



## Serverless Computing – Architectural Considerations & Principles

January 2018



# Serverless Computing – Architectural Considerations & Principles

## Introduction

The last decade has seen tremendous innovation in the enterprise IT space. The primary focus of this technological innovation has been enabling businesses agility, improve resiliency and drive cost efficiencies.

The revolution in the server computing space over the past decade has allowed enterprise IT to develop and deploy scalable software without worrying about the underlying infrastructure. At the same time, cost of technology operations

have also gone down significantly while improving time-to-market for enterprises.

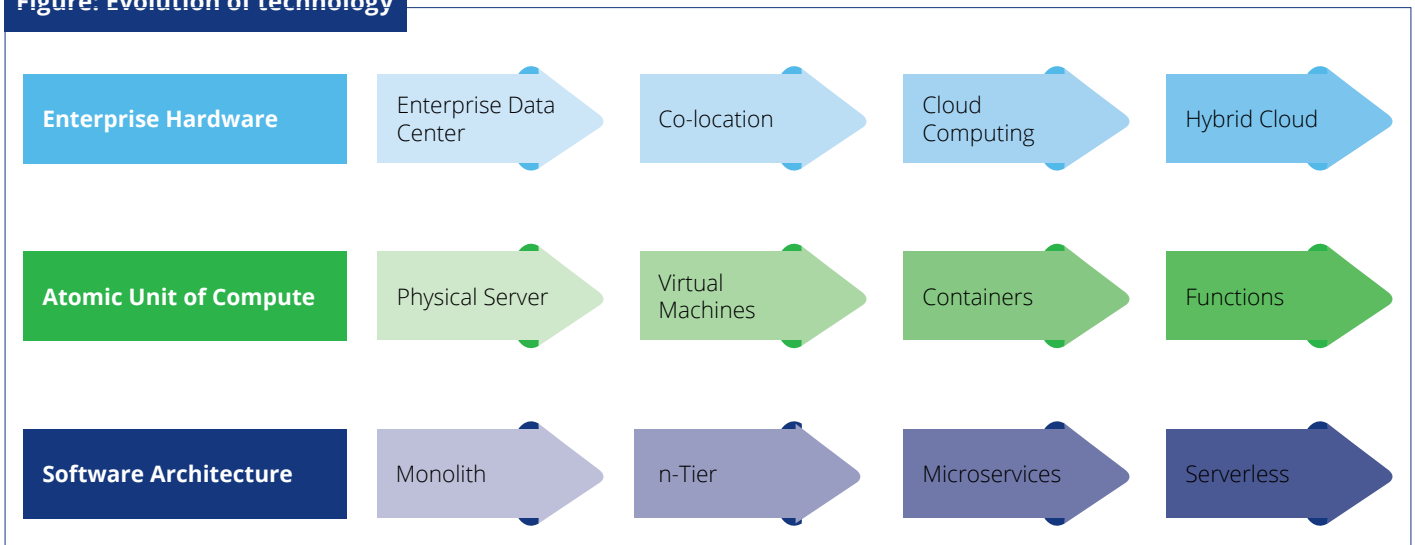
Server computing is now evolving towards even smaller units of scale – from virtual machines to containers to serverless<sup>1</sup>. The technology evolution has been depicted below.

This thought paper provides a vendor neutral summary of industry recommended principles and key considerations for architecting serverless systems.

### Serverless as a concept is not new

Amazon launched its cloud storage service (AWS S3) in 2006, which was also a 'serverless service' that provided unlimited storage at infinite scale without having to maintain servers. However, the differentiation now is that the compute services are becoming 'serverless' thus defining a new paradigm in software architecture.

