


# The future of infrastructure will be containerized

Tech companies and startups are rapidly adopting managed container platforms that let them spend fewer engineering resources on maintaining infrastructure and more on road map priorities that drive business outcomes like growth, competitive advantage, and increased profitability.



---

**How you can reduce infrastructure management to get to market faster**

Page 02

**Why consider managed services based on open standards**

Page 05

**How you can get the biggest bang for your migration buck**

Page 09

## Chapter 1

# How you can reduce infrastructure management to get to market faster

While most tech companies and startups today run in the cloud, many have yet to realize all the benefits of doing so. If you're in the cloud but not on [Kubernetes](#), you're probably leveraging proprietary solutions while building and maintaining your own supplemental custom tooling. You're also leaving a lot on the table in terms of efficiency, running your own workloads on underutilized [virtual machines](#) (VMs), and potentially locking yourself in.

How many tools are you using today to manage and patch your VMs? How do you upgrade your applications? And what's your VM utilization like? What you have right now might not be all that efficient. Things are breaking (outages, scalability issues, etc.) due to weaknesses in your VM architecture, or costs are spiraling out of control, or your infrastructure isn't set up to support lots of things your business needs to do:

- Refactoring/re-architecting an MVP into a scalable solution
- Expanding into additional cloud providers to meet regulatory or customer expectations
- Expanding geographically to reduce latency and provide better experiences to a widespread customer base
- Improving your end-to-end security posture
- Improving the customer experience (e.g. service availability)

**Legacy and technical debt can slow you down.** That's why we're seeing this massive shift to Kubernetes. A modern architecture consisting of managed containers gives you access to proven patterns for running reliable and secure infrastructure. You can speed up your time to market and increase your productivity without sacrificing stability and security - with the added benefit of helping you attract the best technical talent to work on innovation.

You should also be concerned about locking yourself out from the Kubernetes community and its surrounding ecosystem, whose stable innovation defines today's industry standards and best practices. Ultimately -- and this is even more important given current hiring challenges -- you have to decide where you want your engineers to spend their time: maintaining infrastructure and building and maintaining custom tooling, or ticking off your priority list to drive your business forward. What you have today might be working, but your roadmap probably includes items you'd rather it didn't, like repaying technical debt and filling in platform gaps around:

- End-to-end encryption
- Observability (logs, metrics, auto-logging)
- Policy management and enforcement
- High availability and automatic fail-over
- Cost reduction

Kubernetes is open source and platform-agnostic, and offers all the common tooling out of the box to secure and speed up each stage of the build-and-deploy life cycle. It's the sum of all the bash scripts and best practices that most system administrators cobble together over time, presented as one system behind a declarative set of APIs. Everything is automated, the details are hidden, and it's ready to use. Kubernetes can eliminate the vast majority of infrastructure-as-code while shifting your platform to infrastructure-as-data. You don't have to write or maintain code; you tell Kubernetes what you want, not what to do. This is a colossal timesaver when it comes to management overhead.

## Containers are the best way to leverage the compute continuum.



If you want to run traditional workloads, you can do so on a modern platform with Kubernetes by separating apps from VMs and putting them in containers. Adopting container images to package your software will make upgrades of your VMs easier. You can now decouple the lifecycle management of the VM and the lifecycle management of the application, simplifying the VM by removing everything like Ruby, Python, and Java. And by moving it to where it belongs – next to the developer’s application – you can control it all in one spot and keep your machines bare metal.



Managed computing platforms turn cloud services into platforms-as-a-service, giving you the power and flexibility of containers and the convenience of serverless. There’s no server, no cluster configuration, and no maintenance, which means you can see tremendous benefit while retaining control.



For workloads that don’t require much control over cluster configuration, you can let your services provision and manage the cluster’s underlying infrastructure, including nodes and node pools, while you only pay for the workload, not the cluster. In this way, you can eliminate cluster administration while optimizing security and saving potentially substantial amounts of money.



For more cloud-native applications, serverless does the same thing with less work, eliminating underlying infrastructure and serving as the end-to-end host for your applications, data, and even analysis. A serverless platform will let you start running containers to a fully managed environment with minimal complexity and security, performance, scalability, and best practices baked in.

