

The Need For Speed: Fast Data Development Trends

Insights from over 2,400 developers on the impact of “Data in Motion” in the real world

About This Report

The digitization of the world has fueled unprecedented [growth](#) in data, creating huge implications for how enterprises interact with data to create future business opportunities.

Fast data is a new opportunity made possible by emerging technologies and, in many cases, by new approaches to established technologies. Like its big data counterpart, fast data is surrounded by hype and confusion.

To better understand the current state of fast data, Lightbend surveyed 2,457 global developers to get their real-world take on:

- Alignment between fast data and business value
- Impacts on software development and tool choices
- Patterns and challenges facing early adopters

The survey was conducted in June 2017 and represents a wide range of industries and company sizes. It is weighted toward the perspective of developers and architects.

Respondents by Role

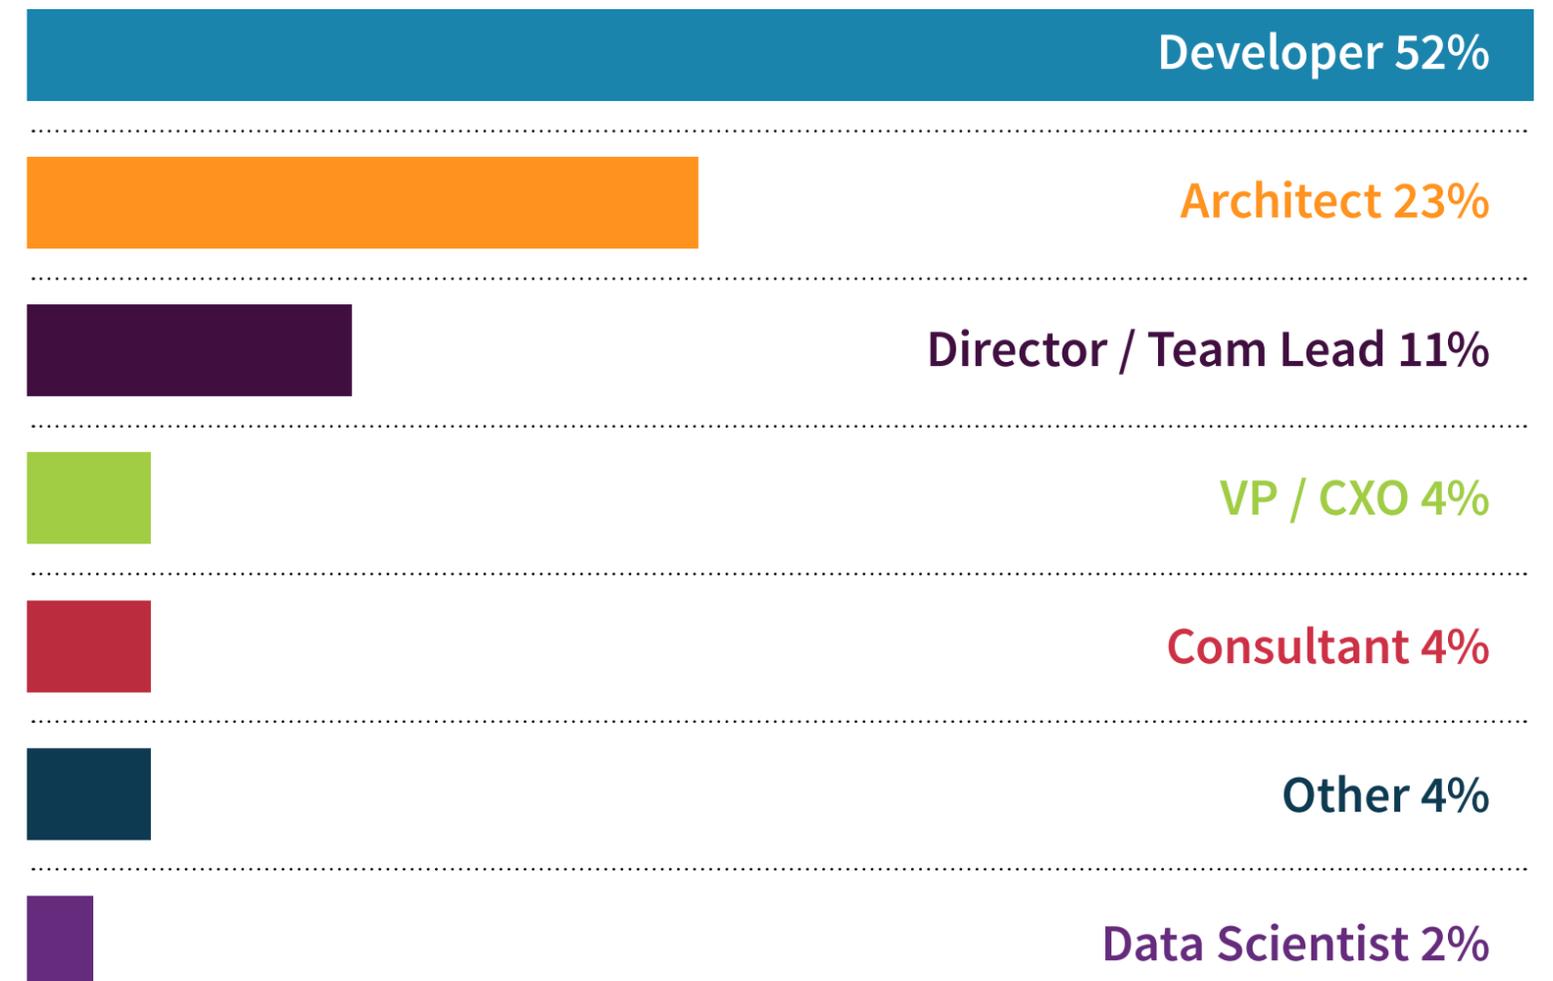


Table of Contents

Executive Summary.....	4
Fast Data Value Is Clear: Senior Management Buys In	5
Batch vs Streaming: Where Speed Really Matters.....	8
Technology Shifts: Fast Data Shakes Up Traditional Stack	13
Last Look: Steps for Success	18

Executive Summary

The big data market is undergoing a rapid transformation from data at rest to data in motion. Analysts indicate the adoption of fast data is happening at a rate three times faster than traditional Hadoop.

Why is the fast data market moving so fast? In this report, we leverage real-world insights from more than 2,400 developers to examine adoption trends across three key areas — business value, use case, and technology.

Finding 01

Fast Data Value Is Clear: Senior Management Buys In

(see page 8)

Unlike its big data counterpart, fast data appears to be more intuitive from a business value perspective. **Sixty percent of senior management can effectively link strategic value to projects when data is in motion.** Management in some industries, however, is getting it faster than others.

Finding 02

Batch vs Streaming: Where Speed Really Matters

(see page 8)

The move to real-time data is accelerating. **Developers say ninety percent of their data processing workloads include a real-time component.** The need for speed increases as use cases climb the maturity curve. Rather than batch *versus* streaming, enterprises will need batch *and* streaming to succeed with fast data.

Finding 03

Technology Shifts: Fast Data Shakes Up Traditional Stack

(see page 13)

Developers are in the driver's seat with regards to tech selection. **Fifty-five percent say they are choosing new frameworks and languages based on fast data requirements.** But where the new ecosystem of streaming engines is concerned, developers and architects say they need guidance to choose the right tools.

