

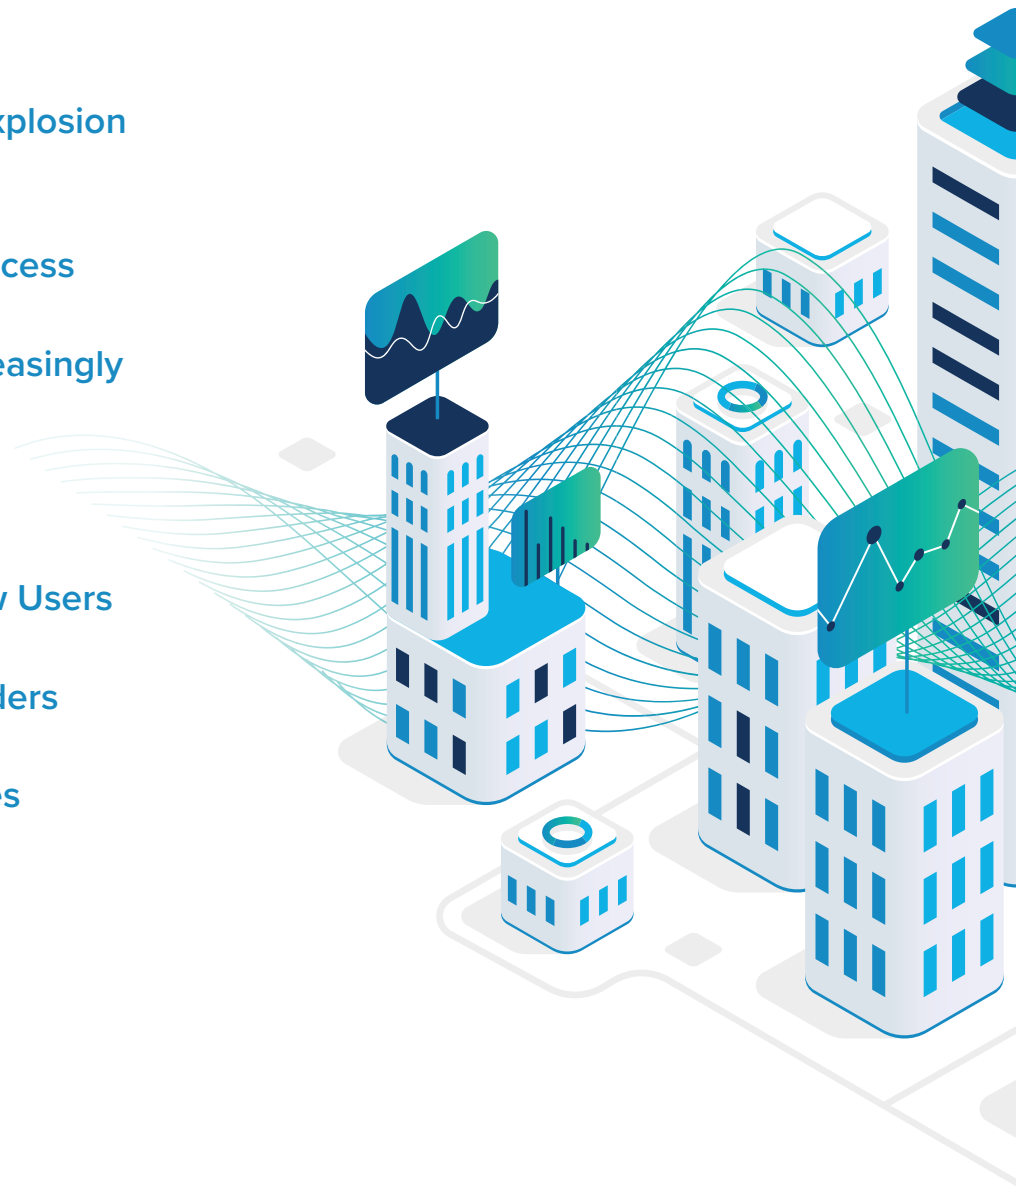


# 2022 Service Mesh Adoption Survey

Insights from technology leaders and practitioners on how microservices, Kubernetes and Istio are being adopted.

# Table of Contents

3	Executive Summary
4	Key Findings
6	The Microservices Explosion
10	API Gateways are Foundational for Success
12	Service Mesh is Increasingly a Core Requirement
14	Service Mesh is Still Emerging and Can Be Confusing to New Users
16	Learn From The Leaders
22	Service Mesh Choices
24	Key Observations & Recommendations
25	Methodology



# Executive Summary

To succeed, companies continually push to do more, faster, with less. Digital transformation and the shift to a cloud-first posture or even fully cloud native development has revolutionized the ability of companies to deliver more products and features faster by unlocking developer productivity. Leading companies practically equate developer productivity and developer happiness with increased revenue and customer satisfaction.

The broad adoption of containers and the emergence of Kubernetes as a de facto standard have pushed the boundaries of cloud native development's impact even further. Containers, though not a prerequisite for microservices, provide a physical vehicle for the mass adoption of microservices. As our recent survey<sup>1</sup> shows, we now see roughly comparable levels of adoption for both microservices and containers.

With all new tech, however, when one problem is solved, another emerges. For Kubernetes, there exists a robust market of service providers to resolve that pain for most companies. And for those who choose to operate Kubernetes in-house, the solutions, skills, and smarts now exist—to say nothing about the maturity of the products—to make that concept orders of magnitude easier than just a few years ago.

Microservices, on the other hand, remain an acute and often underappreciated pain. There has been an explosion of microservices and APIs. Most companies are still sorting out how to deal with this new form of technical debt. As they seek to increase development velocity and make developers more productive, modern enterprises now struggle to deal with the negative externalities that result from the microservice and API sprawl.

