

LEARNING MADE EASY

SUSE Special Edition

Kubernetes Management

for
dummies[®]



Define your
strategy

Learn from industry
leaders

Select the right
platform

Compliments
of



Tom Callway

Peter Smails

Caroline Tarbett

Shannon Williams

About SUSE and Rancher

In December 2020, SUSE completed the acquisition of Rancher Labs — the company behind Rancher, RKE, K3s, and Longhorn. Each of these technologies will remain as free, open source projects. However, in 2H 2021, the fully supported, commercial version of Rancher will be called SUSE Rancher.



Kubernetes Management

SUSE Special Edition

**by Tom Callway, Peter Smails,
Caroline Tarbett, and Shannon
Williams**

for
dummies[®]
A Wiley Brand

Kubernetes Management For Dummies®, SUSE Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2021 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

ISBN 978-1-119-78165-3 (pbk); ISBN 978-1-119-78172-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor:

Elizabeth Kuball

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development

Representative: Matt Cox

Production Editor:

Tamilmani Varadharaj

Special Help: Colleen Diamond

Introduction

Kubernetes is not just the leading container orchestration solution — it has become the standard. Technology employees of all stripes love working with Kubernetes:

- » **Developers** love the extensibility of Kubernetes, which gives them maximum agility and flexibility when delivering cloud-native applications, whether on-premises, in the cloud, or at the edge.
- » **Information technology operations (ITOps) teams** love Kubernetes because it helps boost productivity, reduce costs and risks, and moves organizations closer to achieving their hybrid cloud goals.
- » **CEOs and CIOs** love Kubernetes because it helps significantly increase the agility and efficiency of their software development teams, enabling them to reduce the time and complexity associated with putting differentiated applications into production.

Simply put, Kubernetes makes it easier to manage software complexity. As enterprise applications become more complex, development and IT operations teams need a tool that can orchestrate that complexity. They need a way to launch all the services dependent on these applications, making sure the apps and services are healthy and can connect to one another.

This book provides you with a basic road map to successfully develop a Kubernetes strategy and select the appropriate Kubernetes management solution to address your needs. We include three detailed case studies from companies leading the way with Kubernetes and Rancher as their solution of choice. Our hope is that IT professionals and others who lead an organization of any size can see how other technology leaders have used Kubernetes to improve their operations.

Foolish Assumptions

To help you get the most out of this book, we make some assumptions about you:

- » We assume that you already have some basic familiarity with containers and Kubernetes, or you've at least heard the terms mentioned in passing. We do *not* assume that you know all there is to know about Kubernetes or that you've even heard of Rancher.
- » We assume that regardless of whether you head up your DevOps or IT department or work in an executive capacity, you come to this book ready to solve a problem, and you want to determine whether Kubernetes, Rancher, and other related technologies are the solutions you're seeking.

Icons Used in This Book

This book uses the following icons to highlight paragraphs that have a little something “extra” about them:



TIP

The Tip icon calls your attention to information that may help make your Kubernetes venture lift off a little more smoothly.



REMEMBER

The Remember icon highlights information that is so important, it's worth repeating — and remembering!



WARNING

The Warning icon calls attention to pitfalls you may encounter in your Kubernetes implementation and how to mitigate them.

Beyond This Book

The Rancher blog (<https://rancher.com/blog>), Rancher docs (<https://rancher.com/docs>), and Rancher community links (<https://rancher.com/community>) are chock-full of helpful information — and searchable!

IN THIS CHAPTER

- » Understanding why IT teams love Kubernetes
- » Assessing your progress on your Kubernetes journey
- » Establishing who owns your strategy
- » Setting goals
- » Meeting your standardization and innovation needs
- » Training your teams

Chapter 1

Creating an Enterprise Kubernetes Strategy

To understand why Kubernetes is so popular, you first need to know why containers have risen dramatically in popularity. Containers provide a consistent way to package application components and their dependencies into a single object that can run in any environment. By packaging code and its dependencies into containers, a development team can use standardized units of code as consistent building blocks. The container runs the same way in any environment and allows applications to scale to any size.

Development teams use containers to package entire applications and move them to the cloud without making any code changes. Containers also simplify the process of building workflows for applications that run between on-premises and cloud environments, enabling the smooth operation of almost any hybrid environment.

So, why Kubernetes? Enterprises love Kubernetes because it helps significantly increase their software development teams' agility

and efficiency, enabling them to reduce the time and complexity associated with putting differentiated applications into production. Information technology operations (ITOps) teams love Kubernetes because it helps boost productivity, reduce costs and risks, and move organizations closer to achieving their hybrid cloud goals. Developers love the extensibility of Kubernetes, which gives them maximum agility and flexibility when delivering cloud-native applications.

Simply put, Kubernetes makes it easier to manage software complexity. As enterprise applications become more complex, development and operations teams need a tool to orchestrate that complexity. They need a way to launch all the services dependent on these applications, ensuring the apps and services are healthy and can connect to one another.

Knowing Where You Are on Your Kubernetes Journey

Building an enterprise Kubernetes strategy starts with understanding where Kubernetes is running in your organization and imagining how it's going to change over the next decade.

Over the last several years, accessing Kubernetes has become much easier. Open-source tools make provisioning and upgrading a Kubernetes *cluster* (a set of nodes that run a containerized application) quick and easy. Cloud providers are now offering Kubernetes as a hosted service. Any team using Amazon Web Services (AWS), Google Cloud Platform (GCP), or Microsoft Azure can provision a Kubernetes cluster in minutes.

Organizations that run Kubernetes often approach it the same way as when they built OpenStack or other shared, centralized services. ITOps teams typically use Kubernetes to build large clusters and then offer development teams shared access to them through Kubernetes namespaces. *Namespaces* enable cluster administrators to control access to cluster resources based on usage quotas and resource limits. Namespaces help deliver a reasonably well-isolated experience for each team that needs access to Kubernetes.

COLLABORATING ACROSS TEAMS

Tension can develop between software development teams who need to run Kubernetes in a certain way to accomplish a development goal and an IT department that prioritizes maintaining security and control over how Kubernetes gets implemented.

Development teams want flexibility. Having cluster-level administrative control allows them to configure a cluster to run exactly how they need it to in terms of storage, security policy, or which infrastructure it runs on.

On the other hand, IT teams are especially nervous about clusters that are deployed and left unpatched and unmanaged. They want to centralize the operations and policy around clusters and restrict access to only those teams that require it.

If Kubernetes and containers are going to become the primary platform for running applications across any infrastructure, ITOps must collaborate with developers on a plan and a strategy for Kubernetes that satisfies both their needs.

Other organizations have left it to individual departments or development and IT operations (DevOps) teams to decide how and where to use Kubernetes. These organizations often have dozens of clusters deployed across public clouds and company data centers.



TIP

As you document and understand where Kubernetes is running in your enterprise, be on the lookout for individuals who show existing expertise in containerization. As you progress in building your strategy, developing a team of experts who can administer your Kubernetes clusters and deploy applications to them will be critical to driving adoption.

Defining Who Owns Your Strategy

New technologies like Kubernetes are exciting to work with, and it isn't uncommon for many teams to try to own their company's Kubernetes strategy — individual DevOps teams, shared services groups, central IT, cloud platform, or platform-as-a-service (PaaS) groups.