Figure: Evolution of technology



<sup>1</sup> Evolution of Server Computing: VMs to Containers to Serverless — Which to Use When?, Gartner, June 2017


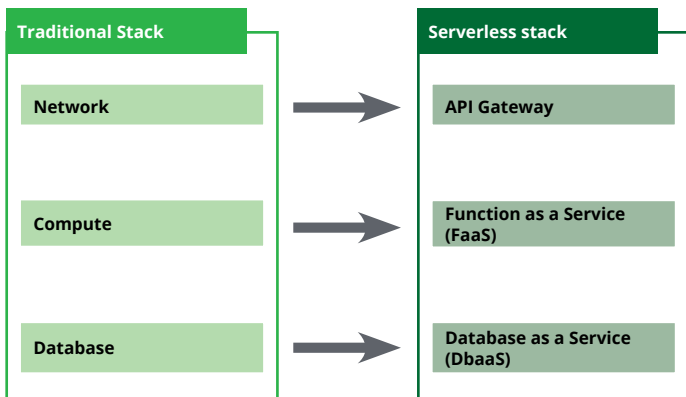


	IaaS	PaaS	FaaS
Unit of deployment	Operating System	Applications	Functions
Provides	Virtual machines packaged with operating systems	Application development platform	Execute code (with business logic) on-demand
Abstracts	Physical hardware	Operating system & middleware	Programming runtime

and Google Functions (launched in 2016 – alpha release). IBM has also released its serverless service called OpenWhisk.

### Re-imagining the technology stack

As depicted below the traditional technology stack for service delivery can be re-imagined to fit the serverless stack across each layers of network, compute and database.



The three key core technology components of serverless computing stack includes the following:

01

**API Gateway:** The API Gateway acts as the communication layer between the frontend and the FaaS layer. It maps REST API endpoints with the respective functions that runs the business logic. With servers out of equation there is no need for deploying and manage load balancers also in this model.

02

**Functions or Function as a Service (FaaS):** This is the layer that executes specific business logic (or code) with the cloud provider providing the level of abstraction in terms of executing the business logic.

03

**Backend as a Service (BaaS):** This is essentially a cloud based distributed NoSQL database which essentially removes database administration overheads.

<sup>4</sup> Relational databases like SQL are not preferred in a serverless architecture because they have a limit on the number of database connections that can be opened simultaneously at a time which can result in scalability and performance challenges

## Benefits of serverless computing

Key attributes of serverless computing include the following:

01

**Fully managed service:** As iterated earlier, it is a fully managed service provided by the cloud service providers and the developers doesn't need to worry about the underlying infrastructure, operating system, middleware, language runtime and its management and dependencies.

02

**Supports event-driven approach:** Functions are triggered based on events. Different cloud services<sup>5</sup> and existing applications that supports a trigger mechanism can initiate and launch a function.

03

**Provides infinite scalability and built-in high availability:** Depending on the user traffic, functions scale horizontally in a completely automatic and elastic fashion which is managed by the cloud service provider.

04

**Less-Ops:** Serverless doesn't necessarily mean NoOps<sup>6</sup> for the service consumers; however, it can definitely mean 'Less-Ops' as operational tasks like debugging, testing, trouble-shooting etc. remains while the infrastructure management is fully outsourced to the cloud service provider.

05

**Pay for execution time:** In a serverless computing model service consumers pay only for the duration of execution of a function and the number of functions executed. When a function is not executed no charge is levied – thus eliminating any idle time. This is a significant benefit over cloud computing where users are charged on an hourly basis for running virtual machines.



<sup>5</sup> For example, AWS services like S3, DynamoDB, Kinesis Streams, SNS and SES etc. can trigger AWS Lambda function. HTTP triggers either from a REST API are also supported

<sup>6</sup> No operations is a concept which refers to an IT environment becoming automated to the extent that the underlying infrastructure is abstracted without need for a dedicated team to manage the same



**Industry incumbents can regain competitive advantage through this smallest form of computing**

Competitive advantage in today’s business landscape is determined by the quality of customer interaction and service, customized products and services and effectiveness

of internal business processes effectiveness. And all of these are driven by technology with architectural design and deployment of software being at the core.

