National Aeronautics and Space Administration

HIGH-END COMPUTING CAPABILITY PORTFOLIO

William Thigpen
NASA Advanced Supercomputing Division

November 10, 2020

Aitken Expansion Benchmarked using LINPACK and HPCG

- In preparation for installing new nodes in the Modular Supercomputing Facility to expand Aitken, HPE engineers ran the LINPACK and HPCG benchmarks on the nodes at the factory to stress test the system and identify any faulty components.
  - The Aitken expansion alone has a theoretical peak performance of 4.7 petaflops (PF); the system achieved 4.101 PF on the LINPACK benchmark and 70.6 teraflops (TF) on the HPCG benchmark.
  - With the new nodes incorporated, the expanded Aitken system would have ranked at #61 on the June 2020 TOP500 list (LINPACK) and at #36 on the June 2020 HPCG list.
  - The benchmarks were run over a three-day period; thorough and extensive diagnostic tests identified eight marginal hardware components that were replaced prior to shipping to Ames.
- The LINPACK and HPCG benchmarks are widely used to evaluate the performance of supercomputing systems and provide two complementary viewpoints on how systems perform on different workloads.
- The new Aitken nodes were handed over to the HECC Systems team on November 6, 2020 for integration and testing.

IMPACT: HECC regularly upgrades its resources to meet NASA's increasing supercomputing requirements and to significantly increase science and engineering results. Running the LINPACK and HPCG benchmarks on new or expanded systems provides a good method to identify and address system issues, thereby improving overall reliability for users.

The TOP500 list, which uses the LINPACK benchmark to rank the performance of the most powerful supercomputers in the world, is released in June and November each year.

New Nodes Significantly Expand HECC GPGPU Capability

- Thirty-eight nodes containing NVIDIA V100 General-Purpose Graphics Processing Units (GPGPUs) were installed into seven racks and handed over to the HECC Systems team on September 24, 2020 for integration and testing with Pleiades.
  - The addition will increase the number of V100 nodes available to users from 19 to 57, bringing the total number of V100 GPGPUs to 236.
  - This is the second GPU augmentation since the original 64 NVIDIA Tesla K40 nodes were installed in 2015.
- GPGPU is a traditional GPU repurposed beyond graphics and used as an alternative to conventional central processing units (CPUs), where the highly parallel nature of graphics processing can benefit codebases that take advantage of such hardware.
  - GPGPU-based pipelines, working in conjunction with traditional CPUs, have proven to be highly complementary hardware environment for scientific computing
  - With the emerging popularity and practical uses of artificial intelligence and machine learning for NASA applications, HECC users will greatly benefit from the expanded GPGPU capability.
- The new GPGPUs are expected to be in general production by November 30, 2020.

IMPACT: The expansion will triple HECC's V100 GPGPU compute capabilities, enabling further practical uses of artificial intelligence and machine learning by HECC scientists and engineers

Pleiades was augmented with 38 new nodes containing NVIDIA V100 NVLink General-Purpose Graphics Processing Units (GPGPUs), expanding HECC's GPGPU capability. *Matt Lepp, HPE*

HECC Hosts NASA GPU Hackathon 2020

- HECC and NVIDIA jointly organized the NASA GPU Hackathon 2020 to bring together application developers and computer experts to help get important NASA applications running effectively on GPU nodes. The Hackathon was a virtual event and took place over four days from September 28 to October 7.
  - Nine teams from three NASA centers (ARC, LaRC, GRC) and supporting institutions participated. They brought a variety of Aerosciences/CFD applications and used several acceleration APIs, including OpenACC, CUDA, and Kokkos.
  - The participants used Microsoft Teams and Slack for collaboration. In addition to using breakout rooms for working with their expert mentors. A key to the success was the pairing of teams with expert mentors from industry and federal agencies.
  - All the teams met together in a daily “scrum” session to present their progress.
  - HECC provided five Pleiades nodes containing NVIDIA V100 GPUs for the teams to use during the event.
- By the end of the Hackathon, most teams had made good progress and achieved considerable performance improvements on both GPUs and CPUs:
  - Working with the mentors allowed a team with no prior GPU experience to complete a first port of a time-critical loop to a GPU.
  - Code review and GPU performance analysis enabled a team of expert CUDA programmers to restructure their algorithm, yielding a 5-times speed-up.
  - Another team was able to speed up some of their CUDA kernels by a factor of 20, which directly translated into their production code.
- HECC will continue to coordinate GPU hackathons and is considering hosting another event in 2021.

IMPACT: The NASA GPU Hackathon allowed the teams to gain experience porting and optimizing their codes for GPUs potentially increasing the performance of their applications.

GPU Hackathon group photo showing 30 of the 60 participants.
Image courtesy of NVIDIA

New Allocation Period for Supercomputer Time Begins for NASA Mission Directorates

- The new allocation period for all NASA mission directorates began October 1, 2020.
  – Mission directorates were awarded approximately 169 million Standard Billing Units (SBUs) for 530 projects.
  – The Aeronautics Mission Directorate (ARMD) awarded just under 52 million SBUs to 148 projects.
  – The Human Exploration and Operations Mission Directorate, Space Technology Mission Directorate, and NASA Engineering and Safety Center collectively awarded more than 45 million SBUs to 67 projects.
  – The Science Mission Directorate awarded just under 72 million SBUs, with a surplus of 7.3 million SBUs to be awarded to future projects quarterly to 315 projects.
- The new allocation period is an opportunity for each organization to assess demands for computing time and to rebalance allocations to meet computing needs for the next year.

IMPACT: NASA programs and projects periodically review the distribution of supercomputer time to assess demand for resources and assure consistency with mission-specific goals and objectives.

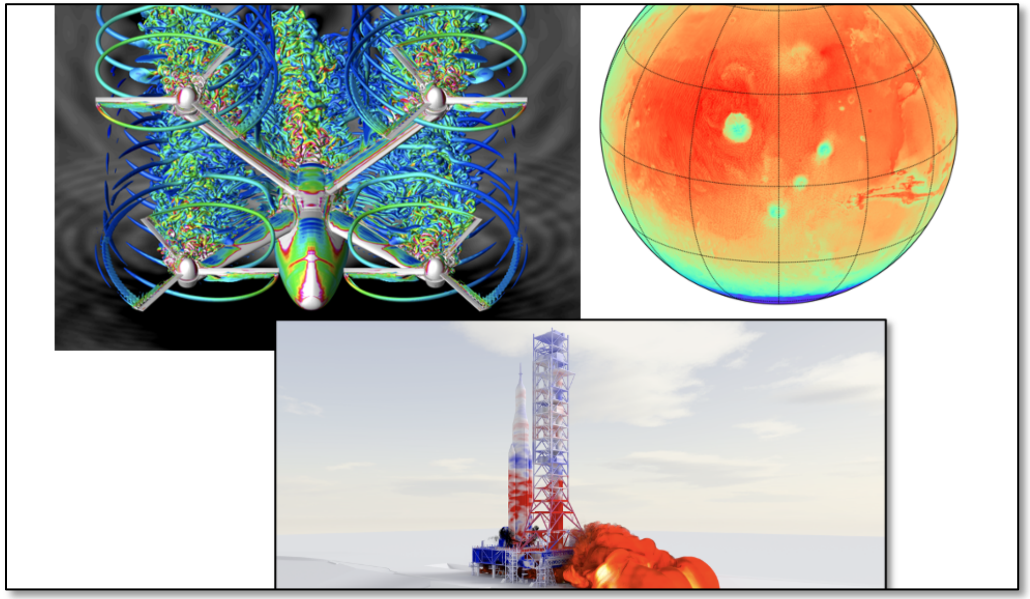

Images from simulations representing the variety of projects supported by HECC resources.

NAS Data Portal Enables Sharing of High-Volume Datasets

- The web-based NAS Data Portal is designed for sharing high-volume computational datasets, to promote collaboration, data dissemination, and sharing of science results.
  - Datasets can be shared publicly or with specific collaborators.
  - Shared dataset files are safely and securely staged to a public interface outside of the HECC security enclave.
  - There are no restrictions on the amount of data in shared datasets.
- Currently, 15 datasets are publicly available on the Data Portal in various disciplines—primarily in Earth and geosciences, planetary geophysics, and heliophysics. The datasets range from 400 kilobytes to 4 petabytes in size and are comprised of anywhere from a single file to 185 million files.
- New user-friendly capabilities are being developed to enable users to self-publish thousands of files per day.
- Near-future enhancements include more hardware servers for greater stability, throughput, and load-sharing.
- The long-term vision for the Data Portal is to serve as a more general data exploration and discovery platform.

IMPACT: The NAS Data Portal has significantly enhanced the dissemination of computational results, freeing researchers from the technical concerns typically associated with sharing massive volumes of data.

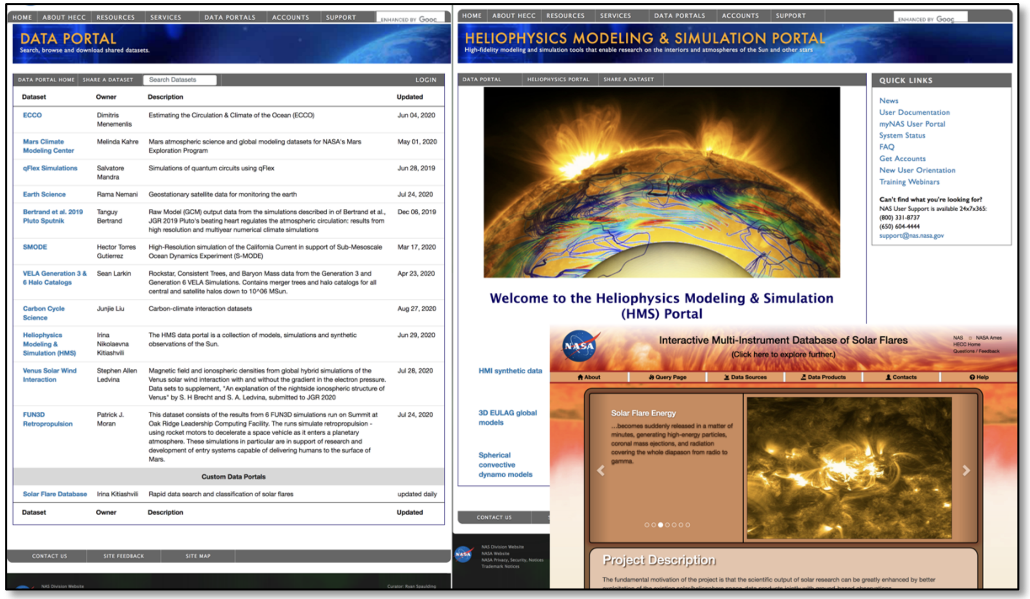

Left: The NAS Data Portal splash page (https://data.nas.nasa.gov) allows access to the individual datasets and customized portals. Right: The Heliophysics and Solar Flare Database data portals are shown. *Glenn Deardorff, NASA/Ames*

Simulating the Water Cycle and Cloud Formation on Mars*

- Using a global climate model (GCM) developed at NASA's Mars Climate Modeling Center, researchers at NASA Ames ran simulations on the Pleiades supercomputer to investigate the water cycle and how water ice clouds affect the Martian climate. Adapted to Mars, the GCM yielded critical insight into the importance of the water cycle and water ice clouds in the Martian climate system.
  - The GCM includes a cubed-sphere finite volume dynamical core developed at the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory (NOAA/GFDL).
  - This modern dynamical core is highly parallelized, scalable, and flexible, which enables global, multi-year, high-resolution modeling of Mars' atmosphere and the ability to investigate a multitude of open questions about water, clouds, and the Martian climate.
- When comparing the GCM results with spacecraft observations, the researchers found that radiative effects of Martian clouds, though thin compared to Earth's clouds, significantly change the thermal structure of the atmosphere and the intensity of its global-scale wind systems. In this way, clouds help control the movement of water around the planet.
- In the distant past when orbital properties were different, Mars may have been much warmer than it is now: the clouds could have provided a strong greenhouse effect and warmed the surface by tens of degrees Kelvin.

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

IMPACT: These studies directly address NASA's goals to understand how bodies in the solar system evolved, and they assist the agency's mission planning activities—including the Mars 2020 Mission, which is scheduled to land on February 18, 2021.

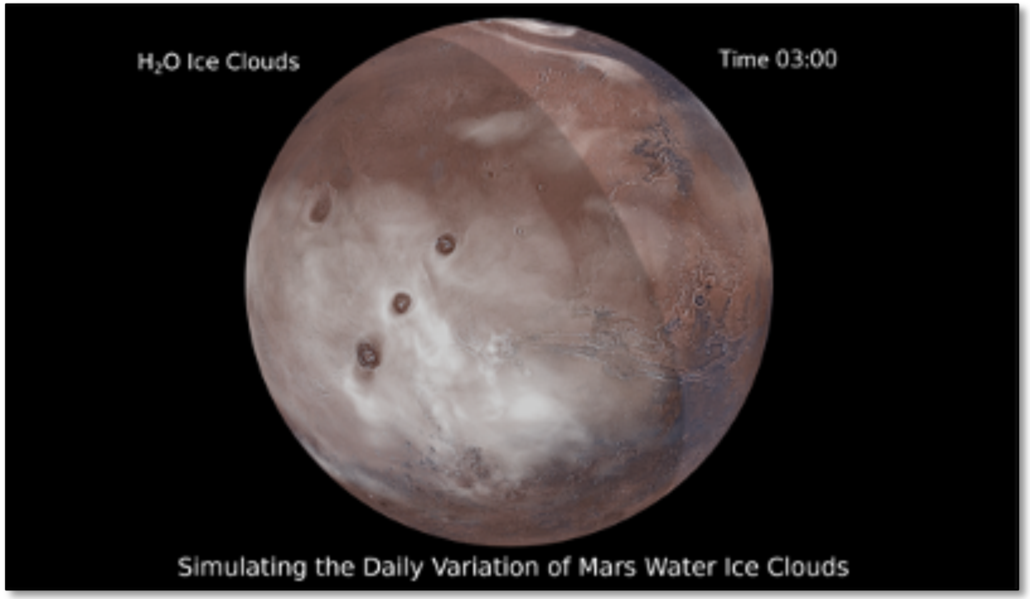

Visualization showing the simulated diurnal cycle of water ice clouds and the environment over the Tharsis Montes volcanoes, during the northern summer on Mars. The simulation was run on NASA's Pleiades supercomputer. *Alex Kling, David Ellsworth, NASA/Ames*

Discovering Distinct Cycles in the Global Solar Dynamo*

- Scientists at the University of Colorado Boulder are running global simulations of the Sun's interior magnetism on the Pleaides supercomputer to better understand distinct features of the global solar dynamo and solar cycle. The project goal is to characterize the nonlinear dynamo responsible for the cycles in solar magnetic activity.
- The team used the open-source, nonlinear magnetohydrodynamic (MHD) Rayleigh code to produce a new class of simulations to simulate a magnetized fluid in a rotating, convecting spherical shell—offering for the first time a systematic, integrated perspective on disparate solar phenomena.
  - The simulations reproduce regular cycles, during which globally connected "wreaths" of magnetism reverse in polarity and migrate toward the equator, two features highly reminiscent of the sunspot cycle.
  - Additionally, there is an asymmetric cycle, during which strong magnetism preferentially occupies one hemisphere at certain longitudes, reminiscent of the "active longitudes" and "active hemispheres" observed on the Sun.
- The simulations build on previous work to understand equatorward propagation in the sunspot cycle, regular polarity reversals, and active longitudes. Future work will include a stably stratified region below the convection zone and a tachocline of shear at the transition between the two regions.

IMPACT: HECC applications, expertise, and resources are required to capture the small-scale, turbulent structures associated with convective solar flows in this unique global solar dynamo simulation.

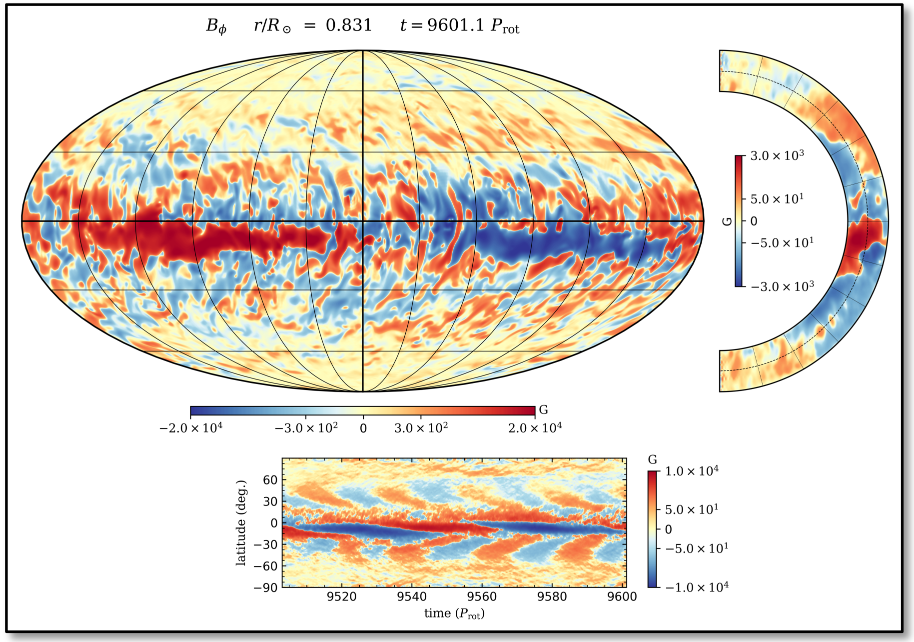

Visualization showing the temporal evolution of the azimuthal magnetic field shown for a global solar dynamo simulation. The dominant polarity reverses regularly, about once every three sunspot cycles. *Loren Matilsky, University of Colorado Boulder and JILA*

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

Papers

- "**An Earth-like Stellar Wind Environment for Proxima Centauri c**," J. Alvarado-Gomez, et al., The Astrophysical Journal Letters, vol. 902, no. 1, October 7, 2020. *
  https://iopscience.iop.org/article/10.3847/2041-8213/abb885
- "**Single-Hemisphere Dynamos in M-Dwarf Stars**," B. Brown, et al., The Astrophysical Journal Letters, vol. 902, no. 1, October 7, 2020. *
  https://iopscience.iop.org/article/10.3847/2041-8213/abb9a4
- "**Direct Numerical Simulation of a Turbulent Boundary Layer on a Flat Plate Using Synthetic Turbulence Generation**," J. Wright, et al., arXiv:2010.04277 [physics.flu-dyn], October 8, 2020. *
  https://arxiv.org/abs/2010.04277
- "**A Case Study on the Origin of Near-Earth Plasma**," A. Glocer, et al., Journal of Geophysical Research: Space Physics, vol. 125, issue 11, October 10, 2020. *
  https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020JA028205
- "**Convectively Driven 3D Turbulence in Massive Star Envelopes. I. A 1D Implementation of Diffusive Radiative Transport**," W. Schultz, L. Bildsten, Y.-F. Jiang, The Astrophysical Journal, vol. 902, no. 1, October 13, 2020. *
  https://iopscience.iop.org/article/10.3847/1538-4357/abb405/meta

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

Papers (cont.)

- “**Pareto-Efficient Combustion Framework for Predicting Transient Ignition Dynamics in Turbulent Flames: Application to a Pulsed Jet-in-Hot-Coflow Flame**,” Q. Douasbin, M. Ihme, C. Arndt, Combustion and Flame, vol. 223, published online October 14, 2020. *
  https://www.sciencedirect.com/science/article/pii/S0010218020304156
- “**Magnetic Reconnection in a Charged, Electron-Dominant Current Sheet**,” S. Lu, et al., Physics of Plasmas, vol. 27, October 14, 2020. *
  https://aip.scitation.org/doi/abs/10.1063/5.0020857
- “**On the Surface Current Measurement Capabilities of Spaceborne Doppler Scatterometry**,” A. Wineteer, H. Torres, E. Rodriguez, Geophysical Research Letters, vol. 47, issue 21, October 14, 2020.
  https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020GL090116
- “**Direct Numerical Simulation of a Turbulent Boundary Layer with Strong Pressure Gradients**,” R. Balin, K. Jansen, arXiv:2010.08577 [physics.fluid-dyn], October 16, 2020. *
  https://arxiv.org/abs/2010.08577
- “**Equilibrium Eccentricity of Accreting Binaries**,” J. Zrake, C. Tiede, A. MacFadyen, Z. Haimen, arXiv:2010.09707 [astro-ph.HE], October 19, 2020. *
  https://arxiv.org/abs/2010.09707

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

Papers (cont.)

- "**The CARMENES Search for Exoplanets Around M Dwarfs**," G. Nowak, et al., Astronomy & Astrophysics, vol. 642, October 19, 2020. *
https://www.aanda.org/articles/aa/abs/2020/10/aa37867-20/aa37867-20.html
- "**A Numerical Study of Hypoxia in Chesapeake Bay Using an Unstructured Grid Model: Validation and Sensitivity to Bathymetry Representation**," X. Cai, et al., Journal of American Water Resources Association, October 20, 2020. *
https://onlinelibrary.wiley.com/doi/abs/10.1111/1752-1688.12887
- "**DeepGalaxy: Deducing the Properties of Galaxy Mergers from Images Using Deep Neural Networks**," M. Cai, et al., arXiv:2010.11630 [astro-ph.IM], October 22, 2020. *
https://arxiv.org/abs/2010.11630
- "**Spitzer Reveals Evidence of Molecular Absorption in the Atmosphere of the Hot Neptune LTT 9779b**," D. Dragomir, et al., The Astrophysical Journal Letters, vol. 903, no. 1, October 26, 2020. *
https://iopscience.iop.org/article/10.3847/2041-8213/abbc70/meta
- "**Phase Curves of Hot Neptune LTT 9779b Suggest a High-Metallicity Atmosphere**," I. Crossfield, et al., The Astrophysical Journal Letters, vol. 903, no. 1, October 26, 2020. *
https://iopscience.iop.org/article/10.3847/2041-8213/abbc71/meta

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

Papers (cont.)

- “**Evolution of a Foreshock Bubble in the Midtail Foreshock and Impact on the Magnetopause: 3D Global Hybrid Simulation**,” C.-P. Wang, X. Wang, T. Liu, Y. Lin, Geophysical Research Letters, published online October 26, 2020. * https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020GL089844
- “**No Evidence for [CII] Halos or High-Velocity Outflows in z>6 Quasar Host Galaxies**,” M. Novak, et al., arXiv:2010.14875 [astro-ph.GA], October 28, 2020. * https://arxiv.org/abs/2010.14875

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

News and Events

- **3D Model Seeks to Explain Mysterious Hexagon Storm on Saturn**, *The Harvard Gazette*, October 5, 2020—The hexagonal flow pattern near the north pole of Saturn is a striking example of turbulent self-organization that has been puzzling scientists for decades. Researchers from Harvard University used a complex 3D model run on the high-performance computing resources at the NASA Advanced Supercomputing facility to better understand the physics behind the phenomenon.
  https://news.harvard.edu/gazette/story/2020/10/3d-model-seeks-to-explain-mysterious-hexagon-storm-on-saturn/
- **Exploration and Scientific Discovery at NASA**, *Scientific Computing World*, October 15, 2020—William Thigpen, assistant division chief of HPC operations for the NASA Advanced Supercomputing Division, discusses the role the facility plays in helping the organization to meet its goals.
  https://www.scientific-computing.com/analysis-opinion/exploration-and-scientific-discovery-nasa

News and Events: Social Media

- **Coverage of NAS Stories**
  - X-59 Sonic Boom Models:
    - NAS: Twitter 4 retweets, 14 favorites
    - NASA Ames: Twitter 28 retweets, 148 favorites; Facebook 116 likes, 23 shares
    - NASA Supercomputing: Facebook 241 users reached, 17 engagements, 6 likes, 2 shares
  - Hexagonal Flow Patterns at Saturn’s North Pole:
    - NAS: Twitter 10 retweets, 22 favorites
    - NASA Supercomputing: Facebook 441 users reached, 35 engagements, 12 likes, 5 shares

