


# DIY or Managed?

## Understanding the True Cost of DIY Kubernetes



Example Report

Summer, 2021

## Table of Contents

Overview	1
Executive Summary	1
COMPANY Goals and Requirements	3
TCO Comparison	3
Initial Kubernetes Deployment	4
Ongoing Kubernetes Management	5
Revenue Impact of Time-to-Market and Customer Onboarding	6
Appendix A — Key Capabilities Needed for Production-Ready Kubernetes	8
Appendix B — Additional Assumptions	9

## Overview

This TCO report was prepared for a recent Platform9 prospect (referred to as COMPANY) to help them understand the costs and other factors involved in managing a Kubernetes deployment internally (DIY) versus using the Platform9 Managed Kubernetes (PMK) service. It outlines their business goals and requirements and estimates their TCO for DIY and managed scenarios looking at both the initial deployment and ongoing management phases.

We use a proprietary TCO calculator which has proved accurate with many customers. The report has been compiled based on information shared by COMPANY with Platform9 to date about the deployment size and current resourcing as well as additional assumptions made by Platform9 to provide a complete cost assessment. These additional assumptions are a mixture of COMPANY-specific assumptions as well as working assumptions informed by Platform9 experience across many enterprises going through similar decision processes. The assumptions are included in the report and its appendices.

Platform9 envisages updating key elements of the TCO modeling based on further engagement with the COMPANY team on their requirements.

## Executive Summary

Companies beginning to scale their Kubernetes production deployments face a choice between pursuing an internal development approach (termed DIY in the rest of the report) or a managed-service option. At smaller deployment scales, managed services typically provide the best TCO, while the largest deployments can be more economical using a DIY approach. This pattern is driven by the large fixed investments required to create and manage robust production environments with 24/7 coverage under the DIY approach.

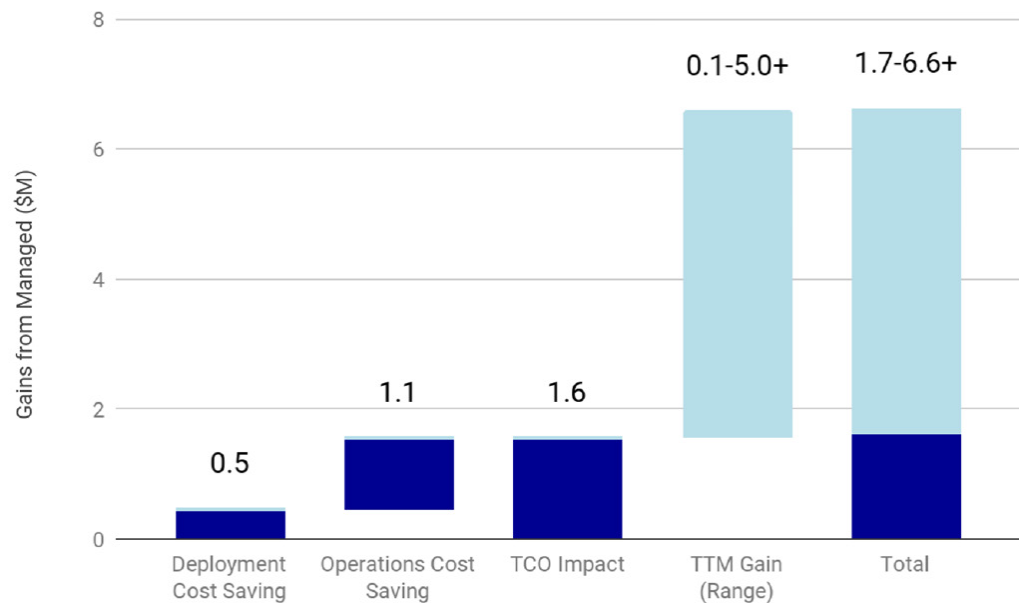
- Based on COMPANY's current deployment scale and our growth estimates, we find that the DIY approach to Kubernetes would increase COMPANY's costs and slow time-to-market, which could negatively impact revenue by slowing new product and customer rollouts.
- We estimate that COMPANY can reduce TCO by 40% or \$1.6M over the next three years by lowering setup costs 88% and ongoing management costs 32% using a managed approach with Platform9 rather than DIY. These reductions are achieved by using an existing, robust, managed Kubernetes platform proven in production by leading enterprises. This also lets COMPANY focus internal engineering resources on revenue-generating applications instead of infrastructure maintenance.

