



2017

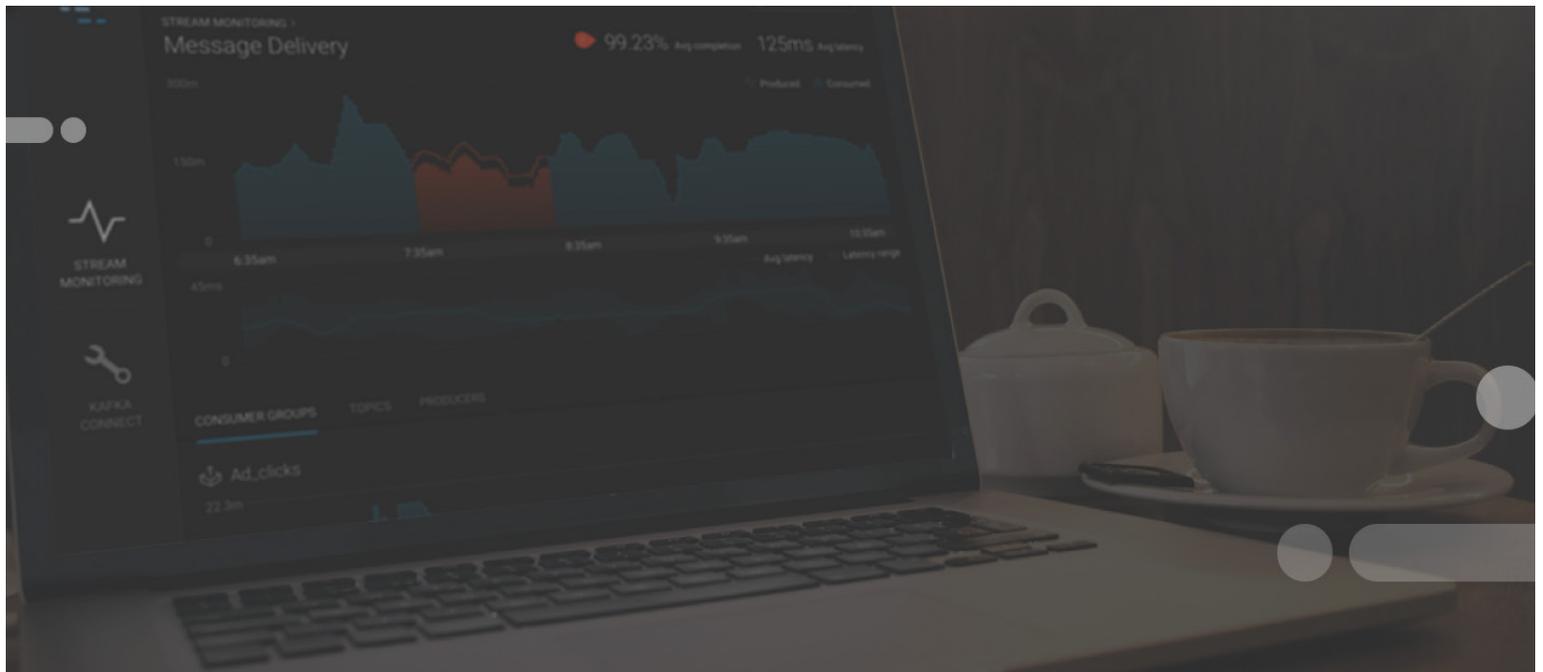
APACHE KAFKA™

REPORT



CONTENTS

Introduction	3
Demographics	4
Executive Summary	5
Key Findings	7
Apache Kafka creates new business opportunities and improves the old ones, too	7
Putting Apache Kafka to use: far beyond pub-sub	8
The impact of Apache Kafka's growing feature list	9
There is a shortage of skilled Apache Kafka engineers	10
Conclusion	11
Who Took the Survey	12
Learn About Confluent	13



INTRODUCTION

Over the past several years, organizations across many industries have discovered, and are filling, an increasingly important gap in their data infrastructure. It sits at the nexus of big data, data integration, and all of their data stores and applications – a gap that is being filled by streaming platforms like Apache Kafka™. Confluent has enjoyed a front row view as companies adopt streaming platforms to create new products, become more responsive to customers and make business decisions in real time.

