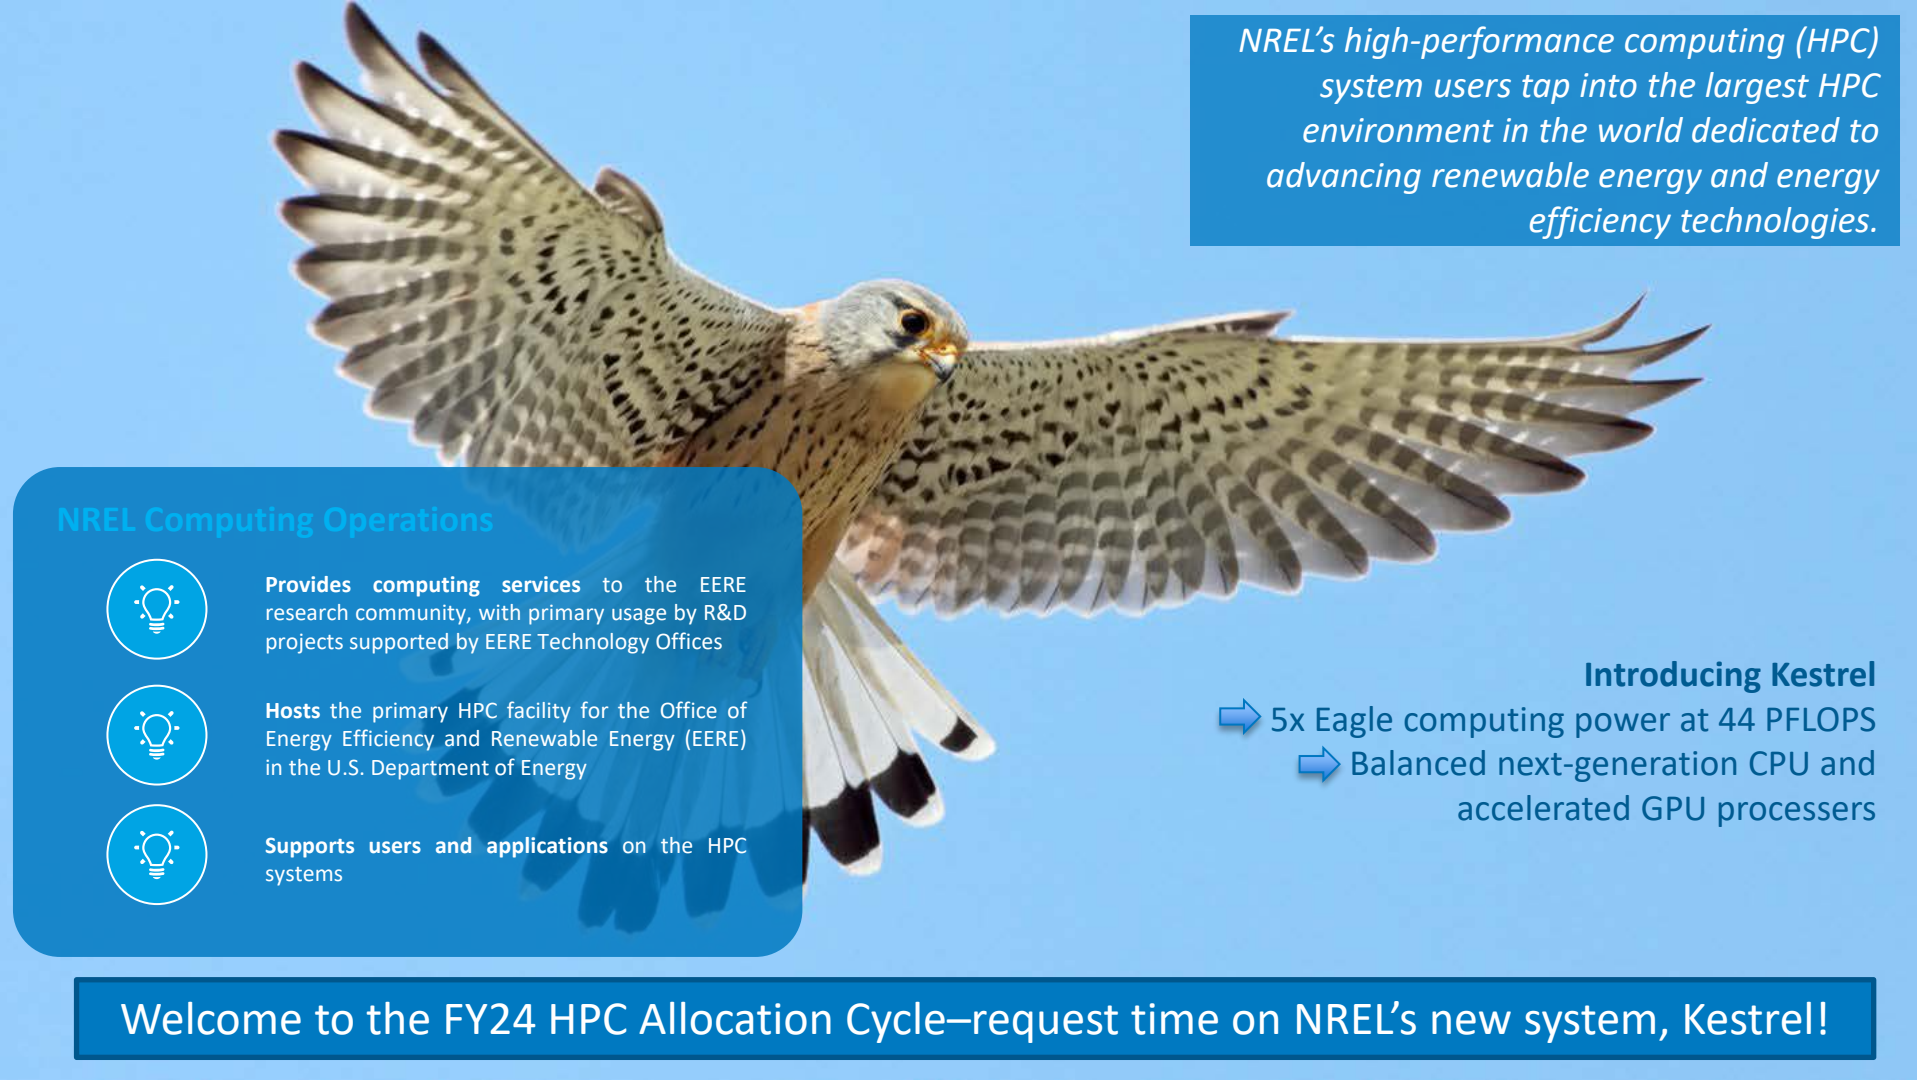




NREL HPC FY24 Allocation Cycle 05/18/2023 & 5/24/2023

Kris Munch
Aaron Andersen



NREL's high-performance computing (HPC) system users tap into the largest HPC environment in the world dedicated to advancing renewable energy and energy efficiency technologies.

NREL Computing Operations



Provides computing services to the EERE research community, with primary usage by R&D projects supported by EERE Technology Offices



Hosts the primary HPC facility for the Office of Energy Efficiency and Renewable Energy (EERE) in the U.S. Department of Energy



Supports users and applications on the HPC systems

Introducing Kestrel

- ➡ 5x Eagle computing power at 44 PFLOPS
- ➡ Balanced next-generation CPU and accelerated GPU processors

Welcome to the FY24 HPC Allocation Cycle—request time on NREL's new system, Kestrel!

FY24 Summary



In case you are just here for the highlights...

- **Allocations will be made on Kestrel for FY24**
- Eagle is only available for existing projects to transition codes and data to Kestrel, or to finish up projects ending in Q1 of FY24
- Kestrel is designed with both CPU nodes and GPU nodes
- Kestrel GPUs are scheduled to come online Q2 (January 2024)
- Requests for FY24 need to pay particular attention to Kestrel's architecture:
 - Requests must differentiate between the CPU request and GPU request
 - Allocation Unit (AU) factors are different from Eagle and distinct between CPU and GPU nodes
 - Give yourself computing time for transitioning codes and workflows, and for ramping up on Kestrel's new architecture
- **Get your requests in early. If you don't already have a login, request an account now to avoid delays!**

FY24 Request Changes

What is different about this year's Allocation request process?

