



Service meshes are
on the rise — but greater
understanding and
experience are required

CNCF conducted a microsurvey of the cloud native community at the end of 2021 to discover how organizations are adopting service meshes.

Cloud native is driving digital transformation, with organizations keen to capitalize on the agility and flexibility it provides to their business and operations. But as more applications and services are deployed using a diverse technology stack, it has become a challenge to deliver and manage performance and availability.


A service mesh provides an answer, creating a dedicated layer that handles service-to-service communications that ensures consistency and reliability of services, security, and observability. Moreover, a service mesh available off-the-shelf, as a community-driven, open source project, means organizations can avoid the challenges and overhead of building their own, thereby reducing the support burden.

It's no surprise that service mesh has become a key component of cloud native infrastructures. Of the 253 survey respondents, 70% run a service mesh in production or development, and 19% in evaluation mode. Some said they are running a service mesh for clients who had high levels of adoption. Those not implementing a service mesh formed a distinct minority — fewer than 10%.

Service mesh adoption is running hand-in-hand with the rollout of Kubernetes clusters. The majority of participants (65%) run or plan to run between two and ten Kubernetes clusters on a service mesh. Another 11% are operating or planning to operate between 11 and 25, with just 10% going further with 26 or more clusters.


The number of projects and products has mushroomed in response to the spread of service mesh. The differences between them vary according to the breadth of features, ease of deployment, use cases (such as edge), or

Do you/your organization currently run a service mesh?



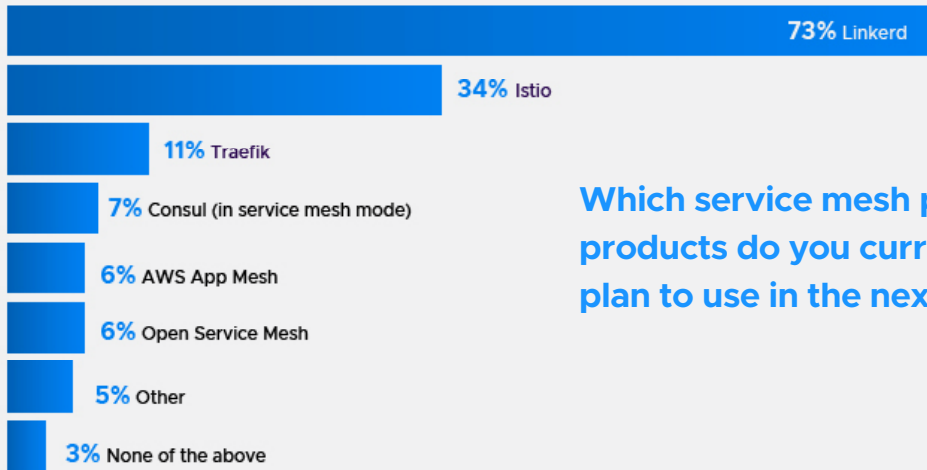
- **60%** Yes, in production
- **19%** We are evaluating a service mesh
- **10%** Yes, in development
- **9%** No

How many production Kubernetes clusters do you expect to run a service mesh on during the next 12 months?



- **9%** 1
- **65%** 2-10
- **11%** 11-25
- **10%** 26 or more
- **5%** None

optimization for different levels of the network stack. Participants picked from a list of 15 projects or products they currently use or plan to use in the next year. Two led the field: Linkerd and Istio with 72% and 34%, respectively — a clear expression of confidence in open source.



Which service mesh projects/ products do you currently use or plan to use in the next 12 months?

It's all about features

We explored the factors influencing people's choices by asking which features and capabilities drive their organization's adoption of service mesh. Security is a top concern, with 79% putting their faith in techniques such as mTLS authentication of servers and clients during transactions to help reduce the risk of a successful attack.

Observability came a close second behind security, at 78%. As cloud infrastructure has grown in importance and complexity, we've seen a growing interest in observability to understand the health of systems. Observability entails collecting logs, metrics, and traces for analysis.

Traffic management came third (62%). This is a key consideration given the complexity of cloud native that a service mesh is expected to help mitigate. As organizations seek to run more sophisticated, blue/green deployment scenarios, traffic management can apply to a range of capabilities, including discovering endpoints and services, controlling API calls between services, and hiding or exposing services.

