FORRESTER®

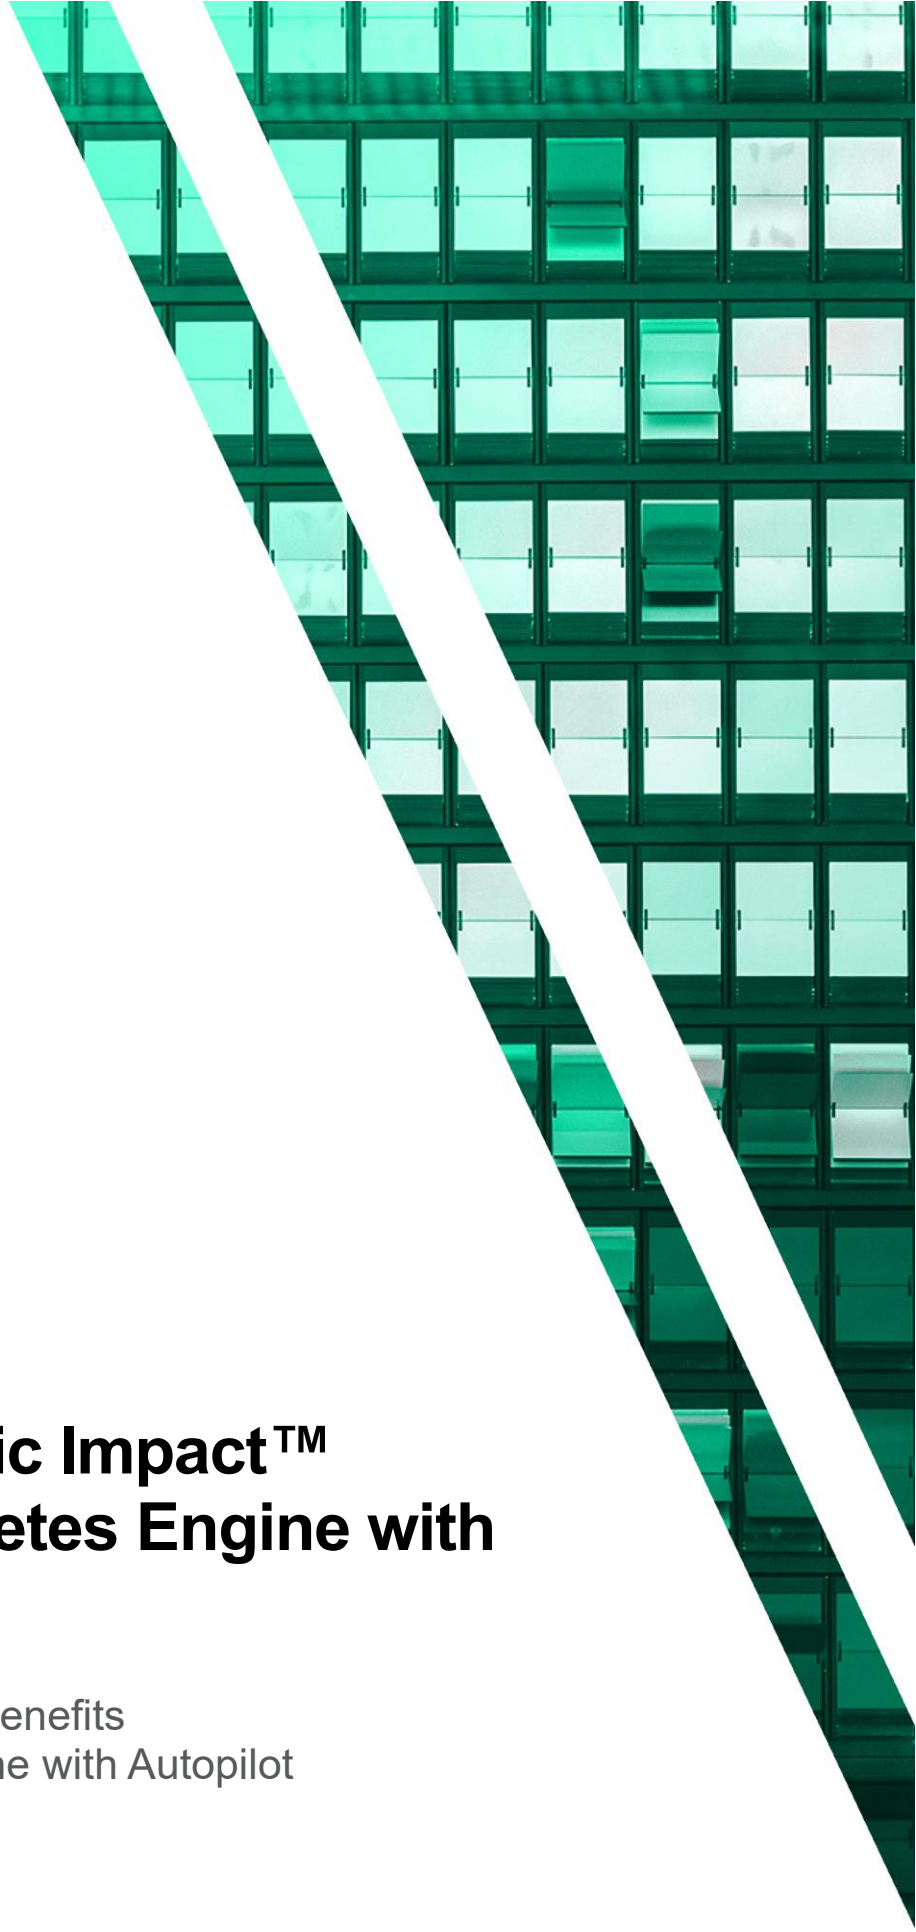

The Total Economic Impact™ Of Google Kubernetes Engine with Autopilot

Cost Savings And Business Benefits
Enabled By Kubernetes Engine with Autopilot

JANUARY 2023

Table Of Contents

Consulting Team: Sam Sexton

Executive Summary 1

The Google Kubernetes Engine with Autopilot Customer Journey 6

Key Challenges 6

Investment Objectives 7

Solution requirements 7

Composite Organization 8

Analysis Of Benefits 9

Infrastructure Provisioning Ops Savings 9

Ongoing Infrastructure Ops Savings 10

Improved Developer Efficiency 12

Infrastructure Cost Savings 14

Container-Based PaaS License Savings 15

Improved Security Productivity 15

Avoided Income Loss From Improved Availability 17

Unquantified Benefits 18

Flexibility 19

Analysis Of Costs 20

GKE Annual Cluster Management Fee 20

Implementation And Ongoing Labor 21

Financial Summary 23

Appendix A: Total Economic Impact 24

Appendix B: Endnotes 25

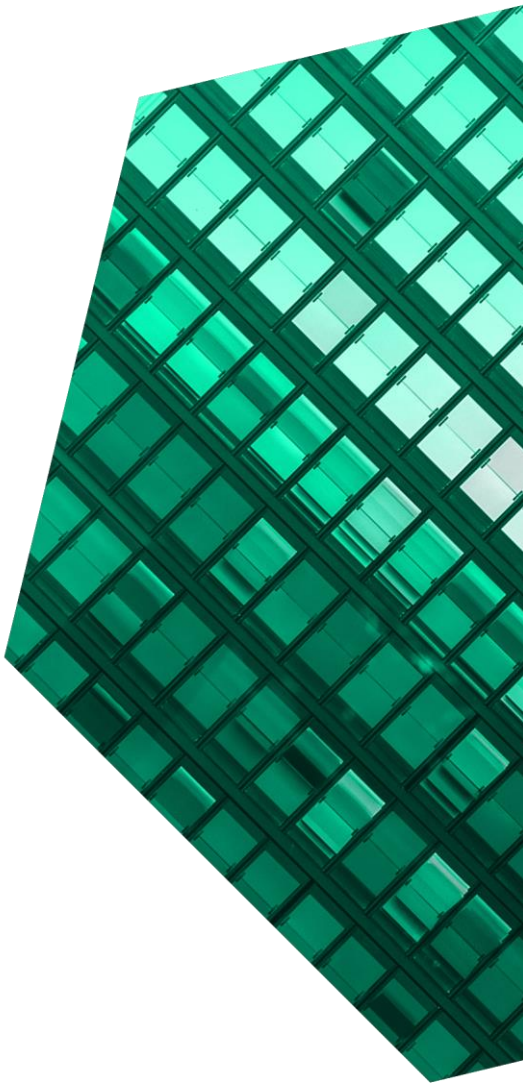

ABOUT FORRESTER CONSULTING

Forrester provides independent and objective research-based consulting to help leaders deliver key transformation outcomes. Fueled by our customer-obsessed research, Forrester's seasoned consultants partner with leaders to execute on their priorities using a unique engagement model that tailors to diverse needs and ensures lasting impact. For more information, visit forrester.com/consulting.

Executive Summary

Software containerization provides organizations with a versatile tool to improve efficiency, performance, and innovation. However, doing so effectively often requires expertise, troubleshooting, and lengthy implementations that can hold back these benefits, especially at scale. This journey doesn't necessarily conclude when organizations implement solutions to improve their containerization because configuration, patching, and inefficiencies can still cost precious time and resources.