Two teams that often lead the Kubernetes strategy are the shared services team (responsible for supporting developers and DevOps) and the central IT team (responsible for computing platforms). Putting either team in charge of Kubernetes strategy provides the following benefits:

- » **Shared services:** The shared services team brings key insights on how an organization is modernizing its approach to application development, as well as the requirements teams have identified they need in a Kubernetes platform. They often understand other key systems that have been built for DevOps, such as continuous integration/continuous delivery (CI/CD) tools, development environments, data services, and application monitoring tools. Whether these teams own the strategy or simply contribute to it, they represent at the very least one of the primary consumers of containers in the organization. They should be a critical part of developing your organization's strategy.
- » **Central IT:** The central IT team, focused on cloud computing and other computing platforms, is also a logical team to lead a Kubernetes strategy. They have a strong understanding of platform operations, infrastructure, security, multi-tenancy, and existing IT investments, and they usually have significant experience running critical projects. A project led by the IT platforms team will benefit from their understanding of the broad requirements of many different teams across a large, complex organization.

Note that projects coming out of central IT often suffer from too little engagement with end users and too much influence from existing technology vendors. These teams often have very little experience with the latest application architectures and benefit enormously from working closely with teams leading innovation around application development.



TIP

Successful teams often bring together talent from across the organization and collaborate to determine requirements. Still, investing in a strategy and building a platform means working within budget constraints, so it's most common for one team to take the lead on delivering on the strategy.

Prioritizing Your Goals

Building an organization-wide Kubernetes strategy means prioritizing your goals for this new technology.

If your team sets out to use Kubernetes as a means to reduce infrastructure costs, you'll probably focus on building big clusters and trying to get as much density as possible out of them.

If your team focuses instead on using Kubernetes to accelerate innovation, you'll take a different approach, emphasizing flexibility and delivering more tooling around Kubernetes, such as monitoring and CI/CD integration.

To prioritize your goals, try to understand the potential of Kubernetes, and imagine how your organization may be using it in the future.

In five years, for example, you may use Kubernetes to do any of the following:

- » **Create microservice-centric applications.** Kubernetes is a great way to run modern, microservice-centric applications. It offers a rich set of functionalities that allow teams to determine how different services within modern applications are run, handle unexpected events, connect with each other, and connect with other applications and application programming interfaces (APIs).
- » **Rapidly deploy Kubernetes clusters.** Today, every major cloud provider has made it easy to deploy Kubernetes clusters within minutes. Teams are continuously building new applications, deploying them to different clouds, and using Kubernetes to run them. Between clusters used for development, staging, and production, and the need to deploy Kubernetes clusters across different data centers and cloud providers, it isn't hard to imagine that even the most well-organized company is still running dozens of Kubernetes clusters.
- » **Move onto the edge.** The same modern application architectures that we think of as cloud-native are now beginning to move out of the data center. Teams building software for factories, hospitals, and stores now want to run applications with rich data analytics and complex

architectures as close to their customers and production facilities as possible. Running applications this way is referred to as “running on the edge.”

- » **Develop for single-node devices.** Even single-node devices such as point-of-sale terminals, outdoor advertising, medical devices, 5G-enabled communication equipment, security cameras, or automobiles now benefit from the ability to deploy and run applications easily using microservices. We’re witnessing the sprawl of tens of thousands of edge deployments, all running as individual Kubernetes clusters, and presenting an API that needs to be managed.

Between clusters running in different clouds, data centers, and the edge, it’s almost certain that your organization will be running more than one Kubernetes cluster. Unless you know you’ll only be running a single application in one location, it probably makes sense to build your Kubernetes strategy with an expectation that you’ll need to be able to easily provision and manage multiple Kubernetes clusters running in many different places.

Weighing Standardization against Innovation

Regardless of who owns your strategy, one of the critical questions that will emerge is how much standardization is possible without impacting innovation. Many teams will have experienced projects around OpenStack and PaaS that struggled to get adoption because users weren’t able to get enough flexibility to deploy the next-generation applications they were building.

With Kubernetes, there is enough flexibility in the platform and the ecosystem to satisfy any team. Exposing that flexibility is critical to delivering value. Any strategy that abstracts away Kubernetes will probably face resistance from your most innovative teams. At the same time, the flexibility of Kubernetes and its ecosystem can be a hindrance to some teams looking for a platform to just run standard apps.

One of the most exciting developments in the Kubernetes space in the past two years has been the emergence of lightweight projects that run on Kubernetes but provide frameworks that simplify application management. These approaches allow containers to “scale to zero” and provide simple declarative languages to build, connect, scale, and monitor services. They can deliver a powerful experience without requiring a deep understanding of the underlying technology, which can benefit teams using CI/CD and stateless applications.

Google Cloud Run (<https://cloud.google.com/run>) is an example of this approach to running containers that simplify some of the complexity of Kubernetes.

As you build your Kubernetes strategy, consider blending the best of a decentralized approach with enough controls and management to ensure compliance and remove repetitive tasks. Try to centralize and automate everyday tasks such as Kubernetes cluster lifecycle management, role-based access control (RBAC) policies, infrastructure management, and other day-2 operations.

At the same time, give your teams options for where they can get access to Kubernetes clusters and whether they can use a shared cluster or a dedicated cluster. Focus primarily on maintaining visibility into all the provisioned clusters, not necessarily forcing teams to use a set of preapproved clusters in a specified way.

Preparing Your Teams

A critical part of any Kubernetes strategy is determining how you'll train your teams to leverage Kubernetes. As we mention earlier, if you find that your enterprise already has some staff members with expertise in containers or Kubernetes, consider how you can incorporate them into your initiative. This doesn't mean necessarily pulling them off their existing work, but perhaps they can work as part of the team setting requirements, evaluating tools, or developing policies.



TIP

Regardless of your team's skill level, you'll almost certainly have team members who need to be trained on either using or administering Kubernetes. Luckily, there is no shortage of Kubernetes training providers and online courses. One example is the Rancher Academy (<https://academy.rancher.com>).

As you build your core team of early Kubernetes admins and users, consider setting a goal to train and certify as many members of your team as possible. The tests are rigorous and help you ensure that you build strong internal knowledge about using containers and Kubernetes.

After you have some initial expertise, you may want to wait to do further training until you're out of the design phase of your strategy and bringing on more teams to work with the specific implementations of Kubernetes your organization is adopting.

- » Exploring the layers of a container management platform
- » Deciding between an on-premises and a hosted platform
- » Understanding the value of experience

Chapter 2

Building an Enterprise-Grade Kubernetes Environment

Chapter 1 covers the basics of building an enterprise Kubernetes strategy that considers how your organization will use Kubernetes over the next five years. We also discuss the importance of maintaining flexibility while at the same time providing centralized controls and management when you build this strategy.

This chapter covers the next step in that process. You find out how IT operations and development teams should examine their options for managing containers at scale across their organization. Analyst firms such as Gartner and Forrester call the class of software that helps organizations do this a *container management platform* (CMP).

In this chapter, we discuss the key components of a CMP and what to look for when choosing a solution. You see examples of how Rancher, the world's most popular open-source CMP, can help implement your Kubernetes strategy in your enterprise.

Getting to Know Container Management Platforms

A CMP contains functional layers that work in concert to deliver all the capabilities you need to build and manage a Kubernetes infrastructure (see Figure 2-1).

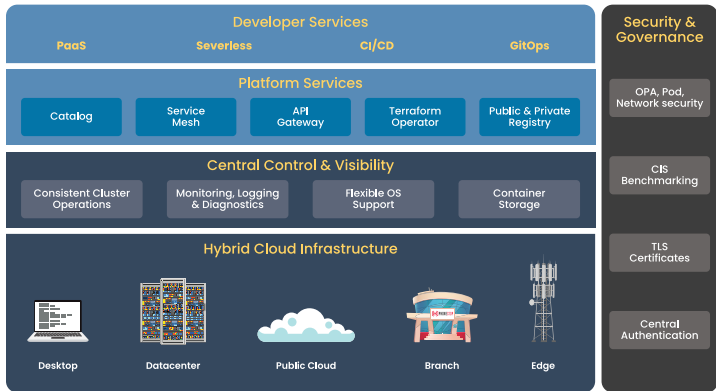


FIGURE 2-1: The anatomy of a container management platform.

