

Building Serverless Applications with Python

Develop fast, scalable, and cost-effective web applications that are always available



Packt>

www.packt.com

By Jalem Raj Rohit

Building Serverless Applications with Python

Develop fast, scalable, and cost-effective web applications that are always available

Jalem Raj Rohit



BIRMINGHAM - MUMBAI

Building Serverless Applications with Python

Copyright © 2018 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Commissioning Editor: Merint Mathew

Acquisition Editor: Sandeep Mishra

Content Development Editor: Rohit Kumar Singh

Technical Editor: Ruvika Rao

Copy Editor: Safis Editing

Project Coordinator: Vaidehi Sawant

Proofreader: Safis Editing

Indexer: Priyanka Dhadke

Graphics: Jason Monteiro

Production Coordinator: Arvindkumar Gupta

First published: April 2018

Production reference: 1190418

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-78728-867-6

www.packtpub.com



`mapt.io`

Mapt is an online digital library that gives you full access to over 5,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Mapt is fully searchable
- Copy and paste, print, and bookmark content

PacktPub.com

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details.

At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Contributors

About the author

Jalem Raj Rohit is an IIT Jodhpur graduate with a keen interest in recommender systems, machine learning, and serverless and distributed systems. Raj currently works as a senior consultant—data science—and NLP at Episource, before which he worked at Zomato and Kayako. He contributes to open source projects in Python, Go, and Julia. He also speaks at tech conferences about serverless engineering and machine learning.

About the reviewer

Sanjeev Jaiswal is a computer graduate from CUSAT with 9 years of industrial experience. He basically uses Perl, Python, AWS, and GNU/Linux for his day-to-day activities. He is currently working on projects involving penetration testing, source code review, security design and implementations in AWS, and cloud security projects.

He is learning DevSecOps and Security Automation currently as well. Sanjeev loves teaching engineering students and IT professionals. He has been teaching in his leisure time for the last 8 years.

Special thanks to my wife, Shalini Jaiswal, for her unconditional support, and my friends Ranjan, Ritesh, Mickey, Shankar, and Santosh for their care and support all the time. Thanks to the people at Packt for the project and the opportunity to learn good stuff from skilled professionals through reviewing the project.

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit authors.packtpub.com and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Table of Contents

Preface	1
Chapter 1: The Serverless Paradigm	5
Understanding serverless architectures	6
Understanding microservices	7
Serverless architectures don't have to be real-time only	9
Pros and cons of serverless	11
Summary	14
Chapter 2: Building a Serverless Application in AWS	15
Triggers in AWS Lambda	16
Lambda functions	21
Functions as containers	22
Configuring functions	24
Testing Lambda functions	34
Versioning Lambda functions	39
Creating deployment packages	44
Summary	49
Chapter 3: Setting Up Serverless Architectures	50
S3 trigger	51
SNS trigger	61
SQS trigger	74
CloudWatch trigger	84
Summary	91
Chapter 4: Deploying Serverless APIs	93
API methods and resources	94
Setting up integration	101
Deploying the Lambda function for API execution	111
Handling authentication and user controls	118
Summary	126
Chapter 5: Logging and Monitoring	127
Understanding CloudWatch	128
Understanding CloudTrail	138
Lambda's metrics in CloudWatch	145
Lambda's logs in CloudWatch	155
Logging statements in Lambda	160

Summary	165
Chapter 6: Scaling Up Serverless Architectures	166
Third-party orchestration tools	167
The creation and termination of servers	175
Security best practices	182
Identifying and handling difficulties in scaling	189
Summary	190
Chapter 7: Security in AWS Lambda	192
Understanding AWS Virtual Private Clouds (VPCs)	193
Understanding subnets in VPCs	200
Securing Lambda inside private subnets	206
Controlling access to Lambda functions	209
Using STS inside Lambda for secure session-based execution	210
Summary	210
Chapter 8: Deploying a Lambda Function with SAM	211
Introduction to SAM	212
CloudFormation for serverless services	215
Deploying with SAM	216
Understanding security in SAM	225
Summary	230
Chapter 9: Introduction to Microsoft Azure Functions	231
Introduction to Microsoft Azure Functions	232
Creating your first Azure Function	235
Understanding triggers	239
Understanding logging and monitoring in Azures Functions	246
Best practices for writing Azure Functions	250
Summary	252
Other Books You May Enjoy	253
Index	256

Preface

Serverless engineering is a new domain of engineering that allows developers to write code and deploy infrastructures without having to worry about maintaining servers. This book explains the concepts of serverless engineering with Python examples on cloud architectures.

Who this book is for

This book is for Python developers who would like to learn about serverless architectures in cloud-based platforms such as Azure and Amazon Web Services (AWS). Python programming knowledge is assumed.

What this book covers

Chapter 1, *The Serverless Paradigm*, introduces the reader to the concepts of microservices and serverless architectures, and provides a clear understanding of the pros and cons.

Chapter 2, *Building a Serverless Application in AWS*, covers AWS Lambda and explains the concepts, workings, and the components involved in the tool in detail. It also explains the nuances involved in security, user controls, and versioning code inside Lambda.

Chapter 3, *Setting Up Serverless Architectures*, goes into further detail about the various triggers in AWS Lambda and how they integrate with the functions. The reader will learn what the event structure of each trigger will look like and how to modify the Lambda function with respect to the type of trigger used.

Chapter 4, *Deploying Serverless APIs*, explores the AWS API Gateway and also teaches the reader how to build efficient, scalable serverless APIs using the API Gateway and Lambda. It goes on to teach the reader how to improve the API by adding authorization and how to set up user-level controls such as throttling of requests.

Chapter 5, *Logging and Monitoring*, presents the concept of logging and monitoring in serverless applications. This is mostly still an unsolved problem in this domain. This chapter guides the reader through setting up logging and monitoring in the AWS environment with Python via custom metrics and logging. This chapter also goes into the details of best practices when it comes to logging and monitoring Lambda functions in Python.

Chapter 6, *Scaling Up Serverless Architectures*, discusses the practice of scaling up serverless architectures for large workloads using several third-party tools. This chapter also teaches the reader how to handle security, logging, and monitoring using the available Python modules.

Chapter 7, *Security in AWS Lambda*, teaches readers to deploy secure serverless applications by leveraging the AWS security features available. This involves having strict controls on components that the application can access, on the users who can handle or access the application, and so on. This also involves understanding AWS virtual private clouds and subnets for an in-depth understanding of the security features and best practices you can follow in AWS Lambda.

Chapter 8, *Deploying a Lambda Function with SAM*, looks at how to deploy Lambda functions as infrastructure as code via the Serverless Application Model, which is a new way of writing and deploying Lambda functions that makes it easier to integrate with other IaaS services, such as CloudFormation.

Chapter 9, *Introduction to Microsoft Azure Functions*, familiarizes the reader with Microsoft Azure Functions, and explains how to configure and understand the components of the tool.

To get the most out of this book

The reader should be comfortable with the Python programming language. So, prior experience with it is expected. Prior experience with cloud-based systems will also be helpful.

Conventions used

There are a number of text conventions used throughout this book.

CodeInText: Indicates code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. Here is an example: "It is to be noted that the meta information should always be included for all SAM, which includes `AWSTemplateFormatVersion` and `Transform`. This would tell `CloudFormation` that the code you have written is an AWS SAM code and a serverless application."

A block of code is set as follows:

```
AWSTemplateFormatVersion: '2010-09-09'  
Transform: AWS::Serverless-2016-10-31
```

Bold: Indicates a new term, an important word, or words that you see onscreen. For example, words in menus or dialog boxes appear in the text like this. Here is an example: "For creating a function, you need to click on the orange **Create a function** button on the right."



Warnings or important notes appear like this.



Tips and tricks appear like this.

Get in touch

Feedback from our readers is always welcome.

General feedback: Email feedback@packtpub.com and mention the book title in the subject of your message. If you have questions about any aspect of this book, please email us at questions@packtpub.com.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you would report this to us. Please visit www.packtpub.com/submit-errata, selecting your book, clicking on the Errata Submission Form link, and entering the details.

Piracy: If you come across any illegal copies of our works in any form on the Internet, we would be grateful if you would provide us with the location address or website name. Please contact us at copyright@packtpub.com with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit authors.packtpub.com.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions, we at Packt can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about Packt, please visit packtpub.com.

1 The Serverless Paradigm

Most probably, if you are reading this book, you have already heard about the serverless paradigm and the terms serverless engineering and serverless architecture. Nowadays, the way developers deploy applications has changed drastically, especially in the domain of data engineering and web development, thanks to **event-based architectural designs**, also called **serverless architectures**.

It is not uncommon to have idle resources and servers in production idle after the server workload has finished, or waiting for the next workload to come. This introduces a bit of redundancy in the infrastructure. What if there was no need for idle resources lying around when there is no workload? What if resources can be created when necessary and be destroyed once the work is done?

At the end of this chapter, you will understand how serverless architectures and functions as a service work, and how you can build them into your existing software infrastructure. You will also learn what microservices are, and decide whether microservices or serverless operations are well-suited for your architecture or not. You will also learn how to build serverless applications with Python on major cloud service providers, such as **Amazon Web Services (AWS)** and **Microsoft's Azure**.

This chapter will cover the following points:

- Understanding serverless architectures
- Understanding microservices
- Serverless architectures don't have to be real-time only
- Pros and cons of serverless architectures

Understanding serverless architectures

The concept of serverless architectures or serverless engineering revolves entirely around understanding the concept of functions as a service. The most technical and accurate definition of serverless computing on the internet is as follows:

*"Serverless computing, also known as **function as a service (FAAS)**, is a cloud computing and code execution model in which the cloud provider fully manages starting and stopping of a function's container **platform as a service (PaaS)**."*

Now, let's go into the details of each part of that definition to understand the paradigm of serverless computing better. We shall start with the term function as a service. It means that every serverless model has a function that is executed on the cloud. These functions are nothing but blocks of code, that are executed depending on the trigger that is associated with the function. This is a complete list of triggers in the AWS Lambda environment:

- Amazon S3
- Amazon DynamoDB
- Amazon Kinesis Streams
- Amazon Simple Notification Service
- Amazon Simple Email Service
- Amazon Cognito
- AWS CloudFormation
- Amazon CloudWatch Logs
- Amazon CloudWatch Events
- AWS CodeCommit
- Scheduled Events (powered by Amazon CloudWatch Events)
- AWS Config
- Amazon Alexa
- Amazon Lex
- Amazon API Gateway
- Other Event Sources: Invoking a Lambda Function On Demand
- Sample Events Published by Event Sources

Now let's understand what manages the starting and stopping of a function. Whenever a function is triggered via one of these available triggers, the cloud provider launches a container in which the function executes. Also, after the function is successfully executed the function has returned something, or if the function has run out of time, the container gets thatched away or destroyed. The thatching happens so that the container can be reused in the event of high demand and whenever there is very little time between two triggers. Now, we come to the next part of the sentence, the function's container. This means that the functions are launched and executed in containers. This is the standard definition of a container from Docker, a company that made the concept of containers very popular:

"A container image is a lightweight, stand-alone, executable package of a piece of software that includes everything needed to run it: code, runtime, system tools, system libraries, settings."

This helps in packaging the code, the runtime environment, and so on of the function into a single deployment package for seamless execution. The **deployment package** contains the main code file for the function, all the non-standard libraries which are required for the function to execute. The creation process of a deployment package looks very similar to that of a virtual environment in Python.

So, we can clearly make out that there are no servers running round the clock in the case of serverless infrastructures. There is a clear benefit for this, which includes not having a dedicated Ops team member for monitoring the server boxes. So the extra member, if any, can focus on better things, such as software research, and so on. Not having servers running through the entire day saves a lot of money and resources for the company and/or personally. This benefit can be very clearly seen among machine learning and data engineering teams who make use of GPU instances for their regular workload. So having on-demand serverless GPU instances running, saves a lot of money without the developers or the Ops team needing to maintain them around the clock.

Understanding microservices

Similar to the concept of serverless, the design strategy, which is the microservice-oriented strategy, has also been very popular recently. This architecture design existed a long time before the idea of serverless came into existence though. Just as we tried to understand the serverless architectures from the technical definition on the internet, we shall try to do the same for microservices. The technical definition for microservices is:

*"Microservices, also known as the **microservice architecture**, is an architectural style that structures an application as a collection of loosely coupled services, which implement business capabilities."*

Planning and designing the architecture in the form of microservices has its fair share of positives and negatives, just like serverless architectures. It's important to know about both, in order to appreciate and understand when and when not to leverage microservices in your existing architecture. Let's look at this and understand the positives of having microservice architectures, before moving over to the negatives.

Microservices help software teams stay agile, and improve incrementally. In simpler terms, as the services are decoupled from each other, it is very easy to upgrade and improve a service without causing the other to go down. For example, in social network software, if the chat and the feed are both microservices, then the feed doesn't have to go down when the software team are trying to upgrade or do minor fixes on the chat service. However, in large monolithic systems, it is difficult to break things up so easily in the way one can do with microservices. So, any fix or upgrade on even a small component of the architecture comes with downtime with the fix taking more time than intended.

The sheer size of the code base of monolithic architectures itself acts as a hindrance progress in the case of any small failures. Microservices, on the other hand, greatly help in boosting developer productivity by keeping code bases lean, so that they can fix and improve the service with very little or no overhead and downtime. Microservices can be much better leveraged via containers, which provide effective and complete virtual operating system environments, processes with isolation, and dedicated access to underlying hardware resources.

However, microservices come with their own bunch of disadvantages and downsides, the major one being having to deal with distributed systems. Now that each service is surviving on its own, the architect needs to figure out how each of them interacts with the others in order to make a fully functional product. So, proper co-ordination between the services and the decisions regarding how services move data between them is a very difficult choice that needs to be taken by the architect. Major distributed problems such as the *consensus*, the *CAP theorem*, and *maintaining the stability of consensus*, and the *connection*, are some issues that the engineer needs to handle while architecting for microservices. Ensuring and maintaining security is also a major problem in distributed systems and microservices. You needs to decide on separate security patterns and layers for each microservice, along with the security decisions necessary for the data interaction to happen between the services.

Serverless architectures don't have to be real-time only

Serverless architectures generally are leveraged as real-time systems as they work as a *function as service* which is triggered by a set of available triggers. However, this is a very common misconception, as serverless systems can be leveraged equally well both as real-time and batch architectures. Knowing how to leverage the concept of serverless systems as batch architectures will open up many engineering possibilities, as all engineering teams don't necessarily need or have real-time systems to operate.

Serverless systems can be batched by leveraging the following:

- The cron facility in triggers
- The concept of queues

Firstly, let's understand the concept of the **cron facility** in triggers. Serverless systems on the cloud have the ability to set up monitoring, which enables the trigger to get triggered every few minutes or hours and can be set as a normal cron job. This helps in leveraging the concept of serverless as a regular cron batch job. In the AWS environment, Lambda can be triggered as a cron via AWS CloudWatch, by setting the frequency of the cron by manually entering the time interval as the input and also by entering the interval in the cron format:

Example	Cron expression
Invoke Lambda function every 5 minutes	 Copy rate(5 minutes)
Invoke Lambda function every hour	 Copy rate(1 hour)
Invoke Lambda function every seven days	 Copy rate(7 days)

One can also leverage the concept of queues when building serverless batch architectures. Let's understand this by setting an example data pipeline. Let's say the system which we intend to build does the following tasks:

1. A user or a service sends some data into a database or a much simpler data store, such as AWS's S3.
2. Once there are more than 100 files in my data store, we'll want to do some task. Let's say, doing some analytics on them, for example, such as counting the pages.

This can be achieved via queues, and this is one of the simpler serverless systems we can consider as an example. So, this can be achieved as follows:

1. The user or the service uploads or sends the data to the data store which we have selected.
2. A queue is configured for the purpose of this task.
3. An event can be configured to S3 buckets or data stores so that as soon as data enters into the store, a message is sent to the queue which we have configured earlier.
4. Monitoring systems can be set to monitor the queue for the number of messages in it. It is advisable to use the monitoring system of the cloud provider you are using so that the system stays completely serverless.
5. Alarms can be set to the monitoring systems, configuring a threshold for these alarms. For example, the alarm needs to be triggered whenever the number of messages in our queue reaches or exceeds 100.
6. This alarm can act as a trigger to the Lambda function which does the analytics by first receiving messages from the queue and then querying the data store using the filename received from the message.
7. Once the analytics are completed on the files, the processed files can be pushed to another data store for storage.
8. After the entire task is completed, the container or the server where the Lambda function has run will be terminated, thus making this pipeline completely serverless.

Pros and cons of serverless

As we now understand what serverless architectures and pipelines look like, how they may be leveraged into existing architectures, and also how microservices help keep architectures leaner and boost developer productivity, we shall look at the pros and cons of serverless systems in detail, so that software developers and architects can make decisions regarding when to leverage the serverless paradigm into their existing systems and when not to.

The positives of serverless systems are:

- **Lower infrastructure costs:** By deploying serverless systems, the infrastructure costs can be greatly optimized, as there would not be a need for servers to be running around the clock every day. As the servers start whenever the function is triggered, and stop whenever the function gets executed successfully, the billing would only be done for that brief time period when the function was running.
- **Less maintenance needed:** By virtue of the preceding reason, there is also no need for continuous monitoring and maintenance of servers. As the functions and triggers are automated, there is almost zero maintenance required for serverless systems.
- **Higher developer productivity:** As the developers don't need to worry about downtime and server maintenance, they can focus and work on better software challenges, such as scaling and designing functionalities.

The remaining part of the book will show you how serverless systems are changing the way software is done. So, as this chapter is intended to help architects decide whether serverless systems are a good choice for their architecture or not, we shall now look at the disadvantages of serverless systems.

The disadvantages of serverless systems are:

- **Time limit of the function:** The function which is whether executed, be it AWS's Lambda or GCP's cloud functions, has an upper time limit of 5 minutes. This makes execution of heavy computations impossible. However, this can be solved by executing a provisioning tool's playbook in nohup mode. This will be covered in detail, later in the chapter. However, making the playbook ready and setting up the container and anything else should be completed within the 5 minute time limit. The container gets automatically killed when the 5 minute limit is exceeded.

- **No control over the container environment:** The developer has no control over the environment of the container that is being created for executing the function. The operating system, the filesystem, and so on, are all decided by the cloud provider. For example, AWS's Lambda functions are executed inside containers that run the Amazon Linux operating system.
- **Monitoring containers:** Apart from the basic monitoring capabilities that are provided by the cloud provider via their in-house monitoring tools, there is no mechanism to do detailed monitoring of the container that is executing the serverless function. This becomes even more difficult when scaling up serverless systems to accommodate distributed systems.
- **No control on security:** There is no control on how the security of the data flow is ensured, as there is very little control over the container's environment. The container can be run in the VPC and subnets of the developer's choice, though, which helps work around this disadvantage.

However, serverless systems can be scaled up to distributed systems for large-scale computations where the developer need not worry about the time limit. As already mentioned, this will be discussed in detail in the upcoming chapters. However, for insight into an intuition on how one can choose serverless systems over monolithic systems for large-scale computations, let us understand some important pointers that need to be kept in mind when taking that architectural decision.

The pointers to be kept in mind when scaling serverless systems to distributed systems are:

- To scale up serverless systems into serverless distributed systems, one must understand how the concept of `nohup` works. It is a **POSIX** command that allows programs and processes to run in the background.
- `Nohup` processes should be properly logged, including both the output and the error logs. This is the only source of information for your processes.
- A provisioning tool, such as **Ansible** or **Chef** or a similar one, needs to be leveraged to create a master-workers architecture which has been spawned via the playbook running in `nohup` mode in the container where the serverless function is being executed.
- It is a good practice to ensure that all tasks that are being executed by the provisioning tool via the master server are properly monitored and logged, as there is no way one can retrieve the logs once the entire setup finishes executing.

- Proper security needs to be ensured by using a temporary credential facility available from the cloud providers.
- Proper closure should be ensured for the system. The workers and the master should self-terminate immediately after the pipeline of tasks finishes executing. This is very important and this is what makes the system serverless.
- Generally, temporary credentials come with an expiry time, which is 3,600 seconds for most environments. So, if the developer is using temporary credentials to execute a task which is supposed to take more than the expiry time, then there is a danger of the credentials getting expired.
- Debugging distributed serverless systems is an extremely difficult task for the following reasons:
 - Monitoring and debugging a nohup process is extremely difficult. Whenever you want to debug one, you have to either refer to the log file created by the process or kill the nohup process by using the process ID, and then manually run the scripts for debugging.
 - As the complete list of tasks executes sequentially in the provisioning tool, there is a danger that the instances may get terminated because the developer has forgotten to kill the nohup process before starting the debugging process.
 - As this is a distributed system, it goes without saying that the architecture should be able to self-heal in the case of any failure or a disaster. An example scenario can be when one of the workers goes down while performing some operation on a bunch of files. The entire bunch of files is now lost, and there is no means of recovery.
 - Another advanced disaster scenario can be when two worker servers go down while performing some operations on a bunch of files. In this case, the developer does not know which files have been executed successfully and which haven't.
- It is a good practice to ensure that all the worker instances receive an equal amount of the load to execute so that the load across the distributed system stays even and time and resources are well optimized.

Summary

In this chapter, we learned what serverless architecture is. Most importantly, the chapter helps architects decide if serverless is the way forward for their team and their engineering, and how to transition/migrate from their existing infrastructure to a serverless paradigm. We also looked at the paradigm of microservices and how they help make lightweight and highly agile architectures. This chapter also went into great detail about when a team should start thinking about microservices and when can they migrate or break their existing monolith(s) into microservices.

We then learned the art of building batch architectures in the serverless domain. One of the most common myths is that serverless systems are only for real-time computation purposes. However, we have learned how to leverage these systems for batch computations too, thus facilitating a whole lot of applications with the serverless paradigm. We looked at the pros and cons of going serverless so that better engineering decisions can be taken accordingly.

In the next chapter, we will cover a very detailed understanding of how AWS Lambda works, which is the core component of serverless engineering in the AWS cloud environment. We will understand how triggers work and how AWS Lambda functions work. You will learn about the concept of leveraging containers for executing serverless functions and the associated computational workload. Following that, we will also learn about configuring and testing Lambda functions, along with understanding the best practices while doing so. We will also cover versioning Lambda functions, in the same way versioning is done in code, and then create deployment packages for AWS Lambda, so that the developer can accommodate third-party libraries comfortably, along with the standard library ones.

2

Building a Serverless Application in AWS

This chapter will introduce the concept of serverless applications using AWS Lambda as the tool of choice. This will help you understand the concept, intuition, and working components involved in a serverless tool. It will also explain the nuances involved in security, user-controls, and versioning code inside Lambda. You will be guided via hands-on tutorials and lessons for understanding and learning to use AWS Lambda. So, it is recommended that you follow along this chapter with a laptop and an AWS account setup to easily execute the given instructions.

This chapter will cover the following topics:


- Triggers in AWS Lambda
- Lambda functions
- Functions as containers
- Configuring functions
- Testing Lambda functions
- Versioning Lambda functions
- Creating deployment packages

Triggers in AWS Lambda

Serverless functions are on-demand computational concepts. So, there has to be an event that needs to trigger a Lambda function so that the entire computational process is started. AWS Lambda has several events which can act as a trigger. Almost all services of AWS can act as AWS Lambda's triggers. Here is the list of services that you can use for generating events for Lambda to respond to:


- API Gateway
- AWS IoT
- CloudWatch Events
- CloudWatch Logs
- CodeCommit
- Cognito Sync Trigger
- DynamoDB
- Kinesis
- S3
- SNS

The triggers page of AWS Lambda looks like this:



Let's take a look at some of the following important and widely-used triggers that are available, and understand how they can be leveraged as FaaS in the serverless paradigm. They are as follows:

- **API Gateway:** This trigger can be used to create efficient, scalable, and serverless APIs. One scenario where a serverless API makes sense would be while building a querying interface for S3. Let us assume that we have a bunch of text files in an S3 bucket. Whenever a user hits the API with a query parameter, which can be some word that we want to search in the text files in the bucket, the API Gateway's trigger will launch a Lambda function that executes the computational logic and workload for executing the query. The Lambda function that we want our API to trigger can be specified at the API creation time. The trigger will be created accordingly in the corresponding Lambda function's console. This is what it looks like:



- **CloudWatch:** Its events mostly help the user in setting the cron scheduling for Lambda. The CloudWatch Logs trigger is useful whenever a user wants to execute a computational workload depending on some keyword in the Cloudwatch Logs. However, the CloudWatch Alarms cannot trigger Lambda directly via the CloudWatch trigger. They have to be sent via a notification system, such as the **AWS Simple Notification Service (AWS SNS)**. This is how you can create a cron execution in AWS Lambda. In the following screenshot, the Lambda function is set to execute every minute:

The screenshot shows the AWS Lambda console's 'Configure function' page. On the left, a sidebar indicates 'Step 4 Review'. The main area is titled 'CloudWatch Events' with a play button icon and the 'Lambda' logo. The configuration form includes:

- Rule:** A dropdown menu set to 'Create a new rule'.
- Rule name*:** A text input field containing 'Lambda-cron'.
- Rule description:** An empty text input field.
- Rule type:** Two radio buttons; 'Event pattern' is unselected, and 'Schedule expression' is selected.
- Schedule expression*:** A text input field containing '1 minute'.

Below the schedule expression field, a note states: 'Self-trigger your target on an automated schedule using Cron or rate expressions. Cron expressions are in UTC.' Examples provided are 'e.g. rate(1 day), cron(0 17 ? * MON-FRI *)'.

- **S3:** This is a document store of AWS. So, whenever a file is added, removed, or changed, an event will be sent to AWS Lambda when added as a trigger. So, if you want to do some computational workload on a file as soon as the file gets uploaded, then this trigger helps to do that. This is what an S3's event structure looks like:

```
{
  "Records": [
    {
      "eventVersion": "2.0",
      "eventSource": "aws:s3",
      "awsRegion": "us-east-1",
      "eventTime": "The time, in ISO-8601 format, for example, 1970-01-01T00:00:00.000Z, when S3 finished the event",
      "eventName": "event-type",
      "userIdentity": {
        "principalId": "Amazon-customer-ID-of-the-user-who-caused-the-event"
      },
      "requestParameters": {
        "sourceIPAddress": "ip-address-where-request-came-from"
      },
      "responseElements": {
        "x-amz-request-id": "Amazon S3 generated request ID",
        "x-amz-id-2": "Amazon S3 host that processed the request"
      },
      "s3": {
        "s3SchemaVersion": "1.0",
        "configurationId": "ID found in the bucket notification configuration",
        "bucket": {
          "name": "bucket-name",
          "ownerIdentity": {
            "principalId": "Amazon-customer-ID-of-the-bucket-owner"
          },
          "arn": "bucket-ARN"
        },
        "object": {
          "key": "object-key",
          "size": "object-size",
          "eTag": "object eTag",
          "versionId": "object version if bucket is versioning-enabled, otherwise null",
          "sequencer": "a string representation of a hexadecimal value used to determine event sequence, only used with PUTs and DELETes"
        }
      }
    }
  ],
}
```

- **AWS SNS:** The SNS service of AWS helps users to send notifications to other systems. This service can also be used for catching CloudWatch Alarms and sending the notifications to a Lambda function for computational execution. This is what a sample SNS event looks like:

Amazon SNS Sample Event

```
{
  "Records": [
    {
      "EventVersion": "1.0",
      "EventSubscriptionArn": eventsubscriptionarn,
      "EventSource": "aws:sns",
      "Sns": {
        "SignatureVersion": "1",
        "Timestamp": "1970-01-01T00:00:00.000Z",
        "Signature": "EXAMPLE",
        "SigningCertUrl": "EXAMPLE",
        "MessageId": "95df01b4-ee98-5cb9-9903-4c221d41eb5e",
        "Message": "Hello from SNS!",
        "MessageAttributes": {
          "Test": {
            "Type": "String",
            "Value": "TestString"
          },
          "TestBinary": {
            "Type": "Binary",
            "Value": "TestBinary"
          }
        },
        "Type": "Notification",
        "UnsubscribeUrl": "EXAMPLE",
        "TopicArn": topicarn,
        "Subject": "TestInvoke"
      }
    }
  ]
}
```



Lambda functions

Lambda functions are the core operating parts of a serverless architecture. They contain the code which is supposed to be executed. These functions are executed whenever the trigger attached to it has been set off. We have already learned about some of the most popular Lambda triggers in the previous section.

