



The Good, the Bad and Kubernetes Ugly

 **CONTAINER
JOURNAL**

In the Kubernetes Beginning...



Kubernetes traces its lineage back to an internal Google project known as Borg, on which multiple applications are centrally managed. When Kubernetes was launched as an open source project, the assumption was that IT teams would run a small number of large clusters that would be shared by lots of clusters.

In reality, the number of Kubernetes clusters being deployed by enterprise IT teams is multiplying. Application teams in many organizations want to keep control of their own infrastructure platforms. As a result, fleets of clusters are now being deployed from the network edge to the cloud.

How can enterprise IT organizations efficiently manage all those Kubernetes clusters without compromising the independence of application teams?

The Kubernetes Good



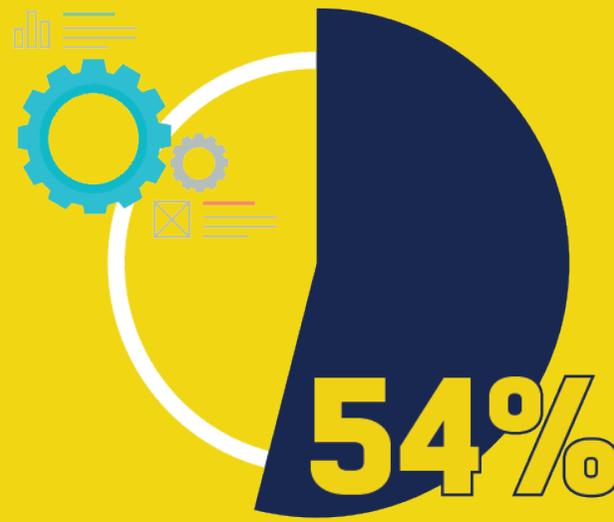
Kubernetes is a central part of the pandemic-accelerated digital transformation strategy for 77% of organizations.¹

And 68% have increased Kubernetes use since the pandemic started.

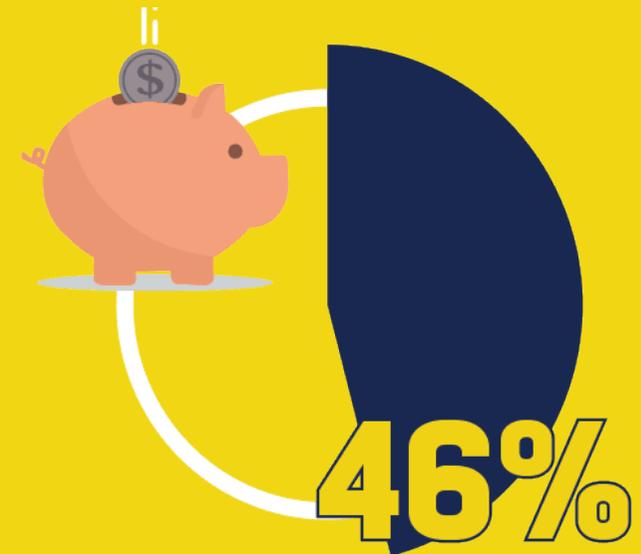
Top reasons for Kubernetes adoption include:



Accelerating deployment frequency



Increasing Automation



Reducing IT costs²

The Kubernetes Bad



The top Kubernetes pain points continue to be:



Setup and configuration



Maintenance



Too many platforms³



The Kubernetes Ugly

Kubernetes is the most powerful IT platform deployed in an enterprise – but it's also the most complicated.

Kubernetes clusters are challenging for IT team to manage, and without tools that abstract away the underlying Kubernetes complexity, developer productivity often declines after transitioning to a Kubernetes environment.



Deployment Debate

Most Kubernetes instances were deployed on top of existing virtual machine environments. But more are deploying Kubernetes on bare-metal servers and then encapsulating legacy applications in open source virtual machines using containers that run on Kubernetes. That approach, some say, greatly reduces the total cost of IT.



Convergence Confusion

Most Kubernetes clusters today run stateless applications, with data typically stored on an external storage platform. But some IT teams are deploying stateful applications directly onto Kubernetes clusters to enable IT teams to unify the compute and storage management⁴. But managing storage adds another layer of complexity.



The software-defined cloud-native storage market is expected to **grow by \$42.79 billion** from 2020–2024.



Kubernetes Insecurity

As is the case with most emerging IT platforms, security is often overlooked in the name of trying to make a platform accessible. Kubernetes is no exception to that rule. Vulnerabilities continue to be regularly uncovered by researchers.⁶



Automation Required

There are not enough IT professionals with the skills required to manage distributed Kubernetes clusters at scale. IT organizations need to invest in tools that enable the average IT administrator to manage Kubernetes clusters.

Sponsored by:



www.fairwinds.com



www.epsagon.com



www.redhat.com



www.codefresh.io

Organizations that have deployed Kubernetes spend **16.5 hours a week** on low-level DevOps tasks.⁷

Thank you for reading

The Good, the Bad and Kubernetes Ugly