The following sections describe the role each layer plays and discuss why that layer’s role is important.

Enabling innovation everywhere

Possibly the most powerful benefit of Kubernetes is its transformational ability to enable innovation everywhere. Using certified distributions, organizations can develop applications once and then deploy and run them wherever they need to run — from desktop to data center, the cloud, and the edge.

The secret to successfully enabling innovation everywhere, of course, is to use a certified Kubernetes distribution, a software package that provides a prebuilt version of Kubernetes.



WARNING

In the world of Kubernetes distributions, one size does *not* fit all. Don’t be fooled by vendors saying that all you need is “their” distribution. Be smart: Select the best distribution for the job based on where your application will be running.

The following sections describe places where you may run your applications.

Public clouds

To reduce the time and complexity of deploying Kubernetes, many organizations choose to deploy their clusters using public-cloud-based infrastructure such as Amazon Web Services (AWS) and Microsoft Azure. Most public cloud providers have developed their own certified Kubernetes distributions that are optimized for that cloud. So, it's important that your CMP supports Kubernetes in any public or private cloud environment and treats popular hosted distributions like Amazon Elastic Kubernetes Service (EKS), Microsoft Azure Kubernetes Service (AKS), and Google Kubernetes Engine (GKE) as first-class citizens.

The support Rancher offers for Amazon EKS is an excellent example of treating a cloud distribution like a first-class citizen. As of publication, Rancher provides full life-cycle management of EKS clusters. This includes importing, provisioning, upgrading, configuring, and securing them directly from within Rancher's unified, intuitive user interface (UI). Table 2-1 shows how Rancher 2.5 supports full life-cycle management of Amazon EKS clusters.

TABLE 2-1 Rancher 2.5 and Amazon EKS Clusters

Life-Cycle Management	Features Required	EKS with Rancher 2.5	EKS Only
Configure and provision	Enterprise Kubernetes management	Provision new clusters through the Rancher graphical user interface (GUI)/application programming interface (API) Enhanced EKS cluster import	AWS Management Console, third-party tools
Manage	Infrastructure management	Enhanced configuration of the underlying infrastructure	AWS Management Console

(continued)

TABLE 2-1 (continued)

Life-Cycle Management	Features Required	EKS with Rancher 2.5	EKS Only
	Visualize Kubernetes resources	New Rancher dashboard, which explores all Kubernetes resources	Kubect!, third-party tools
	Integrated monitoring and logging	Enhanced monitoring (Prometheus) Enhanced logging (Fluentbit/ Fluentd)	Manual install
	Simplified service mesh	Rancher-supported Istio 1.6	Manual
Secure	Centralized tooling and visibility	Centralized role-based access control (RBAC) policy definition Centralized authentication Center for Internet Security (CIS) scanning Open Policy Agent (OPA) Gatekeeper	Kubernetes native
Upgrades	Push-button Kubernetes upgrades	Rancher GUI/API upgrades of created and imported clusters	AWS Management Console, third-party tools
Apps	Easy access to Cloud Native Computing Foundation (CNCF) tools and third-party apps	Rancher-certified packages Custom Rancher catalogs Helm Kubect!	Help, Kubect!

On-premises and hybrid cloud

Most Kubernetes management vendors provide a Kubernetes distribution that can provision and maintain a cluster, including upgrading it to the latest version of Kubernetes. Some Kubernetes distributions also include integrated monitoring, etcd backup and recovery, and infrastructure provisioning and auto-scaling.



TIP

To find out more about etcd backup, see <https://rancher.com/blog/2019/2019-01-29-what-is-etcd>.



REMEMBER

Be sure that your Kubernetes distribution is certified by the CNCF. This certification ensures that the distribution is consistent with upstream Kubernetes and quickly supports the latest features being developed in the community.



TIP

Rancher offers RKE for on-premises datacenter or hybrid cloud uses-cases. RKE is a CNCF-certified Kubernetes distribution that runs entirely within containers. This approach solves the common frustration of installation complexity with Kubernetes by removing most host dependencies and presenting a stable path for deployment, upgrades, and rollbacks. RKE Government is a hardened, FIPS-enabled derivative of RKE. By adopting a compliance-based approach toward security, RKE Government is ideal for applications requiring FIPS compliance.

The edge

Enterprises are increasingly pushing the boundaries of Kubernetes by exploring how to run next-generation 5G, machine learning (ML), and artificial intelligence (AI) workloads at the edge. As a result, Kubernetes-powered applications are popping up everywhere, including retail stores, banks, factories, cars, trains, ships, oil rigs, and wind farms — just to name a few! The inherent complexity and resource overhead of Kubernetes, however, makes “full-size” distributions too big to operate or too complex to manage in these potentially resource-constrained environments.

Today, for these edge environments, Rancher supports a lightweight, highly available, CNCF-certified Kubernetes distribution, called K3s, which is built to run production workloads in unattended, resource-constrained, remote locations, or inside Internet of Things (IoT) appliances. First created by Rancher Labs (acquired by SUSE in 2020) and then donated to the CNCF in 2020, K3s is packaged as a single less-than-100MB binary that reduces the dependencies and steps needed to install, run, and auto-update a production Kubernetes cluster.

Maintaining central control and visibility

Leading CMPs support multi-cloud and multi-cluster use cases and excel at simplifying cluster operations and making them consistent across substrates. Many also include advanced observability tools with minimal additional configuration required.

Simplified cluster operations

You may be wondering whether you're likely to be managing hundreds, thousands, or even millions of containers across your entire infrastructure. If left unchecked, such *container sprawl* can become a big problem.



TIP

When shopping for a CMP, here are some of the key features and capabilities to look for that will help manage cluster sprawl:

» **Intuitive UI:** A CMP needs to enable users to easily provision multiple Kubernetes clusters from core to cloud to edge and make those clusters accessible to all users at any stage in the clusters' life cycles. Look for solutions with an intuitive UI that gives you complete access to everything you need to deploy and manage Kubernetes workloads while still giving power-users full access via command-line interface (CLI) tools.

» **Copy-and-paste YAML:** If you have existing manifests that you want to deploy, you should be able to copy them into your CMP and have the CMP implement them for you. This capability gives you an easy way to use YAML copied from documentation, tutorials, or other users, without needing to first save it and then apply it via `kubectl`.

Note: YAML stands for "Yet Another Markup Language." You can find a useful definition of YAML at <https://developer.ibm.com/technologies/containers/tutorials/yaml-basics-and-usage-in-kubernetes>.

» **Flexible load-balancing configuration:** Load balancing may be an important part of your application deployment, so any CMP should make it easy to select the best ingress controller for your use case. Some examples of ingress controllers include HAProxy, NGINX, and Traefik. You want the flexibility to use the load-balancing configuration that works best for you and your environment.

- » **Visual support for Secrets and ConfigMaps:** Secrets and ConfigMaps are two powerful Kubernetes resources. Your CMP should allow you to define and update both types visually and then select how to map them into workloads — either as environment variables or volumes.
- » **Rolling updates of worker nodes:** As more business-critical applications become containerized, zero downtime maintenance becomes an operational imperative. The best CMPs support rolling updates of multiple worker nodes, making it easy for operators to select and configure an upgrade strategy so that Domain Name System (DNS) and Ingress experience zero downtime during cluster updates.

Monitoring, logging, and diagnostics

Monitoring, logging, and diagnostics are critical to maintaining cluster health. Look for a CMP that lets you quickly add these services to operationalize your clusters.

