

Learning Azure API Management

A beginner's guide to creating and managing APIs

Naman Sinha



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India


ISBN: 978-93-65893-694

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.



Dedicated to

In loving memory of my mother, whose presence I long for each day with all my heart. I miss you more than words can express.

This book is a small tribute to the love you gave me and is a way to keep your memory alive.

About the Author

Naman Sinha is a seasoned software engineer with extensive experience in the tech industry. He has earned numerous awards for his contributions, focusing on innovation in technology throughout his career. Naman specializes in Microsoft technologies, particularly Azure, .NET, and Power BI.

Naman is not just a technologist but also a prolific writer, educator, and mentor. Based out of Bengaluru, he regularly conducts corporate training sessions and tech talks, inspiring the next generation of developers and IT professionals across the globe.

Naman has built a diverse career, starting with a couple of start-ups and then serving renowned tech giants such as Fractal, Philips, and Zeiss, where he has honed his skills across various technologies to build great products. He holds hands-on experience as an Azure DevOps Engineer, Azure Developer, and Azure Data and AI Engineer. He holds a Bachelor of Engineering in Computer Science with a first-class distinction and has been honored by many prestigious institutions for his contributions towards the engineering community.

Passionate about giving back to the tech community, Naman actively contributes to public forums like Stack Overflow, Medium, Quora, Topmate, and GitHub. His insights and expertise are widely shared, benefiting countless professionals and enthusiasts. He has also authored many articles on Medium and KnowledgeHut upGrad on various technologies and global issues. With a popular LinkedIn newsletter, “Discover with Naman,” he continues to empower people and explore the latest in technology.

Naman’s decision to write this book stems from his passion for simplifying complex topics and enabling others to harness the full potential of cloud platforms. Having worked extensively with Azure API Management, Naman has seen firsthand how a well-managed API strategy can transform an organization’s digital capabilities.

About the Reviewer

Sanyam Jain is an accomplished cloud security engineer with a rich and varied background in the cybersecurity industry. His dedication to security transcends professional obligations, showcasing his deep commitment to protecting digital ecosystems.

Throughout his career, Sanyam has held pivotal roles in cloud security, security operations, and security automation, where he has exhibited a comprehensive understanding of cloud security, security automation, and operational security. He continually elevates enterprise security by devising and implementing robust strategies that safeguard critical infrastructure.

Sanyam's proficiency spans numerous security practices, including network security, threat detection, data encryption, and access control. He possesses technical expertise in various cloud platforms, including AWS, Azure, and Google Cloud.

Beyond his hands-on work, Sanyam's innovative research has garnered recognition, being featured in renowned publications such as Forbes, TechCrunch, and over 40 other platforms. This acclaim underscores his authority and credibility in the cybersecurity field.

Educationally, Sanyam holds a Master's degree in Technology from BITS Pilani, achieved with distinction. His professional journey is marked by leading significant projects and initiatives that advance enterprise security in the organizations he has served.

As the India Chapter Lead at CSIRT.Global and a volunteer security researcher, Sanyam mentors upcoming researchers. His relentless pursuit of knowledge and significant contributions to cybersecurity distinguish him as a leading professional committed to pioneering a safer digital world.

Acknowledgement

I owe my deepest gratitude to my parents, *the deceased Menka Sahay and Mr. Rajesh Kumar Sinha*, whose unwavering love and support have been the foundation of everything I've achieved. Your encouragement has been my guiding light, and this book is a testament to the values you've instilled in me.

To my elder parents, *Mrs. Poonam and Mr. Prabhat Kumar*, thank you for your wisdom and the life lessons you have shared. Your guidance has always been a source of strength, and I am forever grateful for the example you have set.

To my fiancée, *Nishika Tiku*, your love, patience, and belief in me have made this journey possible. Thank you for standing by my side through the highs and lows. You have truly outdone yourself in every way, and I could not ask for a better person.

A heartfelt thanks to my family members, friends, and relatives. Whether through kind words, thoughtful advice, or simply being there, your support has meant the world to me. Each of you has contributed to this journey in your special way.

I am profoundly grateful to the following individuals who have been my guides and mentors. *Mr. Bhanu Pratap* introduced APIs to me when I started my career. Since then, there has been no looking back, and your extraordinary guidance has largely shaped this work. *Mr. Varun Sinha*, my cousin and mentor who has always been there for me throughout, your insights and support have been invaluable, and your belief in my abilities has pushed me to strive for excellence.