***We estimate that COMPANY can reduce TCO by 40% or \$1.6M over the next three years by lowering setup costs 88% and ongoing management costs 32% using a managed approach with Platform9 rather than DIY.***

- We find that COMPANY can also reduce time-to-market up to 79% using a managed approach instead of DIY. The time-to-market gains are possible by avoiding the complex infrastructure buildout required by a complete Kubernetes platform. We provide a range of sensitivities for the impact on profitability of this time-to-market advantage. For example, at a \$10M revenue run rate for products on the platform, a nine-month advantage would provide a \$5M one-time profitability benefit.

**COMPANY can also reduce time-to-market up to 79% using a managed approach instead of DIY.**

The chart below shows the TCO benefits (dark areas) as well as a range of the profitability impacts from time-to-market gains (lighter to emphasize the greater uncertainty at this stage) to COMPANY. Time-to-market and revenue considerations, while more intangible, can be the largest drivers of the overall economic impact of a DIY vs managed solution.



While not part of the TCO analysis, we also note that COMPANY is likely to improve quality-of-service through a managed approach that provides access to greater Kubernetes management capabilities as well as a large team of certified professionals who are 100% dedicated to managed Kubernetes support. COMPANY can also avoid the challenging task of recruiting and retaining a large Kubernetes team.

## COMPANY Goals and Deployment Assumptions

Business goals include:

- Support revenue growth for products running on the Kubernetes platform including COMPANY PRODUCT NAMES and potential future products.
- Improve the customer experience by giving customers greater flexibility and ease of access with deploying workloads.
- Focus the team on delivering COMPANY's unique value and minimizing effort on non-strategic capabilities.
- Maximize team productivity.
- Find the lowest TCO.

Requirements include:

- Support for five POPs initially with 20 nodes per POP (approximately 100 nodes initially). We assume that business growth at COMPANY will lead to increases in the number of POPs (we assume they increase by two per year) and a 25% annual increase in nodes per POP. We project 100 nodes in year 1, 175 in year 2, and 281 in year 3.
- Automated deployment: Today, if customers want to download and use COMPANY capabilities, they have to engage the COMPANY team and this may require on-site support. An automated experience would improve this situation.

***COMPANY is likely to improve quality-of-service through a managed approach that provides access to greater Kubernetes management capabilities as well as a large team of certified professionals who are 100% dedicated to managed Kubernetes support.***

## TCO Comparison

This evaluation compares a number of factors involved in deploying and running Kubernetes including:

- Costs to operate a Kubernetes deployment including expenses for people effort and software/systems.
- Time-to-market for a basic production to operate a Kubernetes deployment.
- Costs to enable new customers.
- Time-to-market for new customers.

We project that COMPANY, from a Kubernetes operations perspective alone, would save \$1.6M over the first three years by pursuing a fully managed Kubernetes solution instead of a DIY approach.

*The ongoing management of a Kubernetes deployment is a significant task. Platform9's combination of technology and support transfers the burden from internal teams to Platform9.*

	Setup	Year 1	Year 2	Year 3	Total
<b>DIY Cost Summary</b>					
Build phase	\$532,583				\$532,583
Ongoing costs		\$854,900	\$1,025,880	\$1,231,056	\$3,111,836
Ongoing costs for additional locations			\$82,070	\$180,555	\$262,625
Totals	\$532,583	\$854,900	\$1,107,950	\$1,411,611	\$3,907,045
<b>Platform9 Managed Cost Summary</b>					
License		\$440,000	\$731,500	\$1,112,760	\$2,284,260
Internal labor for setup	\$20,750				\$20,750
Quick start	\$25,000				\$25,000
Training	\$20,000				\$20,000
Totals	\$65,750	\$440,000	\$731,500	\$1,112,760	\$2,350,010
Savings vs. DIY	\$466,833	\$414,900	\$376,450	\$298,851	\$1,557,035
TCO reduction					<b>40%</b>

### Initial Kubernetes Deployment

Based on significant experience, the following summary table compares how many person-months it's likely to take and the labor costs to roll out the COMPANY Kubernetes deployment.

	DIY	Platform9	% Change
Time for rollout	10	2	-79%
FTE needed during setup	8	1.5	-81%
Person-months	77	3	-96%
<b>Labor cost for setup</b>	<b>\$532,583</b>	<b>\$20,750</b>	<b>-96%</b>

Key areas of time investment are estimated below. Each capability is defined in Appendix A.

