
Table of Contents

CADES Support

| | |
|---|-----|
| Introduction | 1.1 |
| Table of Contents | 1.2 |
| Support | 1.3 |
| Access CADES Externally | 1.4 |
| CADES Team | 1.5 |
| CADES Acknowledgement | 1.6 |
| Glossary | 1.7 |

Quick-Start Guides

| | |
|---|-----|
| Launch a VM Instance: Quick-Start | 2.1 |
| SHPC Condos: Quick-Start | 2.2 |

CADES Cloud User Guide

| | |
|---|-------|
| Getting Started with OpenStack | 3.1 |
| Request Your Cloud Allocation | 3.1.1 |
| Manage Your Cloud Allocation | 3.1.2 |
| OpenStack Project Quota | 3.1.3 |
| OpenStack Help | 3.1.4 |
| Launch a VM Instance | 3.2 |
| Log In & Name the VM | 3.2.1 |
| Configure the VM | 3.2.2 |
| Networks & Security | 3.2.3 |
| Key Pair Use | 3.2.4 |
| Access VM Instances | 3.3 |
| Access Your VM Instance Using SSH | 3.3.1 |
| Access Your VM Instance Using PuTTY (Windows) | 3.3.2 |
| Access Your VM Instance Using Horizon | 3.3.3 |
| Add More Users to Your VM Instance | 3.3.4 |
| Manage Your VM Instances | 3.4 |
| Delete a VM Instance from Your Project | 3.4.1 |
| Delete a Volume from Your Project | 3.4.2 |
| Resize a VM Instance | 3.4.3 |
| Add a Volume to a VM Instance | 3.4.4 |

| | |
|--|---------|
| Create a Snapshot | 3.4.5 |
| OpenStack Security Groups | 3.4.6 |
| Modify the Default Security Group | 3.4.6.1 |
| Create a new Security Group | 3.4.6.2 |
| Security Group CIDR Examples | 3.4.6.3 |
| Overview - CADES Cloud Information | 3.5 |
| Available VM Images & Configurations | 3.5.1 |
| Software & Hardware | 3.5.2 |
| Network & Storage | 3.5.3 |
| Additional OpenStack Resources | 3.6 |
| Request Firewall Exception | 3.6.1 |
| Run a Simple Web Server | 3.6.2 |
| SSL - Let's Encrypt | 3.6.3 |
| Install CPUID | 3.6.4 |
| CPUID Hypervisor Codes | 3.6.4.1 |
| CPUID Instance Codes | 3.6.4.2 |

SHPC Condos User Guide

| | |
|---|-------|
| Overview | 4.1 |
| Hardware | 4.2 |
| Storage | 4.3 |
| Software | 4.4 |
| Scheduling Jobs | 4.4.1 |
| Bash Environment Customization | 4.4.2 |
| Modules | 4.4.3 |
| Compilers | 4.4.4 |
| Condo Workflows | 4.4.5 |
| Crystal Workflow | 4.4.6 |
| How to Use | 4.5 |
| Prerequisites | 4.5.1 |
| Request Access to an Allocation | 4.5.2 |
| Access your Allocation | 4.5.3 |
| Execute a Job | 4.5.4 |

User Contributed Tutorials

| | |
|---|-------|
| Note: Community contributed content | 5.1 |
| not officially supported by CADES. | 5.1.1 |
| Launch a Docker Container | 5.2 |
| Launch Shiny within Docker | 5.3 |

| | |
|--|-----|
| Eclipse in CADES HPC | 5.4 |
| Allinea DDT in CADES HPC | 5.5 |

Data Transfer & Storage

| | |
|---|-------|
| Moving Data | 6.1 |
| Graphical Client SFTP | 6.1.1 |
| Globus Data Transfer Tool | 6.2 |
| Globus Endpoints | 6.2.1 |
| Globus Transfers & More | 6.2.2 |
| Globus Command Line Interface | 6.2.3 |
| Scality Object Storage User Guide | 6.3 |
| Scality Advanced Usage | 6.3.1 |
| Scality in a Python Virtual Environment | 6.3.2 |

Contributing to Documentation

| | |
|---|-------|
| Ways to Contribute | 7.1 |
| Contribute with Git | 7.1.1 |
| Git in the Command Line | 7.1.2 |
| CADES Authoring Guide | 7.2 |

CADES Policies

| | |
|---|-----|
| CADES Cloud User Policy | 8.1 |
|---|-----|

CADES User Documentation

Oak Ridge National Laboratory's (ORNL) Compute and Data Environment for Science (CADES) provides an integrated computing infrastructure to deliver data science solutions and workflows to ORNL personnel. CADES provides dedicated computing resources through our **SHPC (Scalable High Performance Computing) Condo** allocations and customizable Software as a Service (SaaS) through our **Birthright Cloud** solution.

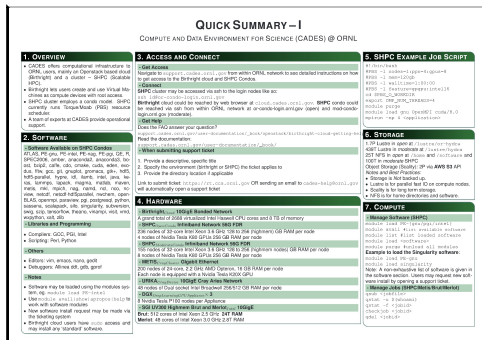
This diverse computing and data ecosystem is bolstered by a CADES support team that employs experts matrixed from different ORNL research directives to provide focused, expert support for a user's scientific computing needs.

Features of this Guide

- The **navigation panel** on the left provides you with a birds-eye view of the content of these pages.
- **GitBook search** at the top of the left-hand side allows you to list only content that matches your keywords.

CADES Quick Summary (Cheat Sheet in pdf format)

- [Download CADES Quick Summary Cheat Sheet](#)



CADES User Documentation Table of Contents

CADES Support

- [Table of Contents](#)
- [Support](#)
- [Access CADES Externally](#)
- [CADES Team](#)
- [CADES Acknowledgement](#)
- [Glossary](#)

Quick-Start Guides

- [Launch a VM Instance: Quick-Start](#)
- [SHPC Condos: Quick-Start](#)

CADES Cloud User Guide

- [Getting Started with OpenStack](#)
 - [Request Your Cloud Allocation](#)
 - [Manage Your Cloud Allocation](#)
 - [OpenStack Project Quota](#)
 - [OpenStack Help](#)
- [Launch a VM Instance](#)
 - [Log In & Name the VM](#)
 - [Configure the VM](#)
 - [Networks & Security](#)
 - [Key Pair Use](#)
- [Access VM Instances](#)
 - [Access Your VM Instance Using SSH](#)
 - [Access Your VM Instance Using PuTTY \(Windows\)](#)
 - [Access Your VM Instance Using Horizon](#)
 - [Add More Users to Your VM Instance](#)
- [Manage Your VM Instances](#)
 - [Delete a VM Instance from Your Project](#)
 - [Delete a Volume from Your Project](#)
 - [Resize a VM Instance](#)
 - [Add a Volume to a VM Instance](#)
 - [Create a Snapshot](#)
 - [OpenStack Security Groups](#)
 - [Modify the Default Security Group](#)
 - [Create a new Security Group](#)
 - [Security Group CIDR Examples](#)
- [Overview - CADES Cloud Information](#)
 - [Available VM Images & Configurations](#)
 - [Software & Hardware](#)
 - [Network & Storage](#)

- [Additional OpenStack Resources](#)
 - [Request Firewall Exception](#)
 - [Run a Simple Web Server](#)
 - [SSL - Let's Encrypt](#)
 - [Install CPUID](#)
 - [CPUID Hypervisor Codes](#)
 - [CPUID Instance Codes](#)

SHPC Condos User Guide

- [Overview](#)
- [Hardware](#)
- [Storage](#)
- [Software](#)
 - [Scheduling Jobs](#)
 - [Bash Environment Customization](#)
 - [Modules](#)
 - [Compilers](#)
 - [Condo Workflows](#)
 - [Crystal Workflow](#)
- [How to Use](#)
 - [Prerequisites](#)
 - [Request Access to an Allocation](#)
 - [Access your Allocation](#)
 - [Execute a Job](#)

User Contributed Tutorials

- [Note: Community contributed content](#)
 - [not officially supported by CADES.](#)
- [Launch a Docker Container](#)
- [Launch Shiny within Docker](#)
- [Eclipse in CADES HPC](#)
- [Allinea DDT in CADES HPC](#)

Data Transfer & Storage

- [Moving Data](#)
 - [Graphical Client SFTP](#)
- [Globus Data Transfer Tool](#)
 - [Globus Endpoints](#)
 - [Globus Transfers & More](#)
 - [Globus Command Line Interface](#)
- [Scality Object Storage User Guide](#)
 - [Scality Advanced Usage](#)
 - [Scality in a Python Virtual Environment](#)

Contributing to Documentation

- [Ways to Contribute](#)
 - [Contribute with Git](#)
 - [Git in the Command Line](#)
- [CADES Authoring Guide](#)

CADES Policies

- [CADES Cloud User Policy](#)

CADES Quick Summary (Cheat Sheet in pdf format)

- [Download CADES Quick Summary Cheat Sheet](#)

[CADES](#) → [User Documentation](#) → [Support](#)

Getting Help

Within this documentation the CADES team and user community have assembled getting started guides, a glossary, and user-created tutorials to help you use resources within our environments.

Sometimes, though, docs are not enough. If you have questions not answered here, or would like to open a trouble ticket, please contact the CADES team directly at ca-des-help@ornl.gov.

You may also join us at <http://ca-des.slack.com> where both CADES operations and community members share tips and can work more interactively together.

Contributing

If you would like to add your own tutorials to this site, correct errors, or expand and clarify content, your contributions are welcome. Please see our [contributing guide](#).

[CADES](#) → [User Documentation](#) → [External Access](#)

Externally Accessing CADES Resources

You may find that you need access to internal resources at ORNL or CADES when you are off-site. This guide is intended to provide guidance on how to access some of the most common resources that our researchers need.

Note: ORNL VPN service is only available to those using ORNL-owned hardware.

If you need assistance, you can [email the CADES team](#) or join our [Slack community](#) (available externally).

CADES External Login Node

CADES maintains an external login node `ca-des-extlogin1.ornl.gov`.

New Users: To log in to this node, you must be granted access via the [XCAMS portal](#).

If you are not sure if you already have a XCAMS account, visit [this site](#) and use the "Forgot your username?" and "Forgot your password?" links to investigate. If you find that you have an account, but need access to the CADES-misc resource, [email us](#).

To login to the external access node use your XCAMS ID to SSH:

```
ssh xcams@ca-des-extlogin1.ornl.gov
```

From here you can access internal resources, including:

- OpenStack Horizon Web Interface
- CADES SHPC Condos
- ORNL Internal Websites
- User Documentation: a copy of CADES user documentation is regularly saved to `tmp/ca-des-user-guide.pdf`

Accessing the OpenStack Horizon Web Interface

To view the OpenStack Horizon Web Interface you may use SSH port forwarding (replace `user` with your user ID and enter your password when prompted):

```
ssh -L 9000:cloud.ca-des.ornl.gov:80 user@ca-des-extlogin1.ornl.gov
```

Then view in your browser: <http://localhost:9000>

Accessing the CADES SHPC Cluster

Open Protection Zone

1. After logging in to the CADES external login node, execute `ssh xcams@or-condo-login.ornl.gov`.
 - Replace `xcams` with the username you registered above.
2. When prompted, enter your XCAMS password.

Moderate Protection Zone

1. After logging in to the CADES external login node, execute `ssh ucams@mod-condo-login.ornl.gov`.
 - Replace `ucams` with your ORNL UCAMS ID.

2. When prompted, enter your UCAMS password.

Accessing ORNL Internal Websites

To view internal ORNL websites you may use SSH port forwarding (replace `USER` with your user ID and enter your password when prompted):

```
ssh -L 9000:portal.ornl.gov:80 user@ca-des-extlogin1.ornl.gov
```

Then view in your browser: <http://localhost:9000>

[CADES](#) → [User Documentation](#) → [CADES Team](#)

How can we help you?

CADES staff is ready to assist you. Choose one of the two ways:

- Send an **email to** ca-des-help@ornl.gov.
- Join our **Slack channel** at <https://ca-des.slack.com/signup>.

We are led by:

- **CADES Director** *Arjun Shankar* shankarm@ornl.gov
- **CADES Leader** *Brian Zachary* zacharybs@ornl.gov

Acknowledgments in Scientific Publications and Presentations

Please acknowledge in your publications the role CADES facilities played in your research. Alerting our communications staff when a paper is accepted is also appreciated.

Sample acknowledgement:

This research used resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725"

You may use any variation on this theme, calling out specific simulations or portions of the research that used CADES resources, or citing specific resources used.

However, the crucial elements to include are:

- The spelled out center name (it's okay to include the acronym, too): Compute and Data Environment for Science (CADES)
- Office of Science and U.S. Department of Energy
- Contract No. DE-AC05-00OR22725

We appreciate your conscientiousness in this matter. Acknowledgement and pre-publication notification helps CADES communicate the importance of its role in science to our sponsors, helping assure the continued availability of this valuable resource.

[CADES](#) → [User Documentation](#) → [Glossary](#)

Glossary

The definitions below are related to this document and are provided for quick reference. A more complete list of definitions can be found in the official [OpenStack documentation](#).

Bash

A UNIX shell used for entering command-line executions. Included with most Linux distributions and macOS. Includes SSH capability.

Horizon

A web GUI front end for OpenStack that is accessed via <https://cloud.cades.ornl.gov>.

Hypervisor

Also known as a *virtual machine monitor*, a hypervisor is software/hardware that creates, runs, and manages virtual machines.

Instance

A virtual machine set up through OpenStack. See "Virtual Machine".

OpenStack

A cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center.

Project

The base unit of "ownership" in OpenStack. All resources in OpenStack should be owned by a specific project. In OpenStack Identity, a project must be owned by a specific domain.

Tenant

A legacy OpenStack term for "Project" that is still used in the Horizon web GUI.

Virtual Machine

An operating system instance that runs on top of a hypervisor. Multiple virtual machines (VM) can run at the same time on the same physical host.

CADES Cloud

Oak Ridge National Laboratory's (ORNL's) Compute and Data Environment for Science (CADES) provides eligible customers with an OpenStack cloud computing solution with customizable virtual machines (VM). This resource, called "CADES Cloud," enables customers in [science and technology directorates](#) to leverage self-service portals to rapidly request these VMs for production, testing, and development. These CADES cloud services are available at no cost to ORNL researchers.

Launch a VM: Quick-Start

Creating and launching a VM Instance is one of the first steps to utilizing your cloud allocation. When launching a VM, you will choose a name for your Instance, which network it will utilize (internal or external), which operating system you would like to run (CentOS or Ubuntu), and which VM flavor (size and specifications) you need for your application.

The following is meant to be used as a minimum-detail quick guide to launching a VM. You may find, at the end of the process, that the VM you have created isn't quite right for your needs. The great thing about using these OpenStack VMs is that you can delete it and start over so that you may change the specifications.

Before you begin, have you [requested your CADES Cloud allocation](#)?

An activation notice will be dispatched to your ORNL email address when your resources are ready for use.

- Navigate to the web interface at <https://cloud.cades.ornl.gov/>.
- Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`
- Navigate to `Project` → `Compute` → `Instances` .
 - Click the `Launch Instance` button, and fill out each section of the resulting dialog (shown in the next sections).

What follows is a series of tabs (along the left of the Horizon "Launch Instance" dialog screen). Fill out the information as it pertains to each tab. You may move freely between tabs without losing progress.

- **Details Tab**
 - `Instance Name` – This name can contain up to 15 alphanumeric characters and a hyphen. **No** special characters are permitted, and the host names are case sensitive.
 - `Availability Zone` – "nova" is the default zone.
 - `Count` – The number of instances to start up at once. If creating multiple Instances, the Instance names will be numerated (instance-1, instance-2, instance-3, and so on).
- **Source Tab**
 - `Select Boot Source` – Set this to `Image`
 - `Create New Volume` – Two options:
 - `Yes` : This creates a virtual disk on CADES's central storage (*recommended*). This type of storage is referred to as a "cinder volume."
 - `No` : A virtual disk is created on a hypervisor; this disk is *not* persistent when the VM is deleted. This type of storage is referred to as an "ephemeral volume." The size of the volume can be increased, by [migrating](#) the instance to a larger flavor size, which can be done yourself without CADES assistance.
 - `Delete Volume on Instance Delete` – Set to `No` if data should persist between Instance restarts.
 - `Volume Size` – Must be greater than or equal to the flavor size and fit within your [allocation quota](#).

Increasing the size of a root volume later will require [emailing CADES support](#).

Note: It is recommended to size the root volume appropriate to your needs, keeping in mind that small root volumes are typically used. Additional volumes (for data, logs, etc.) can be attached to an instance, detached and attached to a new VM, etc. A best practice recommendation is for root volumes to contain only the OS.

- `Device Name` – This should almost always be the default, `vda`.
- Choose from the available options by clicking the `+` next to the desired image.
- **Flavor Tab**
 - Choose the flavor which provides the desired CPU and memory and click `+` to add it to your allocation.
- **Networks Tab**
 - Choose **one** of two routable network configurations, and click `+` to add it to your allocation.
 - `general_extnetwork1`, 128.219.184.0/21 - Available from outside ORNL. However, outward-facing services (e.g., SSH, web server) will require ORNL [firewall exceptions](#).
 - `general_intnetwork1`, 172.22.0.0/20 - Internal to ORNL.
- **Network Ports Tab**
 - No user input required. Skip this tab.
- **Security Groups Tab**
 - No user input required for standard SSH access. Skip this tab.
- **Key Pair Tab**

If you skip this step, the instance will not allow you to log in! See [here](#) for more information.

Note: Before deciding between *Option 1* or *Option 2*, you should check your local machine for an existing key pair ([instructions](#)).

 - **Option 1:** Create a new key pair for this instance.
 - Click the `+ Create Key Pair` button.
 - Enter a name for your new key pair in the resulting dialog.
 - Click `Create Key Pair` to associate this new key pair to your Instance.
 - The private key will be downloaded to your local machine as a `.pem` file.
 - On your local machine, place the `.pem` file in the `~/.ssh/` directory ([instructions](#)).
 - **Option 2:** Use available key pair for this Instance.
 - Choose the desired key pair, and click `+` to associate it with your VM Instance.
- **Configuration Tab**
 - No user input required. Skip this tab.
- **Metadata Tab**
 - No user input required. Skip this tab.

Click `Launch Instance` when you have completed all required sections.

Congratulations! A new instance will be launched. Once fully provisioned, the status will change to "Running," and you can access your VM Instance using SSH ([instructions](#)).

Getting Help

If at any point you feel stuck and need some help figuring out your next move, please contact the CADES support team at ca-des-help@ornl.gov or join our Slack channel at <https://ca-des.slack.com/signup>.

More Details About Launching a VM

1. [Log in to Horizon, name your VM](#)
2. [Choose a flavor, image, and boot source](#)
3. [Set up a security group](#)

4. [Configure a key pair for accessing the VM](#)

Related Tutorials

- [Delete a VM](#)
- [Accessing a VM Using SSH](#)
- [Cloud Overview](#)

CADES SHPC Condos

The CADES Scalable HPC (SHPC) Condos consist of two HPC clusters: one in the ORNL Moderate protection zone (CADES Mod) and one in the ORNL Open protection zone (CADES Open). The protection zones contain and control both the software base and the data produced on those systems.

This section outlines the most basic procedures for procuring and using a CADES SHPC Condo allocation.

Step 1: Request Your SHPC Condo Allocation

You can request a CADES SHPC Condo allocation by clicking on the appropriate XCAMS registration [link for your group](#). If you do not see your group listed, please contact the [CADES team](#).

The XCAMS registration process will ask you to acknowledge the XCAMS User Agreement and register your UCAMS with an LDAP group or create a new XCAMS user ID and the register it with an LDAP group. Use the steps below to enter your request.

1. Navigate to the appropriate [XCAMS registration link for your group](#).
2. Enter your email address (your ORNL address if available) and click `Continue`.
3. Review the XCAMS user agreement, and select `Agree`.
4. Enter your UCAMS ID (or a new XCAMS user name).
5. Enter your UCAMS password (or a new XCAMS password).
6. Click `Submit` to complete the XCAMS request.

When your resources are ready for use, an activation notice will be dispatched to the email address entered above. This process can take up to 24 hours to complete.

Step 2: Access Your SHPC Condo Allocation

Once your request for resources has been approved, you can access the SHPC Condo login nodes using SSH. Open and Moderate protection zones each have their own login node. Choose the login node for your protection zone. See below.

Open Protection Zone

1. Open a Bash terminal (or PuTTY for Windows users).
2. Execute `ssh xcams@or-condo-login.ornl.gov`.
 - Replace "xcams" with the username you registered above.
3. When prompted, enter your XCAMS password.

Moderate Protection Zone

1. Open a Bash terminal (or PuTTY for Windows users).
2. Execute `ssh ucams@mod-condo-login.ornl.gov`.
 - Replace "ucams" with your ORNL UCAMS ID.
3. When prompted, enter your UCAMS password.

Step 3: Execute a Job on Your SHPC Allocation

Now that you have access to your allocation through the login node, it is time to do some work. The tutorial linked below, intended for users who are new to the CADES SHPC Condo environment, outlines the basic steps to setting up and executing a job on the SHPC Condo compute nodes.

How to: [Execute a Job on Your SHCP Condo Allocation](#)

Note: Do not execute jobs on the login nodes; only use the login nodes to access your compute nodes. Processor-intensive, memory-intensive, or otherwise disruptive processes running on login nodes will be killed without warning.

Getting Help

If at any point you feel stuck and need some help figuring out your next move, please contact the CADES support team at ca-des-help@ornl.gov or join our Slack channel at <https://ca-des.slack.com/signup>.

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Getting Started](#)

Getting Started with your CADES Cloud Resources

Using CADES OpenStack resources is meant to be as straightforward as possible to get users up and running quickly. This guide provides walk-throughs guides and detailed information for you to get the most out of these resources.

[What is OpenStack? What are the CADES Cloud Resources?](#)

This guide includes several sections that act both as a step-by-step guide and a quick reference.

In this and the following sections, you will discover how to:

1. [Request and Manage your Cloud Allocation](#)
2. [Launch a Virtual Machine \(VM\)](#)
3. [Access a VM](#)
4. [Manage a VM](#)
5. Learn about [Additional Resources](#)

Before you begin, be sure you meet the prerequisites below.

Prerequisites

To properly utilize your CADES Cloud allocation, you will need a couple of utilities loaded on your local machine. These utilities are free and widely used for this type of application.

- Required: **SSH client**
- Recommended: **Bash terminal**

Note: CADES does not provide support for getting these utilities up and running on your personal computer.

MacOS and Linux

Both macOS and Linux distributions includes a Bash terminal and an SSH client by default. No additional software should be required to access your VM Instance.

Windows Users

► [Click for Details](#)

Next Steps

Before you can use the CADES Cloud resources, you will need to [request a cloud allocation](#). After your request has been approved, you can [manage your resources](#) and [launch a VM](#).

Looking for More Information?

- [Overview of CADES Cloud Resources](#)

[CADES](#) → [User Documentation](#) → [Birthright Cloud User Guide](#) → [Getting Started](#) → [Request Cloud Allocation](#)

Request Your Birthright Cloud Allocation

Any member of a [science and technology directorate](#) can request a Birthright Cloud allocation. This is currently an automated process that takes 1–2 hours to complete. Use the steps below to enter your request.

1. Navigate to ORNL's [XCAMS portal](#). The preceding link will prefill a request for "CADES Birthright Cloud" resources.
2. Ensure your [UCAMS ID](#) is selected.
3. Enter a reason for the request (e.g., "I need a Birthright Cloud allocation.").
4. Click [Next](#) to complete the XCAMS request.

An activation notice will be dispatched to your ORNL email address when your resources are ready for use.

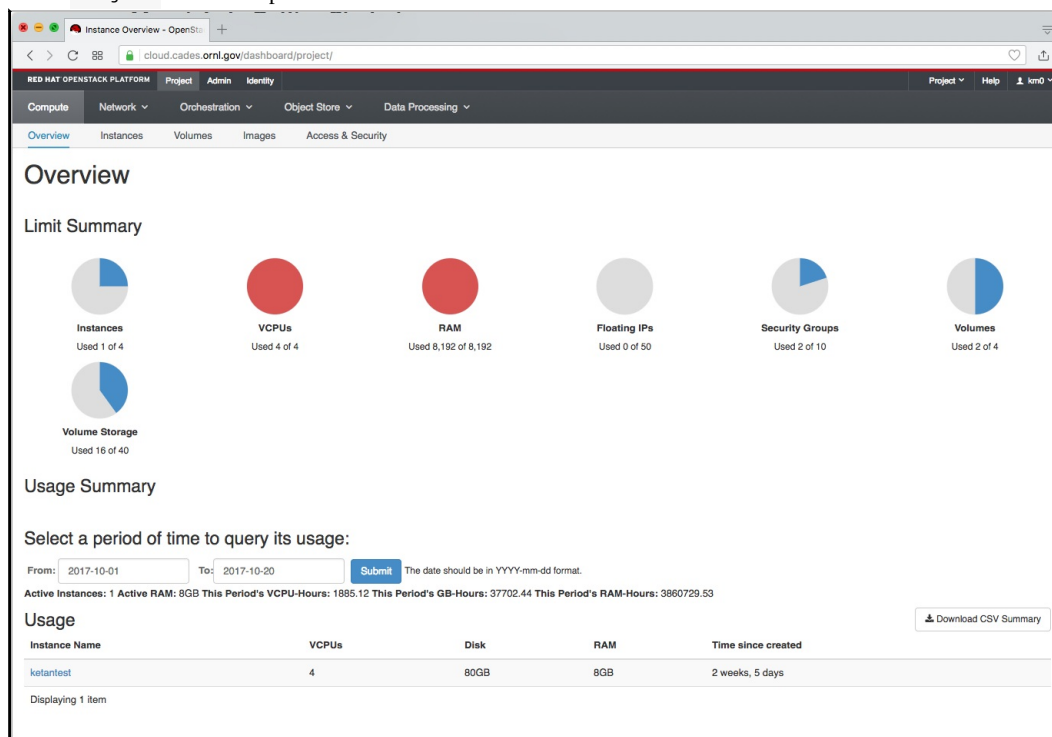
Important Notes for Requesting Your Birthright Cloud Allocation

- Owing to resource constraints, we are currently throttling access to Birthright Cloud allocations. If during registration you get an error that the group is full, please [contact the CADES team](#).
- The registration process can take 1–2 hours. If after 2 hours you are unable to log in, please contact the [CADES team](#) and include the following information in your email:
 - UCAMS ID
 - Contact information
 - Approximate time of your registration attempt

Manage Your Birthright Cloud Allocation

Once your request for resources has been approved, you can manage your allocation using OpenStack's web-based Horizon GUI.

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - o Domain: ornl
 - o User Name: UCAMS ID
 - o Password: UCAMS password
3. Select **Project** from the top left menu to view available resources.



The screenshot shows the OpenStack Horizon GUI for a project. The 'Limit Summary' section contains the following data:

| Resource | Used | Limit |
|-----------------|----------------|-------|
| Instances | 1 of 4 | 4 |
| VCPUs | 4 of 4 | 4 |
| RAM | 8,192 of 8,192 | 8,192 |
| Floating IPs | 0 of 50 | 50 |
| Security Groups | 2 of 10 | 10 |
| Volumes | 2 of 4 | 4 |
| Volume Storage | 16 of 40 | 40 |

The 'Usage Summary' section shows a date range from 2017-10-01 to 2017-10-20. The usage data is as follows:

| Instance Name | VCPUs | Disk | RAM | Time since created |
|---------------|-------|------|-----|--------------------|
| ketantest | 4 | 80GB | 8GB | 2 weeks, 5 days |

From here you can easily manage your allocation through a variety of tools that enable you to:

- o [View Your OpenStack Project Quota](#)
- o [Manage VM Instances](#)
- o [Manage OpenStack Security Groups](#)
- o [Configure SSH Access to VM Instances](#)

Your OpenStack Project Quota

Each CADES Birthright Cloud allocation creates a "Project". Each Project in CADES has a preset resource quota that can be leveraged at the user's discretion.

Request More Resources

If a user requires more resources for their allocation, he or she can submit a proposal to the CADES Resource Utilization Council (RUC) to request a quota increase. This proposal should describe the resources desired (RAM, CPUs, storage, etc.) and the scientific goal and merit of the work being performed using the CADES Birthright Cloud allocation. These requests for increased resources, subject to review, should be directed to the [CADES team](#).

View Your Project Quota

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`
3. Navigate to `Project` → `Overview`.

In this Overview you can see your resource allocation, including:

- **Instances** – The number of VMs you can run at once.
- **VCPUS** – The number of CPU cores you can use across all of your VMs.
- **RAM** – The total amount of RAM you can use across all of your VMs.
- **Floating IPs** – COMING SOON – IP addresses that you can attach and move between Instances. These addresses are in addition to the IP(s) already allocated to your Instance.
- **Security Groups** – Blocks of firewall rules that you can attach to an Instance. All groups start with a default Security Group that contains a basic configuration to get started. You can use this "default" group or you can create your own group. Check out the [Security Groups section](#) for more information.
- **Volumes** – These are the number of block-storage volumes that you can create. You can attach any number of volumes to an Instance. These are also the preferred method of storage in the Birthright cloud environment.
- **Volume Storage** – This is the total storage available to the volumes in your Project.

Once an Instance is launched, the resource utilization appears at the bottom of Overview screen, under the Usage section.

Getting Help

In addition to OpenStack's documentation, the CADES team has assembled a list of [frequently asked questions](#), a [glossary](#), and detailed information to help you get moving on your OpenStack Project.

If you have questions that were not answered in the resources listed here or would like to open a trouble ticket, please contact the CADES team directly at ca-des-help@ornl.gov.

If you have questions about getting your SAFER firewall rules in place, email the SAFER team directly at opsapprovers@ornl.gov.

General OpenStack Usage Questions

The OpenStack community has extensive documentation, a general mailing list, and a mechanism for asking general OpenStack usage questions. These resources can be indispensable for new users, and the CADES support team recommends that new users leverage the [OpenStack documentation](#) and other support features for general OpenStack questions.

| Information | Description |
|--|--|
| OpenStack documentation | OpenStack maintains significant documentation for general OpenStack usage. The user manuals cover OpenStack's functionality in depth, although not all of the features listed in their documentation are available from CADES. |
| ask.openstack.org | General end-user resource with questions and answers. |
| OpenStack mailing list | General questions mailing list, OpenStack operations mailing list, and a developer mailing list. |

Frequently Asked Questions

- *Who can request a CADES Birthright Cloud allocation? Can I have one?*
 - Any member of a science and technology directorate can request a Birthright Cloud allocation. See [list of ORNL's technical organizations](#).
- *What kind of user support does CADES provide for Birthright Cloud allocations?*
 - The CADES team can provide support for issues in OpenStack and can help a user get started with their allocation—up to and including loading and launching a VM Instance from one of the CADES-provided images. Once a user begins adding software and configuring the operating system on their VM, they are more-or-less responsible for any issues that might present themselves in the VM.
- *Does CADES provide backups or support for backing up my VM Instance(s)?*
 - In short, no. CADES does not provide backups of a user's VM Instances and does not currently support this functionality. There are several methods a user could leverage to back up the operating system running on the VM, but any such method must comply with ORNL's best practices, and this backup strategy is the user's responsibility and is beyond the scope of this user guide.
- *Can I have my instance IPs be sequential?*
 - Owing to the dynamic nature of a cloud environment this is not possible. We recommend using `/etc/hosts`, Netreg (DNS), or another mechanism to accomplish host-to-host communication.
- *How do I get access to Lustre on my instance?*

- Lustre access is currently provided on a case-by-case basis. [Contact the CADES team](#) if you wish to add Lustre to your Project.
- *What ORNL file systems can I mount within my Birthright Cloud Instance?*
 - The target file system must be "open science." No moderate/confidential file systems or their respective data can or should be mounted within your VM Instance.

Launch a VM: At-A-Glance

Creating and launching a VM Instance is one of the first steps to utilizing your CADES Cloud allocation. When launching a VM, you will choose a name for your Instance, which network it will utilize (internal or external), which operating system you would like to run (CentOS or Ubuntu), and which VM flavor you need for your application.

Note: When launching a new VM Instance, be aware that your VM Instance name *may* also serve as your DNS host name. This name can contain up to 20 alphanumeric characters and a hyphen. **No** special characters are permitted, and the host names are case sensitive. For example, good: `my-instance-server`, no good: `my_instance&server`. See [RFC 952](#) and [RFC 1123](#) for more information.

If you would like to have a DNS host name for your instance, please submit a ticket to ca-des-help@ornl.gov with the Instance Name and the OpenStack Project ID.

Prerequisites

You will need to have a CADES Cloud allocation before you can launch your VM Instance using the web GUI. The link below will show you how to request your CADES Cloud allocation.

How to: [Request Your CADES Cloud Allocation](#)

There are three primary steps to complete to have a functioning VM Instance. These steps are outlined below. After each section a link is provided where you can find more information and in-depth descriptions about the configuration options.

1. [Log in to Horizon, name your VM.](#)
2. [Choose a flavor, image, and boot source.](#)
3. [Set up a security group.](#)
4. [Configure a key pair for accessing the VM.](#)

Launch a VM: Log In & Naming

Once you receive the email notification that your resources are available, you can login to the Horizon web interface and get started.

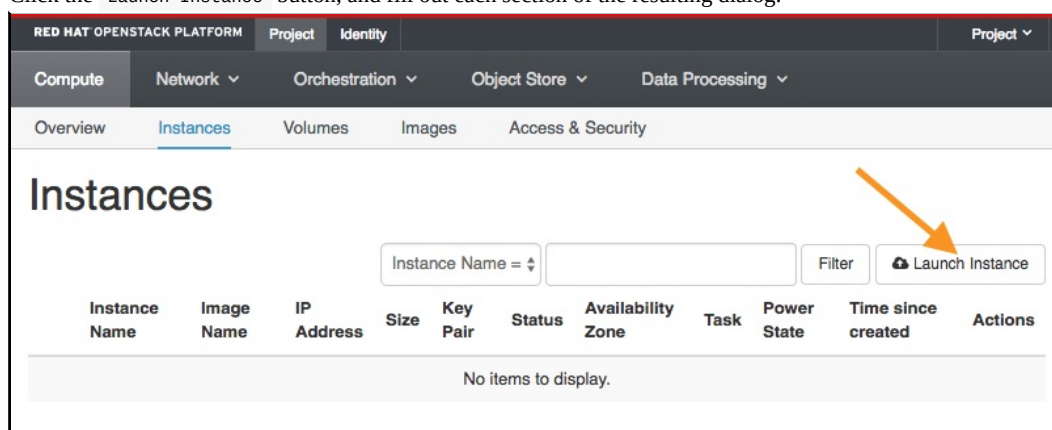
Log In to Horizon

- Navigate to the web interface at <https://cloud.cades.ornl.gov/>.
- Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`

You can now launch a new VM Instance from within the web interface.

Launch an Instance

1. Navigate to `Project` → `Compute` → `Instances`.
2. Click the `Launch Instance` button, and fill out each section of the resulting dialog.



Details Tab – Fill out the required fields for the Details section.

- `Instance Name` – The instance name can contain up to 20 alphanumeric characters and a hyphen. **No** special characters are permitted, and the host names are case sensitive. For example, good: `my-instance`, no good: `my_instance&server`. See [RFC 952](#) and [RFC 1123](#) for more information.

Note: When launching a new VM Instance, be aware that your VM Instance name *may* also serve as your DNS host name. If you wish for a DNS record to be created, matching the instance name you create, email the name of the instance to ca-des-help@ornl.gov with the subject "Request DNS Name Creation for VM."
- `Availability Zone` – "nova" is the default zone. **Do not change unless instructed to do so by a CADES administrator.**
- `count` – The number of instances to start up at once. If using multiple Instances, the Instance names will be numerated (instance-1, instance-2, instance-3, and so on).

Launch Instance

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

Instance Name *
black-mesa

Availability Zone
nova

Count *
1

Total Instances (4 Max)
25%

0 Current Usage
1 Added
3 Remaining

- Once finished, click `next` to proceed to the next tab.

Launch a VM Instance: Flavor and Image

On the next two tabs, you will choose the storage type and size, and the operating system image.

Source Tab – Fill out the required fields for the Source tab.

- **Select Boot Source** – Set this to **Image**
- **Create New Volume** – Two options:
 - **Yes** : This creates a virtual disk on CADES's central storage (*recommended*). This type of storage is referred to as a "cinder volume."
 - **No** : A virtual disk is created on a hypervisor; this disk is *not* persistent when the VM is deleted. This type of storage is referred to as an "ephemeral volume." The size of the volume can be increased, by [migrating](#) the instance to a larger flavor size, which can be done yourself without CADES assistance.
- **Delete Volume on Instance Delete** – Set to **No** if data should persist between Instance restarts.
- **Volume Size** – Must be greater than or equal to the flavor size and fit within your [allocation quota](#).

Increasing the size of a root volume later will require [emailing CADES support](#).

Note: It is recommended to size the root volume appropriate to your needs, keeping in mind that small root volumes are typically used. Additional volumes (for data, logs, etc.) can be attached to an instance, detached and attached to a new VM, etc. You may also consider attaching shared NFS storage rather than adding large volumes to each VM ([email ca-des-help@ornl.gov](mailto:ca-des-help@ornl.gov)). A best practice recommendation is for root volumes to contain only the OS.

- **Device Name** – This should almost always be the default, `vda`.
- **Available** – List of available images. Choose from the available options by clicking the **+** next to the desired image.

We're choosing the Ubuntu image for this example.

Launch Instance ✕

Details

Source *

Flavor *

Networks *

Network Ports

Security Groups

Key Pair

Configuration

Metadata

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image or a volume (if enabled). You can also choose to use persistent storage by creating a new volume. ?

Select Boot Source

Image ▼

Create New Volume

Yes No

Volume Size (GB) *

8

Delete Volume on Instance Delete

Yes No

Device Name

vda

Allocated

| Name | Updated | Size | Type | Visibility |
|--|---------|------|------|------------|
| Select a source from those listed below. | | | | |

▼ Available 2 Select one

Q [Click here for filters.](#)

| Name ^ | Updated | Size | Type | Visibility | |
|-------------------------------|-----------------|---------|-------|------------|---|
| CADES_CentOS-7.3_v20170310_0 | 3/14/17 3:04 PM | 1.43 GB | QCOW2 | Public | + |
| CADES_Ubuntu16.04_v20170111_0 | 1/12/17 9:33 PM | 1.83 GB | QCOW2 | Public | + |

✕ Cancel

< Back

Next >

Launch Instance

- Upon selecting an image, it will move from the *Available* list to the *Allocated* list.

Launch Instance
✕

Details

Source

Flavor *

Networks *

Network Ports

Security Groups

Key Pair

Configuration

Metadata

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image or a volume (if enabled). You can also choose to use persistent storage by creating a new volume. ?

Select Boot Source

Image

Volume Size (GB) *

8

Device Name

vda

Create New Volume

Yes No

Delete Volume on Instance Delete

Yes No

Allocated

| Name | Updated | Size | Type | Visibility | |
|---------------------------------|------------------|---------|-------|------------|---|
| > CADES_Ubuntu16.04_v20170111_0 | 1/12/17 9:33 P M | 1.83 GB | QCOW2 | Public | - |

▼ Available 1 Select one

| Name ^ | Updated | Size | Type | Visibility | |
|--------------------------------|------------------|---------|-------|------------|---|
| > CADES_CentOS-7.3_v20170310_0 | 3/14/17 3:04 P M | 1.43 GB | QCOW2 | Public | + |

✕ Cancel
< Back
Next >
Launch Instance

- Once finished, click `next` to proceed to the next section.

Flavor Tab – Choose an image flavor for your VM Instance.

- Available – List of available images. Choose an image, and click + to add it to your allocation.

Launch Instance ✕

Details *

Source *

Flavor *

Networks *

Network Ports

Security Groups

Key Pair

Configuration

Metadata

Flavors manage the sizing for the compute, memory and storage capacity of the instance ?

Allocated

| Name | VCPUS | RAM | Total Disk | Public |
|---|-------|-----|------------|--------|
| Select an item from Available items below | | | | |

▼ Available 5 Select one

| Name | VCPUS | RAM ^ | Total Disk | Public |
|-------------|-------|---------|------------|--|
| ▶ m1.tiny | 1 | 512 MB | 8 GB | Yes + |
| ▶ m1.small | 1 | 2 GB | 20 GB | Yes + |
| ▶ m1.medium | 2 | 4 GB | 40 GB | Yes + |
| ▶ m1.large | ⚠ 4 | ⚠ 8 GB | 80 GB | Yes + ⚠ |
| ▶ m1.xlarge | ⚠ 8 | ⚠ 16 GB | 160 GB | Yes + ⚠ |

✕ Cancel
< Back
Next >
Launch Instance

Note: The ⚠ indicates that your quota has inadequate resources for the image.

- Upon selection, the chosen image will move to the *Allocated* list.

Launch Instance ✕

Details ^{*}

Source ^{*}

Flavor

Networks ^{*}

Network Ports

Security Groups

Key Pair

Configuration

Metadata

Flavors manage the sizing for the compute, memory and storage capacity of the instance ?

Allocated

| Name | VCPUS | RAM | Total Disk | Public | |
|-----------|-------|--------|------------|--------|---|
| > m1.tiny | 1 | 512 MB | 8 GB | Yes | - |

▼ Available 4 Select one

| Name | VCPUS | RAM [▲] | Total Disk | Public | |
|-------------|---|---|------------|--------|---|
| > m1.small | 1 | 2 GB | 20 GB | Yes | + |
| > m1.medium | 2 | 4 GB | 40 GB | Yes | + |
| > m1.large | ▲ 4 | ▲ 8 GB | 80 GB | Yes | + ▲ |
| > m1.xlarge | ▲ 8 | ▲ 16 GB | 160 GB | Yes | + ▲ |

✕ Cancel
< Back
Next >
Launch Instance

- Once finished, click `next` to proceed to the next section.

Launch a VM: Networks and Security

Networks Tab – Choose a network for your VM Instance.

Note: Contact the CADES team if you require more than one IP from each Network. Additionally, if your needs are not met by following this guide, feel free to [email the CADES team](#) to discuss options.

- **Available** – List of available networks. Choose **one** of two routable network configurations, and click **+** to add it to your allocation.
 - `general_extnetwork1` , 128.219.184.0/21 - Available from outside ORNL. However, outward-facing services (e.g., SSH, web server) will require ORNL firewall exceptions ([instructions](#)).
 - `general_intnetwork1` , 172.22.0.0/20 - Internal to ORNL.

Note: If you wish to run services on your VM Instance that should be available outside of ORNL's network, ensure that you select the External Network option when setting up your VM Instance and that you also add a rule to your Security Group for that particular service.
- For this example, we will choose the external network (`general_extnetwork1`).

The screenshot shows the 'Launch Instance' dialog box with the 'Networks' tab selected. The 'Available' section is expanded, showing a search bar and a table of available networks. The table has columns for Network, Subnets Associated, Shared, Admin State, and Status. Two networks are listed: 'or_provider_general_extnetwork1' and 'or_provider_general_intnetwork1'. An orange arrow points to the '+' button next to the external network option.

| Network | Subnets Associated | Shared | Admin State | Status |
|-----------------------------------|--------------------------------|--------|-------------|----------|
| > or_provider_general_extnetwork1 | or_provider_general_extsubnet1 | Yes | Up | Active + |
| > or_provider_general_intnetwork1 | or_provider_general_intsubnet1 | Yes | Up | Active + |

- **Allocated** – Upon selection, the chosen network will move to the *Allocated* list.

Launch Instance

Details *
Source *
Flavor
Networks
Network Ports
Security Groups
Key Pair
Configuration
Metadata

Networks provide the communication channels for instances in the cloud. ?

▼ Allocated ¹ Select networks from those listed below.

| | Network | Subnets Associated | Shared | Admin State | Status | |
|-------|-------------------------------------|------------------------------------|--------|-------------|--------|---|
| ⇅ 1 > | or_provider_general_extnetw ork1 | or_provider_general_extsub net1 | Yes | Up | Active | - |

▼ Available ¹ Select at least one network

Q Click here for filters.

| | Network ^ | Subnets Associated | Shared | Admin State | Status | |
|---|-------------------------------------|------------------------------------|--------|-------------|--------|---|
| > | or_provider_general_intnetwo rk1 | or_provider_general_intsubn et1 | Yes | Up | Active | + |

✕ Cancel < Back Next > Launch Instance

- Once finished, click `next` to proceed to the next section.

Network Ports Tab – No user input required. Skip this step.

Security Groups Tab - Choose a security group for your VM Instance.

Note: Skipping this step will make your VM Instance unreachable! [See additional documentation on Security Groups.](#)

- `Available` – List of available security groups. Choose the desired Security group, and click `+` to add it to your allocation. The `default` Security Group, used for this example, has the basic services you need to get started. Users can also [create their own custom Security Groups](#).

Launch Instance

Details *
Source *
Flavor
Networks
Network Ports
Security Groups
Key Pair
Configuration
Metadata

Select the security groups to launch the instance in. ?

▼ Allocated

Name ^

Select one or more security groups from the available groups below.

▼ Available ² Select one or more

Q Filter

| | Name ^ | Description | |
|---|---------|------------------------|---|
| > | Custom | Custom security group | + |
| > | default | Default security group | + |

✕ Cancel < Back Next > Launch Instance

- `Allocated` – Upon selection, the chosen Security Group(s) will move to the `Allocated` list.

Launch Instance

Select the security groups to launch the instance in.

▼ Allocated ¹

Name ▲

| | |
|-----------|---|
| > default | - |
|-----------|---|

▼ Available ¹ Select one or more

Q Filter

| Name | Description | |
|----------|-----------------------|---|
| > Custom | Custom security group | + |

✕ Cancel < Back Next > Launch Instance

- Once finished, click `next` to proceed to the next section.

Launch a VM: Key Pair Use

Key Pair Tab – An SSH key pair is required to access your VM Instance.

