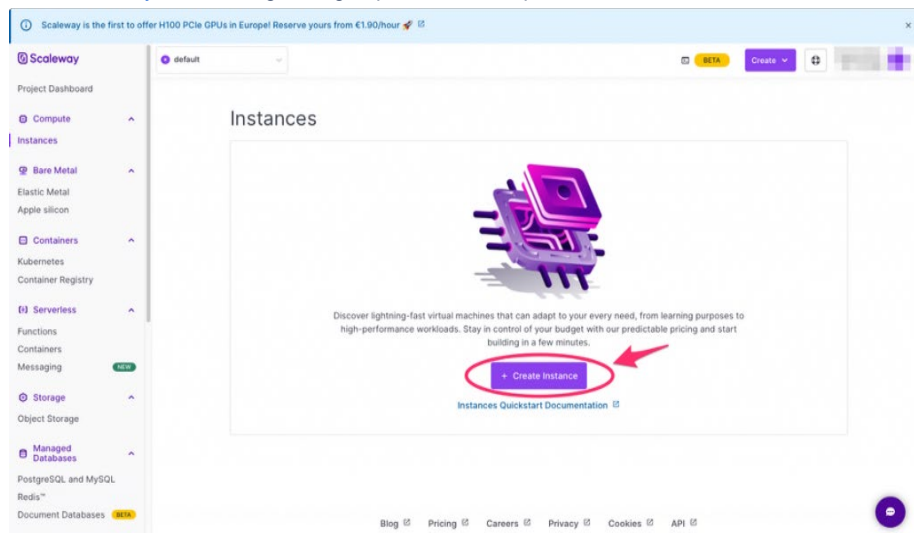
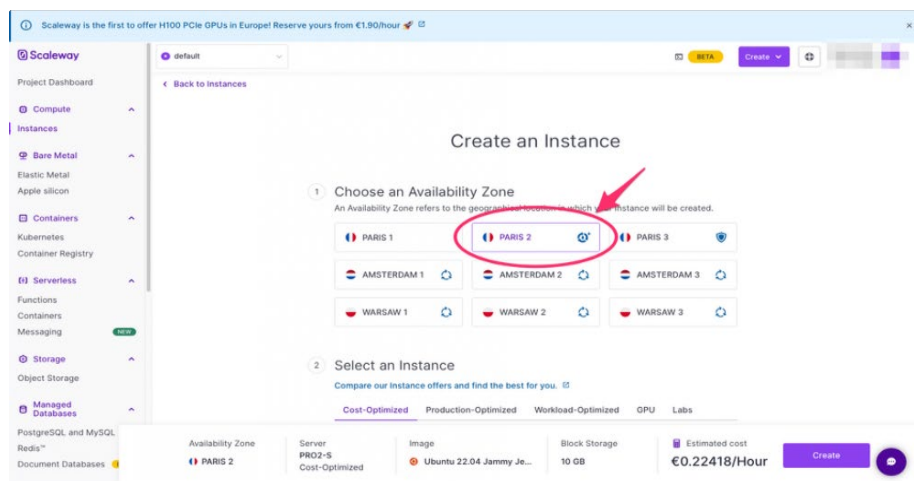


Setting up Ampere Optimized AI Frameworks on Scaleway's COP-ARM Virtual Machines

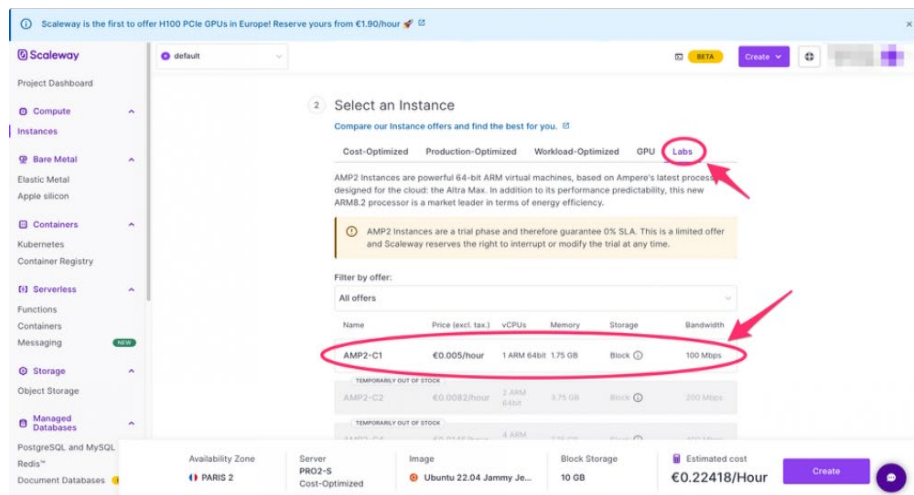
1. Go to scaleway.com and sign-in/sign-up. As a next step, click on **Instances**, then on **Create Instance**.



