



February 2024

# 451 Research Technology Impact on Business Report

## Serverless

# Table of contents

<b>Research overview</b>	<b>3</b>
<b>Research methodology &amp; firmographics</b>	<b>5</b>
Research process	5
Key considerations	6
<b>Firmographics</b>	<b>7</b>
<b>Key findings</b>	<b>8</b>
<b>Cost comparison: Serverless vs. traditional cloud infrastructure</b>	<b>9</b>
<b>Serverless' impact expands at scale: People and time</b>	<b>10</b>
<b>Serverless has varying impacts across organizational and technological criteria</b>	<b>12</b>
Organizational & technological benefits	12
Serverless benefits matrix	15
<b>Serverless stories behind the cost savings</b>	<b>16</b>
Serverless functions deployment analysis: Global apparel brand	17
Serverless functions deployment analysis: Global airline	19
Serverless container platforms deployment analysis: Global streaming provider	21
Serverless container platforms deployment analysis: Global data service	23
<b>Quantifying the multi-year serverless journey</b>	<b>25</b>
<b>Conclusion and considerations</b>	<b>26</b>
<b>About the authors</b>	<b>27</b>

# Research overview

The cost of IT infrastructure is increasingly center stage as organizations look to digital transformation and the advantages of the cloud including scalability, agility and security. For many applications, the cloud has driven down costs immensely versus on-premises deployment, yet many businesses with a traditional operating model on cloud infrastructure have reached a point of diminishing returns for cost savings. The potential of serverless architectures — for both serverless functions and container platforms — to scale based on application demand offers significant opportunities to better optimize resources and recognize the next wave of cloud savings. However, 36% of organizations still see the cost of serverless as a significant barrier to adoption, according to 451 Research's Voice of the Enterprise: Cloud Native, Adoption & Usage 2023 – Serverless survey.

## Research purpose

The purpose of this study is to demystify this cloud cost conundrum by quantifying and analyzing the costs, benefits and impacts of serverless deployments. Based on a detailed examination of 10 real-world deployments, this analysis provides critical insights into expected costs and savings for IT leaders considering or implementing serverless.

Serverless computing is a cloud-based execution model that abstracts away and automates several cloud computing infrastructure management capabilities and tasks. The serverless operating model automates compute provisioning, security patching and several other tasks, relieving developers and operations teams from routine work and empowering them to take on more interesting and challenging responsibilities. The impact on IT operations efficiencies (46%) and developer speed/productivity (43%) is significant; according to S&P Global Market Intelligence 451 Research survey data<sup>1</sup>, these were the top workforce benefits of serverless.

## Study highlights



**35%**

Annual cost savings



**\$105,000**

Savings per serverless FTE



**\$11.5 million**

5-year cumulative net savings



<sup>1</sup> Voice of the Enterprise, Cloud Native, Adoption & Usage 2023 – Serverless



## Automation of IT Tasks

Self-managed			Semi-automated			Fully automated						
Operating models												
On-premises						On cloud			On serverless			
Applications	New applications/features						New applications/features			New applications/features		
	Workflow automation						Workflow automation			Workflow automation		
IT	Back-up/ disaster recovery		Sys admin		Back-up/ disaster recovery		Sys admin		Back-up/ disaster recovery		Sys admin	
	Monitoring	Database	CI/CD		Monitoring	Database	CI/CD		Monitoring	Database	CI/CD	
Security	Patching		Compliance		Patching		Compliance		Patching		Compliance	
Infrastructure	Networking		Storage		Networking		Storage		Networking		Storage	
	Virtualization				Virtualization				Virtualization			
	Operating system		Runtime		Operating system		Runtime		Operating system		Runtime	
	Server management/provisioning				Server management/provisioning				Server management/provisioning			
	Scaling servers/services				Scaling servers/services				Scaling servers/services			
	Setup/configuration				Setup/configuration				Setup/configuration			
	Optimize compute/performance				Optimize compute/performance				Optimize compute/performance			
	Container management				Container management				Container management			

Source: 451 Research conducted in-depth interviews of serverless deployment leaders.

Serverless has performance implications from an IT and infrastructure standpoint. Compute efficiencies are attainable from an event-driven architecture that triggers actions based on events, as well as more efficient auto-scaling that can better match capacity to application usage, reducing operational load, CPU usage and response times. These capabilities especially impact modern applications with high transaction volumes and fluctuating traffic, customer-facing real-time elements requiring minimal latency, stateless systems processing short-term request/response workflows and asynchronous processing of big data such as video content.

We analyzed the infrastructure and workforce impacts of two core serverless implementations: serverless functions and serverless container platforms. Serverless functions are short-lived processes that abstract away most of the infrastructure management, making them ideal for high-transaction-volume, event-driven and short-lived processes that are costly at scale for traditional cloud infrastructure. Serverless container platforms package code and dependencies in containers for longer-running processes and those requiring greater customization or control.