HECC Utilization

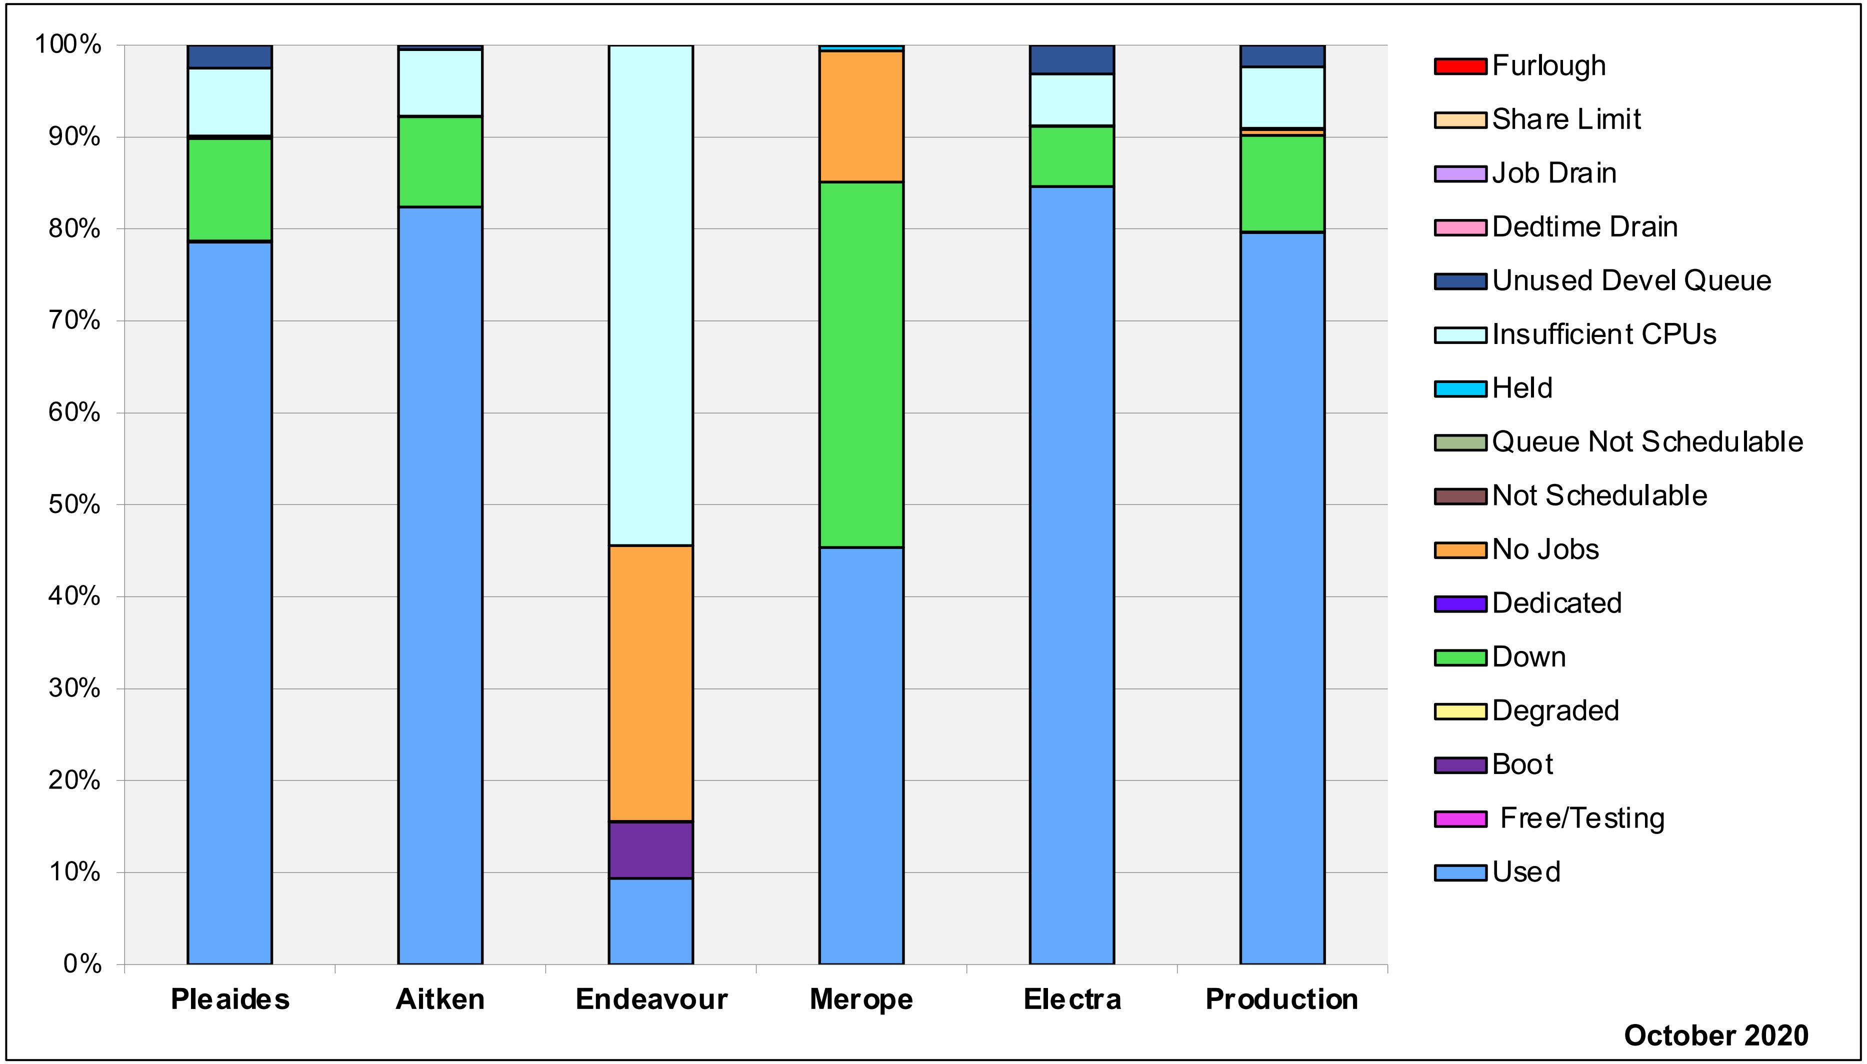

HECC Utilization Normalized to 30-Day Month

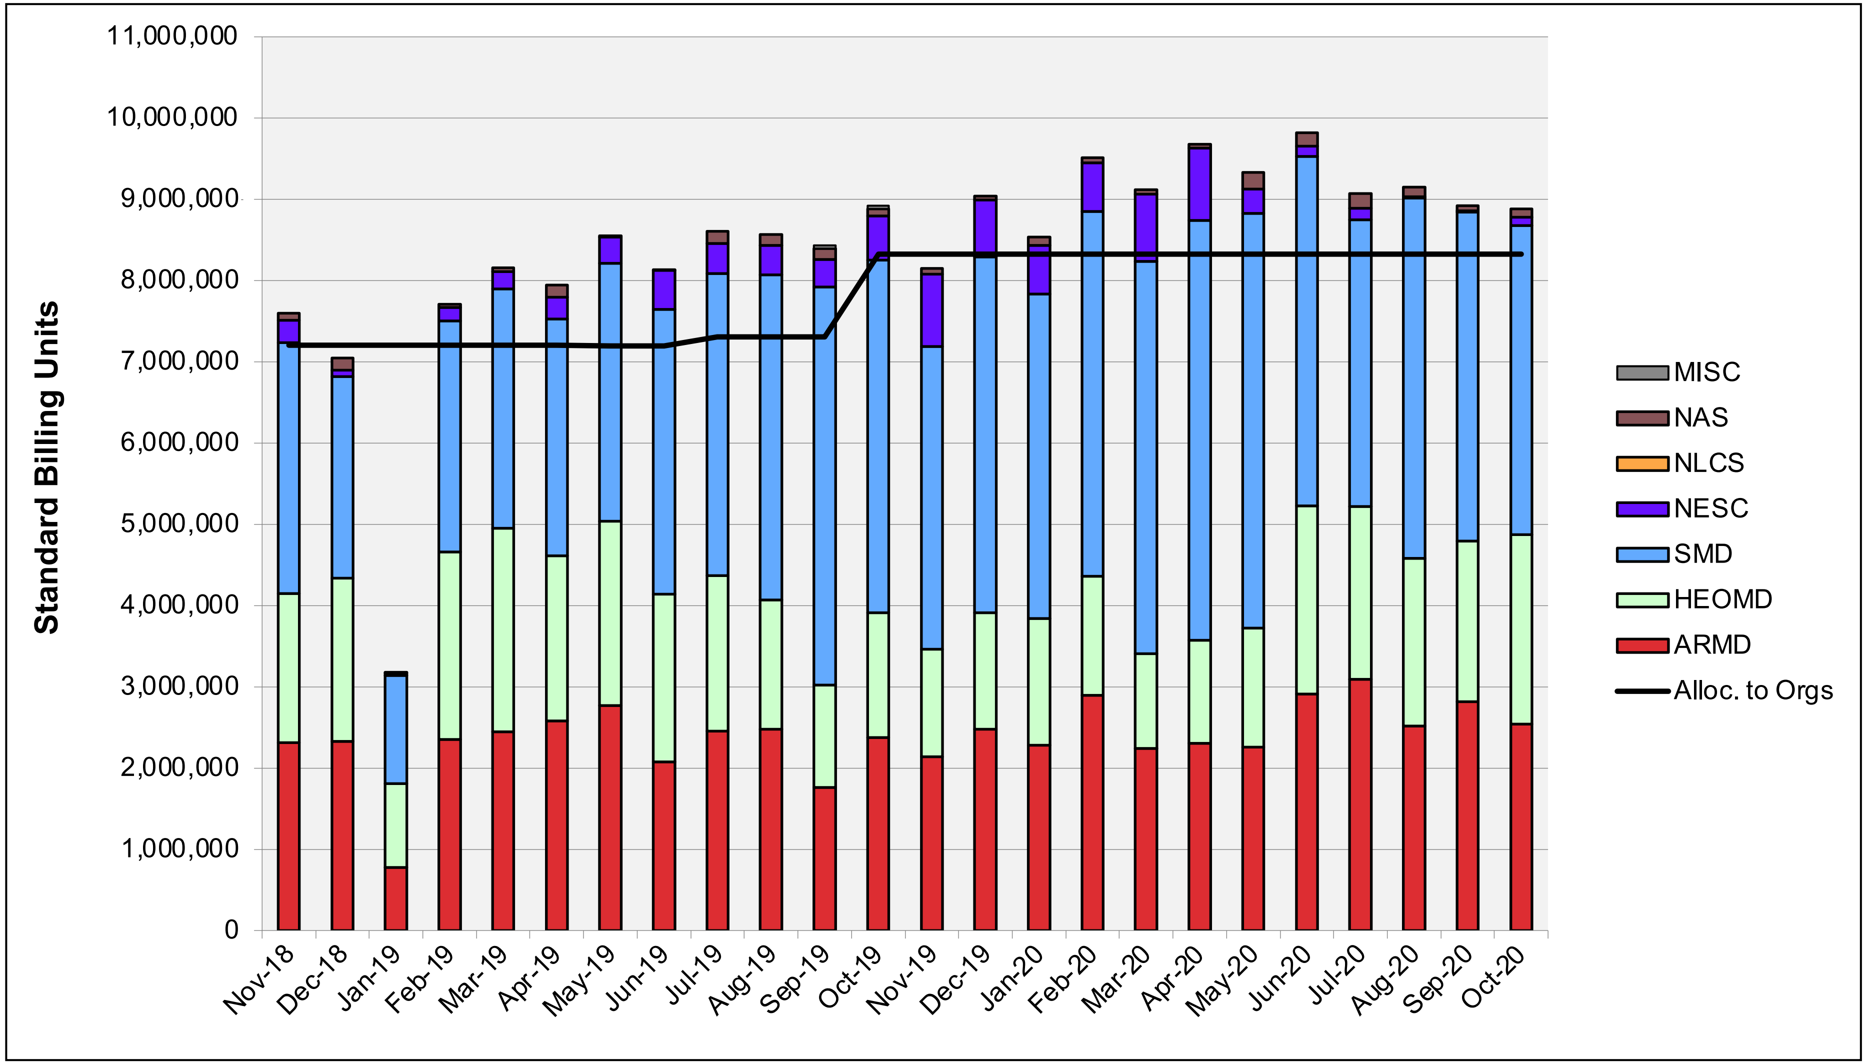

HECC Utilization Normalized to 30-Day Month

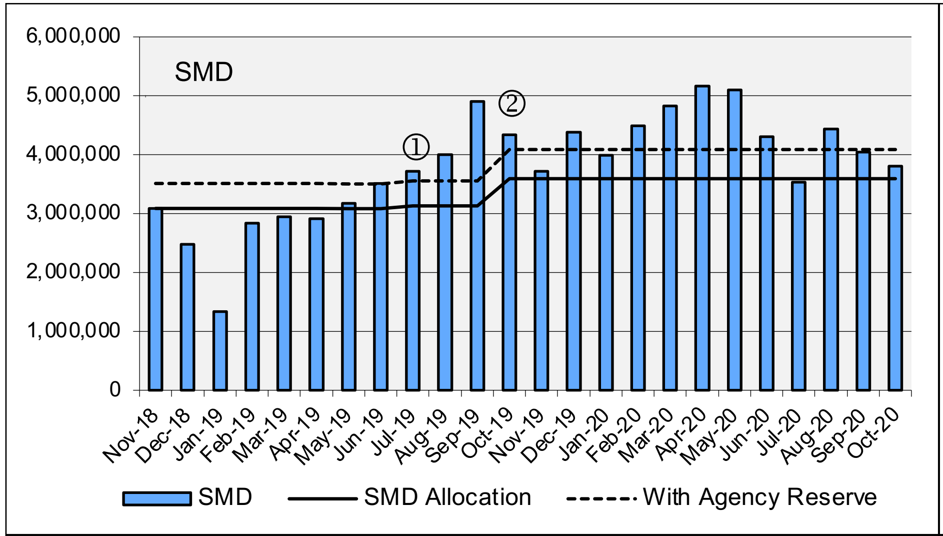

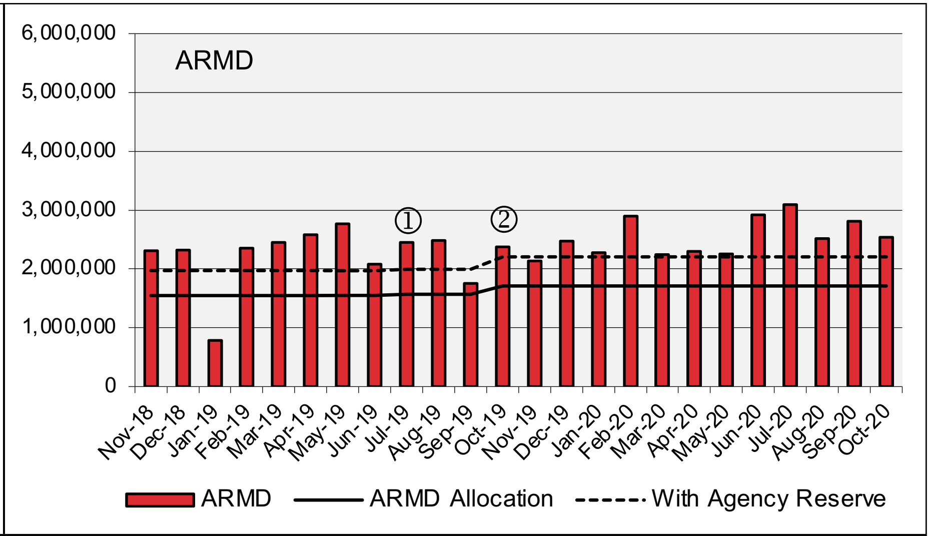

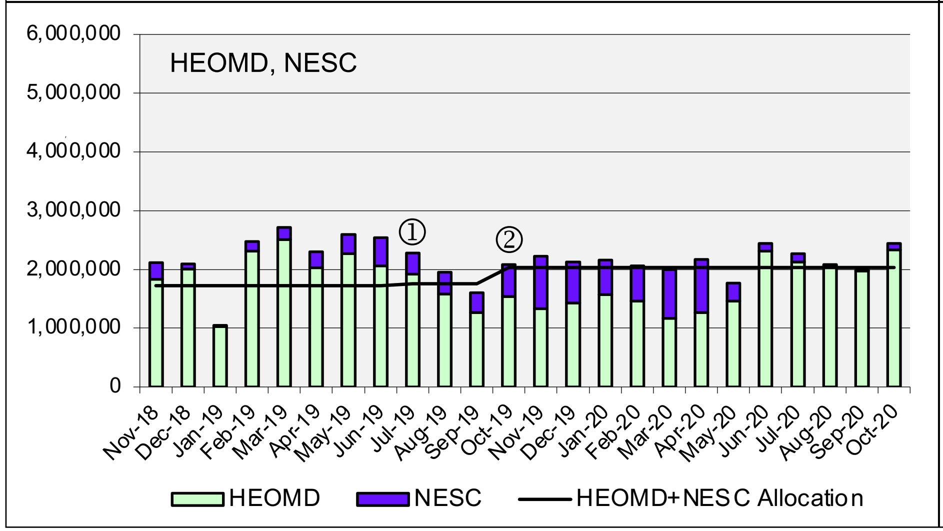

① Skylake Tesla GPU V100 Nodes installed

② 4 Cascade Lake E cells introduced into Aitken

Tape Archive Status

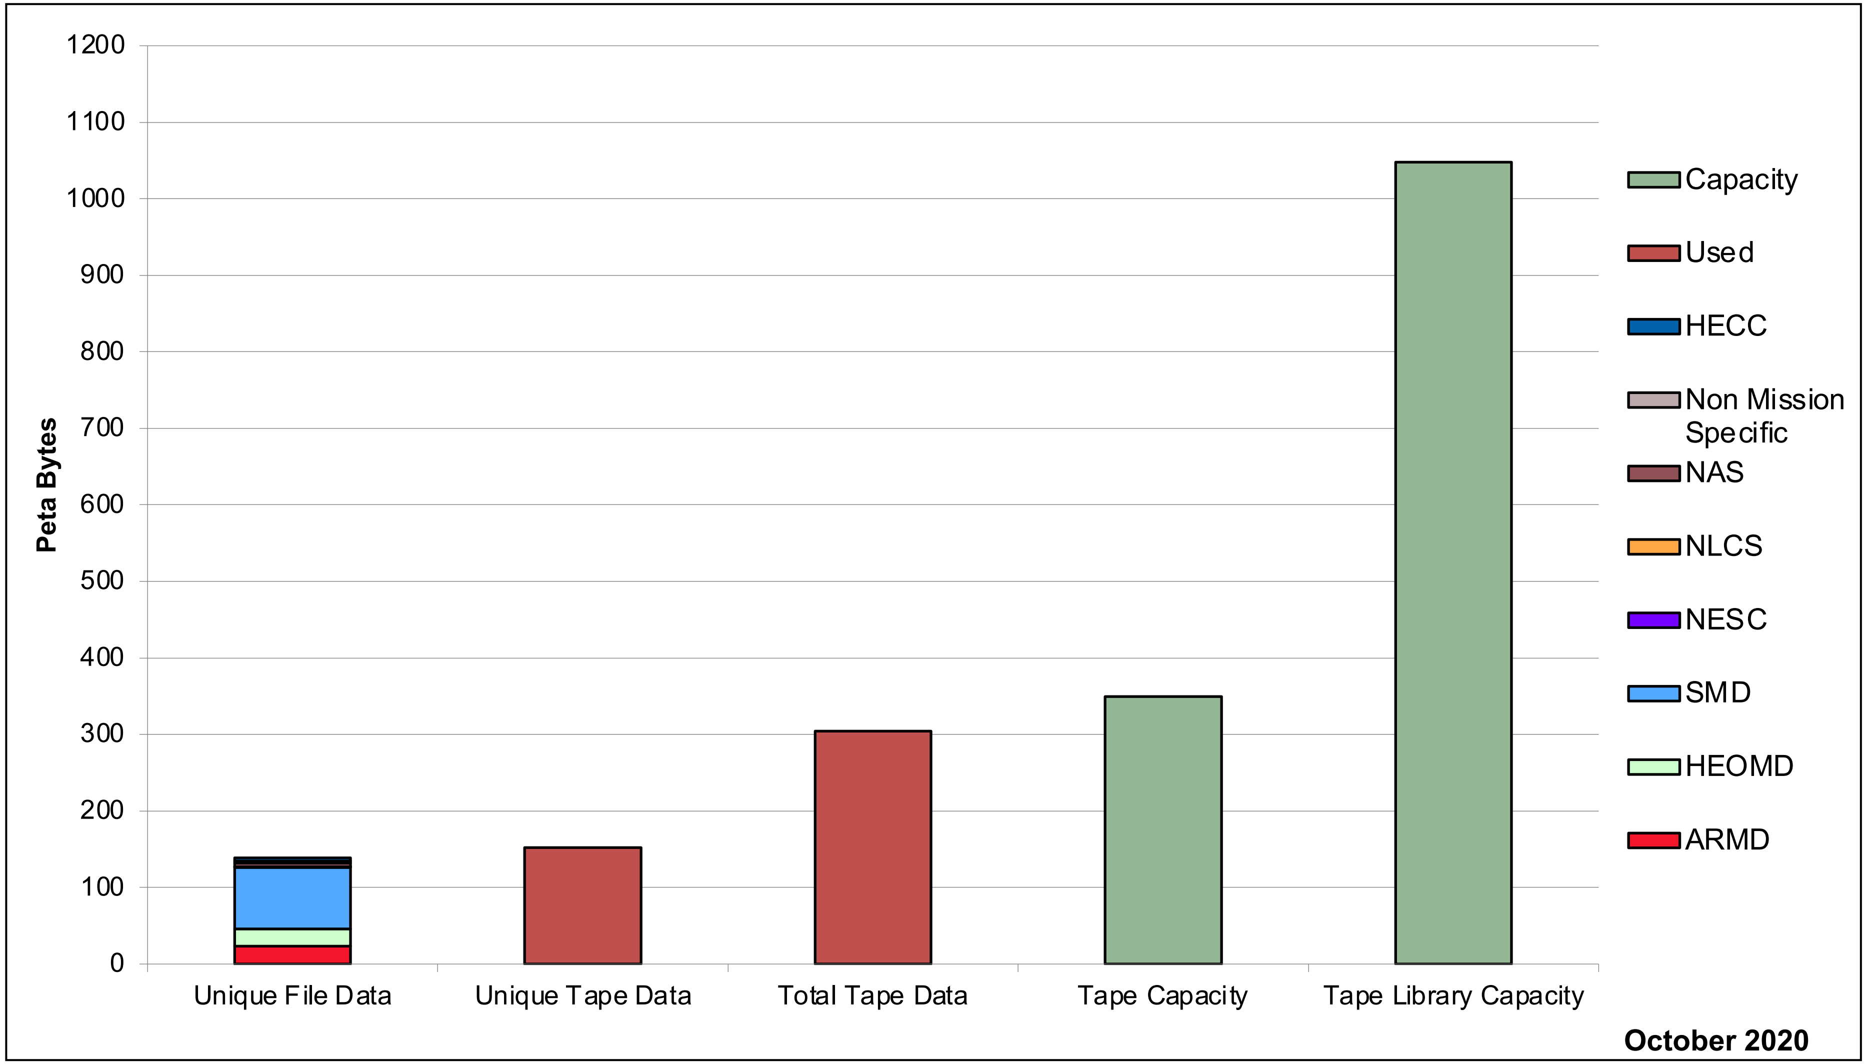

