FORRESTER®

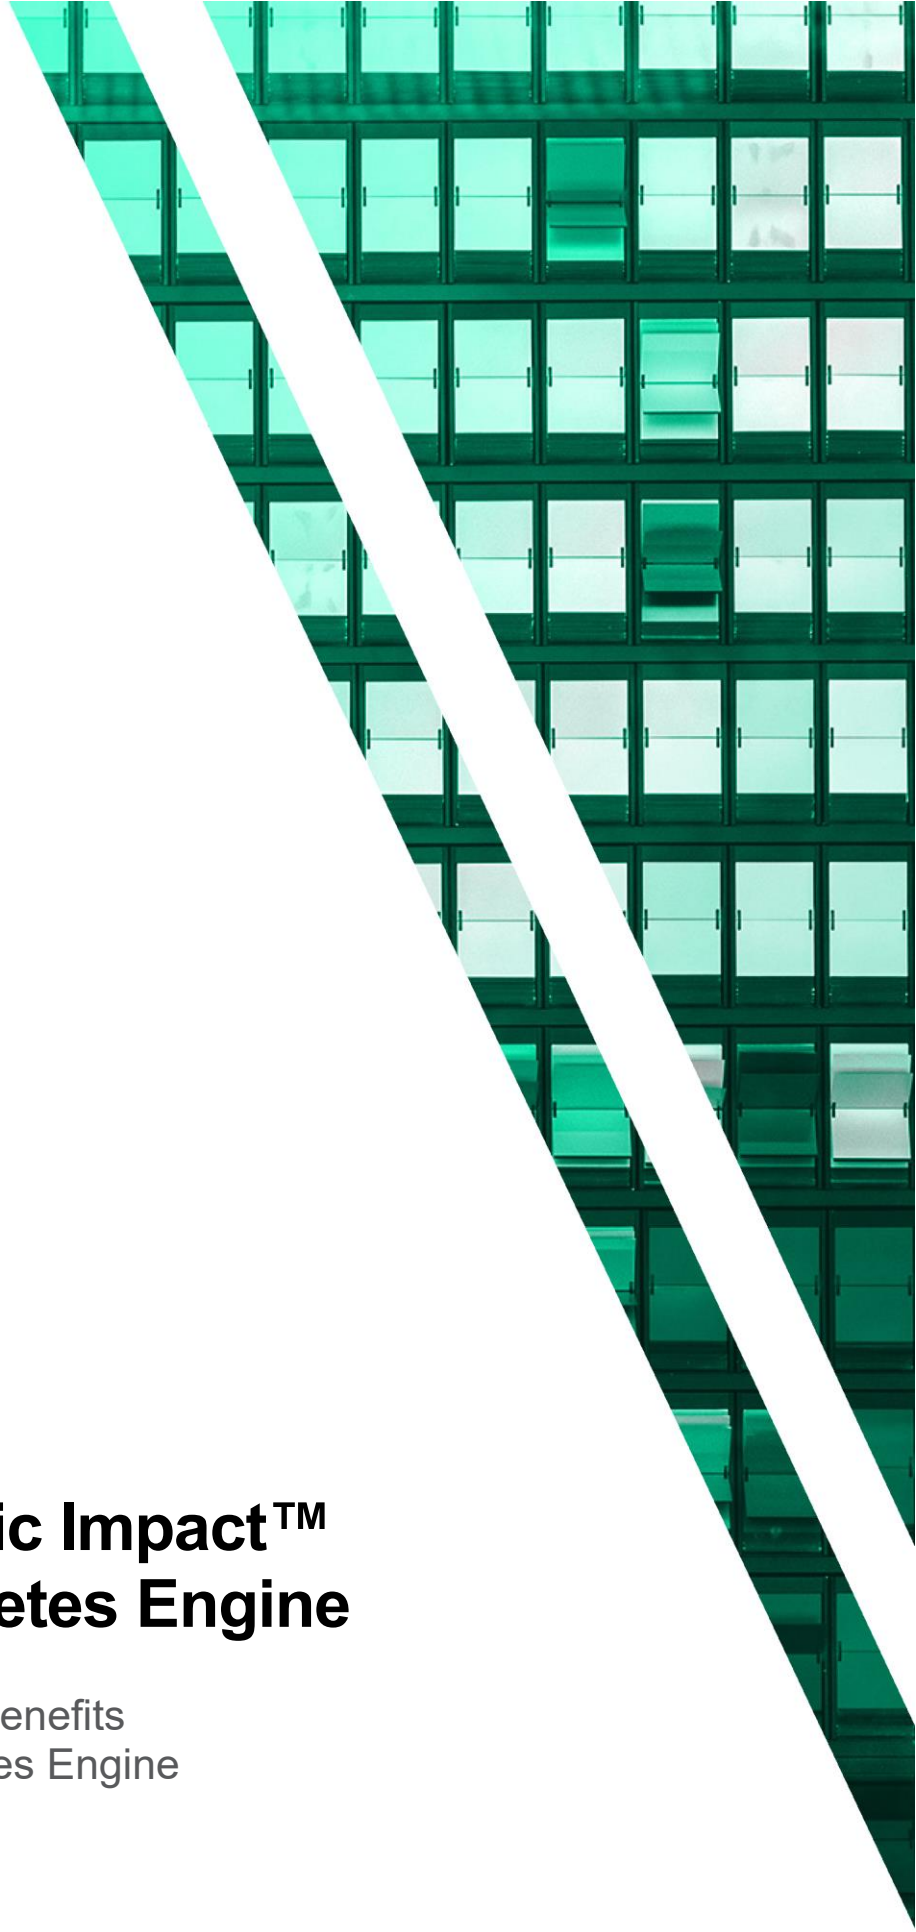

The Total Economic Impact™ Of Google Kubernetes Engine

Cost Savings And Business Benefits Enabled By Google Kubernetes Engine

July 2021

Table Of Contents

Consulting Team: Sam Conway
Sam Sexton

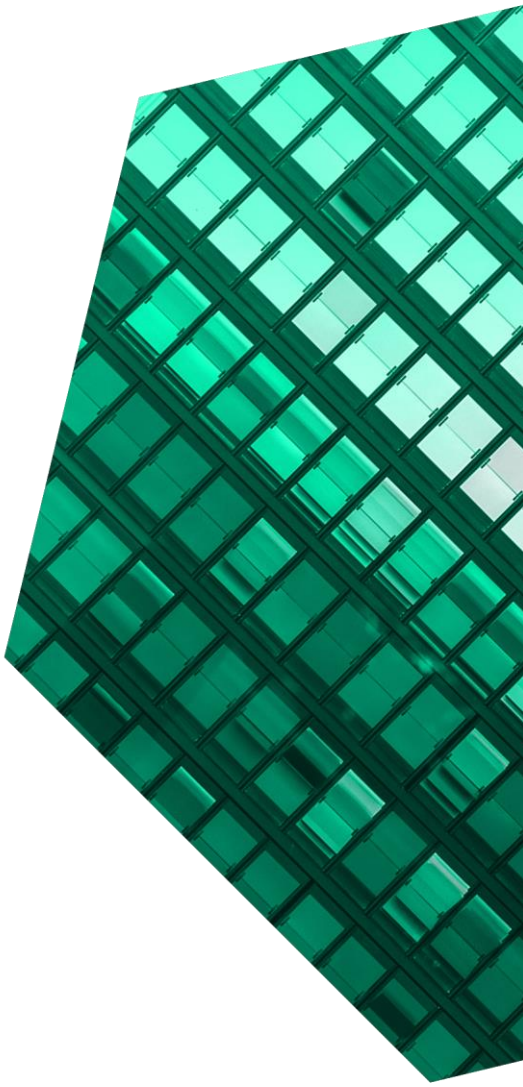

ABOUT FORRESTER CONSULTING

Forrester Consulting provides independent and objective research-based consulting to help leaders succeed in their organizations. For more information, visit forrester.com/consulting.

Executive Summary

Software containers are a versatile tool that enable organizations to accelerate innovation and improve efficiency across a wide breadth of use cases. However, infrastructure and operations professionals are still searching for the optimal method to implement these containers at scale. Organizations get bogged down troubleshooting infrastructure issues related to enterprise-scale container deployment, negating their ability to recognize the benefits of innovation initiatives.

Google Kubernetes Engine (GKE) is an enterprise-grade platform that is used to build, deploy, and run a wide variety of containerized applications. GKE provides secured and stress-free management of containerized applications, allowing organizations to focus on business innovation rather than troubleshooting infrastructure issues.

Google commissioned Forrester Consulting to conduct a Total Economic Impact™ (TEI) study and examine the potential return on investment (ROI) enterprises may realize by deploying GKE.[1] The purpose of this study is to provide readers with a framework to evaluate the potential financial impact of GKE on their organizations.

To better understand the benefits, costs, and risks associated with this investment, Forrester interviewed four customers with experience using GKE. For the purposes of this study, Forrester aggregated the experiences of the interviewed customers and combined the results into a single composite organization.

Prior to using GKE, the customers managed their own container deployments in multiple locations (on-premises and in the cloud) in complex and capital-intensive manners. However, this yielded limited success, causing customers to focus on infrastructure troubleshooting, limit their ability to scale container deployment, and ultimately slow innovation.

KEY STATISTICS

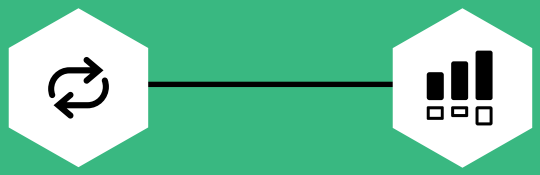

Return on investment (ROI)
182%

Net present value (NPV)
$19.4 million

After the investment in GKE, the customers modernized their application infrastructure and simplified its deployment and management. This resulted in time and cost savings, improved performance, enhanced agility, and accelerated innovation.

Organizations experienced benefits in three main areas: 1) greater productivity, 2) lower costs, and 3) reduced risk.

"GKE is giving us that balance between scalability and cost."
Lead data engineer, life sciences

KEY FINDINGS

Quantified benefits. Risk-adjusted present value (PV) quantified benefits include:

- **Reduced initial environment setup labor requirements by 35%.** GKE drastically reduces the labor required for initial environment setup Kubernetes tasks through automation and smart default features. The composite organization requires 35% less effort to conduct Day 0/1 tasks than it did when managing its own on-premises Kubernetes deployment.
- **Reduced ongoing cluster management time by 75%.** GKE manages much of the ongoing cluster lifecycle activities, including monitoring, tuning, right-sizing, and auto-repairing. Having these operational tasks managed through GKE reduces the time the composite organization spends on ongoing cluster management and optimization by 75%.
- **Improved developer efficiency by 10% to 30%.** GKE frees developers from development and configuration activities, improving their productivity. Additionally, using GKE in the cloud means the organization can quickly provision resources to development teams that would normally wait for on-premises resource provisioning. The composite organization experiences developer productivity improvement of 10% in Year 1, growing to 30% by Year 3.

"We really like the fact that Google — to us and to a lot of people — is considered an engineering company rather than a sales company. To them, it's about technology, the right solutions, and listening to the customers. We really like that."
Global head of cloud, financial services

- **Reduced infrastructure costs by 75%.** With GKE, organizations can utilize autoscaling to maximize infrastructure utilization and reduce costs. Furthermore, preemptible virtual machines (VMs) provide a cost-effective manner for organizations to complete compute-intensive batch work at a lower cost. The composite organization uses GKE's autoscaling capabilities and preemptible VMs to reduce ongoing infrastructure spend by 75%.
- **Eliminated existing container-based PaaS license spend.** GKE provides organizations with a suite of fully managed container management and orchestration capabilities, eliminating the need to maintain existing PaaS licenses. The composite organization eliminates its third-party container-based PaaS licenses, which are valued at $1,500 per application instance.
- **Improved security productivity by 80%.** Google scans containers to discover vulnerabilities or missing patches in Google-managed containers and, if possible, automatically begins the patching and release process. Organizations leverage these capabilities to reduce the time and effort patching teams spend. The composite organization improves the efficiency of its patching team by 80% with GKE.
- **Avoided losing revenue with 97% improvement in availability.** GKE's autoscaling ensures that organizations have adequate resources to meet compute demands, especially during variable peak periods. Additionally, Google site reliability engineers (SREs) monitor clusters and the supporting compute, network, and storage resources. With GKE, organizations have better availability and experience fewer degradation events than when managing their own Kubernetes deployment, reducing the risk of downtime and, therefore, lost revenue. The composite organization experiences a 97%

improvement in availability, resulting in additional revenue.

