

Analyzing the Economic Benefits of Google Cloud Spanner Relational Database Service

By Tony Palmer, Senior Validation Analyst; and Alex Arcilla, Validation Analyst
February 2019

Executive Summary

On-premises relational database management systems (RDBMS) are the traditional heart of many core business operations, but cloud-based relational databases provide the agility and scalability of the cloud, while eliminating many of the onerous tasks associated with maintaining your own infrastructure.

Google Cloud Spanner is a scalable, globally distributed, strongly consistent, highly available, cloud-native relational database platform that provides organizations with a cost-effective, high-performance solution that is easy to deploy, scale, and manage. ESG validated that Cloud Spanner provides cost-effective and agile managed databases that can easily be spun up and down as required and can be optimally configured for individual workloads by development teams. Customers reported that there were significant cost benefits realized from shifting on-premises RDBMS environments to Cloud Spanner, and substantial benefits compared to other cloud-based databases.

This total cost of ownership (TCO) analysis compares an on-premises RDBMS strategy against hosting the same data on Cloud Spanner. In the process, we were able to gain insights comparing Cloud Spanner and other cloud service providers from customers who have experienced using multiple cloud environments.

Our analysis shows that hosting the same data using Google Cloud Spanner is 78% less expensive than using on-premises servers and up to 37% less expensive than using other cloud databases. Additionally, we found dramatic benefits in the ability of customers to enable app development teams to directly manage their own databases without needing IT assistance in planning, procuring, deploying, or managing the environment. This ability enabled greater flexibility and faster innovation in responding to the needs of the business. While cost is always a driving factor, every customer ESG interviewed as part of our research identified business and operational benefits that far outweighed the cost benefits.



Introduction

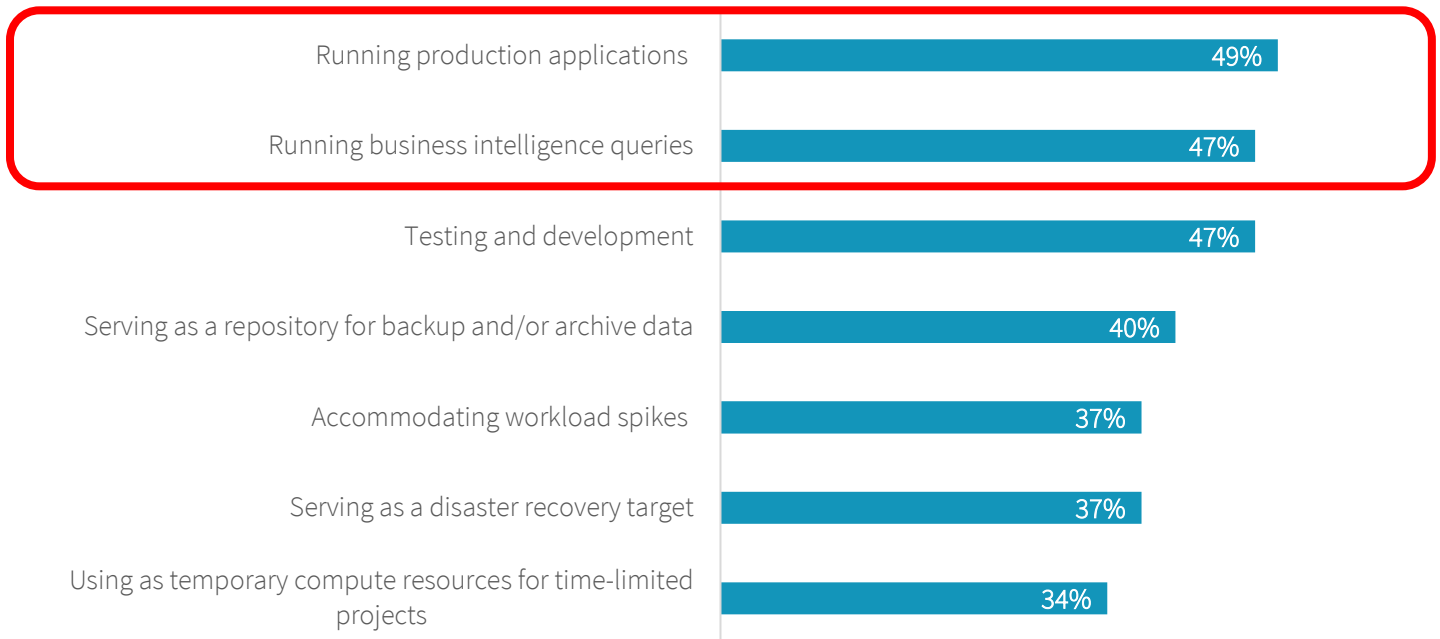
This ESG Economic Value Audit focused on the quantitative and qualitative benefits organizations can expect from leveraging Cloud Spanner instead of deploying on-premises relational database systems like Oracle or MySQL, or leveraging other cloud-based managed database solutions like AWS DynamoDB or Azure CosmosDB. ESG created a modeled scenario that factored in cost of servers, storage, software, support, maintenance, and administration over a three-year period.