Finding
01

Fast Data Value Is Clear

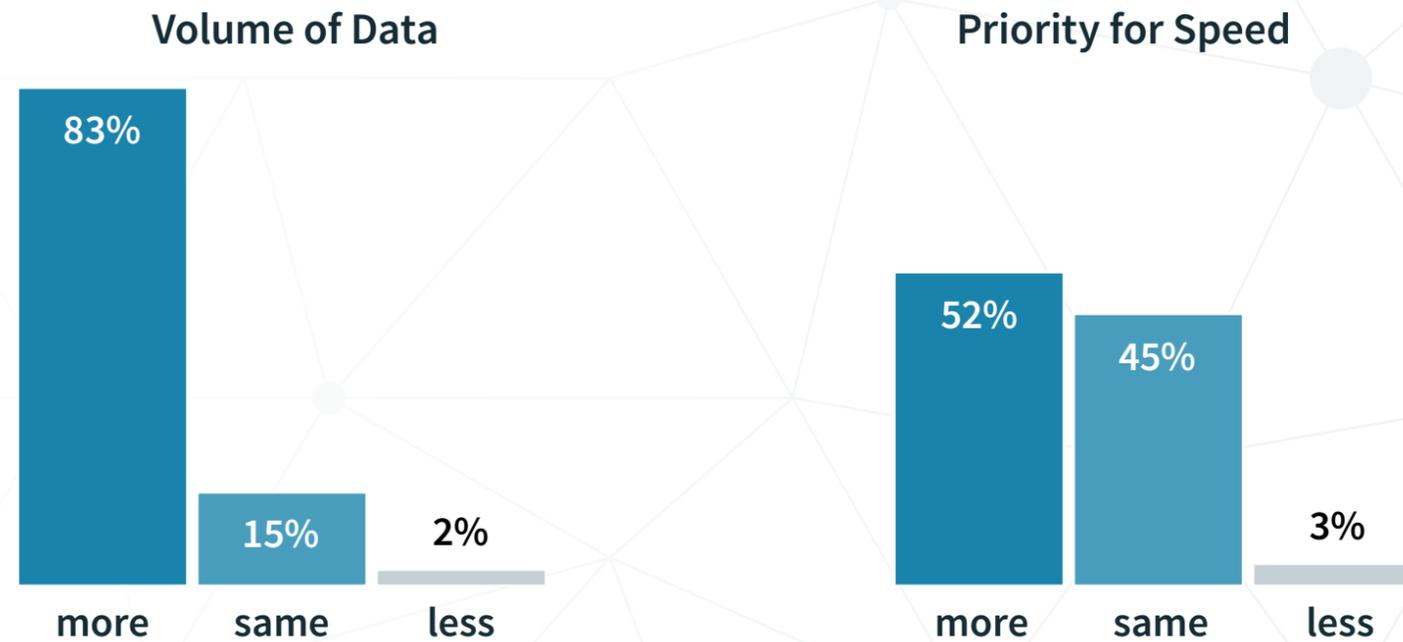
Senior Management Buys In

To compete in the digital era, the necessity is rising for enterprises to use data faster. This need for speed is expanding beyond analytics to applications that adapt to changing conditions, personalize customer engagement, and power the internet of everything.

Consequently, an overwhelming majority of developers reported they are obligated to use data faster today than two years ago.

Urgency to use more data faster is on the rise

*We asked developers to contrast the **volume of data** and **priority for speed** today as to compared with two years ago. Not surprisingly, both are on the rise.*



Market Perspective

“The future of large size data streaming and innovation is more critical than any other innovation for the next decade.”

Dr. Hossein Eslambolchi, Technical Advisor at Facebook

Finding
01

Fast Data Value Is Clear

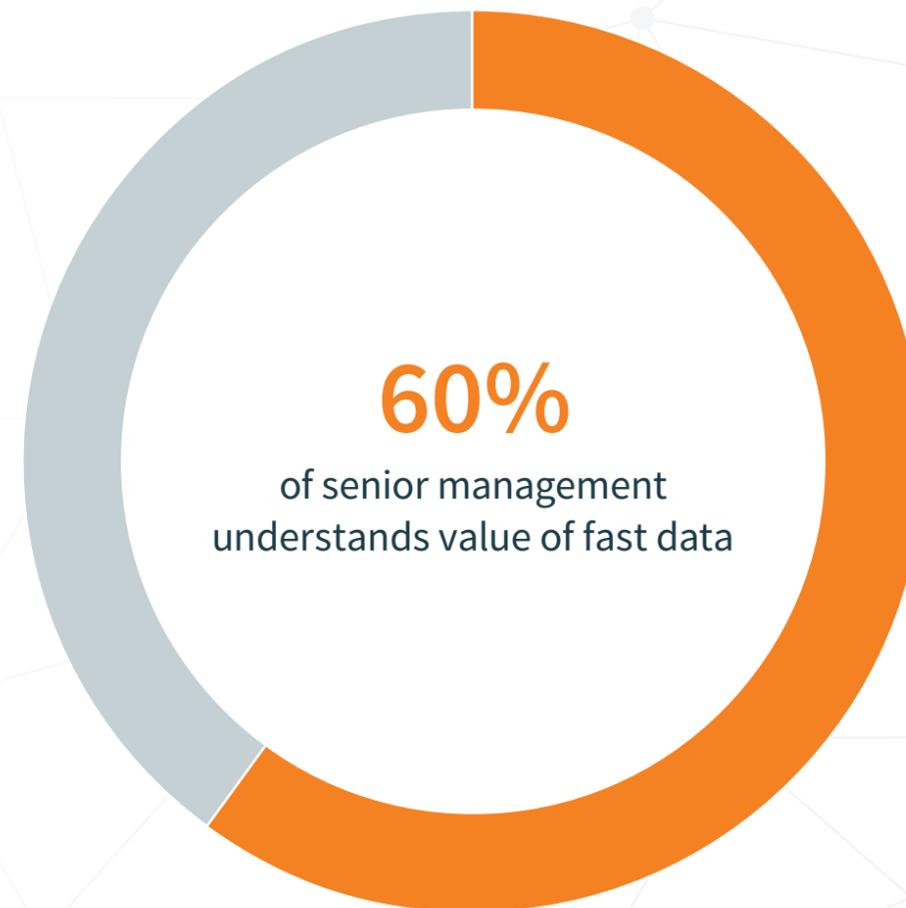
Senior Management Buys In

Unlike its big data counterpart, fast data appears to be more intuitive to senior management from a business value perspective. For years, industry analysts have been [reporting](#) high failure rates for big data projects. Lack of a clearly defined business case and skepticism of internal stakeholders are among the most commonly cited derailleurs of big data initiatives.

By contrast, survey results suggest that senior management can effectively link strategic value to projects when data is in motion.

Value of fast data is well understood

Sixty percent of developers surveyed do not have a challenge getting senior management to understand the value of their fast data projects.



Finding
01

Fast Data Value Is Clear

Senior Management Buys In

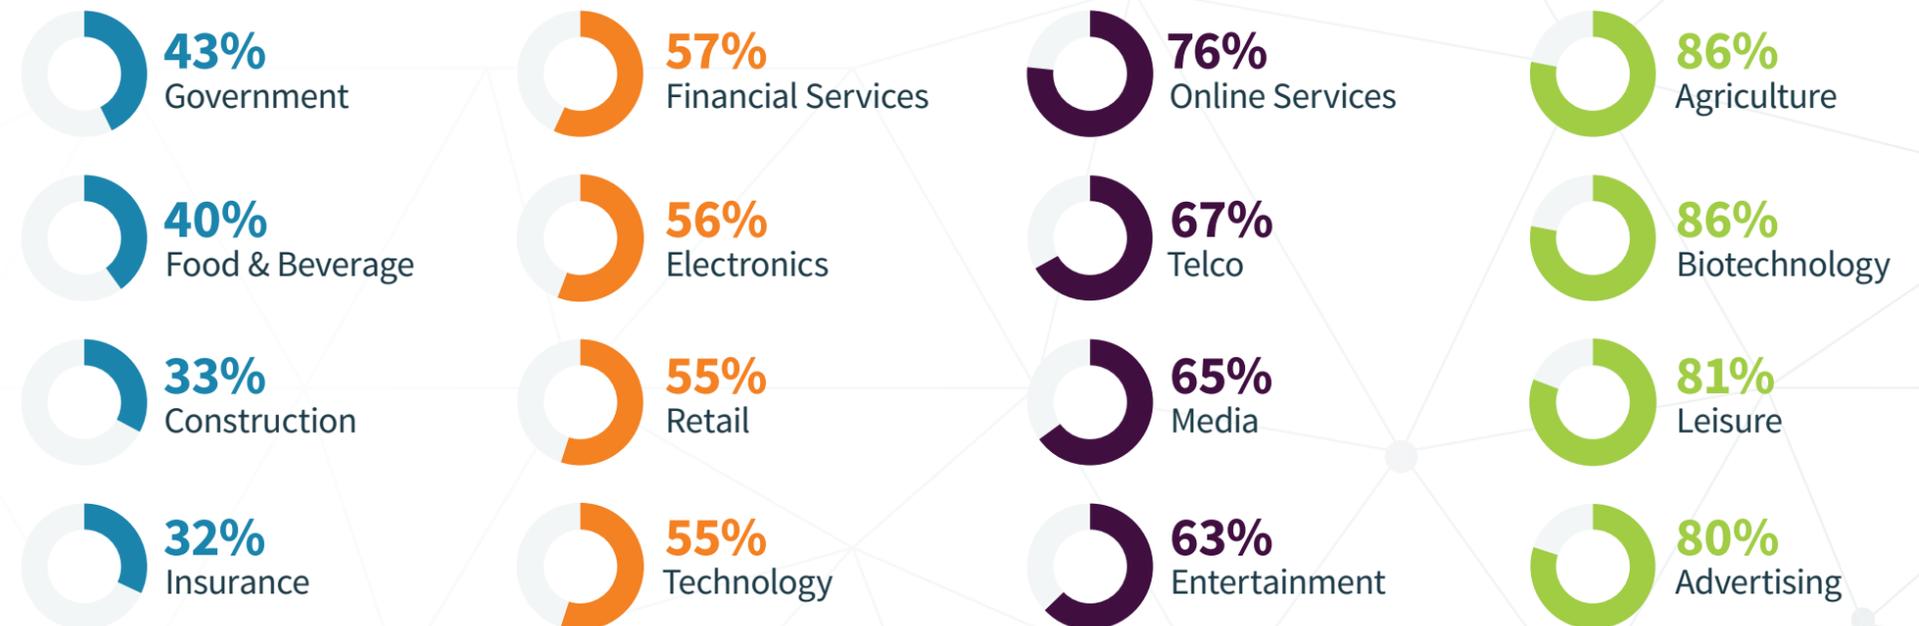
While the majority of developers say their senior management understands the value of fast data, some industries appear to outpace the pack.