**Let’s take a look at the potential implications for your business, and how you can keep your organization ahead of the efficiency curve.**

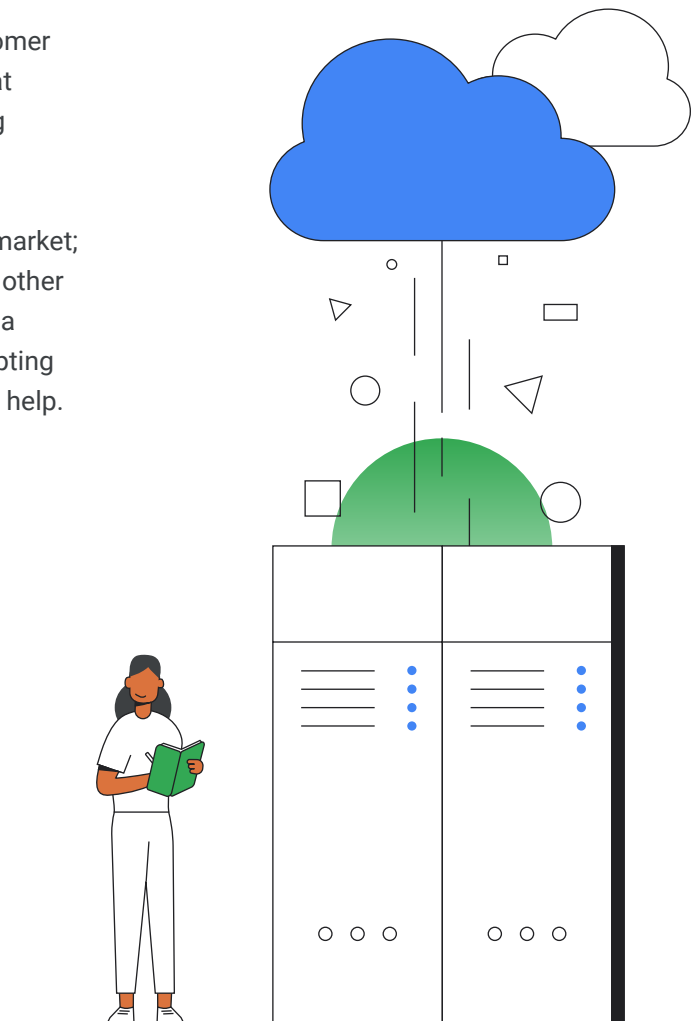
Chapter 2

# Why consider managed services based on open standards

## If you need multicloud, use technology that makes it easier

Some businesses find it necessary to operate in multiple clouds. Data gravity is real, and if you have a global customer base, you'll inevitably find yourself serving customers that want to minimize latency and networking fees by keeping their compute close to where their data actually lives.

In such cases, multicloud can expand your addressable market; you'll wind up supporting managed services and data on other providers anyway. Other companies value multi-cloud as a risk mitigation strategy. In either situation, the key is adopting multicloud correctly, and industry-standard solutions can help.



## 01

First: **make sure you can re-leverage as many of your workflows for getting things done as possible.** By workflows, we mean enabling task automation to create the dataflow architecture between the database and compute. This is where open source is important; if you select a database that isn't open, you'll have a hard time implementing the dataflow architecture and automation. You'll do better to go with something like the Postgres protocol (e.g. [Cloud Spanner](#)) or a managed Postgres database (e.g. [Cloud SQL](#)) on any cloud provider.

## 02

Second: on the compute side, Kubernetes saves substantial time on deployments and automation, so when you're picking a set of technologies, **be sure they work across the boundaries established by your provider.** Boundaries could be different regions or zones, different projects or accounts, and even on-prem or cloud. Don't waste time building and maintaining separate infrastructures for each cloud provider (e.g. one on AWS, one on Google and one on Azure); your engineers will drown trying to keep them on par with each other. If you build once and deploy across multiple clouds, when you need to make updates, you can do so centrally and consistently. Compute stacks like Kubernetes lend a huge advantage to customers that are serious about doing multicloud in a way that's efficient and doesn't require reinventing the wheel every time you want to onboard a new cloud provider.

