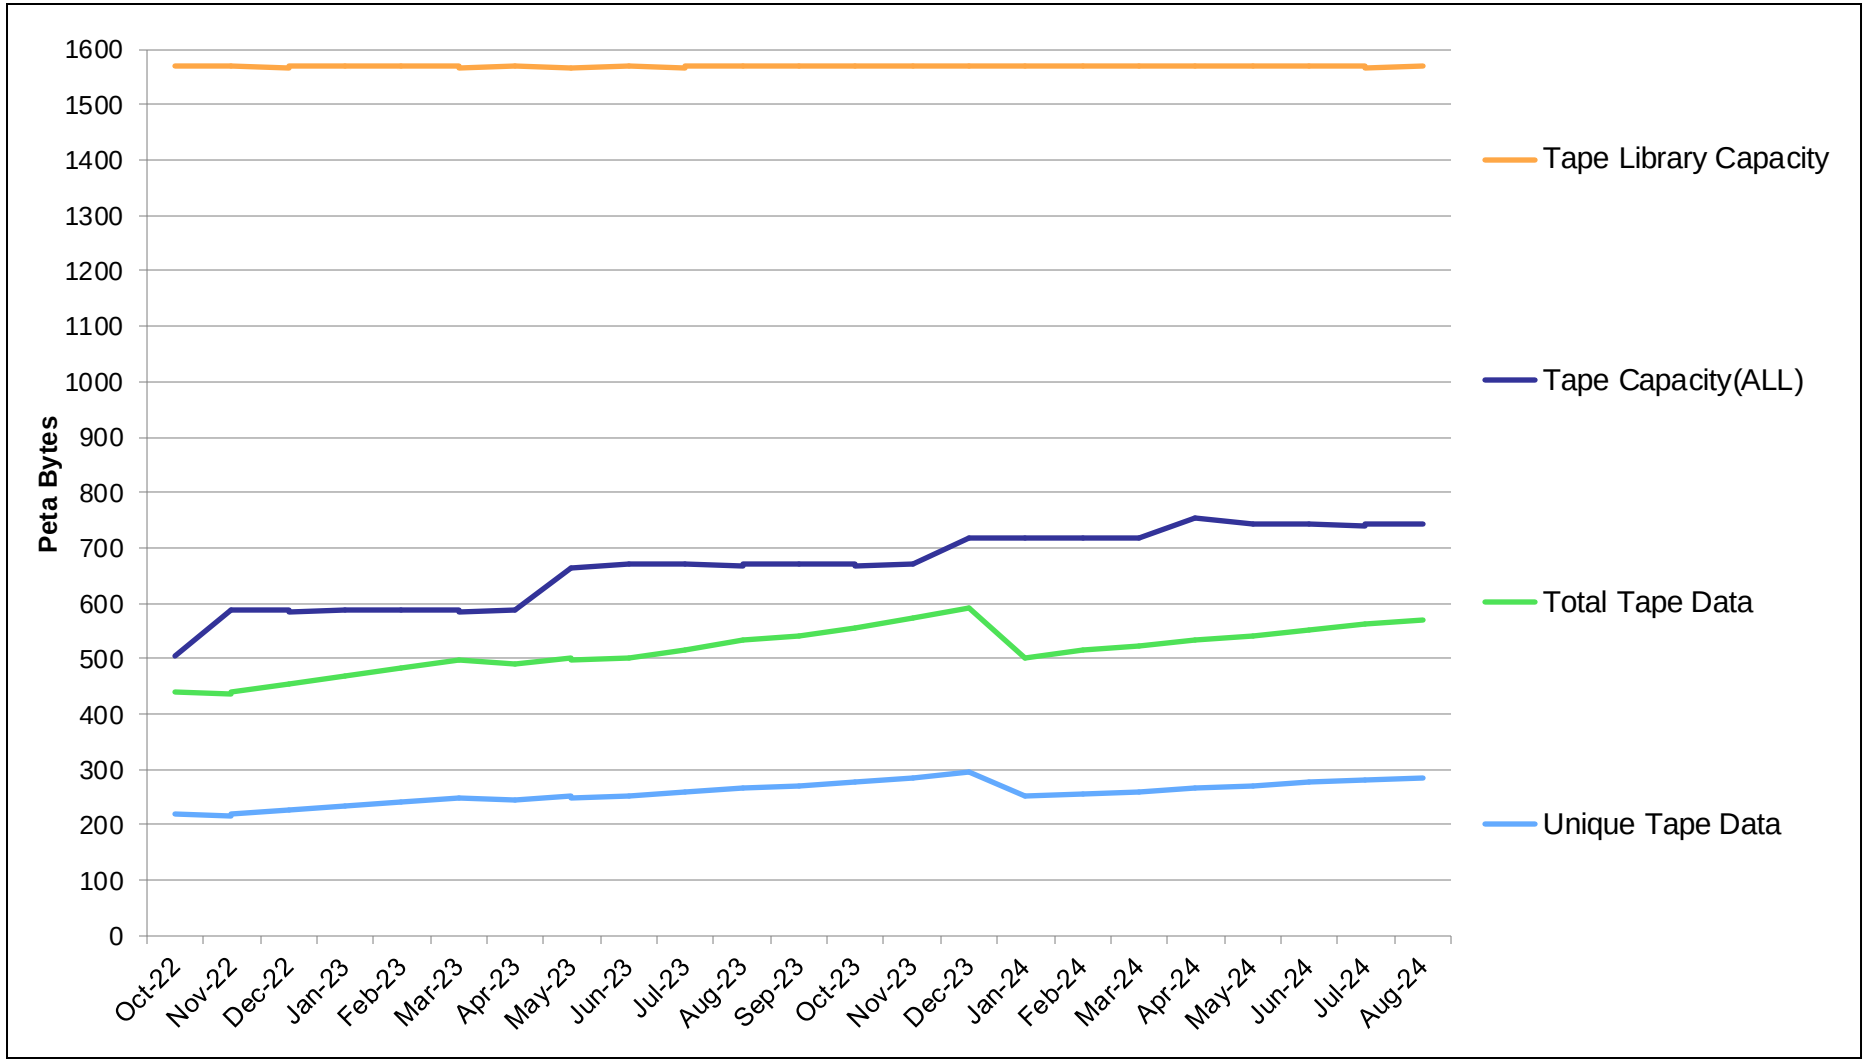
③ Cabeus Milan a100 nodes added

④ Aitken Milan Nodes added

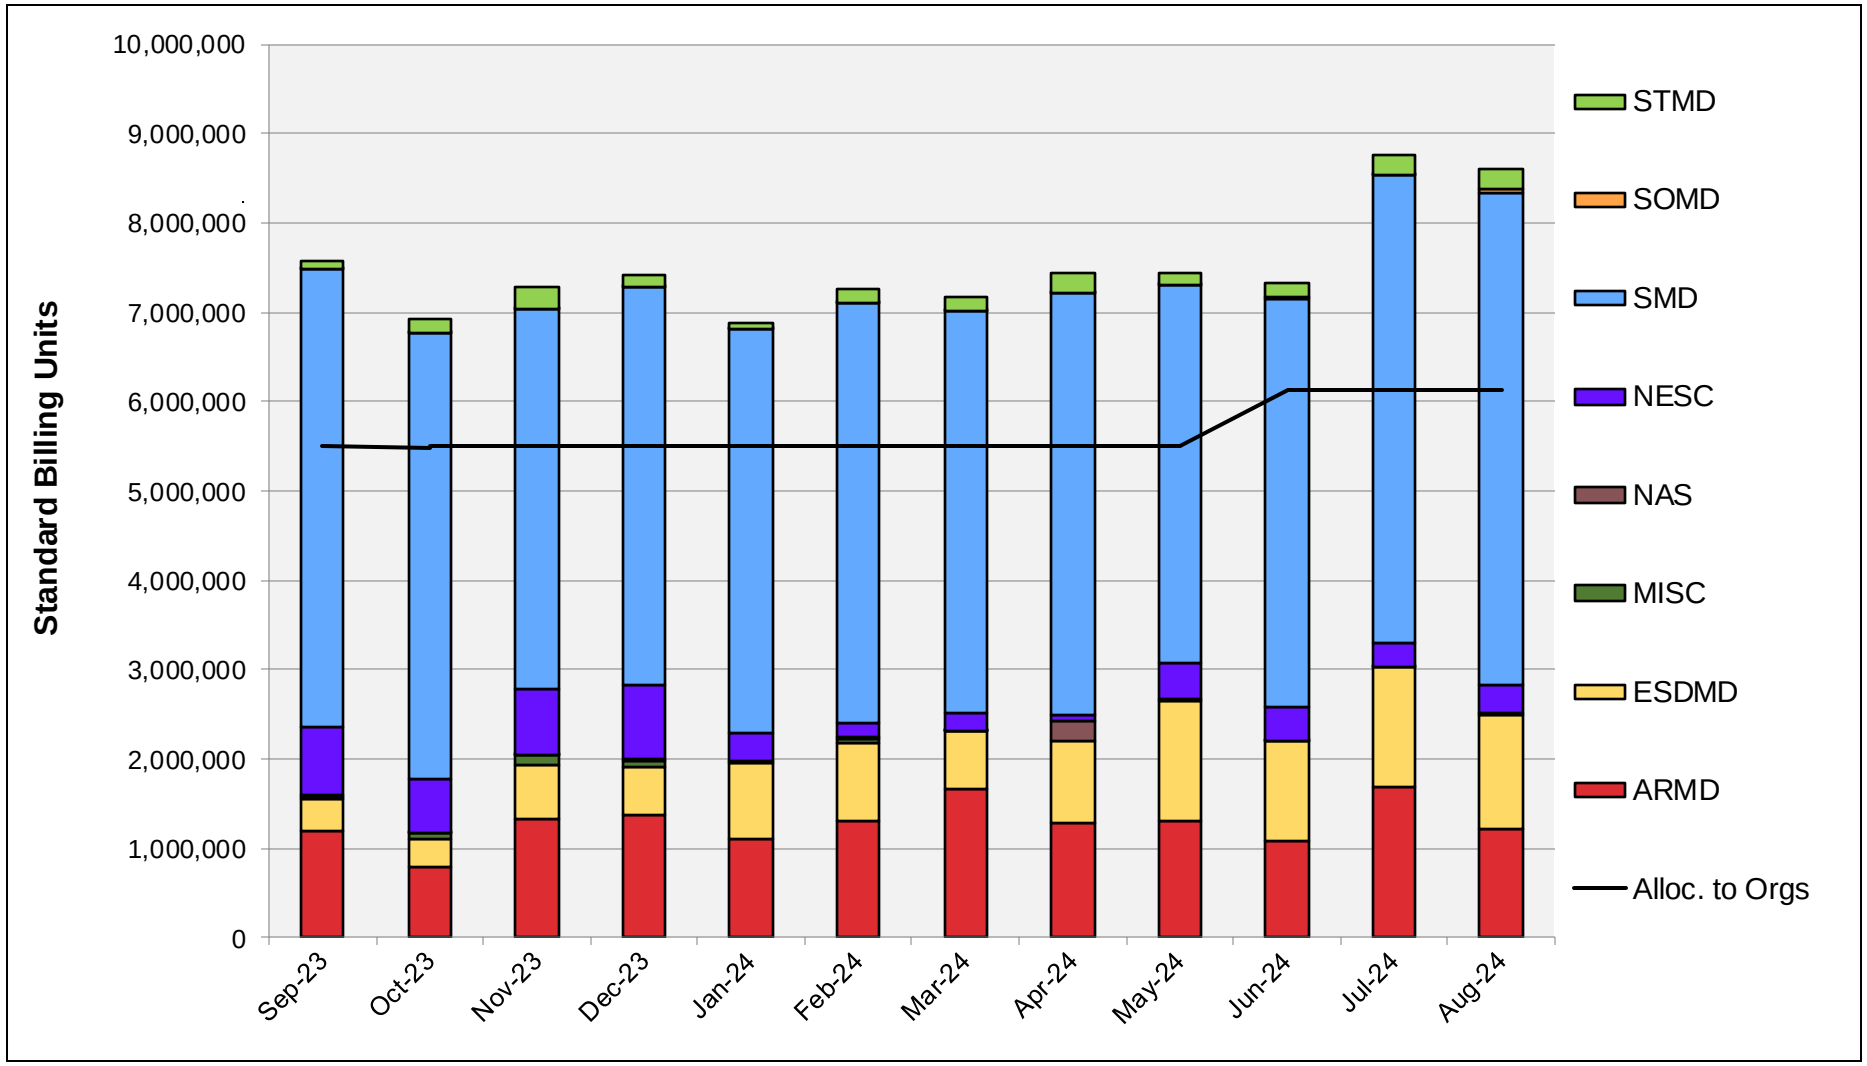
Tape Archive Status



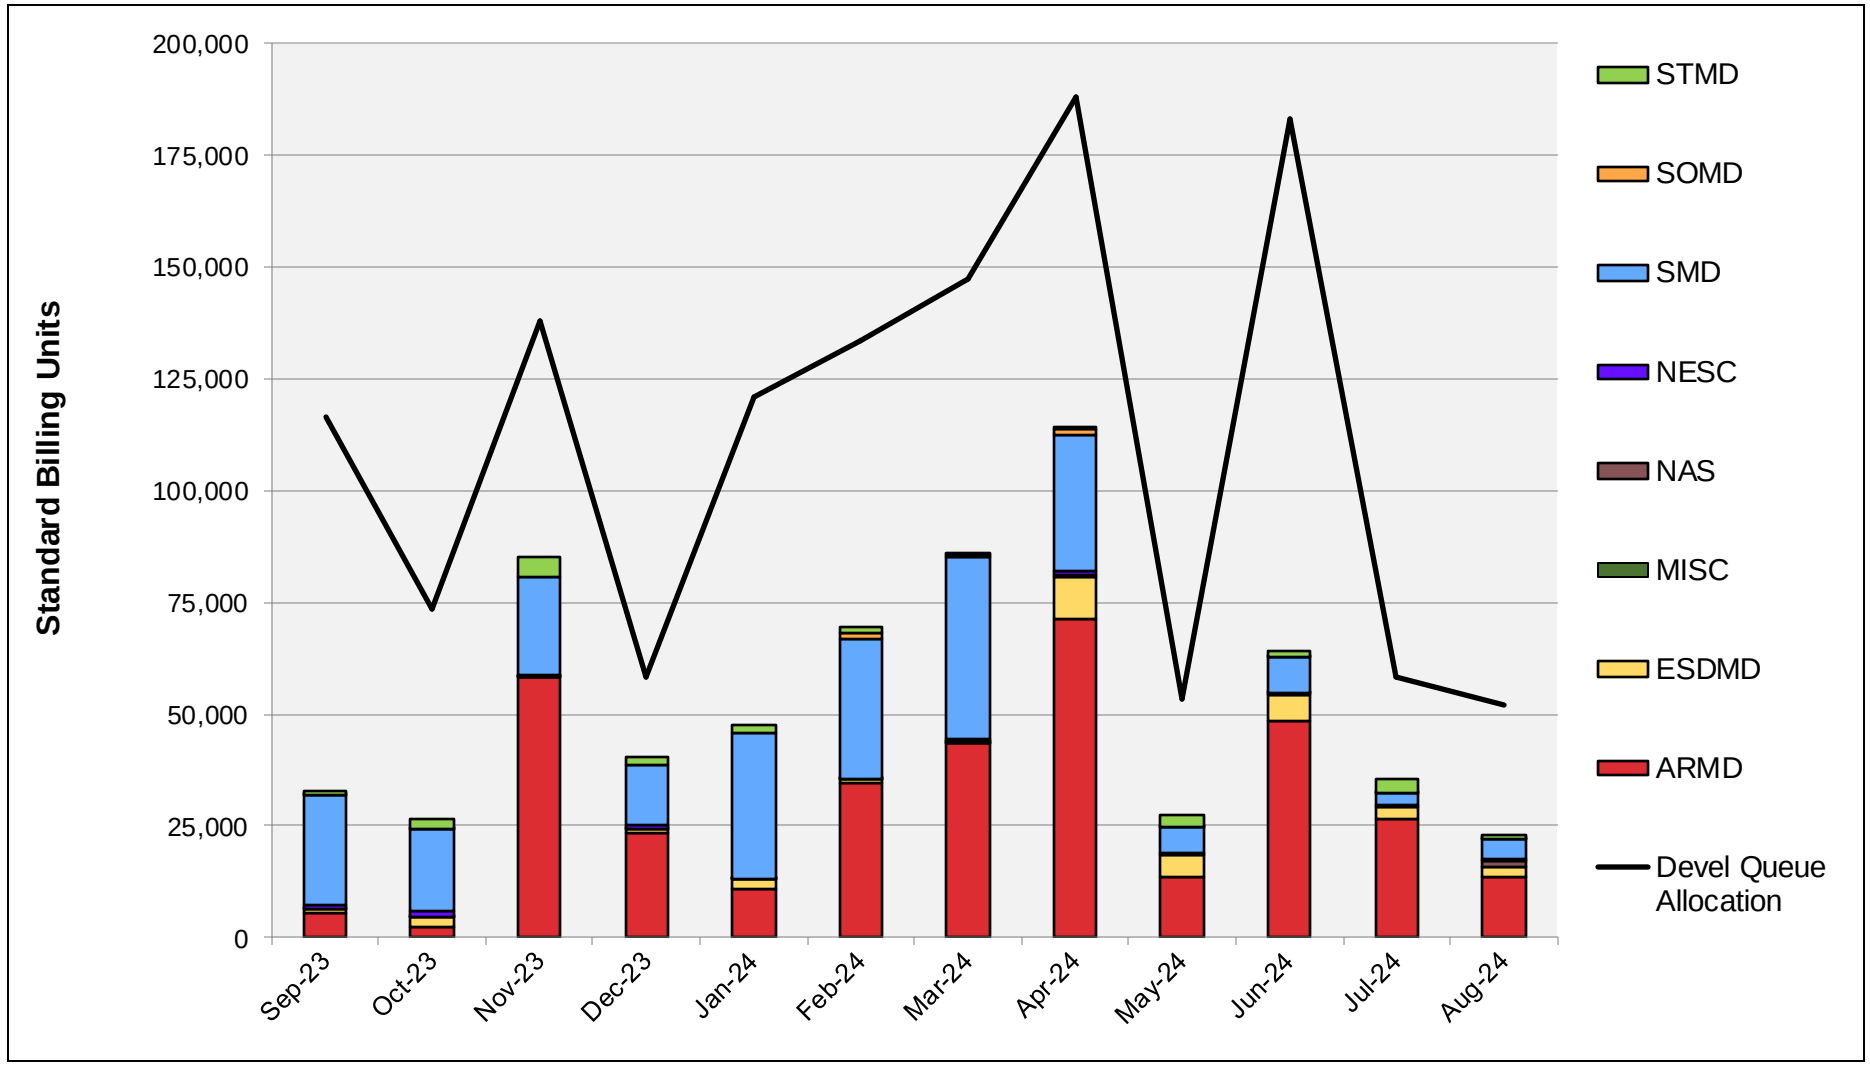
Tape Archive Status



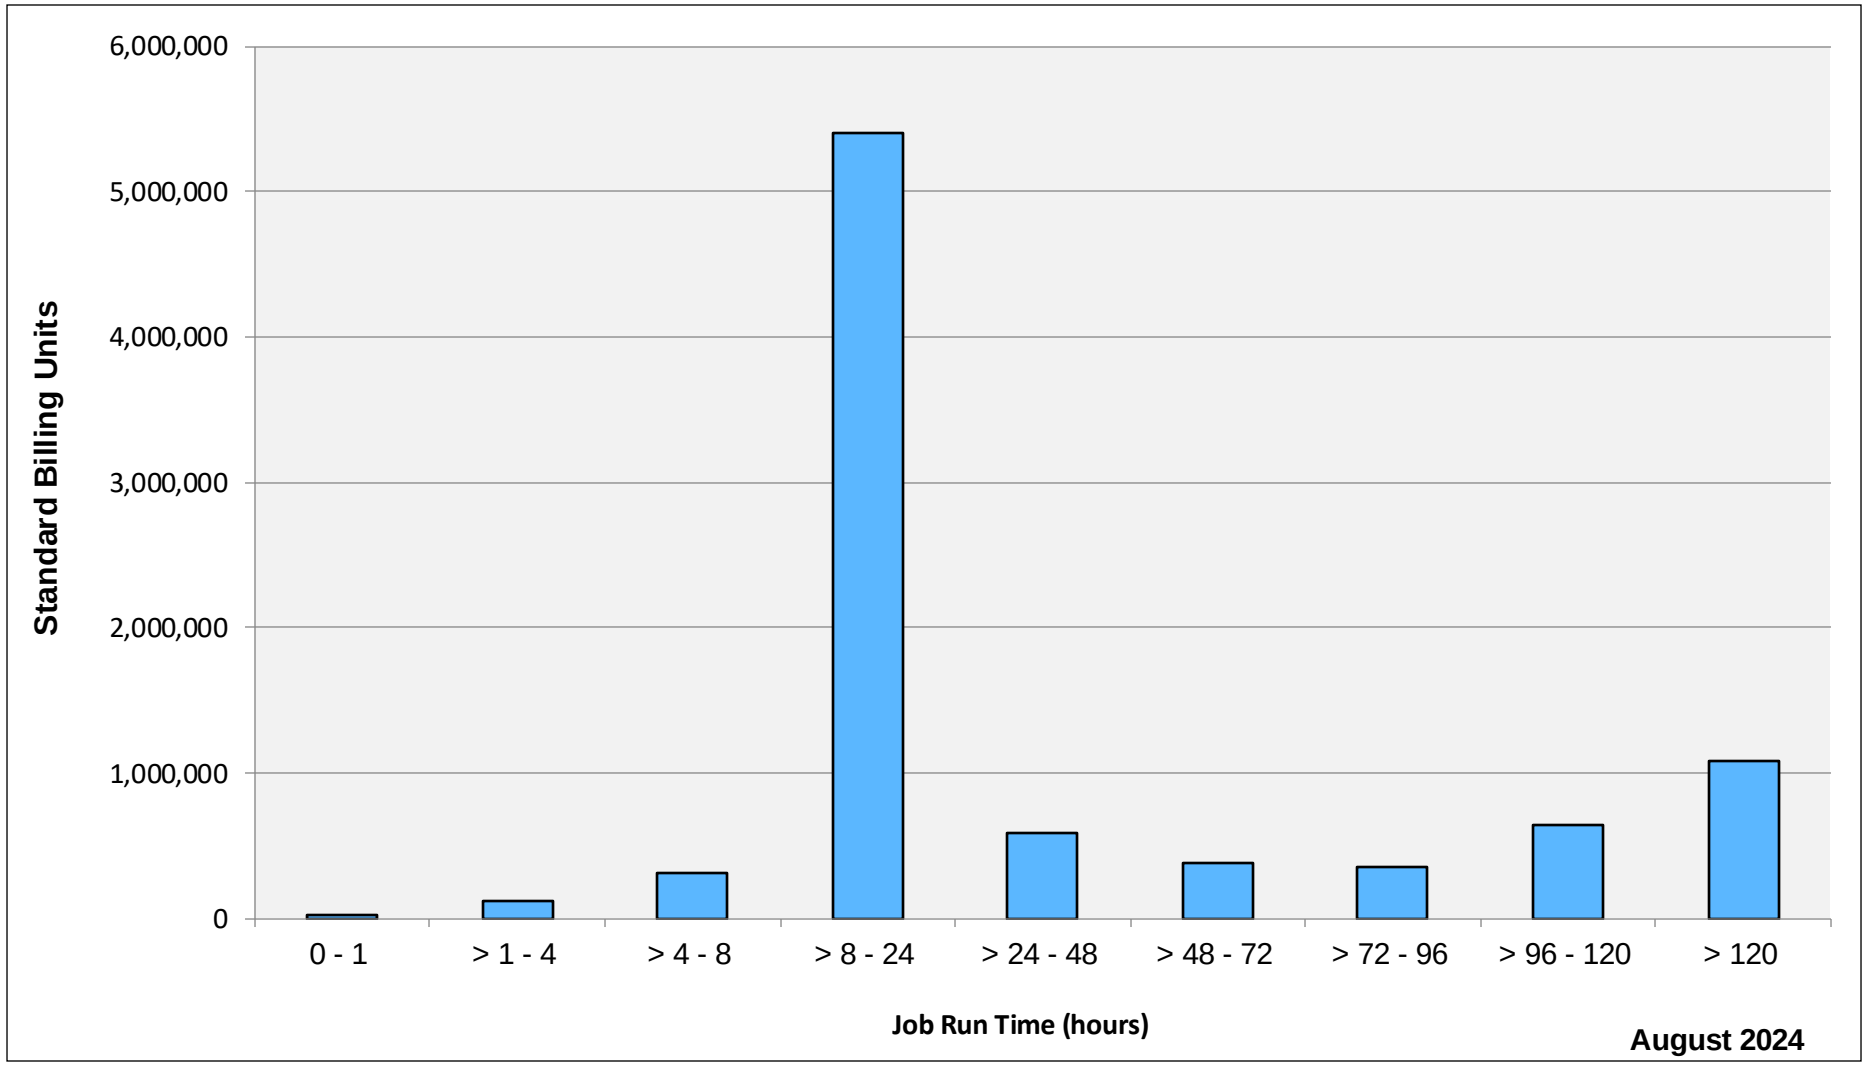
Aitken: SBUs Reported, Normalized to 30-Day Month



