



NREL HPC

FY23 Allocation Webinar

Michael Martin

Webinar

May 17, 2022

FY23 Changes

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

We will bring on a new computer, Kestrel, with 5 x the computing power of Eagle.

Kestrel is designed with a much larger GPU capacity than Eagle: approximately 50% of the computing power comes from GPU nodes.

Due to supply change issues, delivery has shifted...CPU nodes are expected to be available to users mid-FY23, while GPU nodes will be available later in the year (or potentially FY24).



All of this leads to large changes in the allocations process for FY23!

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	2,304 nodes, 2 processors/node
Data Storage	14 PetaBytes Lustre	75 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

FY23 Request Changes

What really changes?

Because **GPUs are 50% of the computational power**, we are going to ask (1) is your code GPU ready, and (2) what would it take to make the code GPU-ready? (Note that not all code is expected to make the shift to GPUs.)

We are going to ask for an explicit transition plan, since Eagle is currently scheduled to shut down at the end of FY23. If a project will not continue past then, it does not need to plan a transition.

As many as 100M AUs may be available *on Kestrel for the 2nd half of the year*. This compares to the 65M AUs available for the entire year on Eagle. NREL and EERE are going to look closely at the “Maximum AUs” request to see which projects may be in a position to use Kestrel.

FY23 Allocation Process Intro

What's Available?

The allocation process covers approximately 65 million “Allocation Units,” or AUs on Eagle, and potentially as many as 100 million AUs on Kestrel.

- Each AU is 1/3 of a node hour on Eagle. Note that even if a project uses just one core on a node, it will be charged for the entire node.
- ~ 85% of this is for EERE-funded projects, ~ 15% is for NREL non-EERE projects: LDRD, SPP, non-EERE DOE (ARPA-E, Office of Science), DoD, etc.
- A second cluster for VTO projects, “Swift,” will be available. *VTO projects requesting time on Eagle will also be considered for time on Swift.*
- A cluster for internal NREL projects, “Vermillion,” will be available. *LDRD projects requesting time on Eagle will also be considered for time on Vermillion.*

FY23 Allocation Process Intro

What projects are eligible?

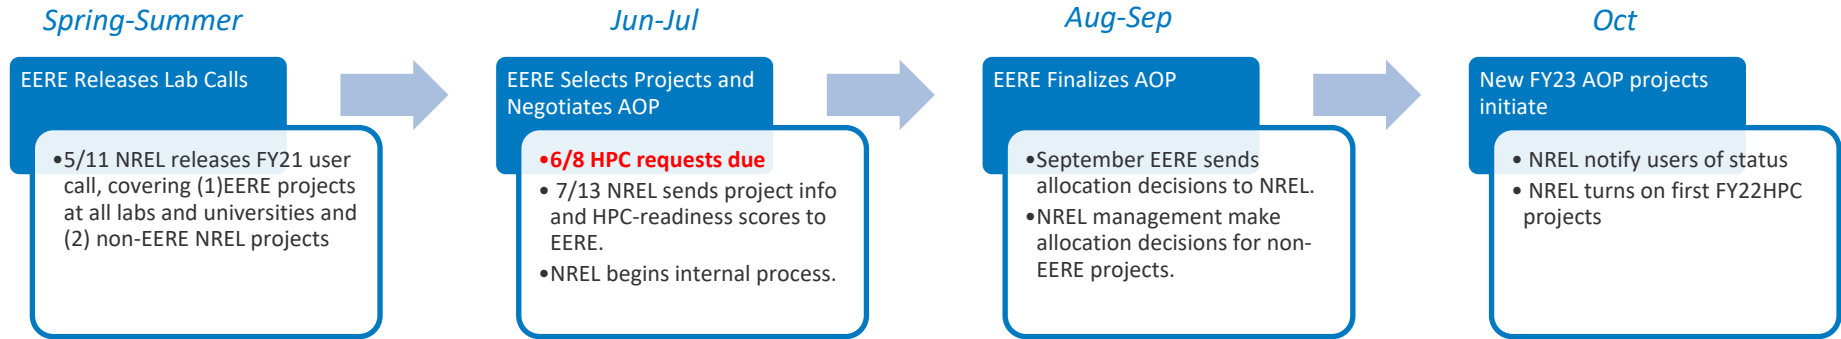
- All EERE-funded projects at all national laboratories and universities. (HPC4EI projects should consult with program management, because some are Office of Fossil Energy-funded.)
- All NREL projects, regardless of funding source.
- Projects should apply for allocations if: (1) funded, (2) a proposal is under consideration, or (3) if a proposal will be submitted during summer 2022.
- One allocation per funded project. *Do not split projects up into multiple allocations, even if the project is scattered across several labs. Avoid using large “umbrella allocations” to cover multiple projects.*

Due to the large demand for resources, the allocation process is not able to support projects that do not have a link to either EERE or NREL, even if they support the clean energy mission.

