

451

Research®

PATHFINDER REPORT

The Journey to Serverless-First

ENTERPRISE USE CASES SHOW THE WAY

COMMISSIONED BY



DECEMBER 2019

©COPYRIGHT 2019 451 RESEARCH. ALL RIGHTS RESERVED.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

ABOUT THE AUTHOR



JAY LYMAN

PRINCIPAL ANALYST,
CLOUD NATIVE AND DEVOPS

Jay Lyman is a Principal Analyst with 451 Research's Applied Infrastructure & DevOps Channel. He covers infrastructure software, primarily private cloud platforms, cloud management and enterprise use cases that center on orchestration, the confluence of software development and IT operations known as DevOps, Docker and containers. Jay's analysis encompasses evolving IT operations and software release models, as well as the technology used to create, deploy and support infrastructure and applications in today's enterprise and service-provider markets. Key areas of research also include OpenStack, PaaS and enterprise end users.

Executive Summary

Serverless is more than a shiny new technology. In a recent survey of enterprises, more than 15% reported using serverless at the team level or above, and 43% said they expect to deploy serverless technology over the next two years (*see Figure 1*). This quantitative research is backed up by qualitative investigation.

Across industries, organizational size and organizational maturity, serverless continues to gain widespread interest and adoption for enterprises and startups alike. The benefits of a serverless approach are compelling – a pay-for-value pricing model is cost-effective, and the elimination of tasks such as server management and scaling enables development teams to spend more time on new products, features and innovation. However, it's not enough to replace traditional infrastructure with functions as a service – benefiting from a serverless operational model requires a new way of thinking about application design. Many organizations treat this learning process as a journey.

While organizations may embark on a serverless journey for different reasons, the way they do so is often similar: Choose a discrete task, build incrementally, and move forward one step at a time. Many choose to eschew refactoring – rebuilding their previous applications – and instead get started with serverless by building net new applications on top of, or in conjunction with, their existing architecture.

Serverless, defined here as an operational model that eliminates the burdens of server and infrastructure management, has been widely adopted by the three companies covered in this paper: iRobot, a robotics-focused technology company; Alma Media, a European media company; and Fender, the American music giant. These companies first considered serverless for a variety of reasons – achieving rapid release cycles, lowering costs, scaling efficiently, and the ability to attract and retain top talent. Each of these companies also arrived at a similar conclusion: the decision to build 'serverless first.' For iRobot, Alma Media and Fender, building serverless-first means opting for serverless technologies as their first choice, which enables their teams to gain expertise with the new technologies quickly and achieve the agility, elasticity and total cost-efficiency benefits from the overall serverless operational model.

This paper offers a closer examination of iRobot, Alma Media and Fender's serverless journeys. It highlights their adoption patterns as diverse customers with divergent starting points but shared goals – finding ways to deliver products faster, save money, scale efficiently and nurture team enthusiasm. We present these case studies to provide actionable insights to companies pursuing those same goals.

Use cases



iRobot

COMPANY	iRobot
INDUSTRY	Consumer robotics and smart home
EMPLOYEES	1,000+
HEADQUARTERS	Bedford, Massachusetts
USE CASE	Serverless for the edge
KEY SERVERLESS SERVICES	IoT platform, API Gateway, compute, database, storage, data streaming, ETL, query service, provisioning

Connected technology company and maker of the Roomba robotic vacuum cleaner, iRobot wanted to scale its technology out without scaling its teams up. This required the company to consider how it could deliver its products and services more efficiently, and how it could reduce the number of services that it had to manage. Concurrently, iRobot needed to maintain or improve reliability and get new features to market quickly.

The company saw serverless as the right fit: serverless met iRobot's technical needs while allowing it to expand customer capabilities and enter new markets without increasing team size or infrastructure. Today, the team size is about 25 people, all of whom get to focus on building applications that drive business value because they don't have to worry about server and infrastructure maintenance. Innovation remains at the heart of iRobot's vision. That innovation is powered by the ability of development teams to run efficiently and increase the rate at which they release new features. The company has been able to enhance its robots and expand into the smart home market without expanding its team.

