

MIND THE GAP

Most Common Cloud Mistakes

MAGDALENA WOJNAROWSKA-PIETRZAK

MIND THE GAP

Most Common Cloud Mistakes

Magdalena Wojnarowska-Pietrzak

Copyright

Title: MIND THE GAP: Most Common Cloud Mistakes

Copyright © 2024 by Magdalena Wojnarowska-Pietrzak

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage retrieval system except for the inclusion of brief quotations in a review, without permission in writing from the author or the publisher.

Author: Magdalena Wojnarowska-Pietrzak

Disclaimer: Opinions shared in this book are my own, and do not reflect or represent any official view of any employer.

Table of contents

COPYRIGHT	3
TABLE OF CONTENTS	4
INTRODUCTION	9
About the Author	11
Purpose of this book	13
LAYOUT OF THIS BOOK	15
PART I – INTRODUCTION TO THE CLOUD	19
CHAPTER 1 – UNVEILING CLOUD COMPUTING	20
CLOUD COMPUTING ESSENTIAL CHARACTERISTICS	22
CLOUD COMPUTING SERVICE MODELS	23
CLOUD COMPUTING DEPLOYMENT MODELS	24
LOGICAL MODEL FOR IDENTIFYING DIFFERENT CLOUD FUNCTIONALITY LA	YERS
	26
WHAT DO YOU NEED TO UNDERSTAND ABOUT A PUBLIC CLOUD?	27
DEBUNK THE CLOUD COMPUTING PORTFOLIO	37
CHAPTER 2 – UNLEASHING THE BENEFITS OF THE CLOUD	40
LESS PAINFUL PLANNING	41
Scale as you need	42
ENTERPRISE-SCALE SECURITY SOLUTIONS	44
DOCUMENTATION	46
LOCATION REQUIREMENTS	48

Table of contents
PAYING FOR WHAT YOU USE
SERVICES AT NO ADDITIONAL COSTS

PART II – AREAS OF COMMON MISTAKES IN THE CLOUD 65

CHAPTER 3 – NAVIGATING CLOUD CAUTIOUSLY

MONITOR YOUR EXPENSES	68
USE WHAT YOU NEED	70
CHALLENGES WITH PROTECTING DATA (PRIVACY AND PERSONAL	
INFORMATION)	72
ALWAYS FAMILIARIZE YOURSELVES WITH SLA OF EACH SERVICE	81
CHECK THE DOCUMENTATION PROVIDED BY CSP	82
LARGE VARIETY OF SERVICES AVAILABLE IN A PUBLIC CLOUD	84
MISTAKES FROM MISUNDERSTANDING THE TECHNOLOGY	88
NOT BEING AWARE ABOUT CLOUD SERVICE PROVIDER SPECIFICS	90
CHANGES IN THE TECHNOLOGY, CLOUD PORTFOLIO OR CLOUD SERVICE	92
INCORPORATING NEW CLOUD SERVICES INTO YOUR LANDSCAPE	94
NOT ADAPTING BEST PRACTICES IN SCOPE OF SECURITY, TECHNOLOGY OR	
REGULATIONS	96
USING CLOUD SERVICES AGAINST ITS INTENDED PURPOSE	98
INTEGRATION WITH THIRD PARTY	99
HANDFUL TIPS FOR CLOUD PLATFORM	101

CHAPTER 4 – DECODING NAMING CONVENTION MISTAKES 106

AMBIGUOUS NAMING CONVENTION	108
NAMING CONVENTION GO DISORDERLY IN ORGANIZATIONS	111
NAMING CONVENTION DOES NOT INCLUDE CLOUD PROVIDER'S LIN	IITATIONS
	116
UNIQUENESS OF RESOURCES NAMES	119
HANDFUL TIPS FOR NAMING CONVENTION	120

50 51

67

CHAPTER 5 – UNRAVELING TAGGING MISTAKES	124
LACK OF TAGGING GUIDELINES	126
OVERCOMPLICATING TAGS	128
Overloading tags	129
TAGS ARE NOT FOLLOWED	131
PROBLEMS WITH NOT WELL-CONSIDERED TAGGING GUIDELINES	133
OUTDATED INFORMATION IN TAGS	135
HANDFUL TIPS FOR TAGGING	137
CHAPTER 6 – COUNTING THE COSTS: MANAGING YOUR	
FINANCES	140
Pay as you go misconception	142
NEGLECTING COST-REDUCTION MECHANISMS IN THE CLOUD	143
LACK OF CLOUD COSTS ANALYSES	146
ORPHANED RESOURCES IN A CLOUD ENVIRONMENT	155
"Over-tiering" your services	158
NOT LEVERAGING "BRING YOUR OWN LICENSE" FOR COST REDUCTION	160
HANDFUL TIPS FOR COSTS OPTIMIZATION	161
CHAPTER 7 – SAFEGUARDING YOUR CLOUD: SECURITY	
MISTAKES	165
NON-STANDARDIZED SECURITY CONTROLS AND BASELINES	168
Overdoing cloud security	171
DO IT THE SAME WAY YOU DID FOR ON-PREMISES	174
LACK OF FAMILIARITY WITH THE CLOUD NATIVE SECURITY TOOLS AND	
SOLUTIONS	176
NOT DISABLING UNNECESSARY SERVICES OR FEATURES	179
Forgetting about secure data disposal	180
IMPLEMENTING SELECTIVE PROTECTION MECHANISMS	181
DISREGARDING A MULTITENANCY RISK	187