I also wish to thank BPB Publications for their unwavering support and professionalism in bringing this book to life. Your dedication to quality and your commitment to this project have been remarkable.

To the reviewers, technical experts, and editors, thank you for your meticulous attention to detail and invaluable feedback. Your expertise has elevated this work to a higher standard.

A special acknowledgement to the readers of this book. Your interest and engagement are what make this endeavor worthwhile. I hope this book serves as a valuable resource on your journey.

Lastly, to everyone else who has contributed to this journey in ways both big and small—thank you. Your support, upliftment, and faith in this project have made all the difference.

Preface

The age of digital transformation and more autonomy has established that there is essentially an expedited demand for system-to-system integration, with data flooding in innumerable volumes. Now more than ever, managing APIs efficiently and securely at scale is of utmost importance. Azure **API Management (APIM)** offers a robust solution to meet these growing needs, empowering organizations to streamline their API ecosystem, optimize performance, and ensure security in their digital landscape.

As a tech lead and an avid follower of Microsoft technologies (more specifically Azure), my experience with APIs so far has been both enriching and challenging. This book encapsulates years of API design, development, and management experience that spans across different sectors. My goal is to provide you with a complete guide to Azure API Management, starting with the basics, which will help beginners, and then advanced methods. For developers who want to get hands-on experience with Azure API Management, solutions architects looking for ways to integrate APIM into enterprise-scale solutions, and everyone in between, this book offers something valuable by way of guidance.

This book has been structured in a manner that starts with the basics of Azure API Management while slowly moving on to cover more advanced concepts by providing crystal clear explanations and real-world use cases. We start with simple questions: what are APIs, and why do we need their management in modern software architectures? We then move on to exploring the features and capabilities of Azure's API Management platform using best practices for implementation in your own projects. Every chapter continues from the previous and provides gradually accumulating insights that give you the essential knowledge to guide your hands through mastery of Azure APIM.

In the process of writing this book, I also thought about the diverse audience it may attract—from students and budding developers to seasoned professionals and enterprise-level engineers. The aim is to make the ideas understandable and offer practical guidance that you can apply directly to your projects.

This book will be a useful road map to your API management journey, helping you navigate the growing world of API Management and maximize what Azure offers. Learn to unveil the patterns that have helped countless companies innovate faster in Azure using Azure API Management.

Chapter 1: Introduction to Azure – This introductory chapter explains the fundamental concepts of cloud with Microsoft Azure as an example. The idea is to introduce the cloud

to a beginner and help the reader understand its need in today's **Information Technology (IT)** industry. The chapter touches upon the various offerings from Microsoft and provides an overview of the various services (resources, as we shall call them moving forward) so that one knows the capabilities of the Azure Cloud platform and the potential problems it is trying to solve for start-ups as well as for the industry.

Chapter 2: API Overview – This chapter introduces **application programming interfaces (APIs)** and shall help one understand the role they play in any software or web application. The intent is to ensure that the readers understand why we need APIs, what value they bring, and how the development of reliable, secure, and faster APIs can make a difference in any software.

Chapter 3: Introduction to Azure API Management – This chapter introduces the Azure API Management (APIM) service offered by Microsoft Azure. We shall discuss the basic features and concepts of API management and get an overview of the service. This will help us get started with the service and guide us in creating a new API management resource in Azure.

Chapter 4: Building My First API – This chapter introduces the Azure Functions service offered by Microsoft Azure. We shall start by creating a simple Azure Function locally with some basic code and eventually publish the same in Azure. This will also help us understand the process of cloning a repository from GitHub and making use of the code. Notably, we shall use the same API code in the following chapters as well.

Chapter 5: APIM in Focus: Navigating the Basics – This chapter introduces the Azure API Management service and a few of the basic options available in it. While there are many things that one can do with API Management, this chapter intends to make the readers comfortable with the resource by navigating to the first few options available in it, most of which will not have a major impact on the resource.

Chapter 6: APIs and Products – This chapter digs deeper into two of the core features of the Azure API Management service. APIs are the fundamental part of the service, on top of which the Products and Subscriptions are built. This is useful for enterprises that aim to build reusable services and wish to expose them in the form of APIs to multiple clients.