This survey focuses on why and how companies are using Apache Kafka and streaming data and the impact it has on their business.

DEMOGRAPHICS

CONFLUENT SURVEYED OVER 350 ORGANIZATIONS FROM 47 COUNTRIES AND A WIDE VARIETY OF INDUSTRIES TO UNDERSTAND THE EVOLVING APACHE KAFKA USER BASE, USE CASES AND DEPLOYMENTS.

Those surveyed range in position from developers and architects to technical management and engineers. 1 in 4 respondents (26%) work for organizations with more than \$1 billion in annual sales, illustrating how quickly this open source technology has gained traction across large enterprises.

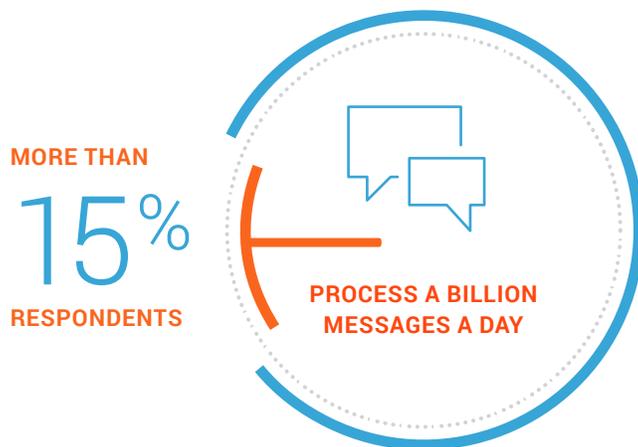
Executive Summary

The problem of managing data in enterprises is an exercise in complexity. As companies add new applications, integrate and modernize existing systems, change to a microservices architecture, and embrace a factor increase in data-producing endpoints, their infrastructure, and the engineers who support it, face an existential question:

Should we pursue each new data project as a unique discrete effort, or find an approach with greater scalability?

For many, Kafka has become the answer to this question. It transforms an ever-increasing number of new data producers and consumers into a simple, unified streaming platform at the center of their organization. It allows any team to join the platform, allows a central team to manage the service, and scales to trillions of messages per day while delivering messages in real time.

Kafka takes on legacy technology across many different areas, including ETL, data warehouses, Hadoop, messaging middleware and data integration technologies, to massively simplify an organization's infrastructure. In many cases, Kafka can replace or augment an existing system to make data more consistently available, faster and less costly to deliver.



SOME PARTICULARLY REVEALING RESULTS FROM THIS YEAR'S SURVEY INCLUDE:

Kafka use is experiencing a surge:

86% of respondents reported that the number of their systems that use Kafka is increasing and a fifth (20%) reported that the number is "growing a lot!" A majority (52%) of organizations have at least 6 systems running Kafka with over a fifth (21%) having more than 20. According to last year's report, only 41% of organizations had at least 6 systems running Kafka and only one-tenth (10%) had more than 20.



86% of respondents reported that the number of their systems that use Kafka is increasing

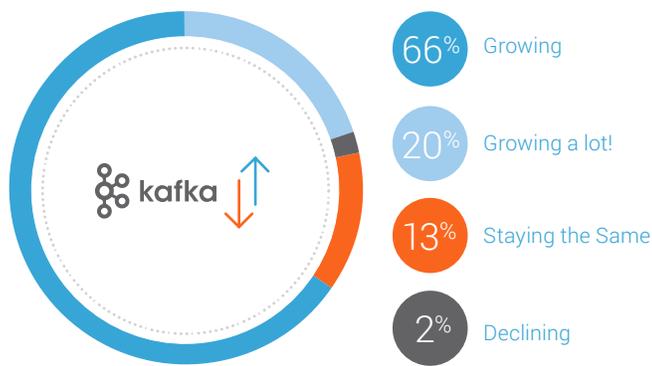


20% reported that the number is "growing a lot!"



