

Serverless: Changing the Face of Business Economics – a Venture Capital and Startup Perspective

April 2018



greylockpartners.



Accel

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Table of Contents

How Serverless Impacts Startups, Investors, and the Business Economics of the Future	4
Greylock Partners: The Serverless Opportunity.....	8
Greylock Partners’ Serverless Investing Thesis.....	9
Greylock Partners Portfolio Companies Spotlight:	10
Nextdoor	10
Skyhigh	11
Madrona Venture Group: Intelligent Applications Powered by Serverless Approaches.....	13
Madrona’s Serverless Investing Thesis	14
Madrona Portfolio Companies Spotlight:	15
Algorithmia.....	15
Pulumi	15
IOpipe.....	16
Snowflake.....	17
Accel: From Mainframe to Serverless and Beyond.....	18
Accel’s Serverless Investing Thesis	19
Accel Portfolio Companies Spotlight:	20
Segment	20
Trifacta	20
Conclusion.....	21
The Contributors	21

How Serverless Impacts Startups, Investors, and the Business Economics of the Future

By TIM WAGNER, *General Manager, AWS Lambda, Amazon API Gateway, and AWS Serverless App Repository at Amazon Web Services* and LINDA LIAN, *AWS Product Marketing Manager*

Startups that adopt serverless architectures find they are more agile and have dramatically lower operating costs, enabling them to focus their resources on business value-add instead of infrastructure. In this white paper we examine the startup experience with serverless through the lens of three of the industry's prominent VCs – Greylock Partners, Madrona Venture Group, and Accel – and take an in-depth look at the portfolio impact serverless has had on each. We examine in detail how services like AWS Lambda dramatically enhance agility and lower operating costs for startups, enabling them to better concentrate their investments on business-critical deliverables and outcomes.

We start below with a Q&A with AWS General Manager of Serverless, Dr. Tim Wagner. We'll discuss why serverless matters for doing business today, and some thoughts about where serverless could take us.



Dr. Tim Wagner – General Manager AWS Serverless

What is serverless and why is it a paradigm shift for business economics?

Serverless computing enables companies to build and run applications and services without thinking about servers. The serverless approach is applicable to nearly any type of application or backend service, and all sizes and maturity of organizations, from startups to multi-national corporations. With serverless compute services such as [AWS Lambda](#), provisioning, scaling, and management of servers are handled by AWS. Building serverless applications means that developers can focus on their core product instead of worrying about managing and operating servers or runtimes, either in the cloud or on-premises. This reduced overhead lets developers reclaim time and energy that can be better spent on developing innovative, scalable and reliable products. It

also allows for the ability to achieve massive scale with minimal human capital, making developers and companies alike more agile in a fast-moving economy.

In addition to accelerating innovation, serverless approaches offer a second advantage that's particularly critical for startups and venture capital portfolios at large: lower costs. Compute power is often a key element of a startup's cost structure: consumer-oriented businesses often run services that are used by hundreds to millions of users 24/7 and need to optimize logistics and sales through deep analytics. Business-to-business companies often require APIs and cross-company services with tight service-level agreements (SLAs). Business-to-employee services and apps strive to make employees productive with real-time information and processing capabilities.

In delivering these experiences, startups struggle with a common problem: They want to “get off on the right foot” with customers by providing a reliable service that scales effectively, while simultaneously being cost-conscious and using their limited resources wisely. AWS Lambda enables a startup to achieve these two seemingly opposing goals easily.

At the heart of serverless computing is a simple but revolutionary idea – that companies only pay for compute when they actually use it. Lambda’s approach to cloud computing disrupted the economics of the industry, because it made computation proportional to the workload. If an application or service has little or no workload (for example, in the middle of the night), there are no charges. If there’s a sudden spike, such as a flash sale, Lambda will instantly scale to accommodate the massive load. In between, the compute is always exactly aligned to customer-and-workload-driven usage, never higher (which would result in low utilization and wasted money) or lower (which would result in low availability and poor end user experiences). Compared to server-based architectures, Lambda’s pay-per-use model can reduce cost of goods sold (COGs) and lead to more competitive margins. It makes it easy for companies to assign costs on a per-customer, per-use basis. Serverless architectures double down on the cloud-enabled software-as-a-service (SaaS) business model by providing recurring revenue and higher gross margins even for companies not yet operating at scale, enabling new models of selling and consuming services.