FY23 Allocation Process Intro

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 non-EERE NREL projects.
- Allocation process is designed to be parallel to funding process.



Updates and off-cycle requests are possible but in-cycle requests are easier to accommodate: request compute time now if you think you will need it!

Technical Readiness Review

Technical Readiness Review: Part of how NREL makes sure EERE has the best possible information for decision-making.

- Is this an HPC task? “Ideal” uses large-scale parallelization capabilities, “Suitable” can run effectively, and “Unsuited” is just not ready. “Suitable” is an acceptable outcome to pass TRR.
- Is the project ready (if appropriate) to use GPUs?
- Is the size of the request justified?
- Is the storage request appropriate? Eagle is not for long-term storage of large data sets...
- Which software? Is the software available?
- NREL staff will follow up with submitter if there are issues.

Process

Request

HPC User Account

- Visit <https://www.nrel.gov/hpc/user-accounts.html> to request account if you don't already have one.

Draft Allocation Request Online

- Starting May 11, 2022.

Submit Allocation Request

- Before midnight **June 8, 2022**.

Review

Triage

- HPC Operations reviews requests for completeness and assigns Technical Readiness reviewer.

Technical Readiness

- CSC Center staff review requests for mission alignment (broadly), method and software compatibility with, and readiness to run on NREL HPC.

Packaging

- HPC Operations prepares a package, one for each EERE Office or DOE crosscutting activity. Package includes overview (hours requested), scorecard (for assignment of hours by project), and each request.
- Packages are transmitted no **later than July 14, 2022**.

Allocation

Apportion

- Allocation authorities meet to determine high-level distribution of hours to each EERE Office or DOE crosscutting activity.

Allocation

- Hours are assigned to each request within each office or DOE crosscutting activity.
- Allocation decisions are planned to be complete by **early September, 2022**.

Communication

- HPC Operations prepares memos in response to each allocation request.
- Memos are planned to be sent by **late September, 2022**.

Implementation

User Accounts

- HPC Operations will contact each PI to associate users with allocations and set up user accounts if necessary.

Full Allocations

- Full-year allocations will be set to start **October 1, 2022**.

Provisional Allocations

- Provisional allocations will be set to start when authorized.

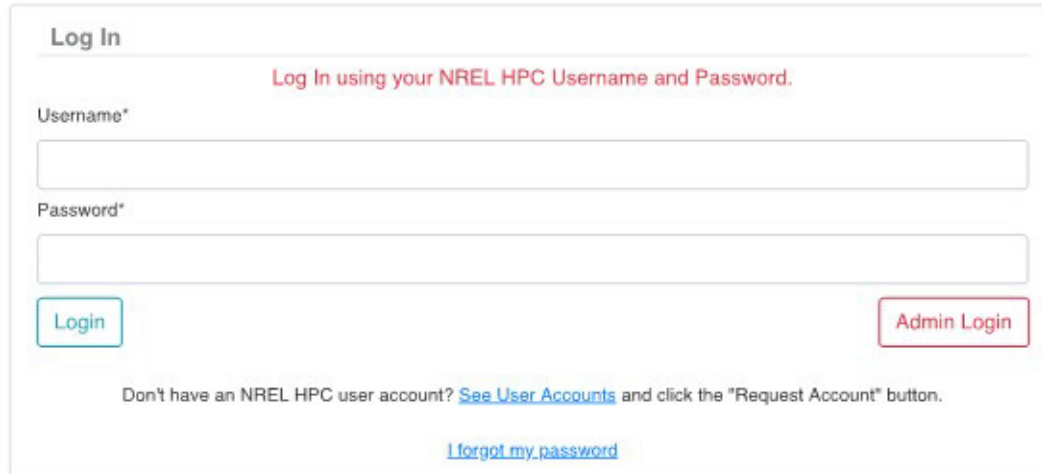
Online Submission: Login

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

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

6/10/2020

Lex



The screenshot shows a web 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 FY23 allocation request.

- Use the FY22 button if you need a pilot allocation to test code, otherwise use FY23.
- 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 FY22 cycle (October 1, 2021 to September 30, 2022) were open from **May 7, 2021** to **June 7, 2021**.

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

Create New FY22 HPC Allocation Request

Create New FY23 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 

FY23

Title 

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

Word count: 17

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: 20

Primary Project Focus 

- ✓ -----
- 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

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.

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

Tell us where HPC fits into the overall program.

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 Project Category Project Type

Funding Status Funding End Date AOP WBS or Contract #

AOP or Project Title DOE/Federal Contact

Word count: 0

Program or Project Tasks

Describe in as much detail as possible how you will use Eagle for specific project tasks.

Character count: 89

Program or Project Milestones

Describe which project milestones are explicitly linked to Eagle use.

