

Computing Resources: A Review on Resources used in the Cloud

Rutu Kiran,
Student, Electronics and Communication
R. V. College of Engineering,
Bengaluru, Karnataka

Dr. Prakash Biswagar,
Professor, Electronics and Communication
R. V. College of Engineering
Bengaluru, Karnataka

Abstract — Cloud computing has become an integral part of any organization that deals with organizing and managing data on a large scale. These organizations rely on cloud providers like GCP, Azure, AWS etc. for computing resources, storage, networking and AI/ML services. Computing power and its performance are one of the key factors for companies that want to migrate to the cloud. As such, cloud providers offer a variety of options in computing resources which have their pros and cons. Hence, it is up to the user to make an informed decision on what resource is best suited for the use case. This paper attempts to give a bird's eye view of the different types of computing resources offered by popular cloud providers, their pros and cons and an overview of their performance.

Keywords— *Virtual machines; cloud computing; containers; clusters; serverless computing*

I. INTRODUCTION

Cloud computing is a blanket term that is used to refer to any kind of hosted service like servers, storage, databases, networking, software, analytics and intelligence [1] that does not require active user management. These services are provided by organizations like GCP, Azure and AWS which can be customized based on the customers' needs. Based on the extent of control the user wants to have on these services, hosting models (IaaS, PaaS, SaaS or FaaS) and the costs incurred for using these services the decision can be made.

Computing resources, more commonly used by the customer, come in different forms and can be tailored to the user's needs. These include virtual machines, their instances which are generated by scale sets or autoscalers, managed Kubernetes cluster services that make use of container instances and serverless computing services.[2] Each of these options have a certain tradeoff between them and perform better for different use cases.

II. KEY CONCEPTS

Virtualization refers to the creation of versions of computer hardware or software resources like storage devices, network resources and computing resources 'virtually', essentially mimicking the behavior of an actual, physical computer. Resource virtualization technologies include Containers and Virtual Machines.[4]

A. Virtual Machines

Virtual machines use only the operating system or software, and the hardware is rented out. A general block diagram of the virtual machine is illustrated in Fig. 1[4]. This gives a general idea about how the virtual machine is coordinated with the host.

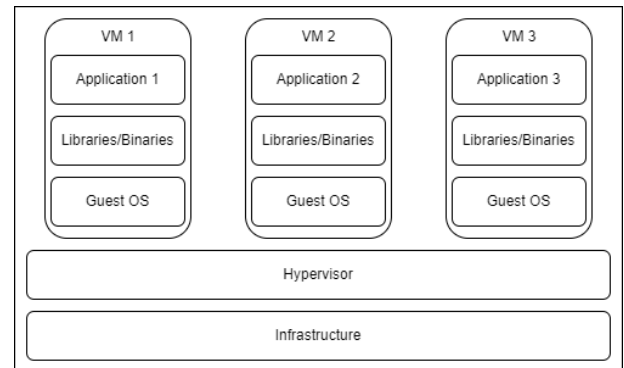


Fig. 1: Virtual Machine schematic

Numerous virtual machines also known as guest machines run on a physical machine also known as a host. Each guest machine is independent of the other and can function with different operating systems, creating an isolated environment and only using the physical machine for hardware resources. To create a more cost efficient and flexible in computing resources, these virtual machine instances are provided as services by cloud providers on the public cloud. These instances are usually used to deploy applications to suit the requirements of processing power one might need to run the application.

The most general use case is to test and run applications in a safe, isolated environment or run it on a different operating system. It is also used for testing operations that might need to run isolated from the host environment. They also support legacy applications which reduce the cost of migrating a system to have a new operating system. Some VMs also have disaster recovery and application provision options.[5]

There are also some disadvantages that need to be considered while running virtual machine instances. The iteration speed is one of them. The process of building and regenerating a virtual machine takes a lot of time. This makes it hard to do modifications. There is also the issue of cost from storage size. Since the OS is different from the host machine and the guest machine itself must store its OS, it might lead to a higher cost when renting out from a service provider.