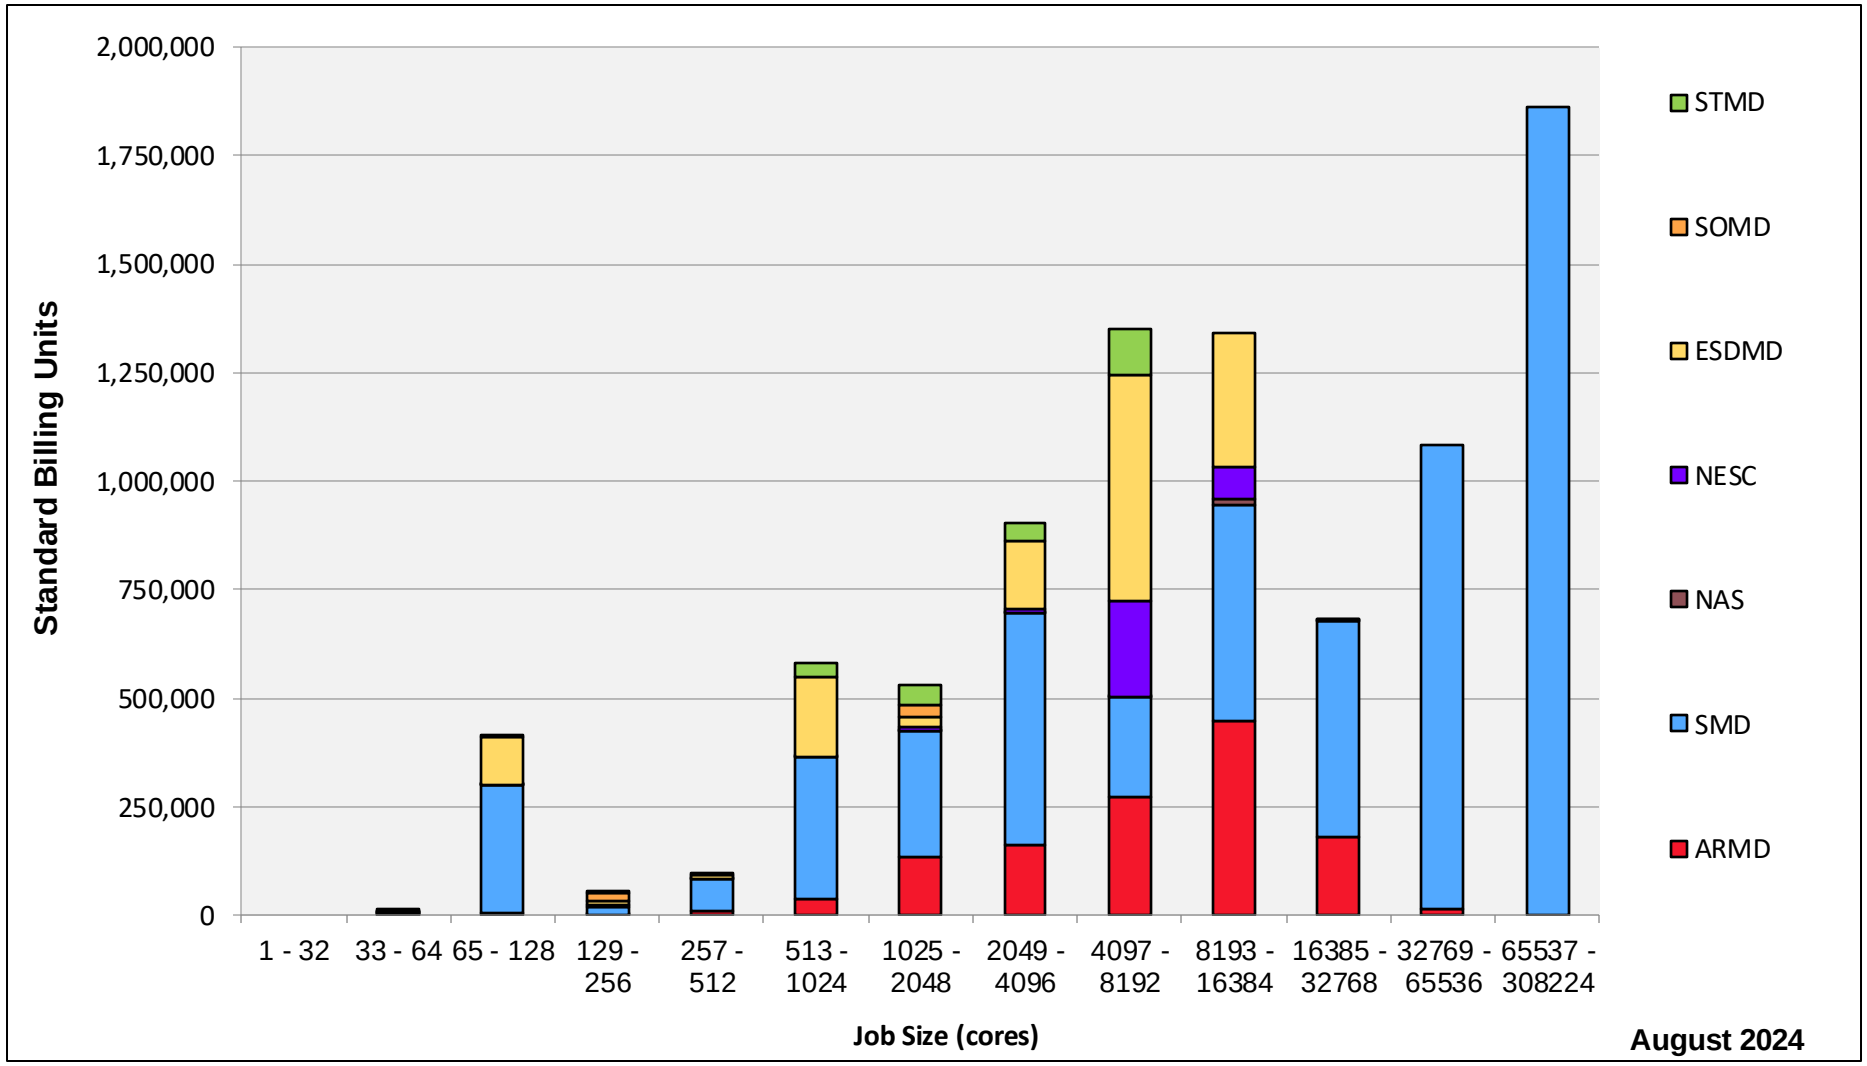
Aitken: Devel Queue Utilization



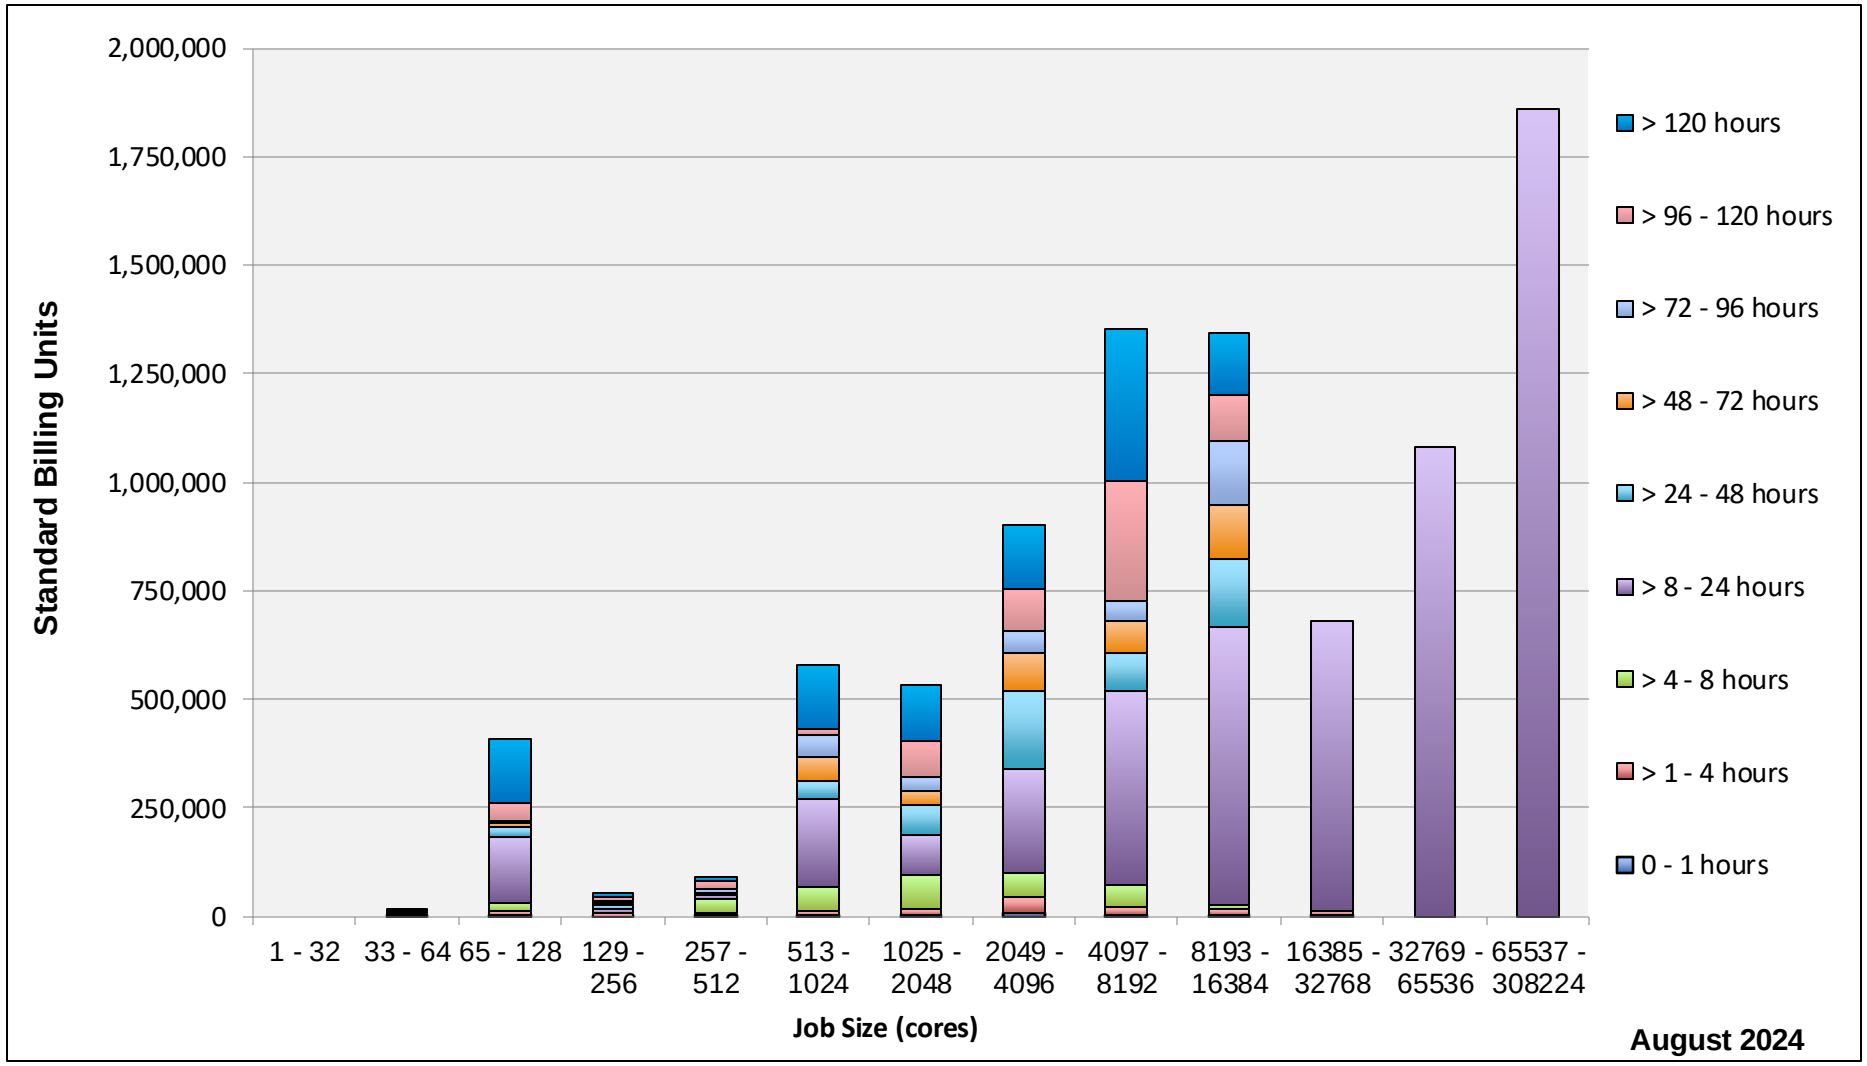
Aitken: Monthly Utilization by Job Length



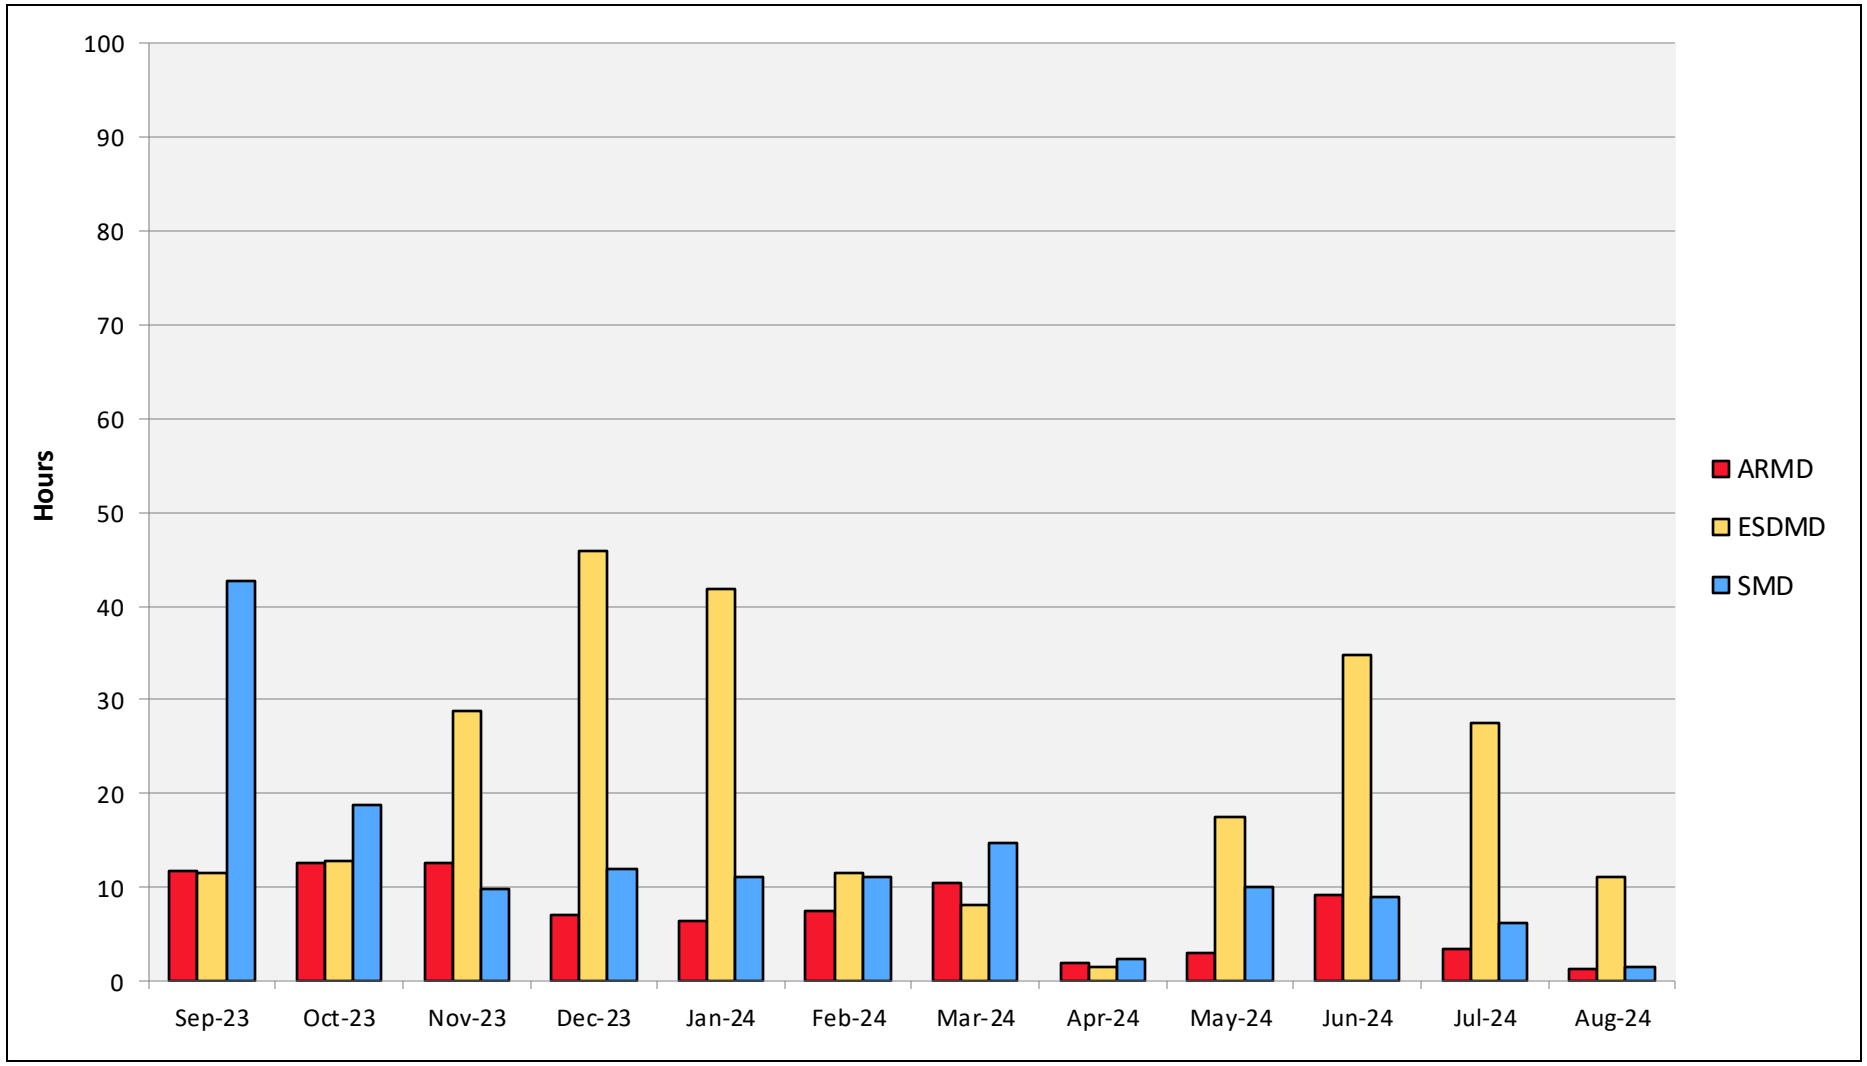
Aitken: Monthly Utilization by Job Size



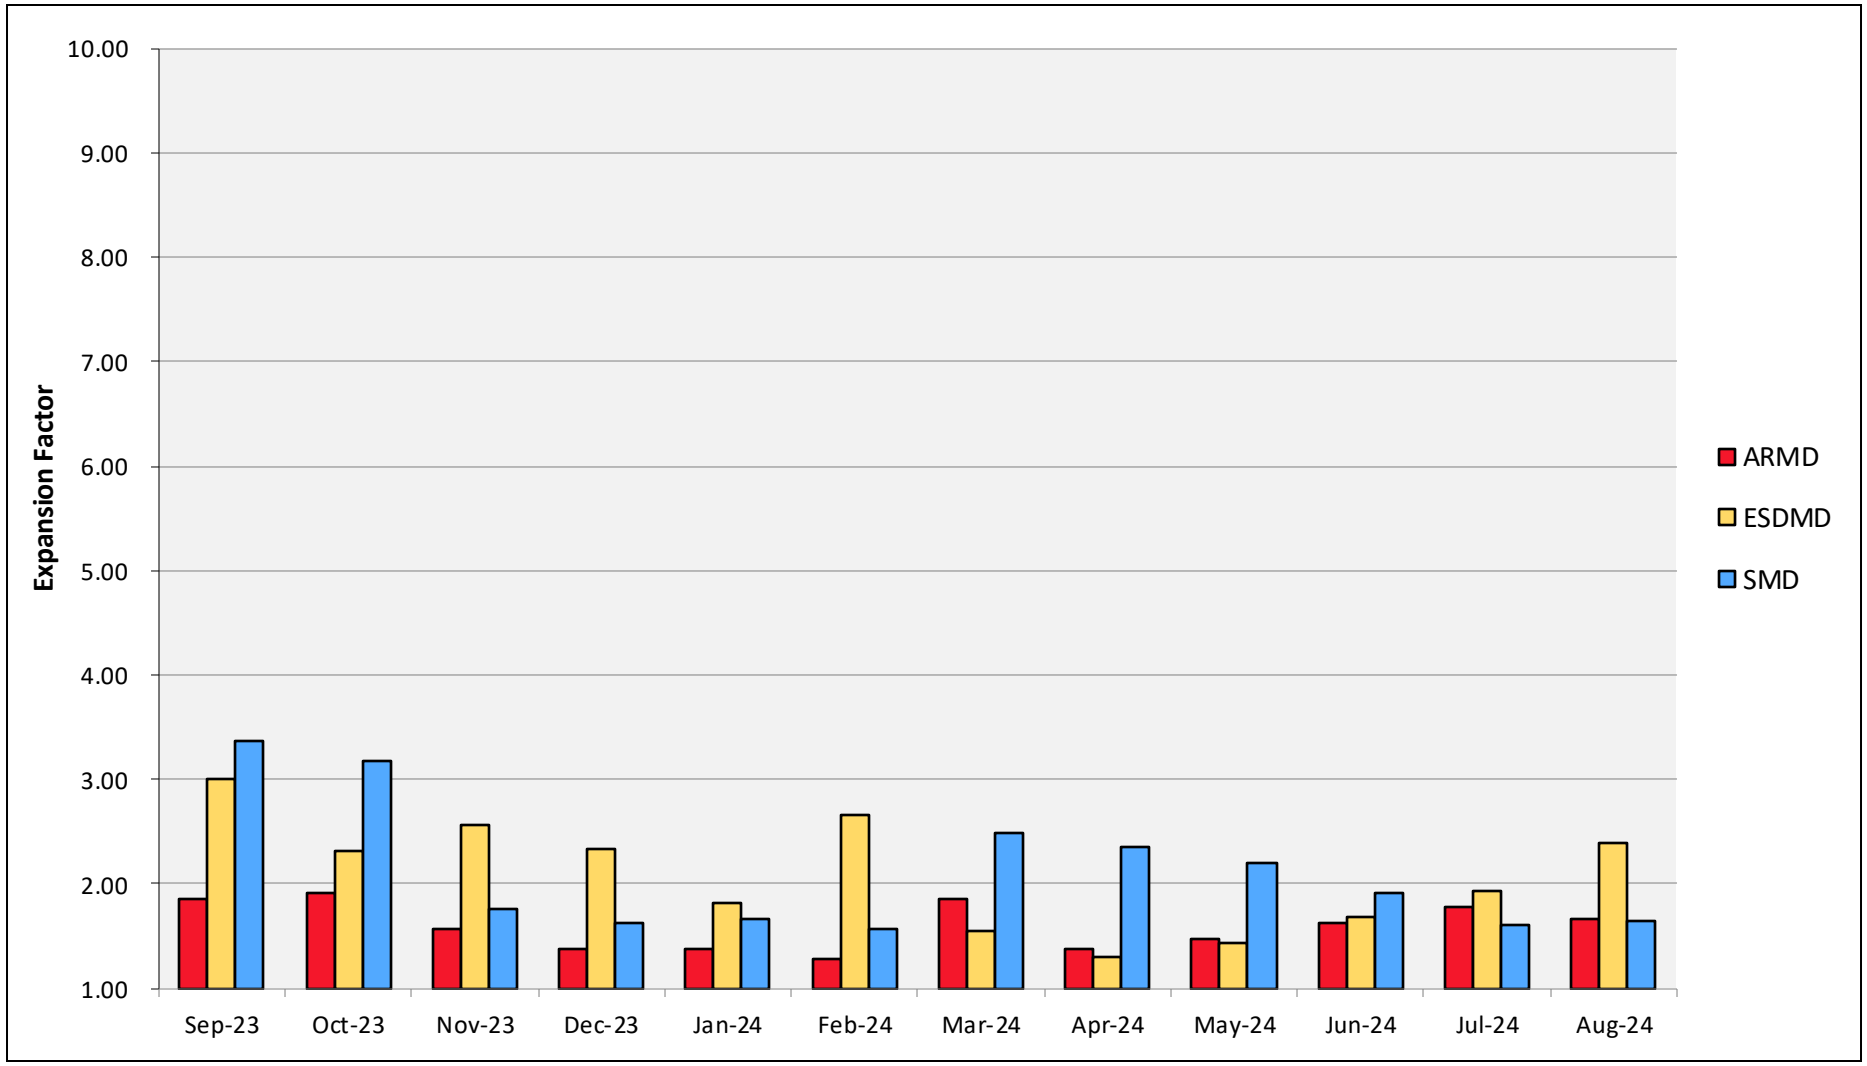
Aitken: Monthly Utilization by Size and Length



