




From infrastructure to innovation

How Google Cloud enables you to shift your focus from maintaining infrastructure to driving business outcomes faster

 Google Cloud

Contents

Streamlining application delivery gets you to market faster	1
Making the most of your data is a competitive advantage	10
Getting things done faster and easier saves you money	22
Google's business, technical and financial partnership helps you grow	24



Google Cloud



Why are 80% of [Fast Company's top 10 most innovative companies of 2020](#) using Google Cloud?

And more than 90% of [Forbes top 20 digital companies](#)?

And 80% of [Y Combinator's top 10 companies](#)?

Tech companies and startups tell us they prefer Google Cloud because our industry-leading app development and smart analytics services help them get things done faster and easier. Google Cloud reduces and eliminates the need for infrastructure maintenance, custom tooling and operations. This frees up engineering to focus on roadmap priorities that drive business outcomes like competitive advantage, growth and increased profitability.

Yes, you really can focus on innovation and let Google Cloud take care of the rest. Let's explore how.



“Google Cloud removes a lot of the operational complexity from our ecosystem. That frees up time. We can iterate quicker on key needs, like data insights and machine learning. Having infrastructure managed for us, with the lower-value details taken away, streamlines our ability to concentrate on what’s important to our users and give them the experiences they know and love about Spotify.”

Tyson Singer,
VP of Technology and Platform, [Spotify](#)

Streamlining application delivery gets you to market faster

Speed is the new currency of business. That's why tech companies and startups are choosing Google Cloud's state-of-the-art containers and microservices: to get from code to production in record time.

By running on Google Cloud, fin-tech startup [Current](#) improved time to market for app development by 400% while eliminating downtime for users of its debit card app. [Idwall](#) improved developer productivity by 30% and [Rakuten Viki](#) freed up engineers from managing infrastructure to focus on scaling the business to serve rapidly growing user demand.

Let's take a closer look at how you can speed up your application delivery.





Get speed and flexibility from containers

Containers give you more flexibility and speed than virtual machines, and Google Cloud takes the hard part - maintenance and operations - out of the equation.

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

- Refactoring/re-architecting an MVP into a scalable solution
- Expanding into additional cloud providers to meet regulatory or customer expectations
- Expanding geographically to reduce latency and provide better experiences to a global customer base
- Improving your end-to-end security posture

Or maybe you want to improve the existing customer experience (eg. service availability) or launch a differentiated experience that will take your company to the next level, but the project is taking longer than you'd like.

Legacy or technical debt can slow you down. But a modern architecture consisting of managed services and containers gives you access to proven patterns for running reliable and secure infrastructure. This helps speed up your time to market and increase your productivity without sacrificing stability and security - with the added benefit of helping you attract the best technical talent to work on innovation.



Use Kubernetes to reduce infrastructure management

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

But you should also be worried about locking yourself out. You're missing out on the innovation of the Kubernetes community and surrounding ecosystem that's setting current industry standards and best practices.

It comes down to where you want your engineers to spend their time -- maintaining infrastructure and building and maintaining custom tooling or ticking off your priority list to drive your business forward. What you have today might be working, but your roadmap probably includes things like repaying technical debt and filling in platform gaps:

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

What are your choices? You can learn and build all the cloud-native patterns yourself from scratch, or you can leverage Kubernetes and shift your creative focus to building your own products and services.

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



Choose Google Kubernetes Engine for a fully managed service


If you're on Kubernetes, how do you currently do configuration management and deal with end-to-end encryption or multi-tenancy? Maybe you've attempted to run it yourself, or tried a managed service like EKS that fell short on the managed part?

Kubernetes is only as good as the infrastructure it runs on, and Google Kubernetes Engine (GKE), the most mature container orchestration service available today, fills the gaps as a fully managed Kubernetes offering. At Google Cloud we believe in containers as the future of packaging and runtime for new and existing applications. While other cloud providers are dividing their attention with multiple container orchestration solutions (eg. Amazon has EKS, ECS and Fargate), Google Cloud has one container platform - Google Kubernetes Engine (GKE). We're all in and it shows.