LAX SECURITY CONTROLS	189
NOT PREPARING FOR POTENTIAL SECURITY BREACHES	204
HANDFUL TIPS FOR CLOUD SECURITY	207

CHAPTER 8 – NETWORK MISTAKES: NAVIGAT	ING THE CLOUD'S
CONNECTIVITY	210

FORGETTING ABOUT CLOUD PLATFORM'S NETWORK SPECIFICS	212
OVERLOOKING NETWORK TRANSFER COSTS	214
TREATING NETWORK LIKE IT IS ON-PREMISE	216
NO GUIDANCE OR POLICIES FOR NETWORK	218
NOT FAMILIAR WITH NETWORK TOOLS AND FEATURES	220
HANDFUL TIPS FOR CLOUD NETWORK	222

CHAPTER 9 – CLOUD MONITORING GONE WRONG: AVOIDING MISSTEPS 224

REUSING ON-PREMISES MONITORING SOLUTION	226
MISSING SECURITY MONITORING AND ALERTING	228
INSUFFICIENT METRIC COLLECTION	230
Ignoring Resource Dependencies	232
STATIC MONITORING CONFIGURATION	233
MISTAKES CONCERNING THE CLOUD NATIVE MONITORING TOOLS	236
PITFALLS WITH ALERTING	238
HANDFUL TIPS ON MONITORING AND ALERTING	245

CHAPTER 10 – UNMASKING THE PERILS OF CLOUD	IDENTITY
MANAGEMENT	247

OVERCOMPLICATING ROLE-BASED ACCESS CONTROLS	249
UNCLEARLY DEFINED ROLES AND RESPONSIBILITIES IN THE CLOUD	251
MISTAKES WITH MANAGING PRIVILEGED IDENTITIES	253

NOT DOING ANY REPORTING OR ACCESS REVIEWS ON PERMISSIONS	259
NEGLECTING ONBOARDING AND OFFBOARDING OF INTERNAL AND EXTERN	IAL
USERS ON YOUR PLATFORM	262
PREVENTING ACCESS TO THE MANAGEMENT PLANE AS A RESULT OF	
CONDITIONAL ACCESS MISCONFIGURATION	266
HANDFUL TIPS FOR THE CLOUD IDENTITY AND ACCESS MANAGEMENT	269

<u>CHAPTER 11 – BEYOND THE MISTAKES: EXPLORING WIDER</u> HORIZONS 271

TECHNOLOGY MISTAKES	272		
NOT CONSIDERING CONFIGURATION DRIFT INADEQUATE DATA BACKUP AND RECOVERY STRATEGIES AUTOMATION MISTAKES ORGANIZATION-LEVEL PITFALLS LACK OF CONSIDERATION FOR EXIT STRATEGY HANDFUL TIPS FOR CHAPTER 11 <u>CLOSING CHAPTER</u>	279 282 286 293 305 309 <u>311</u>		
		ACKNOWLEDGEMENTS	315
		ABBREVIATIONS	316
		INDEX	320

Introduction

"THE ONLY REAL MISTAKE IS THE ONE FROM WHICH WE LEARN NOTHING"

~Henry Ford

Cloud computing has taken the world by storm, becoming increasingly popular with each passing day. And why not? When you encounter a problem, the cloud seems like the perfect solution. But here is the thing: what you see is just the tip of the iceberg. Cloud Service Providers (CSPs) are continuously enhancing their offerings, making it easier than ever for you to embrace the cloud. It is no wonder that cloud computing has emerged as the answer to countless challenges faced by organizations worldwide.

Imagine running a large-scale organization with multiple data centers scattered across the globe. Maintaining such a dispersed infrastructure can be a daunting task. But fear not, for there are companies out there specializing in providing infrastructure designed to meet the highest standards. However, even with their assistance, you still need to procure your own equipment, hire skilled personnel capable of installing and managing it, and establish robust lifecycle management processes. And let's not forget about the crucial aspect of ensuring adequate support. To add to the complexity, the recent global pandemic has exposed the vulnerabilities in supply chains, leading to delays in the delivery of hardware and technology components. Such delays can hinder or severely restrict your organization's expansion plans or even its ability to operate smoothly during demand fluctuations. The above scenario is just a glimpse into the value that CSPs bring to the table. They take care of all the physical aspects, from data center facilities and security to hardware procurement and maintenance. What's more, they operate on a grand scale, with multiple data centers spread across the globe, serving numerous customers. This gives them a significant advantage in negotiating better contracts for the delivery of essential hardware to their data centers.