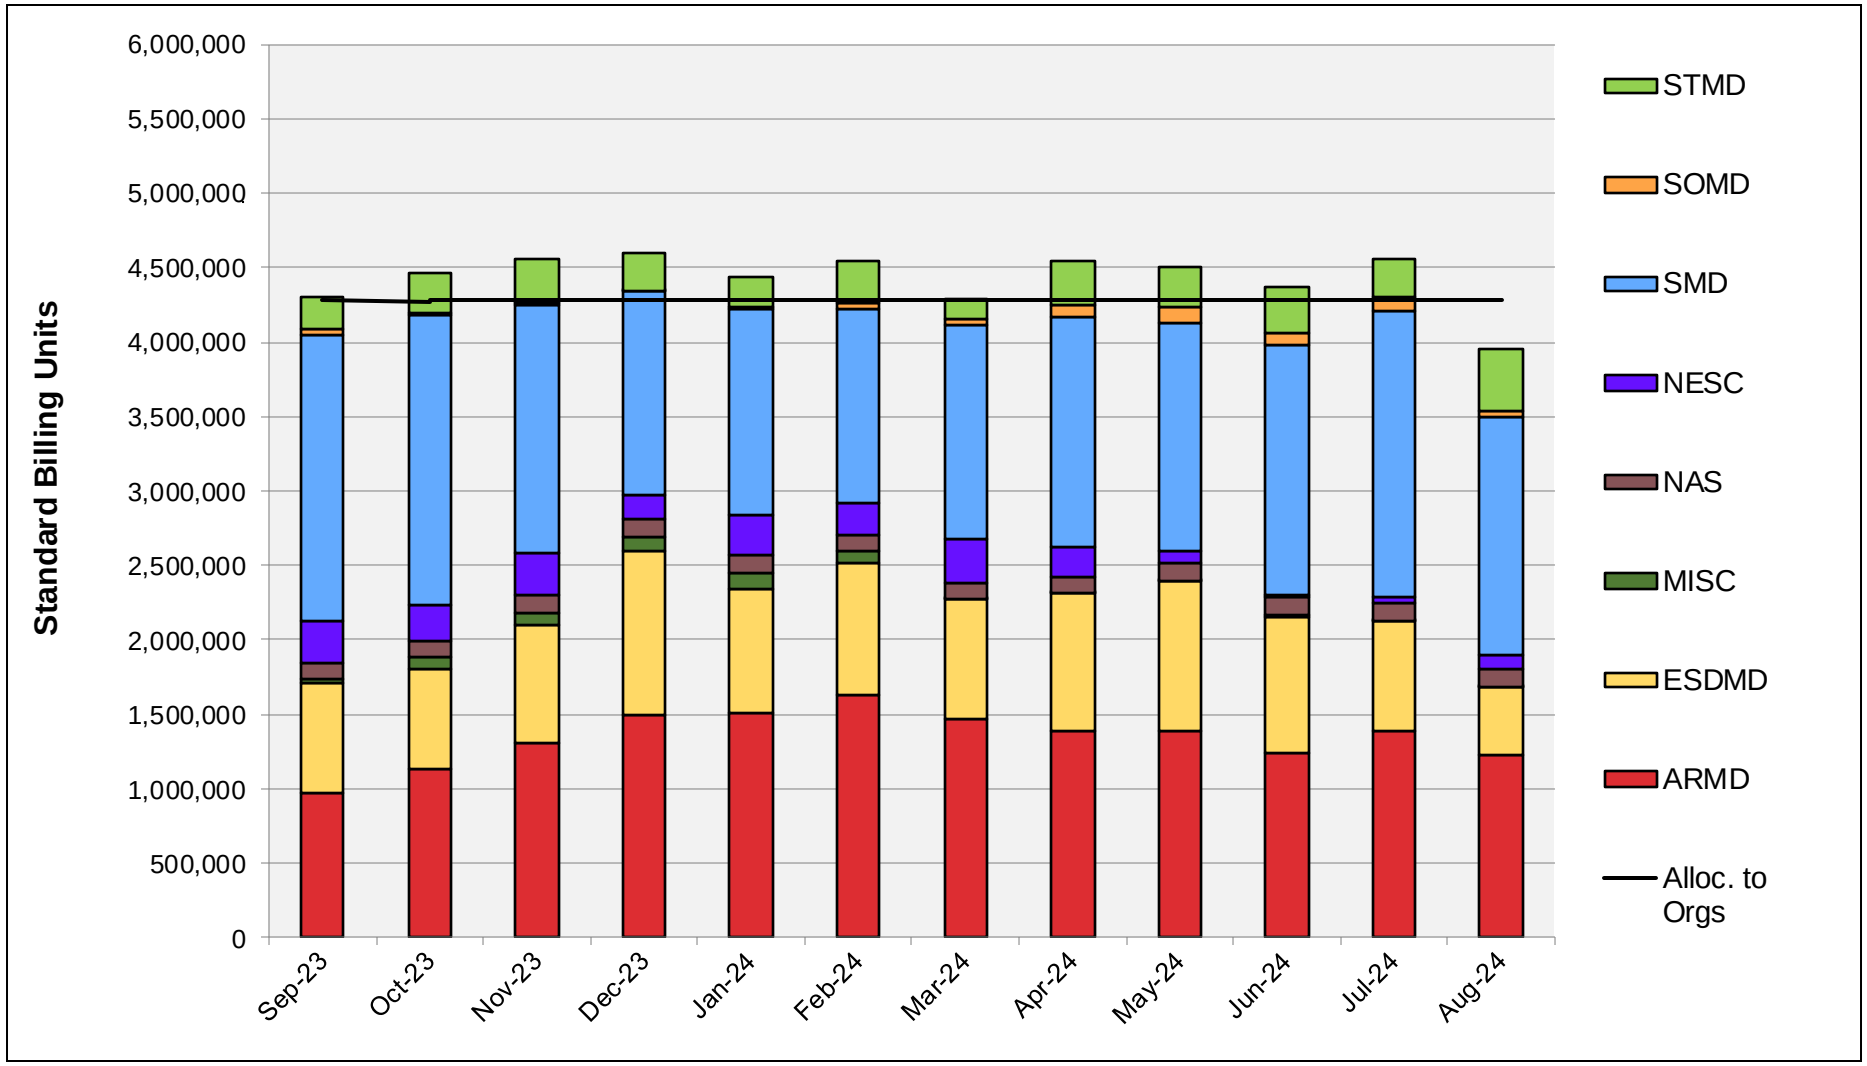
Aitken: Average Time to Clear All Jobs



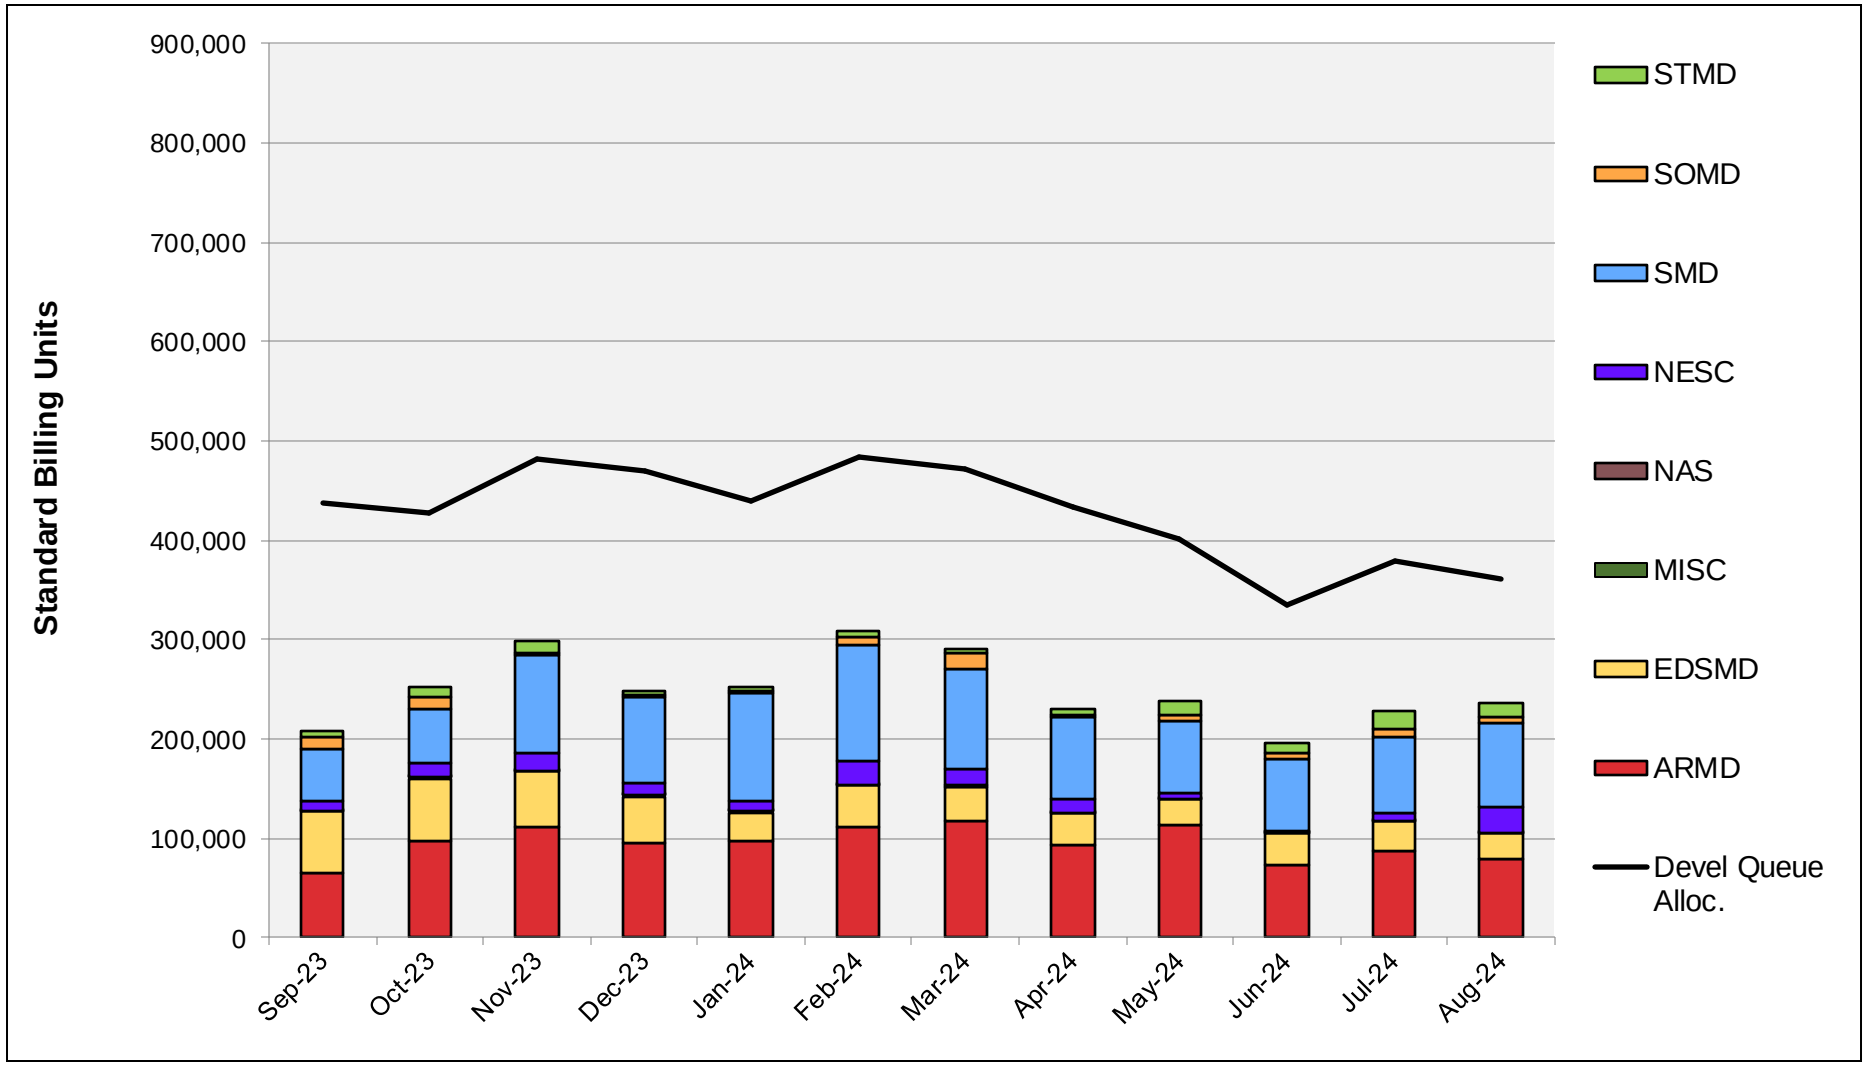
Aitken: Average Expansion Factor



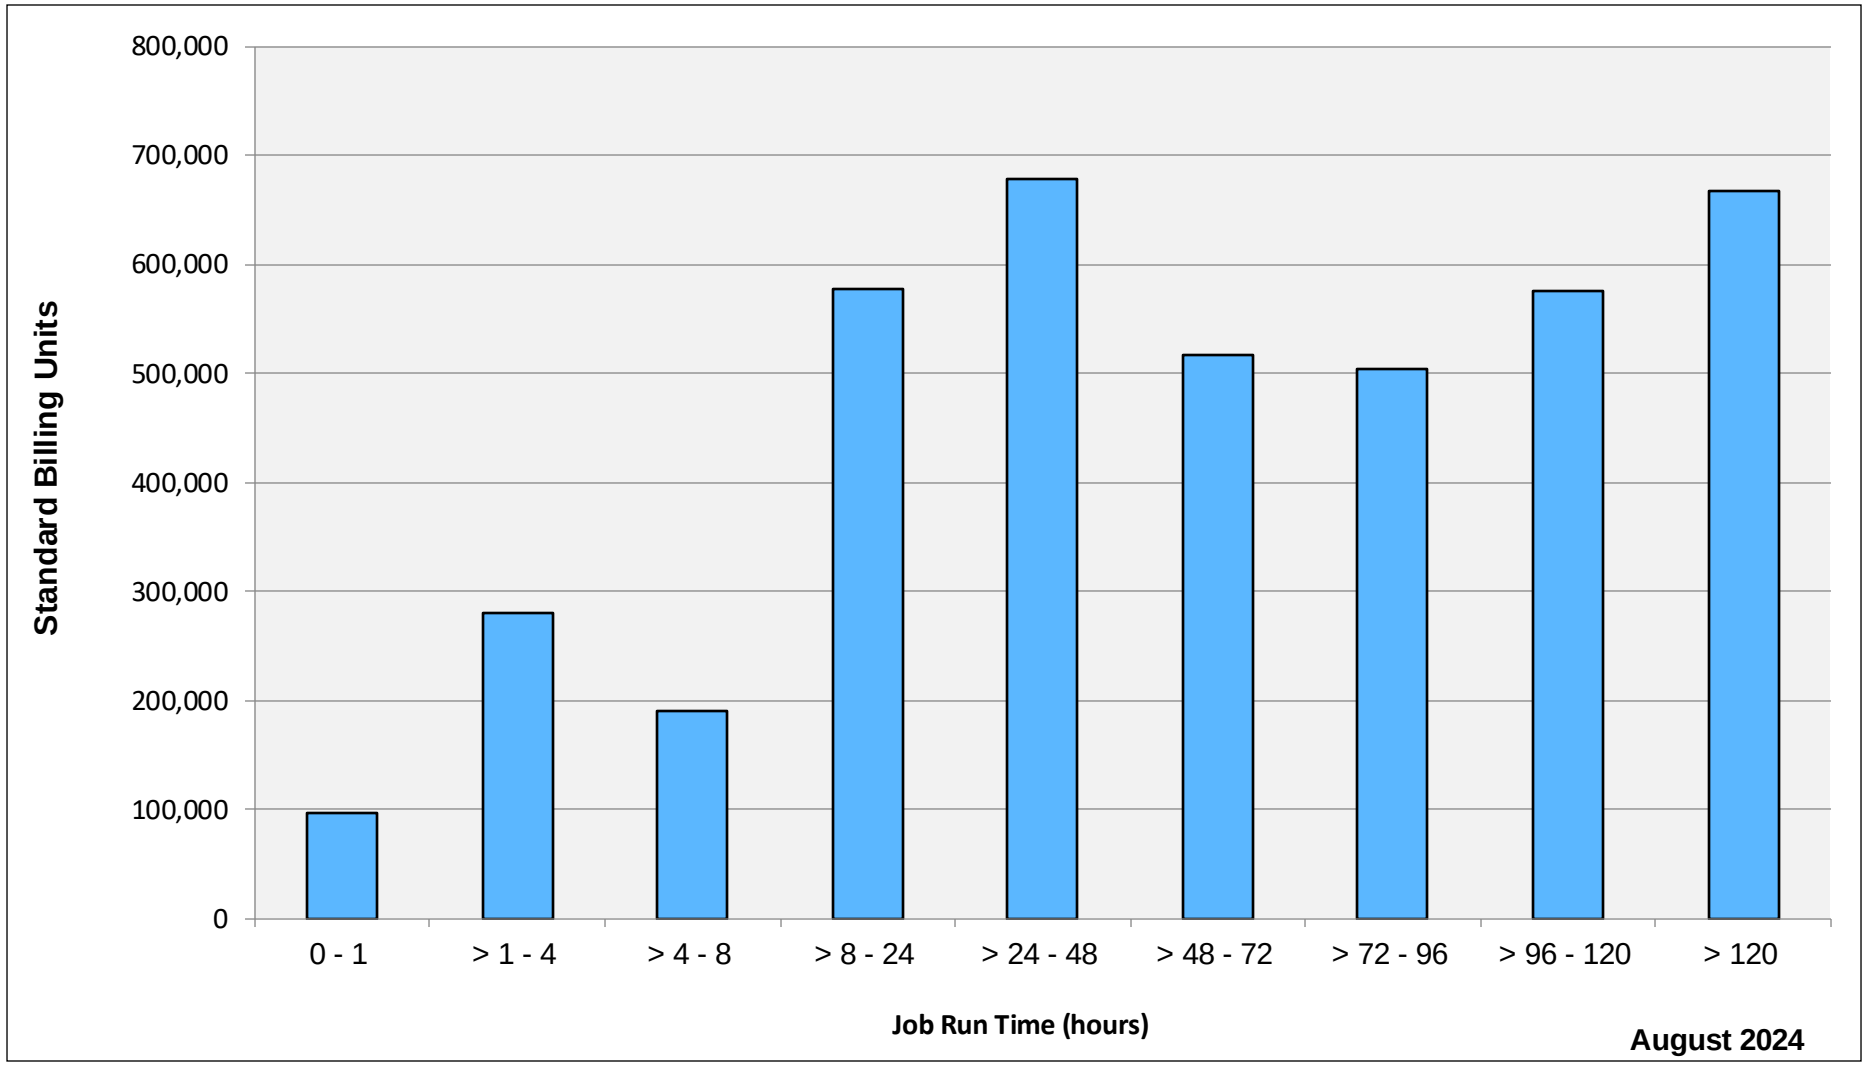
Pleiades: SBUs Reported, Normalized to 30-Day Month



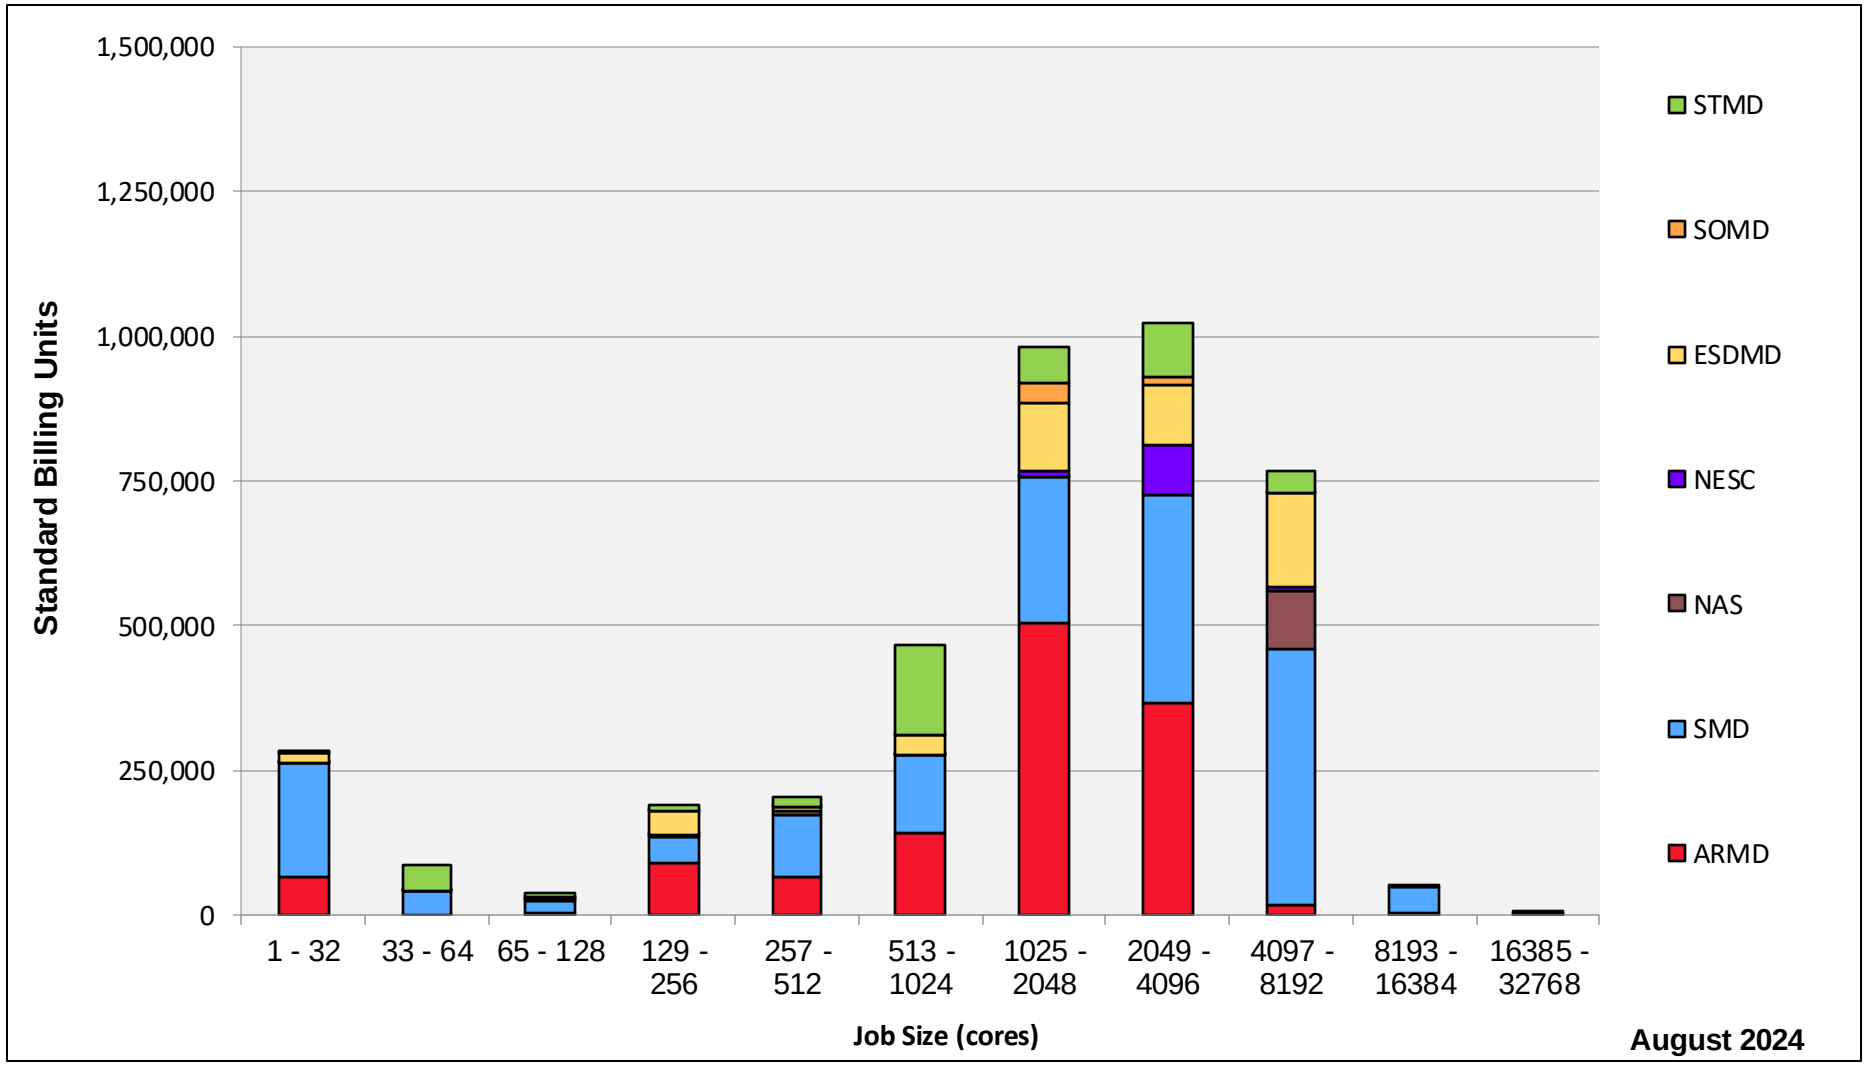
Pleiades: Devel Queue Utilization



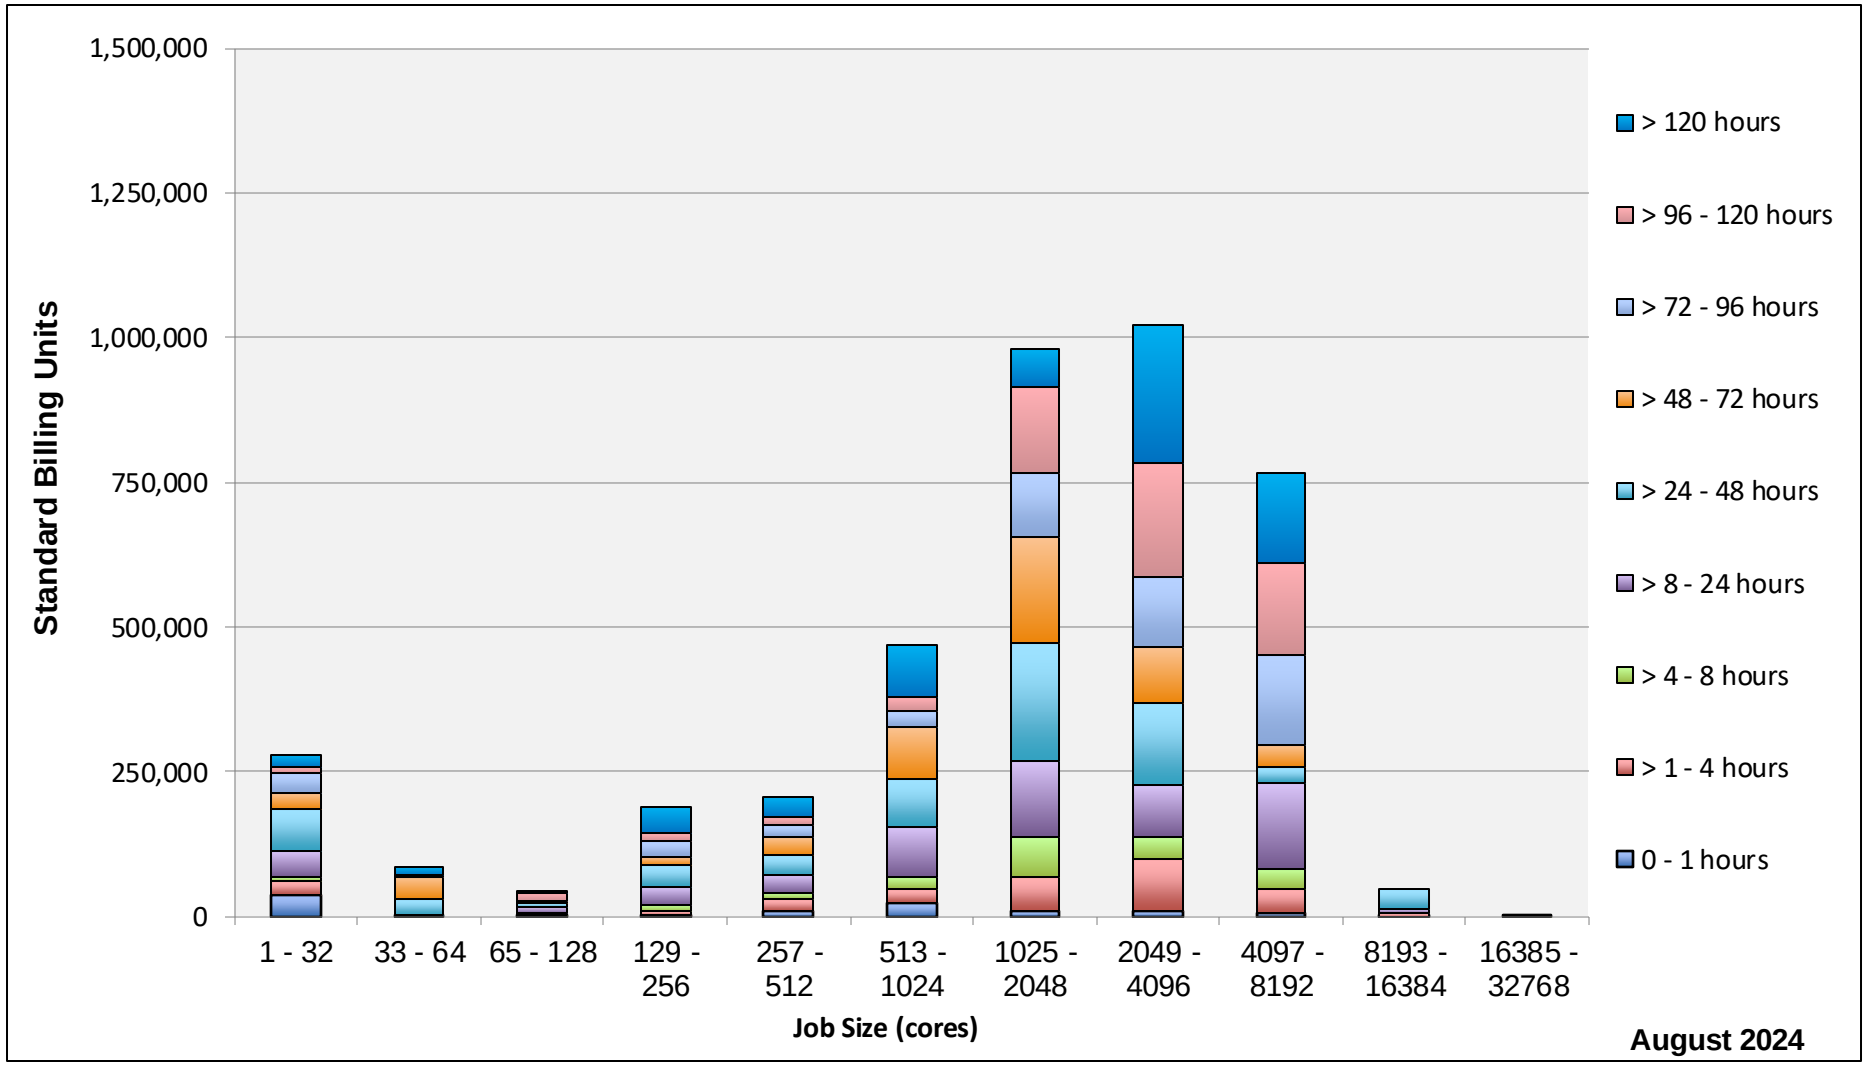
Pleiades: Monthly Utilization by Job Length