Unquantified benefits. Benefits that are not quantified for this study include:

- **Improved customer experience (CX).** Excellent customer experience demands rock-solid app performance, even during high demand. With GKE's cluster autoscaler, nodes are added to the node pool when demand is high to ensure availability. Additionally, Google SREs monitor GKE clusters to ensure they are accessible. By limiting downtime and optimizing performance of underlying application infrastructure, organizations can provide a better CX.
- **Hardened security profile and quick app fixes.** GKE is secure by default with vulnerability scanning and data encryption. It empowers developers to accelerate release velocity, enabling more frequent patching and eliminating vulnerabilities.
- **Enhanced employee experience (EX).** Automation of repetitive manual tasks allows organizations to redeploy employees to more value-add work. GKE also allows developers to focus on development instead of troubleshooting infrastructure issues.
- **Greater scalability and agility to expand into new markets and regions.** With GKE, organizations can quickly spin up a cluster and have a production-ready environment for teams in a matter of hours. This is a significant improvement on traditional provisioning requests, which may require a lengthy approval process and coordination across global teams.

Costs. Risk-adjusted PV costs include:

- **GKE annual cluster management fees.** Costs include usage-based fees for GKE clusters.
- **Implementation and ongoing labor.** Costs include professional services, training, and ongoing platform support.

The customer interviews and financial analysis found that a composite organization experiences benefits of $30,010,749 over three years versus costs of $10,647,758 adding up to a net present value (NPV) of $19,362,991 and an ROI of 182%.

> **"GKE is so low maintenance for us when adding new services, combining services, or deprecating them. It's a very, very fast exercise and means that the large cycle times of trying and doing different stuff is daily — not weekly, monthly, or quarterly."**
> *CTO, advertising*

ROI
182%

BENEFITS PV
$30 million

NPV
$19.4 million

PAYBACK
<6 months

Benefits (Three-Year)

GREATER PRODUCTIVITY
$18.7 million

LOWER COSTS
$17.1 million

REDUCED RISKS
$1.2 million

TEI FRAMEWORK AND METHODOLOGY

From the information provided in the interviews, Forrester constructed a Total Economic Impact™ framework for those organizations considering an investment in GKE.

The objective of the framework is to identify the cost, benefit, flexibility, and risk factors that affect the investment decision. Forrester took a multistep approach to evaluate the impact that GKE can have on an organization.

DUE DILIGENCE

Interviewed Google stakeholders and Forrester analysts to gather data relative to GKE.

CUSTOMER INTERVIEWS

Interviewed four decision-makers at organizations using GKE to obtain data with respect to costs, benefits, and risks.

COMPOSITE ORGANIZATION

Designed a composite organization based on characteristics of the interviewed organizations.

FINANCIAL MODEL FRAMEWORK

Constructed a financial model representative of the interviews using the TEI methodology and risk-adjusted the financial model based on issues and concerns of the interviewed organizations.

CASE STUDY

Employed four fundamental elements of TEI in modeling the investment impact: benefits, costs, flexibility, and risks. Given the increasing sophistication of ROI analyses related to IT investments, Forrester's TEI methodology provides a complete picture of the total economic impact of purchase decisions. Please see Appendix A for additional information on the TEI methodology.

The Google Kubernetes Engine Customer Journey

Drivers leading to the GKE investment

Interviewed Organizations

Industry	Region	Interviewee	Revenue	Employees
Life sciences	Global	Lead data engineer	$43 billion	100,000
Financial services	Global	Global head of cloud	$50 billion	240,000
Advertising	Global	CTO	$170 million	200
Financial services	Global	Software engineer	N/A	1,600

KEY CHALLENGES

Forrester interviewed four decision-makers at organizations using GKE. All organizations had some experience with Kubernetes, primarily managing their own on-premises container deployments. The interviewees' organizations struggled with common challenges, including:

- **Reaching the limits of scalability with legacy platforms.** Interviewees found that their organizations had exhausted the scale of Kubernetes they could deploy and manage themselves. Teams focused more on troubleshooting issues and ensuring availability than harnessing the benefits of containerization.

> **"GKE really reduces our operational burden with the way it's structured and some of the features that are available there."**
> *Lead data engineer, life sciences*

- **Pressure to reduce operating costs.** Interviewees highlighted organizational goals to continually reduce operating costs. While their organizations had already improved workload density with initial container deployments, the organizations sought further ways to improve resource utilization.
- **Desire to focus development and operations teams on more interesting and high-value work.** Interviewees wanted to focus on supporting developers and product innovation — not myopically focusing on configuration and repairs.
- **Difficulty flagging and responding to capacity and availability issues.** As adoption within their organizations grew, interviewees found that availability and degradation events persisted as teams could not keep up with fluctuating demand. Global organizations found that their teams were slow to react to developer requests and resorted to costly overprovisioning to avoid degradation.

> **"All of the telemetry is handled in Google, and we don't have to worry about it."**
> *Software engineer, financial services*

INVESTMENT OBJECTIVES

When evaluating potential solutions, the organizations focused on a provider that would:

- Reduce operations complexity and help consolidate development environments.

- Provide a platform requiring little effort to scale and manage.
- Improve the availability and performance of applications.
- Improve release frequency and time-to-market for development teams.
- Accommodate very large workloads (10,000+ nodes).
- Reduce infrastructure costs.

SOLUTION REQUIREMENTS

The interviewees' organizations invested in GKE due to its unique features and functionality, including:

- Ability to support clusters up to 15,000 nodes.
- Four-way autoscaling.
- Integrated Cloud Logging and Monitoring.
- Security.
- Integration with specialized Google Cloud APIs.
- Fully-managed service and SRE-based operational tooling.

COMPOSITE ORGANIZATION

Based on the interviews, Forrester constructed a TEI framework, a composite company, and a ROI analysis that illustrates the areas financially affected. The composite organization is representative of the four companies that Forrester interviewed and is used to present the aggregate financial analysis in the next section. The composite organization has the following characteristics:

Description of composite. The composite organization is a large, global enterprise with annual revenues of $5 billion and 10,000 employees. The organization develops and maintains both internal-use applications and B2C SaaS solutions.

Prior to investing in GKE, the organization managed its own on-premises container deployment using a container-based PaaS software solution that was licensed per app instance or workload.

Deployment characteristics. The organization uses GKE in Standard deployment mode.