There are distinct differences between on-premise data centers and the cloud services offered by CSPs. The key is to understand the unique characteristics of the cloud and align it with your organization's specific needs. As more and more companies embrace cloud technology, it becomes crucial to grasp the common pitfalls and avoid them at all costs.

Within the pages of this book, I provide an overview of cloud computing, diving into its essential characteristics, benefits, risks, and key aspects to understand. In subsequent sections, I focus on the most prevalent mistakes that organizations encounter on their cloud journey.

You can expect me to shed light on the most common cloud mistakes and provide actionable recommendations on how to steer clear of them. Whether you are a newcomer to cloud computing or have been leveraging its power for years, this book will serve as a resource, guiding through the challenges of cloud adoption and ensuring the success of your cloud projects.

ABOUT THE AUTHOR

I have been working in IT industry in various roles even during my first studies. However, I started in the finance sector. How I shifted my career into IT is interesting. Frankly speaking, during preuniversity education, I was effectively discouraged from pursuing IT career path. As I have always liked math (have, what you can call, "analytical" mind) and I have been intrigued with economics and how this worked. Hence, I have decided to go for finance industry and pursue this as my career path. During my studies, I have chosen "Financial Engineering" Faculty, and this is when my passion for IT emerged.

I have noticed how great part of work you do in finances relies on and can be simplified (or automated) by simply writing a few lines of code. So, during my studies, I have decided to start working in IT sector and began a second faculty in IT on Technical University. Since then, I have been working in various roles, but the one I found myself the best is designing and bringing IT architecture to life. I have been working with cloud since 2014, but started exploring cloud even a bit earlier. Architecting cloud solutions brings me a lot of joy, and I like to transform various needs and requirements into working technical solutions and helping customers leveraging technology in their businesses.

My role as a cloud architect is diverse, requiring a mix of technical expertise and soft skills. This involves understanding customer needs, translating them into specifications, and building solutions that align with expectations. Communication, analytical, and technical skills are crucial in this role, and continuous learning is essential due to the rapid pace of development. I also contribute to building guidelines and architecture patterns for reuse, enabling efficient implementations and standardizing the landscape for effective governance. As an Architect, I embrace the multifaceted nature of the role and enjoy the challenges and opportunities it brings.

Besides doing solutions designs (popular diagrams), I focus on preparing comprehensive documentation (high level and low level designs). Which at the beginning of the project seems unnecessary, because you remember what is in there, but after short time, memory becomes a bit blurry, especially with number of projects and details you have to remember. Also, documentation is a good reference for technical teams on what and how should be implemented – kind of manual.

As an architect, I am also involved in or responsible for building guidelines and providing architecture patterns which can be reused. A set of generic guidance which can give other teams a head start, how to do implementations or engineering a solutions, but also standardize the landscape for more optimal governance and management.

I am fortunate to be involved in a diverse range of cloud projects and initiatives. Cloud architecture and engineering include building new platforms, maintaining and enhancing existing ones, providing consultations on cloud transformation, and engaging with businesses to identify their specific cloud needs. What makes this field truly exciting is the constant variety and uniqueness of each project. While there may be similarities in terms of components or services utilized, every end product is tailored to meet the specific expectations of the customer. One of the fascinating aspects of working in cloud engineering is that you never get bored. Each project presents its own set of challenges and opportunities.

PURPOSE OF THIS BOOK

This book is a must-read for anyone interested in the public cloud computing. It is suitable for Business and Digital Leaders, IT Administrators, IT Architects, IT Engineers, and Managers. This book is especially helpful for anyone involved in cloud migrations, adoptions and operations. Whether you are already immersed in cloud platforms or exploring their possibilities, I have tailored this book to meet your needs. If you are just starting your cloud journey and haven't yet delved into the concept, I've prepared a list of recommended resources at the beginning of Part II. These materials specifically cater to beginners who may feel uncertain about their cloud knowledge.

By highlighting common mistakes, my aim is to raise awareness and shed light on crucial aspects that often go unnoticed. The truth is, these mistakes can happen to anyone. Cloud ecosystems are intricate and multifaceted, and in the following pages, you will discover the multitude of factors that need to be considered in order to structure your cloud presence effectively and safeguard your data and systems.

This book has been a labor of passion, slowly taking shape in my mind throughout my years of working in the IT industry, with a specific focus on public cloud environments. Over the course of my career, I have held various roles, from engineering to architecture, encompassing administrative work, automation scripting, system management, architectural design, and the development of processes and strategies. I have had the privilege of working as an Infrastructure, Enterprise, Solution, and Security Architect for a diverse range of multinational organizations across different industries.