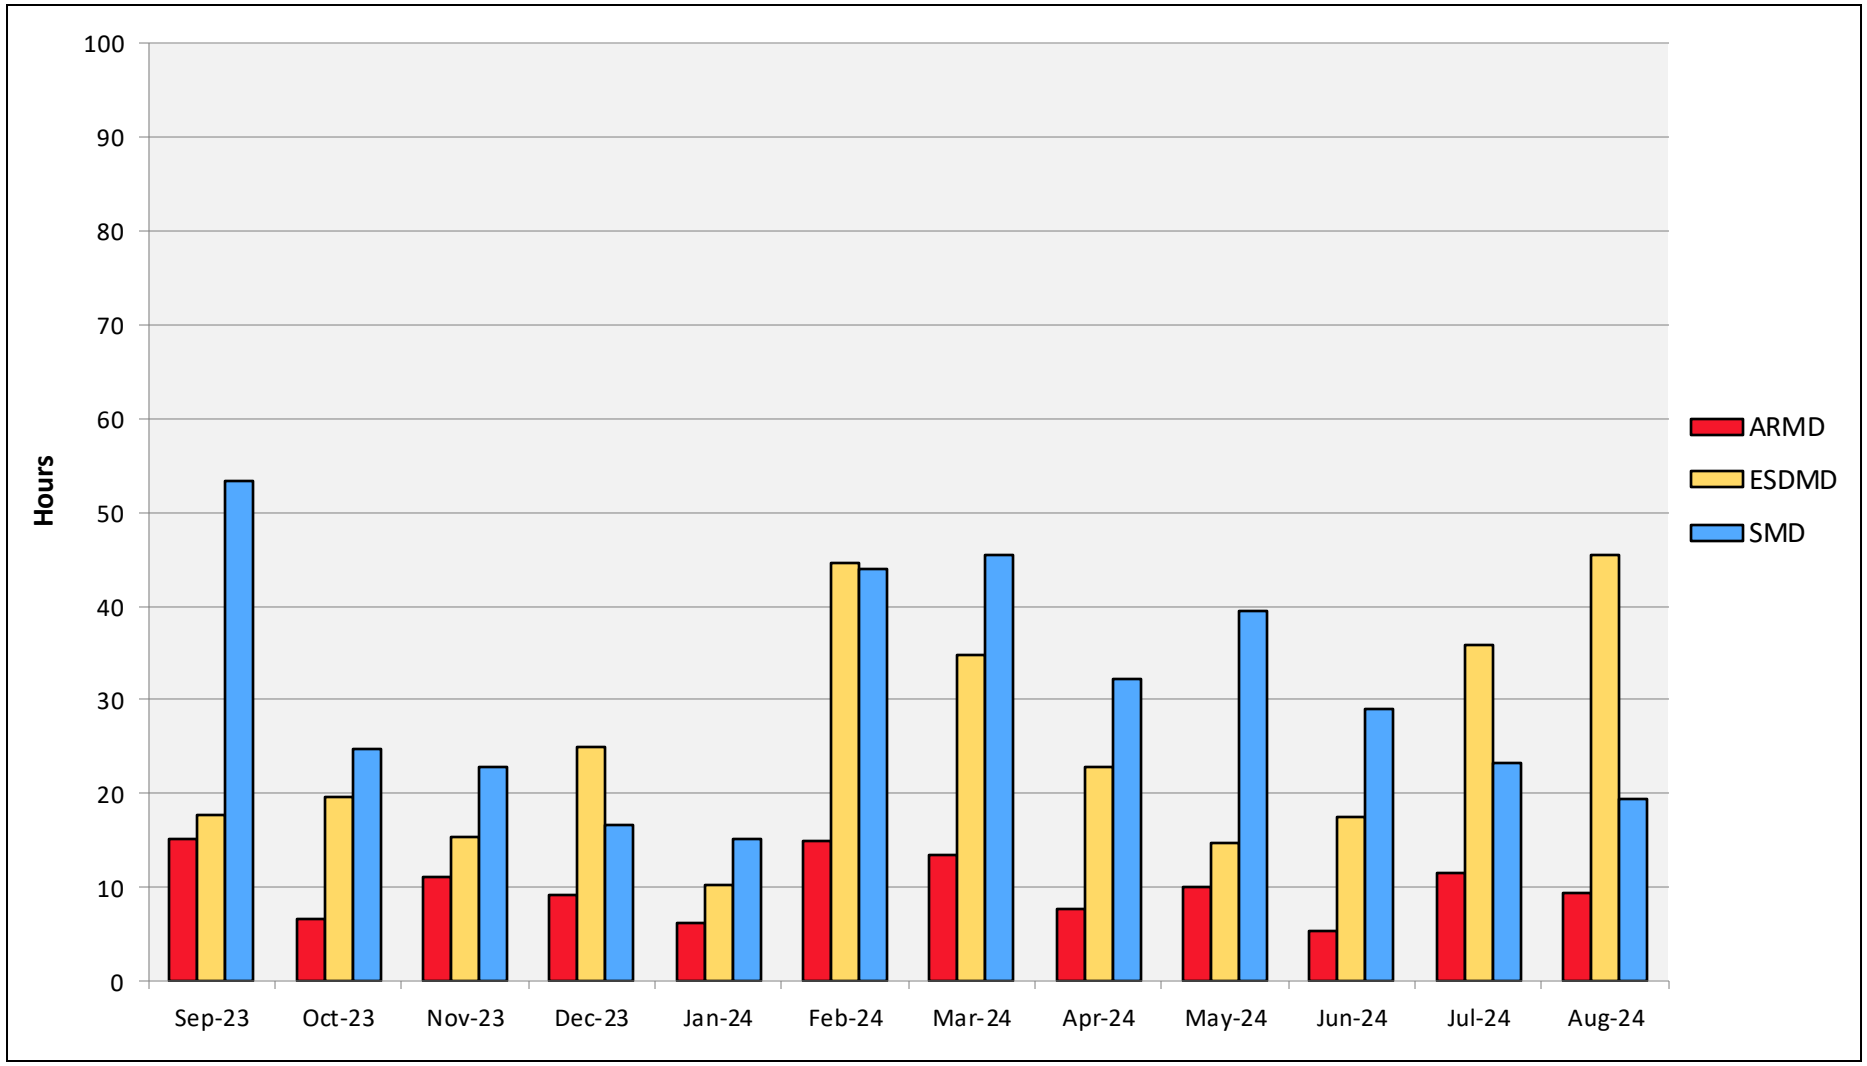
Pleiades: Monthly Utilization by Job Size



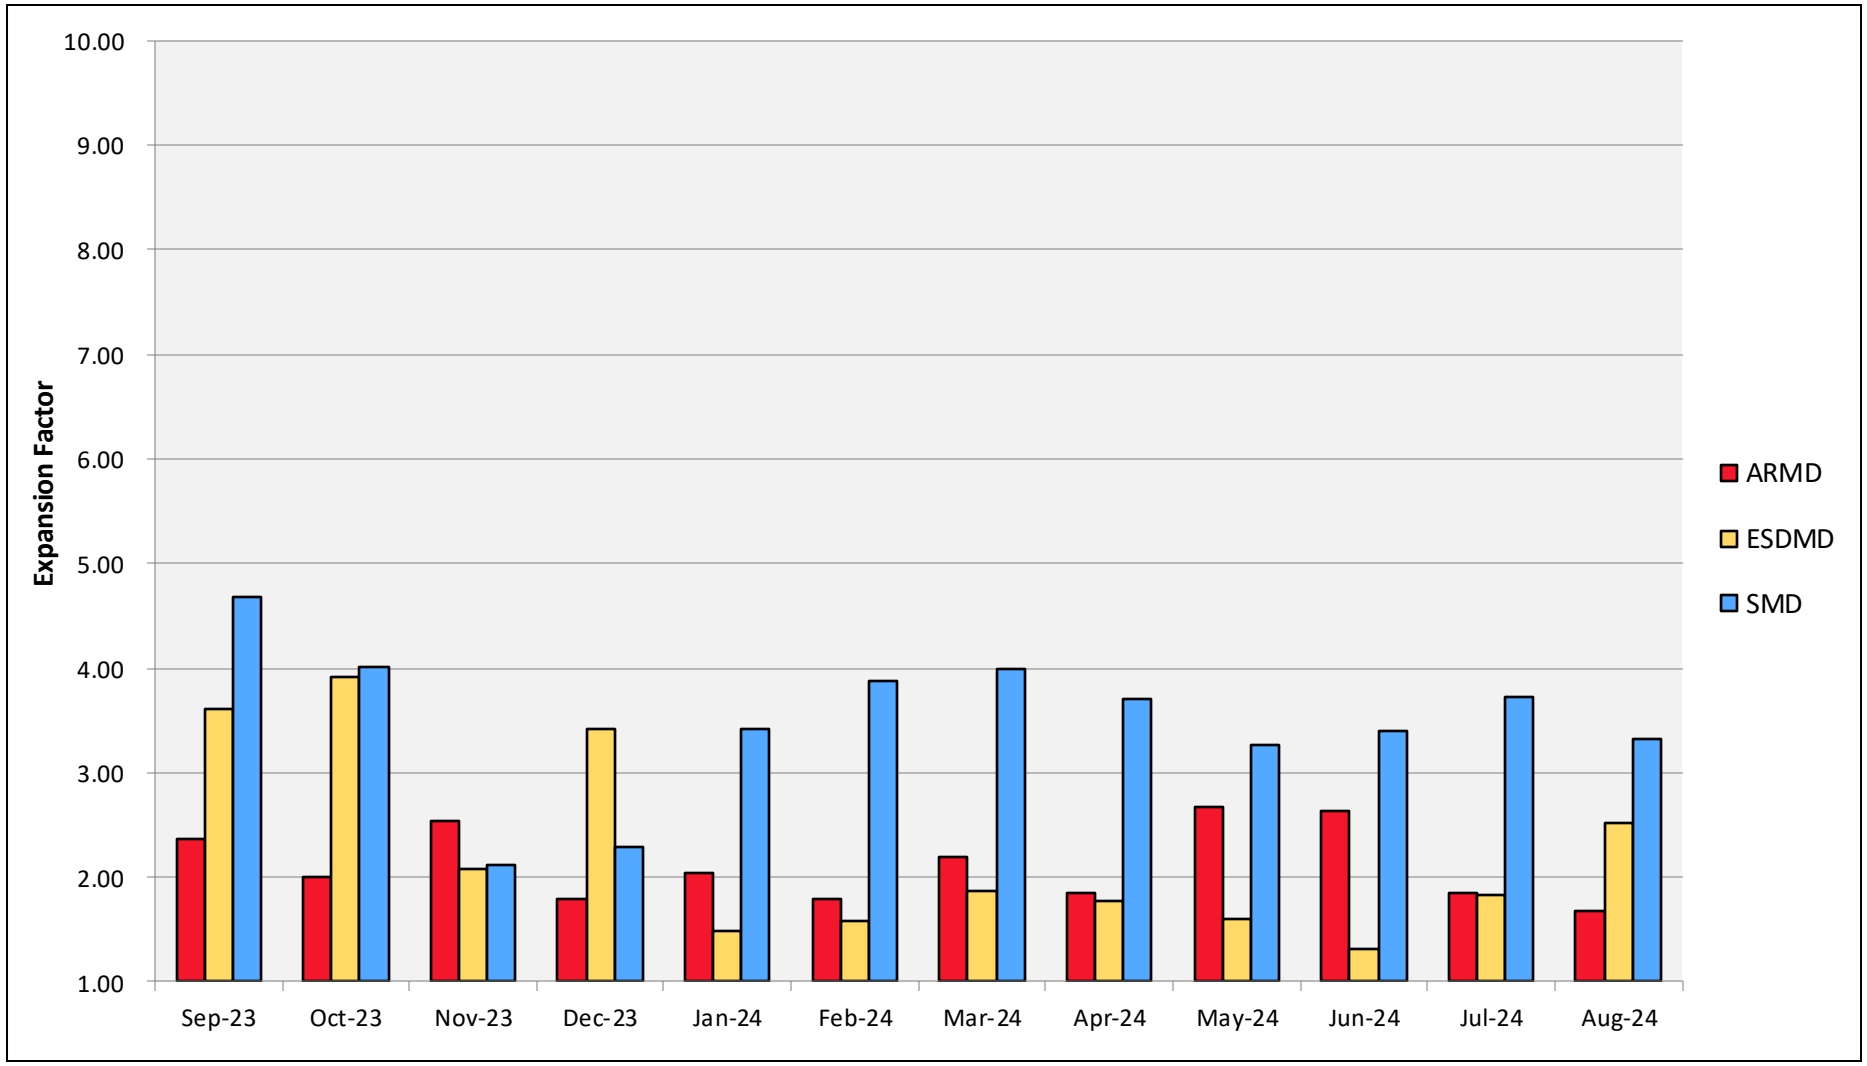
Pleiades: Monthly Utilization by Size and Length



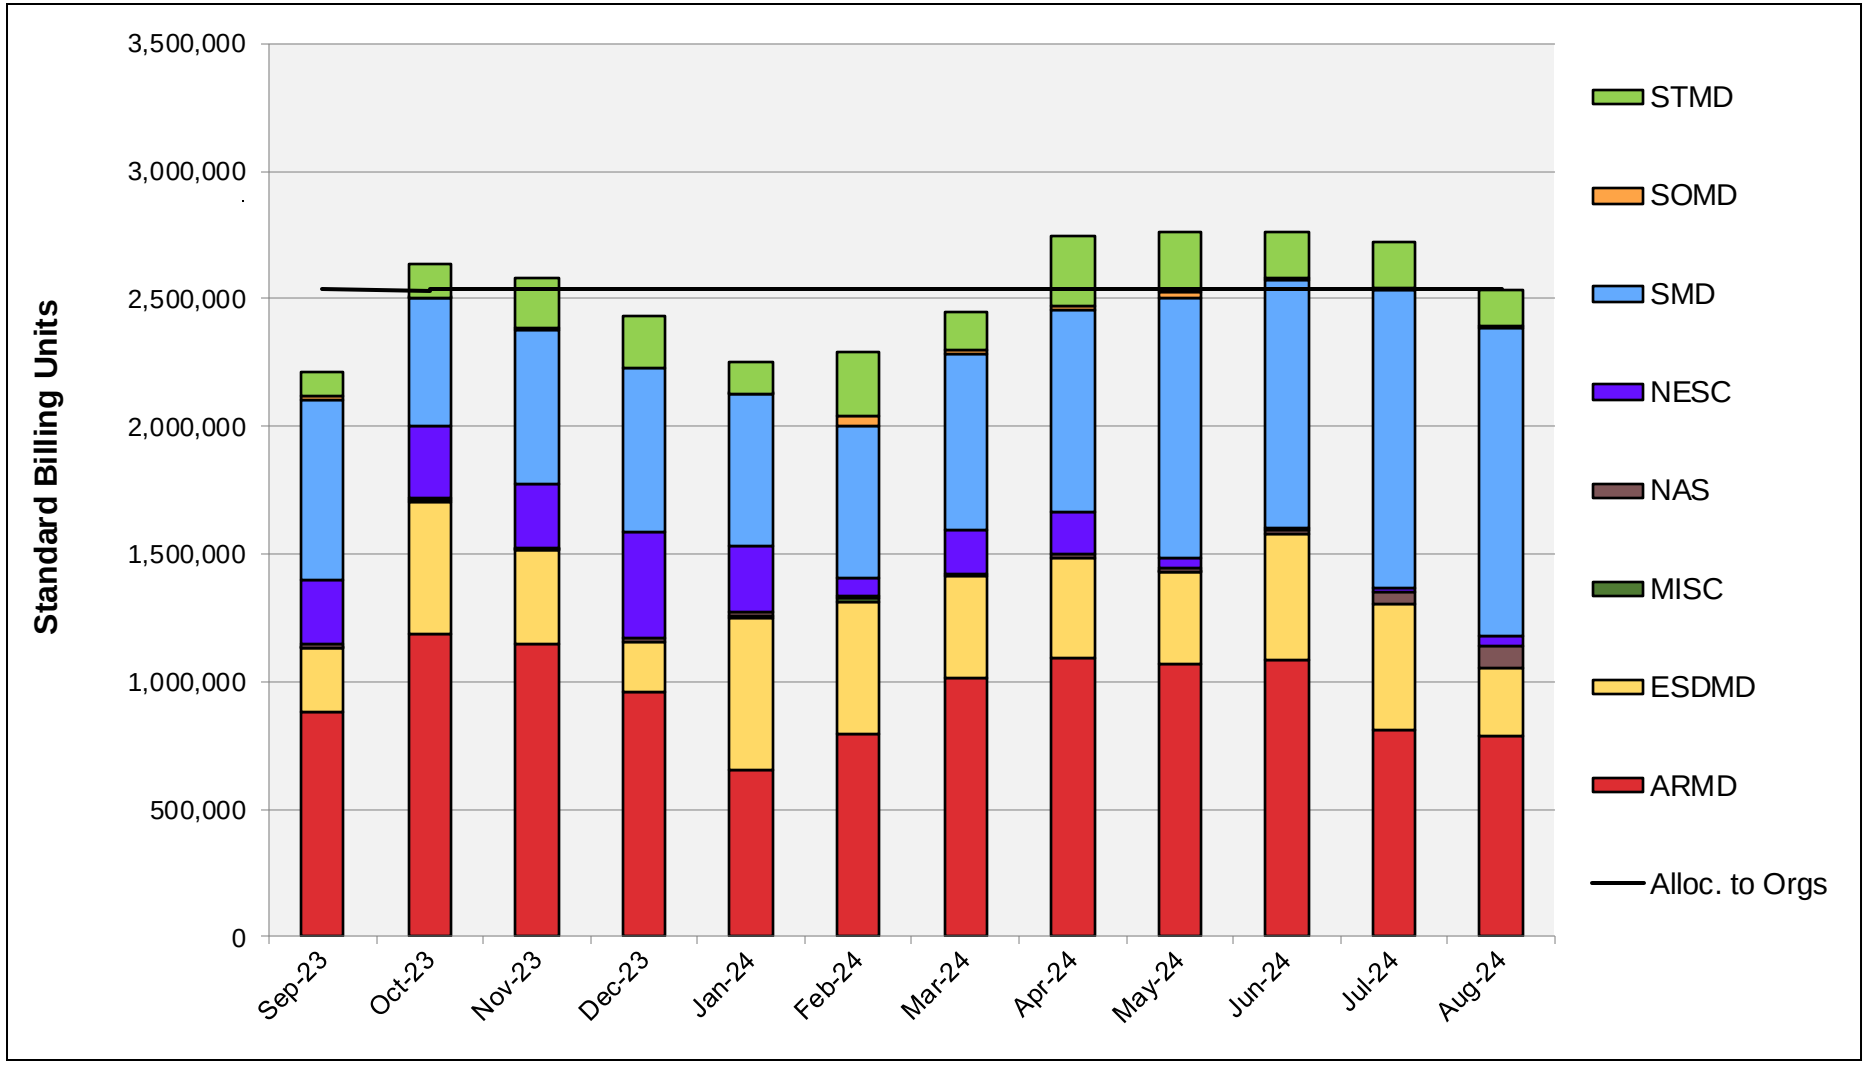
Pleiades: Average Time to Clear All Jobs