Challenges

The database market continues to evolve rapidly, no longer dominated by just a handful of traditional approaches. One of the biggest impacts has been the availability—and increasing popularity—of public cloud options. Businesses are finding new opportunities in cloud-based databases, and not only for dev/test or transient workloads. Over 80% of businesses surveyed by ESG currently use or expect to use public cloud-based production databases. ESG’s 2019 Technology Spending Intentions Survey asked organizations how they have used public cloud infrastructure, and 49% of respondents cited running production applications, with 47% stating business intelligence/data analytics. Respondents could select more than one response. (see Figure 1).¹

Figure 1. Cloud Infrastructure Trends

For which of the following purposes does your organization use cloud infrastructure services (IaaS and/or PaaS)? (Percent of respondents, N=438, multiple responses accepted)



Source: Enterprise Strategy Group

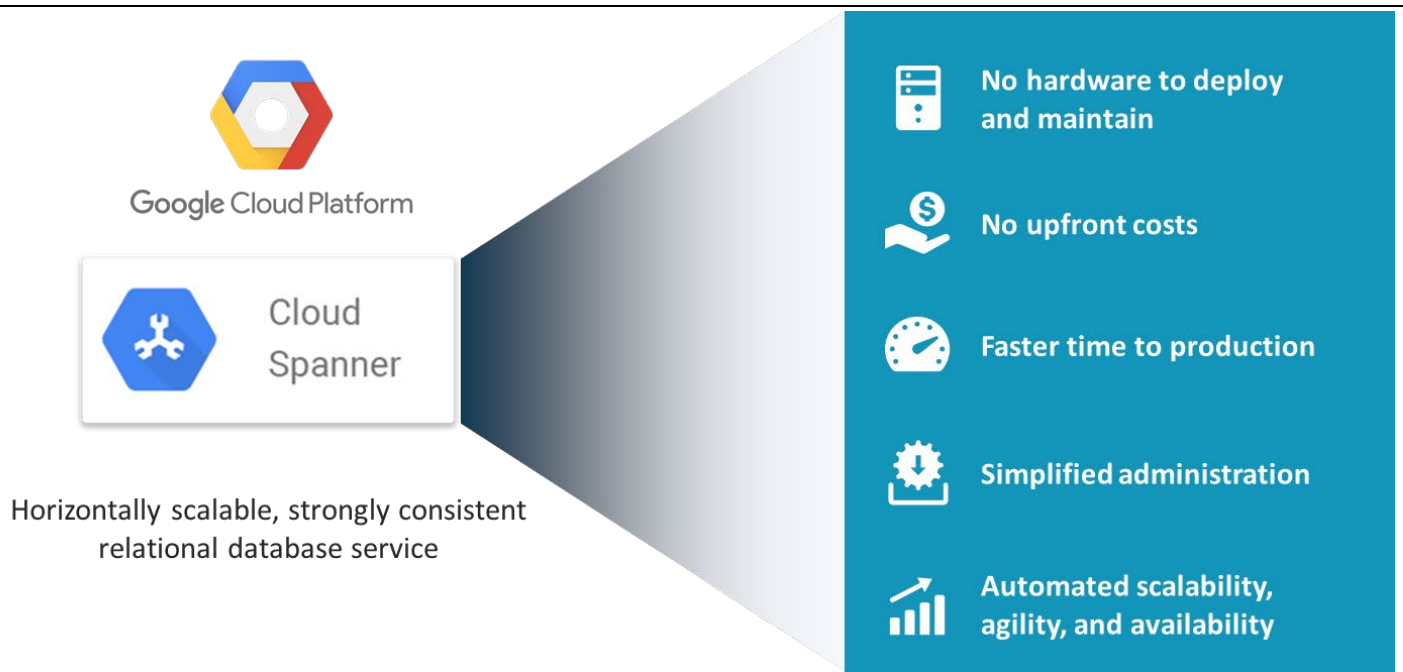
While cloud has been popular for non-production and temporary workloads for some time, its adoption for full production use cases represents a significant attitudinal shift. This trend is unlikely to reverse course soon, as a variety of cloud service providers offer increasingly strong alternatives and deliver enterprise (or better) quality at competitive prices.