- Because **GPUs are ~60% of the computational capacity**, you will be asked for separate request amounts for CPU-node AUs and GPU-node AUs.
- CPU and GPU nodes have **updated and distinct AU factors** – calculate and plan accordingly!
 - 1 CPU node hour = 10 AUs
 - 1 GPU node hour = 100 AUs
- With Kestrel's new capacity, the allocation process may be able to consider projects that are not supported by EERE or NREL – for these projects, an **EERE mission impact statement** is required
- Technical Readiness Reviews will pay particular attention to CPU/GPU needs and supporting computational readiness documentation: be responsive to the TRR team's questions!

Please provide timely responses to Technical Readiness Review team's questions!



Hewlett Packard Enterprise

Cray EX supercomputer



Balanced Capability

Next-generation Intel Sapphire Rapids processors
Accelerated NVIDIA H100 GPU processors



Production Use Summer 2023

Installation Began April 2023



High-Speed Connectivity

Slingshot interconnect dragonfly topology
>75 petabytes of parallel file system storage



Compatibility

Showcase warm-water waste-heat recovery strategy in ESIF

Eagle vs Kestrel Technical Details

	Eagle	Kestrel
Peak Performance	8 Pflops peak	44 Pflops peak
Processors	Intel Xeon-Gold Skylake 18 cores	Intel Sapphire Rapids 52 cores
Nodes	2,114 nodes, 2 processors/node, 36 cores/node	2,304 nodes, 2 processors/node, 104 cores/node
Data Storage	14 PetaBytes Lustre	95 Petabytes Lustre
Standard Nodes	1728@96GB, 288@192GB	2304@ 256GB memory
Accelerated Nodes	50 w/2x NVIDIA Tesla V100	132 w/4x NVIDIA H100 Dual socket AMD Genoa
Large Data Nodes	48@768GB memory	10@ 2TB memory
DAV Nodes	6 w/1x NVIDIA Quadro GV100	8 w/2x NVIDIA A40
Interconnect	InfiniBand EDR 8-Dimensional Enhanced Hypercube	HPE Slingshot Dragonfly
Efficient Computing	4.7 gflops/watt	10.4 gflops/watt

Eagle CPU Node vs Kestrel CPU Node

Eagle CPU Node

Dual socket Xeon Gold Skylake (**18 core**)
Total of **36** CPU cores per node
96 GB of memory

Kestrel CPU Node

Dual socket Intel Xeon Sapphire Rapids (**52 core**)
Total of **104** CPU cores per node
256 GB of DDR5 memory

When making your request for CPU AUs, keep in mind:

- Pay attention to the advances of Kestrel's nodes: more cores and more memory
- Know how your codes/workflows utilize resources
- Be prepared to manage jobs to efficiently utilize node hours
- Be sure to account for development time and ramp time for new workflows
- Consider needs of all project members for large projects

Eagle GPU Node vs Kestrel GPU Node

Eagle GPU Node

Dual socket Xeon Gold Skylake (**18 core**)

Total of 36 CPU cores per node

2 X NVIDIA Tesla V100 GPUs

768 GB of memory

Kestrel GPU Node

Dual socket AMD Genoa (**64 core**)

Total of **128** CPU cores per node

4 X NVIDIA H100 SXM GPUs with 80 GB Memory

384 GB of DDR5 memory

When making your request for GPU AUs, keep in mind:

- Pay attention to the advances of Kestrel's GPU nodes: AMD CPUs, SXM technology
- Know how your codes/workflows utilize resources on GPUs and how codes scale across multiple GPUs
- Kestrel jobs can request 1, 2 or 4 GPUs per node
- Be sure to account for development time and ramp time for new workflows
- Consider needs of all project members for large projects

Available Computing Resources

What is an “Allocation Unit” (AU)?