## 03

Third: **risk management.** Having the ability to run your stack in another environment will help mitigate the risk of your cloud provider going down or starting to compete with your business. To comply with regulations, organizations will pick providers to ensure their business continuity. For example, if you lose operations in one region, you won't experience any downtime with a backup provider.

The multicloud migrations that tend to work well are those that leverage open standards. Take Kubernetes, for example, which offers a provider-agnostic API both for running applications and for configuring and deploying them and for integrating things like security policies, networking, and more. Think of Kubernetes as a multicloud operating system; once it's your layer of abstraction, you usually can hide the differences between most major cloud providers.

## Reduce operational overhead by going fully managed

When you decide to use Kubernetes, you have choices. You can certainly run it yourself; Kubernetes is an open source project, so you can download it and spend years integrating it into your cloud provider or preferred environment.

But if you decide this isn't the best use of your time, you can use a managed Kubernetes offering. If you're on AWS, that would mean EKS. If you're on Azure, that will mean AKS. And if you're on Google Cloud, that will mean [Google Kubernetes Engine](#) (GKE). All those options will give you a common Kubernetes API, so when your team builds out its toolings and workflows you can reuse them across various providers.

But not all managed service offerings are created equal. Kubernetes is only as good as the infrastructure it runs on, and GKE fills the gaps as a mature, fully managed Kubernetes orchestration service. It offers fully integrated IaaS ranging from VM provisioning of [Tau VMs](#), with a 42% better price-performance over comparable general-purpose offerings,<sup>1</sup> autoscaling across multiple zones and upgrades, to creating and managing GPUs, TPUs for machine learning, storage volumes, and security credentials on demand. All you have to do is put your application in a container and choose a system based on your needs.

What if you've chosen AWS as your cloud provider for VMs? Do you need to stick with EKS? At a high level, Kubernetes is equal across all cloud providers; you'll end up with the Kubernetes API. But beneath that API is a cluster, worker nodes, security policies, the whole nine yards - and this is where GKE stands out.

For example, if you still need those other clusters, you can connect them to [GKE Connect](#), which will give you a single place to manage, view, troubleshoot and debug them all, while also centrally managing things like credentials. GKE is best-of-breed Kubernetes because of its end-to-end manageability, not just because of its control plane or its multi-region or multi-zone high availability. GKE can also leverage global load balancers using the centrally managed [Multi Cluster Ingress](#) across multiple clusters and multiple regions.

What if you want the Kubernetes API but not the responsibility of provisioning, scaling, and upgrading clusters? For the majority of workloads, [GKE Autopilot](#) abstracts away the cluster's underlying infrastructure, including nodes and node pools, and you pay only for the workload. GKE Autopilot is all about giving you standard Kubernetes API with all the strong security defaults, so you can focus on your workloads and not the cluster.

<sup>1</sup>Results are based on estimated SPECrate@2017\_int\_base run on production VMs of two other leading cloud vendors and pre-production Google Cloud Tau VMs using vendor recommended compilers. SPECrate is a trademark of the Standard Performance Evaluation Corporation. More information available at [www.spec.org](http://www.spec.org).


## Simplify app delivery with serverless

Kubernetes moved organizations from VMs to a new set of abstractions that allow you to automate operations and focus on your applications. But for more specific workloads (e.g. web and mobile apps, REST APIs backend, data processing, workflow automation), you can simplify even further and optimize your deployment by leveraging the serverless model.

Maybe you're using AWS Lambda, a popular serverless platform that lets you write functions as a service and connect them to all kinds of events. But because you end up connecting to a database and handling security concerns, these functions tend to grow in complexity, some actually bigger than normal applications. So what happens when you have an application that outgrows the simplicity of a function as a service, or an existing application that you want to run in serverless fashion?

Unlike a traditional serverless platform that requires you to rewrite your applications, [Cloud Run](#) offers an approach that helps you reuse your existing containerized application investments. Even though GKE is a managed service, you still have to make some key decisions: what zones to run in, where to store logs, how to manage traffic between different versions of applications, registered domain names, managing SSL certificates.

Cloud Run eliminates all those decisions, letting you run more traditional workloads and avoid cold starts by disabling scaling to zero altogether. If your applications have to always run, Cloud Run also supports that, along with other traditional requirements like NFS, WebSockets, and VPC integrations. But like most traditional serverless platforms, Cloud Run is opinionated and offers features like built-in traffic management and autoscaling.