451 Research dove into the workforce and infrastructure impacts of 10 real-world serverless deployments, examining their influence on costs and key organizational and technical criteria.

# Research methodology and firmographics

## Research process

- 1. Interview recruitment:** 451 Research and the serverless provider recruited and qualified participants for this study.
- 2. Information collection:** Each qualified participant provided key quantitative metrics regarding their serverless application(s) (e.g., IT costs, size of supporting IT staff, number of requests, etc.) and participated in in-depth interviews with 451 Research to affirm data and share additional context surrounding the deployment (e.g., why serverless, progression of deployment, challenges, benefits, etc.).
- 3. Leverage additional research:** 451 Research integrated inputs from other sources (451 Research, serverless provider, public) to fill gaps and develop model assumptions.
- 4. Model development and analysis:** 451 Research scrutinized collected data using relevant proprietary data (including the Cloud Price Index™) and weightings. 451 Research created a composite organization based on the study participants, measured and analyzed infrastructure and workforce costs for both traditional cloud and serverless deployments, and aggregated deployment anecdotes from the interviews, among other information.
- 5. Quality review:** 451 Research analysts and the serverless provider reviewed key findings and insights to ensure validity.



# Key considerations

- This research and analysis are based on a study of 10 organizations that have deployed serverless applications and is not necessarily illustrative of every serverless deployment.
- Not every application is suited for serverless infrastructure, nor are the financial impacts and benefits described here attainable for every organization.

## Model framework: Traditional cloud versus serverless

Costs	Infrastructure		Workforce	Metrics collected
	IT costs (compute, networking, memory, etc.)		Development	<ul style="list-style-type: none"><li>• Annual serverless costs</li><li>• Annual workforce hours</li><li>• FTEs (supporting deployment)</li><li>• FTE hours and costs</li><li>• Transaction/request/response sizes</li><li>• Number of users</li><li>• Duration of deployment</li></ul>
	Supporting systems	Security	Setup	
		Monitoring	Maintenance	
		Backup/DR		

Source: S&P Global Market Intelligence 451 Research.

**Model framework:** 451 Research developed a model that captured the majority of costs associated with deploying and sustaining serverless applications. These costs were apparent across the infrastructure (including supporting systems) and the workforce. We made a few assumptions in our cost models to provide a repeatable and measurable baseline to compare and contrast serverless impacts. These include but are not limited to:

- **Workforce:** \$200,000 annual salary for IT employees working with serverless or referred to as “serverless IT FTEs,” 2,080 annual working hours per IT employee, etc.
- **Infrastructure/supporting systems:** In situations where the research participant could not make estimates or provide hard data on the costs of running the application on traditional cloud infrastructure, we made calculations based on 451 Research Cloud Price Index data.

**Research firmographics overview:** Research study participants encompassed a range of roles, industries and company sizes, and all had more than a year of experience in serverless deployments. Their use of serverless functions and container platforms took up a sizable portion of their organization’s application mixes, in many cases operating alongside applications still running in traditional cloud and on-premises/datacenter environments. The majority are several years into their serverless deployments and have dozens to hundreds of IT employees leveraging, managing and optimizing them daily.










# Firmographics

10 Serverless leaders






~3 Years of serverless deployments

**Roles**  
Lead/principal engineer  
Software development engineer  
CTO  
Senior architect

## Industries

	Airline		Data service		Information security
	Apparel brand		E-commerce		Streaming provider
	Collaboration software		Electronics		
	Cosmetics		Gaming		

## Deployment details

	<b>Infrastructure mix</b>	Serverless functions Serverless container platforms Traditional cloud infrastructure
	<b>Serverless IT FTEs</b>	~38 – Average employees across organizations 5-150 – Range of employees across organizations
	<b>Duration of deployment</b>	~36 months – Average months of deployment 18-54 – Range of months of deployment
	<b>Annual serverless spending</b>	~\$592,000
	<b>Application requests &amp; compute data volume</b>	~3.7 billion – Average monthly requests ~52 TB – Average compute data volume

# Key findings



In aggregate, the organizations analyzed in this study recognized a 35% annual cost reduction from leveraging serverless in comparison to traditional cloud infrastructure.



Most (86%) of these annual cost savings were from workforce-related serverless benefits, primarily from reducing the time and cost of maintenance (by 54% on average), infrastructure setup (48%) and application development (19%).



This impact of serverless adoption on workforce efficiency was noticeable and sizable for both smaller IT organizations (i.e., five teams in our study with fewer than 20 IT employees generated \$1.6 million in annual savings on average) and larger ones (i.e., teams with 100-150 IT employees recorded a range of annual savings from \$7 million to \$14 million).



This impact of scale is noticeable also in terms of time (duration of deployment), where teams with longer experience and more pervasive serverless adoption reported significant savings in later years. The four organizations that deployed serverless for less than 36 months recognized \$2.6 million in annual serverless savings on average, compared to six organizations that deployed longer than 36 months and recognized \$5.6 million in annual serverless savings on average.