Throughout my journey, I have had the honor of collaborating with brilliant experts who shared their invaluable insights and guidance with me. This book is a culmination of all those years spent working in and observing the evolution of public cloud environments. Its purpose is to help you learn from the mistakes of others, strengthen your presence in the public cloud, and build robust, secure, and resilient environments.

Most of the pitfalls discussed in this book are based on my own observations, personal experiences, and the numerous consultations and dialogues I have had over the years. While cloud security is undeniably a crucial aspect, this book is not solely focused on that topic. Nor is it a step-by-step manual for configuring your public cloud environment. There are already plenty of resources available from vendors, cloud providers, and consultancy firms that provide detailed instructions. My aim is to provide you with a deeper understanding of the cloud and offer insights on how to enhance your current situation.

Consider this book your guide on the mistakes commonly made in the cloud, as well as how to avoid them. Its purpose is to accelerate your cloud journey and provide support in overcoming potential challenges you may encounter along the way.

Here is what lies ahead:

1. **Demystifying the Cloud**: This book help in simplifying the intricacies of cloud computing by breaking down the various layers essential for cloud transformation.

- 2. Navigating Risk and Challenges: The public cloud has its share of risks and challenges. This book will help you identifying potential roadblocks and vulnerabilities that may arise in the public cloud.
- Unleash problem-solving skills: No cloud journey happens without its share of challenges. This book equips you to recognize cloud-related challenges, while providing advice to effectively tackle them.

LAYOUT OF THIS BOOK

When I started with the cloud, I knew I had a lot to learn. My first cloud projects and later big cloud transitions, were a continuous learning and uncovering uncharted waters. I made mistakes along the way, which definitely fortified and sharpened my cloud skills.

During my years working with the cloud I was looking for such a guide, describing common mistakes and how they can be prevented. It would definitely mitigate some mistakes I made and alleviate them as well as boost the pace of gathering the cloud knowledge. However, I did not find such a book back then. The idea of writing this particular piece came from consulting projects I was doing, where I was approached by teams to share more about perils in the public cloud and how they can be prevented. The teams were looking for kind of a cheat sheet to simplify their way into the cloud. And after doing another research, I decided to start collecting my notes and experience. And long story short, this is how this book has been initiated.

You can read this book in various ways. You can start from the beginning and read till the end. If you feel comfortable with the cloud, you can skip the Part I and jump directly to the Part II. Also,

feel free to start with exploring the chapters that you are most interested in, and in case there is any reference to prior or onward chapters of the book, you will be informed about that fact.

Each chapter is summarized with my recommendations, in form of handful tips. You can just jump to the tips and in case something catches your eye, you can then read the subchapter about it. You also do not need to read all at once, you can read what you need and later go back and explore other chapters. But remember, the book as a whole provides a comprehensive overview of cloud mistakes within different layers.

These book is a collection of common cloud mistakes. It is composed of two parts, and eleven chapters in total:

Part I – Introduction to the Cloud

Part I covers cloud's essential baselines and insights. This section offers the opportunity to refresh your memory or complement your existing cloud expertise. If you are already familiar with the cloud foundations can skip **Part I**, and go directly to **Part II**.

Chapter 1 – Unveiling Cloud Computing

This chapter provides a n overview of public cloud computing. Covering the fundamental concepts and characteristics that shape the cloud landscape.

Chapter 2 – Unleashing the Benefits of the Cloud

Explore the benefits the cloud has to offer. This chapter uncovers how you can leverage the cloud to your advantage, to make informed decisions and maximize the value of cloud adoption.

Part II – Areas of Common Cloud Mistakes

Part II Focuses on cloud mistakes happening in various cloud layers.

Chapter 3 – Navigating Cloud Cautiously

No infrastructure exists without some challenges, and the cloud is no exception. This section highlights potential stumbling blocks within the cloud platform.

Chapter 4 – Decoding Naming Convention Mistakes

Governance layer of the cloud demands attention to detail, especially when it comes to naming conventions. This chapter covers the pitfalls and how a well-structured naming convention can enhance your cloud ecosystem.

Chapter 5 – Unraveling Tagging Mistakes

Labels and tags play a pivotal role in cloud governance. In this chapter you learn how to avoid common tagging mistakes and streamline your cloud management practices.

Chapter 6 – Counting the Costs: Managing Your Finances

Every investment, including the cloud, comes with financial considerations. This chapter sheds light on cloud areas that, if wrongly approached, can have a significant impact on your financials.

Chapter 7 – Safeguarding Your Cloud: Security Mistakes

Security is paramount, both within and beyond the cloud. This chapter unveils the various security mistakes that can compromise the integrity of your data residing in the cloud.