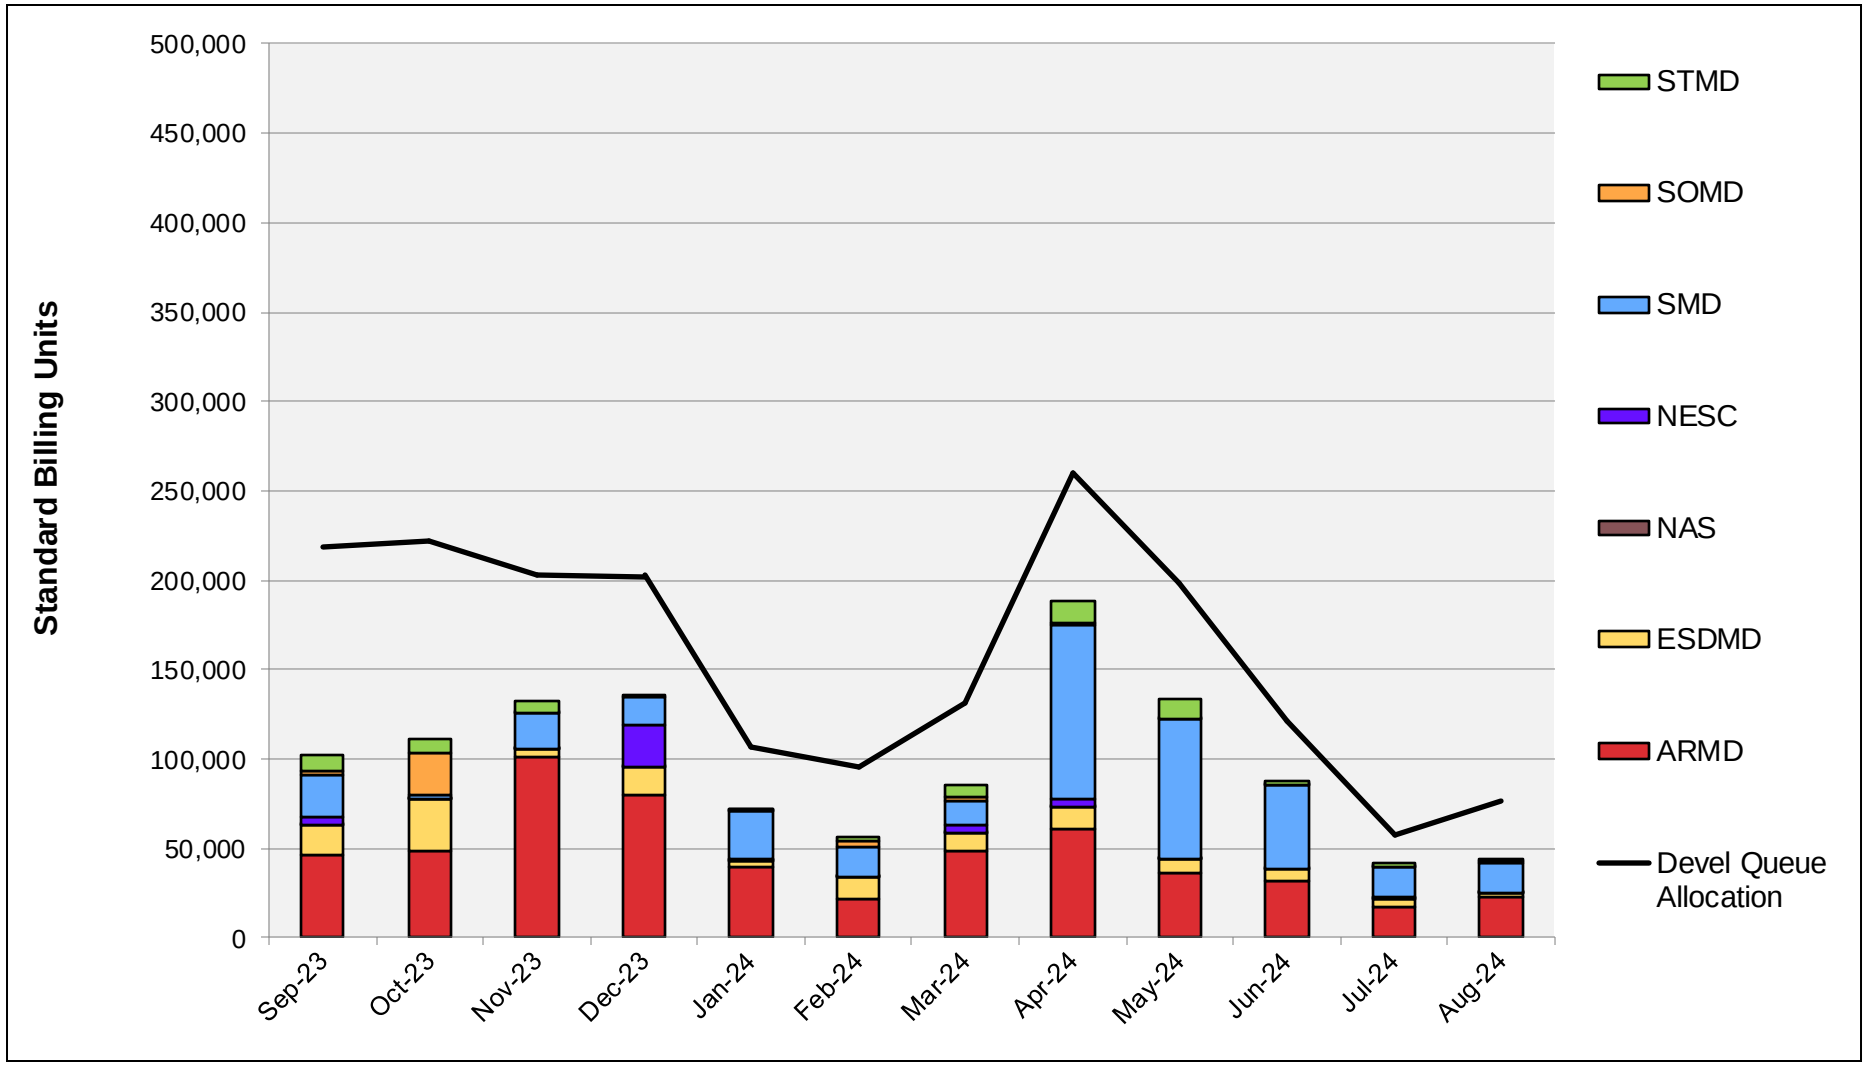
Pleiades: Average Expansion Factor



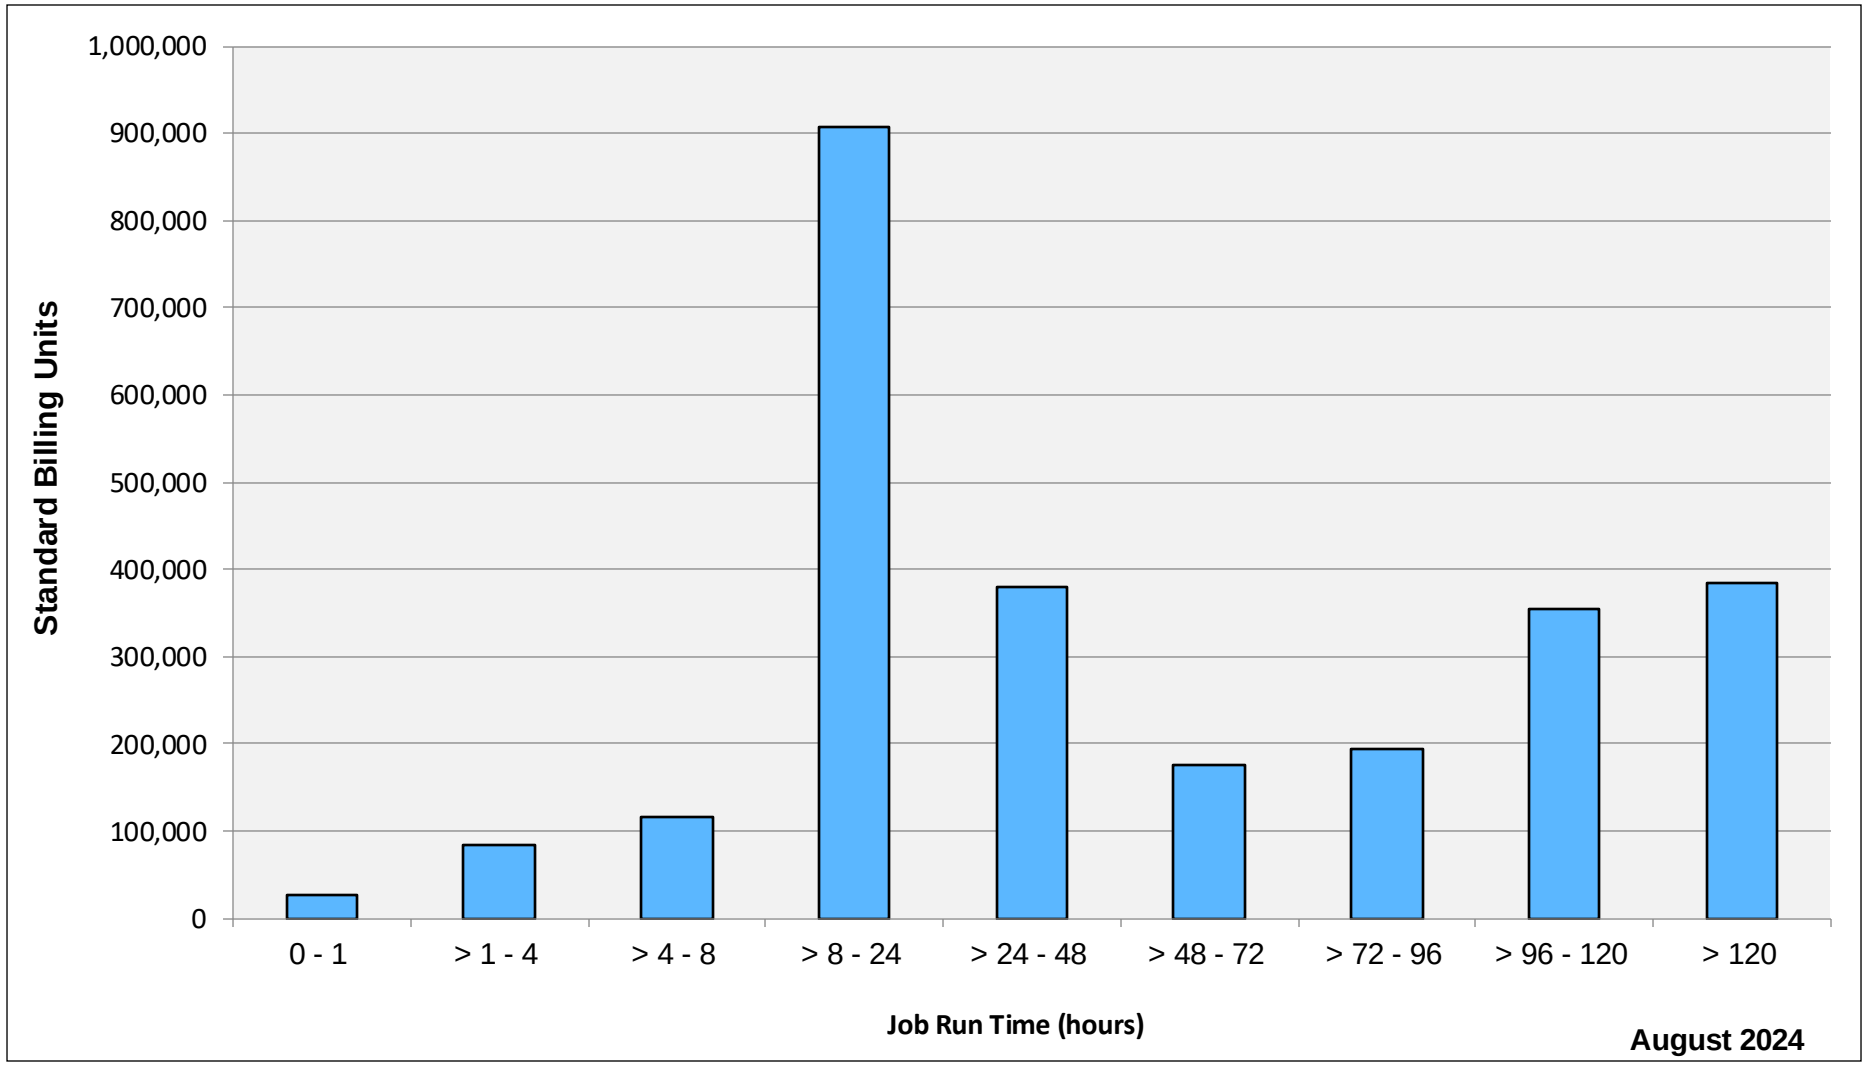
Electra: SBUs Reported, Normalized to 30-Day Month



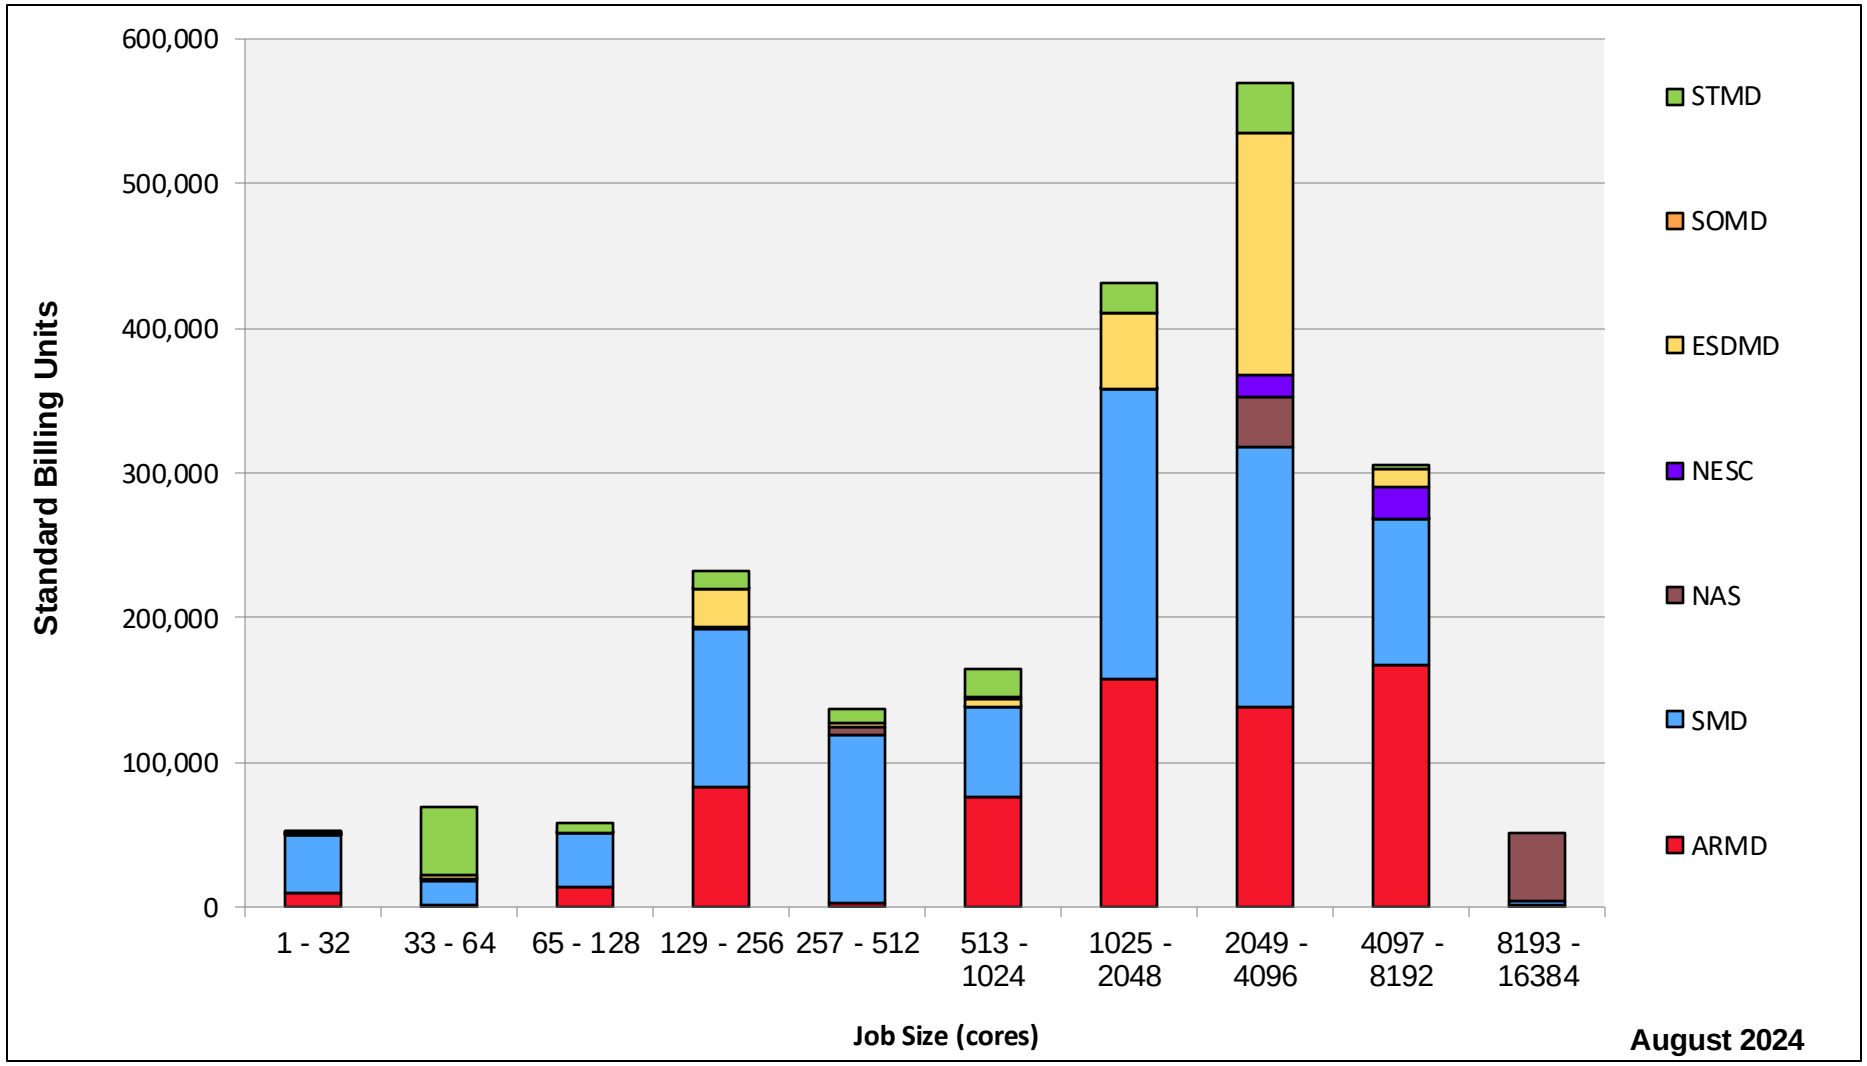
Electra: Devel Queue Utilization



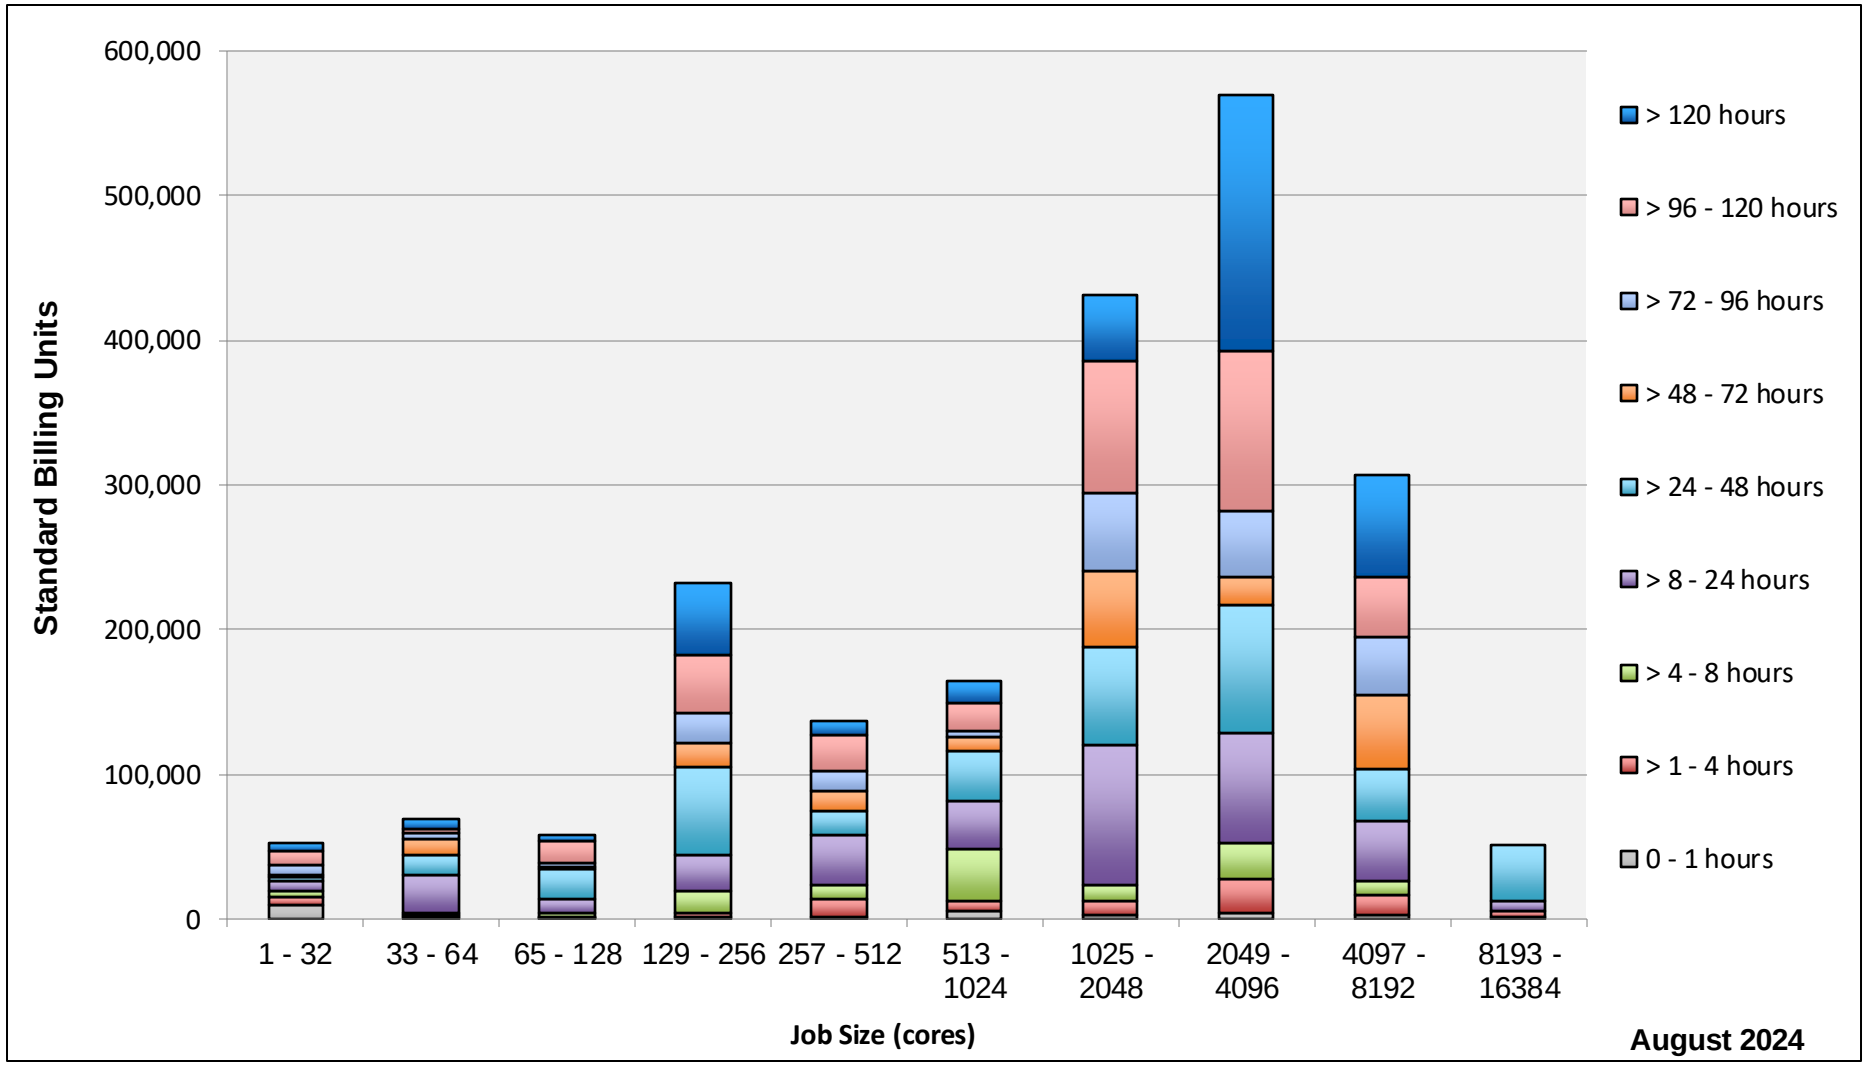
Electra: Monthly Utilization by Job Length



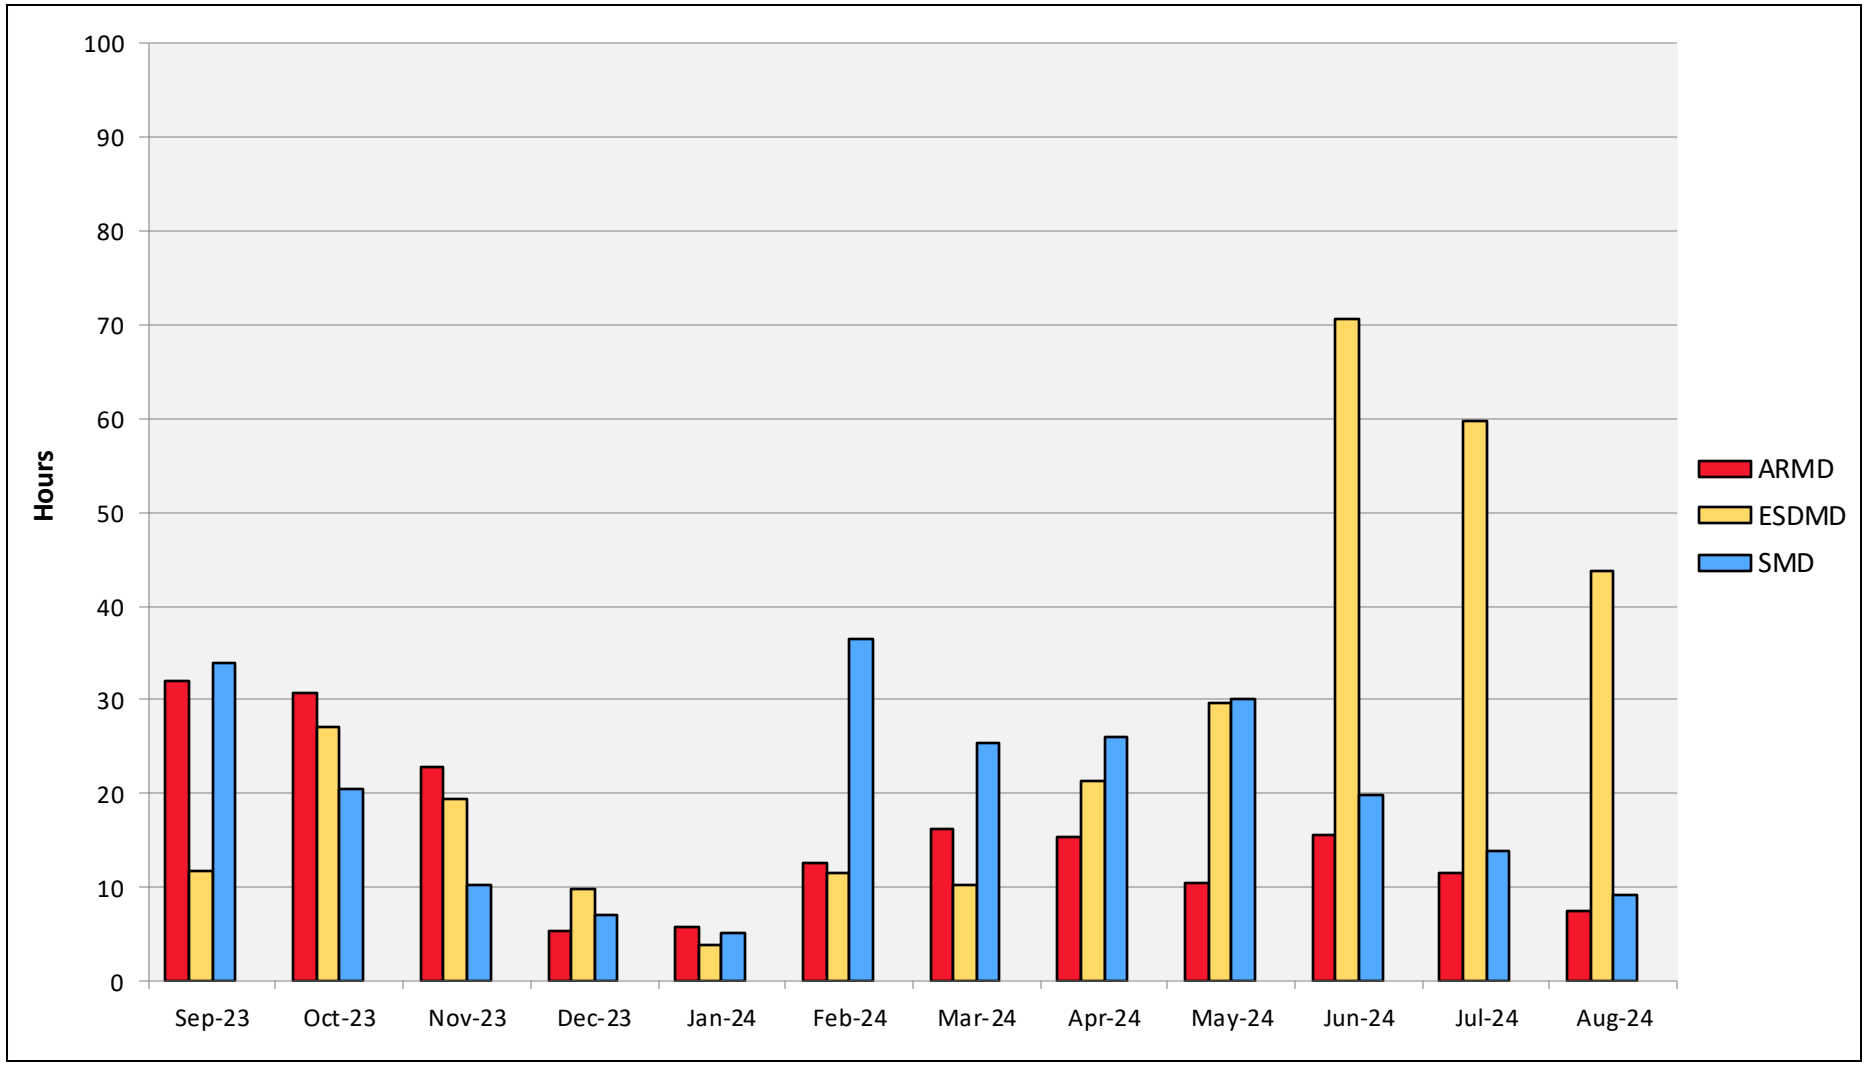
Electra: Monthly Utilization by Job Size