Serverless can have major impacts on key strategic technological and organizational drivers where participants noticed a compounding effect as serverless permeated their organization. However, in many cases, recognizing these significant benefits required overcoming internal headwinds caused by resistance to change, lack of skill sets and complexity from legacy systems and the surrounding environment, among other factors.





# Cost comparison: Serverless vs. traditional cloud infrastructure

451 Research built a composite organization by analyzing and weighting key organizational and deployment characteristics of the 10 organizations participating in the study. From this aggregation, the composite organization experienced a **35% reduction in costs** compared to traditional cloud infrastructure.

The greatest drivers of cost reduction came from time saved in developing and operating new applications and features or updating existing ones. Participants cited three major areas of improvement:

- **Application development:** Reducing workforce time and costs from writing fewer lines of code, dealing with less code complexity and fewer languages, abstracting away manual tasks, and leveraging reusable functions and other pre-built components.
- **Setting up infrastructure:** Reducing workforce time and costs for provisioning infrastructure, including installing, managing and updating existing or new servers, provisioning new services to accommodate fluctuating demand and spikes in traffic, configuring memory allocation and security policies, and installing operating systems and other activities.
- **Maintenance:** Reducing workforce time and costs for maintaining infrastructure and servers, including patching and upgrading applications, operating systems and security systems.

## Composite organization

(Aggregation of 10 organizations analyzed)




**Dedicated IT employees (serverless):** 36  
**Duration of serverless deployment:** 36 months  
**Annual serverless IT spending:** ~\$592,011  
**Monthly transactions/requests:** 3.77 billion  
**Compute data volume:** 52.8 TB

“Our development team has increased application feature output by three times; features that required 8-12 hours of setup and configuration prior to serverless can now be deployed in a few minutes.”


**Senior enterprise architect**  
Global airline, 74,000+ employees

## Serverless achieves 35% cost savings over traditional cloud



### Serverless workforce efficiency drivers

Reduced time and costs by...



Source: 451 Research's in-depth interviews with serverless deployment leaders.


Infrastructure cost reductions were less dramatic than workforce savings, but research participants shared data and anecdotes on the time saved by not having to set up and configure supporting systems (such as cybersecurity, disaster recovery, IT monitoring, etc.).

# Serverless' impact expands at scale and over time

The impact of serverless at scale is noticeable when accounting for the size of IT teams working with serverless. We defined “serverless IT FTEs” as those with daily interactions and management of serverless applications; this primarily included software developers, engineers and architects.

This impact in relation to people may not be surprising given the impact of serverless adoption on workforce efficiency, but the sheer volume of workforce cost savings quickly accumulates even for relatively small serverless IT teams.

## Annual serverless savings grows with larger serverless workforces



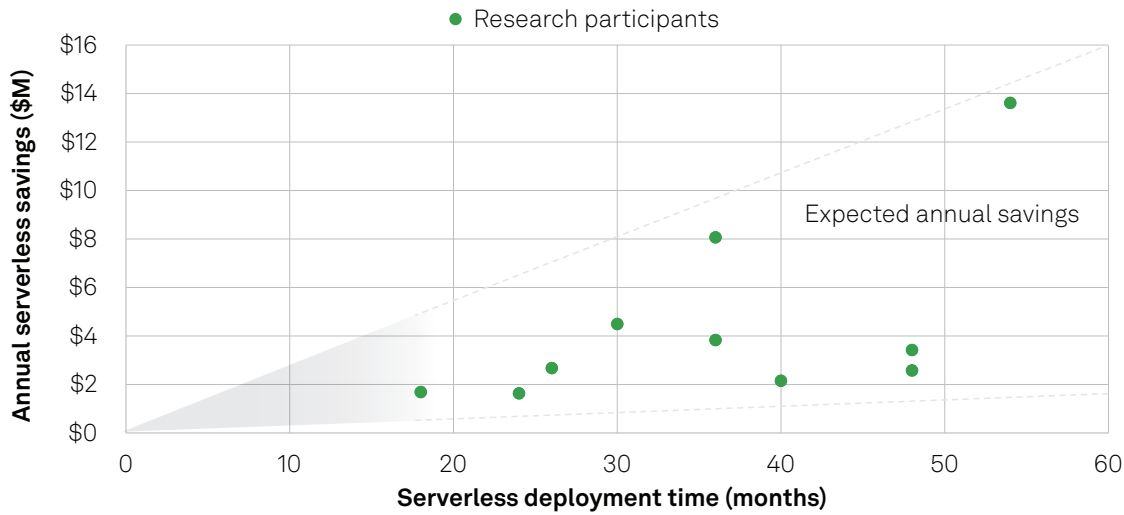
Source: 451 Research's in-depth interviews with serverless deployment leaders.

