Prepare for the Next Disruptive Waves: Serverless & InterCloud

Abstract

Cloud computing is in its infancy and has the potential to provide greater opportunities for cloud providers, vendors, and customers. Cloud computing can run legacy monolithic workloads providing an easy on-ramp for most organizations. However, cloud computing uses a different architecture which offers much greater opportunity and value yet to be realized. On-demand serverless compute, Function as a Service (FaaS), combined with existing cloud services allows the cloud itself to become programmable. Specifically, cloud services are the components for building customer applications. First, developers will program cloud services for a given cloud provider. Next, cloud providers will open their services and allow intercloud communications yielding intercloud services. Ultimately, architects and programmers will design and program the intercloud using InterCloud as a Service (ICaaS).

2021 Predictions

- PaaS usage will exceed IaaS
- Serverless/Microservice/API architectures will standardize and evolve into InterCloud as a Service (ICaaS)
- Cloud/InterCloud Architecture role will supplant DevOps role (think LessOps)

Recommendations

Business needs should drive cloud service adoption
Manage risks for both cloud provider PaaS and vendor PaaS (vPaaS)
Anticipate and investigate Vendor Platform as a Service (vPass) options
Anticipate the reduction of customer IaaS, including use of container services
Architects should be trained and gain experience designing Serverless/FaaS
Budget resources for Serverless/FaaS projects
Anticipate and plan for the expansion of Serverless/FaaS to InterCloud
Data should be encrypted in-motion between cloud providers
InterCloud service selection requires effective risk and opportunity management

Analysis

The introduction of public clouds, modeled after the Service Oriented Architecture (SOA) pattern, was the most important information technology innovation since the adoption of the World Wide Web. Specifically, Infrastructure as a Service (IaaS) disrupted the data center model and created new platforms of services (e.g., SaaS) and capabilities (e.g., Big Data, Analytics, Machine Learning). Virtually every business and government has or will be affected by this change with both positive and negative consequences. Innovations are creating new businesses and reinventing old business models while governments are adapting and responding to demands to meet new levels of service. New opportunities, business survival, improved security, and monetary and time savings are all driving adoption at a rapid pace.

The first iteration of cloud adoption was moving web architecture workloads from on-premises data centers to public and private cloud providers. Many adopters realized both data center and reduced employee headcount cost savings. Some realized more benefits by taking advantage of the cloud's on-demand capabilities (e.g., elasticity and scalability). The next iteration was the convergence of mobile, social, and cloud computing. The birth of the IoT market started a transformation Gartner calls Digital Business, the "creation of new business designs by blurring the digital and physical worlds."

Today, many businesses are adopting a cloud first strategy but are running against similar challenges from their on-premises deployments: overhead associated with creating, deploying, managing, and securing the IaaS stack. IaaS has a shared responsibility model, with cloud providers taking complete ownership and responsibility for the physical data center, operations, and management of the stack up to and including the hypervisor and virtual machine which runs the workload. Customers are responsible for management and security from the client accessing the SaaS application to the operating system and applications providing them.

Managing this workload in an ever increasing on-demand world puts stress on the development, operations, and quality assurance teams to work faster and better. The cultural movement DevOps gained appeal with the goal of fostering collaboration between teams and automating processes (continuous delivery/integration). Ultimately, inventive ways to speed the delivery of IaaS.

Meanwhile, Application Containers (e.g., Docker) exploded on the scene. These containers leverage the Microservices architecture described by NIST as "small, stateless, loosely coupled and isolated processes built around capabilities as opposed to services. Microservices are independently deployable in Application Containers, use less resources and can be created, destroyed, started and stopped far faster than in SOA."

At the same time, visionaries are taking a 30,000 foot view of the cloud landscape and asking "Why are we designing cloud solutions using old architectural methods? Why are people managing software infrastructure instead of focusing on the solution? If we can manage infrastructure for Platform as a Service (PaaS), such as databases (e.g., DynamoDB) and Content Delivery Networks (CDN) (e.g., CloudFront), why can't we do the same with compute?"

Serverless

In 2014, AWS introduced a new managed service for compute called Lambda and the cloud world changed. Lambda, as its name implies, is an anonymous function which runs a snippet of code once and is triggered by an event (e.g., HTTPS request). Simply put, Lambda allows you to run code without provisioning or managing servers. Lambda is a component of a new architectural movement called Serverless (Microservice) Architecture. The abstraction of server management is effectively the removal of IaaS. This adoption of PaaS increases the level of responsibility for the cloud provider and decreases the responsibility of the customer. With PaaS, customers no longer manage and secure infrastructure. Benefits include no servers to manage, continuous scaling, less security to manage, and sub second metering. You pay only for the compute time not when services are idle.

With compute as a service, there are all the components necessary to fulfil the promise of SOA. Organizations can now build web, mobile, and IoT applications using the serverless model. Even multi-tier architectures can be serverless. We can now architect and program the cloud. Architecting and programming the cloud requires a new way of thinking, new patterns of behavior. For one, you don't have a local infrastructure - you can't test it on your PC or Mac or test it at your data center. The serverless model represents a complete cloud commitment and only those with a cloud first strategy will be prepared for the transition.

Vendor Platform as a Service (vPaas)

Cloud providers, especially AWS, have been encroaching on cloud infrastructure vendor solutions (e.g., database, security). With the coming reduction of customer-based IaaS, vendors are especially vulnerable. Therefore, many are creating their own PaaS model for their products. For example, database vendor Redis Labs offers Redis Cloud as a fully-managed cloud service.

Vendor-based IaaS will effectively become Vendor Platform as a Service (vPaaS). Cloud providers will not be able to keep up with the demand for new and innovative services. vPaaS will become a substantial growth market. vPaaS operates on IaaS so cloud providers will share in vendor's revenue stream. Customers must be aware that vendors are responsible for their IaaS stack, not the cloud provider.

With new platforms come new architectural patterns, best practices, and ultimately frameworks. New frameworks are budding such as the Serverless Framework: "The Serverless Framework is an application framework for building web, mobile and IoT applications powered by AWS Lambda, AWS API Gateway, and in the future other Function as a Service providers."

InterCloud

Why should applications be restricted to one cloud? Leading cloud providers recognize the AWS vision and are offering their own compute PaaS solution, Function as a Service (FaaS):

Microsoft Azure Functions
Google Cloud Functions
IBM Bluemix OpenWhisk

When visionaries take a 60,000 foot view of the cloud landscape they see isolated cloud networks operating independently. The next iteration of cloud innovation will be the integration of disparate cloud services functioning as an InterCloud. This rich cross section of cloud services will be programmable via InterCloud as a Service (ICaaS). With ICaaS, architects can select the right cloud provider service depending on a wide range of factors including capability, performance, price, location, and/or availability of the service to satisfy a specific business need.

What about data? InterCloud will allow movement of application logic to where the data is, to any cloud worldwide, while only paying for resources utilized. Alternatively, data consumed by users can be transported to the cloud provider and service of choice to reduce latency.

Database vendor MongoDB understands the benefits of running and communicating between multiple clouds as a service, seamlessly. On June 28th, MongoDB introduced Atlas, an elastic, secure, highly available, and scalable cloud service. It will first be released on AWS with planned releases for Azure and Google Cloud with the ability to replicate data between cloud providers.

Frameworks, such as Serverless, who are building extensible plugins for individual clouds will be well positioned to tackle the challenges with programming the InterCloud. InterCloud will allow access to the best cloud services located where and when needed while only paying for the resources used.

Wayne Scarano, AWS-SA, CISSP, CCSK, SABSA

Cloud/Cybersecurity Analyst

wscarano@sga.com