Google Kubernetes Engine (GKE) is an enterprise-grade platform that is used to build, deploy, and run a wide variety of containerized applications. When used correctly, GKE can aid in the management and security of these applications. With the Autopilot mode of operation engaged, organizations can further save time on configuration, optimization, and security updates by fully automating these tasks.

In 2021, Google commissioned Forrester Consulting to conduct a Total Economic Impact™ (TEI) study that evaluated the potential return on investment (ROI) enterprises can realize by deploying Google Kubernetes Engine ("GKE Standard")[1] The purpose of this study is to update the 2021 findings by evaluating the potential financial impact of Kubernetes Engine with Autopilot on organizations.

To better understand the benefits, costs, and risks associated with this investment, Forrester interviewed four representatives with experience using Google Kubernetes Engine Standard for the original study, and five more customers using GKE Autopilot for this update. For the purposes of this study, Forrester aggregated the interviewees' experiences and combined the results into a single composite organization that is a large, global enterprise with annual revenues of $5 billion and 10,000 employees.

KEY STATISTICS

Return on investment (ROI)
315%

Net present value (NPV)
$27.57M

Security patching time savings with Autopilot
95%

The interviewees noted that prior to using GKE Autopilot, their organizations lacked an effective, unified solution for containerization and they often managed their own deployments both on-premises and in the cloud. This led to wasted time configuring and managing security updates, infrastructure troubleshooting, and limited deployment of containers at scale.

After the investment in GKE, the interviewees' organizations were able to modernize application infrastructure and simplify deployment and management. GKE Autopilot provided further savings by automating previously manual tasks, which improved efficiency and saved time on top of the

broader performance, agility, and accelerated innovation of GKE Standard.