The annual serverless savings can compound as serverless becomes pervasive across larger IT Teams (i.e., a global airline with a massive IT headcount base) and/or the organization is organically growing in correlation to its serverless deployment (i.e., a cybersecurity startup increasing its IT headcount). While the annual serverless savings unquestionably increase with larger IT workforces, more IT employees working with serverless doesn't necessarily mean proportionally greater savings per IT employee. Smaller IT organizations working with serverless are generating notable savings: The five teams with fewer than 20 IT employees on average generated \$1.6 million in annual savings.

On average, research participants recognized \$105,000 in annual serverless savings per serverless IT FTE. Our trendline based on statistical correlation shows that an IT serverless workforce of 20 employees experienced about \$2 million in annual savings on average versus traditional cloud deployment; those with 120 employees can expect nearly \$9 million.

The second key finding related to the scale of serverless deployment was based on time or how long (in months) the deployment has been in production. The gaps in annual savings between organizations widen as more time passes, and some study participants experienced a steeper learning curve for initial serverless adoption with varying levels of success at scale. For example, for an average of 36 months of deployment, an organization may realize between \$1 million and \$10 million in annual serverless savings.

## The growing gap of serverless success over time



Source: 451 Research conducted in-depth interviews of Serverless deployment leaders.

“To make it easier for teams to adopt serverless, we’ve hooked into existing infrastructure, wrote shared modules (i.e., CI/CD, Terraform), ran learning sessions and targeted teams. This has helped us overcome bad habits in coding (custom coding) and optimize resources more efficiently.”

**Principal software engineer**

Global collaboration software provider, 1,000-5,000 employees

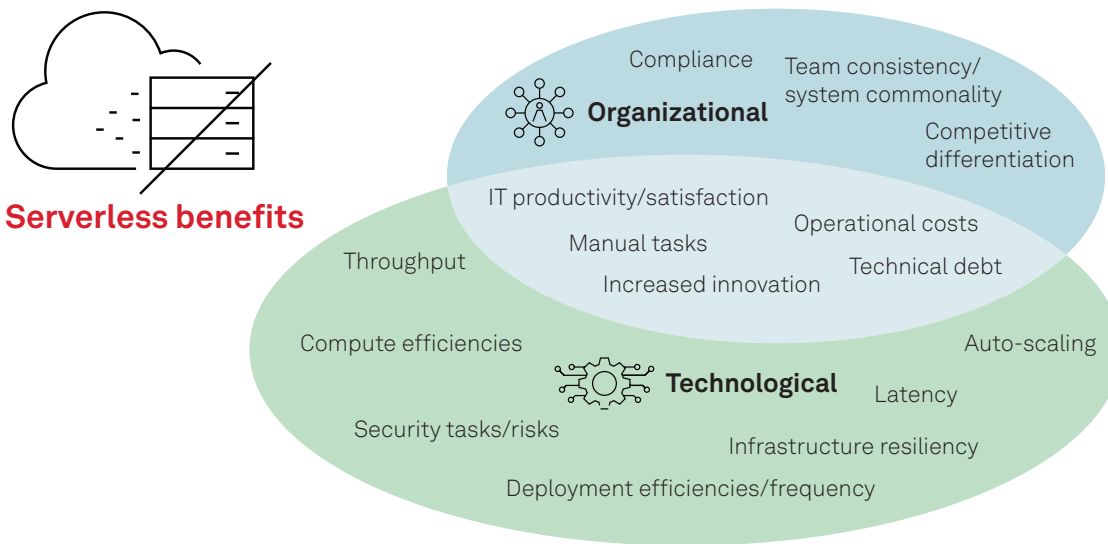
Factors that influence the speed and persistence of serverless savings may reflect the success of internal evangelism efforts, rate of adoption across the organization and the ability to overcome technical challenges (i.e., legacy systems/difficult to migrate, etc.), among several others. Naturally, these obstacles will have less impact as more workloads move onto a common infrastructure and internal expertise develops from implementing best practices, as well as education and upskilling programs. Research participants who were more successful in incubating a serverless culture were less deterred by challenges and recognized greater cost savings in less time.

# Serverless has varying impacts across organizational and technological criteria

While the cost analysis of our composite organization shows measurable financial impact of serverless deployment, additional benefits captured in the in-depth interviews are worth exploring. Although many of the advantages can impact the broader organization, this discussion focuses on areas where the primary beneficiaries were IT teams.

451 Research found that serverless has benefits across organizational and technological parameters.

## Serverless benefits across organizational and technological criteria



Source: S&P Global Market Intelligence 451 Research.

## Organizational and technological benefits

- **Competitive differentiation:** Serverless adoption often plays a role in developing or improving business-critical applications where performance directly contributes to the top and bottom line. In some instances, serverless capabilities unlocked competitive differentiation for successful implementers, whether it was a cyber startup building its threat detection platform using a serverless architecture or a global airline achieving a pricing advantage due to lower transaction costs.
- **Team consistency and system commonality:** Irregularities in code, inconsistent runtimes and tool sprawl can quickly compound technical debt across organizations. Serverless deployments can provide a more unified foundation for multiple teams to build from and foster a more cohesive IT culture.
- **Improved compliance:** Compliance with industry-specific (i.e., HIPAA, PCI) or regional regulations is table stakes for organizations to compete in their given market. Also, with cyberattacks an ever-increasing threat, adherence to security and data protection mandates is a strategic necessity. Serverless mitigates some compliance risks and costs by centralizing cybersecurity services (e.g., automated patching) and shifting some of these responsibilities to the serverless provider.



