

# **Harnessing the Complexity of Kubernetes to Benefit your Business**

# Table of Contents

Your go-to for agility	3
Your business, your benefits	3
What makes it special?	5
Deploying apps the containerized way	5
The complexity of it all	7
Benefiting from AI and ML	10
K8s is great	11



# Your go-to for agility

Kubernetes (aka k8s by many). If you ask anyone associated with IT if they've heard of this technology, you'll be greeted with a resounding and emphatic, "Yes." Not only do IT admins and businesses know about Kubernetes, they find it one of the most exciting technologies to come around in a very long time. However, if you ask those same admins and businesses if they've implemented Kubernetes and the answers will vary.

But why?

Why hasn't every business on the planet plunged into the inviting waters of Kubernetes? Once again, the answers will vary. And, for the most part, the reactions won't surprise you.

- It's too time consuming.
- It's too expensive to implement.
- It's too complicated.
- We don't have the resources.
- Our staff doesn't understand the technology.

Before your company makes that final decision on whether or not to implement Kubernetes, let's take a look at what this technology has to offer.

# Your business, your benefits

If you haven't heard of Kubernetes, or if you have and are still a bit confused as to what this relatively new technology is, let's unpack it in a way that anyone can understand.

First, you must know what a container is. In terms of technology, a container is an application that has been packaged together with

all of the necessary code and dependencies such that it can be deployed on almost any platform (so long as said platform includes a container runtime engine). Containers make deploying apps and services incredibly easy. And because they are portable, a developer can create a containerized application with a development environment (such as their desktop computer) and then move it to a production environment, where it should run as expected.

Kubernetes is much more than that. What Kubernetes does is make it possible for developers and admins to manage numerous aspects of container deployment. With Kubernetes it is possible to orchestrate application deployment, scaling, and management—all from a single point of entry.

Think of it this way. Imagine you have a massive collection of Legos—different colors, shapes, and sizes. Alone you'll eventually be able to take those individual bricks and put them together such that they create a very specific sculpture. However, without guidance that could take some time and you might not get the exact end results you were hoping for. Kubernetes is like the photo and the instructions you get with those Legos—it makes it possible for you to piece those bricks together, in exactly the right configuration, and do it in a much more efficient and reliable way.

However, the problem with that analogy is that it makes Kubernetes seem really easy. Up to a point it is. That Lego photo might give you the basic idea of what you're building, but there are intricacies and inner workings it might not show.

Plus, Kubernetes isn't about building Death Stars with Legos. It's much more.

# What makes it special?

The biggest promise that Kubernetes makes is that it makes it possible for your company to get the absolute most power, flexibility, and reliability from your container deployments. With this particular management platform, you can develop a very specific application (or set of applications), deploy them to a cluster, and then set them to scale as needed so your business benefits from a level of failover and reliability you've probably never experienced. And, with the right tools, much of this can be done automatically.

That's a big deal, especially in today's ultra competitive world of business, where the slightest hiccup in your system can cause your company to fall behind. And with Kubernetes, there are a lot of "moving pieces" that can go wrong.

## Deploying apps the containerized way

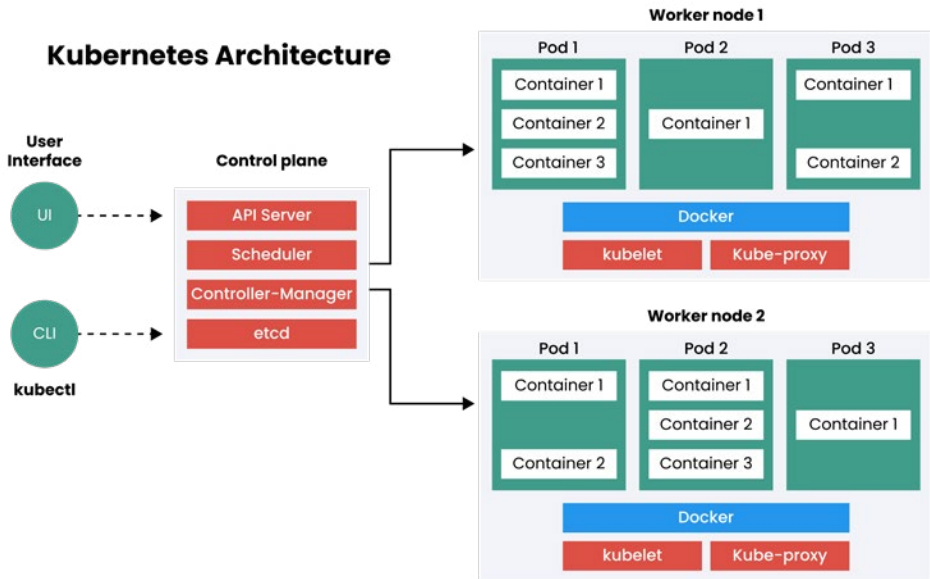
The building blocks of Kubernetes are numerous and can be complicated. Here's a short list of some of the components that come together to help form a Kubernetes deployment:

- **Cluster** - the cluster is a collection of Kubernetes servers (be they bare metal or virtual) that work together such that if one node fails, the deployed application will not go down.
- **Kubernetes Controller** (once called the "Kubernetes Master" but that nomenclature is thankfully changing) - the Kubernetes Controller is the one server in the cluster that is tasked with managing all nodes and various core controllers.
- **Worker Node** - this is a single server (either bare metal or virtual) that serves the cluster. All containers are deployed to the nodes (with the controller handling the orchestration).

- **Pod** - pods are the smallest deployable units that can be managed by the controller. A pod consists of one or more containers that share the same IP address and port.
- **Namespaces** - make it possible to divide resources between users without running into name collision. However, while namespaces do provide some segregation, they are not a complete tenancy solution.
- **API Server** - exposes APIs to the cluster.
- **Scheduler** - watches for newly created pods and selects a node for them to run on.
- **ETCD** - stores the configuration information used by Kubernetes to track information about all of the objects in play in a cluster.
- **Docker/CRI-O** - the container engine that makes it possible to deploy containers.
- **Kubelet** - the agent that runs on each node in the cluster.
- **Kubernetes Proxy Service** - a service that runs on each node and makes it possible for containers within the cluster to be reached from outside (from LAN or WAN).

And that's just the short list. There are quite a few more pieces that make up this rather complicated puzzle. And, to make matters even more complicated, many of these pieces need to be configured manually. And within each piece, there are sub-pieces that must be configured.

So within a Kubernetes cluster, there are seemingly countless optimizations to be made, each of which will affect your cluster in different ways.



## The complexity of it all

Now that you understand the complexity of Kubernetes, let's make it even more so. Pretend you already have your Kubernetes cluster up and running (that is, after all, the easy part). Now you must develop, build, and deploy an application and/or a service to the cluster. How do you do that?

You craft a YAML file such that it declares every aspect of your app or service and then, using a command like `helm` (package manager for Kubernetes), you deploy the service to the cluster. What is a YAML file? A YAML file is a configuration file, written in a human-readable, data-serialization language.

Once you've pieced together your YAML file, what do you do? Within the Kubernetes controller, you'll find a very powerful command, called *kubectl*. That command enables you to deploy to your cluster. Say, for instance, you have a YAML file named `deploy.yaml`. In that file you

have everything put together to deploy your application. How do you deploy it? Simple. From the command line you'd issue something like:

```
kubectl create -f deploy.yaml
```

It really is that simple. But what if you find you need to change a configuration or two within your YAML file? The `kubectl` command has you covered. After you've made your edits, apply them with the command:

```
kubectl apply -f deploy.yaml
```

Kubernetes also has another very useful tool that can help make your life a bit easier. This tool is called Helm, which is a package manager for Kubernetes. Helm streamlines the installation and management of Kubernetes applications.

Once you have Helm installed, you can then work with charts, which are a collection of files that describe a related set of Kubernetes resources. With these charts you can deploy simple or very complex apps and services to your Kubernetes clusters. For example, with Helm you could quickly download the Prometheus chart with the command:

```
helm install stable/prometheus --generate-name
```

That command would download the chart for Prometheus and save it to your local storage. You could then modify the chart to perfectly fit your needs.

Why reinvent the wheel? With the help of Helm charts you can deploy incredibly complicated services and apps to your Kubernetes cluster with minimal effort.