Tape Archive Status

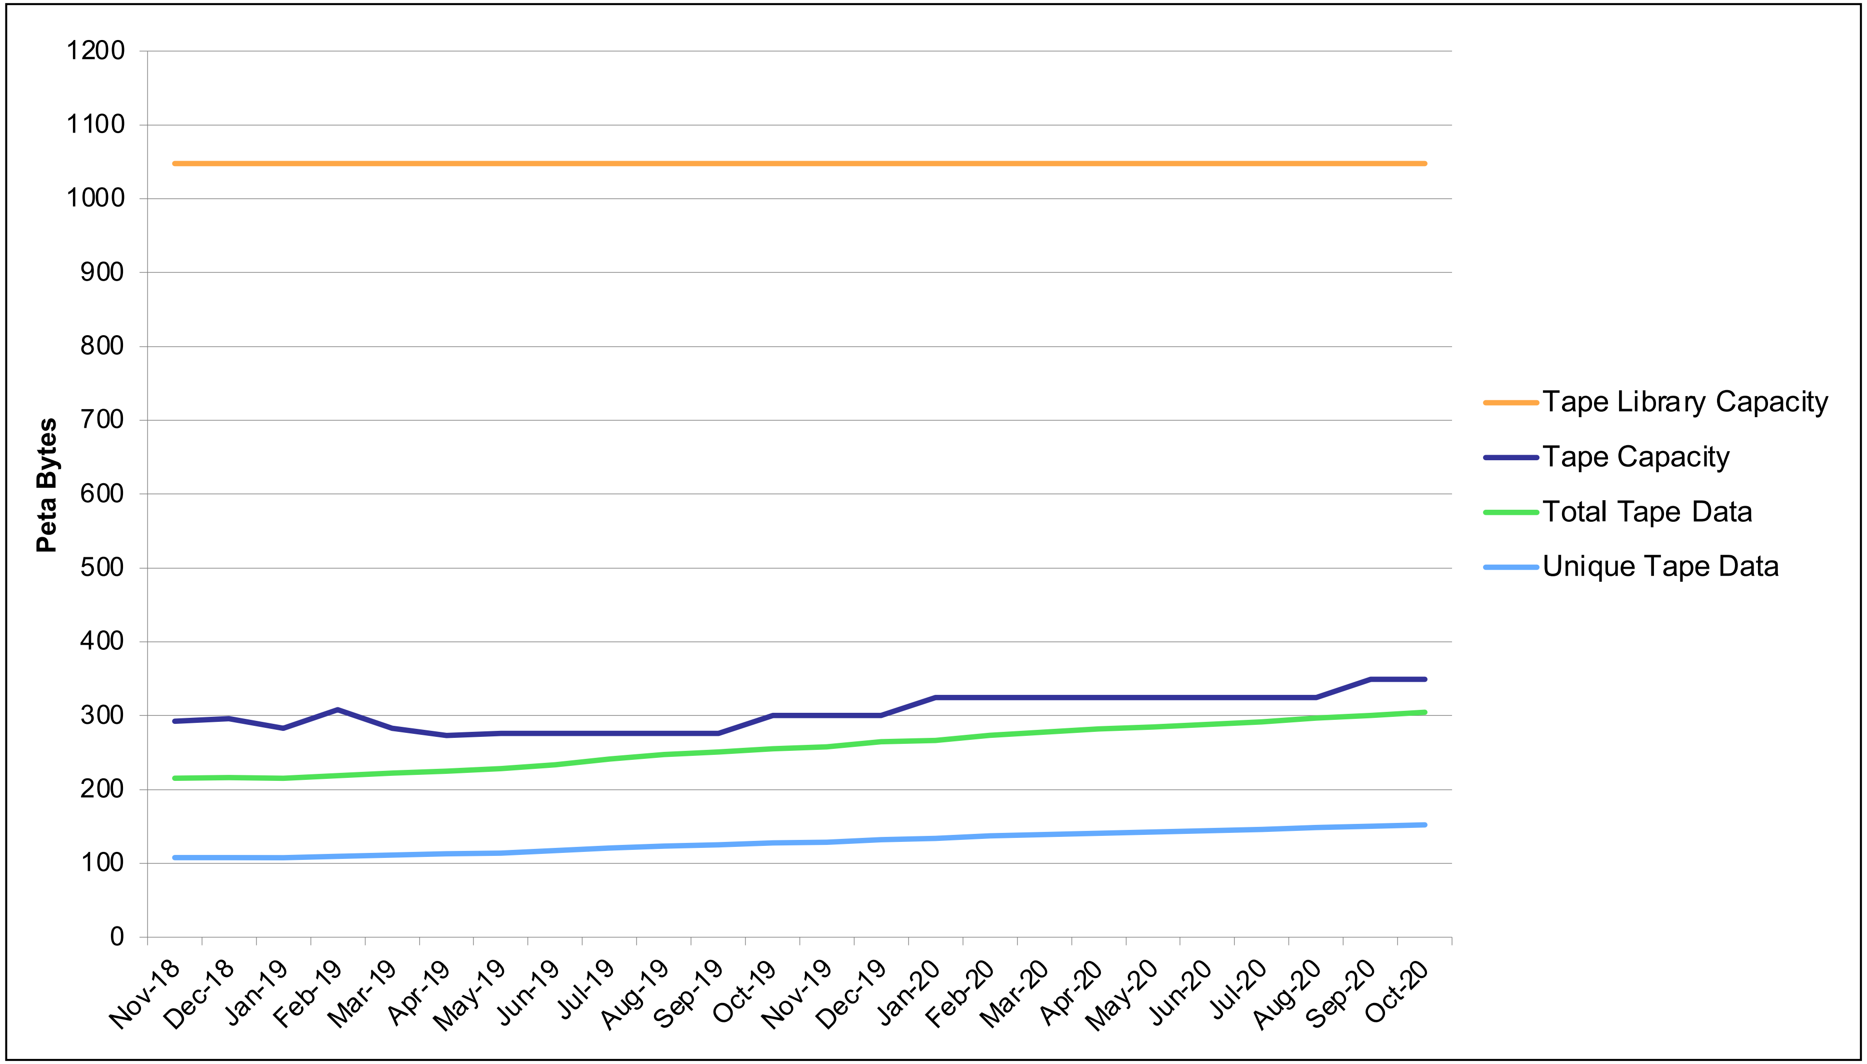

Pleiades: SBUs Reported, Normalized to 30-Day Month

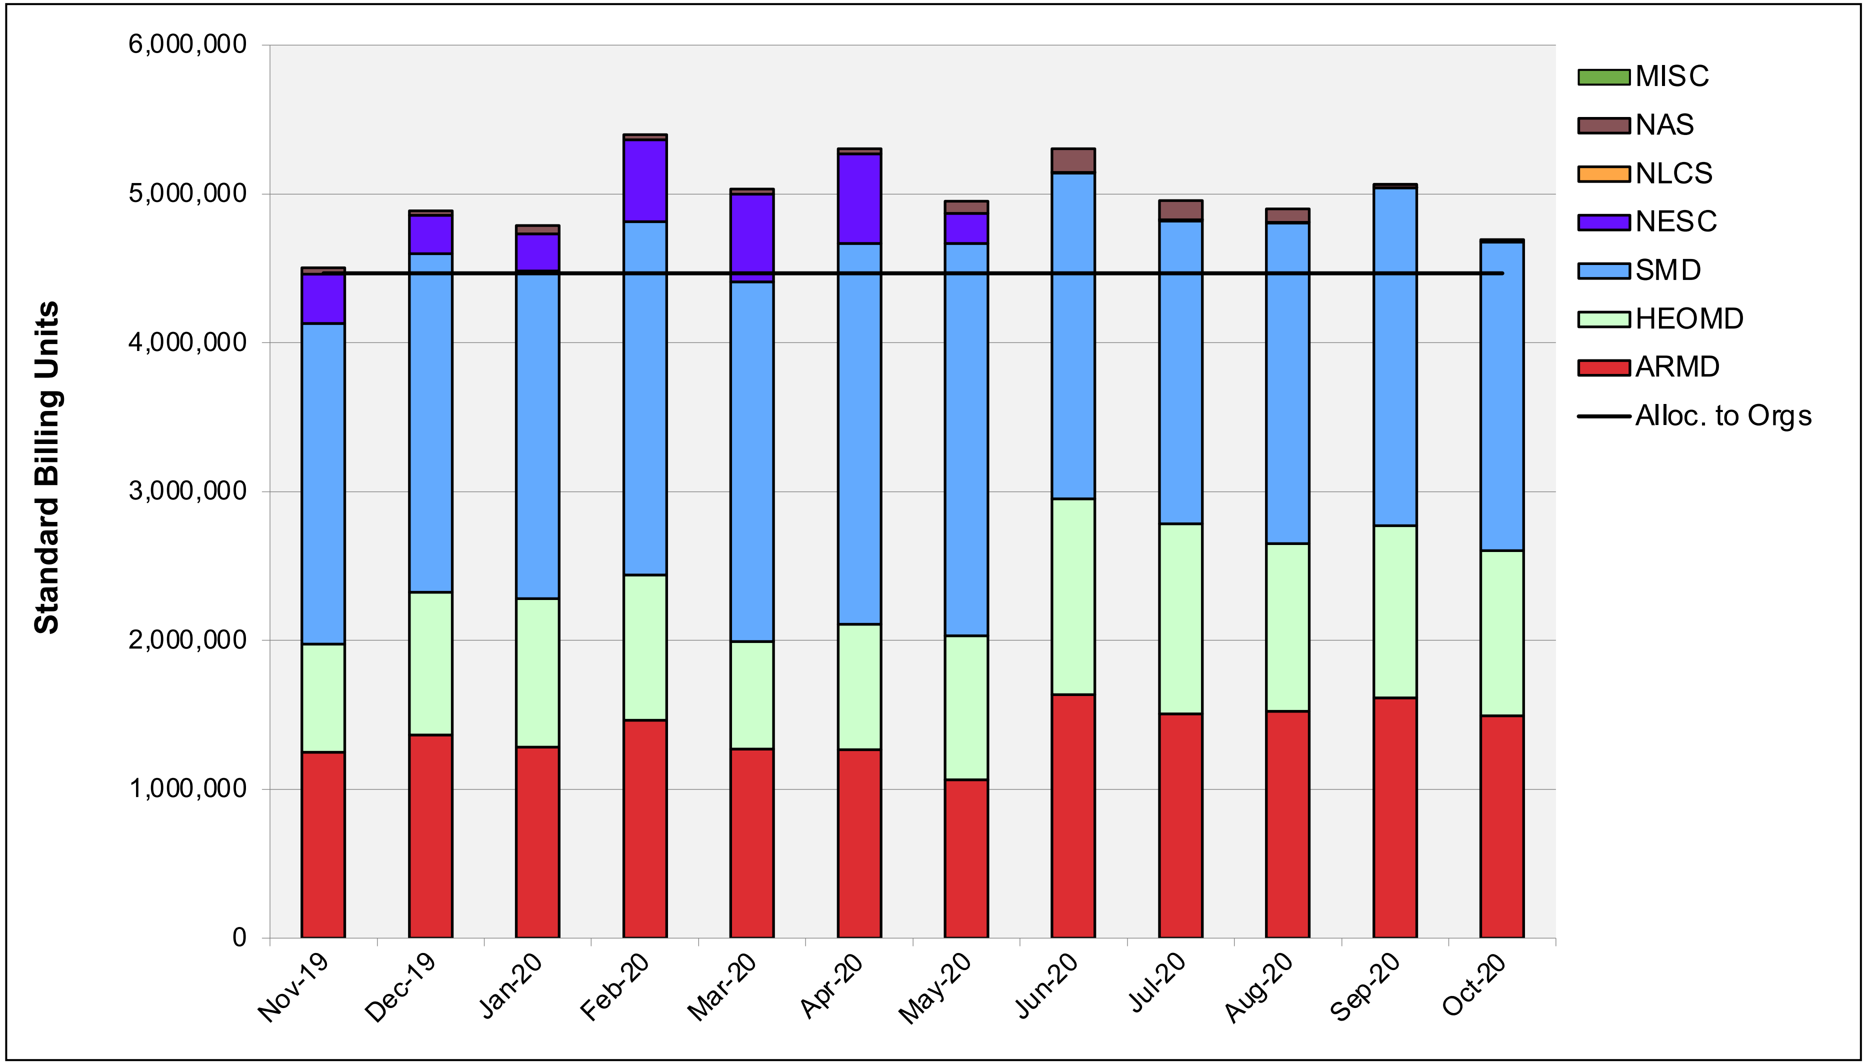

Pleiades: Devel Queue Utilization

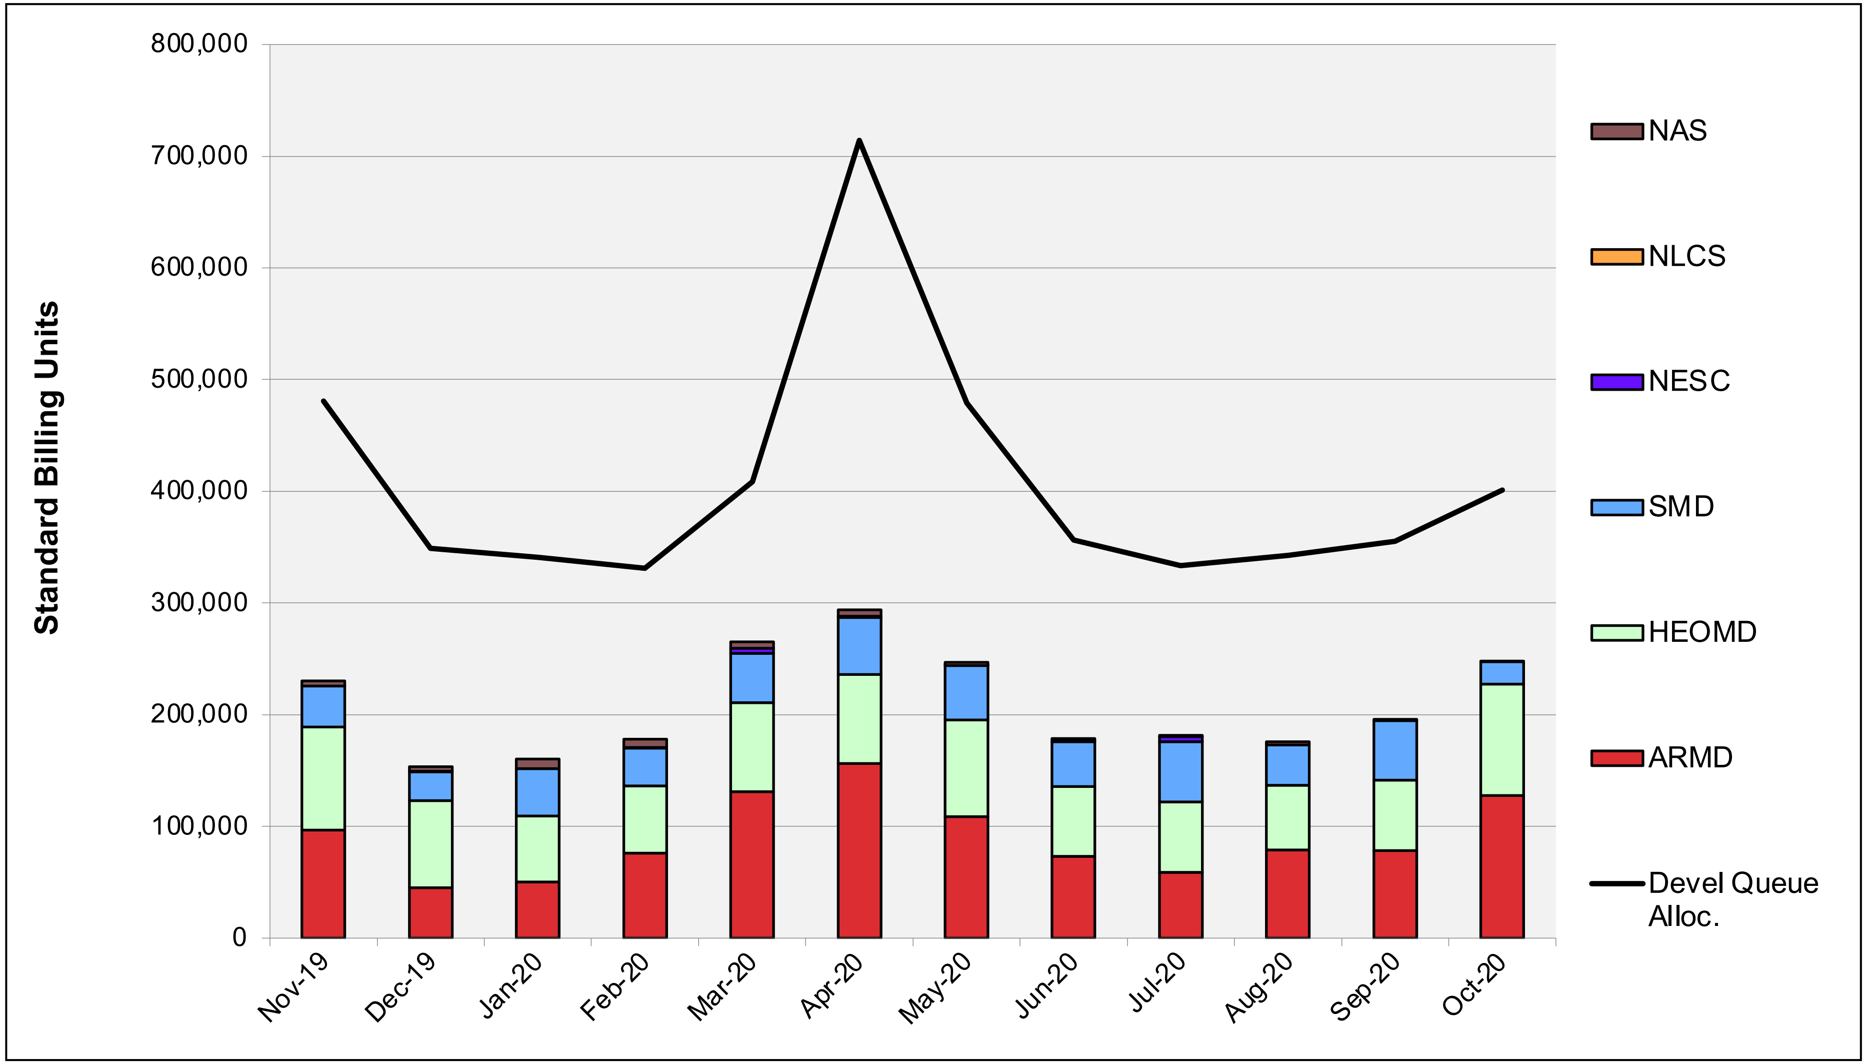

Pleiades: Monthly Utilization by Job Length

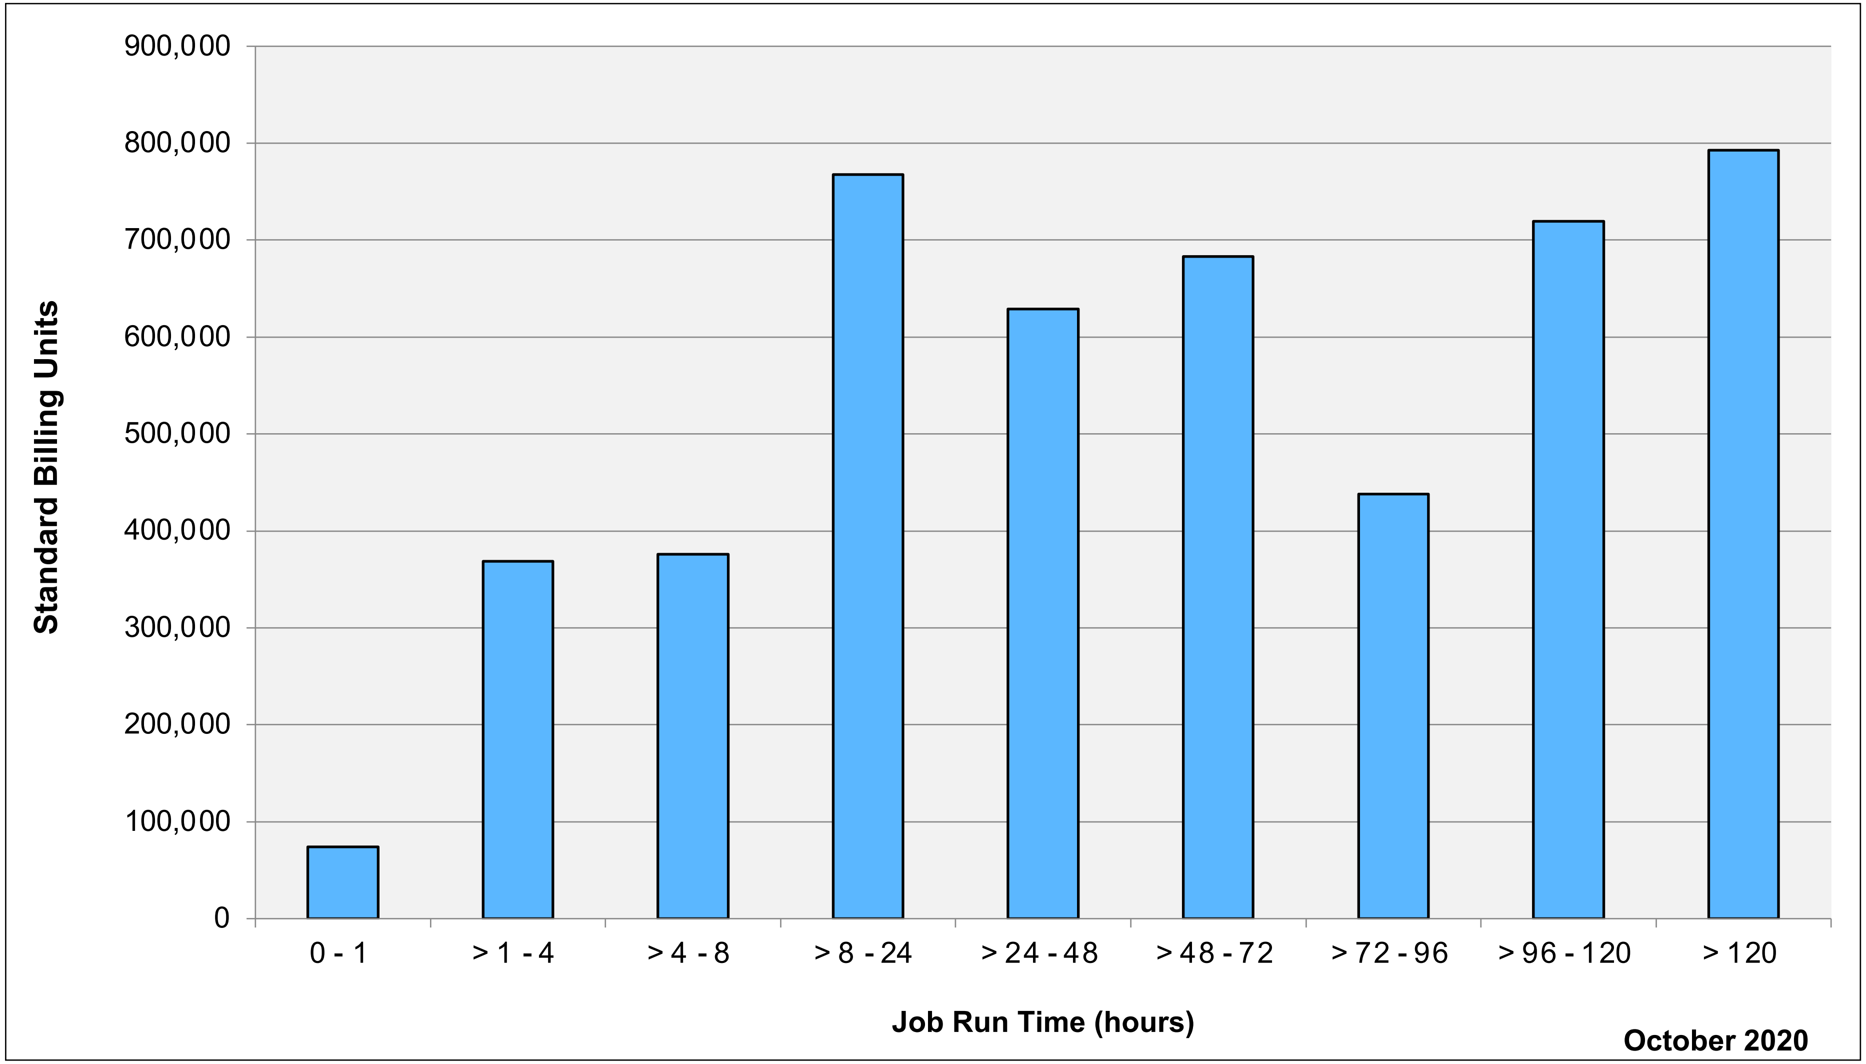

Pleiades: Monthly Utilization by Job Size

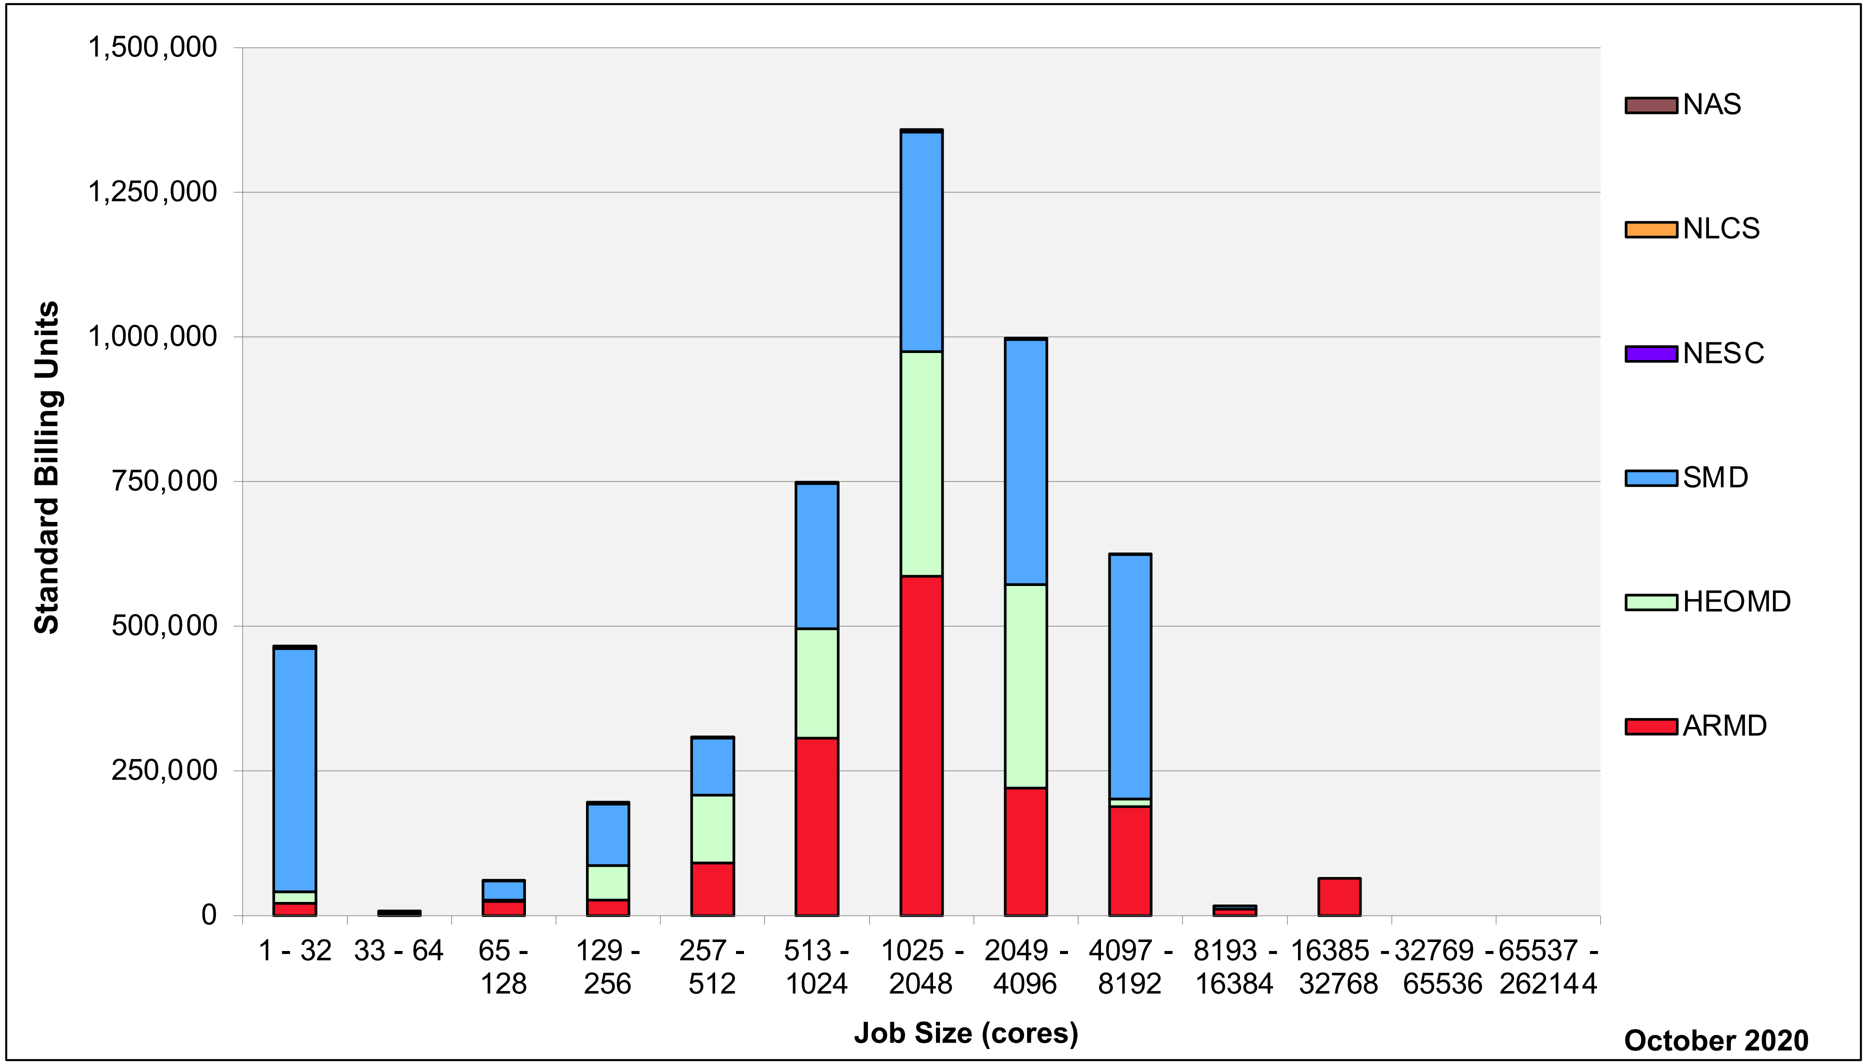

Pleiades: Monthly Utilization by Size and Length

