VU Research Portal

Serverless is More

Van Eyk, Erwin; Toader, Lucian; Talluri, Sacheendra; Versluis, Laurens; Uta, Alexandru; Iosup, Alexandru

published in

IEEE Internet Computing
2018

DOI (link to publisher)

[10.1109/MIC.2018.053681358](https://doi.org/10.1109/MIC.2018.053681358)

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Van Eyk, E., Toader, L., Talluri, S., Versluis, L., Uta, A., & Iosup, A. (2018). Serverless is More: From PaaS to Present Cloud Computing. *IEEE Internet Computing*, 22(5), 8-17. [8481652].
<https://doi.org/10.1109/MIC.2018.053681358>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Serverless is More: From PaaS to Present Cloud Computing

Erwin van Eyk

AtLarge Research Team,
TU Delft,
The Netherlands

Lucian Toader

AtLarge Research Team,
VU Amsterdam,
The Netherlands

Sacheendra Talluri

AtLarge Research Team,
TU Delft,
The Netherlands

Laurens Versluis

AtLarge Research Team,
VU Amsterdam,
The Netherlands

Alexandru Uta

AtLarge Research Team,
VU Amsterdam,
The Netherlands

Alexandru Iosup

AtLarge Research Team,
TU Delft & VU Amsterdam,
The Netherlands

In the late-1950s, leasing time on an IBM 704 cost hundreds of dollars per minute. Today, cloud computing, that is, using IT as a service, on-demand and pay-per-use, is a widely used computing paradigm that offers large economies of scale. Born from a need to make platform as a service (PaaS) more accessible, fine-grained, and affordable, serverless computing has garnered interest from both industry and academia. This article aims to give an understanding of these early days of serverless computing: what it is, where it comes from, what is the current status of serverless technology, and what are its main obstacles and opportunities.

The 1950s saw the emergence of two technologies that are currently shaping the world: containerization in shipping and time-sharing in computing. By allowing shipping to become standardized and automated, the former gave rise to manufacturing and retail ecosystems, and ultimately to the economic phenomenon of globalization¹. By enabling multiple clients to share the same physical infrastructure, time-sharing gave rise to cloud computing and the modern digital ecosystems, which are key drivers for growth in knowledge-based societies².

Whereas few could afford the costs of time-sharing services and paid dearly for simple computer simulations in the late-1950s, today over 80% of companies use the hundreds of services accessible as cloud computing (source: Studies by European

Commission³ and Cloudability (2018)), along with many private individuals. Following with remarkable regularity the evolution observed in the history of containerization, cloud services have adapted to offer better fitting *containers* that require less time to *load* (boot) and provide higher *automation in handling* (orchestrating) containers on behalf of the client. *Serverless computing* promises more: to achieve full-automation in managing fine-grained containers. Already,

IT spending on serverless computing is expected to exceed \$8 billion per year, by 2021⁴. To understand what serverless is and what it can deliver, we trace the evolution of computing technology that has given rise to serverless computing, analyze the current status of serverless technology, and identify the main obstacles and opportunities we see in delivering on its promise.

What is serverless computing? We have proposed the following definition⁵:

Serverless Computing is a form of cloud computing which allows users to run event-driven and granularly billed applications, without having to address the operational logic.

This definition places serverless as a computing abstraction, partially overlapping with platform as a service (PaaS). With serverless, developers focus on high-level abstractions (e.g., functions, queries, events) and build applications that infrastructure operators map to concrete resources and supporting services. This effectively separates concerns, with developers focusing on the business logic, and on ways to interconnect elements of business logic into complex *workflows*; meanwhile, service-providers ensure that the serverless applications are *orchestrated*, that is, containerized, deployed, provisioned, and available on-demand, while only billing the user for the resources used. These separated roles have also emerged for physical containers, with manufacturers and retailers in the role of developers, and shipping lines in the role of service-providers.

Clients of serverless computing could use the Function-as-a-Service (FaaS) model, which we define⁵ as:

Function-as-a-Service (FaaS) is a form of serverless computing where the cloud provider manages the resources, lifecycle, and event-driven execution of user-provided functions.