Since Google donated Kubernetes to the developer community in 2015, GKE has become the best way to consume the cloud. It's the de facto standard for organizations looking for advanced container orchestration, delivering the highest levels of reliability, security, and scalability. IaaS offerings are fully integrated - everything from VM provisioning, autoscaling, and upgrades to creating and managing GPUs, TPUs for machine learning, storage volumes, and security credentials on demand. You simply put your application in a container and then choose which system you want to use based on what flexibility you need.

And for workloads that don't require a lot of control over cluster configuration, like serving up a website or API, you can eliminate cluster administration while optimizing security. When you use GKE Autopilot, GKE provisions and manages the cluster's underlying infrastructure for you, including nodes and node pools, and you only pay for the workload, not the cluster, which can save you a lot of money.

GKE's simplicity and advanced features increase developer and platform team productivity, so you can focus on getting to market faster.





For example: What if you want to process work from a message broker, where each of the message queues need to be dynamic in scaling, and store the results in a database?

If you're currently running on EC2, you've likely signed up to maintain infrastructure for each set of VMs, and you've built some custom tooling/automation on the side but chances are it's ad hoc, non-portable and also requires maintenance.

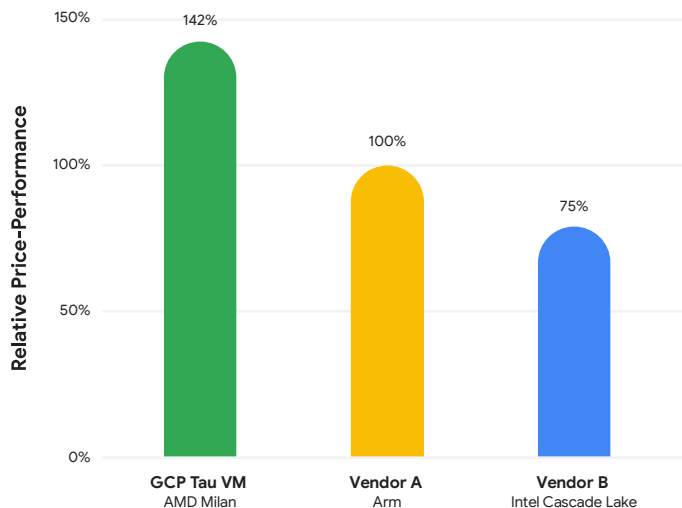
When you use GKE, you can consolidate these workloads into a single Kubernetes cluster by containerizing those applications and using Kubernetes' built-in orchestration to manage them, replacing custom scripts and tooling with RabbitMQ Operator and autoscale with Rabbit MQ Autoscaler. At the end of the day, thanks to GKE's ability to pack workloads from underutilized machines, you have something that's more streamlined, portable, and cost-efficient.

Get leading price-performance with GKE

GKE is the best place to build a modern application, but the compute underpinning it is equally important. Scale-out workloads demand the best combination of performance and price to bring down the cost of delivering applications, all while providing an excellent user experience.

We recently [announced](#) a new virtual machine (VM) family coming to Google Cloud. [Tau VMs](#) extend Compute Engine’s VM offerings with a new option optimized for cost-effective performance of scale-out workloads. T2D, the first instance type in the Tau VM family, is based on 3rd Gen AMD EPYCTM processors and leapfrogs the VMs for scale-out workloads of any leading public cloud provider available today, both in terms of performance and workload total cost of ownership (TCO). The x86 compatibility provided by these AMD EPYC processor-based VMs gives you market-leading performance improvements and cost savings, without having to port your applications to a new processor architecture.

You can easily add Tau VMs to your GKE clusters by specifying the T2D machine type in your GKE node-pools, which helps you optimize price-performance for your containerized workloads.



New Tau VMs

56% higher absolute performance

42% higher price-performance

Compared to general-purpose VMs from any of the leading public cloud vendors

*Results are based on estimated SPECrate@2017_int_base run on production VMs of two other leading cloud vendors and pre-production Google Cloud Tau VMs using vendor recommended compilers. [View testing details here.](#)

SPECrate is a trademark of the Standard Performance Evaluation Corporation. More information available at www.spec.org

Developers find GKE easier and faster to use than other Kubernetes services.

See for yourself!

Create a Kubernetes cluster with just three clicks using the [GCP console](#).