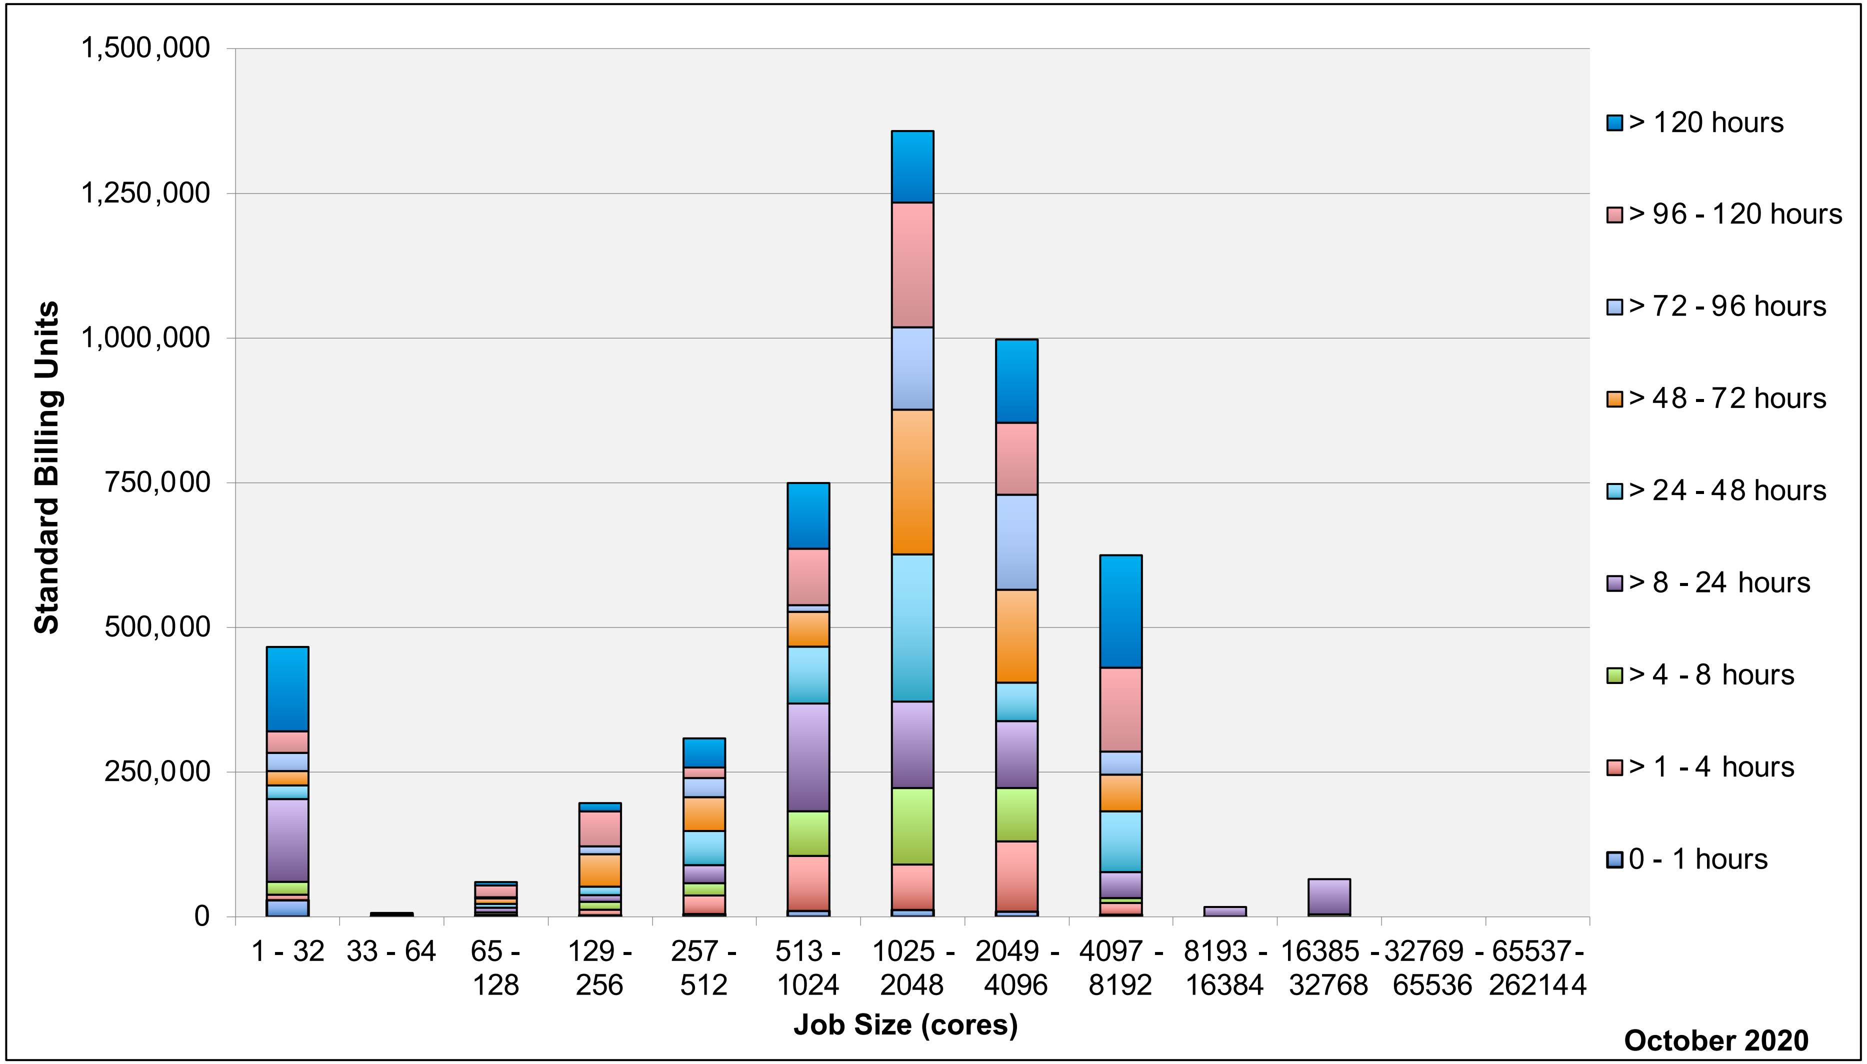

Pleiades: Average Time to Clear All Jobs

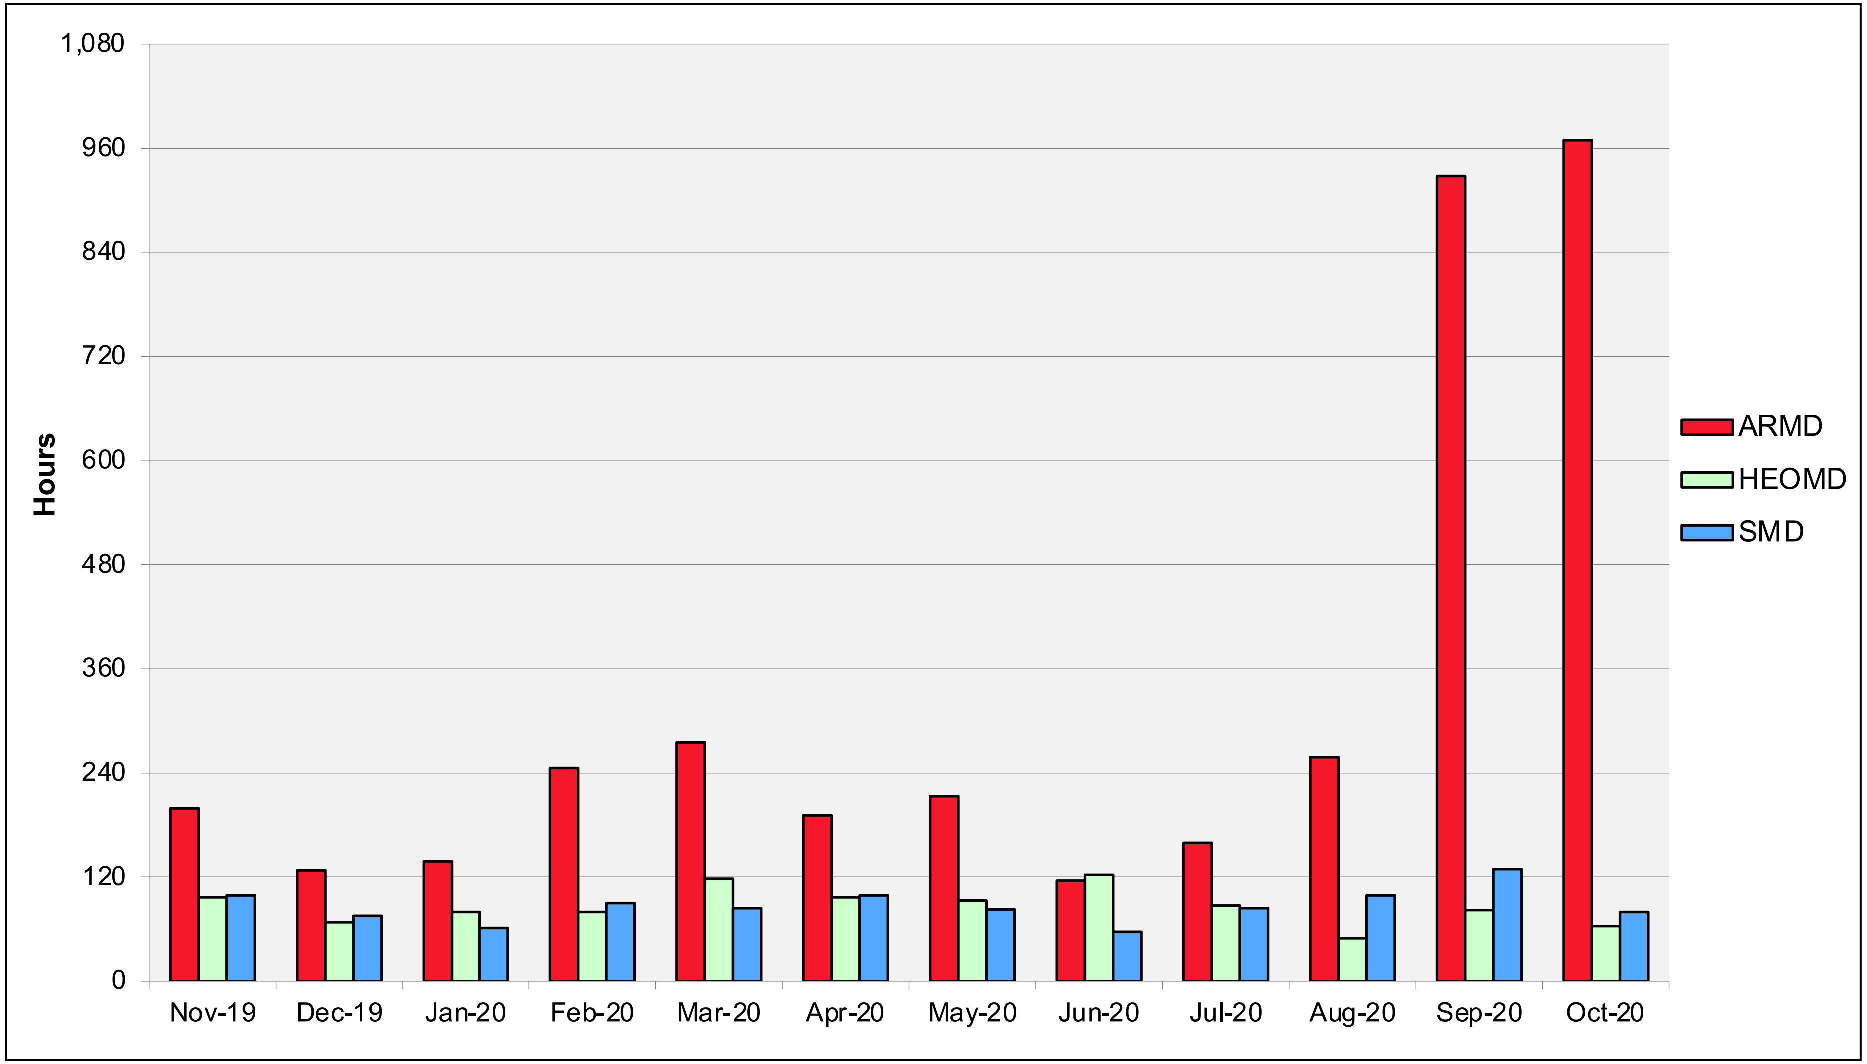

Pleiades: Average Expansion Factor

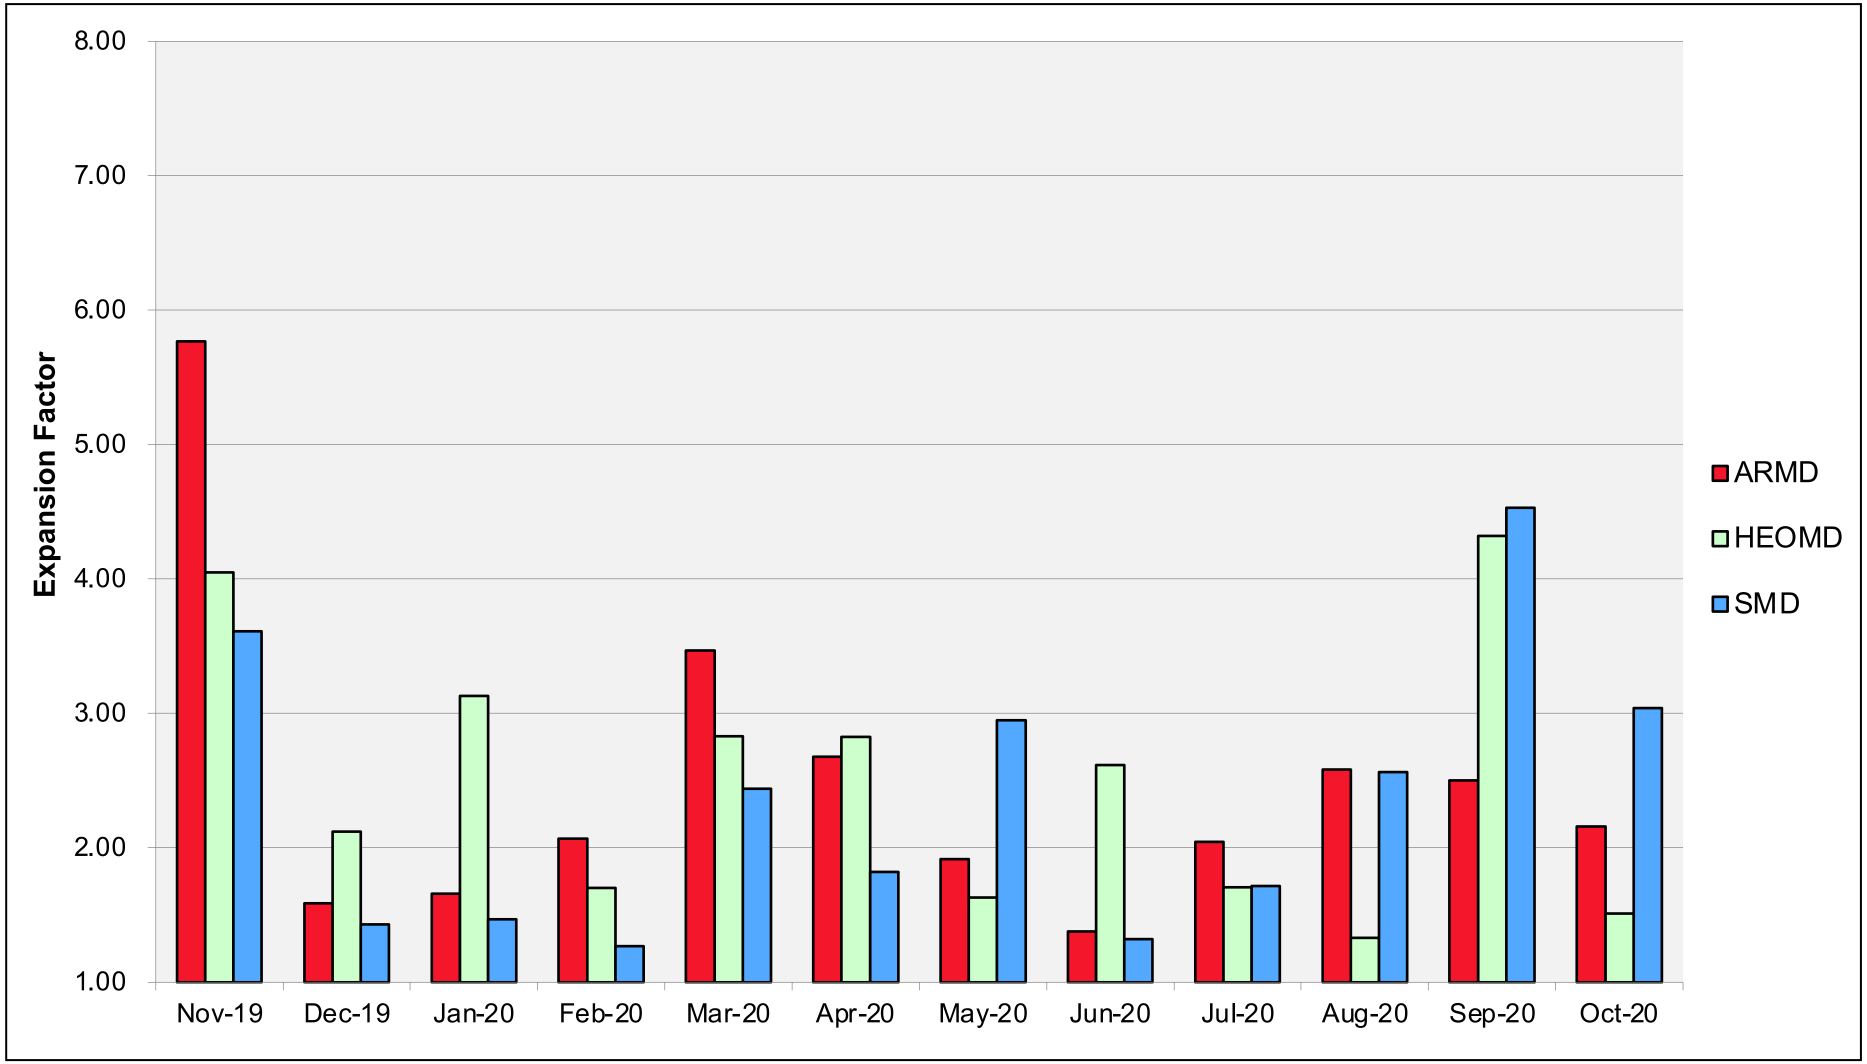

Aitken: SBUs Reported, Normalized to 30-Day Month

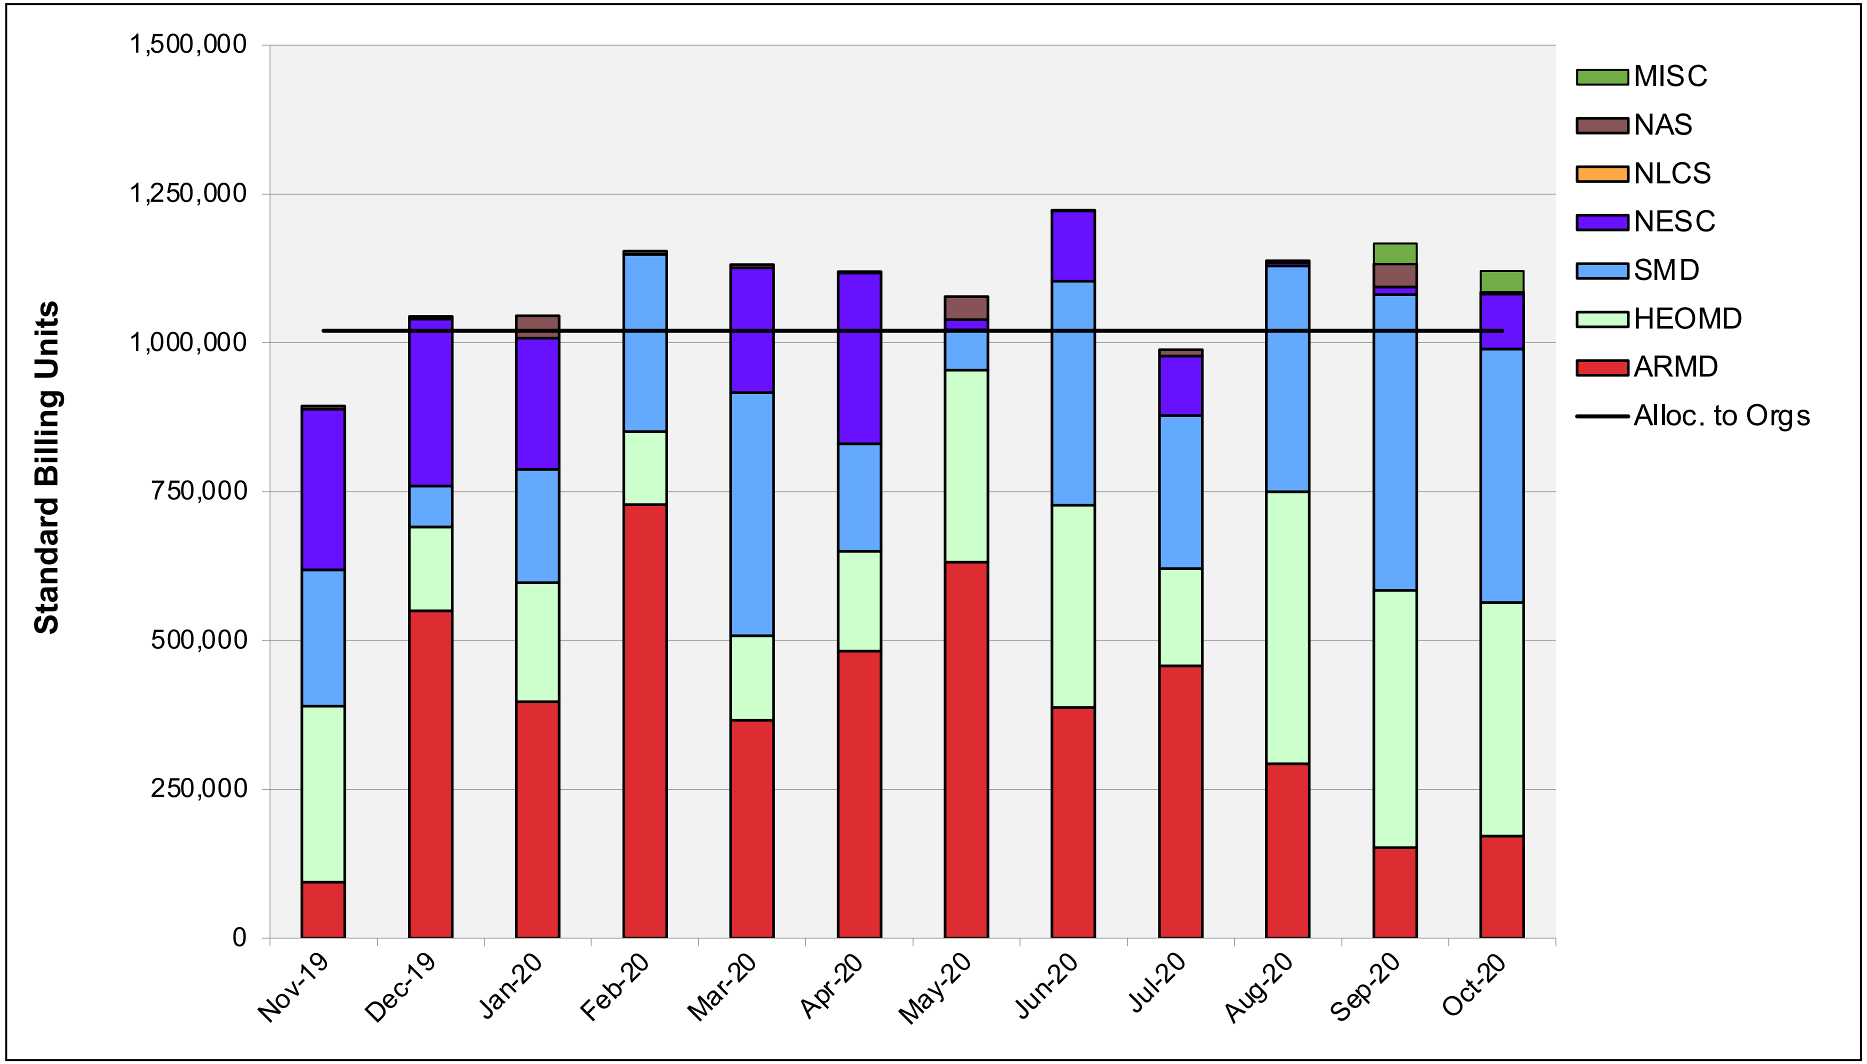

Aitken: Devel Queue Utilization

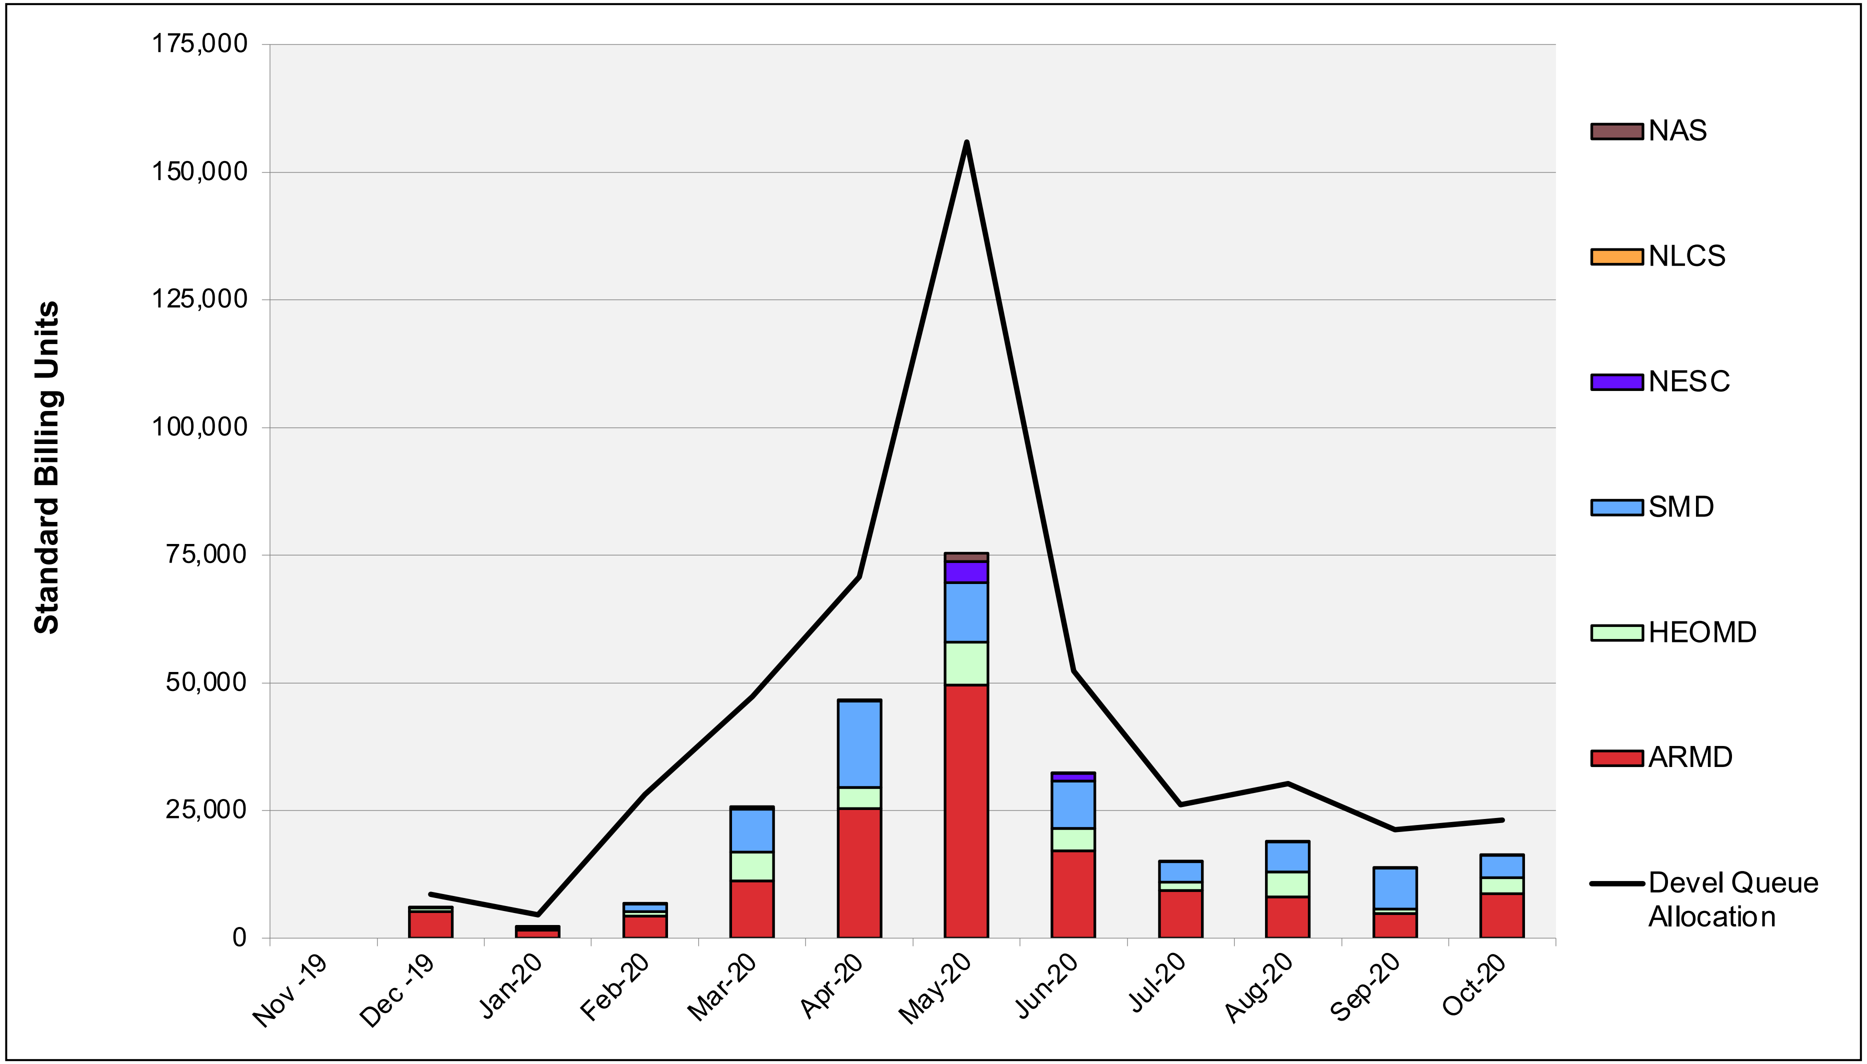

Aitken: Monthly Utilization by Job Length

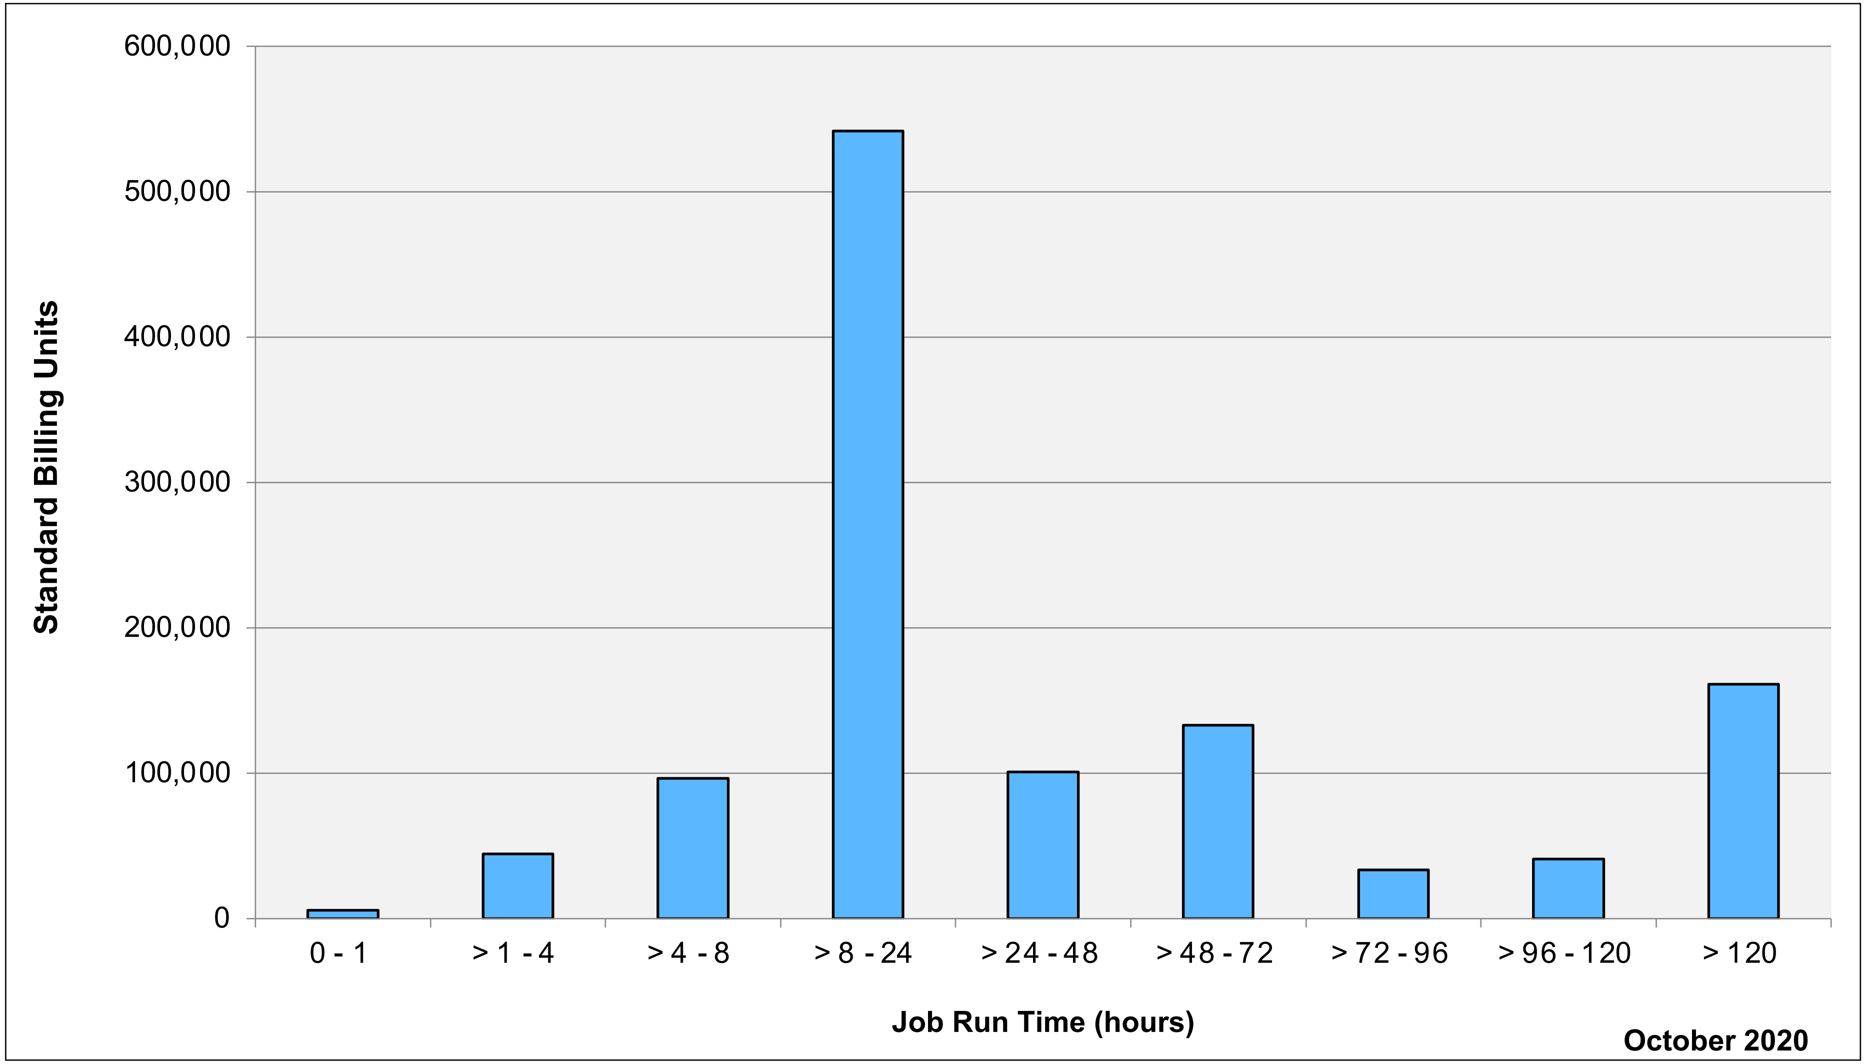

Aitken: Monthly Utilization by Job Size