¹ Source: ESG Research Report, [2019 Technology Spending Intentions Survey](#), February 2019.

The Solution: Google Cloud Spanner

Cloud Spanner is a globally distributed, strongly consistent database service built specifically for the cloud to combine the benefits of relational database structure with non-relational horizontal scale. This combination is designed to deliver high-performance transactions and strong consistency across rows, regions, and continents with a 99.999% availability SLA, no planned downtime, and enterprise-grade security. The service provides relational semantics and handles schema changes as an online operation with no planned downtime. Organizations can leverage existing skills to query data in Cloud Spanner using familiar, industry-standard ANSI 2011 SQL. Cloud Spanner also provides automatic sharding. A database shard is a horizontal partition of data in a database. Shards are held on separate database servers within a Cloud Spanner instance to distribute load. Google Cloud’s security architecture includes encryption by default in transit and at rest, identity and access management, comprehensive audit logging, custom-manufactured hardware, hardware tracking and disposal, and Google’s extensive global network.

Figure 2. Google Cloud Spanner



Source: Enterprise Strategy Group

ESG Economic Validation

ESG completed a quantitative economic analysis of Cloud Spanner. Focus was placed on the economic benefits organizations can expect when leveraging Cloud Spanner compared with on-premises relational database clusters and other cloud-native database offerings, like AWS DynamoDB and Azure CosmosDB.

ESG’s Economic Validation process is a proven method for understanding, validating, quantifying, and modeling the economic value propositions of a product or solution. The process leverages ESG’s core competencies in market and industry analysis, forward-looking research, and technical/economic validation.

ESG conducted in-depth interviews with end-users to better understand and quantify how Cloud Spanner has impacted their organizations, particularly in comparison with previously deployed and/or experienced solutions. In addition to having experience with on-premises relational database solutions, some of the customers interviewed had migrated their other cloud database environments to Cloud Spanner and were able to give feedback on ongoing administration differences between the cloud solutions. The qualitative and quantitative findings were used as the basis for a simple economic model comparing the expected costs of on-premises and various cloud-native managed database services.

Google Cloud Spanner Economic Value Overview

ESG's economic analysis revealed that an effective deployment of Cloud Spanner can provide significant cost, administration, and agility benefits when compared with other competitive on-premises and cloud deployments. Additionally, ESG found the flexibility and cost structure of Cloud Spanner provides savings and benefits even when compared with using alternative cloud deployments. We also found that Spanner provided customers with significant savings and benefits in the following categories:

- **No hardware, software, or infrastructure investment**—Cloud Spanner is a cloud-native, fully managed database service with no hardware to manage. Google keeps the service up and running—including patching and disaster recovery. The drastically reduced need for onsite hardware results in lower upfront costs while the simplified scalability of a Cloud Spanner environment eliminates the need to over-purchase hardware or over-provision cloud services. Additionally, Cloud Spanner has pre-provisioned resources, so costs can be controlled and forecasted.
- **Simplified administration**—ESG validated that administration, maintenance, and operation of a Cloud Spanner environment is easier, faster, and more effective than managing alternative on-premises or cloud-based solutions. Organizations can spend less time managing, and more time adding value. In traditional RDBMS environments, developers build application logic to handle transactions and/or eventual consistency. Because Cloud Spanner offers strong consistency, scale, and relational semantics including transactions, it can make it easier for developers to produce cleaner code. Maintenance is simplified with online schema updates and automatic sharding.
- **Automated scalability, availability, and agility**—Cloud Spanner can be scaled on demand without downtime as the application grows, so there is no need for over-provisioning to reserve capacity ahead of time. Customers reported that Cloud Spanner enabled them to enhance business agility, which allowed them to harness the value of their data to better service existing customers and open new revenue streams. Strong consistency across tables, regions, and continents means that Cloud Spanner can achieve global scale when organizations need it.