“At the heart of serverless computing is a simple but revolutionary idea – that companies only pay for compute when they actually use it.”

Why do we see startups and serverless as a natural match?

Serverless is a great fit for startups, which need to get up and running with low upfront investment and limited human and machine capital. AWS Lambda and other serverless technologies concentrate time and money investments, to the point where even a single talented developer can create massive fixed leverage by generating revenue from a popular serverless app or service.

AWS Lambda’s event-driven, stateless, and functional nature make it a good fit for the diverse array of modern architectural patterns as well as emerging patterns just being developed. Today, many web, mobile, and data processing backends are built serverlessly. Emerging patterns that fundamentally leverage serverless approaches include chat bots, Alexa skills, machine learning, and internet of things (IoT). Importantly, in all of these patterns, using AWS Lambda allows developers to focus on writing business logic instead of infrastructure-related code. This lowers the time to market and tightens innovation cycles by funneling resources to the application level, away from undifferentiated activities like capacity planning, server and operating system upgrades, security patches, and machine monitoring – tasks that can quickly drain the much-needed resources of a small company.

Is the advantage startups get from serverless approaches and AWS Lambda different from the benefits enterprises receive?

The benefits that serverless approaches like Lambda provide aren’t different for a startup, but the fast pace of delivery, the need to deliver results with a minimum amount of new investment (or limited run rate impact to a prior investment), and the focus on business model and technology agility make the decision to adopt serverless even more critical for a startup. While the pressures on an established

enterprise company are similar in nature, startups have a shorter runway in every dimension – time, money, and people – and need to deploy their precious resources with even greater care.

How does buying and selling software change in a serverless world?



*Linda Lian – Product Marketing Manager
AWS Serverless*

Multitenancy was an important technological breakthrough for SaaS. It enabled the same instance of software to be served to multiple tenants or customers, instead of each instance of software living in its own server and database. This consolidated IT resources and lowered costs for SaaS providers (who can benefit from improved margins or elect to pass savings onto their end users, giving them a competitive pricing advantage). Multitenancy is a large part of why higher gross margins are a characteristic of SaaS companies today. It changed the way software was sold and consumed from traditional enterprise licenses locked to CPUs or server counts to the modern approach characterized by users or “seats.”

Similarly, serverless compute is both a new way of building software *and* a new way of consuming and paying for it. Multitenancy on a per-request basis makes it easy to construct SaaS services that offer either a per-

seat software consumption model *or* an even finer-grained per-request model. Imagine business leaders no longer needing to audit, track and manage the proliferation of SaaS tools and associated seats within an enterprise because costs are matched exactly to actual usage. Imagine that metrics like quality of service and guarantee of speed could alter pricing. With serverless, new, disruptive pricing models become possible – we can’t wait to see how companies will innovate in this space to deliver even more cost-effective and user-friendly pricing models!

How are Venture Capital firms reacting to serverless?

Venture Capital funds are interested in the serverless movement for several reasons. From a purely pragmatic perspective, anything that enables their investments to be used more productively by improving the leverage of dollars invested in human and machine capital improves the probability of a successful outcome. Helping portfolio companies adopt serverless generally lowers the burn rate by making startups more efficient.

At a deeper level, building an entire business model around a serverless approach doubles down on the “SaaS Advantage”, providing a startup with an improved cost structure and the opportunity to provide request-level multitenancy, passing those improved economics along to end users, regardless of the domain. The ability to track both the revenue and cost of an end user’s every interaction with a web, mobile, or internet of things (IoT) app lets even tiny startups access the kind of customer analytic power and machine learning enablement

“...building an entire business model around a serverless approach doubles down on the ‘SaaS Advantage.’”

that was formerly restricted to only the largest and wealthiest of Fortune 100 companies.

Finally, venture capital firms are also increasingly investing in the growing serverless ecosystem itself, which is repeating the development of the infrastructure-as-a-service (IaaS) ecosystem, but at an even faster pace. Startups purpose-built to support developer and enterprise adoption of serverless are already present in spaces such as deployment, monitoring, security, and a host of domain-specific frameworks. With the launch of the AWS Serverless Application Repository, startups can now distribute serverless apps and freeware and engage their communities faster, easier, and with greater reach than ever before.