With FaaS, users provide small, stateless functions to the cloud provider, who manage all the operational aspects to run these functions. For example, consider the ExCamera⁶ application: ExCamera uses cloud functions and workflows to edit, transform, and encode videos with low latency and cost. A majority of the tasks in these operations can be executed concurrently, allowing the application to improve its performance through parallelizing these tasks. To deploy ExCamera using the traditional infrastructure as a service (IaaS) model, a user needs to spin up virtual machines (VMs), provision them, orchestrate the workflows, manage resources as needed, and manage the variety of dynamic issues (e.g., faults, inconsistencies). This requires considerable expertise and continuous effort in orchestrating ExCamera, and yet results in significant amounts of underutilized but paid-for resources. Instead, by leveraging serverless computing, ExCamera defers the operational complexity to the cloud provider, using FaaS to manage the operational lifecycle of the individual video tasks.

However promising, serverless computing is still an emerging technology. Understanding how applications such as ExCamera can leverage it requires finding answers to questions such as: *What are the computer technologies underlying serverless? What is the status of the current serverless technology? And, what can we expect from the field of the serverless computing in the foreseeable future?*

We address these questions with a threefold contribution: (1) We identify the concepts leading to serverless computing, with deep historical roots in concrete and abstract innovations in computer science (see Section 2); (2) We analyze the current state of the technology and the complex technological ecosystems it consists of (see Section 3); (3) We identify and analyze important obstacles and opportunities (see Section 4) in this emerging field that we – as a community – need to address to make the serverless promise a reality. We conclude with a forewarning question: can we reproduce the successes and avoid the downsides of physical containerization?

THE LONG ROAD TO SERVERLESS

In this section, we analyze the evolution of computer technology that led to serverless computing – going back to the 1960s. All the breakthroughs indicate that *serverless computing would not have been possible a decade ago*, when it would have missed enabling technologies such as the distinction between IaaS and PaaS (standardized by NIST), fine-grained containerization (e.g.,

Docker), and even a diverse set of applications³. In Figure 1, we distinguish six main dimensions of these critical breakthroughs that together led to the emergence of serverless.

Figure 1: A history of computer science concepts leading to serverless computing.

Containerized Resources

Complementary to time-sharing, virtualization abstracts away the physical machine to reduce the operational effort, and to allow the same physical resources to be shared across multiple applications and users (multiplexing). Although associated in recent memory with VMWare's ESX technology (2001), virtualization was invented much earlier, and used in production in late-1960s IBM mainframes. Virtualization was finally conceptualized, nearly a decade later⁷. With the emergence of the Internet (1990s), these early concepts of virtualization were introduced online to enable shared hosting through virtual private servers. Soon after the release of ESX, cloud computing emerged, making virtual resources available over the Internet.

Like their physical counterparts, (digital) containers protect their content from external abuse. Containers do this by adding a layer of abstraction over the resources provided by the system. FreeBSD added jails – independent partitions of the system with their own directory subtree, hostname, IP address, and set of users. Linux followed with cgroups (2006), a mechanism to group processes and to assign each group separate resources. The Linux Containers project (LXC, 2008) bundled cgroups and kernel namespaces, along with better tooling. Built upon LXC, Docker (2013) offered convenient *container* orchestration, fostering an entire ecosystem based on digital containers.

Serverless computing is the latest result of this long-term process of defining virtualization abstractions, to eliminate concerns related to server provisioning and management. Although it exposes abstract resources to the user (e.g. functions), these are mapped to concrete resources (e.g. containers), continuing the transition from "bare metal" to "bare code".

Code as Functions

The ability to execute arbitrary "cloud functions" is essential to serverless computing.

Technology has emerged and reemerged often for running domain- and context-specific remote functions. We can trace this concept to as early as 1968, when, with IBM's Customer Information Control System (CICS), users were able to associate user-provided programs to transactions. The remote procedure call (RPC, specification in 1976, implementation in 1984) enabled the invocation of arbitrary procedures located in remote systems, over a communication network. Derived from RPC, stored procedures for databases (1980s) and Common Gateway Interface (CGI) scripts for web servers (1990s) aimed to bring support for executing functions to specific domains. Google App Engine – with other PaaS platforms following its example – started allowing users to asynchronously execute *arbitrary tasks* in the background.