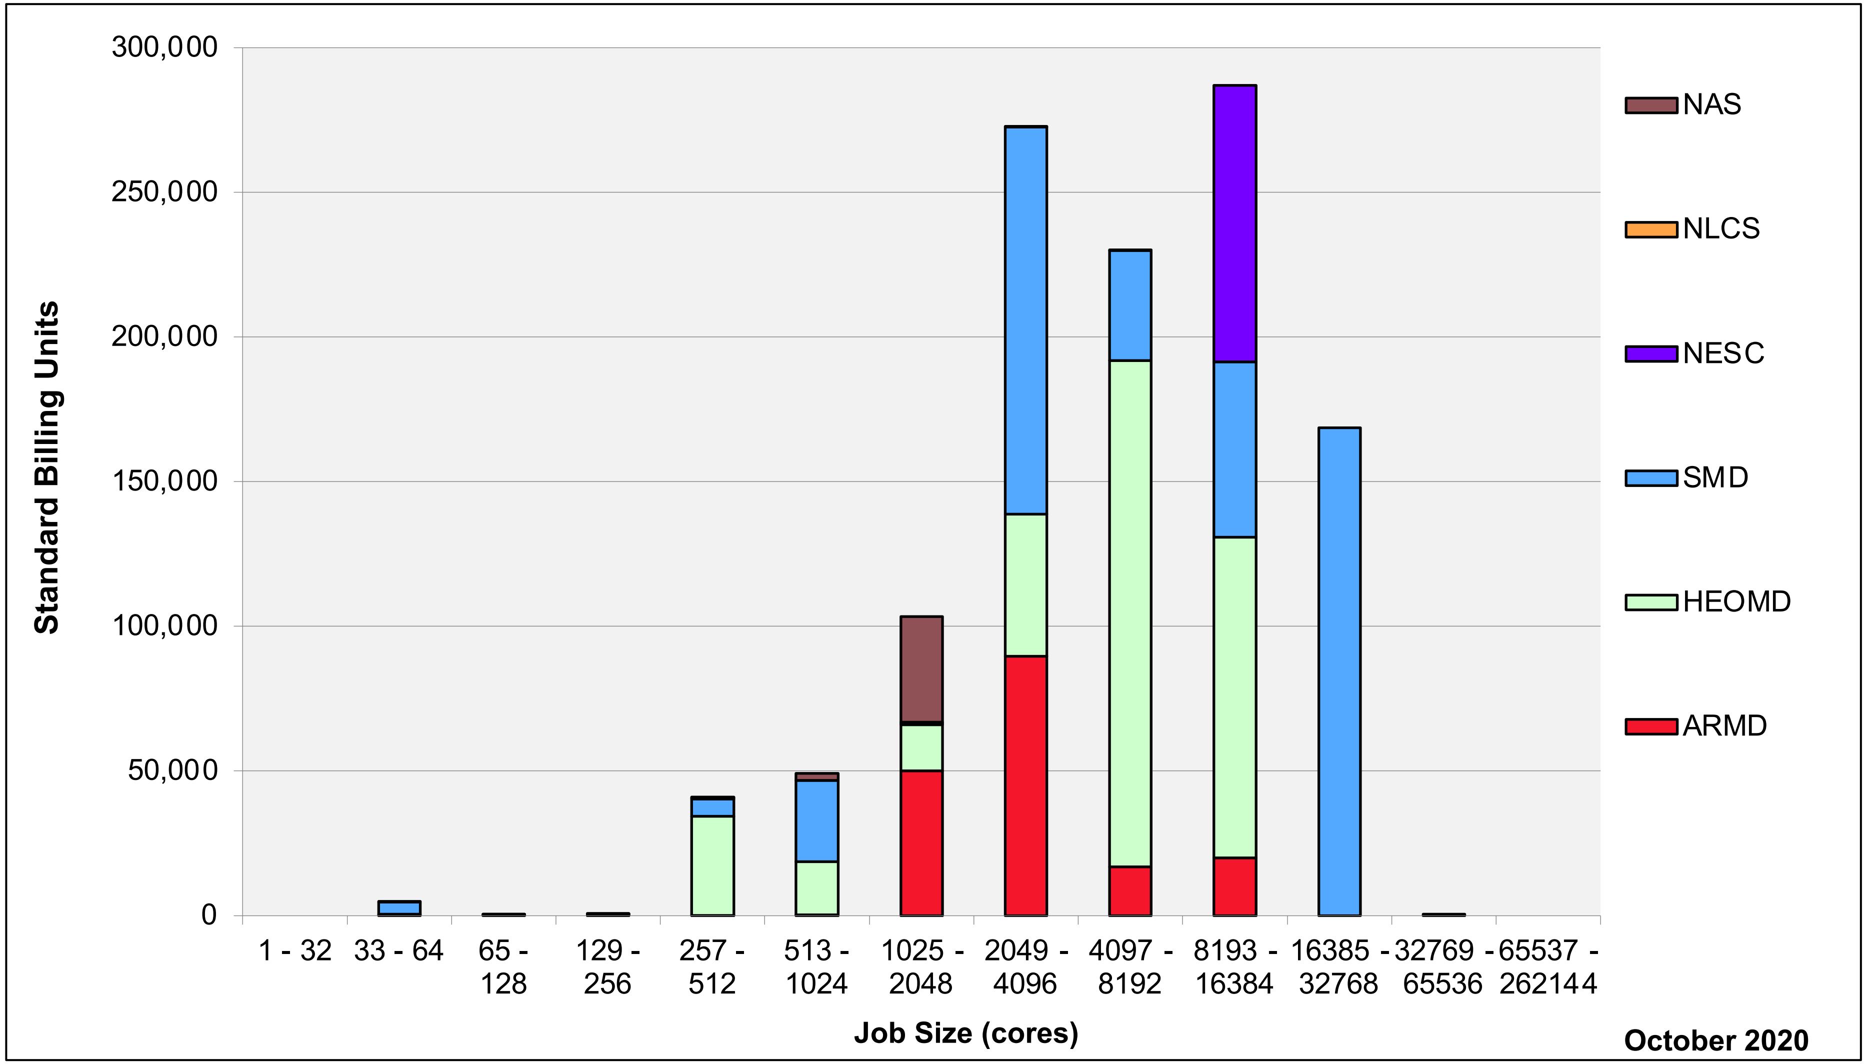

Aitken: Monthly Utilization by Size and Length

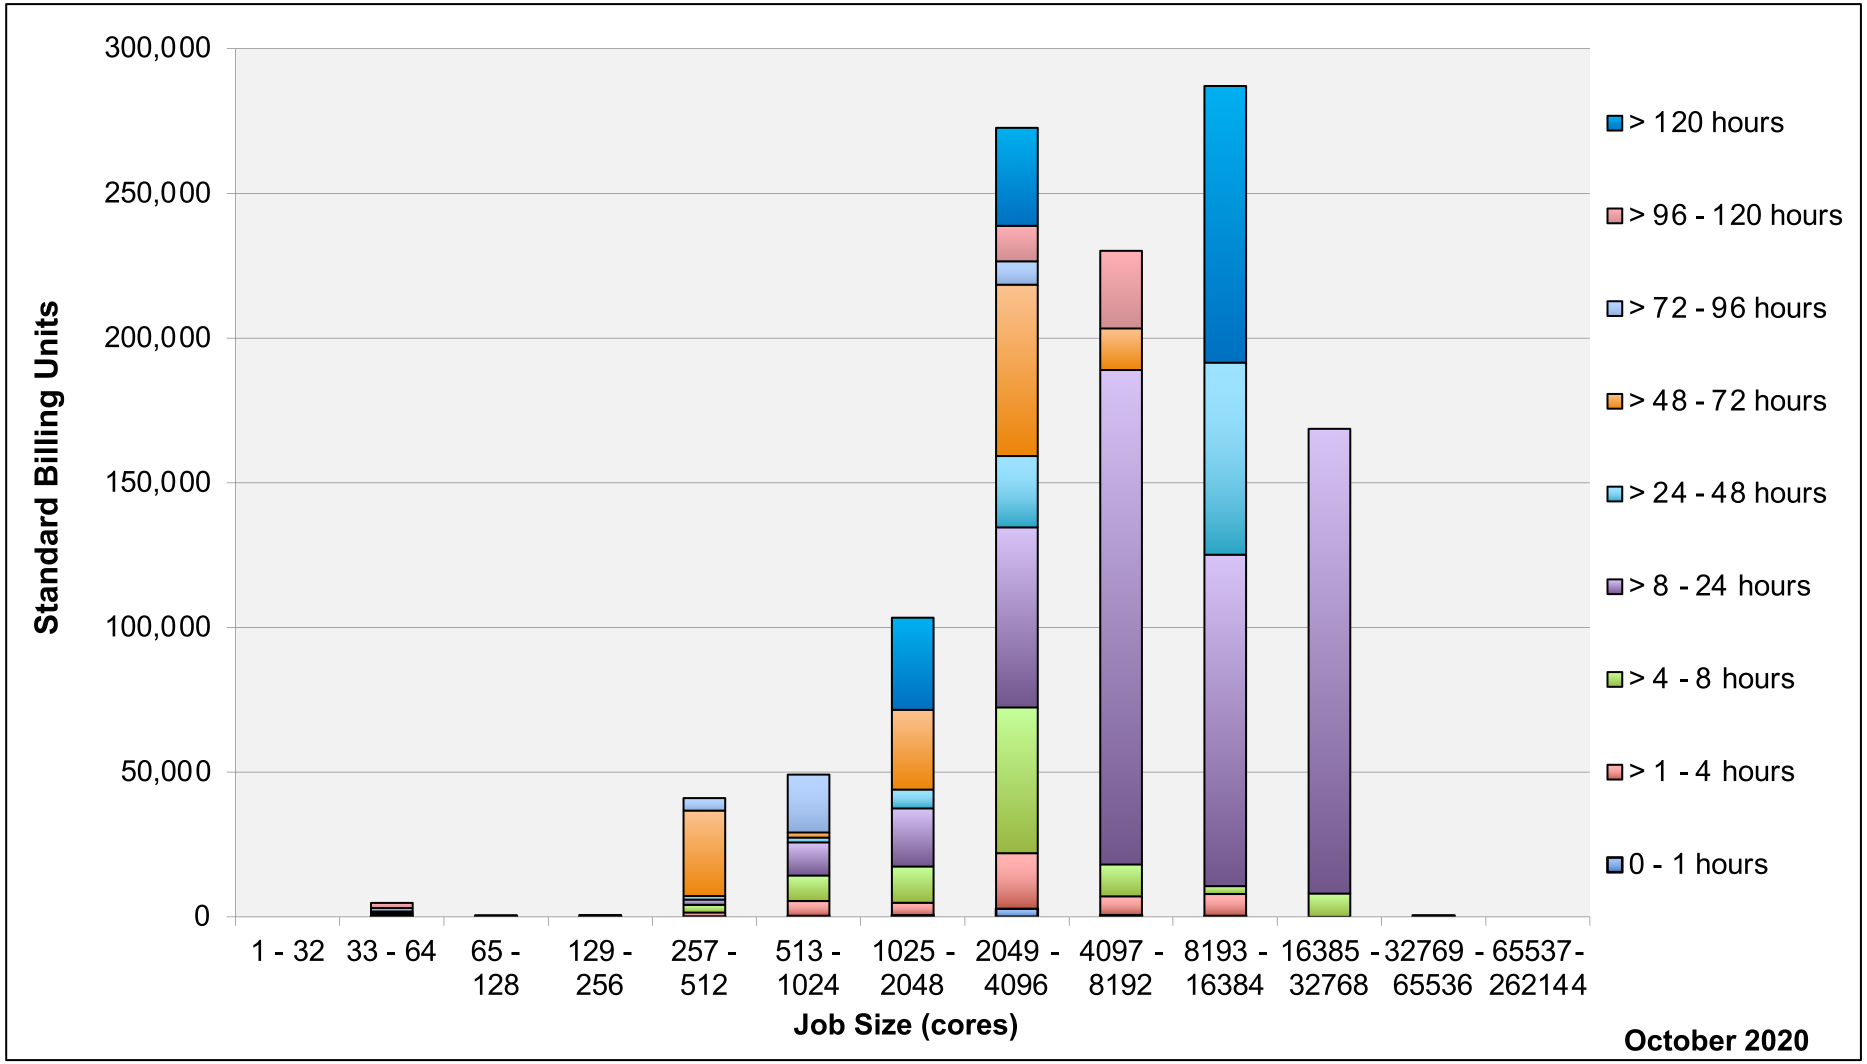

Aitken: Average Time to Clear All Jobs

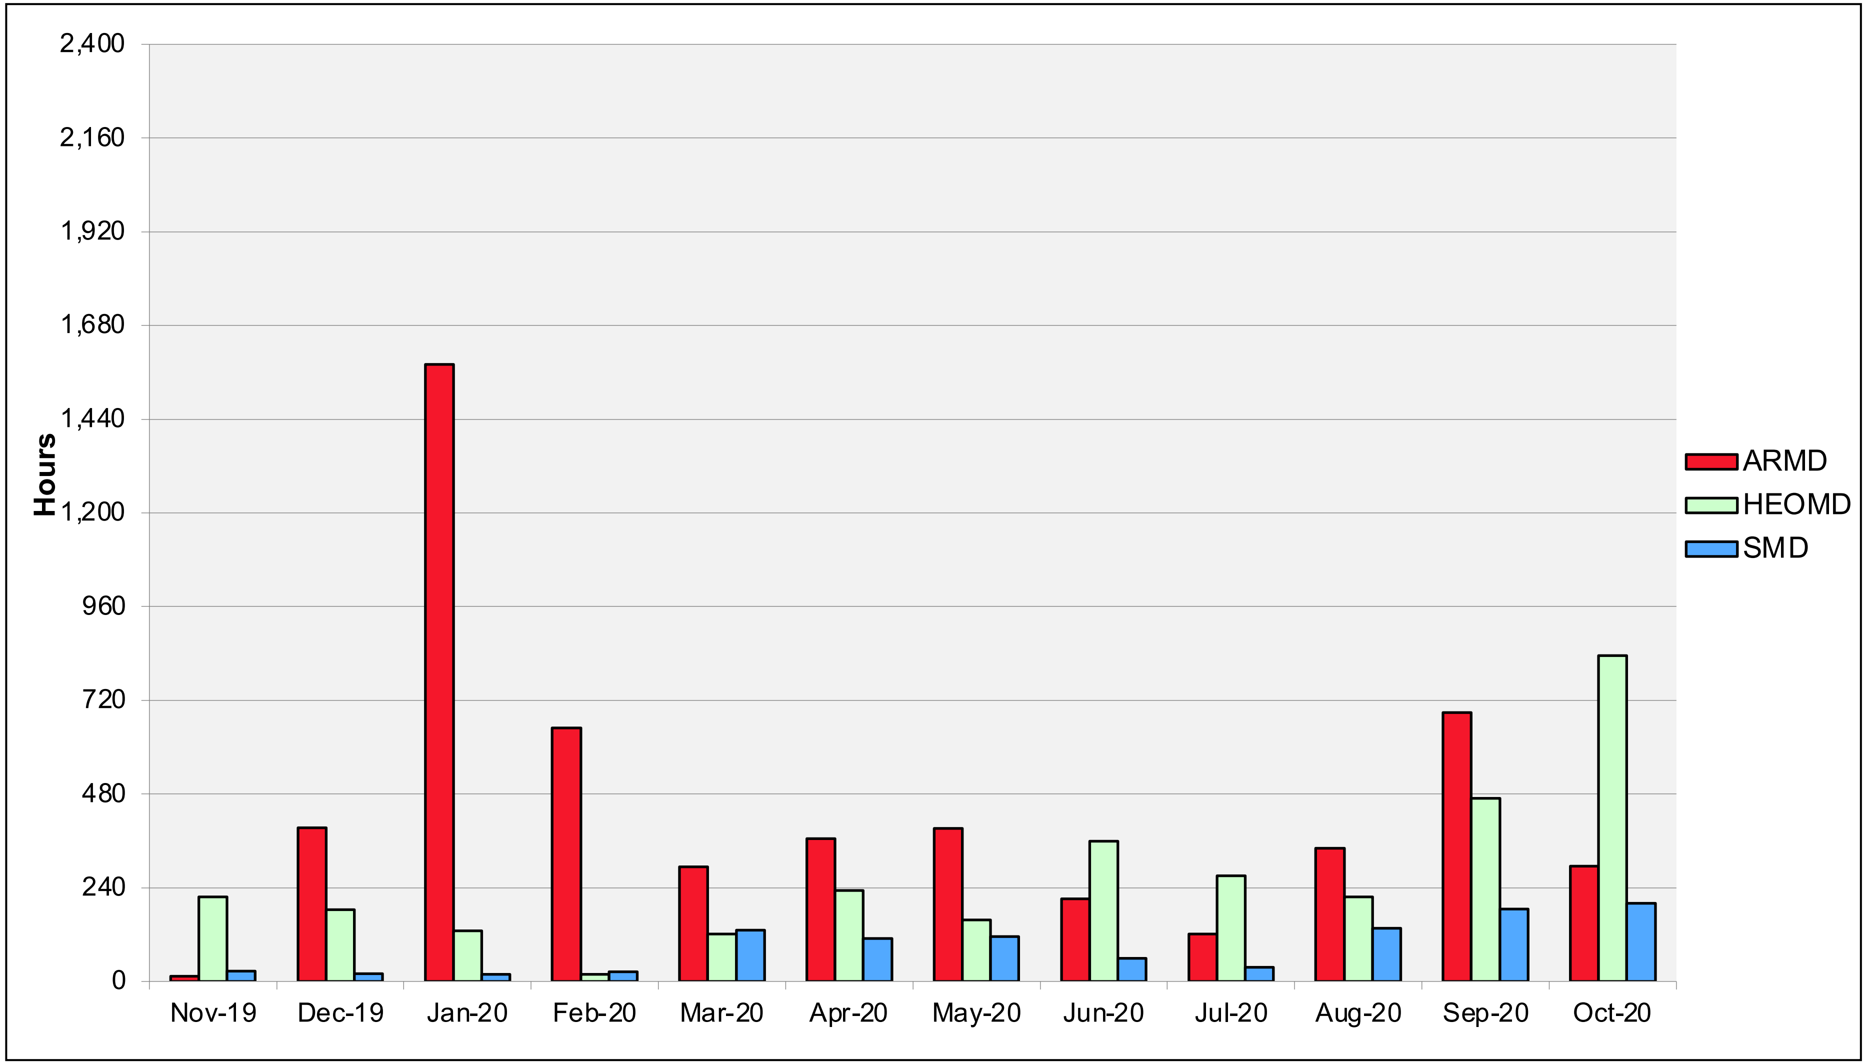

Aitken: Average Expansion Factor

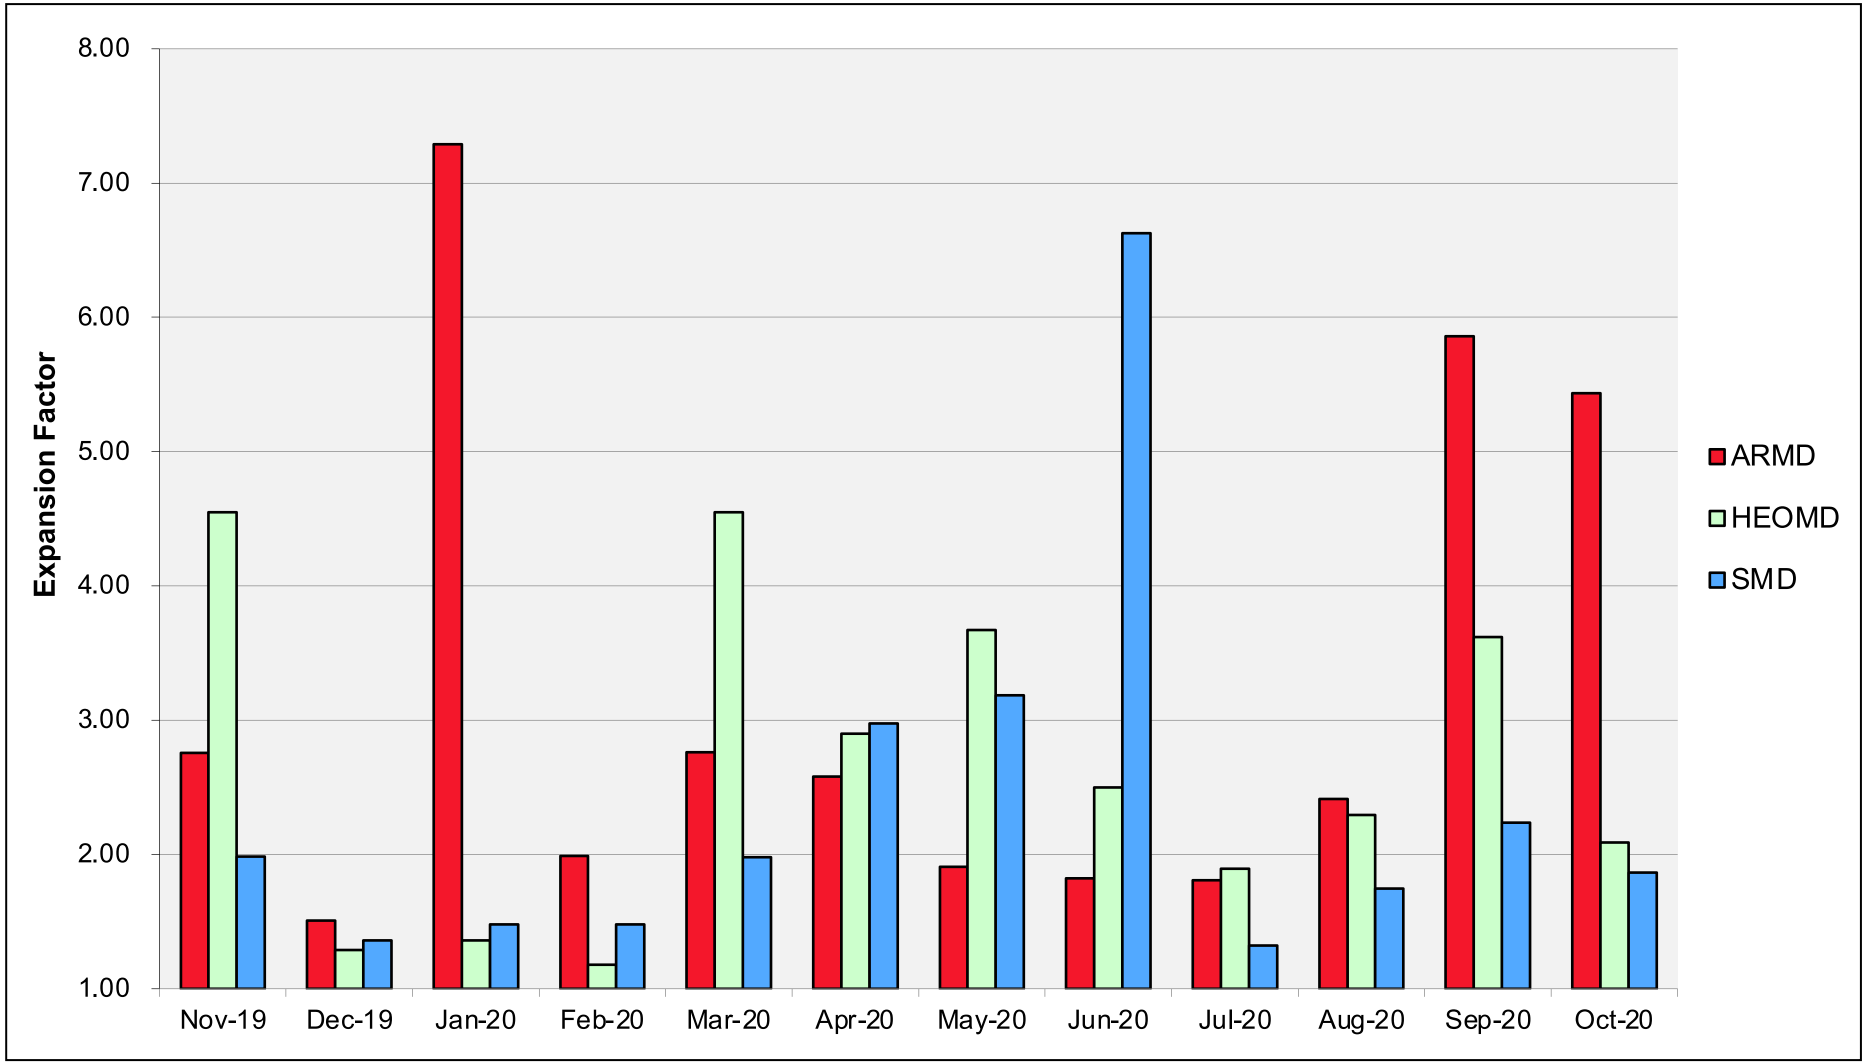

Electra: SBUs Reported, Normalized to 30-Day Month

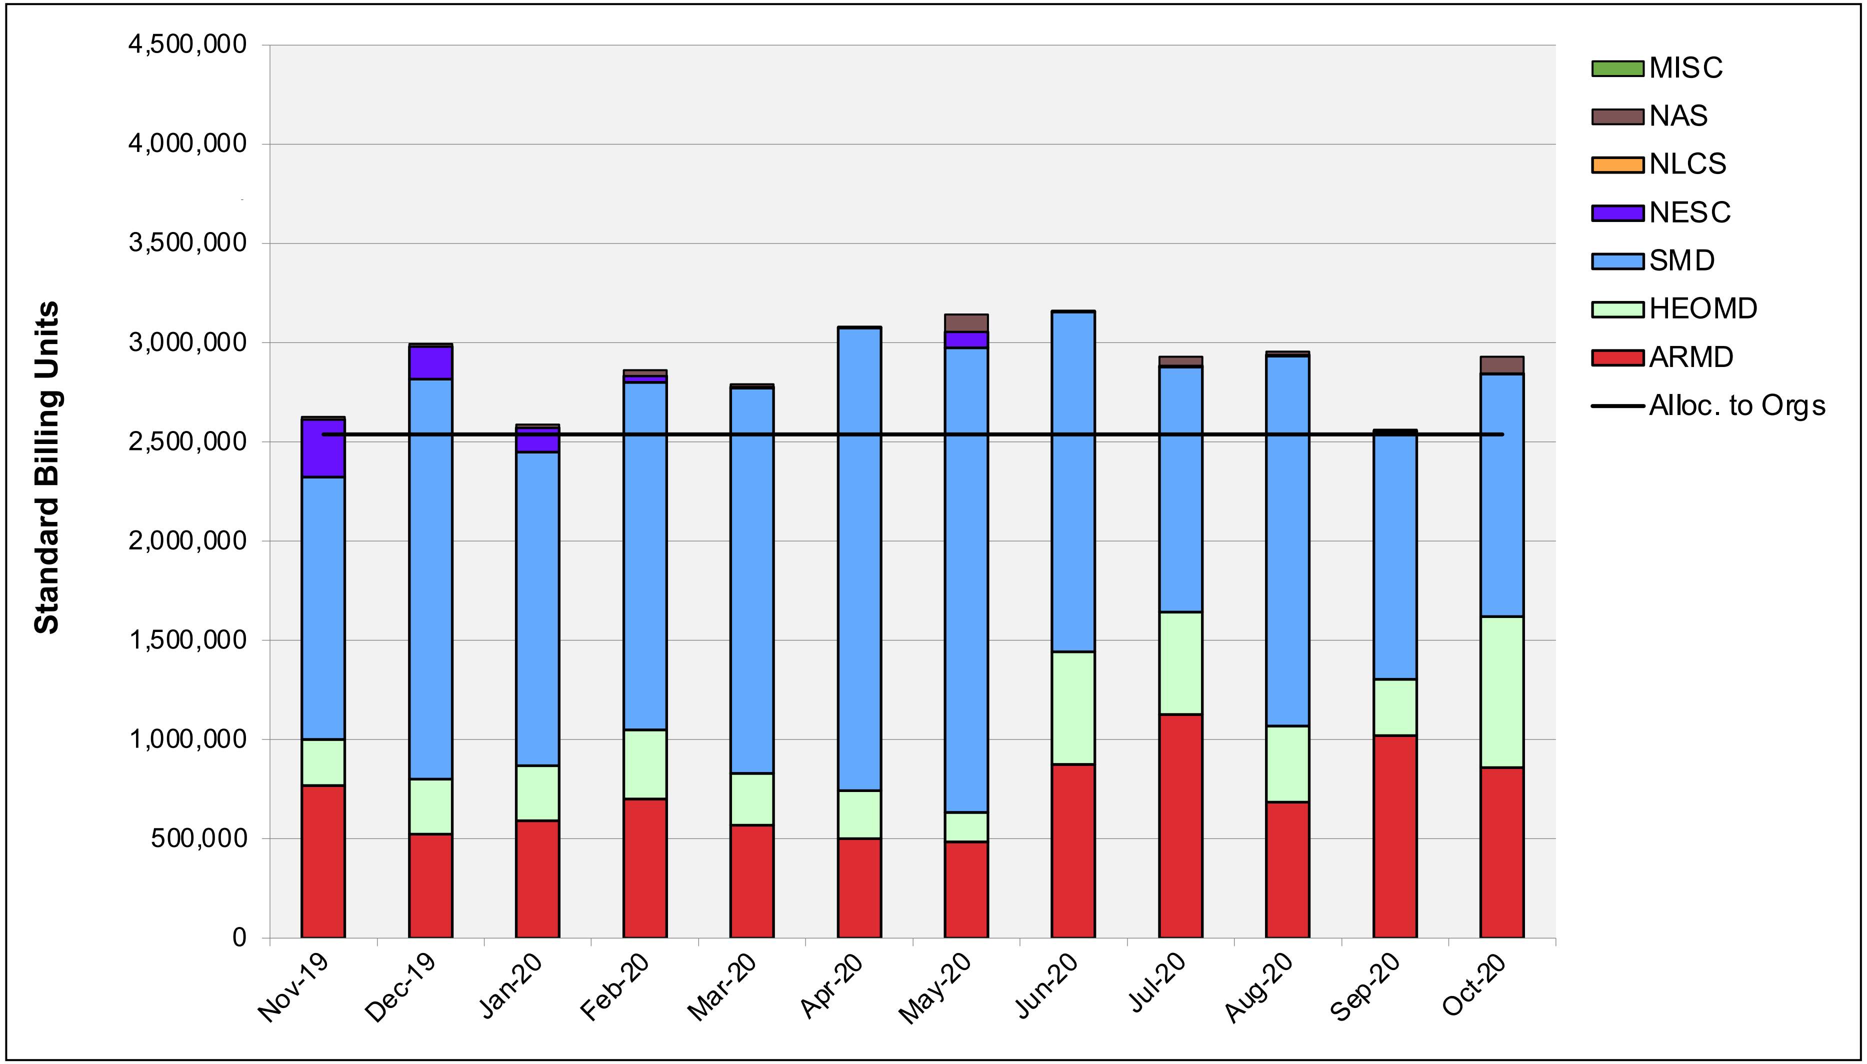

Electra: Devel Queue Utilization

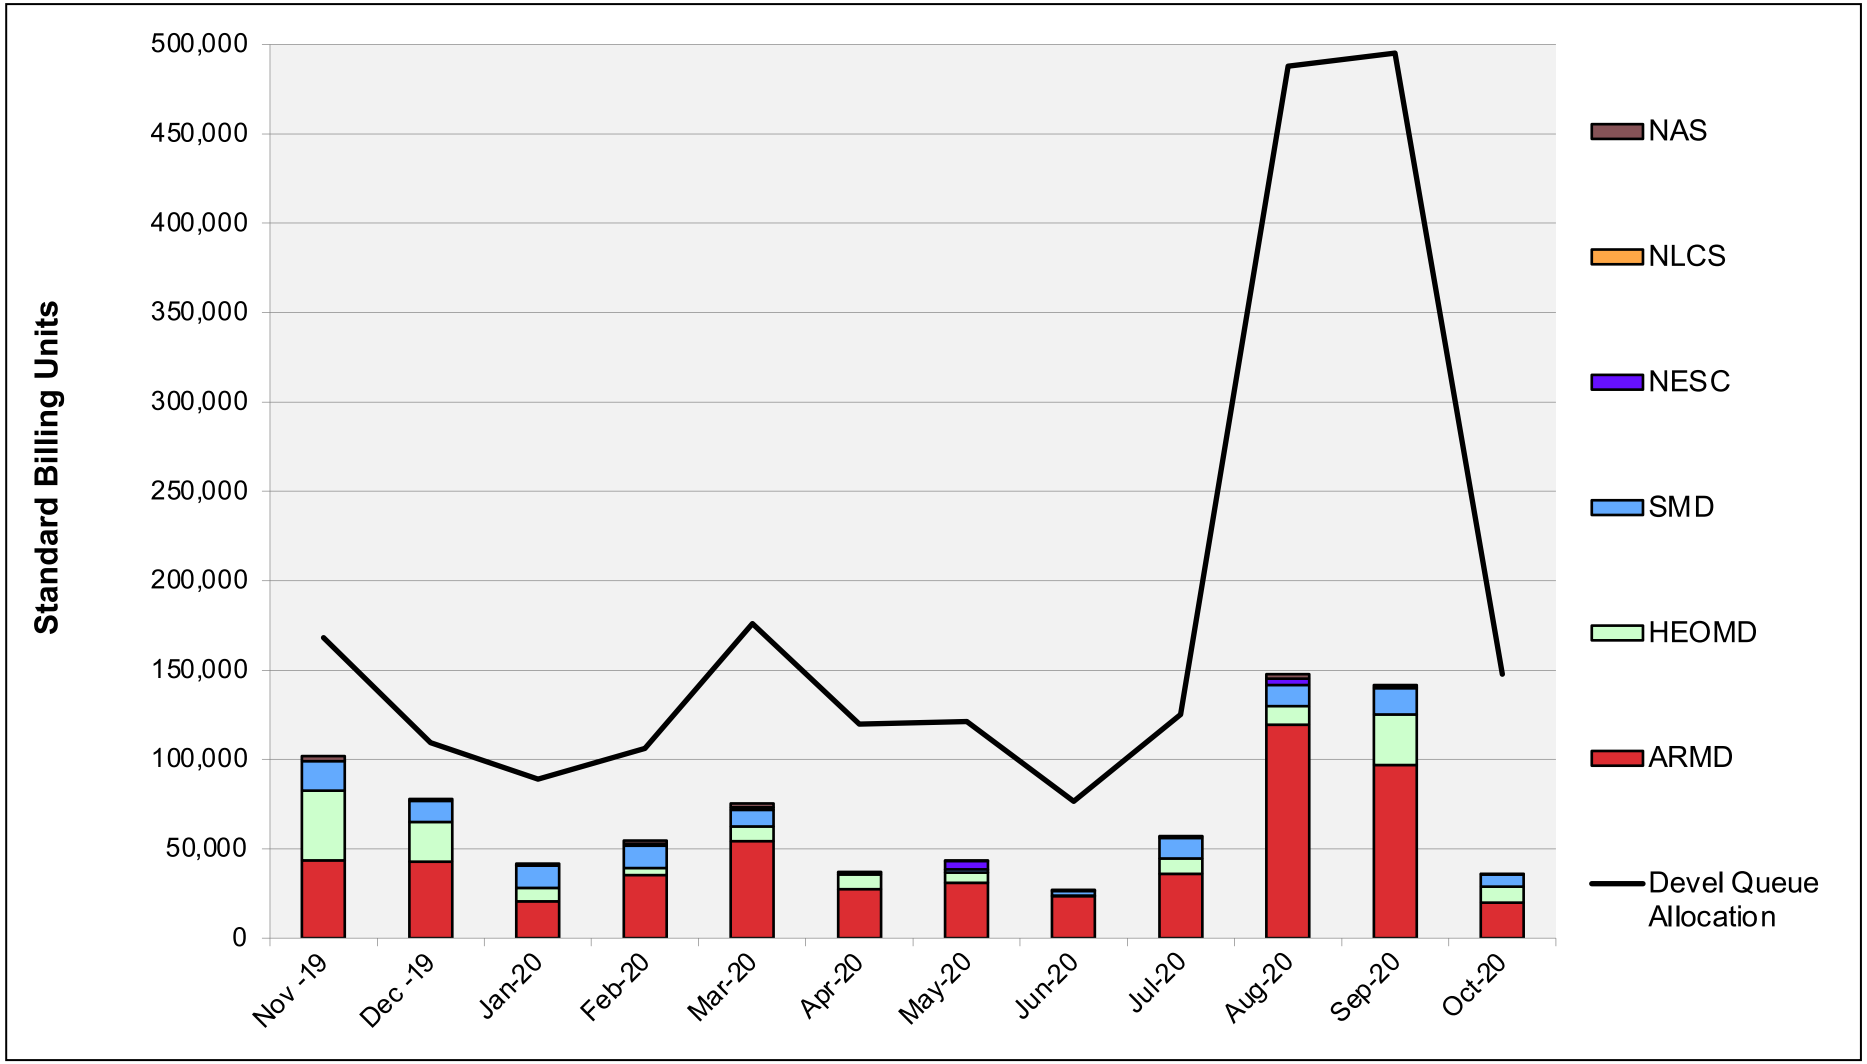