Key assumptions

- **40 traditional infrastructure FTEs required to manage legacy infrastructure provisioning**
- **10 to 20 new clusters created per year**
- **100 to 300 developers supported by GKE**
- **$5.5 million in provisioning costs prior to GKE with 20% annual growth**
- **$1,500 legacy PaaS license cost**
- **20 FTEs on the security patching team**
- **500 to 720 downtime of degradation events per year**
- **$10,000 revenue impact per downtime or degradation event**

Analysis Of Benefits

Quantified benefit data as applied to the composite

Total Benefits						
Ref.	**Benefit**	**Year 1**	**Year 2**	**Year 3**	**Total**	**Present Value**
Atr	Infrastructure provisioning ops savings	$1,512,000	$1,512,000	$1,512,000	$4,536,000	$3,760,120
Btr	Ongoing infrastructure ops savings	$1,090,385	$1,635,577	$2,180,769	$4,906,731	$3,981,419
Ctr	Improved developer efficiency	$510,000	$2,040,000	$4,590,000	$7,140,000	$5,598,122
Dtr	Infrastructure cost savings	$3,712,500	$4,455,000	$5,346,000	$13,513,500	$11,073,347
Etr	Container-based PaaS license savings	$956,250	$1,147,500	$1,434,375	$3,538,125	$2,895,332
Ftr	Improved security productivity	$691,200	$691,200	$691,200	$2,073,600	$1,718,912
Gtr	Avoided income loss from improved availability	$329,800	$395,760	$474,640	$1,200,200	$983,497
	Total benefits (risk-adjusted)	$8,802,135	$11,877,037	$16,228,984	$36,908,156	$30,010,749

INFRASTRUCTURE PROVISIONING OPS SAVINGS

Evidence and data. Prior to adopting GKE, the interviewees' organizations struggled with the complexities of supporting multiple on-premises container deployments, incurring significant labor overhead to provision adequate infrastructure. After adopting GKE, organizations could easily provision GKE clusters with the gcloud command line or open-source infrastructure as code (IaC):

- The CIO of an advertising firm explained: "Everything is [IaC]. It builds the infrastructure automatically with GKE's stuff."
- A software engineer for a financial services firm detailed, "You can do more with fewer people."

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- Prior to investing in GKE, the composite organization requires a team of 40 infrastructure engineers to complete initial Day 0 and Day 1 tasks, including planning and deployment.
- With GKE, the composite organization quickly creates and deploys clusters to the cloud. The ability to easily allocate compute, networking, and storage resources reduces the required headcount for Day 0 and Day 1 activities by 35%.
- The average fully burdened salary of an infrastructure engineer within the composite organization is $120,000.

Risks. Forrester recognizes that infrastructure provisioning savings may vary from organization to organization. Specific considerations include:

- Organizational size and scope of operations.
- Complexity of legacy environment.
- Location and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV (discounted at 10%) of nearly $3.8 million.

Infrastructure Provisioning Ops Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
A1	Traditional infrastructure engineers required for initial infrastructure provisioning	Composite	40	40	40
A2	Reduction in required headcount with GKE	Interviews	35%	35%	35%
A3	Reduction in FTE requirements for initial provisioning	A1*A2	14	14	14
A4	Fully burdened infrastructure FTE salary (annual)	Composite	$120,000	$120,000	$120,000
At	Infrastructure provisioning ops savings	A3*A4	$1,680,000	$1,680,000	$1,680,000
	Risk adjustment	↓10%			
Atr	Infrastructure provisioning ops savings (risk-adjusted)		$1,512,000	$1,512,000	$1,512,000
	Three-year total: $4,536,000			**Three-year present value: $3,760,120**	

ONGOING INFRASTRUCTURE OPS SAVINGS

Evidence and data. GKE offers organizations a managed and highly automated platform to operate container environments. Prior to having GKE, interviewees found their teams spending a significant amount of time on Day 2 tasks for their clusters, such as monitoring, configuring, and repairing. Managing these ongoing daily tasks was not an issue with one or two clusters, but it became overly burdensome and costly to manage at enterprise scale.

With GKE, organizations automated much of their Day 2 operational work, freeing resources to better support developers and ensuring that Kubernetes could be used at scale across the enterprise.

- The CTO of an advertising firm explained: “We have the intelligence in GKE to manage this variable load situation very effectively. I went from this horde of people managing instances and configuring tools and services and all that jazz to a completely automated solution.”
- A software engineer at a financial services firm detailed, “You are getting your people to do more work higher in the stack, instead of working lower in the infra level.”
- The lead data engineer for a life sciences firm stated: “The biggest thing for us on Day 2 in the cluster is observability. That’s where GKE dovetails nicely with the [Google Cloud Platform (GCP)] environment, logging, cloud monitoring, and some of the off-the-shelf capabilities. We can ship logs in a more federated manner — it’s seamless bridging our GKE and GCP facilities.”
- The life sciences firm interviewee also specifically stated that their firm streamlined the process of auto-scaling profiles through GKE’s easy access to CPU utilization and memory utilization metrics.

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite organization creates 10 to 20 clusters per year over the three-year analysis. The number of clusters created increases over time as use of containers expands across the organization.
- In the prior state, the organization requires 12 weeks of work from a team of 10 cloud engineers to properly tune and configure clusters. With GKE, the time to perform these Day 2 management tasks reduces to three weeks.
- The average fully burdened compensation for a Kubernetes engineer is $140,000 and 50% of their time saved is redeployed to value-add efforts.

Risks. Forrester recognizes that infrastructure operation savings may vary from organization to organization. Specific considerations include.

- Organizational size and scope of operations.
- Complexity of legacy environment.
- Location and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of nearly $4 million.

Ongoing Infrastructure Ops Savings					
Ref.	**Metric**	**Source**	**Year 1**	**Year 2**	**Year 3**
B1	New clusters created per year	Composite	10	15	20
B2	Time required to refine cluster for developer needs without GKE (hours)	12 weeks x 10 cloud engineers x 40 hours	4,800	4,800	4,800
B3	Time required to refine cluster for developer needs with GKE (hours)	3 weeks x 10 cloud engineers x 40 hours	1,200	1,200	1,200
B4	Cluster creation and refinement time savings from improved cluster observability and automation (hours)	B1*(B2-B3)	36,000	54,000	72,000
B5	Fully burdened Kubernetes engineer FTE salary (annual)	Composite	$140,000	$140,000	$140,000
B6	Productivity recapture	Assumption	50%	50%	50%
Bt	Ongoing infrastructure ops savings	B4*(B5/2,080)*B6	$1,211,538	$1,817,308	$2,423,077
	Risk adjustment	↓10%			
Btr	Ongoing infrastructure ops savings (risk-adjusted)		$1,090,385	$1,635,577	$2,180,769
	Three-year total: $4,906,731			**Three-year present value: $3,981,419**	

IMPROVED DEVELOPER EFFICIENCY

Evidence and data. With GKE, developers can focus on writing code for the business instead of troubleshooting container platform issues or waiting for resources. GKE supports stateful and stateless apps, serverless apps, and offers application accelerators for common design patterns. It integrates with Kubernetes-native development tools to improve the velocity of innovation.