To solve for this, enterprises are turning to service mesh and API gateways. These technologies are exactly the prescription for their ills. Increasingly, they realize this. They see service mesh and API gateways as exciting solutions. The problem is, they don't know exactly how to make the best use of them or what the best practices for these tools are.

Some companies have figured this out, however. By looking at these leading companies and how they are using service mesh and API gateways, other companies can reap the same benefits and maximize the positive impact to their dev teams and KPIs.

**We hope this [survey](#) and report will help inform companies looking to take advantage of service mesh and API gateways.**



**We now see roughly comparable levels of adoption for both microservices and containers.**

1. In January 2022, we conducted a survey of 704 executives, managers, and practitioners to learn more about their use of microservices, service mesh, Kubernetes, and other related technologies. A fuller methodology report is available at the end of this report.



There is an explosion of microservices



Companies are turning to service mesh as an answer



Service Mesh is relatively new; best practices are not widely known



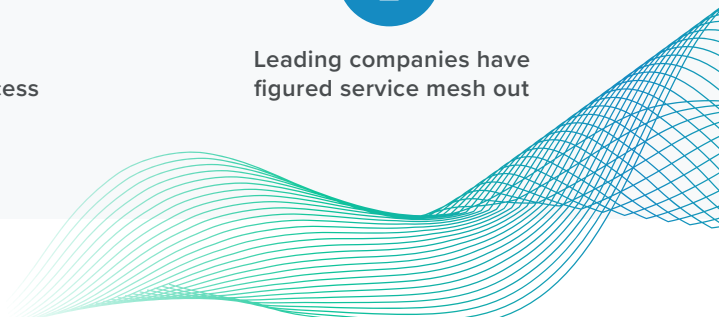
Kubernetes has crossed the chasm



API gateways are foundational for success




Leading companies have figured service mesh out



## Key Findings

Organizations demand more from their development teams and their technology infrastructure than ever before. The increasing pressure to ship more and better features, to improve the customer experience, and do all of this faster requires flexibility, automation, and abstraction. A key technological development enabling this is the explosion of microservices. Managing those microservices has become its own challenge, and organizations are turning to service mesh to solve it.

- 1. There is an explosion of microservices.** Microservices are being used more than in the past. Fully 85% of companies say they are modernizing their apps to a microservices architecture. This transition is both new and large scale; customers still haven't figured out how to get the best benefits from them, let alone how to manage it all.
- 2. Kubernetes has crossed the chasm.** Nearly two-thirds of companies (64%) are using Kubernetes in production to some degree. More importantly, a majority of companies (53%) report at least half of their production workloads running on Kubernetes, and they have figured out how to make the most of it, as satisfaction with Kubernetes is nearly universal (94%).
- 3. API gateways are foundational for success.** API gateways are a stable market. Fully 93% of companies are using or evaluating an API gateway. The benefits (primarily to ensure service reliability) are clearly understood and align to organizations' increasing use of microservices and app modernization strategies.




85% of companies say they are modernizing their apps to a microservices architecture

4. **Companies are turning to service mesh as an answer.** As microservices architectures become a requirement for competitive companies, service mesh has grown in popularity as an organizational tool. Nearly half of all companies (49%) report using a service mesh at some level with a further 38% evaluating a service mesh for use.

5. **Service mesh is relatively new; best practices are not widely known.** Because service mesh is an emerging market, there are fewer established best practices in the way there are for API gateways. The reasons for adopting service mesh are varied with less broadly recognizable trends.


6. **Leading companies have figured service mesh out.** There is a strong correlation between success with microservices and faster, more reliable application development.

- **Success with microservices means faster development cycles.** A majority of organizations (56%) with at least half of their apps on microservices architectures have release cycles that are daily or more frequent. On the other hand, only 17% of those with less than half of their apps on microservices architectures report daily or faster release cycles.
- **Service mesh leads to more reliable apps.** An overwhelming majority of organizations (89%) report very positive impact on app reliability as a result of using service mesh. This includes 44% who say the impact is “transformative.”



■ Using    ■ Evaluating

**87% of companies report using or evaluating a service mesh for use**



**89%**

of organizations report **very positive impact** on app reliability as a result of using service mesh.

**44%**

say the impact is **“transformative.”**

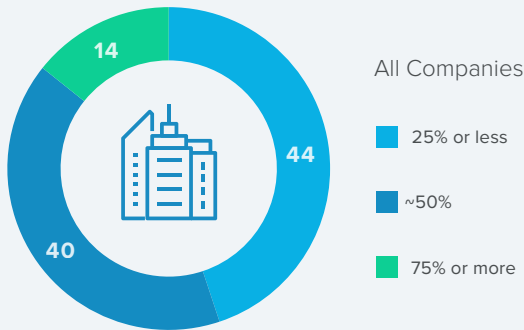
## The Microservices Explosion

Microservice architecture is the new normal. A majority of companies (54%) report that at least half of their apps use a microservices architecture. Moreover, an overwhelming majority (85%) of companies say they are modernizing their apps and moving to a microservices architecture.


**Leading\*** companies are 2x as likely to report 75% of their apps on microservices architecture compared to others

### Percentage of Apps on Microservices Architecture

Q: Approximately what percentage of your apps use a microservices architecture?



### Companies, by Maturity



In parallel with this trend is the now-widespread adoption of Kubernetes and containers. It is no surprise at this point that Kubernetes has “crossed the chasm,” as the most recent [State of Cloud Native Development Report by CNCF](#) stated. Our data reflects this trend as well, with 64% of companies reporting usage of Kubernetes in production at some level, and a further 31% evaluating Kubernetes. Of those using Kubernetes in production, a majority (53%) now report at least half of their production workloads running on Kubernetes.