Electra: Monthly Utilization by Job Length

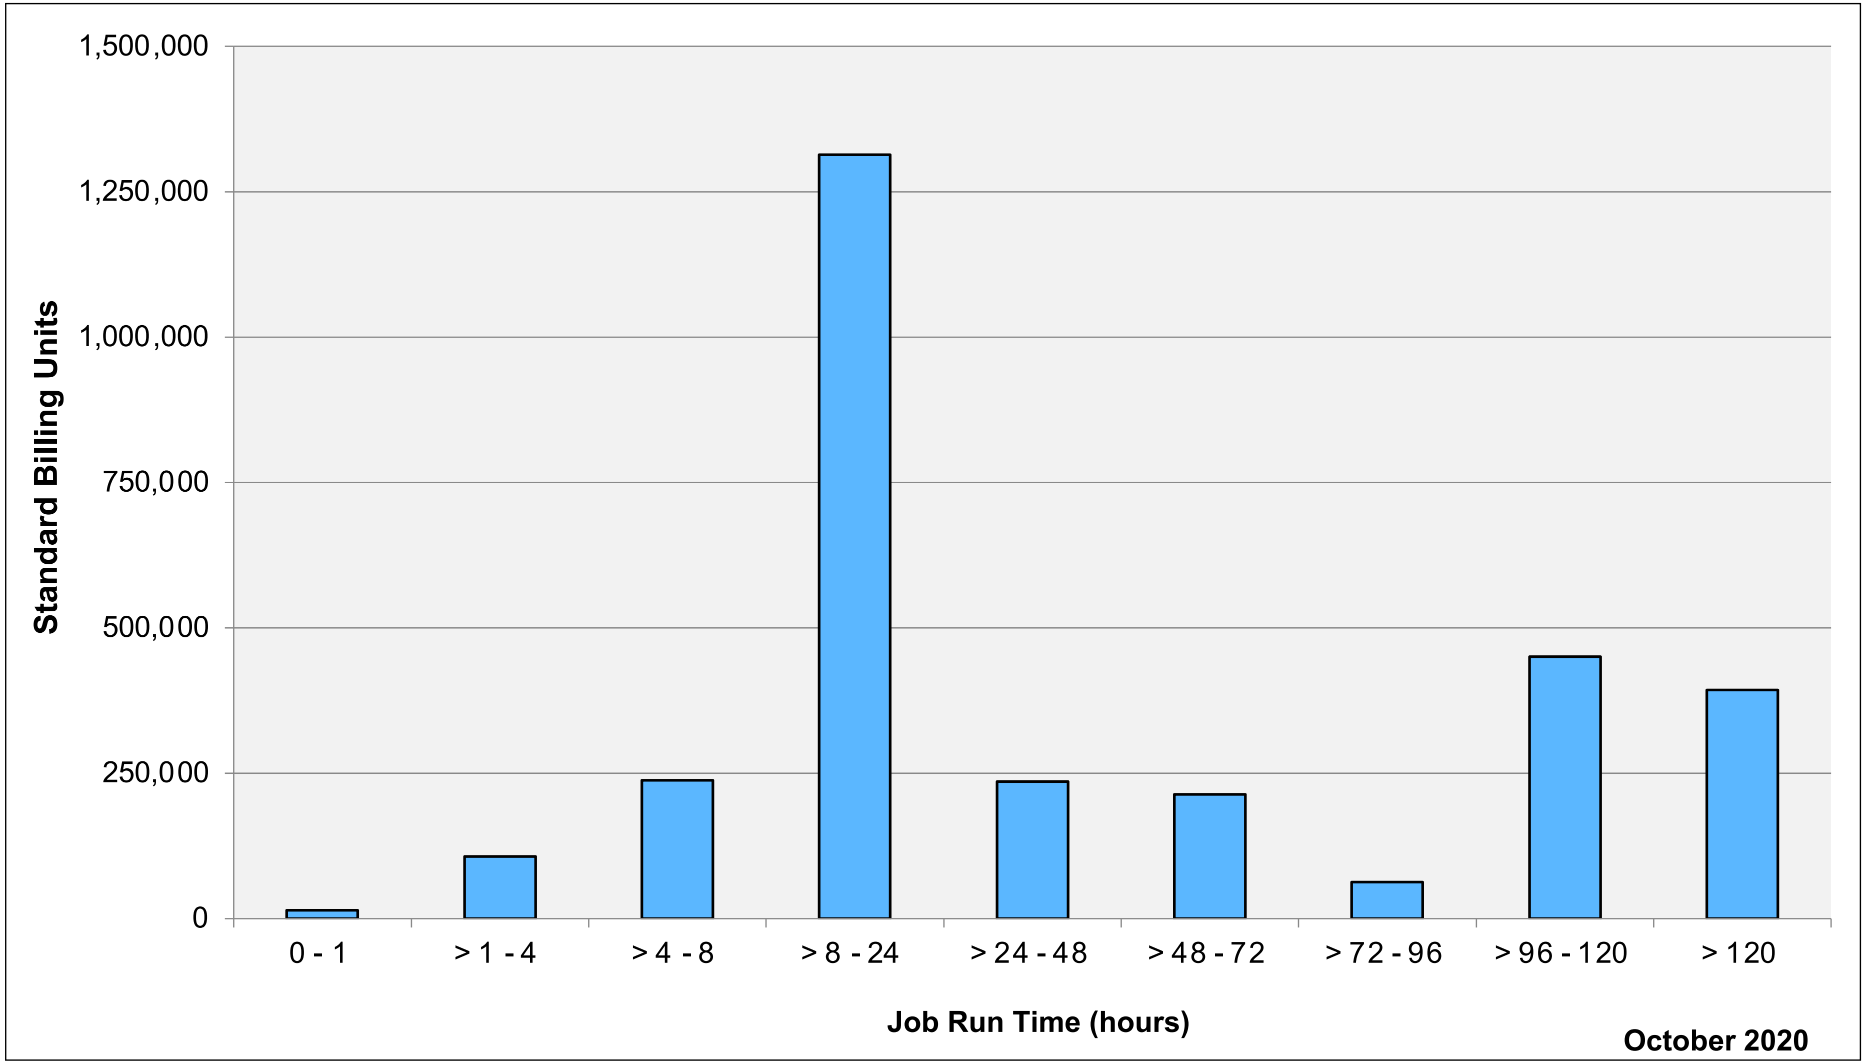

Electra: Monthly Utilization by Job Size

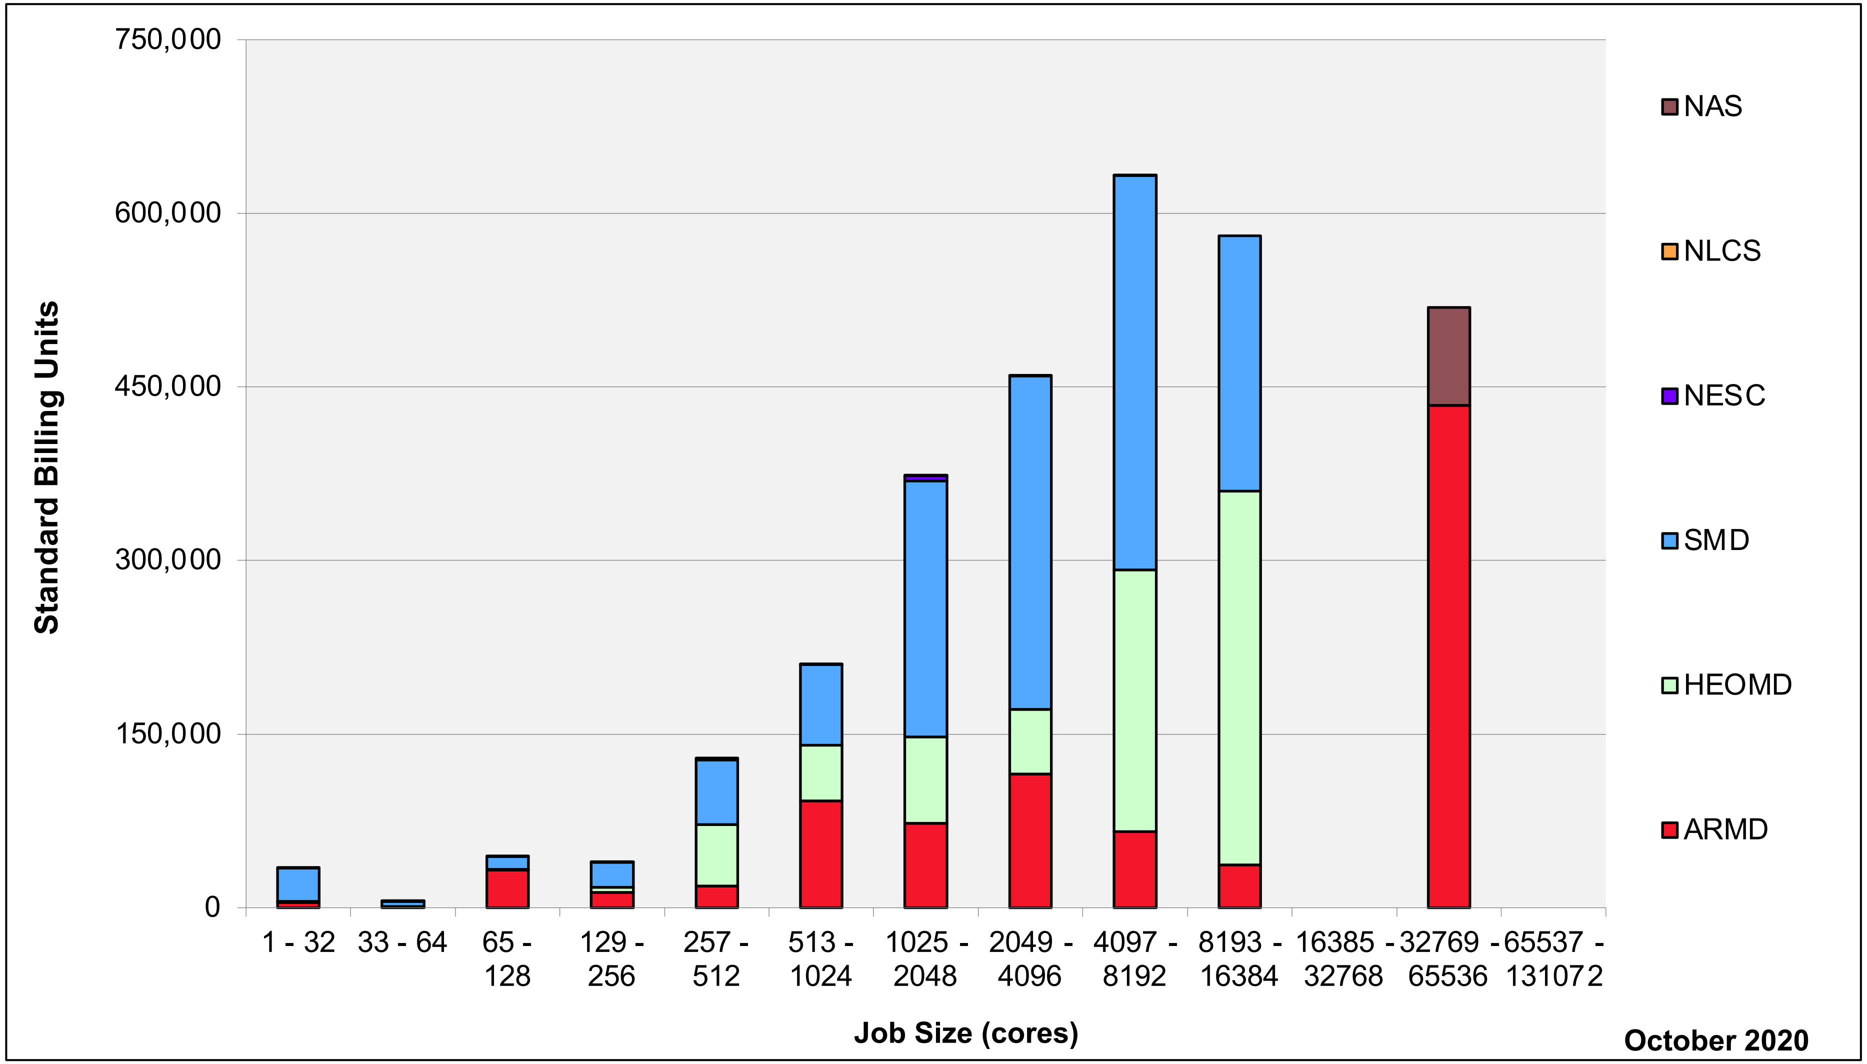

Electra: Monthly Utilization by Size and Length

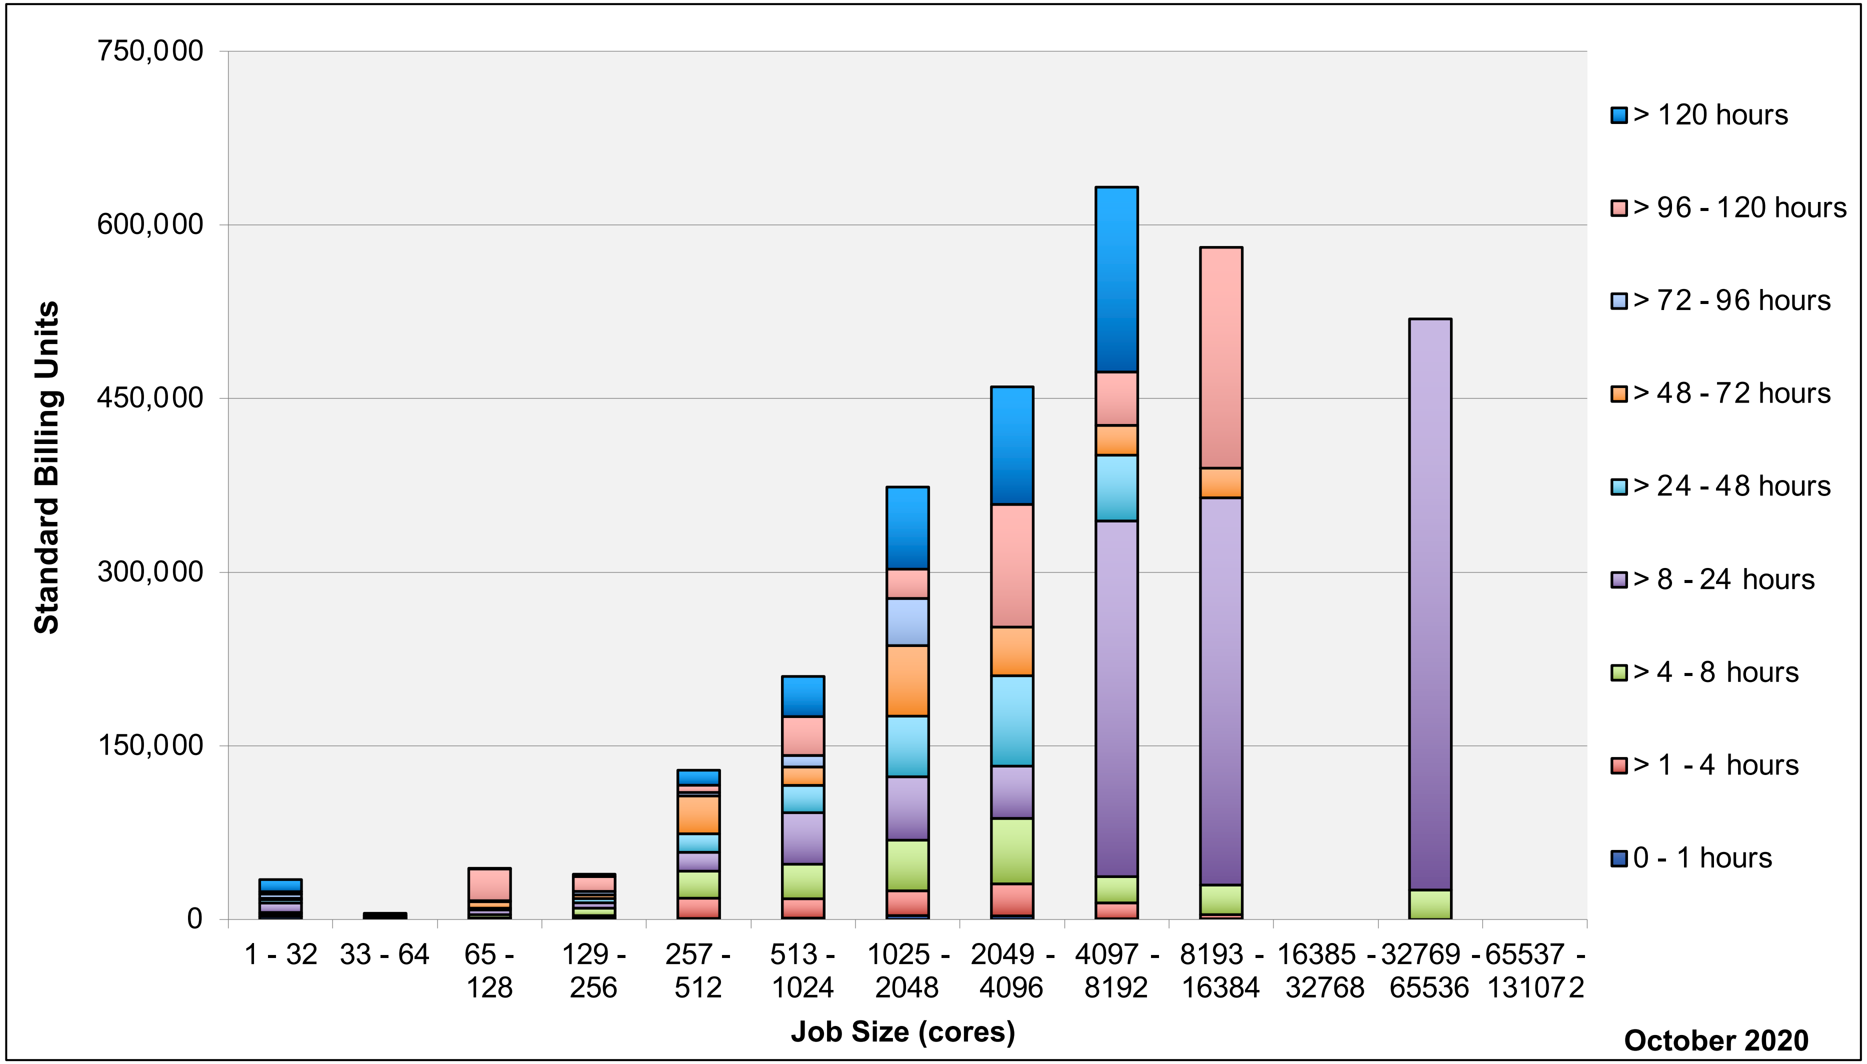

Electra: Average Time to Clear All Jobs

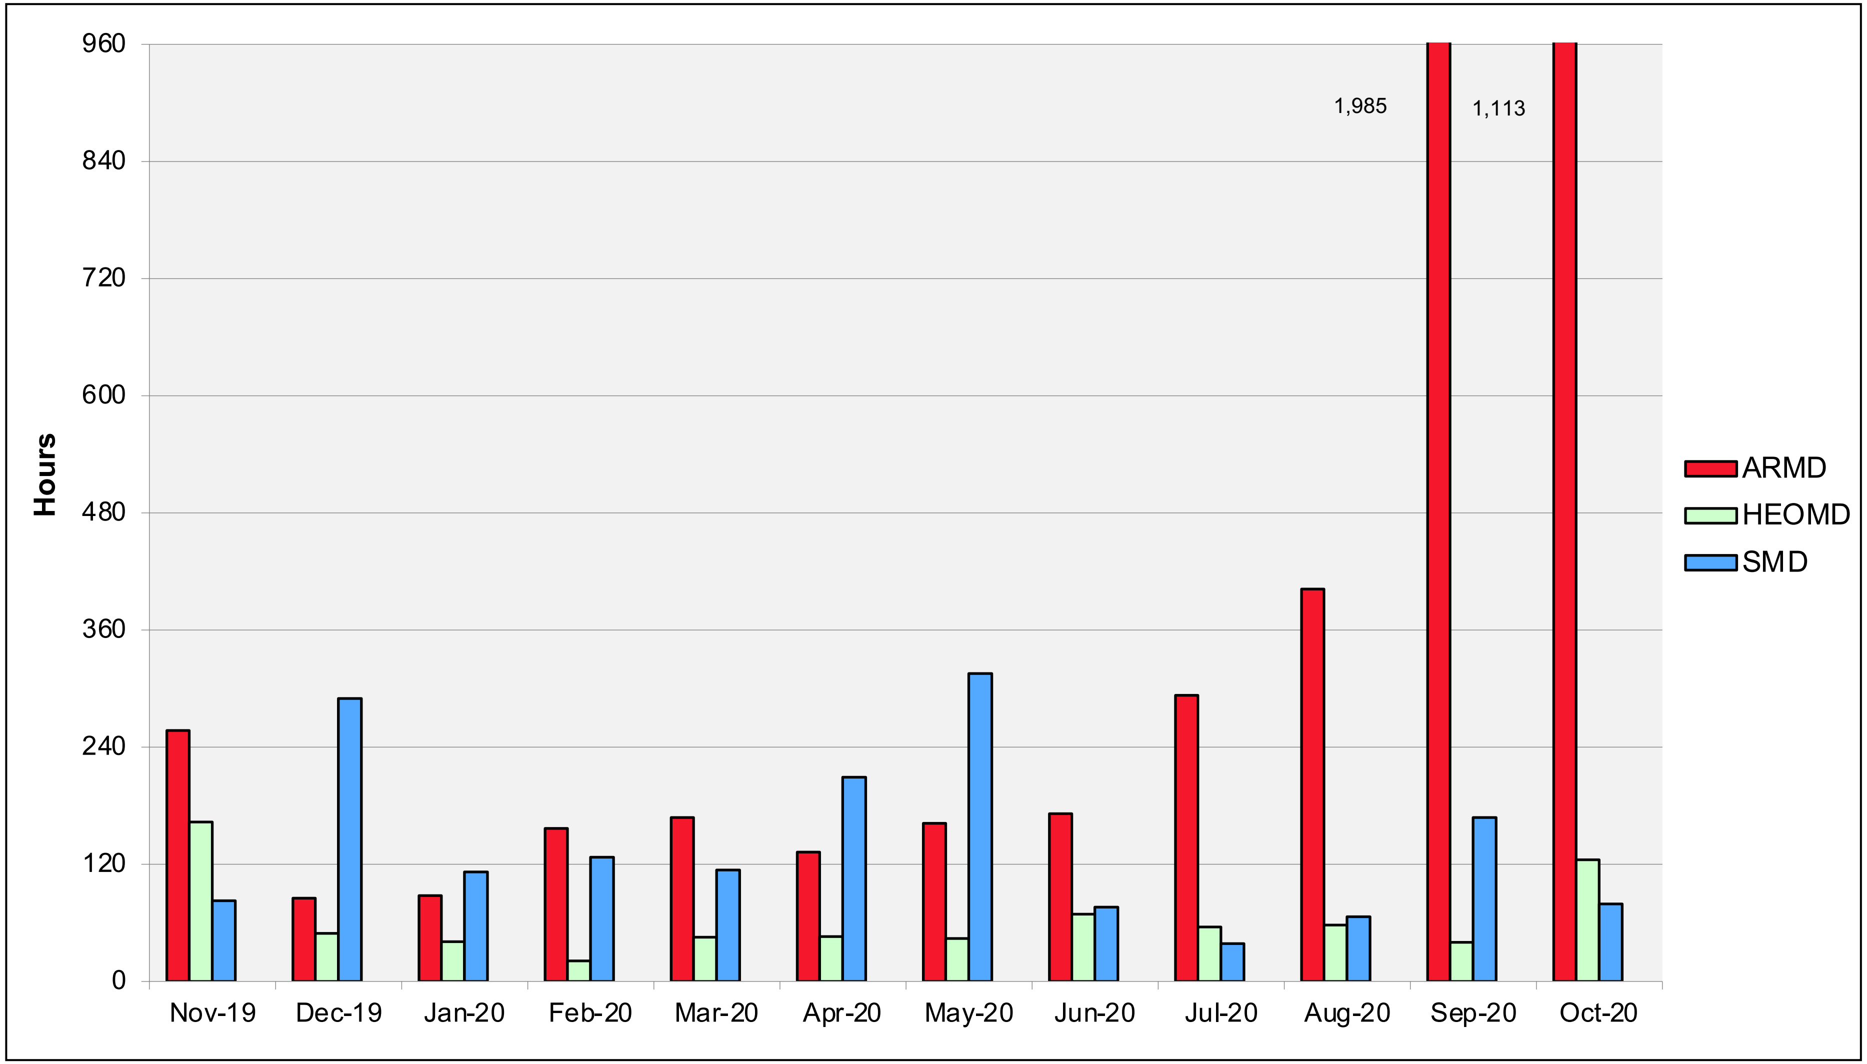

Electra: Average Expansion Factor

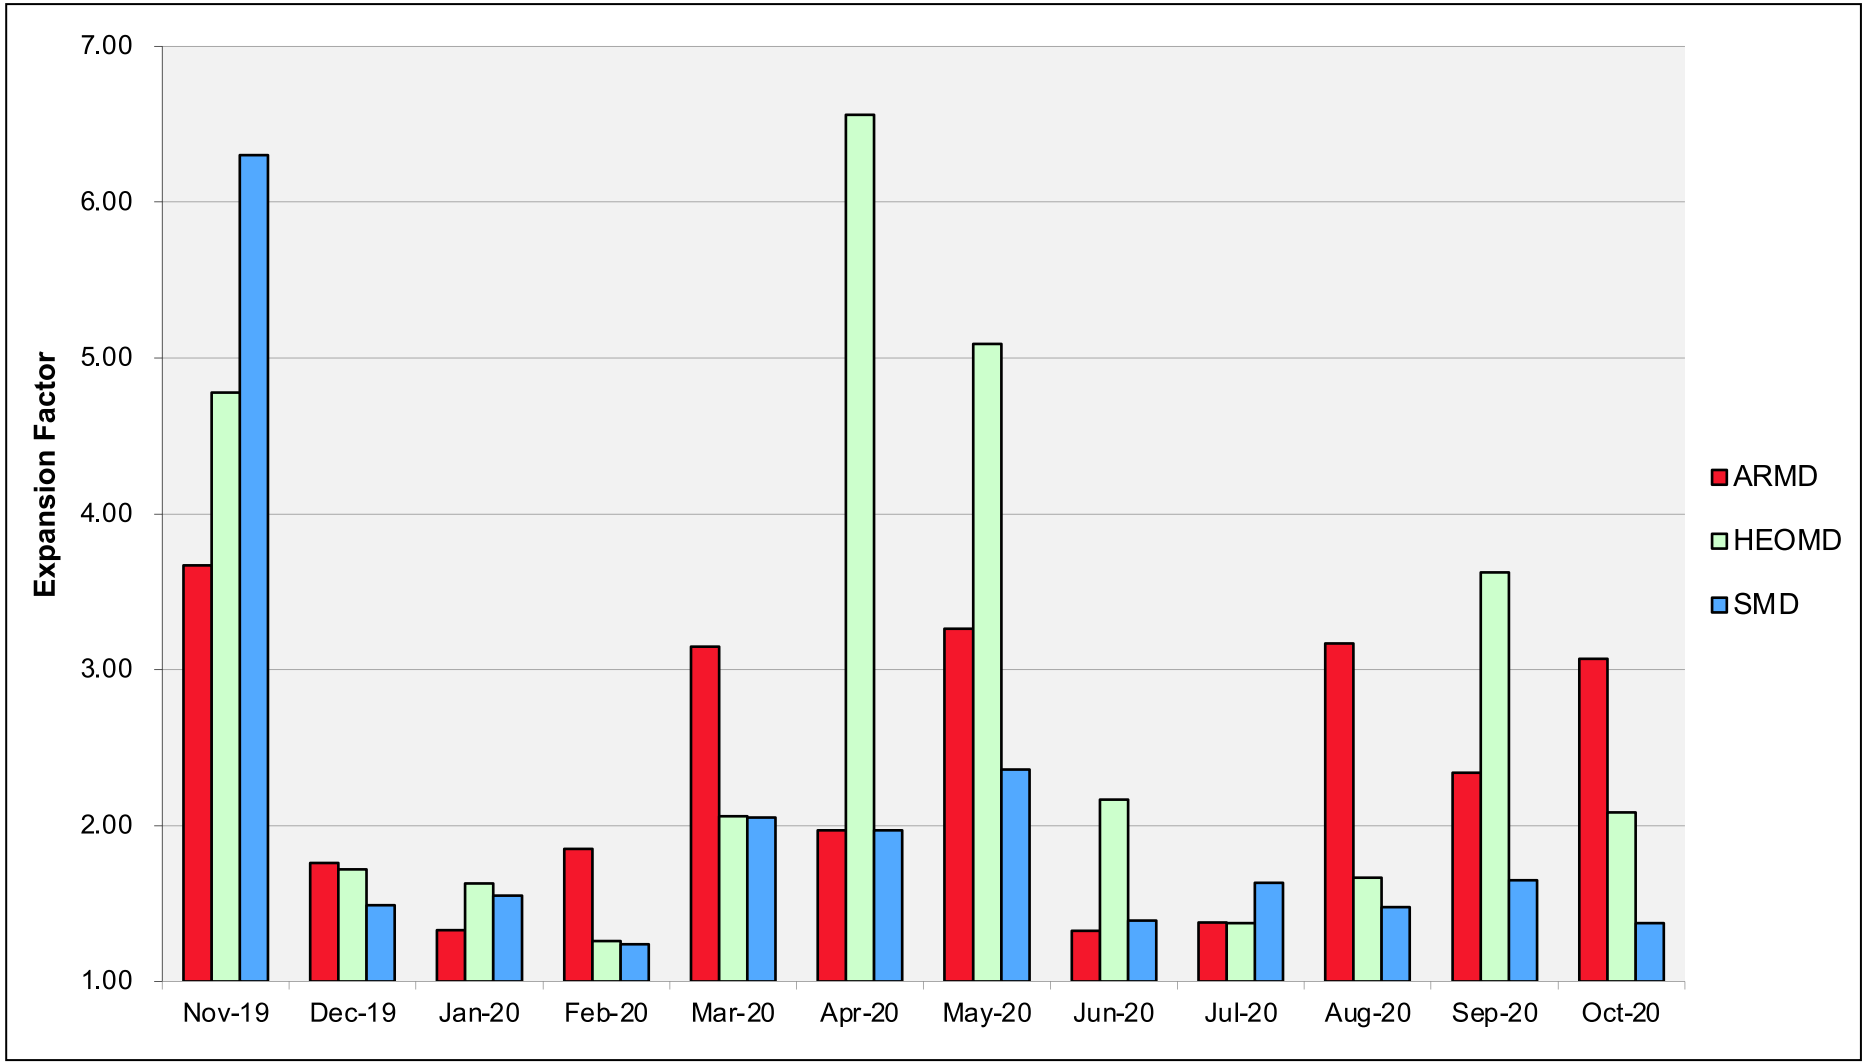

Merope: SBUs Reported, Normalized to 30-Day Month