In contrast to these context-specific implementations, serverless computing aims to provide a full abstraction for arbitrary, event-driven execution of *generic* functions.

Naming and Discovery

Managing and invoking services, including functions, depends on being able to name and discover services. Derived from a long line of technological innovations, current approaches follow the path-breaking concepts of Lightweight Directory Access Protocol (LDAP, 1993) and Uniform Resource Identifier (URI, 1994). LDAP uses naming and properties and enables distributed directory services over TCP/IP. URI provides unique identifiers for resources, encoded as character-strings.

Serverless extends naming and discovery with function versioning and aliases (e.g., offered by AWS). With versioning, it is possible to work with different immutable versions of a function simultaneously. Aliases are mutable pointers to a version and can be used to transition a version from one stage to another (e.g., from development to production) without changing the deployed application.

Functions as Computation

Serverless computing relies on the concept of *function as computation*, which stems from a long tradition of ever-higher-level abstractions and specialization in computer science.

Functional programming⁸ departed from procedural programs, to allow the developer to manage abstract data-types and control-flows, instead of the concrete details of memory and processors. The application of object-oriented principles to distributed systems lead to the creation of DCOM, CORBA, and OSF (1990s). In the 2000s, we climbed up the specialization ladder by contextualizing and interconnecting services, e.g., through the Service Oriented Architecture⁹ (SOA) and REST-based architectures.

These previous developments gradually led to *microservices*: self-contained applications providing specific functions over well-defined protocols¹⁰. Continuing this trend, serverless is development effectively a hyperspecialization of services.

Execution Flows

Serverless computing depends on the ability to coordinate execution flows.

Concurrency¹¹ has been an early and vital model for the evolution of computing allowing multiple processes to make progress at the same time, while remaining under the developer's control.

This model has many applications and many other models are rooted in concurrency, including generalized processes, threads, and actors¹².

Over the past two decades, we have moved towards a declarative form of expressing concurrency. Workflows declare the structure of applications, leaving the concrete execution and synchronization of workflow-tasks to the runtime system; this model has a multitude of applications¹³ and underlies our view of serverless computing.

Events to Trigger Functions

The first computer programs were synchronous, carefully crafted to follow a particular code-path. This model made programs difficult to create and modify, and less robust to changing conditions. Soon, *event sourcing* addressed the need to record, order, and respond to requests for state-changes.

With the proliferation of high-level languages and advanced operating systems, the concept of linking disparate computation together with special communication constructs took hold – with device drivers being early examples of this event-driven programming. With the rise of the Internet, event-driven distributed systems became widely used – with events mapped intuitively to the asynchrony of real-world networks.

Indeed, in modern systems, event-based protocols allow systems within an ecosystem to communicate without excessive dependence on the implementation details of each individual system. Due to its highly networked nature, serverless computing is apt to leverage this idea, through well-defined event-protocols and ways to manage events, for example, by using message queues.

SERVERLESS NOW

In this section, we discuss how the ExCamera application can be developed to use serverless computing. We explain the current and emerging technological ecosystem for serverless, and detail the expected benefits from using serverless: better resource management, scaling, and more insight and control.

The State of Serverless Technology

To execute parts of the workloads using small serverless functions, ExCamera parallelizes video-transcoding using AWS Lambda, which is one of the many FaaS platforms¹⁴. These various platforms differ in focus, target domain, assumed model and architectural decisions. Next to the closed-source FaaS platforms, addressing the lack of insight and vendor lock-in, several open-source platforms have emerged, including Apache OpenWhisk, Fission, and OpenLambda¹⁵.

To address the complexity of working with many different FaaS platforms and their incompatible APIs, the community has focused on informal standardization. Serverless frameworks, e.g., Apex or the Serverless Framework provide a common programming model that enable easier, platform-agnostic development, along with better interoperability of functions.