In the early days of Kubernetes, it was more hype than reality. It was not a panacea for digital transformation as many wanted it to be. It brought its own pain. Today, companies have largely solved that pain. An overwhelming majority of organizations now report a high degree of satisfaction with Kubernetes 94% rated 4 or 5 on a 5-point scale. A similar share (92%) plan to migrate additional workloads to Kubernetes, including a majority of enterprises reporting a very high likelihood of doing so.

## Satisfaction with Kubernetes is high, especially amongst Enterprise orgs


### K8s satisfaction

Q: In general, how satisfied are you with the use of Kubernetes for production workloads in your organization? Very satisfied or not at all satisfied?



### Likelihood of migrating additional workloads to K8s

Q: How likely are you to migrate additional production workloads to Kubernetes in your organization? Very satisfied or not at all satisfied?



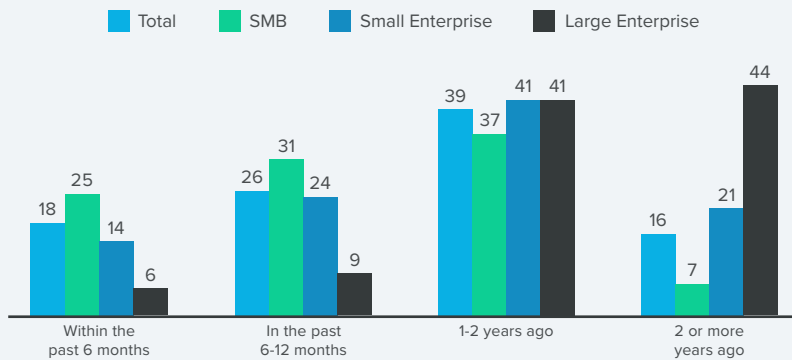
★ ★ The data shows that Large Enterprises are more likely to migrate workloads to Kubernetes (56% instead of 49% rating it very likely).

The widespread adoption of Kubernetes in production is a recent phenomenon. Only 16% of organizations report using Kubernetes for over 2 years, and nearly half (46%) report using Kubernetes for less than a year.

### When organizations started to use Kubernetes by company size

Q: When did your organization start to use Kubernetes?


Kubernetes use start, by company size



Mostly, organizations run Kubernetes on their own, at least to some degree. Only a small minority (21%) exclusively use a service provider. Of those who run their Kubernetes clusters on a service provider for all or some of their Kubernetes workloads, the hyperscalers have the lion’s share of users. Amazon Elastic Kubernetes Service (EKS) and Google Kubernetes Engine (GKE) are used by most, followed by Azure Kubernetes Service (AKS).

### Does your organization run Kubernetes clusters directly or use a service provider?


Q: Does your organization run Kubernetes clusters directly or do you use a service provider?





## EKS and GKE are the leading services in organizations


Q: Which of the following services, if any, does your organization use?



Despite the broad adoption of Kubernetes and microservices, we are still not at a fully mature market. For example, only 5% of companies: a) use microservices and Kubernetes in production, b) have deployed broadly across their organizations, and c) are very satisfied with Kubernetes. In other words, companies are using Kubernetes, in production, at scale. Companies are using microservices, in production, at scale. But only a small minority are doing both and are satisfied. There is still a lot of growing to do. There is still a lot to learn. That some companies have successfully embraced both microservices and widespread production utilization of Kubernetes makes this a superpower for them. The question we ask is—how can we make this a superpower for others?

## How many Kubernetes clusters is your organization running?

Q: Approximately how many Kubernetes clusters is your organization running across all management services?



The challenge today is less about Kubernetes and less about microservices, per se, and more about managing the explosion of services.

To that end, the challenge today is less about Kubernetes and less about microservices, per se, and more about managing the explosion of services. Just like Kubernetes, microservices are not a panacea if you don't know what you're doing. Leading companies have figured this out through the purposeful adoption and deployment of API gateways and service mesh. Most companies recognize this trend and are tackling it but have not quite figured out the "purposeful" part yet.

# API Gateways are Foundational for Success

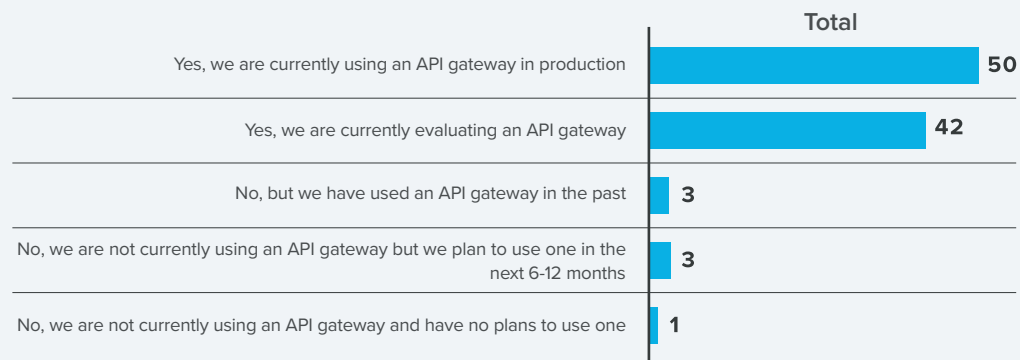


The proliferation of APIs has had profound effects on application development and what companies can deliver to their customers. Enterprises today leverage any number of APIs to empower developers and organizations alike—to create new revenue streams, improve customer experience, or simply streamline operations processes. At this point, API gateways or ingress providers are foundational for a competitive company. **Fully 93% of companies are using or evaluating API gateways or ingress providers** to manage their API requests, security policies, rate limits, etc., and 85% have an API management program in place specifically to increase collaboration and reuse.

And of course they are. How else could they manage? The API economy is not new, and over time, API gateways have developed into a sort of Swiss army knife—one reliable tool with many purposes. This is reflected in how companies think about them and their utility. **The primary role they serve is ensuring “service reliability” with 45% of companies citing this as their primary purpose.** At large enterprises (companies with over 1,000 employees), this increases to 59%.