Next, we will hear from three venture capital funds – Greylock Partners, Madrona Venture Group, and Accel on why they are betting big on serverless, and how their portfolio companies are benefitting from it.

greylockpartners.

By JERRY CHEN, *Partner and* SARAH GUO, *Principal*

During the past five years, we have seen accelerating change in software architectures and development processes. Rising consumer expectations for constantly improving digital experiences and the continued growth of mobile drive the demand side. Supporting this sea change on the supply side is the availability of more flexible and economical public cloud services at every layer of the stack.



Jerry Chen – Partner, Greylock Partners

The increasing pace of change in software development requires organizations to restructure their efforts and split their applications into microservices in order to keep up. The more they can independently evolve parts of their increasingly complex systems at scale, and decouple their innovation (applications and user experience) from infrastructure management, the better positioned they will be to serve their customers.

The promise of serverless is that application developers will no longer have to provision and manage physical servers, operating systems, traditional infrastructure security, or the myriad of other functions. They will be able to consume the compute (or packaged functions) they want, in a more granular and flexible way, with the freedom to use the right language for the job. While we are early in this shift to serverless, we already see many of our forward-leading engineering organizations adopt this mentality, and we believe it represents a significant long-term opportunity for enterprises and the vendors that serve them.

Greylock Partners' Serverless Investing Thesis

Although there is much promise and excitement around serverless, the ecosystem that supports these new, highly distributed architectures is nascent. Unlike Docker in the container world, there is no dominant standard yet. Tooling is still emerging. *We believe the beginning of ecosystems like this is a promising investment opportunity.*



Sarah Guo – Principal, Greylock Partners

We are especially interested in investing in companies that:

Help enterprises evolve their architecture to microservices and serverless. Enterprises manage a complex and shifting range of application environments (on-premises, multiple clouds). The first step is often to offload some parts of application delivery and resilience to services that enable modern architectures. Enterprise solutions include Avi Networks, API proxies like Kong, as well as OSS inter-service meshes like Istio and Linkerd.

Enable better understanding of (and management of) highly distributed systems with varied end-user patterns. The shift to serverless (and to microservices) is not free. Independently evolving services, some of which your development organizations own, some of which are built and run by third parties, must work gracefully together in a tightly-directed concert under many different types of often unpredictable end-user load. This is no small feat, and we believe that understanding these systems will require new approaches beyond traditional unstructured logging and metrics. Structured, dynamically sampled logs, distributed tracing, iterative debugging in production, and proactive tools that draw more conclusions about when business is impacted and what component or build is correlated with a particular behavior change, will augment the tools we have today. Companies like LightStep and Honeycomb are already making strides in this area - but we think analytics of systems is going to get a lot smarter.

“We believe the beginning of ecosystems like this is a promising investment opportunity.”

Support developers by making (everything) easier. We are interested in companies that try to give specific audiences a serverless experience for a small piece of the tech stack, out-of-the-box. Examples would include service marketplaces like RapidAPI, Netlify for front-end web developers, deployment services like Zeit, and Graphcool for working with databases. These companies don't expect users to

adopt serverless from the ground up — they simply use serverless principles to make something easier for developers.

Secure services and applications in a cloud-native way. Companies like SkyHigh and Obsidian help organizations move away from traditional host and network-oriented security models to more application, data and identity-centric models. More granular service (potentially even function-level) controls with usable and maintainable policy models still lie ahead.

“The way a company is architected...is the basis for product velocity – one of the biggest differentiators a startup (or any company) has.”

In addition to looking for these opportunities related to the serverless ecosystem, Greylock Partners evaluates the technical architecture of companies we are looking to invest in. The way a product is architected is often

difficult to change as a company grows, and can be a significant driver of cost at scale. More importantly, architecture is the basis for product velocity – one of the biggest differentiators a startup (or any company) has.

Below, we’ll look at how two of the startups in our portfolio are leveraging serverless technologies to innovate faster and better serve their customers. [Nextdoor](#) is the private social network for your neighborhood. [Skyhigh](#), now part of McAfee, provides total control over all your data in the cloud, so you can confidently leverage the power of the cloud to accelerate your business.