"Being serverless, we have so many fewer tasks to handle," says Ben Kehoe, cloud robotics research scientist at iRobot. "Our operations burden has not increased as our fleet of connected robots has vastly increased in size over the past couple years." Serverless, says iRobot, enables the company to handle spikes in orders without additional personnel. On Christmas day, for example, daily traffic spikes to 20x the average. In traditional architectures, that spike would necessitate the allocation of extra engineering time and resources. With serverless, those spikes become a 'non-event.'

Starting with Serverless

iRobot leveraged serverless to drive new features, capabilities and integrations, including a remote push-to-clean button, Wi-Fi mapping, robot teaming for vacuum and mopping capabilities, enhanced home coverage, and integration with Alexa voice control.

The company first created its own deployment tooling since existing options were not sufficient for its scale and application needs. To do this, iRobot used its cloud provider's services and its own development teams to write shorter deployment templates and to package code. It then used serverless to build permissions and access to robots for registration and analytics and to manage firmware rollouts. It wanted to build scalability in at the outset, so iRobot leveraged its provider's IoT service and built only using its serverless compute platform.

After getting development tooling, templating and packaging right, iRobot simplified its API gateway for serverless-to-database. The new direct approach did not require any new code or scripting, which further reduced what the company had to manage. Adopting serverless has meant fewer tasks to handle and a lower operations burden despite the vastly increased fleet of connected robots over the last two years. Kehoe credits serverless for enabling iRobot to focus solely on expanding and evolving its core business, "We went from being a device company to being a connected device company without becoming a cloud technology company. We're a cloud-enabled features company, which is what matters to us."

What's Next for iRobot

Serverless use cases have spread at iRobot, thanks largely to their initial success and the advocacy of internal champions who educate their teammates about the role serverless played in that success. The company hosts a growing breadth of services for its robots and the applications that support them – managing permissions, storing mission history, pushing notifications and supporting integrations. The company's serverless deployments now include an analytics component that has grown in complexity to generate data insights that enhance customer experiences. Now, iRobot is considering more support for data science using serverless, as well as shifting CI/CD pipelines to serverless.

"We went from being a device company to being a connected device company without becoming a cloud technology company. We're a cloud-enabled features company, which is what matters to us."

"Being serverless, we have so many fewer tasks to handle. Our operations burden has not increased as our fleet of connected robots has vastly increased in size over the past couple years."



Alma Media

COMPANY	Alma Media
INDUSTRY	Media
EMPLOYEES	1,900
HEADQUARTERS	Helsinki, Finland
USE CASE	Spikey workload and low-touch operations
KEY SERVERLESS SERVICES	Compute, API gateway, AppSync, database, storage, data streaming

Alma Media is a digital services and multi-channel media company that operates in Finland, Sweden and in Eastern Central Europe, serving both business and consumer markets.

Within its competitive, advertising-based market, financial prudence is paramount. When one of Alma Media's applications received unpredictable, spikey traffic, the company began investigating a solution that offered right-sized scaling on demand, without worrying about the risk of overprovisioning servers and overpaying for unused resources. With 280 developers who work specifically on software development and are responsible for building and running their own applications, Alma Media was attracted to serverless because it allows focus to be on developing new features instead of operating existing infrastructure. The company saw great benefit in the ability to enable its technical staff to focus on new initiatives and the development and improvement of existing products. "The main thing," says Niilo Ursin, CTO of Alma Talent, "is developing new and better services to our customers. We don't really have people maintaining servers...we don't think that is our objective."

Getting Started with Serverless: Beginning with Scripts and APIs

Alma Media first started moving deployments and operations to AWS in 2012, and the first real shift to serverless started in 2017. The company didn't let the idea of legacy impede it from building new solutions to current challenges.

Alma Media started to build serverless solutions for customer-facing products and scripts for static storage files and background jobs. These maintenance jobs, updates and database workloads were already suited for serverless based on their size and simplicity. Among Alma Media's first larger-scale achievements with serverless in production in 2017 was an API to support its media websites, which today attract 750 million monthly page views. Serverless also enabled Alma Media to sync user settings across 80 services.