## Most companies are using or evaluating API gateways or ingress providers

*Q: Do you use an API gateway or ingress provider?*



## Service reliability is the most important feature of API gateways and ingress providers

Q: Which THREE of the following features of an API gateway or ingress provider are the most important for you? (Select up to three.)

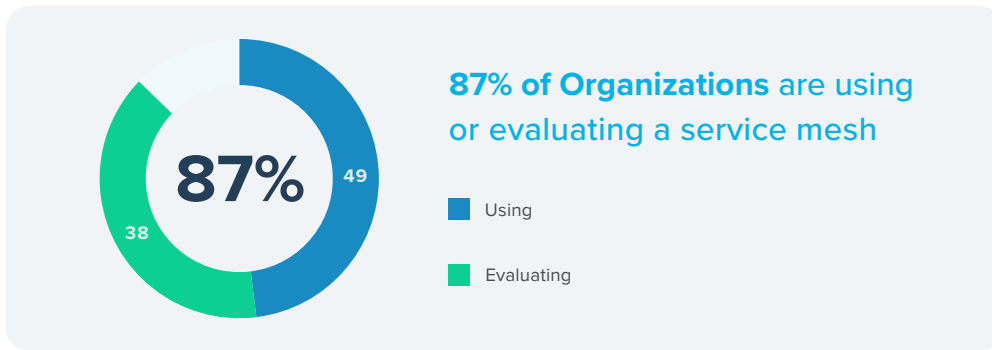


“APIs are a game-changer. We take for granted the kind of customer experience we can provide now with far less effort, less developer time, faster development cycles. Literally hundreds of APIs. We couldn’t possibly manage without an API gateway. Collaboration would be sluggish. Something would fall over. Something would get missed and months later, we’d be scrambling to figure out the cause of failure.”

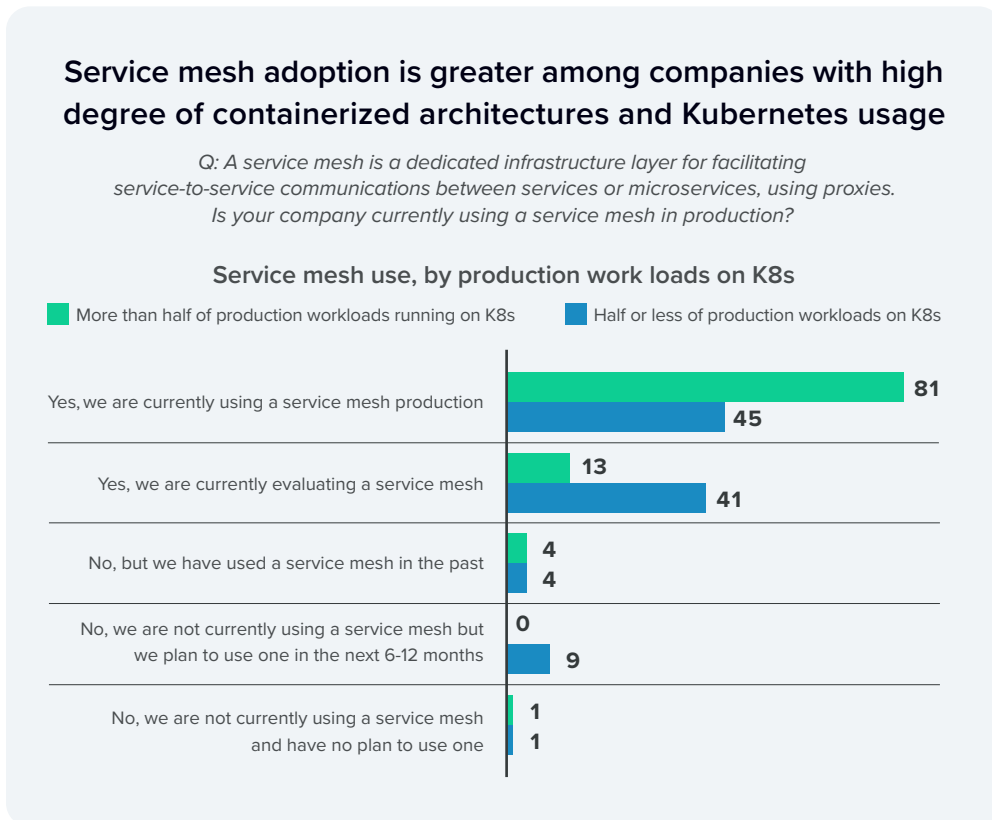
CTO, LARGE FINANCIAL SERVICES ORGANIZATION

## Service Mesh is Increasingly a Core Requirement

In response to the increased usage of microservices, service mesh adoption is increasing significantly. In the past two years, service mesh adoption has grown from 30% (as reported by CNCF in 2020) to 49% in our survey, with a further 38% evaluating a service mesh for adoption.