Serverless computing gives industry incumbents the required computing model to compete with the high-tech new comers in their industry. In this technology approach, software applications can be broken down into individual functionalities or functions in the serverless computing parlay (i.e. a micro-services based architecture) that are portable, cost-efficient and most importantly not bound to a legacy infrastructure footprint. The separation of the application functionalities from the supporting infrastructure provides the greatest opportunity for enterprises for application modernization and remove all development constrains posed

by legacy infrastructure. In fact, **serverless can be the exit strategy for enterprises from their legacy IT footprint.**

**Key considerations for architecting serverless systems**

Serverless computing might seem a lot more attractive over other traditional infrastructure consumption and computational models. However, we at Deloitte believe that without the right software architecture it is not possible to realize the benefits.

**"We are a technology company",**  
 JPMorgan CFO Marianne Lake  
 (Source: Business Insider, 2016)  
 With tech start-ups disrupting almost every industry with technological innovations, industry incumbents are struggling with their legacy IT footprint to innovate at the required pace to keep up with the changing customer

Not all application or service can be delivered in a serverless model – we believe the future enterprise IT landscape will be a hybrid landscape.

In addition, not all application or service can be delivered in a serverless model – we believe the future enterprise IT landscape will be a hybrid landscape. Hence, to ensure that the optimum architecture is in place, we have outlined certain architectural considerations that needs to be kept in mind while architecting serverless systems.

Key architectural considerations for serverless are as follows:

**Functions are ephemeral i.e. persist for a certain period of time only:**

Functions are deployed in a container<sup>7</sup>. These containers are started when a specific event is triggered and remains active for a certain duration beyond which it shuts-down automatically and all resources within the container cease to exist thereafter<sup>8</sup>. Hence, applications with larger processing requirements may not be suitable candidates for serverless.

**Cold start:** Cold start happens when an inactive (i.e. cold) function is invoked for the first time. This is typically the time taken by the underlying container running the function to load the runtime along with the libraries and associated dependencies. Cold starts can increase the execution time significantly and affect application performance if not handled appropriately. One way of addressing this bottleneck is to ensure functions stay alive (i.e. hot) before execution.

**Functions provide limited ability to configure database indexes:**

Use of relational database in serverless architectures creates scalability and performance challenges due to limitations of the number of simultaneous connections that can be opened in relational databases. Instead NoSQL databases can be the preferred choice. In addition, there can be limitations with regard to configuring nodes and indexes.

01

**Functions are stateless:** This implies that any kind of resource within a function will not exist after the function ceases to exist. This is an important consideration and hence appropriate measures must be implemented. For example, storing the session state in a persistent database if the session needs to be preserved.

02

03

**Language support:** Currently different serverless services support different languages for development. For example, AWS Lambda supports Node.js, Python and Java; Azure Functions support C#, Javascript and (preview of F#, Python, Batch, PHP, PowerShell); and, Google Functions supports only Javascript. Hence, selection of the platform may dictate the choice of language for development.

04

05

06

**Functions don't allow file systems level access:** This essentially means features like reading attributes from configuration files or spilling over in-memory cache to disk is not supported. Hence, serverless may not be the right architectural model for applications requiring file system level access or operating system level access.

07

**Comes with built-in logging & monitoring mechanisms:** Serverless services today comes with their own in-built logging and monitoring mechanisms<sup>9</sup>. This can be an architectural and operational shift in case an enterprise has been using custom tools for similar purposes.

<sup>7</sup> Containerization involves splitting applications into small parts, or containers, all sharing the same underlying operating system

<sup>8</sup> For example, AWS Lambda and Azure Functions have maximum execution time (per request) of 300 seconds (i.e. 5 minutes)

<sup>9</sup> For example, AWS Lambda supports CloudWatch and Azure Functions support App Service monitoring



## Recommended serverless architectural principles

We believe designing the appropriate technology architecture, aligned to the business requirements is the foundation for capitalizing on the technology innovations to drive competitive advantage in business. This is the reason why companies like are Netflix, Amazon are successful. CIOs and Senior IT executives will need to realize this fact and the sooner they do its better.

