Numerical Weather Prediction on AWS with WRF

Numerical Weather Prediction on AWS with WRF is an integrated AMI that allows AWS users to quickly set up and run their WRF cases. This guide provides instructions on how to launch EC2 instances and clusters running WRF with AMIs from the AWS Marketplace. These AMIs work either on single instances or in conjunction with AWS ParallelCluster for cluster launch. For users interested in working with a Linux desktop environment, the AMIs already include the extra elements and open the necessary ports; section 4 provides further details on how to create graphical sessions. This guide also includes descriptions on more advanced solutions such as AMI tailoring or storage options, which might be of interest to WRF users.

TABLE OF CONTENTS

Basic options

1. Launching EC2 instances running WRF ……………………………………………….. 2
   a. Account creation …………………………………………………………………………. 2
   b. Subscribing to the AMI ………………………………………………………………… 2
   c. Launching instances …………………………………………………………………… 2
   d. Connection to the instance …………………………………………………………… 3
   e. Running the apps ……………………………………………………………………….. 3
   f. Other precompiled app ……………………………………………………………….. 3
2. Data download, preprocessing and postprocessing ……………………………… 4
   a. Data download …………………………………………………………………………… 4
   b. Preprocessing tools ……………………………………………………………………. 4
   c. Postprocessing tools …………………………………………………………………… 5
3. Running an example case (Washington D.C. area nested simulation) ……… 6
4. Using a Linux desktop environment ………………………………………………… 8

Advanced options

5. Tailoring your AMI ……………………………………………………………………………. 10
6. Storage options ……………………………………………………………………………… 11

AWS offers a very large variety of services, which grow almost on a weekly basis. Many of our customers who are new to cloud environments and AWS use videoconferencing to get extra support. If either you are getting familiar with the AWS environment or wish to learn how to use the AMIs, feel free to contact us (support@odyhpc.com) to schedule any such session or for any other question.

1 - Launching EC2 instances running WRF

The following instructions discuss how to launch EC2 instances running WRF.

a. **Account creation** - Numerical Weather Prediction on AWS & AWS-Graviton2 with WRF can be used by any user with an active AWS account. If you are new to AWS, follow the instructions at https://aws.amazon.com/premiumsupport/knowledge-center/create-and-activate-aws-account/ to create a new account. You will also need to create an IAM user (https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_create.html) and assign the required permissions to create keypairs and launch instances (https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_change-permissions.html). You do not need the access & secret keys for running WRF on a single instance, but you will need them for more complex actions such as cluster launch or downloading data from S3 buckets. Either way, keep these keys safely as required by AWS.

b. **Subscribing to the AMI** - The AMI with the preinstalled software can be downloaded from the AWS Marketplace (https://aws.amazon.com/marketplace/pp/prodview-4axaf4xxxa6te). Make sure to understand the charges for AWS infrastructure and for the AMI.

c. **Launching instances** - Once your subscription is active, you can launch available instances based on your choices of architecture and configuration. More information about launching EC2 instances is available at https://docs.aws.amazon.com/quickstarts/latest/vmlaunch/step-1-launch-instance.html. The AMI is available in most regions and AZs. A few tips for launching instances follow:

   i. Spot instances are available per the usual conditions.

   ii. The AMI weights 75 GiB but have little free space. Most production runs will require far more free space with some of our customers reporting to have used disk space in excess of 1 TB. AWS makes easy to handle large storage needs and users have several options. Two easy options are to increase disk storage at the time of launching instances or to use EBS volumes. The latter can be particularly useful for large datasets to be recycled with different instances and clusters. Section 4 describes these options in greater detail.

   iii. Choose the SG (Security Group) according to you own needs. As a minimum, it should have port 22 open. This is the default configuration.

    iv. Once you have selected the region, you will need to use your own existing keypair or create a new one for launching WRF.