In fourth place was reliability with 56%. Potential issues here include latency, lack of bandwidth, security incidents, the heterogeneous composition of the cloud environment, and changes in architecture or topology. Respondents want a service mesh to overcome these networking and in-service communications challenges.

Which of the following features and capabilities are driving your adoption of service mesh?



Support for multi-cluster communications was of significant interest, with 87% in total classifying it as important. Breaking that figure down revealed where organizations are in their deployment: 50% described multi-cluster communications as “somewhat” or “very” important, while the 37% who said it was not important at the moment but would be in the future may be in the development stage, pilot mode or, running a relatively simple production-level service mesh.

Users expect to implement a broad array of service mesh features in the coming year. The most popular (46%) is header-based routing control, making it easier to implement more sophisticated capabilities such as A/B testing and blue/green deployments.

A cluster of features came next. The ability to detect and manage the presence of shadow traffic/dark traffic will be important for 29%, suggesting a desire to collect system data that is as accurate as possible by accounting for traffic that analytics tools might otherwise miss.

Extending the service mesh to work with non-Kubernetes environments was important for 27%. Kubernetes might have become a fundamental building block of cloud native, but this response indicates the strong presence of alternative environments. Extending a service mesh to encompass those is consistent with the bigger drive to eliminate silos in IT infrastructure and management. Similarly, plug-ins to environments such as WebAssembly were important for 25%.

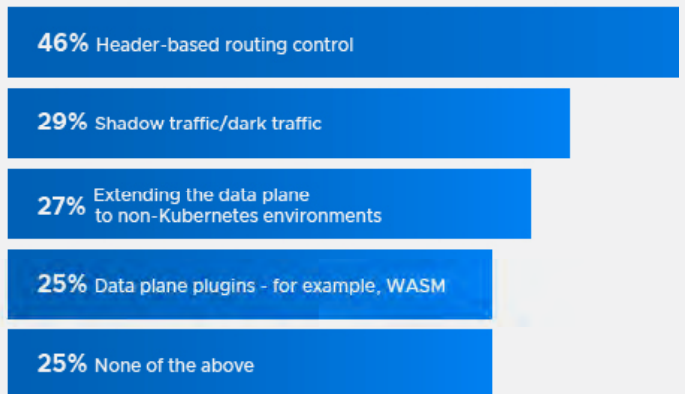
Interestingly 25% also gave “none of the above” as their choice of new feature. Rather, authorization and authentication emerged as the most common response.