For example, Rancher integrates with Prometheus and Grafana to visualize and gain insights into live monitoring data. If you already use third-party monitoring platforms like Datadog, Dynatrace, Sensu, or Sysdig, you can launch them easily from Rancher’s app catalog.

You also need to troubleshoot and debug your clusters by looking for trends in the log stream. And, because the logs are stored outside the cluster, you need to access them even if your cluster fails. To that end, Rancher also integrates with popular logging tools like Elasticsearch, Fluentd, Kafka, Splunk, and Syslog to capture and analyze logs.



TIP

For more information on monitoring and alerting best practices, see “Best Practices for Monitoring and Alerting on Kubernetes” on Rancher’s blog (<https://rancher.com/learning-paths/best-practices-for-monitoring-and-alerting-on-kubernetes>).

Flexible operating system support

Whether you want to create and tear down development/test environments or “lift and shift” legacy applications to the cloud, your CMP needs to provide flexible operating system (OS) support that includes both Linux and Windows. This way, you can bring

the benefits of Kubernetes to all your existing and new containerized applications.



TIP

For more information about how Rancher supports launching Kubernetes on Windows clusters, check out <https://rancher.com/docs/rancher/v2.x/en/cluster-provisioning/rke-clusters/windows-clusters>.

Container-attached storage

With more stateful applications such as database applications becoming containerized, persistent storage has become an important feature for any CMP. As a result, numerous open-source and proprietary persistent storage solutions have become popular in the Kubernetes ecosystem, including Longhorn, OpenEBS, and StorageOS.

Today, Rancher provides out-of-the-box integration with a persistent storage solution (Longhorn), making it easy to provision, secure, and back up highly available container-attached storage in your Kubernetes environment with just a few clicks.



TIP

To read more about how Rancher integrates with Longhorn to create vendor-neutral persistent storage, see “Longhorn Simplifies Distributed Block Storage in Kubernetes” on Rancher’s blog (<https://rancher.com/blog/2020/longhorn-container-storage>).

Ensuring global security and governance

From integrating with popular authentication tools and services to configuring an enterprise-grade RBAC capability, any CMP must ensure the security of your single, multi-cluster, or edge-scale Kubernetes environment.

In addition to platform-level security, your CMP should provide easy access to the vibrant ecosystem of container security technology vendors. These vendors offer specific security capabilities that are worth evaluating as part of your broader implementation of Kubernetes. For example, Rancher provides seamless access to leading security tools, including Aqua Security, NeuVector, Portshift, and Prisma.

The following sections cover some of the specific security capabilities to look for in a CMP.

Centralized authentication and role-based access control

Your CMP should provide centralized authentication and RBAC for all your Kubernetes clusters and users, enabling users to connect to any cluster with one set of credentials stored in the authentication service of your choice — from GitHub to OpenLDAP to Active Directory. Administrators can then grant user/group access to any cluster or project by leveraging custom resource definitions (CRDs) and custom controllers for RBAC.



TIP

For more on centralized authentication, check out <https://rancher.com/docs/rancher/v2.x/en/admin-settings/authentication>. And for more on RBAC, check out <https://info.rancher.com/authentication-authorization-multiple-clusters-kubernetes>.

Transport Layer Security certificates

Your CMP should be able to store TLS certificates to keep them safe. Subsequently, users can deploy resources that use a certificate without being given a copy of the certificate and private key. After it's installed, a certificate's private key should be held securely by the CMP.



TIP

For more on TLS certificates, check out <https://rancher.com/docs/rancher/v2.x/en/k8s-in-rancher/certificates>.

Center for Internet Security benchmarking

The more clusters you manage, the higher your risk of security exposure. To avoid noncompliant clusters, look for a CMP that provides cluster templates, which enable you to apply cluster settings uniformly across many clusters to prevent configuration drift.

Additionally, look for a CMP that provides the ability to automatically scan clusters against CIS, which offers more than 100 benchmarks for validating the security of your clusters.

Open Policy Agent, pod, and network policies

The OPA Gatekeeper project is a policy enforcement mechanism that forces Kubernetes clusters to run and access designated privileges and resources. Gatekeeper helps you ensure compliance with legal and organizational policies by providing the ability to define custom policies using native Kubernetes CRDs. Look for a CMP that uses OPA as a threat prevention mechanism by enabling controlling policies for images, ingress, pods, and namespaces.

Leveraging open, flexible platform services

The following sections detail what to look for when considering open, flexible platform services.

App catalog

Helm is one of the most popular tools to package and deploy applications into your cluster safely. Look for a CMP that extends Helm charts to simultaneously install and upgrade certified applications in multiple clusters from a global application catalog. If you're exploring a multi-cloud or multi-provider solution, this feature ensures flexibility while also giving you confidence that your applications stay consistent.



TIP

For more on deploying applications from catalogs, check out <https://rancher.com/docs/rancher/v2.x/en/catalog/launching-apps>.

Service mesh

Service mesh is designed to eliminate developers' need to write specific code for key application services, including fault tolerance, canary rollouts, A/B testing, monitoring and metrics, tracing and observability, and authentication and authorization. Look for a CMP that integrates with service mesh technologies, including Istio, and traffic/telemetry visualization tools like Jaeger and Kiali.

Application programming interface gateway

A microservice-based architecture running on Kubernetes may have 10 to 100 or more services. Look for a CMP that supports an API gateway to provide a unified entry point for external consumers, independent of the number and composition of internal microservices.

Terraform operator

Infrastructure as code is an important methodology for ensuring that your distributed systems are treated like cattle and not pets. Your Kubernetes clusters are no different. You should be able to automate the provisioning of your Kubernetes clusters and all your apps.

For example, Rancher supports Hashicorp's Terraform, an open-source tool that automates creating and versioning infrastructure and other system components. Terraform is popular because its syntax is easy to use. The tool is also highly customizable with a pluggable provider mode.



TIP

For more on Terraform and Rancher, check out <https://rancher.com/blog/2019/rancher-2-terraform-provider>.

Public and private registry support

Your CMP should support deployment from any public registry. If you also use private registries, you should be able to load the authentication data into your CMP when you deploy workloads that use containers from a private registry and securely pass the authentication information to Kubernetes for use when pulling the images.

Giving developers choice

Kubernetes is a powerful engine with a rich ecosystem of development tools around it. As such, no one best way exists for developers to leverage Kubernetes. Instead, the key is to provide developers the flexibility to use the development tools they want.

For some developers, the answer may be an entirely curated user experience around Kubernetes that incorporates ecosystem tools and delivers a UI to simplify workload management. This is referred to as *platform as a service* (PaaS). A potential drawback to PaaS is limited functionality or lack of flexibility.

Alternatively, developers may have a more do-it-yourself mindset, requiring your CMP to integrate with adjacent technologies. These could include the container engine, overlay networking, automation tooling, container registries, service mesh, monitoring, logging, continuous integration/continuous delivery (CI/CD), and application catalogs.



REMEMBER

The best CMPs support a variety of developer experiences. Some developers work best with a UI that simplifies workload but restricts functionality; others prefer to do more of the integration work themselves and, therefore, work best with a CMP that provides more flexibility.

The following sections describe several of the most popular developer experiences.

Platform as a service

The intention of a PaaS is to eliminate manual IT configuration and help accelerate innovation by getting applications to market faster. Developers can serve themselves and get apps to the cloud in minutes instead of weeks while staying within IT guidelines or relying on scarce IT resources to perform manual configuration each step of the way.

Serverless

Serverless architectures refer to the application architecture that abstracts away server management tasks from the developer and enhances development speed and efficiency by dynamically allocating and managing compute resources. Function as a service (FaaS) is a runtime on top of which a serverless architecture can be built.