KEY FINDINGS

Quantified benefits. Three-year, risk-adjusted present value (PV) quantified benefits for the composite organization include:

- **40% reduction of initial environment setup labor requirements.** GKE Standard reduces or automates the labor required for the composite organization to initially set up Kubernetes tasks via smart default features, saving 35% of setup labor time. This is bolstered by GKE Autopilot, which provides further automation of setting up node pools. The composite organization saves 40% of the effort on Day 0/1 tasks compared to managing its own on-premises Kubernetes deployment.
- **83% reduction of ongoing cluster management time.** GKE Standard manages many of the cluster lifecycle's ongoing activities, from monitoring, tuning to right-sizing and auto-repairing. This saves the composite organization up to 75% of cluster management time. GKE Autopilot further reduces the amount of refining cluster needs, which reduces the composite's ongoing cluster management and optimization time by 83%.
- **45% improvement to developer efficiency.** Prior to implementing GKE, developers had to manually manage several development and configuration duties; GKE automates this, which improves productivity. The composite reaps additional benefits by using GKE in the cloud for faster provisioning of resources for development teams relative to on-premises provisioning; these benefits are further enhanced with GKE Autopilot's automation. The composite organization experiences productivity improvements of 10% per year due to GKE Standard, and another 15% improvement due to Autopilot, which culminates in a 45% improvement in productivity by Year 3.
- **85% reduction of infrastructure costs.** GKE enables the composite organization to use autoscaling to maximize infrastructure utilization and reduce costs by 75%. Preemptible virtual machines (VMs) provide further savings via cost-effective methods of completing compute-intensive batch work at lower cost. With GKE Autopilot, the composite organization increases its cost savings relative to GKE Standard by taking advantage of the "pay-as-you-go" model, and this saves an additional 10% or a total of 85%.
- **Elimination of existing container-based platform-as-a-service (PaaS) license spend.** GKE Standard's suite of managed container orchestration and management capabilities eliminates the composite organization's need to maintain its prior PaaS licenses. The composite realizes savings of $1,500 per application instance by eliminating its third-party, container-based PaaS licenses.
- **95% improvement to security productivity.** With GKE Standard, Google scans the composite organization's containers to discover vulnerabilities or missing patches in Google-managed containers. This allows it to reduce its time and energy spent on hunting for vulnerabilities or patching and reduces the time spent on these tasks by 80%. Because GKE Autopilot takes on complete management of clusters for the composite organization, the organization is also able to completely manage security patching, which vastly reduces the workload of the security team. The composite organization improves the patching team's efficiency by another 15%, for a combined 95% improvement, due to GKE Autopilot.

- **Avoidance of lost revenue with 98% improvement in availability.** GKE's autoscaling improves performance by ensuring that the composite organization has adequate resources to meet compute demands, even during peak periods of usage, further aided by Google's site reliability engineers (SREs) monitoring clusters.

"With one use case, [Autopilot] gave us close to a 50% reduction in spend because we're only paying for what we use."

Engineering manager, Media

This boosts availability by 97%, and the increase in uptime is slightly improved by the workload separation effects of GKE Autopilot, which results in another 1% improvement in availability for a total improvement of 98%.

Unquantified benefits. Benefits that provide value for the composite organization but are not quantified in this study include:

- **Improved customer experience (CX).** Interviewees said few things are more frustrating for their organizations than poor app performance, but providing solid, reliable application performance during periods of high demand is challenging. To ensure availability during peak demand, the GKE cluster autoscaler adds nodes to the node pool. Furthermore, Google SREs monitor GKE clusters to limit downtime and optimize performance, which helps to provide a better CX.
- **Hardened security profile and quick app fixes.** GKE starts with a high level of baseline security due to data encryption and regular vulnerability scans further enhanced with GKE Autopilot, which ensures everything is patched as soon as possible. This enables the composite organization's developers to accelerate their release velocity and quickly patch with little effort and without increasing security risks.
- **Enhanced employee experience (EX).** GKE Standard already automates several repetitive tasks for the composite organization with GKE Autopilot saving even more time on provisioning and patching. This allows the composite to redeploy developers and other employees to higher-value work instead of troubleshooting infrastructure issues.
- **Greater scalability and agility to expand into new markets and regions.** GKE lets the composite organization quickly spin up new clusters, which creates a production-ready environment for its teams in just a few hours. This is much faster than the lengthy approval process the composite's global teams previously required with traditional provisioning.

Costs. Three-year, risk-adjusted PV costs for the composite organization include:

- **$3.8 million in GKE Autopilot annual cluster management fees.** The composite organization pays an hourly per-cluster fee for usage of GKE Autopilot on its clusters.
- **$4.9 million for implementation and ongoing labor.** Costs include professional services, training, and ongoing platform support.

The representative interviews and financial analysis found that a composite organization experiences benefits of $36.32 million over three years versus costs of $8.75 million, adding up to a net present value (NPV) of $27.57 million and an ROI of 315%.

ROI
315%

BENEFITS PV
$36.32M

NPV
$27.57M

PAYBACK
<6 months

Benefits (Three-Year)

“For us, the major difference between Standard and Autopilot is going to be the money that we saved from the node breakage or overage costs and the lack of security personnel.”

— Lead SRE, Financial Technology

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 Autopilot.

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 Autopilot can have on an organization.

DUE DILIGENCE

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

INTERVIEWS

Interviewed four representatives at organizations using GKE Standard and five representatives at organizations using GKE Autopilot to obtain data with respect to costs, benefits, and risks.

COMPOSITE ORGANIZATION

Designed a composite organization based on characteristics of the interviewees' 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 interviewees.

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 with Autopilot Customer Journey

Drivers leading to the GKE Autopilot investment

Interviews

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
Medical technology	APAC	Head of platform engineering	$14 million	200
Media	NA	Engineering manager	$650 million	4,500
Financial technology	NA	Lead SRE	$20 million	200
Education	NA	Assistant professor	N/A	N/A
Health care	Global	Senior leader of IT	$5 billion	10,000

KEY CHALLENGES

Forrester interviewed four decision-makers at organizations using GKE Standard and five decision-makers at organizations using GKE Autopilot a year later. Each interviewee's organization had some prior experience with Kubernetes: Some managed their own on-premises container deployments and others worked with a different PaaS before switching to GKE.

