

Cloud native: Better, faster, cheaper is driving serverless – how much compute will eventually be done this way?

Analysts - William Fellows, Jean Atelsek

Publication date: Thursday, February 6 2020

Introduction

The promise of serverless applications – in which reusable functions and triggers are assembled into software that works independently of the infrastructure that executes it – has attracted a wave of startups, open source projects and cloud providers because the potential benefits are too compelling to ignore: faster development, hands-off provisioning and dramatically lower costs.

Rather than having VMs sitting idly by, burning cash, compute resources and energy, these applications invoke compute resources only when needed, operating on a pay-per-use rather than pay-per-provision basis as virtualized hardware does. Although the technology for building and operating serverless applications at scale still has rough edges, partly due to the lack of an open standard, born-in-the-cloud companies are adopting a 'serverless-first' strategy for many applications given the favorable economics and speed of development that it makes possible.

The 451 Take

Cloud-native approaches to software and service design enable enterprises to act faster, more efficiently and at greater scale: enterprises can go faster with cloud and be more efficient with microservices. Re-platforming to cloud native is therefore an imperative for digital transformation strategies and this will sweep through the market over the next decade, much like the re-platforming to the internet and web in the 1990s and 2000s. The demand for new skills and expertise means incumbent vendors are building, buying and partnering for this, while the lack of available talent is becoming more of a constraint for both end users and vendors than the access to capital. For investors and enterprises alike, this market is crowded and confusing and this is why

clean and simple stories are likely to succeed. Serverless wins here and is the reason its adoption is skyrocketing.

Context

In the market, serverless is used interchangeably with FaaS – functions as a service. FaaS uses the serverless computing concept in which reusable functions and triggers are assembled into software that is executed without concern for or access to, and works independently of, the infrastructure that executes it. In the serverless model – more properly 'provisionless' – the service provider runs the server and dynamically manages the allocation of machine resources. Pricing is based on the actual amount of resources consumed by an application (the execution of the function, not any storage, databases or other attached services) rather than on pre-purchased units of capacity. The infrastructure is rendered invisible to the developer or consumer, improving developer velocity and shortening time to market.

Market

The serverless market has consolidated in the two years since Google's Knative was introduced, which enables FaaS to run on Kubernetes. The open source Knative (also developed by IBM, Red Hat, Pivotal and SAP) simplifies the use of serverless because many operational functions are taken care of by Kubernetes. Various companies developing native serverless offerings altered course and changed product/business plans consequently. Moreover, the concentration of serverless around the major vendors (AWS, Azure, Google, IBM, Oracle and VMware) has made independent approaches less viable. Platform9, for example, is not further developing its Fission FaaS. However, some see that having only big ships which are difficult to steer could be an opportunity – startup Nuweba, for example. The startup is taking on AWS Lambda with a FaaS offering architected from the kernel up for speed and security. It recently added GPU support and claims its technology can spin up containers more than 10 times faster than Lambda, making it a natural environment for compute-intensive AI workloads.

In the meantime, following the acquisition on December 30, 2019, of Pivotal by VMware, Pivotal Function Service (PFS) seems likely to be VMware's serverless offering going forward, although post-acquisition product line decisions will be made public over the next few months. Through a series of 2019 acquisitions in addition to VMware-developed offerings, VMware has a number of serverless frameworks, including PFS from Pivotal (based on the Riff open source project), and Kubeless from its Bitnami acquisition. Dispatch, a VMware-originated open source serverless project, is inactive. VMware has scaled back its contributions to OpenFaaS.

Given the huge explosion in the number of public cloud SKUs and services available and opportunities to customize servers, efforts that organize the buying experience around application needs rather than provider SKUs will be welcomed by customers. The year 2020 will see serverless begin moving out of point projects and proofs of concept and into overall IT strategy as enterprises seek to operationalize the benefits of speed, agility, visibility and management in microservice-based FaaS deployments.


Cross-cloud serverless

Serverless adoption is growing rapidly (see Figure 1 below) and is a core component of many digital transformation projects in the enterprise – some believe all compute will go serverless and 'serverless first' strategies will be a thing in 2020. Regardless whether it is some, most or all compute that goes serverless, it's going to be a big opportunity and there's a huge need for integration of serverless with both Kubernetes-based and legacy applications. Enterprises will want to benefit from the best features of each. This will be enabled by continued momentum for Knative plus an

ecosystem of companies providing ways to take functions/events (such as AWS Lambda, Google Functions and Azure Cloud Functions) and run them on any cloud, or on-premises environments, using Kubernetes and containers as a portability layer. The question is whether there is room or a platform for alternative approaches such as a new 'serverless' programming language like Xqizit's Ecstasy.

Companies offering cross-cloud serverless capabilities include Serverless Inc, TriggerMesh, NowFloats and Kitsune. There are also a number of point solutions covering various aspects of the serverless application lifecycle, notably New Relic and Datadog, which offer multi-cloud metrics and monitoring, as well as IOpipe, Epsagon, Stackery, Dashbird and Thundra.

Figure 1: Adoption of cloud-native technologies



Source: 451 Research's Voice of the Enterprise: DevOps Q1 2019 and Digital Pulse: Budgets and Outlook 2019

Bringing serverless to existing apps

In the serverless sector, the security vendors have cashed out first. Serverless security was a hot M&A theme in 2019 – PureSec went to Palo Alto Networks, Protego Labs went to CheckPoint and Immuno to Trend Micro. Integration or 'glue providers' in the application space could be the next 'must have' serverless capability: first, to support greenfield cloud-native applications across multiple serverless environments; and second, to support enterprises seeking to run existing/legacy applications on serverless without a re-write, such as legacy banking applications or modernizing SAP in place to run on Lambda. Companies such as MongoDB and AWS Aurora are extracting events from SAP and other on-premises environments and will need a way to trigger these in the cloud and serverless. Typically, banks have, when selecting integration offerings, purchased source licenses, knowing that their software supplier will either go out of business or be acquired. Banks are already pushing Java applications into containers and running them on serverless environments inside their datacenters.

Cloud native in 451 Research

451 Research's definition: cloud-native technologies are used in the design or redesign of applications built to run in public, private and hybrid cloud infrastructure. Cloud-native technologies include containers, service meshes, microservices and serverless functions, all of which can be independently updated, controlled, scaled or reconfigured to deliver a coordinated application experience.

Cloud native: Better, faster, cheaper is driving serverless – how much compute will eventually be done this way?

More than 20 analysts from multiple channels are collaborating to create insight on cloud native. Anchoring our research with the 451 Research domain experts (across security, IoT, storage, monitoring, cloud) provides customers with a direct path to understanding the impact and opportunity which the re-platforming to cloud native presents in different sectors. The will enable users, vendors and investors to understand and compare the tools, technologies, services and skills which are available now, how they will develop and how they can be used to deliver digital transformation. We will examine best practices and challenges, key use cases and partner and investment opportunities in each sector. 451 Research's upcoming Serverless Market Monitor report includes analysis of 51 vendors, with individual estimates and forecasts.