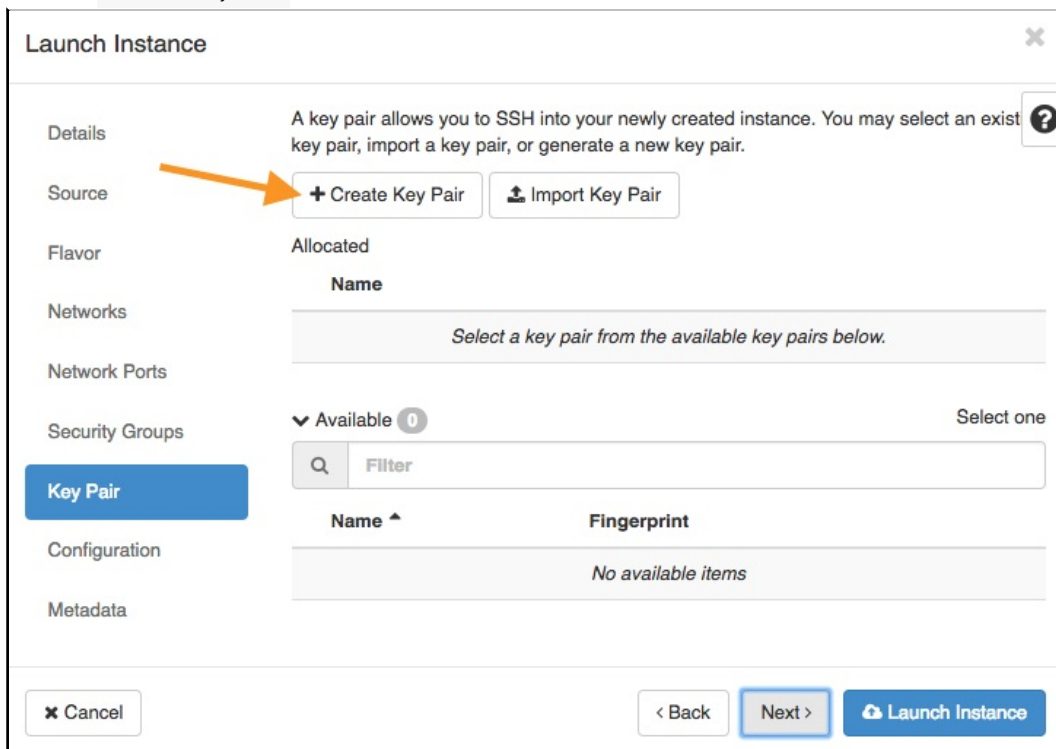
You can create a new key pair for this Instance (**Option 1**) or choose a key pair from the Available list (**Option 2**).

Note: Before deciding, you should check your local machine for an existing key pair ([instructions](#)). If you skip this step, the instance will not allow you to log in!

See the [Access your VM Instance Using SSH documentation](#) for more information.

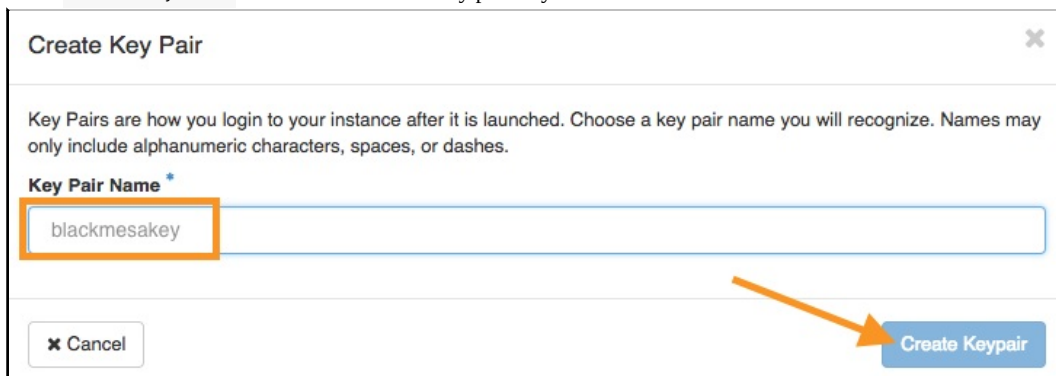
- **Option 1: Create a new key pair for this instance.**

- Click the `+ Create Key Pair` button.



The screenshot shows the 'Launch Instance' dialog box with the 'Key Pair' tab selected. The 'Source' section has two buttons: '+ Create Key Pair' and 'Import Key Pair'. An orange arrow points to the '+ Create Key Pair' button. Below this, there is a section for 'Available' key pairs, which is currently empty, showing 'No available items'. At the bottom, there are buttons for 'Cancel', '< Back', 'Next >', and 'Launch Instance'.

- Enter a name for your new key pair in the resulting dialog.
- Click `create key pair` to associate this new key pair to your Instance.

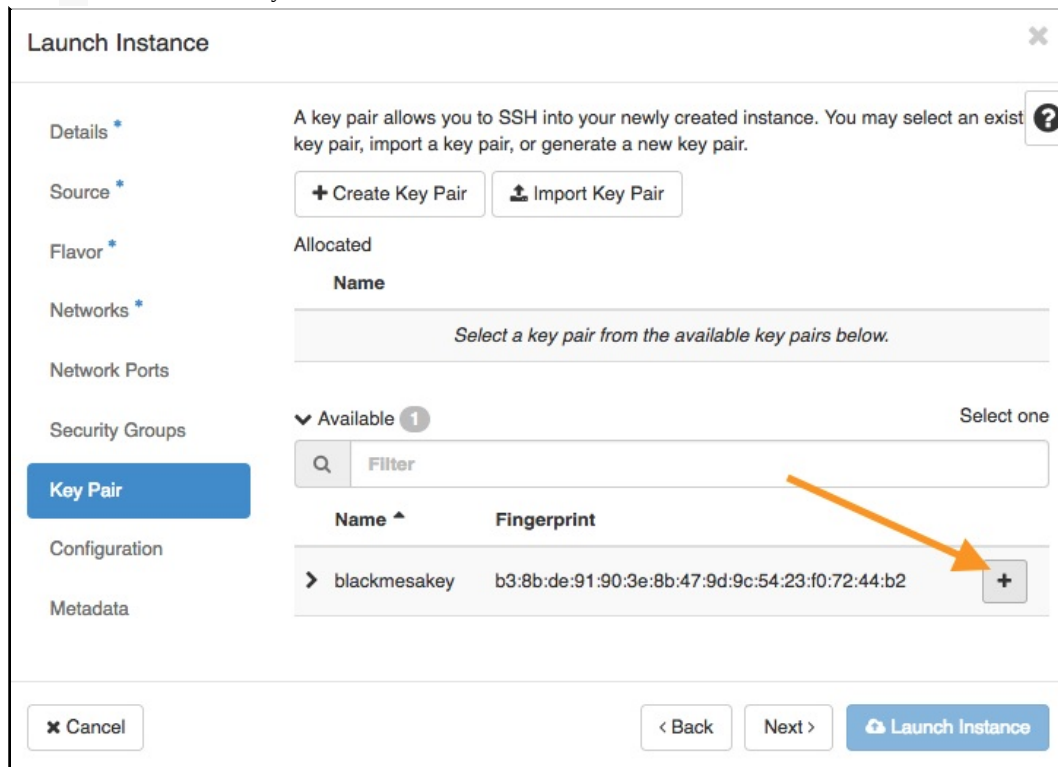


The screenshot shows the 'Create Key Pair' dialog box. The 'Key Pair Name' field is highlighted with an orange box and contains the text 'blackmesakey'. Below the field, there is a 'Create Keypair' button, which is also highlighted with an orange arrow. At the bottom left, there is a 'Cancel' button.

- The private key will be downloaded to your local machine as a `.pem` file.
- On your local machine, place the `.pem` file in the `~/.ssh/` directory ([instructions](#)).
- Once finished, click `next` to proceed to the next tab.

- **Option 2: Use available key pair for this Instance.**

- o Available – List of available key pairs that were previously generated or imported. Choose the desired key pair, and click + to associate it with your VM Instance.



Launch Instance

Details * A key pair allows you to SSH into your newly created instance. You may select an exist key pair, import a key pair, or generate a new key pair.

Source *

Flavor * Allocated

Networks *

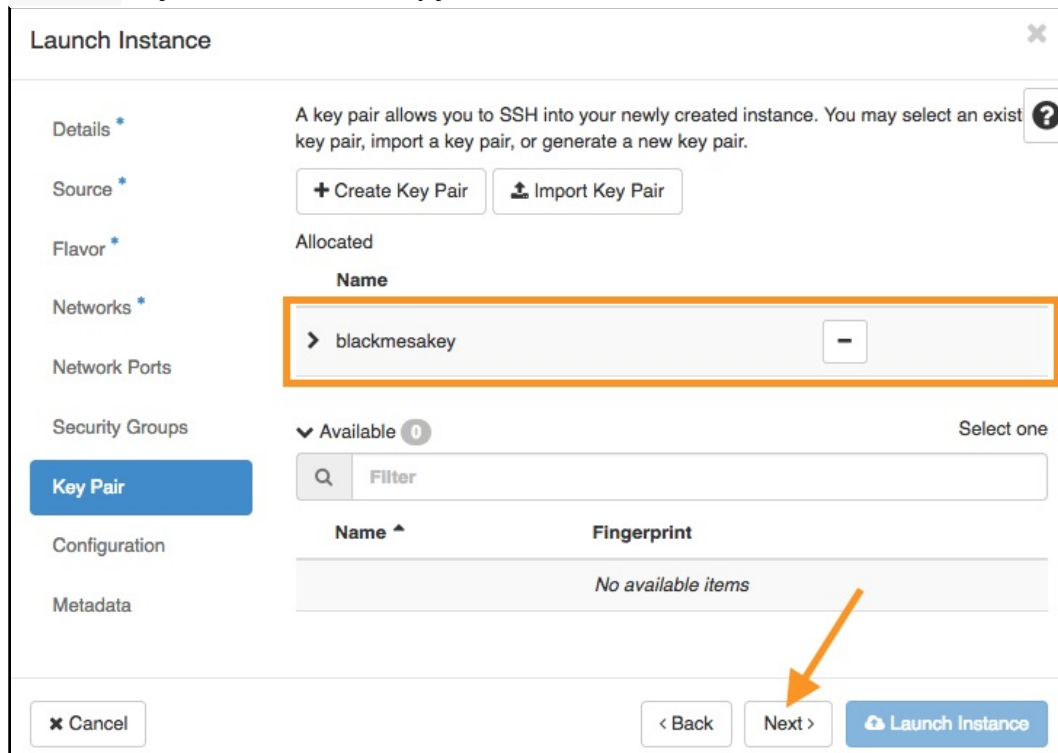
Network Ports

Security Groups Available 1

Key Pair

| Name ^ | Fingerprint | |
|----------------|---|----------------------------------|
| > blackmesakey | b3:8b:de:91:90:3e:8b:47:9d:9c:54:23:f0:72:44:b2 | <input type="button" value="+"/> |

- o Allocated – Upon selection, the chosen key pair will move to the *Allocated* list.



Launch Instance

Details * A key pair allows you to SSH into your newly created instance. You may select an exist key pair, import a key pair, or generate a new key pair.

Source *

Flavor * Allocated

Networks *

Network Ports

Security Groups Available 0

Key Pair

| Name ^ | Fingerprint | |
|--------------------|-------------|--|
| No available items | | |

- o Once finished, click next to proceed to the next tab.

Configuration Tab – This section is not required for deployment and is not currently supported by the CADES team.

Metadata Tab – No user input required. Skip this step.

Click `Launch Instance` when you have completed all required sections. *Congratulations!* A new instance will be launched. Once fully provisioned, the status will change to "Running," and you can access your VM Instance using SSH ([instructions](#)).

Access VM Instances Running in OpenStack

There are several ways you can access your VM Instances.

1. The best way to access your VM Instance is [through an SSH connection](#).
 - If you use Windows: [Access Your VM Instance Using PuTTY](#)
2. You can also [access the VM Instance's terminal using Horizon](#).
 - But you must first [create a new user for this purpose using SSH](#).

Note: UCAMS credentials and SSH key pairs are unrelated, unconnected authentication methods. A user will only be prompted for a UCAMS username and password when he or she logs in to Horizon. All other authentication relies on SSH key pairs or, in the case of creating your own non-UCAMS users, a generic user-created password.

Note: To use the Horizon console, you must first [add a user and password](#) to the Instance's operating system (via SSH) to enable access to the VM without an SSH key.

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Access Your VM Instance](#) → [Access Your VM Instance Using SSH](#)

Access Your VM Instance Using SSH

Once you have created a VM Instance using Horizon, you can access this VM Instance through a secure shell (SSH) using an SSH key pair. The following tutorial will walk you through connecting to your VM Instance through SSH using Linux or macOS.

If you have several users that require access to a single instance, you may request they extract and provide you with their public key, which you may add to your instance's `/home/cades/.ssh/authorized_keys` file. Doing so grants access, as the user `cades`, to that instance. If desired, you may [create specific local accounts](#), with limited sudo roles, etc. rather than use the `cades` user account.

Windows users will have to install an SSH client on their machine, and the CADES team recommends PuTTY for this purpose. The CADES team has compiled a separate tutorial for Windows/PuTTY users, linked below.

How to: [Access VM Instances Using PuTTY \(Windows\)](#)

If you would like to access your VM Instance over SSH from *outside* of the ORNL network, you must first [create a firewall \(SAFER\) exception](#). All local (within the ORNL network) SSH connections are permitted by default. Request a SAFER exception with source 'VISITORS' and your VM(s) as destination.

Table of Contents

- [Add SSH Keys to an Instance](#)
 - [Check for Existing Key Pairs](#)
 - [Generate a New Key Pair](#)
 - [Import Existing Key](#)
- [Connect to Your VM Instance Using SSH](#)
 - [Find Your Instance's IP](#)
 - [Connect to Your Instance](#)
- [Extract Public Key](#)
- [Related Tutorials](#)

Add SSH Keys to an Instance

The first step to adding a key pair is actually to check for existing key pairs on your local machine. If no key pairs exist, then you can [generate a new key](#) in the Horizon web GUI.

If you do have an existing key pair on your machine, then you can [import the public key](#) using the Horizon web GUI. Alternatively, you can create a new (additional) key for your Instance and use a `config` file to manage your SSH credentials.

Check for Existing Key Pairs

1. Open a Bash terminal.
2. Execute `ls -al ~/.ssh`.
3. Check the results of the directory listing.

If the directory listing is empty or the directory is not found, then you do not have an existing SSH key and you should follow the procedure outlined in [Generate a New Key Pair](#).

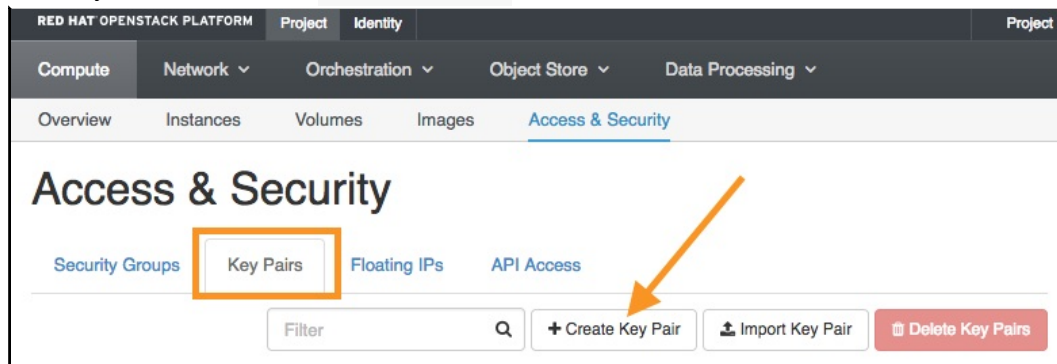
If the directory contains one of the files listed below, then you *do* have an existing SSH key, and you can import the public key into your Instance using the Horizon Web GUI by following the procedure outlined in [Import Existing Key](#).

- id_rsa.pub
- d_dsa.pub
- id_ecdsa.pub
- id_ed25519.pub

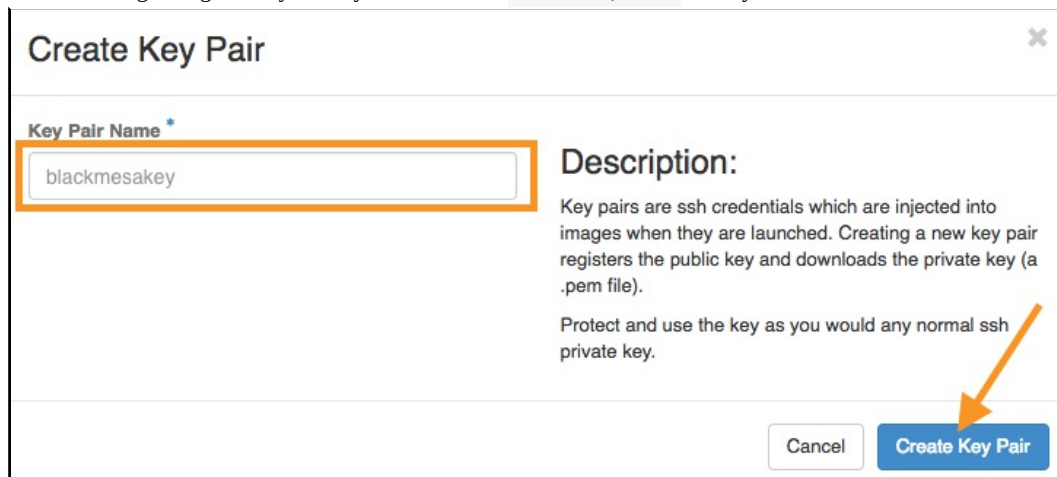
Generate a New Key Pair

Generate the keys

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: ornl
 - User Name: Your three-letter UCAMS ID
 - Password: Your UCAMS password
3. Navigate to Project → Compute → Access & Security → Key Pairs .
4. In the Key Pairs screen, click the + Create New Pair button.




5. In the resulting dialog, name your Key Pair, and click `Create Key Pair` when you're finished.



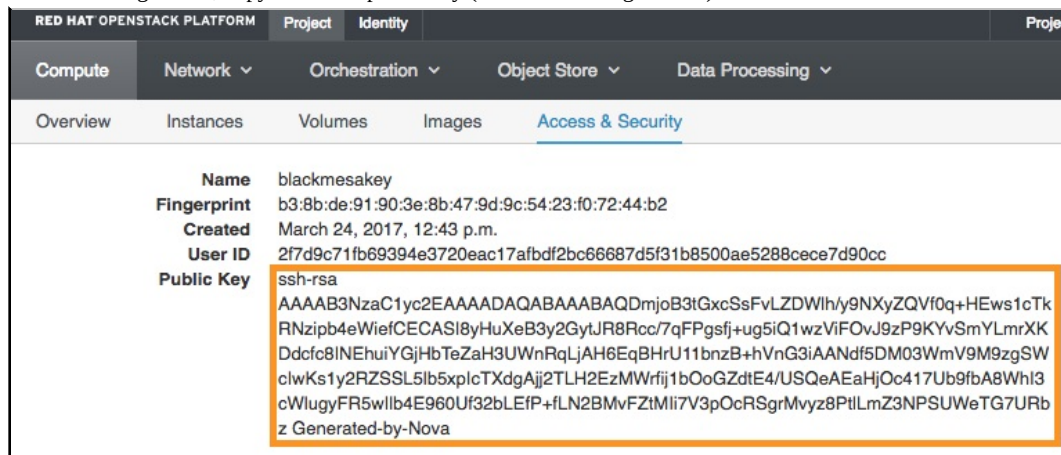
6. The private key will be downloaded to your local machine as a `.pem` file. The public key will be stored in OpenStack.
7. Place the downloaded private key in `~/.ssh/`, e.g., `~/.ssh/blackmesakey.pem`.
8. Secure the private key by setting the permissions to 600 in a Bash terminal.

```
$ chmod 600 ~/.ssh/blackmesakey.pem
```
9. Copy (from Horizon [Option A]) or generate (Bash terminal [Option B]) the public key and place in `~/.ssh/`, e.g., `~/.ssh/blackmesakey.pub`.
 - **Copy public key from Horizon (Option A)**

- Navigate to Project → Compute → Access & Security → Key Pairs
- In the Key Pairs screen, click on the Key Pair you just created (blackmesakey).



- In the resulting screen, copy the entire public key (outlined in orange below).



- Using a text editor, create a new file (e.g., blackmesakey.pub) and paste the public key into this file.
- Save/move this new file to ~/.ssh/blackmesakey.pub .
 - **Generate public key using a Bash terminal (Option B)**
- Open a Bash terminal.
- Ensure your private key is in ~/.ssh/ .
- Use ssh-keygen to generate your public key:

```
$ ssh-keygen -y -f ~/.ssh/blackmesakey.pem > ~/.ssh/blackmesakey.pub .
```

10. Ensure that both your public and your private keys are in ~/.ssh/ .

Note: Having completed the procedure above, you can now connect to your Instance via SSH using a long-form login ([instructions](#)) or you can take these newly generated keys and place them in your local system's default files ([instructions](#)).

Place keys on local machine

If you prefer to leave the Key Pair as it is—as documented above—you can skip the following steps and use a long-form SSH login that specifies the public key's location and file name for each log in attempt. [Click here for instructions](#).

However, if you wish to log in without specifying the key location with each authentication attempt, you can put your public and private keys in their default location on your local system.

Key locations for Linux and macOS:

- Private key should be ~/.ssh/id_rsa .
- Public key should be ~/.ssh/id_rsa.pub .

Note: Since we are generating new keys, these files (and the ~/.ssh/ directory) may not exist on your local machine.

Place the Private key

1. Create a new file using a text editor.
2. Copy your private key from your .pem file and paste it into this new file.

3. Save the new file as `id_rsa` and place it in `~/.ssh/`.
 - o Create the `/.ssh/` directory within `~/` if necessary.
4. Set permissions on this new file using `$ chmod 600 ~/.ssh/id_rsa`.


Place Public key

1. Create a new file using a text editor.
2. Copy your public key from the Horizon web GUI ([instructions](#)).
3. Save the new file as `id_rsa.pub` and place it in `~/.ssh/`.

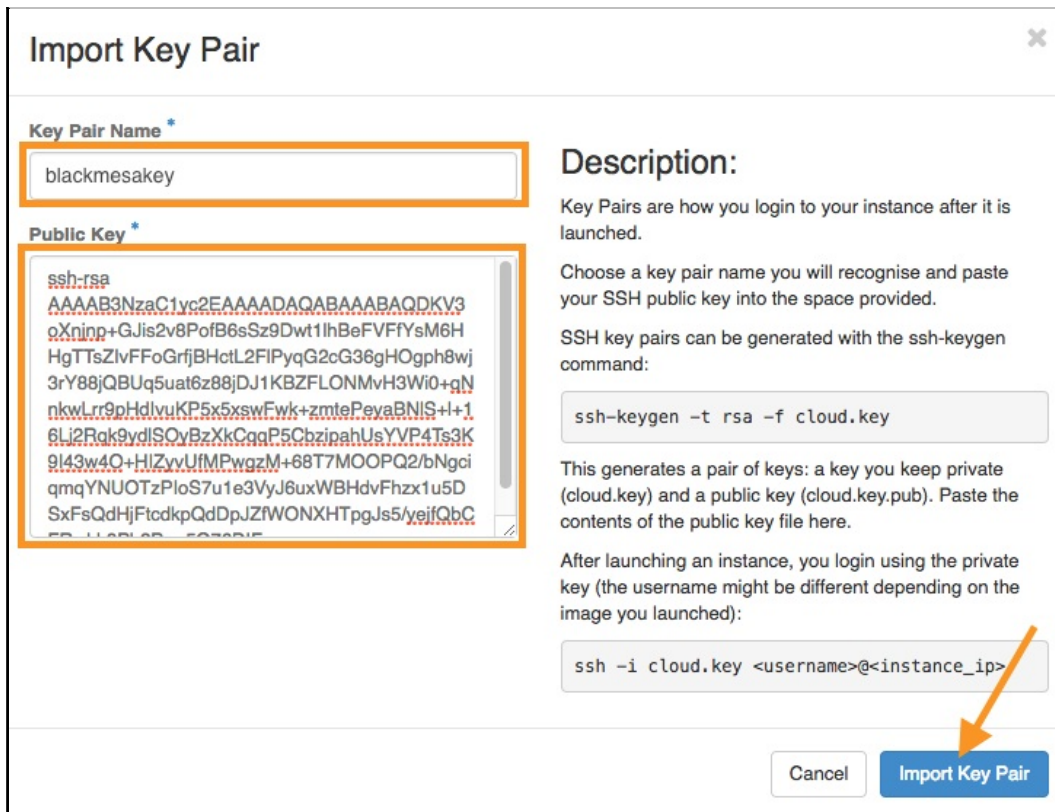
Import Existing Key

If you have an existing key pair that you would like to also use for your OpenStack Project, then you can import the public key using the Horizon web GUI.

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - o Domain: `ornl`
 - o User Name: `Your three-letter UCAMS ID`
 - o Password: `Your UCAMS password`
3. Navigate to `Project` → `Compute` → `Access & Security` → `Key Pairs`.
4. In the Key Pairs screen, click the `Import Key Pair` button.



5. In the resulting dialog, enter a key pair name in the `Key Pair Name` field.
6. Copy your public key from the file (e.g., `~/.ssh/id_rsa.pub`) on your local machine.
7. Paste the public key in the `Public Key` text box.
8. Click the `Import Key Pair` button.



9. Confirm your newly imported key appears in the Key Pair list.

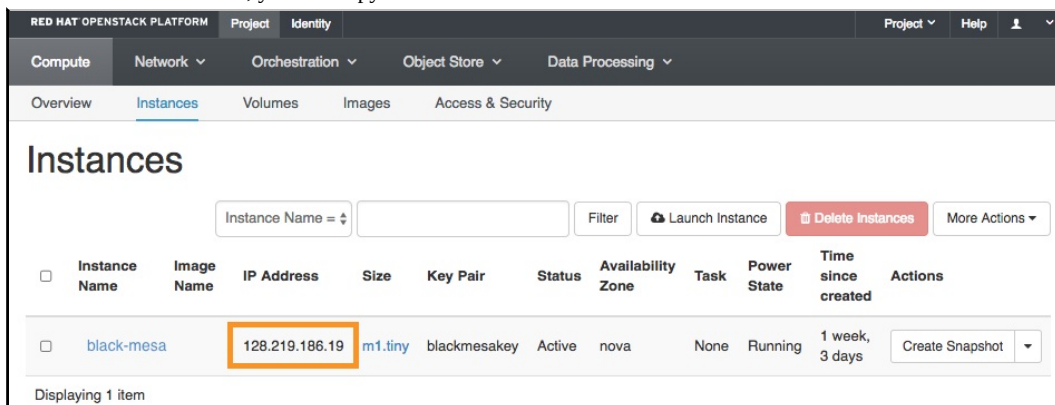
Connect to Your VM Instance Using SSH

All VM Instances have SSH access enabled by default. Use your key pair and the user name "cades" for authentication. SSH to your Instance using the Key Pair you configured above.

Find Your Instance's IP

The IP address of your VM instance is shown in the Instances tab in the Horizon web GUI.

1. Navigate to Project → Compute → Instances .
2. From the Instances screen, you can copy the IP address for the desired VM Instance.



Connect to Your Instance

Once you have determined the IP address of your VM Instance and placed your keys in `~/.ssh/id_rsa` and `~/.ssh/id_rsa.pub`, you can connect using a simple SSH command where "cades" is always the user.

Standard login

1. Open a Bash terminal.
2. Execute `ssh cadés@128.219.186.19`.
 - Replace `128.219.186.19` with the IP address of your own Instance.
3. You should now be connected to your VM Instance via SSH.
 - Use the Bash terminal to install your software-stack and perform work.
 - You can also add user credentials to your VM Instance to grant access to other users.

Note: In a newly created Instance, "cades" is the only user name that will correctly authenticate over SSH (using an SSH key pair). A user can add more users, and add public keys for each, once logged in as "cades" ([instructions](#)).

Long-form login

If you prefer to specify your key location during login (e.g., you downloaded your key from the Horizon web GUI), you can use a long-form SSH login.

1. Open a Bash terminal.
2. Execute `$ ssh -i ~/.ssh/blackmesakey.pem cadés@128.219.186.19`.
 - Replace `~/.ssh/blackmesakey.pem` with the file path of your key.
 - Replace `128.219.186.19` with the IP address of your own Instance.
3. You should now be connected to your VM Instance via SSH.
 - Use the Bash terminal to install your software-stack and perform work.
 - You can also add user credentials to your VM Instance to grant access to other users.

Note: In a newly created Instance, "cades" is the only user name that will correctly authenticate over SSH (using an SSH key pair). A user can add more users, and add public keys for each, once logged in as "cades" ([instructions](#)).

Extract Public Key

If you have several users, you can share your public key with them. New users must login using the username 'cades' to gain access to the instance, or you can [create another user account](#) for them. Once a user extracts the `.pub` file they can safely share with another OpenStack user. The receiving user can add that public key to the instance's `/home/cades/.ssh/authorized_keys` file.

To extract the public key from the private key:

```
openssl rsa -in privkey.pem -pubout > key.pub
```

In this way, through exchanging public keys, users can control who has SSH access to their instances.

Users should refrain from insecurely sending someone `.pem` files as they contains both public and private keys.

Related Tutorials

- [Add More Users to VM Instances](#)
- [Access VM Instances Using PuTTY \(Windows\)](#)
- [Access Your VM Instances Using Horizon](#)

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Access Your VM Instance](#) → [Access Your VM Instance Using PuTTY \(Windows\)](#)

Access Your VM Instance Using PuTTY (Windows)

Once you have created a VM Instance using Horizon, you can access this VM Instance through a secure shell (SSH) using an SSH key pair. For Windows users, the CADES team recommends [PuTTY](#). The [PuTTY MSI installer](#) also includes PuTTYGen, which allows you to convert a `.pem` private key file to a `.ppk` PuTTY key file.

If you would like to access your VM Instance over SSH from *outside* of the ORNL network, you must first [create a firewall \(SAFER\) exception](#) for this purpose. All local (within the ORNL network) SSH connections are permitted by default.

Table of Contents

- [Add SSH Keys to a VM Instance](#)
 - [Generate a New Key Pair](#)
 - [Import an Existing Key](#)
- [Download and Install PuTTY](#)
- [Connect to Your VM Instance Using PuTTY](#)
 - [Find Your Instance's IP](#)
 - [Connect to Your Instance](#)
- [Related Tutorials](#)

Add SSH Keys to a VM Instance

If you do not have any existing SSH key pairs, or you wish to generate a new key pair for your VM Instance, then you should use the Horizon web GUI to [generate a new SSH key pair](#) (Option 1). If you *do* have an existing key pair on your machine, then you can [import the public key using the Horizon web GUI](#) (Option 2).

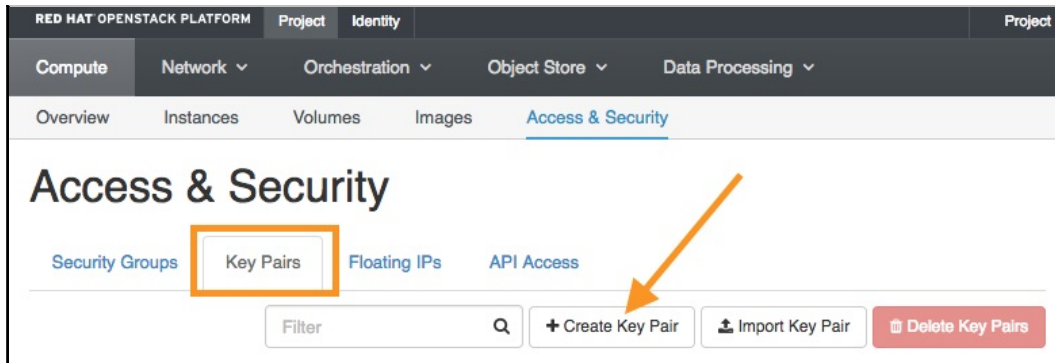
In either case, you will need to log into the Horizon web GUI to associate your SSH key with your VM Instance(s).

Log in to Horizon

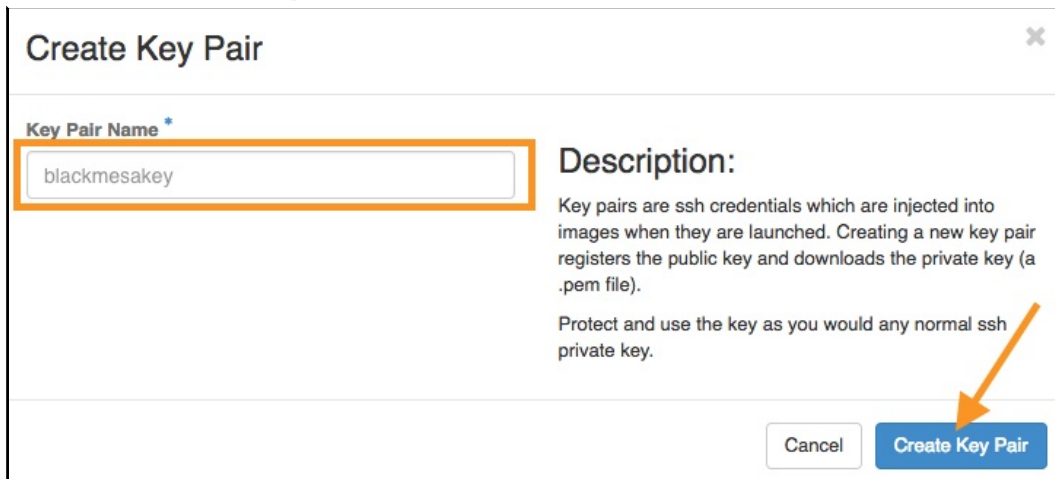
1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`

Option 1: Generate a New Key Pair

1. Within Horizon, navigate to `Project` → `Compute` → `Access & Security` → `Key Pairs`.
2. In the Key Pairs screen, click the `+ Create New Pair` button.



3. In the resulting dialog, name your Key Pair, and click `create key pair` when you're finished. We went with `blackmesakey` in this example.



4. The private key will be downloaded to your local machine as a `.pem` file. The public key will be stored in OpenStack.
5. Place the downloaded private key in a directory of your choosing, e.g., `C:\Users\Username\SSH`.


Note: The private key is in a `.pem` format and has to be converted to a `.ppk` file using PuTTYGen.

Having completed the procedure above, you can now [connect to your Instance using PuTTY](#).

Option 2: Import an Existing Key

If you have an *existing* key pair that you would like to also use for your OpenStack Project, then you can import the public key using the Horizon web GUI.

1. Within Horizon, navigate to `Project` → `Compute` → `Access & Security` → `Key Pairs`.
2. In the Key Pairs screen, click the `Import Key Pair` button.



3. In the resulting dialog, enter a key pair name in the `Key Pair Name` field.
4. On your local machine, locate your public key file and open it with a text editor.
5. Copy the key to your clipboard using the text editor.
6. Back in Horizon, paste the public key into the `Public Key` text box.
7. Click the `Import Key Pair` button.

Import Key Pair

Key Pair Name *
blackmesakey

Public Key *

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDKV3
oXninp+GJis2v8PofB6sSz9Dwt1lhBeFVFFYsM6H
HgTTsZlvFFoGrfjBHctL2FIPyqG2cG36gHOgph8wj
3rY88jQBUq5uat6z88jDJ1KBZFLONMvH3Wi0+qN
nkwLrr9pHdlvuKP5x5xswFwk+zmtPeVaBNIS++1
6Li2Rak9ydlSOyBzXkCqgP5CbzipahUsYVP4Ts3K
9l43w4Q+HIZvUfMPwgzM+68T7MOOPQ2/bNgci
qmqYNUOTzPloS7u1e3VyJ6uxWBHdvFhzx1u5D
SxFsQdHjFtdkpkQdDpJZFWONXHTpgJs5/yejfQbC
5P...L...P...5...3...5...
```

Description:

Key Pairs are how you login to your instance after it is launched.

Choose a key pair name you will recognise and paste your SSH public key into the space provided.

SSH key pairs can be generated with the ssh-keygen command:

```
ssh-keygen -t rsa -f cloud.key
```

This generates a pair of keys: a key you keep private (cloud.key) and a public key (cloud.key.pub). Paste the contents of the public key file here.

After launching an instance, you login using the private key (the username might be different depending on the image you launched):