Whenever a Lambda function is triggered, it creates a container with the respective settings set by the user. We'll learn more about the container in our next section.

The spinning up of containers takes a bit of time, which may result in a latency whenever a fresh invocation of a Lambda function is done, as it takes time to set up the environment and bootstrap the settings mentioned by the user in the **Advanced settings** tab. So, to overcome this latency, AWS thaws a container for some time for reuse in case of another Lambda invocation within the thawing time. So, using a thawed or a ready-made Lambda function helps in overcoming the latency problem. However, the same global namespace of the thawed container would be reused for the new invocation too.

So, if the Lambda function has any global variables that get manipulated inside the function, it is a good idea to convert them into local namespaces, as the manipulated global namespace variables will be reused, leading to faulty execution results of the Lambda function.

The user needs to specify the technical details for the Lambda function in the **Advanced Settings** tab, which include the following:

- **Memory (MB):** This is the maximum memory that the Lambda function needs to be allocated for the purpose of your function. The CPU of the container would be assigned accordingly.
- **Timeout:** The maximum amount of time the function needs to execute before the container gets automatically stopped.
- **DLQ Resource:** This is a dead-letter setting to AWS Lambda. The user can add either an SQS queue or an SNS topic for configuring this. Lambda functions get asynchronously retried for at least five times on failure.
- **VPC:** This enables the Lambda function to access components or services in some particular VPCs. The Lambda function executes in a default VPC of its own.
- **KMS key:** If there are any environment variables entered along with the Lambda function, this helps us encrypt them using an **AWS Key Management Service (KMS)** by default.

The Lambda function's **Advanced settings** page looks like this:

▼ Advanced settings

Memory (MB)*
Your function is allocated CPU proportional to the memory configured.
128 MB

Timeout*
0 min 3 sec

DLQ Resource [Info](#)
Choose the AWS service to send event payload to after exceeding maximum retries.
Select resource ▼

VPC [Info](#)
Select a VPC that your function will access.
No VPC ▼

Enable active tracing [Info](#)
☐

KMS key [Info](#)
Select a KMS key to encrypt the environment variables at rest, or simply let Lambda manage the encryption.
(default) aws/lambda ▼ [Enter value](#)

Functions as containers

For understanding the concept of functions being executed as/inside containers, we need to properly understand the concept of containers. To cite the definition of a container from the Docker documentation (<https://www.docker.com/what-docker>):

A container image is a lightweight, stand-alone, executable package of a piece of software that includes everything needed to run it: code, runtime, system tools, system libraries, settings.

What is available for both Linux and Windows based applications; containerized software will always run the same, regardless of the environment.

Containers isolate software from its surroundings (for example, differences between development and staging environments) and help reduce conflicts between teams running different software on the same infrastructure.

So, the concept of containers is that they are self-sustainable isolated environments just like the containers in a container ship that can be hosted and be worked upon any host OS, the host OS being the host ship in our analogy. The figurative depiction of the analogy would look something like this:




Similar to the aforementioned analogy, AWS Lambda's functions are also launched inside a unique container for each function. So, let us understand this topic in more detail, point by point:

1. The Lambda function can be in the form of a single code file or in the form of a **deployment package**. The deployment package is a zipped file that includes the core function file along with the libraries which would be used by the function. We shall be learning in detail about how to create the deployment package in the *Creating deployment packages* section of this chapter.
2. Whenever a function is triggered or started, AWS spins up an EC2 instance with the AWS Linux operating system for running the function. The configuration of the instance would be dependent on the ones provided by the user in the **Advanced settings** tab of the Lambda function.
3. There is a maximum time limit of 300 seconds, or 5 minutes, for a function to execute successfully, after which the container would be destroyed. So, this needs to be kept in mind while designing the Lambda functions and/or the deployment packages.


Configuring functions

In this section, we will go through the ways of configuring Lambda functions and understand all the settings in great detail. Like in the previous section, we will learn about each configuration and its settings, as follows:


1. You can go to the page of AWS Lambda by selecting it from the drop-down menu that is present in the top-left corner of the AWS console. This can be done as follows:



2. Once the **Lambda** option is selected, it redirects the user to the **AWS Lambda** console, which looks something like this:



3. For creating a function, you need to click on the orange **Create a function** button on the right. This will open a console for the function creation. This looks something like this:




- Let's create a function from scratch in order to understand the configurations better. So, for doing that, click on the **Author from scratch** button on the top-right corner. After clicking it, the user will be directed to Lambda's first-run console, which looks something like this:

The screenshot shows the AWS Lambda console's 'Create function' page, specifically the 'Author from scratch' tab. The breadcrumb navigation at the top reads 'Lambda > Functions > Create function > Author from scratch'. The main section is titled 'Basic information' with an 'Info' link. It contains three required fields: 'Name*' with a text input containing 'myFunctionName', 'Role*' with a dropdown menu showing 'Choose an existing role', and 'Existing role*' with another dropdown menu. Below the 'Existing role*' dropdown is a note: 'You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.' At the bottom left, a note states '* These fields are required.' At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Create function'.


- This page has three configurations which the user can select, which are **Name**, **Role**, and **Existing role**. The **Name** value is where the user can enter the name of the Lambda function. The **Role** value is how you can define permissions in the AWS environment. The **Role** value's drop-down list would contain the following options: **Choose an existing role**, **Create new role from template(s)**, and **Create a custom role**. They can be seen as follows:

This is a close-up of the 'Role*' dropdown menu. The label 'Role*' is at the top, followed by the text 'Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.' The dropdown list is open, showing four options: 'Choose an existing role' (highlighted in light blue), 'Choose an existing role', 'Create new role from template(s)', and 'Create a custom role'. To the right of the dropdown, a partial note is visible: 'a and must have Cloudwatch'.

The **Choose an existing role** option will enable us to select an already existing role with pre-configured permissions. The second option helps the user with creating a role from pre-baked templates. The **Create a custom role** option allows the user to create a role with permissions from scratch. The list of pre-baked roles looks like this:



6. Select one from the pre-baked templates for the sake of this tutorial. By pressing **Create function** in the lower-right part of the screen, we will land on the Lambda function's creation page, which looks similar to this:



7. The preceding page indicates that we have successfully created an AWS Lambda function. We shall now explore the advanced settings of this function. They are present in the lower part of the same console. They will look something like this:

AWS Lambda ×

Dashboard
Functions

► Environment variables

► Tags

▼ Execution role

Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.

Choose an existing role ▼

Existing role
You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.

service-role/Pycontw-Role ▼

▼ Basic settings

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.

128 MB

Timeout [Info](#)
0 min 3 sec

Description

► Network

► Debugging and error handling

We shall now try to understand each of those parts in detail.

- The unfurled **Environment variables** section contains text boxes to enter the key-value pair of environment variables that will be used by our function. One can also optionally mention details on the encryption setting that we want to have for the environment variables. The encryption needs to be done via **AWS KMS (Key Management Service)**. The unfurled settings box of the environment variables looks something like this:

▼ Environment variables

You can define Environment Variables as key-value pairs that are accessible from your function code. These are useful to store configuration settings without the need to change function code. [Learn more](#).

Key	Value	Remove

▼ Encryption configuration

Enable helpers for encryption in transit [Info](#)
☐

KMS key to encrypt at rest [Info](#)
Select a KMS key to encrypt the environment variables at rest, or simply let Lambda manage the encryption.


(default) aws/lambda ▼ Enter value

9. The next settings section is **Tags**. This is similar to the tagging feature of all the available AWS services for easy service discovery purposes. So, similar to all AWS services's tags, this also needs just a key and a value. The unfurled **Tags** section looks something like this:




The screenshot shows the 'Tags' section of the AWS Lambda console. It has a title bar with a dropdown arrow and the text 'Tags'. Below the title bar, there is a descriptive text: 'You can use tags to group and filter your functions. A tag consists of a case-sensitive key-value pair. [Learn more.](#)'. Below this text, there are two input fields: 'Key' and 'Value', and a 'Remove' button.

10. The next section that will be visible after the **Tags** section is the **Execution role** section, in which the user can set the **Identity Access Management (IAM)** role for the execution of the Lambda function. As we have already discussed what IAM roles are previously in the book, we will not be covering that again here. If the user has not set the role when creating the function itself, they can always set that here. The section will be visible in the Lambda console as follows:



The screenshot shows the 'Execution role' section of the AWS Lambda console. It has a title bar with a dropdown arrow and the text 'Execution role'. Below the title bar, there is a descriptive text: 'Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.' Below this text, there is a dropdown menu with the text 'Choose an existing role'. Below the dropdown menu, there is a section titled 'Existing role' with a descriptive text: 'You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.' Below this text, there is a dropdown menu with the text 'service-role/Pycontw-Role'.

11. The next section is the **Basic settings** section, which includes settings such as the memory of the Lambda container, time-out for the container, and the description for the Lambda function. The memory of the container can range from 128 MB to 1,536 MB. The user can choose any value within that range and will be billed accordingly. The time-out can be set from 1 second to 300 seconds, which is 5 minutes. The time-out is the time which the Lambda function and its container would run before being stopped or terminated. The next setting is the **Description** value of the Lambda function, which acts as the metadata of a Lambda function. The section looks like this in the console:



The screenshot shows the 'Basic settings' section of the AWS Lambda console. It includes a memory slider set to 128 MB, a timeout spinner set to 1 second, and a description text box.

▼ **Basic settings**


Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.

128 MB

Timeout [Info](#)
0 min 1 sec


Description

12. The next section is the **Network** section, which is also about the network settings of the Lambda function related to **AWS's Virtual Private Cloud (VPC)** and related subnets. Even if **No VPC** is selected as an option, AWS Lambda runs in its own secure VPC. However, if your Lambda function accesses or deals with any other service which is in a particular VPC or in a subnet, the corresponding information needs to be added in this section so that the network allows traffic from the Lambda function's container. This section looks like this in the console:



The sensitive information in the preceding screenshot, such as the IP address and the ID of the VPC, are masked for security purposes.

13. The next section is the **Debugging and error handling** section. This section enables the user to set up measures for ensuring fault tolerance and exception handling of the Lambda function. This includes the **Dead Letter Queue (DLQ)** settings.
14. Lambda automatically retries failed executions for asynchronous invocations. So, the payloads that were not processed would be automatically forwarded to the DLQ resource. The DLQ settings look like this in the Lambda console:




▼ **Debugging and error handling**

DLQ Resource [Info](#)
Choose the AWS service to send event payload to after exceeding maximum retries.

Select resource ▼

- None
- SNS
- SQS

The user can also enable active tracing for the Lambda functions, which would help in detailed monitoring of the Lambda container. This setting in the **Debugging and error handling** section of the Lambda console looks like this:



▼ **Debugging and error handling**

DLQ Resource [Info](#)
Choose the AWS service to send event payload to after exceeding maximum retries.

Select resource ▼


Enable active tracing [Info](#)

☐

Testing Lambda functions

Just like every other software system and programming paradigm, proper testing of Lambda functions and serverless architectures is very important before deploying into production. We will try to understand the testing of Lambda functions in the following points:

1. In the top-most bar of the Lambda console, one can observe the **Save and test** option, which is represented by an orange button. This button saves the Lambda function and then runs the configured tests on that function. This looks something like this in the console:



2. Also, in the same bar, there exists a drop-down menu that reads **Select a test event....** This contains a list of testing events available for testing Lambda functions. The drop-down looks like this:



- Now, for further configuration of test events for the Lambda function, the user needs to select the **Configure test events** option in the drop-down. This will open a popup with the test events menu, which looks like this:

Configure test event

A function can have up to 10 test events. The events are persisted so you can switch to another computer or web browser and test your function with the same events.

☒ Create new test event
☐ Edit saved test events

Event template


Hello World

Event name


MyEventName

```
1 {  
2   "key3": "value3",  
3   "key2": "value2",  
4   "key1": "value1"  
5 }
```


4. That would open the basic **Hello World** template, which has three pre-configured JSON format test events, or edge cases. However, depending on what the Lambda function does, one can select some other test event. The available list of testing templates can be seen in the **Event template** drop-down menu. The list in the drop-down looks something like this:



5. For example, let's imagine we are building a pipeline that involves the Lambda function getting started whenever an image file is added to an S3 bucket, and the function does some image processing tasks and puts it back to some data store. The test event of the **S3 Put** notification looks something like this:



6. After selecting or creating a test event, the user can select the **Create** option in the bottom-right corner of the event creation console, wherein you shall be asked to enter a name for the event. After entering the necessary details, the user will be re-directed back to the Lambda console. Now, when you check the **TestEvent** drop-down in the Lambda console, you can see the saved test event in the list. This can be verified as follows:



As I have named the event as **TestEvent**, the test is visible by the same name in the events drop-down menu.


6. Additionally, when we take a closer look at the event structure of S3 in the test event, we can observe the meta-details that are being made available to the Lambda function. The event structure looks like this:

```
{
  "Records": [
    {
      "eventVersion": "2.0",
      "eventSource": "aws:s3",
      "awsRegion": "us-west-2",
      "eventTime": "1970-01-01T00:00:00.000Z",
      "eventName": "ObjectCreated:Put",
      "userIdentity": {
        "principalId": "AIDAJDPLRKLG7UEXAMPLE"
      },
      "requestParameters": {
        "sourceIPAddress": "127.0.0.1"
      },
      "responseElements": {
        "x-amz-request-id": "C3D13FE58DE4C810",
        "x-amz-id-2": "FMjUUVURIY8/IgAtTv8xRjskZQpcIZ9KG4V5Wp6S7S/JRWeUwerMUE5JgHvAN0jpd"
      },
      "s3": {
        "s3SchemaVersion": "1.0",
        "configurationId": "testConfigRule",
        "bucket": {
          "name": "sourcebucket",
          "ownerIdentity": {
            "principalId": "A3NL1K0ZZKExample"
          },
          "arn": "arn:aws:s3:::sourcebucket"
        },
        "object": {
          "key": "HappyFace.jpg",
          "size": 1024,
          "eTag": "d41d8cd98f00b204e9800998ecf8427e",
          "versionId": "096fKKXTRTtL3on89fv0.nfljtsv6qko"
        }
      }
    }
  ]
}
```


Versioning Lambda functions

The concept of the **Version Control System (VCS)** is for controlling and managing versions of code. This functionality is available directly from the main Lambda console. Let's try and learn how to version our Lambda functions:

1. The first option in the **Actions** drop-down in the Lambda console is the **Publish new version** option. This option can be seen here:




2. When the **Publish new version** option is selected, the versioning popup of the Lambda console would be seen on the console. This would ask about the name for the new version of your Lambda function. The popup looks something like this:




The screenshot shows the AWS Lambda console interface. At the top, there's a header bar with 'Resource Groups' and a user profile 'Raj Rohit'. A modal dialog box titled 'Publish new version from \$LATEST' is open. The dialog contains a warning message: 'Publishing a new version will save a "snapshot" of the code and configuration of the \$LATEST version. You will be unable to edit the new version's code. Please click to confirm.' Below this is a text input field labeled 'Version description'. At the bottom right of the dialog are 'Cancel' and 'Publish' buttons. The background of the console is dimmed, showing a message 'You do not have any triggers for this function.' with options to '+ Add trigger', 'Refresh triggers', and 'View function policy'.


3. After clicking the **Publish** button, you will be re-directed to the main Lambda console. The successfully created Lambda version in the console looks something like this:



4. In the bottom half of the page, the following message can be noticed: **Code and handler editing is only available for the \$LATEST version.** This means that one can only edit the code in the version named **\$LATEST**. The versioned version of Lambda functions are read-only and cannot be edited and manipulated. When something goes wrong or when the user wants to revert back or refer to a previous version, that version will overlay the **\$LATEST** version to make edits possible. The message looks like this:




- When the **Click here to go to \$LATEST** link is clicked, the user will be re-directed to the **\$LATEST** version of the function, which can be edited and manipulated by the user. The console of the **\$LATEST** version of the Lambda function looks like this:



Creating deployment packages

Lambda functions that have external libraries as dependencies can be packaged as deployment packages and be uploaded into the AWS Lambda console. This is very similar to creating a virtual environment in Python. So in this section, we shall learn and understand the process of creating Python deployment for using in the Lambda functions. We shall try and understand the process of creating deployment packages in detail, as follows:


1. Deployment packages are generally in the format of ZIP packages. The contents of the ZIP package is exactly the same as a normal library of any programming language.
2. The package structure should be such that the library folders and the function file are in the same destination or in the same hierarchy inside the folder structure of the deployment package. The layout looks something like this:




3. The Python libraries can be installed by using the `pip install <library_name> -t <path_of_the_target_folder>` command. This will install the package inside the target folder. This can be done as in the following screenshot:

```
+ ServerlessBook sudo pip install numpy -t TestDir
Password:
The directory '/Users/Dawny33/Library/Caches/pip/http' or its parent directory is not owned by the current user and the cache has been disabled. Please check the permissions and owner of that directory. If executing pip with sudo, you may want sudo's -H flag.
The directory '/Users/Dawny33/Library/Caches/pip' or its parent directory is not owned by the current user and caching wheels has been disabled. check the permissions and owner of that directory. If executing pip with sudo, you may want sudo's -H flag.
Collecting numpy
  Downloading numpy-1.13.3-cp27-cp27m-macosx_10_6_intel.macosx_10_9_intel.macosx_10_9_x86_64.macosx_10_10_intel.macosx_10_10_x86_64.whl (4.6MB)
    100% |#####| 4.6MB 172kB/s
Installing collected packages: numpy
Successfully installed numpy-1.13.3
```

4. Now, when we have the entire deployment package's folder along with the library folders ready, we need to zip all of the folders including the Lambda function file before uploading it into the console. The following screenshot shows how the zipping needs to be done as per the folder hierarchy:



- Now, as the zipped package is ready, we shall be trying to upload the package to the Lambda console for processing. For uploading a Lambda package, we need to select the drop-down list of the **Code entry type** option in the console. The selection looks like this in the Lambda console:



- Once the **Upload a .ZIP file** option is selected, the uploader will become visible, where the user can directly upload the deployment package or even upload it via an S3 bucket. The wizard would look like this in the Lambda console:

Test-Lambda Qualifiers Actions Save TestEvent Save and test

Configuration Triggers Monitoring

▼ Function code

Code entry type: Upload a .ZIP file Runtime: Python 2.7 Handler: index.handler [info](#)

Function package* For files larger than 10 MB, consider uploading via S3.

► Environment variables

► Tags

English (US) © 2008 - 2017, Amazon Internet Services Private Ltd. or its affiliates. All rights reserved. Privacy Policy Terms of Use

7. As mentioned previously, the user can choose to upload the deployment package via an S3 file location too. This wizard looks like this in the Lambda console:

Test-Lambda Qualifiers Actions Save TestEvent Save and test

Configuration Triggers Monitoring

▼ Function code

Code entry type: Upload a file from Amazon S3 Runtime: Python 2.7 Handler: index.handler [info](#)


S3 link URL* Paste an S3 link URL to your function code .ZIP.

► Environment variables

► Tags

English (US) © 2008 - 2017, Amazon Internet Services Private Ltd. or its affiliates. All rights reserved. Privacy Policy Terms of Use

8. The deployment package's naming should be aligned with the values entered in the handler part of the settings. The deployment package's name and the Lambda function file's name are separated by a dot (.) and arranged in that order. This can be explicitly seen in the following screenshot:



`index` should be the name of the Lambda function's file name deployment package. The `handler` function file is the name of the core function handler inside, which is the Lambda function. As AWS's documentation states:

The module-name export value in your function". For example, `index.handler` would call `exports.handler` in `index.py`.

Summary

In this chapter, we have learned the concepts of how triggers work for AWS Lambda and how to select triggers depending on the problem statement and time intervals, in case of cron job triggers. We understood what Lambda functions are, along with understanding their functionalities and settings related to memory, VPCs, security, and fault tolerance. We also learned about the way container reuse is done under the hood specifically for AWS Lambda. Then, we covered event-driven functions and how they are implemented under the hood, the concept of containers, and their uses and applications in the domain of software engineering in general. Most importantly, from the concepts we learned regarding containers, we can now appreciate the options for choosing containers for running the Lambda functions.

After that, we talked about all the configuration settings available in the AWS Lambda dashboard, which are necessary to build and run a Lambda function from start to finish without any settings-related problems. We also learned about and understood the security settings inside Lambda so that the necessary VPC details and security keys settings are taken care of when configuring our Lambda functions. This was followed by testing Lambda functions depending on the choice of trigger selected. We learned what the responses of various AWS services look like, as they are the inputs for the Lambda functions. We then learned how to write custom hand-made tests for custom testing purposes.

Following that, we saw how versioning happens for the AWS Lambda functions. We learned the differences between past and present versions. We also learned that the present version is immutable, unlike the past versions, and also how to revert to past versions without much effort. We also learned how to create deployment packages for functions that have dependencies on external packages, which are not included in Python's standard library. We came across the function code naming nuances, including the filename and the method handler names, followed by the two ways deployment packages can be uploaded to the Lambda console; one being a manual upload and the other being from an S3 file location.

In the next chapter, we will be gaining a detailed understanding of the different triggers available in the Lambda console and how to use them. We will also learn about implementing them in Python code. We will understand the event structures and the responses from different AWS services and use that to build our Lambda functions. We will understand how to integrate each trigger into a Lambda function and do a specific task in Python. Finally, we will also be learning about ideas and best practices on how to move your existing infrastructures to serverless using the serverless paradigm.

3

Setting Up Serverless Architectures

So far, we have understood what the serverless paradigm is, and also how serverless systems work. We have also understood how AWS Lambda's serverless tool works. We have also learned the basics of how triggers work in AWS Lambda as well as a detailed understanding of the system settings and configuration available to the user in the Lambda environment. We have also learned how the Lambda console works, and also how to identify and use various parts of the Lambda console in detail, including code deployment, trigger manipulation, deploying tests in the console, versioning our Lambda function, and also the different settings available.

By the end of this chapter, you will have a clear understanding of all the important triggers available for AWS Lambda and how you can use them to set up efficient Lambda architectures. You will also understand what an event structure is, and what an event structure looks like for some AWS resources, and how you can use that knowledge to write and deploy better trigger-based Lambda architectures.

This chapter will cover the following points:

- S3 trigger
- SNS trigger
- SQS trigger
- CloudWatch Event and Logs trigger

S3 trigger


S3 is the AWS object storage service, where the user can store and retrieve any type of object. In this section, we shall be learning how the S3 trigger works, what the event structure of an S3 event looks like, and also how to make use of them in the learning to build a Lambda function.

We will be building a Lambda function that does the following:


1. Receives a PUT request event from the S3 service
2. Prints the name of the file and other major details
3. Transfers that file to a different bucket

So, let's get started on learning how to use the S3 trigger efficiently. We will be working on this task step-by-step, as follows:

1. Firstly, we need to create two S3 buckets for the task. One will be the bucket where the file will be uploaded by the user. The other will be the one where the file is transferred and uploaded by the Lambda function.
2. The S3 console looks like the following screenshot when there are no pre-existing buckets. You can go there by selecting the **S3** service from the drop-down **Services** menu in the top-left of your AWS console:




3. Let's create two buckets, namely `receiver-bucket` and `sender-bucket`.
4. The `sender-bucket` bucket will be used as the bucket where the user uploads the files. The `receiver-bucket` bucket is the one where the Lambda function uploads the files. So, as per our problem statement, whenever we upload files to the `sender-bucket` bucket, the Lambda function gets triggered and the files get uploaded to `receiver-bucket`.
5. When we click on the **Create bucket** button in the S3 console, we get a dialog that looks like this:



The screenshot shows the 'Create bucket' dialog box in the AWS S3 console. The dialog has a blue header with the title 'Create bucket' and a close button (X). Below the header is a progress bar with four steps: 1. Name and region, 2. Set properties, 3. Set permissions, and 4. Review. The first step, 'Name and region', is currently active. It contains three input fields: 'Bucket name' with a placeholder 'Enter DNS-compliant bucket name', 'Region' with a dropdown menu showing 'US East (N. Virginia)', and 'Copy settings from an existing bucket' with a dropdown menu showing 'You have no buckets' and '0 Buckets'. At the bottom of the dialog are three buttons: 'Create', 'Cancel', and 'Next'.

6. In the preceding dialog, we need to enter the following settings:
 - **Bucket Name:** As the name suggests, we need to enter the name of the bucket we are creating. For the creation of the first bucket, name it `sender-bucket` and name the second bucket `receiver-bucket`.
 - **Region:** This is an AWS region we want the bucket to reside in. You can use the default region for this or the region closest to wherever you are located.
 - **Copy settings from an existing bucket:** This specifies whether we want to use the same settings as in some other bucket in the console for this bucket too. As we do not currently have any other bucket in our console, we can skip this setting by leaving it empty. After this, you can click on the **Next** button in the bottom-right part of the popup.
7. Once we click **Next**, we get redirected to the second tab of the popup, which is the **Set properties** menu and looks like this:




8. In this part of the popup, we need to decide on the following settings:
- **Versioning:** This is relevant if we want to keep multiple versions of the files in the S3 bucket. This setting is required when you need a Git style versioning for your S3 bucket. Note that the storage cost would be included in line with the number of versioned documents.
 - **Server access logging:** This will log all the access requests to the S3 bucket. This helps debug any security breaches and secure the S3 bucket and the files.
 - **Tags:** This will tag the bucket using a *Name:Value* style, the same style of tagging as we learned for Lambda functions.
 - **Object-level logging:** This will use the CloudTrail service of AWS to log all the access requests and other details and activities on the S3 bucket. This will also include CloudTrail costs too. So, use this feature only if you need detailed logging. We shall skip using this for this section.
9. After finishing creating the buckets, the S3 console will look like this, with both the created buckets listed:




10. We have successfully created S3 buckets for our task. Now, we have to create a Lambda function that can recognize an object upload in the `sender-bucket` bucket and send that file to the `receiver-bucket` bucket.

11. While creating the Lambda function, this time choose the **s3-get-object-python** blueprint from the listed options available:



12. Configure the bucket details in the next step. In the **Bucket** section, select the **sender-bucket** bucket and select the **Object Created (All)** option in the **Event type** action. This is because we want to send a notification to Lambda whenever an object gets created in the `sender-bucket` bucket. The completed part of the section will look like this:



The screenshot shows the configuration interface for an S3 bucket trigger in the AWS Lambda console. The interface is titled 's3' and includes a 'Remove' button in the top right corner. The configuration is divided into several sections:

- Bucket:** A dropdown menu showing 'sender-bucket'. The text below reads: 'Please select the S3 bucket that serves as the event source. The bucket must be in the same region as the function.'
- Event type:** A dropdown menu showing 'Object Created (All)'. The text below reads: 'Select the events that you want to have trigger the Lambda function. You can optionally set up a prefix or suffix for an event. However, for each bucket, individual events cannot have multiple configurations with overlapping prefixes or suffixes that could match the same object key.'
- Prefix:** A text input field with the placeholder 'e.g. images/'. The text above reads: 'Enter an optional prefix to limit the notifications to objects with keys that start with matching characters.'
- Suffix:** A text input field with the placeholder 'e.g. .jpg'. The text above reads: 'Enter an optional suffix to limit the notifications to objects with keys that end with matching characters.'