Chapter 7: Users, Groups and Subscriptions – This chapter digs deeper into three more features of Azure API Management service. Users and Groups are important for the API management resource to identify users and manage them effectively with the right privileges. Subscription is useful for enterprises that aim to build reusable services and wish to expose them in the form of APIs to multiple clients with custom needs. For the readers who have not followed this book from the beginning and wish to read only

these topics, a fair idea of API management is a must, and reading the last chapter is recommended to get the most out of this chapter.

Chapter 8: Versions and Revisions – Understanding versions and revisions is crucial for effective API governance in Azure API Management. Versions allow for simultaneous deployment of APIs with different functionalities, while revisions enable iterative improvements without disrupting existing services. Mastering these concepts empowers seamless API lifecycle management in Azure. This chapter will help you with these concepts.

Chapter 9: Policies – This chapter delves into policies in Azure API Management and provides a comprehensive overview of how policies can be leveraged to modify API behavior through configuration. By understanding policy concepts, configuration files, and sections, readers gain the knowledge needed to create and implement policies effectively within Azure API Management for enhanced API control and management.

Chapter 10: Monitoring – This chapter explores the various options under the monitoring section of Azure API Management. It revolves around understanding the metrics using application insights and analytics pane, managing logs using logs, exporting them via diagnostic settings, acting on advisor recommendations, etc. These options together offer a comprehensive toolset for monitoring APIs in Azure API Management or any Azure resource in general.

Chapter 11: Infrastructure and Deployment – This chapter digs deeper into a lot more options available in Azure API Management related to deployments and infrastructure. These options are usually utilized after we have stabilized the APIs deployed in the resource and are vital for smooth operations and achieving enterprise-centric goals.

Chapter 12: Security and Automation – This chapter includes a few less commonly used options in Azure API Management related to security and automation. While managed identities and certificates ensure the secure functioning of the service, the defender for cloud option offers suggestions to improve the infrastructure.

Code Bundle and Coloured Images

Please follow the link to download the *Code Bundle* and the *Coloured Images* of the book:

<https://rebrand.ly/3gvaqxr>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/Learning-Azure-API-Management>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At **www.bpbonline.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 BPB books and eBooks.

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 **business@bpbonline.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 **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights 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.

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 BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Introduction to Azure	1
Introduction	1
Structure	1
Objectives	2
Introduction to cloud.....	2
History and evolution	3
Cloud delivery models.....	4
<i>Infrastructure-as-a-service</i>	7
<i>Platform-as-a-service</i>	7
<i>Software-as-a-service</i>	8
Cloud deployment models	9
<i>Public cloud</i>	10
<i>Private cloud</i>	11
<i>Hybrid cloud</i>	12
<i>Community cloud</i>	13
<i>Multi-cloud</i>	13
Microsoft Azure as a cloud provider.....	13
Azure essentials.....	15
<i>Subscription</i>	16
<i>Resource group</i>	16
<i>Resource</i>	16
<i>Azure management options</i>	17
<i>Azure Resource Manager</i>	18
<i>Tags</i>	18
<i>Policies</i>	19
<i>Role-based access control</i>	19
Conclusion	20
2. API Overview	21
Introduction	21

Structure	21
Objectives	22
The concept	22
Introduction to API	23
Local and remote APIs	24
Types of APIs	24
<i>Public APIs</i>	25
<i>Private APIs</i>	25
<i>Partner APIs</i>	26
<i>Composite APIs</i>	26
API protocols	26
<i>Simple Object Access Protocol</i>	27
<i>Remote Procedure Calls</i>	27
<i>WebSocket</i>	28
<i>GraphQL</i>	28
<i>Representational State Transfer</i>	28
REST APIs	29
<i>Benefits of REST APIs</i>	29
<i>HTTP and REST APIs</i>	30
CRUD operations	30
Basic HTTP methods	31
<i>GET</i>	32
<i>POST</i>	32
<i>PUT</i>	33
<i>DELETE</i>	34
<i>PATCH</i>	34
<i>HEAD</i>	34
<i>OPTIONS</i>	34
<i>CONNECT</i>	34
<i>TRACE</i>	34
HTTP responses	35
<i>HTTP response codes</i>	36
Conclusion	37

