



Managed Kubernetes




Contents

Click to navigate to each section

Page

What is Managed Kubernetes?	3
Case Study - Machine Learning	4
Resource Comparison	5
Fast, Flexible and Reliable Cloud	6
Contact	7



What is Managed Kubernetes?



UpCloud's Managed Kubernetes is a new, fully-serviced system to streamline all processes from development to production. Easily control your containerised applications in a fast, secure, and flexible environment without worrying about downtime.

Kubernetes is an open-source **container orchestration system** that allows users to automate the deployment, scaling and management of their cloud applications.

The standardised system allows users to serve apps across multiple hosts, making it easier to deploy and manage systems across the cloud.

It also ensures peace of mind, with features like self-healing, automatic scaling, and service discovery providing security and reliability.



How does it work?

Kubernetes consists of two layers: the control plane and the data plane. The control plane acts as the 'brain' of the operation, while the data plane does the legwork.

Managing the control layer is often seen as the 'hard' part of Kubernetes systems, which is why we provide a managed service.

Meanwhile, the data plane consists of clusters of 'worker nodes' which are managed by the control plane. Users typically pay by the hour for the cloud resources used by these worker nodes.

UpCloud's Managed Kubernetes solution manages the specialised control plane, while giving users full control of the data plane - allowing them to optimise their data usage without worrying about manually deploying all of the actions contained in the control plane.

Case Study - Machine Learning



Save on cloud resources by serving your pre-trained machine-learning models on Managed Kubernetes

Running a machine-learning system in the cloud is no easy task.

Just setting up the system architecture is a complex process in itself, not to mention the resource-guzzling processing power needed to allow it to run smoothly later on.

And then you need to ensure it stays active 24/7, without downtime or interference.

Training and deploying a machine-learning model is also a cost-heavy process, too.

From the initial investment in resource and opportunity cost to expensive high-end GPUs and storing large datasets in object storage, this is one long process with too many points of failure.

And ensuring the process is as fault-tolerant as possible is going to require meticulous work on an ongoing basis. Because now you've invested the time and energy into building it, you can't make deployment an afterthought.

No one sees your hard work.

Things need to 'just work' for the end user. And this is where a Managed Kubernetes system is vital. By deploying and serving the model, Managed Kubernetes ensures the final step in the process is accessible to everyone.

We know that the most feature-rich environment for training the model might not be the most cost-effective one for actually deploying it to the end user. By uncoupling the two, you can choose the best and most cost-effective approach for each part of the process.

You've already done the hard part by building the model. We think deploying it should be much easier.

Find out why, with UpCloud, it is.

Resource Comparison

Competitor cloud resources



AWS Sagemaker for real-time inference (€1.98/h pricing)*

Development: 1x ml.c7g.2xlarge (8CPU-16GB) @ €0.33/h
 Staging: 1x ml.c7g.2xlarge (8CPU-16GB) @ €0.33/h
 Production: 4x ml.c7g.2xlarge (8CPU-16GB) @ 1.32/h €0.33/h/node



Google Cloud Vertex AI (€4.10/h pricing)*

Development: 1x n1-highcpu-16 @ €0.69/h
 Staging: 1x n1-highcpu-16 @ €0.68/h
 Production: 4x n1-highcpu-16 @ €2.73/h (€0.68/h per node)

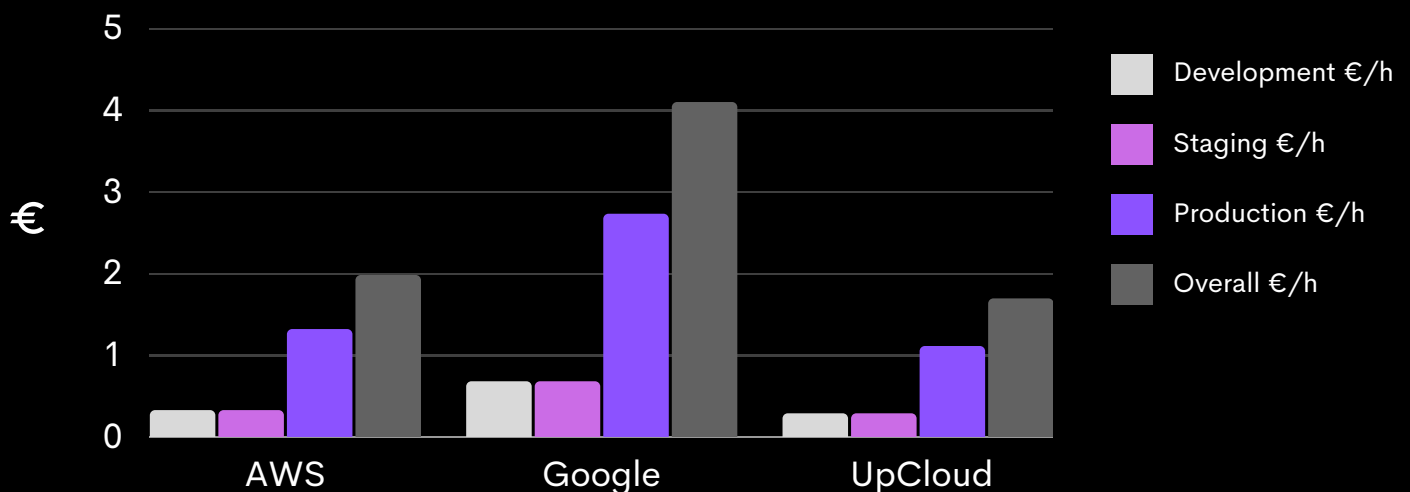


UpCloud Managed resources (€1.69/h pricing)

Development: Kubernetes as low as €0.29/h
 1x 8CPU-16GB (high-cpu) worker nodes
 One development LB
 Staging: Kubernetes as low as €0.29/h
 1x 8CPU-16GB (high-cpu) worker nodes
 One development LB
 Production: Kubernetes as low as €1.11/h
 4x 8CPU-16GB (high-cpu) worker nodes
 One production LB

*Note that these prices do not include traffic costs and assume that model training is already done and available in object storage. Additionally, the hourly pricing for AWS Sagemaker and Google Cloud Vertex AI has been converted from USD to Euro at a rate of 1.05 USD to 1 EURO.

Resource Comparison



Fast, flexible and reliable cloud



Customers worldwide trust our cloud platform to run their business-critical applications and transform their business.



Owned cloud infrastructure

Data centres across Europe, Asia, Australia and the USA



Simply faster

Fastest cloud servers in the market with a proven track record



24/7/365 support

Every day of the year from our skilled teams



50x confidence

Business security with 50x compensation for service outages

**For more information on
UpCloud's Managed
Kubernetes service, or
to speak to us about a
partnership – visit
upcloud.com/contact**