d. **Connection to the instance** – In order to connect to your instance, you must use a SSH client such as PuTTY, MobaXterm or a similar app (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AccessingInstances.html). Your username is ‘ubuntu’ and you enjoy superuser privileges.

e. **Running the apps** – The WPS and WRF apps are installed at the /home/ubuntu/PREPRO/WPS and /home/ ubuntu/WRF-4.4 subdirectories, respectively. Preparing the case files can be completed either with the pre-installed WPS or by directly transferring the input files. After the files are ready, WRF is run with $ mpirun -np N wrf.exe where N is the number of MPI ranks. If the case files and the executable are not in the same directory, you will need to either create a symlink with the file or include the subdirectory when invoking WRF /home/ubuntu/WRF-4.4/main/wrf.exe. We recommend running one MPI rank per core (e.g. it would be mpirun -np 64 ./wrf/exe for a c6i.32xlarge instance) but running in hyperthreaded mode requires adding the ‘–oversubscribe’ flag’ (e.g. mpirun –oversubscribe -np 128 ./wrf/exe).

f. **Other precompiled apps** - In addition to the latest WRF (ARW) release, the AMI incorporates precompiled executables for several WRF related apps:

    i. WRF-Chem with kinetic pre-processor at /home/ubuntu/WRF-Chem.
    ii. WRF Data Assimilation (WRFDA) at /home/ubuntu/DA/WRFDA.
    iii. WRFPLUS at /home/ubuntu/DA/WRFPLUS.
    iv. WRF-v3 is available per request (support@odyhpc.com).

2 – Data download, preprocessing, and postprocessing tools

In addition to the precompiled WRF executables, the AMI includes several pre and postprocessing tools. Because there is a constant flow of new apps and versions, feel free to contact us if you are interested in a particular one even if it is not discussed in the present document.

a. **Data download** -There are many different sources of meteorological data as discussed in the WRF documentation. The most commonly used data are GFS and NAM, with the latter being mostly used by U.S. organizations. As downloading GFS data, particularly from the NCEI dataset, is very time consuming, the AMI include a new script /home/ubuntu/DATA/GFS_download.sh that facilitates the download of GFS data from an S3 bucket. You will only have to select the type of data and dates, and the download will begin automatically. This significantly accelerates data download. The data will be stored at /home/ubuntu/DATA/GFS based on dates. The download of NAM or other meteorological data must follow standard procedures; contact us if you have any question about it.

b. **Preprocessing tools** - The subdirectory /home/ubuntu/PREPRO is set up to perform the usual preprocessing steps before running WRF. For example, and once the data is loaded in the appropriate directory and the namelist.wps file is modified to reflect the target simulation, the sequence of preprocessing commands would read:

```
$ cd /home/ubuntu/PREPRO/WPS
$ ./geogrid.exe
$ ln -sf /home/ubuntu/PREPRO/WPS/ungrib/Variable_Tables/Vtable.GFS
/home/ubuntu/PREPRO/WPS/Vtable
[$ ln -sf /home/ubuntu/PREPRO/WPS/ungrib/Variable_Tables/Vtable.NAM or alternative
/home/ubuntu/PREPRO/WPS/Vtable] for NAM or alternative data
$ ./link_grib.csh /home/ubuntu/DATA/GFS/gfs_3_2021
[$ ./link_grib.csh /home/ubuntu/DATA/NAM/nam_218_2021]
$ ./ungrib.exe
$ ./metgrid.exe
```

Once the met_em… files are generated, they can be either directly copied to the directory (e.g. /home/ubuntu/WRF-4.4/test/em_real) or symlinks can be set up.

c. Postprocessing tools - The AMIs also incorporate several postprocessing apps. This subsection outlines the available apps and any significant detail in their integration on the AMI. However, it does not describe them and the user if referred to the appropriate sources.

NCL (NCAR Command Language) is available for post-processing duties. The official website (https://www.ncl.ucar.edu/) furnishes documentation on the language itself and how to create high-quality graphics from WRF data.