3. Introduction to Azure API Management	39
Introduction	39
Structure	39
Objectives	39
Introduction to API management	40
Key concepts	41
<i>API gateway</i>	41
<i>API lifecycle management</i>	43
<i>Developer portal</i>	44
<i>Security and access control</i>	45
<i>Traffic management</i>	46
<i>Analytics and monitoring</i>	47
<i>Transformation and mediation</i>	47
<i>Integration with Azure services</i>	47
Applications.....	48
<i>API monetization</i>	48
<i>Legacy system modernization</i>	48
<i>Partner integration</i>	48
<i>Multi-cloud environments</i>	48
Getting started	49
Conclusion	53
4. Building My First API	55
Introduction	55
Structure	55
Objectives	56
Background	56
Prerequisites.....	56
<i>C#</i>	57
<i>Integrated development environment</i>	57
<i>Git and GitHub</i>	59
Understanding the workflow.....	60
Cloning a repository from GitHub	61
Running .NET Azure Function locally.....	63

Creating Function App in Azure	68
Deploying .NET Azure Function to Azure.....	74
Conclusion	84
5. APIM in Focus: Navigating the Basics	85
Introduction	85
Structure	85
Objectives	86
Overview of the resource	86
Activity log.....	88
Access control: IAM.....	91
Tags	93
Diagnose and solve problems	94
Events.....	94
Locks	95
Properties	97
Platform migration.....	97
Conclusion	99
6. APIs and Products	101
Introduction	101
Structure	101
Objectives	102
Importing an API from Azure Function	102
Running APIs from API management	106
Products in API management	109
Adding APIs to Products	109
Publishing the developer portal	111
Publishing products to the developer portal	113
Products walkthrough.....	114
Conclusion	117
7. Users, Groups and Subscriptions.....	119
Introduction	119
Structure	119

Objectives	120
Users.....	120
Groups	121
Subscriptions.....	121
<i>Creating subscriptions</i>	121
<i>Authentication and authorization</i>	121
<i>Access control and rate limiting</i>	122
<i>Monetization and billing</i>	122
<i>Developer experience</i>	122
<i>Key considerations</i>	122
<i>Best practices</i>	123
Run API from developer portal: As Guest	123
Run API from developer portal: As Developer	133
Subscription limits and approval	138
Conclusion	143
8. Versions and Revisions	145
Introduction	145
Structure	145
Objectives	145
Revisions	146
Revisions in action	146
Versions.....	154
<i>Types of versions</i>	154
<i>Best practices</i>	155
Versions in action	156
Conclusion	160
9. Policies.....	161
Introduction	161
Structure	161
Objectives	161
Introducing policies.....	162
Types of policies	162

Key features	163
Policy scopes	163
Policies in action	164
<i>Authentication policies</i>	164
<i>Basic authentication</i>	164
<i>Client certificate authentication</i>	165
<i>Authentication with managed identity</i>	167
<i>Access-restriction policies</i>	170
<i>Check HTTP header</i>	170
<i>Ip-filter</i>	172
<i>Rate-limit-by-key</i>	173
<i>Set-usage-quota-by-key</i>	175
<i>Rate limits versus quotas</i>	176
<i>Limit call rate by subscription</i>	176
<i>Transformation policies</i>	179
<i>Set-headers</i>	180
<i>Mask URLs in content</i>	182
Conclusion	183
10. Monitoring	185
Introduction	185
Structure	185
Objectives	186
Analytics	186
Application Insights	189
Alerts and metrics	191
<i>Alerts</i>	191
<i>Creating an alert rule</i>	192
<i>Metrics</i>	194
Diagnostic settings	195
Logs	199
Advisor recommendations	200
Workbooks	201

Conclusion	202
11. Infrastructure and Deployment.....	203
Introduction	203
Structure	203
Objectives	204
Pricing tier.....	204
Locations	205
Scale out.....	206
Gateways	207
<i>Adding a gateway.....</i>	<i>208</i>
External cache	209
Custom domains	211
Network.....	212
Notifications.....	213
<i>Notification templates</i>	<i>215</i>
Management API	216
Repository	218
Conclusion	219
12. Security and Automation.....	221
Introduction	221
Structure	221
Objectives	222
Defender for Cloud.....	222
Managed identities	224
Certificates.....	225
Protocols + ciphers.....	226
Tasks.....	228
Export template	229
Conclusion	230
Index	231-236

CHAPTER 1

Introduction to Azure

Introduction

This introductory chapter explains the fundamental concepts of cloud with Microsoft Azure as an example. The idea is to introduce the cloud to a beginner and help the reader understand its need in today's **Information Technology (IT)** industry. The chapter touches upon the various offerings from Microsoft and provides an overview of the various services (resources, as we shall call it moving forward) so that one knows the capabilities of the Azure Cloud platform and the potential problems it is trying to solve for the start-ups as well as for the industry.