Management is considered a laggard by developers in the insurance industry.

Developers in financial services and retail rank senior management understanding as average, while agriculture and biotechnology management lead the way.

Some industries are getting it faster than others

Does senior management understand the value of fast data? Below is a sample of management buy-in by industry.



Market Perspective

Why does senior management in agriculture get fast data? “Smart agriculture is already becoming more commonplace among farmers, and high tech farming is quickly becoming the standard thanks to agricultural drones and sensors.”

Finding
02

Batch vs Streaming

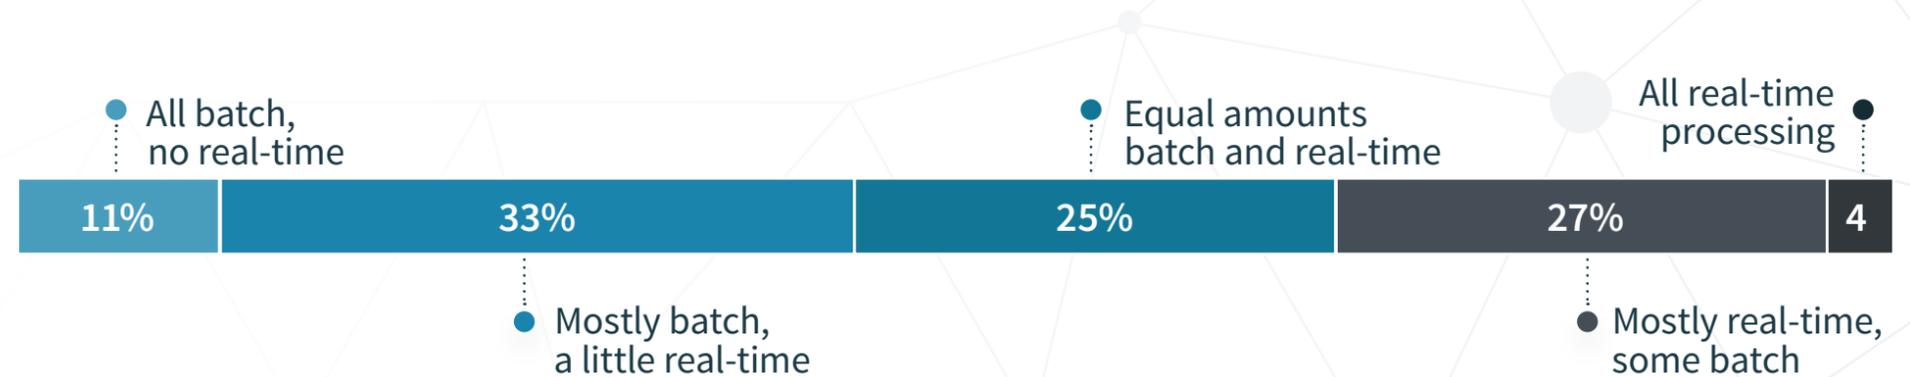
Where Speed Really Matters

When the Internet’s pioneers were struggling to gain control of their ballooning data sets, the “classic” Apache Hadoop architecture solved the primary use case of batch-mode analytics and data warehousing.

While our survey suggests Hadoop may not be relevant for fast data use cases, batch continues to play a role. Of particular note, however, is the role of real-time data—ninety percent of respondent workloads include a real-time component.

Enterprises begin to embrace streaming

Developers say ninety percent of their data processing workloads include a real-time component. Here we see the progression breakdown from batch to real-time.