WRF-python is also available and can be used as in any other Linux system.

UPP (Unified Post Processing) was developed by the National Center for Environmental Protection and comes precompiled with the AMIs. The official website (https://dtcenter.org/community-code/unified-post-processor-upp) provides further information and instructions on how to use it.

IDV (Integrated Data Viewer) can be opened from a graphical session on Intel-powered instances (see section II for details).

ARWpost is available at `/home/ubuntu/POSTPRO`. Once the WRF simulation is complete, ARWpost will convert the output files.

GrADS is a postprocessing app developed at George Mason University is also available at `/home/ubuntu/POSTPRO`. Starting to process some data simply requires opening the tool:

```
$ cd /home/ubuntu/POSTPRO/grads-2.2.2
$ ./grads
```

3. Running an example case (Washington D.C. area nested simulation)

The AMIs come with a test case consisting of a parent domain covering the continental U.S. and a nested domain covering the Washington D.C. area. This case has been previously preprocessed and the initial and lateral boundary condition files (wrfinput_d01, wrfinput_d02 and wrfbdy_d01) are located at the /home/ubuntu/WRF-4.4/test/em_real subdirectory. The test-case, which can be used either as it-is or slightly modified, can help users to either get acclimated with the AWS environment and the AMI itself. Running the test case as-it-is simply requires moving into the subdirectory where the initial and boundary conditions are and invoking the app:

```
$ cd /home/ubuntu/WRF-4.4/test/em_real
$ mpirun -np N ./wrf.exe  (where N represents the number of MPI ranks)
```

Something that always need to be kept in mind when running HPC simulations is memory requirements. WRF is not particularly memory hungry compared to other HPC apps. The memory requirements can fluctuate slightly depending on the number of MPI ranks; for the test case, it requires a minimum of 10 GB, which makes it suitable to run on single instances with 16 GB of memory or higher. If memory is insufficient, the simulation will break down with an error message like:

```
--------------------------------------------------------------------------
Primary job  terminated normally, but 1 process returned
a non-zero exit code. Per user-direction, the job has been aborted.
--------------------------------------------------------------------------
--------------------------------------------------------------------------
mpirun noticed that process rank 2 with PID 0 on node ip-172-31-3-9 exited on signal 9 (Killed).
--------------------------------------------------------------------------
```

If memory is sufficient and everything else goes without a problem, the app will complete the simulation and the output files will be in the /home/ubuntu/WRF-4.4/test/em_real subdirectory unless otherwise indicated. It is easily feasible to modify the test case in many ways to test different aspects of the simulations such as space and time resolutions or the meteorological data used to generate initial and boundary conditions (the data for the Global Forecast System with a resolution of 1 degree for the simulated period of time are also included in the subdirectory /home/ubuntu/DATA/GFS so that it can be easily preprocess, run and compare versus the forecast with NAM data).

The test case has a major domain covering the continental U.S. with a grid of 611 by 371 points and a resolution of 8 x 8 km^2. The vertical resolution is 33 points and the time resolution for this major domain is 45 s. The data are generated for a period of 48 hours starting on January 1st of 2022.