B. Containers

Containers as opposed to virtual machines virtualize the machine at the OS level rather than the hardware level and the dependencies exist on higher stack levels as opposed to virtual machines. This makes them light weight compared to VMs. Containers share the OS kernel of the host, thereby reducing the memory requirement. They are convenient for sharing CPU, memory, storage and network resources [6] and offer applications abstraction from the running environment. A schematic of the abstraction is given below in Fig. 2 [4]

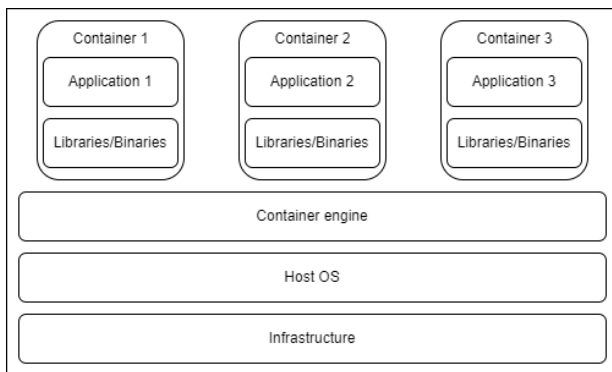


Fig. 2: Container schematic

Containers provide views that are logically isolated from other applications that are running on the host machine and are highly portable. It also creates a clear distinction and separability between the maintenance and management of the application as opposed to logic and dependencies making it easier to manage. These are generally used inside clusters which are managed by a container orchestration system like Kubernetes.

Since containers are lightweight and are at a higher level, modification and iterations are easier and faster. Most hosted containers also have a repository of containers that were previously made which can be easily downloaded and replicated on other systems, making them deploy easily and consistently. Their portability is their most attractive feature making them great for testing small modifications to applications since they can be easily created and destroyed without consuming much time.

Containers also come with certain disadvantages which can make the host system vulnerable. Pre-built containers with public images might pose a security risk for the host machine in case they have certain exploits which can break out of the container and are vulnerable to being taken over. [4] Hence containers are usually run in a protected environment mitigating these security risks.

C. Serverless computing model

Serverless computing is a service (Function as a Service, FaaS and Backend as a Service, BaaS) where on demand allocation of machine resources is allocated by a cloud provider. Serverless computing usually contains packages that are configured by simple parameters and computing resources get spun up on demand. The user uploads this package on to a serverless computing service and the rest is handled by the service offered by the cloud provider.

Serverless computing is similar to containers except that the provisioning, deploying and managing of this package is handled by the service provider. [7] Serverless, like containers have rapid development. They also encourage modularity since each package or function needs to perform a particular task with a limited number of inputs or parameters simplifying the operations that run on these machines independently.

There exist some limitations regarding size, latency and memory usage. This limits the control of the user on the program regarding the customizability in the aspect of scalability, monitoring and debugging due to limitations of the vendor. It can perform its best when it is running small simple applications that run as background processes alongside another containerized application.

III. USE CASES FOR COMPUTING RESOURCES

There is a recent trend seen in companies migrating their business to the cloud. A rise in companies not wanting to handle on premises physical resources is also seen. Except for handling extremely sensitive data, most other operations are handled by outsourcing computing resources provided by cloud providers through the public or private cloud. Depending on their requirements, the user can choose to use certain kinds of computing resources or a combination of these resources to fit their needs. Depending on the use case, certain types of resources work better than the other. To actively avoid incurring unnecessary cost, choosing the right type of computing resource becomes important.

Virtual machines tend to be best when used to run many applications together, to see their performance when they run simultaneously and check for isolation of different applications running in that system. It also is useful for running legacy applications on different operating systems. Any requirement that needs to have its applications logically isolated from one another but have their own OS kernel, virtual machines is the best choice for this.[5]

Containers are used to run more lightweight applications that utilize the host OS but include dependencies like binaries and libraries that are required to run. Owing to their faster booting times, they are used for testing and development of web applications. They are developed to maximize the resources used in the server, increasing the number of applications that can be deployed on the server. They also make deployment and delivery easy, making them a good resource for Agile development and DevOps testing.

Serverless model has computing that is usually done in short intervals whose results are saved into the storage. Since the service provider offers to handle the managing of the resources, there needs to be focus only on the development side by the user. These require much less memory than containers and only require it for processing application logic. Since this model is rapidly scalable and can be used for processes that need to run only for a short period of time, it is generally used for ephemeral applications like data analytics and CI/CD deployment processes.

