

What are MorphoCloud On-Demand Instances?

MC Instances are virtual workstations provided to researchers who temporarily need powerful computers to accomplish 3D digital morphology tasks. For example, you have a large segmentation task that requires more memory than your personal computer has or have a computationally intensive task that benefits having dozens of compute cores, or you want to test a newly published deep-learning model for your specific data, but don't have a big enough GPU, you will benefit from access to MC Instances. These workstations come with open-source 3D Slicer image analysis platform and its SlicerMorph extension and many other extensions preloaded. You simply click the provided URL, login with your credentials and start immediately working on your virtual computer.

When you are done, you can stop (shelve) your instance until the next time. When you are ready to use it again, you can restart (unshelve) your instance, and within a couple minutes you are ready to continue. These virtual workstations are provided through the NSF funded JetStream2 (JS2) Cloud farm (<https://docs.jetstream-cloud.org/>) located at Indiana University.

How much computing power? For how long?

Standard GPU instance has 32 cores (EPYC 7703), 125GB RAM, and a NVIDIA A100 GPU with 40GB of RAM. You are also provided with 100GB storage volume for your data (called MyData volume).

For how long depends on how you quickly consume your usage credits. You have 200h (wall time) of usage credits, and 60 days to use your credits. At any time during those 60 days, you can renew your instance for an additional 60 days (for a total of 120 days starting from the issue creation date). This renewal is granted once. Unless you request a renewal, at the end of 60 days your instance will be deleted, regardless of whether you have unused credits or left or not. You will get email notification prior to the 60-day (or 120-day, if you renewed it) expiration deadline. You can of course reapply and request more credits with a new valid use case.

How can I get started?

We use GitHub's workflow automation (issues page) to create and manage MC instances. To be able to request a MC instance through the issues page you will need:

(1) a GitHub account (it is free);

(2) an ORCID (it is free);

(3) link your ORCID with your GitHub account

(<https://github.blog/changelog/2024-03-13-authenticate-orcid-id/>).

Once you have those, go to

<https://github.com/MorphoCloud/MorphoCloudInstances/issues/new/choose> , and click on the "Get Started" button, which will take you to the "Instance Request Template".

What information do I need to provide to request an instance?

You need to briefly describe who you are, your project and your needs, provide your ORCID and your contact email (this is the email address where MC will send you instance credentials. It can be different from the email you registered on GH with), and choose the instance type (CPU or GPU). Please follow the template provided to create the issue title. We will review your request and will try to approve your request within 24 h after your issue submission. If we need more information, or there is some ambiguity on usage needs, we might get in touch with you via email before the approval. It is important that your ORCID profile gives a sufficient picture about your research/teaching activities related to 3D biological structure. Make sure you enter your ORCID in the form of WWW-XXXX-YYYY-ZZZZ (not the URL).

How do I know I am approved?

When you are approved, MorphoCloud admins will update the issue page by issuing a `/create` command. This will invoke the automated workflow to create your instance, which will continuously update the issue page on its progress and when your instance is up and running, you will receive an email from the MorphoCloudPortal with the access URL and your passphrase, which could take an hour after approval. You can bookmark your issue page, and check for updates. On the right hand side of the issue page, the current status of your issue will be displayed via labels:

● Instance: Extend session test flavor:g3.xl request:approved status:active
#60 opened 3 days ago by muratmaga

● Instance: CT Slicer workshop flavor:g3.xl request:approved status:shelved_offloaded

What commands can I use on the issue page?

To stop your instance, you can issue the `/shelve` command. Shelving will preserve the state of the instance (e.g., any custom software you installed will be there), but not the running computations. It will also stop the credit usage counter. To restart a “shelved” instance, use the `/unshelve` command. Unshelving is typically faster than creating an instance from scratch, your instance will be up and running -in most cases- less than 5 minutes, and you will receive another email with new login credentials. Each time you unshelve your instance a new email with the access credentials and URL will be sent. If you somehow lose this information, you can issue the `/email` command to have the latest login credentials resend to you. `/renew` command will automatically extend your instance for another 60 days for a maximum of 120 days. Always check the Github issue page for your instance to confirm the command executions and see the updates. MC admins have other commands that they can use to recreate your instance or data storage.

How much does it cost to use it?