TIP

For more about serverless frameworks for Kubernetes, see “Evaluation of Serverless Frameworks for Kubernetes (K8s)” on Rancher’s blog (<https://rancher.com/blog/2018/2018-04-23-evaluation-of-serverless-frameworks-for-k8s>).

Continuous integration/continuous delivery

The demands of modern software development, combined with the complexities of deploying to various infrastructure, can make creating applications a tedious process. As applications grow in size and scope, and development teams become more distributed and diverse, releasing software quickly and consistently becomes more difficult. To address these issues, teams must automate their build, test, and release processes using CI/CD pipelines.



TIP

Look for a CMP that supports popular CI/CD pipeline tools like Jenkins to simplify all aspects of the application delivery process. For more on Rancher’s CI/CD support, check out <https://rancher.com/docs/rancher/v2.x/en/project-admin/pipelines>.

GitOps at scale

GitOps is a new way to do Kubernetes cluster management and continuous deployment (CD). It works by using Git as a single source of truth for declarative infrastructure and applications. With Git at the center of your delivery pipelines, developers can use familiar tools to accelerate application deployments and operations tasks to Kubernetes.

Look for a CMP that can apply GitOps best practices to your infrastructure and application code, allowing developers to increase velocity and improve system reliability.



TIP

For more about GitOps in Kubernetes, see “The GitOps Kubernetes Connection” on the Rancher blog (<https://rancher.com/blog/2020/gitops-kubernetes-connection>).

Choosing the Best Deployment Option for Your Needs

Before you can choose your CMP, you need to decide how you'll operate it. You may want to host your CMP within your data center to ensure optimal security. Or, if you don't have the infrastructure or people resources required to host your CMP, you may choose a public-cloud-hosted option. Ideally, your CMP gives you the option to do either.

Rancher, for example, can be deployed on-premises. Rancher is also available in a hosted version so that you don't have to operate the Rancher control plane and can instead focus exclusively on the day-2 operations of your Kubernetes clusters.



REMEMBER

The CMP you choose should depend on how you intend to operate it. You may prioritize security, or you may want to streamline resources.

Learning from Experience

Adopting new technology across a large organization is never easy. As technologists, we get excited when new approaches emerge that can create amazing experiences for our customers. Many of us who have been working in technology for the last 20 years see Kubernetes and containerization as the third phase in a process that started with the emergence of virtualization and expanded with cloud computing.

As you adopt a CMP, be sure to learn from your organization's past successes and failures in adopting these other technologies. If you have team members who were instrumental in rolling out VMware or AWS in your organization, incorporate them into your project and see what insights they can provide to your organization.

Pay special attention to the teams who are already running apps on Kubernetes. Use their expertise to validate that your preferred CMP won't introduce constraints that would keep them from adopting it. Focusing on the early adopters will help you avoid oversimplification and delivering a platform that deviates from mainstream Kubernetes adoption.

As you begin to deploy your CMP, know that you aren't alone, and pay special attention to learning from other organizations adopting Kubernetes. The Kubernetes community is growing rapidly and provides a wealth of real-world advice from teams who have gone through rolling out Kubernetes at either a project or company-wide level.

IN THIS CHAPTER

- » Following Schneider Electric's container-migration journey
- » Analyzing what containers have done for Continental
- » Digging into containers at the Municipal Property Assessment Corporation

Chapter 3

Looking at Case Studies

Reading about high-level strategic approaches and the benefits of a particular platform is undoubtedly useful. Still, nothing beats a detailed description of how customers benefit from the technology and hearing their insights.

This chapter contains case studies of companies of different sizes, from various industries and locations, to illustrate how beneficial Kubernetes and Rancher can be to an organization's digital transformation.

Schneider Electric

Schneider Electric is one of the most innovative and long-established global companies in its market. Founded in the 1800s, the company is a world-leading provider of energy and digital automation solutions for efficiency and sustainability.

Believing access to energy and digital services is a basic human right, Schneider Electric creates integrated solutions for homes, commercial and municipal buildings, data centers, and industrial infrastructure. By putting efficiency and sustainability at the heart of its portfolio, the company helps consumers and businesses make the most of their energy resources.

Head of Global Infrastructure Strategy, Anthony Andrades, is guiding the company through a period of significant transformation. He's building Schneider Electric's strategic vision and analyzing everything the business does from an innovation perspective. His analysis encompasses how their estate of data centers operates, the diverse ways applications are built and run, asset obsolescence, configuration, and cost. Andrades is also responsible for managing the cultural shift associated with large-scale digital transformation. According to Andrades, "After a quarter of a century of technical evolution, we're embarking on one of the most important transformations in our history. By modernizing all our legacy systems to create a cluster of cloud-native microservices, we are becoming more agile and innovative."

Seeing why Schneider Electric chose containers

Schneider Electric had already entered the cloud ecosystem in 2013, with a couple of business-driven projects running quietly in Amazon Web Services (AWS) and Microsoft Azure. When the success of these projects became known, Andrades was drafted in to build on this success and create an enterprise-grade cloud strategy. By 2016, the company had expanded its global AWS footprint and its mission to migrate its infrastructure to the cloud had begun.

The team became aware of Kubernetes a year earlier in 2015 and quickly identified it as a cost-effective way to create the microservices-based, service-oriented architecture that large digital enterprises, like Google, had pioneered. There were some pockets of excellence where Kubernetes was already running, but the picture wasn't consistent. Access control was a major issue. Several customer development teams needed access to clusters, but this was uncontrolled which, in some cases, resulted in the suspension of Docker usage until a rules-based platform as a service (PaaS) was put in place.

The team was already familiar with Rancher so, in early 2018, Andrades carried out an initial successful proof of concept (PoC) with Rancher Labs (acquired by SUSE in 2020) and their

security partner, Aqua. Soon after, the team started using Rancher on top of Kubernetes to provide the access control, identity management, and globalized performance metrics that don't ship with Kubernetes.

Rancher performed so well that Schneider Electric chose Rancher to underpin its container-management platform. In June 2019, the platform was deployed to run 20 nodes, and the painstaking process of application modernization began.

Understanding the problems Schneider Electric is solving

After an extended period of technical evolution, Schneider Electric knew it needed to adopt a strategy that transformed legacy technology and embraced the cultural shift required to reorient its business.

Legacy transformation

Like many established businesses, Schneider has been through 25 years of transformation. Over time, the company has built and deployed thousands of separate services and applications running on Windows Server or Red Hat that must be re-engineered or rebuilt before migrating to the cloud.

Andrades's primary objective is to complete the transformation and migration of all applications within five years. This is no small feat when you consider the volume of applications involved and how different applications require different modernization approaches. In late 2019, the team started the painstaking process of analyzing the entire estate of applications, categorizing each one according to the most appropriate and efficient way to modernize and migrate.

For some key applications, the transformation will be done in stages; the application will be "lifted and shifted" to the cloud, optimized, and made available as a service. Teams will then redesign the application later. Others may be decommissioned entirely and rebuilt as microservices. Static web servers, for example, can easily be converted into S3 buckets. Where two-tier applications are concerned (web front end running a user interface [UI] with a relational database in the back end), the UI would run in a container and the database would be ported to Amazon Relational Database Service (RDS).

In Kubernetes, development teams can deploy multiple clusters, each configured to specific application requirements. In Rancher, the infrastructure team can run each of these bespoke environments side by side via one intuitive platform. Crucially, when used with other solutions, such as Aqua, Rancher becomes a secure and compliant environment for teams — both internal and external — to collaborate. With access control easily configurable in Rancher, the infrastructure team can allow unhindered access to the platform. This approach significantly boosts team innovation.

