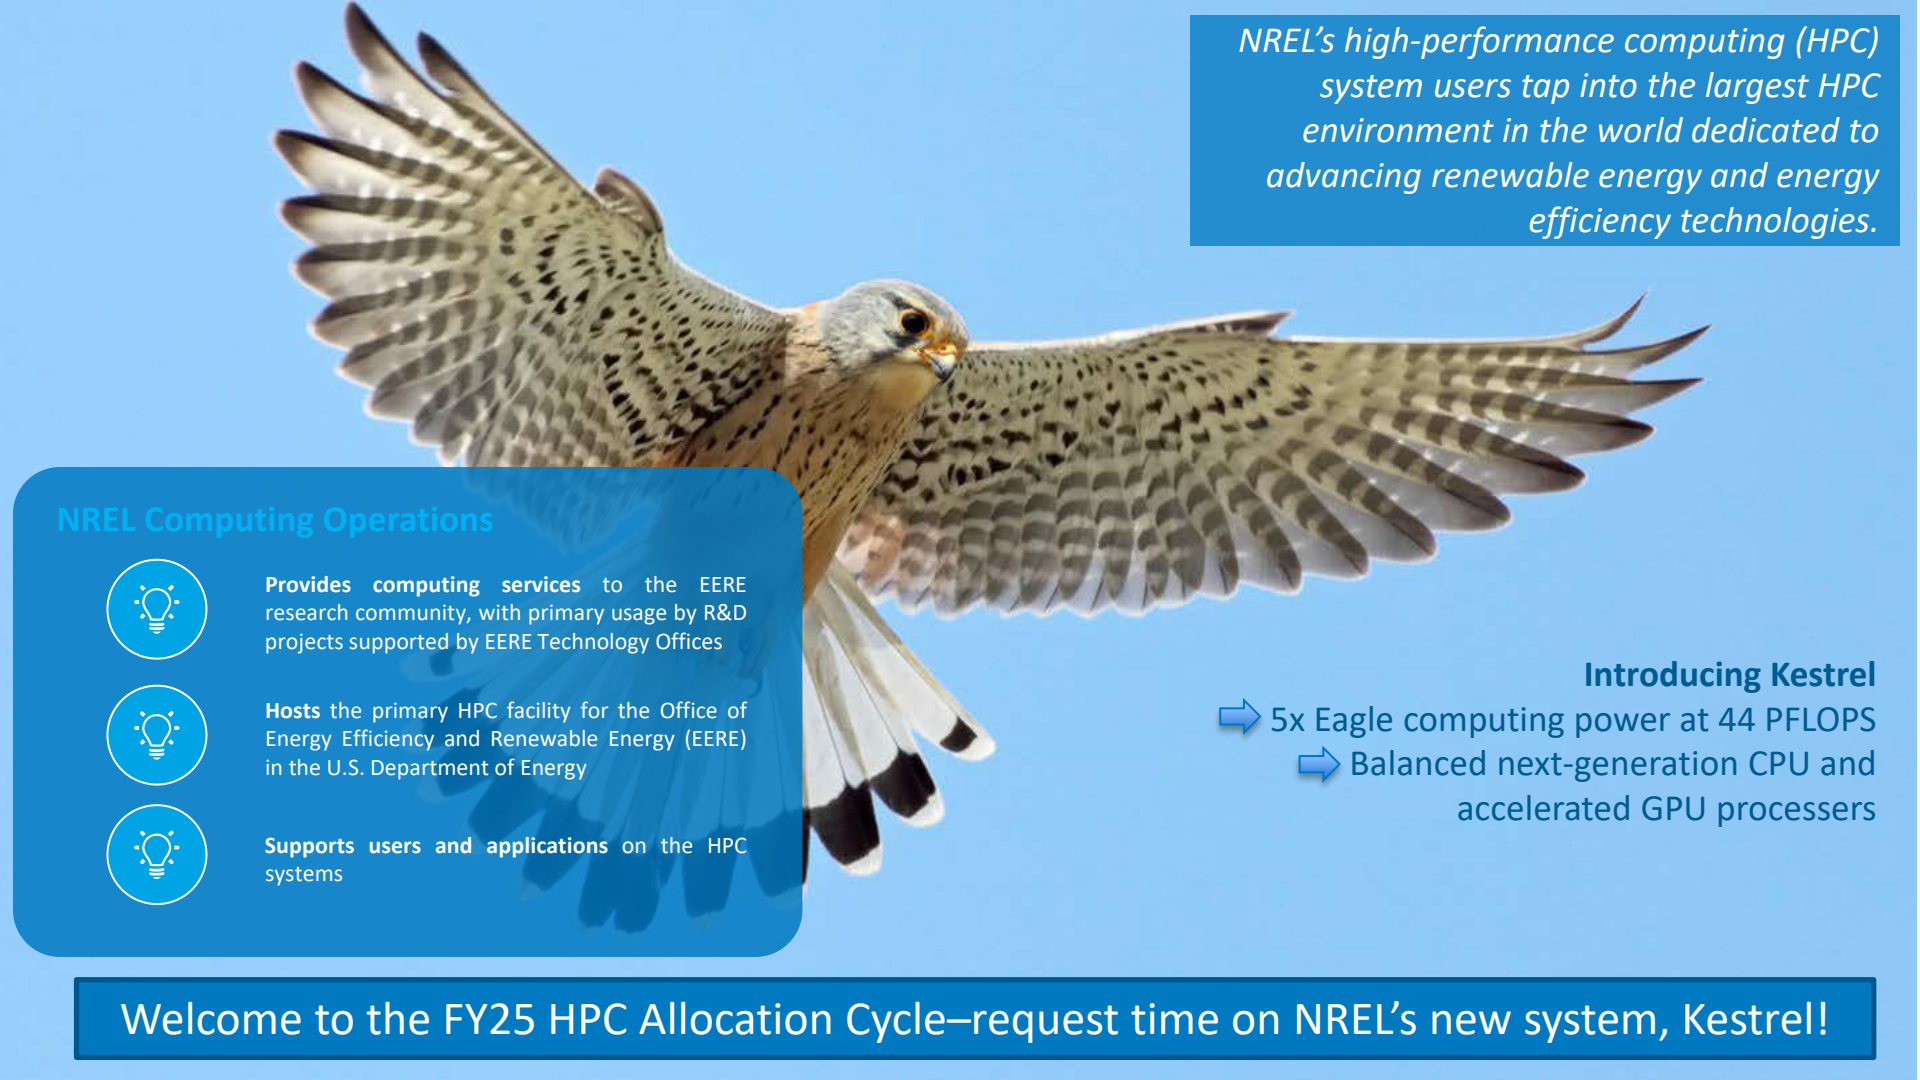




# NREL HPC FY25 Allocation Cycle

05/23/2024 & 5/30/2024

Kris Munch  
David Gaines  
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 FY25 HPC Allocation Cycle—request time on NREL's new system, Kestrel!

# FY25 Summary



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

- **Allocations will be made on Kestrel for FY25**
- Kestrel is designed with both CPU nodes and GPU nodes
- Requests for FY25 need to pay particular attention to Kestrel's architecture:
  - Requests must differentiate between the CPU request and GPU request
  - Give yourself computing time for transitioning codes and workflows to GPUs, and for new projects ramping up on Kestrel
- **Get your requests in early.**
  - **If you don't already have a login, request an account now to avoid delays!**
  - **You should still submit a proposal if you are unsure of funding status!**

# FY25 Request Reminders

## We will have a full year of GPUs available in FY25:

- 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 **distinct AU factors** – calculate and plan accordingly!
  - 1 CPU node hour = 10 AUs
  - 1 GPU node hour = 100 AUs
- With Kestrel's 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!

# Introducing Kestrel in FY25

*Dedicated to renewable energy and energy efficiency research*

2,322 Intel SPR CPU  
Nodes

132 NVIDIA H100  
GPU Nodes

1 TB total local  
node storage

200 Gb/s high speed  
network

**44 Petaflops**

**#68 on Top500\***



**650 TB memory**

**95 PB storage**

**2,454 nodes**

**258,385 processors**

514 Users

371 Projects

>180 High performance  
computing codes

Warm water cooling and  
waste heat recovery

*\*Nov 2023 Top500 result included Kestrel CPU nodes only*

# 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 FY25?

Resources to be allocated on Kestrel include **300** million AUs across both CPU nodes and GPU nodes:

- 191 M AUs on CPU nodes
- 109 M AUs on GPU nodes
- 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.