83% of fast data systems today are **not** running on Apache Hadoop

90% of workloads include a real-time component

Finding
02

Batch vs Streaming

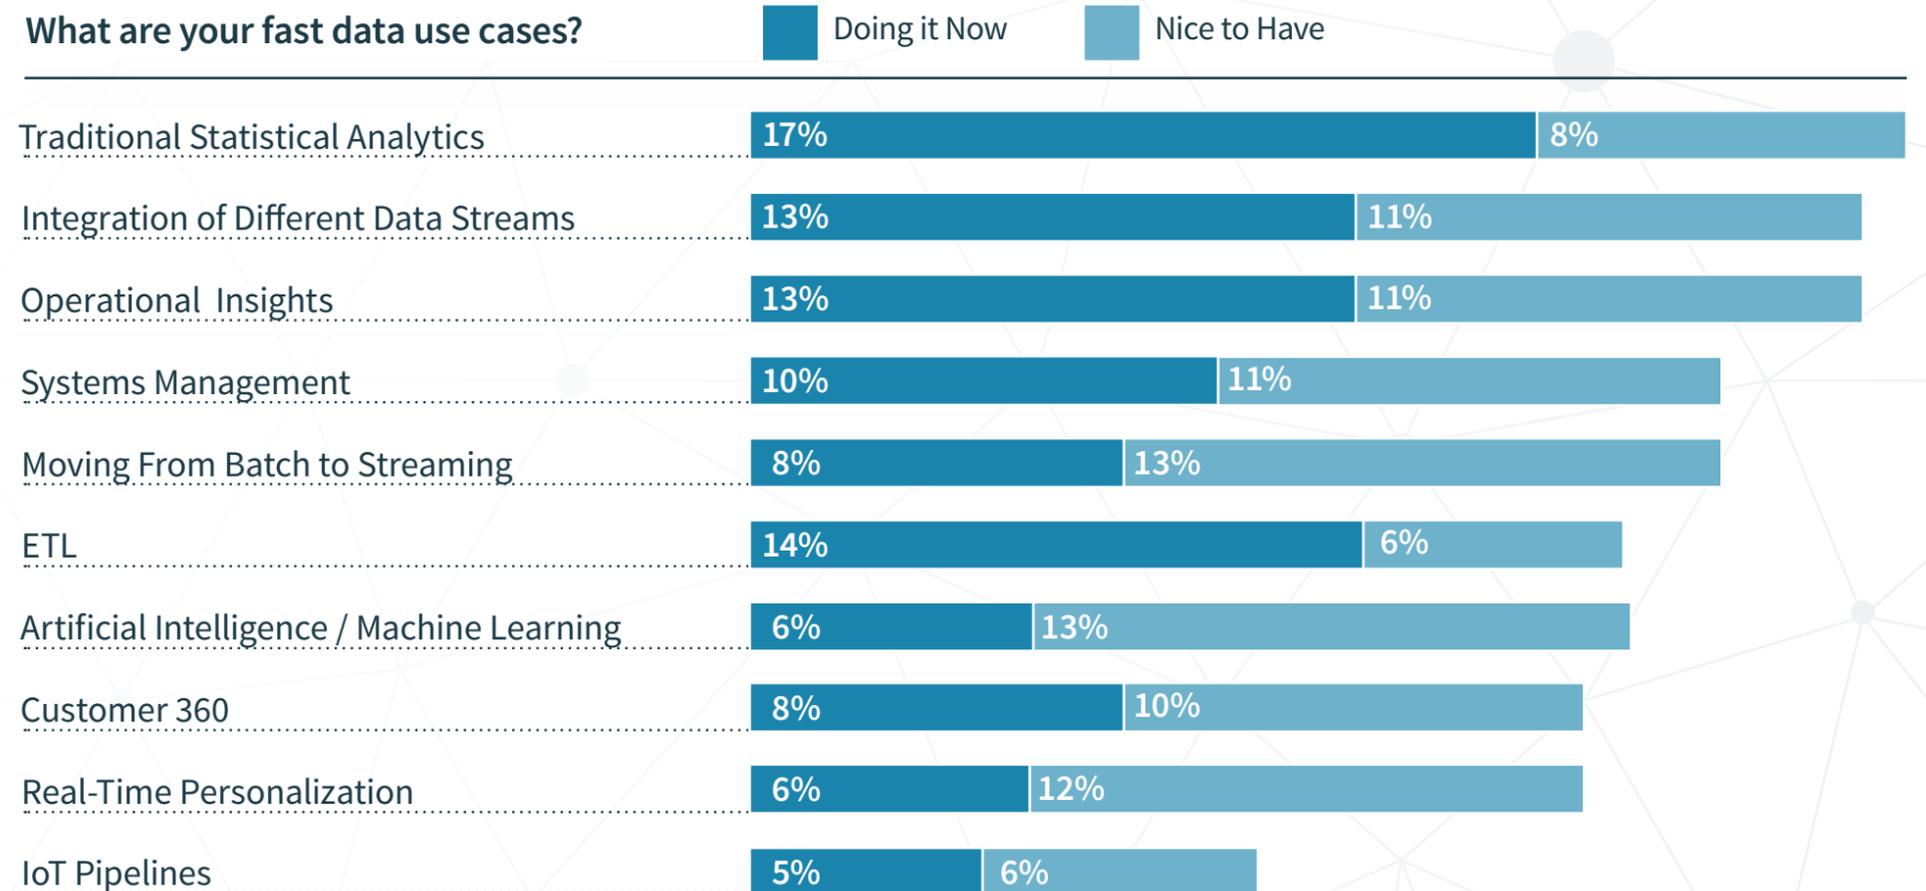
Where Speed Really Matters

Although much more difficult to build than batch, fast data architectures represent the state of the art for powering use cases that are propelling business innovation and competitive advantage.

In this segment of the report, we drill down into fast data uses cases in production to determine the impact on the need for speed.

Fast Data use cases span the maturity curve

From traditional analytics and ETL to advanced machine learning and IoT pipelines, we see the breakdown of fast data uses cases in production and on the horizon.



Finding
02

Batch vs Streaming

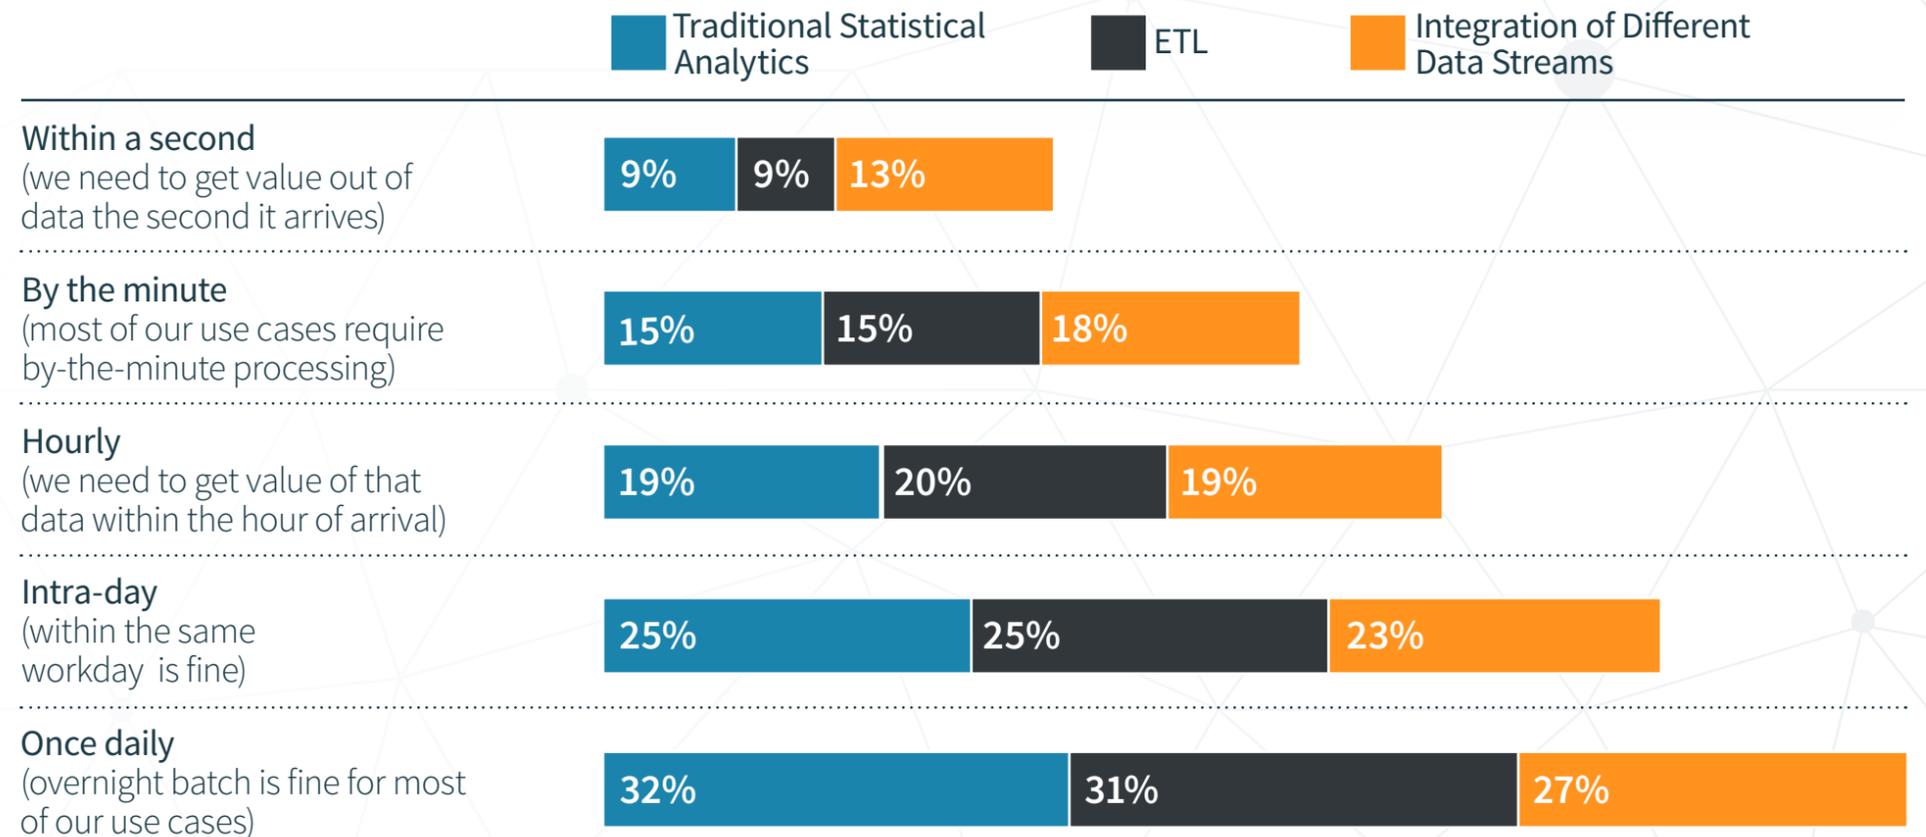
Where Speed Really Matters

Rather than jumping directly into artificial intelligence and machine learning, most enterprises start their fast data efforts by addressing business situations where value does not need to be analyzed immediately.

These use cases are aimed at ingesting the data as it arrives; sometimes applying ETL or data integration techniques in real time; storing the data in a data lake or other data store; and conducting the analytics on the data at rest in a much more compressed time frame: daily, intra-day or hourly.

Faster analysis and ETL are intuitive starting points

Developers cite traditional statistical analysis, ETL, and integration of data streams as top fast data uses cases in production, which have been correlated to the speed progression breakdown below.



