



Hands-on tutorials

Deploy a P4d EC2 UltraCluster



Deploy a P4d EC2 UltraCluster: Hands-on tutorials

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Deploy a P4d EC2 UltraCluster

Overview

What you'll accomplish

Implementation

Congratulations!

i

1

1

2

6

Deploy a P4d EC2 UltraCluster

AWS experience	Intermediate
Time to complete	10 minutes
Last updated	December 8, 2020

Overview

[Amazon EC2 P4d instances](#) deliver the highest performance for machine learning (ML) training and high performance computing (HPC) applications in the cloud. Amazon EC2 P4d instances are deployed in clusters called EC2 UltraClusters that are comprised of the high performance compute, networking, and storage in the cloud. Each EC2 UltraCluster of P4d instances comprises more than 4,000 of the latest NVIDIA A100 GPUs, Petabit-scale non-blocking networking infrastructure, and high throughput low latency storage with FSx for Lustre.

This step-by-step tutorial will help you launch a high performance HPC cluster in the cloud using EC2 UltraClusters of P4d Instances. You will setup the underlying networking for the cluster, deploy FSx for Lustre and P4d cluster, and delete your AWS resources.

Before launching an EC2 UltraCluster it is recommended first to launch a single P4d instance and get familiar with the instance type. Also note in which Availability Zone in your account and region you launched the P4d instance. You will need this information later in the tutorial.

What you'll accomplish

- Login to the AWS Management Console
- Create a private subnet with a NAT Gateway
- Create three security groups for access to the EC2 UltraCluster
- Launch a FSx for Lustre file system
- Launch a cluster of EC2 P4d instances with four EFA ENIs
- Launch a jump host
- Deprovision resources in the EC2 UltraCluster

Implementation

Step 1: Create a private subnet with a NAT Gateway

The EC2 UltraCluster will have multiple elastic network interfaces per instance. We will need to create the instances in a private subnet and route a NAT Gateway through a public subnet with the internet gateway (IGW) attached.

1. Login to the AWS Management Console

When you [click here](#), the AWS Management Console will open in a new browser window, so you can keep this step-by-step guide open. When the screen loads, enter your **user name** and **password** to get started. Then type **VPC** in the search bar and select **VPC** to open the console.

2. Create a private subnet

Create a subnet in your VPC with an available free CIDR range this CIDR range needs to be able to accommodate the number of instances you want to launch * 4.

Subnet settings
Specify the CIDR block and Availability Zone for the subnet.

Subnet 1 of 1

Subnet name
Create a tag with a key of 'Name' and a value that you specify.
p4d_cluster_subnet
The name can be up to 256 characters long.

Availability Zone [info](#)
Choose the zone in which your subnet will reside, or let Amazon choose one for you.
US East (N. Virginia) / us-east-1c

IPv4 CIDR block [info](#)
172.31.128.0/20

Tags - optional

Key	Value - optional	
Name	p4d_cluster_subnet	Remove

[Add new tag](#)
You can add 49 more tags.

[Remove](#)

[Add new subnet](#)

3. Create a NAT Gateway

Create a NAT Gateway by going to NAT Gateways in the side menu launching a gateway in public subnet in the VPC. This will take a few minutes to provision.

NAT gateway settings

Name - optional
Create a tag with a key of 'Name' and a value that you specify.
ec2ultracluster-nat
The name can be up to 256 characters long.

Subnet
Select a public subnet in which to create the NAT gateway.
subnet-ba7f22e5

Elastic IP allocation ID [info](#)
Assign an Elastic IP address to the NAT gateway.
eipalloc-00aa67ecd45a2b23

[Allocate Elastic IP](#)

4. Create a routing table

After provisioning is complete go to **route tables** and create a new route table selecting the VPC that your Gateway was created in. In **Routes** for the route table add a route for the destination **0.0.0.0/0** where the target is the NAT Gateway ID you created earlier.



Destination	Target	Status	Propagated
172.31.0.0/16	local	active	No
0.0.0.0/0	nw-0462742db0c0a0ab	active	No

5. Associate the route table with the subnet

Associate this route table with the private subnet you created earlier, right click on the route table **ID** and choose **Edit subnet associations**.

Step 2: Create security groups for access to the EC2 UltraCluster


We will create 2 security groups with different policies for access for:

- external SSH access
- EFA networking

1. Configure the security groups

In the EC2 Console navigate to the security groups and choose **Create security groups**.

- Choose the VPC used earlier to associate this security group with
- For EFA: For inbound rules add All traffic on all ports in scope of the security group that is being created.
- For EFA: For outbound rules add All traffic on all ports in scope of the security group being create



Type	Protocol	Port range	Source
All traffic	All	All	sg-0504670fb1c555723 (efa-access)

2. Verify the TCP port settings

For the new SSH security group ensure that **TCP port 22** is open inbound with outbound set to **0.0.0.0/0**.

Step 3: Launch a FSx for Lustre file system

As part of the EC2 UltraClusters you will need to launch a FSx for Lustre file system. You can use any process to launch the FSx for Lustre file system but it needs to be launched in the private subnet you created earlier.

1. Create a file system

In the FSx for Lustre console click on **Create file system**.

Choose Amazon FSx for Lustre and click **Next**.

