

Are You a Certified Kubernetes Application Developer?

TechRounder PDF Edition

Live article: <https://www.techrounder.com/technology/are-you-a-certified-kubernetes-application-developer/>

By Vipin PG | Published September 25, 2020 | Updated March 8, 2026 | Format: Article | 7 min read

In brief

Many IT organizations adopting Kubernetes think that Kubernetes is the initial step towards building scalable, innovative applications. However, it is not the complete truth.

Many IT organizations adopting Kubernetes think that Kubernetes is the initial step towards building scalable, innovative applications. However, it is not the complete truth. Kubernetes is one of the most popular container orchestration systems available in the market.

In this blog, you will come across several aspects of Kubernetes based on GPU and learn about the suitability of Kubernetes for your organization. However, before knowing the Kubernetes, let us first understand what Kubernetes is.

What is Kubernetes?

Kubernetes is the system to deploy applications. It effectively utilizes the containerized infrastructure that powers the apps. As a result, Kubernetes can save your organization money because it requires less workforce to manage IT; it makes apps more flexible and performant.

You can operate Kubernetes on-premises or within the public cloud. AWS, Azure, and GCP offer Kubernetes solutions to assist clients in running K8 applications fast and efficiently. If you need portable apps, then Kubernetes can also do it. The benefit of mobile apps is that you can effortlessly move them between various clouds and internal environments.

In short, Kubernetes is the new Linux Operating System of the cloud. Google had launched it, and now it is part of CNCF. It has active engagement and contributes to lots of small and large organizations.

Do enterprises adopt Kubernetes?

Some numbers show that Kubernetes is getting adopted by multiple enterprises. One report states that K8 is growing with a rapid rate on-premise as well as a cloud-based environment. Moreover, over 33% of the AWS cloud enterprises use Kubernetes as their key orchestration solution.

What are the benefits of using Kubernetes?

Here are five necessary benefits of Kubernetes for businesses. In addition, we have gathered some real-world examples to make you understand the concept efficiently.

Improved efficiency of app development/deployment

Kubernetes offers a "Microservice" method of creating apps. Therefore, you can divide your development team into moderate or smaller ones to focus more on a single, smaller microservice. Your teams would be more productive due to the particular focused function of each section.

APIs among these microservices reduce the cross-team communication required to deploy and build. So now, you can scale various moderate or smaller groups of specialists who each help a flotilla of hundreds of machines.

Moreover, it enables your IT teams to manage large apps across lots of containers more effectively by handling many of the nitty-gritty details of maintaining container-based apps. For instance, Kubernetes manages service discovery, assists containers in talking to each other, and manages access storage from multiple providers such as AWS and Microsoft Azure.

Airbnb as an example

It is fantastic to observe the transition of Airbnb from a monolithic to a microservices architecture. They are required to scale continuous delivery horizontally. The aim was to make continuous delivery available to the company's more than 1000 engineers to include more new services.

Airbnb used Kubernetes to support more than 1,000 engineers, concurrently configuring and deploying more than 250 essential services to Kubernetes. The net conclusion is that Airbnb can do more than 500 deploys/day on average.

Helps to cut infrastructure costs

Kubernetes' extraordinary capabilities can help your business reduce infrastructure costs if you run on a massive scale. It builds container-based architecture feasible by packing together applications optimally using your hardware investments and cloud.

When Kubernetes was not in the market, managers usually over-provisioned their infrastructure to handle sudden spikes conservatively. The other reason was to be over-provisioned simply because it was challenging and time-consuming to scale containerized applications manually.

It brilliantly programs and tightly ties containers, considering the available resources.

Besides, Kubernetes automatically scales your application to satisfy business requirements, saving human resources to concentrate on other fertile duties.

There are plenty of examples of clients who have experienced dramatic advancements in cost optimization using K8s.

Spotify for an instance

Spotify is an early K8s user. As outlined in this note, it has significant cost-saving values by using K8s. Getting benefits of K8s, Spotify has recorded 2-3x CPU utilization by applying the orchestration capacities of K8s. The final result converted into better IT spending optimization.

Enhanced Scalability and Availability

Nowadays, an application requires more than one factor to taste success in the market. The victory depends on features, and scalability also plays a significant role in the process. If your application is not scaled well, it would be highly non-performant at the best scale and unavailable in the worst case.

Kubernetes is an orchestration system, a robust management system that "auto-magically" scales and increases app performance. So, for example, you have a CPU-intensive service and a powerful user load that adjusts based on the business (for instance, an event ticketing app that will see moving users and loads earlier to the event and moderate usage at other times).

You require a solution that can upscale the app's infrastructure. So that new machines are automatically generated as the load boosts (more users are ordering tickets) and scale down when the load sinks.