```
ssh -i cloud.key <username>@<instance_ip>
```

Cancel Import Key Pair

8. Confirm your newly imported key appears in the Key Pair list.

Having completed the procedure above, you can now [connect to your Instance using PuTTY](#).

Download and Install PuTTY

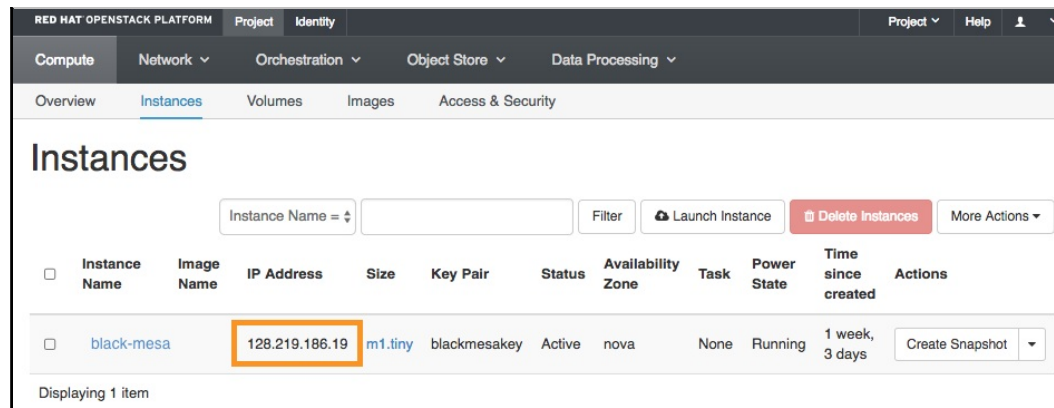
1. Navigate to the [official PuTTY download page](#).
2. Download the appropriate MSI file for your Windows system (32 bit or 64 bit).
 - Alternatively, if you don't want to "install" PuTTY on your system, you can download `putty.exe` and `puttygen.exe` from the "Alternative binary files" list and execute them as needed.
3. Run the PuTTY MSI installer. Note the destination directory.
4. Confirm installation of PuTTY executables.

Connect to Your VM Instance Using PuTTY

Find Your Instance's IP

The IP address of your VM instance is shown in the Instances tab in the Horizon web GUI.

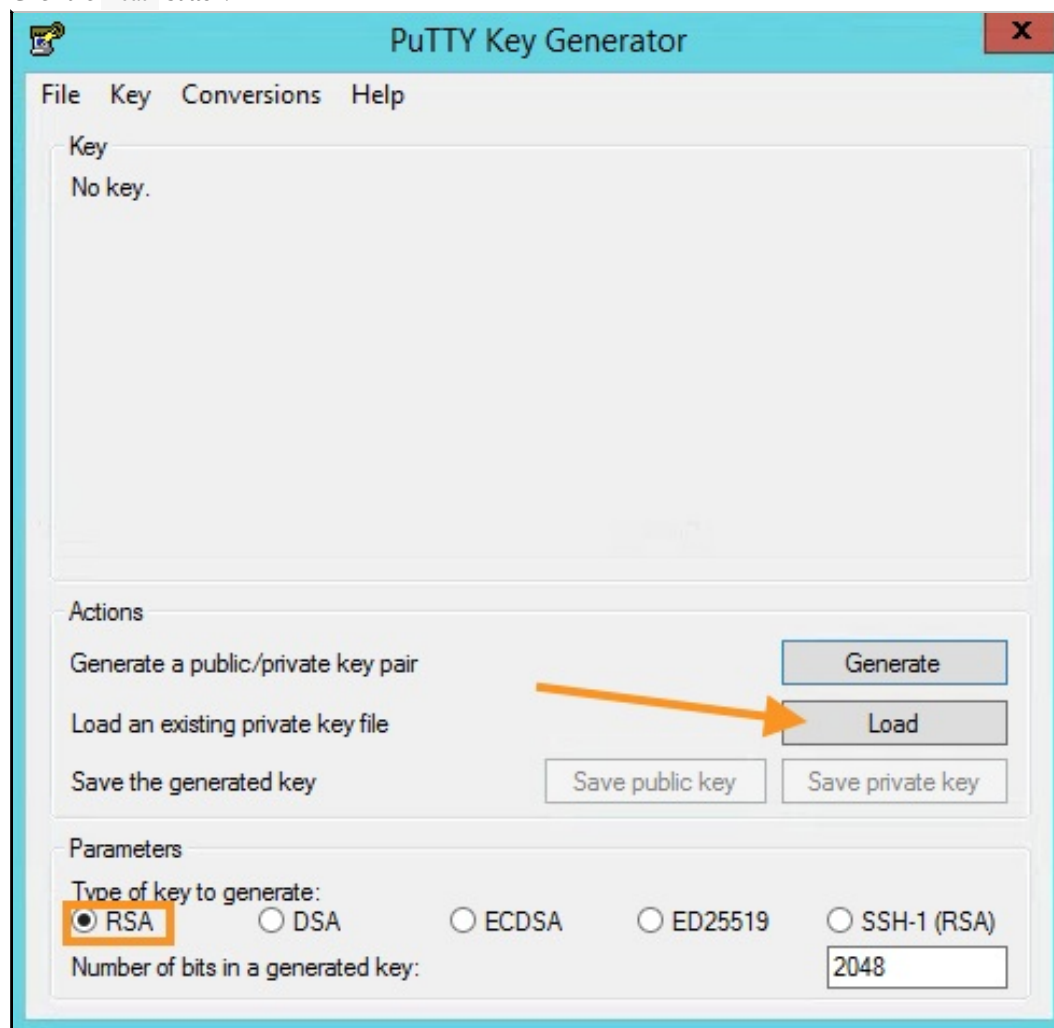
1. Navigate to `Project` → `Compute` → `Instances`.
2. From the Instances screen, you can copy the IP address for the desired VM Instance.




Convert Your Private Key to PuTTY Format

If you generated a key in the `.pem` file format, then you need to convert that key to the `.ppk` format using PuTTYGen.

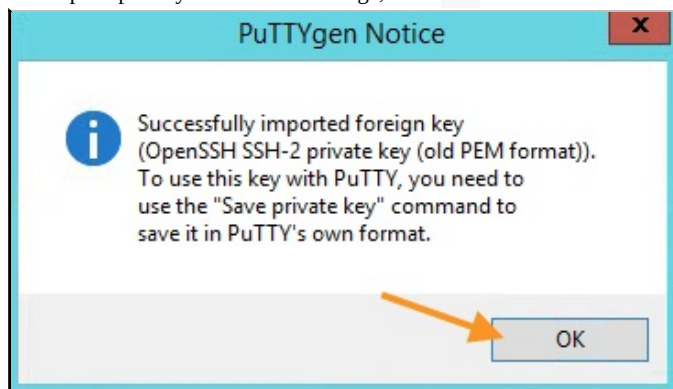
1. Locate and execute `PuTTYGen.exe`.
2. Click the `Load` button.



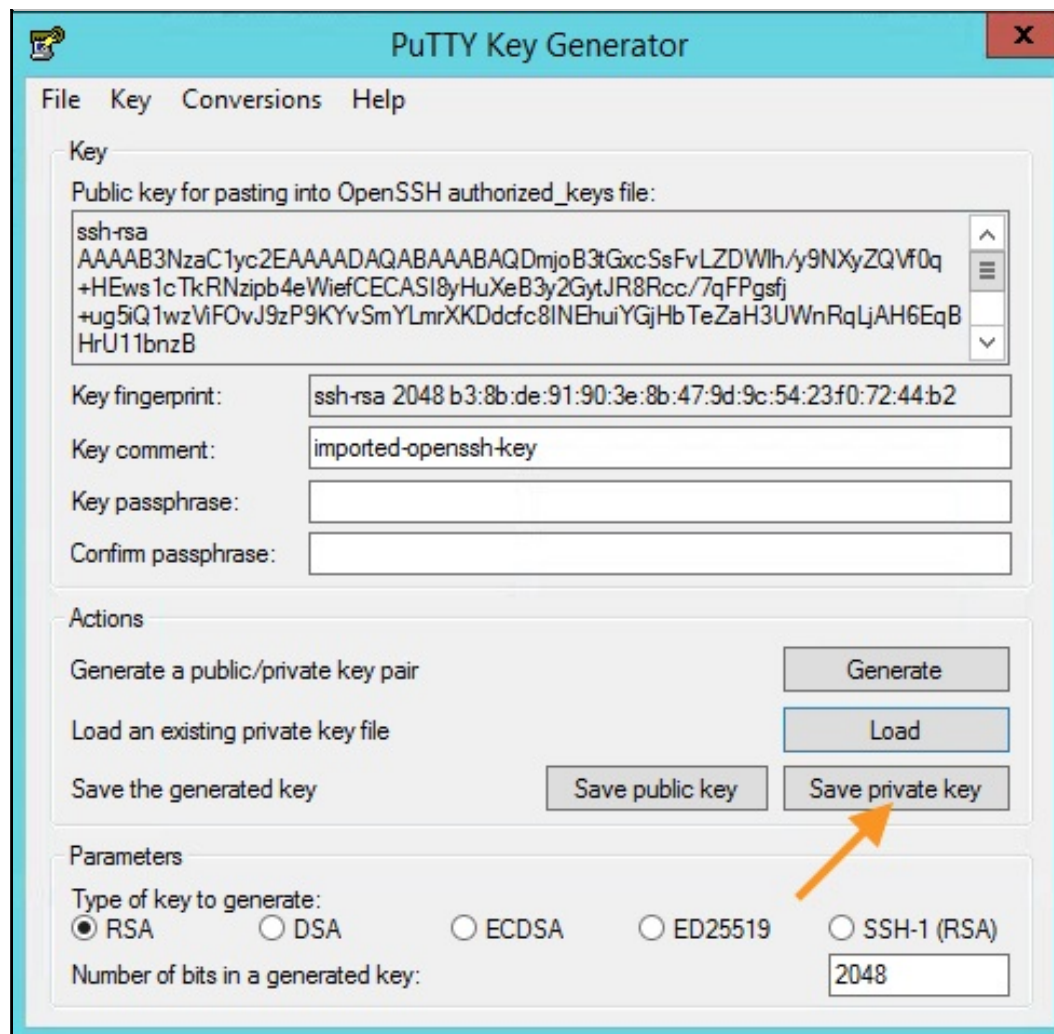
3. In the resulting dialog, navigate to the location of your key file.
4. Select `All Files (*.*)` to show `.pem` files in the dialog.
5. Select the appropriate key file and click `open`.



6. When prompted by the success message, click `ok` .



7. Save the loaded key by clicking `save private key` .



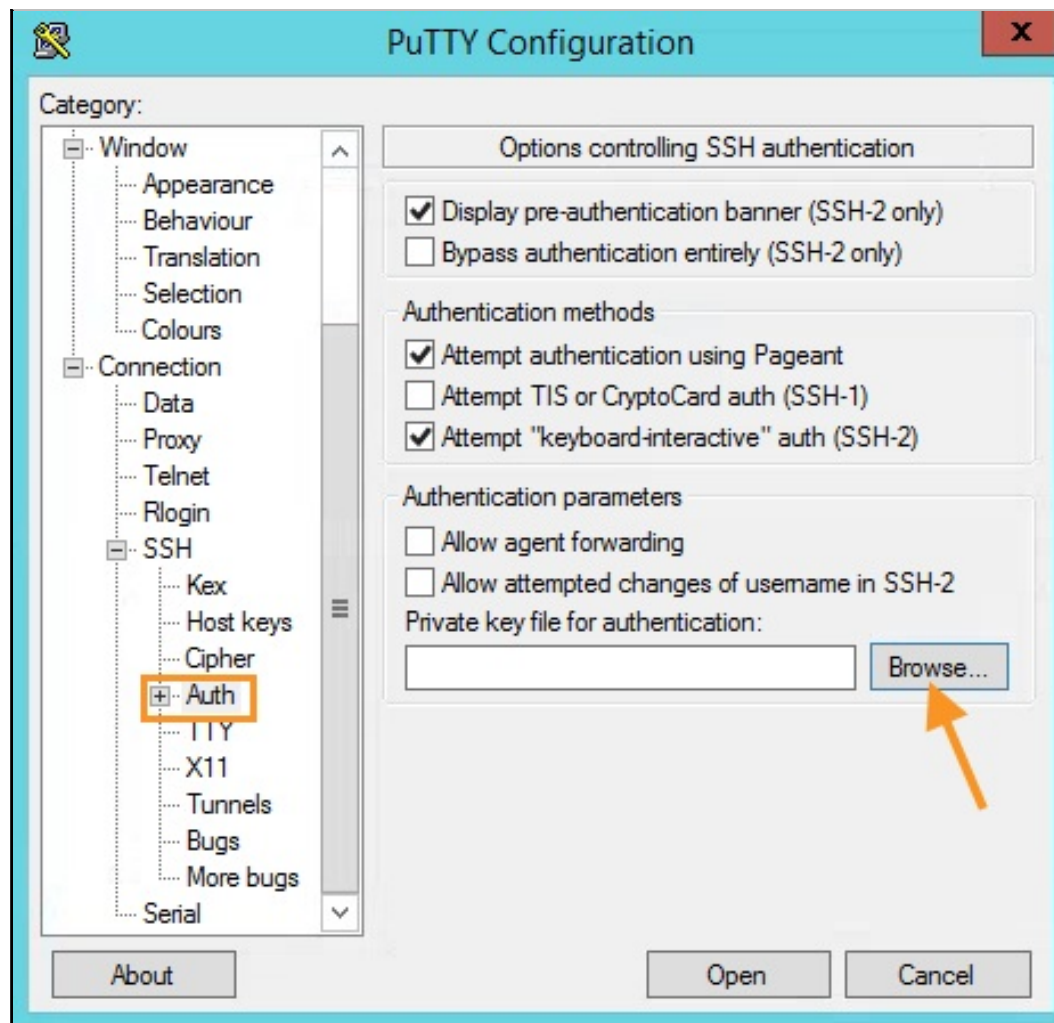
8. Using the resulting dialog, save the key in an intuitive place (e.g., `C:\Users\Username\SSH`).

Your key is now saved as a `.ppk` file that can be loaded directly into PuTTY, and we can now use PuTTY to connect to your VM Instance, without a password, using the "cades" username.

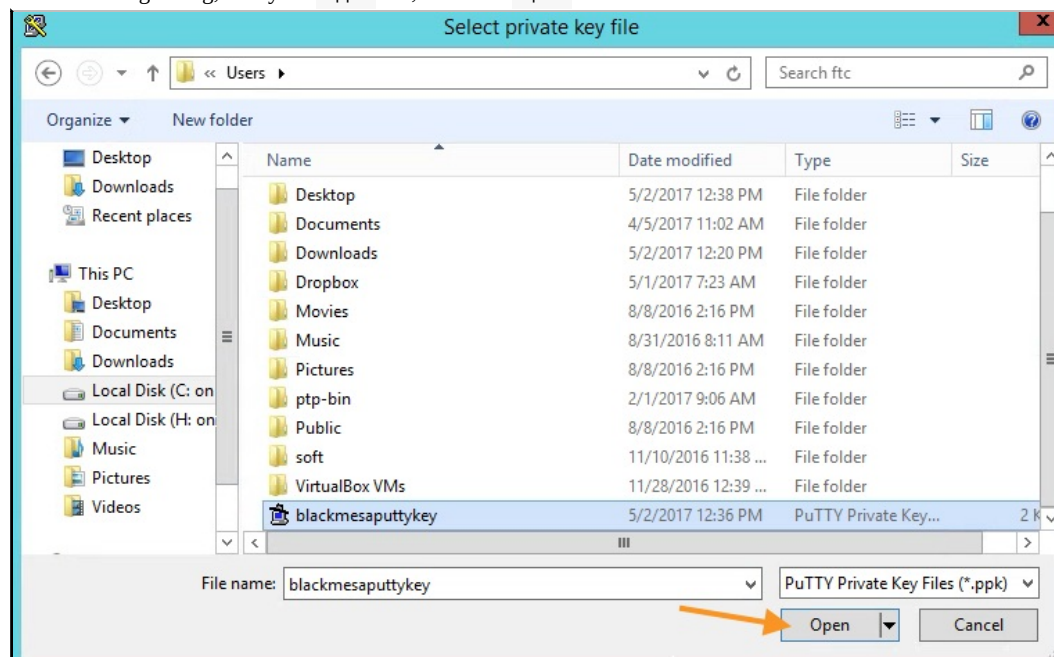
Connect to Your Instance

First we're going to load the `.ppk` key file into PuTTY.

1. Locate and execute `PuTTY.exe`.
2. In PuTTY's navigation pane, go to `Connection` → `SSH` → `Auth`.
3. In the "Options for controlling SSH authentication" screen, click `Browse`.



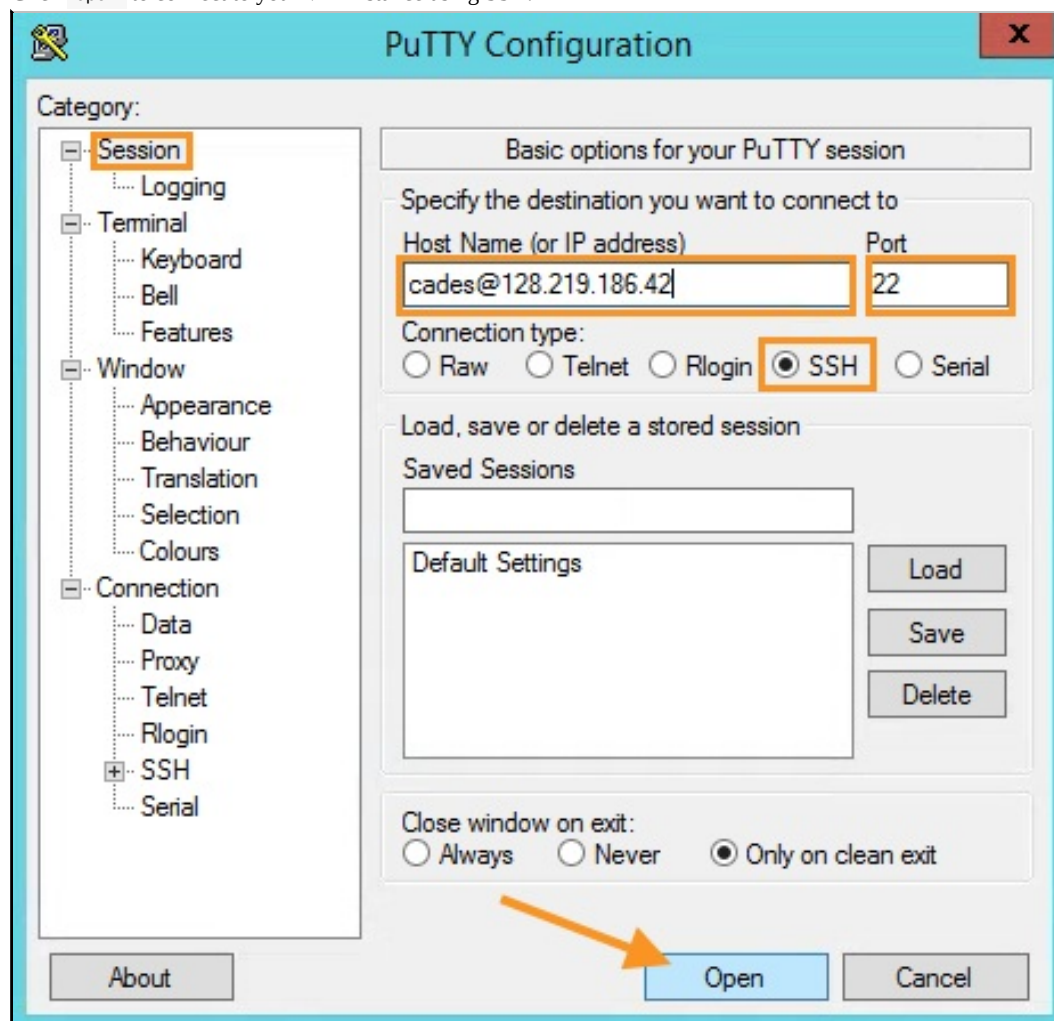
4. In the resulting dialog, find your `.ppk` file, and click `open`.



Your key file is now loaded into PuTTY. Next, we need to add your host information.

1. In PuTTY's navigation pane, click `Session`.

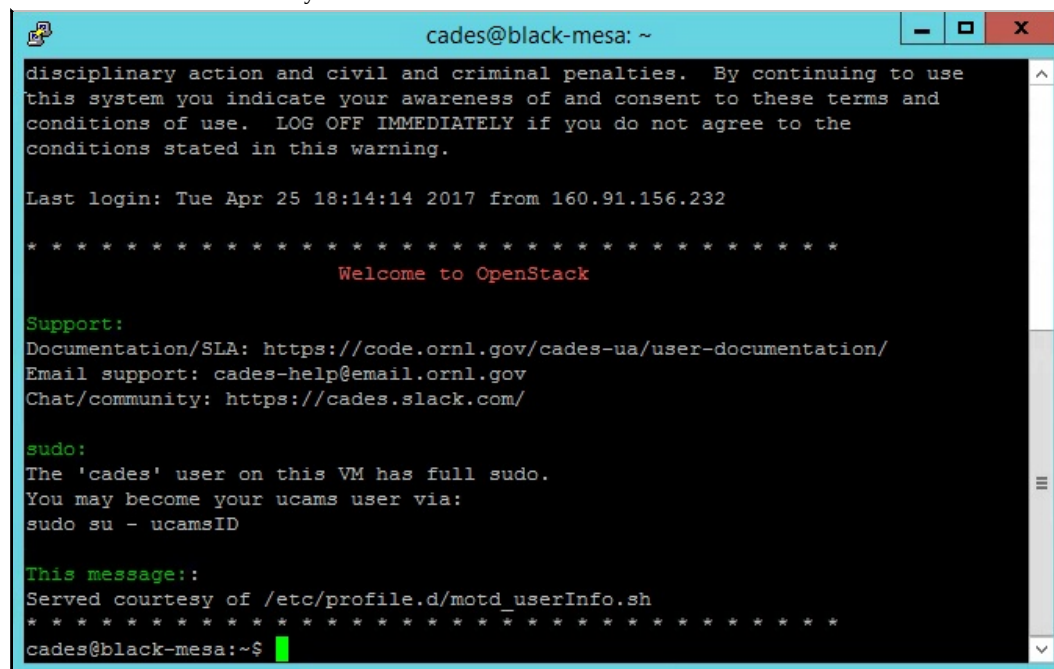
2. In the `Host Name (or IP address)` field, enter the IP address of your VM Instance preceded by "cades" (e.g., `cades@128.219.186.42`).
3. Ensure that the `SSH` radio button is selected and the port is set to `22`.
4. Click `open` to connect to your VM Instance using SSH.



5. If prompted, you can choose to cache the server's fingerprint (click `yes`), not to cache the server's fingerprint for this session (click `no`), or to reject the connection outright (click `cancel`). We're going to click `yes`.



6. You should now have access to your VM Instance's Bash terminal.

A screenshot of a Bash terminal window. The title bar shows "cades@black-mesa: ~" with standard window controls. The terminal content includes a warning about disciplinary action, a login timestamp "Last login: Tue Apr 25 18:14:14 2017 from 160.91.156.232", a "Welcome to OpenStack" message, support information (Documentation/SLA, Email support, Chat/community), and sudo instructions: "The 'cades' user on this VM has full sudo. You may become your ucams user via: sudo su - ucamsID". It also shows "This message: Served courtesy of /etc/profile.d/motd_userInfo.sh" and ends with the prompt "cades@black-mesa:~\$" and a green cursor.

Related Tutorials

- [Add More Users to VM Instances](#)
- [Access Your VM Instances Using Horizon](#)
- [Access Your VM Instances Using SSH](#)

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Access Your VM Instance](#) → [Access Your VM Instance Using Horizon](#)

Access Your VM Instance Using Horizon

The Horizon web GUI also has a built-in console from which you can access your VM Instance once you have [added a generic user over SSH](#). Note that this is not the preferred method, and that the CADES team recommends using [the traditional SSH connection for accessing your VM Instance](#).

Prerequisites

The Horizon console uses a simple username/password authentication model, and cannot leverage SSH key pairs for authentication. This means that you cannot use the console with the "cades" username or your UCAMS credentials.

As a result, to access your VM Instance using the Horizon console, you must first use SSH to create a non-UCAMS user. More info on creating users is available [here](#).

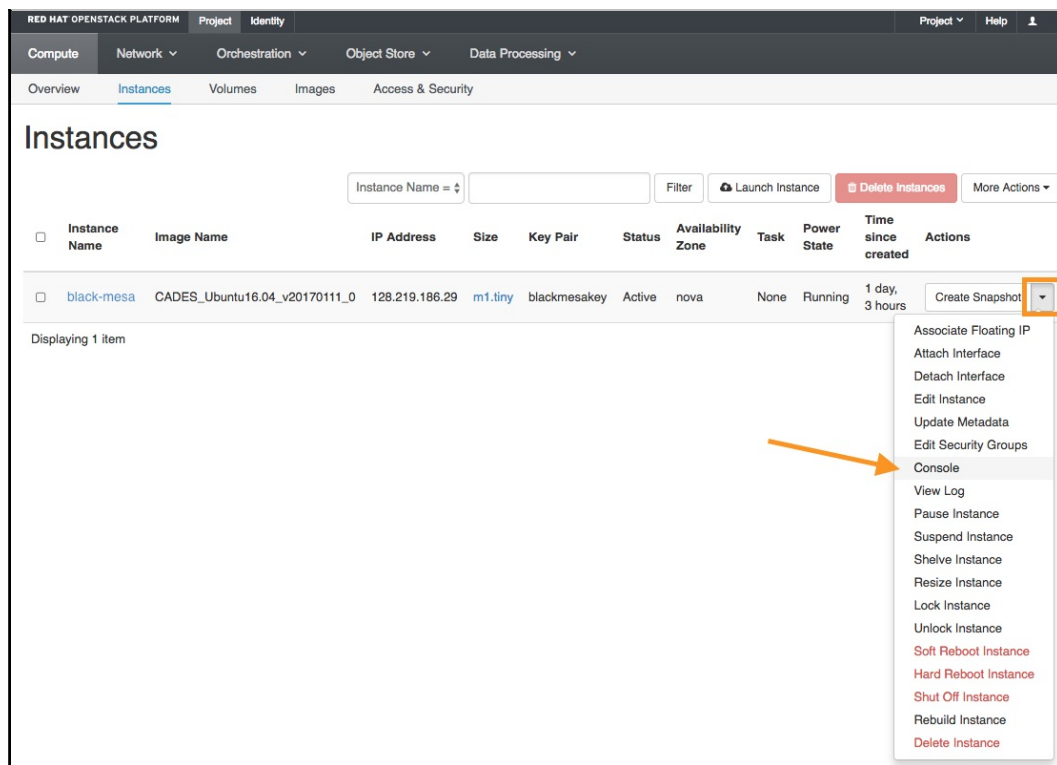
Log in to Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`

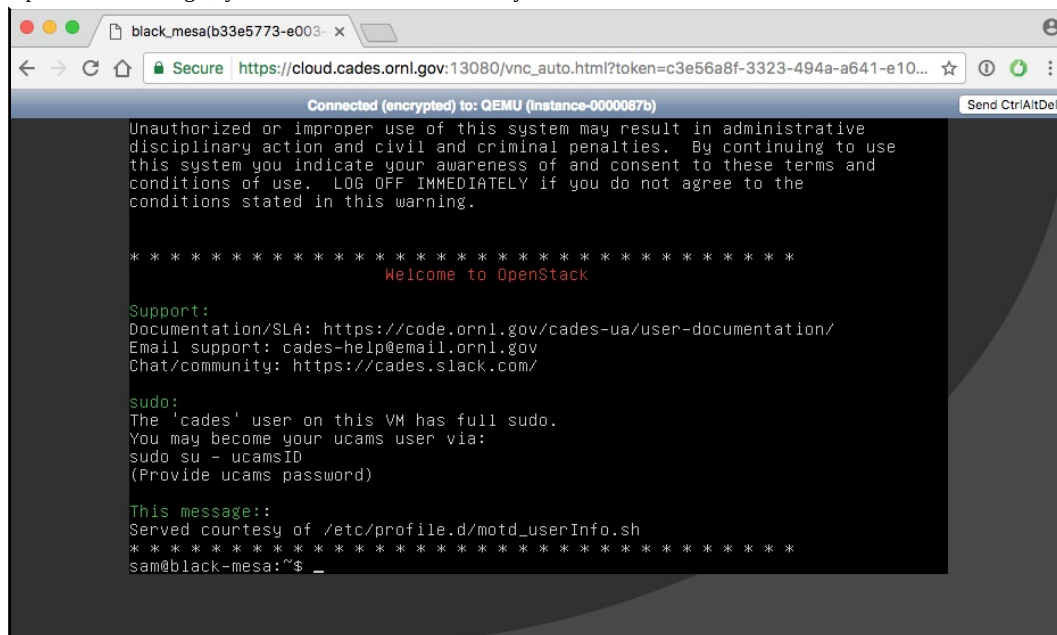
We can now use Horizon to access your VM Instances using the Console option.

Launch Console

1. Navigate to `Project` → `Compute` → `Instances` .
2. Click the `drop-down menu` of the Instance you would like to access.
3. Select `console` from the resulting menu.



4. In the console, input your user credentials.
5. Upon successful login, you can execute commands on your VM Instance.



Note: To use the Horizon console, you must first [add a user and password](#) to the Instance's operating system (via SSH) to enable access to the VM without an SSH key.

Related Tutorials

- [Add More Users to VM Instances](#)
- [Access Your VM Instances Using SSH](#)
- [Access Your VM Instances Using PuTTY \(Windows\)](#)

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Access Your VM Instance](#) → [Add More Users to Your VM Instance](#)

Add More Users to Your VM Instance

Once you have gained [SSH access](#) to your Instance (as "cades"), additional users can be created through the Bash terminal. The CADES OpenStack Instances leverage ORNL's UCAMS system to enable customers to create UCAMS users and home areas on their Instances. The following procedures will enable you to create a new user on your VM Instance and grant access using SSH key pairs.

Add a UCAMS User

Because your Instance can access ORNL's UCAMS system, you can easily add a UCAMS user to your Instance. You can start with your own UCAMS ID to verify functionality and then move on to adding the UCAMS IDs of your collaborators (using their public SSH keys that they must provide to you).

1. Open a Bash terminal.
2. Upload a copy of your (or your collaborator's) public key (`blackmesakey.pub` in this case), `$ scp ~/.ssh/blackmesakey.pub cades@128.219.186.39:/home/cades` . Replace `128.219.186.39` with the IP address of your Instance.
3. Using SSH, log in to your VM Instance as "cades".
4. Once logged in, become root, `$ sudo -s` .
5. Create the home directory for the UCAMS ID, `$ su - UCAMS` . Replace "UCAMS" with the desired UCAMS ID.
6. Use `cat` and `>>` to copy the contents of your public SSH key to `/home/UCAMS/.ssh/authorized_keys` .

SSH commands for the above procedure are provided below. First:

```
$ scp ~/.ssh/blackmesakey.pub cades@128.219.186.39:/home/cades
```

Then:

```
$ ssh cades@128.219.186.39
$ sudo -s
$ su - UCAMS
$ cat /home/cades/blackmesakey.pub >> /home/UCAMS/.ssh/authorized_keys
```

Properly configured, you can now access your VM Instance over SSH using your UCAMS ID (instead of "cades"). Execute the command shown below (replace "UCAMS" with your own UCAMS ID).

```
$ ssh UCAMS@128.219.186.39
```

Note: You will not be prompted for your UCAMS password. This process uses SSH key pairs for authentication.

Add a Non-UCAMS User

In some cases, it may be useful to have a generic user (not affiliated with UCAMS) on your VM Instance. The following procedure outlines how to create such a user.

1. Open a Bash terminal.
2. Upload a copy of your (or your collaborator's) public key (`blackmesakey.pub` in this case), `$ scp ~/.ssh/blackmesakey.pub cades@128.219.186.39:/home/cades` . Replace `128.219.186.39` with the IP address of your

Instance.

3. Using SSH, log in to your VM Instance as "cades".
4. Once logged in, become root, `$ sudo -s`.
5. Add a new user, `useradd USERNAME`.
6. Create a password for this user, `passwd USERNAME`.
 - o Enter desired password. Reenter password to confirm.
7. Create the home directory for the generic user, `$ su - USERNAME`. Replace "USERNAME" with the desired user name.
8. Use `cat` and `>>` to copy the contents of your public SSH key to `/home/USERNAME/.ssh/authorized_keys`. Replace "USERNAME" with the desired user name.

SSH commands for the above procedure are provided below. First:

```
$ scp ~/.ssh/blackmesakey.pub cades@128.219.186.39:/home/cades
```

Then (replacing "USERNAME" with your desired username):

```
$ ssh cades@128.219.186.39
$ sudo -s
$ useradd USERNAME
$ passwd USERNAME
$ su - USERNAME
$ exit

$ sudo -s
$ cat /home/cades/blackmesakey.pub >> /home/USERNAME/.ssh/authorized_keys
```

Properly configured, you can now access your VM Instance over SSH using your newly created username (instead of "cades"). Execute the command shown below (replace "USERNAME" with your new username).

```
$ ssh USERNAME@128.219.186.39
```

See the [Access Your VM Instance documentation](#) for instructions on how to use SSH and Horizon to access your Instance.

Related Tutorials

- [Access Your VM Instances Using SSH](#)
- [Access Your VM Instances Using Horizon](#)

Manage Your VM Instances

The Horizon web GUI enables users to view and manage their VM Instances directly from a web browser. Using the following procedure, users can manage existing instances and create new ones through the Instances management screen.

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: ornl
 - User Name: UCAMS ID
 - Password: UCAMS password
3. Click on the `Project` tab on the top left.
4. Select the `compute` sub tab.
5. Finally, click on the `Instances` sub tab.

From the Instances screen, you can:

- [Launch a new VM Instance from an Image](#)
- [Delete old VM Instances](#)
- [Delete a Volume](#)
- [Resize a VM](#)
- [Create a Snapshot](#)
- [Modify Security Groups](#)

Delete a VM Instance from your Project

At some point you may wish to delete a VM Instance from your project. This can help a user free up space in his or her allocation or clear out old Instances that have outlived their usefulness. Deleting a VM Instance is a simple process, outlined in the procedure below.

Log in to Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - o Domain: `ornl`
 - o User Name: `Your three-letter UCAMS ID`
 - o Password: `Your UCAMS password`

From here, you can manage your VM Instance(s) from within Horizon.

Delete an Instance

1. Navigate to `Project` → `Compute` → `Instances` .
2. Click the drop-down menu of the Instance you wish to delete.
3. Select `Delete Instance` from the resulting menu.

The screenshot shows the Horizon web interface with the 'Instances' page selected. The page has a navigation bar with tabs for 'Overview', 'Instances', 'Volumes', 'Images', and 'Access & Security'. Below the navigation bar, there is a search bar for 'Instance Name' and buttons for 'Filter', 'Launch Instance', 'Delete Instances', and 'More Actions'. The main content area displays a table of instances:

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|-------------------------------------|-------------------------------|----------------|---------|--------------|--------|-------------------|------|-------------|---------------------|--|
| <input type="checkbox"/> glados | CADES_CentOS-7.3_v20170310_0 | 128.219.186.31 | m1.tiny | blackmesakey | Active | nova | None | Running | 5 minutes | Create Snapshot |
| <input type="checkbox"/> black-mesa | CADES_Ubuntu16.04_v20170111_0 | 128.219.186.29 | m1.tiny | blackmesakey | Active | nova | None | Running | 6 hours, 18 minutes | Associate Floating IP Attach Interface Detach Interface Edit Instance Update Metadata Edit Security Groups Console View Log Pause Instance Suspend Instance Shelve Instance Resize Instance Lock Instance Unlock Instance Soft Reboot Instance Hard Reboot Instance Shut Off Instance Rebuild Instance Delete Instance |

Below the table, it says 'Displaying 2 items'. An orange arrow points to the 'Delete Instance' option in the context menu for the 'black-mesa' instance.

4. Click `Delete Instance` in the dialog.
5. Confirm that the Instance is scheduled for deletion.

Note: Deleting a VM Instance does not delete the volume associated with the Instance. To remove both, you must follow the [Delete a VM Instance](#) procedure *and* the [Delete a Volume](#) procedure.

Related Tutorials

- [Delete a Volume](#)
- [Launch a VM Instance](#)

Delete a Volume from Your Project

If you decide to delete a VM Instance, you may still have its associated Volume left behind (i.e., `Delete Volume on Instance` `delete` was set to `No` when you launched the Instance). Use the following procedure if you wish to delete a Volume.

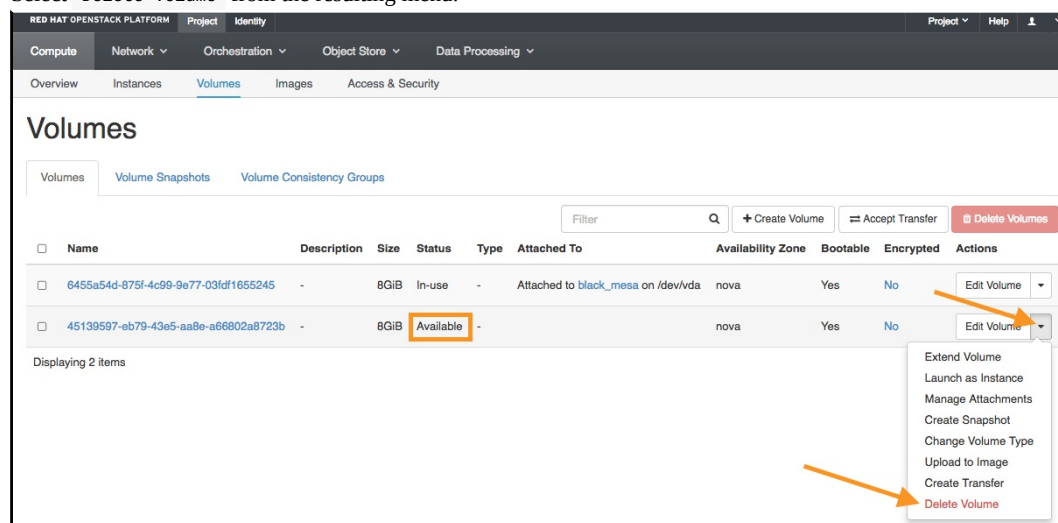
Log in to Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - o Domain: `ornl`
 - o User Name: Your three-letter UCAMS ID
 - o Password: Your UCAMS password

From here, you can delete the Volume from your Project.

Delete a Volume

1. Navigate to `Project` → `Compute` → `Volumes`.
2. In the Volumes screen, ensure that the Volume is not in use by any VM Instance.
3. Click the drop-down menu of the volume you wish to delete.
4. Select `Delete Volume` from the resulting menu.



5. Click `Delete Volume` in the resulting dialog.
6. Confirm that the Volume is no longer listed in the Volumes screen.

Related Tutorials

- [Delete a VM Instance](#)
- [Launch a VM Instance](#)

Resize an Existing VM Instance

Sometimes users may wish to add CPU or Memory resources to an existing VM instance. This can be accomplished easily through the Horizon interface:

1. Navigate to your Instances List at <https://cloud.cades.ornl.gov/dashboard/project/instances/>.
2. Under **Actions** on the right-hand side, select **Resize Instance** for the instance you wish to resize.
3. You will be prompted to select a new **Flavor**.
4. As long as the new flavor fits in your **Allocation**, click **resize** in the bottom right corner and your instance will begin resizing.
5. Your instance will reboot into the new flavor. Once this is complete, click **Confirm Resize or Migration**.

At this point, your instance will be up and running at the size you selected.

Increasing the size of a root volume *may* require [emailing CADES support](#). This might apply to your volume if you did *not* select an ephemeral volume.

CADES → User Documentation → CADES Cloud User Guide → Manage Your VM Instances → Add a Volume to a VM Instance

Add a Volume to an Existing VM Instance

1. Navigate to your Volumes List at <https://cloud.cades.ornl.gov/dashboard/project/volumes/>.
2. At the top right click on `Create Volume`.
3. Fill in the following fields:
 - Name: `user choice`
 - Description: optional and can be left blank
 - Source: No source, empty volume (*this is default*)
 - Type: `Netapp` (*this is default*)
 - Size: Size you need up to your quota, which is displayed on the right
 - Availability Zone: `nova` (*this is default*)
4. Click `Create Volume`.
5. Once created use the drop down () on the right of the volume you just created and choose `Manage Attachments`.
6. Select the instance from the drop down to which you would like to attach your new volume.
7. Next, the volume will have to be partitioned, formatted, and mounted. To begin, [SSH into your VM](#).
8. Check that the volume was assigned by listing the available disks: type `lsblk`. You should see a new disk with the allotted storage amount in the listed output. Example output is shown below.

```
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 8G 0 disk
└─vda1 253:1 0 8G 0 part /
vdb 253:16 0 8G 0 disk
```

- In this case, the disk `vda` is the original disk that has a partition named `vda1`. You may want to create a partition in `vdb` (optional).
9. *Optional:* To create partition of the new disk, type `sudo fdisk /dev/vdb`. This command expects you to enter additional information.
 - To create a new partition, enter `n` then press `Enter`.
 - Choose `p` for primary or `e` for an extended partition. Usually, you will choose `e` then press `Enter`.
 - The next two prompts request space allocations. In typical situations, press `Enter` on these two prompts to select the defaults. *If you would like more information about your options, check the "SIZES" section in the fdisk manual by typing `man fdisk`.*
 - Type `w` and then press `Enter` to write your changes and reboot the system.
 - You will be returned to your bash prompt.
 - To check that the partition was created correctly, type `lsblk` and you should see the new partition labelled `vdb1` (in this example case).

```
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 8G 0 disk
└─vda1 253:1 0 8G 0 part /
vdb 253:16 0 8G 0 disk
└─vdb1 253:17 0 8G 0 part
```

10. Next, the new volume/partition must be formatted. Type `sudo mkfs.ext4 /dev/vdb1`.

Note: OpenStack and/or `fdisk` may impart a default filesystem type on the new volume. In this case, you will be asked if you would like to proceed with the formatting although an existing partition table exists. If you choose to proceed by typing `y`, you will rewrite the partition. **DO NOT PROCEED** if this partition contains data that you need to keep!

11. To check the formatting, type `lsblk -f`. You should see that the `vdb1` entry has the `ext4` filesystem type.
12. To mount the volume for use, you need to create the mountpoint. For example: `mkdir volume`.

13. Then, to mount the volume to the newly-created mountpoint, type `sudo mount /dev/vdb1/ ./volume`.
14. Lastly, to check the mountpoint, type `lsblk -f` and you should see `vdb` is mounted at the location you chose:
`/home/cades/volume`.
15. To avoid having to mount the volume every time you boot your VM Instance, you may set up automounting by viewing the contents of `/etc/fstab`.

- Type `sudo nano /etc/fstab`.
- Check for an existing line of code:

```
/dev/vdb /data auto defaults,nobootwait 0 2
```

- If it is not present, add it. Save your changes.
16. The new storage volume is ready to use.

Create a Snapshot of a VM Instance

A VM snapshot is an instantaneous duplicate of an instance. A snapshot of an instance can be used to back-up data, to create a restore point, or as the basis of an instance and booted up at a later time. A snapshot can be booted as a new VM Instance and contains an image of the state of the filesystem at the moment that the snapshot is taken.

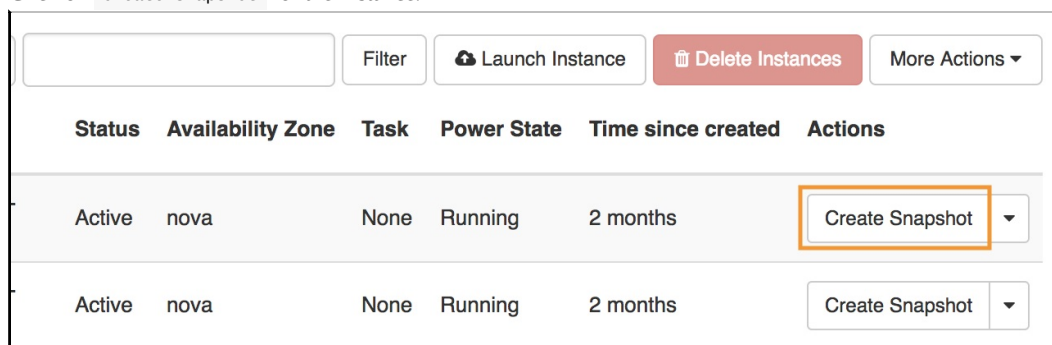
Log in to Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`

From here, you can manage your VM Instance(s) from within Horizon.

Create a Snapshot

1. Navigate to `Project` → `Compute` → `Instances` .
2. Click on `Create Snapshot` for the Instance.



| Status | Availability Zone | Task | Power State | Time since created | Actions |
|--------|-------------------|------|-------------|--------------------|-----------------------------------|
| Active | nova | None | Running | 2 months | Create Snapshot ▼ |
| Active | nova | None | Running | 2 months | Create Snapshot ▼ |

3. In the window that appears, choose a descriptive name for your snapshot and then click `Create Snapshot` .
4. Once the snapshot is created, a list of snapshot will appear. You can later navigate to this list by navigating to `Project` → `Compute` → `Images` .
5. From the `Images` screen you can launch, edit, or delete the snapshot.

Related Tutorials

- [Delete a VM Instance](#)
- [Launch a VM Instance](#)

OpenStack Security Groups

At their core, the OpenStack Security Groups are iptable-based firewalls built around an Instance at the hypervisor level. The Security Groups can be used in conjunction with the OS-level firewalls (e.g., FirewallD, iptables) but do not overlap with them (see [Important Notes](#)).

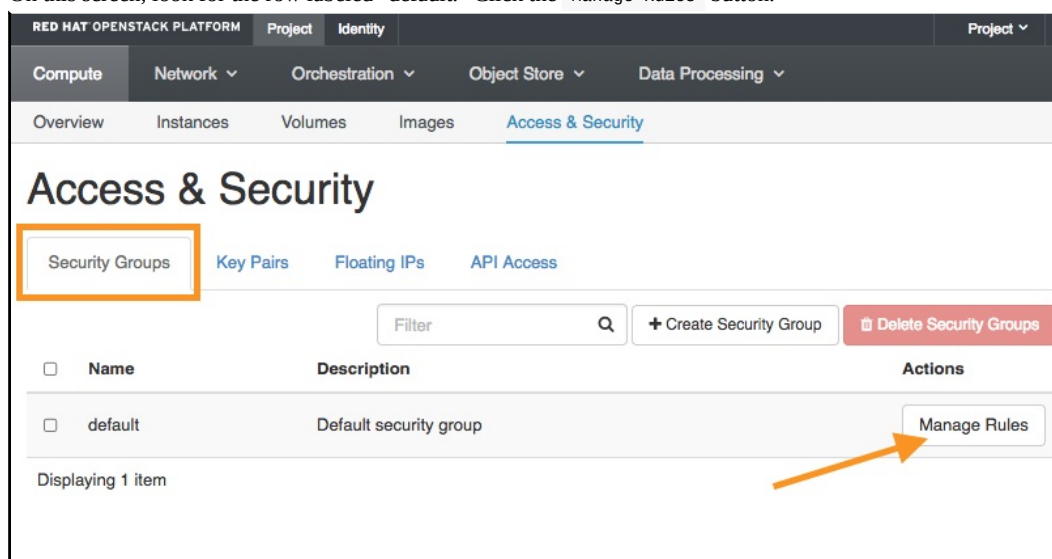
Important Notes for OpenStack Security Groups

- IPV6 is not currently supported in OpenStack.
- Changes to Security Groups take effect immediately.
- Unlike normal Linux firewall rules, the rule order does not matter in OpenStack Security Groups.
- By default, all Instances within the same Project can communicate with each other.
- Using 160.91.8.218:6556 to access ORNL's Check_MK service is allowed but not enabled by default. For monitoring of uptime and basic metrics, please contact the [CADES team](#) for assistance.
- No firewall is enabled in the CADES-provided operating system (OS) images. Instead, we rely on the OpenStack Security Groups. The user is responsible for enabling and configuring extra OS-level firewall rules as desired.
- User-added firewall and iptable rules supersede rules set in OpenStack Security Groups. For example, ingress access enabled by a rule in the OpenStack Security Group that are otherwise blocked at the OS level using the firewall or iptables will be ineffective, and that traffic will still be blocked.
- By default, all newly created Security Groups allow all outbound IPV4 and IPV6 (enabled but not functional). By default, no inbound traffic is allowed.
- The CADES team recommends that users leave existing Security Group rules in place as many of these rules are used by the CADES support team (e.g., for monitoring and metrics).

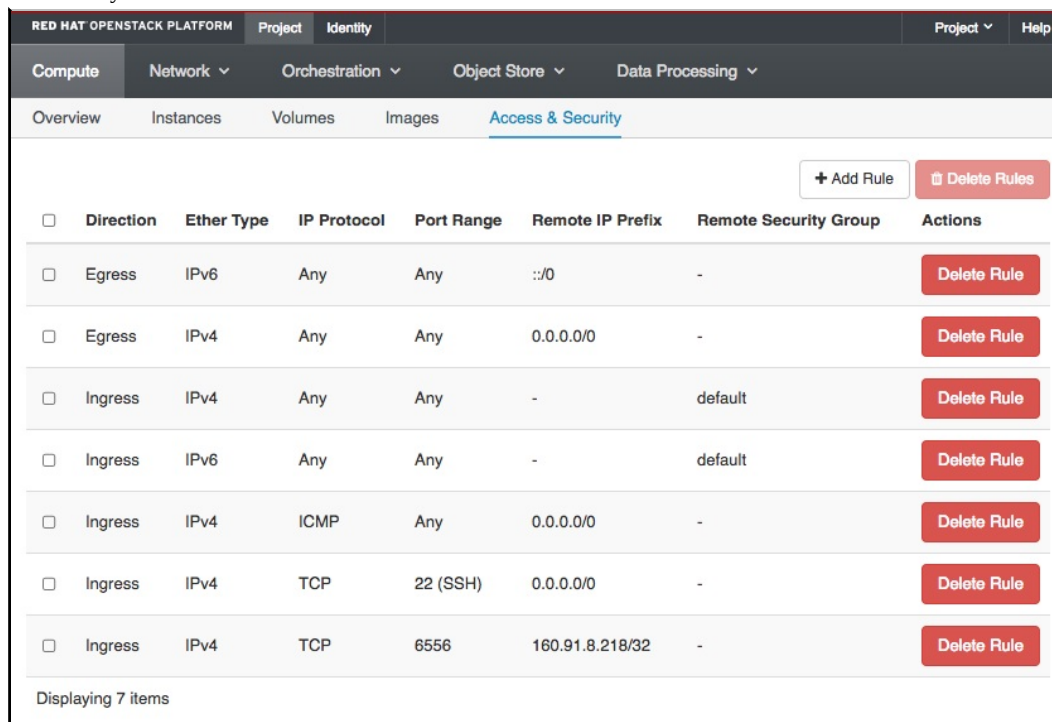
CADES → User Documentation → CADES Cloud User Guide → Manage Your VM Instances → OpenStack Security Groups → Modify Security Groups

Modify the Default Security Group

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: ornl
 - User Name: Your three-letter UCAMS ID
 - Password: Your UCAMS password
3. Navigate to Project → Compute → Access & Security → Security Groups .
4. On this screen, look for the row labeled "default." Click the Manage Rules button.



From here you can create a new rule or remove rules.

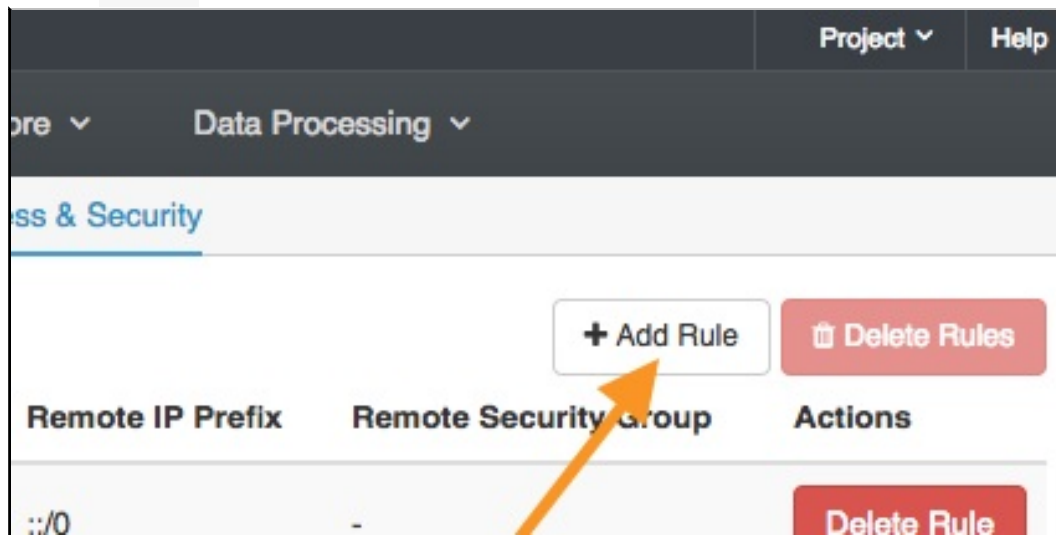


Create a New Rule

You can add a new rule to your Security Group using the built-in rules provided in OpenStack or you can create your own custom rule.

Create a Rule Using the Built-in Rules

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups` .
2. Click the `Add Rule` button.

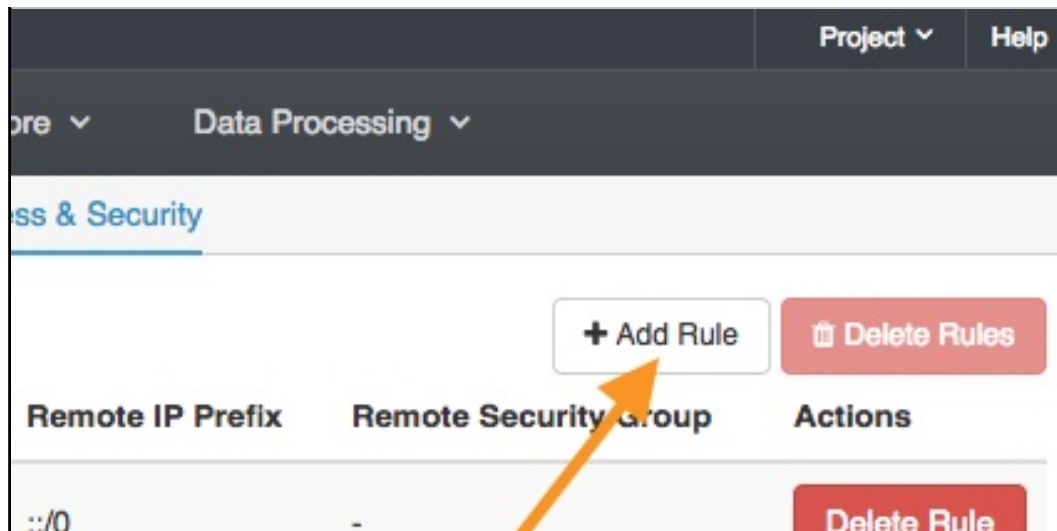


3. In the resulting dialog, click the drop-down field under `Rule` .
4. Choose a rule from the list that fits your needs (DNS, HTTP, HTTPS, etc.).
5. In the `Remote` box directly under `Rule` , choose either `CIDR` or `Security Group` .
 - If you selected `CIDR` , enter the desired inter-domain range in the `CIDR` box directly under the `Remote` box. See [CIDR examples](#).
 - If you select `Security Group` , choose a security group shared by the Instance(s) with which you wish to communicate. You can also do this via their IP addresses using the `CIDR` option.

Note: this option only allows access to the Instances within that security group. This differs from CIDR.
6. Click `Add` at the bottom of the dialog box to implement your rule.
7. Confirm the new rule is displayed in the Manage Rules screen.

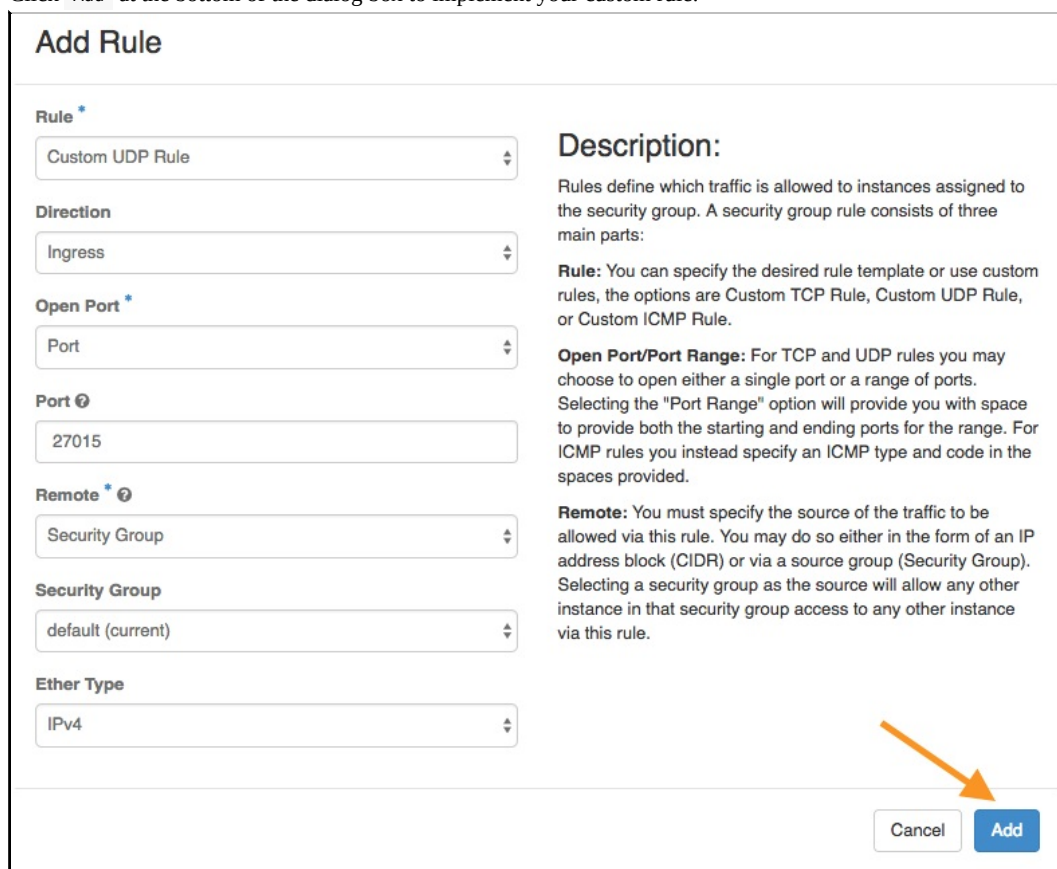
Create a Custom Rule

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups` .
2. Click the `Add Rule` button.



3. In the resulting dialog, click the drop-down field under `Rule`.
4. Choose the rule type from the drop-down list (e.g., `Custom TCP|ICMP|UDP Rule`).
5. Set the preferred direction in the `Direction` field (i.e., `Ingress` or `Egress`).
6. Choose either a single port or a range of ports in the `Open Port` section.
7. Enter the port or port range in the respective field.
8. In the `Remote` box, choose either `CIDR` or `Security Group`.
 - o If you select `CIDR`, enter the desired inter-domain range in the `CIDR` field. See [CIDR examples](#).
 - o If you select `Security Group`, choose a security group shared by the Instance(s) with which you wish to communicate.

Note: This option only allows access to the Instances within that security group. This differs from CIDR.
9. Click `Add` at the bottom of the dialog box to implement your custom rule.



10. Confirm your new custom rule is displayed in the Manage Rules screen.

Remove Rules

Remove an existing rule

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups` .
2. Select the rule that you would like to remove.
3. Click the `Delete Rule` button on the far right of the selected rule.
4. Confirm deletion of the rule.

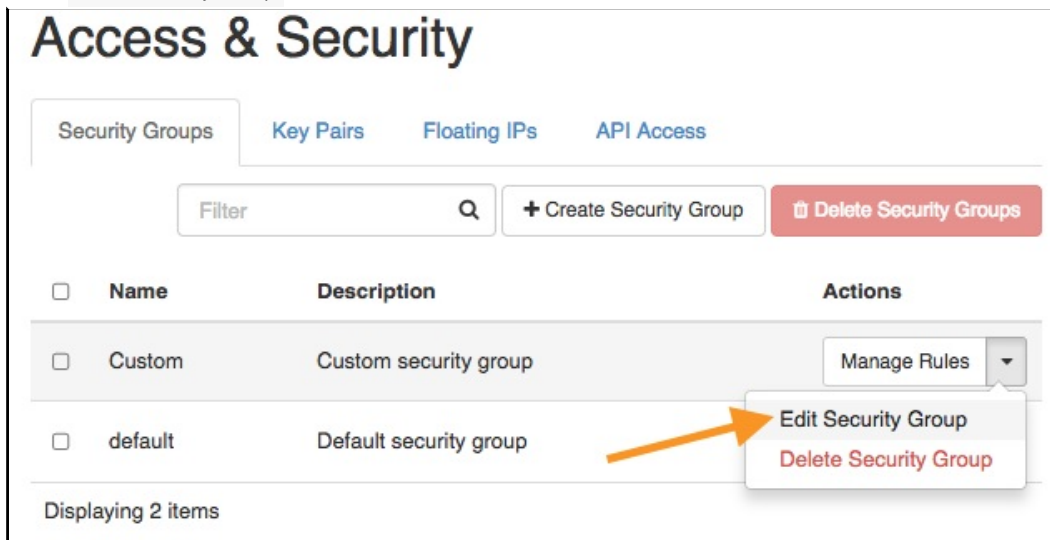
Remove multiple rules

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups` .
2. Toggle the check box next to each rule you would like to remove.
3. Click the `Delete Rules` button in the top-right corner of the Manage Rules screen.
4. Confirm deletion of the rules.

Rename a Security Group

Note: You cannot rename the default security group.

1. In the Security Groups table, select the drop-down menu on the far right of the row.
2. Select `Edit Security Group` .



The screenshot shows the 'Access & Security' console interface. At the top, there are tabs for 'Security Groups', 'Key Pairs', 'Floating IPs', and 'API Access'. Below the tabs is a search bar labeled 'Filter' and two buttons: '+ Create Security Group' and 'Delete Security Groups'. The main content is a table with the following structure:

| <input type="checkbox"/> | Name | Description | Actions |
|--------------------------|---------|------------------------|--|
| <input type="checkbox"/> | Custom | Custom security group | Manage Rules Edit Security Group Delete Security Group |
| <input type="checkbox"/> | default | Default security group | |

At the bottom of the table, it says 'Displaying 2 items'. An orange arrow points to the 'Edit Security Group' option in the dropdown menu for the 'Custom' security group.

3. In the resulting dialog, you can modify the name and description of the user-added Security Group.
4. Click `Edit Security Group` to save your changes.

Edit Security Group ✕

Name *

Description

Custom security group

Description:

Security groups are sets of IP filter rules that are applied to the network settings for the VM. Edit the security group to add and change the rules.

5. Confirm your changes in the Security Groups table.

Delete a Security Group

1. In the Security Groups table, find the Security Group you wish to delete, and select the drop-down menu on the far right of its row.
2. Select `Delete Security Group`.

Access & Security

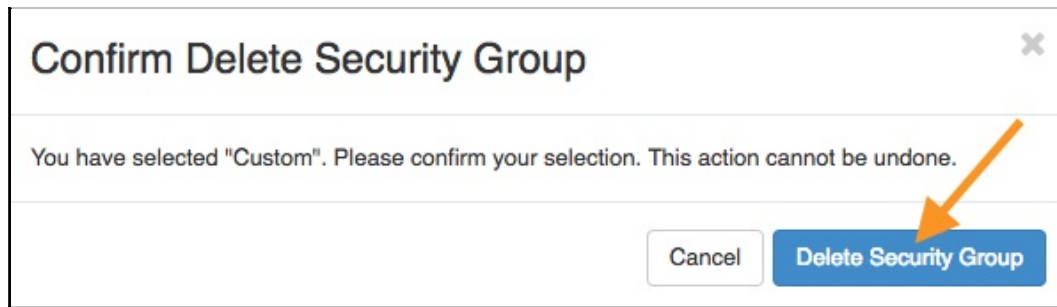
Security Groups

[Key Pairs](#)
[Floating IPs](#)
[API Access](#)

| | Name | Description | Actions |
|--------------------------|---------|------------------------|---|
| <input type="checkbox"/> | Custom | Custom security group | Manage Rules ▼ |
| <input type="checkbox"/> | default | Default security group | Edit Security Group Delete Security Group |

Displaying 2 items

3. Click `Delete Security Group` in the resulting dialog.



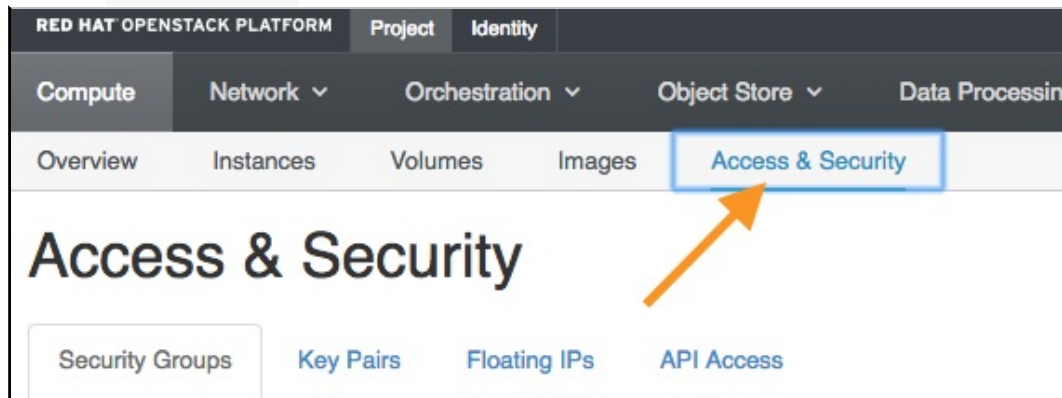
4. Confirm deletion of the Security Group.

CADES → User Documentation → CADES Cloud User Guide → Manage Your VM Instances → OpenStack Security Groups → Create New Security Groups

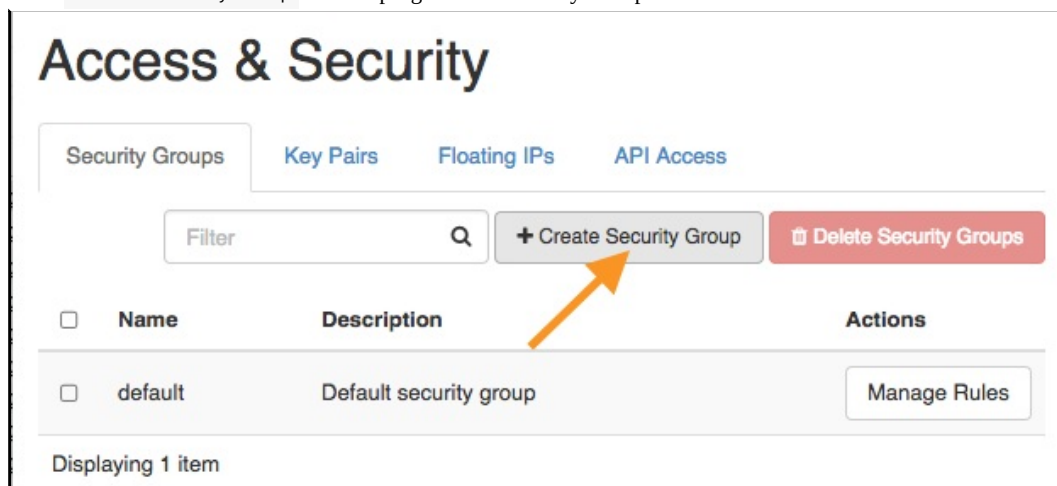
Create a New Security Group

Note: By default, all newly created Security Groups allow all outbound IPV4 and IPV6 (enabled but not functional). By default, no inbound traffic is allowed.

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: ornl
 - User Name: Your three-letter UCAMS ID
 - Password: Your UCAMS password
3. Click on the `Project` tab on the top left.
4. Select the `compute` sub tab.
5. Select the `Access & Security` sub tab.



6. Select the `Security Groups` sub tab to view a table/list of the Security Groups.
7. Click `create security group` at the top right of the Security Groups table.



8. In the resulting dialog, fill out the `Name` and `Description` (optional) fields.
9. Click the `create security group` button to complete the creation of the new Security Group.

Create Security Group ✕

Name *

Description

Custom security group

Description:

Security groups are sets of IP filter rules that are applied to the network settings for the VM. After the security group is created, you can add rules to the security group.

Your new Security Group should now be available in the Security Groups table.

Access & Security

Security Groups

[Key Pairs](#)
[Floating IPs](#)
[API Access](#)

| <input type="checkbox"/> | Name | Description | Actions |
|--------------------------|---------|------------------------|---|
| <input type="checkbox"/> | Custom | Custom security group | Manage Rules ▼ |
| <input type="checkbox"/> | default | Default security group | Manage Rules |

Displaying 2 items

Note: The Security Group must be added to the Instance to take effect. See below.

Add a Security Group to your Instance

- To add the new Security Group to your VM, navigate to `Project` → `Compute` → `Instances`
- Click on the drop-down () menu to the right of the instance to which you would like to attach the new rule, then select `Edit security groups`.
- On the left side of the resulting window are all of the available security groups. On the right is a list of the security groups that are attached to your instance. Find the security group on the left that you would like to add to your instance and click (+).

- Click `save` .

Just as you can with the default Security Group, you can create and manage rules for user-created Security Groups from the Manage Rules screen.

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Manage Your VM Instances](#) → [OpenStack Security Groups](#) → [CIDR Security Examples](#)

Security Group CIDR Examples

When [adding a new rule to a Security Group](#), you can also specify CIDR configurations for each rule. Some examples are provided below.

- `0.0.0.0/0` – This CIDR configuration leaves traffic open to the world. However, other firewalls between the CIDR and the remote machine can still block traffic.
- `216.37.64.68/32` – This CIDR configuration only allows 216.37.64.68 access to the selected port(s). The `/32` is used to specify traffic for only the preceding IP address.
- `192.168.1.0/24` – This CIDR configuration allows all IPs between 192.168.1.1 and 192.168.1.254 access to the selected port(s). The `/24` is used to specify this traffic range.

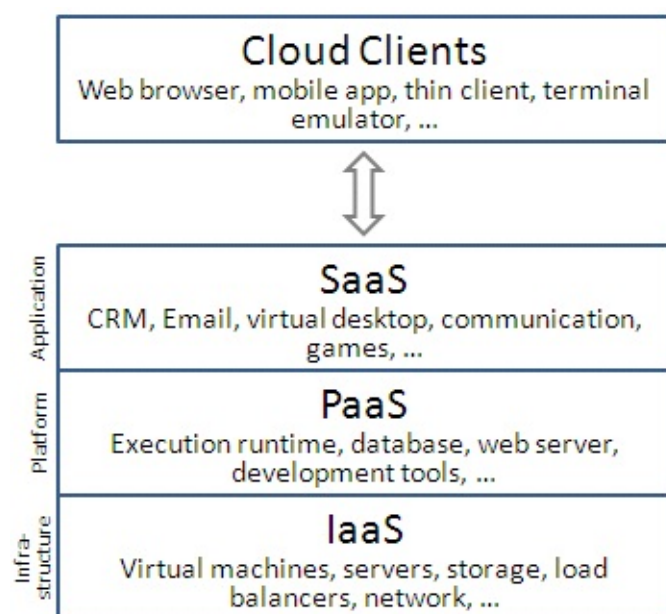
CADES Cloud User Guide

CADES (Compute and Data Environment for Science) provides eligible customers with an OpenStack-based cloud computing solution with customizable Virtual Machines (VM). This resource, called "CADES Cloud", enables customers in [science and technology directorates](#) to leverage self-service portals to rapidly request VMs for production, testing, and development. This documentation will walk you through how to configure and use your CADES Cloud allocation. The CADES Cloud allocations are intended and configured to be integrated within the ORNL network.

What is cloud computing?

Cloud computing provides an efficient pool of on-demand, self-managed virtual infrastructure, consumed as a service.

As shown in the figure below, classic cloud clients interact with three layers in the cloud environment.



First, the Software as a Service (SaaS) layer that presents software tools and frameworks such as emails, enterprise systems to users. Second, the Platform as a Service (PaaS) that presents the runtime services such as web servers, programming environment etc. to user. Third, the Infrastructure as a Service (IaaS) layer that provides hardware and firmware such as storage, drivers and load balancers to users.

The CADES Cloud allocations provide:

- **Self Service** – Through the Horizon web interface, users can create, manage, and delete VMs.
- **Portable** – Operations can be performed using any local ORNL system that provides a Bash terminal and SSH functionality.
- **Elasticity** – Users can create VMs on demand and delete them when they are no longer desired.

What is OpenStack?

[OpenStack](#) is an open-source cloud computing software framework that allows administrators (the CADES team) to create individual "Project" allocations for their customers. The customers/users can then fill these Project allocations with their own VMs without further intervention from CADES administrators—a true self-service implementation.

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Available VM Images](#)

Available VM Images and Configurations

The CADES OpenStack allocations can use either CADES-provided images or user-added images for the VM Instances. However, for full compatibility and best reliability, we strongly recommend that customers use one of the CADES images.

CADES Images

The CADES team currently provides two Linux images for use in OpenStack.

- CentOS 7.x
- Ubuntu 16.04 LTS

Features of CADES Images

CADES images have been preconfigured with the following features and modifications.

- The default user has been changed to "cades" and given full sudo privileges.
- A user named "cades-ops" has been added to ensure that the CADES team has adequate access to provide support.
- Limited Lightweight Directory Access Protocol (LDAP) functionality has been added to enable the use of UCAMS IDs after the initial key-based authentication as the "cades" user. To add your UCAMS ID, simply execute `sudo su - YOUR_UCAMS_ID` in the terminal.
 - You will not be prompted for a password.
 - A local home directory will be created for your user.
 - [Click here for more detailed instructions for adding UCAMS users to your Instance.](#)
- Yum/Apt repos and the Network Time Protocol (NTP) have been configured to use local ORNL resources.
- The minimum disk size required for each image is 8 GB.

Naming Scheme for CADES Images

The CADES naming scheme is based around a `CADES_{$OSRELEASE}_v{$BUILDDATE}_{$RELEASE}` formatting. The `$_RELEASE` field is always a single digit, with a `1` indicating a production version of an image and a `0` indicating a development version of the image.

User Images

While you can run user-provided images in OpenStack, the CADES team strongly recommends that customers use CADES-provided images for best reliability and integration in the ORNL environment. **CADES will not provide support for user-provided images.** If you still want to run a custom cloud image please [contact the CADES team](#) for your request.

Important Notes for Available VM Images

Once an image is launched as a volume, the image is no longer tied to the volume. If the base image is updated, those updates would not propagate to the Instance and vice versa.

Available VM Instance Flavors

CADES Cloud allocations are available in a variety of sizes to fit your needs. The size options can be viewed during VM setup.

If you see that your needs cannot be met by one of our preset configurations, feel free to [email us](#).

Software & Hardware Details

Software Stack

- **Host operating system:** Red Hat OpenStack Platform 9
- **Available VM operating systems:** CentOS 7.x and Ubuntu 16.04 LTS

Hardware Configuration

RAM Information

- **Make:** Samsung
- **Model:** M393A2G40DB0-CPB
- **Speed:** DDR4 2133
- **Error correction:** Registered ECC

CPU Information

The following CPUs are used in the "nova" and "Lustre-OpenStack" Availability Zones.

- **Make:** Intel
- **Model:** Xeon E5-2698 v3
- **Speed:** 2.30 GHz base clock, 3.60 GHz Turbo Boost clock
- **Instruction set (VM Instances):** [CPUID Instance Codes](#)
- **Instruction set (Hypervisors):** [CPUID Hypervisor Codes](#)

Detailed CPU information is also available through CPUID.

How to: [Install CPUID](#)

CPU Layout



Network and Storage Details

Network Configuration

The CADES Cloud allocations consist of two primary network environments or subnets—an external subnet and an internal subnet—described below. While either subnet can be used for a VM Instance, **only one IP address can be allocated to each subnet per VM Instance**.

Note: If you wish to run services on your VM Instance that should be available outside of ORNL's network, ensure that you select the [External Network option](#) when setting up your VM Instance and that you also add a rule to your Security Group for that particular service.

External Network

- `general_extnetwork1`: 128.219.184.0/21.
- **Required for services to be available outside of ORNL's network.**
- Outward-facing services (e.g., a web server) will require ORNL firewall exceptions ([instructions](#)).
- Routed to most "open" networks at ORNL.
- Outbound access is allowed per existing open research firewall exceptions.

Internal Network

- `general_intnetwork1` = 172.22.0.0/20.
- **Services will not be available outside of ORNL's network.**
- Outward-facing services (e.g., a web server) will only be available from within ORNL.
- Routed to most "open" networks at ORNL.
- Outbound access is allowed through existing open research firewall exceptions.

Important Notes for OpenStack Network Design

- IPV6 is not currently supported.
- In all cases, accessing a VM Instance via SSH from outside of ORNL's network requires a SAFER exception to allow inbound traffic on port 22 (SSH).
- SSH access *from* an instance to a destination outside of ORNL's network may route through the ORNL SSH proxy service ([instructions](#)), or the user can request an outbound exception using SAFER.
- If you wish to run services on your VM Instance that should be available outside of ORNL's network, ensure that you also select the [External Network option](#) when setting up your VM Instance.

Storage Configuration

- **NFS** – CADES Projects use a Network File System (NFS).
- **Lustre** – Lustre allocations are available upon request.

Note: No moderate/confidential data should be mounted or copied to your CADES Cloud VM Instance. CADES Cloud VM Instances are for open science.

[CADES](#) → [User Documentation](#) → [CADES Cloud User Guide](#) → [Additional Cloud Resources](#)

Additional Resources

There are some topics that are optional and their use is based on how you chose to use your CADES VMs.

In this section you can:

- [Request Firewall Exception](#)
- [Run a Simple Web Server](#)
- [SSL - Let's Encrypt](#)
- [Install CPUID](#)
 - [CPUID Hypervisor Codes](#)
 - [CPUID Instance Codes](#)

Also, you can peruse the User-Contributed Tutorial section which currently covers the following topics:

- [Launch a Docker Container](#)
- [Launch Shiny within Docker](#)
- [Eclipse in CADES HPC](#)
- [Allinea DDT in CADES HPC](#)

Note: Content contributed by the community is not supported by CADES.

If there is something missing, please [email CADES](#) to let us know. You can also help us by contributing your own content. See the [Contributing](#) section for more information.

ORNL Firewall Configuration (SAFER)

In some cases, a customer may want to run a service on their VM Instance that should be available outside of ORNL's network (e.g., SSH, web server, GitLab service, Docker), which requires an exception in ORNL's firewall. The procedure below will show how to set up a firewall exception for a service running on your VM Instance.

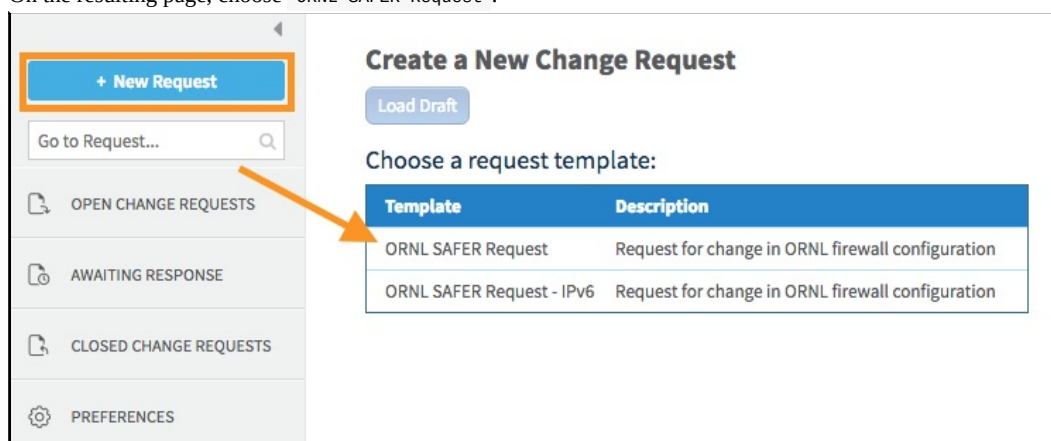
If you have further questions about getting your ORNL (SAFER) firewall rules in place, e-mail the SAFER team directly at security@ornl.gov.

If you wish to run services on your VM Instance that should be available outside of ORNL's network, ensure that you select the [External Network option](#) when setting up your VM Instance and that you also add a rule to your Security Group for that particular service. Read more about the CADES Cloud network design [here](#).

Request a Firewall Exception

For the purposes of this procedure, we will set up an exception for a web server running on port 80.

1. Navigate to <https://safer.ornl.gov>.
2. Log in to the SAFER interface using your UCAMS credentials.
3. Click **+ New Request** at the top left of the screen.
4. On the resulting page, choose **ORNL SAFER Request**.



5. In the resulting request dialog, we need to fill out the following fields:
 - **Subject** – A simple subject will do. We're going with **VM web server**.
 - **Authorization** – Set to **None**.
 - **Change Request Justification** – Provide the reason for your firewall exception.
 - **Expires** – Leave this blank to make this exception indefinite. Otherwise, choose a date for the exception to expire.
 - **Source** – The source IP or IP range (with CIDR notation if range) of the server for which you would like to make an exception. We're going to use the single IP address of our VM Instance, **128 . 219 . 186 . 29**.
 - **Destination** – For this example, we're going to make the source and the destination the same IP address, **128 . 219 . 186 . 29**.
 - **Service** – This can be formatted using the protocol/port (e.g., **TCP/80**) or you can choose from a list of common multi-port services in the drop-down menu. We're using **TCP/80**.
 - **Service Name** – User-defined name of the rule. We're calling ours **blackmesa_web**.
6. Once filled out, click **Next** to submit your request.

Create a New Change Request

[Back](#) [Save Draft](#) [Next](#)

General

Subject: VM web server

Requestor: crawfordst@ornl.gov

Authentication: None
 RSA
 Badge
 UCAMS
 XCAMS
 DCAMS

Expires:

Change request justification: Working with external developers who need to see and provide feedback on the website before it goes live on another server.

E-mail notification list (opt.):

Traffic

Set traffic values [Import traffic from csv](#)

| Source | Destination | Service | Action |
|----------------|----------------|---------|--------|
| 128.219.186.29 | 128.219.186.29 | tcp/80 | Allow |

NAT settings

Service Name: blackmesa_web

[+ Add more traffic](#)

[Back](#) [Save Draft](#) [Next](#)

You will receive an e-mail confirmation of your request. You can also view the status of your exception request at any time by logging into the [SAFER interface](#).

Related Tutorials

- [Run a Web Server on Your VM Instance](#)

Run a Simple Web Server

The following documentation will show you how to launch an Instance using the default Security Group. Once running, we will enable a basic HTTP service on port 80 by adding an additional Security Group.

Launch an Instance Using Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`
3. Navigate to `Project` → `Compute` → `Instances`.
4. Click the `Launch Instance` button, and complete the launch instance wizard.
 - For this exercise use the `default` security group.

If you have never launched a VM Instance before, check out the tutorial linked below before proceeding.

How to: [Launch a VM Instance from an Image](#)

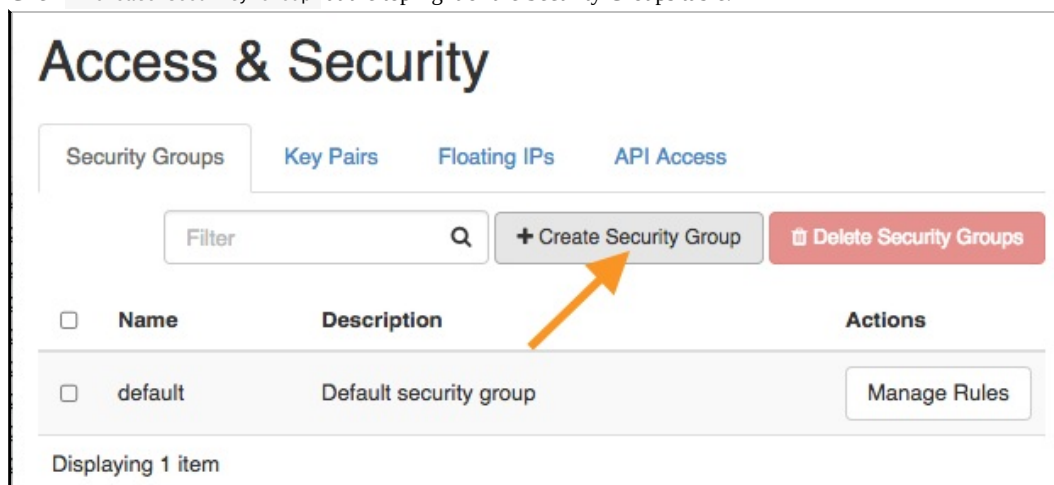
Add Rule for TCP Traffic

Once we have our VM Instance up and running, we need to make sure that the Instance can receive incoming traffic on port 80. We'll create a custom Security Group, add the required rule for HTTP traffic to our new (or existing) Security Group, and assign that Security Group to our VM Instance.

If you require HTTPS support on port 443, see our [SSL - Let's Encrypt](#) documentation.

Create Custom Security Group

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups`.
2. Click `+ Create Security Group` at the top right of the Security Groups table.



3. In the resulting dialog, fill out the `Name` and `Description` (optional) fields.
 - We're using `http-server` for this example.

4. Click the `create Security Group` button.

Create Security Group ✕

Name *

Description

For the simple web server tutorial.

Description:

Security groups are sets of IP filter rules that are applied to the network settings for the VM. After the security group is created, you can add rules to the security group.

Your new Security Group should now be available in the Security Groups table.

Alternatively, you can modify an existing group using the corresponding `Manage Rules` button.

Add Rule to Custom Security Group

1. Find the newly created `http-server` rule on the `Security Groups` tab.

Access & Security

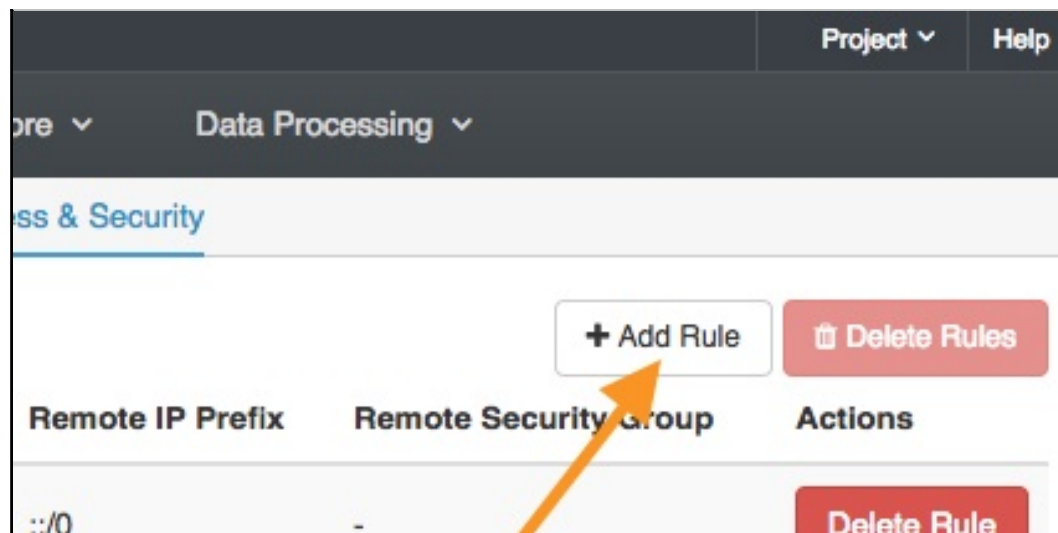
Security Groups **Key Pairs** Floating IPs API Access

Filter

| <input type="checkbox"/> | Name | Description | Actions |
|--------------------------|-------------|-------------------------------------|---|
| <input type="checkbox"/> | Custom | Custom security group | <input type="button" value="Manage Rules"/> ▾ |
| <input type="checkbox"/> | default | Default security group | <input type="button" value="Manage Rules"/> |
| <input type="checkbox"/> | http-server | For the simple web server tutorial. | <input type="button" value="Manage Rules"/> ▾ |

Displaying 3 items

2. Click the `Manage Rules` button.
3. On the Manage Rules screen, click the `Add Rule` button.



4. In the resulting dialog, click the drop-down field under **Rule**.
5. Choose the **HTTP** rule template from the drop-down list.
 - Choosing the HTTP template will automatically set the port to **80** and set the direction to **Ingress**.
6. In the **Remote** box, choose **CIDR** (preferred) and leave the field as **0.0.0.0/0** or enter the desired inter-domain range. See [CIDR examples](#) for more information.
7. Click **Add** at the bottom of the dialog box to implement your custom rule.

Add Rule ✕

Rule *

Remote * ?

CIDR ?

Description:
 Rules define which traffic is allowed to instances assigned to the security group. A security group rule consists of three main parts:

Rule: You can specify the desired rule template or use custom rules, the options are Custom TCP Rule, Custom UDP Rule, or Custom ICMP Rule.

Open Port/Port Range: For TCP and UDP rules you may choose to open either a single port or a range of ports. Selecting the "Port Range" option will provide you with space to provide both the starting and ending ports for the range. For ICMP rules you instead specify an ICMP type and code in the spaces provided.

Remote: You must specify the source of the traffic to be allowed via this rule. You may do so either in the form of an IP address block (CIDR) or via a source group (Security Group). Selecting a security group as the source will allow any other instance in that security group access to any other instance via this rule.

8. Confirm your new custom rule is displayed in the Manage Rules screen.

| <input type="checkbox"/> | Direction | Ether Type | IP Protocol | Port Range | Remote IP Prefix | Remote Security Group | Actions |
|--------------------------|-----------|------------|-------------|------------|------------------|-----------------------|-------------|
| <input type="checkbox"/> | Egress | IPv4 | Any | Any | 0.0.0.0/0 | - | Delete Rule |
| <input type="checkbox"/> | Egress | IPv6 | Any | Any | :::0 | - | Delete Rule |
| <input type="checkbox"/> | Ingress | IPv4 | TCP | 80 (HTTP) | 0.0.0.0/0 | - | Delete Rule |

Displaying 3 items

Add New Security Group to Your VM Instance

1. Navigate to **Project** → **Compute** → **Instances**, and find the Instance on which you would like to run your HTTP server.
2. Click the drop-down menu on your Instance's listing, and select **Edit Security Groups** from the resulting menu.

Overview Instances Volumes Images Access & Security

Instances

Instance Name = Filter Launch Instance Delete Instances More Actions

| <input type="checkbox"/> | Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created | Actions |
|--------------------------|---------------|------------|----------------|---------|--------------|--------|-------------------|------|-------------|--------------------|-----------------|
| <input type="checkbox"/> | aperture | - | 128.219.186.46 | m1.tiny | blackmesakey | Active | nova | None | Running | 3 weeks, 5 days | Create Snapshot |
| <input type="checkbox"/> | black-mesa | - | 128.219.186.42 | m1.tiny | blackmesakey | Active | nova | None | Running | 1 month | Create Snapshot |

Displaying 2 items

- Associate Floating IP
- Attach Interface
- Detach Interface
- Edit Instance
- Update Metadata
- Edit Security Groups**
- Console

3. In the resulting window, click **+** to add the `http-server` security group to your VM Instance.
4. Click **save** to associate your `http-server` security group with your Instance.

Edit Instance

Information * **Security Groups**

Add and remove security groups to this instance from the list of available security groups.

All Security Groups Filter

Custom

http-server

Instance Security Groups Filter

default

Connect to Your Instance Using SSH

1. Open a Bash terminal.
2. Execute `$ ssh cades@128.219.186.42`.

- o Replace `128.219.186.42` with the IP address of your own Instance.
3. You should now be connected to your VM Instance via SSH.

```
*****
Welcome to OpenStack

Support:
Documentation/SLA: https://code.ornl.gov/cades-ua/user-documentation/
Email support: cades-help@email.ornl.gov
Chat/community: https://cades.slack.com/

sudo:
The 'cades' user on this VM has full sudo.
You may become your ucams user via:
sudo su - ucamsID

This message::
Served courtesy of /etc/profile.d/motd_userInfo.sh
*****
cades@black-mesa:~$
```

For more detailed information on connecting to your VM Instance using SSH, please see the SSH tutorials, linked below.

How to: [Access Your VM Instance Using SSH](#)

For Windows users, we have developed a separate tutorial that walks you through connecting to your VM Instance using PuTTY's SSH client.

How to: [Access Your VM Instance Using PuTTY \(Windows\)](#)

Set up Your Web Server

For the web server, we need to create a web directory and populate it with a basic `index.html` page.

1. Execute the following in the VM's Bash terminal.

```
$ mkdir www
$ cd www
```

This will create a `www` directory. Next, we need to open a new file (`index.html`) in the VIM text editor and populate it with some basic HTML markup.

2. In the Bash terminal, execute `$ vim index.html` to create the new file and open it in VIM.
3. Once in VIM, add the following lines to your file.

```
<html>
<head>
  <title>
    Hello world!
  </title>
</head>
<body>
  <h1>Hello world!</h1>
</body>
</html>
```


Troubleshooting

If you get an error when attempting to navigate to your Instance's IP using a browser, ensure that:

- You used the external network option when you configured your Instance.
- You are within the ORNL network or you have added a SAFER firewall exception for your VM Instance's IP address on Port 80 (or whatever port you specified in the Python command).
- You have a Rule in your Security Group that grants access to your VM Instance on Port 80 (or whatever port you specified in the Python command).
- Your Python server is active on your VM Instance.

Do You Need HTTPS?

- [SSL - Let's Encrypt](#)

Related Tutorials

- [Launch a VM Instance](#)
- [Access VM Instance Using SSH](#)

Adding a Security Certificate to Your Website

Obtaining a security certificate from a Certificate Authority enables you to use https on your website hosted by CADES. Utilizing https maintains the confidentiality of the transferred information by using a set of encryption keys. Additionally, this document will cover requirements for Firewall exceptions, cipher specifications, and HSTS preloading conditions.

Caveat: Auto SSL Creation for CADES VMs

In order to expedite ORNL's required compliance with [DHS BOD 18-01](#) as many sites as possible were automatically adjusted to defaulting to HTTPS with a Let's Encrypt certificate.

Directive Overview

Pursuant to [DHS BOD 18-01](#) all federal websites must be migrated to HTTPS. The prescribed implementation pattern is as follows

- Redirect HTTP traffic to HTTPS via [301 Moved Permanently](#)
- HTTPS should be configured with a trusted TLS certificate
- All of the following should be disallowed:
 - SSL v2
 - SSL v3
 - RC4, DES, and 3DES ciphers
- [HSTS](#) header set with `max-age` equal to 1 year.

Table of Contents

Prerequisites

Part 1: Modifying the OpenStack Horizon Security Group

- Add Rule for Secured TCP Traffic
- Include new rule in your Instance's Security Group

Part 2: Obtaining the Security Certificate for your Site

- Accessing your VM via SSH
- Using Certbot's automated client (Ubuntu or CentOS)
- Setting Up Autorenewal

Part 3: Updating Firewall Exception

- Requesting new firewall exception

Part 4: Configuring your Webserver

Part 5: Certificate Auto-Renewal

Prerequisites

- A website hosted by CADES resources
- Bash access to the VM hosting the site
- Enable virtual hosts in your apache config file and ensure the default site is served as a virtual host.
 - [CentOS Apache Virtual Host Configuration](#)
 - [Ubuntu Apache Virtual Host Configuration](#)

Part 1: Modifying the OpenStack Horizon Security Group

View the Instance Using Horizon

1. Navigate to the Horizon web interface at <https://cloud.cades.ornl.gov/>.
2. Log in with your UCAMS credentials.
 - Domain: `ornl`
 - User Name: `Your three-letter UCAMS ID`
 - Password: `Your UCAMS password`
3. Navigate to `Project` → `Compute` → `Instances` .
4. Click on the existing instance that hosts your web server if you'd like to review the settings.

If you have never launched a VM Instance before, check out the tutorial linked below before proceeding.

How to: [Launch a VM Instance](#)

Add Rule for Secured TCP Traffic

When you originally set up your web server, the settings allowed incoming traffic on port 80. Now we need to allow secure traffic. Then we'll make sure the rule is properly assigned to your Instance.

1. Navigate to `Project` → `Compute` → `Access & Security` → `Security Groups` .
2. Find the existing Security Group that contains your http access rules and click `Manage Rules` on the right side of the screen.
3. In the resulting window, click `+ Add Rule` .
4. In the resulting dialog, click the drop-down field under `Rule` .
5. Choose the `HTTPS` rule template from the drop-down list.
 - Choosing the HTTPS template will automatically set the port to `443` and set the direction to `Ingress` .
6. In the `Remote` box, choose `CIDR` (preferred) and leave the field as `0.0.0.0/0` or enter the desired inter-domain range. See [CIDR examples](#) for more information.
7. Click `Add` at the bottom of the dialog box to implement your custom rule.
8. Confirm your new custom rule is displayed in the Manage Rules screen.

More info: [Security Groups](#)

Add Rule to the Instance's Security Group

1. Navigate to `Project` → `Compute` → `Instances` .
2. On the right, click the down arrow () next to `Create Snapshot` .
3. Select `Manage Security Groups` .
4. Click the `+` next to your new HTTPS rule to add it to your Instance.

Part 2: Obtaining the Security Certificate for your Site

The process of obtaining the SSL certificate is automated via the command line using Let's Encrypt and Certbot.

Access your VM via ssh.

1. Open a Bash terminal.
2. Execute `ssh cades@128.219.186.42` .
 - o Replace `128.219.186.42` with the IP address of your own Instance.
3. You should now be connected to your VM Instance via SSH.

```

*****
                Welcome to OpenStack

Support:
Documentation/SLA: https://code.ornl.gov/cades-ua/user-documentation/
Email support: cades-help@email.ornl.gov
Chat/community: https://cades.slack.com/

sudo:
The 'cades' user on this VM has full sudo.
You may become your ucams user via:
sudo su - ucamsID

This message::
Served courtesy of /etc/profile.d/motd_userInfo.sh
*****
cades@black-mesa:~$
```

For more detailed information on connecting to your VM Instance using SSH, please see the SSH tutorials, linked below.

How to: [Access Your VM Instance Using SSH](#)

For Windows users, we have developed a separate tutorial that walks you through connecting to your VM Instance using PuTTY's SSH client.

How to: [Access Your VM Instance Using PuTTY \(Windows\)](#)

Installing Certbot for Ubuntu (see below for CentOS)

Certbot is an automated client that obtains and implements SSL certificates for your website.

1. Obtain the necessary Certbot packages.

```

sudo apt-get update
sudo apt-get install software-properties-common
sudo add-apt-repository ppa:certbot/certbot
sudo apt-get update
sudo apt-get install python-certbot-apache
```

2. Use Certbot's Apache plugins to automate the certificate process.
 - o There are two options. Option A should be used if *cannot* stop and restart your web server. Option B should be used if you *can* start and stop your web server. If you are not sure, choose Option A.
 - o Option A


```
sudo certbot --authenticator webroot --installer apache
```
 - o Option B


```
sudo certbot --authenticator standalone --installer apache --pre-hook "apachectl -k stop" --post-hook
```

```
"apachectl -k start"
```

- You will be asked to enter your domain name (e.g. `example.com`) and your webroot (the folder where your website lives, e.g. `www`).

Installing Certbot for CentOS (see above for Ubuntu)

Certbot is an automated client that obtains and implements SSL certificates for your website.

1. Obtain the necessary Certbot packages.

```
wget https://dl.eff.org/certbot-auto
chmod a+x certbot-auto
```

2. Use Certbot's plugins to automate the certificate process. (Fill in the `/path/to/` with your specific path.)

```
sudo ./path/to/certbot-auto certonly
```

- Select `2` when prompted to place Certbot's files in the webroot directory.
- Please also provide your email address and your domain name, e.g. `example.com`.
- Next, you'll be asked to provide the webroot for your domain (the folder where your website lives), e.g. `www`.
- The next command will stop and restart your server. Please be aware of the temporary service interruption.

```
sudo certbot --authenticator standalone --installer apache --pre-hook "apachectl -k stop" --post-hook "apachectl -k start"
```

Part 3: Updating Firewall Exception (for External-Facing Sites)

1. Navigate to <https://safer.ornl.gov>.
2. Log in to the SAFER interface using your UCAMS credentials.
 - User Name: Your three-letter UCAMS ID
 - Password: Your UCAMS password
3. Click `+ New Request` at the top left of the screen.
4. On the resulting page, choose `ORNL SAFER Request`.

The screenshot shows the SAFER interface. On the left, there is a sidebar with a '+ New Request' button highlighted in orange. Below it is a search bar 'Go to Request...' and a list of request categories: 'OPEN CHANGE REQUESTS', 'AWAITING RESPONSE', 'CLOSED CHANGE REQUESTS', and 'PREFERENCES'. The main content area is titled 'Create a New Change Request' and has a 'Load Draft' button. Below that, it says 'Choose a request template:' followed by a table with two columns: 'Template' and 'Description'. The table contains two rows: 'ORNL SAFER Request' with description 'Request for change in ORNL firewall configuration', and 'ORNL SAFER Request - IPv6' with description 'Request for change in ORNL firewall configuration'. An orange arrow points from the '+ New Request' button to the first row of the table.

| Template | Description |
|---------------------------|---|
| ORNL SAFER Request | Request for change in ORNL firewall configuration |
| ORNL SAFER Request - IPv6 | Request for change in ORNL firewall configuration |

5. In the resulting request dialog, we need to fill out the following fields:
 - Subject – A simple subject will do. We're going with `VM web server`.
 - Authorization – Set to `None`.
 - Change Request Justification – Provide the reason for your firewall exception.

- Expires – Leave this blank to make this exception indefinite. Otherwise, choose a date for the exception to expire.
- Source – The source IP or IP range (with CIDR notation if range) of the server for which you would like to make an exception. We're going to use the single IP address of our VM Instance, 128 . 219 . 186 . 29.
- Destination – For this example, we're going to make the source and the destination the same IP address, 128 . 219 . 186 . 29.
- Service – This can be formatted using the protocol/port (e.g., TCP/443) or you can choose from a list of common multi-port services in the drop-down menu. Please use TCP/443 as the secure traffic port.
- Service Name – User-defined name of the rule. We're calling ours blackmesa_web.

6. Once filled out, click **Next** to submit your request.

Create a New Change Request

General

Subject: VM web server

Requestor: crawfordst@ornl.gov

Authentication: None

Expires: []

Change request justification: Working with external developers who need to see and provide feedback on the website before it goes live on another server.

Traffic

Set traffic values

| Source | Destination | Service | Action |
|----------------|----------------|---------|--------|
| 128.219.186.29 | 128.219.186.29 | tcp/80 | Allow |

Service Name: blackmesa_web

Buttons: Back, Save Draft, Next

You will receive an e-mail confirmation of your request. You can also view the status of your exception request at any time by logging into the [SAFER interface](#).

Part 4: Webserver configuration

Your webserver will need to be configured to do the following:

- Redirect HTTP traffic to HTTPS
- Set up the HSTS header and preloading for external-facing sites (see [ORNL documentation](#))
- Use the TLS certificate designated for your site
- Allow only approved ciphers (see [ORNL documentation](#))

The exact details will vary between different web and application servers, but Mozilla provides a simple interactive tool that can help you get started:

<https://mozilla.github.io/server-side-tls/ssl-config-generator/>

Part 5: Certificate Auto-Renewal

Using a simple script we can automate the certificate renewal process. Although this certificate lasts for 90 days, running this cron job often will ensure your certificate stays up to date.

```
@weekly python -c 'import random; import time; time.sleep(random.random() * 3600)' && certbot renew
```

If you find that you need more guidance setting up automatic renewal, you can visit [this site](#).

Install CPUID

Installing CPUID on your VM Instance is a fairly simple task. Once connected to your Instance using SSH, use the Bash terminal to install and run CPUID. Below you will find a procedure for each CADES-provided operating system.

Prerequisites

To install CPUID on your VM Instance, you must first have access to your Instance through SSH ([instructions](#)).

Ubuntu

1. Update Ubuntu.

```
$ sudo apt-get update
```

2. Install CPUID using the package manager.

```
$ sudo apt-get install cpuid
```

3. CPUID is now installed. To run CPUID, execute the following.

```
$ cpuid
```

CentOS

1. Update CentOS.

```
$ sudo yum check-update
```

2. Install CPUID using the package manager.

```
$ sudo yum install cpuid
```

3. When prompted, confirm that you wish to install CPUID. `$ Is this ok [y/d/N]: y`

4. CPUID is now installed. To run CPUID, execute the following.