- **Reduced operational costs:** By using serverless, organizations were able to reduce infrastructure and workforce operational costs, lowering the extent of overprovisioned and unused cloud infrastructure and decreasing operational overhead for infrastructure management tasks.
- **Reduced manual tasks:** Not having to perform repetitive infrastructure tasks, such as configuring and patching servers and upgrading supporting systems, reduced operational toil and let staff focus on feature and application development.
- **Improved IT productivity/satisfaction:** IT teams are the cornerstone of any technology-first organization, yet their skill sets are in short supply. Dedicating more of these critical resources to high-value activities (new features/applications) and away from repetitive and manual tasks helped reduce operational costs, improve the developer experience and lower developer churn rates.
- **Reduced technical debt:** Organizations were looking to move away from disparate and customized systems ridden with inconsistencies in software and underlying infrastructure. The ability of serverless to abstract away many IT management tasks (provisioning, scaling) reduces the opportunity for coding errors or one-off configuration tweaks that can add to technical debt.
- **Increased innovation:** With serverless, organizations were able to invest time and energy they would have spent performing IT infrastructure chores into innovating products and services, which can have a massive effect on organizational (increased revenue, product differentiation) and technological (new features, deployment frequency) parameters.
- **Auto-scaling:** In interviews, respondents cited auto-scaling features that can rapidly accommodate changes in application demand, including unexpected spikes in traffic, as the top technological benefit. They deemed serverless' ability to more efficiently horizontally scale (add net-new compute, servers, containers etc.) and vertically scale (add to existing compute, servers, containers) based on changing traffic or usage patterns to be a critical advantage.
- **Deployment efficiency/frequency:** Mundane infrastructure management tasks can create bottlenecks for IT teams motivated to deploy more frequently. By working above the infrastructure layer, serverless deployment allowed teams to focus on releases while fostering CI/CD best practices and minimizing dependencies.
- **Security tasks/risks:** The automation of security tasks, including patching and updating systems for compliance or in response to cybersecurity events (e.g., Log4j) can reduce risks and vulnerabilities.

**“With serverless, we don’t have to do provisioning or OS patching. Engineers can spend more time building features.”**

**Senior software development engineer**  
Global streaming provider, 100,000+ employees

**“We can try new and creative things a lot faster. For example, it took two months to develop a new flight opportunity optimization application. Before serverless, this application would have taken 1-2 years to develop.”**

**Senior enterprise architect**  
Global airline, 74,000+ employees

**“So many security features are built in, it was very easy to build an insecure application in traditional cloud infrastructure, but it’s very hard to do it in serverless.”**

**Distinguished software engineer**  
Global apparel brand, 75,000+ employees

- **Compute efficiencies:** Serverless features such as auto-scaling provide an opportunity to improve CPU usage and reduce operational load, overprovisioned infrastructure and unused resources.
- **Throughput and latency:** Many customer-facing and increasingly complex and high-volume web applications demand consistently fast performance. Serverless can enable increased throughput by processing a higher volume of transactions with greater compute requirements. The on-demand compute resources available in serverless can reduce latency during peak traffic periods.
- **Infrastructure resiliency:** The removal of infrastructure management tasks where there may be an array of workforce errors lessens the chance of application failures and downtime. Applications are less likely to fail with serverless' event-driven architecture that can automatically react to unexpected traffic spikes.

**“We improved CPU performance from 5% (on traditional cloud infrastructure) to 30% with serverless.”**

**Senior software development engineer**  
Data service provider, 100,000+ employees

## Serverless benefits matrix

These serverless benefits are theoretically achievable by all. Moreover, 451 Research found that, depending on the organization and its technical and organizational challenges and drivers, the impact of these benefits varied from organization to organization.