At the bottom, a note states: 'Lambda will add the necessary permissions for Amazon S3 to invoke your Lambda function from this trigger. [Learn more](#) about the Lambda permissions model.'

13. Once you have enabled the trigger, Lambda helps you by creating a boilerplate code for the task. All we need to do is write the code to put the object into the receiver-bucket bucket. The boilerplate code can be seen in the **Lambda function code** section:


Lambda function code

Code is pre-configured by the chosen blueprint. You can configure it after you create the function.


Runtime
Python 2.7

```
1 from __future__ import print_function
2
3 import json
4 import urllib
5 import boto3
6
7 print('Loading function')
8
9 s3 = boto3.client('s3')
10
11
12 def lambda_handler(event, context):
13     #print("Received event: " + json.dumps(event, indent=2))
14
15     # Get the object from the event and show its content type
16     bucket = event['Records'][0]['s3']['bucket']['name']
17     key = urllib.unquote_plus(event['Records'][0]['s3']['object']['key'].encode('utf8'))
18     try:
19         response = s3.get_object(Bucket=bucket, Key=key)
20         print("CONTENT TYPE: " + response['ContentType'])
21         return response['ContentType']
22     except Exception as e:
23         print(e)
24         print('Error getting object {} from bucket {}. Make sure they exist and your bucket is in the
25               raise e
26
```

14. When this step has been completed and you have clicked the **Create function** button, you can see the **Triggers** section of the Lambda console, which displays a success message in green at the top:



15. I have uploaded a small image file into the `sender-bucket` bucket. So, now the contents of the **sender-bucket** bucket look like this:



16. As soon as this file had been uploaded, the Lambda function got triggered. The Lambda function code looks like this:

```
from __future__ import print_function

import json
import urllib
import boto3
from botocore.client import Config

print('Loading function')
sts_client = boto3.client('sts', use_ssl=True)

# Assume a Role for temporary credentials
assumedRoleObject = sts_client.assume_role(
    RoleArn="arn:aws:iam::080983167913:role/service-role/Pycontw-
    Role",
    RoleSessionName="AssumeRoleSession1"
)
credentials = assumedRoleObject['Credentials']
region = 'us-east-1'

def lambda_handler(event, context):
    #print("Received event: " + json.dumps(event, indent=2))

    # Get the object from the event and show its content type
    bucket = event['Records'][0]['s3']['bucket']['name']
    key = urllib.unquote_plus(event['Records'][0]['s3']
    ['object']['key']).encode('utf8'))
    try:
        # Creates a session
        session = boto3.Session(credentials['AccessKeyId'],
        credentials['SecretAccessKey'],
        aws_session_token=credentials['SessionToken'],
        region_name=region)

        #Instantiates an S3 resource
        s3 = session.resource('s3',
        config=Config(signature_version='s3v4'), use_ssl=True)

        #Instantiates an S3 client
        client = session.client('s3',
        config=Config(signature_version='s3v4'), use_ssl=True)

        # Gets the list of objects of a bucket
        response = client.list_objects(Bucket=bucket)
```

```

destination_bucket = 'receiver-bucket'
source_bucket = 'sender-bucket'


# Adding all the file names in the S3 bucket in an
array
keys = []
if 'Contents' in response:
    for item in response['Contents']:
        keys.append(item['Key']);

# Add all the files in the bucket into the receiver
bucket
for key in keys:
    path = source_bucket + '/' + key
    print(key)
    s3.Object(destination_bucket,
key).copy_from(CopySource=path)

Exception as e:
    print(e)
print('Error getting object {} from bucket {}. Make sure they
exist and your bucket is in the same region as this
function.'.format(key, bucket))
raise e

```

17. Now, when you run the Lambda function, you can see the same file in the **receiver-bucket** bucket:



SNS trigger

The SNS notification service can be used across multiple use cases, one of which involves triggering Lambda functions. The SNS trigger is popularly used as an interface between the AWS CloudWatch service and Lambda.


So, in this section, we will do the following:

1. Create an SNS topic
2. Create a CloudWatch alarm for our `receiver-bucket` bucket to monitor the number of objects in the bucket
3. Once the objects count reaches 5, the alarm will be set to **ALERT** and the corresponding notification will be sent to the SNS topic that we have just created
4. This SNS topic will then trigger a Lambda function, which prints out a **Hello World** message for us


This will help you understand how to monitor different AWS services and set up alarms for some thresholds for those metrics. And depending on whether the service's metrics have hit that threshold or not, the Lambda function will get triggered.

The process flow for this is as follows:

1. SNS topics can be created from the SNS dashboard. By clicking on the **Create topic** option, you will be redirected to the topic creation dashboard of SNS. The SNS dashboard of AWS looks like this:




The SNS topic creation wizard in the next step looks like this:




The screenshot shows the 'Create new topic' wizard. At the top, it says 'Create new topic'. Below that, a note states: 'A topic name will be used to create a permanent unique identifier called an Amazon Resource Name (ARN)'. There are two input fields: 'Topic name' with a placeholder 'Enter topic name' and an information icon, and 'Display name' with a placeholder 'Enter topic display name. Required for topics with SMS subscriptions.' and an information icon. At the bottom right, there are two buttons: 'Cancel' and 'Create topic'.

In this creation wizard, you can name the SNS topic that you are creating, and add any meta information you want to.


2. Once the topic is created, you can view it in the **Topics** menu, which is on the left of your SNS dashboard. The button looks like this:



Upon clicking the **Topics** tab, a list of topics will be displayed, as shown in the following screenshot:



3. Now that we have successfully created an SNS topic, we shall create a CloudWatch alarm to monitor our S3 bucket for files. The AWS **CloudWatch** dashboard looks something like this:



4. Now, we can go to the **Alarms** page by clicking the **Alarms** button in the list on the left of the dashboard. The AWS **Alarms** page of looks like this:



- Next, click on **Create Alarm** to create an alarm. This will open an alarm creation wizard with multiple options. The wizard looks like this, depending on the services running in your AWS ecosystem:

The screenshot shows the 'Create Alarm' wizard in the AWS Management Console. The title bar says 'Create Alarm' with a close button. Below the title bar, there are two tabs: '1. Select Metric' (active) and '2. Define Alarm'. The main content area is titled 'CloudWatch Metrics by Category'. It shows a summary: 'Your CloudWatch metric summary has loaded. Total metrics: 62'. Below this, there are several categories of metrics listed in a grid:

- EBS Metrics : 18**
 - Per-Volume Metrics : 18
- EC2 Metrics : 24**
 - Per-Instance Metrics : 24
- Lambda Metrics : 12**
 - Across All Functions : 4
 - By Function Name : 4
 - By Resource : 4
- Logs Metrics : 4**
 - Account Metrics : 2
 - Log Group Metrics : 2
- S3 Metrics : 4**
 - Storage Metrics : 4

At the bottom of the wizard, there is a 'Update Graph' button and three view icons (line, area, bar). Below these are four buttons: 'Cancel', 'Previous', 'Next' (highlighted in blue), and 'Create Alarm' (highlighted in blue).

6. As we intend to create an alarm for our S3 bucket, we can go to the **S3 Metrics** tab and ignore the rest of the available metrics. If you click on the **Storage Metrics** option in the **S3 Metrics** category, you will be re-directed to another alarm creation wizard that looks like the following, depending on the number of buckets you have in your S3:

Create Alarm

1. Select Metric

2. Define Alarm

S3

Q Search Metrics

X

1 to 4 of 4 metrics

Storage Metrics X

S3 > Storage Metrics

	BucketName	StorageType	Metric Name
☐	receiver-bucket	AllStorageTypes	NumberOfObjects
☐	receiver-bucket	StandardStorage	BucketSizeBytes
☐	sender-bucket	AllStorageTypes	NumberOfObjects
☐	sender-bucket	StandardStorage	BucketSizeBytes

Update Graph

Time Range

RelativeAbsoluteUTC (GMT)

From: 12.01 hours ago

To: 0 hours ago

Zoom: 1h3h6h12h1d3d1w2w

Select a metric above to view graph

Click a checkbox to select a metric

Click on text to add to search

CancelPreviousNextCreate Alarm

7. If you observe the options in the **Metric Name** column, you will see two options available for each bucket: **NumberOfObjects** and **BucketSizeBytes**. They are self-explanatory and we will only need the **NumberOfObjects** option for the receiver-bucket bucket. So, select that option and click **Next**:

The screenshot shows the 'Create Alarm' console in AWS CloudWatch. The interface is divided into two tabs: '1. Select Metric' (active) and '2. Define Alarm'. A search bar at the top shows 'S3' and 'Search Metrics'. Below the search bar, the 'Storage Metrics' section is expanded, showing a search for 'S3 > Storage Metrics...'. A table lists the available metrics:

BucketName	StorageType	Metric Name
☒ receiver-bucket	AllStorageTypes	NumberOfObjects

Below the table, a graph titled 'Title: NumberOfObjects' is displayed. The graph shows a single data point at 1.00 on the y-axis (labeled 'Count') for the date 11/09 on the x-axis. The graph is set to 'Average' and '1 Day'. To the right of the graph, the 'Time Range' section is expanded, showing 'Relative' time range with 'From: 3.67 days ago' and 'To: 0 days ago'. The 'Zoom' options are '1h', '3h', '6h', '12h', '1d', '3d', '1w', and '2w'. At the bottom right, there are buttons for 'Cancel', 'Previous', 'Next', and 'Create Alarm'.

This will take you to the alarm definition wizard, where you need to specify the details of the SNS topic and the threshold for the alarm. The wizard looks like this:

Create Alarm

1. Select Metric 2. Define Alarm

Alarm Threshold

Provide the details and threshold for your alarm. Use the graph on the right to help set the appropriate threshold.

Name:

Description:

Whenever: NumberOfObjects

is:

for: consecutive period(s)

Alarm Preview

This alarm will trigger when the blue line goes up to or above the red line for a duration of 1 day

Namespace: AWS/S3

BucketName:

StorageType:

Metric Name:

Period:

Statistic: ☒ Standard ☐ Custom

Additional settings

Provide additional configuration for your alarm.

Treat missing data as:

Actions

Define what actions are taken when your alarm changes state.

Notification	Delete
--------------	--------

Cancel Previous Next Create Alarm

8. Add in the details for the threshold and the name of the alarm. The threshold is five files, which means that the alarm will be triggered as soon as the number of files in the corresponding bucket (`receiver-bucket` in our case) reaches a total of five. The wizard looks like this:

Create Alarm

1. Select Metric

2. Define Alarm

Alarm Threshold

Provide the details and threshold for your alarm. Use the graph on the right to help set the appropriate threshold.

Name:

Packet-Test

Description:

Alarm for monitoring number of objects in b

Whenever:

NumberOfObjects

is:

>=

5

for:

1

consecutive period(s)

Alarm Preview

This alarm will trigger when the blue line goes up to or above the red line for a duration of 1 day

NumberOfObjects >= 5

Namespace:

AWS/S3

BucketName:

receiver-bucket

StorageType:

AllStorageTypes

Metric Name:

NumberOfObjects

Period:

1 Day

Statistic:

Standard

Custom

Average

Additional settings

Provide additional configuration for your alarm.

Treat missing data as:

missing

Actions

Define what actions are taken when your alarm changes state.

Notification

Delete

Cancel


Previous

Next


Create Alarm

9. In the **Actions** option, we can configure the alarm to send the notification to the SNS topic that we have just created. You can select the topic from the drop-down list, as follows:

10. Once we have configured the SNS topic, we can click on the blue **Create Alarm** button at the bottom. This will create the alarm that is linked to the SNS topic as a notification pipeline. The created alarm will look like this on the dashboard:



11. Now, we can move on to building the Lambda function for the task. For this particular task, use the **sns-message-python** blueprint while creating our Lambda function:




11. In the previous step, when you have selected the blueprint, you will be asked to enter some meta information regarding your Lambda function, just like we did previously while creating Lambda functions. In the same wizard, you will also be asked to mention the name of the SNS topic. You can specify it here:

The screenshot shows the AWS Lambda console configuration page for an SNS trigger. The page is divided into several sections:

- sns**: A section header with a **Remove** button.
- SNS topic**: A section with the text "Select the SNS topic to subscribe to." Below this is a dropdown menu showing "Packt-Topic" and an **Enter value** button.
- Lambda will add the necessary permissions for Amazon SNS to invoke your Lambda function from this trigger. [Learn more](#) about the Lambda permissions model.**: A section with a heading and a link.
- Enable trigger**: A section with the text "Enable the trigger now, or create it in a disabled state for testing (recommended)." and a checked checkbox.
- Lambda function code**: A section with the text "Code is pre-configured by the chosen blueprint. You can configure it after you create the function."
- Runtime**: A section with the text "Python 2.7".

12. Now that we have selected all the options for the Lambda function correctly, we can now go on to the code. The desired code will look like this:



The screenshot shows the AWS Lambda console's code editor interface. At the top, there are three dropdown menus: 'Code entry type' set to 'Edit code inline', 'Runtime' set to 'Python 2.7', and 'Handler' set to 'lambda_function.lambda_handler'. Below these, the code editor displays the following Python code for `lambda_function.py`:

```
1 from __future__ import print_function
2
3 import json
4
5 print('Loading function')
6
7
8 def lambda_handler(event, context):
9     #print("Received event: " + json.dumps(event, indent=2))
10    message = event['Records'][0]['Sns']['Message']
11    print("From SNS: " + message)
12    print('Hello World')
13    return message
14
```

The code is displayed in a light gray editor with a vertical line number margin on the left. The line containing `message = event['Records'][0]['Sns']['Message']` is highlighted with a blue background. In the bottom right corner of the editor, there are two small icons: a gear for settings and a square with an 'X' for a full-screen view.

The preceding code will display a `Hello World` message whenever the Lambda function gets triggered. This we have completed the setup for this task.

13. To test the preceding setup, you can simply upload more than five files to your `receiver-bucket` bucket and check for Lambda function's execution.

SQS trigger

The **AWS Simple Queue Service (SQS)** is the AWS queue service. This service is similar to the queuing mechanisms that are used generally in software engineering. This enables us to add, store, and remove messages inside the queue.

We will learn how to trigger a Lambda function, depending on the number of messages in a SQS queue. This task will help you understand how serverless batch data architectures can be built and how to build one yourself.

We will do this by monitoring our SQS queue with a CloudWatch alarm and relaying the information to Lambda via an SNS topic, just like we did in the previous task.


So, in this section, we will do the following:

1. Create an SQS queue
2. Create an SNS topic
3. Create a CloudWatch alarm for our SQS queue to monitor the number of messages in the queue
4. Once the messages count reaches 5, the alarm will be set to **ALERT** and the corresponding notification will be sent to the SNS topic we have just created
5. This SNS topic will then trigger a Lambda function, which prints out a `Hello World` message for us

This will help you understand how to monitor queues and build efficient serverless data architectures that are batched, instead of in real time.

The process flow for this is as follows:

1. We will start by creating an AWS SQS queue. We need to go to the SQS dashboard of our AWS account. The dashboard looks like this:



2. Click on the **Get Started Now** button to create an SQS queue. It will redirect you to the queue creation wizard, where you need to enter details such as the name, type of queue, and so on. The queue creation wizard looks like this:

Create New Queue

What do you want to name your queue?

Queue Name ⓘ

Type the queue name.

Region ⓘ US East (N. Virginia)

What type of queue do you need?

Standard Queue

FIFO Queue

High Throughput: Standard queues have nearly-unlimited transactions per second (TPS).

At-Least-Once Delivery: A message is delivered at least once, but occasionally more than one copy of a message is delivered.

Best-Effort Ordering: Occasionally, messages might be delivered in an order different from which they were sent.

First-In-First-out Delivery: The order in which messages are sent and received is strictly preserved.

Exactly-Once Processing: A message is delivered once and remains available until a consumer processes and deletes it. Duplicates are not introduced into the queue.

Limited Throughput: Without batching, FIFO queues can support up to 300 messages per second (300 SendMessage, ReceiveMessage, or DeleteMessage operations per second). If you take advantage of the maximum batching of 10 messages per operation, FIFO queues can

3. You can enter the name of the queue in **Queue Name**. In the **What type of queue do you need?** option, select the **Standard Queue** option. In the options at the bottom, select the blue **Quick>Create Queue** option:

- Allocate tasks to multiple worker nodes: process a high number of credit card validation requests.
- Batch messages for future processing: schedule multiple entries to be added to a database.
- Display the correct product price by sending price modifications in the right order.
- Prevent a student from enrolling in a course before registering for an account.

For more information, see the [Amazon SQS FAQs](#) and the [Amazon SQS Developer Guide](#).

To create a new queue, choose Quick>Create Queue. To configure your queue's parameters, choose Configure Queue.

Cancel Configure Queue Quick>Create Queue

The **Configure Queue** option is for advanced settings. It is not necessary to tweak those settings for this task. This is what the advanced settings look like:

Queue Attributes

Default Visibility Timeout ⓘ

30

seconds ▾

Value must be between 0 seconds and 12 hours.

Message Retention Period ⓘ

4

days ▾

Value must be between 1 minute and 14 days.

Maximum Message Size ⓘ

256

KB

Value must be between 1 and 256 KB.

Delivery Delay ⓘ

0

seconds ▾

Value must be between 0 seconds and 15 minutes.

Receive Message Wait Time ⓘ

0

seconds

Value must be between 0 and 20 seconds.

Dead Letter Queue Settings

Use Redrive Policy ⓘ

☐

Dead Letter Queue ⓘ

Value must be an existing queue name.

Maximum Receives ⓘ

Value must be between 1 and 1000.

Server-Side Encryption (SSE) Settings

Use SSE ⓘ


☐

AWS KMS Customer Master Key (CMK) ⓘ


Data Key Reuse Period ⓘ

This value must be between 1 minute and 24 hours.


5. Once you have created the queue, you will be taken to the SQS page, where all the queues that you have created are listed similarly to the SNS list. This page looks like this:



6. As we have already created an SNS topic in the previous task, we will use the same topic for this purpose. If you haven't created an SNS topic, you can refer to the previous task for instructions on how to create one. The list of SNS topics looks like this:



- Now, we will go to the **CloudWatch** dashboard to create an alarm to monitor our SQS queue and send a notification to Lambda via the SNS topic that we have already created. We can now see the SQS queue metrics in the alarm creation wizard:



- By clicking on the **Queue Metrics** option under **SQS Metrics**, we will be taken to the page where all queue metrics are listed, and we need to select one of them for our alarm:

Create Alarm

1. Select Metric

2. Define Alarm

SQS

Search Metrics

1 to 9 of 9 metrics

Queue Metrics

SQS > Queue Metrics

	QueueName	Metric Name
☐	Packt-Test	ApproximateAgeOfOldestMessage
☐	Packt-Test	ApproximateNumberOfMessagesDelayed
☐	Packt-Test	ApproximateNumberOfMessagesNotVisible
☐	Packt-Test	ApproximateNumberOfMessagesVisible
☐	Packt-Test	NumberOfEmptyReceives
☐	Packt-Test	NumberOfMessagesDeleted
☐	Packt-Test	NumberOfMessagesReceived
☐	Packt-Test	NumberOfMessagesSent
☐	Packt-Test	SentMessageSize

Update Graph

Time Range

RelativeAbsoluteUTC (GMT)

From: 12.04 hours ago

Select a metric above to view graph

CancelPreviousNextCreate Alarm

9. Here, we are interested in the **ApproximateNumberOfMessagesVisible** metric, which gives the number of messages in the queue. It says Approximate, as SQS is a distributed queue and the number of messages can only be determined stochastically.

10. In the next page, after selecting the **ApproximateNumberOfMessagesVisible** metric from the list, the necessary settings can be configured as we did for the **S3 Metrics** in the previous task. The page should look like this:

Create Alarm

[1. Select Metric](#) [2. Define Alarm](#)

Alarm Threshold

Provide the details and threshold for your alarm. Use the graph on the right to help set the appropriate threshold.

Name:

Description:

Whenever: ApproximateNumberOfMessagesVisible

is:

for: consecutive period(s)

Alarm Preview

This alarm will trigger when the blue line goes up to or above the red line for a duration of 5 minutes

Namespace: AWS/SQS

QueueName:

Metric Name:

Period:

Statistic: ☒ Standard ☐ Custom

Additional settings

Provide additional configuration for your alarm.

Treat missing data as:

Actions

Define what actions are taken when your alarm changes state.

[Delete](#)

[Cancel](#) [Previous](#) [Next](#) [Create Alarm](#)

11. In the **Actions** section, configure the SNS topic to which we want to send our notification. This step is also similar to how we configured the SNS topic in the previous task:

Create Alarm

1. Select Metric
2. Define Alarm

Whenever: ApproximateNumberOfMessagesVisible

is:

for: consecutive period(s)

Additional settings

Provide additional configuration for your alarm.

Treat missing data as: ⓘ

Actions

Define what actions are taken when your alarm changes state.

Notification

Whenever this alarm:

Send notification to: [New list](#) [Enter list](#) ⓘ

This notification list is managed in the SNS console.

+ Notification
+ AutoScaling Action
+ EC2 Action

Namespace: AWS/SQS

QueueName:


Metric Name:

Period:


Statistic: ☒ Standard ☐ Custom

Cancel Previous Next **Create Alarm**

12. Once you are satisfied with the metadata and the settings you have configured for the alarm, you can click the blue **Create Alarm** button on the bottom-right side of the screen. That will successfully create an alarm that monitors your SQS queue and sends a notification to the SNS topic that you have configured:



13. We can use the Lambda function that we created in the previous task. Make sure the trigger is the SNS topic that we are using to configure the notification system of the alarm:



14. The Lambda function code for this task is as follows:

```
from __future__ import print_function
import json
print('Loading function')
def lambda_handler(event, context):
    #print("Received event: " + json.dumps(event, indent=2))
    message = event['Records'][0]['Sns']['Message']
    print("From SNS: " + message)
    print('Hello World')
    return message
```

CloudWatch trigger

CloudWatch is the logging and monitoring service for AWS, where logs from most services get stored and monitored. In this section, we will learn how CloudWatch trigger works, how CloudWatch querying works in practice, configuring this in the Lambda function, and also how to make use of this knowledge to build a Lambda function.


So, in this section, we will do the following:

1. Create a CloudWatch log
2. Briefly understand how a CloudWatch log works
3. Create a Lambda function that gets triggered by the CloudWatch trigger


This will help you understand and build resilient and stable serverless architectures.

The process flow for this is as follows:


1. To create a CloudWatch Logs group, click on the **Logs** option to the left of the **CloudWatch** console:




2. Once you are on the AWS CloudWatch **Logs** page, you will see a list of log groups that are already present. The CloudWatch **Logs** page looks something like this:




- Let's go ahead and create a new CloudWatch log. You can see the option to create a new log group from the **Actions** drop-down menu at the top:




- In the next step, you will be asked to name the log group that you are creating. Go ahead and enter the relevant information and click **Create log group**:




5. So, now we have a new log group listed in the list of log groups in our **CloudWatch** console:



6. Once the log group has been created, we can now start working on our Lambda function. So, let's move on to the Lambda console and start creating a new function.
7. From the blueprints, choose the **cloudwatch-logs-process-data** blueprint. The description reads: **A real-time consumer of log events ingested by an Amazon CloudWatch Logs log group**:



8. After selecting the corresponding blueprint option, you will be redirected to the Lambda creation wizard, as usual:




The screenshot shows the AWS Lambda 'Create function' wizard. The breadcrumb trail at the top reads: **Lambda** > **Functions** > **Create function** > Using blueprint cloudwatch-logs-process-data.

The main section is titled **Basic information** with an **Info** link. It contains the following fields:

- Name***: A text input field containing the placeholder text `myFunctionName`.
- Role***: A dropdown menu with the text 'Choose an existing role'. Below it is a smaller dropdown labeled 'Existing role*' with the text 'You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.'


Below the 'Basic information' section is a section titled **cloudwatch-logs** with a **Remove** button on the right. It contains a **Log Group** field with the text 'Please select the CloudWatch Logs log group that serves as the event source. Log Events sent to the log group will trigger your lambda function with the contents of the logs received.' The dropdown menu for this field shows the value `/aws/batch/job`.

9. Just as we did in the previous task, we will also enter relevant information about the log name and other details in the **cloudwatch-logs** pane of the Lambda creation panel:




The screenshot shows the 'cloudwatch-logs' configuration pane in the AWS Lambda console. At the top left is the title 'cloudwatch-logs' and at the top right is a 'Remove' button. The main content area is divided into two sections. The first section, titled 'Log Group', contains the instruction 'Please select the CloudWatch Logs log group that serves as the event source. Log Events sent to the log group will trigger your lambda function with the contents of the logs received.' Below this is a dropdown menu with 'Packt-Test' selected. The second section, titled 'Filter Name', contains the instruction 'Choose a name for your filter.' and a text input field with 'myLambdaTrigger' entered. Below this is a section titled 'Filter Pattern' with the instruction 'Enter an optional Filter Pattern.' and an empty text input field. At the bottom of the pane, there is a note: 'Lambda will add the necessary permissions for Amazon CloudWatch Logs to invoke your Lambda function from this trigger. [Learn more](#) about the Lambda permissions model.' Below this note is a section titled 'Enable trigger' with the instruction 'Enable the trigger now, or create it in a disabled state for testing (recommended).' and an unchecked checkbox.

10. After clicking **Create function**, we will be directed to a **Triggers** page with the success message.



11. So, now we write the Lambda function code to identify the log group and print `Hello World` message:



12. We have now successfully completed another task where we understood how to trigger a Lambda function via AWS CloudWatch Logs. The Lambda function code for this task is as follows:

```
import boto3
import logging
import json
logger = logging.getLogger()
logger.setLevel(logging.INFO)
def lambda_handler(event, context):
    #capturing the CloudWatch log data
    LogEvent = str(event['awslogs']['data'])
    #converting the log data from JSON into a dictionary
    cleanEvent = json.loads(LogEvent)
    print 'Hello World'
    print cleanEvent['logEvents']
```

Summary

In this chapter, we have learned a great deal about how various Lambda triggers work, and how to configure them, set up the triggers, and write Lambda function code to handle the data from them.

In the first task, we learned how S3 events work and how to understand and receive events from the S3 service to AWS Lambda. We have understood how to monitor S3 buckets for file details via their metrics in CloudWatch and then send that notification via AWS SNS to a Lambda functions.

We have also learned how to create SNS topics and how to use them as an intermediate route between several metrics of AWS services from CloudWatch to AWS Lambda.

We have learned briefly about how AWS CloudWatch works. We understood what the metrics of various AWS services, such as S3, SQS, and CloudWatch, look like. We also learned how to set thresholds for CloudWatch Alarms, and how to connect those alarms to notification services, such as AWS SNS.

We learned how AWS CloudWatch Logs work and how to connect and use the CloudWatch trigger in Lambda so it's triggered whenever a new log event is added/received. Overall, we have successfully created new AWS services, such as SQS, CloudWatch Logs, SNS, and S3 buckets in this chapter, and successfully built and deployed three serverless tasks/pipelines.

In the next chapter, we will learn how to build serverless APIs, on which we will perform some tasks just like we did in this chapter, and get a hands-on understanding of how APIs work and, most importantly, how serverless APIs work.