Finding
02

Batch vs Streaming

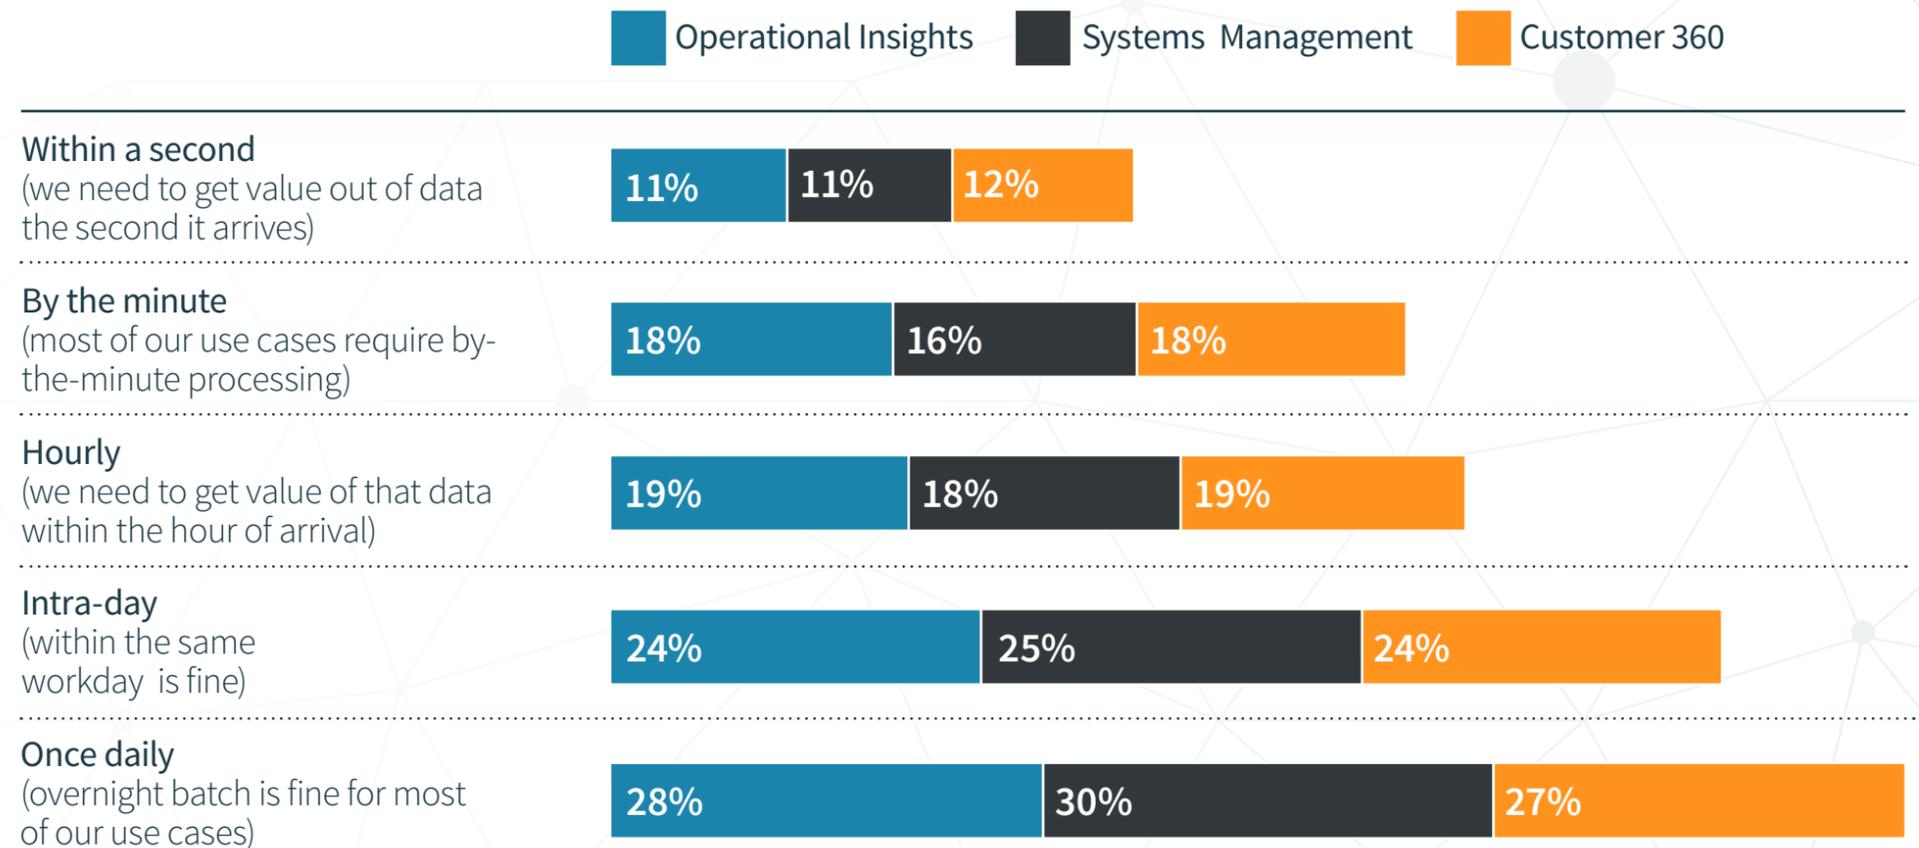
Where Speed Really Matters

Moving up the maturity curve, respondents report use cases that benefit from situational awareness for operations, systems, and customers.

Regardless of industry or environment, situational awareness means having an understanding of what you need to know, what you have control of, and conducting analysis in near real-time to identify anomalies in normal patterns or behaviors that can affect the outcome of a business or process. If you have these things, making the right decision within the right amount of time in any context becomes much easier.

Situational awareness use cases follow

Business functions leading the way for use cases in production include operations, systems, and customers. Here we see the speed progression breakdown.