Below are some of the industry recommended guiding principles or best practices that can be kept in mind while architecting serverless systems:

### Develop single-purpose functions that are stateless:

Since functions are stateless and persists for a limited duration only, it is recommended to write single-purpose codes for function. This limits the execution time of a function which has a direct impact on cost. In addition, single purpose codes are easier to test, deploy and release thus improving enterprise agility. Finally, even though statelessness may be perceived as a limitation, it provides infinite scalability to a platform to handle an increasing number of requests, which otherwise would not have been possible.

01

### Create thicker and powerful frontends:

Executing more complex functionality at the front-end especially through rich client-side application framework helps reduce cost by minimizing function calls and execution times. Completely decoupling back-end logic from the front-end while not compromising on security is one way of doing. This also allows more services to be accessed from front-end resulting in better application performance and richer user experience.

03

**Identify performance bottlenecks:** On-going measurement of performance bottlenecks in terms of identifying which functions are slowing down a particular service is critical to ensure optimal customer experience.

05

**Design push-based, event-driven patterns:** Designing push-based and event-driven architecture patterns where a chain of events propagate without any user input imparts scalability to an architecture.

02

### Incorporate appropriate security mechanism across the technology stack:

Appropriate security mechanisms must be incorporated at the API Gateway layer and also at the FaaS layer. These security mechanisms include features like access controls, authentication, identify and access management, encryption and establishing trust relationship etc.

04

### Leverage third party services:

Serverless being an emerging field existing enterprise tools for various services like logging, monitoring etc. may not be compatible. Choosing the right third party tools for executing the task at hand will be key for enterprises to ensure the benefits of serverless are utilized to the fullest.

06

<sup>10</sup>. Serverless: The future of cloud computing by Peter Sbarski, ServerlessConf, 2016

# Conclusion

With the right architectural considerations and due diligence, serverless may present industry incumbents with an exit strategy to move from legacy infrastructure to adopt public cloud models – this was not observed earlier. However, industry incumbents will need move away from the status-quo and embrace the change. A structured approach with a defined roadmap is required to move away from the current software architecture paradigms which is based on the

legacy monolithic model with the serverless paradigm where the focus isn't on infrastructure but on delivering the required business functionalities which in turn changes the economic model for IT service delivery.

There are concerns though with regard to vendor lock-in and adherence to industry specific compliance requirements in a serverless model. Like any emerging technology, Serverless will go through its cycle

before mainstream adoption happens. Serverless has the potential to change the economic model of IT consumption of enterprises leading to significant cost reductions associated primarily with IT support and maintenance (which can be a significant 50 – 60% of the total IT budget of enterprises), reduce time-to-market and foster innovation to support changing business requirements and provide an edge over competitors.





# Contacts



**Rakesh Barik**

Partner, Technology Consulting Leader  
Deloitte Touche Tohmatsu India LLP  
E-mail: [rakeshbarik@deloitte.com](mailto:rakeshbarik@deloitte.com)



**Anubrata Chakrabarti**

Partner, Consulting  
Technology Strategy & Architecture  
E-mail: [anchakrabarti@deloitte.com](mailto:anchakrabarti@deloitte.com)



**Ritesh Pal**

Manager, Consulting  
Technology Strategy & Architecture  
E-mail: [riteshpal@deloitte.com](mailto:riteshpal@deloitte.com)

# Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. Please see [www.deloitte.com/about](http://www.deloitte.com/about) for a more detailed description of DTTL and its member firms.

The information contained in this material is meant for internal purposes and use only among personnel of Deloitte Touche Tohmatsu Limited, its member firms, and their related entities (collectively, the “Deloitte Network”). The recipient is strictly prohibited from further circulation of this material. Any breach of this requirement may invite disciplinary action (which may include dismissal) and/or prosecution. None of the Deloitte Network shall be responsible for any loss whatsoever sustained by any person who relies on this material.