Google Kubernetes Engine

The industry's most mature and configurable container orchestration service lets you:

- **Scale seamlessly, automatically and high:** GKE's industry-first four-way autoscaling automatically resizes your cluster based on the demands of your workloads—no manual provisioning or overprovisioning required. It also supports an industry-leading 15,000 nodes in a single cluster, handling whatever workload you can throw at it.
- **Minimize management overhead:** Autopilot, a GKE configuration option, provides automated management of your Kubernetes infrastructure, including both control nodes and user nodes, and applies all Google-known secure settings by default.
- **Integrate with developer toolchains and processes:** Tight integration with Google Cloud's built-in DevOps toolchain makes it easy for developers to automatically build, test, deploy, and manage code changes, so applications are built and improved faster.

[17 Media](#) increased their development cycle by 80% using GKE and is now releasing daily rather than every one to two weeks.

Google “leads the pack with unparalleled Kubernetes expertise.”

Forrester New Wave™: Public Cloud Enterprise Container Platforms, Q3 2019



Adopt Cloud Run to run containers without infrastructure

Whereas GKE integrates with Google Cloud, [Cloud Run](#) turns GCP into platform-as-a-service, giving you the power and flexibility of containers and the convenience of serverless.

You can eliminate infrastructure - no server or cluster configuration or maintenance - and focus instead on your applications. And you can go beyond functions and run real-world applications using your favorite languages and frameworks, backed by open industry standards, including OCI container images and knative, so you're not locked in.

Cloud Run lets you, with minimal complexity, start running containers to a fully managed environment with the best of Google Cloud's security, performance, scalability and best practices baked in. As the load on your application changes, for example, Cloud Run automatically increases or decreases the number of running container instances to handle it, including bringing the number (and thus your cost) to zero when it's warranted. You can scale applications to zero, or avoid cold starts and additional latency by keeping applications always running and ready when you need them.

If you want to run traditional workloads, you can do so on a modern platform with GKE. For more cloud-native applications, Cloud Run can do the same thing as GKE but with less work. So you can start with Cloud Run in order to move quickly, while knowing that if you grow out of it and need custom infrastructure, you can move to GKE without reworking your application.



Cloud Run



Develop and deploy highly scalable containerized applications on a fully managed serverless platform:

- **Get from container to production in seconds:** Cloud Run is the fastest way to write, run, and manage applications.
- **Eliminate infrastructure management:** Focus on code while we handle infrastructure management, including security, reliability, redundancy, and availability.
- **Write your code your way:** Cloud Run lets your developers use any language (Go, Python, Java, Ruby, Node.js, and more), any library, any binary. You aren't limited to whichever languages are supported by a given cloud platform's FaaS technology.
- **Use your containers in other environments:** There's nothing Google-specific about the containers in Cloud Run; they'll work unchanged in other environments. They fit into the standard container ecosystem, for example, including the local Docker environment on your desktop and your CI/CD workflow. You can also run them in GKE or in Kubernetes on another cloud platform.

“Cloud Run promises to dramatically reduce the operational complexity of deploying containerized software. The ability to put an automatically scaling service in production with one command is very attractive.”

Jamie Talbot,
Principal Engineer at Mailchimp





Making the most of your data is a competitive advantage

How are you making decisions today? How are you using data to improve the customer experience? Where are your data engineers and scientists spending their time?

Data is crucial to driving innovative product direction and user experiences along with broad go-to-market decisions. If you can harness your data successfully, you can achieve a significant competitive advantage. That's why most technical leaders are under tremendous pressure to do more - modernize and operate at larger and larger scales, justify current and future data costs, and elevate their organizational maturity and decision-making.

But there are challenges with access, storage, inconsistent tools, compliance and security that make it hard to go below the surface and unlock real value from your data.

Maybe you've inherited legacy systems that you're trying to marry with new ones. Should all your data be in one cloud? Across multiple clouds? How do you modernize analytics stacks that have historically been vertically integrated to platforms that can scale horizontally?

Or, rather than processing data in real time, maybe you're currently batching or micro-batching your data. This adds complexity to your architecture with an orchestration system and scheduling and requires maintenance around contention and resilience. The operations overhead from managing and maintaining the batch architecture is expensive, and you're still compromising on data latency.

Lacking easy access to all your data and the ability to process and analyze it as it comes in puts you at a disadvantage. The modern tech stack needs to be a streaming tech stack that keeps up with the scale of your data, uses the most recent data that's available, and incorporates and understands unstructured data. And the most advanced analytics teams have shifted their focus from operation to action, using AI/ML to experiment and operationalize processes.