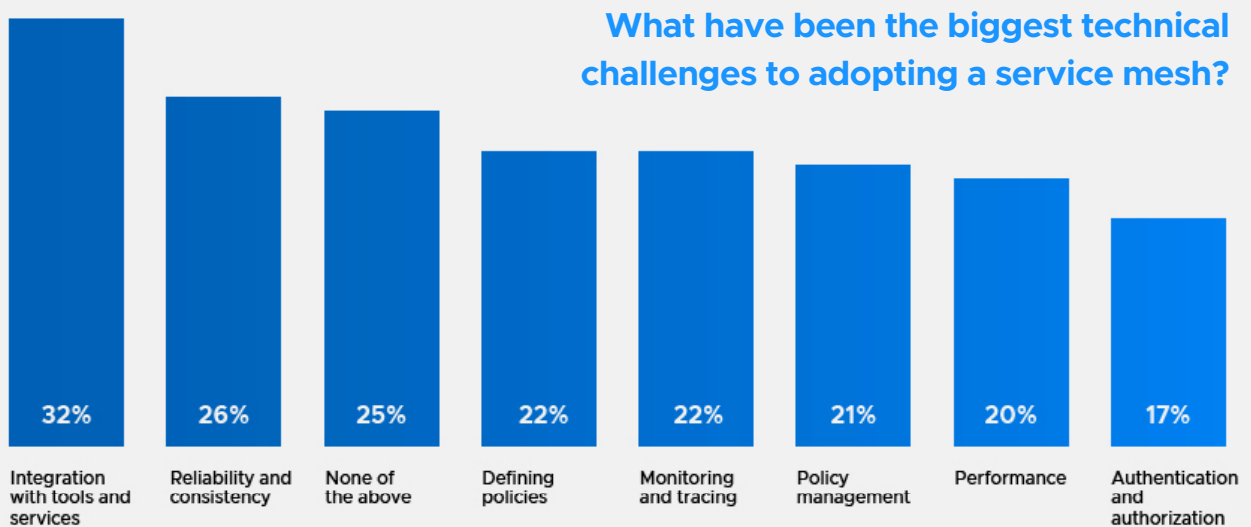
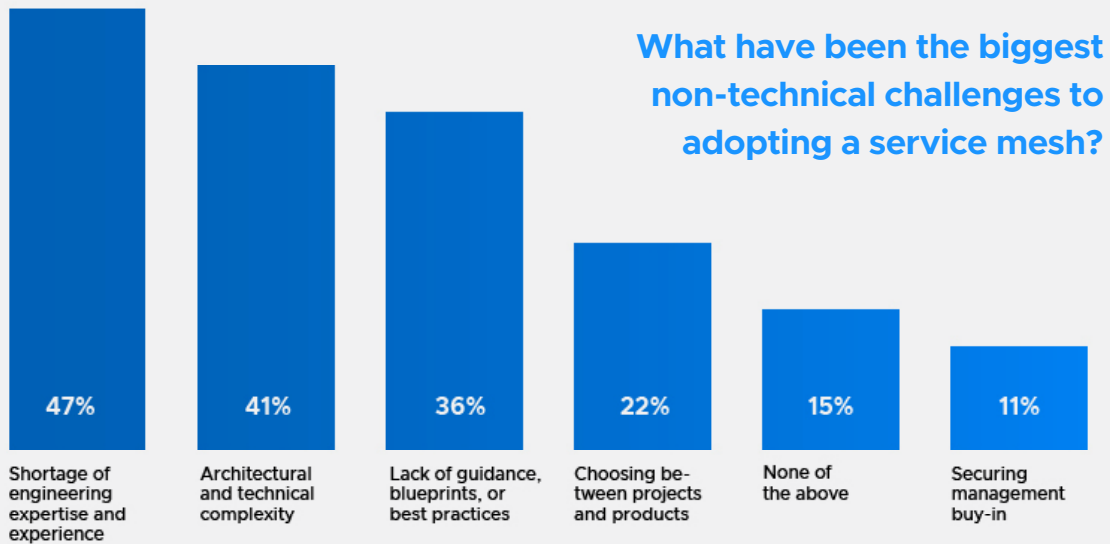
How important is multi-cluster communication to you?



- **37%** Not important today, but will be in the future
- **27%** Very important
- **23%** Somewhat important
- **13%** Not important

Which of the following service-mesh features do you expect to use in the next 12 months?





Challenges remain

There's plenty of interest in service mesh, but there are hurdles to adoption, so we asked survey participants to rank them. First, non-technical challenges: the top three will be familiar to any newer technology – shortage of engineering expertise and experience (47%), architectural and technical complexity (41%), and a lack of guidance, blueprints, and best practices (36%).

When questioned about technical challenges, respondents reported struggling in various areas. Integration topped the list (32%), followed by reliability and consistency (26%), defining policies (22%), monitoring and tracing (22%), and policy management (21%). A quarter of respondents threw in additional challenges under “none of the above.” When asked to explain further, they listed CI/CD integration, difficulty troubleshooting, and problems with specific products.

Methodology

The microsurvey was designed by CNCF and conducted between November and December 2021 among 253 members of the CNCF and Kubernetes communities.

Of 253 respondents:

- just over two fifths (43%) were from Europe
- 30% were from North America
- 17% were from Asia
- the rest (10%) were from Australia and Oceania, South and Central America and Africa.

Just over a fifth (21%) — the single largest group of respondents — represented organizations with 100-499 employees.

- slightly fewer (19%) were from organizations with 10-49 employees
- 13% represented organizations with 50-99
- 10% were from organizations with 500-999 employees, another 10% with 1,000-4,999.
- 18% had more than 5,000 employees
- 8% had a headcount smaller than 10

The most common job function was Site Reliability or DevOps Engineer, specified by 51.38% of respondents.

- 36.36% were software architects
- 25.69% were back-end developers.

47.83% of the respondents worked in the software/technology industry, and 17% — the next largest group — in financial services.