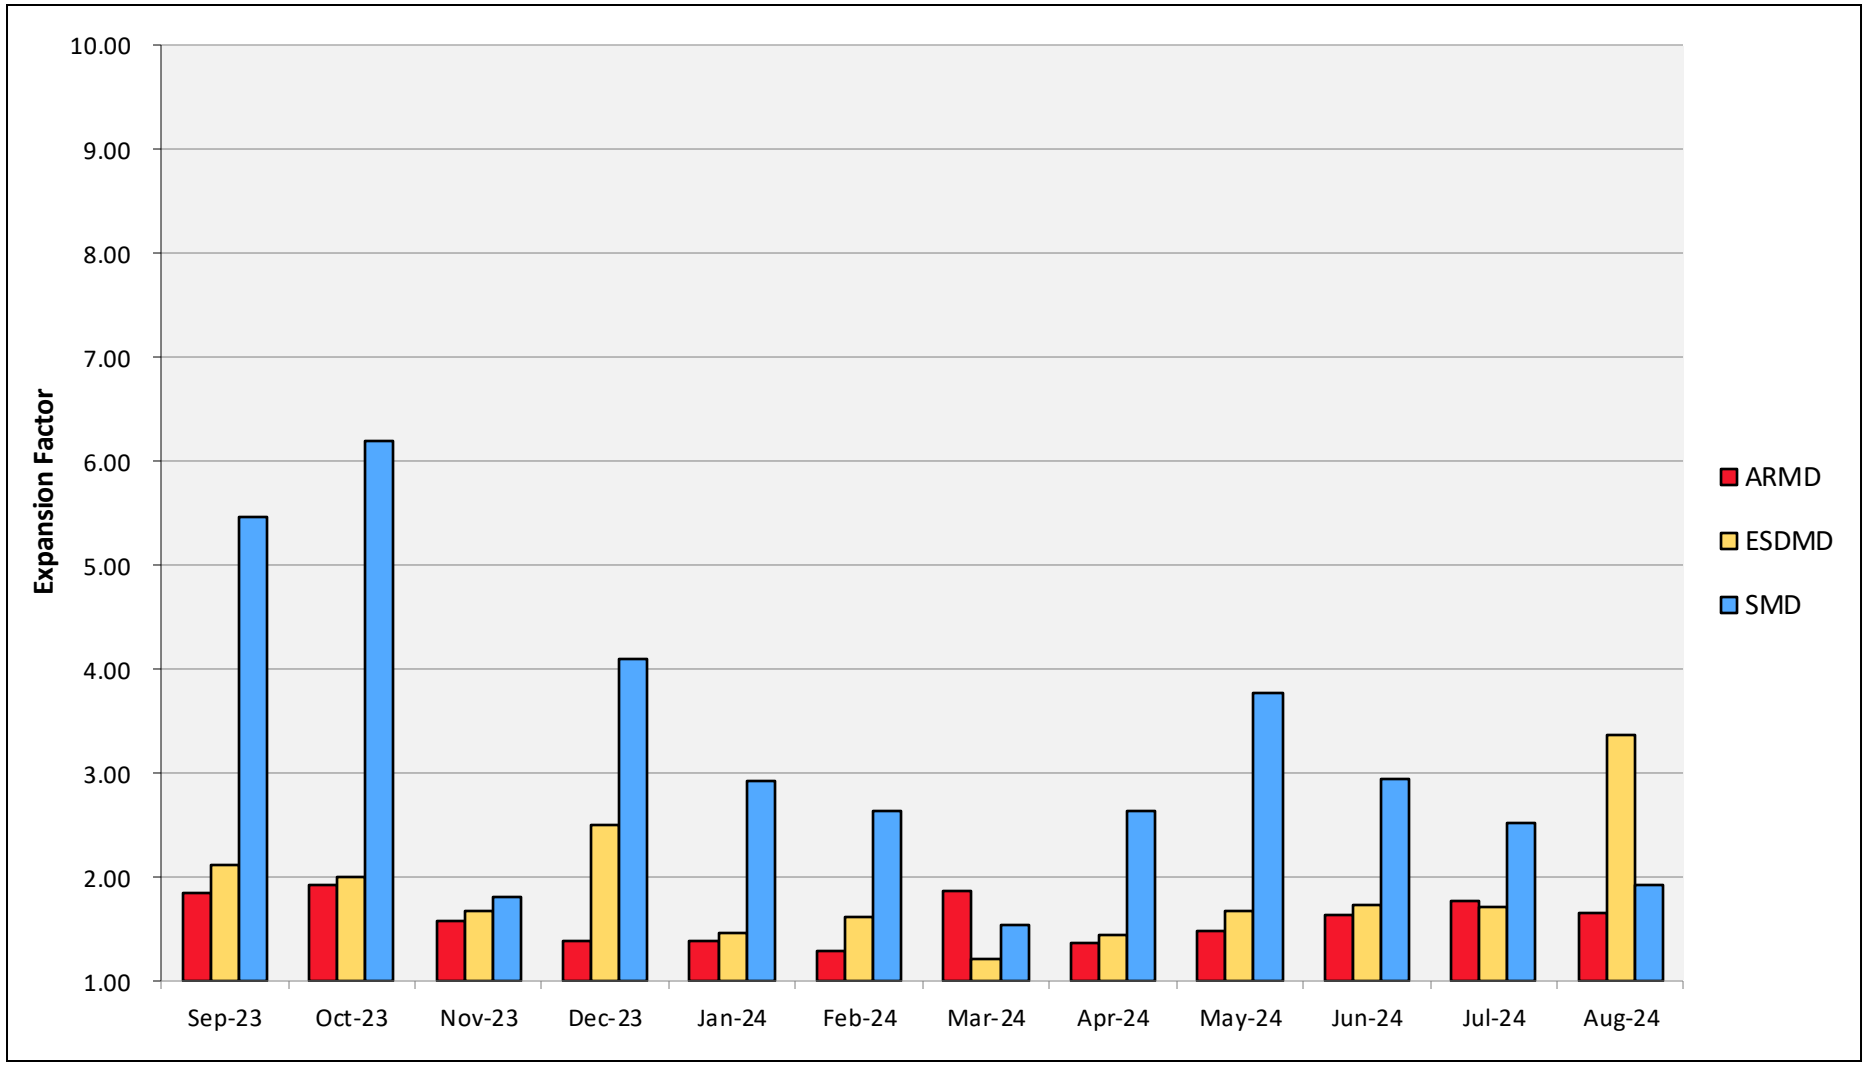
Electra: Monthly Utilization by Size and Length



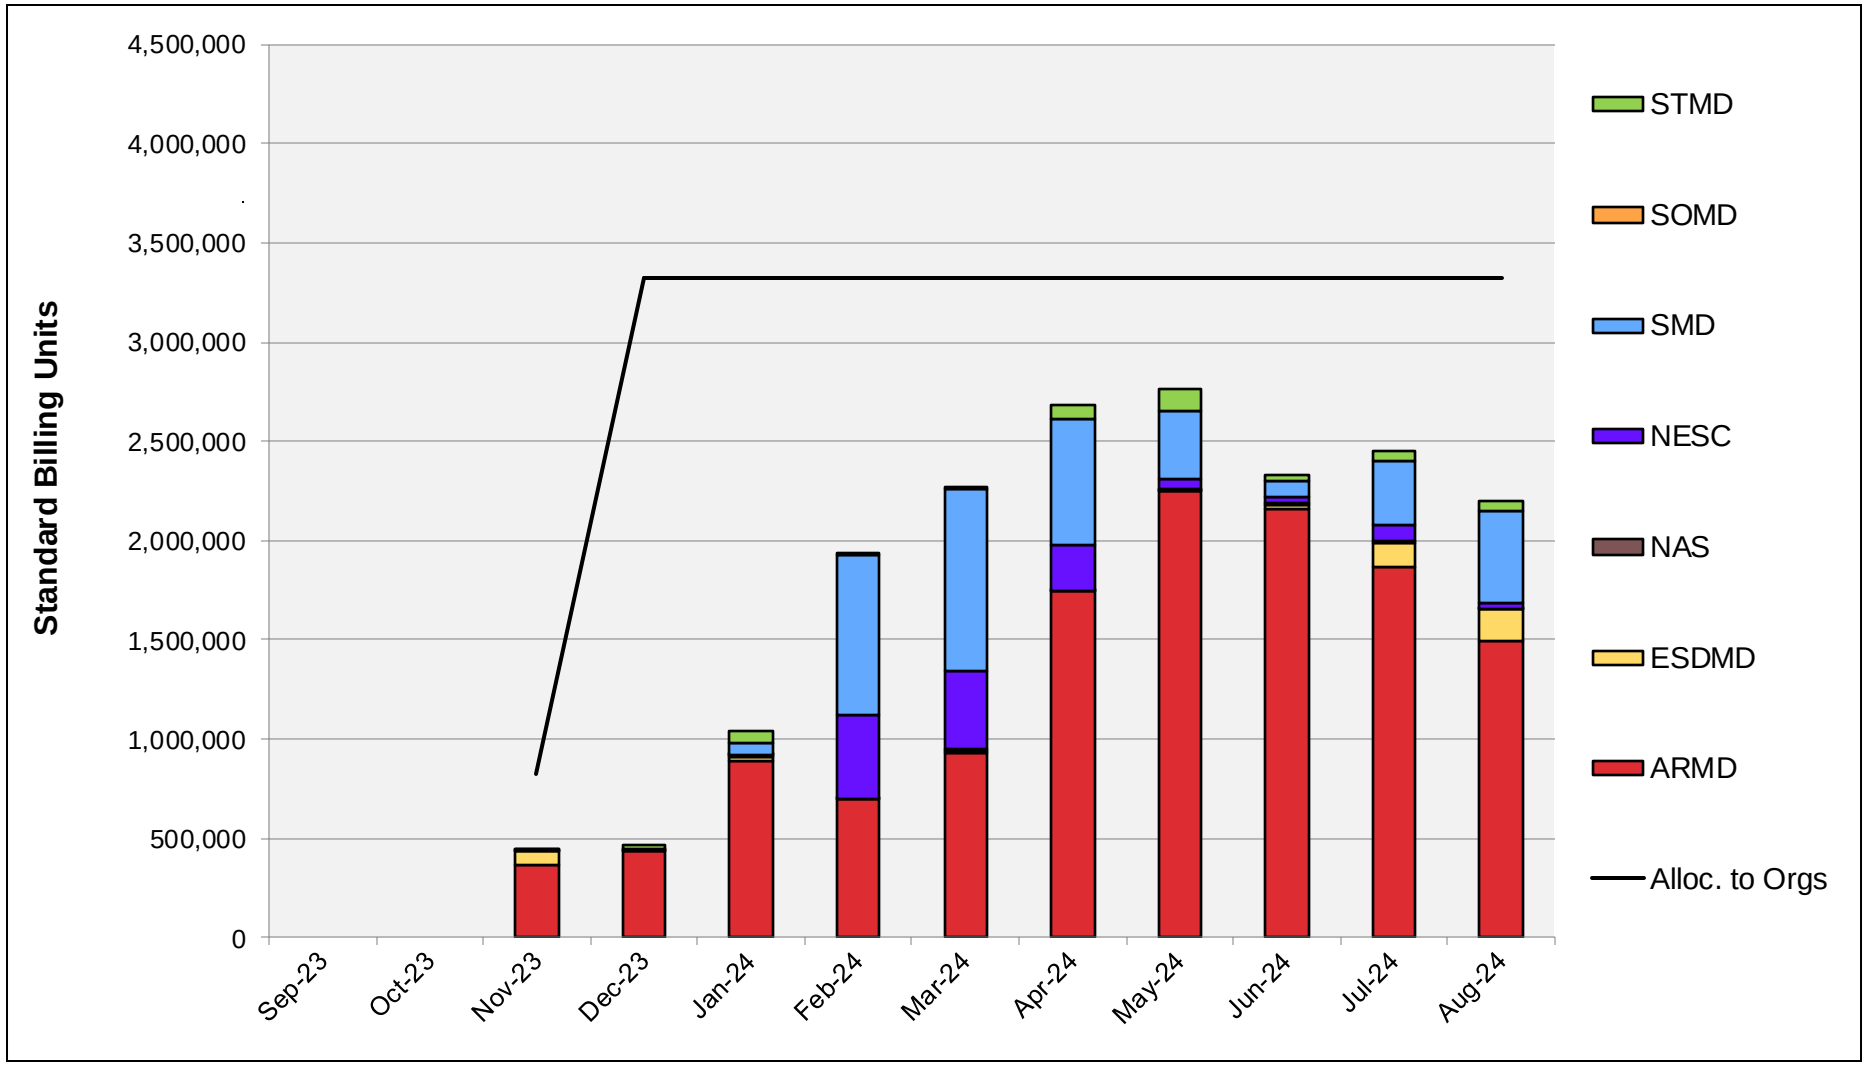
Electra: Average Time to Clear All Jobs



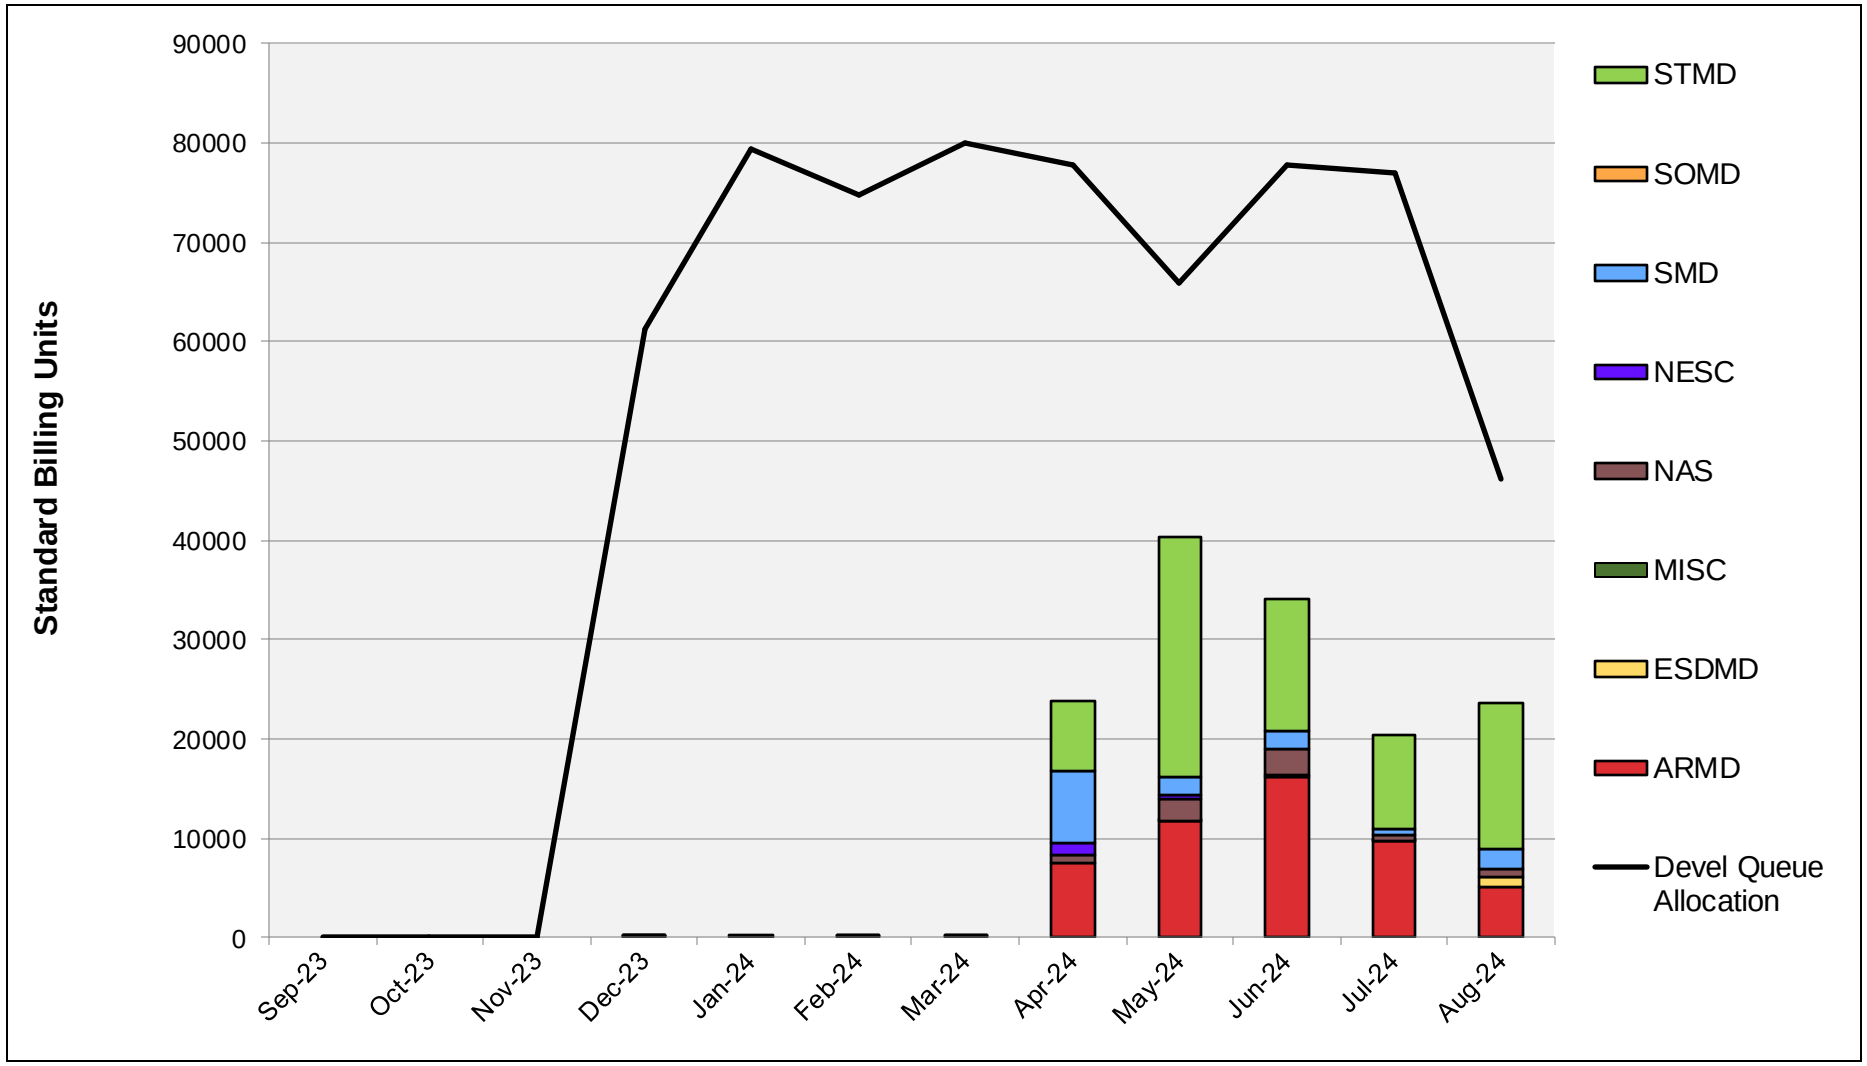
Electra: Average Expansion Factor



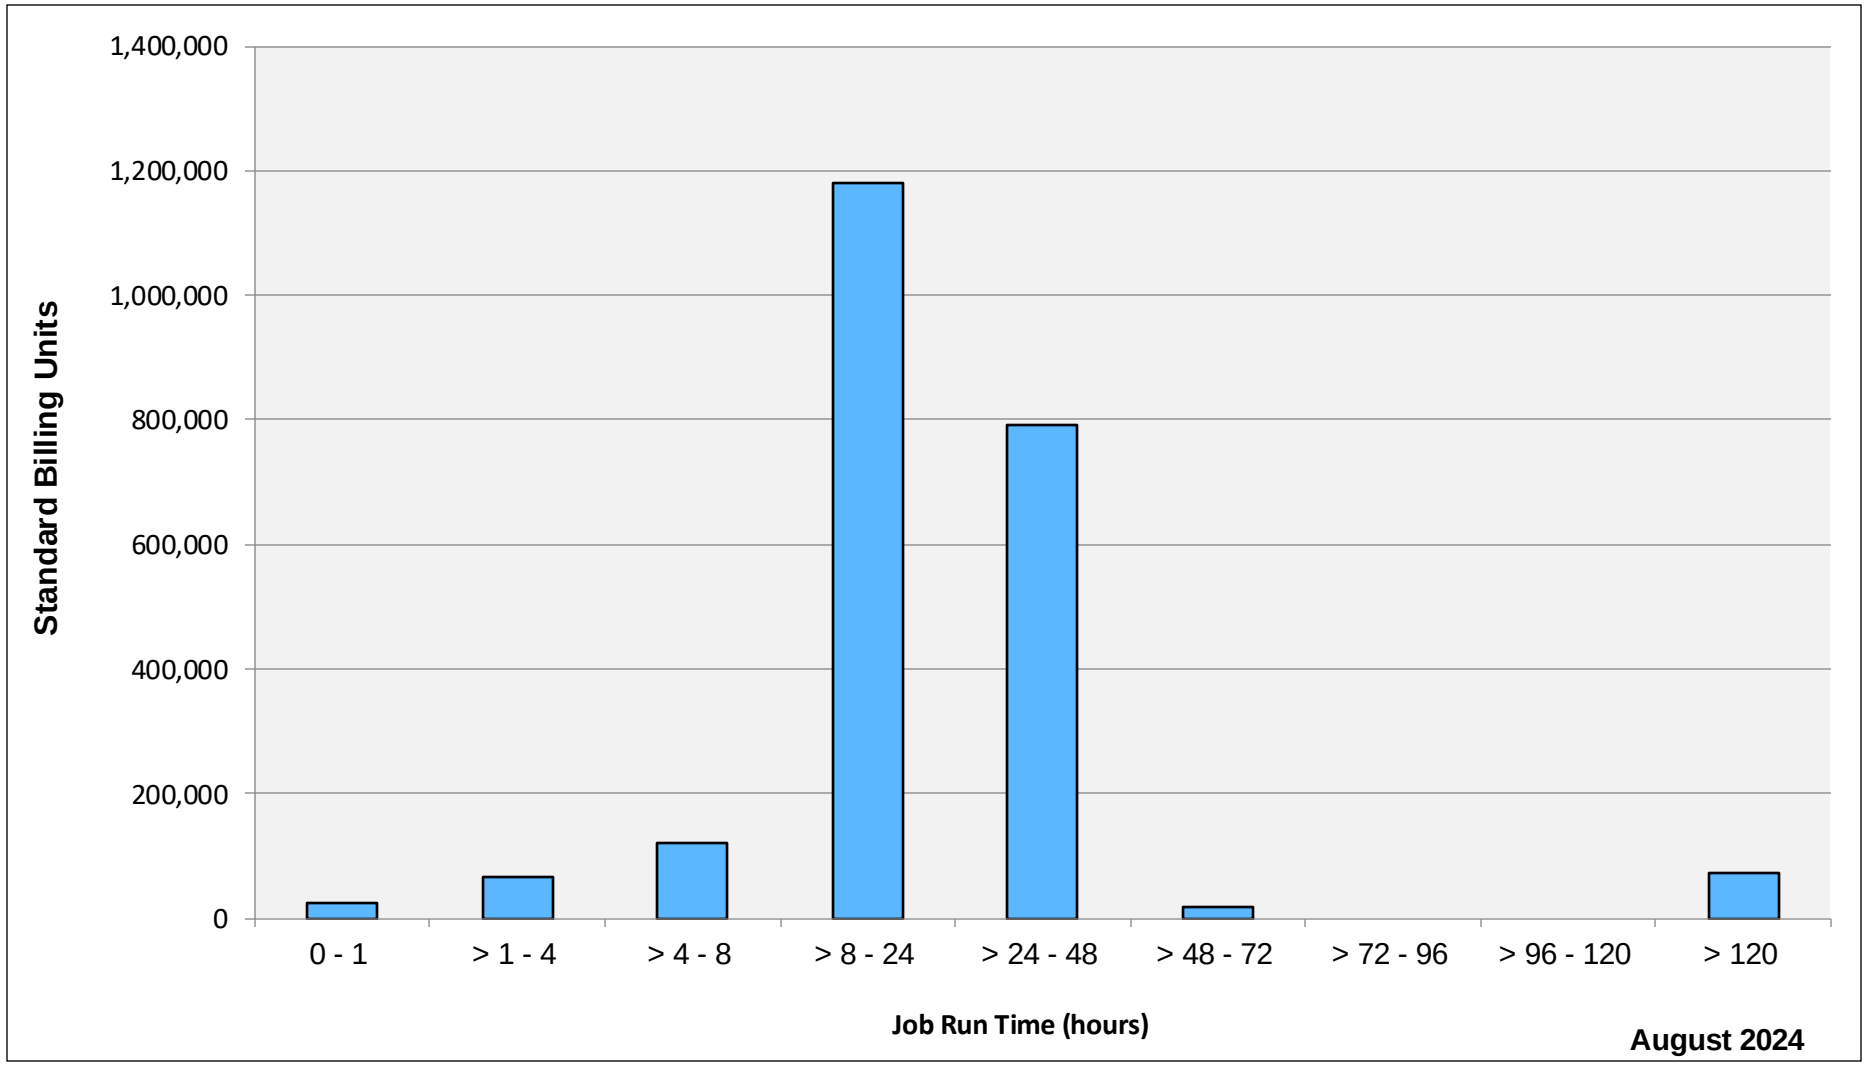
Cabeus: SBUs Reported, Normalized to 30-Day Month