The interviewees noted how their organizations struggled with common challenges, including:

- **Reaching the limits of scalability with legacy platforms.** Many of the interviewees told Forrester that before using GKE, their organization had exhausted the scale of Kubernetes it was able to self-deploy and manage. Those from organizations working with another partner bemoaned the lack of auto-scaling. In either scenario, teams were forced to focus more on troubleshooting issues and ensuring availability than fully realizing the benefits of containerization.
- **Pressure to reduce operating costs.** Interviewees described common organizational goals and pressures to continually reduce operating costs beyond the savings realized during the initial container deployments. This also applied to interviewees from organizations with other PaaS providers.
- **Desire to focus development and operations teams on more interesting and high-value work.** The interviewees said configuration and repair work is tedious, and they wanted to have their organization's developers focus on product innovation instead of basic maintenance. Even those with a more advanced Kubernetes

deployment still had to spend more time on provisioning than they preferred.

- **Difficulty flagging and responding to capacity and availability issues.** Internal adoption within the interviewees' organizations continued to grow, and that led to availability issues, degradation events, and overall reduced performance.
- **Too much time spent on managing security**. The interviewees from organizations with experience with GKE Autopilot specifically cited security as a time-consuming issue before the implementation. Several reported that they were spending too much time manually ensuring high security, or keeping their environment updated via patches.

"[Calculating our compute] was very finger-in-the-wind guesstimates. Before GKE Autopilot, there was a lot of waste that was happening."

Lead SRE, financial technology

INVESTMENT OBJECTIVES

The interviewees' organizations searched for a solution that could:

- Reduce operational complexity and consolidate development environments.
- Provide a platform requiring little effort to scale and manage.
- Improve application availability and performance.
- Improve release frequency and time-to-market for development teams.
- Accommodate very large workloads (i.e., workloads of 10,000 nodes or more).
- Reduce infrastructure costs.
- Reduce manual labor on managing VMs and security.

"[Downtime was] significant to our business since we're a media company. Uptime of our websites drives the engagement by end-users, so any sort of downtime would be a significant impact to our revenue."

Engineering manager, media

SOLUTION REQUIREMENTS

The initial four interviewees invested in GKE Standard due to its unique features and functionality, including:

- The ability to support clusters of 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.

The four interviewees that chose to deploy GKE Autopilot had the following additional solution requirements:

- Automated security patching.
- Further automation of cluster management.

COMPOSITE ORGANIZATION

Based on the interviews, Forrester constructed a TEI framework, a composite company, and an ROI analysis that illustrates the areas financially affected. The composite organization is representative of the eight interviewees, and it 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 Autopilot.

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 with 20% annual growth prior to using GKE**
- **$1,500 legacy PaaS license cost**
- **20 FTEs on the security patching team**
- **500 to 720 downtime or degradation events per year**
- **$10,000 revenue impact per downtime or degradation event**
- **GKE Standard in Year 1**
- **GKE Autopilot deployment**

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	$1,700,000	$1,700,000	$1,700,000	$5,100,000	$4,227,648
Btr	Ongoing infrastructure ops savings	$1,302,462	$1,953,692	$2,604,923	$5,861,077	$4,755,795
Ctr	Improved developer efficiency	$1,250,000	$3,500,000	$6,750,000	$11,500,000	$9,100,301
Dtr	Infrastructure cost savings	$3,973,750	$4,768,500	$5,722,200	$14,464,450	$11,852,583
Etr	Container-based PaaS license savings	$900,000	$1,080,000	$1,350,000	$3,330,000	$2,725,019
Ftr	Improved security productivity	$1,094,400	$1,094,400	$1,094,400	$3,283,200	$2,721,611
Gtr	Avoided income loss from improved availability	$313,600	$376,320	$451,840	$1,141,760	$935,573
	Total benefits (risk-adjusted)	$10,534,212	$14,472,912	$19,673,363	$44,680,487	$36,318,530

INFRASTRUCTURE PROVISIONING OPS SAVINGS

Evidence and data. Before adopting GKE, several of the interviewees' organizations incurred significant labor costs to provision infrastructure for their on-premises container deployments. Adopting GKE made provisioning clusters much easier via the Google Cloud CLI (gcloud CLI) command line and open-source infrastructure as code (IaC), while adopting GKE Autopilot further streamlined the process of provisioning and freed up resources:

- The engineering manager for the media company explained said, "Now our team is able to work with product to focus on adding business value versus tactical concerns."
- The lead SRE of the financial technology company stated, "Setting up clusters [in Autopilot] is much easier because you don't have to terraform any of the node pools."
- An assistant professor said: "You set a couple parameters, you push the button and it's done… Autopilot makes it even easier because you don't have to worry to figure out whether or not you [have] got the right machine types to find your node pools."

"The reason we use Autopilot is that we provision hundreds of clusters, especially for customers, for other customized medical institutions. We need customer specific environment of application."

Head of platform engineering, medical technology

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- Prior to investing in GKE, the composite organization required a team of 40 infrastructure engineers to complete initial Day 0 and Day 1 tasks including planning and deployment.
- GKE Standard lets the composite organization quickly create and deploy clusters to the cloud, which reduces Day 0 and Day 1 storage resources by 35%. GKE Autopilot further boosts this another 5% for a total of 40%.
- The average fully burdened salary of an infrastructure engineer is $125,000.

Risks. Factors that could impact the size of this benefit for organizations include:

- Organizational size and scope of operations.
- The complexity of legacy environment.
- Location and prevailing labor rates.
- The number and type of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV (discounted at 15%) of $4.2 million.