```
$ cpuid
```

Supported CPUID Codes

- CPUID Codes supported by Hypervisor CPU (Not all in Birthright)
 - fpu: Onboard FPU (floating point support)
 - eagerfpu: Non lazy FPU restore
 - vme: Virtual 8086 mode enhancements
 - de: Debugging Extensions (CR4.DE)
 - smx: Safer mode: TXT (TPM support)
 - pse: Page Size Extensions (4MB memory pages)
 - tsc: Time Stamp Counter (RDTSC)
 - constant_tsc: TSC ticks at a constant rate
 - nonstop_tsc: TSC does not stop in C states
 - ptsc: performance time-stamp counter
 - msr: Model-Specific Registers (RDMSR, WRMSR)
 - nodeid_msr: NodeID MSR
 - pae: Physical Address Extensions (support for more than 4GB of RAM)
 - mce: Machine Check Exception
 - cx8: CMPXCHG8 instruction (64-bit compare-and-swap)
 - apic: Onboard APIC
 - x2apic: x2APIC
 - extapic: Extended APIC space
 - sep: SYSENTER/SYSEXIT
 - mtrr: Memory Type Range Registers
 - k6_mtrr: AMD K6 nonstandard MTRRs
 - pge: Page Global Enable (global bit in PDEs and PTEs)
 - mca: Machine Check Architecture
 - smca: Scalable MCA
 - cmov: CMOV instructions (conditional move) (also FCMOV)
 - pat: Page Attribute Table
 - pse36: 36-bit PSEs (huge pages)
 - clflush: Cache Line Flush instruction
 - dts: Debug Store (buffer for debugging and profiling instructions)
 - acpi: ACPI via MSR (temperature monitoring and clock speed modulation)
 - mmx: Multimedia Extensions
 - cxmmx: Cyrix MMX extensions
 - fxsr: FXSAVE/FXRSTOR, CR4.OSFXSR
 - sse: Intel SSE vector instructions
 - misalignsse: indicates if a general-protection exception (#GP) is generated when some legacy SSE instructions operate on unaligned data. Also depends on CR0 and Alignment Checking bit
 - sse2: SSE2
 - ss: CPU self snoop
 - ht: Hyper-Threading
 - tm: Automatic clock control (Thermal Monitor)
 - rtm: Restricted Transactional Memory
 - pbe: Pending Break Enable (PBE# pin) wakeup support
 - syscall: SYSCALL (Fast System Call) and SYSRET (Return From Fast System Call)
 - nx: Execute Disable
 - pdpe1gb: One GB pages (allows hugepagesz=1G)
 - rdtscp: Read Time-Stamp Counter and Processor ID

- o `lm`: Long Mode (x86-64: amd64, also known as Intel 64, i.e. 64-bit capable)
- o `lahf_lm`: Load AH from Flags (LAHF) and Store AH into Flags (SAHF) in long mode
- o `constant_tsc`: TSC ticks at a constant rate
- o `arch_perfmon`: Intel Architectural PerfMon
- o `pebs`: Precise-Event Based Sampling
- o `bts`: Branch Trace Store
- o `rep_good`: rep microcode works well
- o `nopl`: The NOPL (0F 1F) instructions
- o `xtopology`: cpu topology enum extensions
- o `nonstop_tsc`: TSC does not stop in C states
- o `aperfmpperf`: APERFMPERF
- o `eagerfpu`: Non lazy FPU restore
- o `pni`: SSE-3 (“Prescott New Instructions”)
- o `pclmulqdq`: Perform a Carry-Less Multiplication of Quadword instruction — accelerator for GCM)
- o `dtes64`: 64-bit Debug Store
- o `monitor`: Monitor/Mwait support (Intel SSE3 supplements)
- o `ds_cpl`: CPL Qual. Debug Store
- o `vmx`: Hardware virtualization: Intel VMX
- o `smx`: Safer mode: TXT (TPM support)
- o `est`: Enhanced SpeedStep
- o `tm2`: Thermal Monitor 2
- o `ssse3`: Supplemental SSE-3
- o `fma`: Fused multiply-add
- o `cx16`: CMPXCHG16B
- o `xtptr`: Send Task Priority Messages
- o `pdc`: Performance Capabilities
- o `pcid`: Process Context Identifiers
- o `invpcid`: Invalidate Processor Context ID
- o `dca`: Direct Cache Access
- o `sse4_1`: SSE-4.1
- o `sse4_2`: SSE-4.2
- o `x2apic`: x2APIC
- o `movbe`: Move Data After Swapping Bytes instruction
- o `popcnt`: Return the Count of Number of Bits Set to 1 instruction (Hamming weight, i.e. bit count)
- o `tsc_deadline_timer`: Tsc deadline timer
- o `xsaves`: Save Processor Extended States: also provides XGETBY,XRSTOR,XSETBY
- o `avx`: Advanced Vector Extensions
- o `f16c`: 16-bit fp conversions (CVT16)
- o `rdrand`: Read Random Number from hardware random number generator instruction
- o `lahf_lm`: Load AH from Flags (LAHF) and Store AH into Flags (SAHF) in long mode
- o `abm`: Advanced Bit Manipulation
- o `ida`: Intel Dynamic Acceleration
- o `arat`: Always Running APIC Timer
- o `epb`: IA32_ENERGY_PERF_BIAS support
- o `pln`: Intel Power Limit Notification
- o `pts`: Intel Package Thermal Status
- o `tpr_shadow`: Intel TPR Shadow
- o `vnmi`: Intel Virtual NMI
- o `flexpriority`: Intel FlexPriority
- o `ept`: Intel Extended Page Table
- o `vpid`: Intel Virtual Processor ID

- fsgsbase: {RD/WR}{FS/GS}BASE instructions
- tsc_adjust: TSC adjustment MSR
- bmi1: 1st group bit manipulation extensions
- avx2: AVX2 instructions
- smep: Supervisor Mode Execution Protection
- bmi2: 2nd group bit manipulation extensions
- erms: Enhanced REP MOVSB/STOSB
- invpcid: Invalidate Processor Context ID
- cqm: Cache QoS Monitoring
- xsaveopt: Optimized XSAVE
- cqm_llc: LLC QoS
- cqm_occup_llc: LLC occupancy monitoring

Supported CPUID Instruction Set Codes

- CPUID Instruction Set Codes Supported by VM Instances
 - fpu: Onboard FPU (floating point support)
 - eagerfpu: Non lazy FPU restore
 - vme: Virtual 8086 mode enhancements
 - de: Debugging Extensions (CR4.DE)
 - smx: Safer mode: TXT (TPM support)
 - pse: Page Size Extensions (4MB memory pages)
 - tsc: Time Stamp Counter (RDTSC)
 - constant_tsc: TSC ticks at a constant rate
 - nonstop_tsc: TSC does not stop in C states
 - ptsc: performance time-stamp counter
 - msr: Model-Specific Registers (RDMSR, WRMSR)
 - nodeid_msr: NodeId MSR
 - pae: Physical Address Extensions (support for more than 4GB of RAM)
 - mce: Machine Check Exception
 - cx8: CMPXCHG8 instruction (64-bit compare-and-swap)
 - apic: Onboard APIC
 - x2apic: x2APIC
 - extapic: Extended APIC space
 - sep: SYSENTER/SYSEXIT
 - mtrr: Memory Type Range Registers
 - k6_mtrr: AMD K6 nonstandard MTRRs
 - pge: Page Global Enable (global bit in PDEs and PTEs)
 - mca: Machine Check Architecture
 - smca: Scalable MCA
 - cmov: CMOV instructions (conditional move) (also FCMOV)
 - pat: Page Attribute Table
 - pse36: 36-bit PSEs (huge pages)
 - clflush: Cache Line Flush instruction
 - mmx: Multimedia Extensions
 - cxmmx: Cyrix MMX extensions
 - fxsr: FXSAVE/FXRSTOR, CR4.OSFXSR
 - sse: Intel SSE vector instructions
 - misalignsse: indicates if a general-protection exception (#GP) is generated when some legacy SSE instructions operate on unaligned data. Also depends on CR0 and Alignment Checking bit
 - sse2: SSE2
 - ss: CPU self snoop
 - syscall: SYSCALL (Fast System Call) and SYSRET (Return From Fast System Call)
 - nx: Execute Disable
 - pdpe1gb: One GB pages (allows hugepagesz=1G)
 - rdtscp: Read Time-Stamp Counter and Processor ID
 - lm: Long Mode (x86-64: amd64, also known as Intel 64, i.e. 64-bit capable)
 - lahf_lm: Load AH from Flags (LAHF) and Store AH into Flags (SAHF) in long mode
 - constant_tsc: TSC ticks at a constant rate
 - rep_good: rep microcode works well
 - nopl: The NOPL (0F 1F) instructions
 - eagerfpu: Non lazy FPU restore

- o pni: SSE-3 (“Prescott New Instructions”)
- o pclmulqdq: Perform a Carry-Less Multiplication of Quadword instruction — accelerator for GCM)
- o ssse3: Supplemental SSE-3
- o fma: Fused multiply-add
- o cx16: CMPXCHG16B
- o pcid: Process Context Identifiers
- o invpcid: Invalidate Processor Context ID
- o sse4_1: SSE-4.1
- o sse4_2: SSE-4.2
- o x2apic: x2APIC
- o movbe: Move Data After Swapping Bytes instruction
- o popcnt: Return the Count of Number of Bits Set to 1 instruction (Hamming weight, i.e. bit count)
- o tsc_deadline_timer: Tsc deadline timer
- o xsave: Save Processor Extended States: also provides XGETBY,XRSTOR,XSETBY
- o avx: Advanced Vector Extensions
- o f16c: 16-bit fp conversions (CVT16)
- o rdrand: Read Random Number from hardware random number generator instruction
- o hypervisor: Running on a hypervisor
- o lahf_lm: Load AH from Flags (LAHF) and Store AH into Flags (SAHF) in long mode
- o abm: Advanced Bit Manipulation
- o fsgsbase: {RD/WR}{FS/GS}BASE instructions
- o bmi1: 1st group bit manipulation extensions
- o avx2: AVX2 instructions
- o smep: Supervisor Mode Execution Protection
- o bmi2: 2nd group bit manipulation extensions
- o erms: Enhanced REP MOVSB/STOSB
- o invpcid: Invalidate Processor Context ID
- o xsaveopt: Optimized XSAVE

[CADES](#) → [User Documentation](#) → [SHPC Condo User Guide](#) → [Overview](#)

Overview of SHPC Condos

The CADES Scalable HPC (SHPC) Condos consist of two HPC clusters: one in the ORNL Moderate protection zone (CADES Mod) and one in the ORNL Open protection zone (CADES Open).

To get started using the SHPC Condo, check first to see if you are ready by looking over the [prerequisites](#). Then, learn how to [request access](#). Finally, you are ready to [access your Condo allocation](#).

The list of current groups according to divisions at ORNL can be found [here](#).

SHPC Condo Hardware Configuration

The SHPC is a commodity cluster that contains a set of [MPPs \(Massive Parallel Processors\)](#).

A processor in this cluster is commonly referred to as a node and has its own CPU, memory, and I/O subsystem and is capable of communicating with other nodes.

Node Information

- **Make:** Cray
- **Model:** CS400

RAM Information

- **Speed:** DDR4 2133
- **Error correction:** Registered ECC
- **Capacity:** 128–256 GB per node (GPU nodes and high memory nodes have 256 GB of RAM)

CPU Information

- **Make:** Intel
- **Model:** Xeon E5-2698 v3
- **Speed:** 2.30 GHz base clock, 3.60 GHz Turbo Boost clock
- **Capacity:** 2 CPUs per node
- **CPU layout:** [Click to see image.](#)

GPU Information

- **Make:** NVIDIA
- **Model:** Tesla K80 (2 GK210 GPUs on each K80)
- **Speed:** 560 MHz base clock
- **VRAM:** 24 GB of GDDR5
- **Error correction:** Registered ECC
- **Capacity:** 2 Tesla K80s per node, 2 GK210 GPUs per K80 (4 total GK210 GPUs per node)

SHPC Condo Storage Configuration

Lustre

[Lustre](#) is an on-premises, high performance, parallel file system that utilize technologies such as key, value, and set of attributes to compute data in the following environments:

Open Lustre:

- 1.7 PB of temporary computational storage

Your **temporary local storage** is located at: `/lustre/or-hydra/group/xcams`

Replace `group` with *your group name*, and `xcams` with *your XCAMS/UCAMS ID*.

Moderate Lustre:

- 400 TB of temporary computational storage

Your **temporary local storage** is located at: `/lustre/hydra/group/ucams`

Replace `group` with *your group name*, and `ucams` with *your XCAMS/UCAMS ID*.

Note: All data is automatically purged every 2 weeks.

NFS

[NFS \(Network File System\)](#) is a service that allows shared directories and files with others over a network. Home, software, and project directories have been set up on NFS and *are permanently available through the network*.

Open NFS:

- Each user is automatically given 20 GB of permanent NFS storage.

Moderate NFS:

- Each user is automatically given 20 GB of permanent NFS storage.

Note: If your needs differ from what is listed here, be sure to [contact us](#) to discuss options.

SHPC Condo Software Configuration

In this section, we discuss the SHPC Condos software configuration. Our software environment uses Linux environment [modules](#) to perform this configuration. The software modules available to users also contain preconfigured [compiler toolchains](#), or programming environments which include parallel compiler wrappers and associated MPI stacks. There are also [workflow tools](#) that may help with your applications as well.

Job Scheduler

SHPC utilizes Torque/Moab as a resource manager to schedule jobs, where Moab is used as an external scheduler for the PBS resource management system including job queues and the compute resources.

The job scheduler supports a maximum walltime of 48 hours. If you need more time to run a job, please [contact us](#).

Modules

SHPC has more than one hundred software packages installed. Our software environment uses Linux (CentOS 7.x) environment modules to manage versions and dependencies of software packages. When you load a module, it sets the environment variables necessary for running your program.

Modules: Local repository

By default the local repository is used as a source of software installations. To list available modules, type `module avail`. To load a module, use `module load module_name`. Similarly, unload modules by typing `module unload module_name`.

Modules: CVMFS-based repository

A CVMFS (Cern Virtual File System)-based repository is available for use that has several software packages available for use. To use the CVMFS-based repository run the following commands from your login node:

```
source /software/dev_tools/swtree/cs400/modulefiles/switch-modules.sh
```

```
switch_modules oasis
```

After entering the above commands the new repository should be active and the command below will list the software available for use:

```
module avail
```

Similarly `switch_modules local` will bring back the local modules to use.

Additional information on SHPC modules may be found [here](#).

Compiler ToolChains

Depending on the application/code you are working on, you might choose a specific compiler to achieve the best performance of your programs. Compiler toolchains such as GNU, Intel, PGI and NAG are already installed to work with other libraries. [See here](#) for more information on SHPC compilers

Workflow Tools

Workflow tools orchestrate multi-stage computations. Several [workflow tools](#) are available on SHPC.

Scheduling Jobs

SHPC utilizes Torque/Moab to manage jobs that users submit to various queues on a computer system. Each queue represents a group of resources with attributes necessary for the queue's jobs. You can see the list of queues that SHPC has by typing `qstat -q`. `batch` is the default queue.

Note: Do not run jobs on the login nodes. All jobs launched from those nodes will be terminated without notice.

Listing jobs

To list all jobs:

```
qstat
```

To refine the list of jobs to only those submitted by a user:

```
qstat -u UID
```

To further refine the list of jobs, the following command will list jobs submitted by a user and which are running.

```
qstat -u UID -s -r
```

To obtain the status of a job, run the following command using the job's ID number (this is provided at time of job submission).

```
qstat -f job_ID
```

You can also use `checkjob job_ID` to show the current status of the job.

Submitting a job

To submit a job, use the `qsub` command, followed by the name of your submission file. A Job ID will be provided. You may want to make note of the ID for later use.

```
qsub your_script
```

Deleting a job

Note: Be aware that deleting a job cannot be undone. Double check the job ID before deleting a job.

Users can delete their jobs by typing the following command.

```
qdel job_ID
```

To delete all the jobs of a user:

```
qdel $(qselect -u UID)
```

Related Information

- [Execute a Job](#)

Customizing Your Environment in SHPC

The SHPC environment may be customized to suit your needs.

Project-Specific Environment Variables

Some projects have environment modules that will prepare the environment with the specific needs of the project. To list the available project modules, type `module avail`. At the top of the output is a section titled `/software/tools/modules`. The environment modules begin with `env/`. To load one of these environments:

```
module load env/cades-bsd
```


Modules

Modules are a utility which allow users to load and manage applications and their versions. The modules of software packages allow you to dynamically modify your user environment by using “modulefiles.” Each modulefile contains the information needed to configure the shell for an application. After the module software package is initialized, the environment can be modified on a per-module basis using the module command, which interprets modulefiles. Typically, modulefiles instruct the module command to change or set shell environment variables such as `PATH` , `MANPATH` , and others. The modulefiles can be shared by many users on a system.

Note: Some modules cannot be used simultaneously, such as an Intel compiler and a GNU compiler ([information on compilers](#)). If you attempt to *load* a module that is incompatible with a currently-loaded module, you will be prompted of the conflict. To avoid the error, you may have to *unload* or *switch* modules.

Summary of Module Commands

| Command | Description |
|----------------------------|---|
| <code>module list</code> | Lists modules currently loaded in a user's environment. A module is considered loaded when its associated modulefile has been executed and the user's environment contains the changes from the modulefile. |
| <code>module avail</code> | Lists all available modules on a system. |
| <code>module show</code> | Shows environment changes that will be made by loading a given module. |
| <code>module load</code> | Loads a module. |
| <code>module unload</code> | Unloads a module. |
| <code>module help</code> | Shows help for a module. |
| <code>module swap</code> | Swaps a currently loaded module for an unloaded module. |

Available Modules

To see a list of available modules, type

```
module avail
```

Note: If you need a module that is not available, please [contact us](#).

You can check for the existence of a module and its versions using `module avail <module-name> .`

```
$ module avail cuda
----- /software/dev_tools/swtree/cs400/modulefiles -----
cuda/6.5      cuda/7.5(default)      cuda/8.0
```

Working with Modules

When you load a module, your environment is modified to use a specific software package. To load a module:

```
module load vmd
```

To verify your module has loaded, you can type `module list` .

To display information about the attributes of the module such as the size of the module, the compiler or the source from which the module was created, etc., use the following command:

```
module display your_module
```

Removing and Switching Modules

Unloading a module will avoid conflict and/or messages of failure due to different versions or dependencies.

```
module unload PE-gnu/1.0
```

Switching between different module versions can accomplish the task of having to load, unload and load modules in multiple steps. In the following example, `cuda/7.5` is currently loaded. After running the command, `cuda/7.5` is *unloaded* and `cuda/8.0` is *loaded*.

```
module switch cuda/7.5 cuda/8.0
```

You can unload all the modules on your environment, by executing the module purge command:

```
module purge
```

Related Information

- [Environment Customization](#)

Compiler Toolchains on SHPC Condos

SHPC supports four *programming environment (PE)* modules to easily switch between compilers. Each programming environment contains the full set of compatible compilers and libraries.

These compilers are: [GNU Collection Compiler \(GCC\)](#), the [Intel compiler](#), [The Portland Group \(PGI\)](#), and the [Numerical Algorithms Group \(NAG\)](#).

Note: You cannot use more than one `PE-module` at the same time. For example, if you are working with GNU and then you decide to work with the Intel compiler, first unload the `PE-gnu` module and then load `PE-intel`.

The GNU Compiler Suite

To load the GNU module:

```
module load PE-gnu
```

You can check which modules are loaded in your system by typing:

```
$ module list
Currently Loaded Modulefiles:
  1) gcc/5.3.0      2) openmpi/1.10.3  3) xalt/0.7.5      4) PE-gnu/1.0
```

To display information about the module, such as the size, the compiler, or the source from which the module was created, etc., use the following command:

```
$ module display PE-gnu
-----
/software/dev_tools/swtree/cs400/modulefiles/PE-gnu/1.0:

module-whatis    PE-gnu defines the environment needed to build

                  applications using GNU compiler suites on this system.
conflict        PE-gnu PE-intel PE-pgi
setenv          PE_NAME GNU
setenv          PE_CC mpicc
setenv          PE_CXX mpic++
setenv          PE_FORTRAN mpif90
prepend-path    PATH /software/dev_tools/swtree/cs400-centos7.2_pe2016-08/PE/1.0/noarch/bin
module          load xalt
-----
```

You can switch between the two versions of PE-gnu v1.0 and PE-gnu v2.0:

```
$ module switch PE-gnu/1.0 PE-gnu/2.0
$ module list
Currently Loaded Modulefiles:
  1) gcc/5.3.0      2) openmpi/2.1.1   3) PE-gnu/2.0      4) xalt/0.7.5
```

The Intel Compiler Suite

If you are working with another module, first you need to unload it.

```
module load PE-intel
```

You can see what the module provides with the commands `module list` and `module display`.

```
$ module list
Currently Loaded Modulefiles:
  1) intel/16.0.1    2) openmpi/1.10.3  3) xalt/0.7.5      4) PE-intel/1.0
```

```
module display PE-intel
-----
/software/dev_tools/swtree/cs400/modulefiles/PE-intel/1.0:

module-whatis    PE-intel defines the environment needed to build

                  applications using Intel compiler suites on this system.
conflict        PE-gnu PE-intel PE-pgi
setenv          PE_NAME INTEL
setenv          PE_CC mpicc
setenv          PE_CXX mpic++
setenv          PE_FORTRAN mpif90
prepend-path    PATH /software/dev_tools/swtree/cs400_centos7.2_pe2016-08/PE/1.0/noarch/bin
module          load xalt
-----
```

The Portland Group Compiler Suite

If you are working with another module, first you need to unload it.

```
module load PE-pgi
```

You can see what does the module provides with the commands `module list` and `module display`.

```
$ module list
Currently Loaded Modulefiles:
  1) pgi/15.7.0     2) openmpi/1.10.3  3) xalt/0.7.5      4) PE-pgi/1.0
```

```
$ module display PE-pgi
-----
/software/dev_tools/swtree/cs400/modulefiles/PE-pgi/1.0:

module-whatis    PE-pgi defines the environment needed to build

                  applications using PGI compiler suites on this system.
conflict        PE-gnu PE-intel PE-pgi
setenv          PE_NAME PGI
setenv          PE_CC mpicc
setenv          PE_CXX mpic++
setenv          PE_FORTRAN mpif90
prepend-path    PATH /software/dev_tools/swtree/cs400_centos7.2_pe2016-08/PE/1.0/noarch/bin
module          load xalt
-----
```

The Numerical Algorithm Group Compiler Suite

If you are working with another module, first you need to unload it.

```
module load PE-nag
```

You can see what the module provides with the commands `module list` and `module display`.

```
$ module list
Currently Loaded Modulefiles:
  1) nag/6.0      2) mpich/3.2    3) xalt/0.7.5   4) PE-nag/1.0
```

```
$ module display PE-nag
-----
/software/dev_tools/swtree/cs400/modulefiles/PE-nag/1.0:

module-whatis    PE-nag defines the environment needed to build

                  applications using NAG Fortran compiler on this system.
conflict        PE-gnu PE-intel PE-pgi
setenv          PE_NAME NAG
setenv          PE_CC mpicc
setenv          PE_CXX mpic++
setenv          PE_FORTRAN mpif90
prepend-path    PATH /software/dev_tools/swtree/cs400_centos7.2_pe2016-08/PE/1.0/noarch/bin
module          load xalt
-----
```

Related Information

- [Environment Customization](#)
- [Modules](#)

Running Scientific Computational Workflows

Overview

Workflows offer benefits of automation and efficient orchestration (eg. data parallel execution) of multi-stage computation. Furthermore, they are powerful reproducibility and portability tools for science and engineering applications.

Typically, a workflow is written in a high level language that is offered and understood by a workflow management software or simply a workflow tool.

Workflow tools available on Condos

We currently offer support for the following workflow tools on SHPC:

1. Nextflow
2. Makeflow
3. Swift

A brief description about each of the aforementioned workflow tools is provided below:

Nextflow

[Nextflow](#) is a favored workflow tool among Singularity container users. Similarly, it is popular among users from the Biosciences domain.

Makeflow

The [Makeflow](#) workflow system uses a Makefile like language to define workflows that may be deployed and executed over clusters and clouds.

Swift

[Swift](#) uses a C-like syntax to define workflows. Swift is capable of stitching computational steps defined in the workflow as a true HPC workflow that uses the Message Passing Paradigm of parallel computation using the MPI libraries and its own load balancer.

Note: While Nextflow and Makeflow require additional configuration if you wish to run them on compute nodes, Swift can run directly on compute nodes by simply plugging it into a job definition script just like any other MPI application.

Example Workflows

Hello World

Nextflow

```
#!/usr/bin/env nextflow

params.str = 'Hello world!'

process splitLetters {

    output:
    file 'chunk_*' into letters mode flatten

    """
```

```

    printf '${params.str}' | split -b 6 - chunk_
    ""
}

process convertToUpper {

    input:
    file x from letters

    output:
    stdout result

    ""
    cat $x | tr '[a-z]' '[A-Z]'
    ""
}

result.subscribe {
    println it.trim()
}

```

Save the above code in a file, eg. `hello.nf`. To run the workflow on open condo login node, do the following:

```

$ module purge
$ module load PE-gnu
$ module load java/1.8.0_131
$ module load nextflow
$ nextflow run hello.nf

```

You should see output similar to the following:

```

N E X T F L O W ~ version 0.27.6
Launching `nextflow_example.nf` [insane_meucci] - revision: 5319db7b93
[warm up] executor > local
[f9/cb98ba] Submitted process > splitLetters
[94/6ed3f3] Submitted process > convertToUpper (1)
[cb/506a85] Submitted process > convertToUpper (2)
HELLO
WORLD!

```

Makeflow

A "Hello World" in Makeflow would look something like so:

```

ECHO=/bin/echo

hello.txt:
    $ECHO 'Hello World!' > hello.txt

```

Save the above code in a file, say `hello.mkf` and run it on the open condo like so:

```

$ module load PE-gnu
$ module load cctools/6.2.7
$ makeflow hello.mkf

```

If all goes well, the output should look like so:

```

parsing hello.mkf...
local resources: 32 cores, 128833 MB memory, 6893119 MB disk
max running local jobs: 32

```

```
checking hello.mkf for consistency...
hello.mkf has 1 rules.
recovering from log file hello.mkf.makeflowlog...
makeflow: hello.txt is reported as existing, but does not exist.
starting workflow...
submitting job: /bin/echo 'Hello World!' > hello.txt
submitted job 123822
job 123822 completed
nothing left to do.
```

And you should see a new file called `hello.txt` in your current working directory.

Swift

A Swift Hello World workflow looks like so:

```
import io;

printf("Hello world");
```

Swift uses two steps to workflow execution: compile and run.

Load the swift module on condo like so:

```
$ module purge
$ module load PE-gnu
$ module load java/1.8.0_131 mpich/3.2
$ module load swift
```

Compile and run the workflow like so:

```
$ stc hello.swift
```

The above step will produce a TCL file called `hello.tic`. Run the TCL file like so:

```
turbine -n 2 hello.tic
```

If all goes well, you should see the following output:

```
Hello world
```

General remarks

1. Note that the above workflows will run on login nodes. In order for them to run over compute nodes, more configuration is needed.
2. Note that `nextflow` expects absolute paths for data and executables since it works in its own temp directory. Please adjust the paths to where you choose to run the workflow.

Where to go from here?

- Use the [Crystal Workflow](#) with these workflow tools.
- Come talk to us at CADES if you think one or more of your applications will benefit with the help of the aforementioned workflow tools.

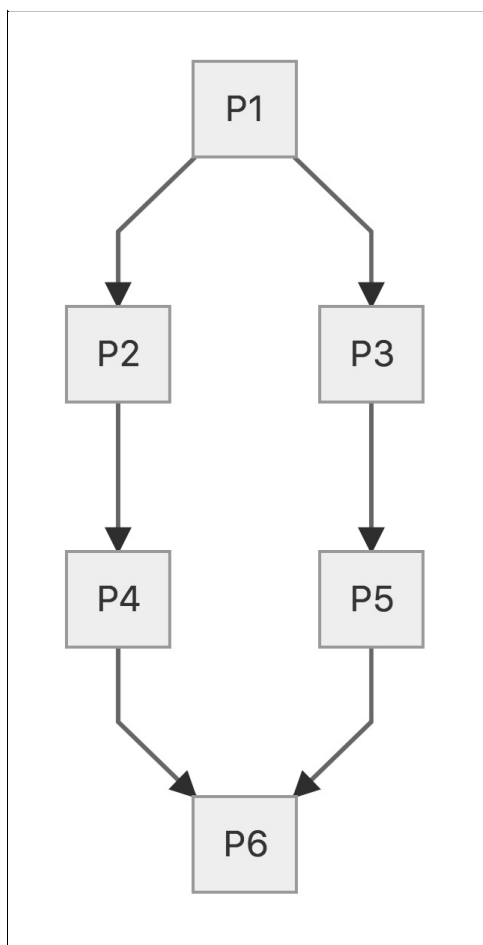
Crystal Workflow

CrystalFlow is a hypothetical workflow with low-medium complexity that adequately illustrates the benefits and characteristics of scientific computational workflows.

Note: The code, executables, and test data for the crystal workflow is available on ORNL's public [GitLab](#).

About the Crystal Workflow

The workflow is a crystal shaped graph as shown in the figure below.



In the above workflow, each of the boxes represent process and arrows represent the dependency between connected processes. For example, process `P1` produces a data file that is consumed by processes `P2` and `P3`.

Each of the 6 processes of this workflow are implemented in C and bash. Either may be used in the workflows shown below. The initial input file is pre-prepared. The code and data for these workflows are available on open SHPC condo at `/software/T/B/D`. A `Makefile` will build the C executables found in the directory named `c`. The following sections show how this workflow may be orchestrated using each of the three workflow management tools.

Nextflow

The following code snippet shows how the above workflow would be expressed in Nextflow.

```
#!/usr/bin/env nextflow

in1 = file('/home/km0/crystalworkflow/shell/inputs/in1.txt')

process p1 {
  input:
  file in1
  output:
  file 'out1.txt' into out1
  """
  ~/crystalworkflow/shell/p1/p1.sh $in1 'out1.txt'
  """
}

process p2 {
  input:
  file out1
  output:
  file 'out2.txt' into out2
  """
  ~/crystalworkflow/shell/p2/p2.sh $out1 'out2.txt'
  """
}

process p3 {
  input:
  file out1
  output:
  file 'out3.txt' into out3
  """
  ~/crystalworkflow/shell/p3/p3.sh $out1 'out3.txt'
  """
}

process p4 {
  input:
  file out2
  output:
  file 'out4.txt' into out4
  """
  ~/crystalworkflow/shell/p4/p4.sh $out2 'out4.txt'
  """
}

process p5 {
  input:
  file out3
  output:
  file 'out5.txt' into out5
  """
  ~/crystalworkflow/shell/p5/p5.sh $out3 'out5.txt'
  """
}

process p6 {
  input:
  file out4
  file out5
  output:
  file 'out6.txt' into out6
  """
  ~/crystalworkflow/shell/p6/p6.sh $out4 $out5 'out6.txt'
  """
}
```

Assuming the above workflow is saved in a file named `crystal.nf`, it could be run as follows:

```
$ module purge
$ module load PE-gnu
```

```

$ module load java/1.8.0_131
$ module load nextflow
$
$ nextflow run crystal.nf

N E X T F L O W ~ version 0.27.6
Launching `crystal.nf` [thirsty_allen] - revision: e3b42d107d
[warm up] executor > local
[db/d513da] Submitted process > p1
[89/e16494] Submitted process > p2
[c3/9d4ddd] Submitted process > p3
[0d/5406b9] Submitted process > p4
[cf/4b94bb] Submitted process > p5
[c2/3bae00] Submitted process > p6

```

Makeflow

The following code snippet shows how the crystal workflow would be implemented using Makeflow.

```

P1=../shell/p1/p1.sh
P2=../shell/p2/p2.sh
P3=../shell/p3/p3.sh
P4=../shell/p4/p4.sh
P5=../shell/p5/p5.sh
P6=../shell/p6/p6.sh

../shell/p1/out1.txt:
    $P1 ../shell/inputs/in1.txt ../shell/p1/out1.txt
../shell/p2/out2.txt:
    $P2 ../shell/p1/out1.txt ../shell/p2/out2.txt
../shell/p3/out3.txt:
    $P3 ../shell/p1/out1.txt ../shell/p3/out3.txt
../shell/p4/out4.txt:
    $P4 ../shell/p2/out2.txt ../shell/p4/out4.txt
../shell/p5/out5.txt:
    $P5 ../shell/p3/out3.txt ../shell/p5/out5.txt
../shell/outputs/out6.txt:
    $P6 ../shell/p4/out4.txt ../shell/p5/out5.txt ../shell/outputs/out6.txt

```

Assuming the above workflow is saved in a file named `crystal.mkf`, it could be executed like so:

```

$ module purge
$ module load PE-gnu
$ module load cctools/6.2.7
$ makeflow crystal.mkf
parsing crystal.mkf...
local resources: 32 cores, 128833 MB memory, 6593404 MB disk
max running local jobs: 32
checking crystal.mkf for consistency...
crystal.mkf has 6 rules.
starting workflow...
submitting job: ../shell/p6/p6.sh ../shell/p4/out4.txt ../shell/p5/out5.txt ../shell/outputs/out6.txt
submitted job 37132
submitting job: ../shell/p5/p5.sh ../shell/p3/out3.txt ../shell/p5/out5.txt
submitted job 37133
submitting job: ../shell/p4/p4.sh ../shell/p2/out2.txt ../shell/p4/out4.txt
submitted job 37134
submitting job: ../shell/p3/p3.sh ../shell/p1/out1.txt ../shell/p3/out3.txt
submitted job 37135
submitting job: ../shell/p2/p2.sh ../shell/p1/out1.txt ../shell/p2/out2.txt
submitted job 37136
submitting job: ../shell/p1/p1.sh ../shell/inputs/in1.txt ../shell/p1/out1.txt
submitted job 37137
cat: ../shell/p3/out3.txt: No such file or directory
p3 completed.

```

```

p5 completed.
p4 completed.
job 37135 completed
p1 completed.
p2 completed.
job 37134 completed
job 37133 completed
job 37136 completed
job 37137 completed
p6 completed.
job 37132 completed
nothing left to do.

```

Swift

The following code snippet shows the Swift implementation of the crystal workflow. Note that the Swift implementation invokes the C version of executables but it can equally invoke the bash version.

```

import io;

app (file out) p1 (file inp){ "../c/p1/p1" inp out }

app (file out) p2 (file inp){ "../c/p2/p2" inp out }

app (file out) p3 (file inp){ "../c/p3/p3" inp out }

app (file out) p4 (file inp){ "../c/p4/p4" inp out }

app (file out) p5 (file inp){ "../c/p5/p5" inp out }

app (file out) p6 (file inp1, file inp2){ "../c/p6/p6" inp1 inp2 out }

file in1 = input("../c/inputs/in1.txt");

file out1 <"../c/p1/out1.txt"> = p1(in1);
file out2 <"../c/p2/out2.txt"> = p2(out1);
file out3 <"../c/p3/out3.txt"> = p3(out1);
file out4 <"../c/p4/out4.txt"> = p4(out2);
file out5 <"../c/p5/out5.txt"> = p5(out3);
file out6 <"../c/outputs/out6.txt"> = p6(out4,out5);

```

Assuming the above program is saved in a file called `crystal.swift`, it may be run on Open SHPC like so:

```

$ module purge
$ module load PE-gnu
$ module load java/1.8.0_131 mpich/3.2
$ module load swift

$ stc crystal.swift

$ turbine -n 2 crystal.tic
../c/p1/out1.txt
../c/p3/out3.txt
../c/p2/out2.txt
../c/p5/out5.txt
../c/p4/out4.txt
../c/outputs/out6.txt

```


How to Use the SHPC Condo

This section will walk you through the primary steps that are required to get you started using the SHPC resources. If, at any time, you have trouble, do not hesitate to reach out to us via [email](#).

- [Prerequisites](#)
- [Request Credentials for an Allocation](#)
- [Access to your Allocation](#)
- [Execute a Job on an Allocation](#)

[CADES](#) → [User Documentation](#) → [SHPC Condo User Guide](#) → [How to Use](#) → [Prerequisites](#)

Prerequisites

To properly utilize SHPC Condos, you will need a couple of utilities loaded on your local machine. These utilities are free and widely used for this type of application.

- Required: **SSH client**
- Recommended: **Bash terminal**

Note: CADES does not provide support for getting these utilities up and running on your personal computer.

MacOS and Linux

Both macOS and Linux distributions includes a Bash terminal and an SSH client by default. No additional software should be required to access SHPC Condos.

Windows Users

Windows does not have a native SSH client or a native Bash terminal. A few solutions are linked below.

- Option 1: [PuTTY](#) - SSH client and Bash environment for Windows.
- Option 2: [Git Bash](#) – Part of the Git for Windows environment includes Git Bash, which provides a light weight ssh client.
- Option 3: [Cygwin](#) – If you wish to have Bash-style functions on your Windows machine, then you should consider installing Cygwin, which ports the Portable Operating System Interface (POSIX) system calls and environment to Windows.

Request Access for SHPC Condo Allocation

You can self-request access with the procedure below. If you are unsure which group to request, or otherwise need assistance, please [contact us](#).

1. Use the appropriate group link below and enter your email address (your ORNL address, if available) and click `Continue`.
2. Review the XCAMS user agreement, and select `Agree`.
3. Enter your UCAMS ID (or a new XCAMS user name).
4. Enter your UCAMS password (or a new XCAMS password).
5. Click `submit` to complete the XCAMS request.
6. The activation notice will be dispatched to the email address entered above. This process can take up to 24 hours to complete.

SHPC Condo Groups

To access SHPC Condos, users will need to be added into an appropriate group. Find your group in the table below, and click on the respective UCAMS/XCAMS registration URL, as outlined in the instructions above. *NSED access exists in the Moderate Protection Zone.*

| Division Name | Division Approver | PBS Directives |
|---|---|---|
| Computing and Computational Sciences Directorate (CCSD) | Jayson Hines (hinesjb@ornl.gov) | <code>#PBS -W group_list=cades-ccsd</code> <code>#PBS -A ccsd</code> |
| Spallation Neutron Source (SNS) | A.J. (Timmy) Ramirez-Cuesta (ramirezcueaj@ornl.gov) | <code>#PBS -W group_list=cades-virtues</code> <code>#PBS -A sns</code> |
| Center for Nanophase Materials Sciences (CNMS) | Bobby Sumpter (sumpterb@ornl.gov) | <code>#PBS -W group_list=cades-cnms</code> <code>#PBS -A cnms</code> |
| Climate Change Science Institute (CCSI) | Dali Wang (wangd@ornl.gov) | <code>#PBS -W group_list=cades-ccsi</code> <code>#PBS -A ccsi</code> |
| Energy Dissipation to Defect Evolution (EDDE) | Malcolm Stocks (stocksgm@ornl.gov) | <code>#PBS -W group_list=cades-edde</code> <code>#PBS -A edde</code> |
| Biosciences Division (BSD) | Bob Cottingham (cottinghamrw@ornl.gov) | <code>#PBS -W group_list=cades-bsd</code> <code>#PBS -A bsd</code> |
| Nuclear Science and Engineering Directorate (NSED) Also, please send an email to notify the CADES team . | Jeff Banta (bantajp@ornl.gov) | <code>#PBS -W group_list=cades-nsed</code> <code>#PBS -A nsed</code> |

Note: If you do not see your group listed, please [contact the CADES team](#) and include:

- **Subject:** Help with SHPC Condo Registration
- **Email body:** UCAMS ID or XCAMS ID, contact information, reason for requesting an SHPC Condo allocation, and the name of your directorate and division.

Access Your SHPC Condo Allocation

After your [access request](#) has been approved and you have installed the [prerequisites](#), you can log in to the Open Protection Zone or the Moderate Protection Zone.

Open Protection Zone

1. Open a Bash terminal (or see [here](#) if you need more help).
2. Execute `ssh ucams@or-condo-login.ornl.gov` . Replace `ucams` with your *UCAMS/XCAMS ID*.
3. When prompted, enter your password.

Moderate Protection Zone

1. Open a Bash terminal (or see [here](#) if you need more help).
2. Execute `ssh ucams@mod-condo-login.ornl.gov` . Replace `ucams` with your *UCAMS/XCAMS ID*.
3. When prompted, enter your password.

By default, `/home` directories should be automatically created for you when logging into SHPC Condos.

You can run the following command on your terminal to see your files:

```
# ls -lhr /home/user
total 20K
drwxr-xr-x 2 user users 4 Apr 6 12:11 Test1
-rw-r--r-- 1 user users 982 Apr 6 12:11 setup.py
-rw-r--r-- 1 user users 1.5K Apr 6 12:11 readme.txt
-rw-r--r-- 1 user users 77 Apr 6 12:11 paralleltestpy2.py
```

Replace the word `user` with your *UCAMS/XCAMS ID*.

- The `ls -lhr /home/user` command will show the whole list and details of the files that a **user** has.

Execute a Job on Your SHCP Condo Allocation

The tutorial below shows you how to run Wes Kendall's basic "hello world" program, written in C, using the message passing interface (MPI) to scale across the SHPC Condo compute nodes [1]. This tutorial is intended for users who are new to the SHPC Condo environment and leverages a portable batch system (PBS) script and a C source code.

Note: Do not execute jobs on the login nodes; only use the login nodes to access your compute nodes. Processor-intensive, memory-intensive, or otherwise disruptive processes running on login nodes will be killed without warning.

Prerequisites

- Access to the login node (XCAMS/UCAMS authorization).
- A PBS script that specifies your conditions/variables and calls the binary/script you would like to execute on the compute nodes.
- A Binary file, source, or script for the problem that you would like to run on the compute nodes.

Files

The steps below walk you through building a PBS script and compiling a C binary from source. However, if you wish to download the files used in this to tutorial files directly, you can do so using the links below.

- [hello-world.pbs](#) – PBS script used for the hello world batch job
- [myMPIhw.c](#) – hello world C source code
- [hello-world](#) – hello world C binary (already compiled)

Step 1: Connect to Your Allocation

Open and Moderate protection zones each have their own login node. Choose the login node for your protection zone.

Note: The Open protection zone can be accessed either using either XCAMS or UCAMS credentials. However, the Moderate protection zone requires an ORNL UCAMS ID.

Open Protection Zone

1. Open a Bash terminal (or PuTTY for Windows users).
2. Execute `ssh xcams@or-condo-login.ornl.gov`.
 - Replace "xcams" with your XCAMS or UCAMS ID.
3. When prompted, enter your XCAMS or UCAMS password.

Moderate Protection Zone

1. Open a Bash terminal (or PuTTY for Windows users).
2. Execute `ssh ucams@mod-condo-login.ornl.gov`.
 - Replace "ucams" with your UCAMS ID.

3. When prompted, enter your UCAMS password.

Once you have connected to the login node, you can proceed to Step 2 and begin assembling your PBS script.

Step 2: Create Your PBS Script

Below is the PBS script we are using to run an MPI "hello world" program as a batch job. PBS scripts use variables to specify things like the number of nodes/cores used to execute your job, estimated wall time for your job, and which compute resources to use (e.g., GPU vs. CPU). The sections below feature an example PBS script for SHPC Condo resources, show you how to create/save your own PBS script, and show you how store the PBS script on an SHPC Condo file system.

Check out the [official Torque documentation](#) for a complete list of PBS variables.

Example PBS script

Here is an example PBS script for running a batch job on a SHPC Condo allocation. We break down each command in the section below.

```
#!/bin/bash
#PBS -N MyMPIhw
#PBS -M YourEmailHere@ornl.gov
#PBS -l nodes=1:ppn=16
#PBS -l walltime=0:00:2:0
#PBS -W group_list=cares-birthright
#PBS -A birthright
#PBS -l qos=std
#PBS -q gpu
module purge
module load PE-gnu
module list
cd $PBS_O_WORKDIR
pwd
mpirun hello-world
```

Download this script (with explanatory comments) [here](#).

PBS Script Breakdown

Here, we break down the essential elements of the above PBS script.

First, we're going to set the script type: Bash.

```
#!/bin/bash
```

Next, you need to set the job name. Make it short and simple because your output files will share this name. We're going with MyMPIhw.

```
#PBS -N MyMPIhw
```

You can add your email address if you would like errors to be emailed to you directly.

```
#PBS -M YourEmail@ornl.gov
```

Set your node spec, including the number of nodes and processors per node that you want to use to run your job. In this case, we're using one node and 16 cores per node.

```
#PBS -l nodes=1:ppn=16
```

Tell PBS the anticipated runtime for your job, where `walltime=HH:MM:S`. The example below has the walltime set to 2 minutes.

```
#PBS -l walltime=0:00:2:0
```

Specify your LDAP group. The full list of SHPC Condo LDAP groups is [here](#). We're using `ca-des-birthright` in this case.

```
#PBS -W group_list=ca-des-birthright
```

You also need to specify your account type. We're also using `birthright` in this case.

```
#PBS -A birthright
```

Now we can set the quality of service (QOS). We can set this to `burst` or `std`.

Burst jobs allow a user to leverage more nodes/cores/GPUs than may be in their formal allocation. However, in exchange for this "resource burst" flexibility, your burst job may be preempted if the rightful owner of those resources needs them to complete his or her own jobs.

In most cases, a user will simply run a job with the QOS set to `std`.

```
#PBS -l qos=std
```

All SHPC Condo nodes have GPUs. Since this example uses a SHPC allocation, we're going to specify the use of GPUs.

```
#PBS -q gpu
```

Next we need to load the modules required for executing our batch job. First thing we'll do is clear any modules currently loaded that might result in a conflict.

```
module purge
```

With a clean slate, we can now load our programming environment using `module load`. For this particular example, all we need is the `PE-gnu` module, which loads OpenMPI, GCC, and XALT.

```
module load PE-gnu
```

The next line confirms the modules that were loaded.

```
module list
```

With our environment loaded, the PBS script now sets the working path. In this example, our binary will be launched from the same directory as our PBS script. The results from the binary will also be placed here.

```
cd $PBS_O_WORKDIR
```

Confirms current working directory.

```
pwd
```

Finally, the last line of the PBS script calls MPI to run our `hello-world` binary. You can replace "hello-world" with the file name of whatever binary you wish to execute on the compute nodes.

```
mpirun hello-world
```

Procedure

Now that we've covered the basics of a PBS script in the context of an SHPC Condo, let's talk about actually creating and using the script on your allocation.

When creating and editing your PBS script, you have two basic options. **Option 1:** Create and edit your PBS script on your local machine and upload it to the Lustre path using secure copy (`scp`). **Option 2:** Create and edit your PBS script directly on the compute node (from Lustre storage) using Vi.

Option 1: Create PBS Locally and Upload to Lustre Storage

1. Using your favorite text editor, create your PBS script on your local machine or [download the pre-made example script](#).
2. Use the `scp` command to copy your PBS script from the source machine to the SHPC file system (Lustre storage in the Open protection zone in this case).

```
scp /path/to/hello-world.pbs xcams@or-condo-login.ornl.gov:/lustre/or-hydra/cades-birthright/xcams/hello-world.pbs
```

Replace `"/path/to/hello-world.pbs"` with the path of your script, and replace `"xcams"` with your XCAMS/UCAMS ID.

You may be prompted for your XCAMS/UCAMS password if you have not copied your public SSH key to the login node.

With the PBS script in place, you can now move on to [compiling your hello world C code](#).

Option 2: Create Your PBS Script Directly on the Compute Node

1. From the login node, change your working directory to the desired file system. We're going to use our Lustre allocation for this example.

```
cd /lustre/or-hydra/cades-birthright/ucams
```

Replace `"ucams"` with your own UCAMS/XCAMS user ID.

2. Use Vi to create and edit your PBS script.