Key Features for Establishing Production-Ready Kubernetes	Initial Cost (# Person Months)
Provisioning of clusters	6
Monitoring	5
Upgrading	6
Security — compliance	9
Security — RBAC	2
Security — authn/z	2
HA and healing	8
Load balancing	4
Site operations and reliability (client)	3
Private registry support	6
Multi-cluster management (e.g., AWS, GKE, Azure, bare metal)	18
Multi-cluster management	8
<b>Total # of person months</b>	<b>77</b>

We understand COMPANY would use a team of approximately eight people to support the production rollout requiring a ten-month period. On the other hand, supporting a Platform9 production deployment would take approximately three person months from the COMPANY team with an average of 1.5 FTEs needed for support from the COMPANY team for two months.

### Ongoing Kubernetes Management

The ongoing management of a Kubernetes deployment is a significant task. Platform9's combination of technology and support transfers the burden from internal teams to Platform9. Platform9 serves a large base of Kubernetes users and its team is able to use extensive automation and a shared team of Certified Kubernetes Administrators (CKAs) to deliver a production-ready environment that has both equivalent or higher service levels and lower cost than customers could achieve on their own.

In the following table, we estimate that COMPANY would require a team of at least 10.3 FTEs to support its Kubernetes environment in year 1. Using Platform9, these resources would be free to work on other activities. Task definitions are provided in Appendix A.

**Platform9 has invested over 150 person-years in developing the IP needed to support a managed Kubernetes service. This exceeds that of most Kubernetes users, and Platform9 continues to invest in improvements through a large engineering team.**

Ongoing Staffing Projections in FTEs: Year 1	
Provisioning of K8s clusters	0.75
Monitoring	0.45
Upgrading	0.50
Security — compliance	1.00
Security — RBAC	0.10
Security — authn/z	0.10
HA and healing	0.50
Load balancing	0.25
Site operations and reliability	4.00
Private registry support	0.15
Heterogeneous environment support (e.g., AWS, GKE, Azure, bare metal)	2.00
Multi-cluster management	0.50
<b>Total FTE</b>	<b>10.30</b>

After year 1, additional staffing would be needed to support growth in the environment. Below is the projected growth in FTEs:

	Year 1	Year 2	Year 3
FTEs in DIY Scenario	10.3	14.4	20.2

In addition to the direct costs of deploying and operating Kubernetes, there are several additional considerations in evaluating a DIY vs. Platform9 path:

- Improved service through a focused quality-of-service — Managed service providers can provide superior service because they are completely focused on one thing. If internal teams find their effort spread across multiple priorities, quality-of-service and its improvement rate can suffer.

***Platform9 enables simpler deployments into new locations compared to a manual setup and fully automated deployments after locations have been brought online.***

- Improved service through cumulative experience — Platform9 ranks among the top companies in the world in terms of cumulative experience in managing Kubernetes. Platform9 has operated Kubernetes environments exceeding more than 6000 nodes and over 475 regions providing Platform9 a broad knowledge base of solutions to common and less-common Kubernetes challenges. Companies new to Kubernetes would find this level of expertise difficult if not impossible to replicate.
- IP development on Kubernetes — While Kubernetes is fast becoming a universal control plane, it is not a plug-and-play solution. Components require skilled integration to provide a superior operator and developer experience. Platform9 has invested over 150 person-years in developing the IP needed to support a managed Kubernetes service. This exceeds that of most Kubernetes users, and Platform9 continues to invest in improvements through a large engineering team.
- Talent scarcity — There has been an extremely tight market for Kubernetes talent. Platform9 has found that companies outside a small number of predominantly consumer internet companies (e.g., Uber, Lyft) struggle to recruit and retain a full team of leading-edge Kubernetes practitioners. Using a managed service provides access to top talent without the associated cost of managing talent. In addition, if talent cannot immediately be recruited, the deployment timeline will also be delayed.
- Security — Platform9 places a major emphasis on ensuring operational security when delivering Kubernetes as SaaS. Platform9's solutions keep enterprise infrastructure and workloads secure. We have designed numerous security measures within our infrastructure and product (for example, automatic security patching), as well as security options available to administrators, to ensure Platform9 functions as a secure platform to manage your infrastructure. This level of dedicated security investment is above what many Kubernetes customers can support.

The challenges of creating a DIY deployment have led many observers to suggest that most companies should pursue a managed solution.