Infrastructure Provisioning Ops Savings

Ref.	Metric	Calculation	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	Additional reduction with Autopilot	Interviews	5%	5%	5%
A4	Total reduction to required headcount	A2+A3	40%	40%	40%
A5	Reduction in FTE requirements for initial provisioning	A1*A4	16	16	16
A6	Fully burdened infrastructure FTE salary	Composite	$125,000	$125,000	$125,000
At	Infrastructure provisioning ops savings	A3*A4	$2,000,000	$2,000,000	$2,000,000
	Risk adjustment	↓15%			
Atr	Infrastructure provisioning ops savings (risk-adjusted)		$1,700,000	$1,700,000	$1,700,000
	Three-year total: $5,100,000			**Three-year present value: $4,227,648**	

ONGOING INFRASTRUCTURE OPS SAVINGS

Evidence and data. Interviewees said that prior to adopting GKE, their organizations' teams spent more time than ideal on burdensome and manual Day 2 cluster management tasks such as monitoring, configuring, and repairing.

Switching to GKE Standard automated much of the organizations' Day 2 operational work, and the additional handover of duties to Google by activating GKE Autopilot helped further save time. This resulted in better developer support and ensured Kubernetes could be used at scale enterprise wide.

- The head of platform engineering for the medical technology firm told Forrester that, “Standard is nice, but we still need to throw people at upgrading, maintenance, etcetera…with Autopilot, we just need to care about the workloads we deploy.”
- The engineering manager for the media organization explained: “[Autopilot has] definitely freed our engineers to not have to pay as much attention to how quickly or how particularly they get their code out – now they’re able to quickly display their code and get it to the main branch within minutes.”
- The assistant professor at the university stated: “Having the flexibility to iterate and try and tinker quickly is literally the difference between success and failure… It becomes the lifeblood that we rely on because of what it can enable. And GKE Autopilot is one of those tools, right?”

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- The composite creates 10 to 20 clusters per year during the three-year analysis period. The number of clusters created increases over time as container usage expands across the organization.
- In its prior state, the organization required 12 weeks of work from a team of 10 cloud engineers to properly tune and configure clusters. With GKE Standard, the time to perform these Day 2 management tasks is reduced to about three weeks or by 75%.
- By operating with GKE Autopilot, the composite organization further reduces the time to perform Day 2 management tasks by about another week relative to GKE Standard or by another 8%.
- The average fully burdened compensation for a Kubernetes engineer is $160,000, and they redeploy 50% of their time saved to value-add efforts.

> **“A pretty big factor in moving over to GKE, and GKE Autopilot in particular, is the offloading of the operational aspect of Kubernetes from our predecessor, which was on standard VMs that we were responsible for…moving over to GKE and using Google’s Container-Optimized OS allowed us to get a lot of these benefits out of box.”**
>
> *Head of platform engineering, medical technology*

Risks. Factors that could impact this benefit for organizations include:

- Organizational size and scope of operations.
- The complexity of the legacy environment.
- Location and prevailing labor rates.
- The number of remaining Day 2 tasks to be automated with GKE Autopilot.
- The number and type of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of $4.8 million.

Ongoing Infrastructure Ops Savings

Ref.	Metric	Calculation	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 savings on refining clusters for developer needs with GKE	Interviews	75%	75%	75%
B4	Additional time savings via Autopilot	Interviews	8%	8%	8%
B5	Total time savings on refining clusters for developer needs	B3+B4	83%	83%	83%
B6	Cluster creation and refinement time savings from improved cluster observability and automation (hours)	B1*B2*B5	39,840	59,760	79,680
B7	Fully burdened Kubernetes engineer FTE salary	Composite	$160,000	$160,000	$160,000
B8	Productivity recapture	Assumption	50%	50%	50%
Bt	Ongoing infrastructure ops savings	B6*(B7/2080)*B8	$1,532,308	$2,298,462	$3,064,615
	Risk adjustment	↓15%			
Btr	Ongoing infrastructure ops savings (risk-adjusted)		$1,302,462	$1,953,692	$2,604,923
	Three-year total: $5,861,077		**Three-year present value: $4,755,795**		

IMPROVED DEVELOPER EFFICIENCY

Evidence and data. Prior to implementing GKE, organizations had difficulty scaling resources automatically, which meant meeting changing demands required manual effort. Their continuous integration (CI) or continuous delivery (CD) pipelines were slowed down, and organizations could not fully embrace microservices to refactor apps easily. App release cycles were slowed down.

Interviewees said with GKE, developers can take advantage of stateful and stateless apps, serverless apps, and common design patterns with application accelerators. GKE Autopilot further enhanced productivity by enrolling all clusters in a GKE release channel to keep nodes running on the latest qualified versions. Autopilot also provided savings with managed infrastructure and security, eliminating the need for developers to perform duties related to these tasks outside of their job descriptions.

- The lead SRE for the financial technology firm said GKE Autopilot's managing of tertiary aspects of the process which helped developers save time. They said: "when you're implementing a new feature, you don't plan for extra time for maintenance or node outages or security – you just expect it to work."
- The senior leader of IT for the health care organization discussed how GKE Autopilot saved

developers' time and effort. They said: "The resizing, deleting, any sort of cluster inefficiency that we had with our port scaling, that got significantly better with Autopilot and that's where we saw more savings."

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- 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 with GKE Standard improves from 10% in the first year to 30% by the third year as the organization fully takes advantage of GKE's automated scaling and management capabilities. Autopilot improves this further by 15% each year for a total improvement of 45%.
- The fully burdened compensation for a developer within is $125,000, and they redeploy 50% of their time saved to value-add efforts.

Risks. Factors that could impact the size of this benefit for an organization include:

- Organizational agility and CI/CD practices.
- Legacy environment complexity and baseline developer productivity.
- The degree to which managed infrastructure, security, and release channel assistance from Autopilot can eliminate remaining inefficiencies.
- Location and prevailing labor rates.
- The number and type of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 20%, yielding a three-year, risk-adjusted total PV of $9.1 million.