# How you can get the biggest bang for your migration buck

## Thinking through your migration logic

Suppose you're not using containers at all and you're wondering where to start. Here's the pragmatic approach for adopting containerization.

### 01

The first reason to adopt containers is **solving the packaging problem**. There's a lot of work happening today around producing reproducible artifacts and understanding what's actually inside our software; or, as the industry calls it, the ["secure software supply chain."](#) We believe an ideal way to do that is to leverage container images that hold your application code and runtimes and their dependencies. One benefit of containers is that they can be deployed to VMs, reducing the complexity of deploying and maintaining software across your machines.

### 02

The second reason to adopt containers is **orchestration**. Managing VMs comes with tons of management and maintenance overhead. If you're like most companies, your team executes dozens or hundreds of steps to manage your infrastructure, and even if those steps are automated, those automation tools will still require ongoing maintenance. And that's assuming you're using industry standard automation tools such as Terraform, Puppet, Chef, and Ansible. The maintenance overhead is even worse for homegrown or custom tools.

### 03

The third reason to adopt containers is **efficiency**. On top of the maintenance burden, a majority of organizations only hit 5-10% CPU utilization, not to mention memory and storage. That's a lot of wasted resources. Many teams are building even more custom tooling to close this gap by implementing things like autoscaling and bin packing, thus squandering more operational time. This leads to overspending and a surprisingly high cloud bill.

Very few organizations are able to successfully build their own orchestration platform, which is why leveraging an open source platform like Kubernetes, Mesos, or Nomad is a common choice. But if you want a platform that helps you realize benefits like reduced maintenance, industry-standard best practices, and deep integration with the rest of your cloud provider, you're going to want to go with a managed service like GKE to maximize your containers' potential value.



## The fundamentals of getting your migration right

At this point you might be asking: does migrating to containers really amount to a bunch of downtime without any value? The last thing you want to do is hit the pause button on future development, right?

To answer this, let's walk through how to get it right in a way that leverages your cloud experience to date.

It may seem daunting to migrate applications away from VMs and into containers, especially if it means moving all your compute to a new cloud provider. But the reality is, you don't have to go wholesale; you can do this one application at a time, starting with VMs and moving one or two applications over to Kubernetes, where they can live side by side on the same network and talk to the VMs. You don't have to embark on a be-all-and-end-all transition; you can move slowly from one platform to the other.

Staying put on VMs might make sense for a few apps, like big, heavy databases that benefit little from Kubernetes. And that's fine; it's okay to mix and match. But what we've seen is that most customers get a lot of mileage out of moving their applications and open source projects into the Kubernetes landscape. Prioritize the applications that will yield the greatest return by moving to GKE; you don't have to move everything at once.

Kubernetes still runs on top of the same Linux VMs that you're probably using today. What you're really getting is a more streamlined and consistent workflow that incorporates a lot of things that should be on your infrastructure road map, rather than what could be a collection of homegrown scripts and automation. Onboarding new team members also becomes much easier when you're using an industry-standard tool like Kubernetes versus having to teach them all the custom ways that your company goes around it.

At this point, you have a path forward to adopting containers slowly but pragmatically, thus saving your team time in terms of administration overhead and saving your company money on compute costs.

# Ready to take your next steps?

We've talked a lot here about containers and cloud migrations - whether you're all-in on Kubernetes and looking for a better offering, or want a different approach that fits your existing applications and road map. Whatever your app development path, the managed containers option minimizes your infrastructure costs while maximizing your team's ability to focus on building great products. The future of infrastructure is containerization. It's time to ensure your engineers and your business are set up to be as successful as they can be.



To learn more about how Google Cloud's industry-leading managed services can help you make the most of your modern apps journey, please get in touch.

[Talk to an expert](#)



## Further reading

- [Kubernetes to Cloud Native: Jumpstart your journey developing on GKE](#)
- [Containers in Cloud - A better way to develop and deploy applications](#)
- [The past, present, and future of Kubernetes with Eric Brewer](#)
- [Re-architecting to cloud native: an evolutionary approach to increasing developer productivity at scale](#)