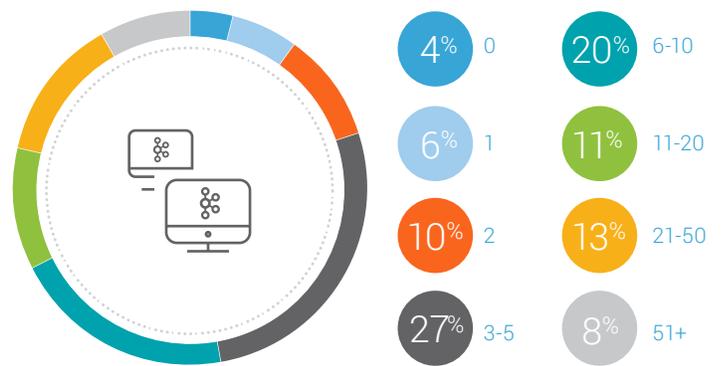
52% of organizations have at least 6 systems running Kafka.

IS THE NUMBER OF SYSTEMS USING KAFKA GROWING, DECLINING OR STAYING THE SAME?



Note: Only a single option could be selected

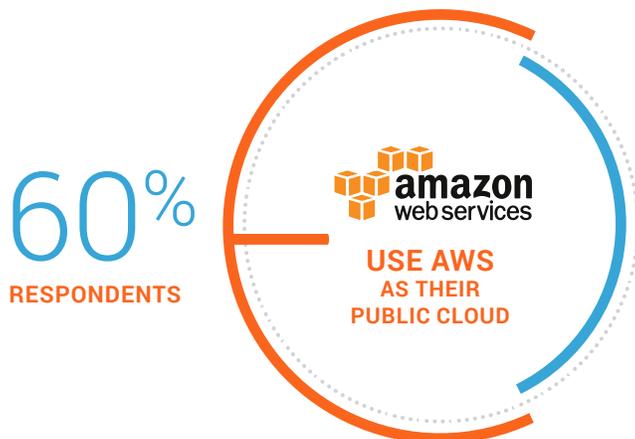
HOW MANY SYSTEMS OR APPLICATIONS USE KAFKA TODAY?



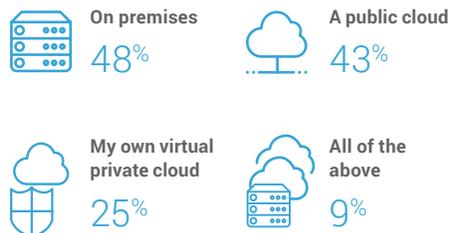
Note: Only a single option could be selected

Kafka is broadly used in the cloud:

Kafka is used by organizations in some combination of virtual private clouds (34%), public clouds (52%), and on premises (57%). Nearly one-third (32%) of respondents who use Kafka in the cloud have at least 6 Kafka applications in the cloud.



IN WHICH OF THESE ENVIRONMENTS ARE YOU USING KAFKA TODAY?