That's where Google Cloud's data analytics and AI capabilities come in. We offer a comprehensive set of services, from data management to analysis to providing the right information in the right moment to take action. Google's Data Cloud is designed to be simple and powerful and span the vast majority of customer needs. It's why so many technology companies and startups rely on Google Cloud - to reduce infrastructure management so they can focus on using data to drive competitive advantage.

Google is a Leader in The Forrester Wave™: Cloud Data Warehouse, Q1 2021, and the The Forrester Wave™: Streaming Analytics, Q2 2021



“Migrating from Redshift to BigQuery has been game-changing for our organization. We’ve been able to overcome performance bottlenecks and capacity constraints as well as fearlessly unlock actionable insights for our business.”

Spencer Aiello,
Tech Lead and Manager, Machine Learning
at Discord

“With our new data pipeline and warehouse, we are able to personalize access to large volumes of data that were not previously there. That means new insights and correlations and, therefore, better decisions and increased revenue for our customers.”

Jean-Yves Simon,
VP Product, AB Tasty

Google Cloud

Choose simplicity AND scalability in your data platform

Chances are you have a lot of data at your disposal right now, it's growing exponentially and you want to maintain or increase ROI while keeping up with volume. Maybe you anticipate how much data you expect to have in the future (eg. a terabyte) and design your systems to process that amount while knowing that if it grows beyond what you expect, you're looking at a wholesale system migration. Or you choose a data warehouse that can scale to your anticipated future data size, but increasing processing needs make it complex to manage.

Smaller systems have traditionally been simpler, but today, with platforms like Google BigQuery, you don't need to choose between ease of use and highly scalable systems. BigQuery's serverless architecture eliminates the need for cluster management, and its ability to handle massive scale for both compute and storage means you don't have to worry about the size of data exceeding your technical capacity again.



BigQuery automatically scales to thousands of cores in seconds, letting you run analytics at scale with up to 34% lower three-year TCO than [AWS Redshift](#), [Azure SQL Data Warehouse](#) and [Snowflake](#). All while ensuring full business continuity for your organization with no planned downtime with the industry-leading uptime SLA of 99.99%.



Decouple storage and compute to stay flexible and keep costs down

Any data management system that combines compute and storage means you'll have to scale up compute to deal with increasing data volumes, even if you don't need it, which can be expensive. So you might find yourself making compromises, like only storing the last 12 months of data in your analytics warehouse. You might also choose not to include data because you don't have an immediate use case for it, only to find that you're unable to test a hypothesis down the road because the data isn't there, and would require a new pipeline to get started.

Other systems get you halfway there, letting you scale and pay for compute and storage independently, but still require you to manually set up, scale and optimize clusters. If you don't want to spend your time managing infrastructure, Google Cloud's serverless, multi-cloud data warehouse, BigQuery, offers you enhanced reliability, performance and built-in data protection.

Beyond cost and management, our customers also care about flexibility. When your data changes, how long does it take for you to notice and react? Queries on systems like Redshift have to be optimized to be efficient, and that limits the amount of experimentation you can do, so you might only extract and pull in data when you suspect there might be a problem. The compromises you make due to the lack of compute/storage separation and the need to optimize your data warehouse tie one hand behind your back.

With Google Cloud, you don't need to plan queries in advance or index your data sets. Because we decouple storage and compute, you can land data without worrying it's going to drive up your querying costs, and your data scientists can experiment without having to worry about clusters or sizing their data warehouses to try new ideas through ad hoc queries.



Make data-driven decisions in real time

BigQuery's native support for ingesting streaming data makes it easy to capture data in real time without worrying about managing your own streaming buffer, and makes that data immediately available for analysis using SQL and low latency querying using BI Engine. This means your team can be vastly more efficient, react in real time, and keep up with the speed of your business.

Along with BigQuery's easy-to-use Streaming API and BI Engine, you can use Cloud Pub/Sub with Cloud Dataflow to build even more scalable and resilient streaming pipelines. Cloud Pub/Sub's serverless event messaging system makes it simple to publish and consume asynchronous messages with no managerial overhead, while Dataflow's serverless approach means you no longer need to manage server clusters to run your data pipelines, removing operational overhead from data engineering workloads. Resource autoscaling paired with cost-optimized batch processing capabilities means Dataflow offers virtually limitless capacity to manage your seasonal and spiky workloads without overspending.



