GlobalLogic®

Cloud Sandboxes:

How to Train Your Engineers To Go Cloud-Native

Authored by:
Surbhi Nijhara and Nitin Unni

Contents

Executive Summary

> "Line-of-business leades everywhere are bypassing IT departments to get applications from the cloud (known as software as a service, or SaaS) and paying for them like they would a magazine subscription. And when the service is no longer required, they can cacel that subscription with no equipment left unused in the corner."
>
> ***Daryl Plummer, Managing Vice President and Distinguished Analyst at Gartner***

In today's world of software development and delivery, the public cloud is omnipresent. Each and every facet of the software development life cycle has dependencies on public cloud platforms. For IT and software product development organizations, the need to train and equip our workforce to understand, manage, and use cloud platforms effectively is greater than ever before.

While the ease of initial access to public cloud services is extremely simple and beneficial, its business model of pay-per-usage is tricky. Untrained hands can quickly and easily rack up significant costs. Moreover, having the necessary guardrails for security and resource access is also crucial in public cloud adoption and training. In this context, creating and managing cloud sandboxes is an interesting area that is emerging as a key component of the training experience.

A sandbox is a virtual environment in which software developers can test new codes without impacting live sites or programs. Sandboxes are often used as part of virtual training labs to train software developers in new skills by allowing them to practice without risk to actual products with the same virtual environment provisioned in a public cloud.

This report outlines the need for cloud sandboxes and some of the characteristics of this infrastructure. It highlights some of the best practices and archetypes for a well-run cloud sandbox from a practitioner perspective. It also identifies some of the future trends and possibilities in this space and calls to action different aspects IT organizations should consider when starting on this training journey.

The Need for Cloud Sandboxes

> "When I played in the sandbox, the cat kept covering me up."
>
> ***Rodney Dangerfield***

Cloud Sandboxes are used for various purposes by organizations that adopt cloud platform engineering and operations. With the foray of cloud computing in almost every technology space and the ever-growing spectrum of cloud services, companies are required to provide a new type of experience. Now, they are at the helm of providing an environment to required teams – one where users can learn the new offerings while prototyping possible solutions to business problems.

It can be argued that existing development and/or test environments can be used as environments to build and innovate. However, this is not a recommended practice as it is simply not always feasible. One reason is that the development and test environments must be specific to a given project or business problem. Another is that a stronger discipline is generally needed for the management to learn and build the prototype blocks. Therefore, the usage policy should be clearly defined.

Similarly, sandboxes are also used during engineers' training drills on cloud services. No training is complete without sufficient hands-on training in which a real cloud (and not a simulated cloud account) should be provisioned for the learners.

Sandbox cloud accounts, therefore, must have different access control strategies when compared to a software development lifecycle environment.

Our Cloud Sandbox Setup and Learnings

In this section, we try to distill our experience of running a DevOps and cloud enterprise-wide training while setting up cloud sandboxes on our own for the attendees.

We evaluated some of the available commercial offerings for cloud sandboxes, such as Lab Playgrounds from Cloudacademy and qwiklabs. However, we decided to set this up for the following reasons.

Our DevOps training course modules include several AWS services based on the current business needs, which fall into the below categories:

- Config Management and Infrastructure-as-Code (IaC)
- SDLC Automation
- Monitoring and Logging
- Policies and Standards automation
- Incident and Event Response Ecosystem
- High Availability, Fault Tolerance and Disaster Recovery
- Containers and COntainer Orchestration management

Refer to the AWS Cloud Sandbox figure below for a comprehensive list of services in each category.

Our research revealed that the required services are not offered in a unified catalog in the commercial sandboxes.

Our 3-month training curriculum included self-paced, online courses followed by instructor-led classroom training which included mini-projects based on real-use cases of our customers.

We wanted to give the same sandbox environment experience consistently throughout the learning journey.

Based on the learning journeys and various course curriculums in the training, we built an outline of the expected usage for the cloud environment and its services. We also went through a rough top-down cost modeling based on standard AWS pricing calculators to get an idea of the average spend. Our objective in creating this cloud sandbox was a judicious mix of flexibility with just enough governance.