Improved Developer Efficiency

Ref.	Metric	Calculation	Year 1	Year 2	Year 3
C1	Number of developers	Composite	100	200	300
C2	Improved productivity with GKE	Interviews	10%	20%	30%
C3	Additional improvement with Autopilot	Composite	15%	15%	15%
C4	Total improved productivity	C2+C3	25%	35%	45%
C5	Fully burdened developer annual salary	Composite	$125,000	$125,000	$125,000
C6	Productivity capture	Assumption	50%	50%	50%
Ct	Improved developer efficiency	C1*C4*C5*C6	$1,562,500	$4,375,000	$8,437,500
	Risk adjustment	↓20%			
Ctr	Improved developer efficiency (risk-adjusted)		$1,250,000	$3,500,000	$6,750,000
	Three-year total: $11,500,000			**Three-year present value: $9,100,301**	

INFRASTRUCTURE COST SAVINGS

Evidence and data. Interviewees said GKE Standard provides their organizations with four-way cluster autoscaling, which enables them to only pay for what they need by maximizing infrastructure usage while minimizing downtime. Before using GKE, the organizations needed to overprovision resources without precise visibility into cluster usage to avoid exceeding capacity. Interviewees also said GKE Autopilot can further reduce the risk of exceeding capacity by automatically managing pod bin packing and automatically setting preconfigured default values for newly deployed workloads.

- The engineering manager for the media organization told Forrester, "Now that we have moved over to managed service with Autopilot, we have the ability to autoscale for unplanned events."
- The lead SRE for the financial technology organization said Autopilot helped reduce infrastructure costs. They said, "The major difference between Standard and Autopilot is going to be the money that you save from the node breakage or overage costs."

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- The composite organization's baseline infrastructure cost for on-premises container workloads is $5.5 million in Year 1. As the organization and its usage of GKE 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. GKE Autopilot further reduces these costs by 10% for a total reduction of 85%.

Risks. Factors that could impact this benefit for an organization include:

- Baseline infrastructure spend.
- The types of workloads, variability of demand, and fault tolerance.
- The additional savings accrued with GKE Autopilot's automatically managed pod bin packing and preconfigured default workloads.
- The number and type of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of $11.9 million.

Infrastructure Cost Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
D1	Annual provisioning costs prior to GKE	Composite	$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%
D3	Additional reduction in run costs with Autopilot	Interviews	10%	10%	10%
D4	Total reduction in run costs	D2+D3	85%	85%	85%
Dt	Infrastructure cost savings	D1*D4	$4,675,000	$5,610,000	$6,732,000
	Risk adjustment	↓15%			
Dtr	Infrastructure cost savings (risk-adjusted)		$3,973,750	$4,768,500	$5,722,200
	Three-year total: $14,464,450			**Three-year present value: $11,852,583**	

CONTAINER-BASED PAAS LICENSE SAVINGS

Evidence and data. Interviewees said GKE is a fully managed Kubernetes service that automatically provides capabilities their organizations would normally need to pay an additional subscription to operate themselves. The interviewees from both the original study and the new set of interviewees described how investing in GKE made their prior licenses in container-based PaaS redundant. However, because Autopilot is another operational mode for GKE, it did not provide significant impact in this area.

The lead data engineer for a life sciences organization told Forrester, "[GKE] helps us balance costs with operational burden [by making sure that] we only pay for what we use."

Modeling and assumptions. For the composite organization, Forrester assumes the following:

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

Risks. Factors that could impact the size of this benefit for an organization include:

- Legacy third-party PaaS usage.
- The number of application instances.
- A lack of impact of GKE Autopilot on previously existing licenses.

Results. To account for these risks, Forrester adjusted this benefit downward by 20%, yielding a three-year, risk-adjusted total PV of $2.7 million.

Container-Based PaaS License Savings

Ref.	Metric	Calculation	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	↓20%			
Etr	Container-based PaaS license savings (risk-adjusted)		$900,000	$1,080,000	$1,350,000
	Three-year total: $3,330,000			**Three-year present value: $2,725,019**	

IMPROVED SECURITY PRODUCTIVITY

Evidence and data. Interviewees said Google scans GKE-managed clusters as a baseline to discover Kubernetes vulnerabilities and missing patches and automatically begins the patching and release process. They said with GKE Autopilot, clusters have all of these benefits plus fully automated security patching throughout the process. Many of the interviewees said this was by far the most impactful benefit of switching to Autopilot.

- The head of platform engineering for the medical technology organization told Forrester that "Autopilot is more secure by default…the major thing is that Autopilot is patched automatically."

- The engineering manager for the media organization said, “[With Autopilot, platform engineers are] able to shift their focus from security work to provide more business value or implement more robust guard rails or systems or templates that can be reasonable versus the day-to-day toil of operations.”
- The lead SRE for the financial technology organization told Forrester: “As far as security goes, Autopilot makes it so that someone cannot compromise a node and then move horizontally in your stack once they’ve taken over a service. … They can’t just easily sidestep by going down to the VM and up into another container.”
- The assistant professor at the education organization said GKE Autopilot is “insanely secure” and went on to say, “Security is an afterthought because … everything is fully secure.”

“I think the most important thing I can say about Autopilot is that I’m able to run [IT Security] or [End Node] security all by myself for the company, easily… Typically a company would have four to six people doing the work I’m doing.”

Lead SRE, financial technology

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- The composite organization has a team of 20 FTEs dedicated to patching and vulnerability management.
- GKE Standard’s automated scanning and starting of the patching process reduces team effort by 80%. GKE Autopilot fully automating the patching process further reduces team effort by another 15% for a total reduction of 95%.
- Patching teams are located in low-cost regions and team members earn a fully burdened annual rate of $80,000.

“A big difference between GKE Standard and Autopilot is the lack of security personnel from keeping the nodes secure and making sure that the containers don’t have permissions to access the VMs in the case of, for example, an RC.”