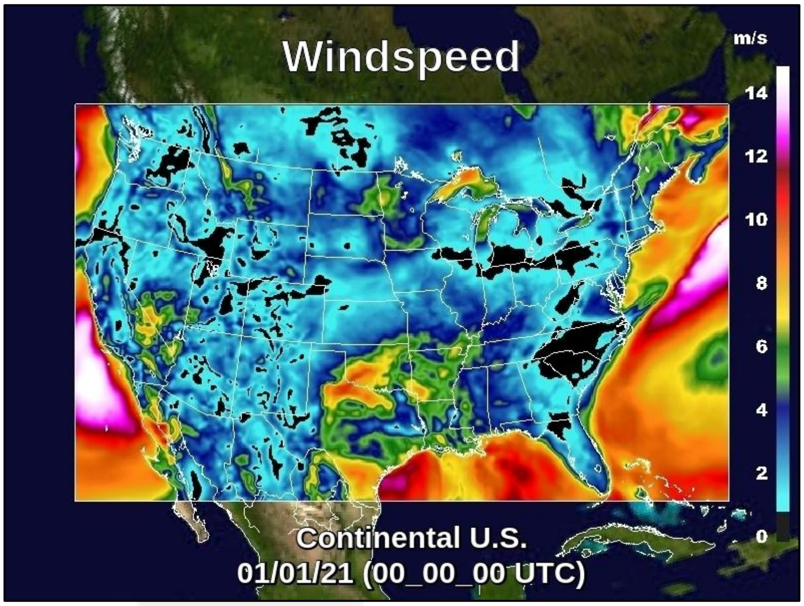

The nested domain of the test case centers around Washington D.C. and covers most of states of Maryland, Virginia, and Delaware and a small of portion of New Jersey, Pennsylvania, and West Virginia. The grid is 200 by 200 points with a resolution of 1,600 x 1,600 m^2 (approximately one square mile). The time-step for this nested domain is 9 s.

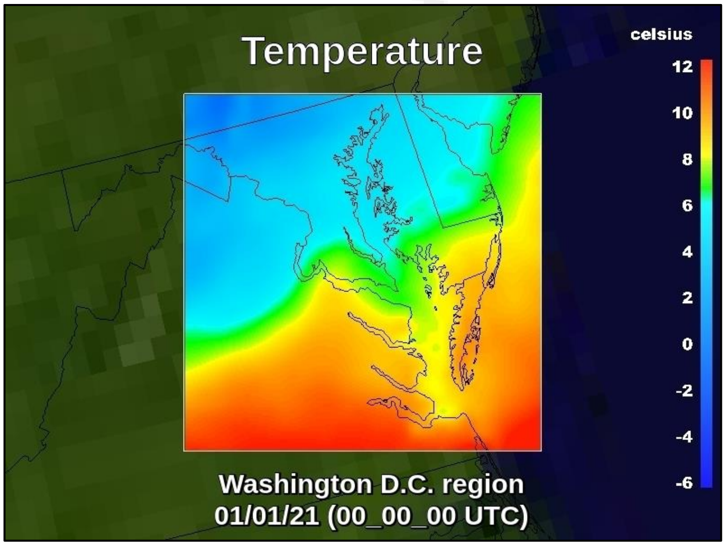

4. Using a Linux desktop environment

WRF users exploring the option of migrating to the public cloud cite the lack of a graphical environment as one of the major drawbacks associated with this option. To overcome this challenge with a simple procedure, the AMIs are fully ready to start a Linux desktop environment session. Launching the AMI will also open the necessary ports so that users do not need to modify SGs or worry about other networking setups.

The AMIs utilize a high-performance display app named NICE-DCV (https://docs.aws.amazon.com/dcv/index.html) and currently owned by AWS. In the case of single instances, the only steps to initialize the desktop environment are to start the server and initialize the session[1] from a terminal:

```
$ sudo systemctl start dcvserver
$ dcv create-session custom_name_for_your_session
```

To check if the session is running and its characteristics, you can simply type:

```
$ dcv list-sessions
$ dcv create-session custom_name_for_your_session
```

The session will be available in most web browsers at https://External_IP_for_your_instance:8443/#custom_name_for_your_session.
The first 2 screens will inform you of a potential certificate issue; this is normal as the app uses a self-certificate with NGINX.