No impact	Some impact	Impact	Major impact

	Auto-scaling	Deployment frequency & innovation	Infrastructure management & optimization	Supporting systems (i.e., security)	Other
Global airline	Auto-scaling unlocks differentiation by enabling lower airline ticket prices through reduced cost per transaction	Increased application and feature output by 2x-3x	Reduced maintenance costs by \$13 million or 135,200 hours	Supporting migration of relational monolithic databases to event-driven architecture	Improved ease for development and developer satisfaction
Global data aggregation service	Faster scaling speeds & processing massive data volume	20% productivity improvement	Reducing costs from operating at 30% CPU (versus 5% before)	Leveraging real-time IT monitoring	Greater resiliency & simplified development environment
Infosec software startup	Scaling enables differentiation (i.e., faster threat detection)	Majority of time spent building new features	"I don't have to do it at all"	Simpler to manage & improved security	Increased competitive differentiation by being 'born-in-the-cloud'
Collaboration software provider	Optimized scaling activities including fluctuations in traffic (i.e., work week) & reduced latency by 50x	Reduced operational overhead and automated launch of application instances & deployments	Reduced infrastructure complexity and unit costs per request by 70%	Automation from monitoring and configurations (i.e., alarms)	Improved code quality and consistency across teams
Global streaming service	Scale-up for live sporting events was a major strategic initiative	Gained 80% more time in application development	Reduced maintenance by 16,640 hours or \$1.6 million	Real-time monitoring of key IT metrics (latency, errors, etc.)	Application latency and throughput improved, reducing customer-facing errors
Global apparel brand	Scaling for millions of endpoints on mobile applications	Gained ~29% more time in application build	Reduced infrastructure set-up time by 90%, and maintenance costs by 67%	Minimize security risks and time responding to security events (i.e., Log4j)	Greatly simplified development environment

Source: 451 Research conducted in-depth interviews of serverless deployment leaders.

451 Research distilled insights from six research participants' outcomes into key serverless benefits and applied weighting to reflect the impact those had on their organization. For example, minimizing security risks was the most critical serverless impact for the global apparel brand, and auto-scaling to enable competitive differentiation was most important for the information security startup.

# Serverless stories behind the cost savings

451 Research dove deep into the impacts of serverless through analysis of four deployments, focusing on specific deployment details, use cases, challenges, serverless key benefits and an analysis comparing the cloud and serverless costs for infrastructure and workforce.

## Organizations using serverless functions



Global apparel brand



Global airlines

## Organizations using serverless container platforms




Global streaming provider









Global data service provider







## Serverless functions deployment analysis

Respondent firmographics					
	Distinguished software engineer		75,000+ employees		<ul style="list-style-type: none"><li>• 5,000 IT employees</li><li>• 120 Working w/serverless</li></ul>
Deployment					
	Use cases: e-commerce, high/low traffic APIs, advertising, batch jobs		36 months into deployment		<ul style="list-style-type: none"><li>• 50 million monthly requests</li><li>• 37.5 TB compute data volume</li></ul>

## Deployment scenario and challenges


A global apparel brand began planning for serverless in 2018. The development team focuses on integrating data from the brand's activity applications with third-party partner services (e-commerce, advertising, etc.) via both high- and low-traffic APIs. The workload includes heavy data flows between applications, millions of requests a day for a single endpoint and daily batch jobs (which became the first serverless project). Prior to adopting serverless, the development team was burdened with infrastructure management tasks including server/VM/OS patching. Additionally, a major security incident (Log4j) sent the team scrambling to apply fixes. Other challenges included development environment complexity, long development lead times and other infrastructure tasks.

## Serverless key impacts and benefits


The top impacts and benefits the global apparel brand recognized from implementing serverless include:

1. Reduced workforce costs for infrastructure management tasks (maintenance, setting up infrastructure, etc.) by \$7.8 million.
2. Significantly reduced time in security-related daily activities (patching, app development, etc.) and in responding to incidents. Other teams spend 4-8 hours a week on infrastructure management; they spend close to zero.
3. Rapidly increased the number of deployments per year by greatly reducing lead time (from 20 hours to less than four hours) and spending "95% of work on feature work, with 5% on infrastructure."
4. A simpler development environment with the ability to use Python and other high-level languages to write code for serverless function invocations and no need to hire additional DevOps staff.
5. Greater compliance with organizational mandate to lower costs from over-provisioned cloud infrastructure.

Infrastructure cost comparison



Workforce cost comparison









Location	Cloud	Serverless
Infrastructure	\$4M	\$1.8M
Workforce	\$24.8M	\$17M
Impact	\$10M in annual cost savings	





## Serverless functions deployment analysis

Respondent firmographics					
	Senior enterprise architect		74,000+ employees		<ul style="list-style-type: none"> <li>• 1,000 IT employees</li> <li>• 150 working w/serverless</li> </ul>
Deployment					
	Use cases: Critical airline operations, promotions, reservation systems		54 months into deployment		<ul style="list-style-type: none"> <li>• ~40 million monthly requests</li> <li>• ~56 TB compute data volume</li> </ul>

## Deployment scenario and challenges


A global airline began its serverless journey in 2018 with the goal to run 80% of its workloads on serverless by 2028. Development pods of about 150 IT employees are responsible for developing and maintaining key applications such as airline operations (ground control, scheduling) and consumer-facing services (web ticketing, promotions). Prior to adopting serverless, half of these teams were consumed with maintenance activities (updating libraries, patching, etc.), which delayed application feature releases. The traditional cloud infrastructure in place was incapable of efficiently scaling for spikes in traffic, both anticipated (promotional campaigns) and unanticipated (flight changes). The cost per transaction/request was very high because dedicated servers were provisioned for peak capacity to ensure applications were resilient enough to meet demand.