Structure

The chapter covers the following topics:

- Introduction to cloud
- History and evolution
- Cloud delivery models
- Cloud deployment models
- Microsoft Azure as a cloud provider
- Azure essentials

Objectives

After reading this chapter, the reader will be able to explain the concepts relevant to the cloud to anyone new to it. They will be able to wisely choose the right type of services from Azure as a cloud provider for the organization's needs and fulfil them. One would also be well-versed with the various cloud deployment models and cloud delivery models available for consumers in the market.

Introduction to cloud

If you watch movies on any of the **over-the-top (OTT)** platforms like Netflix or Amazon Prime Video, or if you use an application like Microsoft OneDrive or Google Photos to store files such as images, videos, and any other kinds of data, you are already a user of cloud. To simplify further, most of the on-demand services that you are using these days by making use of the internet are outcomes of cloud computing services offered by some of the top technology companies in the world, like Microsoft, Amazon, and Google. While Microsoft's Cloud platform is called **Azure**, Amazon and Google offer **Amazon Web Services (AWS)** and **Google Cloud Platform (GCP)**, respectively.

Imagine you have started your business, and you initially kicked off with a one-person company in a single room with just one computer doing it all for you. You do not need anything other than the internet and electricity to cater to your business needs, apart from some other business-specific essentials. Now, as you scale up, you will require more people, more storage, more computing, and hence more devices.

You could still manage the storage with extra **hard disk drives (HDDs)** and **solid-state drives (SSDs)** till you go beyond a certain point. But soon, what you might realize is that buying more stuff corresponds to more costs not only in terms of purchase but maintenance as well because you will now need a technical team to handle the disaster/fault recovery cases. Also, it is not always true that you utilize whatever you purchase to the fullest. This underutilization is another overhead. This is exactly the problem that a lot of tech companies often face as they install more and more servers on their premises. The following is a popular hologram of cloud computing (refer to *Figure 1.1*):



Figure 1.1: Cloud computing

To solve this, assume you have a service that offers you storage for all your data, compute for all that you are doing using your local devices/machines, and infrastructure to manage all the overheads such as maintenance, security, recovery, tech team, etc., for a nominal fee. Yes, this is exactly what the cloud service providers have to offer. An added advantage that would follow is that you pay only for what you use, that is, the pay-as-you-go pricing model gives you the flexibility to immediately scale up or down according to your needs, thereby avoiding any overheads of incurring additional costs for underutilized resources.

Importantly, what one would need to leverage these services is the **internet**. This is because all of the services that the cloud provider offers are hosted remotely, and one needs to connect to them using the internet and pay as per consumption.

As mentioned earlier, in addition to Microsoft, Amazon, and Google, today, companies like IBM, Alibaba, Oracle, Salesforce, and Apple have also been investing heavily in this area and thus providing cloud services.

A simple example of how it impacts end users like us would be the fact that the number of movies that we can watch on our phones/laptops no longer depends on the storage of our devices, unlike a decade ago. With OTT platforms hosting all of these remotely and streaming them for us on an on-demand basis, the cloud has been a game changer. We can access and restore all our data, be it movies, songs, pictures, videos, documents, etc., by using an active subscription from the respective cloud provider and internet, from anywhere across the world within minutes.

History and evolution

Cloud computing technology has advanced over time. The 1950s saw the introduction of large-scale mainframes with enormous processing capability, which is when the idea of cloud computing first emerged. The practice of time-sharing and resource pooling was developed to take advantage of mainframes' computational capability. Multiple users could access the same storage layer and CPU power from any terminal, thanks to open terminals, which were created with the sole intention of facilitating mainframe access. With the introduction of **virtual machine (VM)** technology in the 1970s, mainframes could support multiple virtual systems or virtual machines on a single physical node.

The VM technology improved the shared access mainframe application from the 1950s by enabling numerous distinct compute environments to coexist on the same physical hardware. Even though the resources were shared, each VM ran guest operating systems that operated as if they had their RAM, CPU, and storage.