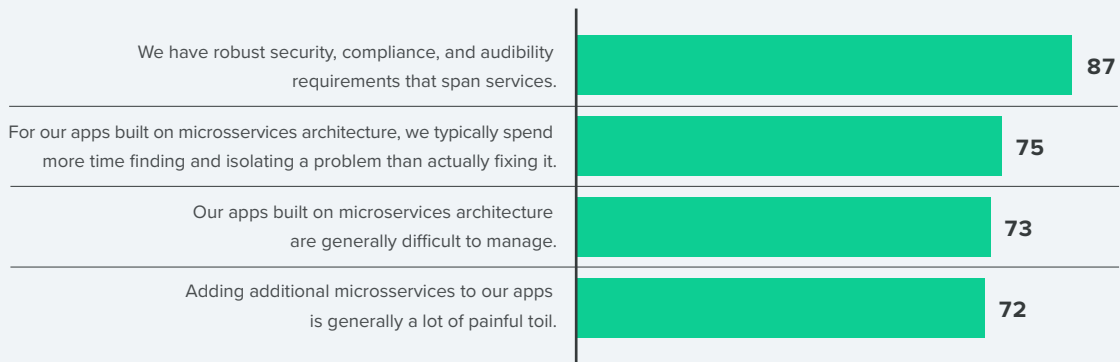
As one might expect, service mesh adoption is greater among companies with a high degree of containerized architectures and Kubernetes usage. Among those organizations with more than half of their production workloads running on Kubernetes, 81% use a service mesh, compared to only 45% of those organizations with half or less of their production workloads on Kubernetes.



Service mesh adoption is likely to continue on this trajectory. A large majority of companies building apps on microservices architecture exhibit behaviors that suggest great value in an appropriate and purposeful deployment of a service mesh. Whether it is service-spanning requirements or a general difficulty managing the volume and diversity of services in their microservices architecture, organizations' needs for the capabilities of a service mesh are manifest.

### Microservices benefit from service mesh deployment

*Q: Please indicate your level of agreement with the following statements.  
Use a 5 point scale in which 1 means you do not agree at all, and 5 means you completely agree.*

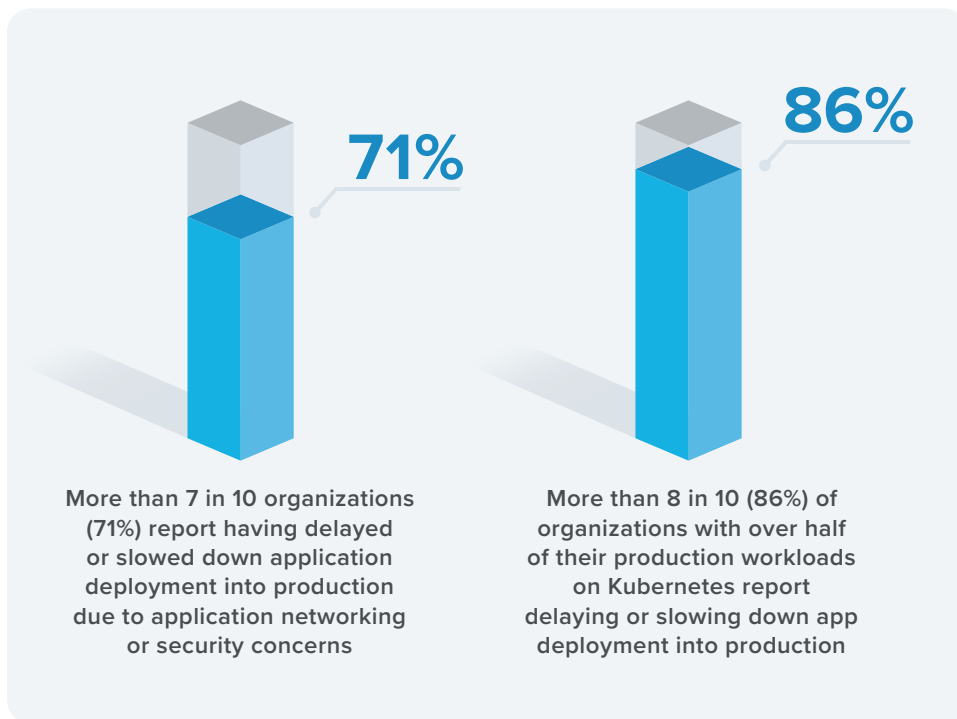


*Percentage of companies responding with a 4 or 5*

## Service Mesh is Still Emerging and Can Be Confusing to New Users

Unlike API gateways, the precise utility of service mesh is not yet clear to all organizations, even to those who have begun using them. Moreover, the best practices, usage-patterns or even necessity of service mesh isn't clear to all users.


They know they have pain— more than **7 in 10 organizations (71%) report having delayed or slowed down application deployment into production** due to application networking or security concerns. This share increases for those deploying broadly to Kubernetes. More than **8 in 10 (86%) of organizations with over half of their production workloads on Kubernetes report delaying or slowing down app deployment into production.** Additionally, 74% of organizations say managing the application networking for multiple clusters is time consuming and error prone.



Most companies that are aware of service mesh use it and agree that a service mesh provides value, but the value varies among organizations and groups. Across organizations the benefits accrued from service mesh are extremely varied. Improvements to application reliability and security—the two most-reported benefits—don't even receive top marks from a quarter of organizations.

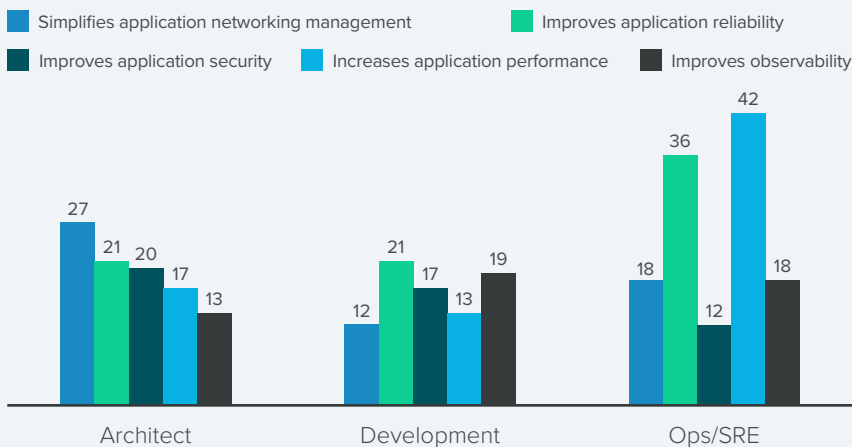
## Improvements in application reliability and security are most valuable service mesh benefits

Q: Which TWO of the following benefits are most valuable for you from a service mesh?