Based on this, we created an exemplar design for the AWS Cloud sandbox which is outlined in the diagram below:

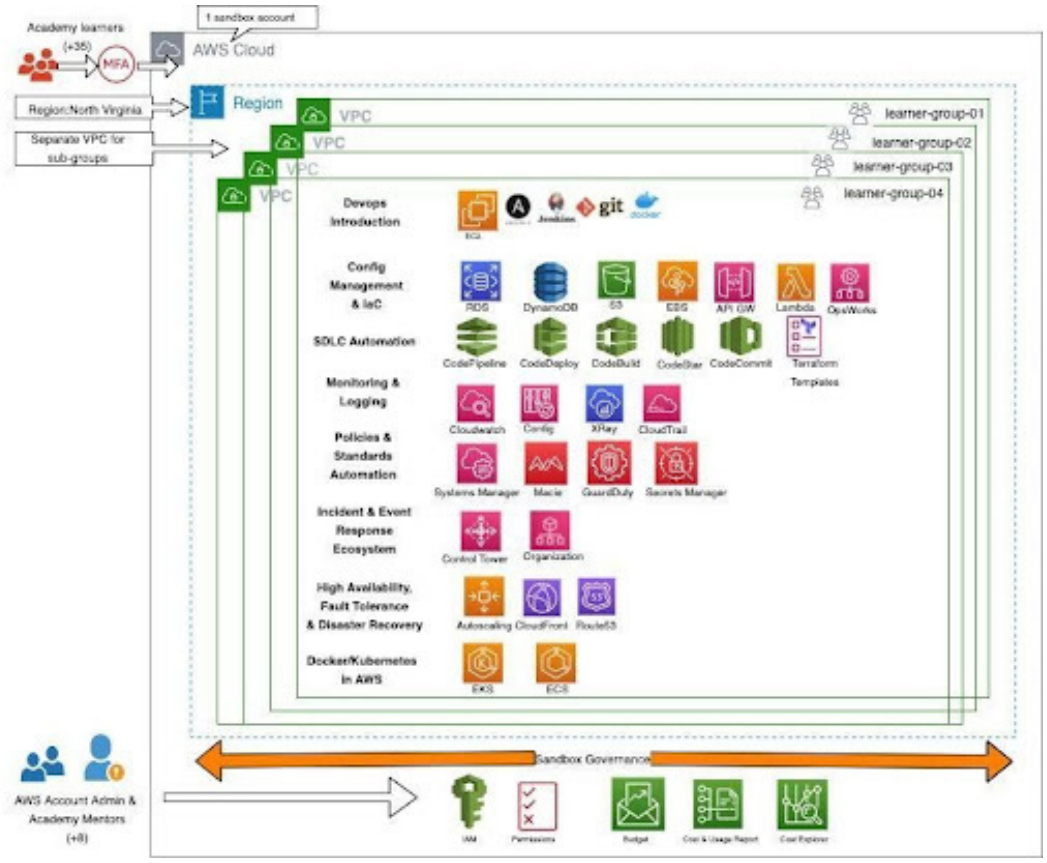

Figure 1: Exemplar Design Outline for the AWS Cloud Sandbox

For approximately 25 to 30 learners, a single learning sandbox account was created where 20+ governance controls were applied.

Key Policy	Key Policy
Access restricted to only One AWS Region:	Including this important restriction ensures that all learners are using the same regions and the instantiated resources are not scattered.
Choice of AWS Region	For sandbox purposes, we preferred to choose AWS North Virginia region as the costs in this region are the cheapest when compared to other regions.
Choice of EC2 Instance	It becomes imperative to choose which AWS instances should be used in the sandbox. The purpose of our sandbox was DevOps training, hence a general purpose type was sufficient as it is the most cost-effective.
Type of General Purpose Instance	Over the years, AWS has improved and increased the choices amongst the general purpose instances. AWS provides T2 and T3 families of general-purpose instances where T3 and T3a are the next generations. It is the more cost effectiveT3a which offers a further 10% in cost savings compared with other types.
Instance Class	Low configuration instance classes i.e micro, small, and medium were the only classes allowed.
Database Class	RDS offers Standard classes, Memory-optimized classes, and Burstable classes to choose from. For reasons just like EC2 instances, an RDS service allowed with burstable class is sufficient for use, hence only using a T3 micro, T3 small, or a T3 medium.
VPC Locking	In a single account per region, a maximum of 5 VPCs can exist. While a single VPC can be used by all participant learners, it was deemed a better strategy to create the maximum possible VPC while VPCs between different groups of users. A learner IAM user group was restricted to accessing a single specific VPC.
Single S3 Bucket:	A single S3 private bucket is a decent strategy to provide to the learners. All learners can use the same bucket by creating individual folders. This approach protects novice learners while creating public buckets that accidentally store sensitive information in the public buckets.
User Onboarding:	IAM users were created with passwords that need to change on the initial login. When setting up a policy that allows users to create and download their own AWS access keys, the AWS keys are never sent over a wire, nor are they shared on a drive. They are never seen by anyone except the learner himself. The AWS keys are by far the most critical asset to secure this strategy.