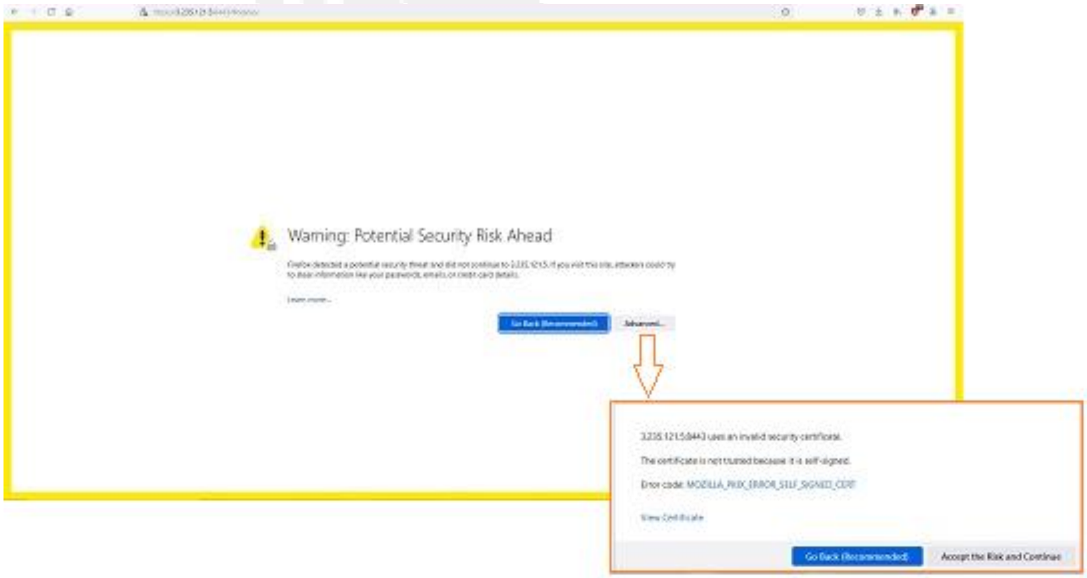

Warning: Potential Security Risk Ahead

Firefox detected a potential security threat and did not continue to 3.235.121.5. If you visit this site, attackers could try to steal information like your passwords, emails, or credit card details.

Learn more...

3.235.121.5:8443 uses an invalid security certificate.

The certificate is not trusted because it is self-signed.

Error code: MOZILLA_PKIX_ERROR_SELF_SIGNED_CERT

View Certificate

[1] Graphical sessions consume resources quickly, particularly memory. Although it is possible to start a session with 4 GB of memory, we recommend using an instance with a minimum memory of 8GB or even more if any visualization app is going to be used.

After clicking on ‘Accept Risk and Continue,’ the GNOME welcome screen will appear:

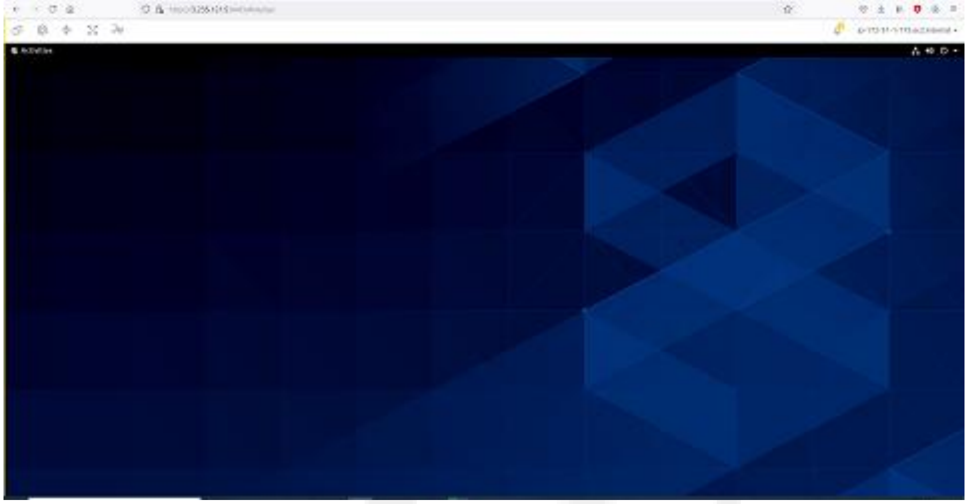

If, for example, IDV is installed and opened, the app can be used as with any other graphical desktop:

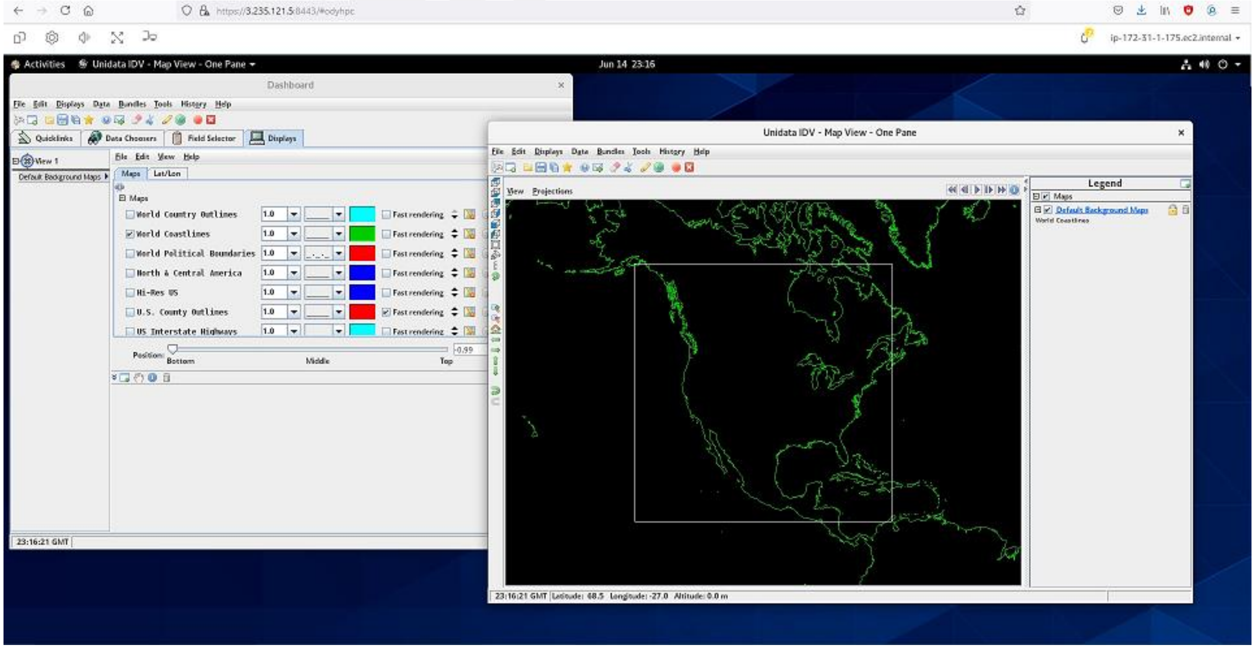

5. Tailoring your AMI

The AMI is built with the latest stable version of the app and with the case discussed in section 3. However, this is a relatively simple case, and most users will want to compute much more complex cases. It is possible to use the provided WPS to build the files necessary a new case or to simply transfer them. If this new case is going to be reused several times or you wish to store it long-term, we strongly recommend using AWS EBS (Elastic Block Storage) and AWS S3 (Simple Storage Service), respectively.

Before launching the instance and building the new case, it is important to assess how much space will be needed. Most of our custom AMIs have little free space (around 10 GiB) but, at the time of launching an instance, it is feasible to increase the available space. Using the console provides you with this feature in Step 4. You can select between increasing the size of the root volume or adding a new volume. The former is easy to select, and no further step is needed, whereas the latter requires mounting the new volume but has the advantage of dedicated EBS to that case (see the storage options section for more details).

Once the new case is available and any other tailoring is complete, you can build your own AMI by clicking on 'Creating Image' as shown in the screenshot or using the CLI.