Digging deeper, we see that application architects value simplified application networking the most, followed by the common benefits of reliability, security, performance and observability. Not surprisingly, application development teams value application reliability and observability as their top choices. What is surprising is the extent to which operations teams and SREs put on improved application performance and reliability. By a wide margin, operations teams value these two benefits over all the other benefits provided by a service mesh.

## Service mesh benefits vary widely between groups



This is a common story with rapidly evolving technologies organizations are newly deploying. Teams see value in many places, but the value varies by application architecture, organizational experience with underlying technologies, and the different teams within the organization, each with their own requirements and goals. And as with any relatively “new” technology claiming to solve your organization’s pain, service mesh is in its “foggy visibility” stage.

## Learn From The Leaders

As with so many new trends, some organizations figure this out before others. Often, they are solving for a specific need and have to patch together solutions with in-house development before commercial or open source solutions exist upon which to build. Whether these solutions themselves become standards (e.g., Kubernetes) or if they are simply a trailblazer worth following, it behooves companies to learn from others who were here first.

However, in doing so, we should be selective. We ought not turn simply to all organizations deploying a new solution to a problem similar to our own. Rather, we should look specifically to those who are using the solution purposefully and succeeding as a result. How are leading companies selecting and deploying these technologies? What features are they looking for, and what benefits do they gain as a result? How are they winning, and how can we learn from them so we win, too?


Before we answer those questions, let's look at the impact service mesh is having on the Leaders. At the highest level, we must ask—is the headache of microservices worth it? The answer from Leaders is plain—while microservices can create pain, service mesh is a powerful and relatively easy salve for that pain.

### Determining what is a Leading company

To segment companies, we categorized them on a continuum from those who are “Leading” to those who are merely “Starting” on their journey to modernizing their application development and data-driven capabilities. To create this continuum, we partnered with ClearPath Strategies and used their model which has been developed and rigorously validated through numerous corporate research projects. Clear Path developed this model using over 75 variables and both factor and conjoint analyses. The result was a set of 10 variables which were most strongly correlated with companies that exhibit successful patterns that lead to business outcomes, including increase in revenue, customer satisfaction, and developer productivity.

Using the model to analyze our survey results, ClearPath found **15% of the companies are in the “Leading” category while 26% are in “Advancing”, 35% in “Aspiring” and 24% in “Starting”.**

### Breakdown of companies

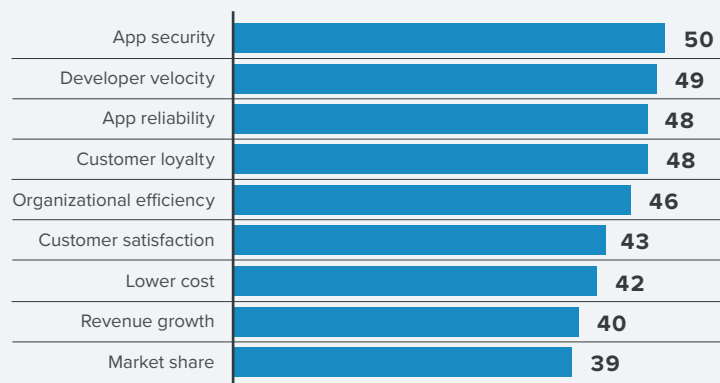




In nearly every category we survey, from “app security” and “app reliability” to more business-oriented outcomes such as “customer loyalty” and “customer satisfaction,” large shares of Leading companies describe the impact of service mesh as transformative.

### Transformative positive impact of service mesh for Leaders

Q: How would you describe the impact of a Service Mesh on each of the following?



So how are these Leaders using service mesh? How is it so transformative for so many? And what are they using it for? There are five basic requirements for service mesh:

- 1. Connectivity - Routing and traffic management.** The fundamental challenge of the microservices explosion is the amount of configuration required for the sheer volume of API calls between services. As the number of services grow, the task of managing routing, load balancing, circuit breaking and resiliency becomes increasingly complex and automation becomes critical. Leaders want to go beyond traffic management and want the ability to set policies across different sets of services as well as features like service discovery, global failover and multi-cluster & multi-mesh management.
- 2. Security.** Leaders view security as one of the top priorities, which is no surprise. Security has always been a basic requirement, but as more companies deploy more to the cloud, more frequently, with more complexity and dependencies, the challenge and priority of security has risen to a new level of urgency. What distinguishes the Leaders is their definition of security. For them, “basics” not only include authentication, authorization and encryption, but expands to areas like credential management and role-based access control (RBAC).

“Our transition to microservices has been a little bumpy, but there was never any question about the direction. Microservices, containers, Kubernetes. The complexity started to slow us down. **Service mesh was an obvious and easy win. It just worked. Reliability was up. Monitoring was easy. Updates were basically instantaneous. We did more in weeks than we had done in the previous half year, and all our KPIs jumped up.**”

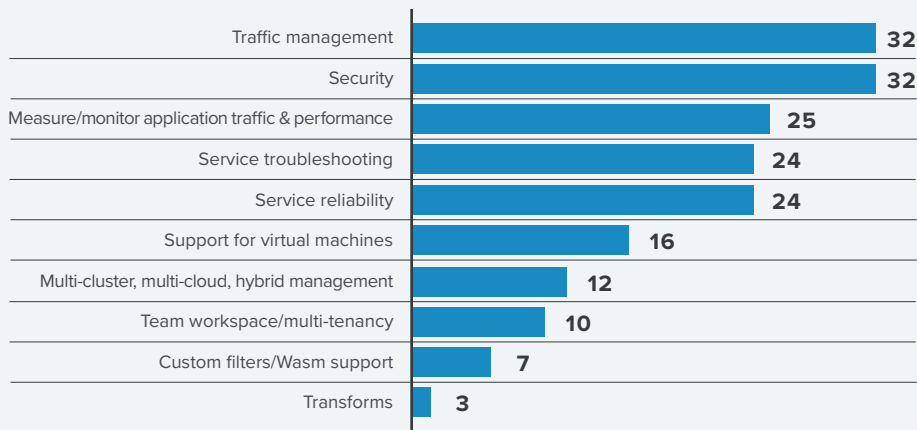
CLOUD ARCHITECT, GLOBAL MANUFACTURING ORGANIZATION