Prior to using GKE, organizations could not scale resources easily or automatically to meet changing demands, slowing down continuous integration (CI) and continuous delivery (CD) pipelines and distracting developers with care and feeding of their K8s environments. Additionally, organizations could not fully embrace microservices to refactor apps easily and accelerate app release cycles:

- The global head of cloud for a financial services firm explained: “You can have all this great stuff, but if you don’t have time to do it and all you’re doing is care and feeding, it’s kind of defeating the purpose because you’re not going to be able to use the cool tools you have. It’s like buying a Ferrari and you’re going up and down your driveway because you can’t go anywhere else.”
- The CTO of an advertising firm detailed: “Time-to-market has decreased massively. Smooth operation of our exchange requires significant predictive modeling — some really big machine-learning requirements. The entire cycle of model development, testing, retraining, and then constantly retraining used to take weeks. Now, it’s done six to 12 times per week. With Kubernetes, we can scale the GCP AI platform appropriately based on request volume — and we’ve created a number of unique metrics within GKE to indicate scaling is required.”
- The lead data engineer at a life sciences firm stated: “The biggest thing I was excited about with Kubernetes and GKE was the autoscaling.

That was something we did not have before. The autoscaling features like the node pools in GKE have been super valuable both from the operation perspective but also translates forward to the developers. Because for our developers — what they build, they maintain and they support their first line and last line of support on APIs. So, any issues with scaling and they're the ones in the fire dealing with it."

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite has 100 developers supported by their Kubernetes deployment in Year 1. This expands to 300 by Year 3 with wider enterprise-adoption.
- Developer productivity improves from 10% in the first year to 30% over the three-year period as the organization refines best practices and takes full advantage of GKE's automated scaling and management capabilities.
- The fully burdened compensation for a developer within the organization is $120,000 and 50% of their time saved is redeployed to value-add efforts.

"Node pools have helped considerably with tailoring and shaping node pools that are specific to either use cases or to teams and allowing their API to grow."
Lead data engineer, life sciences

Risks. Forrester recognizes that developer productivity improvements may vary from organization to organization. Specific considerations include:

- Organizational agility and CI/CD practices.
- Legacy environment complexity and baseline developer productivity.
- Location and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of $5.6 million.

Improvement in developer efficiency

10% to 30% over three years

Improved Developer Efficiency					
Ref.	**Metric**	**Source**	**Year 1**	**Year 2**	**Year 3**
C1	Number of developers	Composite	100	200	300
C2	Improved productivity with GKE	Interviews	10%	20%	30%
C3	Fully burdened developer salary (annual)	Composite	$120,000	$120,000	$120,000
C4	Productivity capture	Assumption	50%	50%	50%
Ct	Improved developer efficiency	C1*C2*C3*C4	$600,000	$2,400,000	$5,400,000
	Risk adjustment	↓15%			
Ctr	Improved developer efficiency (risk-adjusted)		$510,000	$2,040,000	$4,590,000
	Three-year total: $7,140,000			**Three-year present value: $5,598,122**	

INFRASTRUCTURE COST SAVINGS

Evidence and data. GKE provides organizations with innovative four-way cluster autoscaling, allowing organizations to pay only for what they need by minimizing infrastructure downtime and maximizing infrastructure use. Autoscaling starts workloads and their underlying infrastructure before demand increases and shuts them down when demand decreases. Prior to using GKE, organizations would need to overprovision resources without precise visibility into cluster usage and with the fear of exceeding capacity if demand spiked. Interviewees elaborated:

- The CTO of an advertising firm explained: "On-prem you are provisioning the peak. So, you're just buying a lot of capex that sits around unused most of the day."
- The global head of cloud for a financial services firm detailed: "The greatest key feature of Kubernetes is the automatic scaling. Somebody can do horizontal pod autoscaling when they are busy one month and lower their cost versus in the legacy days where they had to have all the VMs they could possibly need at any given point without impacting service."
- The lead data engineer at a life sciences firm stated: "Autoscaling has helped us balance costs. Prior to [GKE], we either had to overprovision or just be willing to under provision and jump in and play support when something wasn't scaling."

Additionally, organizations took advantage of GKE's preemptible VMs to further reduce delivery costs of containerized workloads:

- The lead data engineer at a life sciences firm explained: "Another GKE feature that we made use of is the ability to create preemptible VMs. We've created node pools that preemptible VMs backs and seen an immediate cost savings for processes that can tolerate that preemption."
- The CTO of an advertising firm stated: "Around 80% of [our] instances managed by Kubernetes are preemptive…without GKE, we could not do that, and it has completely flipped our delivery cost structure."

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite organization's baseline infrastructure cost for on-premises container workloads is $5.5 million in Year 1. As the organization and usage grow, costs increase by 20% per year.
- The organization reduces its annual costs to deliver container workloads by 75% with autoscaling, improved utilization and preemptive VMs.

Risks. Forrester recognizes that infrastructure savings may vary from organization to organization. Specific considerations include:

- Baseline infrastructure spend.
- Types of workloads, variability of demand, and fault tolerance.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of nearly $11.1 million.