Virtualization thus became a technology driver and a huge catalyst for some of the biggest evolutions in communications and computing. Even 20 years ago, physical hardware was quite expensive. With the internet becoming more accessible and the need to make hardware costs more viable, servers were virtualized into shared hosting environments, virtual private servers and virtual dedicated servers using the same types of functionalities provided by the virtual machine operating system.

So, for example, if a company needed n number of physical systems to run an application, they could take one physical node and split it into multiple virtual systems. This was enabled by hypervisors. A **hypervisor** is a small software layer that enables multiple operating systems to run alongside each other, sharing the same physical computing resources. It also separates the virtual machines, logically assigning each slice of underlying computing power, memory, and storage, preventing VMs from interfering with each other. If, for example, one operating system suffers a crash or security compromise, others can keep working. As technologies in hypervisors improved and they could share and deliver resources reliably, some companies decided to make cloud benefits accessible to users who did not have an abundance of physical servers to create their cloud computing infrastructure. Since the servers were already online, the process of spinning up a new instance was instantaneous. Users could now order cloud resources they needed from a larger pool of available resources, and they could pay for them on a pay-as-you-go basis.

This pay-as-you-go or utility computing model became the key driver for cloud computing taking off. This allowed companies as well as individual developers to pay for the computing resources as and when they used them, just like units of electricity. This allowed them to switch to a more cash-flow-friendly **Operational Expense (OpEx)** model from a **Capital Expense (CapEx)** model. This model appealed to all sizes of companies. Those who have little or no hardware and even those who have lots of hardware, because now, instead of making a huge capital expenditure on hardware, they could pay for compute resources as and when needed. It also allowed them to scale workloads during usage peaks and scale down when usage subsided. This gave rise to modern-day cloud computing. The impact of the evolution of the cloud has been immense.

Cloud delivery models

Customers these days need flexible pricing and support models to cater to their business needs. This is offered to them by various delivery models, as explained in this section. These distinct models give them control over the various resources that they seek from the cloud provider and draw clear distinctions between the accountability of the same.

To understand this better, let us go back to our example of setting up a business and analyze the requirements step-by-step to develop a consolidated stack of resources. Let us assume that we need to host an application that we have built for our consumers. To build an entire infrastructure for the same, we would need the resources as explained:

1. The first thing we need to do is purchase some servers. This comprises the hardware that will run the application and may include CPUs, memory, motherboards, etc. Therefore, we get the first deck in our stack (refer to *Table 1.1*):

Servers	Memory, CPU, Motherboard
---------	--------------------------

Table 1.1: Server components

2. Next, we need our infrastructure to connect to the external world and with the resources present across the globe. This calls for networking infrastructure that includes routers, switches, the Internet, etc. Our stack, therefore, shall look like *Table 1.2*:

Networking	Routers, Switches, Internet
Servers	Memory, CPU, Motherboard

Table 1.2: Added networking components

3. Since the servers essentially need storage, these can be SSDs, HDDs, disks etc. The Server, Storage, and Networking together fulfil the hardware needs of the business. With these three fundamentals of hardware in place, our stack now looks like *Table 1.3* below:

Storage	SSD, HDD, Disks
Networking	Routers, Switches, Internet
Servers	Memory, CPU, Motherboard

Table 1.3: Added storage components

4. Now that we have the hardware ready and to best utilize these physical components, we need some virtualization software so that we can run multiple virtual machines on a single physical node (hardware machine). This will help us host multiple applications and maintain proper separation between each of them, even if the underlying physical infrastructure remains the same. Each of the virtual machines that we create to host multiple applications would need an operating system like Windows or Linux. Hence, we get our next deck in the stack as shown below (refer to *Table 1.4*):

Operating systems	Windows, Linux
Virtualization	Virtual Machines
Storage	SSD, HDD, Disks
Networking	Routers, Switches, Internet
Servers	Memory, CPU, Motherboard

Table 1.4: Added OS and virtualization components

5. A typical application would usually need additional software called middleware to support the application and requires a runtime as well for hosting the application in real-time. Runtime can be an IIS web container for a web application, and it can be Docker for a container-based application. The runtime is usually managed by the customer. With these two added layers, our final stack looks as shown with the application and data represented on the top (refer to *Table 1.5*):