Greylock Partners Portfolio Companies Spotlight:

Nextdoor

At Nextdoor, data is the lifeblood of the organization. Whether influencing decisions on how to improve product features, fueling analysis of their users’ behavior, or monitoring the health of Nextdoor systems, the Nextdoor platform ingests, transforms, and analyzes massive volumes of data. Over three billion events per day are collected and delivered through Nextdoor’s data pipeline, which has recently been completely rebuilt in a serverless architecture.

Using AWS services including [Amazon Kinesis](#), [AWS Lambda](#), [Amazon SNS](#), and [Amazon S3](#), Nextdoor’s operations team was able to replace the legacy data pipeline based upon Apache Flume and custom ETL software. In so doing, the team improved reliability of the pipeline, reduced their maintenance overhead significantly, and improved delivery times of data to the Redshift-based data warehouse and S3-based Data Lake.

Apache Flume, which served them well for over four years and through daily volumes up to 1.5B events, had a number of operational challenges that led to reliability issues and even data loss in some cases. Nextdoor’s small operations team of only three engineers found themselves consumed with responding to CPU alerts, backed-up queues, and reports of lost or corrupted data. Tuning and configuring the pipeline was time consuming, and the pipeline wasn’t elastic enough to respond to surges or spikes of traffic.

As the team designed the new system, their goals were to minimize Nextdoor's operational overhead, provide a scalable, elastic solution with minimal configuration, and ensure that they could meet our stated SLAs for data delivery times and data integrity.

Kinesis replaced Flume outright, relieving Nextdoor of the management overhead of operations. Kinesis scales well horizontally, offers a simple stream-based interface to data, and with Kinesis Firehose, provides an out-of-the-box solution for streaming data into AWS S3. Using SNS, the team was able to automatically trigger processing of incoming data records as they arrive, apply transformations by invoking Lambda functions written in Python, and deliver the modified records to S3 and ElasticSearch without operating any server instances themselves.

The net result was a marked improvement in reliability and speed of delivery. 99.95% of records processed through the Nextdoor data pipeline appear in their S3-based Data Lake in under 5 minutes. Kinesis provides a robust streaming solution that has eliminated data loss entirely. S3 provides a 99.999999999% durability guarantee for stored objects. By migrating to a serverless platform, Nextdoor can continue to operate their infrastructure efficiently and reliably with a very small team, freeing resources up to focus on deploying product enhancements for Nextdoor's members.

Skyhigh

Ever since Skyhigh launched its Skyhigh for IaaS product in Q1 CY 2017, the company has seen a steady increase in the number of large customers migrating their workloads to AWS for a variety of reasons, including business agility, costs, and risk mitigation. At the same time, customers are looking improve their security posture. Skyhigh continues to collaborate with customers to help them take advantage of the agility provided by AWS, while also providing the deep visibility and comprehensive protection of customer data and user activities so they can more confidently leverage cloud for their businesses.

To help customers achieve the goals mentioned above, Skyhigh uses [AWS Lambda](#) functions to empower DevSecOps to not just customize security configuration policies, but also help implement an auto-remediation workflow to resolve incidents. AWS Lambda helps Skyhigh to implement auto-detection and auto-remediation workflows in near-real time for IaaS usage, and increasingly PaaS usage, where traditional approaches that require host agents or software instrumentation for security controls will not work.

Skyhigh's security configuration monitoring feature allows customers to apply more than 50 policies, adhering to the CIS AWS Foundations Benchmark, in order to audit all AWS services in use. Customers can create a policy within Skyhigh to automatically check for any configuration changes by importing a Lambda function that checks for the custom rule specific to the needs of the DevSecOps team within an enterprise.

To illustrate this use case, one of Skyhigh's enterprise customers, a large financial services firm, has a dedicated team writing hundreds (if not thousands) of Lambda functions in order to build custom rules to monitor configurations on all AWS services and applications deployed. This customer uses Skyhigh's custom policy engine, powered using Lambda, to force Amazon S3 to serve content only if the request originates from a specific IP range and HTTPS/SSL is in use. DevSecOps has also created a Lambda

function to deny accesses to AWS CloudTrail buckets for requests originating from outside the defined IP range and if the server-side encryption does not use AES 256 as the encryption scheme, apart from a number of other customized rules.