Finding
02

Batch vs Streaming

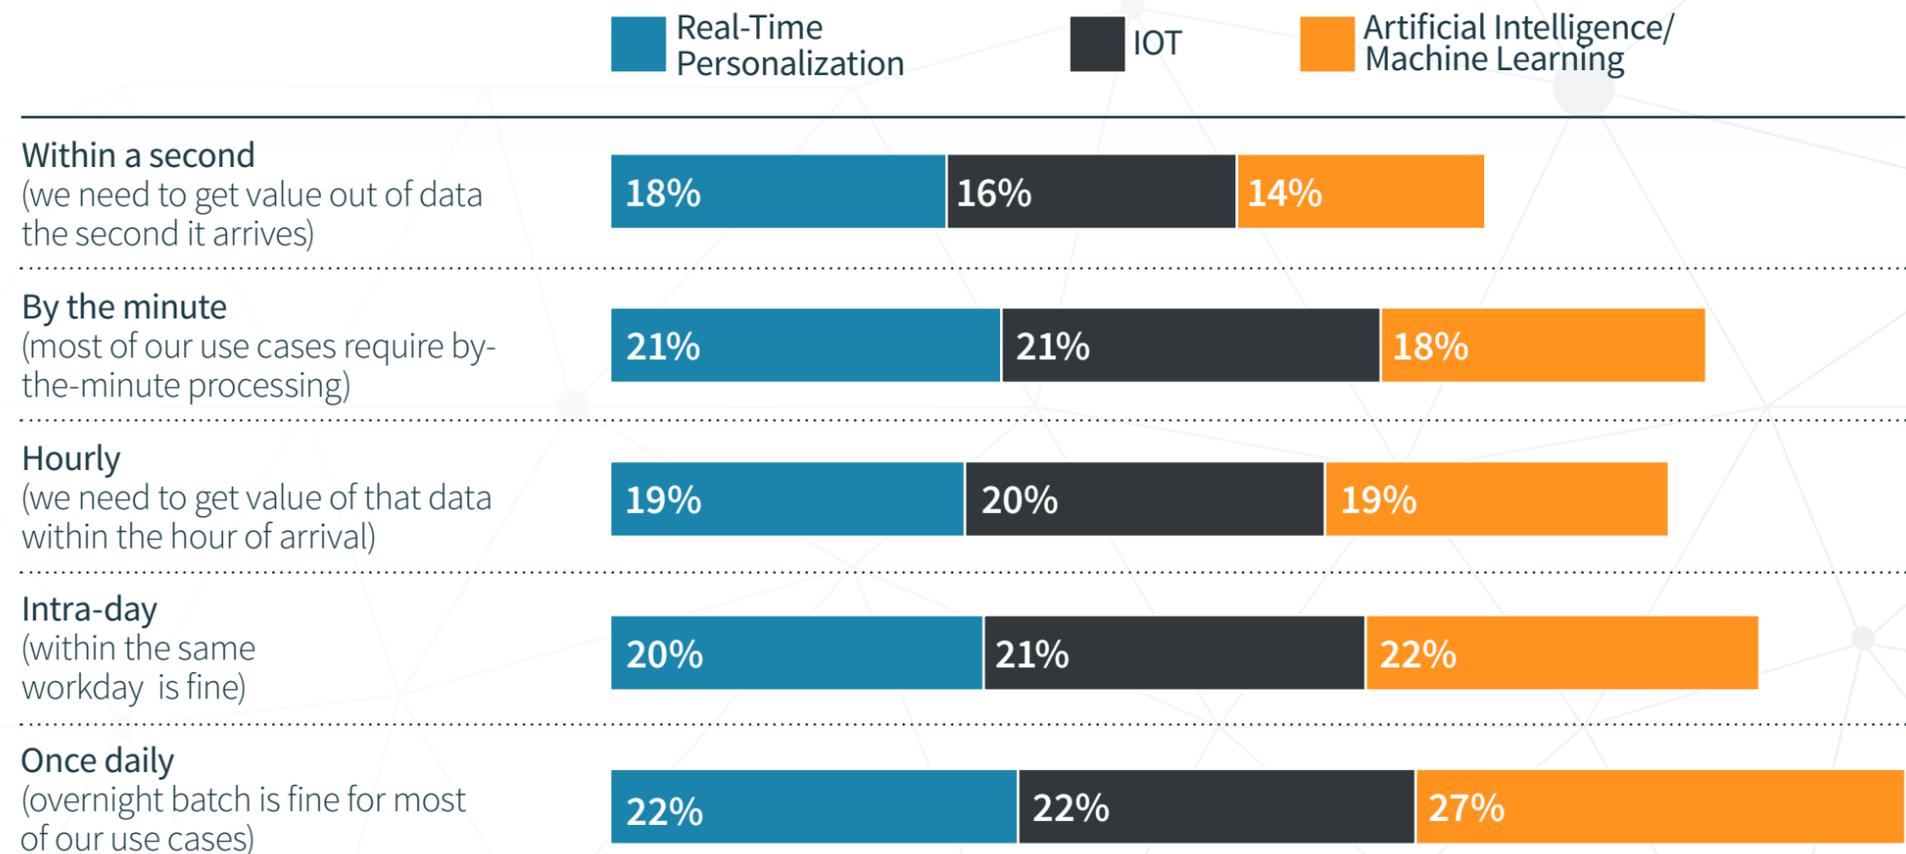
Where Speed Really Matters

Advanced streaming use cases in production are beginning to leverage machine learning to adapt to changing market and environmental conditions. Model updates are performed in predictable batch processes or delivered through continuous intra-day updates. To properly train models, an enterprise needs “an unearthly amount of data” as Neil Lawrence, a member of Amazon’s AI team and professor of machine learning at the University of Sheffield, [puts it](#).

Rather than batch *versus* streaming, enterprises will need batch *and* streaming to succeed with advanced fast data use cases.

Personalization, IoT, and ML are nascent

Not surprisingly, advanced use cases are just beginning to gain a foothold in the enterprise. Here we see the need for speed increase.



Finding
03

Technology Shifts

Fast Data Shakes Up Traditional Stack

With the rising demand to use more data, faster, developers and architects are beginning to favor new frameworks and languages that handle data more effectively than traditional tools.

Traditional systems of record, however, are not being disregarded by developers. Thirty percent are modernizing aging systems to take advantage of new data requirements.

Data requirements are influencing tech selection

Developers are in the driver's seat with regards to tech selection. Fifty-five percent say they are choosing new frameworks and languages based on fast data requirements.

3%

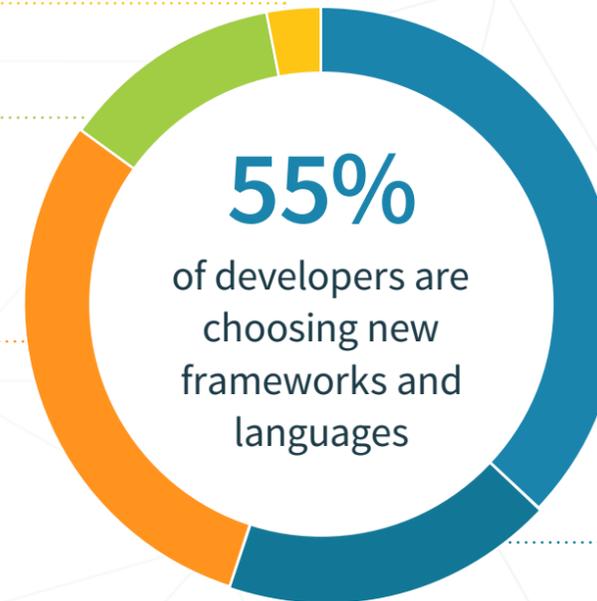
Other

12%

We're prioritizing new engineering hiring based on prior data science or data engineering experience

30%

We're modernizing old systems specifically to be more compatible with our new data requirements



37%
We're choosing new frameworks based on their ability to handle data more effectively

18%
We're choosing new languages based on their ability to handle data more effectively

Market Perspective

“Those responsible for modernizing app infrastructure, [Gartner advises], should ‘retain Java EE servers for existing legacy applications, but use lighter-weight Java frameworks for digital business application development projects or evaluate other language platforms.’”

Finding
03

Technology Shifts

Fast Data Shakes Up Traditional Stack

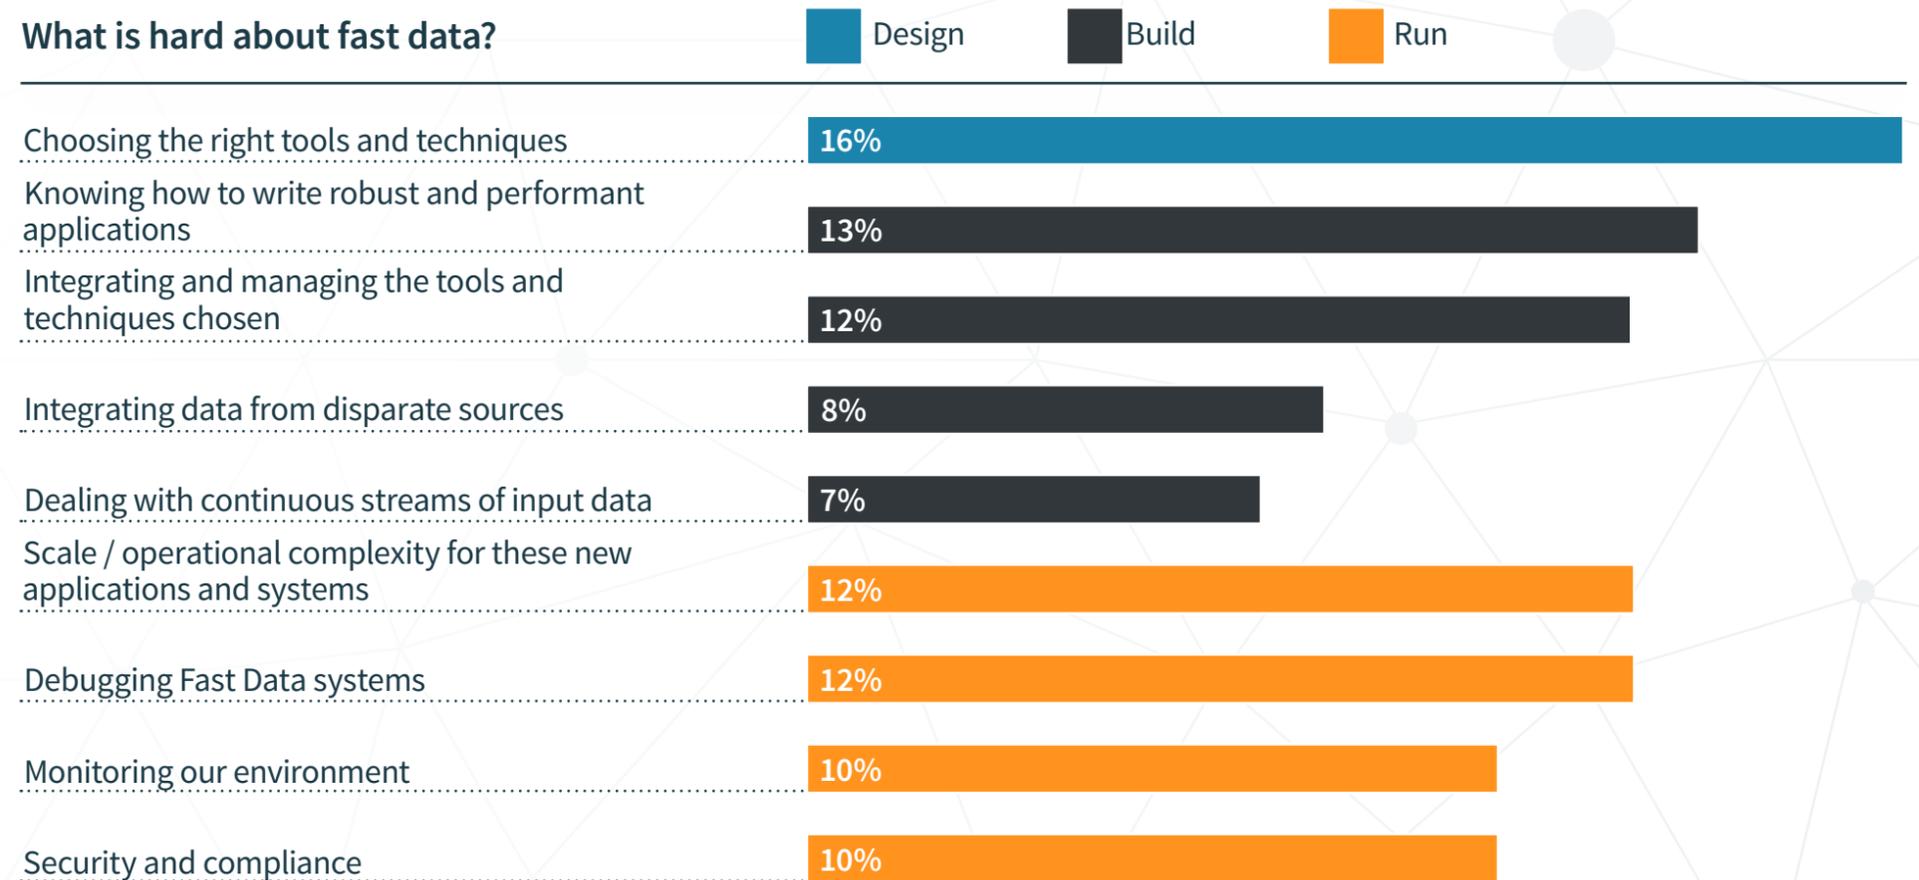
In adopting fast data, respondents appear to be more confident working with disparate data sources and continuous streams of input than operationalizing their systems in production. Integrating, scaling, debugging, and monitoring are posing challenges for developers.