Much serverless tooling already exists to allow developers to defer non-essential tasks, sparking the emergence of a *diverse ecosystem* of (serverless) services, from serverless databases to monitoring to security. For example, workflow engines, such as Azure Logic Apps, Fission Workflows, and PyWren¹⁶, abstract away the complexity of networking in the compositions of higher-order functions and services.

Figure 2: A case for serverless computing: higher utilization, finer granularity, and more detailed control than with container-based, or self-hosted computing.

Benefits of Current Serverless Technology

Serverless computing promises more value than other cloud operations: equal or better performance while reducing the operational costs of applications. This has led industry – rather than academia – to drive the initial development and adoption of this paradigm. Figure 2 illustrates the main case for serverless computing.

Benefit 1. Improved resource management: In the traditional cloud model, the user is responsible for selecting and deploying the concrete resources. To avoid overburdening the user with options, the range of options is generally limited to large, multi-functional resource types (e.g., VMs or containers). Applications rarely fit these resources, and, to mitigate the overhead incurred by the large, general-purpose resources, applications are coarse-grained. As illustrated in Figure 2a, coarse-grained applications lead to inaccurate autoscaling decisions, causing severe under- or over-provisioning. In contrast, serverless computing means applications are fine-grained, which means the cloud provider can more closely match abstract resource-demand to actual system-resources.

Benefit 2. More insight and control: In the traditional model, the user is responsible for deploying, monitoring, and other operational tasks related to the lifecycle of coarse-grained applications (see Figure 2b), but many cloud-users do not have the necessary expertise. Moreover, the operators lack context, so they have to take autoscaling decisions without accurate profiling or insights from the deployed applications. With serverless, the increased responsibility for the operator gives more insight and control. Operators select the resources; deploy and provision resources; implement and control the monitoring of resource usage, workload intensity, and application behavior; and can auto-scale or migrate the application. They can profile and model the granular services comprising the serverless application (Figure 2c), offer this information to users, and improve the decisions made with these insights.

Benefit 3. Granular scaling: In the traditional model, applications consist of large, multi-functional VMs with multi-minute provisioning times. These VMs act as black-boxes, and thus are difficult to model and predict for operators. Although applications are typically bottlenecked by only one of the resources in one part of the application, the operators can only scale the entire application to resolve the bottleneck. Eliminating some of these issues is possible, but requires the user to re-architect the application, typically as microservices. The effectiveness depends highly on the user's expertise¹⁰ – most cannot benefit. With serverless, the operator can better scale the individual, granular services or functions, using deep insights. The contrast between Figures 2b and 2c illustrates this situation.

Other benefits: Other reasons promote the adoption of serverless computing: the shift from capital expenses to operational expenses more accurately aligns the costs to the actual business processes; the independent services allow teams to choose the right tools, and dependencies for a

use case without impacting other parts of the system and organization; and, the high-level abstraction allows software developers to iterate on these distributed systems faster while limiting the need for extensive expertise of distributed systems.

PERSPECTIVES ON SERVERLESS

Although serverless computing already offers many benefits (see previous section), many obstacles could inhibit further adoption. In joint work with the SPEC RG Cloud Group¹⁷, we have identified over twenty detailed challenges and opportunities for serverless computing^{5, 18}. Here, we identify the top-five obstacles and opportunities arising from them (Table 1).

First, the *fine-granularity* for expressing computation adds significant overhead to the resource management and scheduling layers. To overcome this, we envision significant research efforts invested in co-scheduling and orchestration for workflows of functions. Moreover, from a user perspective, we need new tools for navigating the cost-performance trade-offs to explore the complexity of fine-grained pricing models.

Second, *data privacy* is of importance for the clients and non-trivial to offer by the providers, for example, ensuring full GDPR-compliance¹⁹. With the fine-grained nature, serverless computing allows for more enhanced access control, function-level auditing and provenance for seamless and efficient GDPR-compliance.