3. **High-Availability - Service reliability.** Leaders also look to service mesh to add reliability features to their cloud native applications. For example, Istio (one of the leading deployments used by Leaders) uses a sidecar proxy to cache data, so the data plane remains separate from the control plane. This provides higher resilience should a pod fail.
4. **Observability - Monitoring traffic and performance.** Alongside the traffic management, Leaders expect a service mesh to provide a robust, yet clean set of monitoring metrics as a single source of truth for all microservices. Latency, traffic, errors and saturation—the four “golden signals”—allow teams to improve service reliability in the most toil-free, low-friction way.
5. **Service troubleshooting.** Because service mesh is relatively new, Leaders look to deployments with a robust set of troubleshooting tools. Most companies just starting on their service mesh journey will get by with basic troubleshooting tools and configuration debuggers. More advanced Leaders (and in particular larger enterprises with larger microservice footprints), find themselves needing more sophisticated features, including multi-cluster and hybrid management debuggers.

**Latency, traffic, errors, and saturation—the four “golden signals”—allow teams to improve service reliability in the most toil-free, low-friction way.**

### Table stakes service mesh features for Leaders


*Q: Which THREE of the following features of service mesh would you consider basic requirements from your service mesh? (Select up to three.)*



Service mesh products live in an ecosystem of development and operational technologies and systems. For Leaders, the most common integration is connecting their service mesh and API gateway, often with the goal of leveraging the same data plane technology across both environments. Integrating operational metrics into tools like Prometheus, Jaeger, and/or Datadog was the second most common integration requirement followed by integrating the service mesh with access control systems as part of an overall security strategy.

### Service mesh integrations for Leaders



Q: Which of the following technologies is your organization integrating/planning to integrate into their service mesh? (Select all that apply.)



Beyond these basic features, Leaders also can show us the cutting-edge and next frontier feature sets. We ask these Leaders which features they would want from a service mesh that are not included in their current deployments. DLP, WAF, serverless integration and multi-cluster observability rise to the top.

### Wishlist service mesh features for Leaders

Q: From this different list, which THREE of the following features of service mesh are NOT currently in your product or are you not currently using, but WOULD be valuable to you?



Interestingly, all of these features are available, depending on the service mesh deployment organizations choose. It is important for companies looking to adopt and deploy a service mesh to consider their potential long-term needs as well, rather than simply reaching for what is easy/available.

Beyond these features, there are a number of stark contrasts between what Leaders value in a service mesh as compared to organizations that are at the Starting or Aspiring point in their service mesh adoption journey. These differences can often be attributed to their greater experience, larger environments and higher operations requirements than companies just starting out.

- Leaders operate in larger environments with multiple clusters. Not surprisingly, they view multi-cluster management as a critical feature by a 2 to 1 margin over companies at the Starting phase.
- By an even larger 3.5 to 1 margin, Leaders understand the time, effort and potential issues that arise when managing traffic in large, multi-cluster environments.
- Leaders prefer an Istio-based service mesh by almost a 3 to 1 margin.

**Fully 45% of Leaders strongly agree with the statement “Our organization prefers an Istio-based service mesh over alternative service mesh architectures/products.”**

### Leaders need the things Istio provides


Their world is more complex, and they need the added power/features

*Q: please indicate your level of agreement with the following statement.*


*Use a five point scale in which one means you do not agree at all, and five means you agree completely.*

Agree Completely


■ Leaders ■ Aspiring ■ Laggards



Our organization prefers an Istio-based service mesh over alternative service mesh architectures/products



Managing the application networking for multiple clusters is time consuming and error prone

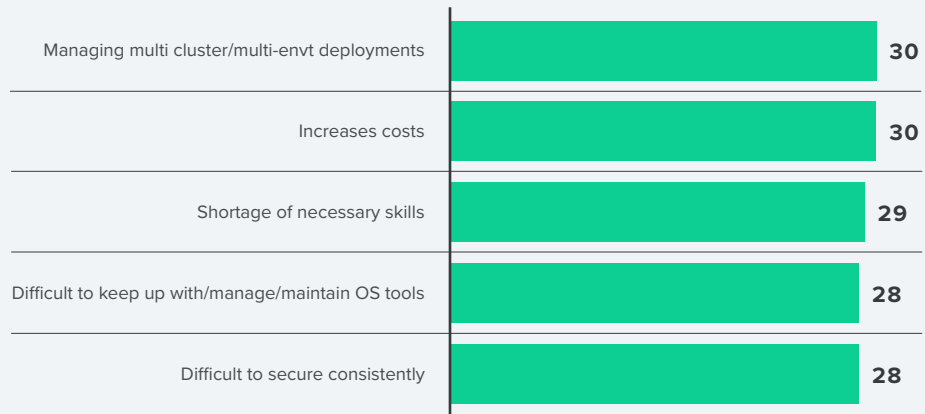


Have you ever delayed or slow down application development into production due to application networking or security concerns

Mostly, orgs using either Istio or non-Istio solutions report similar benefits. Leaders are looking at and deploying a range of service mesh options, but the challenges these Leaders continue to face are not solvable by all service mesh options.