## Serverless key impacts and benefits


The top impacts and benefits the global airline recognized from implementing serverless include:

1. Reduced workforce costs by \$13 million, led by reducing time and costs for maintenance activities by 67%.
2. Flexible auto-scaling for anticipated and variable traffic spike days contributed to reducing infrastructure costs by ~20%.
3. Serverless enabled lower transaction costs, allowing the airline to charge less for tickets and differentiate from the competition.
4. Development team has increased application feature output by three times; features that required 8-12 hours of setup and configuration prior to serverless can now be deployed in a few minutes.
5. Faster experimentation enables the airline to innovate and take advantage of market opportunities. For example, the airline developed a new flight opportunity optimization application in two months; before serverless, the company estimates this application would have taken 1-2 years.


Infrastructure cost comparison



Workforce cost comparison



Location	Cloud	Serverless
Infrastructure	\$2.7M	\$2.1M
Workforce	\$30M	\$17M
Impact	\$13.6M in annual cost savings	







## Serverless container platforms deployment analysis

Respondent firmographics					
	Senior software development engineer		100,000+ employees		<ul style="list-style-type: none"><li>• 25 on team working with serverless</li></ul>
Deployment					
	Use cases: Streaming service device support, user experience		26 months into deployment		<ul style="list-style-type: none"><li>• ~45 million monthly requests</li><li>• ~11.25 TB compute data volume</li></ul>

### Deployment scenario and challenges


A global streaming provider began its serverless deployment in 2021. The team of 25 responsible for managing the client experiences on specific streaming device types had been operating on a centralized traditional cloud infrastructure leveraged across the organization where they experienced significant challenges related to scalability, ownership of fixes, new feature development and noisy neighbors that impacted performance. They spent substantial time and resources on infrastructure management tasks (provisioning, OS patching, etc.) as a result. Additionally, overprovisioned compute infrastructure capacity, including unused containers, created a costly footprint.

### Serverless key impacts and benefits

The client experience team at the global streaming service recognized several impacts and benefits from implementing serverless:


1. Reducing the workforce’s time on infrastructure management saves the team \$1.6 million in annual costs.
2. More efficient auto-scaling and other serverless features enabled the team to cut its infrastructure costs by ~\$1 million by reducing unused containers and over-provisioned capacity.
3. Another key auto-scaling benefit is leveraged for managing the significant uptick in traffic from hosting live sporting events, yielding greater resiliency and performance at lower costs.
4. Operating services independently of the centralized cloud infrastructure lessens disruptions to other teams and services when addressing incidents or upgrades.

Infrastructure cost comparison



**Cost savings**  
**Infrastructure:**  
Saved \$858,000 or 35%  
**Supporting systems:**  
Saved \$218,000 or 35%

Workforce cost comparison









**Cost savings**  
**Maintenance:**  
Saved \$1.2M or 60%  
**Setup infrastructure:**  
Saved \$300,000 or 30%  
**Application development:**  
Saved \$100,000 or 5%

Location	Cloud	Serverless
Infrastructure	\$3M	\$2M
Workforce	\$5M	\$3.4M
Impact	\$2.7M in annual cost savings	





## Serverless container platforms deployment analysis

Respondent firmographics					
	Senior software development engineer		100,000+ employees		<ul style="list-style-type: none"><li>• 13 on team working with serverless</li></ul>
Deployment					
	Use cases: Product telemetry platform		18 months into deployment		<ul style="list-style-type: none"><li>• ~52.6 million monthly requests</li><li>• ~52.6 TB compute data volume</li></ul>


### Deployment scenario and challenges

A team for a global data service provider is responsible for aggregating all telemetry and usage data from millions of customers for real-time monitoring and business reporting to deliver insights to an array of internal stakeholders. Several inefficiencies added to costs when constantly aggregating this massive amount of data using traditional cloud infrastructure. The team of 13 implemented multiple use cases based off a serverless architecture including an edge ingestion service, data publisher and data retrieval API.

### Serverless key impacts and benefits

1. Reduced maintenance costs by 30%, application development by 20% and cost of setting up infrastructure by 20%.
2. Improved CPU performance from 5% (on traditional cloud infrastructure) to 30% with serverless.
3. Improved resiliency and minimized downtime from large-scale events.
4. The edge ingestion service optimizes different transaction/request types and sizes through grouping and batching, saving significantly on infrastructure costs.
5. Simplifying deployment mechanisms for coding lowers development costs and reduces the chance of outages from coding errors.
6. Serverless reduced time for security patching tasks.


Infrastructure cost comparison



Workforce cost comparison



Location	Cloud	Serverless
Infrastructure	\$3M	\$2M
Workforce	\$2.6M	\$2M
Impact	\$1.7M in annual cost savings	






# Quantifying the multi-year serverless journey

The need to more tightly align IT architectures with digital transformation initiatives and business strategy pushes these technology decisions to the forefront. As with previous or alternative investments of compute infrastructure, serverless is often a multi-year strategic investment.