Lower Hardware Investment/Upfront Costs

Building and maintaining an on-premises database environment requires significant upfront hardware, software, and personnel costs. Some cloud database services still require the provisioning of disk (cloud storage) separately from the database service. Customers have found that by deploying on or migrating to Cloud Spanner they could:

- **Reduce or eliminate upfront costs**—Shifting to Cloud Spanner eliminates the need for the server and storage hardware required for on-premises environments. When compared with other database solutions (both on-premises and cloud-based), Cloud Spanner provided less costly and more predictable planning, purchasing, and deployment.
- **Eliminate the need to overbuy capacity**—The ability to effectively spin up databases in a Cloud Spanner environment in minutes eliminates the need to project, procure, deploy, and maintain excess capacity to handle spikes, seasonal requirements, and company growth. The per-node/per-hour billing of Cloud Spanner allows customers to pay for what they need in the moment instead of funding future growth. The virtually unlimited capacity of Google Cloud allows Cloud Spanner customers to plan and sell without concern for the ability to store and access data. Customers reported that the ability to quickly deploy apps without provisioning reduces cost while giving them just-in-time access to the exact amount of resources needed to complete the job.

- **Achieve simplicity**—For companies opting for on-premises solutions, the forecasted data storage and data processing requirements necessitate complex planning in both hardware needs and the physical space to house that hardware. Hardware costs go far beyond the acquisition of the actual servers. Companies are forced to plan and fund real estate to house servers and storage arrays while keeping them powered, cooled, and administered. Cloud Spanner eliminates this need as it allows you to focus on gaining insight from your data instead of planning for storage of your data. Billing is simplified also—there’s no need for large upfront costs or unplanned expenses if the environment grows beyond forecasted requirements for data storage and processing power. Cloud Spanner allows for a level of expense predictability, enabling organizations to plan for other value-added projects.

“We used to budget millions of dollars per year for our RDBMS environment including hardware, software, and staffing. Even worse, we sometimes had to tell our business groups to wait weeks on a project while we sourced and deployed more capacity to handle their needs. This was completely eliminated when we migrated to Cloud Spanner. We reduced our capital and operational expenditures costs by nearly 80% and made our DBAs completely self-sufficient.” Data Engineering Manager, Global 500 Pharmaceutical Firm

Simplified Administration

Customers reported a 90% reduction in administration costs when moving on-premises RDBMS operations to Google Cloud Spanner and up to a 78% savings in administrative costs when shifting from other cloud services to Cloud Spanner. In addition to the elimination of many of the tasks associated with an on-premises RDBMS cluster, the simplicity of administering Cloud Spanner allows less experienced and less expensive resources to complete most of the simplified tasks. For example, customers we interviewed reported that on-premises resources took an average of weeks to months to provision and configure clusters for applications while creating the same capability with Cloud Spanner takes minutes.

“We had one employee working 10 hours a week purely focused on Apache Cassandra. This was a high salaried employee. This was eliminated when we migrated to Cloud Spanner. We eliminated administration costs as a budget line item.” CTO, Software Analytics Company

The majority of savings came from reduction in IT administration, which typically requires tasks such as diagnosis and resolution of hardware and software issues, upgrading, planning for maintenance shutdowns, implementation of updated security measures, and periodic monitoring. Databases are managed entirely by individual development teams without requiring training on database concepts. Cloud Spanner enabled the organizations interviewed by ESG to focus on data, not databases.

In traditional on-premises environments, there are requirements for procurement, deployment, configuration, scaling, tuning, testing, updates, and monitoring of equipment. The same amount of capacity can be administered with Cloud Spanner and requires only minimal time for configuration.

Customers interviewed by ESG who have experience with both Cloud Spanner and other cloud offerings reported reductions in the time and complexity of administrative tasks when comparing the cloud-based services. One reported that they needed a high-level administrator who spent over 25% of their time just making sure the cloud environment was properly configured and available. This need was eliminated when the organization migrated to Cloud Spanner.

Simplified administration also comes from automated and transparent sharding. While sharding is long-established and well understood, in practice, it is complex, and often achieved through custom coding. This can be inflexible and extremely difficult to maintain, as evidenced by one customer interviewed by ESG who described a two-year process of adding a new

shard to a sharded MySQL environment. Cloud Spanner removes the need for manual sharding, since databases are automatically sharded and distributed with no downtime and without performance impact to the application.