The project is in its infancy, but Andrades already sees benefits daily. He has a mammoth task ahead: If he is to reach his five-year migration goal, he must automate a host of basic processes, such as role-based access control (rbac), namespace as a service (NaaS), authentication, application catalog, and more. Rancher takes care of these functions, dramatically reducing the deployment workload. According to Andrades, developers don't need to worry about security or operational processes. They can bring their pipelines and repositories with them and run their workloads seamlessly while Rancher and Aqua guardrail the security controls.

Andrades and the team appreciate that they don't need to worry about the underlying infrastructure. If a problem occurs, they receive a notification. If they want to check the clusters' status quickly, they can check the dashboard to ensure that everything is "green." They no longer have to keep checking performance, workload status, or resource usage — Rancher removes the manual burden. This, Andrades believes, has freed teams to think more creatively.

Over the last year, the team has successfully migrated four key applications and is now managing them in clusters via the Rancher platform. This success has prompted the team to extend its use of the Rancher platform and double the number of nodes running in the cloud.

Cultural transformation

In addition to leading the technical transformation, Andrades is responsible for managing the cultural shift among Schneider Electric's development teams that a move to containers and the cloud requires.

For some who have been working in technology for the last couple of decades, a shift to a cloud-native existence is a big one. Long-ingrained development methodologies, baked into the fabric of the infrastructure, are as hard to modernize as the technology itself — particularly when it may appear that the technology is replacing significant parts of the job.

Andrades's focus, therefore, is to excite and galvanize the company around the opportunity every developer has to build new, disruptive skills. The range of experience spans experts through to complete novices. His mission is to globalize the existing pockets of excellence by bringing the company together to hear their stories and take a closer look at how they're succeeding with Kubernetes. By sharing detailed technical expertise and best practice, along with a sense of long-term value, Andrades and his team believe they'll carry the business along the journey with them.

What's next for Schneider Electric?

Schneider Electric's relationship with Rancher looks set to continue to grow in the future. The team recently renewed its support contract and has doubled its usage of the Rancher platform. This deepening of the relationship illustrates the confidence that Andrades and his team have in the platform, the support they receive, and the long-term value the alliance will bring to Schneider Electric, its customers, and the wider European energy sector.

The benefits include the following:

- » Reduced deployment and management time with automation.
- » Improved security posture with Aqua integration, RBAC, and NaaS.
- » Increased rate of innovation.
- » Established a growing business case for Kubernetes in the European energy sector.

Continental

Continental develops pioneering technologies and services for the sustainable and connected mobility of people and their goods. Founded in 1871, it offers safe, efficient, intelligent, and affordable solutions for vehicles, machines, traffic, and transportation. Headquartered in the German city of Hanover, Continental has grown exponentially during the last 150 years to become a global brand. The manufacturing giant is now present in 585 locations in 59 countries and markets and has 241,000 employees worldwide.

Continental's manufacturing infrastructure team exists to capitalize on the most disruptive technologies, serve application teams better, and drive innovation. With 12 years of experience working as their Infrastructure Team Lead, Roland Paxián took a long-term and global view of technology innovation. Three main objectives drive his strategy — creating efficiency, maintaining the highest quality standards, and achieving these goals in a systematic and networked way. Paxián believes this strategy helps Continental transform novel ideas into mass production faster. According to Paxián, “Manufacturing processes are under scrutiny as companies seek to gain market share through digital transformation. The process of modernizing machinery in harmony with operational software is not a simple one. Adopting a microservices-based containerization strategy removes some of the complexity.”

Seeing why Continental chose containers

Continental has always been a forward-thinking organization and, naturally, digital transformation and modernization are a major focus. However, in manufacturing, the practical implications of modernizing the now decades-old architecture — estates of legacy machinery and all the software that runs within them — cannot be underestimated.

Continental had run a virtualized infrastructure for many years, which suited them well, but management and maintenance became problematic over time. If teams wanted to implement a new feature or upgrade an application, this would be time and resource-intensive.

When containers emerged six years ago, the team saw an opportunity to streamline infrastructure management and started to investigate. It took a few years for a serious discussion to begin around the value of containerization, but when it did, it wasn't long before Kubernetes was identified as the most flexible way to get the containerization strategy off the ground.

In 2017, plans started to crystallize. At that time, the primary consideration was whether to move applications to the cloud or remain in the data center. Without a doubt, running Kubernetes in the cloud would be relatively simple; it was easy to spin up clusters in AWS and Azure. However, it became clear that latency would be an issue where some critical applications were concerned. Factory machinery requires millisecond response times, so some systems would need to remain in the data center. As such, Continental needed a hybrid cloud and on-premises methodology.

If the team wanted to use Kubernetes on-premises and in the cloud, it would need to engineer and support its own solution — something that would have taken time. Kubernetes offered the right container orchestration methodology, but Paxián and his team needed a way to run multiple clouds and on-premises deployments side-by-side in one platform. That's where Rancher came in.

After a short PoC in late 2018, which saw the team evaluate several Kubernetes management options, Rancher emerged as the most suitable platform to help modernize and unify Continental's legion of manufacturing applications. Since Rancher was formally selected in early 2020, growing demand has come from Continental's many manufacturing teams. Paxián's focus is shifting to make the platform safely available to hundreds of application development teams across the world.

Understanding the problems Continental is solving

Continental faced two challenges. First, it wanted to develop the infrastructure platform, and then it tried to roll out the service to hundreds of teams of developers in 45 locations worldwide. Partnering with Rancher helped the company achieve these objectives.

Legacy transformation

The primary driver for adopting a cloud-native, container-centric strategy was the urgent need to transform Continental's manufacturing infrastructure into an agile, cloud-native, and platform-based architecture. It had to be heterogeneous — flexible enough to run on-premises and cloud workloads together with any vendor, via a central UI.

For Continental's application developers, the change couldn't come soon enough. Application deployment and maintenance had become resource-intensive over the years. Everything was handled manually, from design to build to deployment and management — and this rigorous process repeated for each new development. The infrastructure team would encounter a host of problems if it needed to implement a new feature or simply upgrade an application. If an application developer needed an environment to develop something new, it would take time to fulfill the request, which slowed the pace of innovation.

Importantly, many production lines run 24 hours a day, seven days a week. If a line needed to be upgraded or an issue resolved, taking it out of service would cost the company dearly. Paxián needed an environment that would allow him to develop and maintain manufacturing applications without affecting productivity.

Managing his Kubernetes-based infrastructure platform in Rancher, Paxián has created a highly agile and scalable application framework, which has removed complexity and significantly reduced management overheads. His new containerized architecture allows him to run applications in separate clusters, with development, test, and production environments already in place. If they need a place to spin up new containers to try new ideas, they can create them in minutes.

If an application needs to be updated, a feature needs to be added, or maintenance needs to be performed, this can be done using Rancher without halting production lines. The team no longer requires costly maintenance downtime during upgrade periods. Updates are centralized and installed in a couple of clicks, which has reduced the management burden and improved overall productivity. Paxián estimates management time has been reduced by 75 percent. Because the platform promotes a cloud-native approach to building and deploying new services, applications can be created as microservices. These microservices are

highly portable between on-premises and cloud environments, making resource allocation and scaling more predictable.

Creating a global infrastructure platform

Now that Paxián and his team have developed the infrastructure platform, their focus is on rolling out the service to hundreds of teams of developers in 45 locations worldwide. The project has progressed rapidly. Now thousands of developers can access the new containerized platform via a single pane of glass.