Chapter 8 – Network Mistakes: Navigating the Cloud's Connectivity

The cloud's network layer presents its own unique challenges. This chapter addresses common problems that can arise and provide practical guidance to ensure a robust and reliable cloud networking infrastructure.

Chapter 9 – Cloud Monitoring Gone Wrong: Avoiding missteps

Efficient resource monitoring should not be omitted, whether in the cloud or on-premises. This chapter dives into the monitoring of cloud resources and highlights its common pitfalls.

Chapter 10 – Unmasking the Perils of Cloud Identity Management

While you may be familiar with roles and responsibilities, the cloud introduces additional factors to consider. This chapter delves into identity-related challenges specific to the public cloud.

Chapter 11 – Beyond the Mistakes: Exploring Wider Horizons

In this final chapter covers a wide range of topics, from technology mistakes to automation pitfalls, organizational challenges, and the often-overlooked aspect of devising a cloud exit strategy.

Enjoy the Journey!

Part I – Introduction to the Cloud

Chapter 1 – Unveiling Cloud Computing

"The advance of technology is based on making it fit in so that you don't really even notice it, so it's part of everyday life."

~ BILL GATES, MICROSOFT

This chapter covers an introduction to the cloud, to ensure that everyone is equipped with the fundamental knowledge of what the cloud truly is, including its various service and deployment models.

But let's not stop there—let's dive deeper and explore the essential aspects of the public cloud. What is cloud computing, you ask? While the notion of the cloud being "someone else's computer" may serve as a concise explanation, it does not cover the aspect fully.

When I first delved into cloud technology, the entire concept was relatively new. Amazon officially introduced AWS in 2006, although the concept itself dates back to 2002. Microsoft unveiled 'Project Red Dog' in 2008, which in 2014 was officially published and renamed to Azure. Initially, cloud services resembled more of a virtual machine rental service, featuring offerings like virtual machines (no surprise there), disks, and networks. I was fortunate to embark on this journey at its early stages, witnessing firsthand the cloud's evolution and transformations. This early involvement allowed me to gradually acclimate to the cloud.

The current array of cloud services is broader than what was offered in the past. While getting accustomed with the cloud over a few years made the learning process more manageable, it's entirely feasible to familiarize yourself with the cloud now. In those earlier days, my learning path was paved with trial and error, learning valuable lessons from each mistake. Nowadays, with learning portals, communities, forums, and resources like this book at your disposal, you have a significant advantage in expanding your cloud knowledge.

CLOUD COMPUTING ESSENTIAL

CHARACTERISTICS

Let's fortify cloud understanding with insights from the National Institute of Standards and Technology (NIST). In their publication, the NIST SP 800-145¹ standard, they have outlined five key traits that define the very essence of the cloud:

- 1. **Broad Network Access**: Users can access cloud services from any corner of the globe, using devices such as mobile phones, laptops, or workstations.
- Rapid Elasticity: Cloud service providers employ mechanisms to automatically scale their resources in response to demand. This means that as a cloud customer, you can swiftly provision or decommission cloud services as per your needs.
- 3. **Measured Service**: Cloud service providers diligently monitor and analyze resource usage, ensuring that you are only charged for the services you utilize.
- 4. **On-Demand Self-Service**: With the cloud's on-demand selfservice feature, you can instantly spin up new resources with the aid of automation provided by the Cloud Service Provider.
- 5. **Resource Pooling**: This concept, often referred to as multitenancy, enables multiple customers to simultaneously

¹ According to National Institute of Standards and Technology (NIST) "Cloud Computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." <u>https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf</u>

harness the power of the same underlying infrastructure, seamlessly and securely.

Now that you are familiar with the cloud traits, let's navigate other aspects of the cloud landscape.

CLOUD COMPUTING SERVICE MODELS

Within cloud computing, you encounter three primary service models, as described in NIST SP 800-145:

- 1. Infrastructure as a Service (IaaS): It is a virtualized landscape where fundamental infrastructure resources, such as storage, networks, and essential computing elements, are delivered to you by the Cloud Service Provider. Unlike on-premise virtual machines, the ease of IaaS lies in the fact that you relinquish the need of managing the underlying infrastructure and hypervisor, leaving it under the control of the CSP. With IaaS, you have the power to deploy applications directly onto an operating system that you manage and govern. Customize your virtual machine and fine-tune the operating system itself. Azure Virtual Machines (VMs) and AWS Elastic Compute Cloud (EC2s) serve as examples of this model.
- 2. Platform as a Service (PaaS): Here, the provider assumes the responsibility of handling the underlying infrastructure, including the operating system. You can choose the framework that perfectly aligns with your application code. While your configuration options are somewhat limited to adaptations within the application hosting layer, the convenience and efficiency of PaaS are undeniable. Azure App Service and AWS Beanstalk are examples of this model.