Character count: 69

Most Critical Milestone

Online Submission: Computational Resources Requested

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	AUs	/projects	MSS	Use Pattern
FY23	100000	5	5	Development in Q1, production in Q2 or lat
FY24	200000	10	10	Distribute equally across 4 quarters
FY25	200000	20	20	Distribute equally across 4 quarters

Minimum AUs

50000

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: 129

Maximum 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. For FY23, it is important to think big because of the large increase in available AUs.

Character count: 246

Awards are for FY23 but EERE wants long-term plan. AUs are consistent across machines and from year to year.

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

Use as needed.

We will look at this when allocating Kestrel.

AU Estimator

Why Use Patterns?

Eagle is a resource that must be scheduled: If nothing runs and the computer is unused, the AUs aren't kept in a bank. They "expire" & can't be made up (like an empty seat on an airplane or table at a restaurant.)

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 Aus.

Computational Readiness (Part 1)

Allows NREL check to make sure code is ready.

Should add up to 100%.

Helps show EERE how Eagle is used.

Software Requested

Name	Code Source	Estimated Use	Essential?	Platform/Environment	
ANSYS Fluent	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	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

Need to help understand Eagle vs. Kestrel

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

What are all of these?

Platform/Environment: Does the code run mostly on the conventional CPU nodes, does it use GPU nodes, or does it use the highly specialized DAV (Data Analysis and Visualization) nodes? If you don't know, the answer is almost certainly CPU nodes (the default.)

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 (Pt 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 (Pt 4)

Tell us your plan for moving from one machine to another (and "It's all ok" is acceptable!)

-
- Move to Kestrel as soon as possible.
- Move to Kestrel gradually.
- ✓ Delay moving to Kestrel (explanation required).
- Project ending; prefer not to transition.
- Other (explanation required).

Character count: 438

Kestrel Transition Plan ⓘ

Delay moving to Kestrel (explanation required).

Kestrel Transition Plan Explanation ⓘ

Explain to us why you want to do what you wish to do when it comes to moving from Eagle to Kestrel.

Percentage of Allocation Required to Run on Eagle ⓘ

20%

Computational Readiness (Pt 5)

Absolutely critical to showing that your request is realistic.

Only required if you are doing something unusual.

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? EERE prefers allocations over 1 million AUs also apply for other computing.

AU Request Explanation

Ideally this should be built from the "bottom up."

We will need to run 500 cases of our simulation. Each case takes 10 hours on 8 nodes.

$500 \text{ cases} \times 10 \text{ hours} \times 8 \text{ nodes} \times 3 \text{ AUs/node hour} = 120,000 \text{ AUs}$.

Character count: 206

I/O Requirements

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

Character count: 103

Storage Request Explanation (only required if over 10 TB)

If you need large amounts of storage, especially for machine learning, explain the size of your request.

Character count: 104

Other Computational Resources

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

Submitting Your Request

Before submission:

- You can save as many times as you want, but submit only once, so please check and recheck.
- Keep in mind that the number of AUs available will increase mid-year when Kestrel comes on-line & set your “high end” request accordingly.

After submission:

- Your project will undergo an initial review to make sure it is complete and 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 Swift cluster.
- If project changes, 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.

After the Allocation Decision

Decisions:

- Your decision may occur very close to the start of FY22, so be patient!
- Projects that do not receive an allocation will receive a pilot allocation of up to 50,000 AUs.
- There is *very little room for adding AUs to projects*. 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.

Dealing With Resource Restrictions

Because of limited resources in Q1 and Q2, justifying the size of the allocation request remains:

- Realistically justify number of AUs required per case and number of cases. You can request a pilot allocation request if you need to figure out how many AUs will use on Eagle.
- Large allocations (greater than 1 million AUs) should show that they are seeking additional resources.
- If you were low (or high) in your AU usage in FY21 allocations, be responsive when NREL operations asks for updates at the end of Q3.
- Pay attention to *when* you will use allocation and pick appropriate schedule. (“Distribute evenly across 4 quarters,” “Development in Q1, production in Q2 or later,” etc.)
- Balance this with thinking big once Kestrel is available...

After the Allocation Decision

What if I find out after allocations are done that I need HPC?

- Eagle will be completely booked, but there are opportunities to get time on the machine.
- Projects requesting time after the allocation process has concluded will be given pilot allocations of 50,000 AUs.
- Larger requests will be considered for additional AUs on a “space available” basis on the quarter boundaries.

Closing Comment

Eagle is overbooked and will almost certainly remain so for Q1 and Q2- but Kestrel is coming!

Providing the best possible data, including how use of HPC fits into milestones, and showing your 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.

This information also *feeds the planning process for buying Kestrel's successor*, which hopefully will reduce overbooking.

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

Thank you

www.nrel.gov

NREL/PR-2C00-82917

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.