Note: Respondents could select multiple options

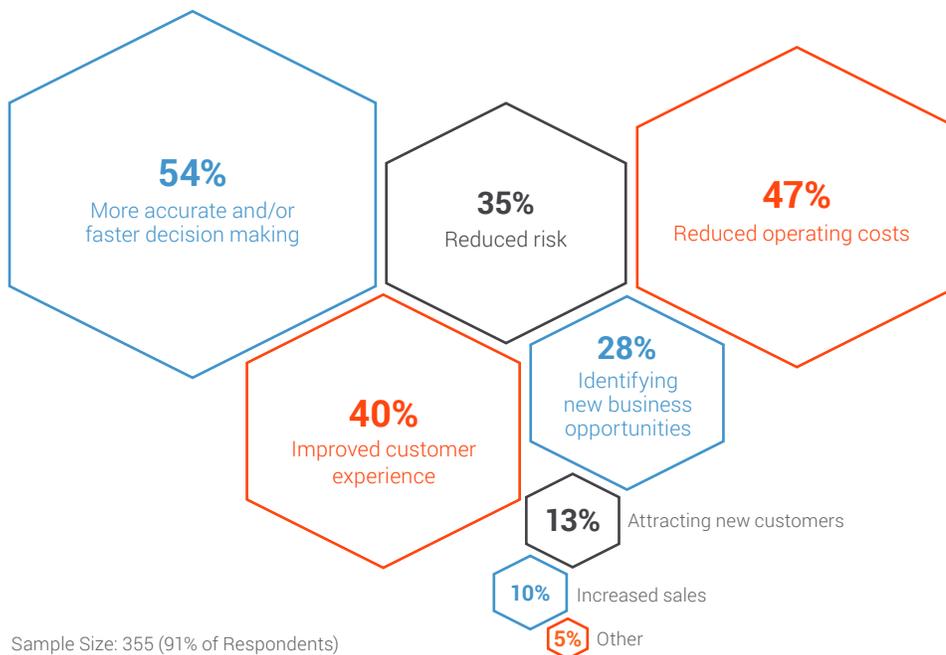
KEY FINDINGS

KAFKA CREATES NEW BUSINESS OPPORTUNITIES AND IMPROVES THE OLD ONES, TOO

While we've seen organic growth of Apache Kafka over the years and the downloads to prove it, this survey quantifies what we see from customers and adopters: streaming data enables organizations to do what they couldn't before:

- Kafka enables companies to create new market opportunities. Because data is available, shared and immediate, companies can create new products and significantly transform existing ones. As Kafka is deployed in more mission-critical infrastructures, a majority (54%) of surveyed organizations say that their business can make more accurate and/or quicker decisions thanks to Kafka.
- In addition to creating new opportunities, companies leverage Kafka to be more efficient and transform existing processes. Of those surveyed, other business benefits worth noting are reduced operating costs (47%) and improved customer experience (40%).

**WHAT VALUE DOES YOUR BUSINESS GET FROM USING KAFKA AND STREAMING PLATFORMS?
(SELECT ALL THAT APPLY.)**



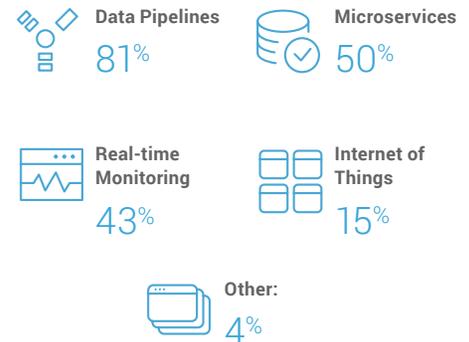
PUTTING APACHE KAFKA TO USE: FAR BEYOND PUB-SUB

Kafka can be used in a variety of different ways for many different use cases. Last year we found a surge of companies adopting streaming platforms. With this approach, companies build mission critical, real-time applications that power their core business – all the way from small to large-scale use cases that handle millions of events per second.

Now, companies are solving a new problem with Kafka: microservices. While microservices involve many independent services, the goal is broader than simply running them across different machines. It's about facing up to a world that is, itself, inherently distributed. Not in some narrow technical sense, but rather as a broad ecosystem composed from many people, many teams and many programs, all of which need the agility that microservices affords them.

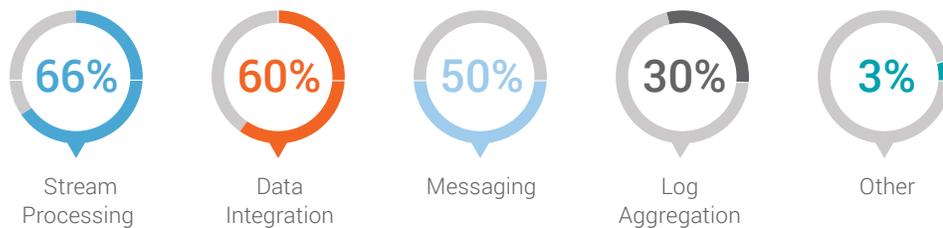
This year, we're seeing organizations use Apache Kafka in many ways: two-thirds (66%) use it for stream processing and three out of five (60%) use it for data integration. The most common use case for Kafka is data pipelines (81%), while half (50%) are already using it for microservices.