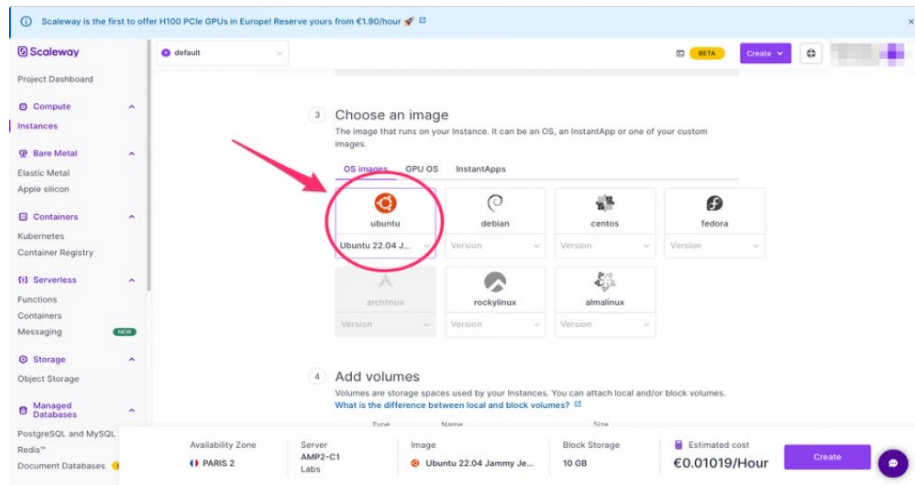
2. Select your **Availability Zone** - as of November 3rd, 2023, COP-ARM instances are available in **PARIS 2** zone only.



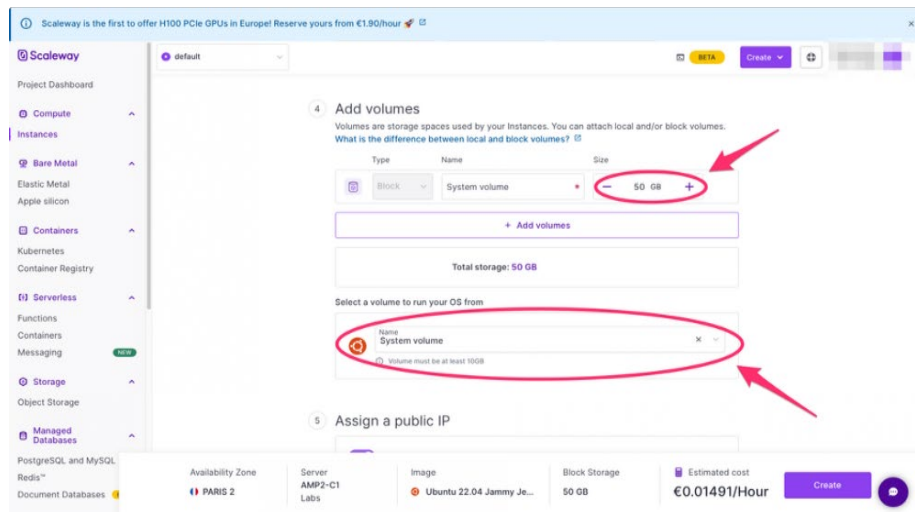
3. Click on **Labs** and select your **COP-ARM** offer accordingly to your requirements.



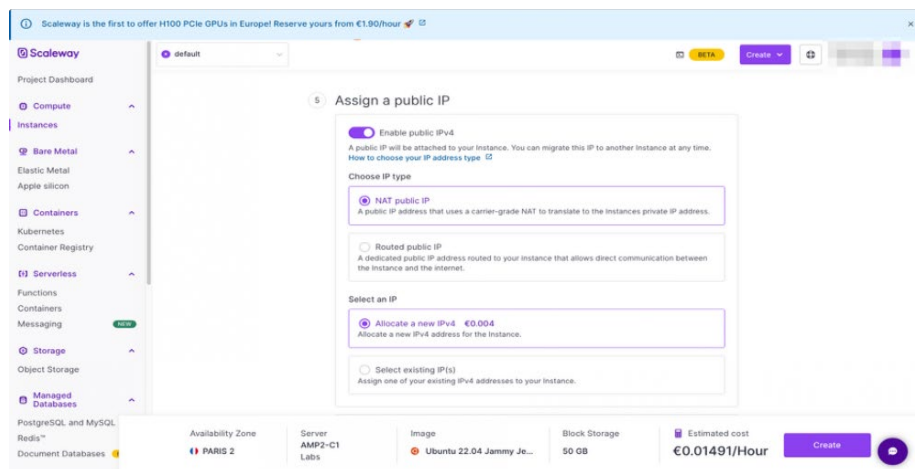
4. In **OS images** section, select **Ubuntu 22.04 Jelly Fish**.