Break down data silos

Many tech companies and startups create silos by storing data separately across departments and business units, with each team owning their data. This means that whenever you want to do analysis that spans departments, you have to identify some way to break down those silos, probably by running extraction (ETL) pipelines to get data and land it into your data warehouse. But departments that own the data often have little incentive to maintain it (the semantics, etc) in the pipeline, so over time it becomes out of date and less useful.

Beyond organizational silos, many companies today have adopted a multi-cloud strategy based on departmental preference, capability alignment, and regulatory pressure. Many companies also deal with the reality of legacy data lake and data warehouse investments that reside on-premise. Today's multi-cloud, hybrid-cloud reality requires another level of sophistication in managing and accessing siloed data.

Moving to a centralized warehouse or a comprehensive data fabric speeds up your ability to get access to high quality data across departments, clouds, and on-premise systems that can be used to solve specific business problems (like product performance or customer behaviour) and gives you the agility to query it on the fly.

BigQuery provides a central location from which users across the organization can manage, secure, access, and share data assets and insights. For example, you can land all your data in BigQuery and provide reusable functions, materialized views and even the ability to train ML models without any data movement. This means even non-technical domain experts, and partners and suppliers with permission, can easily access and use SQL to query the data using familiar tools like spreadsheets and dashboards.

The 'hub and spoke' analogy is appropriate here. BigQuery is the hub that contains your data. Then there are spokes - reporting tools, dashboards, ML models, web applications, recommendation systems etc - reading the data live from BigQuery without having to make a copy of the data. Looker, for example, helps you visualize your data and integrate it into users' daily workflows. This approach lets you improve your data's usability, security and quality, all at the same time.

Simplify access to all your data

Historically, unstructured and semi-structured data was best served by data lakes, while structured data fit best in data warehouses. This separation created technological silos that made crossing the format divide difficult; you'd store all your data in a data lake because it's cheaper and easier to manage, then move to a warehouse so you could use analytics tools to extract insights.

The increasingly popular "lake house" merges these two worlds into a unified environment for all types of data. With Google Cloud, you can use BigQuery as both your data warehouse and your data lake. Using BigQuery's Storage API you can access storage directly to power workloads usually associated with the data lake. Because data can be stored in BigQuery as a single source of truth, fewer copies of data need to be created and maintained. Instead, downstream processing can be done via SQL transformations that are stored in logical views, without having to move data around.

Ease of use matters. If you can get results from your queries in a centralized location in 30 seconds, rather than 30 minutes or 3 hours, you're more likely to use data to aid your decision-making.



Use AI/ML to experiment faster and operationalize workloads

How fast are your data scientists able to experiment?

Chances are they need to stop development and operationalize in order to evaluate their experiments. They develop and iterate on a model using historical data, then hand the model off to the engineers, who often completely rewrite it to incorporate into the production system and test on new data. Then they wait, iterate on their model, and productionize again. This cycle involves a lot of stop-and-go and re-writing code, with coordination between teams and errors introduced along the way. The result is that your data scientists aren't experimenting as much as they could be because they aren't agile; it can take a long time to conduct experiments, and it makes it hard to predict how long a project will take and whether it will be successful, let alone how long it will take to get it into production.

If you're serious about differentiating based on data, you want to extract the highest value you can from the data you're collecting. To do that, you want your data science teams to be as productive as possible and not miss the chance to build a model because infrastructure is too hard. So the quality of pre-built platforms is critical to avoiding interruptive, ad hoc development.

The key to agility is to do end-to-end experiments early and often. Google Cloud Vertex Pipelines give you a history of experiments so you can look back, compare against benchmarks and endpoints and A/B test with shadow models. Because the code is containerized, the same code can be used between development and production systems. Data scientists work in Python, and the production engineers get containers that are fully encapsulated. Both teams can standardize on operationalizing the models with Vertex predictions - and you can move quickly.