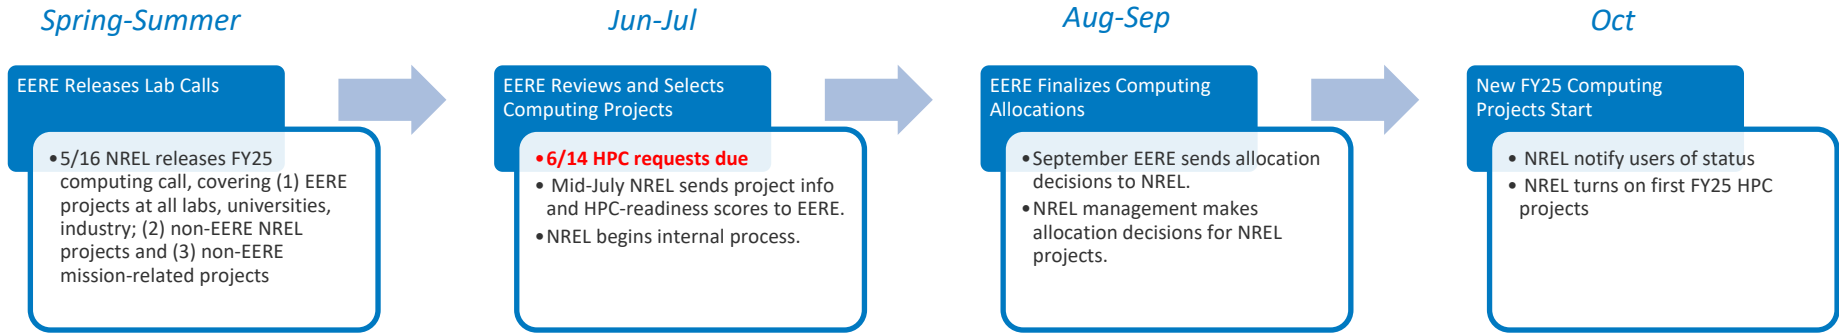
Projects should apply for allocations if: (1) funded, (2) a proposal is under consideration, or (3) if a proposal will be submitted during summer 2024.

*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.*

# FY25 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!

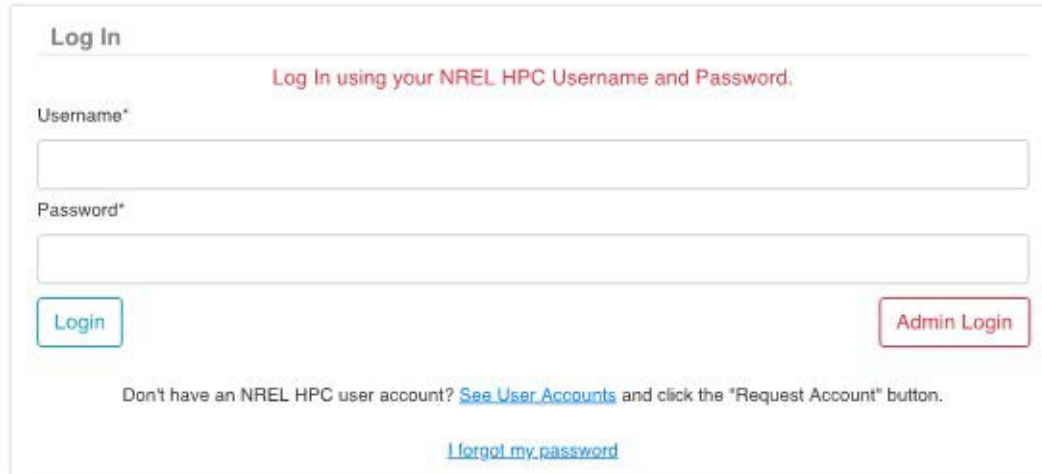
# Lex Allocations: Key Reminders

- Projects should apply for allocations if: (1) funded, (2) a proposal is under consideration, or (3) if a proposal will be submitted during summer 2024.
- If you are updating a FY24 allocation *where you are the lead* there is an option to copy the information over as a starting point. YOU can copy existing allocation from this year to next year.
- We have implemented a new-pilot form which will streamline the allocation process for pilots (projects needing less than 150k AUs)
- We have consolidated and simplified the allocation request forms, by integrating the software request and AU calculator

# Online Submission: Login

Go to [hpcprojects.nrel.gov](https://hpcprojects.nrel.gov) to prepare your submission.