Kubernetes allows that skill by scaling up the app as the CPU usage goes up a specified threshold - for this case, 90% on the current machine. Moreover, when the load decreases, the system can scale back the application, optimizing infrastructure utilization.

Auto-scaling Kubernetes has not been defined as just infrastructure metrics; any metric-resource utilization metrics. Also, custom metrics can be used to trigger the scaling process.

Hybrid and Multi-cloud Flexibility

Kubernetes and containers help you recognize the promise of hybrid and multi-cloud flexibility, which is the most significant advantage. Nowadays, enterprises are operating multi-cloud environments. Not just it, it seems that they will continue doing that in the future too.

Kubernetes system helps them simplify to run any app on any public cloud service or any mixture of public and private clouds. It allows you to place the right workloads on the correct cloud and helps you prevent vendor lock-in.

It will also help you get the best fit, using the top features and holding the advantage to move when it makes sense, all help you realize more Return Over Investment (short and longer-term) from your IT investments.

If you require more informative data to verify the multi-cloud and Kubernetes made-in-heaven story, this report will help you. The Sumo Logic Continuous Intelligence Report states a very impressive upward leaning on K8 adoption based on the number of cloud platforms businesses use and 86% of consumers on all three using controlled or native Kubernetes solutions.

AWS should not be worried. However, it could be an initial indication of a level playing field for Azure and GCP. The reason why is apps used on K8s can be quickly ported across conditions.

Seamless migration to the cloud

It doesn't matter if you are rehosting, re-platforming, or refactoring. Kubernetes has everything you require. Kubernetes presents a more ideal and prescribed way to port your application from on-premise to cloud environments since K8s operate consistently across all domains, on-premise, and clouds such as AWS, Azure, and GCP.

Instead of dealing with all the cloud environment's variations and complexities, enterprises can develop a more guided path:

Move apps to Kubernetes on-premise

At this place, you concentrate more on re-platforming the applications to containers and getting them under Kubernetes orchestration.

Migrate to a cloud-based Kubernetes model

You will get plenty of choices here. You can run natively or choose an organized Kubernetes environment from the cloud vendor.

Now, your app is in the cloud, and you can begin to optimize your app to the cloud environment and its service. When to use Google Compute Engine managed Instance Autoscale vs. Kubernetes (GKE) Engine.

Google Kubernetes Engine (GKE) is the perfect option for managed Kubernetes services. It will help your enterprise if you are finding for a container orchestration platform that provides improved scalability and configuration flexibility.

GKE offers total control over every factor of container orchestration, from networking to storage to how you set up observability and supporting stateful application use cases. However, if your application does not require the upper level of cluster configuration and analyzing, then a fully managed Google Compute Engine could be the perfect system for your enterprise.

Fully Managed Google Computing Engine is a decent serverless platform. It is suitable for stateless containerized microservices that don't need Kubernetes features.

Why use Google Compute Engine?

The managed serverless compute engine offers tons of features and advantages.

The simple process of deploying of microservices

You can deploy a containerized microservice with just one command. You don't need to seek any extra service-oriented configuration.

Easy and Unified developer experience

You will notice that each microservice has been implemented as a Docker image, GCE's deployment unit.

Support for code written in any language

Google Compute Engine is container-based. So, you can quickly write code in the desired language, using any binary and framework.

Scalable Serverless Execution

A microservice deployed into controlled GCE scales automatically based on the number of incoming applications, outwardly having to configure a full-fledged Kubernetes cluster. Managed Google Compute Engine estimates zero if there are no requests, i.e., uses no resources.

Google Compute Engine has two configurations: a managed Google Cloud service and a Google Compute Engine for Anthos. If you're currently operating Anthos, GPE for Anthos can deploy containers into your cluster, enabling access to custom machine models, further networking support, and Kubernetes to improve your GPE services. In addition, you can create and manage both managed GCE services and GKE clusters from the console and the command line.

The Perfect Tool for the Job

Both tools, Cloud Run and GKE, are potent for different user requirements. However, you need to know your functional and non-functional service conditions, such as the ability to scale to zero or the strength to control detailed configuration, before choosing one over the other.

Moreover, you can use both at the same time. For example, an enterprise could have challenging microservice-based apps that need excellent configuration functions of GKE. Some that do not but still need to leverage Google Compute Engine's easy-to-use and scalable features.

Final words

Kubernetes is the best system for IT professionals, CIOs, and CxOs. There are plenty of benefits of deploying applications on it. However, to get better output from Kubernetes, businesses should consult with solution providers who can secure and monitor Kubernetes applications.

References

1. private-ai.com - 2022 / 05 - <https://www.private-ai.com/2022/05/31/how-to-autoscale-kubernetes-pods-based-on-gpu/>