- An “AU” is a defined amount of compute time on the system, which is consistent, theoretically, between machines. An “AU” on Eagle is theoretically the same amount of compute as an “AU” on Kestrel
- On Kestrel: 1 CPU node hour = 10 AUs and 1 GPU node hour = 100 AUs
- When calculating your requested AUs, be mindful of the Kestrel architecture and how it is different from Eagle

What’s Available for FY24?

Resources to be allocated on Kestrel include 274 million AUs across both CPU nodes and GPU nodes:

- 192 M AUs on CPU nodes
- 82 M AUs on GPU nodes (starting Q2)
- 95 Petabytes of storage

How are resources allocated?

- 85% allocated by EERE
- 15% allocated by NREL for NREL projects: LDRD, SPP, non-EERE DOE (ARPA-E, Office of Science), DoD, etc.

Eligibility

What projects are eligible?

- EERE-funded projects at all national laboratories and universities
- EERE-mission impact projects supported by academia or industry, regardless of funding source
- NREL projects, regardless of funding source.

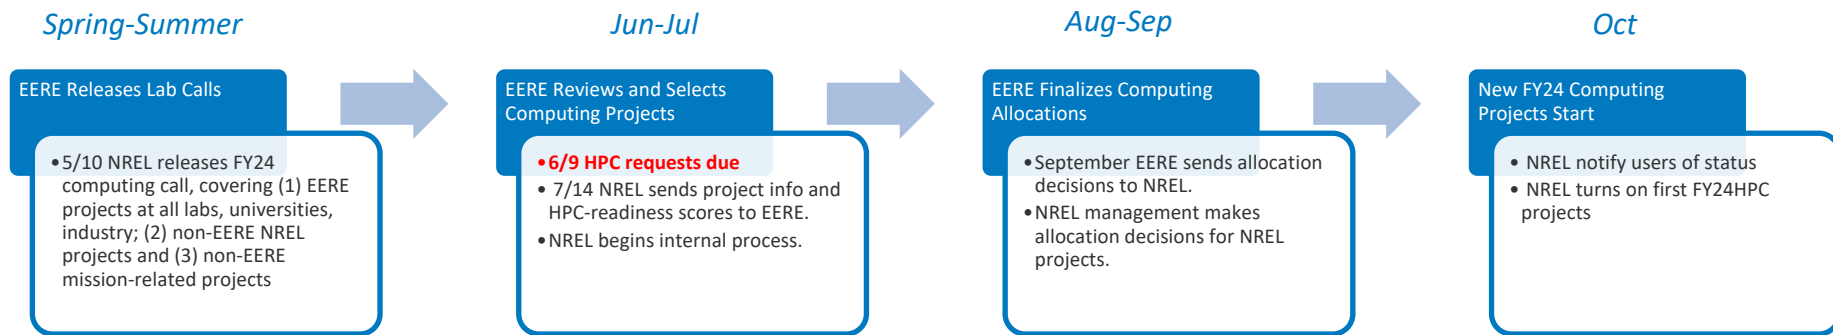
Projects should apply for allocations if: (1) funded, (2) a proposal is under consideration, or (3) if a proposal will be submitted during summer 2023.

One allocation per funded project! *Do not split projects up into multiple allocations, even if the project is scattered across several labs. Do not use large “umbrella allocations” to cover multiple projects.*

FY24 Allocation Process Schedule

NREL is not the decision maker for EERE projects: our role is to provide EERE with the information needed to make decisions.

- A separate NREL decision-making process is in place for NREL projects.
- Allocation process is designed to be parallel to funding process.



Request compute time now if you think you will need it! Updates and mid-year requests are possible, but in-cycle requests are easier to accommodate

Technical Readiness Review

NREL provides a TRR of each computing request -> enabling EERE decisions by providing the best possible information on the technical request. The TRR includes assessment of:

- Is the project suitable for HPC?
 - “Ideal” (uses large-scale parallelization capabilities), “Suitable” (can run effectively), and “Unsuited” (not ready or able to run on HPC).
 - “Suitable” is an acceptable outcome to pass TRR.
- Is the project ready (if appropriate) to use GPUs?
- Is the size of the compute and storage request justified?
- Which software will be used? Is the software available?
- NREL staff will follow up with submitter if there are issues.