MC instances do not cost you anything, provided that you have a relevant use case. All of this is made possible by funding from the National Science Foundation, as well as the National Institutes of Health to us and other parties supporting this infrastructure. As such, we ask you to acknowledge grants from our funders (NSF DBI/2301405, NICHD HD104435) that made MorphoCloud project possible, as well as grants for the JS2 and its Exosphere interface, both of which we rely on to make it happen (NSF OAC/2005506, NSF TI/2229642, JS2 and Exosphere respectively) in your presentations, posters, preprints, publications and such like. Here is a generic acknowledgement text you can copy and paste:

“This work relied on cyberinfrastructure which is supported by grants from National Science Foundation (MorphoCloud: DBI/2301405; JetStream2: OAC/2005506; Exosphere: TI/2229642) and National Institutes of Health (MorphoCloud: NICHD HD104435).”

Of course any tool (e.g., 3D Slicer, SlicerMorph, R, etc..) you are using to accomplish your research goals on MC, you should also acknowledge and cite.

Are there any usage restrictions?

Only a few. No project/data that is considered “Human Subject Research” (HSR) is allowed on MC. For HSR classification, we follow National Institutes of Health guidelines. If you are uncertain whether your project can be categorized as HSR, please go to <https://grants.nih.gov/policy/humansubjects/hs-decision.htm> and take the survey. Because MC is using the JS2cloud farm, you should also review their

acceptable usage policies:

<https://docs.jetstream-cloud.org/general/policies/#acceptable-use-of-jetstream2>.

In a nutshell, JS2 does not allow commercial research, or provide any mechanism to secure EPHI/HIPAA protected data.

If you have a use case that is not predominantly focused on interactive processing of 3D digital morphological data (e.g., pure bioinformatics), we might reject your request, since there are probably much better services through the ACCESS program that will be more relevant to you.

I deleted/overwrote my research data, can you recover that for me?

Sorry, there is no backup of data in MC instances. It is your responsibility to back them up in a timely manner. Remember, with a few clicks you can easily bring your data to your local computer (or upload from your local computer to MC instance). See <https://docs.jetstream-cloud.org/ui/exo/exo-filetransfer/> for more instructions on file transfer between your local computer and remote instance. If you want to transfer data in bulk, you can also use an SFTP client (use the SSH credentials provided in the email).

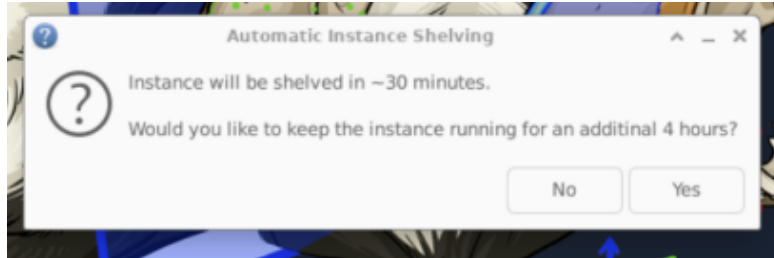
Also, you should always keep your data in your designated storage space, which is located at `/media/volume/MyData`, and which shows up as a shortcut icon on your desktop. That's because if your instance state becomes unusable/inaccessible for any reason, we might have to delete and recreate it. In that case everything that is NOT in your storage volume will be lost. You have 100GB of storage in MyData folder. By default anything you upload to the instance via drag and drop or SFTP will be saved under your MyData folder.

Also note that once your instance has been deleted after the 60 days has been passed, there is no way to recover the data, or the instance status. Please pay attention to the remainder about nearing deadlines and take action.

How long is each session?

Everytime you unshelve an instance, your session is limited to 4h by default. Once the 4h limit is reached, your instance will be automatically shelved, unless you actively intervene and respond to the warning that there is an impending shelving in the next 30 minutes (see the screenshot below). You can choose to continue your session for another 4h, and keep doing that as long as you want to use the system interactively.

Alternatively you can renew your session at any time (without waiting for the popup) by clicking the Extend Instance Session icon on your desktop.



We implemented auto-shelving because: (1) people forget to “shelve” their instance after they are done and burn through their credits without using the system very much; (2) we can only keep a limited number of instances running concurrently. If your instance is stopped (shelved), you can go to your specific issue page and type /unshelve your instance to restart it, but the tasks that will be running at the time of shelving will be

I probably need way more than 200h of usage credits.

We are glad to hear that, but this infrastructure is meant for transient usage of resources (e.g., students who may need to collect measurements from some 3D datasets for a class project. A short course, or a workshop. Make a high-resolution rendering of a 3D datasets that doesn't fit to your own computer...).

If you run out of credits, or you need some more time to complete your project (or starting a new project), by all means submit a new request. However, if you know you will be using these routinely, we encourage you to get your dedicated allocation from the ACCESS program, <https://access-ci.org/about/get-started/>.

ExploreACCESS program makes it very easy to obtain credits, often the approval comes within 24h of submission.