In fact, using BigQuery ML, domain experts can often test the feasibility of an idea by training custom models using only SQL without needing additional experience with traditional data science tools. This means you can experiment in a production-like system and conduct feasibility studies in days instead of months. Using Looker, you can also easily create consistent data models on top of all your data and use LookML to query data, so everyone in the organization can create easy-to-read reports and dashboards to explore patterns of data.

To drive real value in production, systems must be able to ingest, process, and serve data, and machine learning must drive personalized services, based on the customer's context and in real time. But a continuously running production application demands that models be constantly retrained, deployed and checked for security. And for incoming data, you have to perform data preprocessing and validation to make sure there are no quality issues, followed by feature engineering and model training with hyperparameter tuning.



To make it easy to orchestrate and manage these multi-phase machine learning workflows and run them reliably and repeatedly, Google Cloud has integrated data science and machine learning. MLOps tooling and automated workflows enable rapid, continuous delivery and simplify management of models to production. There's a single workflow and vocabulary for all our AI products, regardless of the layer of abstraction, and you can easily interchange custom and AutoML-trained models as they're now leveraging the same format and technical foundation.

For example: What if you want to apply anomaly detection to live, unbounded data streams to combat fraud? With Google Cloud, you'd generate a sample data stream to simulate common network traffic and ingest it to Pub/Sub, then create and train an anomaly detection model in BigQuery using BigQuery ML K-means clustering and Mask PII using Data Loss Prevention (DLP). Afterward, you'd apply the model to live data for real-time detection using Dataflow and use Looker to create a dashboard, alerts, and actions to handle the identified events.



BigQuery

Setting up an advanced data analytics environment doesn't have to be hard, with Google Cloud's comprehensive suite of data analytics tools:

- **Rapidly analyze structured and unstructured data:** BigQuery seamlessly integrates with open-source business tools and data science, and you can securely share real-time insights within your organization to inform decisions and accelerate innovation.
- **Run predictive and real-time analytics:** Stream data directly into BigQuery for real-time analytics and operationalize machine learning models right in the data warehouse using BigQuery ML and without the need to move data.
- **Reduce your IT overhead:** As a serverless data warehouse, BigQuery lets you get up and running in minutes and run petabyte-scale queries, without ever having to worry about system provisioning, maintenance, fault tolerance or performance tuning.




Google Cloud

“We really experienced a difference with BigQuery, compared to other data analytics solutions we had tried before. It has sped up the data analysis process with a 10-fold increase in credit risk model deployment speed. We don’t see slowdowns in performance or issues such as bottlenecks, even when running lots of data.”

Palm Phuwarat

Chief Product and Data Officer, Credit OK



20



“The combination of Looker and Google Big Query is powerful, allowing us to get data-hungry analysts essential information much faster. After we moved to Google BigQuery, query time was reduced exponentially. It’s an astonishing difference, allowing us to run 300 queries per day.”

Sam Chase,

Tech Lead, Data Operations, Blue Apron



Google AI/ML

Developers and data scientists alike can use a rich set of Google Cloud ML services to make sense of your data and tackle your hardest problems:

- **Access pre-trained models for language translation, image recognition, and many other areas, all via API.** Your developers get the benefit of machine learning without having to hire hard-to-find data scientists. Google Cloud's models are widely viewed as more accurate than those of other cloud platforms. For example, Google Cloud's models support more language variants than any other, so your startup can be global from day one.
- **Create new ML models quickly, increasing time to insight or market:** Data scientists and developers can create ML models in BigQuery using a variant of familiar SQL and train models directly where your data is stored, without moving data out of a data warehouse. Cloud AutoML also saves you time by automating the model searching process.
- **Leverage flexible, advanced ML tools.** For your team's data scientists, AI Platform includes support for deep learning, distributed model training, model hosting, and more.
- **Do deep learning with TensorFlow and Cloud TPUs.** As the creator of this open-source framework, we drive today's broadest and deepest TensorFlow ecosystem, including hosting the TensorFlow Summit. Google Cloud also lets you train your models on Google tensor processing units (TPUs), which provide a speed advantage for running TensorFlow.

“Offloading some of the heavy lifting in developing machine learning models to Cloud AutoML is a huge accelerator, because the work can be done in days instead of weeks.”