Table 1. Simplified Management Helps Lower Administration Cost

IT Task	On-Premises RDBMS	Other Cloud Databases	GCP Cloud Spanner
Planning and Research	\$ \$ \$	\$	\$
Justification & Procurement	\$ \$ \$	\$	\$
Deployment	\$ \$ \$	\$	\$
Configuration	\$ \$ \$	\$	\$
Tuning & Optimizing	\$ \$ \$	\$ \$	\$
Daily Job Administration	\$ \$ \$	\$ \$	\$
Troubleshooting Hardware	\$ \$ \$		
Updates and Maintenance	\$ \$ \$		
Capacity Planning	\$ \$ \$	\$	
Hardware Refresh	\$ \$ \$		

Source: Enterprise Strategy Group

Automated Scalability, Availability, and Agility

Data is only valuable when it can be effectively utilized to help meet business needs. ESG’s customer interviews uncovered several ways in which replacing on-premises RDBMS clusters with Cloud Spanner had helped to make the business more agile and better enable business processes through:

“We were able to empower development teams with the ability to run their own environments, which shrank deployment times from weeks to minutes.” Data Engineering Manager, Global 500 Pharmaceutical Firm

- **Faster time to production**—Customer interviews showed a marked improvement in the ability to bring production workloads online in a Cloud Spanner environment. One customer reported, “Spanner’s ability to set up and tear down quickly is critical to us. Interoperability with Cloud DataFlow makes Spanner a force multiplier. We used to use (Apache) Cassandra for our supply chain data, with IT spending a few hours a week managing and maintaining the environment—that time has been reduced to zero. We were able to empower development teams with the ability to run their own environments, which shrank deployment times from weeks to minutes.”
- **Simpler scaling**—Deploying databases to production is only part of the challenge. ESG has found that customers were able to act faster on volatile performance/capacity requirements in a Cloud Spanner environment than they could in the past using on-premises RDBMS or some cloud solutions. A customer who migrated from an on-premises RDBMS cluster said, “Scaling was always a challenge, because to scale meant adding a new database instance to the cluster and all the complexity that comes with that.” “Scaling was a non-trivial exercise in server, database, storage, and network management,” said another customer.

“Cloud Spanner’s short learning curve meant that our data engineering team was able to port existing apps and begin coding new apps very quickly. Instant copies of databases vastly improved developer productivity.” VP Engineering, Consumer Services Firm

- **Faster application development**—Customers reported substantial value in the ability to act quickly on opportunities. The agility brought by Cloud Spanner improved developer productivity by eliminating the need to code for database copies or strong consistency. One customer declared: “Cloud Spanner’s short learning curve meant that our data engineering team was able to port existing apps and begin coding new apps very quickly. Instant copies of databases vastly improved developer productivity.”

Figure 3. Improved Business Agility with Cloud Spanner

Cloud Spanner Advantage



- Faster time to production
- Simple performance and capacity scaling
- Faster app development



- Instant scalability of clusters
- No up-front investment
- Simpler business justification

Benefit to the Business



- Faster and more profitable decisions
- Instant response to business needs
- Increased productivity and revenue



- Ability to act on new opportunities
- Faster success/failure of initiatives
- Capitalization of new revenue streams

Source: Enterprise Strategy Group

ESG Analysis

ESG created a three-year TCO/ROI model that compares the costs and benefits of satisfying a modeled organization’s relational database requirements with Cloud Spanner versus with an on-premises RDBMS cluster or other cloud services. The model leveraged the information collected through vendor-provided material, public and industry knowledge of economics and technologies, and the results of customer interviews. ESG’s interviews with customers who have recently made the transition, combined with experience and expertise in economic modeling and technical validation of IT solutions, helped to form the basis for our modeled scenario.

The modeled organization started with 35 database instances and 10 TB of combined production, test, and dev data hosted on-premises. For the cloud-based solutions they looked at, they required a minimum of 10 TB of storage that was protected and could be accessed by all development teams. In one case, the cloud service required provisioning of disks from their cloud storage service separately from provisioning of the databases. In comparison, the on-premises RDBMS deployment required at least three times as much storage to provide redundancy and capacity for test/dev.