WHAT ARE YOUR CURRENT USE CASES FOR KAFKA?



Note: Respondents could select multiple options

WHAT WAYS ARE YOU USING APACHE KAFKA TODAY? (SELECT ALL THAT APPLY.)

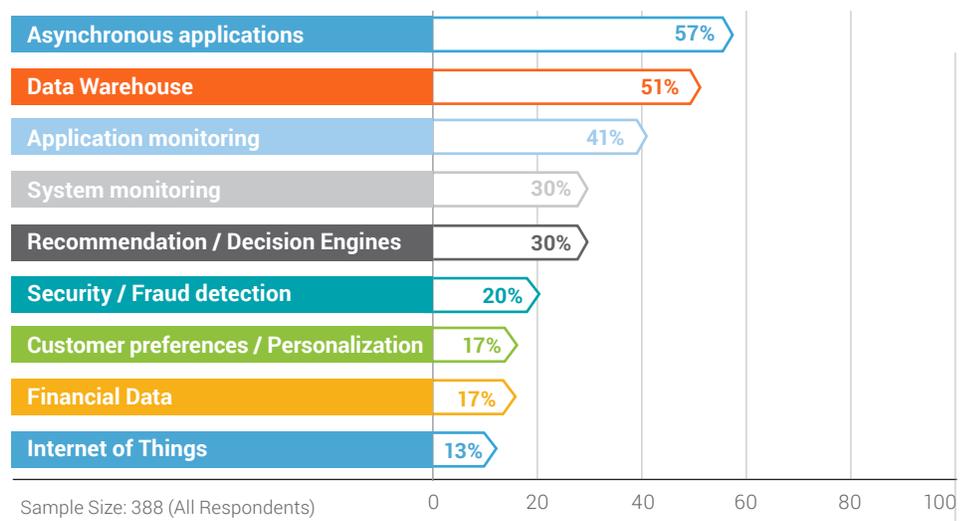


Sample Size: 168 (43% of Respondents)

Similar to last year's report, three-fourths (75%) of organizations have applications connected to their Kafka systems. These applications process data from sensors, websites, analytics and monitoring tools, etc. and share the information so the right teams can process the data they need to make decisions.

The types of applications that a majority of respondents connect to with Kafka are asynchronous applications (57%) and data warehouse (51%). Organizations also have application monitoring (41%), system monitoring (30%), or recommendation/decision engines (30%) powered by Kafka.

WHAT TYPE OF APPLICATIONS ARE POWERED BY KAFKA WITHIN YOUR ORGANIZATION? (SELECT ALL THAT APPLY.)



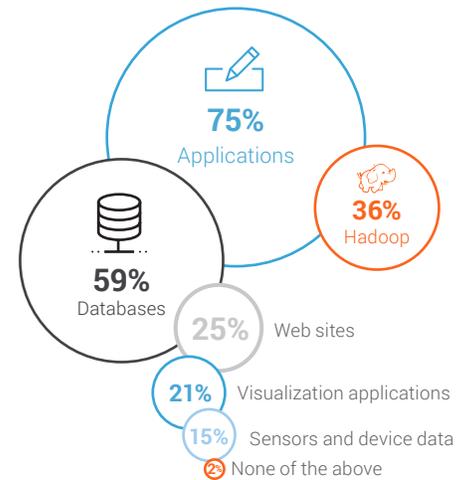
THE IMPACT OF KAFKA'S GROWING FEATURE LIST

The popularity of the Kafka Connect API has grown significantly over the past year. The Kafka Connect API, included in Kafka, makes it easy to add new datastores to your data pipelines without having to write the interfaces from scratch.