Erik Andrejko,
Co-founder and CTO, wellio

Getting things done faster and easier saves you money

Being able to get things done faster and easier with Google Cloud can result in significant savings across your development, operations and infrastructure, whether its money-saving recommendations, competitive pricing for compute, network and storage, or tools that are inherently more efficient:



Efficiency with infrastructure that instantly scales up and down in response to traffic and growth. Startups surveyed by IDC saw a 41% increase in efficiency when using Google Cloud across development, operations and infrastructure.



Billing transparency to control costs with the ability to forecast and optimize spend with AI-powered recommendations and billing reports in the Cloud Console.



A simpler pricing model. Other cloud providers require up-front commitments, such as a promise to use five VMs for a year, in order to earn discounts. Google Cloud doesn't do this. Instead, sustained use discounts automatically reduce the per-second cost of a VM the longer you run it. You don't need to make an upfront commitment to save money. Further, preemptible resources are an easy way to run non-time-sensitive jobs cost effectively.




Right-sized compute. Compute Engine custom machine types allow you to rightsize your resources so that you have exactly the right amount of CPU and memory that you need—no wasted resources!



Serverless economics. Cloud Run, Firebase, BigQuery, and many other Google Cloud services are all serverless, letting you pay only for what you use and scaling down to zero when you can.

Tech companies and startups around the world are seeing real savings with Google Cloud:



Kapten saw

2x

Sustained volume for the same price, by adjusting nodes and pods.