Lead SRE, financial technology

Risks. Factors that could impact the size of this benefit for an organization include:

- Team size and location.
- Baseline effort and frequency of patching.
- The number and type of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of $2.7 million.

Improved Security Productivity					
Ref.	**Metric**	**Calculation**	**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	Additional time savings with Autopilot	Interviews	15%	15%	15%
F4	Total time savings for patching	F2+F3	95%	95%	95%
F5	Fully burdened salary for patch and vulnerability management FTE	Composite	$80,000	$80,000	$80,000
F6	Productivity recapture	Assumption	80%	80%	80%
Ft	Improved security productivity	F1*F4*F5*F6	$1,216,000	$1,216,000	$1,216,000
	Risk adjustment	↓10%			
Ftr	Improved security productivity (risk-adjusted)		$1,094,400	$1,094,400	$1,094,400
	Three-year total: $3,283,200			**Three-year present value: $2,721,611**	

AVOIDED INCOME LOSS FROM IMPROVED AVAILABILITY

Evidence and data. Prior to engaging with GKE, the interviewees' organizations lacked the tools in their legacy environments to effectively automate tasks and scale, which reduced availability and increased downtime issues. They said GKE's autoscaling, cluster logging, monitoring, and managed control pane ensure their organizations' services remain available. With the improved security from Autopilot, interviewees reported a slight further improvement to their organizations' operations in this area.

- The head of platform engineering for the medical technology firm told Forrester: "Sometimes, if you have hundreds of clusters, things break, and you have to recreate them. … But since we migrated to Autopilot, we haven't experienced those issues."

Modeling and assumptions. For the composite organization, Forrester assumes the following:

- In its previous environment, the composite organization experienced 500 downtime or degradation events, which grew at 20% with growth of the company.
- The organization improves availability by 97% with GKE. This improves by another 1% by using GKE Autopilot.
- Each downtime or degradation event costs the organization an average of $10,000 in revenue.
- The organization's average margin is 8%.

Risks. Factors that could impact this benefit for an organization include:

- The baseline number of availability events in legacy environment.
- The types of services and related revenue.
- Average margins.

- The degree to which GKE Autopilot further improves stability.
- The number and types of clusters that can be operated in GKE Autopilot.

Results. To account for these risks, Forrester adjusted this benefit downward by 20%, yielding a three-year, risk-adjusted total PV of $935,573.

Avoided Income Loss From Improved Availability

Ref.	Metric	Calculation	Year 1	Year 2	Year 3
G1	Number of downtime or service degradation events per year without GKE	Composite 20% growth	500	600	720
G2	Improvement in availability and service with GKE	Interviews	97%	97%	97%
G3	Additional improvement in availability and service with GKE Autopilot	Interviews	1%	1%	1%
G4	Total improvement in availability and service	G2+G3	98%	98%	98%
G5	Avoided events involving downtime or degradation of service	G1*G4	490	588	706
G6	Revenue impact per event	Composite	$10,000	$10,000	$10,000
G7	Margin	Composite	8%	8%	8%
Gt	Avoided income loss from improved availability	E3*E4*E5	$392,000	$470,400	$564,800
	Risk adjustment	↓20%			
Gtr	Avoided income loss from improved availability (risk-adjusted)		$313,600	$376,320	$451,840
	Three-year total: $1,141,760			**Three-year present value: $935,573**	

UNQUANTIFIED BENEFITS

Interviewees mentioned the following additional benefits that their organizations experienced but were not able to quantify:

- **Improved customer experience.** Better application performance means fewer outages and delays. This, in turn, means fewer headaches and a better experience for customers.
- **Hardened security profile and quick app fixes.** Both GKE Standard and GKE Autopilot provide security benefits alongside the labor savings with automatic patching and quick app fixes.
- **Improved employee experience (EX).** Use of GKE Standard meant an elimination of several manual tasks. Subsequently, this further reduced manual efforts on cluster maintenance and security patching. GKE Autopilot continued this trajectory of superior employee experience by reducing tedious manual tasks and instead letting employees conduct value-add work or improve their skillset by working with cutting-edge technology.
- **Greater scalability and agility to expand into new markets and regions.** GKE's ability to scale can provide opportunities in new markets where public cloud faces different regulatory challenges. The global head of cloud for a financial services

organization said: “In [a] 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.”

FLEXIBILITY

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

- **Deploying additional Google Cloud technologies.** The use of GKE can facilitate expanding other Google Cloud technologies (e.g., machine learning, AI, vertical-specific APIs) within an organization. The engineering manager for the media organization said, “Google is providing roadmaps and flexibility for future cases that we may not have yet.”
- **Avoiding vendor lock-in.** Performing containerization via Kubernetes can help reduce the risk of vendor lock-in or reengineering in the future.

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	$1,255,800	$1,506,960	$1,883,700	$4,646,460	$3,802,310
Itr	Implementation and ongoing labor	$2,805,846	$861,385	$861,385	$861,385	$5,390,000	$4,947,982
	Total costs (risk-adjusted)	$2,805,846	$2,117,185	$2,368,345	$2,745,085	$10,036,460	$8,750,292

GKE ANNUAL CLUSTER MANAGEMENT FEE

Evidence and data. Interviewees from organizations using GKE said their firms incur usage-based fees.

Modeling and assumptions. For the composite organization, Forrester assumes that:

- The composite organization uses GKE Autopilot list pricing. With this pricing, it incurs an annual cost based on the total number of clusters, nodes, node cores, and total usage.
- The composite organization has 50 clusters in Year 1, 60 clusters in Year 2, and 75 clusters in Year 3.
- Each cluster has an average of 50 nodes, with each node having an average of four cores.
- Pricing may vary. Contact Google for more information.

Risks. Factors that could impact this cost for organizations include:

- The number of clusters, nodes, and node cores.
- The ability to run clusters in GKE Autopilot.
- The efficiency of the environment prior to implementing GKE Autopilot.