The composite organization built from the 10 serverless deployments examined experienced a return on investment as soon as the end of year 1. In aggregate, we found serverless annual savings had a linear growth trajectory with compounding net annual benefits, whereas serverless infrastructure costs only grew a few percentage points each year.

## Five-year serverless impact outlook



Source: 451 Research's in-depth interviews with serverless deployment leaders.

Over a five-year period, the composite organization recognized about \$15 million in total cumulative savings and \$3.6 million in total cumulative costs, resulting in \$11.5 million in total net savings over five years.

These net savings grow as more teams adopt serverless and realize the advantages to their infrastructure and the workforce supporting them. However, this composite view conceals a great deal of variation that can be influenced by the success of internal evangelism, organizational openness to change, the extent of legacy infrastructure and applications, dependencies, environmental complexity, availability of internal skillsets, budget and compliance constraints, and fears of vendor lock-in, among other factors.

**“Currently, around 20% of the organization is running on serverless containers. We estimate this number will be 40% by mid-2024 and 60% by end of 2024.”**

**Principal software engineer**

Global collaboration software provider, 1,000-5,000 employees

# Conclusion and considerations

The impact of serverless is unquestionably substantial for the research participants analyzed in this study. The attainable benefits across the workforce and infrastructure clearly compound at scale and over time.

For net-new applications, a serverless approach requires architecting with smaller blocks of code. For organizations migrating from traditional environments, refactoring applications into small single-purpose containers (microservices) offers a head start. Evaluating the feasibility of using functions first can pay off down the road: Moving from functions into containers is easier than going the other way around. Serverless functions can be a great option for high-traffic, latency-sensitive applications. Also, despite their ephemeral nature, the streamlined nature of serverless processing confers benefits in terms of application security.

Deciding between serverless functions and a serverless container platform for a given application involves a mix of considerations and experimentation, but a “both/and” approach can work. Broadly speaking, short-running services with no need for in-memory processing are a good fit for functions; a serverless container platform is better for longer-running and bursty functions and will be a lighter lift for teams migrating from traditional cloud infrastructure.

Serverless isn’t a one-size-fits-all solution; organizations need to develop talent in managing this novel deployment option, and different application shapes and sizes will have different outcomes. Certain applications are not as suitable for serverless operation. These include steady-state workloads that won’t benefit from on-demand scaling, large and/or legacy complex applications that haven’t yet been refactored for cloud-native operation, software with strict dependencies on underlying hardware and workloads with intense caching and/or monitoring needs.

The organizations in this study have navigated internal obstacles across their people, processes and technologies to drive serverless success. As more applications are built or optimized for cloud-native operation, the case for serverless is growing. Those starting from a greenfield investment perspective without organizational and technical blockers should face few objections to exploring the potential for serverless deployment.



# About the authors



## **David Immerman**

### **Consulting Analyst**

David Immerman is a consulting analyst at S&P Global Market Intelligence. He is responsible for executing on a range of custom research initiatives and development of thought leadership across technology sectors including industrial IoT, digital transformation, edge computing, AI/machine learning and fintech, among others, and verticals such as manufacturing and automotive.



## **Steven Moretti**

### **Consultant**


Steven Moretti is a delivery consultant at S&P Global Market Intelligence, where he is responsible for designing, managing and delivering consulting and custom research projects. Before joining the Consulting team, Steven spent a year as a 451 Research Associate where he began contributing to custom research projects and has since developed a specialty in quantitative research, specifically modeling and model design.



## **Jean Atelsek**

### **Senior Research Analyst**

Jean Atelsek is a senior research analyst working across the Cloud & Managed Services Transformation channel and digital economics unit of 451 Research, a technology research group within S&P Global Market Intelligence. She covers vendors and technologies that manage or optimize public and private cloud total cost of operations, performance or consumption.



## **Jay Lyman**

### **Senior Research Analyst**

Jay Lyman is a senior research analyst with the 451 Research Cloud Native and Applied Infrastructure & DevOps channels at S&P Global Market Intelligence. He covers software development, hybrid and multicloud infrastructure management and orchestration, and enterprise use cases that focus on the confluence of software development and IT operations known as DevOps.



## CONTACTS

**Americas:** +1 800 447 2273

**Japan:** +81 3 6262 1887

**Asia Pacific:** +60 4 291 3600

**Europe, Middle East, Africa:** +44 (0) 134 432 8300

[www.spglobal.com/marketintelligence](http://www.spglobal.com/marketintelligence)

[www.spglobal.com/en/enterprise/about/contact-us.html](http://www.spglobal.com/en/enterprise/about/contact-us.html)

Copyright © 2024 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, [www.standardandpoors.com](http://www.standardandpoors.com) (free of charge) and [www.ratingsdirect.com](http://www.ratingsdirect.com) (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at [www.standardandpoors.com/usratingsfees](http://www.standardandpoors.com/usratingsfees).