Of course, some applications are engineered for the cloud and others are engineered to reside on-premises, closer to production lines. Running in Rancher, the new infrastructure platform provides a consistent framework for application development while allowing teams to configure and secure them for specific conditions. It then allows teams to deploy to any environment and run these clusters side-by-side via the Rancher UI.

This has major benefits for distributed teams. Having a flexible approach allows teams to develop applications with the manufacturing use case clearly in mind and in compliance with local regulations. Teams can choose, for example, to use data centers in highly regulated regions or where processing must take place within the production lines themselves.

In just six months, the team has rolled out the platform to nine regions in Europe and Asia. Paxián believes this is critical for an organization like Continental with a global workforce. For the first time, teams that may be separated by geography and business unit can work together in a unified and consistent way. More important, they can do it safely, within a rules-based domain. By adopting a platform approach to infrastructure management, Paxián has created a scalable, agile framework where collaboration and cooperation can reign. This result would've been impossible before.

The significance of Paxián's strategy has been profound. By working together under a common methodology, projects are completed faster and developments are consistent and created according to defined rules. The platform is accessible 24 hours a day, seven days a week, with access tightly monitored in Rancher.

Looking ahead: Continental's long-term cost reductions

In manufacturing, it's common to find large and resource-intensive servers running next to shop floors. Designed for use with specific machinery, these servers are expensive to run and environmentally outdated. In the long term, by engineering manufacturing applications to be more cloud-native, the infrastructure team will reduce these costs by moving applications to the cloud and the data center.

Where compute resources are still needed in production lines, Internet of Things (IoT) solutions like K3s will allow the team to run lightweight versions of Kubernetes directly on machinery. Hardware transformation always takes time, but Paxián believes by putting the right infrastructure in place now, the path to wider transformation will be smoother.

The benefits include the following:

- » Reduced migration time by 80 percent.
- » Reduced management time by 75 percent.
- » Reduced upgrade time by 80 percent.
- » Reduced costs moving on-site server applications to data center/cloud.
- » Global management platform for 45 regions.

Municipal Property Assessment Corporation

The Municipal Property Assessment Corporation (MPAC) is the largest assessment jurisdiction in North America. It assesses more than five million properties in Ontario, worth \$3 trillion, in compliance with the Assessment Act and provincial regulations. The MPAC's property assessments are the foundation of Ontario's property tax system, which generates \$30 billion annually for municipalities to supply local services.

Every day, thousands of property owners access the organization's external property-valuation application, AboutMyProperty. Assessors use a workflow system to update information, and property owners use AboutMyProperty to view their property profiles, assessment information, and comparable properties in the area. With different valuations for industrial versus residential and commercial properties, for example, secure data processing and analysis at scale are top priorities.

For the cloud operations and infrastructure team, led by IT director Gopi Balasingam and senior infrastructure architects Chruz Cruz, David Zheng, and Ken Tam, security, resilience, and cost-efficiency are major preoccupations. As custodians of public data, the company must ensure that its technical infrastructure is modern and robust. This directive has hastened MPAC's journey to the cloud, Kubernetes, and Rancher. As Balasingam said, "As a government-funded organization, with a clear civic duty, we have a responsibility to choose the technologies that will drive great agility and the greatest efficiencies. That's why we work with Kubernetes and Rancher."

Seeing why MPAC chose containers

Like many companies, MPAC's infrastructure has run on-premises in the data center for many years. As cloud computing matured, the team decided to migrate its data center infrastructure to AWS. Moving to the cloud made sense in those early days when costs were low. Consequently, the aim was to migrate its estate of stand-alone machines as quickly as possible with little to no disruption.

In 2017, the team migrated its entire application ecosystem by lifting and shifting Spring Boot and Java applications from the old on-premises environment to the cloud, running on stand-alone compute instances. Doing this manually was time-consuming and onerous. Scaling and analytics, for example, required lots of manual intervention, which was highly inefficient. It also started to get expensive. On paper, 5 cents an hour didn't seem like much, but when the team added 300 and 400 hosts as stand-alone instances, the costs started to mount.

The team considered containerization as a method to streamline running workloads in the cloud and to reduce costs.

MPAC had an estate of Spring Boot applications that were easy to transplant into containers and then into Kubernetes. Some applications ran in stand-alone Docker containers, load-balanced with Elastic Load Balancing (ELB). Applications weren't self-healing, and the team had to write scripts to do rolling deployments. When comparing this methodology to running containers in Kubernetes, there was no contest. Kubernetes was much more agile and "self-aware," and it soon became the team's one-stop shop.

MPAC trialed several management options — experimenting initially with Mesosphere, Docker Swarm, Tectonic (CoreOS), and Rancher 1.6. With Rancher 2.0 some distance away, and with a need to move quickly, the team opted to work with Kubernetes Operations (kops). At that time, kops was the standard management tool for Amazon-related Kubernetes clusters, and it allowed MPAC to keep its data management systems in Canada — where Amazon Elastic Kubernetes Service (EKS) didn't yet have a presence.

Kops performed well, but before long, the team soon noticed it had gaps, particularly in upgrade management and maintenance functionality. Typically, upgrades and maintenance took three to five days, and if the team wanted to do a security patch, they were at the mercy of the kops release cycle. This was particularly problematic when it came to achieving ISO 27001 certification — the team needed an added layer of security to prove it was on top of patch management to meet certification requirements.

With the launch of Rancher 2.3, MPAC realized many of these issues would be resolved. In February 2020, the team conducted a successful two-month PoC. They ran a small non-production environment and were so excited with the results that the project went into full production right away.

Understanding the problems MPAC is solving

MPAC uses Rancher to increase the efficiency of AWS workloads and applications and democratize its management of Kubernetes.

Achieving major cloud efficiencies

AWS has been an integral part of MPAC's infrastructure for more than eight years. By early 2017, the company had closed all on-premises and hosted data centers; the focus was to migrate to AWS at speed. MPAC deployed all workloads quickly in the AWS U.S. East region (Virginia). Then, with data residency concerns in mind, the company migrated its workloads from the AWS U.S. East region to the new AWS Canada (Central) Canadian regional service.

The team loved (and still loves) the ease of use, flexibility, and tooling inherent in AWS, but as time went on, Cruz and the team noticed costs were accumulating. At a macro level, costs looked low, but on closer analysis in Rancher, an accumulation over time of small over-subscriptions and over-resourcing resulted in a substantial monthly bill. Rancher brought a level of operational visibility to MPAC's AWS-based Kubernetes containers that allowed the team to monitor and identify inefficiencies — and take immediate action.

Suddenly, Cruz and the team could see what kind of resources were truly required to run the business and could scale this analysis down to individual applications. By taking the simple action to monitor individual processes and find tiny resource inefficiencies, the team estimates MPAC's monthly AWS bill has been reduced by 40 percent — a significant saving.

Transforming Kubernetes management

The team knew Rancher would enable a repeatable, predictable Kubernetes deployment strategy that could be supported collectively throughout the business. Senior Architect David Zheng knew Kubernetes inside and out but was the only one with in-depth knowledge — a burden on one person. Cruz wanted every team member to be able to manage MPAC's Kubernetes clusters, whether in a typical deployment scenario or during upgrade and patching cycles.

What Rancher has brought is a central, unified, and intuitive Kubernetes management methodology, which has democratized the use of containers across the business. For the first time, IT and development teams can work side by side, with full visibility of cluster performance, spinning up and tearing down new instances in minutes.

Whereas in kops, upgrades and maintenance took three to five days, in Rancher it now takes just a few hours. Upgrades can take place more regularly and patch management is no longer at the mercy of the kops release cycle. Why is this important? As a public service, MPAC's compliance relies on its systems being fully updated at all times. Overall, update and patch management times have been reduced by more than 80 percent. Finally, cluster deployment and scaling in Rancher is dramatically improved — with highly variable workloads, the team is now able to scale MPAC's five clusters from a few nodes to hundreds of nodes in minutes.