IV. PERFORMANCE AND PRICING

There are several metrics that can be used to look at the performance of these computing resources. Some of them can include CPU load, efficiency, data throughput and latency.

[10] These metrics help compare the different cloud providers and their instance types. It would be an unfair comparison if it is done between virtual machines, containers and serverless models for the same use case as we have seen in the previous sections.

Several studies have attempted to make this comparison between Virtual machines and containers [11] and Serverless environments and virtual machines [9], [8]. While these studies try to give a chance for the resources to compete evenly, it is hard to judge which resource is better, since even though they all belong to the same class of computing resources their use cases are drastically different. There are several opinions in the tech community [5] that all these resources need not be competitive with each other and that they can be used to complement each other. This gives the user more flexibility in choosing the right resource for their application.

There are several services that help the customer gauge which resource might help them in their migration process and get their applications up and running in the cloud. Some cloud providers offer their customers cloud literacy programs to help make the transition to cloud easier. Cloud users can also estimate the incurring cost to deploy and use these resources using price calculators which can estimate the cost incurred based on the setup, expansion and requirements of the user. Several factors like the type of operating system, the location of the datacenter, the type of machine family the instance belongs to, the threads per core and the boot disk size also affect the price. The cloud provider provides these options to the user to have control over the computing resources used.

V. CONCLUSION

Migrating to the cloud enables easier scalability and flexibility to the user and enables them to have greater elasticity, self service provisioning and redundancy making their ecosystem and applications more robust. This also increases the focus on performance, usage and stability of their application. Understanding the security risk and making an informed decision on the cloud service and provider to use

will largely impact the performance of the application on the cloud. While this study tries to lay out the different computing resources that are available for use, its intention is to help the user take into consideration that even though some resources that look feasible, they may not be the right fit for their application. Using the resource that will best fit the user's requirements will help boost the capability of the application.

REFERENCES

- [1] Microsoft Azure, "What is cloud computing, A beginner's guide." <https://azure.microsoft.com/en-in/resources/cloud-computing-dictionary/what-is-cloud-computing/#cloud-computing-models> [Online; accessed 01-June-2022]
- [2] Microsoft Azure, "Azure Compute—Virtualisation and Scalability" <https://azure.microsoft.com/en-in/products/category/compute/> [Online; accessed 29-May-2022]
- [3] Amazon AWS, "Serverless Computing – Amazon Web Services" <https://aws.amazon.com/serverless/> [Online; accessed 01-June-2022]
- [4] I. Buchanan, Containers vs. virtual machines <https://www.atlassian.com/microservices/cloud-computing/containers-vs-vms> [Online; accessed 03-June-2022]
- [5] VMware, Virtual Machines <https://www.vmware.com/topics/glossary/content/virtual-machine.html> [Online; accessed 07-June-2022]
- [6] Google LLC, What are containers <https://cloud.google.com/learn/what-are-containers> [Online; accessed 05-June-2022]
- [7] Will Ezell, Virtual Machines vs Containers vs Serverless Computing: Everything You Need to Know <https://www.dotcms.com/blog/post/virtual-machines-vs-containers-vs-serverless-computing-everything-you-need-to-know> [Online; accessed 07-June-2022]
- [8] BBVA Labs, Economics of Serverless <https://www.bbva.com/en/economics-of-serverless/> [Online; accessed 12-June-2022]
- [9] Lee, H., Satyam, K., & Fox, G. (2018, July). Evaluation of production serverless computing environments. In 2018 IEEE 11th International Conference on Cloud Computing (CLOUD) (pp. 442-450). IEEE.
- [10] Ferreira, A. P., & Sinnott, R. (2019, December). A performance evaluation of containers running on managed kubernetes services. In 2019 IEEE International Conference on Cloud Computing Technology and Science (CloudCom) (pp. 199-208). IEEE.
- [11] Potdar, A. M., Narayan, D. G., Kengond, S., & Mulla, M. M. (2020). Performance evaluation of docker container and virtual machine. *Procedia Computer Science*, 171, 1419-1428.
- [12] Google LLC, Google Cloud Pricing Calculator, <https://cloud.google.com/products/calculator> [Online; accessed 15-June-2022]