The biggest hurdle, however, surfaces earlier in the software development lifecycle: choosing the right tools and techniques.

Choosing the right tools ranks as top challenge

We asked developers what's hard about fast data. The responses were roughly split between the build and run phases of a project. The design phase, however, appears most problematic with choosing the right tools ranking as the top challenge.

What is hard about fast data?



Finding
03

Technology Shifts

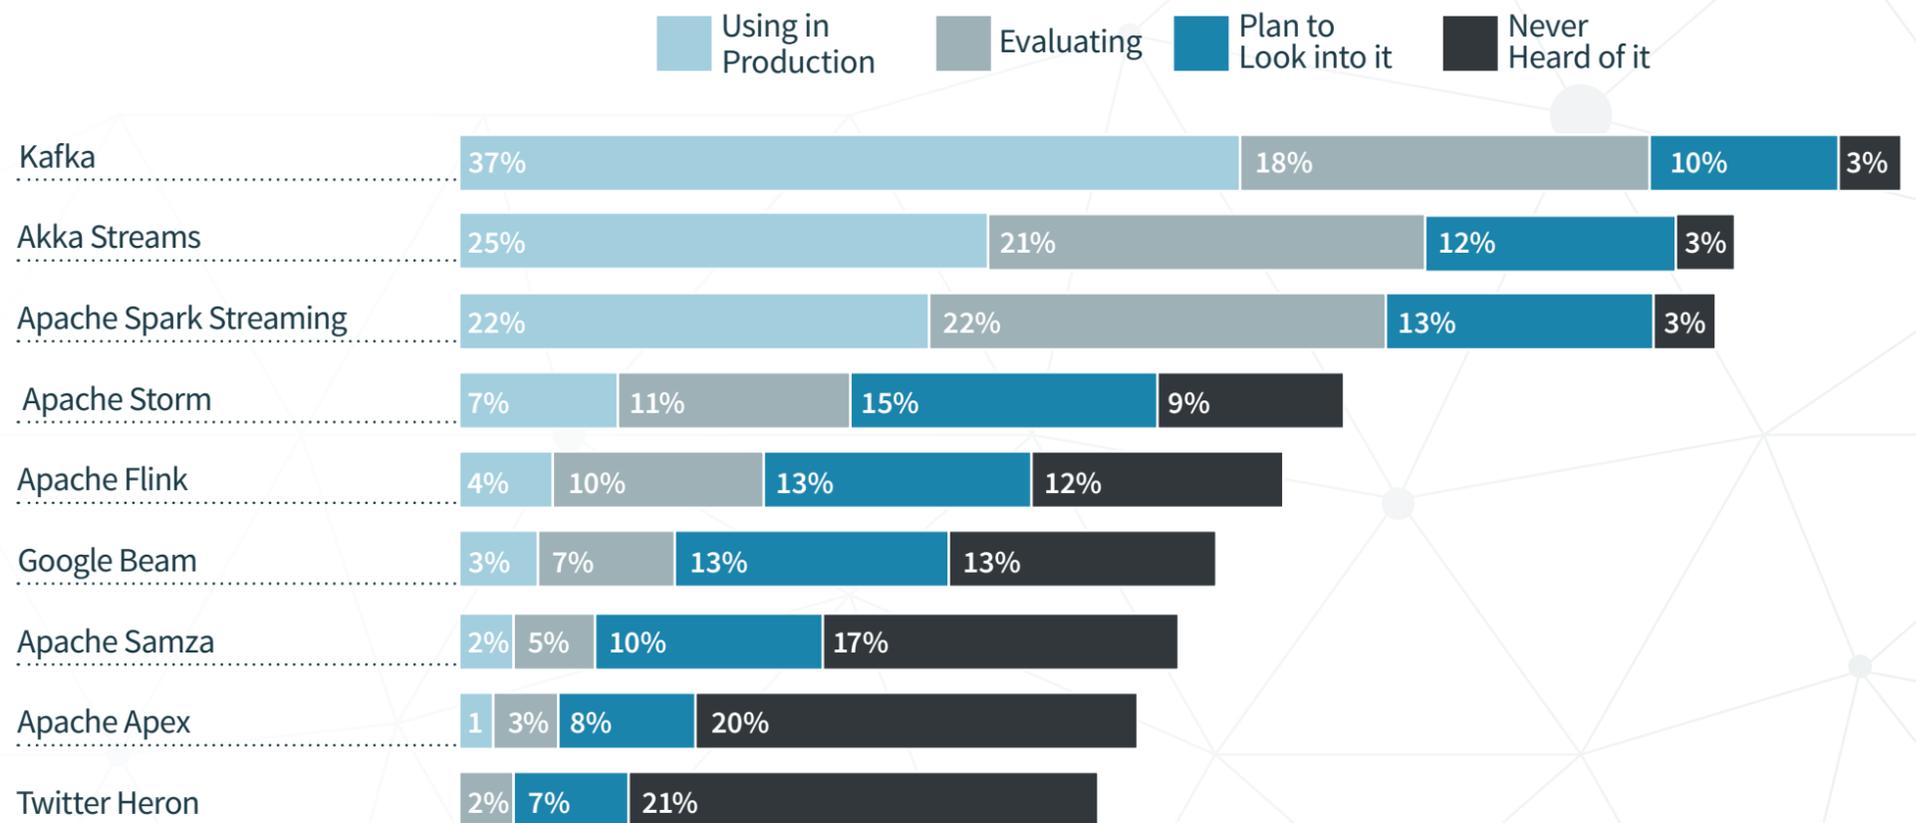
Fast Data Shakes Up Traditional Stack

Choosing the right tools and techniques for fast data can be daunting as the emerging ecosystem of streaming frameworks is constantly shifting and not fully understood by developers and architects.

For most enterprise uses cases, developers will need to mix and match tools based on tradeoffs between latency, volume, transformation, and integration.

Emerging ecosystem is not fully understood

New fast data tools are emerging at a rapid rate. Here we see a progression from experience to awareness across the most popular technologies in the current ecosystem.



