

# AWS Requirements

This document is intended to explain the Requirements of an Amazon Web Services (AWS) Instance in order to run different Nextthink Appliances considering our Hardware Requirements.

## Introduction

AWS uses Elastic Compute Cloud (EC2) which allows users to manage and launch "Instances", which are their Virtual Machines in the Cloud. EC2 allows flexibility and scalability for your environments, which means you can modify and improve the Instance resources easily.

There are different Instance Types and categorized in five main groups:

- **General Purpose** (T2, M4, M3)
- **Compute Optimized** (C4, C3)
- **Memory Optimized** (X1, R4, R3)
- **Accelerated Computing Instances** (P2, G2, F1)
- **Storage Optimized** (I2, D2)

These Instances are oriented to different use cases, and they charge you for the usage depending on the balance each has. For example, Memory optimized Instances will charge you less money when using RAM in our Instance as this type is intended for in-memory Databases, which uses a lot of RAM all the time.

For more information and full details on the Instance Types, please visit: <https://aws.amazon.com/ec2/instance-types/>

## AWS Requirements

- An Amazon account configured for virtual machine (VM) creation from imported VHDs
- An import service role
- A local or on-premise CentOS machine with the Amazon client installed. To acquire the Amazon client, refer to the [Get started with the AWS CLI](#) documentation from Amazon
  - This can be achieved via the Nextthink Appliance
- An S3 bucket within your AWS account
- An existing AWS environment or familiarity with the Amazon network configuration
- A security group configured according to the Nextthink connectivity requirements
  - This is required for secure access to the machine instance

## Nextthink Requirements

The information below takes the Hardware Requirements reference table from our documentation, and contains the EC2 instance types that are closest to fit the On Premise Official Hardware Requirements.

Therefore, there might be some Instances that have more or slightly less resources than the actually needed by the Appliance. But they must be chosen since there is no other Instance type that suits for some specific Appliance setups.

### Portal Appliance:

Max devices	Max complexity	Memory	Data disk size	Details (90 days)	CPU cores	Minimal Amazon AWS requirements
5k	2000	12 GB	120 GB	120 GB	2	r7a.large with gp3 disks
10k	2000	13 GB	200 GB	240 GB	4	r7a.large with gp3 disks
20k	4000	17 GB	400 GB	440 GB	4	r7a.large with gp3 disks
50k	12,000	23 GB	600 GB	900 GB	6	r5d.2xlarge with gp3 disks
100k	40,000	41 GB	1.2 TB	1.4 TB	6	r5d.2xlarge with gp3 disks
150k	60,000	59 GB	2 TB	2 TB	8	r5d.2xlarge with gp3 disks

### Engine Appliance:

Max Events	Max devices / with Web & Cloud	Max entities	Minimal Amazon AWS requirements
20M	500 / 500	20	c6a.large with gp3 disks
50M	3k / 2k	100	c6a.2xlarge with gp3 disks
50M	5k / 3-4k	250	c6a.2xlarge with gp3 disks
100M	10k / 6-8k	100	c6a.4xlarge with provisioned iops SSD io1
100M	10k / 6-8k	500	c6a.4xlarge with provisioned iops SSD io1

200M	10k / 8k	100	c6a.4xlarge with provisioned iops SSD io1
200M	10k / 8k	500	c6a.4xlarge with provisioned iops SSD io1
>200M	ask	ask	ask

**Special considerations:**

- Many AWS EC2 Instance types far exceed the Minimum Hardware requirements, and some slightly lack CPU requirements. This is due to the way they are predefined. Customer may evaluate that is the best for them, always considering that choosing an instance lacking number of CPU cores might present slight slowness or small performance issues.
- Ideally, we recommend to use the instance types that far exceed our hardware Requirements, to avoid having performance issues or experiencing slowness when running investigations.
- Regarding the IOPS / Disk throughput, the bigger is the hard disk, the higher will be the IOPS. As specified on [this page](#), Amazon defines an iops balance that is filled at 100 IOPS for a 33.33 GB disk and increase by 3 IOPS / GB for general purpose SSD.  
For example, a disk of 500 GB will acquire  $(500-33.33)*3 + 100 = 1500$  IO per seconds and be able to burst up to 3000 IOPS during IO intensives processes such as the engine database backup.