### Top service mesh challenges for Leaders

Q: What would you say are the three biggest challenges with using a service mesh? Select up to THREE.



## Service Mesh Choices





No service mesh has all the features an organization needs. For the most part, orgs using either Istio or non-Istio solutions report similar benefits. As noted above, Leaders prefer an Istio-based solution by 3:1 over those at the earlier stages of their service mesh journey.

This distinction is partly based on immediate needs. Leaders are looking at and deploying a range of service mesh options, but the challenges these Leaders continue to face are not solvable by all service mesh options. Those companies which are just starting out may not need the same things today. The question is whether to plan for tomorrow’s needs today, or simply to worry about tomorrow tomorrow.

Use cases among users of the most common service mesh deployments—Istio, Linkerd, and Kuma—vary significantly. Apart from the more-or-less-universal shortage of skills, the pain different users of different service meshes are solving for varies greatly.

### Biggest Challenges for Istio, Linkerd, and Kuma users

Q: What would you say are the *THREE* biggest challenges with using a service mesh




■ Istio   ■ Linkerd   ■ Kuma

Not only do the use cases vary, but the features users wish they had in their service mesh also vary greatly. Here again, there is a universal need for VM support, which is not surprising when we think about enterprises confronted with the decision of whether and what to migrate to cloud.


### Most desired features for Istio, Linkerd, and Kuma users

Q: THREE of the following features of a service mesh are NOT currently in your product


Istio Linkerd Kuma




### LINKERD



### Kuma








In the end, there is no single choice for service mesh. What matters is knowing what you want to achieve with it and considering what you may want it to do for you in the future.



## Key Observations & Recommendations

One constant during the age of digital transformation has been to follow the Leaders. Technology moves too fast for most organizations or practitioners to keep up. Rather than reinvent the wheel, it is often easiest simply to look at what Leaders are doing, find relevant use cases, and apply the best practices to your own environment.

The early days of broad service mesh adoption in response to the microservices explosion and adoption of Kubernetes in production is no different. Leaders who have purposefully committed to broad deployment of container-based architectures have made great use of service mesh technology to solve for the headaches of managing application reliability, security and observability. Here are some trends and best practices we have observed from those Leaders that other organizations could adopt to enjoy transformative impact on app reliability, increased developer productivity, and customer satisfaction.

-  **The market adoption** of microservices and Kubernetes provides a strong foundation to layer on service mesh to solve a broad set of microservice-related issues.
-  **Use service mesh** for its delivery of cross-functional capabilities to infrastructure, security and application teams. Its capabilities should align with platform teams that have those responsibilities.
-  **Success with microservices** tends to align with organizations that release software more frequently and have good CI/CD and GitOps practices.
-  **As more microservices are deployed**, API gateway and service mesh functionality will eventually overlap. Platform teams should look to enable consistent technology, management and policies across both of those functions.
-  **As most companies leverage hybrid and multi-cloud strategies**, they should look at service mesh offerings that enable consistency in any cloud or Kubernetes environment.



## Methodology

This survey was conducted and produced by [ClearPath Strategies](#), a strategic consulting and research firm for the world's Leaders in technology, business, and progressive politics. Below is the detailed methodology of the survey conducted for this report.

### Respondent selection

The survey included 704 respondents sourced from a leading global online panel provider. They were selected from the panel based on geographic and role-based quotas, as well as screening questions based on role/title, decision-making role, and company size. Selected respondents were further screened based on knowledge of general technology issues and basic understanding of microservices to ensure respondents were sufficiently informed to provide useful information for this study.

### Roles and Titles


The survey included 21% of respondents who were developers, 25% who were CIO/CTO levels, and 39% who were VP/Director level. Respondents were asked to select which role—from a list of 16 options—most closely described their primary responsibility, even if no one was quite right or even if they performed more than one of these roles.

### Decision-making level

Respondents in the survey were mostly key decision makers in licenses and new technologies (84%). The remaining respondents either directly influenced these decisions (12%) or provided necessary feedback to the decision-making (4%).

### Company size

The survey included respondents from a range of company sizes, including 49% from companies with 250-999 employees, 44% from companies with 1,000 to 9,999 employees, and 7% from companies with 10,000 or more employees.



**84% of respondents in the survey were mostly key decision makers in licenses and new technologies.**

## Geographic representation

The survey included respondents from the US (42%), the UK (20%), Singapore (15%), Germany (10%), Canada (5%), Brazil (3%), Italy (2%), Japan (2%), France (1%), and Australia (1%).

## Industry representation

Although no industry-level quotas were deployed, we monitored the data to ensure that no single industry was over-represented in the data. The final breakdown of respondents by industry is as follows: IT (software, hardware, services) 20%, Retail and Wholesale 17%, Financial services 16%, Manufacturing 15%, Transportation and logistics 9%, Health care 7%, Light industry 5%, Infrastructure/construction 2%, Primary goods 2%, Government 1%, Education 1%, Non-financial services services 1%, Telecommunications/ISP/Web hosting 1%, Media 1%, Utilities 1%, and Other 1%.

## Margin of error

It is technically impossible and improper to list a margin of error for a survey of this type. The respondents for this sample were drawn from an online panel with an unknown relationship to the total universe, about which we also do not know the true demographics. As such, the exact representativeness of this, or any similarly produced sample, is unknowable.

[Talk to an Expert](#)



[solo.io](#)

[contact us](#)

## About Solo.io

Solo.io, the modern service connectivity company, delivers application programming interface (API) infrastructure software that makes it easy for your architects and engineers to manage application traffic. As you move to cloud, microservices, Kubernetes containers, and serverless functions, you need a secure and reliable approach to application networking, with unified observability and control. Solo builds on open source Envoy Proxy and Istio to give you comprehensive API gateways and service meshes that work everywhere, at any scale.