There was a 15-point increase in organizations using the Kafka Connect API over last year (37% in 2017 vs 12% in 2016). While a majority (59%) of respondents have databases connected to their Kafka clusters, only 36% use the Kafka Connect API with Hadoop/HDFS, which is a 4-point drop from last year.

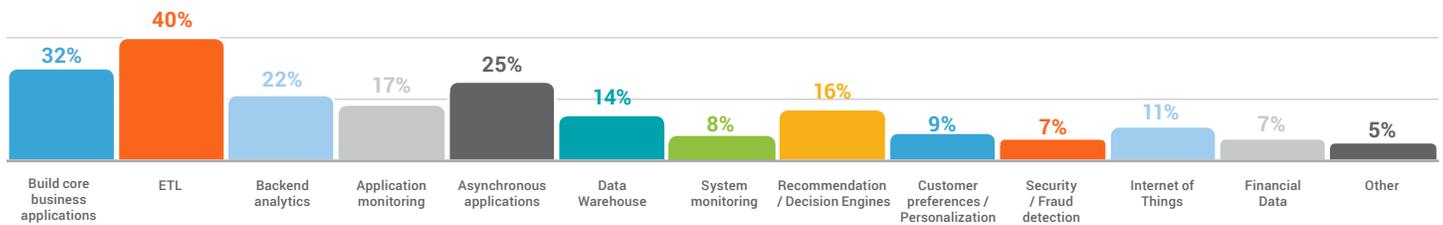
Introduced last year, the Kafka Streams API integrates into your streaming application to build and execute powerful stream processing functions, without needing to set up and run separate infrastructure. Despite its recency, most (89%) organizations are already familiar with the Kafka Streams API. The most common uses for the Kafka Streams API are ETL (40%), building core business applications (32%), and asynchronous applications (25%).

WHAT SOURCES OR SINKS ARE CONNECTED TO YOUR KAFKA CLUSTERS? (SELECT ALL THAT APPLY.)



Note: Respondents could select multiple options

WHAT ARE YOU USING THE KAFKA STREAMS API FOR? (SELECT ALL THAT APPLY.)



Note: Respondents could select multiple options

As we mentioned earlier, microservices is the new way to build custom applications. This is made much easier with the Kafka Streams API as developers have the ability to work with data in flight. Of the organizations that have microservices, 28% use the Kafka Streams API to manage them. The Kafka Streams API is most frequently utilized by developers (85%), but architects (48%) and application teams (43%) also use it.

THERE IS A SHORTAGE OF SKILLED KAFKA ENGINEERS

According to a Dice report, people with Kafka skills receive one of the highest salaries in the technology market. However, despite the salary and the growth of Kafka within organizations, three-quarters (75%) of respondents find it difficult to find the right talent with Kafka skills.



HAVE YOU FOUND IT DIFFICULT TO FIND THE RIGHT TALENT WITH KAFKA SKILLS?



Note: Only a single option could be selected

RESPONDENTS ARE VERY SATISFIED WITH KAFKA

An factor in the growing popularity of Kafka is that it meets the expectations of its users. Overall satisfaction with Kafka remains high with 70% of organizations saying that they are either very or completely satisfied with it.

HOW SATISFIED ARE YOU WITH KAFKA?



Note: Only a single option could be selected

CONCLUSIONS

The results of this survey demonstrate the increasing use of Apache Kafka for a broadening set of business and technical objectives. Many companies are implementing the distributed streaming platform for more accurate and faster decision making, reduced operating costs, improved customer experiences and reduced risk.

With demands from customers to react and respond in real time and legacy technologies holding companies back, it's no wonder companies have turned to Apache Kafka to embrace the power of streaming data for competitive advantage.