5. Adjust the size of your **System volume** and pick this volume as a **Volume to run your OS from**.



6. Leave IP settings in their default state - assign a public IPv4 address to the Instance.



7. Name your Instance arbitrarily.

The screenshot shows the Scaleway console interface. On the left is a sidebar with navigation links: Project Dashboard, Compute, Instances, Bare Metal, Containers, Serverless, Storage, and Managed Databases. The main panel displays the details for an instance named 'ampere-ai' (AMP2-C1). A red arrow points to the 'Public IP' field, which is circled in red. The instance is in a 'Running' state. Below the instance information, there are tabs for Overview, Attached volumes, Images, Private Networks, Advanced settings, and More. The instance information table includes details like Status (Running), Type (AMP2-C1), From image (Ubuntu 22.04 Jammy Jellyfish), Availability Zone (PAR 2), Cores (1 ARM 64bit), RAM (1.75 GB), Local storage (Block), Bandwidth (100 Mbps), Instance ID, Image ID, Volumes (1), Block Storage (50 GB), Public IP (51.159.158.92), Private IP (10.204.188.53), IPv6, Gateway, and Netmask.

8. Click on **Add SSH key** and paste your pre-generated SSH key.

The screenshot shows the 'SSH keys' section in the Scaleway console. A red circle highlights the '+ Add SSH key' button. Below this, there is a section for 'Estimated cost' which provides a summary of the cost estimation based on the configuration. The cost table shows the following items and their costs:

Item	Configuration	Cost
Availability Zone	PARIS 2	€0.00
Server	AMP2-C1 Labs - 1 ARM 64bit - 1.75 GB Memory	€0.00501
Image	Ubuntu 22.04 Jammy Jellyfish	€0.00
Block Storage	50 GB	€0.0059
Flexible IP	Yes	€0.004
Total		€0.01491

9. Validate cost estimation and click on **Create Instance**.

The screenshot shows the 'Create Instance' button at the bottom of the console, which is circled in red. The button is labeled 'Create Instance' and is located below the 'Estimated cost' section. The cost table from the previous screenshot is visible above the button.

10. Wait for the instance to start and then copy its **Public IP** address.

11. SSH into the system by specifying user **ubuntu** @ public IP address of your Instance.

```
Last login: Fri Nov  3 17:19:08 on ttys002
~ % ssh ubuntu@51.159.158.92
```

12. After accessing the Instance, run:

```
1 sudo apt-get update && sudo apt-get install -y docker.io
```

To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

```
ubuntu@ampere-ai:~$ lscpu
Architecture:          aarch64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                1
On-line CPU(s) list:   0
Vendor ID:             ARM
Model name:            Neoverse-N1
Model:                1
Thread(s) per core:    1
Core(s) per socket:    1
Socket(s):             1
Stepping:              r3p1
BogoMIPS:              50.00
Flags:                 fp asimd evtstrm aes pmull sha1 sha2 crc32 atomics fphp asimdhp cpuid asimdrrm lrcpc dcpop asimddp ssbs
NUMA:
NUMA node(s):          1
NUMA node0 CPU(s):     0
Vulnerabilities:
Itlb multihit:         Not affected
L1tf:                  Not affected
Mds:                   Not affected
Meltdown:              Not affected
Mmio stale data:       Not affected
Retbleed:              Not affected
Spec store bypass:     Mitigation; Speculative Store Bypass disabled via prctl
Spectre v1:            Mitigation; __user pointer sanitization
Spectre v2:            Mitigation; CSV2, BHB
Srbds:                 Not affected
Tsx async abort:       Not affected
ubuntu@ampere-ai:~$ sudo apt-get update && sudo apt-get install -y docker.io
```

13. Allow the docker daemon to install. In the meantime go to [Ampere AI DockerHub](#) and select your **Ampere optimized framework** of choice.