Skyhigh's security capabilities, in conjunction with AWS Lambda, helps customers expedite their time-to-market by standardizing on consistent policy enforcement and monitoring and also increases business agility by securely leveraging the AWS platform. The serverless compute ecosystem continues to grow, and the future holds great promise for cloud security applications to take advantage of the disruptive functionality and cost model of Lambda.

Madrona Venture Group: Intelligent Applications Powered by Serverless Approaches



By S. SOMASEGAR, *Managing Director* and DANIEL LI, *Principal*

The progression over the last 20 years from on-premises servers, to virtualization, to containerization, to microservices, to event-driven functions and now to serverless computing is allowing software development to become more and more abstracted from the underlying hardware and infrastructure. The combination of serverless computing, microservices, event-driven functions and containers truly form a distributed computing environment that enables developers to build and deploy at-scale distributed applications and services. This abstraction between applications and hardware allows companies and developers to focus on their applications and customers – not worrying about scaling, managing, and operating servers or runtimes.



S. Somasegar – Managing Director, Madrona

“We believe that microservices and serverless functions will form the fabric of the intelligent applications of the future.”

In today’s cloud world, more and more companies are moving towards serverless products like [AWS Lambda](#) to run application backends, respond to voice and chatbot requests, and process streaming data because of the benefits of scaling, availability, cost, and most importantly, the ability to innovate faster because developers no longer need to manage servers. We believe that microservices

and serverless functions will form the fabric of the intelligent applications of the future. The massive movement towards containers has validated the market demand for hardware abstraction and the ability to “write once, run anywhere,” and serverless computing is the next stage of this evolution.

Madrona’s Serverless Investing Thesis

Today, developers can use products like AWS Lambda, S3, and API Gateway in conjunction with services like Algorithmia, to assemble the right data sources, machine learning models, and business logic to quickly build prototypes and production-ready intelligent applications in a matter of hours. As more companies move towards this mode of application development, we expect to see a massive amount of innovation around AI and machine learning, application of AI to vertically-focused applications, and new applications for IoT devices driven by the ability for companies to build products faster than ever.

Nevertheless, with the move towards containerization and serverless functions, it can be much harder to monitor application performance, debug applications, and ensure that applications have the correct security and policy settings. For example, SPIFFE (Secure Production Identity Framework for Everyone) – www.spiffe.io provides some great context for the kinds of identity and trust related work that needs to happen for people to be able to build, share, and consume micro-services in a safe and secure manner.

For all the above-mentioned reasons, Madrona has made several investments in companies building tools for microservices and serverless computing in the last year and we are continuing to look for opportunities in this space as the cloud infrastructure continues to evolve rapidly.

Below are four of the startups in our portfolio and how they are building tools to enable developers and enterprises to adopt serverless approaches, or leveraging serverless technologies to innovate faster and better serve their customers.



Daniel Li – Principal, Madrona

ALGORITHMIA

[Algorithmia](#) empowers every developer and company to deploy, manage, and share their AI/ML model portfolio with ease. Algorithmia began as the solution to co-founders Kenny Daniel and Diego Oppenheimer's frustrations at how inaccessible AI/ML algorithms were. Kenny was tired of seeing his algorithms stuck in an unused portion of academia and Diego was tired of recreating algorithms he knew already existed for his work at Microsoft.

Kenny and Diego created Algorithmia as an open marketplace for algorithms in 2013 and today it services over 60,000 developers. From the beginning, Algorithmia has relied on serverless microservices, and this has allowed the company to quickly expand its offerings to include hosting AI/ML models and full enterprise [AI Layer](#) services.

AI/ML models are optimally deployed as serverless microservices, which allows them to quickly and effectively scale to handle any influx of data and usage. This is also the most cost efficient method for consumers who only have to pay for the compute time they use. This empowers data scientists to consume and contribute algorithms at will. Every algorithm committed to the Algorithmia Marketplace is named, tagged, cataloged, and searchable by use case, keyword, or title. This has enabled Algorithmia to become an AWS Lambda [Code Library Partner](#).