4

Deploying Serverless APIs

So far, we have come a long way in our journey of learning about serverless applications and building serverless engineering. We have learned what the serverless paradigm actually is, how the AWS Lambda function works, understanding the internals of AWS Lambda, along with a detailed understanding of how several triggers work. We have also done several mini projects around experimenting with triggers and deploying them as end-to-end serverless pipelines.

In this chapter, you will be learning how to build efficient and scalable serverless APIs, using the AWS Lambda and AWS API Gateway services. We will start with understanding how the API Gateway works, instead of diving directly to building the serverless API. After that, we will understand how API Gateway and AWS Lambda integrate with each other. And finally, we will be creating and deploying a fully functional serverless API, as part of your learning from this chapter.


This chapter covers the following topics:

- API methods and resources
- Setting up integration
- Deploying the Lambda function for API execution
- Handling authentication and user controls


API methods and resources

In this section, we will be learning about the API service of AWS, which is the API Gateway, and understanding the components and settings available in the console for the user who is creating APIs. We will go through all of the components and understand the API Gateway better. The steps to create the serverless APIs are as follows:


1. We will start by opening the API Gateway console, which looks like this:



2. In the API Gateway console, click on the **Get Started** button to start creating an API. It will take you to an API creation wizard with a popup saying **Create Example API**:




3. Once you click on the **OK** button, you will be redirected to a page where the **Example API** is shown, from which you can get an idea of what an API response looks like:




The API we are building in this example is for a pet store and for maintaining the pets inside the store. By going through the API, you will see what the bits and pieces of an API looks like. The API looks like this:




4. Once you click on the **Import** button at the end, you will be redirected to the **PetStore (b7exp0d681)** API page that we have just created. The API page with all the components looks like this:




5. The resources in this API are the **GET** and **POST** resources, where you can add pets and view the pets, which are available as a list. The list of resources from the API we have created is as follows:




6. By clicking on the first **GET** resource, we can see a detailed execution flow from the client to the endpoint and back to the client. The execution flow of the resource looks like this:




7. Now, if we click on the **POST** resource, we will find a similar model execution flow for the **POST** resource. It looks very similar to that of the **GET** resource, however, here the API endpoint is mentioned as a URL, as we are trying to retrieve the result from it. The execution model looks as follows:




In the API Gateway, there is something called **Stages**, which can be used as versioning models for an API. Some common names for **Stages** in practice are **test**, **development**, and **production**. The **Stages** menu looks like this:



7. When you click on the **Create** option, it will open a creation wizard for the stage. This looks as follows:




8. You can select any name for the **Stage name** value, and add the **Stage description** value according to the name you have assigned and the purpose you have in mind for this stage. Before that, you need to deploy the API that you have created. This can be selected in the **Actions** drop-down menu as the **Deploy API** button, as follows:



9. In the next menu, you can choose the **Stage name** and other details, before finally clicking on the **Deploy** button, which will deploy your API with that particular stage. This can be seen as follows:

A screenshot of the 'Deploy API' dialog box. The dialog has a title bar 'Deploy API' with a close button. Below the title bar, there is a text prompt: 'Choose a stage where your API will be deployed. For example, a test version of your API could be deployed to a stage named beta.' Below this, there are four labeled input fields: 'Deployment stage' (a dropdown menu showing '[New Stage]'), 'Stage name*' (a text input field with 'Prod' entered), 'Stage description' (a text input field), and 'Deployment description' (a text input field). At the bottom right, there are two buttons: 'Cancel' and 'Deploy'.

The deployed stage would look as follows:




Setting up integration


As we now understand how the AWS API Gateway service works at a basic level, we will move on to use that knowledge for building an end-to-end project which involves deploying a completely serverless API.

In this section, we will be building and deploying a completely serverless API function from scratch, along with learning the internals and other implementation details of the AWS Lambda—AWS API Gateway integrations. We will be building the serverless API step-by-step. So, follow along with the steps in this order. The procedure is as follows:


1. Firstly, we will start by creating a new API. This can be done via the Lambda console which looks like this:




2. Once you have clicked on the **+Create API** button, you will be redirected to the API creation wizard, where you will be asked to enter the name and description of the API you are intending to build. For now, I have entered the name as `TestLambdaAPI`. However, you are free to add whatever name and description you would like to enter. The API creation console looks like this:



3. Once you click on the **Create API** button, you will be redirected to the page of the API you have created. The API page would look similar to this:




4. Now that we have successfully created an API, we will now go ahead and create resources in the API. You can do that by clicking on the **Create Resource** option in the **Actions** drop-down menu:




- This would open up a resource creation wizard where you can add the name and resource path of the API resource which we are intending to build. After creating the resource, click on the **Create Resource** button for your settings for the API resource to be created accordingly. For the sake of this tutorial, I have named it `LambdaAPI`. However, you can give it any name you want. The API creation wizard looks like this:

The screenshot shows the 'New Child Resource' wizard in the AWS API Gateway console. On the left, a sidebar shows the 'Resources' section with a tree view containing a root resource '/'. The main panel is titled 'New Child Resource' and includes a sub-header 'Use this page to create a new child resource for your resource.' Below this, there are two checkboxes: 'Configure as proxy resource' (checked) and 'Enable API Gateway CORS' (unchecked). The 'Resource Name*' field is filled with 'LambdaAPI' and the 'Resource Path*' field is filled with '/ lambdaapi'. A help text block explains that path parameters can be added using brackets, such as '{username}' for a parameter named 'username'. At the bottom, there is a '* Required' label and two buttons: 'Cancel' and 'Create Resource'.


The resource that you have just created is now live in the API console; you can see it under the **Resources** section:




- You can create versions of a resource or even just a resource under a resource. Let's go ahead and create one. For this, you need to click on the resource that you have already created. Then, click on the **Create Resource** option in the drop-down menu in the **Actions** menu:




7. This would open up a similar resource creation wizard under the resource which we have already created. You can name that resource as `version1` or just as `v1` which is a regular software practice. I have named it `v1`. However, you can name it whatever you want to:




Now, we have a resource named `v1` under the already existing resource, `/lambdaapi`. We can see this under our **Resources** section. So, now the resources hierarchy of our API looks like this:




8. We will be creating a serverless API for getting and querying the list of pets in a pet store. So, the following steps will be aligned accordingly. The API should return the name of the pets. So, we will have a new resource for pets for that purpose. We will be creating a resource for this under the `/v1` resource:




9. The resulting hierarchical structure for our API looks like this, after adding the `/pets` resource under the `/v1` resource:




10. Now, we will add a custom resource which enables us to query the API. By custom, we mean that any string can be added to the resource when sending a request to this API, and the API would send back a request after checking and querying for that string via a Lambda code. The custom resources can be differentiated from the normal ones, as they can be created with curly braces. The following screenshot will help you understand how to create them:




11. After clicking on the **Create Resource** button, the new custom child resource for `/pets` will be created. The hierarchy of the resources is now as follows:




12. The overall structure of the API looks like this, as specified in the top-right part of the following screenshot:




13. Now, we will add methods to this custom resource. As we will only be querying the list of pets, we will only add the **GET** method. This can be done by clicking on the `{type}` resource and clicking on **Create Method** in the drop-down **Actions** menu in the top panel:



14. This would create a small drop-down style menu under the **{type}** resource where you can select a method from the available methods:



15. We need to select the **GET** option from the available options. This would look as follows:



16. After selecting the **GET** option and clicking on the small tick button beside it, we will have created the **GET** method under our **{type}** resource. The hierarchy now looks like this:




Deploying the Lambda function for API execution

In this section, we will have a look at the steps to deploy the Lambda function:

1. The details of the **GET** method can also be seen on the right-hand side of the API console, when you click on that method. The details look as follows:

The screenshot displays the AWS API Gateway console interface for configuring a new method. On the left, a 'Resources' sidebar shows a tree structure: '/' expanded, then '/lambdaapi', '/v1', '/pets', and finally '/{type}' with the 'GET' method selected. The main panel is titled '/lambdaapi/v1/pets/{type} - GET - Setup'. Below the title, it says 'Choose the integration point for your new method.' The 'Integration type' section has four radio buttons: 'Lambda Function' (selected), 'HTTP', 'Mock', 'AWS Service', and 'VPC Link'. Below this, the 'Use Lambda Proxy integration' checkbox is unchecked. The 'Lambda Region' is shown as a dropdown menu. The 'Use Default Timeout' checkbox is checked. A 'Save' button is located at the bottom right of the main panel.

2. In the **GET** method console, click on the **Lambda Function** option. Select any one region depending on your preference. I have chosen **us-east-1** as the region as shown in the following screenshot:



3. As expected, it says we do not have a Lambda function in that region. So, we need to go ahead and create one. Click on the **Create a Lambda Function** link. This will take you to the Lambda creation console which we are already comfortable with:

Lambda > Functions > Create function

Create function

Author from scratch ☒
Start with a simple "hello world" example.


Blueprints ☐
Choose a preconfigured template as a starting point for your Lambda function.


Author from scratch [Info](#)

Name*

Runtime*

Node.js 6.10 ▼

Role*

Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.

Choose an existing role ▼

4. From here, choose the **keyword : hello-world-python** blueprint from the list of blueprints:

Lambda > Functions > Create function

Create function

Author from scratch ☐
Start with a simple "hello world" example.


Blueprints ☒
Choose a preconfigured template as a starting point for your Lambda function.


Blueprints [Info](#) Export

ⓘ

keyword : hello-world-python ⓘ

< 1 >

hello-world-python ☐
A starter AWS Lambda function.
python2.7

hello-world-python3 ☐
A starter AWS Lambda function.
python3.6

Cancel Configure

5. In the next console, choose the basic information for the Lambda function as we have done in the previous chapters:

The screenshot shows the 'Basic information' tab in the AWS Lambda console. It includes a 'Name*' field with the placeholder text 'myFunctionName'. Below it is a 'Role*' section with a description: 'Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.' There is a dropdown menu labeled 'Choose an existing role'. At the bottom, there is an 'Existing role*' section with a description: 'You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.' and another dropdown menu.


6. After adding the relevant details, click on the orange **Create function** button. That will take you to the page of the Lambda function you have just created. The code can be edited from there onwards:

The screenshot shows the configuration page for a Lambda function named 'serverless-api'. At the top, there is a success message: 'Congratulations! Your Lambda function "serverless-api" has been successfully created. You can now change its code and configuration. Click on the "Test" button to input a test event when you are ready to test your function.' Below this, there are tabs for 'Configuration' and 'Monitoring'. The 'Configuration' tab is active. On the left, there is a list of 'Add triggers' including API Gateway, AWS IoT, Alexa Skills Kit, Alexa Smart Home, CloudFront, CloudWatch Events, and CloudWatch Logs. In the center, there is a dashed box with the text 'Add triggers from the list on the left'. On the right, there is a list of resources that the function's role has access to, including 'All', 'AWS CloudFormation', and 'Amazon CloudWatch Logs'. At the bottom right, there is a dashed box with the text 'Resources the function's role has access to will be shown here'.

7. In the function's code, use this code instead of the one which is provided along with the blueprint:

```
1 def lambda_handler(event, context):
2     mobs = {
3         "Sea": ["GoldFish", "Turtle", "Tortoise", "Dolphin", "Seal"],
4         "Land": ["Labrador", "Cat", "Dalmatian", "German Shepherd",
5                 "Beagle", "Golden Retriever"],
6         "Exotic": ["Iguana", "Rock Python"]
7     }
8
9     return {"type": mobs[event['type']]}
```

8. We are now done with tweaking the function code. Now, you can go ahead and save the function:



9. Now, head back to the API Gateway console to the **GET** method page. Here, under the Lambda functions in the **us-east-1** region, I start getting the Lambda function which I have just created (**serverless-api**) as an option:

/lambdaapi/v1/pets/{type} - GET - Setup

Choose the integration point for your new method.

Integration type ☒ Lambda Function ⓘ

☐ HTTP ⓘ

☐ Mock ⓘ

☐ AWS Service ⓘ

☐ VPC Link ⓘ

Use Lambda Proxy integration ☐ ⓘ

Lambda Region us-east-1

Lambda Function s ⓘ

Use Default Timeout **serverless-api**

10. On clicking **Save**, you will see a popup asking you to confirm that you are giving API Gateway permission to invoke your Lambda function, you can acknowledge it by clicking on **OK**:


Add Permission to Lambda Function ✕

You are about to give API Gateway permission to invoke your Lambda function:

arn:aws:lambda:us-east-1:080983167913:function:serverless-api

Cancel OK

11. After clicking on **OK**, you will be redirected to the data flow page of the **GET** method, that looks like this:



Handling authentication and user controls


After deploying, next we will discuss how to handle the authentication and user controls. The steps are as follows:

1. Now that we have successfully created the skeleton of our serverless API, we will now work on the nitty-gritty details which are needed to make it a fully functional API. We will start with applying the mapping templates. This can be done in the **Integration Request** menu. Clicking on the **Integration Request** link will take you to a console which looks like this:


The screenshot shows the 'Integration Request' configuration page in the AWS Lambda console. The breadcrumb navigation at the top reads: **Method Execution** /lambdaapi/v1/pets/{type} - GET - Integration Request. Below the breadcrumb is a header instruction: 'Provide information about the target backend that this method will call and whether the incoming request data should be modified.' The configuration options are as follows:

- Integration type:** A radio button selection with 'Lambda Function' selected. Other options are 'HTTP', 'Mock', 'AWS Service', and 'VPC Link'. Each option has an information icon.
- Use Lambda Proxy integration:** A checkbox that is currently unchecked, with an information icon.
- Lambda Region:** A text field containing 'us-east-1' with a copy icon.
- Lambda Function:** A text field containing 'serverless-api' with a copy icon.
- Invoke with caller credentials:** A checkbox that is currently unchecked, with an information icon.
- Credentials cache:** A text field containing 'Do not add caller credentials to cache key' with a copy icon.
- Use Default Timeout:** A checkbox that is currently checked, with an information icon.
- URL Path Parameters:** A section header with a right-pointing arrow icon.


2. If you scroll down a bit in the same console, you will notice the **Body Mapping Templates** section at the end:




3. Clicking on the **Body Mapping Templates** will unfurl the options available in that particular section:




4. Select the second option which says **When there are no templates defined (recommended)**. And then, click on the **Add mapping template** option and add `application/json`, and click on the small grey tick symbol beside it:



5. After clicking the small grey tick symbol beside it, the **Body Mapping Templates** section space will look like this:



6. Now, in the template textbox, add the following code and click the **Save** button underneath the text box:



The screenshot shows a web-based template editor. At the top, there is a text input field containing the text "application/json". Below this is a label "Generate template:" followed by a dropdown menu. The main area is a code editor with a light gray background and a vertical line number margin on the left showing 1, 2, and 3. The code is a JSON object:

```
{  
  "type": "$input.params('type')"  
}
```


 The third line is highlighted. At the bottom right of the code editor is a small icon with two arrows pointing outwards. Below the code editor are two buttons: "Cancel" and "Save".

7. So, after all these steps, the resulting **Body Mapping Templates** section will look like this:

The screenshot shows the 'Body Mapping Templates' section in the AWS Lambda console. It features three radio buttons for 'Request body passthrough': 'When no template matches the request Content-Type header', 'When there are no templates defined (recommended)' (which is selected), and 'Never'. Below these is a 'Content-Type' dropdown menu currently showing 'application/json'. An 'Add mapping template' button is located below the dropdown. A text input field also contains 'application/json'. Under the 'Generate template:' label, a code editor displays a JSON snippet:

```
1 {  
2   "type": "$input.params('type')"  
3 }
```

8. Now, going back to the **Method Execution** page, we can see the **TEST** option on the left with a lightning bolt symbol beneath it:



9. Clicking on the **TEST** button on the left-side in the **Client** section and above the thunderbolt option will take you to a page where you can test the API that you've just created:

← Method Execution

/lambdaapi/v1/pets/{type} - GET - Method Test

Make a test call to your method with the provided input

Path

{type}

Value

Request: /lambdaapi/v1/pets/

Status: 200

Latency: 56 ms

Response Body

Query Strings

No query string parameters exist for this method. You can add them via Method Request.

Headers

No header parameters exist for this method. You can add them via Method Request.

Stage Variables

No [stage variables](#) exist for this method.

Request Body

Request Body is not supported for GET methods.

```
{
  "stackTrace": [
    [
      "/var/task/lambda_function.py",
      9,
      "lambda_handler",
      "return {\\"type\\": mobs[event['type']]}"
    ]
  ],
  "errorType": "KeyError",
  "errorMessage": "u'"
}
```

Response Headers

```
{ "X-Amzn-Trace-Id": "sampled=0;root=1-5a39d711-b4d8112d8b149c35528c3b64", "Content-Type": "application/json" }
```

- Now, let's type `Exotic` in the textbox below `{type}` and click on the **Test** button at the bottom. If everything goes right, we should see the list of all the exotic pets we have entered in the function code of our Lambda function:

← Method Execution /lambdaapi/v1/pets/{type} - GET - Method Test

Make a test call to your method with the provided input

Path
Request: /lambdaapi/v1/pets/Exotic
Status: 200
Latency: 134 ms
Response Body

{type}
Exotic

Query Strings
No query string parameters exist for this method. You can add them via Method Request.

Headers
No header parameters exist for this method. You can add them via Method Request.

Stage Variables
No [stage variables](#) exist for this method.

Request Body
Request Body is not supported for GET methods.

```
{
  "type": [
    "Iguana",
    "Rock Python"
  ]
}
```

Response Headers


```
{
  "X-Amzn-Trace-Id": "sampled=0;root=1-5a39d787-7a48979b982f6f2240353ce1",
  "Content-Type": "application/json"
}
```

Logs

Execution log for request test-request
Wed Dec 20 03:22:47 UTC 2017 : Starting executi

- And rightly so, we did get the list of all of the exotic pets in the catalog. So, this brings this chapter to an end, where you have learned how to build a fully fledged serverless API from scratch, including how to deploy it.

12. In addition, if you want to add additional security settings, such as **Authorizations** and **API Key Required**, you can do it in the **Method Request** menu:



Summary

In this chapter, we have learned how to build a completely serverless API from scratch. We have also learned how to add more resources and methods for the API, as well as how to deploy it successfully to multiple stages of development and how to add additional security settings such as authorization and API keys for authentication purposes.

We then learned how to associate a Lambda function with our API Gateway's API service for handling the computational tasks of our API.

In the next chapter, we will be learning about logging and monitoring serverless applications. In that chapter, we will learn about the logging and monitoring services of AWS such as CloudWatch Metrics, CloudWatch Logs, and CloudWatch Dashboards in detail, and try to set them up for our serverless applications. We will also learn how to create a logging and monitoring pipeline from AWS Lambda to these monitoring tools using some AWS services.

5

Logging and Monitoring

We have learned about the concepts of serverless architectures and understood the basics and the internals of AWS's serverless service, AWS Lambda. We have also created some example serverless projects to understand the concepts better. During the course of our learning, we have also learned the basics of several other AWS services, such as alarms, SNS, SQS, S3 buckets, and CloudWatch.

In this chapter, we will learn about how to do the logging and monitoring for the serverless systems that we are building. Logging and monitoring software code and systems are very important, as they help us with the telemetry and disaster recovery. Logging is a process where we store the logs emitted by our code or by our architecture as a whole. Monitoring is a process where we closely monitor the activities, status, and health of the components and processes in our code or architecture.

So, you will be learning how to set up and understand the monitoring suite of AWS Lambda, which is closely integrated with the monitoring service of AWS, the CloudWatch Dashboards. We will also learn about the logging service of AWS, the CloudWatch Logs service. Finally, we will also learn about and understand the distributed tracing and monitoring service of AWS, the CloudTrail service.

This chapter covers the following topics:


- Understanding CloudWatch
- Understanding CloudTrail
- Lambda's metrics in CloudWatch
- Lambda's logs in CloudWatch
- Logging statements in Lambda

Understanding CloudWatch

As mentioned earlier, CloudWatch is the logging and monitoring service of AWS. We have already looked at and learned about the CloudWatch Alarms, which are a sub-feature of CloudWatch. We will now learn about the graphing suite of the service. Almost every service in the AWS environment has a way to send its logs and metrics to CloudWatch for logging and monitoring purposes. Each service might have several metrics which can be monitored, depending on the function.

Similarly, AWS Lambda also has some metrics, such as the invocation count, the invocation's running time, and so on, which it sends to CloudWatch. It is also helpful to note that the developers can also send custom metrics to CloudWatch. So in the following steps, we shall be learning about the different parts and functions of AWS CloudWatch corresponding to AWS Lambda:

1. Firstly, let us see what the CloudWatch console looks like and also get a feel for it by navigating around the console. Browse to `console.aws.amazon.com/cloudwatch/`:




2. As we can see, there is a lot of information in the CloudWatch console. So, we shall now try to understand each component one after the other. In the left side, we can see a list of options, which includes **Dashboards**, **Alarms**, **Billing**, and so on. We shall try to understand all of them and their functionality as part of understanding the CloudWatch console.


3. A dashboard here is a panel of CloudWatch Metrics that the user can configure. For example, a user might want to have a particular set of server (EC2) metrics at a single place to be able to monitor them better. This is where AWS CloudWatch Dashboards come into play. When you click on the **Dashboards** option on the left, you can see the **Dashboards** console, which looks like this:




4. Let us go ahead and create a new dashboard by clicking the blue **Create dashboard** button on the top left-hand side of the console. The following box appears:




- This will take you to the next step, where you will be asked to select a widget type for the dashboard. There are four types of widgets which are currently available. The widget selection screen looks like this:




- For the sake of this tutorial, I am choosing the **Line** style widget. You can choose whatever widget would fit your graphing style and the monitoring you need to do. Once you select a widget style and click the blue **Configure** button, you will be redirected to a wizard where you will be asked to add a metric as shown in the following screenshot:




7. Select one of the available metrics at the bottom and it will be added to the widget. Once you are done with selecting the metrics, click on the blue **Create widget** button in the lower-right part of the page as shown in the following screenshot:



8. Now, you can see the dashboard that you have just created in the Dashboards section:



9. We have successfully learned and created an AWS CloudWatch Dashboard. We will now move on to learning about CloudWatch Events. We have already learned about CloudWatch Alarms in the previous chapters, looking at both their functionality and how to create and work with them.
10. Click on the **Events** link in the **CloudWatch** menu on the left. You will be redirected to the page of CloudWatch Events, as shown in the following screenshot:



11. Once you click on the blue **Create rule** button, you will be redirected to the **Events** creation wizard, which looks like this:

Step 1: Create rule

Create rules to invoke Targets based on Events happening in your AWS environment.

Event Source
Build or customize an Event Pattern or set a Schedule to invoke Targets.

☒ Event Pattern ⓘ ☐ Schedule ⓘ

Build event pattern to match events by service

Service Name

Event Type

▼ Event Pattern Preview [Copy to clipboard](#) [Edit](#)

```
{ }
```

Targets
Select Target to invoke when an event matches your Event Pattern or when schedule is triggered.

[+ Add target*](#)

12. There can be two types of events, namely **Event Pattern** and **Schedule**, each of which have different purposes. Here we will only learn about the **Schedule** type, as it comes in handy for scheduling Lambda functions:

Step 1: Create rule

Create rules to invoke Targets based on Events happening in your AWS environment.

Event Source
Build or customize an Event Pattern or set a Schedule to invoke Targets.

☐ Event Pattern ⓘ ☒ Schedule ⓘ

☒ Fixed rate of

☐ Cron expression

[Learn more about CloudWatch Events schedules.](#)

► Show sample event(s)


Targets
Select Target to invoke when an event matches your Event Pattern or when schedule is triggered.

[+ Add target*](#)

* Required


[Cancel](#) [Configure details](#)

13. The rate can be either set in terms of **Minutes**, **Hours**, or **Days**, or can be set as a cron expression, whichever way you are comfortable with. Now, the target needs to be selected. The target can be any valid Lambda function, as shown in the following drop-down menu:




The screenshot shows the 'Targets' section of the AWS Lambda console. It includes a heading 'Targets' and a sub-heading 'Select Target to invoke when an event matches your Event Pattern or when schedule is triggered.' Below this is a 'Lambda function' dropdown menu. The 'Function*' dropdown is open, showing 'Select function' as the current selection. A search bar is visible, and the results list 'serverless-api' as the selected function. Below the dropdown are links for 'Configure versioning' and 'Configure input'. At the bottom is an 'Add target*' button.

14. Once you have selected the function, you can click on the blue **Configure details** at the bottom. It will take you to the **Configure rule details** page as shown in the following screenshot:




The screenshot shows the 'Step 2: Configure rule details' page. It has a heading 'Step 2: Configure rule details' and a sub-heading 'Rule definition'. Below this are fields for 'Name*' and 'Description'. The 'Name*' field is empty, and the 'Description' field is empty. Below these fields is a 'State' section with a checkbox labeled 'Enabled' which is checked. Below the 'State' section is a blue box containing the text: 'CloudWatch Events will add necessary permissions for target(s) so they can be invoked when this rule is triggered.' At the bottom left is a '* Required' label. At the bottom right are three buttons: 'Cancel', 'Back', and 'Create rule'.

15. Once you enter the name and the description of the rule that you want to create, you can click on the blue **Create rule** button at the bottom. This will successfully create an event, and the same will be reflected in your CloudWatch console:



We have successfully added a cron event for a Lambda function which means that Lambda will be invoked at regular intervals, as specified by the user in the settings of the event.


16. Now, we shall try to understand the **Logs** feature of AWS CloudWatch. This is where the Lambda functions store their logs. You can click on the **Logs** link in the menu on the left-hand side to access the console of CloudWatch Logs:




17. We can see the complete list of logs for all of the Lambda functions we have ever created throughout the course of the book. When you click on a log group, you can find more details about it, and also options for customization. Each log stream is an invocation of the Lambda function that the log is associated with:




18. You can also make use of the additional functionality provided by CloudWatch for handling the logs data, which can be seen in the drop-down **Actions** menu in **Log Groups**:



19. Finally, we will wrap up by exploring and learning about the CloudWatch Metrics. The metrics console can be accessed by clicking on the **Metrics** option on the left-hand side of the CloudWatch console:



20. You can select any option in the menu at the bottom for graphing the metrics. For the purpose of this tutorial, I have added a Lambda metric, which is the number of errors in the function, `serverless-api`:




Understanding CloudTrail

CloudTrail is another monitoring service of AWS where you can look at all of the events and trails that have happened in your AWS account. This service is a bit more detailed than the CloudWatch service in how it records and stores the events and trails.

So, we shall explore and learn about this service in the following steps:

1. The AWS CloudTrail's dashboard can be accessed at

`console.aws.amazon.com/cloudtrail/`:



CloudTrail

Dashboard

Event history

Trails

Welcome to CloudTrail

With CloudTrail, you can view events for your AWS account. Create a trail to retain a record of these events. With a trail, you can also create event metrics, trigger alerts, and create event workflows. [Learn more](#)

[Create trail](#)

Recent events

These are the most recent events recorded by CloudTrail. To view all events for the last 90 days, go to Event history.

	Event time	User name	Event name	Resc
▶	2018-01-15, 10:05:58 PM	root	RunInstances	EC2 :
▶	2018-01-15, 09:33:39 PM	root	RunInstances	IAM I
▶	2018-01-15, 09:01:20 PM	root	RunInstances	IAM I
▶	2018-01-15, 08:29:01 PM	root	RunInstances	IAM I
▶	2018-01-15, 08:07:53 PM	serverless-api	CreateLogStream	

[View all events](#)

What's new

December 12, 2017
[AWS CloudTrail Enhances Event History View and Search](#)

[View all updates](#)

Learn more

[Pricing](#)
[Documentation](#)
[Forums](#)
[FAQs](#)

2. The list of events in your AWS account can be seen on the left-hand side of the **CloudTrail** menu when you click on the **Event history** button. The **Event history** page looks like this:

Welcome to CloudTrail

With CloudTrail, you can view events for your AWS account. Create a trail to retain a record of these events. With a trail, you can also create event metrics, trigger alerts, and create event workflows. [Learn more](#)

[Create trail](#)

Recent events

These are the most recent events recorded by CloudTrail. To view all events for the last 90 days, go to Event history.