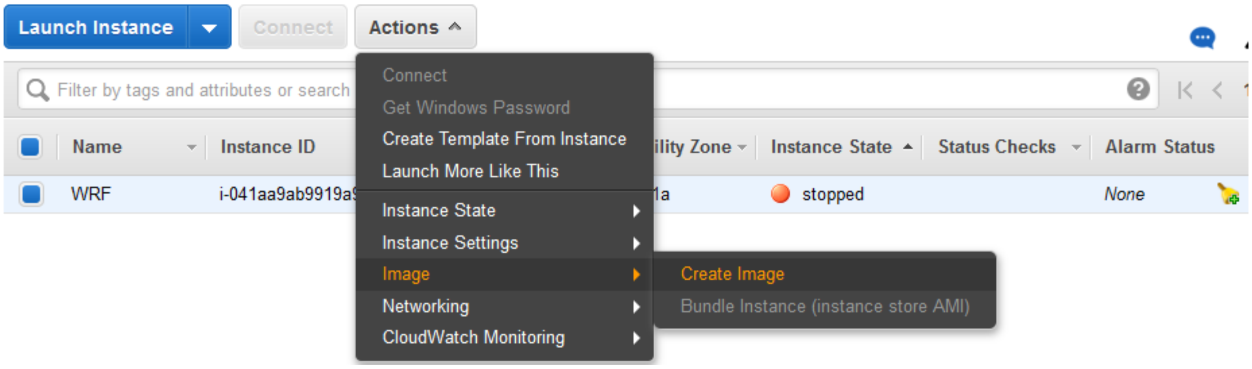

Name	Instance ID	ility Zone	Instance State	Status Checks	Alarm Status
WRF	i-041aa9ab9919a	1a	stopped		None

The new AMI is available only to IAMs in your account. It has the same features as the original odyhpc AMI including its availability for launching clusters. However, and unlike the odyhpc AMI, it is available only in the region where created.

6. Storage options

One of the many advantages of using AWS infrastructure is the availability of many storage solutions meeting different needs and criteria. It is not the purpose of this section to cover every available option, which has its own voluminous documentation, but to describe a few options of interest to WRF users. The two most common options are EBS (Elastic Block Store) and S3 (Simple Storage Service), which should meet most demands. Two more advanced options include the use of Lustre filesystems for high performance I/O attached to instances or clusters, and Amazon EFS (Elastic File System), which can provide a general filesystem attached to several instances or clusters.

The use of EBS is the most common as AWS AMIs are imprinted in this format at the time of instance creation. The AMI contains, the OS, apps, and dependencies, and constitute the root volume associated to the instance. In the case of AMIs containing WRF, the minimum root volume is usually around 75 GB with little free space. When launching an instance, increasing the available disk space can be accomplished either through increasing the root volume size (the system will not accept a smaller size) or by attaching another EBS to the instance as shown in the screenshot.

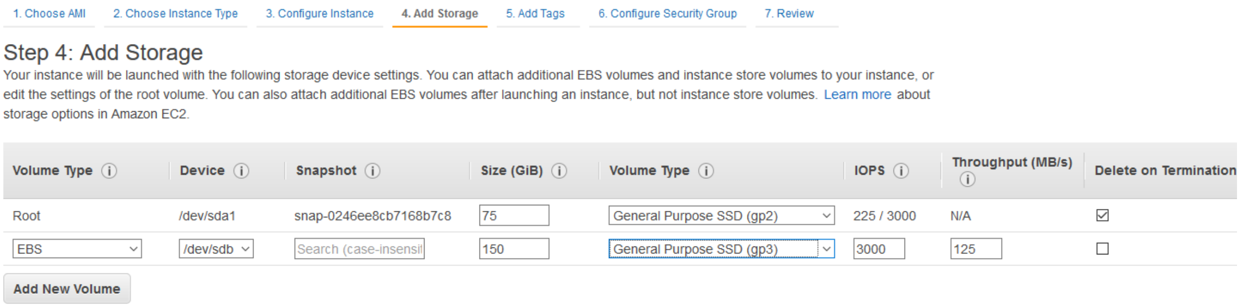

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. Learn more about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination
Root	/dev/sda1	snap-0246ee8cb7168b7c8	75	General Purpose SSD (gp2)	225 / 3000	N/A	☑
EBS	/dev/sdb	Search (case-insensi	150	General Purpose SSD (gp3)	3000	125	☐