In addition to the [Algorithm Marketplace](#), Algorithmia uses the serverless AI Layer to power two additional services: [Hosting AI/ML Models](#) and [Enterprise Services](#) where they work with government agencies, financial institutions, big pharma, and retail. The AI layer is cloud, stack, and language agnostic. It serves as a data connector, pulling data from any cloud or on premises server. Developers can input their algorithms in any language (Python, Java, Scala, NodeJS, Rust, Ruby, and R), and a universal REST API will be automatically generated. This allows any consumer to call and chain algorithms in any combination of languages. Operating under a Kubernetes-orchestrated Docker system allows Algorithmia's services to operate with the highest degree of efficiency.

As companies add AI/ML capabilities across their organizations, they have the opportunity to escape the complications that come with a monolithic application and begin to implement a serverless microservice architecture. Algorithmia provides the expertise and infrastructure to help them be successful.

Pulumi

[Pulumi](#) is a stealth-mode startup that is fundamentally reimagining how developers build and manage modern cloud systems. By using programming languages and tools that developers are already familiar with - rather than DSLs, templates and home-grown solutions - Pulumi's customers can focus on delivering business value rather than getting mired in tedious infrastructure configuration. The recent rise in serverless computing, containers, and managed cloud infrastructure has enabled this new way of programming cloud applications and services.

As an example, one of Pulumi’s Enterprise customers was able to replace their cloud infrastructure templates to 1/100th the amount of code in a language the team already knew. This enabled a shift away from disjointed Dev and DevOps teams towards a combined Engineering organization, empowering them to meet business demands and realize significant productivity gains. The resulting system highlights the best of what the modern cloud has to offer: dozens of AWS Lambdas for event-driven tasks (replacing a costly and complex queuing system), containers that can run flexibly in ECS, Fargate, and/or Kubernetes, and managed AWS services like Amazon CloudFront, Elasticsearch, and ElastiCache. As a bonus, the system now runs at a fraction of the cost before migrating to Pulumi and the team has been able to spin up entirely new environments in minutes where it used to take weeks.

Pulumi fundamentally believes that the real magic is in serverless, containers, and managed infrastructure living together in harmony. Each has its own strengths: containers are great for complex stateful systems, often taking existing codebases and moving them to the cloud; serverless functions are perfect for ultra-low-cost event- and API-oriented systems; and managed infrastructure lets you focus on your application-specific requirements, instead of reinventing the wheel by manually hosting something that your cloud provider can do better and cheaper. Arguably, each is “serverless” in its own way because infrastructure and servers fade into the background. This disruptive sea change has enabled Pulumi to build a single platform and management suite that fully enables this entire spectrum of technologies.

IO|pipe

[IOpipe](#) co-founders Erica Windisch and Adam Johnson went from virtualizing servers at companies like Docker, to going “all in” on serverless in 2016. Erica and Adam identified serverless as the next revolution in infrastructure, coming roughly 10 years after the launch of AWS EC2. With a shift in computing moving towards a serverless world, there are new challenges that emerge. From dozens of production Lambda user interviews, Erica and Adam identified that one of the major challenges in adopting serverless was a lack of visibility and instrumentation. In 2016, Erica and Adam co-founded IOpipe to focus on helping companies build, ship, and run serverless applications, faster.

IOpipe is an application operations platform built for serverless architectures running [AWS Lambda](#). Through the collection of high fidelity telemetry within Lambda invocations, users can quickly correlate important data points to discover anomalies and identify issues. IOpipe is a cloud-based SaaS offering that offers tracing, profiling, metrics, logs, alerting, and debugging tools to power up operations and development teams.

IOpipe enables developers to debug code faster by providing real-time visibility into their functions as they develop them. Developers can dig deep into what’s really happening under the hood with tools such as profiling and tracing. Once they’re in production, IOpipe then provides a rich set of observability tools to help bubble up issues before they affect end-users. IOpipe saw customers who previously spent days debugging tough issues in production now able to find the root cause in minutes using IOpipe.

Since launching the IOpipe service in Q3 of 2017, the company has seen customers ranging from SaaS startups to large enterprises enabling their developers to build and ship Lambda functions into

production at an incredibly rapid pace. What previously took one customer 18 months can now be done in just 2 months.

IOpipe works closely with AWS as an advanced tier partner, enabling AWS customers to embrace serverless architectures with power-tools such as IOpipe.