Spreading the Serverless Mindset

After seeing initial success, the company expanded its use of serverless to more APIs and databases. As it progressed with its serverless ReactJS server-side rendering, Alma Media broadened the APIs, sites and services that it deployed on serverless infrastructure. With deployment times accelerating and teams able to prototype and ship features faster, the company also observed noticeable enthusiasm for serverless across its development teams, which, in turn, has enabled recruitment and retention of top engineering talent.

With Alma Media's success came the desire to share. Development teams saw the significance of sharing serverless knowledge and best practices with internal teams and outside organizations, with the goal of helping engineers and companies understand how to best apply serverless strategies. Today, many teams in Alma Media continue to build with a serverless-first mentality.

Alma Media has always heavily believed in infrastructure as code, which is especially essential in serverless development. In 2019, it started using AWS Cloud Development Kit for developer-friendly infrastructure development.

Recent noteworthy additions to Alma Media's serverless use include managed data service for GraphQL and serverless relational databases. It credits serverless growth to starting small, beginning with a limited number of services and applications, and focusing on solving today's problems rather than past issues. Alma Media's progression from weekly deployments to multiple daily production releases and the considerable time and resource savings were proof enough for the company to do so.

"In the serverless model, we only pay for what we use – we don't pay for any feature or QA environments. When it's not used, we don't pay for it."

"The main thing is developing new and better services for our customers. We don't really have people maintaining servers...we don't think that is our objective."



Fender

COMPANY	Fender
INDUSTRY	Music
EMPLOYEES	2,600
HEADQUARTERS	Scottsdale, Arizona and Hollywood, California
USE CASE	Serverless from the ground up
KEY SERVERLESS SERVICES	Compute, database, storage, network connection, distributed search and analytics, data streaming, data lake, monitoring, software deployment, CDN, identity and access control

In 2016, Fender launched Fender Digital, a digital arm focused on using data to build products that educate and empower musicians and aspiring players. The Fender Digital team launched Fender Tune, a digital tuning application that allows players to tune their guitars, basses and ukuleles using their mobile phone. Tune was shortly followed by Fender Tone and Fender Play, both digital tools and learning applications tightly coupled with Fender's famous physical products – guitars, amplifiers and other instruments.

Fender Digital built its full suite of flagship applications using serverless technologies and a microservices architecture. When considering which architectural approach to take, the organization wanted to imbue its teams with the same creativity that Fender, the historical company, is known for. Fender Digital wanted to give its engineers the space to innovate, prototype quickly and test applications without losing development cycles. For Fender Digital, serverless was a natural choice because it offered its team the ability to “maximize its time for trying new things and experimenting with new ideas,” says Josh Couch, VP of engineering at Fender.

In addition to the pace of innovation that serverless enables, Fender Digital looked to serverless as a way to scale efficiently and manage its team size. Says Couch, “One of the main goals for us was scalability. We realized we were going to be a fairly small team, so we had to be *really* efficient... serverless really seemed like it would be a true, scalable architecture for us.” The approach has proven successful – Fender Digital estimates its team would need to be at least 50% larger if it had to operate without serverless infrastructure. Serverless also allows the team to focus on what really matters to the business: empowering players to continue learning how to play their instruments. The idea of focusing only on what matters to its customers is possible, the company says, because the team doesn't have to spend any time deploying and managing servers as would be required with legacy infrastructure.

Fender Digital also credits serverless technologies with reducing its costs and retaining excellent software engineers. Thanks to fewer manual processes to manage and minimal infrastructure to track, Fender's cloud provider bill is the same or less than it was a year ago, even though it has twice as many paying users and offers 40% more content. Finally, the company notes that serverless has lent the development team great cachet. Fender Digital credits its use of serverless as a basis for attraction and retention at Fender because it allows for creative thinking and offers compelling design challenges that inspire its engineers.