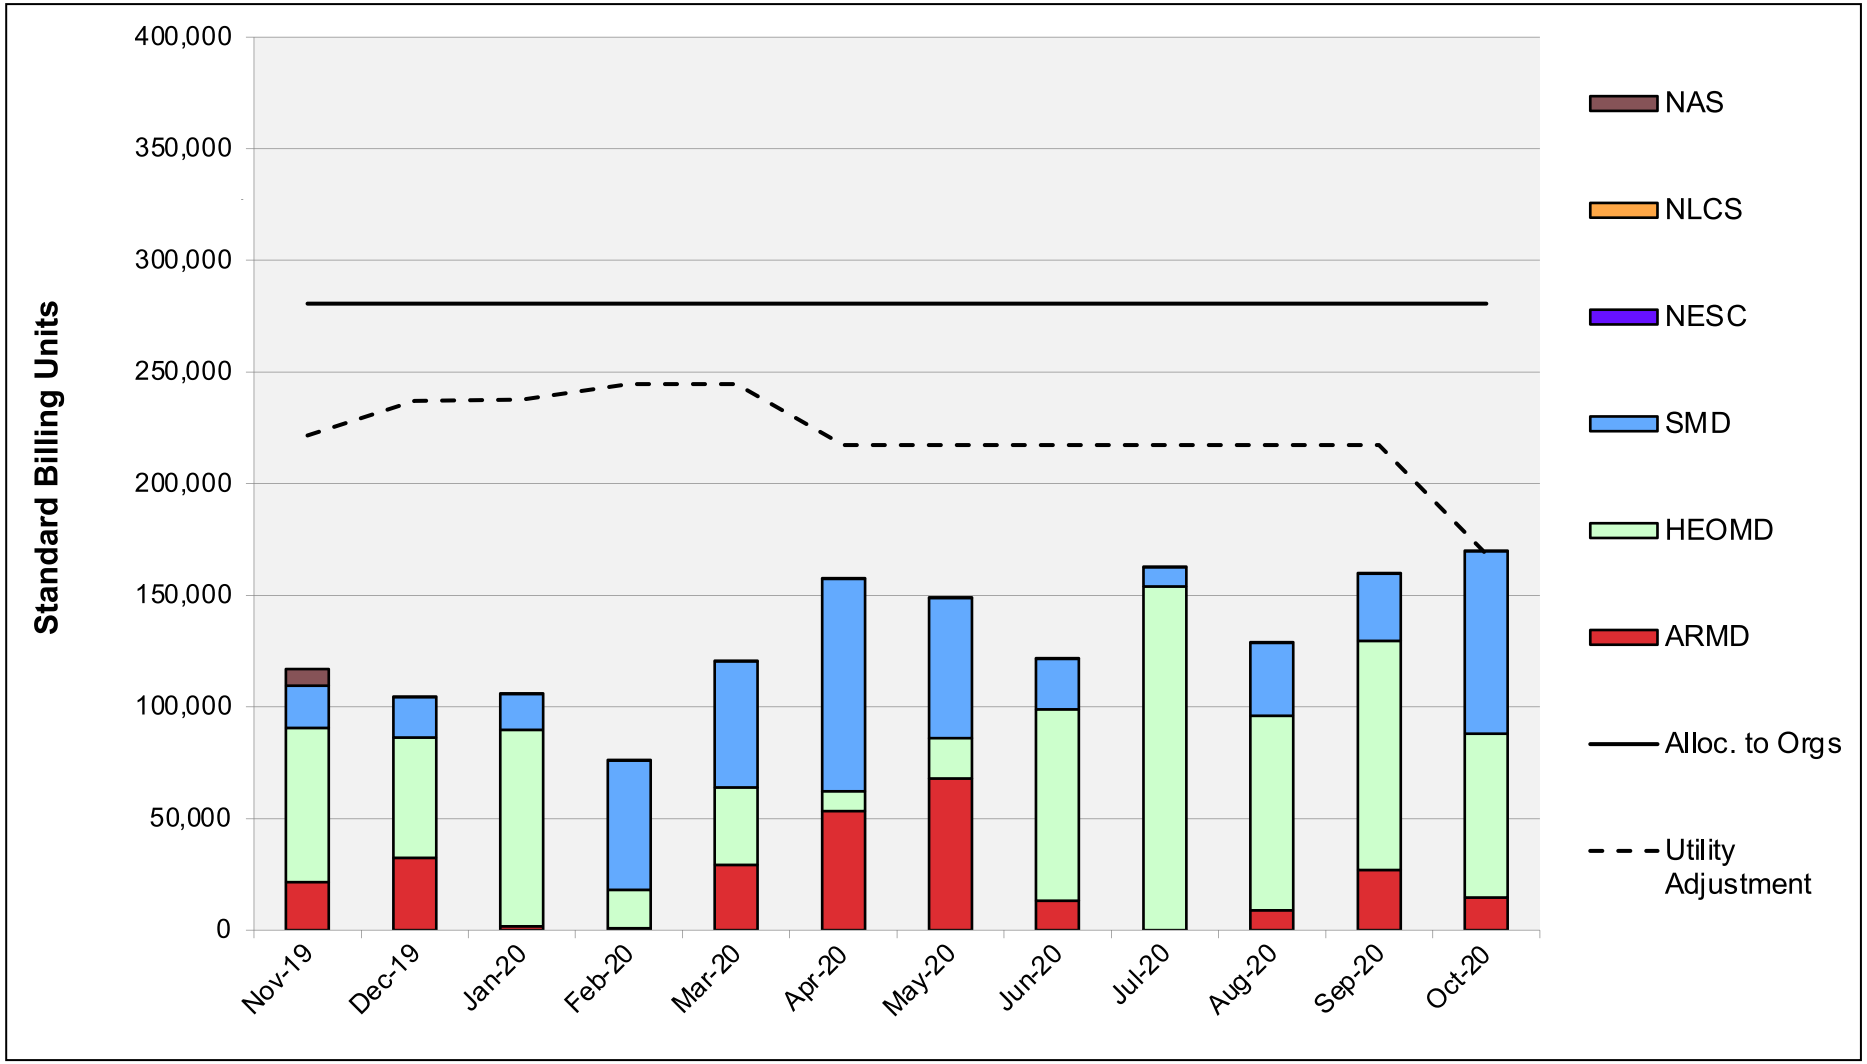

Merope: Monthly Utilization by Job Length

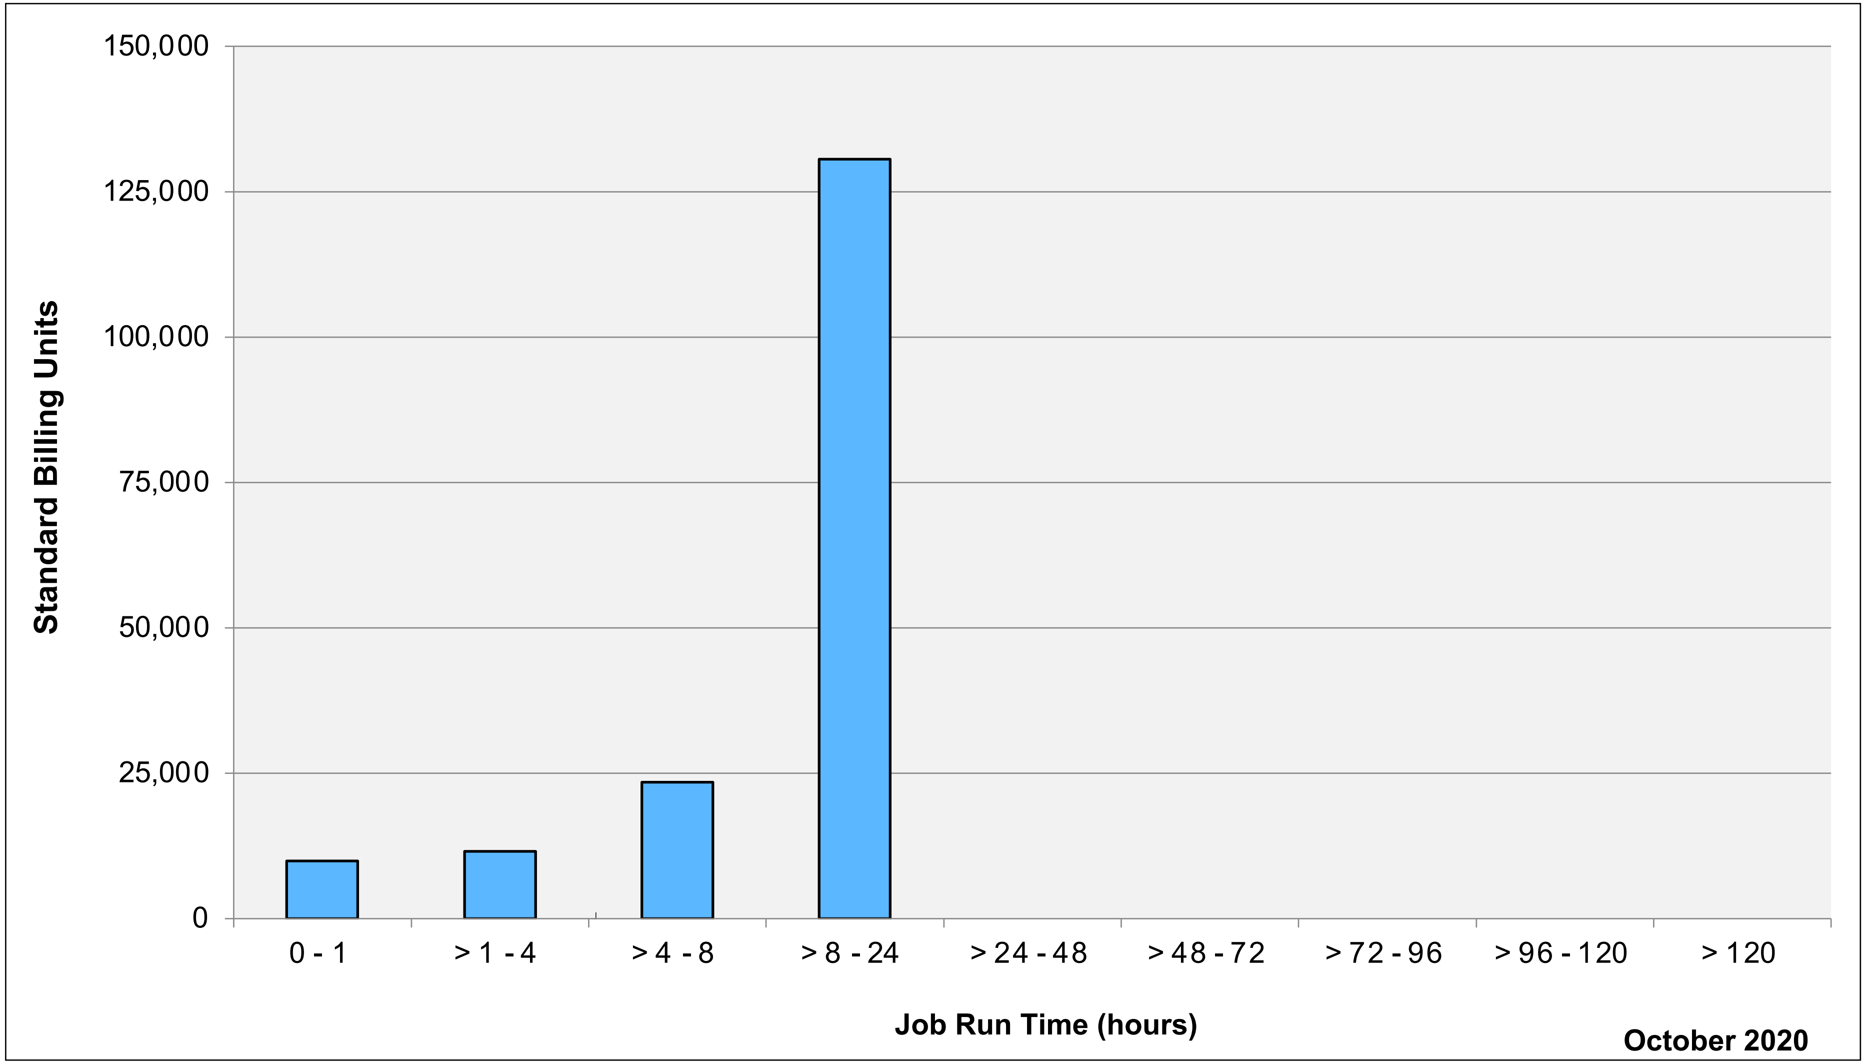

Merope: Monthly Utilization by Job Size

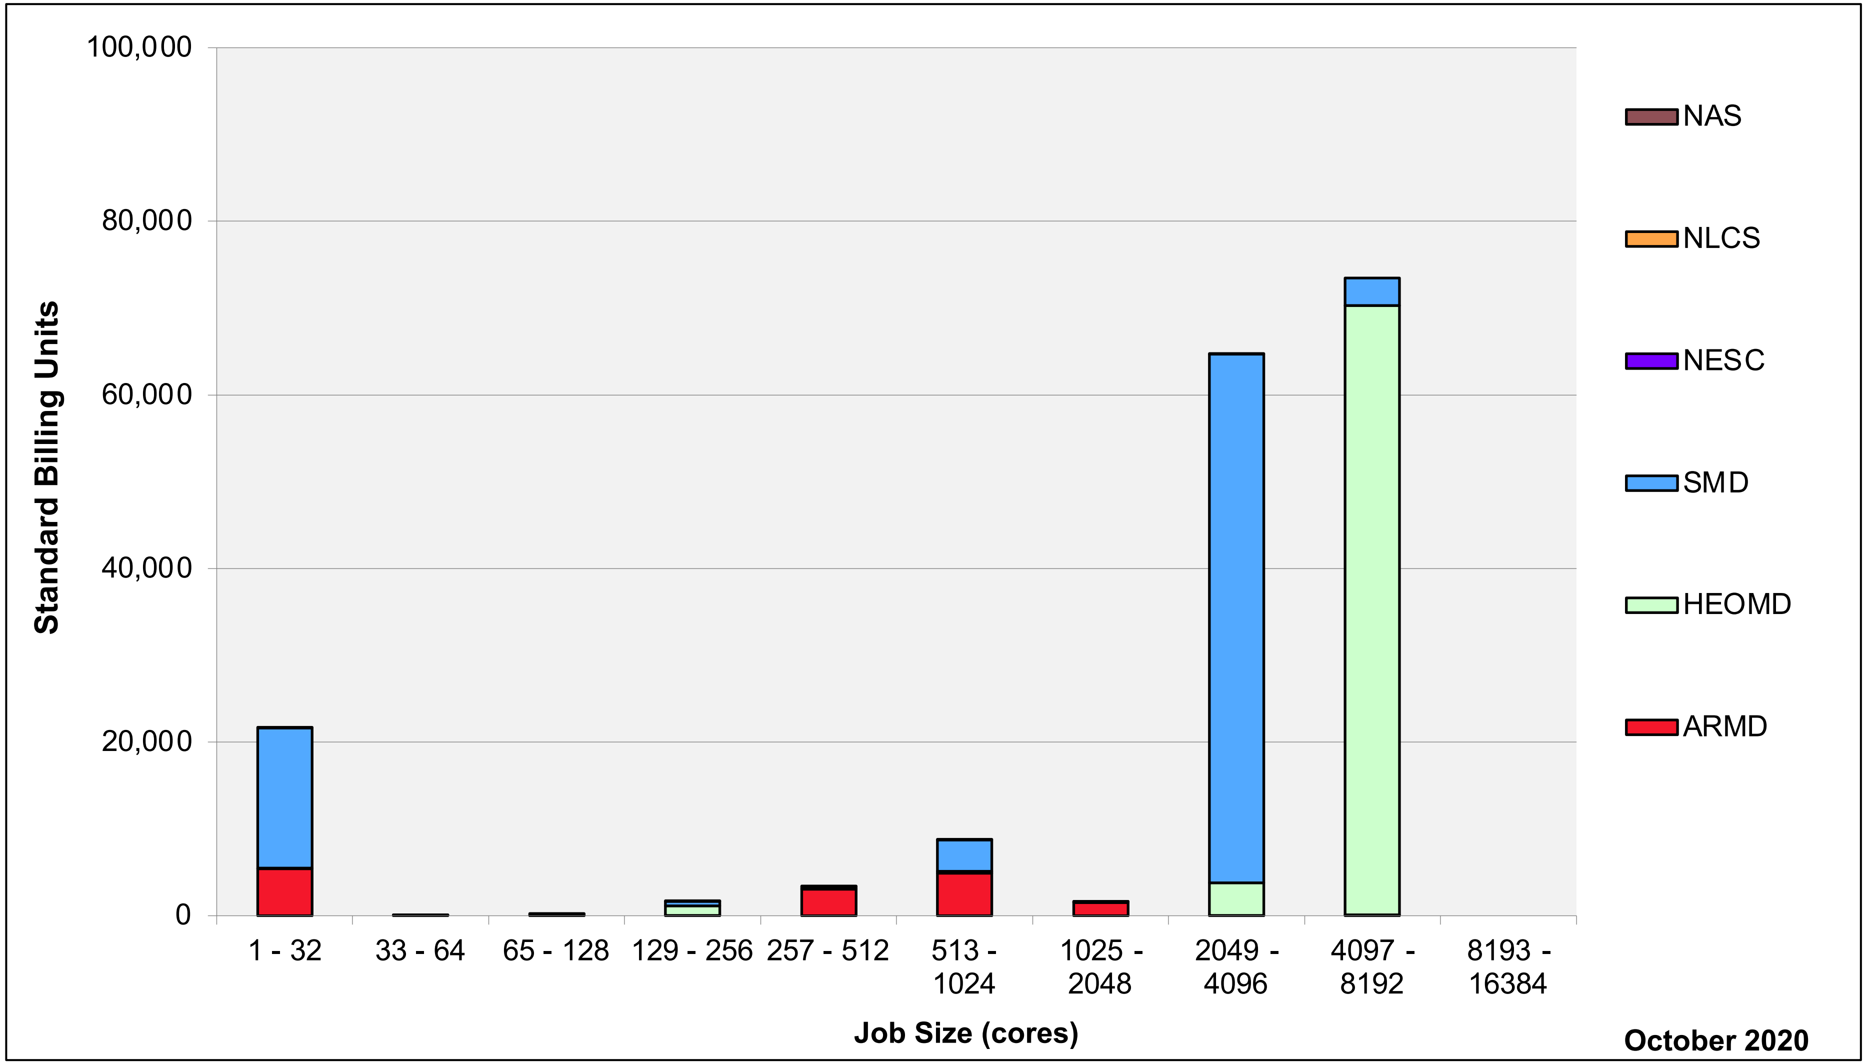

Merope: Monthly Utilization by Size and Length

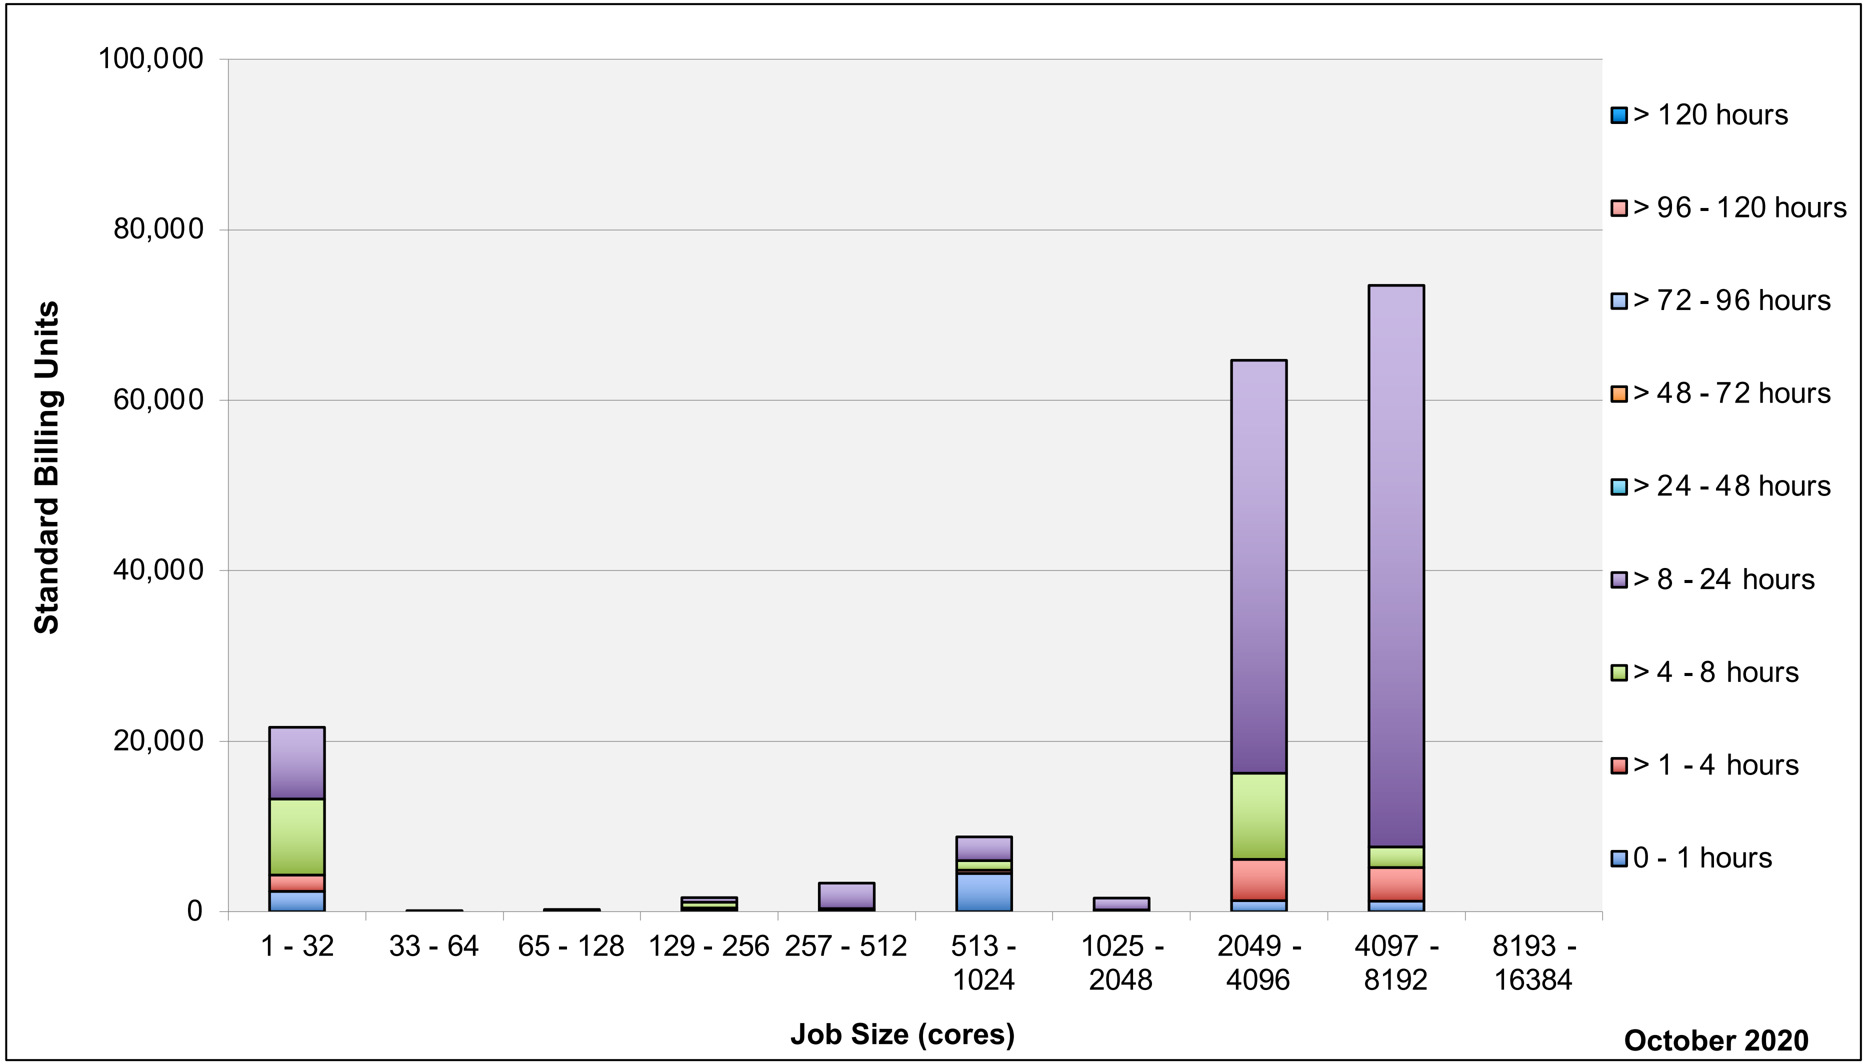

Merope: Average Time to Clear All Jobs

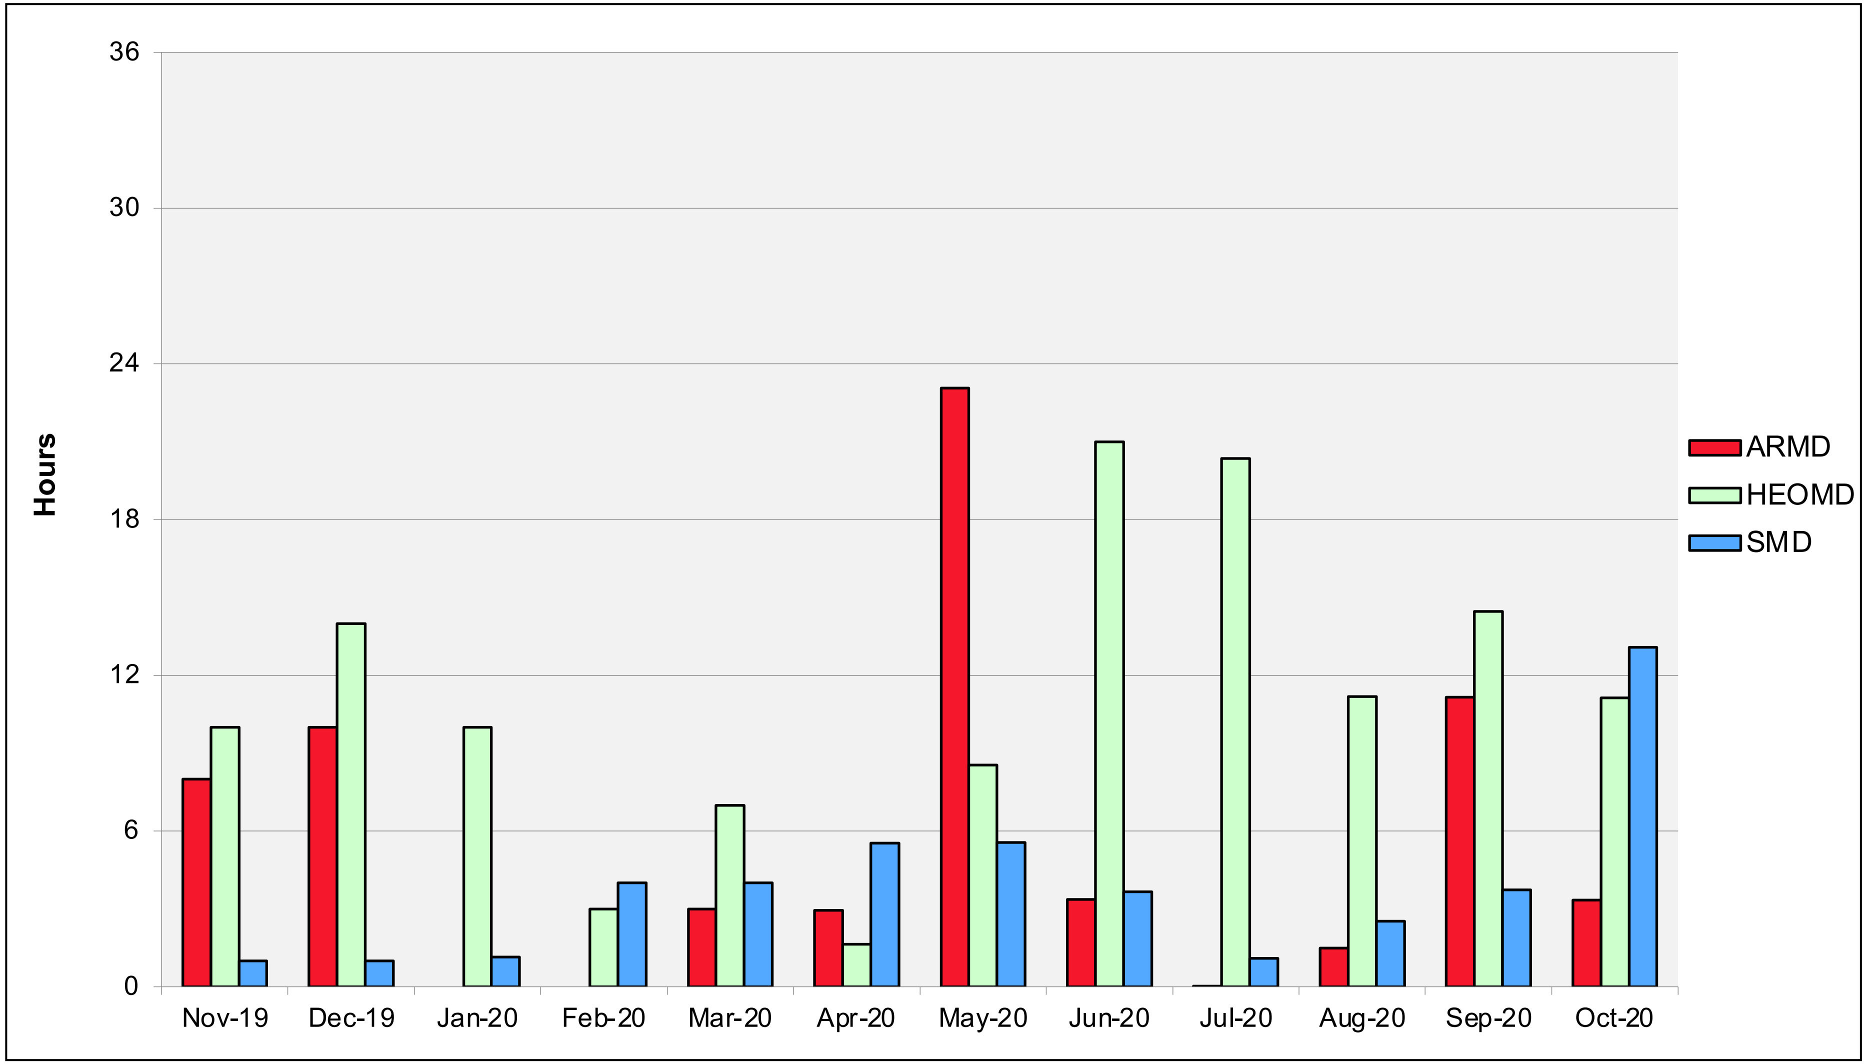

Merope: Average Expansion Factor

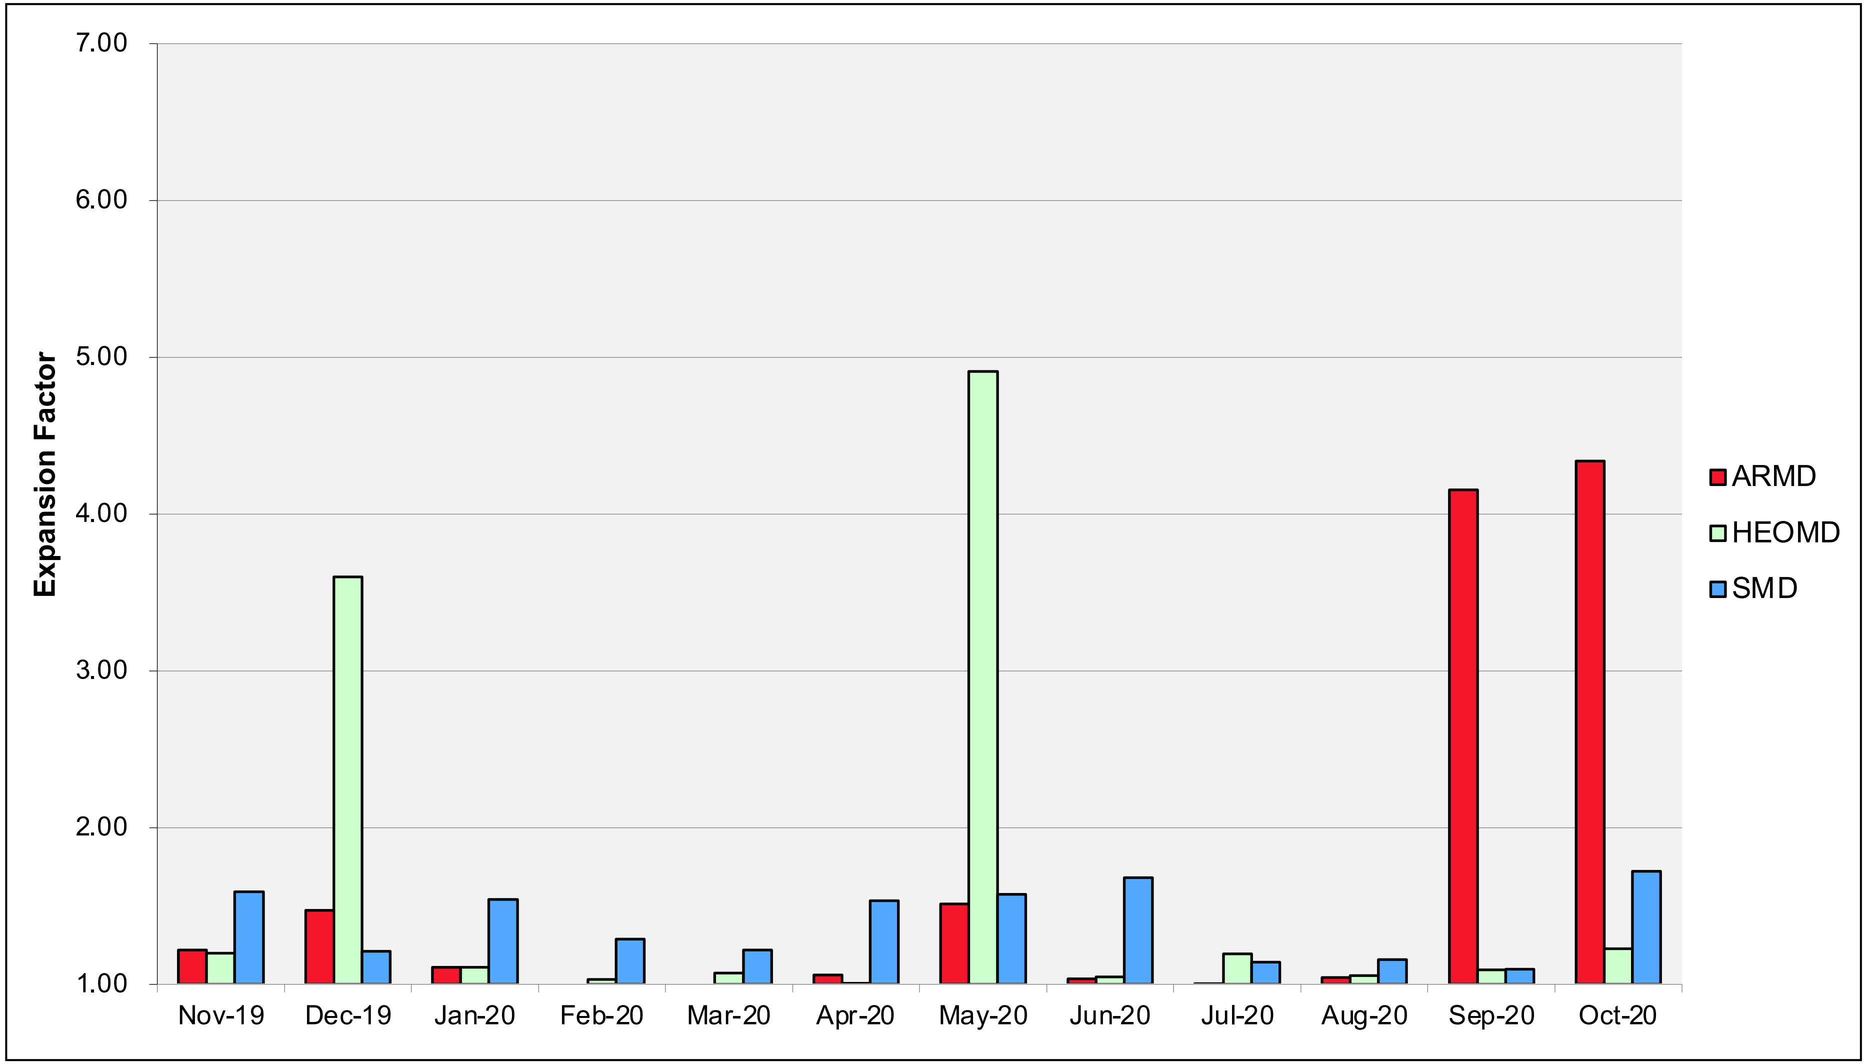