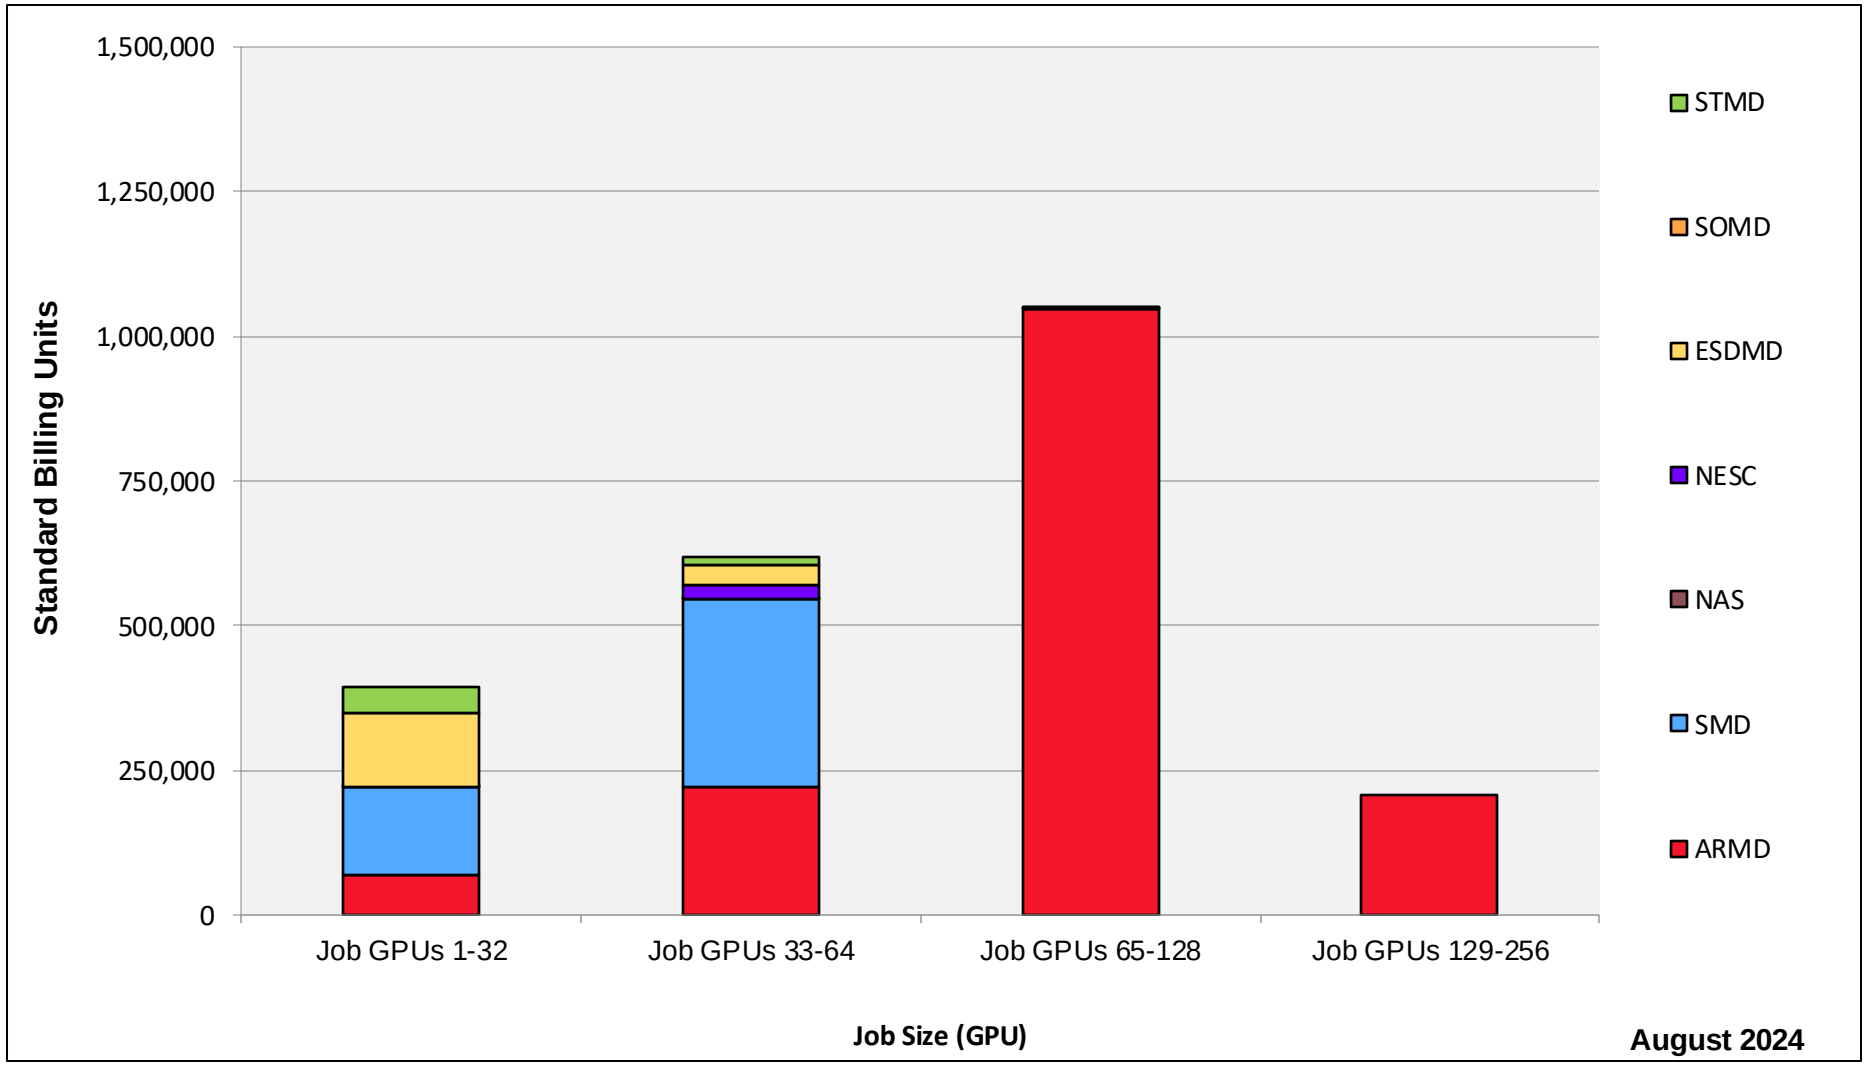
Cabeus: Devel Queue Utilization



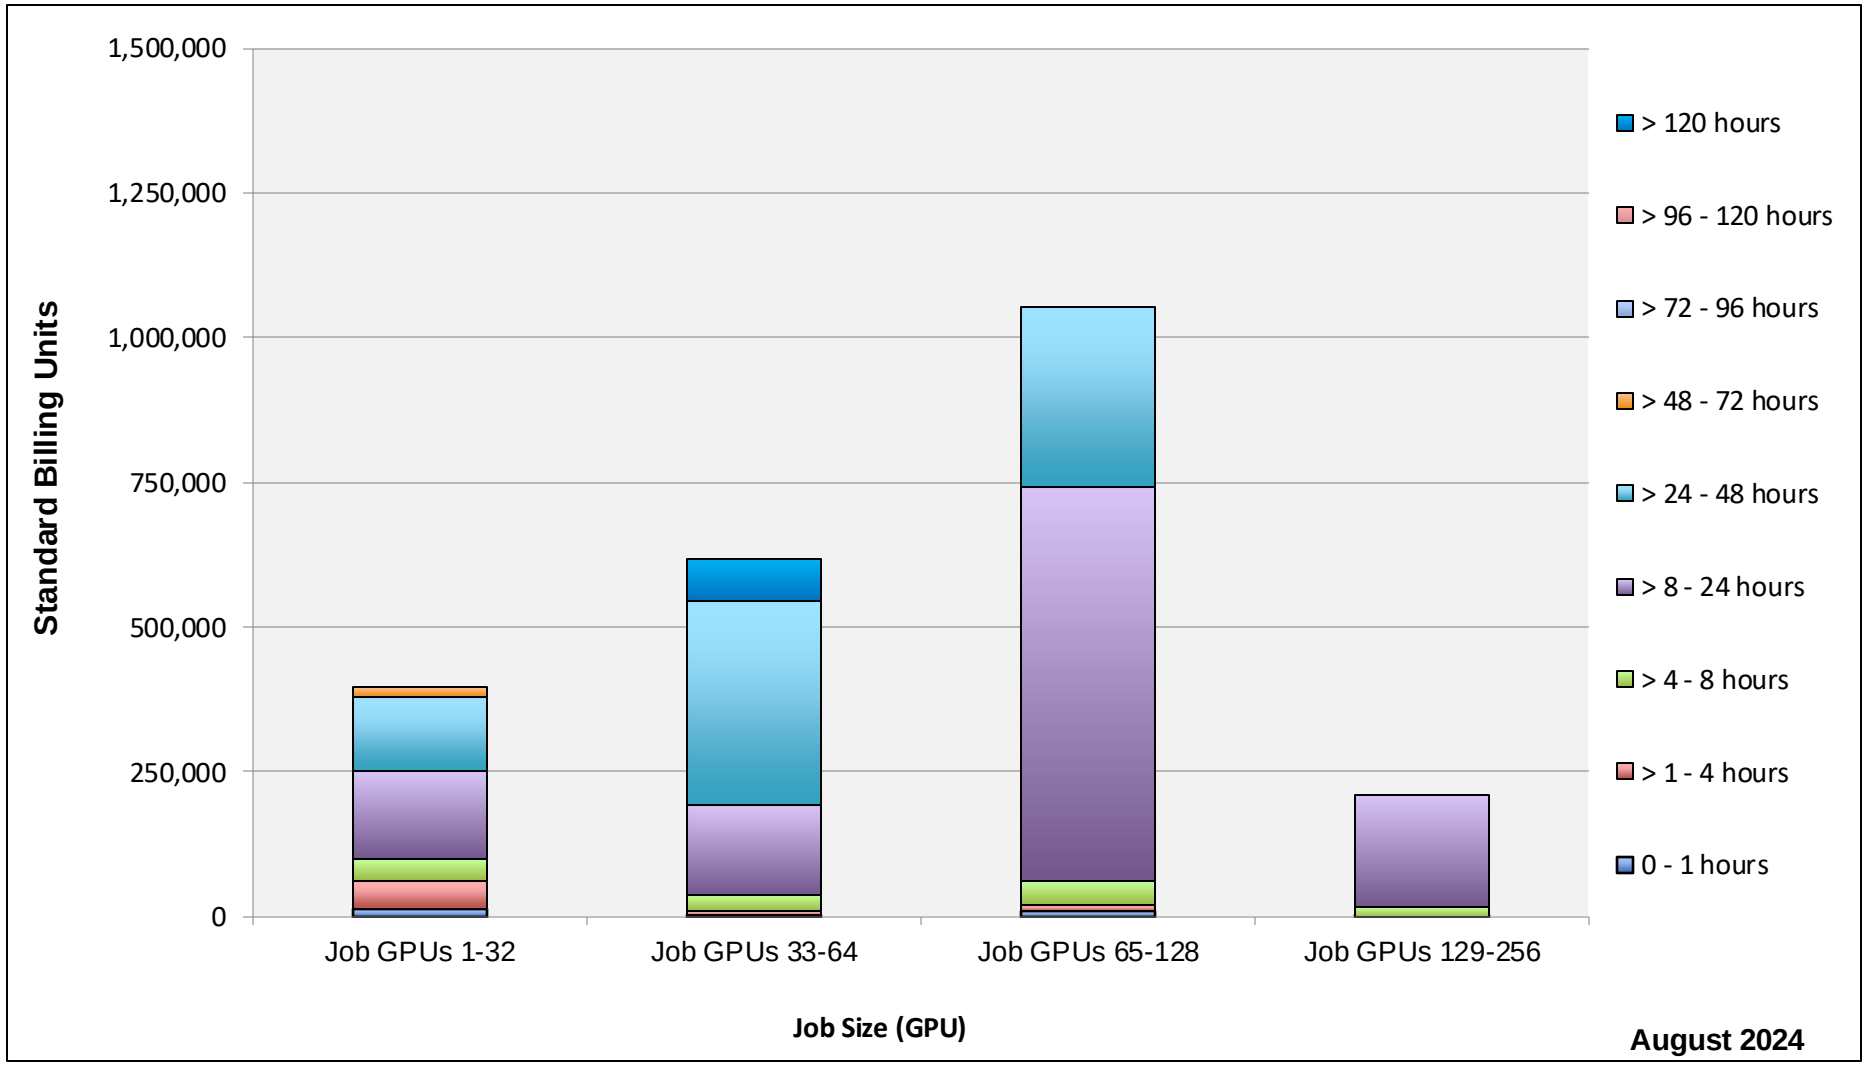
Cabeus: Monthly Utilization by Job Length



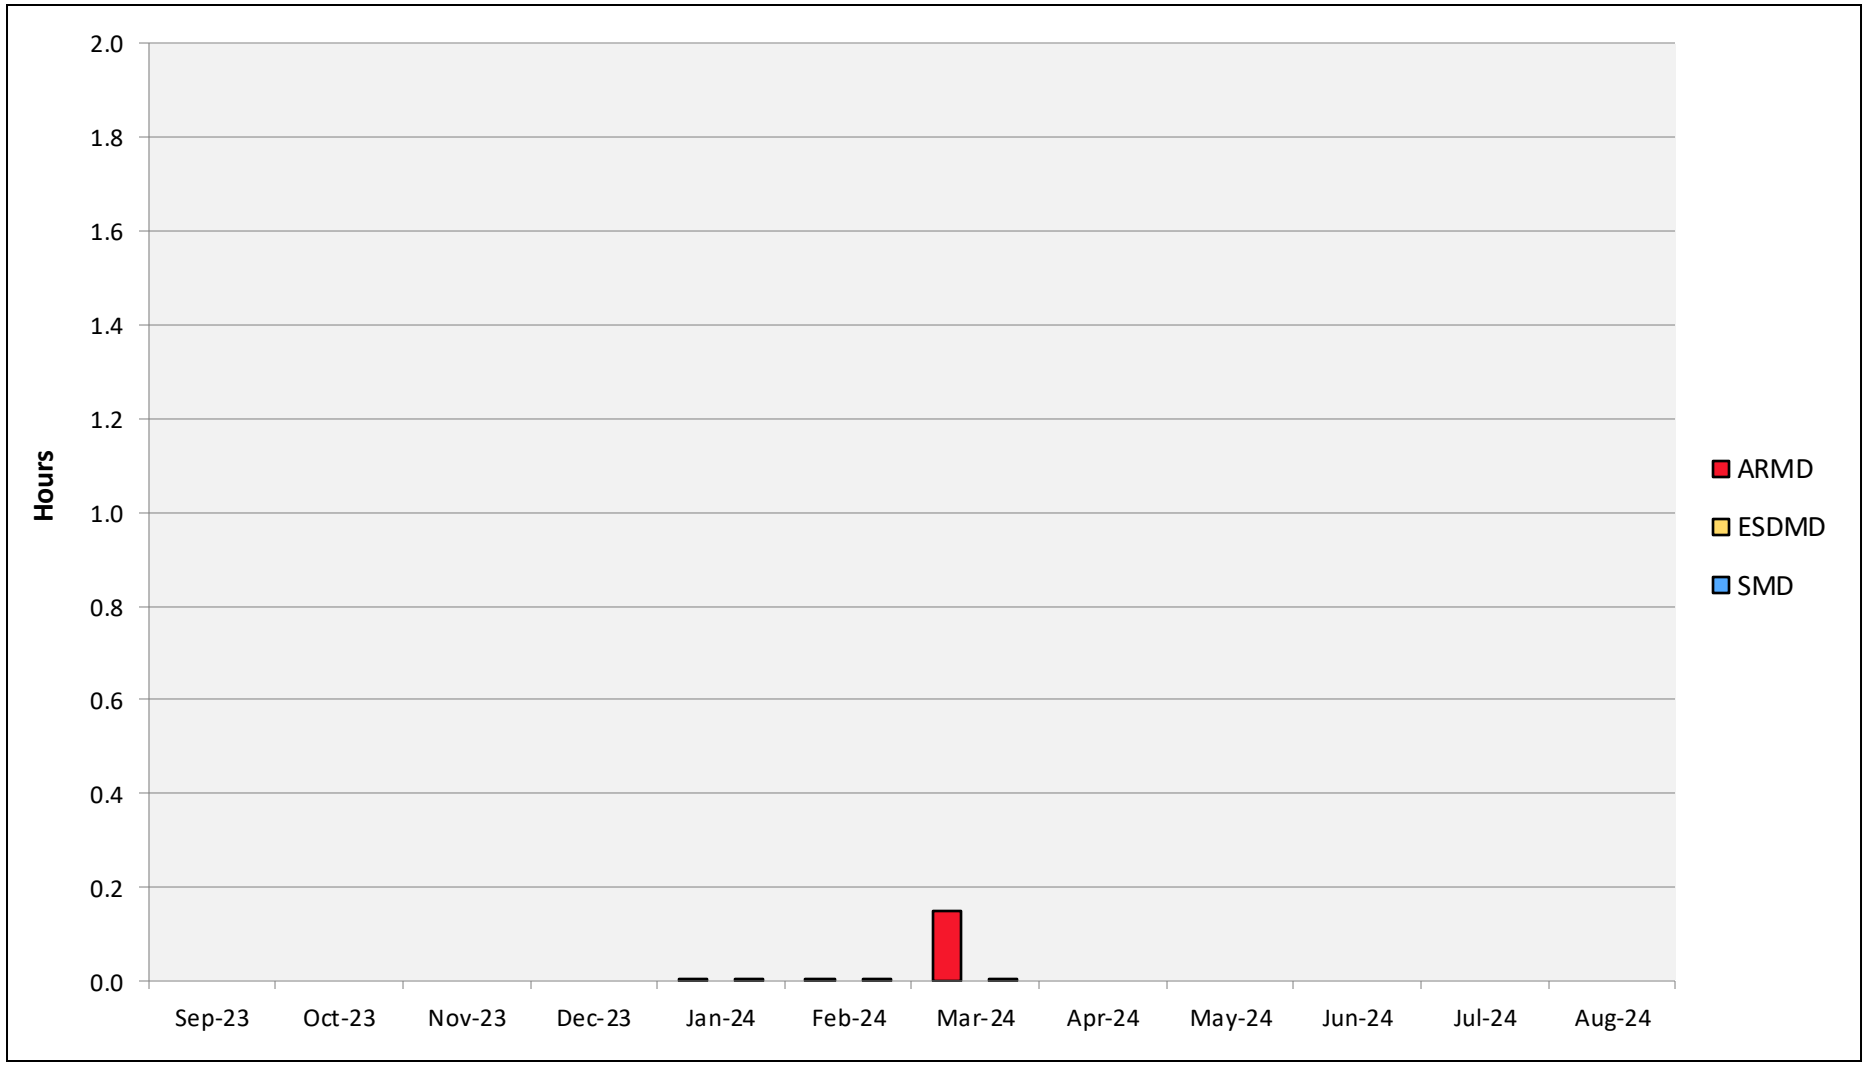
Cabeus: Monthly Utilization by Job Size



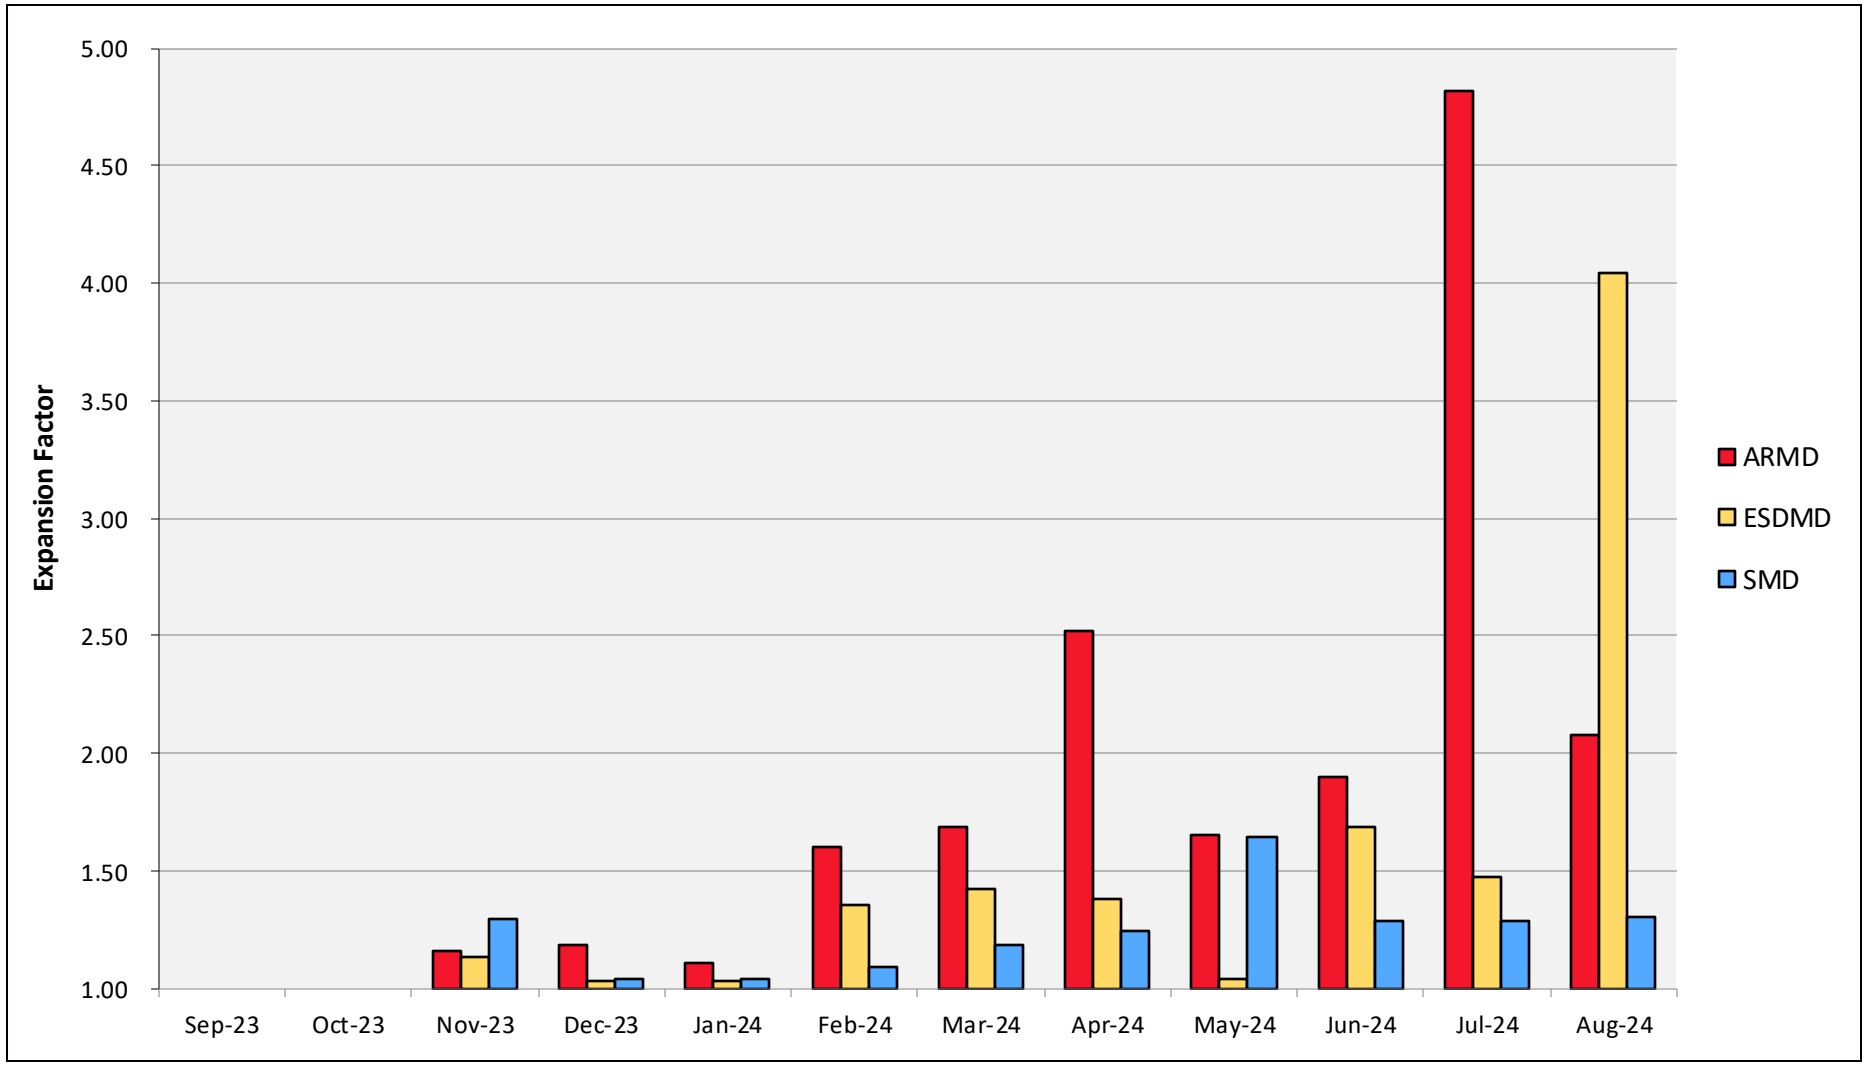
Cabeus: Monthly Utilization by Size and Length