These YAML files become even more complicated, because there are so many options that can be defined. Such a configuration could look like this:

```
## exporter resource limits & requests
## Ref: https://kubernetes.io/docs/user-guide/compute-resources/
##
resources: {}
  limits:
    cpu: 200m
    memory: 50Mi
  requests:
    cpu: 100m
    memory: 30Mi
## PodDisruptionBudget settings
## ref: https://kubernetes.io/docs/concepts/workloads/pods/disruptions/
##
podDisruptionBudget:
  enabled: false
  maxUnavailable: 1

## gateway resource requests and limits
## Ref: http://kubernetes.io/docs/user-guide/compute-resources/
##
resources: {}
  limits:
    cpu: 10m
    memory: 32Mi
  requests:
    cpu: 10m
    memory: 32Mi

## Use a StatefulSet if replicaCount needs to be greater than 1 (see below)
##
replicaCount: 1
```

The above example is very basic, just a few select lines from a 1600 line example of a popular solution. Now imagine you have a very complex Kubernetes configuration file, that includes a number of application definitions, each of which defines resource usages specifications. Such a YAML could wind up with hundreds of definitions. Crafting such a file is challenging, **but not nearly as challenging as optimizing that configuration** to not make your application or service perform only best suit your needs, but to make sure that deployment is cost effective. This is especially important when your Kubernetes deployment is on public cloud, where you are charged based on the resources you use.

# Benefiting from AI and ML

Consider this: You have crafted an incredibly complex YAML file for the deployment of a crucial application to a Kubernetes cluster running on AWS. However, the application isn't running efficiently or effectively. You open the YAML file and your first inclination might be to increase the RAM and CPU resources that are available to the application. You re-deploy only to find the application still isn't running properly. What do you do? Wash, rinse, repeat. In the end, you could wind up with a YAML file that hands over a large amount of resources to your applications. That equates to an exponentially more costly deployment.

When your Kubernetes application configurations (such as startup parameters, environmental variables, Java JVM tuning) grow exponentially in size, the idea of optimizing them becomes challenging. With hundreds upon hundreds of configuration options, each of which you can "dial in" to help improve your deployment, it can be almost impossible to know which option to change.

You can increase the CPU for the database component, only to find it makes little to no difference. Go back and up the memory for the web server. Nothing. Change the CPU for the database back to its original value and up the memory option. Still no change. Increase the CPU for the web server and you might see a slight uptick in performance. However, what you might not see is that all of a sudden you've exceeded your resource limits your provider offers and will be billed for the overages.

Back to the drawing board.

Had this been a simple Kubernetes YAML file (say with only a limited number of components and few options), this back and forth optimization might not take so long. But when those YAML entries number in the hundreds, you're looking at an incredibly time

consuming (and frustrating) endeavor. That's how complicated a Kubernetes deployment can become.

But what if you could use Artificial Intelligence and Machine Learning to test your configurations to perfectly tune your application such that it will not only run with a high level of efficiency and a low cost of delivery? If that were possible, why would you trust that optimization to first-generation management tools that are both imprecise and manual? Such tools often miss opportunities for huge performance gains and cost reductions. Think of it this way: Your deployed application is like a living, breathing animal and each component of that application is critical to the life of that animal. What is the long-term impact of that animal not getting enough water, food, or sleep, or carrying around too much weight or having a buildup of some harmful chemical in its system. Diagnosing each one of those components individually can be pretty simple. But because they work in harmony, the complexity rises exponentially.

Believe it or not, such a service is available. StormForge (with their StormForge Ops tool) uses Machine Learning to optimize your Kubernetes configurations in ways the human mind is simply not capable of doing. In a matter of hours, StormForge Ops can help you find the perfect configuration that will have your deployment humming and your bottom line smiling. Not only is StormForge delivering a game-changing technology, it's led by a strong team, with deep experience in building world class organizations and ensuring customer and partner success.

## **K8s is great**

There are a number of ways to make Kubernetes a part of your business. How you do so will depend on the scale at which you plan to deploy. If you're a small company (or a developer) looking to kick the tires of this technology, you can always deploy a Kubernetes

cluster to your data center. This would require a minimum of 3 servers (be they bare metal or virtual) and could serve as a solid development environment to create the apps and services you are looking to deploy.

If, however, your company exists within the realm of enterprise, chances are pretty good you'll need the ability to seriously scale to meet the ever-growing demand your business will surely experience. If that sounds like you, your best bet would be to sign up for the likes of Amazon Web Services, Azure, or Google Cloud. These services make it easy to deploy a Kubernetes Cluster and scale it out to meet (or exceed) your needs.

Just remember, when you're expanding your container deployments to scale, you're going to want to ensure those configurations are optimized for performance and cost effectiveness. Don't just look at it from the myopic view of performance, as there are other considerations that can help make that containerized application serve you better.

**[Sign up for the StormForge platform](#)** to automatically identify your optimal app configuration, free.