[Snowflake](#) was founded in 2012 by industry veterans in the data warehousing space to build a new data warehouse product from the ground up taking advantages of the capabilities unlocked by cloud platforms like AWS. Snowflake's unique architecture, which separates compute and storage, allows a customer to scale their infrastructure easily to meet peak demand while also scaling down and only paying for the actual usage. Snowflake allows customers to solve the biggest most challenging data processing problems while simplifying the process of managing their infrastructure and capacity.

One of the important initiatives for Snowflake was to provide access to not only the biggest customers but smaller organizations and individuals. In order to build our self-service infrastructure, we chose [AWS Lambda](#) to manage the account signup process and managing our internal information for our customers across all the systems. AWS Lambda was very easy to get started, highly reliable by design and had minimal operational costs for us. Given that account signups are bursty by nature, AWS Lambda was able to quickly accommodate a burst in demand without requiring us to maintain a large 24X7 infrastructure which is lightly utilized.

Since our initial use of AWS Lambda, we have additional uses in monitoring our systems and in our data engineering pipelines. The simplicity, elasticity and availability of AWS Lambda has made it very appealing and a critical part of Snowflake as we continue our growth.

The future of cloud computing is not just migrating existing workloads but also enabling a new architecture enabled by services like AWS Lambda and Snowflake which allow customers to take advantage of near limitless storage and compute infrastructure provided by AWS and build new capabilities which have been extremely difficult to achieve before.

Accel: From Mainframe to Serverless and Beyond

Accel

By PING LI, *Partner*

We view serverless computing as the current endpoint in an IT continuum that has been shifting since the days of the mainframe. We believe the driving force behind these shifts is a desire to make life easier for developers, who capitalize on better abstractions to work faster and build better applications. The advent of cloud computing, which began in earnest with the launch of AWS in 2006, kickstarted an era of unprecedented new tools for developers — from minimizing operational overhead with IaaS to bypassing operations altogether with serverless services like [AWS Lambda](#).



Ping Li – Partner, Accel

One of the major shifts during this time has been the rise of cloud-native architectures, exemplified by technologies such as containers and Kubernetes. They set the stage for widespread serverless computing by abstracting away server- and OS-level details, and simplifying the process of building and managing distributed applications. It's a continuum along which developers and operations teams become more accustomed to increased automation and abstraction, and more comfortable breaking applications into simple, easy-to-manage microservices, APIs, and functions.

The result is that developers are free to target the right tools for the right tasks and to easily build applications that span any number of different services. A single application today could consist of long-running containers managed by Kubernetes, event-driven serverless functions managed by Lambda, and API-based cloud services for artificial intelligence or big data processing. The connective tissue between all these different approaches — and a turn almost inconceivable even 15 years ago — is that developers are building these scalable, flexible and powerful applications without ever configuring a server.

“Looking forward...application development will continue to evolve toward a place of primarily serverless computing...”

Looking forward, it seems logical that application development will continue to evolve toward a place of primarily serverless computing and away from concepts such as RAM, storage, and CPU capacity. We also expect startup business models to develop as they’re able to further minimize IT costs and energy while targeting new markets and new end-user devices/platforms.

The internet of things will help drive this evolution, inspiring developers to continue pushing the boundaries of application architectures and capabilities. IoT is a particularly natural fit for serverless architectures given their contemporaneous maturation and the distributed and event-driven nature of data from sensors and other edge devices.

IoT itself is part of a broader shift from server-side to client-side computing that serverless platforms like Lambda and AWS Greengrass can enable. Truly capitalizing on opportunities such as IoT requires adding functionality to client devices to reduce latency and simplify both programming and application architectures.

Accel’s Serverless Investing Thesis

At Accel, one of the primary tenets of our investment thesis over the past decade is that the public cloud is commoditizing underlying infrastructure such as compute, storage and networking. Therefore, we seek startups delivering higher-value workflows that abstract away those layers so developers can build better applications, faster. In short: Our focus is less on improving infrastructure or operations (although that is a result of technologies like Kubernetes), and more about enabling developers to focus on application delivery.

“Our focus is...about enabling developers to focus on application delivery.”