Reduction in run costs with autoscaling, improved utilization, and preemptible VMs

75%

Infrastructure Cost Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
D1	Annual provisioning costs prior to GKE	Y1: Composite Y2 &Y3: 20% YoY growth	$5,500,000	$6,600,000	$7,920,000
D2	Reduction in run costs with autoscaling, improved utilization, and preemptible VMS	Interviews	75%	75%	75%
Dt	Infrastructure cost savings	D1*D2	$4,125,000	$4,950,000	$5,940,000
	Risk adjustment	↓10%			
Dtr	Infrastructure cost savings (risk-adjusted)		$3,712,500	$4,455,000	$5,346,000
	Three-year total: $13,513,500			**Three-year present value: $11,073,347**	

CONTAINER-BASED PAAS LICENSE SAVINGS

Evidence and data. GKE is a fully managed Kubernetes service, automatically providing capabilities that organizations would normally pay an additional subscription to operate themselves. Interviewees highlighted that after investing in GKE, their organizations' existing legacy container-based PaaS licenses became redundant. The interviewees elaborated:

- The global head of cloud for a financial services firm stated that GKE's improved usage of hardware has led to a decrease in licensing for his organization.
- The life sciences interviewee stated that GKE "helps us balance costs with operational burden" by making sure that "we only pay for what we use"

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite organization has 50 to 75 clusters over the three-year period. Each cluster hosts and average of 15 application instances.
- The organization paid $1,500 per instance in PaaS licensing.

Risks. Forrester recognizes that PaaS license savings may vary from organization to organization. Specific considerations include:

- Legacy third-party PaaS usage.
- Number of application instances.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of nearly $2.9 million.

Container-Based PaaS License Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
E1	Eliminated third-party PaaS licenses (number of application instances)	Composite	750	900	1,125
E2	Cost per license	Composite	$1,500	$1,500	$1,500
Et	Container-based PaaS license savings	E1*E2	$1,125,000	$1,350,000	$1,687,500
	Risk adjustment	↓15%			
Etr	Container-based PaaS license savings (risk-adjusted)		$956,250	$1,147,500	$1,434,375
	Three-year total: $3,538,125			**Three-year present value: $2,895,332**	

IMPROVED SECURITY PRODUCTIVITY

Evidence and data. Google scans containers to discover Kubernetes vulnerabilities and missing patches in all clusters GKE manages. If fixes are available, the scanner automatically begins the patching and release process. GKE's trained SREs ensure that all supported GKE clusters are patched and updated without customer intervention. Patching a vulnerability may involve upgrading to a new GKE version on the control plane, user node, or both. GKE nodes are auto-upgraded by default. Interviewees went into more detail:

- The global head of cloud for a financial service stated: "Every 28 days, you're addressing vulnerabilities and, if you've got your pipelines, you aren't doing anything to do that … There's always going to be one-off patching that you don't want to wait 28 days for, but we can find out what you're running and save time patching."
- A software engineer at a financial services firm explained: "We don't have to think about VM image patching anymore. It's all done through Google APIs. We still have the security team doing some audits when trading is done on these pipelines."

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite organization has a team of 20 FTEs dedicated to patching and vulnerability management.
- Automated scanning and patching reduces team effort by 80%.
- Patching teams are located in low-cost geographies with a fully burdened rate of $60,000.

Risks. Forrester recognizes that security productivity improvements may vary from organization to organization. Specific considerations include:

- Team size and location.
- Baseline effort and frequency of patching.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of over $1.7 million.

Reduction in patching effort:

80%

Improved Security Productivity					
Ref.	**Metric**	**Source**	**Year 1**	**Year 2**	**Year 3**
F1	Security FTEs dedicated to patching and vulnerability management	Composite	20	20	20
F2	Time savings for patching with GKE	Interviews	80%	80%	80%
F3	Fully burdened salary for patch and vulnerability management FTE (annual)	Composite	$60,000	$60,000	$60,000
F4	Productivity recapture	Assumption	80%	80%	80%
Ft	Improved security productivity	F1*F2*F3*F4	$768,000	$768,000	$768,000
	Risk adjustment	↓10%			
Ftr	Improved security productivity (risk-adjusted)		$691,200	$691,200	$691,200
	Three-year total: $2,073,600			**Three-year present value: $1,718,912**	

AVOIDED INCOME LOSS FROM IMPROVED AVAILABILITY

Evidence and data. Interviewees' organizations struggled with availability and downtime issues as their legacy environments lacked scalability and automation. While some organizations served primarily internal development teams, those who also supported customer-facing applications risked delivering a poor customer experience and a potential loss of revenue without GKE. With autoscaling, cluster logging and monitoring, and a managed control plane, organizations ensure that services are always available and eliminate downtime or degradation events. The interviewees provided examples:

- The advertising firm risked losing revenue for every second that ads were not served. The CTO stated: "If the service degraded, you lose gravity. If we slow down, revenue drops. We've eliminated that — and not through overprovisioning."
- One financial services firm supported their consumer banking website and application with GKE. The head of global cloud explained that any slowdowns or disruptions in availability lead to customer complaints and churn.
- A second financial services firm used Kubernetes deployment to support high-frequency trading algorithm modeling. A software engineer in the firm explained that time-sensitive releases needed to be finished before the subsequent day's trading. Any delay due to unavailable clusters would cause them to revert to the previous day's model and potentially lose revenue from not having the most up-to-date data.

Modeling and assumptions. Forrester modeled the financial impact for the composite organization based on the following assumptions:

- The composite organization experiences 500 downtime or degradation events in its previous environment, which grows at 20% with growth of the company.
- The organization improves availability by 97% with GKE.

- Each downtime or degradation event costs the organization an average of $10,000 in revenue.
- The organization's average margin is 8%.

Improvement in availability and service with GKE:

97%

Risks. Forrester recognizes that availability improvement benefits may vary from organization to organization. Specific considerations include:

- Baseline number of availability events in legacy environment.
- Types of services and related revenue.
- Average margins.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of $983,497.

Avoided Income Loss From Improved Availability

Ref.	Metric	Source	Year 1	Year 2	Year 3
G1	Number of downtime or service degradation events per year without GKE	Y1: Composite Y2 and Y3: 20% YoY growth	500	600	720
G2	Improvement in availability and service with GKE	Interviews	97%	97%	97%
G3	Avoided events involving downtime or degradation of service (rounded)	G1*G2	485	582	698
G4	Revenue impact per event	Composite	$10,000	$10,000	$10,000
G5	Profit margin	Composite	8%	8%	8%
Gt	Avoided income loss from improved availability	G3*G4*G5	$388,000	$465,600	$558,400
	Risk adjustment	↓15%			
Gtr	Avoided income loss from improved availability (risk-adjusted)		$329,800	$395,760	$474,640
	Three-year total: $1,200,200			**Three-year present value: $983,497**	

UNQUANTIFIED BENEFITS

Additional benefits that customers experienced but were not able to quantify include:

- **Hardened security posture.** The global head of cloud in financial services stated: "My cybersecurity people absolutely love the other added features on GCP's side that I can use in GKE for image management, signing my images, and more. How do I know if that [dynamic-link library (DLL)] or that binary had malicious content in it? It was tough to do that, and now Google gives you those tools…Globally Google is connected anywhere you want, but they've never been breached because of their security, and that's the workload identification and everything else."
- **Access to new market opportunities.** The same interviewee also spoke about how access to a new market would have been impossible without Google's approach, explaining that: "In this new market, public cloud isn't cleared at all for financial institutions. But GKE gave me the ability to spin up a cluster for them. It was a huge opportunity in Asia Pacific."
- **Improved customer experience (CX).** With GKE's cluster autoscaler, nodes are added to the node pool when demand is high to ensure availability. Additionally, Google SREs monitor GKE clusters to ensure they are accessible. By limiting downtime and optimizing performance of underlying application infrastructure, organizations can provide a better CX.
- **Enhanced employee experience.** Automation of repetitive manual tasks allows organizations to redeploy employees to more value-add work. GKE also allows developers to focus on development instead of troubleshooting infrastructure issues. Working with cutting-edge technology allows developers to grow their skill sets, advance their careers, and can help organizations recruit top-tier talent.

FLEXIBILITY

The value of flexibility is unique to each customer. There are multiple scenarios in which a customer might implement GKE and later realize additional uses and business opportunities, including:

- **Deploying additional Google Cloud technologies.** Using GKE can facilitate expansion of other Google Cloud technologies within an organization. Taking advantage of machine learning, AI, and vertical-specific APIs can provide additional benefits within an organization.
- **Avoiding vendor lock-in**. Performing containerization with Kubernetes as a base helps reduce the risks of lock-in or excess future reengineering.
- **Accelerating GKE benefits with Autopilot.** Autopilot is a new mode of operation in GKE designed to further reduce the costs of managing and optimizing clusters. With Autopilot, GKE provisions and manages both the Kubernetes control plane and the customers' underlying infrastructure in a completely hands-off experience. For further analysis of Autopilot's impact on an organization, see Appendix B.

Flexibility would also be quantified when evaluated as part of a specific project (described in more detail in Appendix A).

Analysis Of Costs

Quantified cost data as applied to the composite

Total Costs

Ref.	Cost	Initial	Year 1	Year 2	Year 3	Total	Present Value
Htr	GKE annual cluster management fee	$0	$1,649,174	$1,979,206	$2,474,253	$6,102,633	$4,993,900
Itr	Implementation and ongoing labor	$2,600,769	$1,227,692	$1,227,692	$1,227,692	$6,283,846	$5,653,858
	Total costs (risk adjusted)	$2,600,769	$2,876,867	$3,206,898	$3,701,945	$12,386,479	$10,647,758

GKE ANNUAL CLUSTER MANAGEMENT FEE

Evidence and data. Organizations using GKE will incur usage-based fees. Based on customer interviews, Forrester has used Standard mode list pricing, in which an organization accrues a management cost of $0.10 per cluster per hour, irrespective of size or topography.

In Standard mode, GKE uses Compute Engine instance worker nodes in the cluster. You are billed for each of those instances according to Compute Engine's pricing until the nodes are deleted. Compute Engine resources are billed on a per-second basis with a one-minute minimum usage cost.

GKE includes a Service Level Agreement (SLA) that's financially backed, providing availability of 99.95% for the control plane of Regional clusters, and 99.5% for the control plane of Zonal clusters.

Modeling and assumptions. When modeling the composite Forrester has assumed:

- The composite organization has 50 clusters in Year 1, 60 in Year 2, and 70 in Year 3.
- An average of 50 nodes per cluster.

Risks. Management costs may vary based on a number of variables, including:

- Size of deployment and usage.
- Variable ancillary costs, such as storage and egress.

Results. To account for these risks, Forrester adjusted this cost upward by 10%, yielding a three-year, risk-adjusted total PV (discounted at 10%) of nearly $5 million.

GKE Annual Cluster Management Fee						
Ref.	**Metric**	**Source**	**Initial**	**Year 1**	**Year 2**	**Year 3**
H1	GKE annual cluster management fee		$0	$1,499,250	$1,799,278	$2,249,321
Ht	GKE annual cluster management fee	H1	$0	$1,499,250	$1,799,278	$2,249,321
	Risk adjustment	↑10%				
Htr	GKE annual cluster management fee (risk-adjusted)		$0	$1,649,174	$1,979,206	$2,474,253
	Three-year total: $6,102,633			**Three-year present value: $4,993,900**		

IMPLEMENTATION AND ONGOING LABOR

Evidence and data. While deploying GKE is light touch, organizations will experience labor costs for planning, design, change management, training, and ongoing management.

Interviewees' organizations invested in professional services for planning, design, and training. And, while GKE and GCP offer a number of off the shelf capabilities and integrations, some organizations chose to invest internal labor hours to build custom visualizations and dashboards.

Modeling and assumptions. In modeling the composite organization, Forrester assumes:

- The composite organization initially invests $1 million in professional services for planning, design, training, and support.
- Initial training is three weeks (60 hours) with one week's worth of training per year to refresh skills and understand new features and capabilities. The fully burdened rate for engineers undergoing training is $120,000 per year.
- The organization has a dedicated team of five SREs with a fully burdened salary of $160,000.

Risks. The costs associated with implementation, training, and ongoing support may vary based on a number of variables including:

- Size and scope of operations.
- Baseline skill sets.
- Organizational agility.
- Prevailing labor rates.

Results. To account for these risks, Forrester adjusted this cost upward by 5%, yielding a three-year, risk-adjusted total PV of nearly $5.7 million.