3. Software as a Service (SaaS): It is the most CSP-managed service model. Accessible through a web browser or a programming interface, the SaaS applications offer a seamless user experience. While you relinquish the management of underlying infrastructure, you still retain the power to customize the application within the boundaries set by the provider. Examples of SaaS offerings can be Office 365, Google Workplace, or Dropbox.

It is worth noting that the cloud landscape may present other variations, such as the DBaaS (Database as a Service). However, remember that these variations are derivatives from the three foundational types. E.g. Database as a Service is a derivative from laaS and PaaS. Some of laaS features are available to you, whereas some of the laaS responsibilities are taken care of by the CSP. Making it more PaaS-like experience.

CLOUD COMPUTING DEPLOYMENT

MODELS

In the NIST SP 800-145 standard, you can find four cloud Deployment Models:

 Private Cloud: this model refers to the cloud infrastructure which is designated for the sole purpose of a single organization. This infrastructure can reside on-premises or exist off-premises. From an operational perspective, it can be managed and administered by the organization, or these activities can be partially or even fully outsourced.

- 2. Public Cloud: An infrastructure built, managed, and operated by Cloud Service Providers (like AWS, Microsoft Azure, GCP). This platform offers services to multiple customers. The key distinction between the public cloud and its private counterpart lies in its multitenant nature. In the public cloud, countless customers from different organizations come together, sharing the resources of the underlying infrastructure. The logical isolation ensures each tenant's data remains secure and separated.
- Community Cloud: Within this model, the infrastructure is exclusively dedicated to a specific group of like-minded consumers. Whether it manifests on-premises or offpremises is a choice made by one or more organizations. Examples are platforms like Microsoft Xbox, Sony PlayStation, or Steam.
- 4. **Hybrid Cloud**: This model is a combination of at least two separate cloud infrastructures, which act as individual entities, but still are connected to enable communication between applications and transmissions of data. An example can be extending the on-premise data center to the cloud.

LOGICAL MODEL FOR IDENTIFYING

DIFFERENT CLOUD FUNCTIONALITY LAYERS

The layers of the cloud computing model introduced by the CSA STAR², published in their "Security Guidance For Critical Areas of Focus In Cloud Computing v4"³ document.

- Infostructure This layer encapsulates the vital resources related to the storage and management of resources. The example of the service can be file storage.
- Applistructure The next layer consists of the applications that utilize the underlying cloud infrastructure and its services. The example can be capabilities like notification services.
- 3. **Metastructure** This layer is unique for cloud environments, covering the protocols and mechanisms acting as an interface between the infrastructure layer and the other layers.
- Infrastructure In this layer, the elements of compute, storage, and network converge to create a foundation for the cloud.

² Cloud Security Alliance (CSA) Security, Trust, Assurance and Risk (STAR) – <u>https://cloudsecurityalliance.org/star/</u>

³ <u>https://cloudsecurityalliance.org/artifacts/security-guidance-v4/</u>

WHAT DO YOU NEED TO UNDERSTAND

ABOUT A PUBLIC CLOUD?

In the previous subchapters, I laid the foundation by exploring the definition of cloud, its essential characteristics, and the unique model types that define cloud computing.

The most popular public cloud providers (Microsoft⁴, Amazon⁵, Google⁶) offer fully managed, globally-present, redundant, and secure infrastructure that can be deployed via self-service portals. But how does this infrastructure differ from traditional on-premise services?

The key differences are related to cost charging models, administrative responsibilities, data privacy, security and availability of computational resources.

Mostly OPEX

In the cloud most of services you purchase are charged based on your usage. Therefore, you rely mostly on operational expenditures (OPEX). However, in the cloud, there are possibilities to explore capital expenditures (CAPEX) and leverage them to your advantage.

One such avenue is the ability to strategically utilize CAPEX by making advance purchases, such as acquiring licenses in bulk. Often, organizations can achieve group discounts that surpass

⁴ <u>https://azure.microsoft.com/en-us/resources/cloud-computing-</u> dictionary/what-is-azure

⁵ <u>https://aws.amazon.com/what-is-aws/</u>

⁶ <u>https://cloud.google.com/docs/overview</u>

direct purchases from Cloud Service Providers. This opens the door to the option known as "Bring Your Own License". Yet, it is worth to note that not all services may support external licenses, so a thorough check is essential.

Another viable route to transition from OPEX to CAPEX is by embracing the concept of "reservations". This entails reserving the required resources in advance for a defined period, such as popular options like 1-year or 3-year reservations. Do keep in mind that the terminology for this option may differ across cloud providers.

While there may be additional paths to explore within the CAPEX, the OPEX is dominant in the cloud. Understanding the interplay between OPEX and CAPEX is essential to navigate the financial landscape and optimize your cloud investments.

Responsibilities in the public cloud

As you delve into the cloud, you will quickly notice the contrasts between on-premises datacenters and the public cloud.