Current saw

7x

Growth in user base

60%

Reduced total cloud hosting costs



Takopedia saw

60%

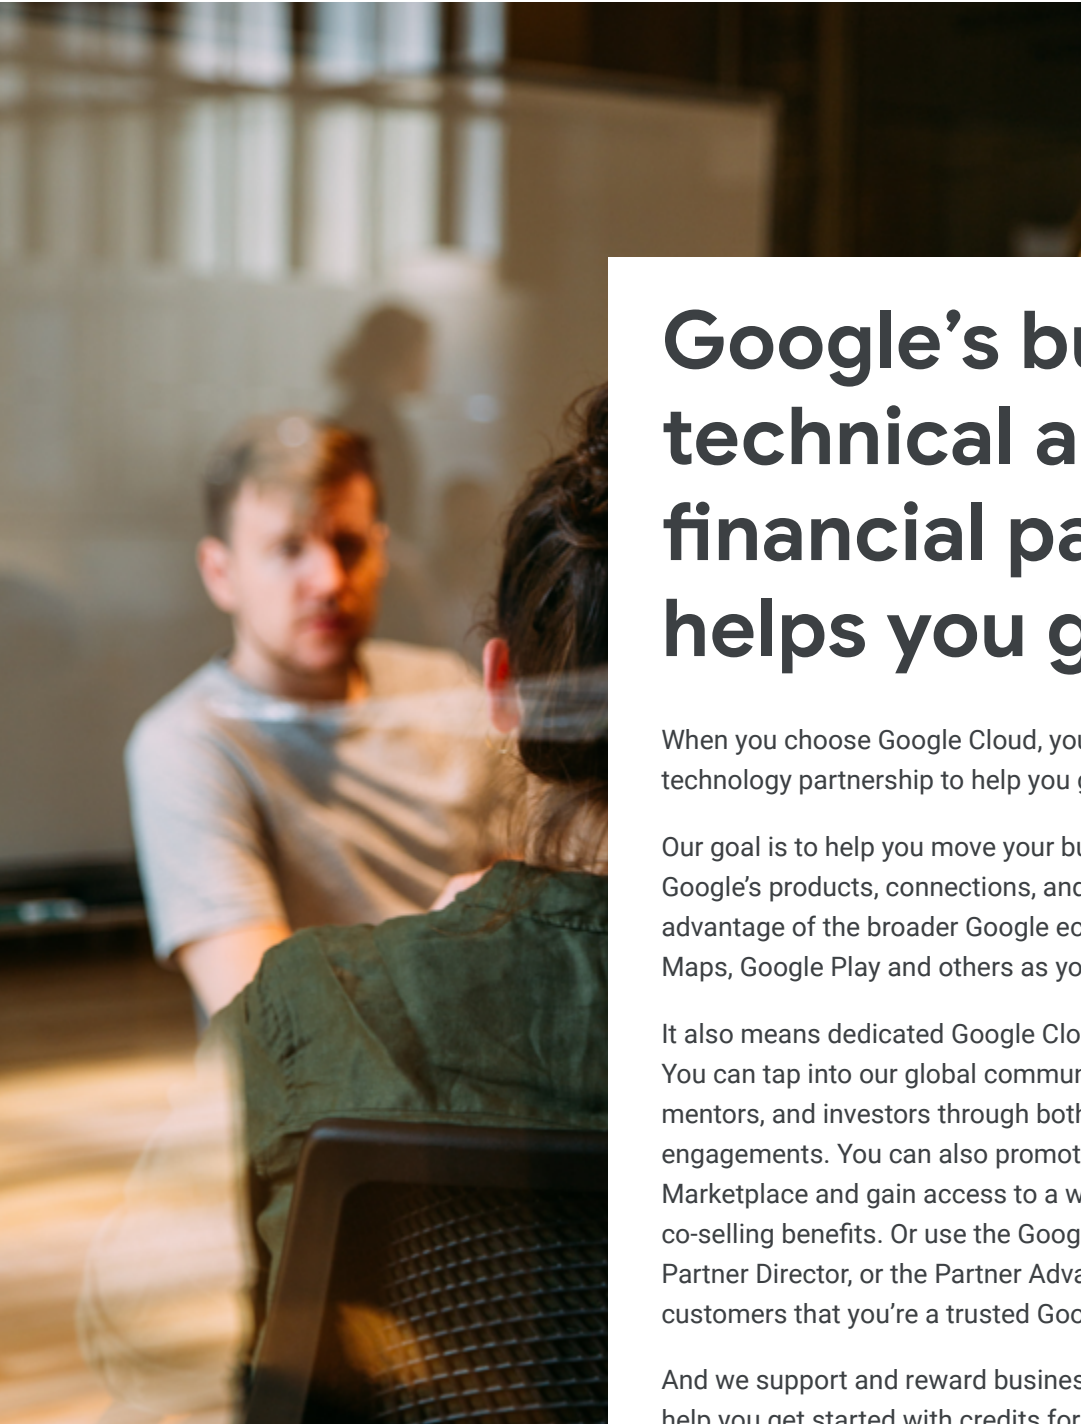
Reduction in error resolution times

30x

Automated scaling

90%

Reduced operating costs



Google's business, technical and financial partnership helps you grow

When you choose Google Cloud, you're choosing a business and technology partnership to help you go to market.

Our goal is to help you move your business forward with the best of Google's products, connections, and best practices. That means taking advantage of the broader Google ecosystem like Google Ads, Google Maps, Google Play and others as you grow and develop your business.

It also means dedicated Google Cloud go-to-market opportunities. You can tap into our global community of founders, technical experts, mentors, and investors through both one-on-one and large scale engagements. You can also promote your solution in the Google Cloud Marketplace and gain access to a wide range of co-marketing and co-selling benefits. Or use the Google Workspace Marketplace, the Partner Director, or the Partner Advantage badge to show prospective customers that you're a trusted Google Cloud partner.

And we support and reward businesses built on Google Cloud. We help you get started with credits for Google Cloud, Firebase, Google Workspace, Marketplace solutions, and cloud support plans. And as you grow with us, so do the incentives – like marketing benefits, dedicated support, and more.



Google Cloud

On the technical front, Google Cloud experts can help you get the most from your implementations, so you can build better products and scale more efficiently from the get-go. Google Cloud offers technical training resources and workshops to upskill your team, strategic planning and implementation guidance from Google engineers, and help with piloting, planning and migrating your cloud infrastructure.



“We saw a difference with Google Cloud from the very beginning because the interactions felt like a strategic relationship. Google gave us startup credits and a lot of face-to-face support, which we hadn’t experienced with other cloud providers.”

Jim Hanifen,
Head of Product, [Brandfolder](#)

“Working shoulder-to-shoulder with Google Cloud and partner engineers in an everything-is-code environment enabled us to easily build new products and features, adapt to changing customer, market or business circumstances, and move on quickly to new projects and activities.”

Firman Gautama,
VP Infrastructure & Security, [Tiket.com](#)



It's time to move

Innovate faster with Google Cloud, technology that makes it faster and easier for you to get things done and a partner committed to your success. Don't get left behind - join the tech companies and startups that are choosing Google Cloud. [Contact sales now.](#)