Getting Started with Serverless: Fender Tune

Fender Digital invested research and development effort into choosing the architectural approach and technologies it would use to build its suite of flagship digital services. The company began with a proof of concept built on serverless. Next, Fender Digital's leaders took time to ensure that they could monitor and manage services, throughput, bottlenecks, tracing and alerting for serverless, which was a critical component to get right because Tune had to connect with Fender Connect, the company's authenticated single sign-on service. Finally, leaders had to ensure that their applications were secure and that their developers could work with them efficiently. With its built-in security and reliability and its ability to scale on demand and handle high, bursty workloads, serverless emerged as the clear architectural choice, transforming Fender Digital's premier application from proof of concept to full production workload in months. Fender Digital's success with Tune validated its serverless decision, and the company quickly followed with its Tone and Play applications – fully built using serverless and capable of handling terabytes of content.

Embracing serverless technology enabled Fender to ideate, prototype and build applications for delivering educational, interactive video content that requires scaling to terabytes of data, all with a small development team – a feat the company believes would have proven impossible without the adoption of serverless.

“One of the main goals for us was scalability. We realized we were going to be a fairly small team, so we had to be really efficient.... serverless really seemed like it would be a true, scalable architecture for us.”

“I do think that (serverless) helps us retain great engineers – to keep challenging them, keep giving them things to learn and improve and push themselves on.”

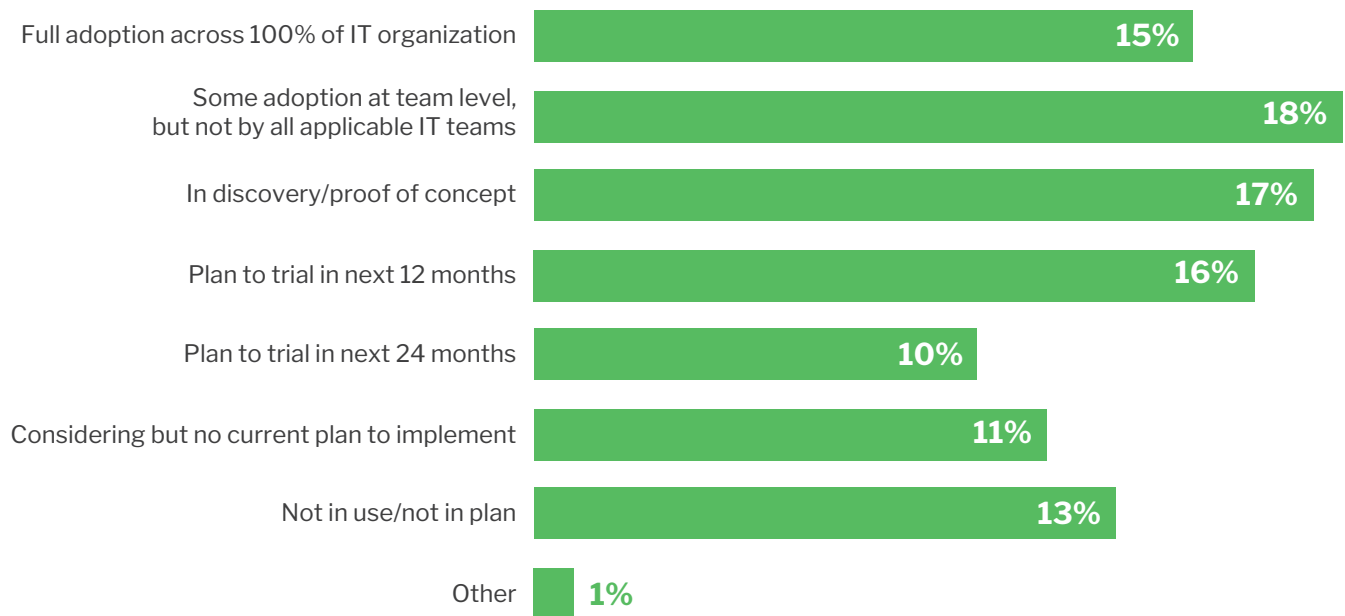
Technological Discussion

Serverless in the Enterprise

Enterprise usage of serverless technologies has grown immensely in the five years since the release of AWS Lambda, the market's first serverless compute engine. Serverless usage continues to grow as organizations move from discovery phases to production workloads. A recent 451 Research survey indicates that serverless technologies are broadly used at the organizational level for 15% of respondents and at the team level for 18% of respondents. The data further illustrates that 43% of enterprises are in discovery/proof-of-concept or planning phases, and are poised to deploy serverless more broadly over the next two years (see Figure 1).