Endeavour: SBUs Reported, Normalized to 30-Day Month

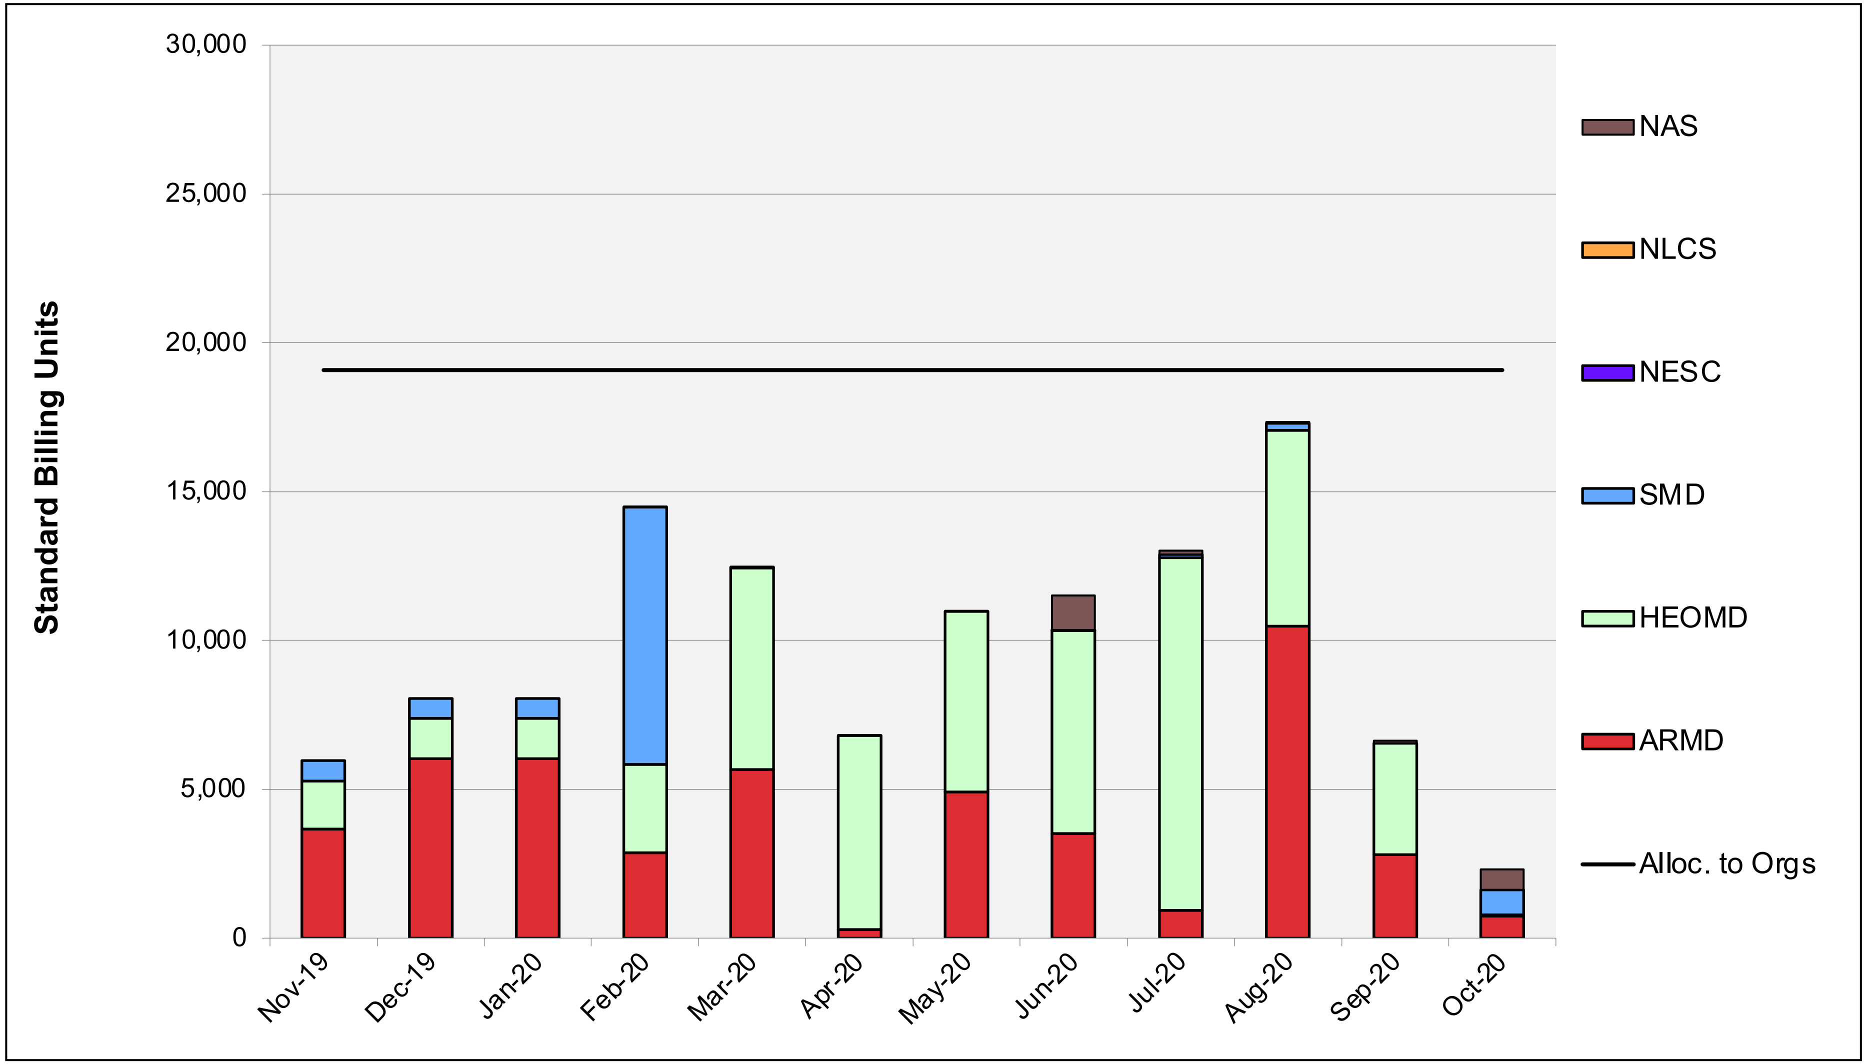

Endeavour: Monthly Utilization by Job Length

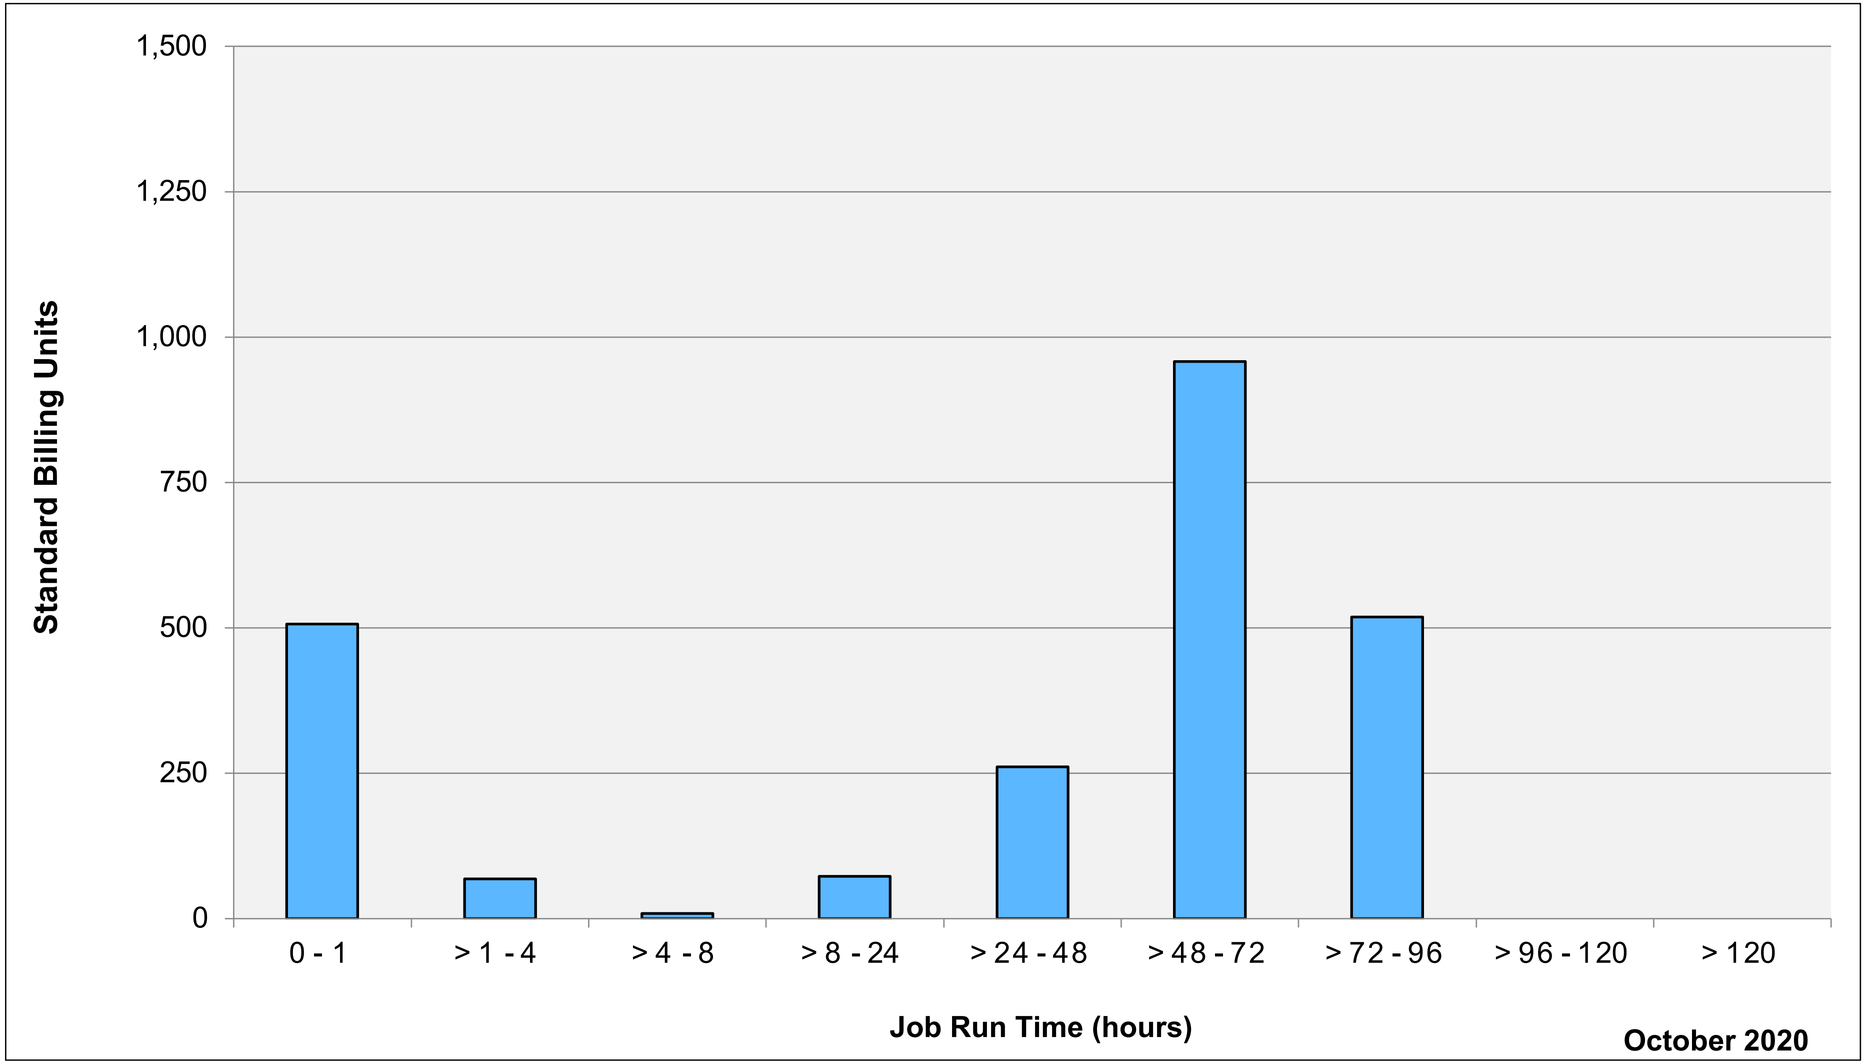

Endeavour: Monthly Utilization by Job Size

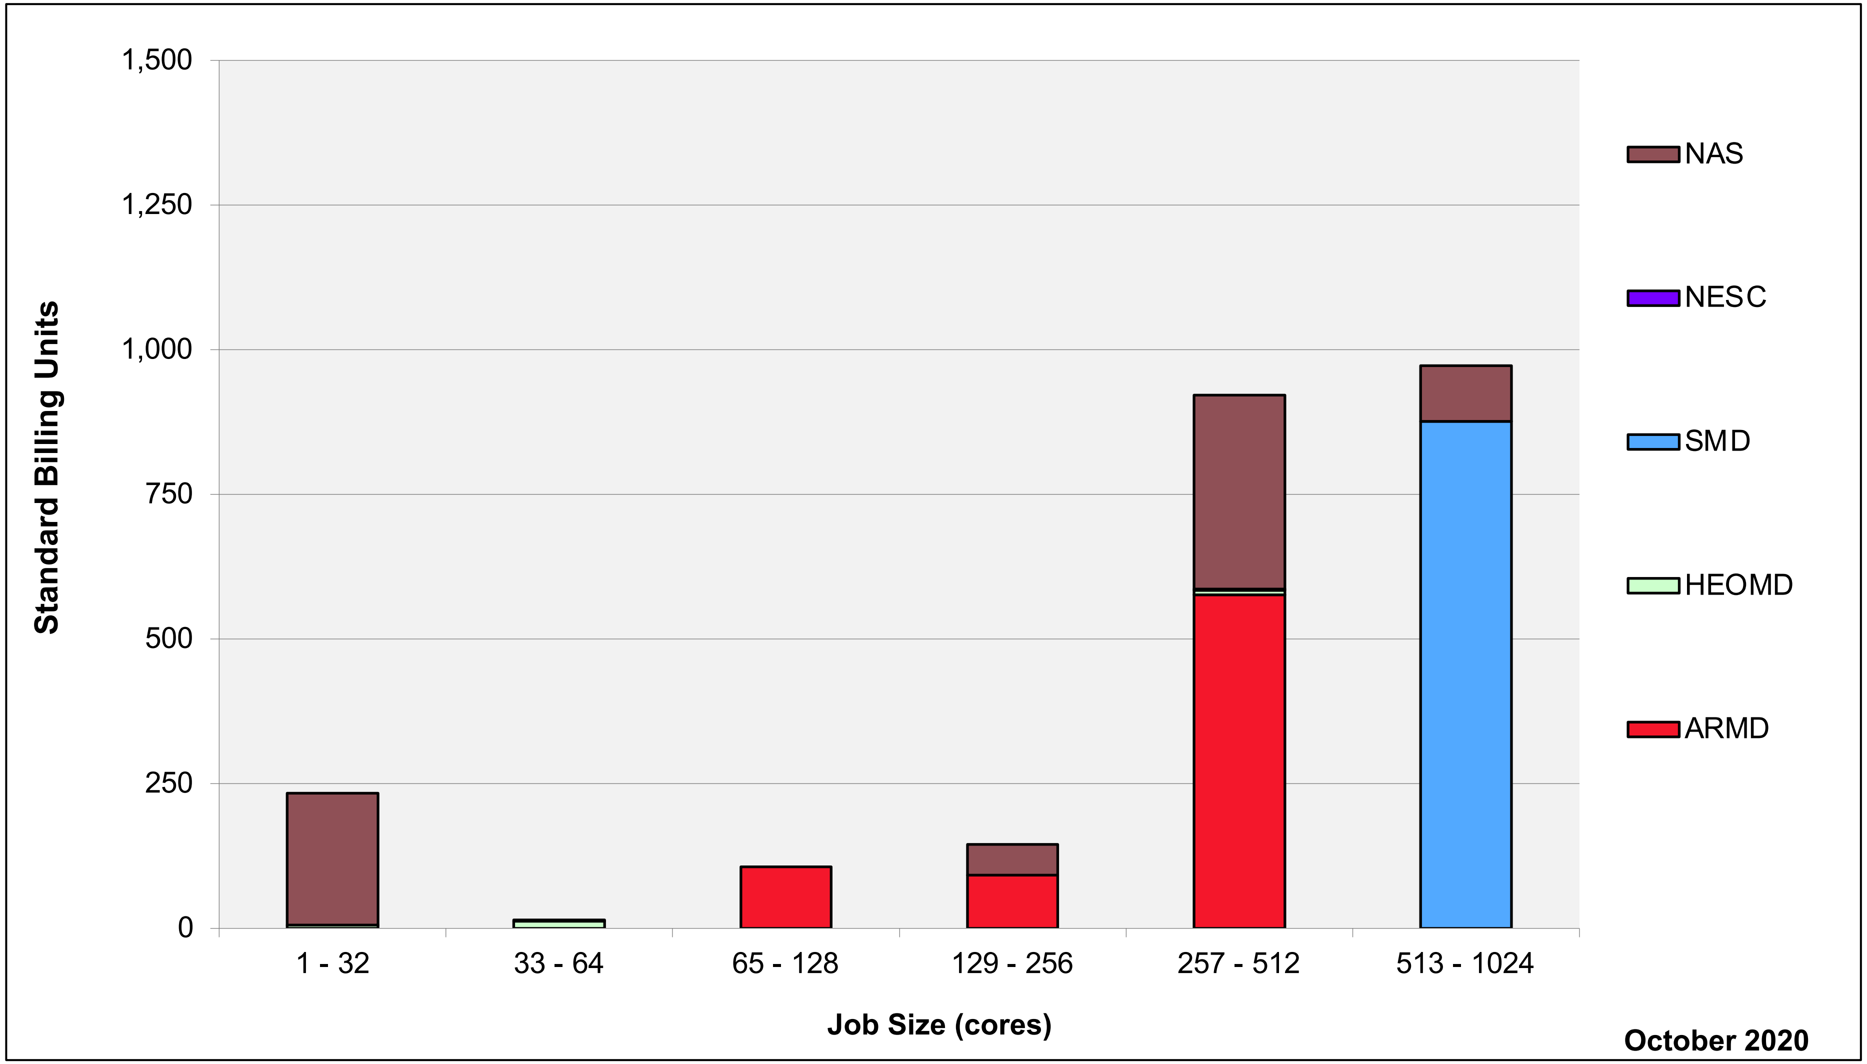

Endeavour: Monthly Utilization by Size and Length

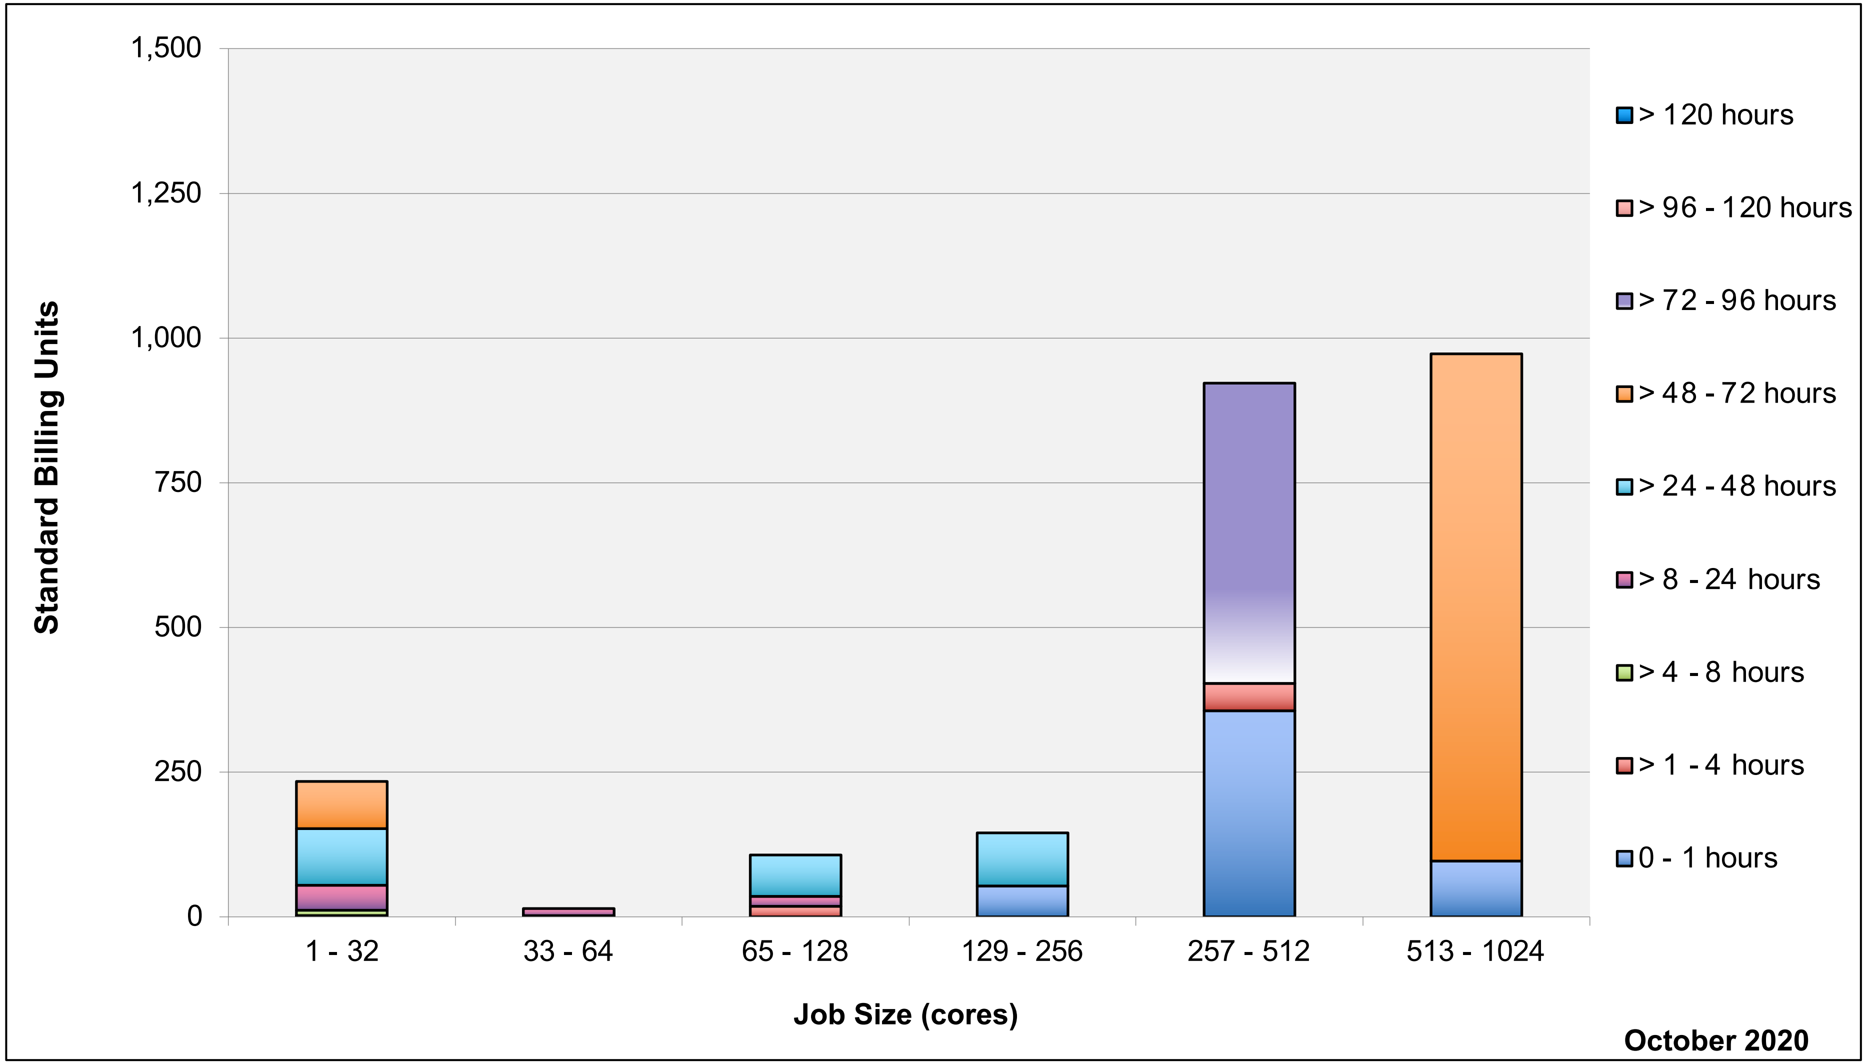

Endeavour: Average Time to Clear All Jobs

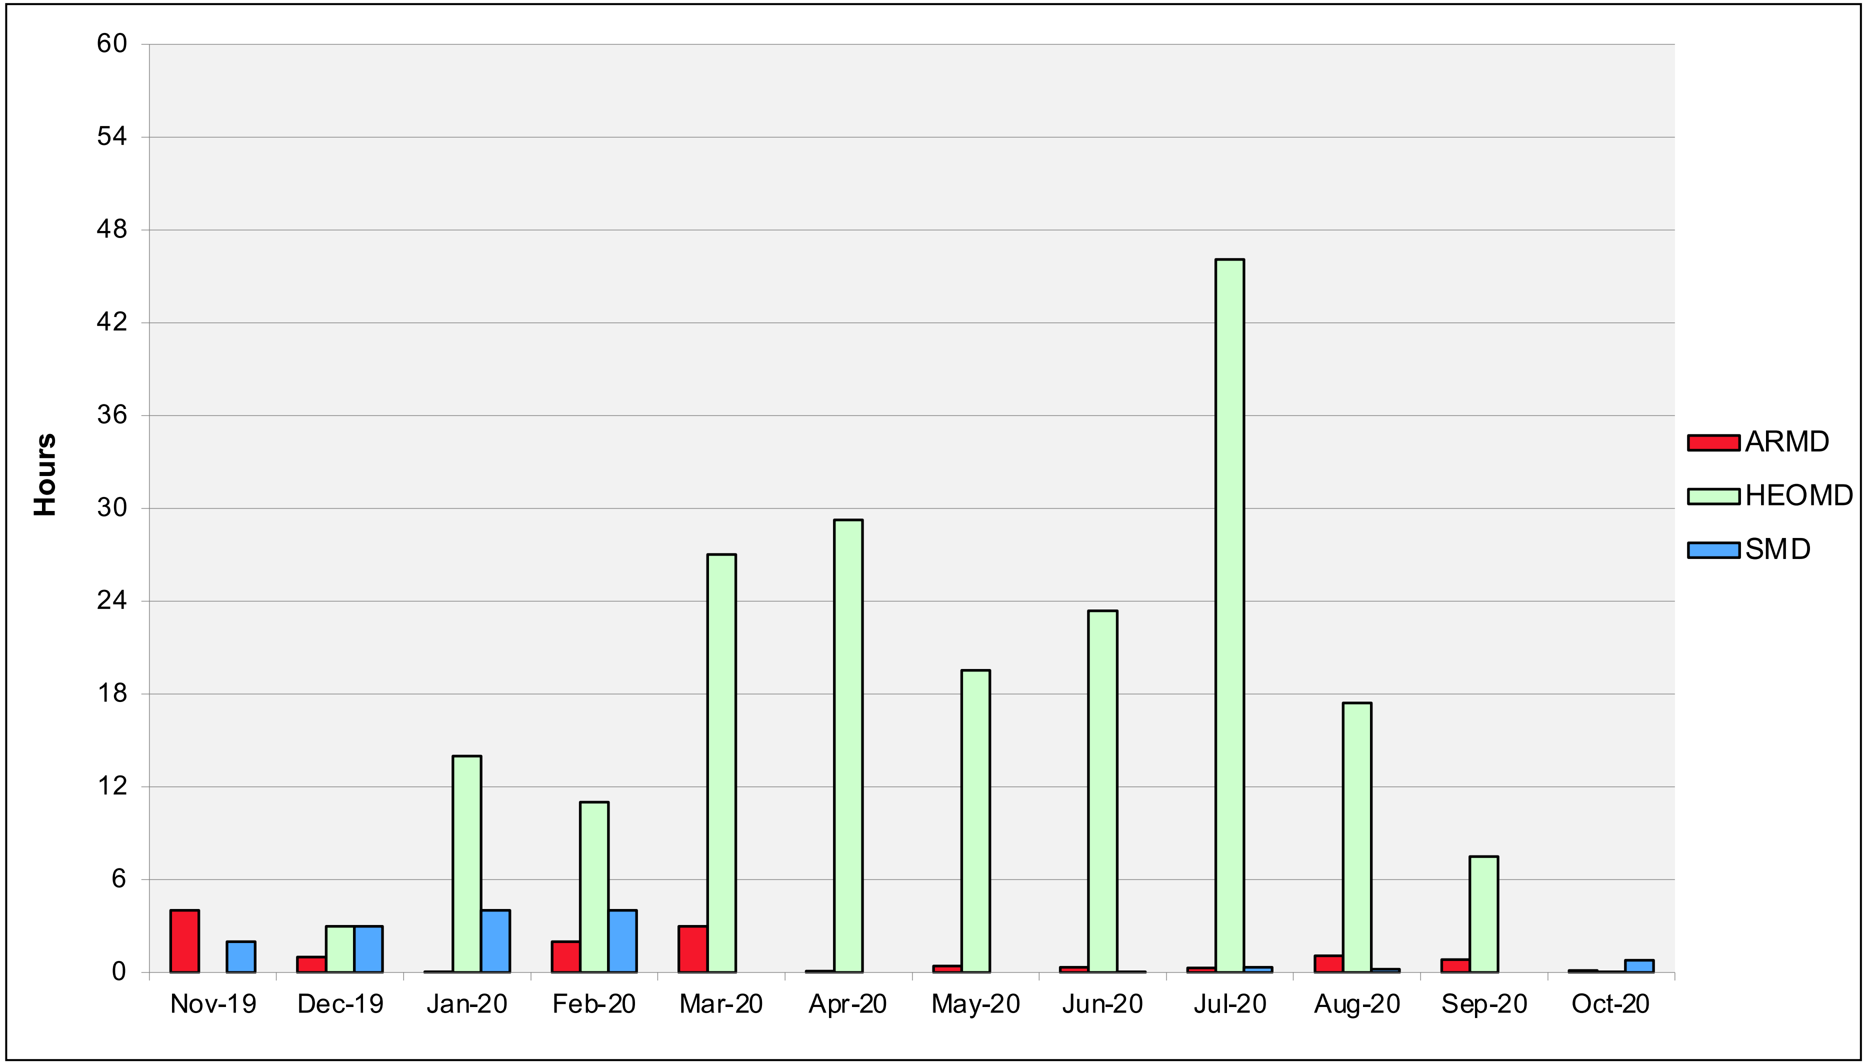

Endeavour: Average Expansion Factor