	Event time	User name	Event name	Resource
▶	2018-01-15, 10:26:31 PM	root	ConsoleLogin	
▶	2018-01-15, 10:05:58 PM	root	RunInstances	EC2 :
▶	2018-01-15, 09:33:39 PM	root	RunInstances	IAM I
▶	2018-01-15, 09:01:20 PM	root	RunInstances	IAM I
▶	2018-01-15, 08:29:01 PM	root	RunInstances	IAM I

[View all events](#)

What's new

December 12, 2017
 AWS CloudTrail Enhances Event History View and Search
[View all updates](#)

Learn more

[Pricing](#)
[Documentation](#)
[Forums](#)
[FAQs](#)

- The third functionality of CloudTrail is the trails. The user can set up trails for their AWS services, such as Lambda. The trails that have been set up can be found on the **Trails** dashboard. This can be accessed by going to the **Trails** console by clicking on the **Trails** option in the menu on the left-hand side:

Trails

Deliver logs to an Amazon S3 bucket. CloudTrail events can be processed by one trail for free. There is a charge for processing events with additional trails. For more information, see [AWS CloudTrail Pricing](#).

[Create trail](#)

Name ▲	Region ▼	S3 bucket ▲	Log file prefix ▲	CloudWatch Logs Log group ▲	Status ▲
No trails have been added					

Learn more

[Pricing](#)
[Documentation](#)
[Forums](#)
[FAQs](#)

4. Now, let us understand how to create a trail in the **CloudTrail** dashboard. You can go to the main dashboard of CloudTrail and click on the blue **Create trail** button. This will take you to the trail creation wizard:

Create Trail

Trail name*

Apply trail to all regions ☒ Yes ☐ No ⓘ

Management events

Management events provide insights into the management operations that are performed on resources in your AWS account. [Learn more](#)

Read/Write events ☒ All ☐ Read-only ☐ Write-only ☐ None ⓘ

Data events

Data events provide insights into the resource operations performed on or within a resource. Additional charges apply. [Learn more](#)

S3 **Lambda**

You can record S3 object-level API activity (for example, GetObject and PutObject) for individual buckets, or for all current and future buckets in your AWS account. Additional charges apply. [Learn more](#)

Showing 0 of 0 resources

Learn more

- [Pricing](#)
- [Documentation](#)
- [Forums](#)
- [FAQs](#)

5. You can enter the details of your trail here. You can leave the default options as they are for the **Apply trail to all regions** and the **Management events** options:

Create Trail

Trail name*

Apply trail to all regions ☒ Yes ☐ No ⓘ

Management events

Management events provide insights into the management operations that are performed on resources in your AWS account. [Learn more](#)

Read/Write events ☒ All ☐ Read-only ☐ Write-only ☐ None ⓘ

Learn more

- [Pricing](#)
- [Documentation](#)
- [Forums](#)
- [FAQs](#)

6. Now, moving on to the next setting, select the **Lambda** option and click on the **Log all current and future functions** in the options list. This will ensure that all of our Lambda functions are logged properly with CloudTrail:

Data events

Data events provide insights into the resource operations performed on or within a resource. Additional [charges](#) apply. [Learn more](#)

S3

Lambda

You can record Invoke API operations for individual functions, or for all current and future functions in your AWS account. Additional [charges](#) apply for trails that include data events. [Learn more](#)

Filter by function name and region or search by keyword ✕

⏪ < Viewing 1 - 1 of 1 functions > ⏩

Function (1 selected) ▾	Region ▾
☒ Log all current and future functions ⓘ	All regions
☒ serverless-api	US East (N. Virginia)

[+ Add function](#)

7. Now, in the final **Storage location** option, select an **S3 bucket** for storing the CloudTrail logs. This can be an already existing bucket or you can also ask CloudTrail to create a new bucket for this purpose. I am using an existing bucket:

Storage location

Create a new S3 bucket ☐ Yes ☒ No

S3 bucket* ⓘ


[▶ Advanced](#)

* Required field


Additional charges may apply ⓘ

Create

8. After all of the details and settings have been configured accordingly, you can click on the blue **Create trail** button to create the trail. Now, you see the trail you have just created in your **CloudTrail** dashboard as shown in the following screenshot:



9. Now, when you click on the trail that you have just created, you can see all of the details with which it has been configured as shown in the following screenshot:



10. You can also notice a very interesting option that enables you to configure CloudWatch Logs along with SNS to notify you of any specific activities, for example when there is an error in a Lambda function:

Data events provide insights into the resource operations performed on or within a resource. Additional [charges](#) apply. [Learn more](#)

S3 **Lambda**

You can record S3 object-level API activity (for example, GetObject and PutObject) for individual buckets, or for all current and future buckets in your AWS account. Additional [charges](#) apply. [Learn more](#)

Configure

▼ Storage location 

S3 bucket [receiver-bucket](#) **✓ Last log file delivered** 2018-01-16, 7:51 am

Encrypt log files No

Enable log file validation Yes

Publish to SNS No

▼ CloudWatch Logs 

Configuring delivery to CloudWatch Logs enables you to receive SNS notifications from CloudWatch when specific API activity occurs. Standard CloudWatch and CloudWatch Logs [charges](#) will apply. [Learn more](#).

Configure

11. And finally, you can also add tags to the trail, just like you can with the rest of your AWS services:

▼ Tags 

Key	Value
No tags added	

12. Additionally, let us understand how to configure CloudWatch Logs for our trail. So, for this you need to click on the blue **Configure** button in the **CloudWatch Logs** section above the **Tags** section:

The screenshot shows the 'CloudWatch Logs' section of the AWS CloudTrail console. It includes a heading 'CloudWatch Logs' with a dropdown arrow. Below it, a paragraph explains that configuring delivery to CloudWatch Logs enables receiving SNS notifications from CloudWatch when specific API activity occurs, and that standard CloudWatch and CloudWatch Logs charges will apply, with a 'Learn more' link. A form field labeled 'New or existing log group*' contains the text 'CloudTrail/DefaultLogGroup' and an information icon. At the bottom, there is a '* Required field' label, a 'Cancel' button, and a blue 'Continue' button.


13. When you click **Continue**, it takes you to the creation wizard where you need to configure the permissions accordingly with your IAM role settings. For the purpose of this tutorial, I have selected the **Create a new IAM Role** option as shown in the following screenshot:

The screenshot shows the 'AWS CloudTrail will deliver CloudTrail events associated with API activity in your account to your CloudWatch Logs log group' page. It includes a heading 'AWS CloudTrail will deliver CloudTrail events associated with API activity in your account to your CloudWatch Logs log group'. Below it, a paragraph explains that in order to successfully deliver CloudTrail events to your CloudWatch Logs log group, CloudTrail will assume the role you are creating or specifying, and that assuming the role grants CloudTrail permissions to two CloudWatch Logs API calls: 1. CreateLogStream: Create a CloudWatch Logs log stream in the CloudWatch Logs log group you specify, and 2. PutLogEvents: Deliver CloudTrail events to the CloudWatch Logs log stream. A 'Hide Details' link is present. The 'Role Summary' section shows the 'Role Description' as 'AWS CloudTrail will assume the role you create or specify to deliver CloudTrail events to your CloudWatch Logs log group'. The 'IAM Role' dropdown is set to 'Create a new IAM Role'. The 'Role Name' field contains 'CloudTrail_CloudWatchLogs_Role'. A 'Hide Policy Document' link is present. The 'Policy Document' section shows a JSON policy snippet:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AWSCloudTrailCreateLogStream20141101",
      "Effect": "Allow",
```

 An 'Edit' link is next to the policy document. At the bottom, there are 'Cancel' and 'Allow' buttons.

14. After you have finished configuring the IAM role settings, you can click on the blue **Allow** button at the bottom. After a couple of seconds of validation, the CloudWatch Logs get configured, which you can see in the same **CloudWatch Logs** section here:




Lambda's metrics in CloudWatch

As we have learned and understood how the CloudWatch and the CloudTrail services work with respect to logging and monitoring, we shall move on to try and implement them for our Lambda function(s). In this section, you will learn about the types of metrics that Lambda possesses, which are monitored by CloudWatch, and how to create a dashboard with all those metrics.


Similar to previous sections in this chapter and book, we shall try and understand the concepts in the form of the following steps:

1. When you navigate over to your AWS Lambda console, you will see the Lambda function which you have already created, in the list of available functions:



2. When you click on the function, you will see two available options on the top: **Configuration** and **Monitoring**. Navigate to the **Monitoring** section. You will see a dashboard of metrics, which contains the following:

- **Invocations**
- **Duration**
- **Errors**
- **Throttles**
- **Iterator age**
- **DLQ errors**



Invocations and duration




Errors and Throttles



Iterator age and DLQ errors

3. Let us understand each of them in detail. The first metric is the **Invocations** metric, which has the time on the x axis and the number of invocations of the Lambda function on the y axis. This metric helps us understand when and how many times our Lambda function has been invoked:



When you click **Jump to Logs**, it takes you to the CloudWatch Logs console of the Lambda invocations, which looks like this:


CloudWatch > Log Groups > /aws/lambda/serverless-api > All streams

Expand all ☒ Row ☐ Text   


Filter events 2018-01-16 (03:00:00) - 2018-01-17 (03:00:00) 

Time (UTC +00:00)	Message	Show in stream
2018-01-16		
No older events found for the selected date range. Adjust the date range.		
03:14:23	START RequestId: f7ae6bfe-fa6a-11e7-a72b-df03dec5c65c Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...
03:14:23	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9	2018/01/15/[\$LATEST]9c6446e...
03:14:23	END RequestId: f7ae6bfe-fa6a-11e7-a72b-df03dec5c65c	2018/01/15/[\$LATEST]9c6446e...
03:14:23	REPORT RequestId: f7ae6bfe-fa6a-11e7-a72b-df03dec5c65c Duration: 0.43 ms Billed Dura	2018/01/15/[\$LATEST]9c6446e...
03:16:38	START RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...
03:16:38	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9	2018/01/15/[\$LATEST]9c6446e...
03:16:38	END RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a	2018/01/15/[\$LATEST]9c6446e...
03:16:38	REPORT RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Duration: 0.43 ms Billed Durr	2018/01/15/[\$LATEST]9c6446e...
03:17:44	START RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...
03:17:44	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9	2018/01/15/[\$LATEST]9c6446e...
03:17:44	END RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a	2018/01/15/[\$LATEST]9c6446e...
03:17:44	REPORT RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Duration: 0.55 ms Billed Durr	2018/01/15/[\$LATEST]9c6446e...
03:19:40	START RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...
03:19:40	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9	2018/01/15/[\$LATEST]9c6446e...
03:19:40	END RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a	2018/01/15/[\$LATEST]9c6446e...
03:19:40	REPORT RequestId: aa8aff17-fa6b-11e7-95c4-e53c3d61044a Duration: 0.43 ms Billed Durr	2018/01/15/[\$LATEST]9c6446e...
03:21:39	START RequestId: 5d7685f0-fa6c-11e7-91e9-7352335265f1 Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...
03:21:39	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9	2018/01/15/[\$LATEST]9c6446e...
03:21:39	END RequestId: 5d7685f0-fa6c-11e7-91e9-7352335265f1	2018/01/15/[\$LATEST]9c6446e...
03:21:39	REPORT RequestId: 5d7685f0-fa6c-11e7-91e9-7352335265f1 Duration: 15.43 ms Billed Du	2018/01/15/[\$LATEST]9c6446e...
03:22:33	START RequestId: 5d7685f0-fa6c-11e7-91e9-7352335265f1 Version: \$LATEST	2018/01/15/[\$LATEST]9c6446e...


And when you click on the **Jump to Metrics** option, it will take you to the CloudWatch Metrics dashboard of that particular metric, which gives you a much more customized and granular graph of the same metric, which looks like this:




- The second metric in the Lambda's monitoring dashboard is the **Duration** metric, which tells you the duration of each invocation of our Lambda function. It also has time as the X axis, and the duration time in the Y axis in the unit of milliseconds. It also tells you the maximum, average, and the minimum duration of your Lambda function over a period of time:




- Again, clicking on the **Jump to Logs** button will take you to the same page as that of the previous metric. Clicking on the **Jump to Metrics** button will take you to the CloudWatch metric page of the **Duration** metric, which looks like this:




6. The third metric is the **Errors** metric, which helps us keep a look out for errors in our invocations of the Lambda function. The Y axis is the number of errors while the X axis is the timeline:




7. The CloudWatch Dashboard of the same metric can be seen by clicking on the **Jump to Metrics** link:




8. The fourth metric is **Throttles**. This metric counts the number of times your Lambda functions have been throttled, which means the number of times the concurrent executions of the functions have breached the set limit of 1,000 per region. We won't encounter this metric very frequently as the Lambda functions which we build as examples in this book stay well within the concurrency limits:




9. By clicking on the **Jump to Metrics** link, we can also see this metric in our CloudWatch Metrics dashboard:




10. The fifth metric is the iterator age. This is only valid for functions which are triggered by the DynamoDB stream or the Kinesis stream. It gives the age of the last record which is processed by the function:




The **Jump to Metrics** link takes you to the CloudWatch Metrics dashboard of this metric:



- The sixth and the last metric is the **DLQ errors** metric. This gives the number of errors that occurred while sending messages (failed event payloads) to a dead letter queue. Most often the errors would be caused due to faulty permission configurations and timeouts:



The **Jump to Metrics** link will take you to the CloudWatch Metrics dashboard of the same metric:



Lambda's logs in CloudWatch

So far, we have learned about and understood the metrics of AWS Lambda in great detail. Now, we will move on to understanding the logs of the Lambda functions. As always, we will try to understand them via the following steps:

1. Logs for AWS Lambda functions are stored in CloudWatch's Logs service. You can access the CloudWatch Logs service by going to the **Logs** dashboard by clicking on the main **CloudWatch** dashboard.

- When you click on the logs of the **serverless-api**, **/aws/lambda/serverless-api**, in the list, we go to the log stream of the serverless API, which looks like this:

CloudWatch > Log Groups > Streams for /aws/lambda/serverless-api

Search Log Group Create Log Stream Delete Log Stream

Filter: Log Stream Name Prefix x

Log Streams 1-27

☐	Log Streams	Last Event Time
☐	2018/01/21/[\$LATEST]9d5e23b7f6a347e683f75b77ce34e08c	2018-01-21 23:09 UTC+5:30
☐	2018/01/21/[\$LATEST]4a04a6e78e834caea812a8b2dbcf22ff	2018-01-21 22:12 UTC+5:30
☐	2018/01/21/[\$LATEST]fb05e685dc6b4b2f82863098ed22236f	2018-01-21 15:02 UTC+5:30
☐	2018/01/20/[\$LATEST]5c1e67dd6fb849d3910115dcbf7809eb	2018-01-21 07:47 UTC+5:30
☐	2018/01/20/[\$LATEST]e774acd73a484686b18bf400bad1c863	2018-01-20 23:56 UTC+5:30
☐	2018/01/20/[\$LATEST]94cfc6283ed447afa4f563abb0c67618	2018-01-20 16:09 UTC+5:30
☐	2018/01/19/[\$LATEST]0fd88f1694aa44b48716f62effbcafd3	2018-01-20 08:49 UTC+5:30
☐	2018/01/19/[\$LATEST]5af5bbec461347a4bcad21d502a84ba0	2018-01-20 02:19 UTC+5:30
☐	2018/01/19/[\$LATEST]a046fbb4b5874ab7a1db9671648f93e4	2018-01-19 19:47 UTC+5:30
☐	2018/01/19/[\$LATEST]dbe9115136c54aafbe76c8245879dc0a	2018-01-19 12:59 UTC+5:30
☐	2018/01/18/[\$LATEST]b84dd6e6b09048e6b8e2ca26d6a5b3c0	2018-01-19 06:49 UTC+5:30
☐	2018/01/18/[\$LATEST]0710fd96881a480cba01e642e45b76f7	2018-01-18 22:56 UTC+5:30
☐	2018/01/18/[\$LATEST]afca54726a7a4dea9cc361eb8034cf6e	2018-01-18 16:09 UTC+5:30
☐	2018/01/17/[\$LATEST]a3fcbaf23038403aac8a4a9efe57e8ea	2018-01-18 09:09 UTC+5:30
☐	2018/01/17/[\$LATEST]57e2e7b470214695afaf5d4e7fdb5056	2018-01-18 02:51 UTC+5:30
☐	2018/01/17/[\$LATEST]3aeb54346fff4913bbd3dd28d3bc59e3	2018-01-17 20:47 UTC+5:30
☐	2018/01/17/[\$LATEST]bcc1e1b5dc024cce9c3e3e8d4d0a6194	2018-01-17 14:39 UTC+5:30
☐	2018/01/16/[\$LATEST]06efd549538d4f38b73be1756cf3de5d	2018-01-17 06:57 UTC+5:30

3. Each log stream here is a Lambda invocation. So, whenever your Lambda function is invoked, it creates a new log stream here. If the invocation is a part of Lambda's retry process, then the logs for that particular invocation will be written under the most recent log stream. A single log stream can contain several details. But firstly, let us look at what a particular log stream looks like:

CloudWatch > Log Groups > /aws/lambda/serverless-api > 2018/01/21/[\$LATEST]9d5e23b7f6a347e683f75b77ce34e08c

Expand all Row Text


Filter events all 30s 5m 1h 6h 1d 1w custom

Time (UTC +00:00)	Message
2018-01-21	
17:36:38	END RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1
17:36:38	REPORT RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1 Duration: 2.28 ms Billed Duration: 100 ms Memory Size: 128 MB Max
17:37:41	START RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1 Version: \$LATEST
17:37:41	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsi
17:37:41	END RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1
17:37:41	REPORT RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1 Duration: 9.11 ms Billed Duration: 100 ms Memory Size: 128 MB Max
17:39:28	START RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1 Version: \$LATEST
17:39:28	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsi
17:39:28	END RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1
17:39:28	REPORT RequestId: a271db0a-fed1-11e7-9937-6bdcf38ac0c1 Duration: 0.70 ms Billed Duration: 100 ms Memory Size: 128 MB Max
17:41:38	START RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Version: \$LATEST
17:41:38	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsi
17:41:38	END RequestId: 55641971-fed2-11e7-8d21-f3a179883953
17:41:38	REPORT RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Duration: 12.96 ms Billed Duration: 100 ms Memory Size: 128 MB Me
17:42:34	START RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Version: \$LATEST
17:42:34	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsi
17:42:34	END RequestId: 55641971-fed2-11e7-8d21-f3a179883953
17:42:34	REPORT RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Duration: 13.31 ms Billed Duration: 100 ms Memory Size: 128 MB Me
17:44:22	START RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Version: \$LATEST
17:44:22	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsi
17:44:22	END RequestId: 55641971-fed2-11e7-8d21-f3a179883953
17:44:22	REPORT RequestId: 55641971-fed2-11e7-8d21-f3a179883953 Duration: 16.19 ms Billed Duration: 100 ms Memory Size: 128 MB Me

- Also, if you look closely, you can observe that Lambda's logs also give out information about the duration of the Lambda function's invocation, the duration for which it is billed for, and also the memory used by the function. These metrics help in understanding our functions' performance better and for further optimization and fine tuning:

Filter events		all 30s 5m 1h 6h 1d 1w custom ▾
Time (UTC +00:00)	Message	
2018-01-21		
	REPORT RequestId: 718212e7-fecd-11e7-a74a-933c4948e6b0 Duration: 3.09 ms Billed Duration: 100 ms Memory Size: 128 MB Max Memory Used: 20 MB	
▶ 17:09:22	START RequestId: 718212e7-fecd-11e7-a74a-933c4948e6b0 Version: \$LATEST	
▶ 17:09:22	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsf	
▶ 17:09:22	END RequestId: 718212e7-fecd-11e7-a74a-933c4948e6b0	
▼ 17:09:22	REPORT RequestId: 718212e7-fecd-11e7-a74a-933c4948e6b0 Duration: 0.79 ms Billed Duration: 100 ms Memory Size: 128 MB Max	
	REPORT RequestId: 718212e7-fecd-11e7-a74a-933c4948e6b0 Duration: 0.79 ms Billed Duration: 100 ms Memory Size: 128 MB Max Memory Used: 20 MB	
▶ 17:11:38	START RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Version: \$LATEST	
▶ 17:11:38	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsf	
▶ 17:11:38	END RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f	
▼ 17:11:38	REPORT RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Duration: 5.70 ms Billed Duration: 100 ms Memory Size: 128 MB Max	
	REPORT RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Duration: 5.70 ms Billed Duration: 100 ms Memory Size: 128 MB Max Memory Used: 20 MB	
▶ 17:12:44	START RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Version: \$LATEST	
▶ 17:12:44	'type': KeyError Traceback (most recent call last): File "/var/task/lambda_function.py", line 9, in lambda_handler return {"type": mobsf	
▶ 17:12:44	END RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f	
▼ 17:12:44	REPORT RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Duration: 8.73 ms Billed Duration: 100 ms Memory Size: 128 MB Max	
	REPORT RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Duration: 8.73 ms Billed Duration: 100 ms Memory Size: 128 MB Max Memory Used: 20 MB	
▶ 17:14:53	START RequestId: 245110b6-fece-11e7-9d04-a31c7282a49f Version: \$LATEST	

5. There are several columns in CloudWatch Logs for you to select from, which are not shown in the preceding screenshots. These are the available options:




So, when you select more of those, you will see them in your dashboard as columns. These come in handy when you're doing a much more fine-grained debugging of our Lambda functions:

Log Streams	Last Event Time	Last Ingestion Time	Stored Bytes	Creation Time
☒ 2018/01/21/[\$LATEST]9d5e23b7f6a347e683f75b77ce34e08c	2018-01-21 23:39 UTC+5:30	2018-01-21 23:39 UTC+5:30	0 bytes	2018-01-21 22:13 UTC+
☐ 2018/01/21/[\$LATEST]4a04a6e78e834caea812a8b2dbcf22ff	2018-01-21 22:12 UTC+5:30	2018-01-21 22:12 UTC+5:30	0 bytes	2018-01-21 15:02 UTC+
☐ 2018/01/21/[\$LATEST]fb05e685dc6b4b2f82863098ed22236f	2018-01-21 15:02 UTC+5:30	2018-01-21 15:02 UTC+5:30	0 bytes	2018-01-21 07:47 UTC+
☐ 2018/01/20/[\$LATEST]5c1e67dd6fb849d3910115dcbf7809eb	2018-01-21 07:47 UTC+5:30	2018-01-21 07:47 UTC+5:30	0 bytes	2018-01-20 23:56 UTC+
☐ 2018/01/20/[\$LATEST]e774acd73a484686b18bf400bad1c863	2018-01-20 23:56 UTC+5:30	2018-01-20 23:56 UTC+5:30	0 bytes	2018-01-20 16:09 UTC+
☐ 2018/01/20/[\$LATEST]94cfc6283ed447afa4f563abb0c67618	2018-01-20 16:09 UTC+5:30	2018-01-20 16:09 UTC+5:30	0 bytes	2018-01-20 08:50 UTC+
☐ 2018/01/19/[\$LATEST]0fd88f1694aa44b716f62effbcfd3	2018-01-20 08:49 UTC+5:30	2018-01-20 08:49 UTC+5:30	0 bytes	2018-01-20 02:20 UTC+
☒ 2018/01/19/[\$LATEST]5af5bbec461347a4bcd21d502a84ba0	2018-01-20 02:19 UTC+5:30	2018-01-20 02:20 UTC+5:30	0 bytes	2018-01-19 19:47 UTC+
☐ 2018/01/19/[\$LATEST]a046fbb4b5874ab7a1db9671648f93e4	2018-01-19 19:47 UTC+5:30	2018-01-19 19:47 UTC+5:30	0 bytes	2018-01-19 12:59 UTC+
☐ 2018/01/19/[\$LATEST]db9115136c54aafbe76c8245879dc0a	2018-01-19 12:59 UTC+5:30	2018-01-19 13:00 UTC+5:30	0 bytes	2018-01-19 06:50 UTC+
☐ 2018/01/18/[\$LATEST]b84dd6e6b09048e6b8e2ca26d6a5b3c0	2018-01-19 06:49 UTC+5:30	2018-01-19 06:49 UTC+5:30	0 bytes	2018-01-18 22:56 UTC+
☐ 2018/01/18/[\$LATEST]0710fd96881a480cba01e642e45b76f7	2018-01-18 22:56 UTC+5:30	2018-01-18 22:56 UTC+5:30	0 bytes	2018-01-18 16:08 UTC+
☐ 2018/01/18/[\$LATEST]afca54726a7a4dea9cc361eb8034cf6e	2018-01-18 16:09 UTC+5:30	2018-01-18 16:09 UTC+5:30	0 bytes	2018-01-18 09:09 UTC+
☐ 2018/01/17/[\$LATEST]ja3fcba723038403aac8a4a9efe57e8ea	2018-01-18 09:09 UTC+5:30	2018-01-18 09:10 UTC+5:30	0 bytes	2018-01-18 02:50 UTC+
☐ 2018/01/17/[\$LATEST]57e2e7b470214695afaf5d4e7fdb5056	2018-01-18 02:51 UTC+5:30	2018-01-18 02:51 UTC+5:30	0 bytes	2018-01-17 20:48 UTC+
☐ 2018/01/17/[\$LATEST]3aeb54346fff4913bbd3dd28d3bc59e3	2018-01-17 20:47 UTC+5:30	2018-01-17 20:47 UTC+5:30	0 bytes	2018-01-17 14:38 UTC+
☐ 2018/01/17/[\$LATEST]bcc1e1b5dc024cce9c3e3e8d4d0a6194	2018-01-17 14:39 UTC+5:30	2018-01-17 14:39 UTC+5:30	0 bytes	2018-01-17 06:57 UTC+
☐ 2018/01/16/[\$LATEST]06efd549538d4f38b73be1756cf3de5d	2018-01-17 06:57 UTC+5:30	2018-01-17 06:57 UTC+5:30	0 bytes	2018-01-16 23:13 UTC+

Logging statements in Lambda

Logging your comments and errors clearly is always a good software practice. So, we shall now understand how to log from inside of Lambda functions. There are broadly two ways of logging inside Lambda functions. We shall now learn and understand them via examples from the following steps:

1. The first way is to use Python's `logging` library. This is widely used as a standard practice for logging in Python scripts. We shall edit the code we have written previously for the serverless API and add in the logging statements in it. The code will look like this:




The code which is in the preceding screenshot is as follows:

```
import logging
logger = logging.getLogger()
logger.setLevel(logging.INFO)
def lambda_handler(event, context):
    mobs = {
        "Sea": ["GoldFish", "Turtle", "Tortoise", "Dolphin", "Seal"],
        "Land": ["Labrador", "Cat", "Dalmatian", "German Shepherd",
        "Beagle", "Golden Retriever"],
        "Exotic": ["Iguana", "Rock Python"]
    }

    logger.info('got event{}'.format(event))
    logger.error('something went wrong')

    return 'Hello from Lambda!'
    #return {"type": mobs[event['type']]}
```


- Now, when you run the Lambda function after saving, you can see a successful execution statement in green color, which looks like this:



- When you click on the **Details** option, you can see the logging statements being executed clearly:




4. The next way of logging statements is by simply using the `print` statements in Python. It is the most common way of printing out logging statements in Python scripts. So, we shall add a `Hello from Lambda` print statement in our function code and see if we get the logs in our Lambda execution or not:




The code for this Lambda function is as follows:

```
def lambda_handler(event, context):
    mobs = {
        "Sea": ["GoldFish", "Turtle", "Tortoise", "Dolphin", "Seal"],
        "Land": ["Labrador", "Cat", "Dalmatian", "German Shepherd",
                 "Beagle", "Golden Retriever"],
        "Exotic": ["Iguana", "Rock Python"]
    }
    print 'Hello from Lambda!'
    return 1
    #return {"type": mobs[event['type']]}
```

5. When we click on **Test** for executing the code, we should see a green color message, which indicates a successful execution:



6. Again, just like we did previously, clicking on the **Details** toggle will give you the complete execution logs:



7. We can see the `Hello from Lambda` message too. Out of the two available logging options for our Lambda functions, it is always preferable to use the first option which is via the Python's logging module. This is because that module gives greater flexibility and helps you differentiate between info, error, and debug logs.

Summary

In this chapter, we have learned about the monitoring and the logging capabilities of AWS. We also learned about the available monitoring and logging tools inside the AWS environment. We have also learned how to monitor our Lambda functions and how to set up logging for our Lambda functions.

We have learned about the logging and monitoring practices that are followed by the industry and the various ways one can log statements in Python from inside of a Lambda function.

In the next chapter, we will learn how to scale up our serverless architectures to become distributed and to be able to handle massive workloads while still preserving the positives of a serverless setup.

6

Scaling Up Serverless Architectures

So far, we have learned how to build, monitor, and log serverless functions. In this chapter, we will be learning concepts and engineering techniques that will help scale up serverless applications to be distributed, and that will also enable them to handle heavy workloads with high standards of security and throughput. In this chapter, we will also use some third-party tools, such as Ansible, to scale up our Lambda functions. We will be scaling up our Lambda functions to spawn a distributed serverless architecture, which will involve spawning multiple servers (or instances in the AWS environment). You therefore need to keep that in mind while following the examples mentioned in this chapter.

This chapter assumes a working knowledge of a provisioning tool, such as **Ansible**, **Chef**, and so on. You can quickly read up on or refresh your knowledge of these on their respective sites, where they have quick tutorials. If not, then you can safely skip this chapter and move on to the next.


This chapter consists of five sections, which cover all of the basics of scaling up serverless architectures and will set you up for building bigger, complex serverless architectures:

- Third-party orchestration tools
- The creation and termination of servers
- Security best practices
- Difficulties of scaling up
- Handling difficulties


Third-party orchestration tools

In this section, we will learn and become versed in the concept of infrastructure provisioning and orchestration. We will be exploring a couple of tools, namely Chef and Ansible. Let's get started by following these steps:


1. We will begin with getting introduced to Chef. You can visit the official website of Chef at <https://www.chef.io/chef/>:



2. Chef has a very good set of tutorials for getting your hands dirty. These are organized in the form of mini 10 to 15 minute tutorials for easy consumption. Head over to <https://learn.chef.io/> to access them:




- For getting started with infrastructure provisioning and orchestrating, you can refer to the Chef documentation here: <https://docs.chef.io/>. The page looks like this:




4. You can refer to the **AWS Driver Resources** page in the documentation to understand how to interact with various AWS services via Chef at: https://docs.chef.io/provisioning_aws.html. The page looks like this:

The screenshot shows the 'AWS Driver Resources' page on the Chef documentation website. The page has a navigation sidebar on the left with links like 'Overview', 'Chef', 'Getting Started', 'Concepts', 'Setup', 'Cookbook Reference', and 'Chef DK'. The main content area is titled 'AWS Driver Resources' and includes a 'Warning' box stating that this functionality is available with Chef provisioning and is packaged in the Chef development kit. Below the warning, it explains that Chef provisioning is a collection of resources for creating machines and machine infrastructures using the chef-client. It also mentions that Chef provisioning is built around two major components: the **machine** resource and drivers. A driver-specific resource is defined as a statement of configuration policy that describes the desired state, declares the steps needed to bring that state, specifies a resource type (package, template, or service), lists additional details (properties), and is grouped into recipes. The page lists several driver-specific resources available for Amazon Web Services (AWS) and Chef provisioning: `aws_auto_scaling_group`, `aws_cache_cluster`, `aws_cache_replication_group`, and `aws_cache_subnet_group`. A 'Table Of Contents' on the right lists various AWS resources and their sub-sections (Syntax, Properties, Examples).


5. You can also refer to the **aws Cookbook** for the same purpose, too. This resource has very good documentation and APIs for interacting with several AWS services. The URL of this documentation is <https://supermarket.chef.io/cookbooks/aws>. The page looks like this:




6. A detailed description of the cookbook can be seen when you scroll down, directly after the title of the cookbook:




7. One other good tool for provisioning and orchestrating software resources is Ansible. This helps software engineers write code for automating several parts of their infrastructure via *yaml scripts*. Similar to the Chef environment, these scripts are called **cookbooks**.
8. We will be using this tool for learning how to provision our infrastructure in the subsequent sections. The documentation for Ansible can be found at <http://docs.ansible.com/>:




9. The product, **ANSIBLE TOWER**, is out of scope for this book. We will be learning and be working with **ANSIBLE CORE**, which is the flagship product of Ansible and its parent company, Red Hat.
10. Ansible has a very helpful video for helping you better understand and make sense of the tool. It can be accessed when you click on the **Quick Start Video** link in the documentation page:



11. After watching the video, you can proceed to understand the product from the documentation itself. The complete documentation of Ansible can be accessed at: <http://docs.ansible.com/ansible/latest/index.html>:



12. The EC2 module is the one we will be using for provisioning and orchestrating our AWS EC2 instances. This part of the documentation has a very clear explanation and demonstration of starting up and terminating EC2 instances, along with adding and mounting volumes; it also enables us to provision our EC2 instances into our own specific **Virtual Private Cloud (VPC)** and/or in our own **Security Groups (SGs)**. The EC2 documentation screen looks like this:



13. You can find this at the following URL of Ansible Core's documentation: http://docs.ansible.com/ansible/latest/ec2_module.html. When you scroll down further, you can see several examples of how to use the EC2 module of Ansible for various tasks concerning AWS EC2 instances. Some of them can be seen as follows:

Examples

```
# Note: These examples do not set authentication details, see the AWS Guide for details.

# Basic provisioning example
- ec2:
    key_name: mykey
    instance_type: t2.micro
    image: ami-123456
    wait: yes
    group: webserver
    count: 3
    vpc_subnet_id: subnet-29e63245
    assign_public_ip: yes

# Advanced example with tagging and CloudWatch
- ec2:
    key_name: mykey
    group: databases
    instance_type: t2.micro
    image: ami-123456
    wait: yes
    wait_timeout: 500
    count: 5
    instance_tags:
        db: postgres
    monitoring: yes
    vpc_subnet_id: subnet-29e63245
    assign_public_ip: yes


# Single instance with additional IOPS volume from snapshot and volume delete on termination
- ec2:
    key_name: mykey
    group: webserver
```

🔍 Search this site

The creation and termination of servers

In this chapter, we will learn how to use some third-party tools that will help us in building the required architecture. Like all of the sections in this chapter, the information will be broken down into steps:

1. The first tool we will be learning about is Ansible. It is a provisioning and orchestrating tool, that helps in automating several parts of an infrastructure. Depending on when you are reading this book, the Ansible project's homepage (<https://www.ansible.com/>) will look something like this:



2. The installation process for Ansible is different for different operating systems. The instructions for some popular operating systems are as follows:

- **For Ubuntu:**

```
sudo apt-get update
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get update
sudo apt-get install ansible
```

- **For Linux:**

```
git clone https://github.com/ansible/ansible.git
cd ./ansible
make rpm
sudo rpm -Uvh ./rpm-build/ansible-*.noarch.rpm
```

- **For OS X:**

```
sudo pip install ansible
```

3. Now, we will understand the concept of **nohup**. So, you don't need to have a persistent SSH connection to the server for making a `nohup` command run, therefore we will be using this technique for running our master-server architecture (to know more about `nohup` refer to: <https://en.wikipedia.org/wiki/Nohup>).



Let's look at its definition on Wikipedia (from the time of writing this book), **nohup** is a POSIX command to ignore the HUP (hangup) signal. The HUP signal is, by convention, the way a terminal warns dependent processes of logout.

4. We will now learn how to provision servers from Ansible, SSH into them, run a simple `apt-get update` task in them, and terminate them. From this, you will learn how to write Ansible scripts, as well as understand how Ansible handles the provisioning of cloud resources. The following Ansible script will help you understand how to provision an EC2 instance:

```
- hosts: localhost
  connection: local
  remote_user: test
  gather_facts: no

  environment:
    AWS_ACCESS_KEY_ID: "{{ aws_id }}"
    AWS_SECRET_ACCESS_KEY: "{{ aws_key }}"

    AWS_DEFAULT_REGION: "{{ aws_region }}"

  tasks:
- name: Provisioning EC2 instances
  ec2:
    assign_public_ip: no
    aws_access_key: "{{ access_key }}"
    aws_secret_key: "{{ secret_key }}"
```

```
region: "{{ aws_region }}"
image: "{{ image_instance }}"
instance_type: "{{ instance_type }}"
key_name: "{{ ssh_keyname }}"
state: present
group_id: "{{ security_group }}"
vpc_subnet_id: "{{ subnet }}"
instance_profile_name: "{{ Profile_Name }}"
wait: true
instance_tags:
    Name: "{{ Instance_Name }}"
delete_on_termination: yes
register: ec2
ignore_errors: True
```


The values in the `{{ }}` brackets need to be filled in as per your convenience and specifications. The preceding code will create an EC2 instance in your console and name it, as per the specification which is given in the `{{ Instance_Name }}` section.

5. The `ansible.cfg` file should include all of the details which give instructions about the control path, the details regarding the forwarding agent, and also the path to the EC2 instance key. The `ansible.cfg` file should look like this:

```
[ssh_connection]
ssh_args=-o ControlMaster=auto -o ControlPersist=60s -o
ControlPath=/tmp/ansible-ssh-%h-%p-%r -o ForwardAgent=yes

[defaults]
private_key_file=/path/to/key/key.pem
```


6. When you execute this code using `ansible-playbook -vvv < name-of-playbook >.yaml`, you can see the EC2 instance being created in your EC2 console:



7. Now, we will terminate the instance which we have just created via Ansible. This will also be done in an Ansible script, similar to how we provisioned the instance. The following code does this:

```
tasks:
  - name: Terminate instances that were previously launched
    connection: local
    become: false
    ec2:
      state: 'absent'
      instance_ids: '{{ ec2.instance_ids }}'
      region: '{{ aws_region }}'
    register: TerminateWorker
    ignore_errors: True
```

8. So, now you can see the instance being terminated in the console. Note that the code is the same up until the tasks, such as provisioning and terminating instances, so you can copy and paste from the provisioning task:



So, we have successfully learned how to provision and terminate EC2 instances via an Ansible script. We will use this knowledge for provisioning and will be terminating EC2 instances at the same time.

9. Making a small change to the provisioning code in the yaml script we used previously, we can provision multiple servers (EC2 instances) at the same time, by simply adding the `count` parameter. The following code will provision the number of instances mentioned in the *jinja template*, beside the `count` parameter. In our example, it is `ninstances`:

```
- hosts: localhost
  connection: local
  remote_user: test
  gather_facts: no

environment:
  AWS_ACCESS_KEY_ID: "{{ aws_id }}"
  AWS_SECRET_ACCESS_KEY: "{{ aws_key }}"
```

```

AWS_DEFAULT_REGION: "{{ aws_region }}"

tasks:
- name: Provisioning EC2 instances
  ec2:
    assign_public_ip: no
    aws_access_key: "{{ access_key }}"
    aws_secret_key: "{{ secret_key }}"
    region: "{{ aws_region }}"
    image: "{{ image_instance }}"
    instance_type: "{{ instance_type }}"
    key_name: "{{ ssh_keyname }}"
    count: "{{ ninstances }}"
    state: present
    group_id: "{{ security_group }}"
    vpc_subnet_id: "{{ subnet }}"
    instance_profile_name: "{{ Profile_Name }}"
    wait: true
    instance_tags:
      Name: "{{ Instance_Name }}"
    delete_on_termination: yes
    register: ec2

```

10. Now, as we have our Ansible script ready, we will now use it to start our infrastructure from the Lambda function. For that, we will make use of our knowledge of nohup.
11. In your Lambda function, all you need to do is to write the logic for creating a server, and then do some basic installations using the library, *paramiko*, and then run the Ansible script in a nohup mode, as shown here:

```

import paramiko
import boto3
import logging

logger = logging.getLogger(__name__)
logger.setLevel(logging.CRITICAL)
region = 'us-east-1'
image = 'ami-<>'
ubuntu_image = 'ami-<>'
keyname = '<>'

def lambda_handler(event, context):
    credentials = {}
    k = paramiko.RSAKey.from_private_key_file("<>")
    c = paramiko.SSHClient()
    c.set_missing_host_key_policy(paramiko.AutoAddPolicy())
    logging.critical("Creating Session")

```

```

        session = boto3.Session(credentials['AccessKeyId'],
                                credentials['SecretAccessKey'],
                                aws_session_token=credentials['SessionToken'],
                                region_name=region)
        logging.critical("Created Session")
        logging.critical("Create Resource")
        ec2 = session.resource('ec2', region_name=region)
        logging.critical("Created Resource")
        logging.critical("Key Verification")

        key = '<>'
        k = paramiko.RSAKey.from_private_key_file(key)
        c = paramiko.SSHClient()
        c.set_missing_host_key_policy(paramiko.AutoAddPolicy())
        logging.critical("Key Verification done")
        # Generate Presigned URL for downloading EC2 key from an
        S3 bucket into master
        s3client = session.client('s3')

        # Presigned url for downloading pem file of the server from an
        S3 bucket
        url = s3client.generate_presigned_url('get_object',
        Params={'Bucket': '<bucket_name>', 'Key':
        '<file_name_of_key>'},
        ExpiresIn=300)
        command = 'wget ' + '-O <>.pem ' + '"' + url + '"'
        logging.critical("Create Instance")
        while True:
            try:
                logging.critical("Trying")
                c.connect(hostname=dns_name, username="ubuntu", pkey=k)
            except:
                logging.critical("Failed")
                continue
            break
        logging.critical("connected")

        if size == 0:
            s3client.upload_file('<>.pem', '<bucket_name>',
            '<>.pem')
        else:
            pass
        set_key = credentials['AccessKeyId']
        set_secret = credentials['SecretAccessKey']
        set_token = credentials['SessionToken']

        # Commands to run inside the SSH session of the server
        commands = [command,

```

```
"sudo apt-get -y update",
"sudo apt-add-repository -y ppa:ansible/ansible",
"sudo apt-get -y update",
"sudo apt-get install -y ansible python-pip git awscli",
"sudo pip install boto markupsafe boto3 python-dateutil
futures",
"ssh-keyscan -H github.com >> ~/.ssh/known_hosts",
"git clone <repository where your ansible script is>
/home/ubuntu/<>/",
"chmod 400 <>.pem",
"cd <>/<>; pwd ; nohup ansible-playbook -vvv provision.yml >
ansible.out 2> ansible.err < /dev/null &"]

# Running the commands
for command in commands:
    logging.critical("Executing %s", command)
    stdin, stdout, stderr = c.exec_command(command)
    logging.critical(stdout.read())
    logging.critical("Errors : %s", stderr.read())
    c.close()
return dns_name
```


Security best practices

Ensuring high-level security has always been a major problem for microservices. There are multiple levels of software that you need to keep in mind while designing the security layers. The engineers need to define the security protocols for each of the services and then also define the protocols for the data interaction and transfer between each service.

You have to keep all these aspects in mind before architecting distributed serverless systems, where (almost) each Ansible task is a microservice. In this section, we will understand how to architect the security protocols, and also monitor them using some of AWS's built-in services.


We will go through a step-by-step understanding of how to write security protocols for our serverless architectures:

1. Firstly, whenever you are creating a session inside your AWS Python scripts using **Boto**, try to create temporary credentials using the **AWS Secure Token Service (STS)**, which creates temporary credentials for a specific period of time:



You can look at the documentation of the STS at: <https://docs.aws.amazon.com/STS/latest/APIReference/Welcome.html>.

2. The **AssumeRole** API of the STS service enables programmers to assume IAM roles into their code:




The screenshot shows the AWS Security Token Service (STS) API Reference page for the `AssumeRole` action. The page is titled "AWS Security Token Service" and "API Reference (API Version 2011-06-15)". The left sidebar contains a navigation menu with sections like "Welcome", "Actions", "Data Types", "Common Parameters", and "Common Errors". Under "Actions", the `AssumeRole` action is selected. The main content area shows the `AssumeRole` action details, including a description: "Returns a set of temporary security credentials (consisting of an access key ID, a secret access key, and a security token) that you can use to access AWS resources that you might not normally have access to." It also includes an "Important" note: "You cannot call `AssumeRole` by using AWS root account credentials; access is denied. You must use credentials for an IAM user or an IAM role to call `AssumeRole`." and a paragraph explaining cross-account access and federation.



You can find its documentation on the following page: https://docs.aws.amazon.com/STS/latest/APIReference/API_AssumeRole.html

3. The Python version of this can be referred to, in the `boto3` documentation:



The screenshot shows the Boto3 documentation page for the STS Client. The page is titled "Boto 3 Docs 1.5.36 documentation" and "Docs / Available Services / STS". The left sidebar contains a "TABLE OF CONTENTS" with sections like "Quickstart", "A Sample Tutorial", "Code Examples", "User Guides", and "Available Services". Under "Available Services", the "STS" service is selected. The main content area shows the "STS" service details, including a "Table of Contents" with links to "Client" and "Paginators". The "Client" section is expanded, showing the `class STS.Client` and a description: "A low-level client representing AWS Security Token Service (STS):". Below the description, there is a code snippet showing how to import and create the client:

```
import boto3

client = boto3.client('sts')
```



This documentation can be found here: <http://boto3.readthedocs.io/en/latest/reference/services/sts.html>.

4. Scrolling down, you can find the usage of the **AssumeRole** API in Python:

Boto 3 Docs 1.5.36 documentation

Request Syntax

```

response = client.assume_role(
    RoleArn='string',
    RoleSessionName='string',
    Policy='string',
    DurationSeconds=123,
    ExternalId='string',
    SerialNumber='string',
    TokenCode='string'
)

```

Parameters

- **RoleArn (string) -- [REQUIRED]**
The Amazon Resource Name (ARN) of the role to assume.
- **RoleSessionName (string) -- [REQUIRED]**
An identifier for the assumed role session.
Use the role session name to uniquely identify a session when the same role is assumed by different principals or for different reasons. In cross-account scenarios, the role session name is visible to, and can be logged by the account that owns the role. The role session name is also used in the ARN of the assumed role principal. This means that subsequent cross-account API requests using the temporary security credentials will expose the role session name to the external account in their CloudTrail logs.

The regex used to validate this parameter is a string of characters consisting of upper- and

5. Proper care should be taken so that the data exchange between microservices and/or between the microservices and other AWS resources happens securely with authentication. For example, the developer can configure S3 buckets to restrict actions such as unencrypted uploads, downloads, and insecure file transfers. The bucket policy can be written as follows to ensure all of these things are taken care of:

```

{
  "Version": "2012-10-17",
  "Id": "PutObjPolicy",
  "Statement": [
    {
      "Sid": "DenyIncorrectEncryptionHeader",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::<bucket_name>/*",


```

```


        "Condition": {
            "StringNotEquals": {
                "s3:x-amz-server-side-encryption": "aws:kms"
            }
        }
    },
    {
        "Sid": "DenyUnEncryptedObjectUploads",
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::<bucket_name2>/*",
        "Condition": {
            "Null": {
                "s3:x-amz-server-side-encryption": "true"
            }
        }
    },
    {
        "Sid": "DenyNonSecureTraffic",
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::<bucket_name>/*",
        "Condition": {
            "Bool": {
                "aws:SecureTransport": "false"
            }
        }
    },
    {
        "Sid": "DenyNonSecureTraffic",
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::<bucket_name2>/*",
        "Condition": {
            "Bool": {
                "aws:SecureTransport": "false"
            }
        }
    }
}

```

6. Once you have finished writing the bucket policy, you can update it in the **Bucket Policy** section of S3:



7. **AWS Config** provides a very useful interface for monitoring several security threats and helps in efficiently avoiding or catching them. The dashboard of **AWS Config** looks like this:



8. You can see that the dashboard shows **2 non-compliant resource(s)** which means that two of my AWS resources are not complying with the rules that I have put into config. Let's have a look at these rules:

AWS Config Rules

Rules represent your desired configuration settings. AWS Config evaluates whether your resource configurations comply with relevant rules and summarizes the results in the following table.

[Add rule](#) [Filter](#) [Status](#)

Rule name	Compliance	Edit rule
s3-bucket-ssl-requests-only	2 noncompliant resource(s)	Edit rule

This means that we have two AWS S3 buckets which do not have SSL requests turned on via the bucket policy. Once you click on the **Rules** link, you can see more details which include the bucket(s) names, and also the timestamps at which these configuration changes have been recorded:

Rules > Rule details

s3-bucket-ssl-requests-only

Description: Checks whether S3 buckets have policies that require requests to use Secure Socket Layer (SSL).

Trigger type: Configuration changes

Scope of changes: Resources

Resource types: S3 Bucket

Config rule ARN: arn:aws:config:us-east-1:config-rule/config-rule-nsbvxj

Parameters: null

Overall rule status: Last successful invocation on February 26, 2018 at 12:36:39 AM ✔
Last successful evaluation on February 26, 2018 at 12:36:40 AM ✔

Resources evaluated

Click on the [⌕](#) icon to view configuration details for the resource when it was last evaluated with this rule.

Resource type	Config timeline	Compliance	Last successful invocation	Last successful evaluation	Manage resource
S3 Bucket	receiver-bucket	Noncompliant	February 26, 2018 12:36:39 AM	February 26, 2018 12:36:40 AM	⌕
S3 Bucket	sender-bucket	Noncompliant	February 26, 2018 12:36:39 AM	February 26, 2018 12:36:40 AM	⌕

Identifying and handling difficulties in scaling

Scaling up distributed serverless systems comes with its own set of engineering roadblocks and problems, and the fact that the concept of serverless systems is still in a very infantile stage, means that most of those problems are still unsolved. But, that shouldn't stop us from trying to solve and work around these roadblocks.

We will try and understand some of these roadblocks, and also learn how to solve or work around them, as discussed here:

- This is more of an architect's mistake rather than a roadblock. However, it is important to address this as one too many architects/software engineers fell and fall into the overestimation or the underestimation trap. The problem we will try to address is the exact number of instances you have to launch when scaling up. In most self-hosted MapReduce-style systems, it is taken care of out of the box.
- This problem can be taken care of, by properly benchmarking the workloads beforehand on different types of instances, and scale accordingly. Let's understand this by taking an example of a machine learning pipeline. Thanks to our benchmarking efforts, we already know that an *m3.medium* instance can handle 100 files in 10 minutes. So, if my workload has 202 files and I want it to be completed in close to 10 minutes, I would like to have two such instances for handling this. Even if we don't know the workloads in advance, we can write a Python script for getting that number from wherever the data is, be it an SQS queue pointer, or S3, or some other database; and that number can be entered into the Ansible script and make the playbook run.
- As we have already learned about handling security in huge serverless systems, we will keep this short. There are several complex data movements happening inside a large distributed serverless workload. Using proper security protocols and monitoring them, as mentioned in detail in the previous security section, will help in overcoming this problem.
- Logging is a major problem in distributed serverless systems, which is also still unsolved completely. As the systems and containers are destroyed once the workload has been completed, logging has been a very difficult task to undertake. There are several ways you can log the workflow. The most popular ones are logging every Ansible task separately, and one where the last Ansible task is to zip up the logs and send the zipped file to a data store, such as S3 or Logstash. The last one is the most preferred way as it captures the execution flow better, as the entire log trace is in a single file.

- Monitoring is similar to logging. Monitoring these systems is also mostly an unsolved problem. As the servers are all terminated once the workload is run, we can't poll for historic logs from the servers, and latency also will not be tolerated, or more precisely, will not be possible. Monitor every task of Ansible by having a task after each, that sends a custom metric to CloudWatch upon a condition that the previous task has executed successfully or not. This will look something like this:

```
- name: OnDemandProvision on success
  command: aws cloudwatch put-metric-data --metric-name OnDemandProvision[M] --namespace Ansible --value 1
  when: ec2|succeeded
- name: OnDemandProvision on failure
  command: aws cloudwatch put-metric-data --metric-name OnDemandProvision[M] --namespace Ansible --value 0
  when: ec2|failed
```

- Debugging trial runs can also become very frustrating, very fast. This is because, if you are not quick, the entire system can be terminated before you even get a chance to look at the logs. Also, Ansible emits very verbose logs while debugging, which might seem overwhelming when spawning several instances.
- Some basic Unix hacks can help in handling these problems. The most important one is to monitor the tail of the log file, about 50 lines or so. This helps in not getting overwhelmed by the huge amount of logs, and it also keeps an eye on the execution of the Ansible notebook.

Summary

In this chapter, we have learned how to scale up our serverless architecture(s) to being massively distributed serverless infrastructure(s). We have learned how to build on our existing knowledge of building and deploying Lambda infrastructures to handle massive workloads.

We have learned to use the concept of `nohup` to use our Lambda function as a launch board for building a master-worker architecture that takes parallel computing into account. We have learned how to leverage configuration and orchestration tools, such as Ansible and Chef, to spawn and orchestrate multiple EC2 instances.

The knowledge gained from this chapter will open doors for building many complex infrastructures which can handle data and requests, both in terms of size and speed. This will allow you to operate multiple microservices closely intertwined together. This will also help you to build MapReduce-style systems and interact with other AWS services, seamlessly.

7 Security in AWS Lambda

We have learned how to build and configure serverless functions in AWS Lambda. We have learned how to scale them up using third-party tools. We have also had a close look at how microservices work and how to ensure security in them, while ensuring resilience and speed.

In this chapter, we will understand security in the AWS environment, keeping in mind our Lambda functions. We will understand how services, such as AWS VPCs, security groups, and subnets work, with respect to Lambda functions.


This chapter covers the following topics:

- Understanding AWS VPCs
- Understanding subnets in VPCs
- Securing Lambda inside private subnets
- Controlling access to Lambda functions
- Using STS inside Lambda for secure session-based execution

Understanding AWS Virtual Private Clouds (VPCs)

In this section, we will understand AWS VPCs. **VPCs** are a very common component in the security layers of the AWS environment; they are isolated parts of the cloud where users can host their services and build their infrastructures. VPCs are the first layer of security. We will try to understand VPCs in the context of Lambda functions, in the form of bullet points, given here:

1. VPCs can be created and modified in the AWS's VPC service dashboard, which looks like this:



2. Now, let's quickly learn how to create a VPC of our own. For that, click on **Create VPC**. You will see a pop-up box which asks you to enter more meta information for your new VPC:

Create VPC [X]

A VPC is an isolated portion of the AWS cloud populated by AWS objects, such as Amazon EC2 instances. You must specify an IPv4 address range for your VPC. Specify the IPv4 address range as a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.0.0/16. You cannot specify an IPv4 CIDR block larger than /16. You can optionally associate an Amazon-provided IPv6 CIDR block with the VPC.

Name tag [] ⓘ

IPv4 CIDR block* [] ⓘ

IPv6 CIDR block* ☒ No IPv6 CIDR Block ⓘ
☐ Amazon provided IPv6 CIDR block

Tenancy [Default] ⓘ

[Cancel] [Yes, Create]

3. The **Name tag** box needs to have the name of the VPC. The **IPv4 CIDR block** is where you enter your IP range for classless inter-domain routing. Then, you can choose whether you want an IPv6 CIDR block or not. You can also select the **Tenancy** settings; this defines how your EC2 instances run within your VPC, and the resource sharing accordingly:

Create VPC [X]

A VPC is an isolated portion of the AWS cloud populated by AWS objects, such as Amazon EC2 instances. You must specify an IPv4 address range for your VPC. Specify the IPv4 address range as a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.0.0/16. You cannot specify an IPv4 CIDR block larger than /16. You can optionally associate an Amazon-provided IPv6 CIDR block with the VPC.

Name tag [Test-VPC] ⓘ

IPv4 CIDR block* [172.31.0.0/16] ⓘ

IPv6 CIDR block* ☒ No IPv6 CIDR Block ⓘ
☐ Amazon provided IPv6 CIDR block

Tenancy [✓ Default] ⓘ
[Dedicated]

[Cancel] [Yes, Create]

4. We have successfully created our VPC with the necessary settings and with the `Test-VPC` name. We can see this in our dashboard with all the corresponding meta settings:

The screenshot shows the AWS VPC Dashboard. On the left, there's a sidebar with navigation links like 'Your VPCs', 'Subnets', 'Route Tables', etc. The main area displays a table of VPCs. Below the table, the 'Summary' tab for 'vpc-56a60a2d | Test-VPC' is selected, showing various configuration details.

Name	VPC ID	State	IPv4 CIDR	IPv6 CIDR	DHCP options set	Route table	Network
vpc-030c8b65	vpc-030c8b65	available	172.31.0.0/16		dopt-3075f057	rtb-59d51020	acl-72723
Test-VPC	vpc-56a60a2d	available	172.31.0.0/16		dopt-3075f057	rtb-c9c013b5	acl-ce2c1

vpc-56a60a2d | Test-VPC

Summary | CIDR Blocks | Flow Logs | Tags

VPC ID: vpc-56a60a2d Test-VPC	Network ACL: acl-ce2c15b5
State: available	Tenancy: Default
IPv4 CIDR: 172.31.0.0/16	DNS resolution: yes
IPv6 CIDR:	DNS hostnames: no

5. You can also see a summary of the VPC which includes the IPv4 settings, the **Network Access Control List (ACL)** settings, the **Dynamic Host Configuration Protocol (DHCP)** options, and also the DNS settings, all of which can also be configured later according to our needs. You can also see IPv4 CIDR blocks under the next **CIDR Blocks** tab:

This screenshot shows the 'CIDR Blocks' tab for the 'vpc-56a60a2d | Test-VPC'. It displays a table with the IPv4 CIDR block and its status.

CIDR	Status	Status Reason
172.31.0.0/16	associated	

6. We can also create VPC flow logs, which log traffic and data movements in and out of the VPC. This will promote better log management, ensuring security, and better monitoring. Currently, flow logs have not been set up:



7. To create VPC flow logs, you can just click on the **Create Flow Log** button at the bottom. This will open up a flow log creation wizard where you can enter the details for the various settings, accordingly. The creation wizard looks like this:

The screenshot shows the 'Create Flow Log' wizard. It includes a title bar with a close button. The main content area has a description: 'Flow logs enable you to capture IP traffic flow information for the network interfaces in your resources. [Learn more about flow logs.](#)'. Below this are several fields: 'Resources' (vpc-56a60a2d), 'Filter*' (All), 'Role*' (empty), 'ARN' (arn:aws:iam::080983167913:role/), and 'Destination Log Group*' (empty). Each field has an information icon. A note below the 'Role*' field states: 'If you have not setup IAM permissions for the destination CloudWatch Account you will need to do so to use Flow Logs. [Set Up Permissions](#)'. At the bottom, there is a legend for '*' indicating required fields, and two buttons: 'Cancel' and 'Create Flow Log'.

8. Once all the details have been entered, you can go ahead and click on the **Create Flow Log** option at the bottom, which will create the VPC flow log with the desired settings:

Create Flow Log

Flow logs enable you to capture IP traffic flow information for the network interfaces in your resources. [Learn more about flow logs.](#)

Resources vpc-56a60a2d

Filter* All

Role* CloudTrail_CloudWatchLogs_Role

If you have not setup IAM permissions for the destination CloudWatch Account you will need to do so to use Flow Logs. [Set Up Permissions](#)

ARN arn:aws:iam::080983167913:role/CloudTrail_CloudWatchLogs_Role

Destination Log Group* /aws/lambda/Test-Lambda

*: Required

Cancel Create Flow Log

9. Once created, you can see the newly created VPC flow log under the **Flow Logs** tab, as shown here:

vpc-56a60a2d | Test-VPC

Summary

CIDR Blocks

Flow Logs

Tags

You can create flow logs on your resources to capture IP traffic flow information for the network interfaces for your resources. [Learn more about flow logs.](#)

Create Flow Log


Flow Log ID	Filter	CloudWatch Logs Group	IAM Role ARN	Creation Time	Status	Inherited From
fl-b908cdd0	ALL	/aws/lambda/Test-Lambda	arn:aws:iam::080983167913:role/CloudTrail_CloudWatchLogs_Role	March 8, 2018 at 8:37:24 AM UTC+5:30	Active	-

10. Now, let's understand VPCs from AWS Lambda's point of view. Just like any other AWS resource, Lambda functions can also be hosted inside VPCs. You can see that setting under the **Network** section of your AWS Lambda function. It looks like this:



The screenshot shows the 'Network' section of the AWS Lambda console. It includes a 'VPC' label with an 'Info' link, a prompt to 'Select a VPC that your function will access.', and a dropdown menu currently displaying 'No VPC'.

11. From the drop-down list, you can select a VPC in which you want to host your Lambda function:



This screenshot shows the 'Network' section with the VPC dropdown menu open. The menu lists 'No VPC' as the selected option, followed by a search bar and a list of available VPCs: 'Default vpc-030c8b65 (172.31.0.0/16)' and 'vpc-56a60a2d (172.31.0.0/16) | Test-VPC'.

12. Once you select a VPC, it will further ask you for details regarding subnets, security groups, and so on, as shown in the following screenshot. We will learn about them in the sections following this, so, we will configure the VPC for our Lambda function later:

Network

VPC [Info](#)
Select a VPC that your function will access.

vpc-56a60a2d (172.31.0.0/16) | Test-VPC ▼

Subnets*
Select the VPC Subnets that Lambda should use to set up your VPC configuration.
Format: "subnet-id (cidr-block) | az name-tag".

▼

Security Groups*
Select the VPC Security Groups that Lambda should use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The table below will show the inbound and outbound rules for the security groups you have selected.

▼

ⓘ When you enable VPC, your Lambda function will lose default internet access. **If you require external internet access for your function, ensure that your security group allows outbound connections and that your VPC has a NAT gateway.**

Inbound rules

Outbound rules

< 1 >

Security group ID	Ports	Source
None	-	-

Understanding subnets in VPCs


In this section, we will learn about and understand AWS subnets, which are subparts of AWS VPCs. VPCs can be further divided into multiple subnets. These subnets can either be public or private, depending on the security needs of your architecture. We will look at the concept of subnets from the point of view of AWS Lambda functions.

We will perform the following steps:

1. You can go to the **Subnets** menu via the VPC page itself. You need to click on the **Subnets** option under the **Your VPCs** option on the left:

The screenshot displays the AWS VPC Dashboard. On the left, the 'Virtual Private Cloud' section is expanded, showing 'Your VPCs' and a list of resources: Subnets, Route Tables, Internet Gateways, Egress Only Internet Gateways, DHCP Options Sets, Elastic IPs, Endpoints, Endpoint Services, NAT Gateways, Peering Connections, Security, Network ACLs, and Security Groups. The main area, titled 'Resources', shows a list of resources in the US East (N. Virginia) region: 2 VPCs, 0 Egress-only Internet Gateways, 2 Route Tables, 0 Elastic IPs, 0 Endpoints, 3 Security Groups, 0 VPN Connections, 0 Customer Gateways, 1 Internet Gateway, 6 Subnets, 2 Network ACLs, 0 VPC Peering Connections, 0 Nat Gateways, 0 Running Instances, 0 Virtual Private Gateways, and 1 DHCP Options Set. A 'Create VPN Connection' button is visible at the bottom. The right sidebar shows 'Service Health' for Amazon VPC and Amazon EC2, both indicating 'Service is operating normally'. Below this, 'Additional Information' links are provided for VPC Documentation, All VPC Resources, Forums, and Report an Issue.

- 2. This will take you to the subnets console, where you will see some already existing subnets. These are the default subnets for each availability zone in your region:



- 3. Now, to create a new subnet, you need to click on the blue **Create Subnet** button on the top-left side of the console. In the creation wizard, you will be asked to enter the following details—the name of the subnet, the VPC you want to place it in, availability zones, and also preferred IPv4 CIDR blocks. I have placed this subnet inside the VPC we created in the previous section:

Create Subnet

Use the CIDR format to specify your subnet's IP address block (e.g., 10.0.0.0/24). Note that block sizes must be between a /16 netmask and /28 netmask. Also, note that a subnet can be the same size as your VPC. An IPv6 CIDR block must be a /64 CIDR block.

Name tag ⓘ

VPC ⓘ

VPC CIDRs

CIDR	Status	Status Reason
172.31.0.0/16	associated	

Availability Zone ⓘ

IPv4 CIDR block ⓘ

[Cancel](#) [Yes, Create](#)

- When you click on the **Yes, Create** button on the bottom-right side of the creation wizard, the new subnet is created. You can see it listed in the list of your subnets on your console:

VPC Dashboard

Filter by VPC:

Virtual Private Cloud

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

DHCP Options Sets

Elastic IPs

Endpoints

Endpoint Services

NAT Gateways

Peering Connections

Security

Network ACLs

Security Groups

[Create Subnet](#) Subnet Actions

Search Subnets and their pro...

<< 1 to 7 of 7 Subnets >>

Name	Subnet ID	State	VPC	IPv4 CIDR	Available IPv4	IPv6 CIDR
	subnet-5702876b	available	vpc-030c8b65	172.31.32.0/20	4091	
☒ Test-Subnet	subnet-949564de	available	vpc-56a60a2d Test-VPC	172.31.0.0/16	65531	
	subnet-7aedc333	available	vpc-030c8b65	172.31.0.0/20	4091	
	subnet-65ccb548	available	vpc-030c8b65	172.31.48.0/20	4091	
	subnet-0f146254	available	vpc-030c8b65	172.31.16.0/20	4091	
	subnet-2b61df27	available	vpc-030c8b65	172.31.80.0/20	4091	
	subnet-5df86138	available	vpc-030c8b65	172.31.64.0/20	4091	

subnet-949564de | Test-Subnet

Summary Route Table Network ACL Flow Logs Tags

Subnet ID: subnet-949564de | Test-Subnet Availability Zone: us-east-1b

IPv4 CIDR: 172.31.0.0/16 Route table: rtb-c9c013b5

IPv6 CIDR: State: available Network ACL: acl-ce2c15b5

Default subnet: no

5. Now, we will fill in the security settings for our Lambda function with our VPC and subnets, which we have just created. Currently, this is what the **Network** setting for AWS Lambda looks like:

serverless-api

Qualifiers

Network

VPC [Info](#)

Select a VPC that your function will access.

vpc-56a60a2d (172.31.0.0/16) | Test-VPC ▼

Subnets*

Select the VPC Subnets that Lambda should use to set up your VPC configuration. Format: "subnet-id (cidr-block) | az name-tag".

▼

Security Groups*

Select the VPC Security Groups that Lambda should use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The table below will show the inbound and outbound rules for the security groups you have selected.

▼

ⓘ

When you enable VPC, your Lambda function will lose default internet access. **If you require external internet access for your function, ensure that your security group allows outbound connections and that your VPC has a NAT gateway.**

Inbound rules

Outbound rules

6. After adding in the required settings, which are the details of the VPC, subnet and security groups, the **Network** settings of our Lambda function will look like this:

The screenshot shows the 'Network' configuration page for a Lambda function named 'serverless-api'. The page has a 'Qualifiers' tab in the top right corner. The 'Network' section is active and contains the following settings:

- VPC** [Info](#): Select a VPC that your function will access. The dropdown menu shows 'vpc-56a60a2d (172.31.0.0/16) | Test-VPC'.
- Subnets***: Select the VPC Subnets that Lambda should use to set up your VPC configuration. Format: "subnet-id (cidr-block) | az name-tag". The dropdown menu shows 'subnet-949564de (172.31.0.0/16) | us-east-1b Test-Subnet' with a red border around the input field. Below the dropdown, a warning message states: 'We recommend you select at least 2 subnets for Lambda to run your functions in high availability mode.'
- Security Groups***: Select the VPC Security Groups that Lambda should use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The dropdown menu shows 'sg-9a19c5ec (default)'.

At the bottom of the form, there is a blue information box with an 'i' icon that reads: 'When you enable VPC, your Lambda function will lose default internet access. If you require external internet access for your'.

...

serverless-api

Qualifiers

Select the VPC Security Groups that Lambda should use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The table below will show the inbound and outbound rules for the security groups you have selected.

sg-9a19c5ec (default) X

When you enable VPC, your Lambda function will lose default internet access. **If you require external internet access for your function, ensure that your security group allows outbound connections and that your VPC has a NAT gateway.**

Inbound rules

Outbound rules

< 1 >


Security group ID	Ports	Source
sg-9a19c5ec	All	sg-9a19c5ec

- After setting up your network settings for your Lambda function, click on the orange **Save** button on the top-right of your Lambda console to save those settings to your Lambda function.


Securing Lambda inside private subnets

Private subnets are subnets that are not open to the internet. All of their traffic is routed via the public subnet in the same VPC using the concept of route tables. Let's understand how to position our Lambda functions inside private subnets to add an extra layer of security:

1. Subnets created in the AWS console are not private by default. Let's evaluate and confirm this by going through the details of the subnet that we just created:



2. Clicking on the **Route Table** tab will show us the routing settings of our subnet, which basically tells us what kind of traffic is allowed into it:



3. In the **Network ACL** tab, you can see the network rules assigned for our subnet. Here, we can see that the subnet is open to all traffic (0.0.0.0/0). So, in order to make our subnet private, we need to fix this:

subnet-949564de | Test-Subnet

Summary Route Table **Network ACL** Flow Logs Tags

Edit

Network ACL: [acl-ce2c15b5](#)

Inbound:

Rule #	Type	Protocol	Port Range / ICMP Type	Source	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY

Outbound:

Rule #	Type	Protocol	Port Range / ICMP Type	Destination	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY

4. Go to the **Network ACLs** console by clicking on the link to the left side of your console. You will arrive at the following page:

VPC Dashboard

Filter by VPC:

Virtual Private Cloud

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

DHCP Options Sets

Elastic IPs

Endpoints

Endpoint Services

NAT Gateways

Peering Connections

Security

Network ACLs

Security Groups

Create Network ACL Delete

Search Network ACLs and the X

<< 1 to 2 of 2 Network ACLs >>

Name	Network ACL ID	Associated With	Default	VPC
☒	acl-ce2c15b5	1 Subnet	Yes	vpc-56a60a2d Test-VPC
☐	acl-72723b14	6 Subnets	Yes	vpc-030c8b65

acl-ce2c15b5

Summary **Inbound Rules** Outbound Rules Subnet Associations Tags

Allows inbound traffic. Because network ACLs are stateless, you must create inbound and outbound rules.

Edit

View: All rules

Rule #	Type	Protocol	Port Range	Source	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY

- Now, click on the blue **Create Network ACL** button to create a new ACL. Select our VPC and then enter a name for the ACL in the creation wizard:

Create Network ACL

A network ACL is an optional layer of security that acts as a firewall for controlling traffic in and out of a subnet.

Name tag ⓘ

VPC ⓘ

[Cancel](#) [Yes, Create](#)

- Now, in the inbound rules of the new ACL, add in the following rule. In the **Source** section, add the IPv4 setting of any of your public subnets and click **Save**:

acl-cc727db7 | Test-ACL

[Summary](#) [Inbound Rules](#) [Outbound Rules](#) [Subnet Associations](#) [Tags](#)

Allows inbound traffic. Because network ACLs are stateless, you must create inbound and outbound rules.

[Edit](#)

View:

Rule #	Type	Protocol	Port Range	Source	Allow / Deny
1	SSH (22)	TCP (6)	22	172.31.48.0/20	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY

7. Now, replace the ACL of our current subnet with the new one that will make our subnet a private subnet:

The screenshot shows the AWS VPC Dashboard. On the left, the 'Subnets' link is selected under the 'Virtual Private Cloud' section. The main area displays a list of subnets. The subnet 'subnet-949564de' is selected, and the 'Network ACL' tab is active. The 'Network ACL' configuration shows the 'Test-Subnet' associated with 'vpc-56a60a2d | Test-VPC'. The 'Network ACL' is 'acl-cc727db7 | Test-ACL'. The configuration is as follows:

Inbound:						
Rule #	Type	Protocol	Port Range / ICMP Type	Source	Allow / Deny	
1	SSH (22)	TCP (6)	22	172.31.48.0/20	ALLOW	
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY	

Outbound:						
Rule #	Type	Protocol	Port Range / ICMP Type	Destination	Allow / Deny	
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY	

Now, we have our Lambda function in a private subnet, making it more secure.

Controlling access to Lambda functions

We have gone through all the security settings needed to ensure that our Lambda functions and our serverless architectures are secure. So, an engineer working on serverless systems should keep the following points in mind while designing their infrastructure from a security point of view:

- The VPC and the subnet settings can be added under the **Network** section of the Lambda function.
- It is recommended that the Lambda function is placed across at least two subnets for fault tolerance purposes. However, this is not compulsory.
- If you are placing your Lambda function inside a private subnet, you need to ensure that the private subnet is receiving the appropriate traffic from your public subnet(s) in that VPC. If not, then the Lambda function is essentially locked out.

Using STS inside Lambda for secure session-based execution

While accessing other AWS services and components from inside your Lambda functions, you can make use of **AWS's Simple Token Service (STS)** to ensure session-based access, which will essentially add an extra layer of security. As we have already discussed, and learned how to use, STS credentials in our code, we will skip over to the documentation links.

The official documentation of AWS STS will help you understand how session-based access works: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html.

And this is the *Boto3 Python Documentation* for using STS credentials inside Python code: <http://boto3.readthedocs.io/en/latest/reference/services/sts.html>.

Summary

In this chapter, we have learned how security works in Lambda functions in a deep-dive mode. We have understood how VPCs and subnets work in the AWS environment. We have learned to create a VPC and also created public and private subnets. This will give you a better understanding of how security works from the whole of the AWS perspective.

We have also learned how to place your Lambda functions inside the VPCs and the subnets we have created throughout this chapter. We understood how to handle and route traffic inside our VPCs and subnets.

Finally, we also learned how to implement better security in our Python code using session-based access to other AWS components, thereby placing security in the control of the developer.

In the next chapter, you will learn about the **Serverless Application Model (SAM)** and how to write SAM models and deploy your Lambda applications through them.

8

Deploying a Lambda Function with SAM

So far, we have learned about Lambda functions and how to build them. We have learned that a Lambda function has a definite set of triggers that would trigger the function to carry out a particular task. The task is written as a Python module and the script is what we call a function. We have also learned about the different settings of Lambda functions, which include its core settings and also other settings, such as security and network.

There is also another alternative to creating and deploying Lambda functions, which is the **AWS Serverless Application Model (AWS SAM)**. This format is based on the concept of **infrastructure as code**. This concept is inspired by **AWS CloudFormation**, which is a form of infrastructure as code.

We will be learning about AWS CloudFormation and using that knowledge to understand and build AWS SAM models for creating Lambda functions. We will be covering the following concepts in this chapter:

- Deploying Lambda functions
- Using CloudFormation for serverless services
- Deploying with SAM
- Understanding security in SAM


Introduction to SAM

In this section, we will learn about SAM, which will help us build and deploy serverless functions:

1. As mentioned earlier, SAM is about writing infrastructure as code. So, this is what a Lambda function would be described as in SAM:

```
AWSTemplateFormatVersion: '2010-09-09'
Transform: AWS::Serverless-2016-10-31
Resources:
  < Name of function >:
    Type: AWS::Serverless::Function
    Properties:
      Handler: < index.handler >
      Runtime: < runtime >
      CodeUri: < URI of the bucket >
```

2. In this block of code, we enter the details—the name of the function, and the URI of the S3 bucket where our code package is hosted. In the same way that we named the index and the handler in our Lambda settings, we need to enter those details here, too. The `index.handler` is the file in which our function code is located. The `Handler` is the name of the function in which our Lambda logic is written. Also, the `Runtime` is user-defined. You can select from all the available languages that are supported by AWS Lambda. The scope of this book is limited to the Python language, so we will stick to either of the available Python versions:



3. We can also add environment variables, as shown here, in our Lambda function, too. These can be very easily edited and configured just as we add, update, and/or delete code, which is an added advantage of the infrastructure as code style of building infrastructures:

```

AWSTemplateFormatVersion: '2010-09-09'
Transform: AWS::Serverless-2016-10-31
Resources:
  PutFunction:
    Type: AWS::Serverless::Function
    Properties:
      Handler: index.handler
      Runtime: < runtime >
      Policies: < AWSLambdaDynamoDBExecutionRole >
      CodeUri: < URI of the zipped function package >
      Environment:
        Variables:
          TABLE_NAME: !Ref Table
  DeleteFunction:
    Type: AWS::Serverless::Function
    Properties:
      Handler: index.handler
      Runtime: nodejs6.10
      Policies: AWSLambdaDynamoDBExecutionRole
      CodeUri: s3://bucketName/codepackage.zip
      Environment:
        Variables:
          TABLE_NAME: !Ref Table
  Events:
    Stream:
      Type: DynamoDB
      Properties:
        Stream: !GetAtt DynamoDBTable.StreamArn
        BatchSize: 100
        StartingPosition: TRIM_HORIZON
  DynamoDBTable:
    Type: AWS::DynamoDB::Table
    Properties:
      AttributeDefinitions:
        - AttributeName: id
          AttributeType: S
      KeySchema:
        - AttributeName: id
          KeyType: HASH
      ProvisionedThroughput:
        ReadCapacityUnits: 5
        WriteCapacityUnits: 5

```

```
StreamSpecification:
  StreamViewType: streamview type
```

4. The preceding SAM code invokes two Lambda functions that point to an AWS DynamoDB table. The entire SAM code is an application that consists of a couple of Lambda functions. You need to enter the necessary details for making this work. The `Runtime` needs to be updated with either of the available Python runtimes. The corresponding policy for dealing with the DynamoDB tables needs to be updated in the `Policies` section. The `CodeUri` section needs to be updated with the S3 URI of the code package.
5. It is to be noted that the meta information should always be included for all SAM, which includes the `AWS::Serverless::SimpleTable` attribute. This would tell CloudFormation that the code you have written is an AWS SAM code and a serverless application. The two lines are as follows:

```
AWS::Serverless::SimpleTable:
  Transform: AWS::Serverless-2016-10-31
```

6. If your serverless function needs to access a single table of DynamoDB, you can start by creating a DynamoDB table via your SAM function itself using the `SimpleTable` attribute. This can be done as follows:

```
AWS::Serverless::SimpleTable:
  Transform: AWS::Serverless-2016-10-31
  Resources:
    < TableName >:
      Type: AWS::Serverless::SimpleTable
      Properties:
        PrimaryKey:
          Name: id
          Type: String
        ProvisionedThroughput:
          ReadCapacityUnits: 5
          WriteCapacityUnits: 5
```

7. Now, we will learn how to create a Lambda function with a trigger. As we are already using DynamoDB for the examples, we will use the same as a trigger in this step. The SAM code for this would look as follows:

```
AWS::Serverless::Function:
  Transform: AWS::Serverless-2016-10-31
  Resources:
    < Name of the function >:
      Type: AWS::Serverless::Function
      Properties:
```

```

    Handler: index.handler
    Runtime: < runtime >
    Events:
      Stream:
        Type: DynamoDB
        Properties:
          Stream: !GetAtt DynamoDBTable.StreamArn
          BatchSize: 100
          StartingPosition: TRIM_HORIZON
  < Name of the table >:
    Type: AWS::DynamoDB::Table
    Properties:
      AttributeDefinitions:
        - AttributeName: id
          AttributeType: S
      KeySchema:
        - AttributeName: id
          KeyType: HASH
      ProvisionedThroughput:
        ReadCapacityUnits: 5
        WriteCapacityUnits: 5

```

CloudFormation for serverless services

In this section, we will learn how CloudFormation can be used to build and deploy Lambda functions. We will do the following:

1. We will write a **CloudFormation** template for a Lambda function that periodically pings a website and gives an error if there is any failure in the process. The **CloudFormation** template for this is as follows:

```

AWSTemplateFormatVersion: '2010-09-09'
Transform: 'AWS::Serverless-2016-10-31'
Description: 'Performs a periodic check of the given site,
erroring out on test failure.'
Resources:
  lambdacanary:
    Type: 'AWS::Serverless::Function'
    Properties:
      Handler: lambda_function.lambda_handler
      Runtime: python2.7
      CodeUri: .
      Description: >-
        Performs a periodic check of the given site,
        erroring out on test failure.

```

```
MemorySize: 128
Timeout: 10
Events:
  Schedule1:
    Type: Schedule
    Properties:
      Schedule: rate(1 minute)
Environment:
  Variables:
    site: 'https://www.google.com/'
    expected: Search site.
```

2. There is a lot of syntax in this CloudFormation snippet. We will now try to understand it in a bit more detail:
 1. In the first three lines that contain the meta details of the Lambda function, we have the following line—`Transform:`
`'AWS::Serverless-2016-10-31'`. This line is used to define the resources that a user will be using/accessing, through a **CloudFormation** template. As we are using a Lambda function, we have specified it as `Serverless`.
 2. We have also defined the memory size that our function will be using. It is similar to how we learned to view and change the memory settings in the Lambda's console.
 3. `Timeout` is the amount of time the Lambda function can keep retrying before considering the attempt as a failure.

You can also see that we have added environment variables to our Lambda function that will be stored in the Lambda container and used when needed by the system.

Deploying with SAM

In this section, we will learn how to deploy the SAM applications. We have already learned what SAM applications and code look like, so we will learn how to deploy them via AWS CloudFormation:

1. Firstly, let's set up our local environment for deployment purposes, and then start by installing `awscli` from `pip`:

```

(venv) → Desktop pip install awscli
collecting awscli
  Downloading awscli-1.14.63-py2.py3-none-any.whl (1.2MB)
    100% |#####| 1.2MB 325kB/s
collecting boto3==1.9.16 (from awscli)
  Downloading boto3-1.9.16-py2.py3-none-any.whl (4.1MB)
    100% |#####| 4.1MB 183kB/s
collecting colorama==0.3.7, >=0.2.5 (from awscli)
  Downloading colorama-0.3.7-py2.py3-none-any.whl
collecting docutils==0.10 (from awscli)
  Downloading docutils-0.14-py2-none-any.whl (543kB)
    100% |#####| 552kB 404kB/s
collecting rsa==3.5.0, >=3.1.2 (from awscli)
  Downloading rsa-3.4.2-py2.py3-none-any.whl (46kB)
    100% |#####| 51kB 637kB/s
collecting PyYAML==3.12, >=3.10 (from awscli)
  Downloading PyYAML-3.12.tar.gz (253kB)
    100% |#####| 256kB 835kB/s
collecting s3transfer<0.2.0, >=0.1.12 (from awscli)
  Downloading s3transfer-0.1.13-py2.py3-none-any.whl (59kB)
    100% |#####| 61kB 204kB/s
collecting jmespath<1.0.0, >=0.7.1 (from boto3==1.9.16->awscli)
  Downloading jmespath-0.9.3-py2.py3-none-any.whl
collecting python-dateutil<2.7.0, >=2.1 (from boto3==1.9.16->awscli)
  Downloading python_dateutil-2.6.1-py2.py3-none-any.whl (194kB)
    100% |#####| 194kB 358kB/s
collecting pyasn1==0.1.3 (from rsa==3.5.0, >=3.1.2->awscli)
  Downloading pyasn1-0.4.2-py2.py3-none-any.whl (71kB)
    100% |#####| 71kB 468kB/s
collecting futures<4.0.0, >=2.2.0; python_version == "2.7" or python_version == "2.6" (from s3transfer<0.2.0, >=0.1.12->awscli)
  Downloading futures-3.2.0-py2-none-any.whl
collecting six>=1.5 (from python-dateutil<2.7.0, >=2.1->boto3==1.9.16->awscli)
  Downloading six-1.11.0-py2.py3-none-any.whl
building wheels for collected packages: PyYAML
  Running setup.py bdist_wheel for PyYAML ... done
  Stored in directory: /Users/Danny33/Library/Caches/pip/wheels/2c/f7/79/13f3a12cd723892437c0cfbde1230ab4d82947ff7b3839a4fc
successfully built PyYAML
installing collected packages: jmespath, six, python-dateutil, docutils, boto3, colorama, pyasn1, rsa, PyYAML, futures, s3transfer, awscli
successfully installed PyYAML-3.12 awscli-1.14.63 boto3-1.9.16 colorama-0.3.7 docutils-0.14 futures-3.2.0 jmespath-0.9.3 pyasn1-0.4.2 python-dateutil-2.6.1
rsa-3.4.2 s3transfer-0.1.13 six-1.11.0
(venv) → Desktop

```

2. Next, you will need to configure your AWS environment using your credentials:

```

(venv) → Desktop aws configure
AWS Access Key ID [*****OP2Q]:
AWS Secret Access Key [*****W0V]:
Default region name [None]:
Default output format [None]:

```

3. You will need to enter the following details to make sure your AWS environment is successfully configured:

- Your AWS Access Key
- Your AWS Secret Key
- The default region in which you want to operate
- The default output format in which you want your data

4. Now, let's try to deploy a simple Hello World Lambda application via SAM. We will have two code files for this. One is the Python file and the other is the template `yaml` file.

5. We will use the default `Hello World` example for Python, as we are trying to understand how SAM deployments work instead of stressing too much about the code for now. The Python script will be as follows:

```
import json
print('Loading function')
def lambda_handler(event, context):
    #print("Received event: " + json.dumps(event, indent=2))
    print("value1 = " + event['key1'])
    print("value2 = " + event['key2'])
    print("value3 = " + event['key3'])
    return event['key1'] # Echo back the first key value
    #raise Exception('Something went wrong')
```

6. We will use a basic template `yaml` file for the SAM function too, whose only job is to define its meta information and to run the Python script that is mentioned previously. The template `yaml` file will look like this:

```
AWSTemplateFormatVersion: '2010-09-09'
Transform: 'AWS::Serverless-2016-10-31'
Description: A starter AWS Lambda function.
Resources:
    helloworldpython3:
        Type: 'AWS::Serverless::Function'
        Properties:
            Handler: lambda_function.lambda_handler
            Runtime: python3.6
            CodeUri: .
            Description: A starter AWS Lambda function.
            MemorySize: 128
            Timeout: 3
```

7. Now, we will package the SAM template we just created, using the command line. The instructions for packaging the code are as follows:

```
aws cloudformation package --template-file template.yaml --
output-template-file output.yaml --s3-bucket receiver-bucket
```

You get the following output:

```
(venv) + SAM
(venv) + SAM
(venv) + SAM
(venv) + SAM aws cloudformation package --template-file template.yaml --output-template-file output.yaml --s3-bucket receiver-bucket
Uploading to 22067de83ab3b7a12a153fbd0517d6cf 670 / 670.0 (100.00%)
Successfully packaged artifacts and wrote output template to file output.yaml.
Execute the following command to deploy the packaged template
aws cloudformation deploy --template-file /Users/                               /SAM/output.yaml --stack-name <YOUR STACK NAME>
(venv) + SAM
(venv) + SAM
(venv) + SAM
```

8. This will create an output `yaml` file that needs to be deployed, as mentioned in the preceding trace. The `output.yaml` file looks like this:

```
AWSTemplateFormatVersion: '2010-09-09'
Description: A starter AWS Lambda function.
Resources:
  helloworldpython3:
    Properties:
      CodeUri: s3://receiver-
bucket/22067de83ab3b7a12a153fbd0517d6cf
      Description: A starter AWS Lambda function.
      Handler: lambda_function.lambda_handler
      MemorySize: 128
      Runtime: python3.6
      Timeout: 3
    Type: AWS::Serverless::Function
  Transform: AWS::Serverless-2016-10-31
```


9. Now, as we have packaged the SAM template, we will now deploy it. We will use the instructions shown in the trace when we did the packaging for the deployment process. The instructions for deployment are as follows:

```
aws cloudformation deploy --template-file
/Users/<path>/SAM/output.yaml --stack-name 'TestSAM' --
capabilities CAPABILITY_IAM
```


This will give you the following output:

```
(venv) + SAM
(venv) + SAM
(venv) + SAM aws cloudformation deploy --template-file /Users/                SAM/output.yaml --stack-name 'TestSAM' --capabilities CAPAB
ILITY_IAM
Waiting for changeset to be created..
Waiting for stack create/update to complete
Successfully created/updated stack - TestSAM
(venv) + SAM
```


10. We can head over to the **CloudFormation** console to look at the template we just deployed. The deployed template will look something like this:




11. In the **Template** tab shown here, we can see both the original template and the processed template. The original template can be seen by selecting the first radio button:




12. The processed template can be seen by selecting the second radio button under the **Template** tab at the bottom:




13. If we head over to the **Lambda** console, we will see the newly created Lambda function via SAM with the corresponding name given:




14. Clicking on the **Functions** will give us more information about it. It also mentions the SAM template and the **CloudFormation** template from which it was created:




15. Let's create basic tests for the Lambda function. The test creation console can be opened by clicking on the **Test** button:




16. Now, once the tests have been created, you can again click on the **Test** button. This will run the testing with the updated test cases. The logs from a successful run will look like this:




17. Now, let's go through each component of the Lambda function properly. The **Configuration** shows the triggers and the logging settings of our Lambda function. We are logging into the CloudWatch service of AWS:



18. We can also see the invocation metrics in the **Monitoring** option in the **Lambda** console. We can see exactly one Lambda invocation:



19. You can see the code files in the **Function code** section. You can see the folder structure in the left-hand corner of the interactive code editor that contains both the `template.yaml` file and the function code:




20. And further below, you can see the pre-existing environment variable named `lambda:createdBy`, and also the timeout setting we mentioned in our template.

Understanding security in SAM

So far, we have learned how to write, build, package, and deploy Lambda functions using the SAM. We will now understand how security works inside them:

1. You can scroll to the bottom of the Lambda console to see the network and security settings, where the VPC and the subnet details are mentioned:



2. Now, we will add in the network settings, which include the security groups and the subnet IDs:

```
AWSTemplateFormatVersion: '2010-09-09'
Transform: 'AWS::Serverless-2016-10-31'
Description: A starter AWS Lambda function.
Resources:
  helloworldpython3:
    Type: 'AWS::Serverless::Function'
    Properties:
      Handler: lambda_function.lambda_handler
      Runtime: python3.6
      CodeUri: .
      Description: A starter AWS Lambda function.
      MemorySize: 128
      Timeout: 3
      VpcConfig:
        SecurityGroupIds:
          - sg-9a19c5ec
        SubnetIds:
          - subnet-949564de
```

3. Now, package and deploy the newly updated SAM template like we did in the previous section:

```
(venv) ➔ SAM vim template.yaml
(venv) ➔ SAM aws cloudformation package --template-file template.yaml --output-template-file output.yaml --s3-bucket receiver-bucket
Uploading to f94a786f2b5b4a8d25bbf5da5a6c7527 1158 / 1158.0 (100.00%)
Successfully packaged artifacts and wrote output template to file output.yaml.
Execute the following command to deploy the packaged template
aws cloudformation deploy --template-file /Users/Dawny33/Desktop/Work/ServerlessBook/SAM/output.yaml --stack-name <YOUR STACK NAME>
(venv) ➔ SAM
(venv) ➔ SAM
(venv) ➔ SAM aws cloudformation deploy --template-file /Users/Dawny33/Desktop/Work/ServerlessBook/SAM/output.yaml --stack-name 'TestSAMSec' --capabilities CAPABILITY_IAM
Waiting for changeset to be created..
Waiting for stack create/update to complete
Successfully created/updated stack - TestSAMSec
```

4. Now you will see the corresponding network and security settings, once you have packaged and deployed the **CloudFormation** template after the corresponding edits. The **Network** section looks as follows:

TestSAM-helloworldpytho... Throttle Qualifiers Actions TestEvent Test Save

Network

VPC [Info](#)
Select a VPC that your function will access.
vpc-56a60a2d (172.31.0.0/16) | Test-VPC

Subnets*
Select the VPC Subnets that Lambda should use to set up your VPC configuration. Format: "subnet-id [cidr-block] | az name-tag".
subnet-949564de (172.31.0.0/16) | us-east-1b Test-Subnet

Security Groups*
Select the VPC Security Groups that Lambda should use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The table below will show the inbound and outbound rules for the security groups you have selected.
sg-9a19c5ec (default)

Debugging and error handling

DLQ Resource [Info](#)
Choose the AWS service to send event payload to after exceeding maximum retries.
None

Enable active tracing [Info](#)
☐

When you enable VPC, your Lambda function will lose default internet access. If you require external internet access for your function, ensure that your security group allows outbound connections and that your VPC has a NAT gateway.

5. You can also see the inbound rules of your corresponding security groups that are linked with the VPC in your **Network** settings:

TestSAM-helloworldpytho... Throttle Qualifiers Actions TestEvent Test Save

When you enable VPC, your Lambda function will lose default internet access. If you require external internet access for your function, ensure that your security group allows outbound connections and that your VPC has a NAT gateway.

Inbound rules Outbound rules

Security group ID	Ports	Source
sg-9a19c5ec	All	sg-9a19c5ec

Concurrency

Unreserved account concurrency **1000**


☒ Use unreserved account concurrency
☐ Reserve concurrency

Auditing and compliance

This function's invocations can be logged by CloudTrail for operational and risk auditing, governance, and compliance. Visit the [CloudTrail console](#) to get started.

Feedback English (US) © 2008 - 2018, Amazon Internet Services Private Ltd. or its affiliates. All rights reserved. Privacy Policy Terms of Use

6. You can also see the completed **CloudFormation** template in your console with the updated network and security settings, which means that deployment has been successful:



7. You can also see the original template under the **Templates** option in the bottom corner of the console:

The screenshot shows the AWS CloudFormation console with the 'TestSAMSec' stack selected. The 'Template' tab is active, showing the following resource definition:

```

Resources:
  helloworldpython3:
    Properties:
      CodeUri: s3://receiver-bucket/f94a786f2b5b4a8d25bbf5da5a6c7527
      Description: A starter AWS Lambda function.
      Handler: lambda_function.lambda_handler
      MemorySize: 128
      Policies: AdministratorAccess
      Runtime: python3.6
      Timeout: 3
      VpcConfig:
        SecurityGroupIds:
          - sg-9a19c5ec
        SubnetIds:
          - subnet-949564de
  
```

8. The processed template can be found by selecting the **View processed template** option beside the original template option at the bottom of the console:

The screenshot shows the AWS CloudFormation console with the 'TestSAMSec' stack selected. The 'Template' tab is active, and the 'View processed template' option is selected. The processed template is displayed as a JSON string:

```

{"AWSTemplateFormatVersion":"2010-09-09","Description":"A starter AWS Lambda function.","Resources":{"helloworldpython3Role":{"Type":"AWS::IAM::Role","Properties":{"ManagedPolicyArns":["arn:aws:iam::aws:policy/service-role/AWSLambdaBasicExecutionRole","arn:aws:iam::aws:policy/AdministratorAccess"],"AssumeRolePolicyDocument":{"Version":"2012-10-17","Statement":[{"Action":["sts:AssumeRole"],"Effect":"Allow","Principal":{"Service":["lambda.amazonaws.com"]}}]},"helloworldpython3":{"Type":"AWS::Lambda::Function","Properties":{"Code":{"S3Bucket":"receiver-bucket","S3Key":"f94a786f2b5b4a8d25bbf5da5a6c7527"},"Description":"A starter AWS Lambda function.","VpcConfig":{"SubnetIds":["subnet-949564de"],"SecurityGroupIds":["sg-9a19c5ec"]},"Tags":[{"Value":"SAM","Key":"lambda:createdBy"},"MemorySize":128,"Handler":"lambda_function.lambda_handler","Role":{"Fn::GetAtt":["helloworldpython3Role","Arn"]},"Timeout":3,"Runtime":"python3.6"}}}}}
  
```

Summary

In this chapter, we learned how to deploy Lambda functions as infrastructure as code via SAM, which is a new way of writing and deploying Lambda functions. This makes it easier to integrate with other IaaS services, such as CloudFormation. We also learned about the AWS CloudFormation service, which is the service that allows and facilitates infrastructure as code. We also learned how security works inside SAM code and how to configure VPC and subnet settings.

In the next chapter, you will be introduced to Microsoft Azure functions, along with configuring and understanding the components of the tool.

9 Introduction to Microsoft Azure Functions

So far, we have learned how to build serverless functions and serverless architectures using Python in the AWS environment. We have also learned about the settings and environment of the AWS Lambda tool in great detail. We shall now learn and explore its counterpart from Microsoft Azure Functions.


In this chapter, you will learn how Microsoft Azure Functions work, what the Microsoft Azure Functions console looks like, and how to go about understanding the settings in the console. This chapter is divided into the following sections:

- Introduction to Microsoft Azure Functions
- Creating your first Azure Function
- Understanding triggers
- Understanding logging and monitoring
- Best practices for writing Microsoft Azure Functions


Introduction to Microsoft Azure Functions

Microsoft Azure Functions is the Azure counterpart of AWS's Lambda service. In this section, we will learn how to locate and navigate the Microsoft Azure Functions console. So, let's start by performing the following steps:


1. You can locate the Azure Functions app by navigating to the **All services** tab on the left menu and typing out the function filter. You will now notice the Microsoft Azure Function's service under the name, **Function Apps**:



2. Once you click on that, you will be re-directed to the **Function Apps** console. For now, it will be empty if you haven't created any functions. The console will look something like this:



3. Now, let's start with creating an Azure Function. To do so, we need to click on the **Create a resource** option on the left menu, then click on the **Compute** option from that list, and then select the **Function App** option from the subsequent list of options:



Microsoft Azure Functions come under the list of **Compute** resources on the dashboard. In the following sections, we will learn how to create Microsoft Azure Functions and also understand the different kinds of triggers and how they work.

Creating your first Azure Function

In this section, we will learn how to create and deploy an Azure Function. We will go through the process step by step in order to understand how each section of an Azure function works:

1. When you click on the **Functions App** in the menu, you will be re-directed to the **Function App** creation wizard, as shown in the following screenshot:


The screenshot shows the Azure portal's 'Function App' creation wizard. The left sidebar contains a navigation menu with options like 'Create a resource', 'All services', 'FAVORITES', 'Dashboard', 'All resources', 'Resource groups', 'App Services', 'Function Apps', 'SQL databases', 'Azure Cosmos DB', 'Virtual machines', 'Load balancers', 'Storage accounts', 'Virtual networks', and 'Azure Active Directory'. The main panel is titled 'Function App' and 'Create'. It contains several configuration fields: 'App name' (text input with placeholder 'Enter a name for your App'), 'Subscription' (dropdown menu showing 'Free Trial'), 'Resource Group' (radio buttons for 'Create new' and 'Use existing'), 'OS' (tabs for 'Windows' and 'Linux (Preview)'), 'Hosting Plan' (dropdown menu showing 'Consumption Plan'), 'Location' (dropdown menu showing 'Central US'), and 'Storage' (radio buttons for 'Create new' and 'Use existing'). At the bottom, there is a 'Pin to dashboard' checkbox and a 'Create' button. A 'Automation options' link is also visible.

2. Add the required information in the wizard accordingly. Choose **Linux (Preview)** as the OS. Then, click on the blue **Create** button at the bottom of the wizard:


The screenshot shows the Azure portal interface for creating a new Function App. The left sidebar contains navigation links: 'Create a resource', 'All services', 'FAVORITES', 'Dashboard', 'All resources', 'Resource groups', 'App Services', 'Function Apps', 'SQL databases', 'Azure Cosmos DB', 'Virtual machines', 'Load balancers', 'Storage accounts', 'Virtual networks', and 'Azure Active Directory'. The main panel is titled 'Function App' and 'Create'. It contains the following fields and options:

- App name:** packt-test (with a green checkmark and .azurewebsites.net domain)
- Subscription:** Free Trial (dropdown menu)
- Resource Group:** packt-test (radio buttons for 'Create new' and 'Use existing')
- OS:** Linux (Preview) (selected, with Windows as an alternative)
- App Service plan/Location:** ServicePlan97e39185-afa5(West ...)
- Storage:** packttestabb8 (radio buttons for 'Create new' and 'Use existing')
- Application Insights:** On (toggle switch)
- Pin to dashboard:** (checkbox)
- Create** button and **Automation options** link at the bottom.


3. Clicking on the **Automation options** at the bottom will open up a validation screen for automating Function deployments. This is not needed for this chapter. This will simply validate your Azure Function:




4. Once you click **Create**, you will see the deployment in progress under the **Notifications** menu:



- Once it has been successfully created, it will be reflected in your notifications list with a green-colored notification:



- Clicking on **Go to resource** will take you to the newly created Azure Function. The function console will look like this:




We have successfully created an Azure Function. We will cover in more detail triggers, monitoring, and security in the forthcoming sections of this chapter.


Understanding triggers

In this section, we will look at how triggers work in Azures Function applications. We will also learn about the different types of triggers and their purpose. Perform the following steps:


1. In the left menu, click on the (+) symbol beside the **Functions** option for adding, removing, or editing a trigger:




2. You will be taken to the function creation console, which looks like this:




3. Azure does not have a lot of support for Python. So, in this console, let's choose a custom function of our own. Click on **Custom function** under the **Get Started on your own** option at the bottom:




4. In the function creation wizard, enable the **Experimental Language** option in the right menu. Now, you will be able to see the **Python** option in the available languages:




5. There are two triggers that are available for the Python language. One is the **HTTP trigger** and the other is the **Queue trigger**, as seen in the following screenshot:



6. The **HTTP trigger** will trigger the function whenever it receives an HTTP request. When you click on it, you will notice options for adding different HTTP-related settings, such as authorization and name:



7. The next trigger is the **Queue trigger**. This will trigger the function whenever a message is added to the queue. We have done the same in AWS Lambda in one of our previous chapters too:



Understanding logging and monitoring in Azures Functions

In this section, we will learn and understand the monitoring and logging mechanisms available to the user in Microsoft Azure Functions. Perform the following steps:


1. By clicking on the **Monitor** option under the function, we can access the monitoring suite of that particular Azure Function:

The screenshot displays the Azure Portal interface for the 'Test-Packt - HttpTriggerPython31' function. The left sidebar shows the 'All resources' view with a list of resources including 'CentralUSPlan', 'Test-Packt', and 'testpacktb6bf'. The main content area is divided into three sections:


- Pay-As-You-Go**: A summary section showing the success count since April 1st as 1, and the error count as 0.
- Invocation log**: A table showing the function's execution history. The table has columns for 'Function', 'Status', 'Details: Last ran', and '(duration)'. A single entry is shown for 'HttpTriggerPython31 (Method: POST, Uri: ...)' with a status of '✓' and a duration of 'a day ago (1,780 ms)'.
- live event stream**: A link to view the live event stream.

The 'Monitor' option is highlighted in the left sidebar, indicating the current view.


2. The monitoring suite for the function that we created looks like this:



3. Now, click on the **Open Application Insights** option at the top of the menu. This will take you to the detailed monitoring page:



- 4. If you scroll down, you will see the function-specific metrics, such as the server response times and request performance. This is very useful as it means we don't need separate dashboards for monitoring all these statistics:




Now that we have learned about Microsoft Azure Functions logging and monitoring, let's go through some best practices.


Best practices for writing Azure Functions

We have learned how to create, configure, and deploy Microsoft Azure Functions. We will now learn about the best practices for using them:


- Microsoft Azure Functions don't have a huge support for Python like AWS Lambda. They have a very limited set of Python-based triggers. So, you need to write custom functions for most of them. Developers need to keep that in mind before taking a decision on using Microsoft Azure Functions. The languages supported by Microsoft Azure Functions are **C#**, **F#**, and **JavaScript**:



- The experimental languages that are supported by Microsoft Azure Functions are **Bash**, **Batch**, **PHP**, **TypeScript**, **Python**, and **PowerShell**:



- Make sure you use the security settings properly to secure your functions. You can find all the settings you need in the **Platform features** options:



- Finally, use monitoring as much as possible, as it is crucial to log and monitor serverless functions. We have already gone through the monitoring details and the corresponding settings.

Summary

In this chapter, we learned about Microsoft Azure Functions and how to build them. We learned about the various functionalities available, along with the available triggers for the Python runtime. We also learned and experimented with the logging and monitoring capabilities of Microsoft Azure Functions along with understanding and experimenting with the experimental features of Azure such as the additional runtimes apart from the standard set of languages it offers out of the box.

Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:



Serverless computing in Azure with .NET Sasha Rosenbaum

ISBN: 978-1-78728-839-3

- Understand the best practices of Serverless architecture
- Learn how to deploy a Text Sentiment Evaluation application in an Azure Serverless environment
- Implement security, identity, and access control
- Take advantage of the speed of deployment in the cloud
- Configure application health monitoring, logging, and alerts
- Design your application to ensure cost effectiveness, high availability, and scale



Building Serverless Architectures

Cagatay Gurturk

ISBN: 978-1-78712-919-1

- Learn to form microservices from bigger Softwares
- Orchestrate and scale microservices
- Design and set up the data flow between cloud services and custom business logic
- Get to grips with cloud provider's APIs, limitations, and known issues
- Migrate existing Java applications to a serverless architecture
- Acquire deployment strategies
- Build a highly available and scalable data persistence layer
- Unravel cost optimization techniques

Leave a review - let other readers know what you think

Please share your thoughts on this book with others by leaving a review on the site that you bought it from. If you purchased the book from Amazon, please leave us an honest review on this book's Amazon page. This is vital so that other potential readers can see and use your unbiased opinion to make purchasing decisions, we can understand what our customers think about our products, and our authors can see your feedback on the title that they have worked with Packt to create. It will only take a few minutes of your time, but is valuable to other potential customers, our authors, and Packt. Thank you!

Index

A

- access
 - controlling, to Lambda functions 209
- Amazon Web Services (AWS) 5
- Ansible
 - reference link 171, 175
- API execution
 - Lambda function, deploying for 111, 112, 114, 115, 116, 117
- API Gateway 17, 94
- API resources 94
- AssumeRole API
 - reference link 183
- authentication
 - handling 118, 119, 120, 121, 122, 123, 124, 125
- AWS cookbook
 - reference link 169
- AWS Key Management Service (KMS) 21, 29
- AWS Lambda, triggers
 - about 6, 16
 - API Gateway 17
 - AWS Simple Notification Service (AWS SNS) 20
 - CloudWatch 18
 - S3 19
- AWS Secure Token Service (STS) 183
- AWS Serverless Application Model 211
- AWS Simple Notification Service (AWS SNS) 18
- AWS Simple Queue Service (SQS) 74
- AWS SNS 20
- AWS's Virtual Private Cloud (VPC) 32
- Azure Function
 - best practices 250, 251, 252
 - creating 235, 237, 238
 - logging 246, 247, 248, 249
 - monitoring 246, 247, 248, 249

- triggers, working 239, 241, 242, 243, 244, 245

C

- Chef
 - reference link 167
- CloudFormation
 - about 211
 - used, for building Lambda functions 215, 216
 - used, for deploying Lambda functions 215, 216
- CloudTrail
 - about 138, 140, 141, 142, 143, 144
 - reference link 138
- CloudWatch trigger
 - about 84
 - working 85, 86, 87, 88, 89, 91
- CloudWatch
 - about 18, 128
 - functions 128, 129, 130, 132, 133, 134, 135, 136
 - Lambda metrics 145, 146, 148, 150, 152, 153, 154, 155
 - Lambda's logs 155, 156, 157, 158, 159, 160
- cookbooks 171
- cron facility
 - in triggers 9

D

- Dead Letter Queue (DLQ) 32
- deployment package
 - about 7, 23
 - creating 44, 45, 46, 48
- Docker
 - reference link 22
- Dynamic Host Configuration Protocol (DHCP) 195

E

event-based architectural designs 5

F

function as a service (FAAS) 6

I

Identity Access Management (IAM) 30

integration

setting up 101

L

Lambda function

about 21

access, controlling to 209

Advanced settings page 22

configuring 24, 26, 27, 28, 30, 31, 32

deploying, for API execution 111, 112, 114, 115, 116, 117

executing, as containers 22, 23

logging statements 160, 162, 163, 165

securing, inside private subnets 206, 207, 208, 209

testing 34, 35, 38

versioning 39, 40, 41, 42, 43

Lambda's logs

in CloudWatch 155, 156, 157, 158, 159, 160

Lambda's metrics

in CloudWatch 155

logging statements

in Lambda 160, 162, 163, 165

logging

in Azure Function 246, 247, 248, 249

M

metrics, Lambda

in CloudWatch 145, 146, 148, 150, 152, 153, 154

microservice architecture 7, 8

Microsoft Azure Functions 232, 233, 234

Microsoft's Azure 5

monitoring

in Azure Function 246, 247, 248, 249

N

Network Access Control List (ACL) 195

nohup

reference link 176

P

platform as a service (PaaS) 6

Q

queues 9, 10

S

S3 19

S3 trigger

about 51

working 51, 52, 53, 54, 55, 56, 57, 58, 59

scaling

difficulties, handling 189, 190

difficulties, identifying 189, 190

secure session-based execution

STS, using for 210

Security Groups (SGs) 173

security

best practices 182, 183, 185, 187

serverless API

building 102, 103, 104, 106, 107, 108, 110

creating 94, 96, 97, 98, 100, 101

Serverless Application Model (SAM)

about 2, 212, 214

applications, deploying 216, 217, 218, 219, 220, 221, 222, 223, 224

security 225, 226, 228

serverless architectures

about 5, 6, 7

advantages 11

disadvantages 11, 12

reasons, for debugging difficulties 13

scaling, considerations 12

servers

creating 175, 178, 179

terminating 175, 178, 179

SNS trigger

about 61

working 61, 62, 63, 65, 66, 67, 68, 69, 71, 72, 73

SQS trigger

about 74

working 75, 77, 79, 81, 82, 83, 84

STS

using, inside Lambda for secure session-based execution 210

T

third-party orchestration tools

Ansible 171, 173, 174

Chef 167, 169

triggers

working 239, 241, 242, 243, 244, 245

U

user controls

handling 118, 119, 120, 121, 122, 123, 124, 125

V

Version Control System (VCS) 39

Virtual Private Cloud (VPCs)

about 173, 193, 195, 196, 197, 198, 199

subnets 200, 201, 203, 204