Additional performance/storage capacity was reserved to handle seasonal bursts and additional load as demanded by the business. While the cloud-based solutions could spin up new nodes only when required, the on-premises RDBMS solution

Why This Matters


On-premises relational databases require the purchase and maintenance of physical servers, usually on a three- or four-year lifecycle plan. This requires a planning and procurement cycle that can be limiting to most businesses.

Many cloud-based RDBMS services allow for dynamic provisioning of capacity. Cloud Spanner adds a level of flexibility when looking at costing models because of Cloud Spanner’s speed in spinning up, copying, and scaling databases.

required deployment of the “worst case” scenario from a performance and capacity standpoint as well as the purchase of additional nodes at the time of initial deployment.

In addition to gaining feedback from the actual environments of the customers we interviewed, we obtained validation of expected administration time spent managing each solution. ESG asked customers about the strengths and weaknesses of each solution and gave them the opportunity to provide guidance that allowed us to model the administrative hours required for an on-premises deployment as well as the expected reduction in hours expected with Google Cloud Spanner and other cloud databases.

Figure 4. Three-year TCO Calculation



Source: Enterprise Strategy Group

ESG found that the Cloud Spanner solution could satisfy the needs of the modeled organization at a total cost that was 78% lower than an on-premises RDBMS deployment with no upfront costs, and up to 37% lower than other cloud database deployments.

The Cloud Spanner solution greatly reduced or eliminated upfront costs and maintenance, support, and infrastructure spending (power/cooling/floorspace). And because administration and maintenance were greatly simplified and did not require trained server, storage, or network administrators, cost of administration was reduced by 90% versus an on-premises deployment, and up to 80% versus other cloud-based managed database offerings. ESG believes benefits will be realized to some degree by most organizations looking to lower cost and complexity, while getting the most out of their relational database deployment. Every organization’s requirements are different, and your particular savings may vary; not every organization needs an enterprise-grade, horizontally scalable, globally-distributed, and strongly consistent database service. For smaller organizations with less demanding requirements, Cloud Spanner might not be the best fit.

The Bigger Truth

While there are clearly many shades of gray for cloud-based databases, the benefits offered by cloud service providers in this space are becoming better understood and more highly valued by buyers. Cloud services now match and often surpass on-premises efforts, in terms of unit costs of resources as well as a wide variety of qualities that enable businesses to spend more time managing and extracting value from their data, rather than maintaining their databases.

In the long term, it is unlikely that many businesses will be able to—or want to—devote the extraordinary energy required to design, purchase, integrate, and manage their infrastructure. Even if they did, their chances of outperforming clouds for relational databases are slim. The economies of scale and magnitude of investment among cloud service providers are compelling.

ESG validated that Cloud Spanner provides cost-effective and agile relational database clusters that can be easily spun up and down as required and can be completely managed by individual dev teams. ESG validated the many benefits of Cloud Spanner with Google customers and found that Cloud Spanner not only provided their operations with significantly faster time to production, but also provided substantial cost savings without the need for a long-term contract or an upfront investment. More importantly, Cloud Spanner freed administrators to work on higher value initiatives and enabled businesses to scale performance and capacity at a speed that was previously impossible—positively impacting revenue.

ESG has performed economic evaluations of many of GCP's offerings such as Google Compute Engine, BigQuery, Dataproc, and Advanced Networking services and has found that with each offering, Google provides its customers with products that offer simple and flexible solutions with fair and highly visible on-demand pricing. Cloud Spanner offers an excellent opportunity for organizations to take advantage of the true agility and cost savings of a cloud-based architecture while leveraging their existing expertise. If your organization is looking to improve business agility with a scalable, cloud-based solution that is high-performing and cost-effective, ESG recommends that you give Google Cloud Spanner serious consideration.

Appendix

Table 2. Three-year Database TCO Model

	On-premises DIY	Cloud Database A	Cloud Database B	Google Cloud Spanner
Hardware/software (includes storage)	\$3,319,800	-	-	-
Operations (power, cooling, infrastructure, planning, installation, and migration)	\$509,909	-	-	-
On-demand cloud costs	-	\$1,372,596	\$1,130,483	\$540,954
Cloud storage	-	\$616,674	\$300,508	\$778,446
Maintenance/support contracts	\$768,640	-	-	-
Administration	\$1,530,000	\$782,870	\$313,650	\$153,000
Total	\$6,128,349	\$2,772,140	\$1,744,641	\$1,472,400

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2019 by The Enterprise Strategy Group, Inc. All Rights Reserved.