In this example, we are adding 150 GB of disk space with the caveat that our choice is of the type gp3, which offers better performance than the standard gp2 type. A discussion on the technical specifications of different EBS types is available, for example, at https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html. The advantage of using this extra EBS volume (instead of simply increasing the root volume) is that the data in this filesystem can be used separately from the instance including its attachment to other instances or clusters. This extra space is not immediately available, but it needs to be mounted. An example of how to mount this EBS volume on the data subdirectory would be:

```
$ lsblk                              ! To check the name of the volume
$ sudo file -s /dev/nvme1n1
$ sudo mkfs -t xfs /dev/nvme1n1
$ sudo mkdir data
$ sudo mount /dev/nvme1n1 data
$ sudo chown ubuntu:ubuntu data
```

After creating or running any case, and to reuse or store this block, we can simply unmount it `$ sudo umount -d /dev/nvme1n1`, detach the volume and create a snapshot from the console volume menu.

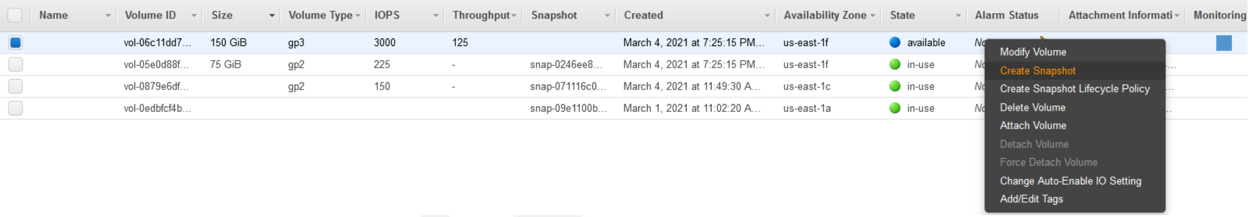

□	Name	Volume ID	Size	Volume Type	IOPS	Throughput	Snapshot	Created	Availability Zone	State	Alarm Status	Attachment Informati	Monitoring
■		vol-06c11dd7...	150 GiB	gp3	3000	125		March 4, 2021 at 7:25:15 PM...	us-east-1f	● available	No		■
□		vol-05e0d88f...	75 GiB	gp2	225	-	snap-0246ee8...	March 4, 2021 at 7:25:15 PM...	us-east-1f	● in-use	No	...	
□		vol-0879e6df...		gp2	150	-	snap-071116c0...	March 4, 2021 at 11:49:30 A...	us-east-1c	● in-use	No	...	
□		vol-0edbfcf4b...					snap-09e1100b...	March 1, 2021 at 11:02:20 A...	us-east-1a	● in-use	No	...	

Also, it is important to notice that we did not select 'Delete on Termination' previously so terminating the instance will not delete the volume and the meter will keep running until its deletion from the console or alternatively the CLI.

S3 is more adequate for long-term storage, backup, and access to large databases. Among the latter, the OpenData_list file provides a list with available datasets for WRF use. To retrieve this data, the command `aws s3 cp s3://ncar-cesm-lens//filename --no-content` will copy it to your instance. To access non-public buckets belonging to your organization, it is necessary to configure your credentials with `$ aws configure`, which will ask for the IAM credentials (aws_access_key_id and aws_secret_access_key). There is plenty of online help and discussions on AWS S3, plus the official documentation (https://docs.aws.amazon.com/s3/index.html), which is voluminous.