Third, in modern clouds *performance* suffers from significant variability due to resource contention, virtualization and congestion overheads; Issues which are – with its granular nature – only amplified in serverless computing. However, the increased insight and control over the operational lifecycle, provides cloud providers with opportunities to minimize these performance issues by being able to more accurately monitor, profile, and schedule these fine-grained services.

Fourth, *data-intensive* applications are not naturally expressed in the – stateless – FaaS paradigm. We envision for the future the design and implementation of fine-grained, data-centric, serverless programming models. One promising research direction is investigating distributed promises in serverless environments.

Finally, the *API jungle* generated by the fast-evolving serverless APIs, frameworks, and libraries represents an important obstacle for software lifecycle management, and service discovery and brokering. To overcome this, significant effort must be invested in multi-cloud API standardization, interoperability, and portability to avoid lock-in and to enable seamless service discovery.

Obstacle	Opportunity
Fine-granularity and cost	Non-trivial resource management, workflows of functions, orchestration, fine-grained "pay-per-use" pricing, optimizing cost-performance trade-offs.
Data Privacy	Fine-grained access control and function-level auditing and provenance, full GDPR compliance.
Performance	Fine-grained scheduling and resource management, new performance models and fairness mechanisms that help reduce resource contention and performance variability.
Data-intensive applications	Fine-grained data-centric programming models.
API jungle	Service discovery and brokering, fine-grained software life-cycle management, standardized multi-cloud APIs, interoperability, portability, multi-modal services.

Table 1: Obstacles and opportunities for serverless computing.

CONCLUSION

Serverless computing is a promising technology, with a burgeoning market already formed around it. By analyzing the computer technology leading to it, we conclude that this model could not have appeared even a decade ago. Instead, it is the result of many incremental advances, spanning diverse domains: from the increasingly more granular resource abstractions, to the emergence of abundant amounts of resources available near-instantly, to the reduction of costs and complexity of distributed applications.

Current serverless technology offers its customers fine billing granularity, detailed insight and control, and the affordable ability to run arbitrary functions on-demand. However, this technology has not been demonstrated beyond selected, convenient applications. We identify several obstacles and opportunities, and argue that industry and academia must work together. Can we make serverless computing available for many, without the drawbacks of the technology and processes underlying physical containerization?

ACKNOWLEDGMENTS

This work is supported by the Dutch projects Vidi MagnaData, by the Dutch Commit and the Commit project Commissioner, and by generous donations from Oracle Labs, USA.

REFERENCES

1. M. Levinson. *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton University Press, 2016. Second Edition.
2. <http://business.nasdaq.com/marketsite/2017/Cloud-Computing-Industry-Report-and-Investment-Case.html>
3. European Commission. Uptake of Cloud in Europe. Digital Agenda for Europe report. Publications Office of the European Union, Luxembourg., Sep 2014.
4. <https://www.marketsandmarkets.com/Market-Reports/function-as-a-service-market-127202409.html>
5. E. van Eyk, A. Iosup, S. Seif, and M. Thommes. The SPEC cloud group's research vision on FaaS and serverless architectures. In *Proceedings of the 2nd International Workshop on Serverless Computing (WoSC '17)*. ACM, New York, NY, USA, 1-4. DOI: <https://doi.org/10.1145/3154847.3154848>.
6. S. Fouladi, R. S. Wahby, B. Shacklett, K. Balasubramaniam, W. Zeng, R. Bhalerao, A. Sivaraman, G. Porter, and K. Winstein. Encoding, fast and slow: Low-latency video processing using thousands of tiny threads. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation (NSDI'17)*. USENIX Association, Berkeley, CA, USA, 363-376.
7. G. J. Popek and R. P. Goldberg. Formal requirements for virtualizable third generation architectures. *Commun. ACM*, 17(7):412–421, 1974.
8. J. McCarthy. Recursive functions of symbolic expressions and their computation by machine, part I. *Commun. ACM*, 3(4):184–195, 1960.
9. N. Josuttis. *SOA in Practice: The Art of Distributed System Design*, chapter 1, page 7. O'Reilly Media, Inc., 2007.
10. R. Heinrich, A. van Hoorn, H. Knoche, F. Li, L. E. Lwakatare, C. Pahl, S. Schulte, and J. Wettinger. Performance engineering for microservices: Research challenges and directions. In *ACM/SPEC ICPE Workshops 2017*, pages 223–226, 2017.
11. E. Dijkstra. *Cooperating Sequential Processes*. Department of Mathematics, Eindhoven Technological University, 1965.