As a public organization, MPAC was also highly focused on security. To achieve ISO 27001, the team needed a reproducible artifact to prove the architecture met mean time to recovery (MTTR) requirements. Achieving an accurate reading in kops was difficult — too many nuances and issues arising along the way. For example, in kops, there was a requirement to hand off hard-coded access tokens, which could be shared among team members. A better access control method was needed, and Rancher brought this functionality. Automated RBAC has reduced complexity while adding a layer of security to the infrastructure.

Importantly, Rancher has improved MPAC's overall security posture. Rancher's built-in security features — CIS benchmarking, RBAC, monitoring and alerting capabilities — provide additional reassurance and help the team maintain compliance in line with its civic responsibilities.

Looking ahead: Freedom of choice

Having a technically agnostic environment will become increasingly important to MPAC. MPAC's Kubernetes landscape is a heterogeneous one. Currently, the team runs a Rancher Kubernetes

Engine (RKE) cluster, an EKS cluster imported into Rancher, and two AWS Linux clusters also imported into Rancher. Rancher gives MPAC the freedom to use EKS alongside RKE, Google Kubernetes Engine (GKE), and any other technology, for that matter. The team believes this agnostic, open-source approach will further boost innovation and drive even greater efficiencies.

The benefits include the following:

- » Reduced cloud costs by 40 percent.
- » Reduced cluster deployment by 85 percent.
- » Reduced update/patch management time by 80 percent.

IN THIS CHAPTER

- » Downloading Rancher products
- » Exploring the Rancher documentation and other training resources
- » Joining the Rancher community via Slack and YouTube
- » Earning Rancher and Kubernetes certification
- » Scheduling a demo with a Rancher and Kubernetes expert

Chapter 4

Ten Ways to Get Started with Kubernetes and Rancher

Rancher Labs (acquired by SUSE in 2020) provides plenty of resources to help you get started with Kubernetes and Rancher. This chapter highlights the resources available via Rancher's website and beyond, including the Rancher Slack, Rancher's YouTube channel, and free Cloud Native Computing Foundation (CNCF) webinars. Rancher also provides training and certification opportunities. The last section of this chapter shows how to arrange for a demo with a Rancher and Kubernetes expert.

Read the Rancher Documentation

Great open-source projects have great technical documentation, and Rancher is no different. The Rancher docs site is free for everyone to access and covers everything from Rancher 2.x,

to Rancher Kubernetes Engine (RKE), and K3s. Get started at <https://rancher.com/docs>.

Download and Install Rancher

To install Rancher in any Kubernetes cluster, follow the Rancher quick-start guide at <https://rancher.com/quick-start>.

After installing Rancher, you can access the Rancher user interface (UI) by opening a browser and going to the hostname or address where you installed Rancher. The UI will guide you through setting up your first cluster.

Download and Install K3s

Advanced users may want to deploy a lightweight Kubernetes distribution on low-resource hardware. Rancher Labs pioneered the development of K3s for this purpose. K3s was donated to the CNCF in August 2020. To install K3s, follow the simple quick-start guide at <https://k3s.io> or read the technical documentation at <https://rancher.com/docs/k3s/latest/en/>.

Install Longhorn Using Rancher

Storage shouldn't be complicated. With Longhorn it's not. Longhorn gives your teams a reliable, simple, and easy-to-operate persistent storage solution for any Kubernetes cluster. Deployed with a single click from the Rancher application catalog, Longhorn provides you the ability to secure, provision, and back up your storage across any Kubernetes cluster.

Like K3s, Longhorn started life as a Rancher Labs project but was donated to the CNCF in October 2019. To find out more about Longhorn, visit <https://longhorn.io> or read the Longhorn technical documentation (<https://longhorn.io/docs>).

Join the Rancher Users Slack Community

Join the Rancher Users Slack to ask questions and learn from other community members, share your experiences using Rancher, and stay up to date on future events and training sessions <https://slack.rancher.io>.

Check Out the Rancher YouTube Channel

If you're starting out and don't have much time to read documentation, check out Rancher's huge library of recorded meetups, the Master Classes, Office Hours videos, and introductory training on Rancher's YouTube channel at <http://youtube.com/c/rancher>.

Register for the Rancher Academy

Earning Rancher certification through the Rancher Academy helps Kubernetes practitioners demonstrate their knowledge and competence with Kubernetes and Rancher. The Rancher Academy is designed to ensure that members of the SUSE Rancher community have the most relevant and up-to-date skills on the industry's most widely adopted Kubernetes management platform. Earning Rancher certification empowers graduates to be at the forefront of the cloud-native way of doing business, which is agile, open-source oriented, and maniacally focused on fast access to innovation.

Enroll for free to earn a highly valued Certified Rancher Operator: Level 1 qualification at <https://academy.rancher.com>.

Sign Up for a Cloud Native Computing Foundation Webinar

The CNCF regularly runs free webinars covering the following topics:

- » Application and development
- » Continuous integration/continuous delivery (CI/CD)
- » Customizing and extending Kubernetes
- » Machine learning and data
- » Observability
- » Security identity and policy
- » Serverless
- » Service mesh
- » Storage

Sign up for the next CNCF webinar at www.cncf.io/webinars/.

Attend a Rancher Rodeo

Rancher Rodeos are free, in-depth online workshops designed to give DevOps and IT teams the hands-on skills they need to deploy and manage Kubernetes everywhere.

During these virtual hands-on workshops, Rancher's technical experts introduce Rancher, Docker, and Kubernetes, and walk through the steps for deploying a Kubernetes cluster.

Sign up for a Rancher Rodeo at <https://rancher.com/rodeos/>.

Schedule a Demo

Discuss your requirements with a Rancher and Kubernetes expert by completing the online form at <https://rancher.com/request-a-demo>.

Compare the market-leading Kubernetes management platforms including Rancher

Forward-thinking organizations are adopting platforms like Rancher to help streamline the deployment and management of their Kubernetes clusters.

In this latest Buyer's Guide, we compare Rancher with the three competitive Kubernetes management platforms: Red Hat OpenShift, VMware Tanzu, and Google Anthos.

Some of the capabilities we review include:

- Ease of install, configuration & maintenance
- Multi-cloud and multi-cluster support
- Security, policy and user management
- Storage and edge support



DOWNLOAD THIS FREE GUIDE AT: [RANCHER.COM/BUYERSGUIDE](https://rancher.com/buyersguide)

Brought to you by  **SUSE**

Build an enterprise-grade Kubernetes environment

As enterprise applications become more complex, development and operations teams need a tool to orchestrate that complexity. Kubernetes is that tool, allowing enterprises to deploy, scale, and manage containerized applications anywhere. But aligning your business to take full advantage of Kubernetes requires careful consideration. This handy guide walks you through that process, from evaluating where your company stands now, to what to look for when selecting a Kubernetes management platform, to real-world examples of how Kubernetes can drive innovation from core to cloud to edge.

Inside...

- Assess your progress on the Kubernetes journey
- Identify the best Kubernetes management platform for your use case
- Learn from real-world examples of what other companies have done
- Understand how Rancher can help you become more agile and competitive



Tom Callway is Global Director of Product Marketing. **Peter Smails** is Vice President, Product & Solutions Marketing. **Caroline Tarbett** is head of Customer Marketing. **Shannon Williams** is COO of the Global Customer Organization.

Cover Image: © dan_prat / Getty Images

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-78165-3
Not For Resale

for
dummies[®]
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.