## **Revenue Impact of Time-to-Market and Customer Onboarding**

In this section we provide a brief assessment of how accelerated Kubernetes service deployment times at an overall and individual customer-service level can improve COMPANY revenue expectations and customer experience.

### **Accelerated Deployment of Revenue-Generating Services**

Customers seeking to launch new capabilities or services often find that time-to-market and revenue considerations become more important than optimization of Kubernetes costs. To provide a guide to assess these impacts for COMPANY's project, we include below a sensitivity table. It shows the profitability impact (based on gross margin) given ranges for the revenue impact of the services and how much faster a service launch could take place. For example, at a \$10M revenue run rate, a nine-month advantage would provide a \$5M one-time profitability benefit. A range of gains from \$0.1M – \$5M appears in the chart in the Executive Summary.

Profitability Impact of Accelerated Deployment (gross margin basis)						
Annualized Revenue						
		\$1M	\$10M	\$100M	\$250M	\$500M
Acceleration of Deployment	1 month	0.1	0.6	5.6	14.0	27.9
	3 months	0.2	1.7	16.8	41.9	83.8
	6 months	0.3	3.4	33.5	83.8	167.6
	9 months	0.5	5.0	50.3	125.7	251.4
	12 months	0.7	6.7	67.1	167.6	335.3
Note: Assumes gross margin of NN% based on COMPANY 10-Q for the three months ended DATE.						

### Automation of New Customer Onboarding

COMPANY wishes to improve the customer experience for new customers onboarding to the platform through automated deployment. Platform9 enables simpler deployments into new locations compared to a manual setup and fully automated deployments after locations have been brought online. A fully automated platform also provides the ability to accelerate customer PoCs and shorten COMPANY enterprise sales cycles.





## Appendix A — Key Capabilities Needed for Production-Ready Kubernetes

- Provisioning of K8s clusters — Standing up of clusters needs to be carried out using config mgmt tool like Puppet, Chef. Cluster stand up will need to be carried out for each supported environment.
- Monitoring — Monitoring capabilities for Kubernetes infrastructure needs to be built to check for downtime. Processes need to be put in so that output from monitoring activities can be understood and acted upon by support staff.
- Upgrading — Orchestration code for performing rolling upgrades for production clusters without incurring API or service downtime.
- Upgrade schedules need to be planned carefully to keep pace with Kubernetes versioning while ensuring roll out of security features appropriately.
- Security/compliance — Includes cost of understanding and patching the system.
- Security/RBAC — Role-based access control needs to be implemented.
- Security/authn/z — Authentication and authorization needs to be carried out ideally by integrating with identity management systems (e.g. Okta).
- HA and healing — Both hardware systems and Kubernetes clusters need to be highly available. Faulty HW/VMs need to be replaced without incurring downtime using a VIP service.
- Load balancing — Application-level load balancing needs to be established (one-time cost to implement). Recurring costs for enhancement and maintenance need to be taken into account.
- Site operations and reliability — Operations support rotations and tracking of key health metrics.
- Private registry support — Support for deploying Docker private registry on-prem w/ upgrade and update support.
- Heterogeneous environment support — Integration with multiple infrastructures including bare metal , VMware/KVM, and public clouds needs to be supported. Support for each cloud provider requires a different integration.
- Multi-cluster management — Support for deploying multi-cluster environments such as application catalogs.

## Appendix B — Additional Assumptions

<b>Initial Deployment Size</b>	
Number of locations	5
Average number of nodes per POP	20
<b>Rate of Increase</b>	
Number of nodes annual rate of increase per existing location	25.0%
Variable labor impact of node increase	80%
Number of new locations added per year	2
Percentage FTE increase added per new location	20.0%
<b>Basic Investment Inputs</b>	
Fully burdened cost of one engineer for one year	\$ 83,000
<b>Platform9 Investment Inputs</b>	
Platform9 cost per mode (annual) — year 1	\$ 4,400
Platform9 cost per mode (annual) — year 2	\$ 4,180
Platform9 cost per mode (annual) — year 3	\$ 3,960
Maintenance % of software license	0.0%
Implementation quick start	\$ 25,000
Annual training	\$ 20,000
<b>Platform9 Investment Calculations</b>	
Platform9 license fees - Year 1	\$ 440,000
Platform9 maintenance fees (annual)	\$ -
Platform9 implementation services	\$ 25,000