APACHE KAFKA IS A DOMINANT STREAMING DATA SOLUTION FOR ALL ENVIRONMENTS

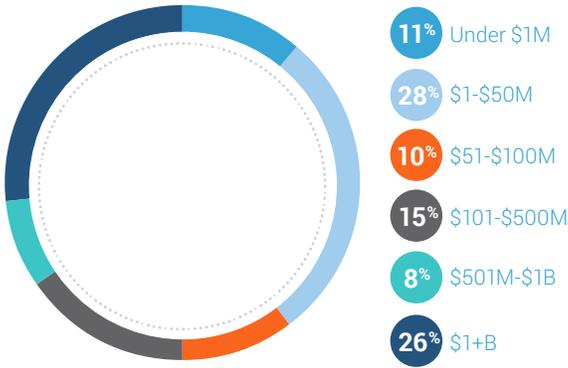
With the rise of microservices and the cloud, Kafka is best-suited to meet the needs of these distributed environments. This survey demonstrates companies are not only actively deploying these architectures, but they are looking and implementing technologies that help them scale.

KAFKA SKILLS ARE IN DEMAND

Companies are increasing their use of Kafka and are looking for people with the right skillset. With this demand, it's hard to imagine a world without streaming platforms at the heart of every company in the very near future.

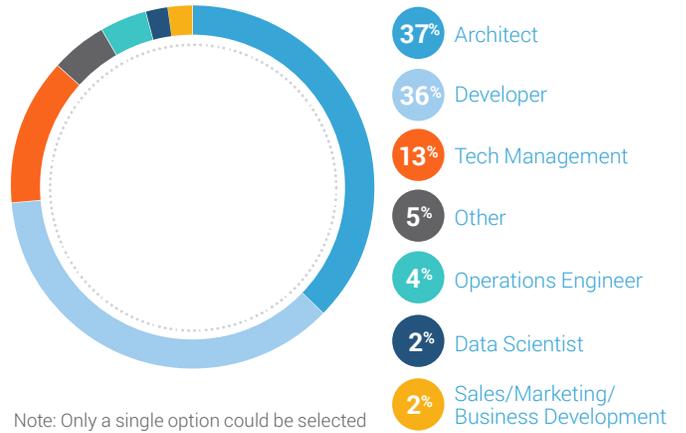
WHO TOOK THE SURVEY?

COMPANY SIZE



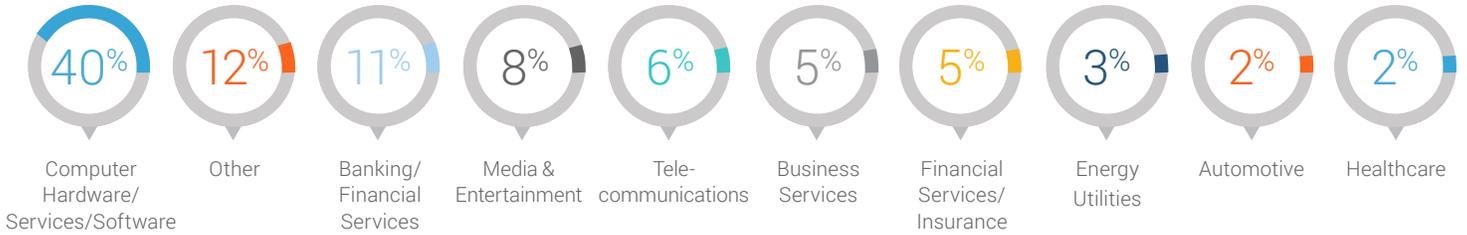
Note: Only a single option could be selected

RESPONDENT ROLE



Note: Only a single option could be selected

INDUSTRIES



LOCATIONS





ABOUT CONFLUENT

Confluent, founded by the creators of open source Apache Kafka™, provides the only streaming platform that enables enterprises to maximize the value of data. Confluent Enterprise empowers leaders in industries such as retail, logistics, manufacturing, financial services, technology and media, move data from isolated systems into a real-time data pipeline where they can act on it immediately. Backed by Benchmark, Index Ventures, and Sequoia, Confluent is based in Palo Alto, California.

To learn more about how Confluent Platform and Apache Kafka can help your organization, visit: www.confluent.io