Please provide timely responses to Technical Readiness Review team's questions!

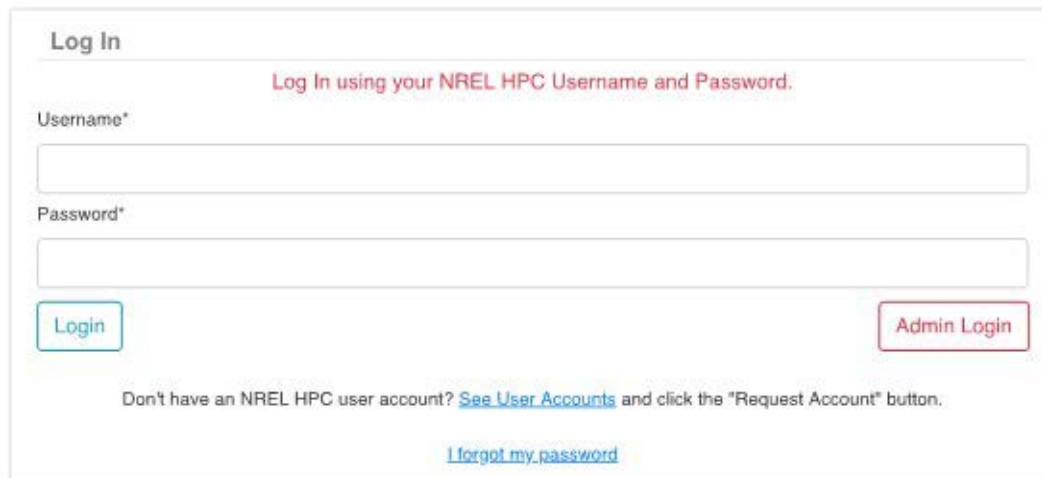
Online Submission: Login

Go to hpcprojects.nrel.gov to prepare your submission.

Get an account ASAP if you plan to submit to avoid last-minute delays!

6/19/2020

Lex



The screenshot shows a web form titled "Log In". Below the title is a red instruction: "Log In using your NREL HPC Username and Password." There are two input fields: "Username*" and "Password*", each with a corresponding text box. Below the "Username*" field is a "Login" button with a blue border. To the right of the "Password*" field is an "Admin Login" button with a red border. At the bottom of the form, there is a link: "Don't have an NREL HPC user account? [See User Accounts](#) and click the "Request Account" button." and another link: "[I forgot my password](#)".

Online Submission: Landing Page

Once logged in, you can submit a FY24 Allocation Request.

- Use the FY23 button if you need a pilot allocation to test code right now, otherwise use FY24.
- If you are updating a FY23 allocation *where you are the lead* there is an option to copy the information over as a starting point.

Allocation Requests

Allocation Requests for the FY23 cycle (October 1, 2022 to September 30, 2023) were open from **May 11, 2022** to **June 8, 2022**.

Allocation Requests for the FY24 cycle (October 1, 2023 to September 30, 2024) are open from **May 10, 2023** to **June 9, 2023**.

Create New **FY23** HPC Allocation Request

Create New **FY24** HPC Allocation Request

Pilot Project and Out-of-Cycle requests may be submitted any time. Pilot allocations are limited to 50,000 allocation units.

Online Submission: Project Information

Request Information

There is strong demand for HPC resources. All requests will be subject to technical review and prioritization based on alignment to mission priorities and technical merit. Allocations are not guaranteed.

Handle ⓘ*

chipcfd

Fiscal Year ⓘ*

FY24

Title ⓘ*

Simulation of on-chip cooling for thermal management of electric vehicles. (Consider using the same title as funded project.)

Word count: 18

Short Abstract ⓘ

This is a 30 word or less description of the project and work to be done. It is included in the one page summary slide we provide to decision makers.