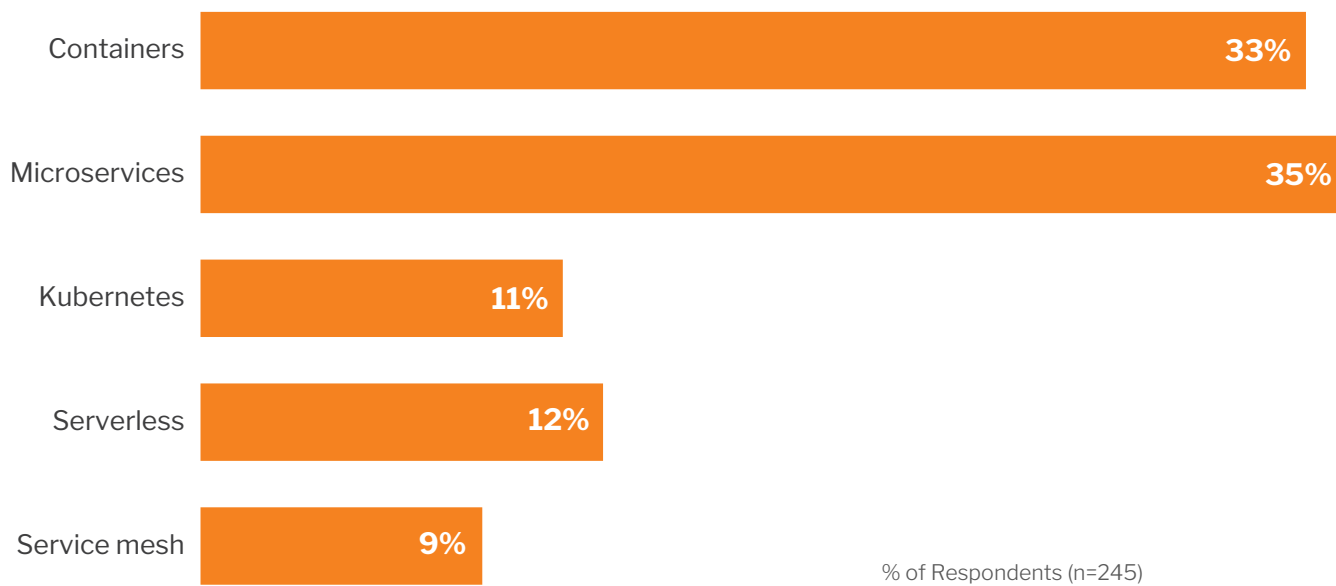
Figure 1: Level of serverless adoption

Source: 451 Research's Voice of the Enterprise: DevOps Q3 2019



451's Voice of the Enterprise: DevOps Q1 2019 survey also found that cloud-native technologies, including serverless technology, are of critical importance to enterprise DevOps teams. Survey results indicate that enterprise DevOps teams are heavy users of these technologies – containers and microservices most prominently, followed by serverless and then Kubernetes (see Figure 2). These results suggest that despite its market nascence in comparison to containers and Kubernetes, serverless infrastructure plays a significant role and is well understood by enterprise DevOps teams.

Figure 2: Most critical cloud-native technologies
451 Research's Voice of the Enterprise: DevOps Q1 2019



Recommendations

While the organizations profiled in this paper were first inspired to use serverless technologies for different purposes and workloads, their foundational reasons for choosing serverless infrastructure were the same: enabling a faster time to market; eliminating unnecessary operational burdens, server management, and overhead; seamlessly scaling to meet demand; reducing costs; and empowering teams to move quickly. These organizations began building with serverless infrastructure from different starting points – iRobot was addressing specific IoT and edge use cases, Alma Media was transitioning from traditional infrastructures to the cloud, and Fender Digital was building net new applications. Nevertheless, they arrived at outcomes that share a common thread: their work with serverless services ultimately led them to adopt a serverless-first strategy for future development work across their organizations.