2. Configure the file system

Complete the form with the following parameters:

- For **Deployment & storage type** choose **Scratch,SSD**
- For **Throughput per unit of storage** choose **200 MBs/TB**
- For **Storage Capacity** choose **2.4TiB**
- For **Virtual Private Cloud:** choose **VPC of the private subnet created earlier**
- For **VPC Security Groups** choose **Choose the EFA security group you created earlier**
- For **Subnet** choose. the private subnet you created earlier

Create file system

File system details

File system name - optional [Info](#)

 Maximum of 256 Unicode letters, whitespaces, and numbers, plus + - = _ . /

Deployment and storage type [Info](#)
 Select a deployment type and storage type to fit your workload requirements

☐ Persistent, SSD
☐ Persistent, HDD
☒ Scratch, SSD

☒ with SSD cache

Throughput per unit of storage [Info](#)
 Throughput (MB/s) per unit of storage (TiB)

☐ 50 MB/s/TiB
☐ 100 MB/s/TiB
☒ 200 MB/s/TiB (up to 1.3 GB/s/TiB burst)

Storage capacity [Info](#)
 TiB
 Supported sizes: 1.2 TiB or increments of 2.4 TiB

Throughput capacity [Info](#)
 Throughput capacity = Storage capacity (TiB) * 200 MB/s/TiB
 480 MB/s

Network & security

Virtual Private Cloud (VPC) [Info](#)
 Specify the VPC from which your file system is accessible.

VPC Security Groups [Info](#)
 Specify VPC Security Groups to associate with your file system's network interface.


☒ sg-050407dfb1c555725 (efa-access) [X](#)

The VPC Security Groups associated with your file system's network interfaces must allow inbound Lustre traffic (TCP ports 988, 1021-1023).

Subnet [Info](#)
 Specify the subnet in which your file system's network interface resides.

3. Choose an S3 bucket

Choose an S3 bucket for data ingestion. The dataset for this tutorial is the BERT dataset. If you don't have it, we can use synthetic benchmarks.



▼ Data Repository Import/Export - optional

☒ Import data from and export data to S3 [Info](#)

When you create your file system, your existing S3 objects will appear as file and directory listings. After you create your file system, how do you want to update it as the contents of your S3 bucket are updated?

☒ Update my file and directory listing as objects are added to my S3 bucket

☐ Update my file and directory listing as objects are added to or changed in my S3 bucket

☐ Do not update my file and directory listing when objects are added to or changed in my S3 bucket

Import bucket

The name of an existing S3 bucket

Import prefix - optional [Info](#)

The prefix containing the data to import

Export prefix [Info](#)

The prefix to which data is exported

☒ A unique prefix that FSx creates in your bucket

☐ The same prefix that you imported from (replace existing objects with updated ones)

☐ A prefix you specify

4. Verify the file system was created

Wait until the FSx cluster is in the **Available** state.

Note the dnsname and mountname of the cluster.

Step 4: Launch a cluster of EC2 P4d instances with four EFA ENIs

We can launch the compute layer of the EC2 UltraCluster. You can use the Deep Learning AMI v36 for support for P4d or create your own. You will need to install the FSx client drivers.

In the EC2 management console, select in the EC2 Dashboard to launch an instance.

- Launch an instance
 - a. Select the AMI with A100 support as well as have the FSx client driver installed.
 - b. For **Instance Type** choose **p4d.24xlarge**
 - c. For the instance details choose the number of instances you want in the count
 - d. Choose the VPC and private subnet created earlier.
 - e. Select a placement group created as a cluster.
 - f. For **network interfaces** add 3 more network interfaces with **Elastic Fabric Adapter** selected
 - g. Set the **NetworkCardIndex** for each EFA adapter to **0,1,2,3**.
 - h. Add any relevant tags in the next screen for the Security Group section choose the security groups created earlier for SSH and EFA access.

- i. Launch the instance and confirm they have 4 private IP addresses per node.

Network Interfaces ⓘ

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses	IPv6 IPs	Network Card Index	Elastic Fabric Adapter
eth0	New network interface	subnet-0b015d21	Auto-assign	Add IP	Add IP	0	☒
eth1	New network interface	subnet-0b015d21	Auto-assign	Add IP	Add IP	1	☒
eth2	New network interface	subnet-0b015d21	Auto-assign	Add IP	Add IP	2	☒
eth3	New network interface	subnet-0b015d21	Auto-assign	Add IP	Add IP	3	☒

Step 5: Launch a jump host

Since the cluster is a private subnet. We need to launch a jump host in the public subnet to be able to access the P4d instance in the EC2 UltraCluster.

1. In the EC2 Console launch an EC2 instance, for example t3a.xlarge, in a public subnet of the VPC.
2. Attach the security groups created earlier.
3. After the instance has launched, you can ssh into the instance and then ssh into one of the p4d.24xlarge nodes in the cluster.

(Optional) Delete resources in the EC2 UltraCluster

You can easily delete the EC2 P4d cluster from the EC2 console and the FSx for Lustre file system from the FSx console. In fact, it is a best practice to delete resources you are no longer using so you don't keep getting charged for them.

Congratulations!

You have just launched a P4d instance in the EC2 UltraCluster. With this cluster you can run large scale distributed deep learning workflows with the best practices for compute and storage.

EC2 UltraClusters is an optimized placement strategy for the EC2 P4d instances and FSx for Lustre file system. EC2 UltraClusters are supported in managed services such as Amazon Elastic Kubernetes Service (Amazon EKS). Follow [examples on Github](#) to launch an EC2 UltraCluster with containers using Amazon EKS.