Word count: 30

Abstract ⓘ

This is a longer (up to 200 words) description of the project that will be available to decision makers as supplementary information.

Word count: 22

Primary Project Focus ⓘ

- ✓ Aerospace
- Materials Science
- Integrated Energy Systems
- Fluid Dynamics
- Forecasting
- Manufacturing
- Other

Project Focus Explanation ⓘ

New/Continuing Project ⓘ

New/Existing Workflow ⓘ

Keywords ⓘ

Replace default.
Avoid using the
year and "HPC."

Without this title,
EERE will not be
able to match your
request to a
funded project.

These are EERE
focus areas for
categorizing HPC
use.

Use info buttons
for more
information.

Do you have this
code and
workflow up and
running on
Eagle?

Use similar
keywords to what
you would use for
a journal.

Online Submission: Contact Information

- Need HPC point person and PI.

Contacts			
HPC Lead Email ⓘ	HPC Lead First Name	HPC Lead Last Name	HPC Lead Institution
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
HPC Alternate Lead Email ⓘ	HPC Alternate Lead First Name	HPC Alternate Lead Last Name	HPC Alternate Lead Institution
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
PI Email ⓘ	PI First Name	PI Last Name	PI Institution
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

This is the person we will contact.

May be contacted by EERE for program reasons.

Funding Information and Project Goals

Funding Information and HPC Goals

This section should provide the information needed to link the allocation to a project that is funded, under consideration, or expected to be submitted for funding. The descriptions given in "Program or Project Tasks" and "Program or Project Milestones" should provide a clear case for the use of HPC resources to complete this project.

Funding Area ●	Funding Source ●	Funding Status ●	Funding End Date ●
VTO		Requested	
Project Category ●	Project Type ●	OMS # ●	LDRD Identifier ●
EERE	AOP		
AOP or Project Title ●	AOP WBS/Contract # ●	DOE/Federal Contact ●	
Word count: 0			
EERE Mission Impact ●			
Describe how this project impacts the EERE mission.			
Character count: 52			
Program or Project Tasks ●			
Describe in as much detail as possible how you will use HPC resources for specific project tasks.			
Character count: 97			
Program or Project Milestones ●			
Describe which project milestones are explicitly linked to HPC use.			
Character count: 68			
Most Critical Milestone ●			

Even if a project is not linked to a specific EERE office, if it is in the technology area, it should be listed.
Example: An ARPA-E project in wind energy should have WETO listed.

Tell us where HPC fits into the overall program.

If the funding proposal has not yet been submitted, list "Requested."

NEW: Any projects not supported by EERE MUST include an impact statement to show alignment with the EERE mission.

Online Submission: Computational Resources

Computational Resources Requested

Requestors should take all possible care to provide an accurate estimate of the size and timing of their resource request. The size of the allocation request should be justified based on the number and length of runs given in the *AU Request Explanation* section of the *Computational Readiness* section below. Requests for more than 1 million AUs in particular will need both a strong *AU Request Explanation*, and a strong explanation of the value of the work in the *Funding Information and HPC Goals* sections. Requestors should also provide the best possible estimate of the timing of their use of resources, since, shifting unused resources between quarters may not be possible.

FY ●	Total AUs ●	CPU AUs ●	GPU AUs ●	/projects ●	MSS ●	Use Pattern ●
FY24	250000	100000	150000	5	5	Development in C ▾
FY25	320000	100000	220000	10	10	Distribute equally ▾
FY26	350000	100000	250000	20	20	Distribute equally ▾

Minimum Total AUs ●

70000

Implication of Minimum AUs ●

This should be the bare minimum number of AUs the project needs to accomplish its milestones, etc., with little margin for error.

Character count: 130

Maximum Total AUs ●

200000

Implication of Maximum AUs ●

This should be the amount of AUs that will enable the project to do the best possible science, through larger simulation data sets or higher fidelity results. |

Character count: 150

AU Estimator

Awards are for FY24 but EERE wants long-term plan.

If you don't fit any of these patterns, comment in "Approach."