Throughout these case studies, three themes remained consistent. First, serverless technologies freed developers and teams from manual tasks and management toil, enabling them to focus on innovation. Second, organizations were able to achieve efficient management at massive scales with smaller teams. Third, serverless helped these companies attract and retain top technology talent. Based on the learnings from these case studies and our other research, we believe that organizations can successfully initiate or expand their own use of serverless by considering the following steps and guidance.

PHASE 1: GETTING STARTED WITH SERVERLESS

Start small. Organizations don't need to rewrite everything for serverless to get started. It's more important to start with one serverless project – one small API or feature that enables teams to familiarize themselves with the products and their capabilities – rather than spend too much time planning for a broad beginning. Other tips from organizations interviewed for this paper include:

- Research the market and consider the types of applications and services to be run on serverless before starting or expanding usage.
- Build a proof-of-concept project that is limited in scope to allow teams to gain familiarity with serverless. For example, automating routine IT tasks is often a good starting point.
- Consider starting with small APIs and new features. These are less intimidating and can be ideal for new serverless deployments.
- Share learnings at internal team meetings, engineering deep dives and during all-hands meetings. This internal evangelization helps teams onboard faster and share best practices, helping to eliminate redundant challenges.
- Engage in the broader serverless community by attending the growing number of conferences and meetups focused on serverless education.

PHASE 2: EXPANDING YOUR SERVERLESS USE

Use your initial serverless proof of concept as a starter kit for additional serverless applications. Once organizations have successfully built a proof of concept, they can build on what they've learned. Teams should identify components of their existing architecture that can be easily transformed into serverless, or evaluate net new features they're seeking to implement and consider how they can be built using serverless. Other ideas from organizations interviewed for this paper include:

- Use a proof of concept to better understand what the costs will be, and how to configure serverless components for maximum efficiency.
- Prepare for data migration and database integration needs. These may require some initial time and resources, so it's important to consider them.
- Leverage provider services and supporting integrations to build security into serverless deployment. Important tooling includes authentication and identity management.
- Look for existing tooling and best practices from providers and from the robust serverless community.
- Establish company-wide governance, deployment and operational tooling practices as applied to serverless technologies.
- Continue promoting a learning culture within the organization, offering developers the opportunity to experience and share serverless learnings.

PHASE 3: MOVING INTO A SERVERLESS-FIRST STATE

Apply serverless to processes such as CI/CD; use cases like media and document workflows, data processing, and IoT; and end-user-facing applications. These workloads can further streamline the development process, reduce costs and prove the serverless model in practice. Additional advice from organizations interviewed for this paper includes:

- Use production workloads as a benchmark and proof point of serverless success.
- Build an audit mechanism for new application architectures to evaluate the serverless implementation.
- Consider contributing to the community and amplifying your organization's engineering success by writing blog posts or sharing reference architectures of your own.
- Evaluate the kind of skills your organization will need to scale your serverless implementation and build programs to train existing talent and/or hire new talent

Conclusions

We expect serverless to be more widely deployed over the next two to three years. In the longer term, we expect deeper crossover and integration of cloud-native technologies and methodologies, including serverless, with adjacent trends including data analytics and data science, artificial intelligence and machine learning, and IoT/edge computing.

These technological frontiers, already trending toward an expected norm for many organizations, highlight both a challenge and an opportunity for enterprises. The challenge is that enterprises must leverage existing teams and technology while seeking new talent and opportunities to use emergent technology to facilitate growth, efficiency and cost reductions. For enterprises, serverless offers a solution to those challenges, enabling teams to accelerate product and feature delivery, reduce costs, eliminate undifferentiating work, focus on the user experience and attract top talent. Based on the experiences of the companies interviewed for this paper, we anticipate headlines of serverless success stories to become more prominent and mainstream than they are today.

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

© 2019 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such.

451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.



NEW YORK

Chrysler Building
405 Lexington Avenue,
9th Floor
New York, NY 10174
+1 212 505 3030



SAN FRANCISCO

505 Montgomery Street,
Suite 1052
San Francisco, CA 94111
+1 212 505 3030



LONDON

Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 203 929 5700



BOSTON

75-101 Federal Street
Boston, MA 02110
+1 617 598 7200