Implementation And Ongoing Labor

Ref.	Metric	Source	Initial	Year 1	Year 2	Year 3
I1	Professional services	Interviews	$1,000,000			
I2	Developer/engineer FTEs trained	Composite organization	160	160	160	160
I3	Hours of training required	Interviews	160	40	40	40
I4	Average trainee salary (annual)	Industry benchmarks	$120,000	$120,000	$120,000	$120,000
I5	Dedicated SRE team	Interviews		5	5	5
I6	Fully burdened SRE salary (annual)	Industry benchmarks		$160,000	$160,000	$160,000
It	Implementation and ongoing labor	I1+(I2*I3*(I4/2,080))+(I5*I6)	$2,476,923	$1,169,231	$1,169,231	$1,169,231
	Risk adjustment	↑5%				
Itr	Implementation and ongoing labor (risk-adjusted)		$2,600,769	$1,227,692	$1,227,692	$1,227,692
	Three-year total: $6,283,846			**Three-year present value: $5,653,858**		

Financial Summary

CONSOLIDATED THREE-YEAR RISK-ADJUSTED METRICS

Cash Flow Chart (Risk-Adjusted)

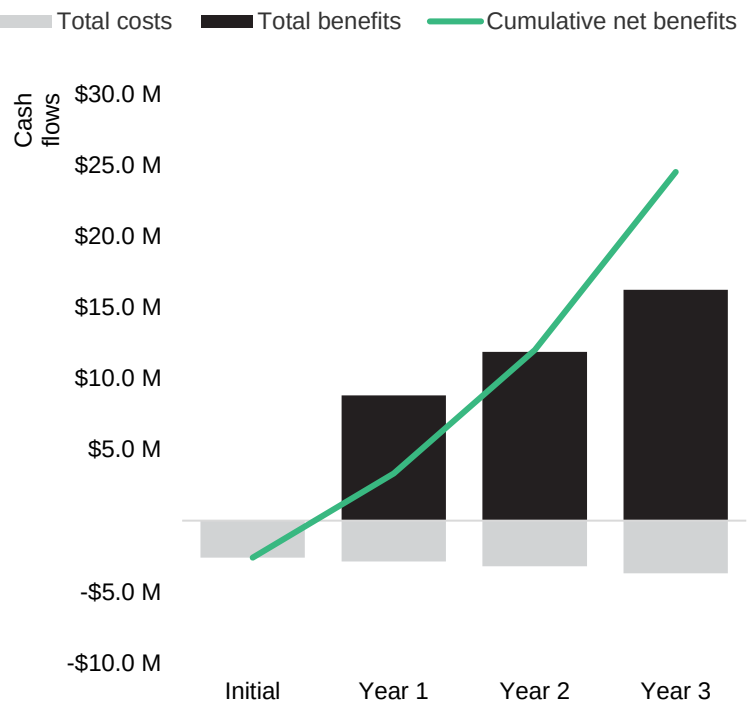

The financial results calculated in the Benefits and Costs sections can be used to determine the ROI, NPV, and payback period for the composite organization's investment. Forrester assumes a yearly discount rate of 10% for this analysis.

These risk-adjusted ROI, NPV, and payback period values are determined by applying risk-adjustment factors to the unadjusted results in each Benefit and Cost section.

Cash Flow Analysis (Risk-Adjusted Estimates)

	Initial	Year 1	Year 2	Year 3	Total	Present Value
Total costs	($2,600,769)	($2,876,867)	($3,206,898)	($3,701,945)	($12,386,479)	($10,647,758)
Total benefits	$0	$8,802,135	$11,877,037	$16,228,984	$36,908,156	$30,010,749
Net benefits	($2,600,769)	$5,925,268	$8,670,139	$12,527,039	$24,521,676	$19,362,991
ROI						182%
Payback						<6 months

Appendix A: Total Economic Impact

Total Economic Impact is a methodology developed by Forrester Research that enhances a company's technology decision-making processes and assists vendors in communicating the value proposition of their products and services to clients. The TEI methodology helps companies demonstrate, justify, and realize the tangible value of IT initiatives to both senior management and other key business stakeholders.

TOTAL ECONOMIC IMPACT APPROACH

Benefits represent the value delivered to the business by the product. The TEI methodology places equal weight on the measure of benefits and the measure of costs, allowing for a full examination of the effect of the technology on the entire organization.

Costs consider all expenses necessary to deliver the proposed value, or benefits, of the product. The cost category within TEI captures incremental costs over the existing environment for ongoing costs associated with the solution.

Flexibility represents the strategic value that can be obtained for some future additional investment building on top of the initial investment already made. Having the ability to capture that benefit has a PV that can be estimated.

Risks measure the uncertainty of benefit and cost estimates given: 1) the likelihood that estimates will meet original projections and 2) the likelihood that estimates will be tracked over time. TEI risk factors are based on "triangular distribution."

The initial investment column contains costs incurred at "time 0" or at the beginning of Year 1 that are not discounted. All other cash flows are discounted using the discount rate at the end of the year. PV calculations are calculated for each total cost and benefit estimate. NPV calculations in the summary tables are the sum of the initial investment and the discounted cash flows in each year. Sums and present value calculations of the Total Benefits, Total Costs, and Cash Flow tables may not exactly add up, as some rounding may occur.

PRESENT VALUE (PV)

The present or current value of (discounted) cost and benefit estimates given at an interest rate (the discount rate). The PV of costs and benefits feed into the total NPV of cash flows.

NET PRESENT VALUE (NPV)

The present or current value of (discounted) future net cash flows given an interest rate (the discount rate). A positive project NPV normally indicates that the investment should be made, unless other projects have higher NPVs.

RETURN ON INVESTMENT (ROI)

A project's expected return in percentage terms. ROI is calculated by dividing net benefits (benefits less costs) by costs.

DISCOUNT RATE

The interest rate used in cash flow analysis to take into account the time value of money. Organizations typically use discount rates between 8% and 16%.

PAYBACK PERIOD

The breakeven point for an investment. This is the point in time at which net benefits (benefits minus costs) equal initial investment or cost.

Appendix B: GKE Autopilot Benefits

- **Optimize production like a K8s expert.** Google creates clusters according to battle-tested and hardened best practices that yield optimized configuration and security posture ready for production, helping you reduce the GKE learning curve and get to production faster.
- **Stronger security posture**. Google takes full responsibility for nodes and infrastructure security, eliminating infrastructure security management tasks.
- **Reduce Day 2 ops.** Google monitors the entire Autopilot cluster — control plane, worker nodes, and core Kubernetes system components — ensuring your pods are always scheduled.
- **Improve resource efficiency.** Google takes full responsibility for optimizing your resource consumption. You pay only for your Pods, not for nodes.
- **Still Kubernetes.** Industry-first to provide fully managed Kubernetes service that implements full Kubernetes' API.

To find out more about the expanded benefits of Autopilot click here.

Appendix C: Endnotes

[1] Total Economic Impact is a methodology developed by Forrester Research that enhances a company's technology decision-making processes and assists vendors in communicating the value proposition of their products and services to clients. The TEI methodology helps companies demonstrate, justify, and realize the tangible value of IT initiatives to both senior management and other key business stakeholders


FORRESTER®