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



The screenshot shows a login form titled "Log In". At the top, it says "Log In using your NREL HPC Username and Password." Below this are two input fields: "Username\*" and "Password\*". There are two buttons: a blue "Login" button and a red "Admin Login" button. At the bottom, there is a link for "I forgot my password" and a note: "Don't have an NREL HPC user account? [See User Accounts](#) and click the 'Request Account' button."

# Online Submission: Landing Page

Once logged in, you can submit a FY25 Allocation Request.

- Use the **FY24** button if you need a pilot allocation to test code right now, otherwise use **FY25**.
- If you are updating a **FY24** allocation *where you are the lead* there is an option to copy the information over as a starting point.

The screenshot displays two columns of information. The left column is for FY24 (current year), with a teal button labeled 'Create New FY24 HPC Allocation Request'. The right column is for FY25 (next year), with a teal button labeled 'Create New FY25 HPC Allocation Request'. A red oval highlights the right column. At the bottom, a note states: 'Pilot Project and Out-of-Cycle requests may be submitted any time. Pilot allocations are limited to 150,000 allocation units.'

FY24 (current year)	FY25 (next year)
Allocation Requests for the FY24 cycle (October 1, 2023 to September 30, 2024) were open from	Allocation Requests for the FY25 cycle (October 1, 2024 to September 30, 2025) are open from
<b>May 10, 2023 to June 8, 2023</b>	<b>May 16, 2024 to June 14, 2024</b>
<a href="#">Create New FY24 HPC Allocation Request</a>	<a href="#">Create New FY25 HPC Allocation Request</a>

*Pilot Project and Out-of-Cycle requests may be submitted any time. Pilot allocations are limited to 150,000 allocation units.*

# Online Submission: Request Type

If pilot (less than 150k AU) select Pilot Request

## Request Type

Select if your request is a pilot. Pilot requests may be requested at any time in the year and enable smaller projects to gain experience with our HPC systems. If approved, pilots are awarded a fixed rate of 150,000 AUs and 10 TB of storage. Fewer fields are required for pilot requests, so it is recommended to select the request type before completing the form. For guidance on estimating your project's AU requirements, please see [System Resource Allocation Units](#).

Pilot Request

## Pilot Request Option

- smaller projects can request a pilot which is less than or equal to 150k AUs
- pilot sized projects can be granted at anytime throughout the year
- new pilot form simplifies and streamlines allocation request form
- resources default to 150k AU and 10TB storage for requested resources

If your project is a pilot sized (less than 150k AU), use the 'Pilot Request' option to streamline the process

# Online Submission: Project Information

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.

Awards are for FY25,  
but do you expect  
this project to  
continue in future  
fiscal years?

**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  Fiscal Year

Title  Word count: 5

Short Abstract  Word count: 0

Abstract  Word count: 0

Primary Project Focus   Project Focus Explanation

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

New/Continuing Project  New/Existing Workflow

Expected Duration of Project

Keywords

Use info buttons for  
more information.

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




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



# Online Submission: Contact Information

HPC Lead, HPC Alternate, and PI are required. HPC Lead and Alternate will require an NREL HPC account.

These are the people we will contact about HPC issues.

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

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 ●	
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."

All projects must include an impact statement to show alignment with the EERE mission.

# Online Submission: Computational Resources Pt 1

## Computational Resources & Software 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 is estimated based on the number and length of runs given in the *Software Resources* section below. Requests for more than 1 million AUs, in particular, will need a strong *Computational Approach* and explanation of the value of the work in the *Funding Information and HPC Goals* sections. Awarded AUs will not be restricted to the software or proportions in the estimate as we understand these can vary throughout the duration of the project. Requestors should, however, provide the best possible estimate of the timing of their use of resources throughout the year (the Use Pattern), since saving unused resources for future quarters may not be possible.