```
vi hello-world.pbs
```

3. Create your PBS script within Vi or paste the contents of your PBS script into Vi.
4. When finished, hit `Esc` on your keyboard to exit the input mode.
5. Enter `:x!` into Vi's command line, and press `Return` to save your file and return to the Bash shell.

With the PBS script in place, you can now move on to [compiling your hello world C code](#).

Step 3: Compile the C Program from Source

Below is Wes Kendall's simple "hello world" C program that utilizes MPI to run the job in parallel [1]. We will need to compile this source code on one of the compute nodes (you can also download the compiled binary below if you prefer).

MPI Hello World Source Code

```
#include <mpi.h>
#include <stdio.h>
int main(int argc, char** argv) {
    // Initialize the MPI environment.
    MPI_Init(NULL, NULL);
    // Get the number of processes.
    int world_size;
    MPI_Comm_size(MPI_COMM_WORLD, &world_size);
    // Get the rank of the process.
    int world_rank;
    MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);
    // Get the name of the processor.
    char processor_name[MPI_MAX_PROCESSOR_NAME];
    int name_len;
    MPI_Get_processor_name(processor_name, &name_len);
    // Print off a hello world message.
    printf("Hello world from processor %s, rank %d"
           " out of %d processors\n",
           processor_name, world_rank, world_size);
    // Finalize the MPI environment.
    MPI_Finalize();
}
```

Download this source [here](#).

Download the compiled binary [here](#).

Procedure

When creating and editing your `hello-world.c` source code, you have two basic options. **Option 1:** Create and edit your source code on your local machine and upload it to the Lustre path using secure copy (`scp`). Compile the binary on the compute node.

Option 2: Create and edit your source code directly on the compute node (from Lustre storage) using Vi. Compile the binary on the compute node.

Option 1: Create Source Code Locally and Upload to Lustre Storage

1. Using your favorite text editor, create and edit your `hello-world.c` source code on your local machine or download the pre-made source [here](#).
2. Use the `scp` command to copy your `hello-world.c` source code from your local machine to the SHPC file system (Lustre storage in the Open protection zone in this case).

```
scp /path/to/hello-world.c xcams@or-condo-login.ornl.gov:/lustre/or-hydra/cades-birthright/xcams/hello-world.c`
```

Replace `"/path/to/hello-world.c"` with the path of your source code, and replace `"xcams"` with your XCAMS/UCAMS ID.

You may be prompted for your XCAMS/UCAMS password if you have not copied your public SSH key to the login node.

3. Load the MPI compiler using the PE-gnu module.

```
module load PE-gnu
```

4. Compile the C source into a binary.

```
mpicc -o hello-world hello-world.c
```

5. Use `ls -al` to verify the presence of the `hello-world` binary in your working directory.

With the C code compiled into a binary (`hello-world`), we can now schedule and run the job on our compute nodes.

Option 2: Create and Edit Source Code Directly on Compute Node

1. Ensure that you are still in your working directory (`/lustre/or-hydra/cades-birthright/ucams`) using `pwd` .
2. Use Vi (`vi`) to create your C source file within your working directory.

```
vi hello-world.c
```

3. Paste the hello world C code into Vi.
 - o Hit `Esc` on your keyboard to exit the input mode.
 - o Enter `:set paste` into Vi's command line, and press `Return` to enter paste mode.
 - o Paste the C code into Vi.
4. When finished, hit `Esc` on your keyboard to exit the paste/input mode.
5. Enter `:x!` into Vi's command line, and press `Return` to save your file and return to the Bash shell.
You now have a C source file that you can compile.
6. Load the MPI compiler using the PE-gnu module.

```
module load PE-gnu
```

7. Compile the C source into a binary.

```
mpicc -o hello-world hello-world.c
```

8. Use `ls -al` to verify the presence of the `hello-world` binary in your working directory.

With the C code compiled into a binary (`hello-world`), we can now schedule and run the job on our compute nodes.

Step 4: Run the Job

1. Before proceeding, ensure that you are still in your working directory (using `pwd`) and that you still have the PE-gnu module loaded (using `module list`).
 - o We need to be in the same path/directory as our PBS script and our C binary. Use `ls -al` to confirm their presence.
 - o PE-gnu also loads OpenMPI, GCC, and XALT. Use `module list` to confirm their presence. If necessary, use `module load PE-gnu` to reload the module(s).
2. Use `qsub` to schedule your batch job in the queue.

```
qsub hello-world.pbs
```

This command will automatically queue your job using Torque and produce a six-digit job number (shown below).


```
143295.or-condo-pbs01
```

You can check the status of your job at any time with the `checkjob` command.

```
checkjob 143295
```

You can also stop your job at any time with the `qdel` command.

```
qdel 143295
```

3. View your results.

Once your job completes, Torque will produce two output/data files. These output/data files, unless otherwise specified in the PBS script, are placed in the same path as your binary.

One file (`myscript.ojobnumber`) contains the results of the binary you just executed, and the other (`myscript.ejobnumber`) contains any errors that occurred during execution.

Replace "myscript" with the name of your script and "jobnumber" with your job number.

You can view the contents of these files using the `more` command followed by the file name.

```
more MyMPIhw.o143295
```

Your output should look something like this, with one line per processor core (16 in this case):

```
Hello world from processor or-condo-c136.ornl.gov, rank 3 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 4 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 6 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 11 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 7 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 14 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 2 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 5 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 8 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 9 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 10 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 12 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 13 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 15 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 0 out of 16 processors
Hello world from processor or-condo-c136.ornl.gov, rank 1 out of 16 processors
```

4. Download your results (using the `scp` command or an SFTP client) or move them to persistent storage.

Works Cited

1. Wes Kendall, "MPI Hello World," *MPI Tutorial*, accessed June 14, 2017, <http://mpitutorial.com/tutorials/mpi-hello-world/>.

[CADES](#) → [User Documentation](#) → [User-Provided Tutorials for CADES Cloud](#) → [Launch a Docker Container](#)

Tutorial contributed by Drew Schmidt. **Note:** User-provided tutorials are not supported by CADES.

Background

[Docker](#) is a [container](#) architecture and ecosystem. A [linux.com](#) [article](#) nicely summarizes Docker as follows:

Docker is a tool that can package an application and its dependencies in a virtual container that can run on any Linux server. This helps enable flexibility and portability on where the application can run, whether on premises, public cloud, private cloud, bare metal, etc.

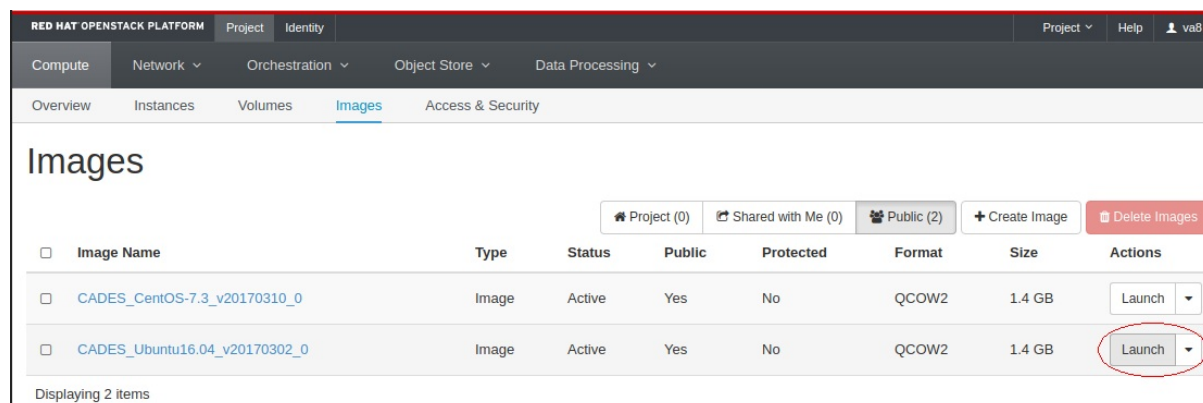
Containers have somewhat similar goals to a virtual machine (vm). However, a Docker container is not a vm. You are probably aware that vm's have some performance overhead compared to running things natively. However, it is worth noting that the applications that run inside of Docker containers actually run *natively*. Your Docker containers share the kernel with their host operating system. So there is no double overhead in running a container inside our vm. However, we still suffer some performance penalty by having virtualized in the first place.

Many of the applications you will be interested in deploying are already configured for very easy use with Docker. You can find public repositories of many of your favorite applications set up on [Docker Hub](#).

First Steps

We will assume that you are already reasonably familiar with the [CADES Cloud](#) system. If not, consider reading the [user documentation](#). Make sure you read the section titled: [Launch a VM Instance from an Image](#)

However, you get there, launch an Ubuntu 16.04 vm.



The screenshot shows the 'Images' page in the Red Hat OpenStack Platform. The page has a navigation bar with 'Project' and 'Identity' tabs, and a sub-navigation bar with 'Compute', 'Network', 'Orchestration', 'Object Store', and 'Data Processing'. The 'Images' section is active, showing a table of images. The table has columns for 'Image Name', 'Type', 'Status', 'Public', 'Protected', 'Format', 'Size', and 'Actions'. Two images are listed: 'CADES_CentOS-7.3_v20170310_0' and 'CADES_Ubuntu16.04_v20170302_0'. The 'Launch' button for the Ubuntu image is circled in red.

| Image Name | Type | Status | Public | Protected | Format | Size | Actions |
|-------------------------------|-------|--------|--------|-----------|--------|--------|---------|
| CADES_CentOS-7.3_v20170310_0 | Image | Active | Yes | No | QCOW2 | 1.4 GB | Launch |
| CADES_Ubuntu16.04_v20170302_0 | Image | Active | Yes | No | QCOW2 | 1.4 GB | Launch |

For now, you can keep the setup very basic, just following the instructions outlined in the "Launch a VM Instance" page linked above.

You may eventually be interested in more complicated configurations. For example, you may need to modify the security group details (say, for example, you want to run a docker container that runs a web server). For now, we will ignore those details. However, we provide a more complicated example in our Shiny tutorial.

Install Docker

Next, you need to [ssh to your new vm](#). I named my instance `t1`, so when I login, it shows me as `ca-des@t1`. Your prompt will show `ca-des@whatever-you-named-your-vm`.

Having logged in, it's time to install Docker. The official Docker documentation [provides a lot of useful information](#) to this end. Below we summarize only the steps outlined in that article. If you wish to understand an individual step or if something goes wrong, please refer to the article.

Otherwise, run:

```
sudo apt-get install \
  apt-transport-https \
  ca-certificates \
  curl \
  software-properties-common

curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo apt-key add -

sudo apt-key fingerprint 0EBFCD88

sudo add-apt-repository \
  "deb [arch=amd64] https://download.docker.com/linux/ubuntu \
  $(lsb_release -cs) \
  stable"

sudo apt-get update
sudo apt-get install -y docker docker.io
```

And if all is well, you should have Docker installed on your vm.

Run a Test Container

While still ssh'd to your CADES Cloud vm, you can test that your setup is working correctly by running:

```
sudo docker pull hello-world
sudo docker run hello-world
```

If all goes well, you will have a small "hello world"-like output and return to your terminal, and should look something like this:

```
ca-des@t1:~$ sudo docker run hello-world

Hello from Docker!
This message shows that your installation appears to be working correctly.

To generate this message, Docker took the following steps:
 1. The Docker client contacted the Docker daemon.
 2. The Docker daemon pulled the "hello-world" image from the Docker Hub.
 3. The Docker daemon created a new container from that image which runs the
    executable that produces the output you are currently reading.
 4. The Docker daemon streamed that output to the Docker client, which sent it
    to your terminal.

To try something more ambitious, you can run an Ubuntu container with:
 $ docker run -it ubuntu bash

Share images, automate workflows, and more with a free Docker ID:
 https://cloud.docker.com/

For more examples and ideas, visit:
 https://docs.docker.com/engine/userguide/

ca-des@t1:~$
```

That's it!

[CADES](#) → [User Documentation](#) → [User-Provided Tutorials for Birthright Cloud](#) → [Launch Shiny within Docker](#)

Tutorial contributed by Drew Schmidt. **Note:** User-provided tutorials are not supported by CADES.

Background

[Shiny](#) is a web app framework for the statistical programming language R. And while that might sound crazy, it actually works very well!

Shiny is generally pretty easy to develop locally, but deploying it in Shiny Server can be a bit of a headache if you go about it the wrong way. This guide should help make the process as painless as possible.

First Steps

For simplicity, we will be using Shiny inside of a Docker container. Since shiny is a web app, that might lead you to think this process is going to be considerably more complicated than running it natively. However, I can assure you that this is not the case, particularly since the [Rocker Project](#) has handled most of the configuration details for us. They provide a [shiny container](#) already set up for business. However, if you wish to install Shiny server natively for some reason, then you will find some useful details in their [Dockerfile](#).

So the first step is to set up an Ubuntu 16.04 vm on the Birthright Cloud with Docker configured. See the Docker tutorial to learn how to set up Docker.

Open Port(s)

Since Shiny is a web app, we need to open some ports on the vm so that we can access it from a laptop. We also have to do something similar for the Docker container so it can communicate with the vm, but that is actually much easier (and explained in the next section).

We will be using port 80. If you need to open another port, you can, but it's not as simple because port 80 is a pre-configured rule set for you. The [Run a Simple Web Server](#) guide from the official Birthright Cloud documentation is quite helpful here.

The condensed version is you want to navigate to Access & Security and:

1. click Create Security Group and give it a name (I called mine `shiny`)
2. click Manage Rules
3. click Add Rule
4. select http (this will open port 80; if you want another port, set up a custom tcp rule)
5. navigate to your instance (under the Instances tab)
6. click the triangle button next to Create Snapshot, select Edit Security Groups
7. add your new security group by clicking the blue "plus" button

That may sound like a lot or feel a bit overwhelming, but it's not so bad. Give it a try!

Run an Existing Shiny Container

Example 1: k-means Demo

Ok, we're finally ready to start talking Shiny. Assuming you've got your vm with Docker running, we'll first run some example (already configured and built) apps.

The first is a simple demonstration of k-means, and is an official [Shiny Gallery](#) example. It is [already set up](#) on Docker Hub, so standing this up is a breeze. Simply ssh to your vm and run:

```
sudo docker pull wrathematics/shinykmeans
sudo docker run -i -t -p 80:3838 wrathematics/shinykmeans
```

If you chose a port other than 80 when setting up the security groups, then you will need to change the `80` above accordingly.

Now just point your web browser to the IP of your vm, which you can find from the Compute -> Instance tab on the Birthright Cloud dashboard. Example screenshot below (with actual IP addresses masked)

The screenshot shows the 'Instances' page in the Red Hat OpenStack Platform. The table below is a representation of the data shown in the image:

| Instance Name | Image Name | IP Address | Size | Key Pair | Status | Availability Zone | Task | Power State | Time since created |
|-----------------------------|-------------------------------|-------------------------------------|----------|----------|--------|-------------------|------|-------------|---------------------|
| <input type="checkbox"/> t1 | CADES_Ubuntu16.04_v20170302_0 | <input type="checkbox"/> [Redacted] | m1.small | mk | Active | nova | None | Running | 2 hours, 46 minutes |

So if your ip is `1.2.3.4`, then you just need to go to `http://1.2.3.4`. If you chose to use a port other than 80, then you will need to append a colon `:` and that port to the end of your ip address. So if you chose port `5555`, then you would go to `http://1.2.3.4:5555`.

Example 2: Plot Builder

Something that shows off the power of R and Shiny a bit more is the app "ggplotwithyourdata". For your convenience, this too has been [set up](#) on Docker Hub, and is similarly easy to use:

```
sudo docker pull wrathematics/ggplotwithyourdata-docker
sudo docker run -i -t -p 80:3838 wrathematics/ggplotwithyourdata-docker
```

As before, just point your web browser to your vm's url and you're good to go.

Deploying Your Own Shiny App

Fortunately or unfortunately (depending on your perspective), this is the most tricky part. In truth, this goes a bit beyond the scope of this document. Only you really know the dependencies of your app, so we can't tell you exactly what to do.

But we can give some general advice:

1. You will need to create your own Dockerfile. This guide is also not the right place to learn all about the various options of Dockerfiles; [this is](#). However, generally speaking, a Dockerfile is not that far removed from a shell script. So knowing how to build the app natively is important (and really, no different) in understanding how to get it to build in the container.

2. See how things are done on existing Dockerfile configurations. The k-means configuration is available [here](#) and the plot one [here](#). These have been deliberately kept fairly simple.
3. Installing R packages from source is often much harder and always *much* more time consuming than installing binary packages. When you run `install.packages()` or `devtools::install_github()` or the like, you are installing a source package. However, there are many binary packages available in `apt`. These packages all have the prefix `r-cran-`, so you should check what is available before going the `install.packages()` route. However, sometimes the `r-cran-` packages are out of date, so check before you install.
4. You can run multiple apps in the same Docker container. You just place each one in its own subdirectory of `/srv/shiny-server`. So if you have apps `foo` and `bar` that you want to host in the same container (say they have very similar dependencies), you would put them in `/srv/shiny-server/foo/` and `/srv/shiny-server/bar`. Then you would view them at `http://1.2.3.4/foo` and `http://1.2.3.4/bar` (replace the fake ip with your real one of course!).
5. For troubleshooting Shiny Server problems, see the [Shiny Server Open Source Administrator's Guide](#).

[CADES](#) → [User Documentation](#) → [User-Provided Tutorials for Birthright Cloud](#) → [Using Eclipse IDE](#)

Tutorial contributed by Fengming Yuan. **Note:** User-provided tutorials are not supported by CADES.

Using the Eclipse IDE

NOTE: Here is an example of create new fortran project, e.g. PFLOTRAN

STEP 1. Log into CADES

by X-windows from your local terminal (e.g. XQuartz on Mac OSX, or Cygwin-X on Windows OS).

```
ssh -X or-condo-login02.ornl.gov
(input ucams/xcams id and pwd)
```



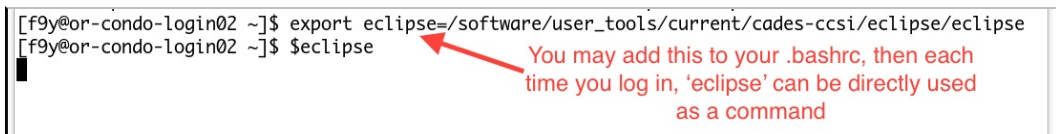
```
f9y@or-condo-login02:~
bash-3.2$ ssh -X or-condo-login02.ornl.gov
key_load_public: invalid format
f9y@or-condo-login02.ornl.gov's password:
Last login: Tue Oct 3 10:43:00 2017 from dyn01015919673.dz.ornl.gov
[f9y@or-condo-login02 ~]$
```

STEP 2. START eclipse

(note: you may create a link, such as optional command in the following, or add `export`

`eclipse=/software/user_tools/current/cades-ccsi/eclipse` in `.bashrc`)

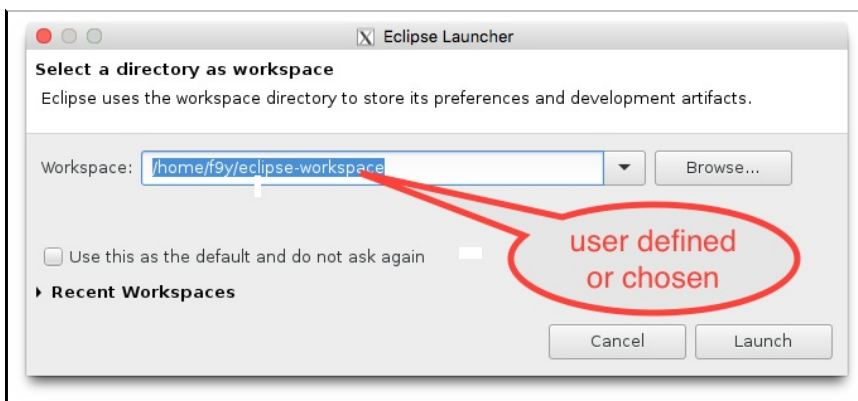
```
$ export eclipse=/software/user_tools/current/cades-ccsi/eclipse
(note: this is optional, otherwise, you can directly start the program)
```



```
[f9y@or-condo-login02 ~]$ export eclipse=/software/user_tools/current/cades-ccsi/eclipse/eclipse
[f9y@or-condo-login02 ~]$ $eclipse
```

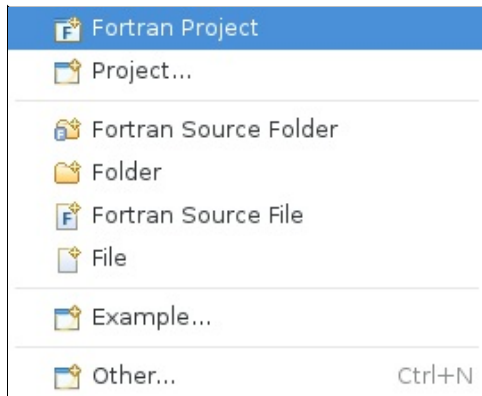

You may add this to your `.bashrc`, then each time you log in, 'eclipse' can be directly used as a command

```
$eclipse
#(this will start ECLIPSE in GUI shown in your desktop/laptop)
```

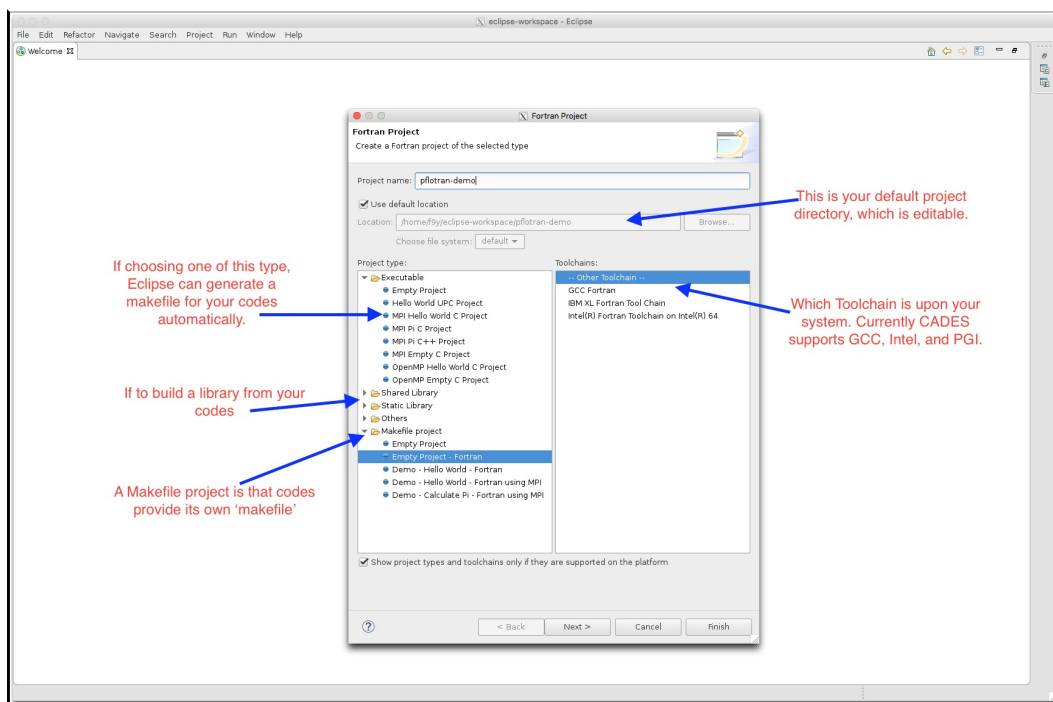


STEP 3. Create a project

WELCOME page (if first time) & FILE DROP-DOWN menu:

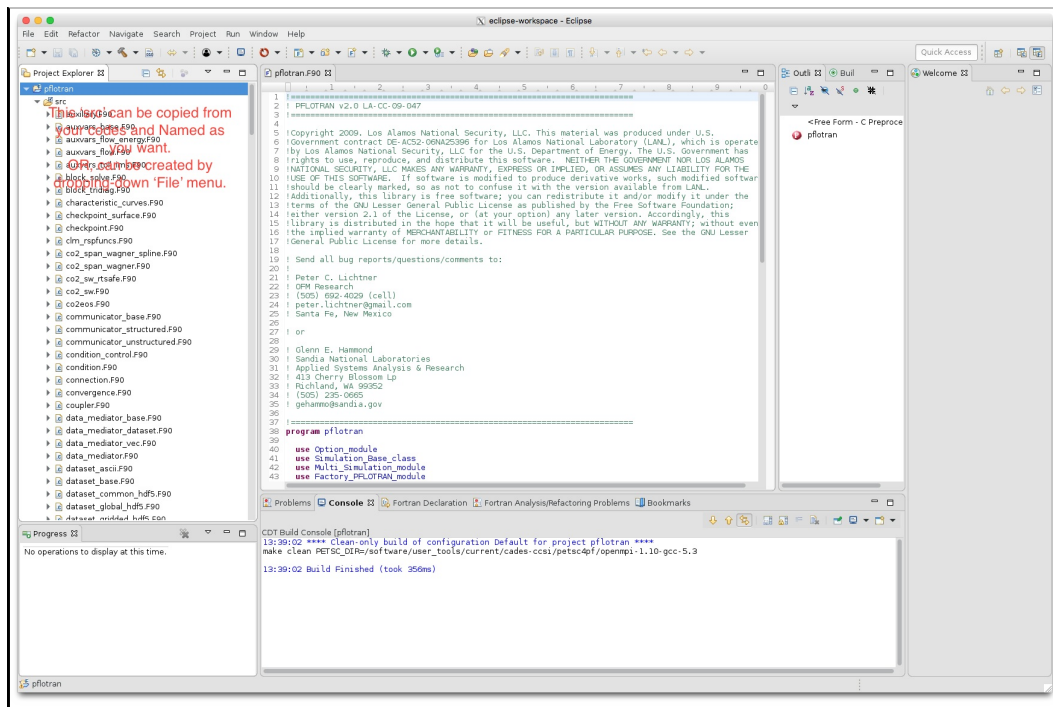


NEW ... PROJECT configuration:

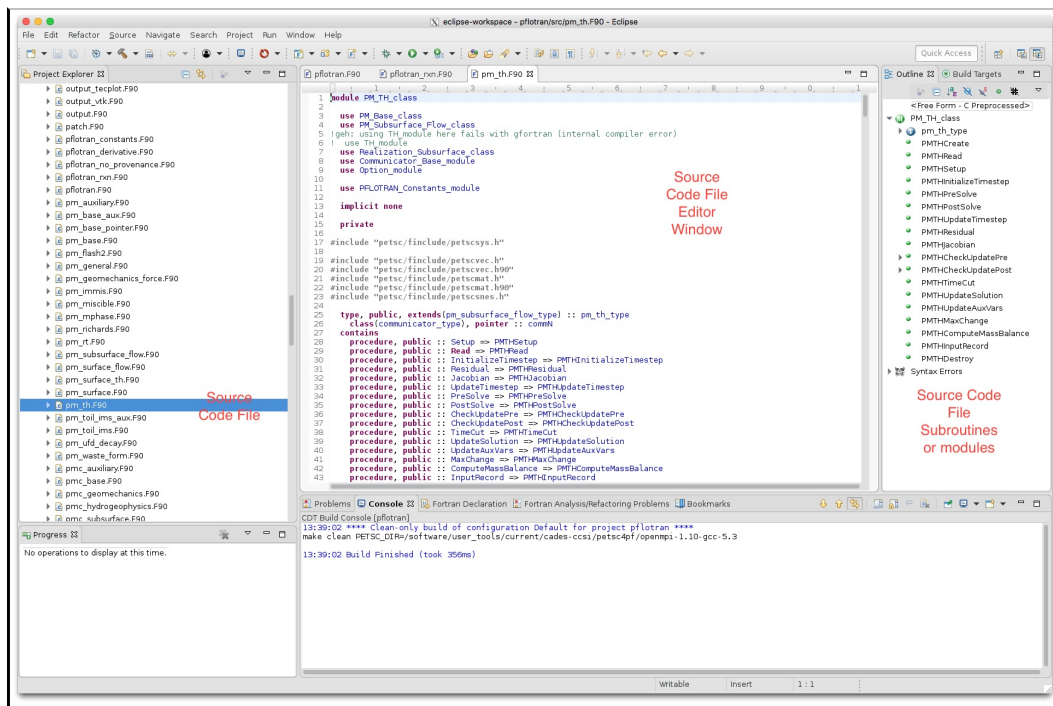


GENERIC Interface/windows of ECLIPSE:

(note on how to add/create source codes)



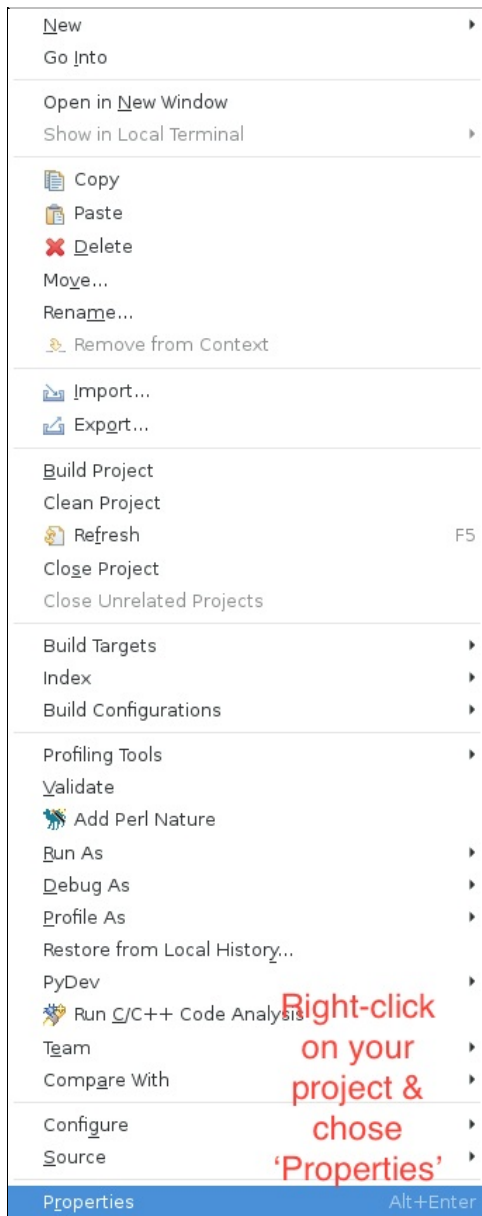
MORE on windows...



STEP 4. Build project

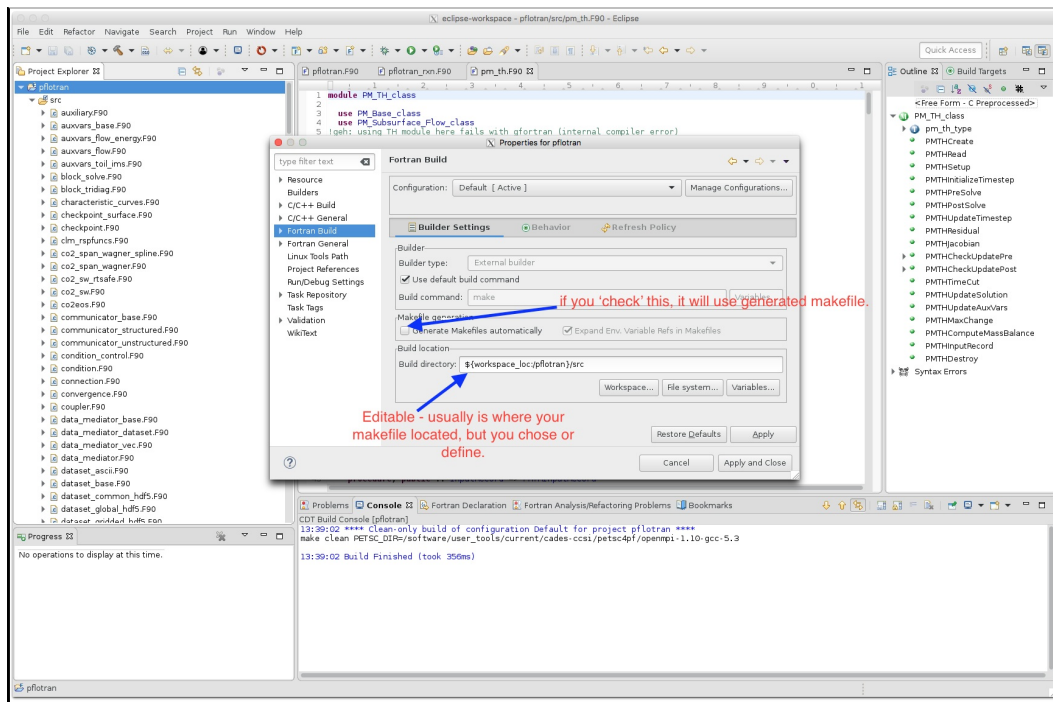
4.1 Project Properties Editing

Right-click on the Project in 'Project Explorer' window ...



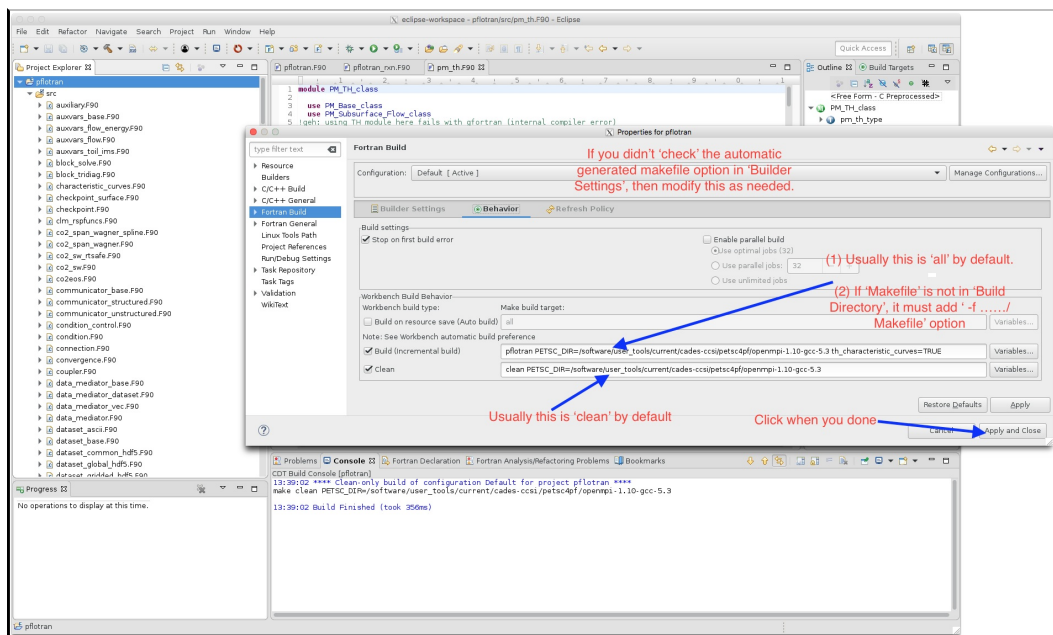
4.2 Building settings

Editing as what you NEEDED ... e.g. Build Directory (NOTE: you may browse your file system)



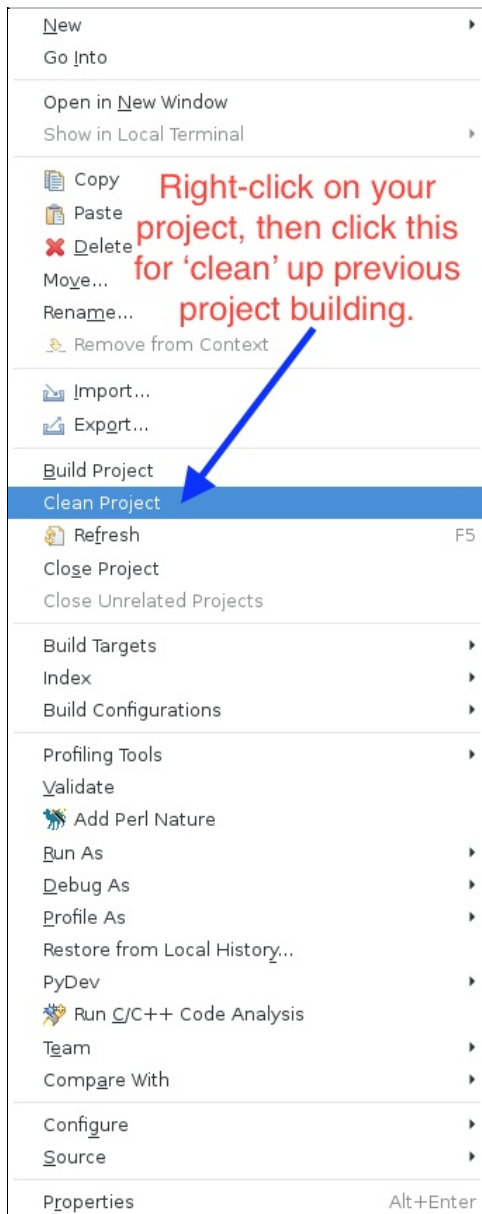
4.3 Building behavior editing

AGAIN, as NEEDED.



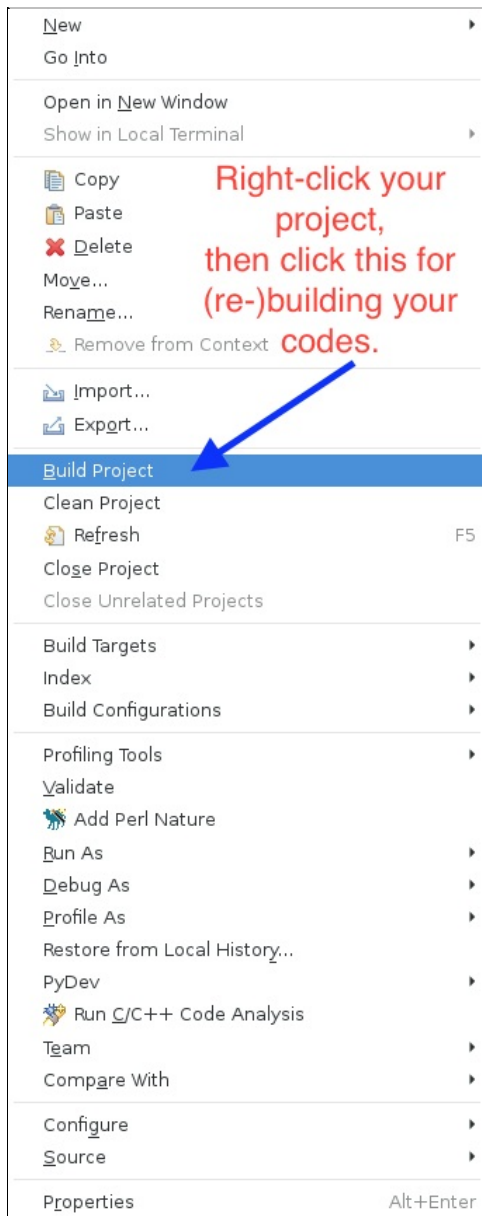
4.4 Clean up (previous) build (OPTIONAL)

IF you would like to, clean previous build of project, by Right-click project, followed by clicking on 'Clean Project'

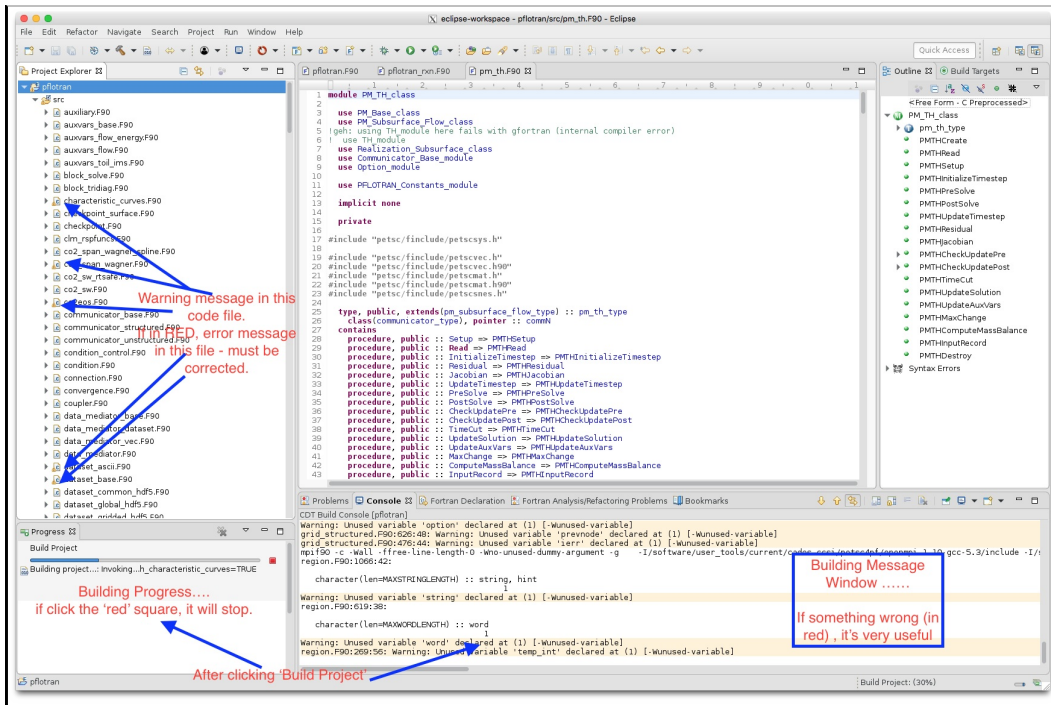


4.5 Build your project

Right-click on your project, then click on 'Build Project'

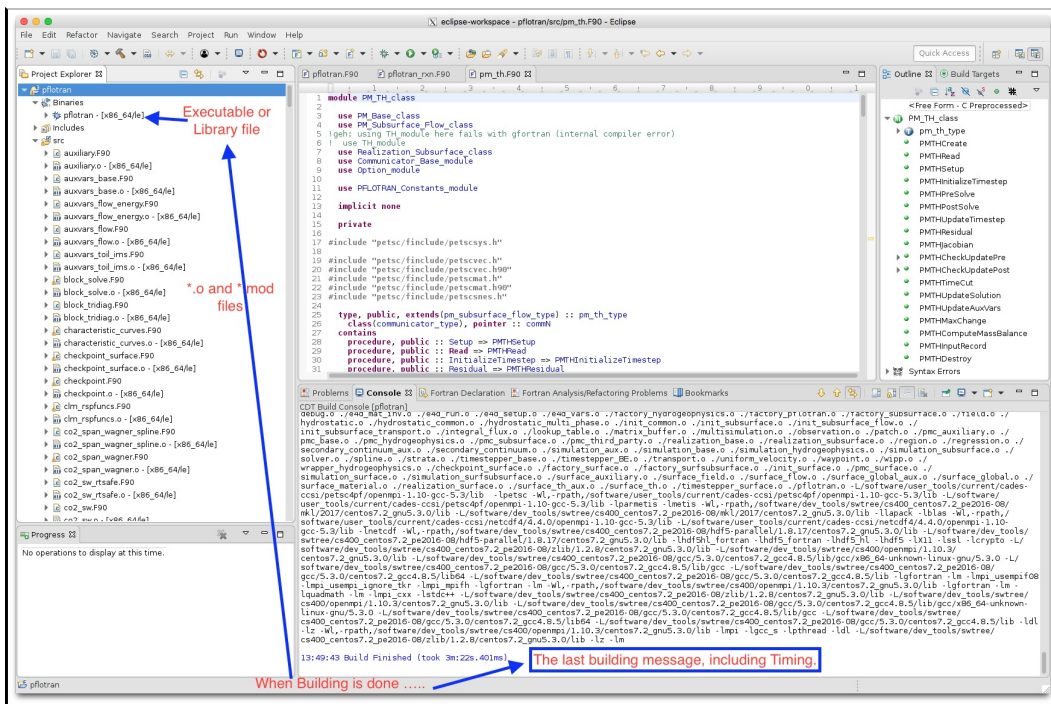


Building progress



After Building Successfully.

IF NOT, clicking the 'error' code file and editing in the 'editor window'



STEP 5. Run project

(NOTE: it's NOT allowed to directly run a program from the login node on CADES!)

Run your executables in terminal, OR, include your built library (as USUALLY you do)

CADES → User Documentation → User-Provided Tutorials for CADES Cloud → Using Allinea DDT

Tutorial contributed by Fengming Yuan. **Note:** User-provided tutorials are not supported by CADES.

Allinea Forge/DDT Client

I. Allinea DDT Client

The Client software is free from: <https://www.allinea.com/products/forge/download#remote-client>.

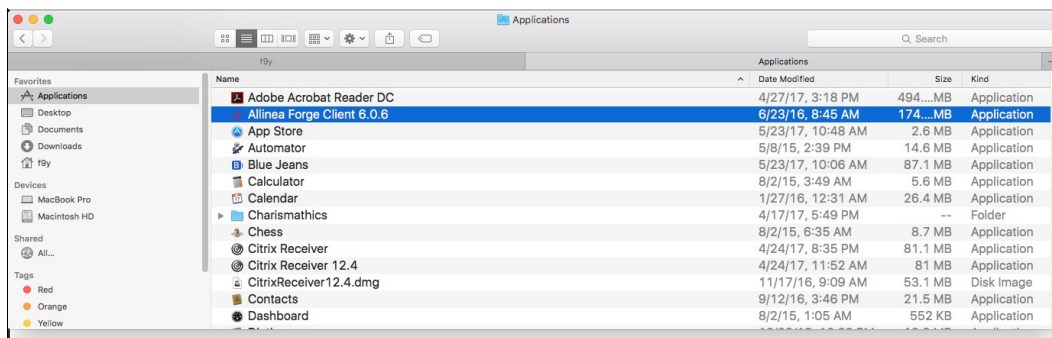
The version for matching with CADES' server-end is 6.0.6 (Currently).

II. Configuration of Remote Server

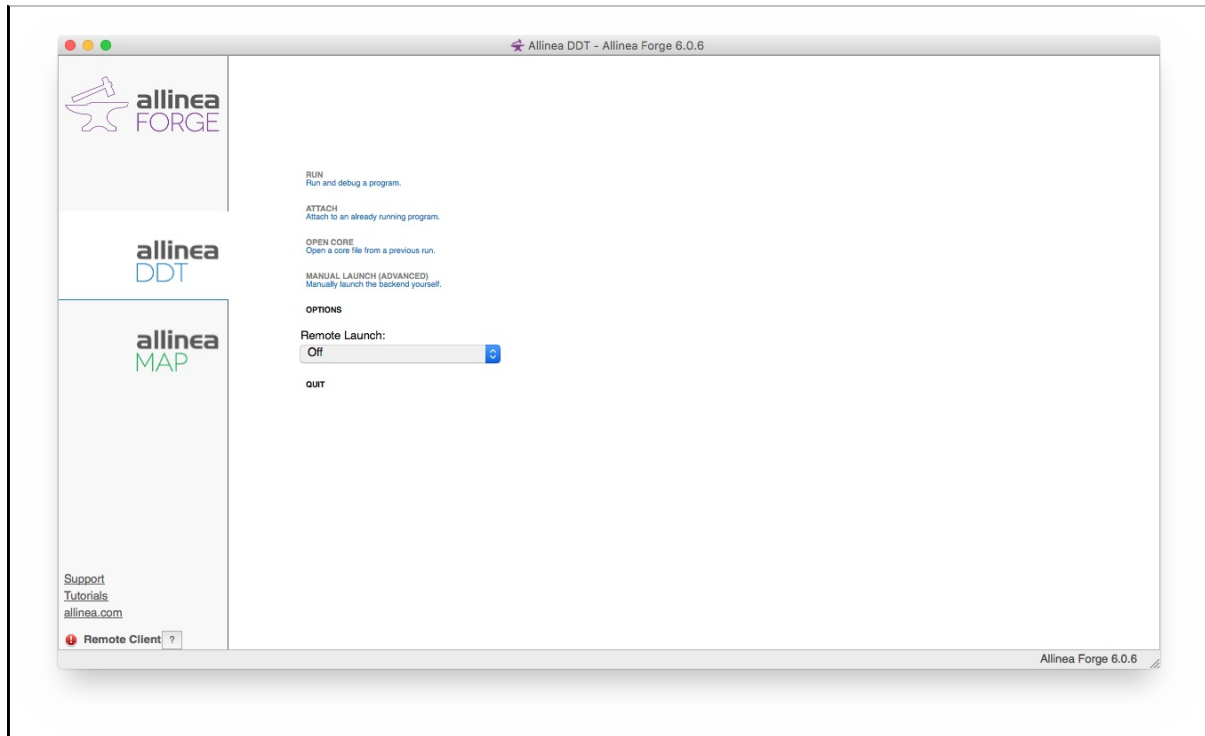
Currently, CADES has 2 versions of Allinea DDT/Forge, the following is tested with v.6.0.6

II-1. Start Allinea Forge Client locally

Allinea Forge/DDT Client application installed

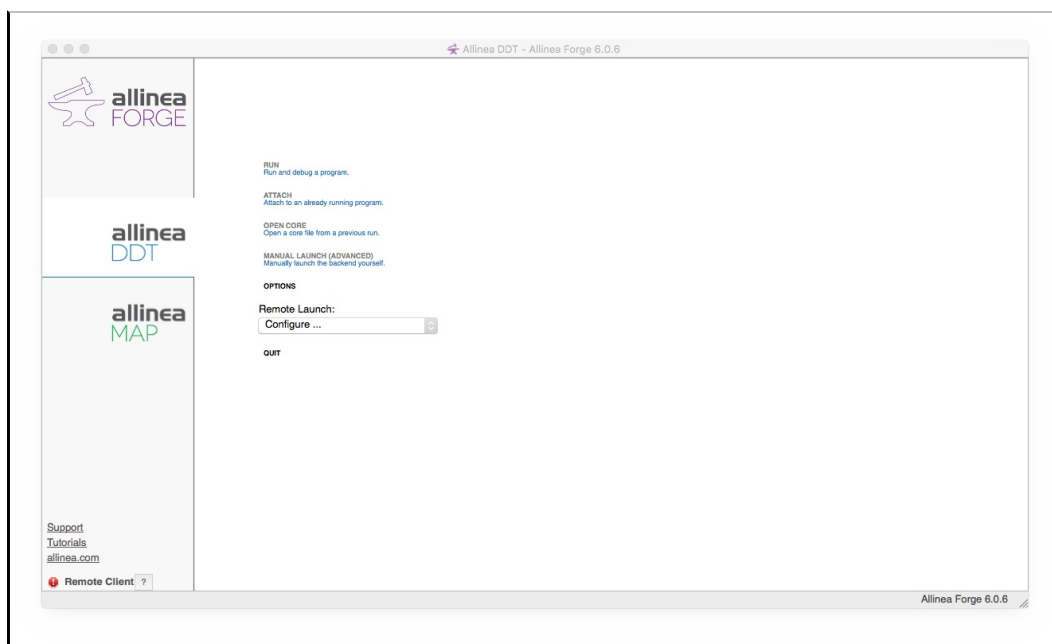


On your screen

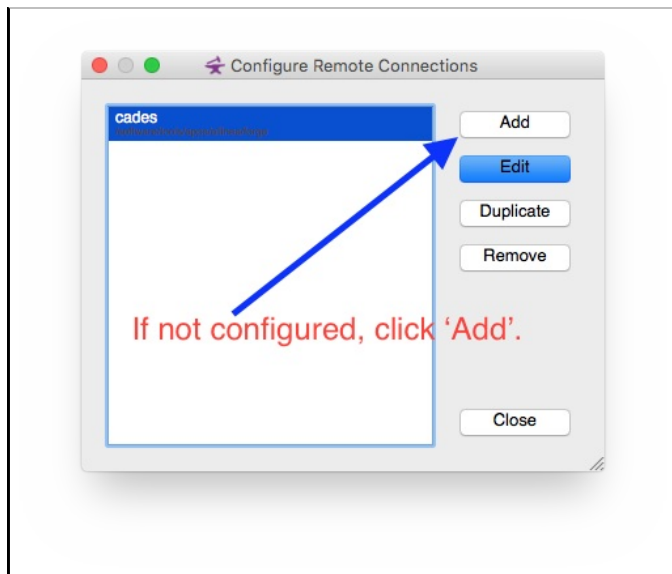


II-2. Configure Remote Launching

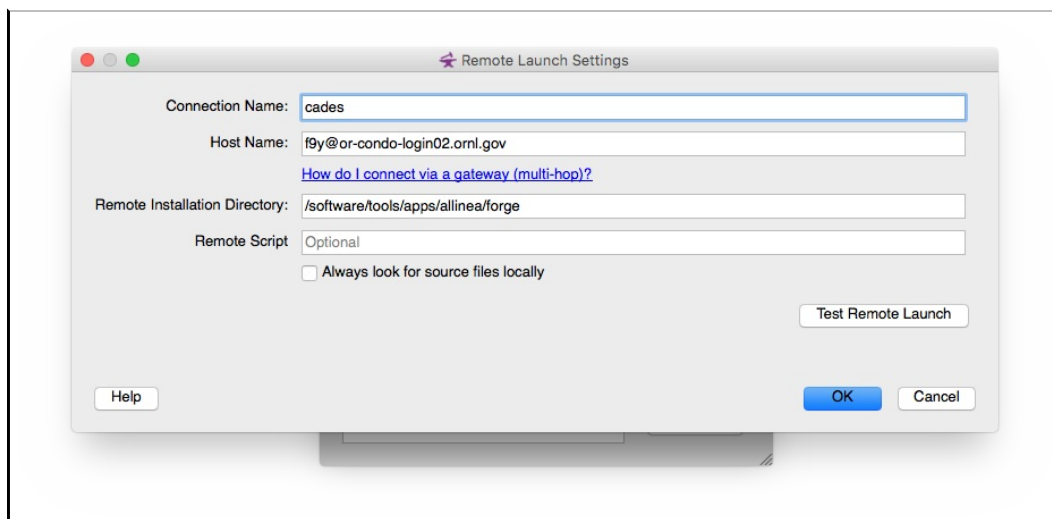
Drop-down configure



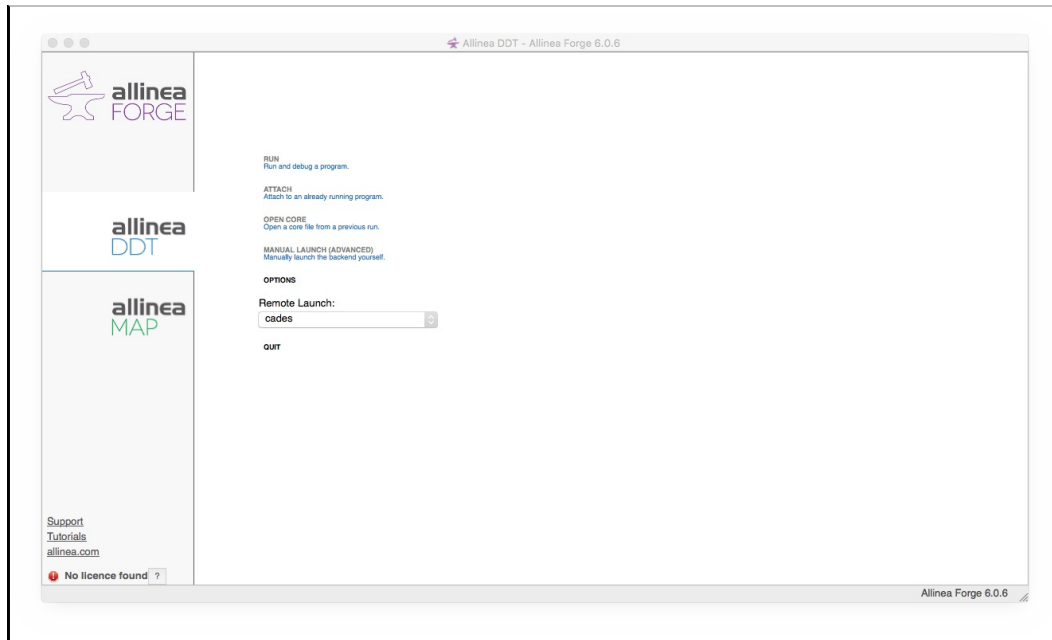
Add or Modify



Remote Launch Settings.....



AND,



II-3. DDT template on CADES

You have to configure a DDT template file, to be used by Allinea Forge Client, as following:

e.g., here a file named as `ddt_template.qtf`, which MUST be put in your home directory.