When it comes to developer tools and other “enterprise” IT plays, we want to help fix the impedance mismatch between how quickly developers can build new features and fix bugs, and how quickly that code makes it into production. Movements like agile development and continuous integration/delivery have taken companies so far, but we look at serverless computing, container orchestration, API services and other abstractions as the bridge between where we’re at today and where we want to be within the next few years.

And outside of these IT-focused, and largely business-to-business startups, we look for startups in all spaces that have figured out how to take advantage of containers, serverless platforms, and API-based services to build and improve their applications more quickly and effectively. Startups that don't

leverage these tools will move too slowly and inefficiently, leaving them at risk of being disrupted by faster-moving competitors or unable to react to opportunities/threats with enough haste. While we don't typically evaluate application architectures for potential portfolio companies, we do consider whether they are leveraging abstractions and tools that allow them to focus on business innovation rather than on managing IT.

“Startups that don’t leverage these tools will move too slowly and inefficiently, leaving them at risk of being disrupted...”

Below, we’ll look at how two innovative startups in our portfolio are leveraging serverless technologies to innovate faster and better serve their customers.

Accel Portfolio Companies Spotlight:



The ability to use technology to solve business problems is critical to any company. The world is changing from companies building static software solutions to developers delivering living and ever-evolving services to help businesses take full advantage of the technologies they’re bringing into their organizations. [Segment](#) is one such company.

Segment provides a service to help businesses collect customer data in a single hub for later use in analytics, marketing, and for other purposes. Today, more than 15,000 companies rely on Segment to process 80 billion end-user actions a month. Segment serves companies of all sizes across the world, including Fortune 500 enterprises like Intuit, Reuters, IBM, and Gap.

The advent of serverless patterns has helped fast-growing companies like Segment enjoy the efficiency of being able to go from code to production deployment in seconds. Managing its production infrastructure is critical to building a strong company and to that end, Segment has been hyper-focused on how to best optimize their approach. The company has open sourced its approach to help other technology startups create a production-ready architecture on AWS, which is outlined [here](#).



[Trifacta’s](#) data wrangling solutions run on an innovative serverless architecture based on working sets of data, rich AI-assisted visual interfaces running in the browser, and a domain-specific language (DSL) that compiles the browser behavior into the APIs of serverless big data infrastructure. Trifacta is designed to provide data analysts with an immersive, AI-assisted visual user experience for assessing and preparing data for analysis. Trifacta achieves anyscale performance by separating (1) design-time user experience

from (2) the operational execution of jobs over big data. Both use serverless principles to minimize ops costs and achieve anyscale performance.

At design time, interactive exploration and the authoring of data transformation recipes happen in the user's browser. Serverless storage API calls fetch a working set of data (an entire file up to moderate size, or a representative sample of big data) and deliver it to a visual browser-based UX implemented with WebAssembly technology.

For operational execution, the user's visual interaction in the browser is captured in a reusable recipe for data transformation, which is represented internally in Trifacta's Wrangle DSL. When it is time to run the recipe on more data, Trifacta compiles the Wrangle code into the API of the cloud provider's serverless infrastructure for big data. At that point, the heavy lifting of executing the transformation logic is done via commodity serverless infrastructure. Even for the biggest of data sets, execution cost is on-demand, and the cloud provider can handle operational concerns.

Conclusion

Serverless approaches reduce or eliminate problems such as idle servers that drain a company's balance sheet without offering value, the cost of building and operating fleets of servers, and server software, that distract and detract from the business of creating differentiated customer value. These benefits are equally applicable to an early stage startup as they are to a large multi-national enterprise. Furthermore, serverless services such as AWS Lambda offer a pay-per-use pricing model that can enable new technology-consumption models that are matched perfectly to usage. We look forward to seeing the new startups and innovative technologies that will be built on serverless technologies, as well as the growth of the community building tools and platforms that further enable the adoption of serverless.

To learn more and read whitepapers on related topics, see [Serverless Computing and Applications](#).

Learn more about [Greylock Partners](#).

Learn more about [Madrona Venture Group](#).

Learn more about [Accel](#).

The Contributors

The following individuals and organizations contributed to this document:

- Aaron Kao, Amazon Web Services
- Abhinav Tiwari, Amazon Web Services
- Ajay Nair, Amazon Web Services
- George Huang, Amazon Web Services
- Larry Yu, Accel
- Stacey Ngo, Greylock Partners