Hackathon time! Join us for the Docker AI/ML Hackathon now through November 7th. [Sign up now](#)

dockerhub

Search Docker Hub

Sign inSign up

Ampere Computing AI

Community User

Ampere Computing AI

Santa Clara, CA

<https://solutions.amperecomputing.com/solutions/ampere-ai>

Joined November 25, 2022

Repositories

Starred

Displaying 1 to 3 repositories

amperecomputingai/tensorflow

430

0

By [amperecomputingai](#) · Updated a month ago

Image

amperecomputingai/pytorch

604

2

By [amperecomputingai](#) · Updated a month ago

Image

amperecomputingai/onnxruntime

43

0

By [amperecomputingai](#) · Updated a month ago

Image

Hackathon time! Join us for the Docker AI/ML Hackathon now through November 7th. [Sign up now](#)

dockerhub

Search Docker Hub

Sign inSign up

Explore / amperecomputingai/pytorch

amperecomputingai/pytorch

☆

Pulls 604

By [amperecomputingai](#) · Updated a month ago

Ampere® optimized PyTorch (<https://solutions.amperecomputing.com/solutions/ampere-ai>)

Image

Overview

Tags

Ampere® optimized PyTorch Docker images

Ampere® Altra®, with high performance Ampere Optimized Frameworks, offers the best-in-class Artificial Intelligence inference performance for standard frameworks including TensorFlow, PyTorch and ONNXRT. Ampere optimized PyTorch inference acceleration engine is fully integrated with PyTorch framework. PyTorch models and software written with PyTorch API can run as-is, without any modifications.

Starting container

```
sudo docker run --privileged=true --name torch-ai0 --network host -it amperecomputingai/pytorch:late
```

Please refer to the [Ampere optimized PyTorch documentation](#) for additional details.

Make sure to visit us at [Ampere Solutions Portal](#)!

Support

Please contact us at ai-support@amperecomputing.com

LEGAL NOTICE

By accessing, downloading or using this software and any required dependent software (the "Ampere AI Software"), you agree to the terms and conditions of the software license agreements for the Ampere AI Software, which may also include notices, disclaimers, or license terms for third party software included with the Ampere AI Software. Please refer to the [Ampere AI Software EULA v1.6](#) or other similarly-named text file for additional details.

Docker Pull Command

```
docker pull amperecomputingai/pyto...
```



```
1 sudo docker run --privileged=true --name ampere-torch --network host -it amperecomputingai/pytorch:latest # o

Selecting previously unselected package ubuntu-fan.
Preparing to unpack .../7-ubuntu-fan_0.12.16_all.deb ...
Unpacking ubuntu-fan (0.12.16) ...
Setting up dnsmasq-base (2.86-1.1ubuntu0.3) ...
Setting up runc (1.1.7-0ubuntu1~22.04.1) ...
Setting up dns-root-data (2021011101) ...
Setting up bridge-utils (1.7-1ubuntu3) ...
Setting up pigz (2.6-1) ...
Setting up containerd (1.7.2-0ubuntu1~22.04.1) ...
Created symlink /etc/systemd/system/multi-user.target.wants/containerd.service → /lib/systemd/system/containerd.service.
Setting up ubuntu-fan (0.12.16) ...
Created symlink /etc/systemd/system/multi-user.target.wants/ubuntu-fan.service → /lib/systemd/system/ubuntu-fan.service.
Setting up docker.io (24.0.5-0ubuntu1~22.04.1) ...
Adding group `docker' (GID 121) ...
Done.
Created symlink /etc/systemd/system/multi-user.target.wants/docker.service → /lib/systemd/system/docker.service.
Created symlink /etc/systemd/system/sockets.target.wants/docker.socket → /lib/systemd/system/docker.socket.
Processing triggers for dbus (1.12.20-2ubuntu4.1) ...
Processing triggers for man-db (2.10.2-1) ...
Scanning processes...
Scanning processor microcode...
Scanning linux images...

Running kernel seems to be up-to-date.

Failed to check for processor microcode upgrades.

No services need to be restarted.


No containers need to be restarted.

No user sessions are running outdated binaries.

No VM guests are running outdated hypervisor (qemu) binaries on this host.
ubuntu@ampere-ai:~$ sudo docker run --privileged=true --name torch-ai0 --network host -it amperecomputingai/pytorch:latest
```

```
1 python3 -c "import torch" # or 'onnxruntime', or 'tensorflow'
```

```
c81ba129bfa8: Pull complete
48547d582fee: Pull complete
b6ee8912bdd5: Pull complete
876c32dc7fd1: Pull complete
e163b5bdc616: Pull complete
464d98940764: Pull complete
5091cb71e717: Pull complete
758a390adbc6: Pull complete
cc28376d26f8: Pull complete
Digest: sha256:7d8e83b1cec60b106120dbdb1d657d5bb1aedc0c247573d7ad0396b432341d5a
Status: Downloaded newer image for amperecomputingai/pytorch:latest
```



Thank you for choosing AI0!

Please visit us at <https://solutions.amperecomputing.com/solutions/ampere-ai>

For quick start please run:

```
cd /workspace/ai0-examples/; bash start_notebook.sh
root@ampere-ai:/workspace# python3 -c "import torch"
```

16. Congratulations - you have successfully deployed Ampere optimized framework on your COP-ARM instance! In case of any issues please reach out to us via email: ai-support@amperecomputing.com