Finding
03

Technology Shifts

Fast Data Shakes Up Traditional Stack

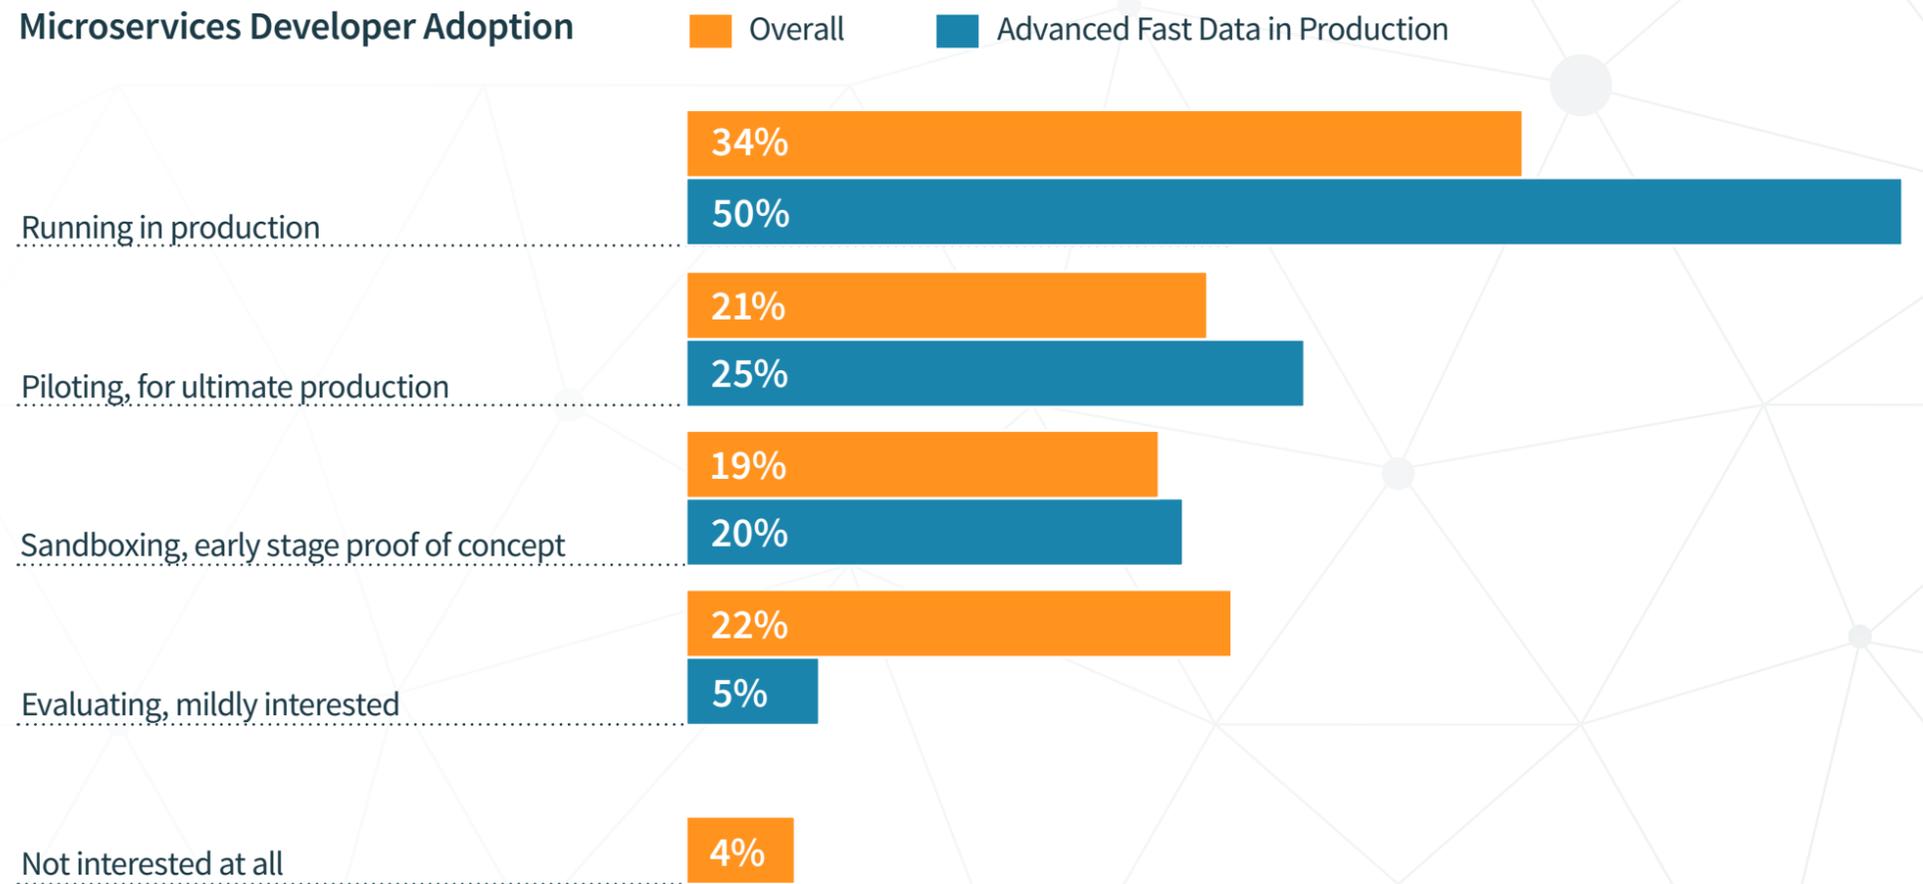
An area of fast data architecture that appears to be more well understood by developers is microservices.

The survey reveals fifty-five percent of developers overall are using or plan to use microservices in production. For developers already running advanced fast data uses cases in production, microservices adoption climbs to seventy-five percent.

Fast data and microservices go hand in hand

A key characteristic of fast data architectures is the use of microservices for streaming applications. Here we see the progression of microservices adoption as advanced fast data use cases move into production.

Microservices Developer Adoption



75%

of developers with advanced fast data use cases in production rely on microservices

Finding
03

Technology Shifts

Fast Data Shakes Up Traditional Stack

The survey reveals a shifting of perceived microservices value as advanced fast data use cases move into production. The most dramatic is developer velocity dropping from the number one position to number four.

The shift in priorities during the production phase of a project is not surprising. DevOps agility and predictability were valued most, followed by elasticity and dynamic scaling, when running advanced fast data use cases in production.

Microservices value shifts in production use cases

Developers are reaching for microservices to improve agility, elasticity, and more. Here are the top factors driving adoption overall contrasted by advanced fast data use cases in production.

Microservices Developer Adoption

	Value of Microservices Overall	Value of Microservices for Advanced Fast Data
Increase development velocity / new releases	1 [26%]	4 [14%]
DevOps agility / predictability for faster, safer deployments	2 [25%]	1 [31%]
Improve elasticity / for scaling up down more dynamically	3 [23%]	2 [28%]
Evaluating, mildly interested	4 [13%]	3 [17%]
Not interested at all	5 [13%]	5 [11%]

Last Look

Steps for Success

Fast data adoption is growing at a rate three times faster than Apache Hadoop. Successful companies are harnessing new, emerging technologies to deliver business value today. Here are steps you can take to start preparing for success in your organization.

01 | Think Ecosystem-First

Fast data apps blend a variety of tools and techniques. Select an ecosystem of engines that can handle shifting tradeoffs between latency, volume, transformation, and integration.

02 | Invest in Training and Development

Help your developers and architects close the skills gap with education. Once your team has been empowered to understand the fundamentals of microservices and streaming, innovation with fast data can follow.

03 | Build DevOps Mindset and Muscle

With microservices and streaming becoming intertwined, the need to embrace DevOps is paramount for fast data success.

04 | Rethink Production Readiness

Begin fast data projects with the end in mind. Implement tools for debugging and monitoring microservices and streams well before entering production.

05 | Prepare for an Ever-More Inter-Connected World

Predicting peaks in streaming applications is nearly impossible. Implement software tools and techniques that deliver resilience and enable elasticity in an ever-more inter-connected world.

Thank you to our partners!

We appreciate you amplifying our call to action for this charity survey,
which generated \$5000 towards the Kojics Foundation.

O'REILLY[®]

 **MESOSPHERE**


CAKESOLUTIONS
ENTERPRISE SOFTWARE SERVICES

skills matter

 **DZone**

 **SOFTWAREMILL**



Lightbend (Twitter: [@Lightbend](https://twitter.com/Lightbend)) provides the leading Reactive application development platform for building distributed systems. Based on a message-driven runtime, these distributed systems, which include microservices and streaming fast data applications, can effortlessly scale on multi-core and cloud architectures. Many of the most admired brands around the globe are transforming their businesses with our platform, engaging billions of users every day through software that is changing the world. For more information on Lightbend, visit www.lightbend.com.

Lightbend, Inc.
625 Market Street
10th Floor
San Francisco, CA 94105

www.lightbend.com