12. C. Hewitt, P. B. Bishop, and R. Steiger. A universal modular ACTOR formalism for artificial intelligence. In Proceedings of the 3rd international joint conference on Artificial intelligence (IJCAI'73). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 235-245.
13. N. Russell, W. M. P. van der Aalst, and A. H. M. ter Hofstede. Workflow Patterns: The Definitive Guide. MIT Press, 2016.
14. Survey of the CNCF serverless WG: <https://github.com/cncf/wg-serverless>
15. S. Hendrickson, S. Sturdevant, T. Harter, V. Venkataramani, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau. Serverless computation with openLambda. In Proceedings of the 8th USENIX Conference on Hot Topics in Cloud Computing (HotCloud'16). USENIX Association, Berkeley, CA, USA, 33-39.
16. E. Jonas, Q. Pu, S. Venkataraman, I. Stoica, and B. Recht. Occupy the cloud: distributed computing for the 99%. In Proceedings of the 2017 Symposium on Cloud Computing (SoCC '17). ACM, New York, NY, USA, 445-451. DOI: <https://doi.org/10.1145/3127479.3128601>.
17. <https://research.spec.org/working-groups/rg-cloud.html>
18. E. V. Eyk, A. Iosup, C. L. Abad, J. Grohmann, and S. Eismann. A SPEC RG cloud group's vision on the performance challenges of FaaS cloud architectures. In ICPE, pages 21–24, 2018.
19. <https://www.eugdpr.org>

ABOUT THE AUTHORS

Erwin van Eyk is an M.Sc. student at the Delft University of Technology, the Netherlands, where he works on serverless function and workflow scheduling. He is leading the Serverless working group of the SPEC RG Cloud Group focussing on performance evaluation and comparison of FaaS platforms. Contact him at E.vanEyk@atlargo-research.com or <https://erwinvaneyk.nl>.

Lucian Toader is an M.Sc. student at Vrije Universiteit Amsterdam, the Netherlands, where he studies modern distributed systems. His work on massivizing computer systems led him to serverless. Contact him at L.Toader@atlargo-research.com.

Sacheendra Talluri is an M.Sc. student at the Delft University of Technology, the Netherlands. In the spring of 2018, he was a research intern at big data company Databricks, working on resource management and scheduling across the memory-storage stack. Contact him at S.Talluri@atlargo-research.com.

Laurens Versluis is an Ph.D. student at Vrije Universiteit Amsterdam, the Netherlands, where he studies modern distributed systems. His work on massivizing computer systems focuses on resource management and scheduling, with applications in cloud computing. Contact him at L.F.D.Versluis@vu.nl.

Dr. Alexandru Uță is a post-doctoral researcher at Vrije Universiteit Amsterdam, the Netherlands, where he studies modern distributed systems. His work on massivizing computer systems focuses on resource management and scheduling, with applications in cloud computing and big data. Contact him at A.Uta@vu.nl.

Prof.dr.ir. Alexandru Iosup is a tenured Full Professor and University Research Chair at the Vrije Universiteit Amsterdam, the Netherlands, where he leads the Massivizing Computer Systems group. He is also Associate Professor with the Distributed Systems group at TU Delft, the Netherlands, where he received his Ph.D. in 2009. His work has received numerous awards, including the Netherlands ICT-Researcher of the Year (2016), Netherlands Teacher of the Year (2015), and several SPEC SPECTacular community-awards (the last in 2017). He is a member of the Young Academy of the Royal Academy of Arts and Sciences of the Netherlands. He is elected Chair of the SPEC Research Cloud Group. In his spare time, he contributes to training legal refugees in the Netherlands. You can contact Alexandru by email [A.Iosup@vu.nl], visiting (check <http://atlargo.science>), or via Twitter [@AIosup].