```
#
# Name: Generic Torque
#
# WARNING: If you install a new version of Allinea Forge to the same
#          directory as this installation, then this file will be overwritten.
#          If you customize this script at all, please rename it.
#
# submit: qsub
# display: qstat
# job regexp: (.+)
# cancel: qdel JOB_ID_TAG
# show num_nodes: yes
#
# WALL_CLOCK_LIMIT_TAG: {type=text,label="Wall Clock Limit",default="00:30:00",mask="09:09:09"}
# QUEUE_TAG: {type=text,label="Queue",default=debug}
## Allinea Forge will generate a submission script from this by
## replacing these tags:
##      TAG NAME          |      DESCRIPTION          |      EXAMPLE
## -----
## PROGRAM_TAG           | target path and filename  | /users/ned/a.out
## PROGRAM_ARGUMENTS_TAG | arguments to target program | -myarg myval
## NUM_PROCS_TAG         | total number of processes | 16
## NUM_NODES_TAG         | number of compute nodes   | 8
## PROCS_PER_NODE_TAG    | processes per node        | 2
## NUM_THREADS_TAG       | OpenMP threads per proc   | 4
## DDT_DEBUGGER_ARGUMENTS_TAG | arguments to be passed to ddt-debugger
## MPIRUN_TAG            | name of mpirun executable | mpirun
## AUTO_MPI_ARGUMENTS_TAG | mpirun arguments         | -np 4
## EXTRA_MPI_ARGUMENTS_TAG | extra mpirun arguments    | -x FAST=1

#!/bin/bash
#PBS -S /bin/bash
#PBS -m ae
#PBS -j oe
#PBS -M yuanf@ornl.gov
#PBS -N acme_debug
```

```

#####PBS -q QUEUE_TAG
#PBS -q batch
#PBS -l nodes=NUM_NODES_TAG:ppn=PROCS_PER_NODE_TAG
#PBS -l walltime=WALL_CLOCK_LIMIT_TAG
#PBS -W group_list=ca-des-ccsi
#PBS -A ccsi
#PBS -l qos=std
#PBS -l naccesspolicy=singlejob
export OMP_NUM_THREADS=1
#PBS -V
#PBS -o PROGRAM_TAG-ddt.output
#PBS -e PROGRAM_TAG-ddt.error
## The following line will use exactly the same default settings that
## Allinea Forge uses to launch without the queue.
module load env/ca-des-ccsi
module load vasp
AUTO_LAUNCH_TAG
## Replace the above for more complex situations - such as for passing unusual
## parameters to mpirun, like machine files or processes per node -- below is
## an example.
##
## if test "MPI_TAG" = "mpich 1 standard" ; then
##   MPIRUN_TAG -tv -np NUM_PROCS_TAG PROGRAM_TAG PROGRAM_ARGUMENTS_TAG
## else
##   if test DEBUG_STARTER_TAG -eq 1 ; then
##     DDT_CLIENT_TAG MPIRUN_TAG -np NUM_PROCS_TAG -machinefile $PBS_NODELIST PROGRAM_TAG PROGRAM_ARGUMENTS_TAG
##   else
##     MPIRUN_TAG -np NUM_PROCS_TAG -machinefile $PBS_NODELIST DDT_DEBUGGER_TAG PROGRAM_ARGUMENTS_TAG
##   fi
## fi

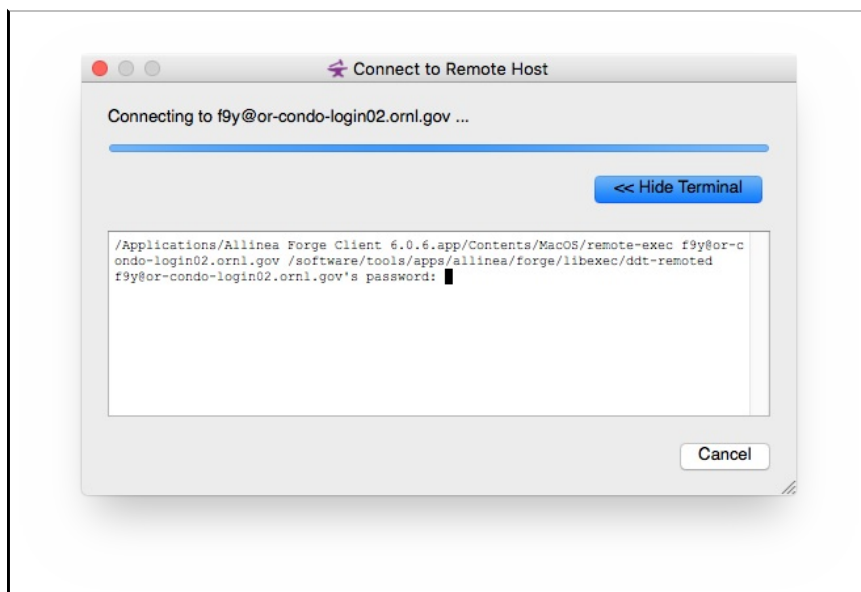
```

NOTE: (1)Please modify user information and others as NEEDED; (2) TIP: # or ## implies comments, EXCEPT for #PBS

III. Step-by-Step Instruction on Debugging

III-1. Start 'Run and debug a program'

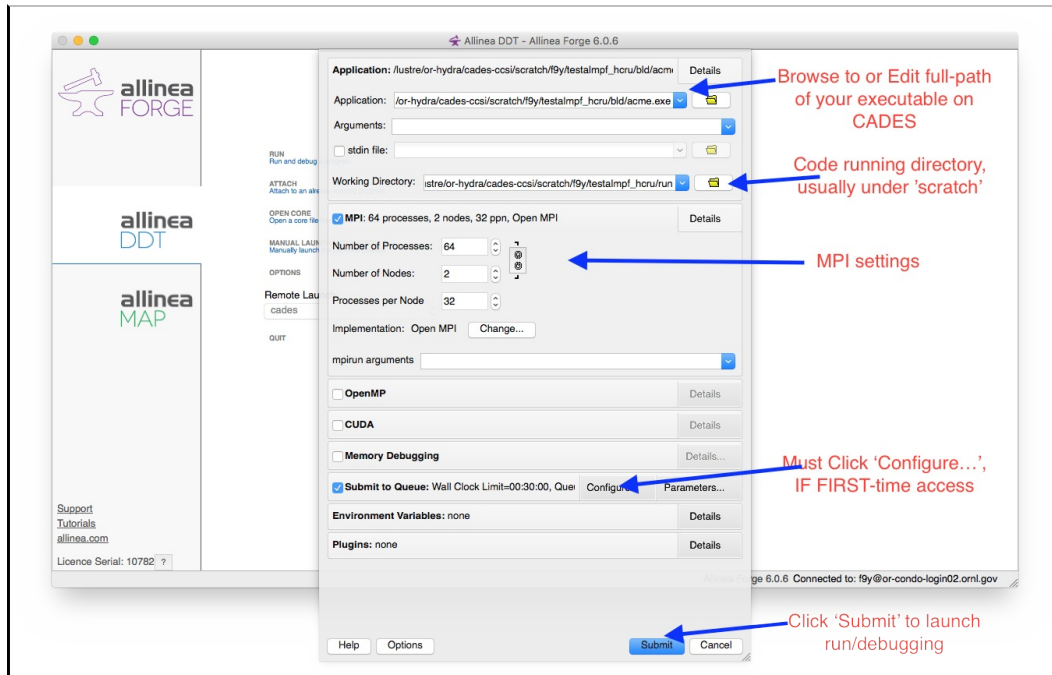
While 'Remote Launch: ca-des' is chosen, click "Run and debug a program" under "RUN"



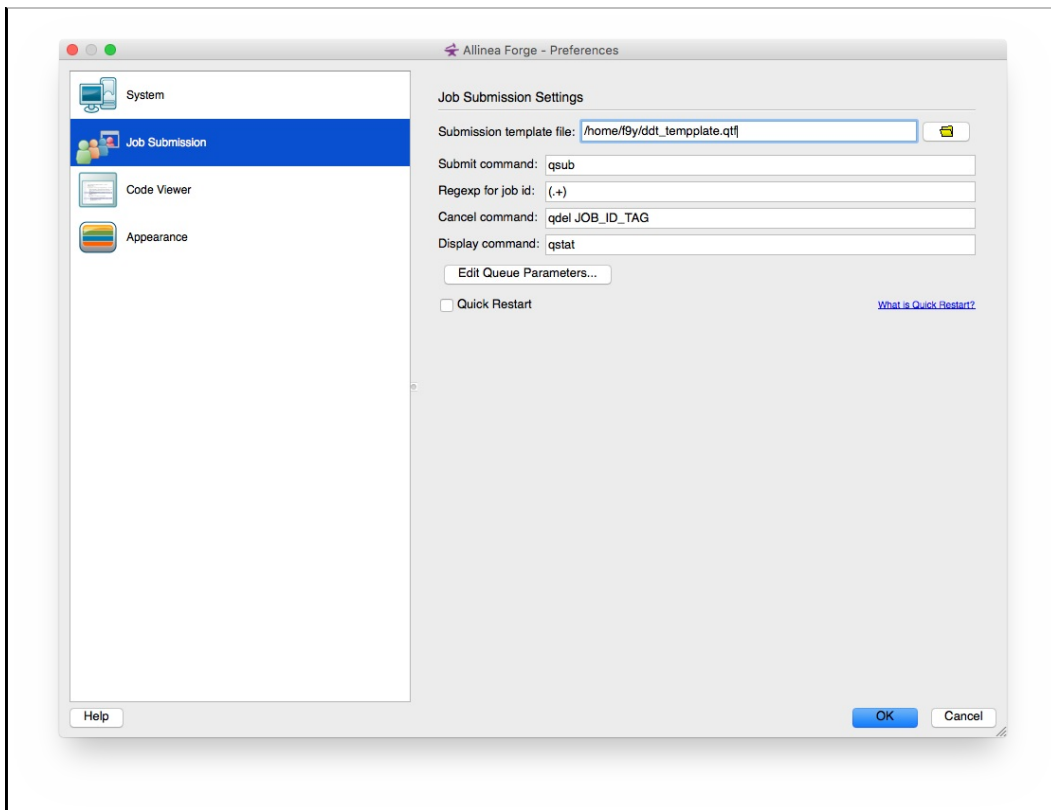
input your pwd, and ENTER

III-2. Configure Allinea DDT job-submission to CADES

After successfully connecting to CADES, the following will pop up. THEN, you may add or edit a job-submission configuration




IMPORTANT: if this is your first time to access CADES for submitting a job to run & debug, you MUST Click 'Configure...' under 'Submit to Queue'. And then add the above template *.qtf file (in your home directory) and other settings, AS Following:



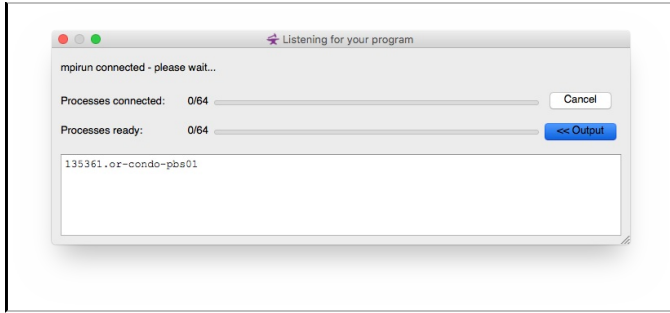
NOTE: This 'configuration' Window also can be used to configure other settings, as SHOWN. You may want to have some inspections by clicking each one*

III-3. Submit job and Run

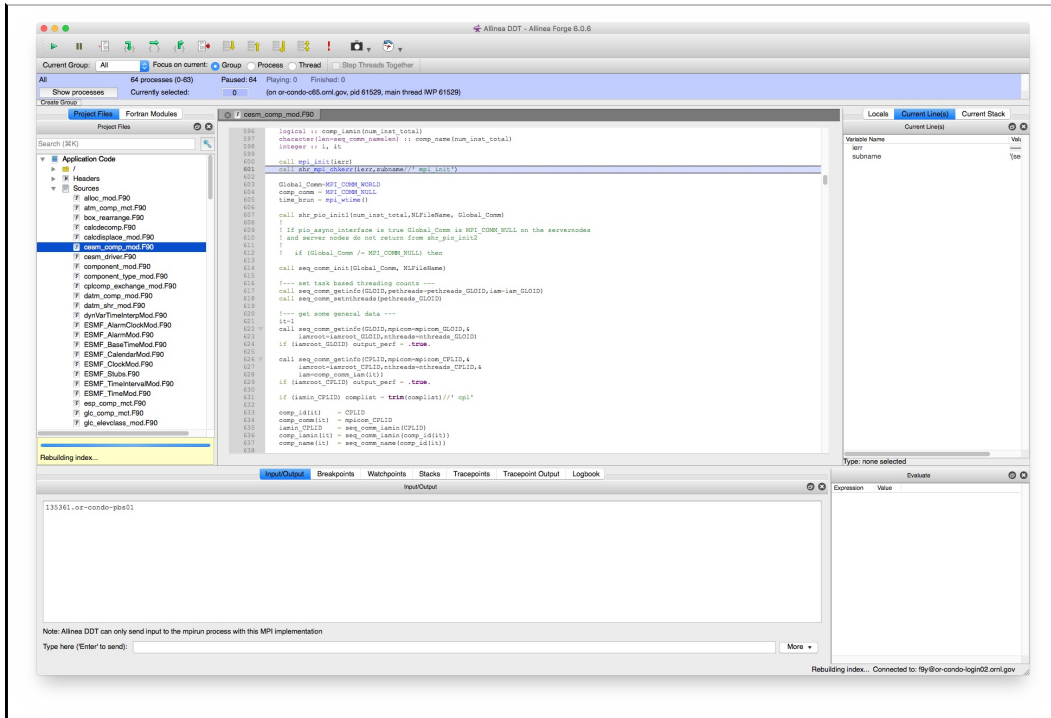
AFTER clicking 'SUBMIT' button



ONCE your job status is as "R", WINDOW will be switching and showing as(OTHERWISE, the above window would be hanging on)

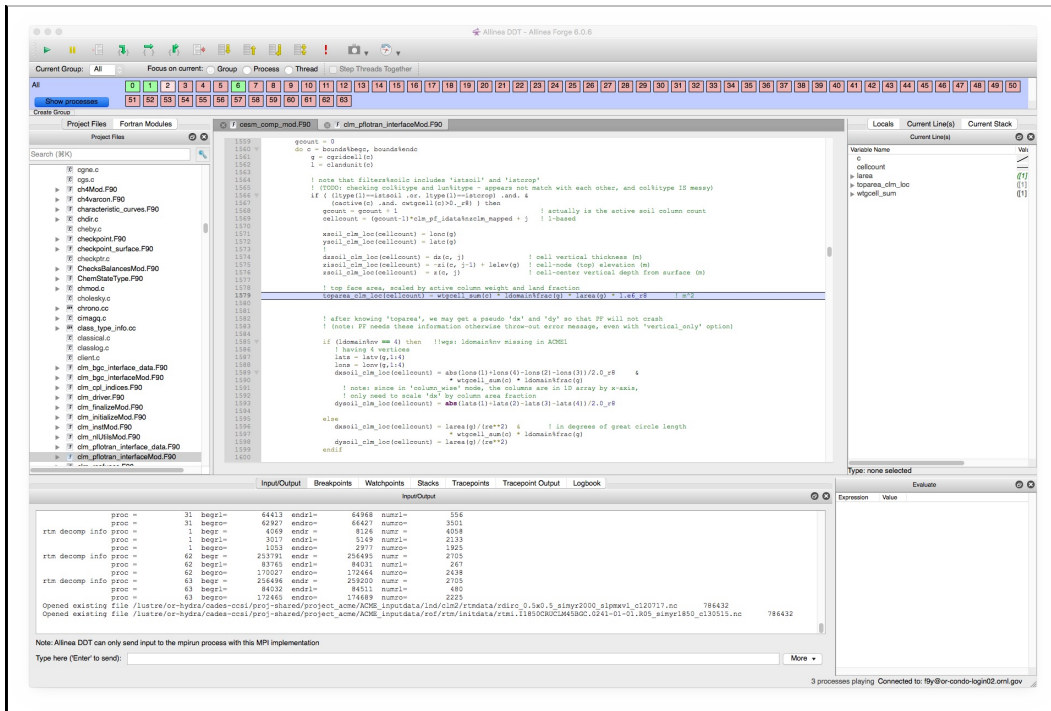


WHEN all requested Processes (e.g. here 64/64) connected and ready, the window will be changing into the run/debug views as FOLLOWING, and PAUSE at the main program entry point (usually).....

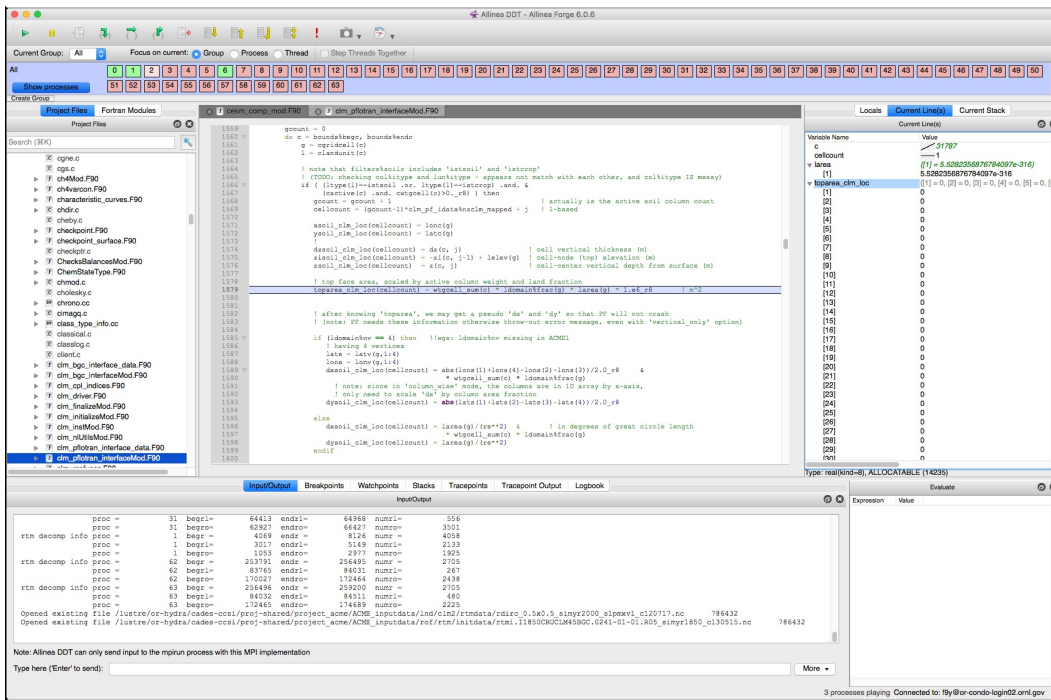


III-4. Debug The Program

PAUSE at Breakpoint(s)



CHECK variables: current line or locals for current subroutine, OR do what you want



HIT a BUG,, (NOT SO FUN?)

The screenshot displays the Allinea DDT - Alinea Forge 6.0.6 IDE. The interface is divided into several panes:

- Project Files:** A tree view on the left showing the project structure with files such as `selectg.c`, `sem.c`, `separator.c`, `seq.c`, `serial.c`, `set_exponent_10.c`, `set_exponent_15.c`, `set_exponent_14.c`, `set_exponent_18.c`, `sl.c`, `sfbasic.c`, `slmic`, `SFMMainMod.F90`, `slp-exceptions.c`, `SFPParamMod.F90`, `sfnegi.c`, `sftype.c`, `slwindoc`, `shape_function.F90`, `shape_116.c`, `shape_14.c`, `shape_18.c`, `shared_ptr.cc`, `shelic`, `shelloc`, `shelloc.c`, `shelloc.c`, `shelloc.c`, and `sl_class_type_info.cc`.
- Main Editor:** Displays a message: "A source file could not be found in its original location: /usr/local/.../signal.c". A "Browse..." button is visible.
- Locals:** A panel on the right showing local variables: `_func` (value: "PetscSignalHandlerDefault"), `err` (value: 59), `ptr` (value: 0), `sig` (value: 11), and `SIGNAME`.
- InputOutput:** A bottom panel showing a stack trace from PETSC:


```
(*) PETSC ERROR: likely location of problem given in stack below
(*) PETSC ERROR: ----- Stack frames -----
(*) PETSC ERROR: Note: The EXACT line numbers in the stack are not available.
(*) PETSC ERROR: INSTEAD the line number of the start of the function
(*) PETSC ERROR: is given.
(*) PETSC ERROR: ----- Error Message -----
(*) PETSC ERROR: Signal received
(*) PETSC ERROR: See http://www.mcs.anl.gov/petsc/documentation/faq.html for trouble shooting.
(*) PETSC ERROR: Petsc Release Version 3.7.3, unknown
(*) PETSC ERROR: /usr/local/condes-codes/condes/branches/fly/testingf_hcrs/bld/cme.exe
(*) PETSC ERROR: Configure options --known-level-dcache-size=32768 --known-level-dcache-ssize=64 --known-level-dcache-ssize=8 --known-sizeof-char=1 --known-sizeof-void-p=8 --x
(*) PETSC ERROR: # User provided function line # in unknown file
```
- Status Bar:** Shows "3 processes playing Connected to fly@or-conds-logn02.ornl.gov".

Moving Data

CADES Data Transfer Nodes (DTNs) allow for speedy movement of large data sets into and out of ORNL's network. There are several transfer tool/protocol options to choose from to fit your needs.

1. [Globus](#) has a web interface or command line tools that you can use to transfer data between your personal endpoints or securely share access to your data. *This is the preferred method of transfer for CADES.*
2. Secure copy (scp) via the command line to and from storage locations, including local computers.

```
scp username@remote-host1.gov:/path/to/directory/example.txt username@remote-host2.gov:/path/to/directory/
```

3. Secure (or SSH) file transfer protocol (SFTP) can be used to transfer files between two remote storage locations (similar to scp) but also allows the user to list directories and see content. You can use SFTP as long as you have SSH access to that host.

```
sftp username@remote_hostname_or_IP
```

4. [Graphical clients \(SFTP\)](#) will allow you to use a graphical user interface with drag-and-drop capabilities. CADES maintains documentation for CyberDuck and WinSCP.

For Linux users, there is no clear recommendation for SFTP clients. No one free client supports all of CADES storage services and behaves consistently. However, [Cloud Explorer](#) supports all of CADES services and *typically* behaves predictably on Linux systems. See [here](#) for a how-to guide using CloudExplorer and Scality.

5. [rsync](#) or [rclone](#) (supports s3) are other command line utilities that may suit your data transfer needs. At this time, CADES does not offer support for these tools.

Transferring Files (SFTP) with a Graphical Client

Graphical file transfer clients can be used to move data between your local machine and remote storage locations. Once you install the client on your computer and set up the remote connection, you may move folders and files between your computer and the remote storage using a drag-and-drop method.

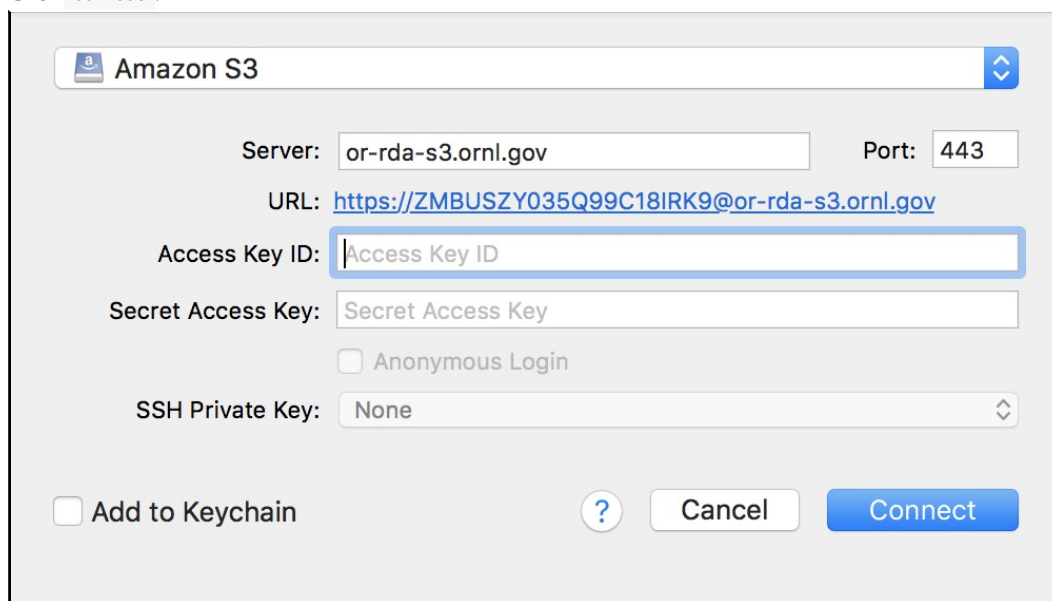
Note: It is impractical to maintain documentation on every storage system that CADES offers. These examples are chosen to be representative of our services. If you need help connecting to a different storage service, please [email CADES](#).

CyberDuck (macOS and Windows)

Download Cyberduck [here](#) and run the installation.

AWS S3 - Scality

- To set up a new connection, click on the `open connection` button in the top left of the window.
- In the dropdown menu of the resulting window, select `Amazon S3`.
- For Scality, change the server field to `or-rda-s3.ornl.gov`.
- Paste your Access Key ID and Secret Access Key that was generated when you signed up for the AWS S3 service.
- Click `connect`.



The screenshot shows a dialog box titled "Amazon S3" with the following fields and options:

- Server:** `or-rda-s3.ornl.gov`
- Port:** `443`
- URL:** `https://ZMBUSZY035Q99C18IRK9@or-rda-s3.ornl.gov`
- Access Key ID:** `Access Key ID`
- Secret Access Key:** `Secret Access Key`
- Anonymous Login**
- SSH Private Key:** `None`
- Add to Keychain**
- Buttons:** `?`, `Cancel`, `Connect`

OpenStack Virtual Machine

- To set up a new connection, click on the `open connection` button in the top left of the window.
- In the dropdown menu of the resulting window, select `SFTP (SSH File Transfer Protocol)`.
- **Server:** the IP address of your virtual machine
- **Username:** `ca-des`
- **Password:** leave blank
- Select your SSH key from the dropdown menu. Be sure to choose the SSH key that allows you to access your OpenStack virtual machine.
- Click `connect`.

CADES OR Condo SHPC, NFS, and Lustre

- To set up a new connection, click on the `open connection` button in the top left of the window.
- In the dropdown menu of the resulting window, select `SFTP (SSH File Transfer Protocol)`.
- Server: `or-condo-login.ornl.gov`
- Username: your UCAMS ID (UID)
- Password: your UCAMS password
- Select your SSH key from the dropdown menu. Be sure to choose the SSH key that allows you to access the CADES OR SHPC Condo login node.
- Click `connect`.

NFS user home directory path: `~/home/UID/`

Lustre storage path: `~/lustre/or-hydra/`

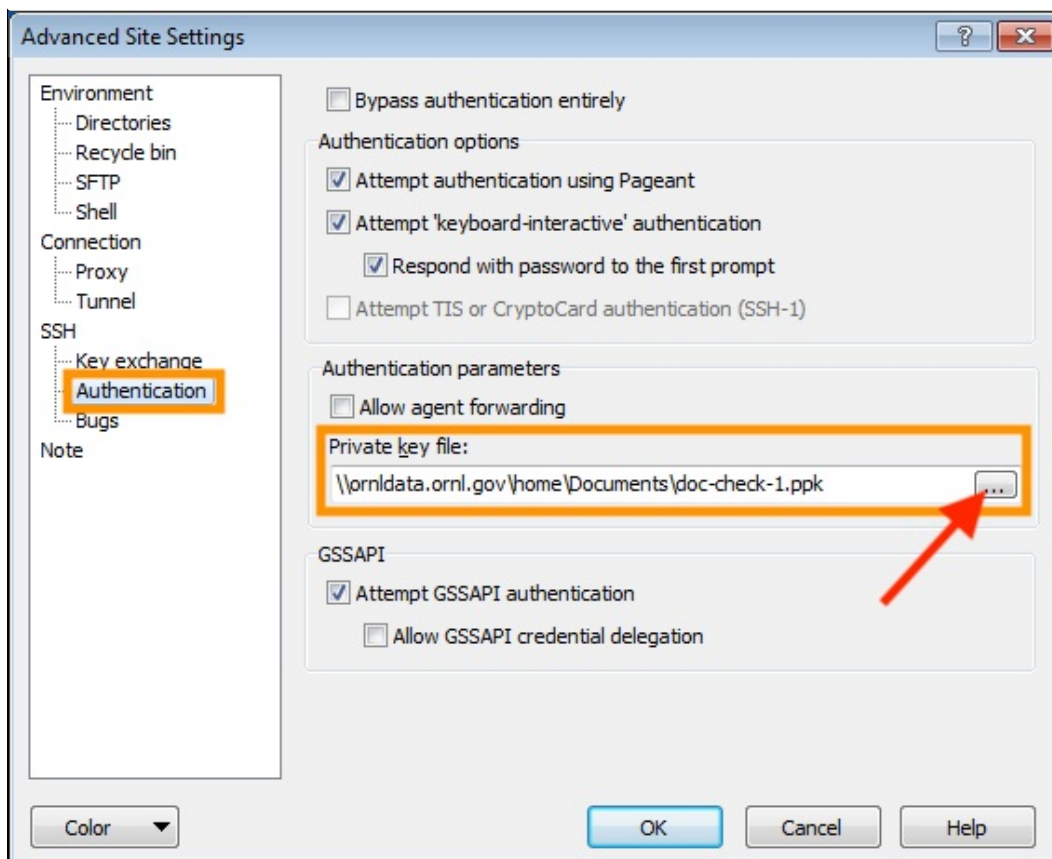
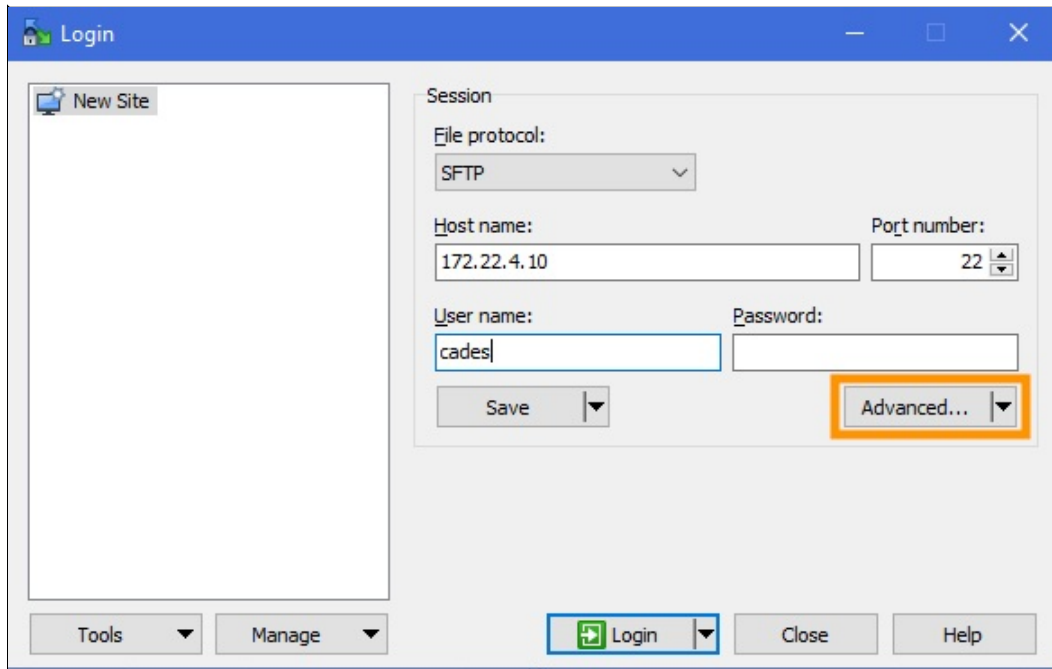
WinSCP (Windows)

Download WinSCP [here](#) and run the installation.

:bangbang: In cases where an SSH key is required for access, you must store the path to the key in WinSCP for each connection.

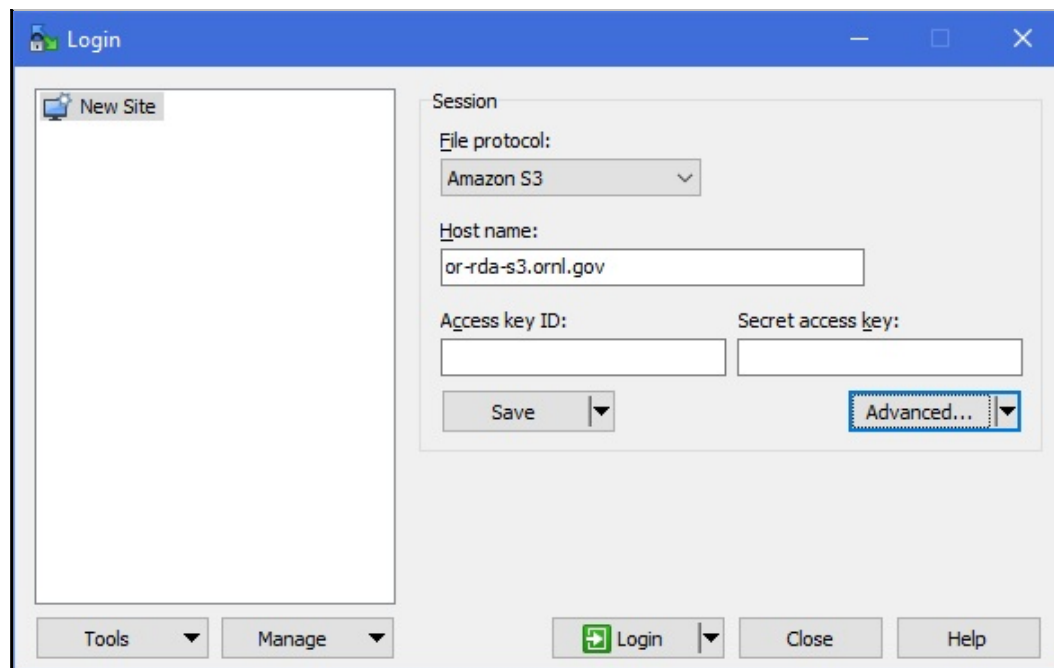
To store the key, enter the connection information that you will find in the steps below. Then, click the `Advanced...` button.

Provide the path to your SSH private key.



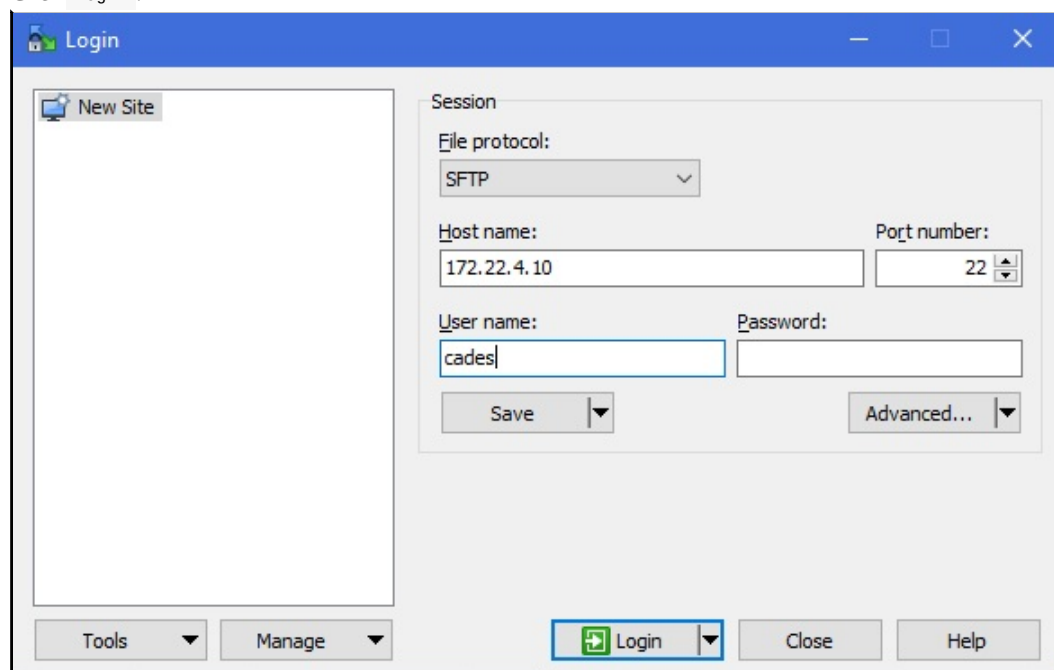
AWS S3 - Scalify

- To set up a new connection, click on `New Site` in the top left of the window.
- In the `File protocol` dropdown menu on the right, select `Amazon S3`.
- Host name: `or-rda-s3.ornl.gov`
- Paste your Access Key ID and Secret Access Key that was generated when you signed up for the AWS S3 service.
- Click `Login`.



OpenStack Virtual Machine

- To set up a new connection, click on **New Site** in the top left of the window.
- In the **File protocol** dropdown menu on the right, select **SFTP**.
- **Host name:** the IP address of your virtual machine
- **User name:** cades
- **Password:** leave blank
- Click **Login**.



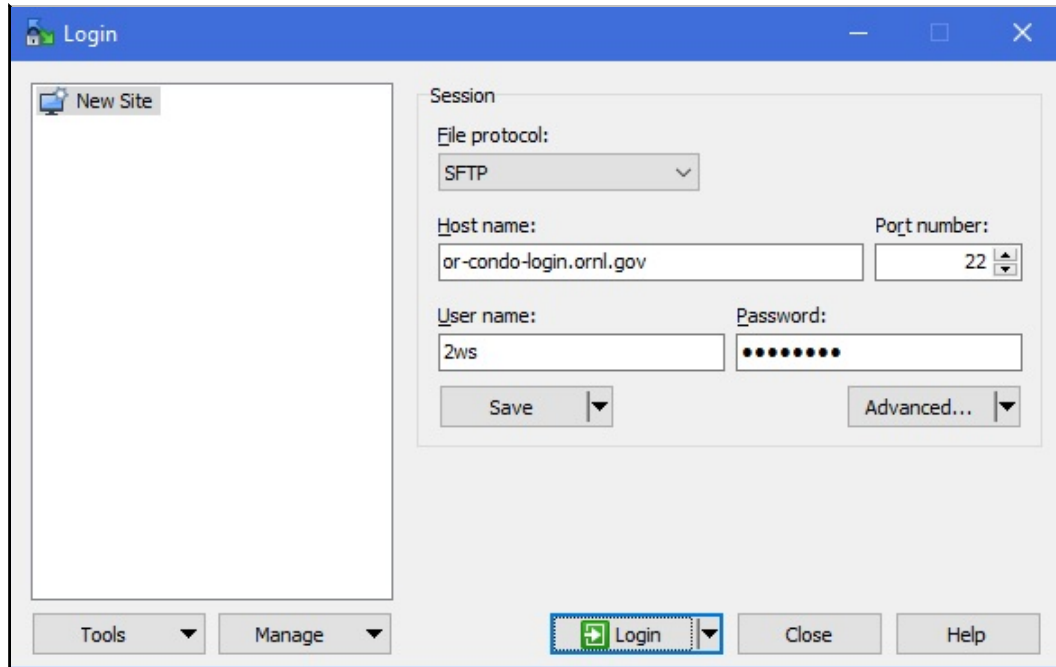
CADES OR Condo SHPC, NFS, and Lustre

- To set up a new connection, click on **New Site** in the top left of the window.
- In the **File protocol** dropdown menu on the right, select **SFTP**.
- **Host name:** the IP address of your virtual machine

- Username: your UCAMS ID (UID)
- Password: your UCAMS password
- Click **Login** .

NFS user home directory path: ~/home/UID/

Lustre storage path: ~/lustre/or-hydra/



Related Tutorials

- [Scality Object Storage User Guide](#)
- [Globus Data Transfer Tool](#)
- [Access VM Instances](#)

[CADES](#) → [User Documentation](#) → [Globus Overview](#)

Globus Overview

Globus is a powerful data transfer tool that has a wide range of support for popular storage systems and a simple graphical user interface. Using Globus is as easy as 1, 2, 3:

1. [Set-up your Globus Account](#)
2. [Find or Set-up Endpoints](#)
3. [Transfer your Files & More](#)

Getting Started and Signing In

Globus is primarily used via its web interface, though it is possible to download a personal client or use command line tools.

1. Navigate to the Globus website (<https://www.globus.org/>) and click `Log in`.
2. Select your organization `Oak Ridge National Laboratories` from the drop-down menu and select `continue`.
3. Use your UCAMS username and password to log in.
 - If you have an existing Globus account, you may choose to link them at this time, or skip to the next step, by clicking `No thanks, continue`.
4. Accept the user agreement and `continue`. The next screen will ask you to `Allow` permissions.

For a list of common Endpoints or if you'd like to learn how to use Globus Endpoints, [click here for our guide](#).

Note: If you can not login to the DTN, but can to other CADES systems (like the login nodes), your account may have been temporarily blocked on the DTN. This occurs, for example, upon too many failed password attempts. If so, [email CADES](#).

Globus Endpoints

Globus Endpoints are storage systems to which you have access. Once an Endpoint is located or created Globus saves the location for you so you do not need to repeatedly search type paths.

| Endpoint Search Term(s) | Storage System | Path | Description |
|-------------------------|----------------|----------------------------|--|
| CADES OR | NFS | /~/ | CADES open research, user home directory |
| CADES OR | NFS | /data/ | CADES open research, NFS project directories |
| CADES OR | Lustre | /lustre/or-hydra/ | CADES open research, project directories. High-performance, temporary storage. |
| NCCS Open DTN | Lustre | See here | Requires NCCS Open (XCAMS) account |
| OLCF ATLAS | OLCF DTN | /path/to/project/file/data | OLCF-managed NFS and Lustre storage system. |

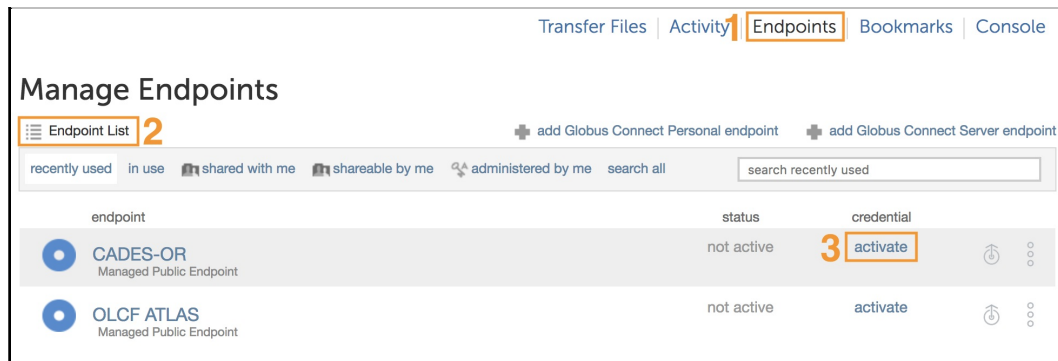
Note: If you're having trouble finding an existing Endpoint, [email the CADES team](#).

Setting Up Endpoints

Note: AWS S3 Scality storage is not yet supported on Globus, but will be in the future.

1. Click in the Endpoint box on the left side and search for CADES-OR (CADES Open Research).
2. You will be redirected to enter your UCAMS credentials.

- o Authenticating the Endpoint with your credentials is known as *Endpoint Activation* and can be done when adding and using an Endpoint for the first time, or can be completed by navigating to the "Manage Endpoints" screen as shown in the following image (Endpoints → Endpoint List → activate).



3. Once the endpoint is set you can modify the path to point to your file/data. In this example, we will connect to Lustre storage:
`lustre/or-hydra/cades-ops/proj-shared`
4. On the right side, set the endpoint. We will use OLCF Titan's file system. Search for `OLCF ATLAS`.
5. Again, you may adjust the path. Your home directory is default.

Creating an Endpoint on your Personal or Work Computer

It is easy to use your personal or work computer as a Globus Endpoint. Follow the instructions below.

Note: You may need to [create a firewall exception](#) for the Globus Personal Client. For configuration instructions, please consult the [details](#) on the Globus site.

1. Choose a descriptive name for your endpoint and click `Generate Setup Key`.
2. Copy the Setup Key. You will paste this into the software during setup.
3. Navigate to the [Globus Personal Connect webpage](#) to download the client onto your personal (or ORNL-owned) computer.
4. Click on the name of your operating system to obtain detailed instructions for installing the client and setting up the Endpoint.

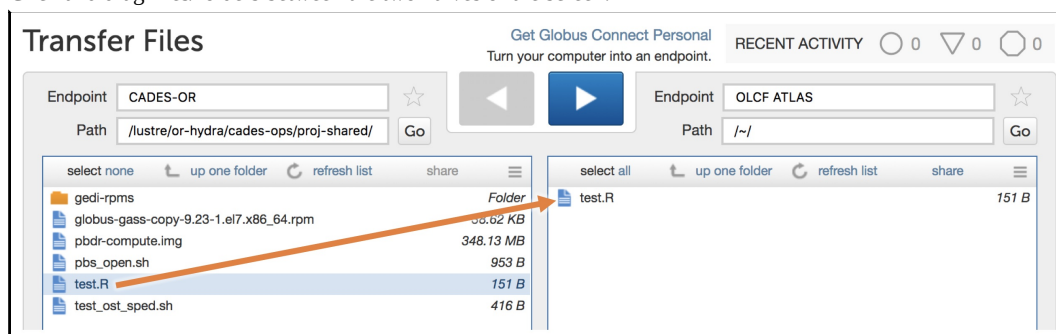


5. Once the client is installed, launch the program. You will be prompted to paste your setup key. **Note:** The Globus Personal Endpoint Client may produce errors if you are connected to the ORNL network via VPN.
6. Now you may use the Globus web interface or the [command line interface](#) to search for your new endpoint using the name you provided in step 1.

Globus Transfers & More

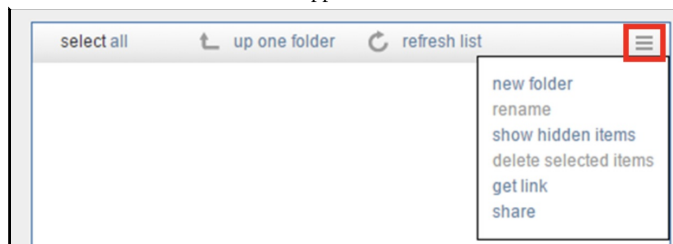
Globus File Transfers

1. Find the endpoints (on the left **and** right of the screen) you wish to use according to the [endpoints](#) instructions.
2. Modify the paths to the data you wish to transfer. *For this example, we will move a file from CADES Lustre storage to OLCF Atlas.*
3. Click and drag files/folders between the two halves of the screen.

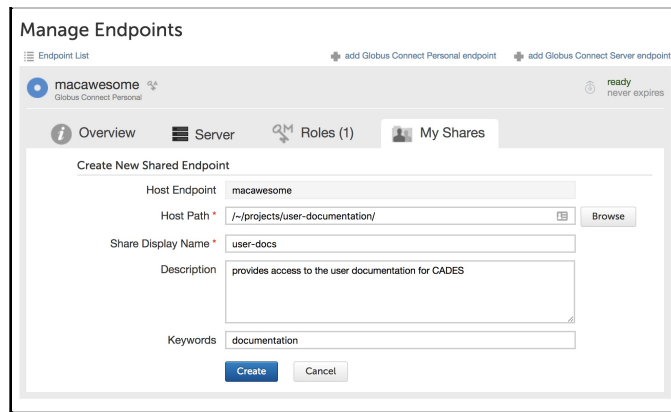


Additional Features

- **Create a Folder:** Globus also supports the creation of folders from within the browser interface.



- **Sharing Endpoints:** You can share endpoints with anyone who has a Globus account. If you are sharing from a managed endpoint (i.e. CADES OR) he or she will also need to have the proper credentials to access that resource.
 - Note:** Shared endpoints can only be created on personal endpoints if you have a subscription service through Globus. On CADES resources, shared endpoints may be requested by [emailing CADES](#).
 - *Sharing a personal endpoint:*
 - Navigate to the endpoint list that is administered by yourself: [here](#).
 - Click on the endpoint you would like to manage.
 - On the resulting screen, click the **My Shares** tab. Then click **+ Add Shared Endpoint**.
 - Fill out the required information, as shown below.



- Follow the instructions for configuring the shared endpoint. *For additional information of sharing files and endpoints, see the [Globus documentation](#).*

Globus Command Line Interface (CLI)

If you wish to utilize the Globus transfer tools from the command line, you can download the Globus Command Line Interface (CLI). It is available as a Python package.

Installing the Required Tools

Since the tool is a Python package, you will need Python installed, as well as the pip installer.

- Ubuntu:

```
sudo apt-get install python
sudo apt-get install python-pip
export PATH=~/.local/bin:$PATH
echo 'export PATH=~/.local/bin:$PATH' >> "$HOME/.bashrc"
```

- CentOS:

```
sudo yum install python
sudo yum install python-pip
export PATH=~/.local/bin:$PATH
echo 'export PATH=~/.local/bin:$PATH' >> "$HOME/.bashrc"
```

- macOS:

```
sudo easy_install python
sudo easy_install pip
export PATH=~/.local/bin:$PATH
echo 'export PATH=~/.local/bin:$PATH' >> "$HOME/.bashrc"
```

Some versions of Python will not be installed in `~/.local`. If you have trouble getting `globus` commands to execute, try the following commands to change the path:

```
GLOBUS_CLI_INSTALL_DIR="$(python -c 'import site; print(site.USER_BASE)')/bin"
echo "GLOBUS_CLI_INSTALL_DIR=$GLOBUS_CLI_INSTALL_DIR"
```

```
export PATH="$GLOBUS_CLI_INSTALL_DIR:$PATH"
echo 'export PATH="$GLOBUS_CLI_INSTALL_DIR":$PATH' >> "$HOME/.bashrc"
```

- Windows:

- The Windows package manager "Chocolatey" is recommended for installation. See [here](#) for Chocolatey installation instructions.
- To install Python and `pip`, see [here](#).

- All Operating Systems:

To install the Globus CLI, use the following command: `pip install --upgrade --user globus-cli`.

Optional: if you wish to use the Globus CLI from within a python virtual environment, see [instructions here](#). Otherwise, you may continue using this guide.

To start, you will need to log in to Globus: `globus login`. Follow the instructions to get logged in. A browser window may appear. To make sure that your login was successful, type `globus get-identities 'go@globusid.org'`. A successful output will look something like this: `c698d42e-d274-11e5-bf75-1fc5bf53bb56`.

Globus CLI Basics

- Endpoint Search

```
$ globus endpoint search 'CADES OR'
ID | Owner | Display Name
-----|-----|-----
57230a10-7ba2-11e7-8c3b-22000b9923ef | cades@globusid.org | CADES-OR
```

- Endpoint Management

- *Use variables for endpoint IDs:* Endpoint IDs are cumbersome. You cannot rename them, but you can store them as variables. For example:

```
epCADESOR=57230a10-7ba2-11e7-8c3b-22000b9923ef
```

Now you can use the variable to display information and manage files (with truncated output):

```
$ globus endpoint show $epCADESOR
Display Name: CADES-OR
ID: 57230a10-7ba2-11e7-8c3b-22000b9923ef
Owner: cades@globusid.org
Activated: True
Shareable: True
Department: CADES
Organization: Oak Ridge National Lab
Department: CADES
Visibility: True
Default Directory: /~/
Force Encryption: False
Managed Endpoint: True
```

- *Make a directory:*

```
globus mkdir $epCADESOR:~/example_dir
```

- *List the contents of a directory:*

```
$ globus ls $epCADESOR:~/
example_dir/
cades-user-guide.pdf
hello-world.c
hello-world.pbs
```

- *File transfer between endpoints:*

- First, search for a second endpoint. Then set that endpoint as a Bash variable.

```
$ globus endpoint search 'OLCF ATLAS'
ID | Owner | Display Name
-----|-----|-----
ef1a9560-7ca1-11e5-992c-22000b96db58 | olcf@globusid.org | OLCF ATLAS
$ epATLAS=ef1a9560-7ca1-11e5-992c-22000b96db58
```

- Make a single file transfer.

```
globus transfer $epCADESOR:/cades-user-guide.pdf $epATLAS:~/cades-user-guide.pdf \
--label "user-guide"
```

- Make a batch transfer.

```
$ globus transfer $epCADESOR:/example_dir/ $epATLAS:~/ \
--batch --label "CADES Batch" < in.txt
```


Related Tutorials

- [Globus Endpoints](#)
- [Graphical SFTP](#)

Scality User Guide

Glossary

AWS - Amazon Web Service

S3 - Simple Storage Service. Amazon's object storage service and, more generally, the protocol it uses.

Scality - An on-premises, object storage as a service, data archiving solution available to CADES users. This environment uses the same `aws` commands and S3 protocol as interacting with Amazon's S3 service. The mechanics of storing and retrieving data into either Amazon or Scality are functionally quite similar.

Table of Contents

1. [Getting Started](#)
2. [Basic Operations](#)
3. [Moving Files with a Graphical SFTP Client](#)

1. Getting Started

Requesting Access

- Internal ORNL (UCAMS) users may self-request access to the object storage resource [here](#).
 - Requests are normally approved within 24 hours after which you may log in to receive your access token.
- External (XCAMS) users **without** an existing account may create one [here](#).
 - Outside users who already have an XCAMS account should email ca-des-help@ornl.gov and request access to the object storage resource.
 - Users who have forgotten their XCAMS user ID or password may recover them [here](#).

Required software

The AWS Command Line tool is used for interaction with the storage service, and can be scripted for automated workflows. Installing the AWS CLI is summarized below, and you may consult the official AWS CLI [install guide](#).

CADES SHPC Users

- The `aws` client is provided via a software module, though you may install a [local version](#) in your home directory if you wish.
- From the SHPC login nodes:

```
-bash-4.2$module load python/3.6.1
-bash-4.2$aws --version
aws-cli/1.14.14 Python/3.6.1 Linux/3.10.0-327.4.4.el7.x86_64 botocore/1.8.18
```

Windows Users

- Download from <https://aws.amazon.com/cli/>.

OSX Users

- See AWS macOS [instructions](#).

Linux/OpenStack Users

You may encounter issues if `awscli` and the `awscli-plugin-endpoint` are installed from different sources e.g. one from your distribution's package manager (`apt` or `yum`) and one from `pip` . Installing both via `pip` usually allows them to work together well.

Note: It is recommended to install the components in a Python virtual environment, the instructions for which are available [here](#).

If you wish to install system-wide (as root) you may do so with `pip` via:

```
sudo pip install awscli
sudo pip install awscli-plugin-endpoint

# Endpoint Plugin
sudo pip install awscli-plugin-endpoint
# or
# pip install awscli-plugin-endpoint --user
```

Logging In

- Log in to the Scality authentication endpoint [here](#).
 - A dropdown presents options for each CADES group of which you are a member
 - For personal access: Choose **scalityiamuser**
 - For project (shared) access: Choose the appropriate group

Auth Tokens

You must log in and acquire an authentication token to interact with the storage environment. The token may be used for interactive or automated workflows until it expires, at which point it must be renewed.

You may have several active tokens, for accessing buckets owned by different projects, etc. and switch between them via the `--profile` option (below).

Retrieving the token and placing it into your environment currently requires a few steps. Progress is being made on a scripted method to automate these steps.

Edit your `~/.aws/credentials` file (or run `aws configure` and paste the appropriate values into the prompts) to create:

```
[default]
aws_access_key_id = <accessKey from JSON string>
aws_secret_access_key = <secretKeyValue from JSON string>
Default region name = <leave blank or us-east-1>
Default output format = <leave blank or text or json>
```

Re-authenticating / Token Expiry

Your AWS-CLI (command line interface) tool should be able to connect to the local Scality S3 instance until the expiration time listed in the JSON string.

When the access keys have expired the `aws` commands will produce an error message similar to the one below. Simply log in again and run `aws configure` again.

```
An error occurred (InvalidAccessKeyId) when calling the ListBuckets operation: The AWS access key Id you provided does not exist in our records.
```

Initial Profile Configuration

The `aws` commands need to know which endpoint to go to. By default the external Amazon S3 service is assumed, so we will change this to use the on-premises storage instead.

Set the default endpoint by running the below command for the configuration profile that you are going to use, which is typically `default`.

```
aws configure set plugins.endpoint awscli_plugin_endpoint
aws configure --profile default set s3.endpoint_url http://or-rda-s3.ornl.gov
aws configure --profile default set s3api.endpoint_url http://or-rda-s3.ornl.gov
```

Note: The first command enables the "endpoint" plugin, which allows easy switching between interacting with multiple internal (Scality) identities or external (AWS) accounts by passing a `--profile` argument. Your `~/.aws/config` and

`~/.aws/credentials` must have profiles and credentials defined for each identity.

Further information on configuring multiple named profiles:

<https://docs.aws.amazon.com/cli/latest/userguide/cli-multiple-profiles.html>

<https://github.com/wbinglee/awscli-plugin-endpoint>

2. Basic Operations

Integrated User Manual

The AWS-CLI tool has help text integrated into it. To invoke this, use `aws help`. To get detailed help about supported features, build your command line and post-pond `help` to the command. As an example if you want help with the S3 copy command, type:

```
aws s3 cp help
```

General Format

As we are dealing with the S3 service we will almost always be specifying one of two commands to run: `aws s3` or `aws s3api`.

Create a New Bucket

Buckets are storage areas similar to Unix volumes or Windows drives. With every `s3` command a bucket must be specified.

Create a new bucket for yourself or shared project. Do not use special characters, other than dashes or underscores.

Note: To provide shared access to other members of specific CADES groups, ensure you create the bucket using the authentication token appropriate for that group. To create a user-private bucket use the authentication token provided when selecting `scalityiamuser`. See section 1 for information on obtaining multiple tokens.

```
aws s3 mb s3://myproject
```

Listing S3 Buckets

Use the `s3api` command with `list-buckets` to display buckets visible to you:

```
aws s3api list-buckets
```

Example output: In this example the important string to note is the association with Buckets → Name:

```
{
  "Buckets": [
    {
      "Name": "eos",
      "CreationDate": "2017-08-23T21:58:17.405Z"
    }
  ],
  "Owner": {
    "DisplayName": "ornl.gov",
    "ID": "463e4bdd134ec2543672faef1066710cc90be348a18af455456518ec1dfd0818"
  }
}
```

Copying Files Into and Out of S3

Copying files into S3 is very similar to copying files on the Unix command line or SCP. The `aws` command is used, along with the endpoint specification, both common to all operations.

We specify the S3 service and that we want to copy files. The direction can either be local → S3 or S3 → local file system, simply by reversing the order.

```
aws s3 cp <local filename> s3://<bucket>/<remote filename>
```

Example:

```
aws s3 cp largefile s3://cades-8d73a078-94c6-4a73-a668-345fc6ee8618/largefile
```

Optionally adding `--profile` may be used to specify the named profile and matching credentials to be used.

Syncing Files

The S3 service provides a capability similar to that of the `rsync` command. Similar to the copy command the direction of synchronization can be either to S3 or from S3. The `<local directory>` can be relative or absolute. This is significantly faster if you have a moderate number of files.

```
aws s3 sync <local directory> s3://<bucket>/directory
```

Example:

```
aws --quiet s3 sync /home/xok/project/S3/scality_s3/benchmark/large s3://cades-8d73a078-94c6-4a73-a668-345fc6ee8618/large
```

When the `sync` operation is used a line is updated with the current command statistics. Above we see the optional parameter `--quiet`. This suppresses the update statistics output. This is useful when capturing command output as the progress bar normally fills log files with a large amount of unintelligible output.

Removing Files

Removing a single file:

```
aws s3 rm s3://<bucket>/<filename>
```

Example:

```
aws s3 rm s3://cades-8d73a078-94c6-4a73-a668-345fc6ee8618/largefile
```

Removing a directory

With the addition of the `--recursive` option an entire directory can be removed. Example:

```
aws s3 rm --recursive s3://cades-8d73a078-94c6-4a73-a668-345fc6ee8618/large_directory
```

List Files

To list the files in a bucket, type:

```
aws s3 ls <bucket>
```

Example:

```
aws s3 ls cades-8d73a078-94c6-4a73-a668-345fc6ee8618
```

3. Moving Files with a Graphical SFTP Client

CADES maintains documentation for CyberDuck and WinSCP clients. See [here](#).

Advanced Scality Operations & FAQ

1. Advanced Operations

Writing Data Directly From an Application

If your application would be sped up by skipping writing data to disk and instead writing directly to S3 this is possible in a sizable number of programming languages. Of relevance to scientific computing are the C++, Python, and Java SDKs. For a complete list please see the AWS Tools page (<https://aws.amazon.com/tools/>). We have tested the Python interface and have found it to be highly performant. Example Python script that puts the contents of the data string into a file called 'test.txt'. This works for serializable objects.

```
#!/usr/bin/env python
import boto3
s3 = boto3.resource('s3')
data = 'This is some test data in a string for S3'
s3.Bucket('cades-8d73a078-94c6-4a73-a668-345fc6ee8618').put_object(Key='test.txt', Body=data)
```

2. FAQ

Connection Hangs

Check that either the default s3 url has been set (see [Getting Started](#) section) or that the below option is added to the `aws` command :

```
--endpoint-url=http://or-rda-s3.ornl.gov
```

Example:

```
aws --endpoint-url=http://or-rda-s3.ornl.gov s3 ls cades-8d73a078-94c6-4a73-a668-345fc6ee8618
```

InvalidAccessKeyId Error

If your command or program exits with the below error, it means that it is either time to update your credentials or that they were not added to the `~/.aws/credentials` file correctly.

```
An error occurred (InvalidAccessKeyId) when calling the PutObject operation: The AWS Access Key Id you provided does not exist in our records.
```

awscli_plugin_endpoint

If you see the error:

```
ModuleNotFoundError: No module named 'awscli_plugin_endpoint'
```

You need to run the following command:

```
module load python
```


Setting up a Python Virtual Environment

Virtual environments are primarily useful for moments when you need to reconfigure or install software but need to be mindful of the influence on other programs and settings.

1. Install the virtual environment: `sudo apt-get python-virtualenv` or `sudo yum python-virtualenv`
2. Tell your system where to store the environment: `virtualenv $HOME/myEnv`
3. Activate the environment: `source $HOME/myEnv/bin/activate`

Use your environment as you normally would use your system.

When you are done, type `exit` to close the environment.

[CADES](#) → [User Documentation](#) → [Contributing](#)

Ways to Contribute

Would you like to make things better? There are a few ways you can contribute to improving our documentation and adding user-created tutorials or content.

1. Email your suggestions to the team <mailto:ca-des-help@ornl.gov>
2. Join our community on Slack! It's friendly. <http://ca-des.slack.com>
3. Want to change things? Feeling adventurous? Want to git savvy?
See [instructions for our git workflow](#) to branch our documentation repository and hack away. You got this.
4. We've made note of a few things to keep in mind while creating user content. You can find them in our [authoring guide](#).

[CADES](#) → [User Documentation](#) → [Contributing](#) → [Git Workflow](#)

Recommended Workflow for Git and Atom

GitLab is a popular platform to share code, store software solutions, and host documentation.

ORNL provides two GitLab servers <https://code.ornl.gov> and <https://code-int.ornl.gov>, the later being accessible only inside of ORNL.

Access to GitLab repositories is controlled by project owners. You may login and create your own projects and repositories, and share them with others.

While there are many text editors to choose from, Atom is recommended due to its ability to be customized and integrated with GitLab/Git.

Install Atom: <https://atom.io/>

Would you prefer not to use Atom? Here is git documentation for [git in the command line](#).

CADES User Documentation

Documentation published to CADES users is available within Gitlab at <https://code-int.ornl.gov/cades-ops/user-documentation> and users are encouraged to contribute to improving the material, or providing user created tutorials to share with the community.

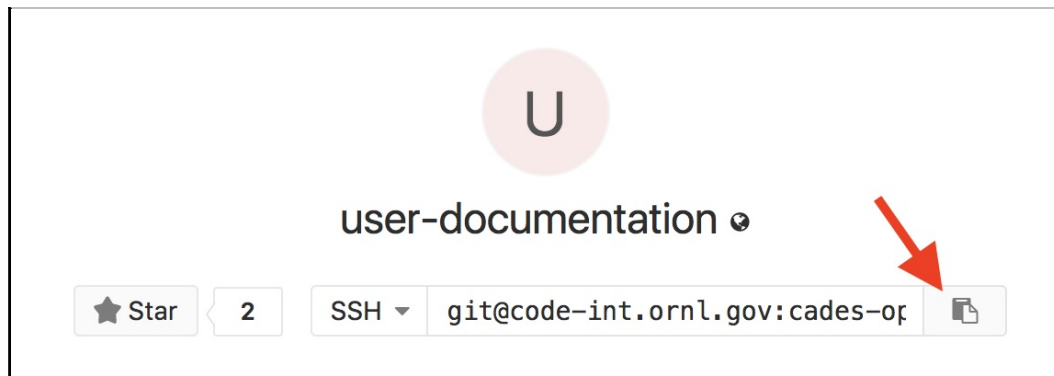
Configuring Atom and Git

Atom has several packages which enhance the user experience and some of them you'll need for the workflow. The packages can be installed by entering the `Settings` screen and choosing `Install`. There are thousands of packages, so try some out and have fun.

1. Necessary Packages:
 - `git-plus` (integrates Git)
 - `todo-show` (provides visuals for TODO and FIXME tags)
2. Optional but Recommended Packages:
 - `linter` (aids in code validation, will also need `linter` package for each language)
 - `minimap` (displays thumbnail version of document along with your location within it)
 - `git-checkout` (allows checking out remote branches within Atom)
3. Now that Atom is set-up, you'll need to make sure Git knows who you are. You'll only need to complete this step if you've never used Git on your machine before.
4. Open a terminal window.
5. If you need to install Git, see [here](#) for detailed instructions for popular operating systems.
6. To configure Git, input your user name and email as below:
 - `git config --global user.name "ab1"` (3-letter UCAMS username)
 - `git config --global user.email "nameab@ornl.gov"` (ORNL email)

Connecting Atom to a Repository

1. Navigate to the GitLab repository in your web browser. For this example, we'll use the `user-documentation` repository. Copy the SSH address to your machine's clipboard.



- Now, back inside Atom, open the `Command Palette`. On Mac, press `shift + command + p`. In Windows/Linux, press `control + shift + p`.
- Type `git clone` and press `enter`.
- Paste the ssh address in the resulting window. You may also modify the location of the local folder.
- Wait for the repository to clone.
- Now you can see the files have populated into the folder you specified. These files represent a local copy, to which you will make changes.
- Before you leave GitLab's webpage, consider adding your ssh key to your profile so you will not be prompted for credentials after every commit. To add your public ssh key to GitLab:
 - Click on your user image in the top-right of the GitLab window.
 - Select `Settings`.
 - On the left, click `ssh keys`.
 - Paste your **public** ssh key in the box, provide a title, and save by clicking `Add key`.

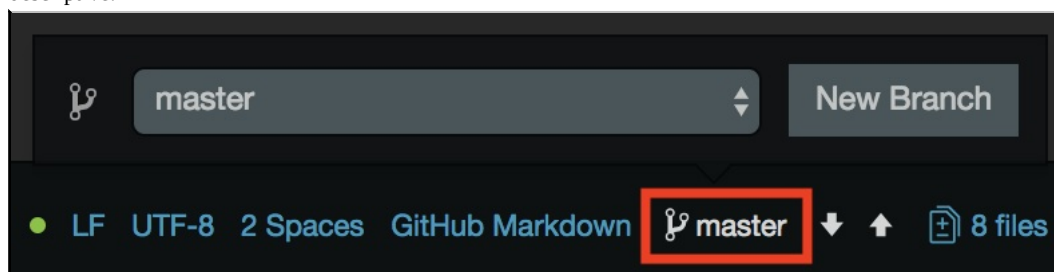
Note: You have now checked out the master branch of the remote repository. You may lack permission to push your changes to master, or may wish not to do so.

Working from Branches

At this point you likely either want to create a new branch and add your contributions there, or checkout a different branch you or someone else has already created. Each of these options are shown next. Unless you have reason otherwise you should choose one of these, rather than attempting to work from the master branch.

Create a New Branch

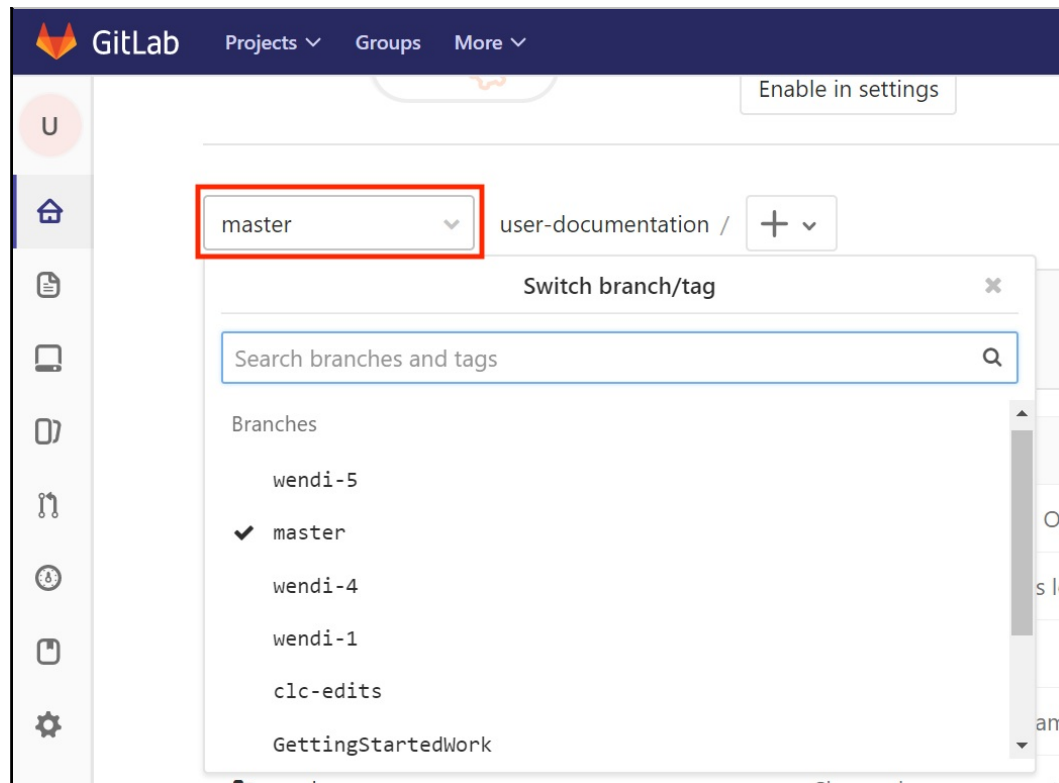
- Create a git branch by clicking on the `master` button on the bottom-right of the Atom window. Name the branch something descriptive.




Checkout an Existing Branch

We will need the name of the remote branch we wish to work on

- The GitLab project page displays a droplist with the name of available branches.




- Open the Atom command palette and search for `git checkout checkout` (requires `git checkout` plugin having been installed).




You may also open the checkout dialog directly using the hotkeys `ctrl + alt + shift + c` (or `ctrl + opt + shift + c` on Mac).

- The checkout dialog is a list of branches to checkout or switch to. Using the arrow keys, highlight the `custom` entry and hit enter. Now type in the name of the remote branch name. You may find this in GitLab.



- A notice will display if the checkout was successful. You may then switch between branches using the branch selector in the bottom-right toolbar.



Note: If checking out a remote branch within Atom using the 'git checkout' plugin, ensure you read the directions carefully. Do *not* click on the 'custom' branch, use the arrow keys and hit enter and then supply the name of the remote branch you wish to checkout. Lastly, if you make a mistake while typing in the branch name, you will end up creating a new branch with the typo. Be

sure to correctly type the branch name.

Command Line Branching

Rather than using the checkout dialogs in Atom, you may also list and checkout remote branches using the git command line tools.

List remote branches:

```
$ git branch -r
origin/GettingStartedWork
origin/HEAD -> origin/master
origin/master
origin/wendi-5
origin/user-contributions
```

Checkout a remote branch from the command line. `git checkout --track origin/wendi-5` after which you may select the branch within Atom.

GitLab GUI Editing

You don't have to use Atom for editing. You can hit the `edit` button in GitLab and edit directly, and preview before committing. Note that only repository owners can edit directly this way - otherwise you can create a branch and edit your branch directly.

Uploading Your Changes

1. Make changes to the files as needed.
2. You can open the Git window by pressing `^ + (`. Here you view unstaged and staged changes.
 - o `Unstaged` means the files are not ready to be committed.
 - o `Staged` means the files are ready to be committed.
3. Stage all of the files that you'd like to commit to the branch. This is accomplished by selecting the `+` symbol next to each file. Alternatively, on the top-right of the Git window, there is a `Stage All` button.
4. Commit your changes either within the Git window, or by entering `git commit` in Atom's command palette. Enter a commit message that will help you and others understand what changes were made. Then click `commit`.
5. Push changes to GitLab by typing `git push` in the command palette, or using Atom's up/down Git arrows located on the bottom-right of the window.

Note: If you get an error after typing `git push` that says "No upstream branch" open your terminal and navigate to the local copy of the repository. Then type `git push --set-upstream origin name_of_branch`. From then on you should be able to use the command palette to type `git push` or use the up/down arrows on the bottom-right of the window.

Creating a Merge Request

Merging your branch into the master branch, thereby makes your changes appear in the final version of the files.

You can send a merge request using the GitLab GUI.

1. From the left menu panel in Gitlab, select `Merge Request` then the green `New merge request` button.
2. Select your branch in the "Source Branch" side.
 - o Target branch will be `master`
 - o Click `compare branches`.
3. On the next screen the only thing needed is simply:
 - o Assign to: `< Project Owner, etc. >`

- Click `Submit merge request`

Previewing Changes

When you push a branch, your changes will get built in a "Review Environment" at http://user-documentation-stf011.granite.ccs.ornl.gov/_review/ and will allow you to click on the name of your branch.

A link to this review environment will be available on the Pull Request page that is created for your branch. This will allow us to take a quick glance at changes before we merge them into production.

Note: Preview sites are a prototype in the workflow. This feature may not always work and is offered as an unsupported convenience.

Git Workflow from the Command Line

There are many reasons one would prefer to work from the command line. Regardless of your reasons, here is how to contribute to the CADES documentation using only command line tools.

This guide is adapted from [GitLab's documentation](#).

It is assumed that users of this guide understand basic git/version control principles. To read more, visit [this page](#).

Install Git and Set-up

1. First, check to see if git is installed.

```
git --version
```

- o To install and/or update git using your package manager:

- CentOS, RedHat:

```
sudo yum install git
sudo yum update git
```

- Debian, Ubuntu:

```
sudo apt-get install git
sudo apt-get update git
```

- MacOS, use [Homebrew](#):

```
/usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/master/install)"
brew install git
brew upgrade git
```

- Windows: download [Git for Windows](#) and install it.

2. Set-up git with your access credentials to GitLab with the following commands:

```
git config --global user.name "your_username"
git config --global user.email "your_email_address@example.com"
```

- o You can review the information that you entered during set-up: `git config --global --list`

3. (*Optional*) Prior to cloning the repository, consider adding your ssh key to your GitLab profile so you will not be prompted for credentials after every commit. To add your public ssh key to GitLab:

- o Click on your user image in the top-right of the GitLab window.
- o Select `Settings`.
- o On the left, click `ssh keys`.
- o Paste your **public** ssh key in the box, provide a title, and save by clicking `Add key`.

Using Repositories and Branches

1. Clone an existing repository.

```
git clone git@code-int.ornl.gov:ca-des-ops/user-documentation.git
```


- If you have already cloned the repository but are returning to your local version after a while, you'll want to make sure your local files are up to date with the branch. You can pull updates from *master* or from *branch_name*.

```
git pull origin branch_name
```

2. You're ready to make edits using your favorite text editor. You will need to create a new branch or checkout an existing branch that will later be merged into the master branch. When naming branches, try to chose something descriptive.

- To create a branch: `git checkout -b branch_name`
- To list existing branches: `git branch -r`
- To checkout an existing branch: `git checkout --track origin/branch_name` or `git checkout branch_name`
 - Note: You may only have one branch checked out at a time.

3. When you are satisfied with your changes, commit them to your branch by adding and committing the changes.

```
git add --all
git commit -m "descriptive text about your changes"
```

4. After committing the edits, you'll want to push the changes to GitLab. If the following produces an error, see below the code snippet for common solutions. The structure of this command is `git push <remote> <branch>`.

```
git push
```

- Upstream error: `git push --set-upstream origin branch_name` or `git push -u origin branch_name`

Creating a Merge Request

At this time, GitLab does not natively support submissions for merge requests via the command line.

You can send a merge request using the GitLab GUI.

1. From the left menu panel in Gitlab, select `Merge Request` then the green `New merge request` button.
2. Select your branch in the "Source Branch" side.
 - Target branch will be *master*
 - Click `compare branches`.
3. On the next screen the only thing needed is simply:
 - Assign to: `< Project Owner, etc. >`
 - Click `submit merge request`

[CADES](#) → [User Documentation](#) → [Contribute](#) → [Authoring Guide](#)

Authoring Guide for CADES

Perhaps you've got some how-to documents tucked away in folders that you'd like to share with the CADES community. Or maybe you've discovered a way of doing things that would benefit other users.

You can submit your user guides for publication within the [CADES documentation site](#)! See the [contributing](#) page for instructions.

We've assembled here the fundamental authoring guidelines for CADES user documentation.

Document and Content Preferences

- Documents should be created using [markdown](#).
- Oak Ridge National Laboratory (ORNL) uses the [Chicago Manual of Style \(CMOS\)](#) as a basic style guide.
- Define the first instance of every acronym in each document. Ensure that the long form is not repeated after it is defined.
- Buttons and links that the user should "click" should go in `code`. For example, "Next, click the `Manage Rules` button."
- Put `📝` in front of NOTES. Renders:
- Use `▾` for those "carrot" drop-down menus. Renders:
- For headings: only use title case for the first three heading levels, `#`, `##`, and `###`. The remaining header levels should be sentence case.

Pictures and Images

Screenshots and images cannot be resized using markdown. Therefore, we embed `.html` that will be rendered when we publish the tutorial to the documentation site.

- Images and screenshots are stored in a folder `./screenshots/`.
- Files should be named descriptively. For example, use names such as `adding-IP-address.png` instead of `image03.png`.
- To remain consistent with other images in tutorials, please use the following `.html` code to resize, add a border, and open in a new browser tab when clicked. Note that you'll need to change the file name twice in the following code.

```
<a target="_new" href="screenshots/ssh_import_pub_key.png"></a>
```

Other Considerations

- Have you redacted sensitive information from text and images?
- Have you removed information that is protected by copyright?
- Are you using a specific version of your software and have you included in the documentation?

Related Topics

- Using a [Git Workflow](#) for creating user content.

Birthright Cloud User Policy

Oak Ridge National Laboratory's (ORNL) Compute and Data Environment for Science (CADES) now provides eligible customers with an OpenStack cloud computing solution with customizable virtual machines (VM). This new resource, called "Birthright Cloud," enables customers in science and technology directorates to leverage self-service portals to rapidly request these VMs for production, testing, and development.

Cloud computing provides an efficient pooling of on-demand, self-managed virtual infrastructure, consumed as a service. The OpenStack platform used here is an open-source cloud computing software solution that allows the creation of individual "Project" allocations for each user. Users can then fill these Project allocations with their own VMs without further intervention from CADES administrators—a true self-service implementation.

The CADES OpenStack Birthright Cloud allocations provide:

- **Self Service** – Through the Horizon web interface, users can create, manage, and delete VMs.
- **Portability** – Operations can be performed using any local system that provides a Bash terminal and SSH functionality.
- **Elasticity** – Users can create VMs on demand and delete them when they are no longer needed.

Disclaimers

If a concept or feature is not explicitly described within this policy, then it is not explicitly supported by the CADES team.

The only official copy of this document is this online electronic version found on <http://support.cades.ornl.gov>.

This policy is subject to change.

Acceptable Use

Computers, software, and communications systems provided by CADES are to be used for work associated with, and within the scope of, an approved project. The use of CADES resources for personal or non-work-related activities is strictly prohibited. All computers, networks, email, and storage systems are property of the US Government. Any misuse or unauthorized access is prohibited and is subject to criminal and civil penalties. CADES systems are provided to users without any warranty. CADES will not be held liable in the event of any system failure or data loss or corruption for any reason, including, but not limited to: negligence, malicious action, accidental loss, software errors, hardware failures, network losses, or inadequate configuration of any computing resource or ancillary system.

User Responsibilities

All Birthright Cloud users must comply with ORNL security rules and with the following:

- All operating system patches must be applied according to ORNL patching requirements.
- If user-sourced software images are uploaded, the user is responsible for keeping a copy of the image in case of accidental deletion or corruption.
- No moderate/confidential data should be mounted or copied to the VMs. Open science only.
- VMs should be removed from your OpenStack Project when they are no longer needed.
- VM operating systems must be updated or migrated before they reach an end-of-life development status.

Application for Resources

Birthright Cloud allocations are available to ORNL research and technical staff, by request, through CADES. The request is made through the ORNL XCAMS portal and requires your UCAMS ID. An activation notice will be sent when your resources are ready for use. CADES reserves the right to throttle access to Birthright Cloud allocations as resource constraints require.

How to: [Request Your CADES Birthright Cloud Allocation](#)

Authentication and Authorization

Users can access their Birthright Cloud allocation using a web-based GUI called "Horizon." See the Birthright Cloud user guide for details.

How to: [Manage Your OpenStack Project in Horizon](#)

Users are prohibited from changing or circumventing access controls to allow themselves or others to perform actions outside of their authorized privileges. In the event that an account is compromised, users must notify the CADES support team (ca-des-help@ornl.gov) immediately.

Users should also promptly inform the CADES support team (ca-des-help@ornl.gov) of any changes in their contact information (email, phone, affiliation, etc.).

The CADES team reserves the right to terminate accounts if any terms of this policy are violated.

Note: DO NOT share your credentials, passwords, private keys, or certificates, with anyone.

Account Access Maintenance

ORNL staff who have been granted a Birthright Cloud allocation have indefinite access to these resources for the duration of their time at ORNL and/or for as long as these resources are made available through CADES.

As underlying technologies and platforms change, users may be required to perform account access maintenance as needed.

Access at the End of a Project

When a user leaves ORNL, their Birthright Cloud allocation will close out, which results in the termination of account access and deletion of any remaining VMs running in their allocation. The user should move or save any data that he or she wishes to keep **before leaving ORNL**.

Computing Policy

CADES provides public VM images for Birthright Cloud customers. These images have been customized for better integration into both the ORNL environment as well as the user's scientific workflow, and are the only images fully supported by CADES. CADES will not provide support for user-provided images. If you still want to run a custom cloud image, or if you would like to inquire about migrating an image from an existing VM resource, please contact the CADES support team (ca-des-help@ornl.gov) for your request.

Because of the highly heterogeneous hardware architecture of CADES resources (in terms of processors, network interconnects, and disk technologies), and the fact that some of these resources are shared or may be overcommitted, CADES provides no guaranteed minimum performance level. However, if your application requires a certain level of performance that you have not

been able to obtain using your Birthright Cloud allocation, please contact the CADES support team (ca-des-help@ornl.gov) for assistance.

Storage Policy

The Birthright Cloud allocations, like other cloud solutions available on the market, provide a fungible resource that is subject to certain reliability constraints. Valuable data should be stored on a secondary storage solution, not exclusively on a user's Birthright Cloud VM Instance.

No moderate/confidential data should be mounted or copied to your Birthright Cloud VM Instance. Birthright Cloud VM Instances are for open science.

Storage Allocation

From inception, each Birthright Cloud allocation has a set storage quota. CADES reserves the right to change this storage quota at any time. See the [Birthright Cloud User Documentation](#) for details.

If a user requires more storage for their allocation, he or she can submit a proposal to the CADES Resource Utilization Council to request a storage quota increase. This proposal should describe the amount of storage desired and the scientific goal and merit of the work being performed using the CADES Birthright Cloud allocation. These requests will be reviewed by the council on a case-by-case basis.

Data Retention

When a project ends and a Birthright Cloud allocation is closed out, account access is terminated, and any remaining VMs and their associated data are deleted. Users are responsible for moving or saving any data that they would like to keep **before their project ends and their allocation is closed out**.

Backups

The Birthright Cloud allocations, like other cloud solutions available on the market, provide a fungible resource that is subject to certain reliability constraints. Valuable data should be stored on a secondary storage solution, not exclusively on a user's Birthright Cloud VM Instance.

At this time, there are no CADES-supported provisions for automatic backups of the VMs or their data. The user is responsible for backing up data and instances to their desired secondary storage solution.

Purge Policy or Quota

From inception, each Birthright Cloud allocation has a set quota, and CADES reserves the right to change this storage quota at any time. See the Birthright Cloud user documentation for details.

How to: [View your OpenStack Project Quota](#)

If a user requires more resources for their allocation, he or she can submit a proposal to the CADES Resource Utilization Council to request a quota increase. This proposal should describe the resources desired (RAM, CPUs, storage) and the scientific goal and merit of the work being performed using the CADES Birthright Cloud allocation. These requests will be reviewed on a case-by-case basis.

Special Requests and Policy Exemptions

Users can request policy exemptions by contacting the CADES support team (ca-des-help@ornl.gov). Requests are subject to review by the CADES Resource Utilization Council.

Acknowledging CADES

The following acknowledgment should be included in publications and presentations that contain work performed using CADES resources.

This research used resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.