Results. To account for these risks, Forrester adjusted this cost upward by 15%, yielding a three-year, risk-adjusted total PV (discounted at 10%) of $3.8 million.

GKE Annual Cluster Management Fee						
Ref.	**Metric**	**Calculation**	**Initial**	**Year 1**	**Year 2**	**Year 3**
H1	GKE annual cluster management fee		$0	$1,092,000	$1,310,400	$1,638,000
Ht	GKE annual cluster management fee	H1	$0	$1,092,000	$1,310,400	$1,638,000
	Risk adjustment	↑15%				
Htr	GKE annual cluster management fee (risk-adjusted)		$0	$1,255,800	$1,506,960	$1,883,700
	Three-year total: $4,646,460			**Three-year present value: $3,802,310**		

IMPLEMENTATION AND ONGOING LABOR

Evidence and data. Interviewees said while the actual deployment of GKE is light touch, it still incurs labor costs for planning, design, change management, training, and ongoing management. Interviewees told Forrester that their organizations invested in professional services for planning, design, and training. However, they also said the implementation and ongoing labor of GKE Autopilot is marginal relative to GKE Standard while enabling savings in ongoing maintenance.

- The engineering manager for the media organization told Forrester, “Generally speaking, deploying Autopilot is as quick as deploying Standard as long as you understand the limitations.”
- The lead SRE for the financial technology organization said: “It’s actually easier to implement Autopilot than GKE Standard… Everything works basically the same if you’re not doing anything that requires you to have access to the nodes.”

Modeling and assumptions. For the composite organization, Forrester assumes:

- The composite organization initially invests $1 million in professional services for planning, design, training, and support.
- Initial training is four weeks (160 hours) with one week’s worth of training per year to refresh skills and provide understanding of new features and capabilities. GKE Autopilot requires an additional day of training during both initial training and ongoing training.
- The fully burdened annual rate for an engineer who undergoes training is $120,000.
- The organization only needs 40% of the SRE workforce of GKE Standard. This requires two full-time SREs.
- The fully loaded annual salary of an SRE is $170,000.

Risks. Factors that could impact this cost for organizations include:

- The size and scope of operations.
- Baseline skill sets.
- Organizational agility.
- Prevailing labor rates.
- The ease of the GKE Autopilot deployment.

Results. To account for these risks, Forrester adjusted this cost upward by 10%, yielding a three-year, risk-adjusted total PV of $4.9 million.

Implementation And Ongoing Labor

Ref.	Metric	Calculation	Initial	Year 1	Year 2	Year 3
I1	Professional services		$1,000,000			
I2	Developer/engineer FTEs trained	Composite organization	160	160	160	160
I3	Hours of training required		168	48	48	48
I4	Average trainee salary		$120,000	$120,000	$120,000	$120,000
I5	Dedicated SRE team			2	2	2
I6	Fully burdened SRE salary			$170,000	$170,000	$170,000
It	Implementation and ongoing labor	I1+(I2*I3*(I4/2080))+(I5*I6)	$2,550,769	$783,077	$783,077	$783,077
	Risk adjustment	↑10%				
Itr	Implementation and ongoing labor (risk-adjusted)		$2,805,846	$861,385	$861,385	$861,385
	Three-year total: $5,390,000			**Three-year present value: $4,947,982**		

Financial Summary

CONSOLIDATED THREE-YEAR RISK-ADJUSTED METRICS

Cash Flow Chart (Risk-Adjusted)

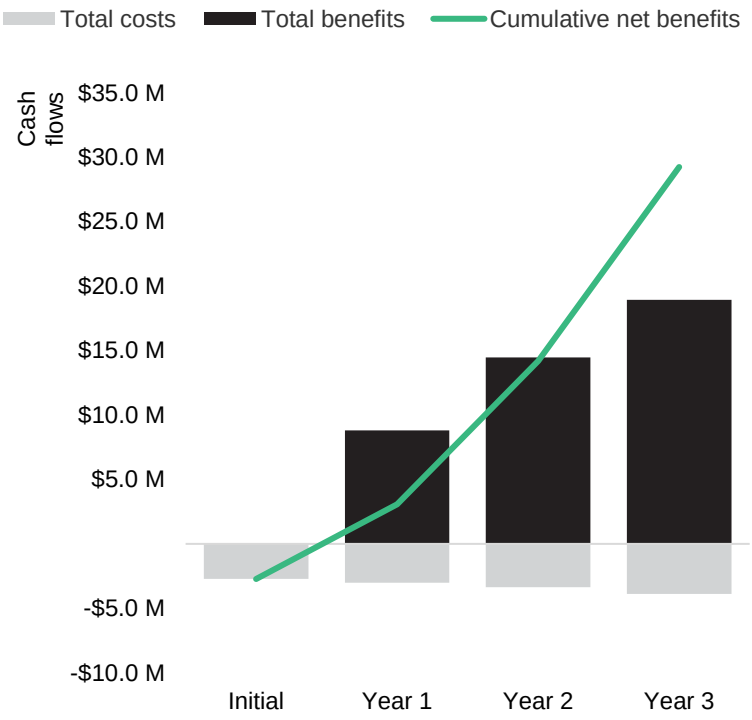

The financial results calculated in the Benefits and Costs sections can be used to determine the ROI, NPV, and payback period for the composite organization's investment. Forrester assumes a yearly discount rate of 10% for this analysis.

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	($2,805,846)	($2,117,185)	($2,368,345)	($2,745,085)	($10,036,460)	($8,750,292)
Total benefits	$0	$10,534,212	$14,472,912	$19,673,363	$44,680,487	$36,318,530
Net benefits	($2,805,846)	$8,417,027	$12,104,568	$16,928,278	$34,644,027	$27,568,238
ROI						315%
Payback period (months)						<6

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: 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®