Key Policy	Key Policy
Multi-Factor Authentication (MFA):	Enable MFA for AWS console logins. This is a simple yet effective best practice that enhances the security of AWS accounts. It is a better strategy for providing an AWS account when enabling federated logins and adopting a single sign-in strategy.
Monthly CostBudget:	The sandbox cost needs to be accounted for in the overall planning and set up the learning platform for the employees. At the setup service/user level, customize the set up at user level (will need automation). Pre-configure actions that can trigger the implementation of the IAM or SCP policies, or stop the target EC2 or RDS running instances in your account. Create a total monthly cost budget for each AWS account you use.
Use Cost AllocationTags:	For a sandbox, the cost allocation tags strategy needs to be enhanced as the learner, i.e. IAM user is an imperative consumer of AWS services. Every service was enforced to be tagged with the learner name which helps to analyze and monitor costs in the IAM users category.
Usage Costs Monitoring:	Beginner-level DevOps learners can leave instances of EC2 and RDS in a running state even when they are not being used. There could be a large number of sizable files stored in S3 which are again not in use. To mitigate these high costs, Amazon CloudWatch alarms were created to monitor our estimated charges.
Restricted IAM Access:	Identity Access Management (IAM) is by far the most critical and sensitive cloud offering. It is imperative to provide the only required actions to learners who change their console passwords by letting them generate their AWS programmatic keys.
Pre-Create the Service and Service-Linked roles	Depending on the learning use case, the services require permissions to access other services. The IAM entity should provide permissions to create the service role, hence delegating the permissions on your behalf. For services that do not support service-link roles, a role should be pre-created and assigned to all the learners.

All of the above security and cost governance controls were applied using appropriate IAM policies and roles.

Our Analysis of Available Cloud Sandbox Offerings

As we highlighted earlier, we evaluated some of the available commercial offerings for cloud sandboxes like Lab Playgrounds from Cloudacademy as well as qwiklabs. This section covers a list of our detailed analysis and reasons why we finally decided to set this up on our own.

1. Available cloud vendor sandboxes are available along with specific learning subscriptions and courses. They are not available in either tailor-made or stand-alone capacity.

2. Sandboxes are available only with the course subscriptions. They are similar to black boxes for users, where there is less control and flexibility. In addition, they are time-bound, meaning they are available for a few minutes to an hour depending upon the course topic.

3. The sandbox users are bound to the vendor who is given specific steps while learning a module. There is no option to try alternative methods.

4. Needless to mention that the sandbox service is a costlier option. There is no visibility or data available to derive the actual costs on cloud resources usage.

5. There are certain sandbox providers like Microsoft, where the sandbox instance has a quota per day per user. This blocks the learner and directly impacts the learning period and motivation.

6. There is often a need to revise the topic learned via the previously attempted labs exercise. Commercial sandbox offerings do not allow multiple invocations of the completed lab which therefore restrain the learner from revising or revalidating the topic better.

Conclusions and Calls to Action

'DevOps begin at home'.

We followed this philosophy in our belief and practice while embarking on implementing the DevOps sandbox from scratch. For training purposes, we set an example by adopting the DevOps culture even during the learning phase.

Creating our own sandbox model as a service has let our enthusiastic learners explore DevOps capabilities while fitting in with our company's needs and the nature of our company's business. It also offered a good level of control and flexibility to accommodate shifts in the agile world.

Mobilizing ourselves to create our own sandboxes has not only given us the landscape of increasing cost efficiency benefits but also developed maintainability and improved reproducibility while risk-averse models can now be extended to include our business client products.

With this model setup, we are confident and in a strong position to adopt the DevOps sandbox with other public clouds like Azure and GCP platforms for the purpose of training our engineers to 'ops' in these clouds, too.

Happy Sandboxing!

About the Authors

Surbhi Nijhara is a Senior Solution Architect at GlobalLogic. She has a deep technical skill in platform architecture, infrastructure, and cloud computing and is a lead consultant in various customer advisories for thinking strategically about cloud solutions to business, product, and technical challenges in enterprise-grade digital transformations. She is also a chief mentor for the in-house Cloud and DevOps competency academy.

Nitin Unni is Director of Engineering and a domain enterprise architect at GlobalLogic. He has been involved in large digital transformations throughout his career and has engaged in various ROI analyses, technology evaluations, and solution blueprinting for enterprise customers. In his charter at GlobalLogic, he is responsible for evolving Cloud, DevOps, and DeepQA competencies and maturity in complex AD&D programs

References

AWS documentation: https://docs.aws.amazon.com/index.html

GlobalLogic®

GlobalLogic, a Hitachi Group Company, is a leader in digital product engineering. We help our clients design and build innovative products, platforms, and digital experiences for the modern world. By integrating our strategic design, complex engineering, and vertical industry expertise with Hitachi's Operating Technology and Information Technology capabilities, we help our clients imagine what's possible and accelerate their transition into tomorrow's digital businesses. Headquartered in Silicon Valley, GlobalLogic operates design studios and engineering centers around the world, extending our deep expertise to customers in the automotive, communications, financial services, healthcare & life sciences, media and entertainment, manufacturing, semiconductor, and technology industries.



www.globallogic.com