Unlike self-managed datacenter, the public cloud is a domain designed, built, and managed by Cloud Service Providers (CSPs). This means that the physical infrastructure layer, encompassing datacenter facilities, hardware components, and more, is expertly handled by the CSP. Accessing your purchased cloud resources is facilitated through a programming interface or management plane provided by the CSP. The cloud services you deploy come with a base configuration that seamlessly communicates with other services within the CSP's infrastructure, fostering nearinstant integration with a vast array of offerings. However, do not forget to familiarize yourself with the capabilities, features, and your own responsibilities as a customer when utilizing these predefined services.

Contrastingly, in your on-premises data center, you are responsible for every aspect of the infrastructure, from the facility itself to the physical and virtual infrastructure, as well as application and data layers. On the other hand, in the cloud, the cloud provider delivers you the infrastructure services, in the form of three distinct service models: IaaS, PaaS and SaaS.

In all service models, the CSP bears full responsibility for preparing the core infrastructure, including networking, storage, services, and virtualization, to support customer workloads. Simultaneously, as a cloud customer, you assume complete accountability for your stored and processed data, as well as all governance aspects tied to cloud computing.

In the IaaS model, you rely on the CSP's service portfolio and utilize virtualized networking to establish seamless communication between workloads. Here, shared responsibility becomes a vital consideration.

> Let's zoom in on the operating system aspect as an example. While the CSP offers ready-to-use images, the necessary customization remains your responsibility. Additionally, you have the option to upload your own images to the cloud, which can later be leveraged in the laaS environment. Security is another area of shared responsibility. The CSP takes charge of safeguarding their physical infrastructure, while

you must fulfill your contractual, regulatory, legal, or jurisdictional obligations by implementing necessary adjustments within the cloud environment. This may involve special hardening and configuration of the operating system in line with best practices and security guidelines.

Considerations around protecting private and sensitive data stored and processed in the cloud introduce further areas of shared responsibility.

Shifting the focus to the PaaS model, you no longer have access to the underlying operating system. Instead, you gain access to services that provide the frameworks for running your code. Here, your responsibility includes sharing the burden of security, runtime, data, governance, and communication patterns for your application services with the CSP.

The SaaS model places the least responsibility on your shoulders, primarily focusing on data and governance. While security and application aspects are shared between you and the CSP, the extent of your involvement depends on the cloud service model and the boundaries set by the vendor's application framework. Nevertheless, you remain ultimately liable for safeguarding the integrity, confidentiality, and availability of your data.

It is important to clarify that the distribution of responsibilities does not imply negligence on the part of the CSP or the cloud customer. Rather, it represents a division of duties (as presented in Figure 1). In an on-premises environment, you retain full responsibility for your infrastructure, while in the cloud, certain responsibilities are shifted to the CSP. However, complete transfer of all responsibilities is not possible, and areas like security remain shared. Additionally, as the data owner, you hold ultimate accountability for any unauthorized data disclosures.

Figure 1 Comparison of responsibilities in different cloud computing service models and on-premise⁷

Legal Liability for the data

In every service model, whether it is Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS),

⁷ Please note, that most common shared responsibility matrix covers three main cloud service models. However, in some cases you may encounter Function as a Service (FaaS). More details you can read in Google Article. https://cloud.google.com/learn/paas-vs-iaas-vs-saas

the responsibility for safeguarding our data falls squarely on our shoulders.

Consider this: the ultimate legal liability for any unauthorized disclosure or illicit access to your data cannot be transferred—it remains solely with you as the data owner. This realization underscores the utmost importance of implementing robust controls and measures to fortify your data security in the cloud. No matter the circumstances, you cannot transfer this responsibility elsewhere.

Security

Cloud Service Providers (CSPs) diligently address numerous security concerns. However, you should acknowledge that, you hold the responsibility of implementing essential security controls and measures to meet jurisdictional, legal, contractual, and regulatory obligations. While the CSP offers security controls for the physical infrastructure and provides customers with solutions and tools, it is up to you to determine which specific security measures are necessary to fulfill our regulatory requirements. The CSP contributes to the security objectives by offering security solutions and mechanisms, while also adhering to general security best practices to safeguard their infrastructure and customers against adversaries.

As you learned in previous subchapter (*responsibilities in the public cloud*, p.28), the distribution of responsibilities varies across different cloud service models. In Infrastructure as a Service (IaaS), you have more control and responsibility, followed by Platform as a Service (PaaS), and Software as a Service (SaaS), where the CSP assumes a greater share of security responsibilities.

This hierarchy is primarily influenced by the customizable layers available for cloud customers to configure according to their specific needs.