Use as needed.

Estimating AUs Needed

Estimating CPU AUs:

- CPU nodes have **104 cores**
- CPU nodes are allocated as **whole nodes**
- Estimate the “node hours” you need (hours of compute per node times number of nodes), making sure to efficiently fill the node with work

Example: 200 node hours = 200 X 10 AUs

CPU AUs Needed =
Node Hours X 10 AUs

Estimating GPU AUs:

- GPU nodes have **4 NVIDIA GPUs**
- GPU nodes can be allocated as **partial nodes**: 1, 2 or 4 GPUs per node
- Estimate the “node hours” you need (hours of compute time per node times number of nodes), and the number of GPUs per node that your code will use

Example: 200 node hours of full GPU node = 200 X (4 GPUs X 25 AUs)

Example: 200 node hours of 2-GPU nodes = 200 X (2 GPUs X 25 AUs)

GPU AUs Needed =
Node Hours X #GPUs per node X 25 AUs

Why do we ask for a Usage Pattern?

Computing time is a resource that must be scheduled: an AU is a unit of time – as time passes, fewer AUs are available on the machine overall.

We need to have some idea of how you plan to use your AUs:

- Distribute equally across 4 quarters: 25% each quarter, *for ongoing projects.*
- Development in Q1, production in Q2 or later: 10% in Q1, 30% each in Q2-Q4, *designed for projects that are starting off and need time to develop their code.*
- Start in 2nd Quarter: 33% each in Q2-Q4, *designed for projects with late starts.*
- Use in first half of FY: 45% each in Q1 and Q2, 5% each in Q3 and Q4, *designed for projects with mid-year end dates or early milestones.*
- Use in second half of FY: 5% each in Q1 and Q2, 45% each in Q3 and Q4, *designed for projects with mid-year start dates or late milestones.*

The closer your use matches your pattern, the better priority you will have and the less likely you are to lose unused AUs.

Computational Readiness (Part 1)

Allows NREL
check to
make sure
code is
ready.

Should add
up to 100%

Helps show
EERE how
Kestrel is
used.

Software Requested

Name	Code Source	Estimated Use	Essential?	Platform/Environment	
ANSYS Fluent x	Proprietary/commercial code	50%	Yes	CPU	
MPI	OpenMP	GPU Readiness	GPU Effort	GPU Quantity	GPU Improvement
Yes	Unknown	Partially	Impossible	6	4x
delete					
Name	Code Source	Estimated Use	Essential?	Platform/Environment	
3D visualization x	Internal NREL code	50%	Yes	DAV	
MPI	OpenMP	GPU Readiness	GPU Effort	GPU Quantity	GPU Improvement
Yes	Yes	Yes	Zero/Low	4	N/A
delete					
add					

Demographics

Size of runs ⓘ

5-10 nodes

Duration of runs ⓘ

< 24 hours

Is this project using/supporting ML/AI? (Check all that apply) ⓘ

☒ Project uses Machine Learning/Artificial Intelligence

☐ Project generates data sets to be used for ML/AI

Computational Readiness (Part 2)

Name	Code Source	Estimated Use	Essential?	Platform/Environment	
ANSYS Fluent x	Proprietary/commercial code	50%	Yes	CPU	
MPI	OpenMP	GPU Readiness	GPU Effort	GPU Quantity	GPU Improvement
Yes	Unknown	Partially	Impossible	6	4x

Where and how will your codes run?

Platform/Environment: Does the code run mostly on the standard CPU nodes, does it use GPU nodes, or does it use the highly specialized DAV (Data Analysis and Visualization) nodes?

GPU Readiness: Can the code currently run on GPUs?

GPU Effort: If the user has some control over the code, how hard is it to move it over to GPUs? If it is commercial code, you have no control and should say “impossible.”

GPU Quantity: How many GPUs does it use?

Computational Readiness (Part 3)

Let us know if you will use DAV nodes and your overall computational approach.

Will this project use DAV nodes? ⓘ

Yes

DAV Node Use Explanation ⓘ

If you expect to use the specialized DAV (Data Analysis and Visualization) nodes please explain why and what for.