**Kestrel CPU Nodes:** Non-integer numbers represent fractional usage of node (e.g., using 52 cores on a single node run would be represented as 0.5)

2048 Dual Socket Intel Xeon Sapphire Rapids (104 cores) and 256 GB RAM,  
256 Dual Socket Intel Xeon Sapphire Rapids (104 cores), 256 GB RAM, and 1.7TB local disk  
10 Dual Socket Intel Xeon Sapphire Rapids (104 cores), 2TB RAM, and 5.8TB local disk

**Kestrel GPU Nodes:** Non-integer numbers represent fractional usage of node (e.g., using 6 GPUs would be represented as 1.5)

132 Dual Socket AMD EPYC 9554 (128 cores), 360GB RAM, and 4x NVIDIA H100 80GB HBM3 per node

Name	Code Source	MPI	OpenMP
VASP x	Proprietary/commercial code	Yes	Yes

CPU Nodes	Run Time (hours)	# of runs	CPU Node Hours	CPU AUs
0.5	4	1000	2000.0	20000

GPU Nodes	Run Time (hours)	# of runs	GPU Node Hours	GPU AUs
.25	1	1000	250.00	25000

GPU Readiness	GPU Effort	GPU Improvement
Yes	Zero/Low	4x

delete

add

Reminders and instructions for allocation process, as well as brief details regarding Kestrel node architecture

**NEW:** Software request used to calculate node hours and AUs request

### Key Points:

- non-integers allowed for nodes and run time to represent fractional node usage (node hours and AUs auto-calculated)
- can assign both GPU and CPU hours to software (or only one)
- name and additional information required (additional details later slides)
- total for all software is automatically used to populate resource request

# Online Submission: Computational Resources Pt 2

Name	Code Source	MPI	OpenMP
VASP	Proprietary/commercial code	Yes	Yes

CPU Nodes	Run Time (hours)	# of runs	CPU Node Hours	CPU AUs
0.5	4	1000	2000.0	20000

GPU Nodes	Run Time (hours)	# of runs	GPU Node Hours	GPU AUs
.25	1	1000	250.00	25000

GPU Readiness	GPU Effort	GPU Improvement
Yes	Zero/Low	4x

FY	Total AUs	CPU AUs	GPU AUs	/projects	MSS	Use Pattern
FY25	45000	20000	25000			Distribute equally a

Minimum Total AUs	Implication of Minimum AUs
40000	This should be the bare minimum number of AUs the project needs to accomplish its milestones, etc., with little margin for error.

Maximum Total AUs	Implication of Maximum AUs
100000	This should be the amount of AUs that will enable the project to do the best possible science, through large simulation datasets or higher fidelity results.

Total AUs, CPU AUs, and GPU AUs auto-calculated from ALL software entries

Min/Max total and implication is required and is used in the event we are unable to grant full request (or able to grant additional AUs)

Use pattern describes how you will work throughout the year (details next slide); if you don't fit available options, comment in 'Comp. Approach'.

'/projects' is disk space requested; justification is required if greater than 10TB

# 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 2<sup>nd</sup> 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.

# Online Submission: Computational Resources Pt 3

## Where and how will your codes run?

Code Source: Is this Internal/Proprietary/Open-Source code? NREL, DOE, or Commercial?

MPI/OpenMP: Does this code utilize MPI or OpenMP for parallelization?

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 Improvement: If this code can run on GPUs, what is the expected performance improvement?

Name	Code Source	MPI	OpenMP	
VASP x	Proprietary/commercial code	Yes	Yes	
CPU Nodes	Run Time (hours)	# of runs	CPU Node Hours	CPU AUs
0.5	4	1000	2000.0	20000
GPU Nodes	Run Time (hours)	# of runs	GPU Node Hours	GPU AUs
.25	1	1000	250.00	25000
GPU Readiness	GPU Effort	GPU Improvement		
Yes	Zero/Low	4x		

# Computational Readiness (Part 1)

Let us know if your project supports ML/AI and how?

Let us know if you are using DAV nodes and provide explanation if so.

Describe what you are doing with Kestrel, and how you will use the computational resources.

## Computational Readiness

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
- Project involves generative AI or large language models (LLMs)
- Project develops new ML/AI methods or algorithms

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.

Character count: 113

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.

Character count: 437  
NREL | 23

# Computational Readiness (Part 2)

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

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

## 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 more than 10TB, explain the size of your request.

Character count: 61

## Other Computational Resources ⓘ

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

Character count: 103



# Submitting Your Request

## Before submission:

- You can save as many times as you want, but submit only once, so please check and recheck.
- If your project is a pilot, please check the pilot request form.

## 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 you submit, email [hpc-requests@nrel.gov](mailto: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, 2024:

- Your decision may occur very close to the start of FY25, 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](mailto: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, 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>
- HPC Documentation: <https://nrel.github.io/HPC/>

# Thank you

---

[www.nrel.gov](http://www.nrel.gov)

NREL/PR-5F00-90033

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

**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.**