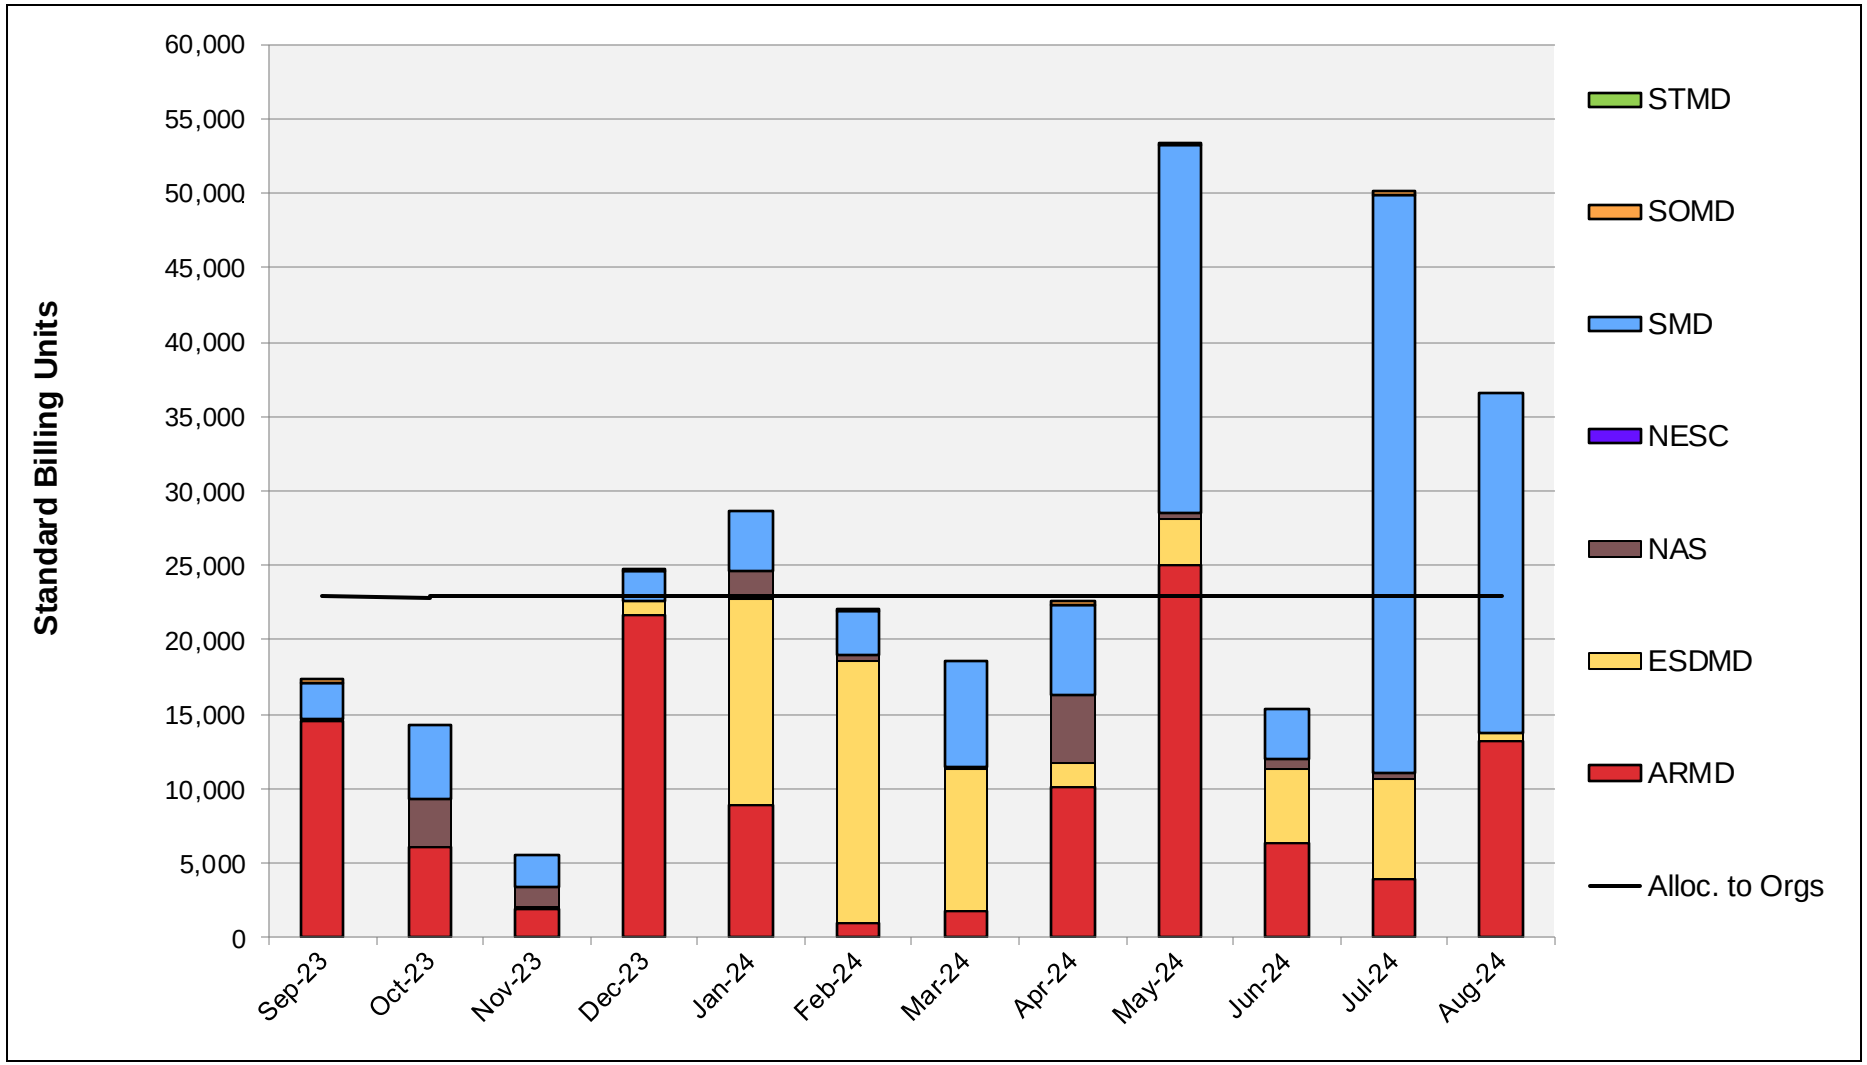
Cabeus: Average Time to Clear All Jobs



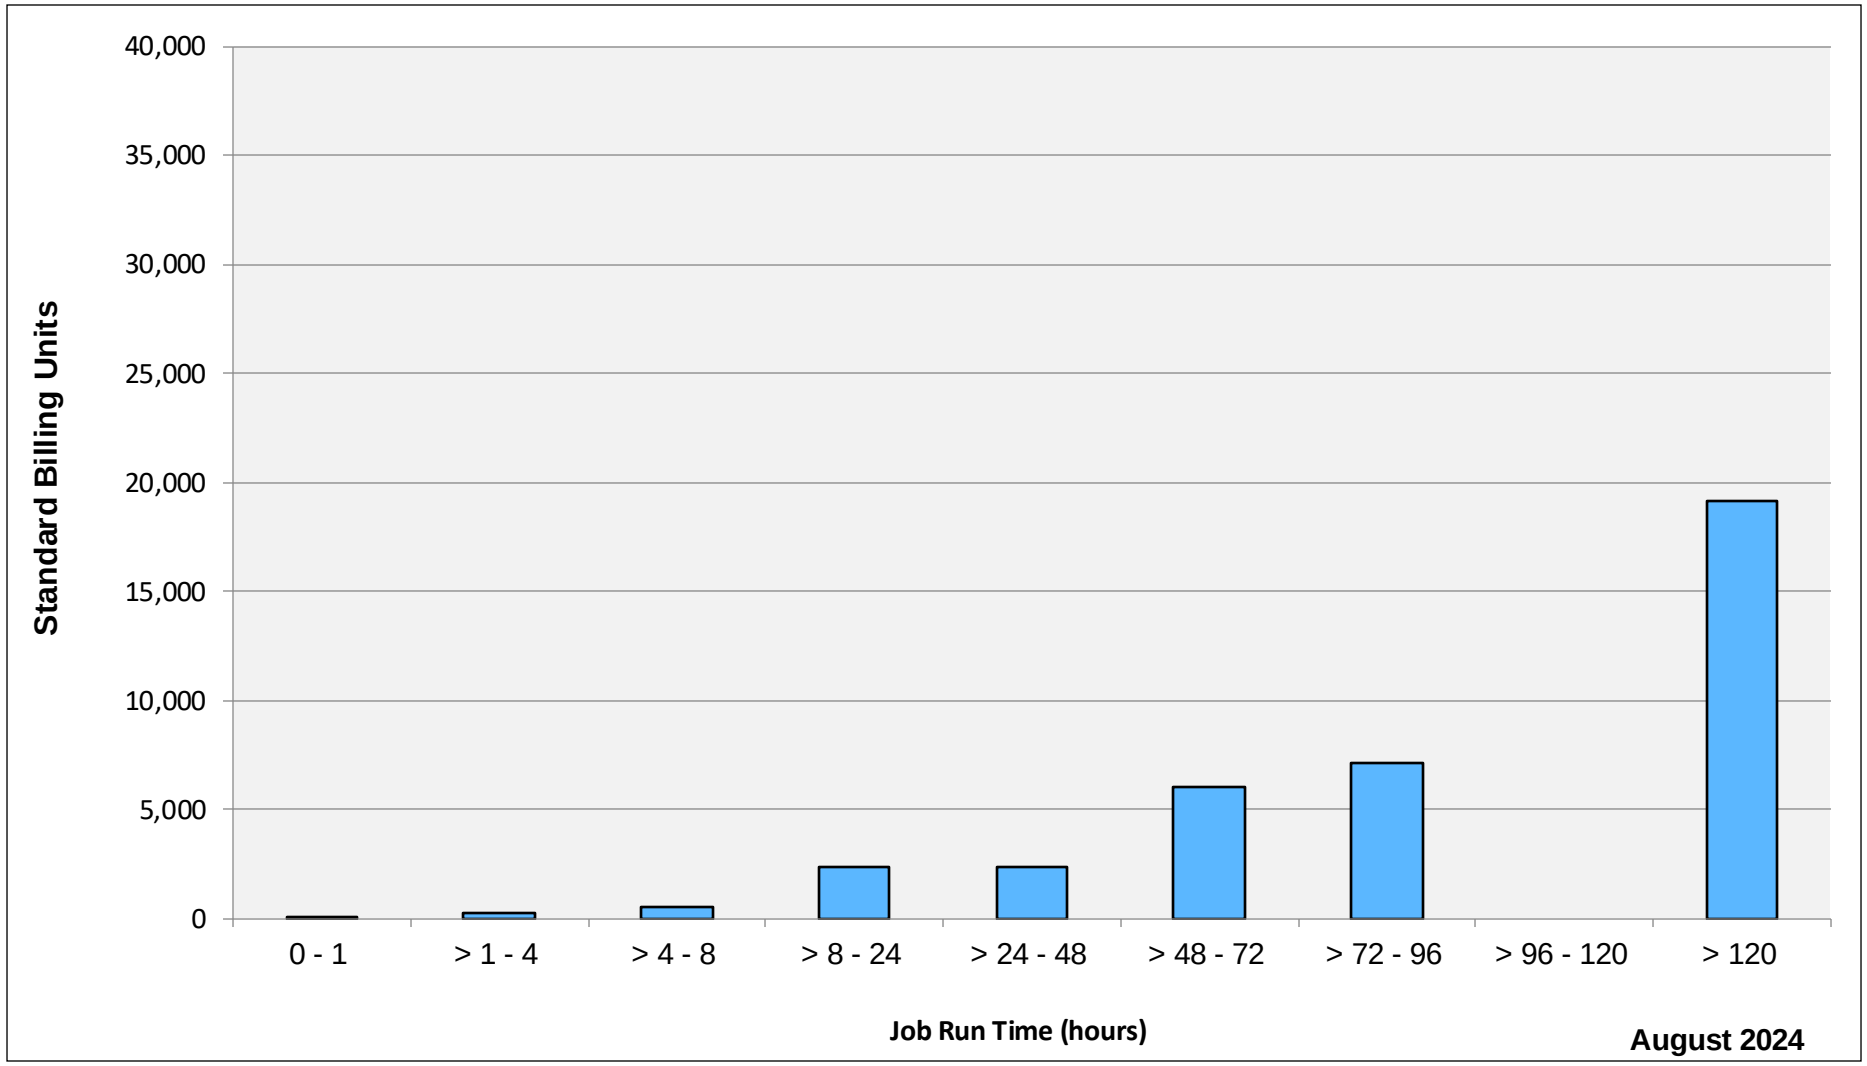
Cabeus: Average Expansion Factor



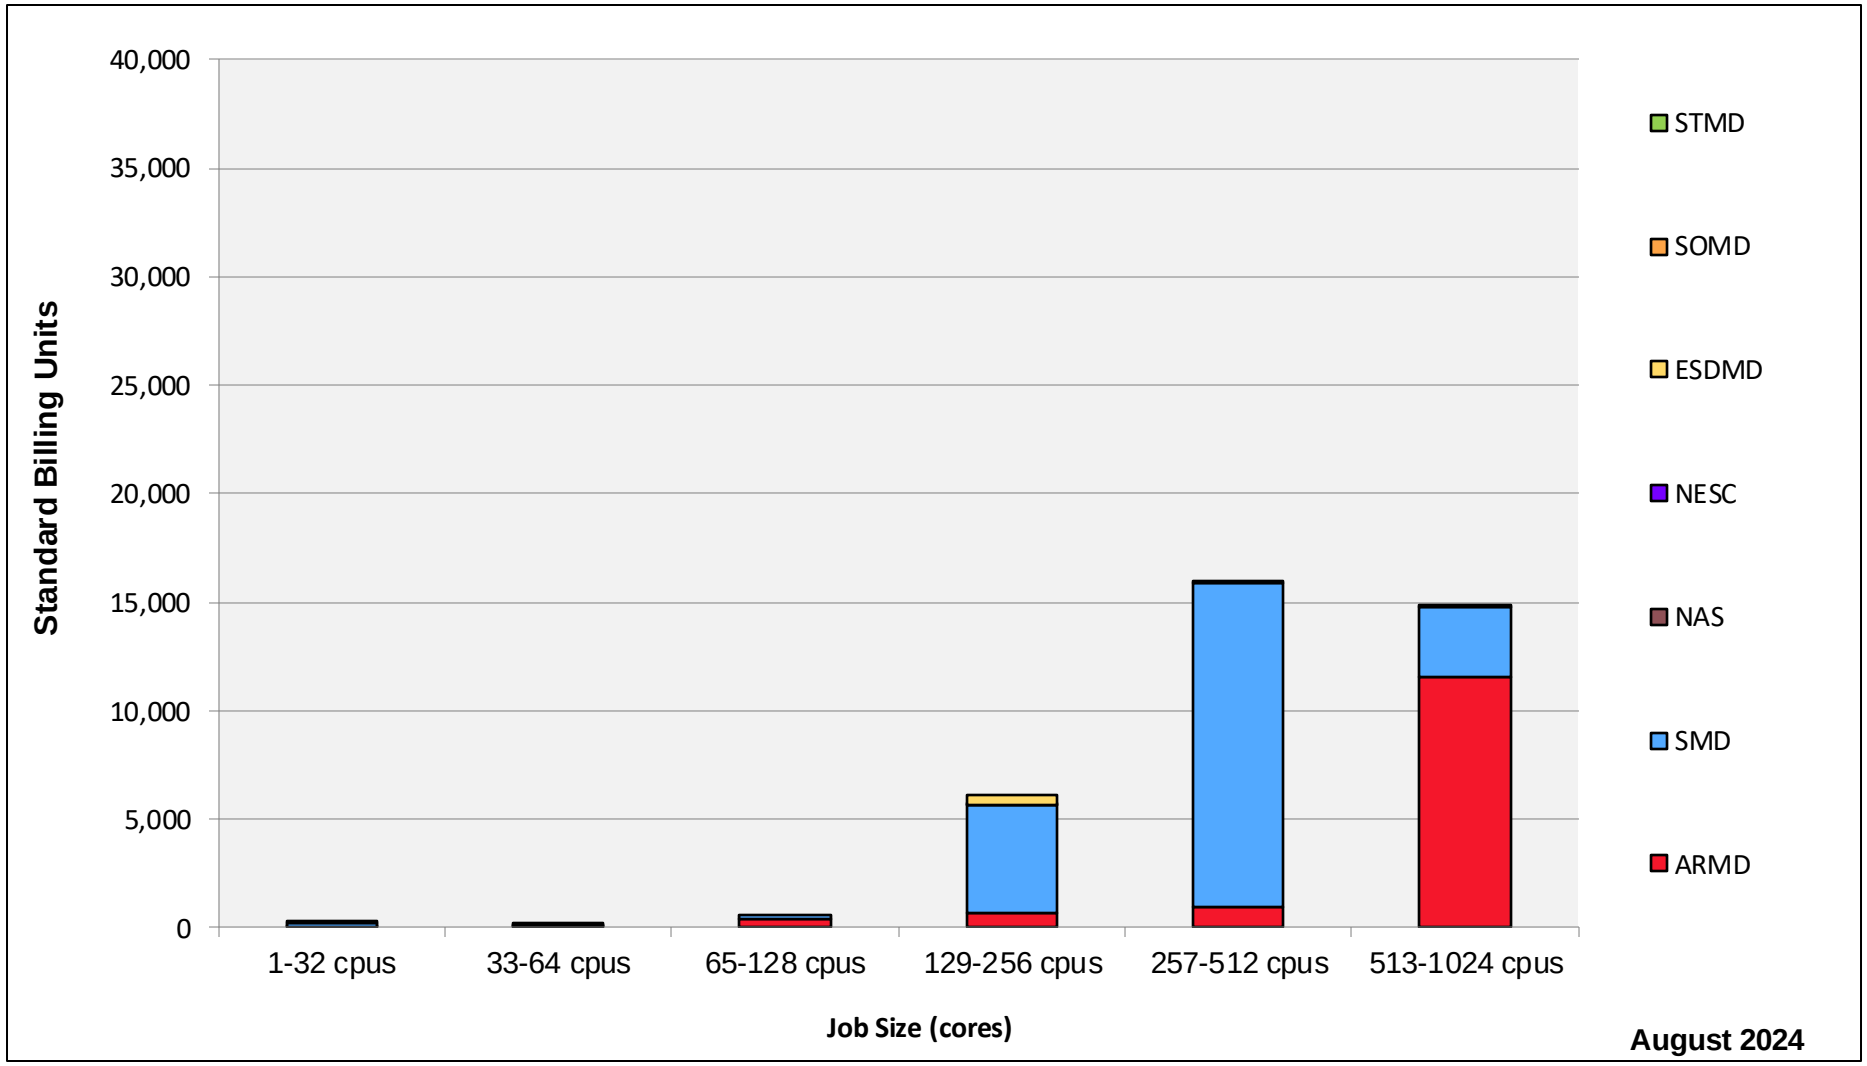
Endeavour: SBUs Reported, Normalized to 30-Day Month



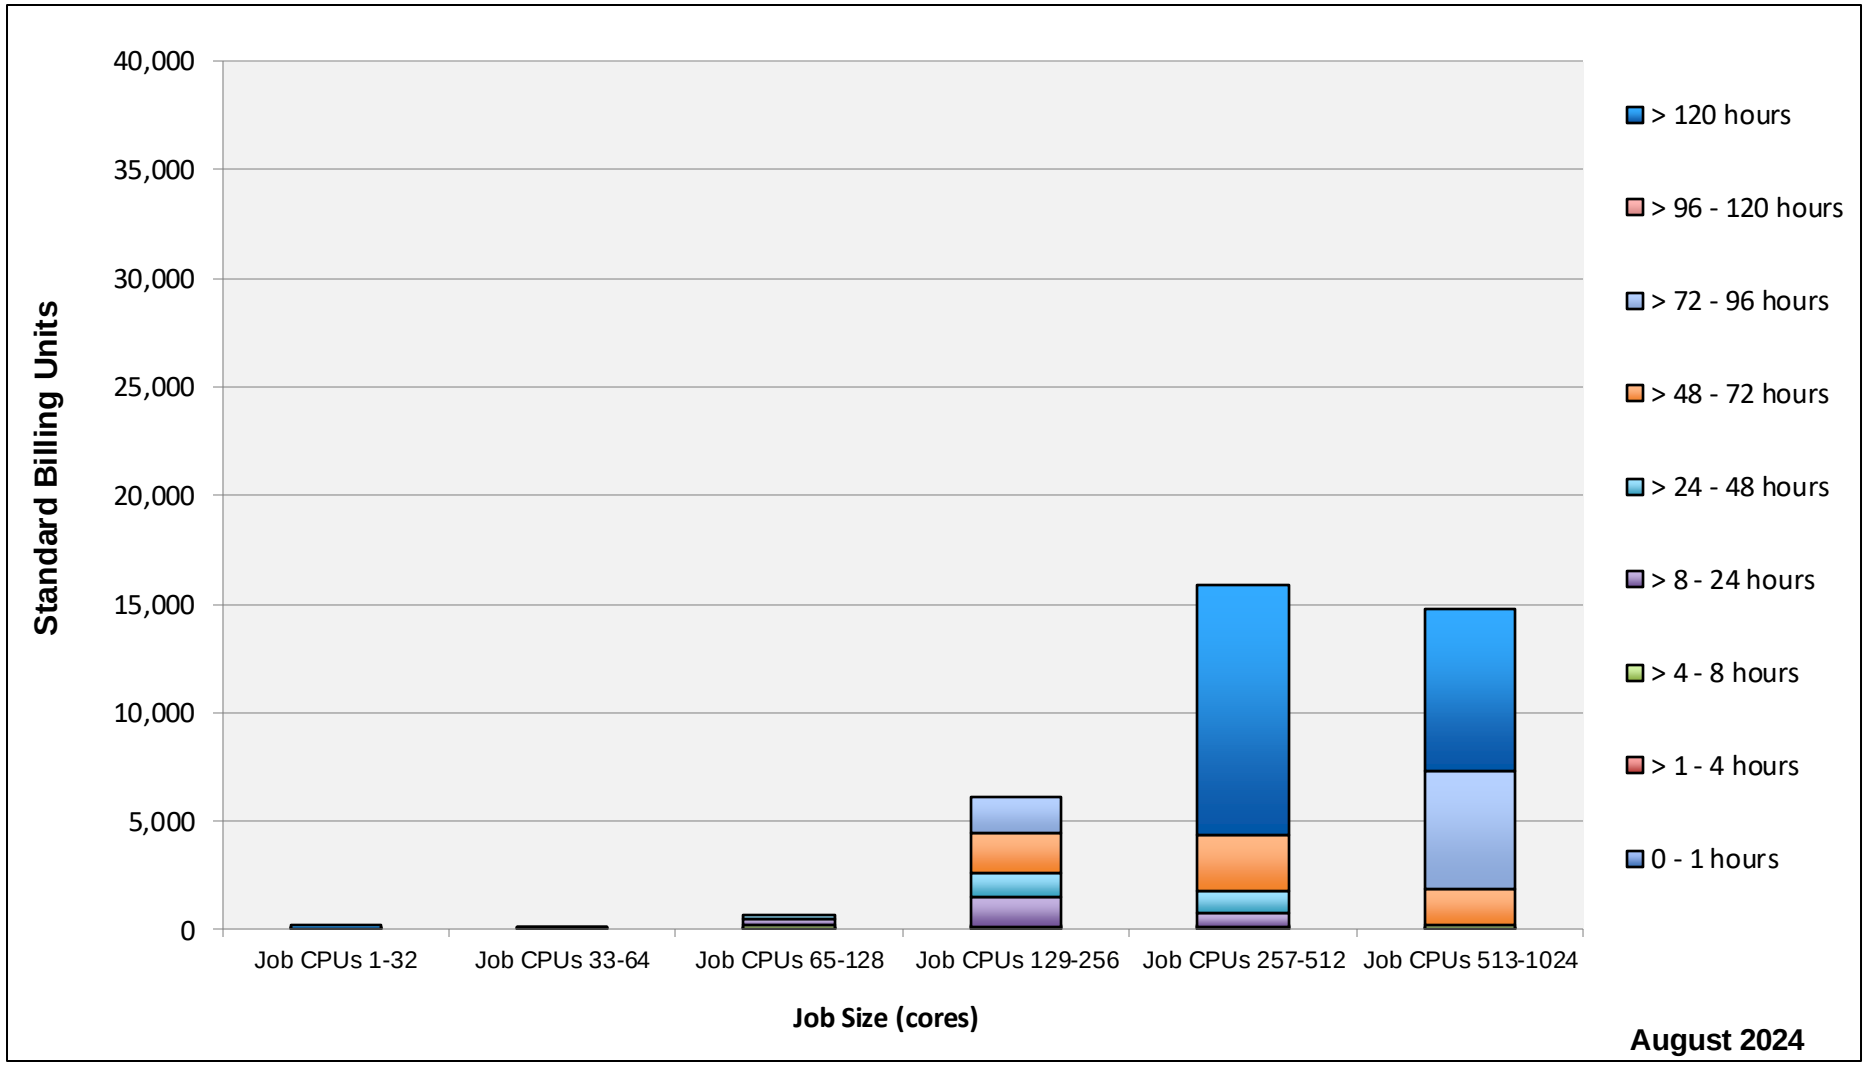
Endeavour: Monthly Utilization by Job Length



Endeavour: Monthly Utilization by Job Size



Endeavour: Monthly Utilization by Size and Length



Endeavour: Average Expansion Factor