Computational Approach ⓘ

Give an overall explanation of what you are trying to do:

Example:

In this project, we will perform a large number of simulations of various battery chemistries using the LAMPPS molecular dynamics code to build a training set. We will then use Tensorflow-based machine learning to model the data and attempt to predict additional chemistries that may improve durability. We will re-run promising chemistries using molecular dynamics.

Computational Readiness (Part 4)

Absolutely critical to showing that your request is realistic.

AU Request Explanation

Ideally this should be built from the "bottom up."
The AU estimator can help build the estimate.
250 CPU Cores per run \approx 3 CPU Nodes
CPU Run time(10) * Number of CPU Nodes(3) * CPU AUs Per Node-Hour(10) * Number of Runs(500) = CPU AU Estimate: 150000

Character count: 441

I/O Requirements

This will be important if you expect to move large amounts of data in and out of Kestrel in a short time.

Character count: 106

Storage Request Explanation (only required if over 10 TB)

If you need large amounts of storage, explain the size of your request.

Character count: 72

Other Computational Resources

List other HPC or cloud resources that you have access to. There is no need to list desktop computing

Character count: 101

Are you applying for, or do you have access to other resources, such as other DOE computing facilities or cloud computing?

Do you have large data sets or output files that you need to keep?

Submitting Your Request

Before submission:

- You can save as many times as you want, but submit only once, so please check and recheck.

After submission:

- Your project will undergo an initial review to make sure it is complete and then a *Technical Readiness Review* to ensure it is appropriate for NREL HPC. If you are asked questions, please respond as quickly as possible so NREL can submit to EERE on time.
- VTO users will be asked for follow up to determine if the project can run on VTO's Swift cluster.
- If your project changes after your submit, email hpc-requests@nrel.gov and we will do our best to update request. If appropriate, we will notify EERE that the request has changed.

Allocation Decisions

Decisions on Allocations will be communicated by October 1, 2023:

- Your decision may occur very close to the start of FY24, so be patient!
- Projects that do not receive an allocation can receive a Kestrel pilot allocation of up to 150,000 AUs.
- There is *very little room for adding AUs to projects mid-year*. If you do not believe you can complete project milestones, etc., within your allocation, talk to your project sponsor *immediately*. If the scope of your project changes, you should ensure you have the AUs to accommodate the changes.

Mid-year Changes

New projects after allocations are decided:

- Can receive Kestrel pilot allocations of 150,000 Aus -> put in a request anytime throughout the year
- Larger requests will be considered for AUs on a “space available” basis on the quarter boundaries.

Changes to projects during the year – email hpc-requests@nrel.gov ASAP if:

- If your project does not get funded and you no longer need the AUs
- Your project has changed and you need less or more AUs

We will do our best to accommodate mid-year requests and changes

Closing Comments

- Get an account now if you don't already have one
- Get your request in early, ask questions if you need help
- Provide the best possible information concerning your request
- Take the time to provide the best possible estimate of your computing needs, particularly your CPU and GPU needs

Providing the best possible data, including how use of HPC fits into your milestones, and showing your detailed estimate of use, helps NREL and EERE plan the most effective use of the machine, and makes it more likely your request will be filled.

Please take the time to give the most accurate information possible!

Submitting a Request and More Information

Go here to enter your request and to get a new account set up:

<https://hpcprojects.nrel.gov/>

- Latest version of the call: <https://www.nrel.gov/hpc/resource-allocation-requests.html>
- Learn more about The Office of Energy Efficiency and Renewable Energy here: <https://www.energy.gov/eere/office-energy-efficiency-renewable-energy>
- Additional information about Kestrel is available here: <https://www.nrel.gov/hpc/kestrel-system-configuration.html>.
- Additional information about Allocation Units is available here: <https://www.nrel.gov/hpc/system-resource-allocation-unit.html>.

Thank you

www.nrel.gov

NREL/PR-5F00-86370

NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, operated by the Alliance for Sustainable Energy, LLC.