> When transitioning to the cloud, trust between the Cloud Service Provider and the customer is crucial. This trust can be formalized through contractual agreements, but it is important to note that not all responsibilities can be transferred to the CSP. As the old saying goes, "trust but verify". However, in the cloud computing, the verification process can be challenging. CSPs typically do not grant unrestricted access to their datacenter facilities, prioritizing security and the protection of intellectual property. Thus, physically inspecting the datacenters may not be feasible.

Fortunately, there are alternative methods to assess the reliability of a CSP's security measures. CSPs are subject to regular audits to ensure compliance with industry standards. By examining the certifications obtained by the CSP, such as SOC⁸,

⁸ SOC (System and Organizational Controls) reports are performed in accordance with Statement on Standards for Attestation Engagements (SSAE) 16

PCI DSS⁹, and various ISO¹⁰ standards, can assess their commitment to security. The cloud provider's website often provides information on their compliance with these standards. Additionally, the CSA STAR¹¹ registry offers a wealth of valuable data on cloud security for different CSPs. It presents a framework for evaluating CSPs, consisting of three levels of assurance: self-assessment, attestation conducted by an independent party, and results from continuous monitoring performed by a certified third party.

Costs

When it comes to costs, there are distinct differences between onpremise infrastructure and the cloud. As you explored in the *Cloud computing essential characteristics* subchapter (p.22), the cloud offers a feature called measured service. This means you are billed based on your actual usage, a concept that sets it apart from onpremise setups. In an on-premise environment, where you own the servers, data center, and other components, you incur costs regardless of whether you utilize them. This means that even if a virtual machine (VM) remains idle or turned off, your department or cost center is still charged.

⁹ Payment Card Industry Data Security Standard (PCI DSS) – is an information security standard used to handle credit cards from major card brands. The standard is administered by the Payment Card Industry Security Standards Council, and its use is mandated by the card brands. It was created to better control cardholder data and reduce credit card fraud. (source of explanation: <u>https://en.wikipedia.org/wiki/Payment Card Industry Data Security Standard</u> rd)

¹⁰ https://www.iso.org/home.html

¹¹ <u>https://cloudsecurityalliance.org/star</u>

In the cloud, however, the scenario changes dramatically. Let's consider the same VM example. If you deploy a VM but later switch it off, freeing up the compute resources such as RAM and CPU, you stop incurring charges for those resources. In essence, you only pay for the storage required to store the data and configuration of your resources. On-premise environments require payment regardless of whether machines are active or not, leading to a tendency to leave them running constantly. Unfortunately, this oversight can result in substantial costs, particularly when a large fleet of machines remains unattended for extended periods.

While cost optimization is critical, it is essential to note that unattended workloads in both the cloud and on-premise environments can pose security threats. Addressing security concerns, hardening systems, and applying patches remain necessary, regardless of the chosen infrastructure model.

In the cloud, if you fail to decommission resources you no longer require, you will continue to be charged for them. Embracing cloud services requires a shift in mindset. Remember to release resources that are no longer in use. Fortunately, the cloud offers various automation mechanisms that allow you to apply controls and streamline this process.

Automation by design

The way cloud platform is designed is determined by its essential characteristics, shaping the experience and capabilities it offers to customers. To provide resource pooling options, metered services, and more, Cloud Service Providers (CSPs) must craft and construct their cloud service offerings. This involves implementing the necessary technical and security controls that allow customers to deploy resources independently, without constant intervention from the CSP.

The CSP has the physical data center facilities equipped with necessary hardware and infrastructure. Building upon this foundation, the CSP constructs a platform layer, which is then delivered to customers in the form of a user-friendly portal or programmable interface. Within this management plane, customers can deploy or decommission services at their convenience. The underlying automation infrastructure, running beneath the customer-facing portal, streamlines the entire ordering process.

> Customers can simply choose from the extensive array of services in the CSP's portfolio, configure their preferences in the "order form" or select appropriate features, sizes, tiers, and more, and almost instantaneously obtain their pre-defined products (although the complexity of certain services may extend the deployment time slightly).

Cloud automation, however, encompasses more than just quick service deployment for numerous customers. It aligns with the very essence of cloud characteristics. As highlighted in *the logical model for identifying different cloud functionalities* subchapter (p.26), the metastructure emerges as a unique layer specific to the cloud environment. This layer serves as the interface between the infrastructure layer and other interconnected layers, granting users the ability to remotely manage their cloud resources.

Nevertheless, bear in mind that while this automation empowers you as a cloud customer, it does not translate into autonomous access to every aspect of the cloud platform's infrastructure. Boundaries are established by the cloud provider. The comprehensive scope remains firmly under the control of the CSP, while you are granted access to manage the logically isolated portion of the cloud environment assigned to you.

DEBUNK THE CLOUD COMPUTING

PORTFOLIO

Cloud portfolio directly influences various aspects of your cloud journey, including design, architecture, automation, deployment, and maintenance. I present here different categories of cloud services you encounter in the cloud:

1. **Identity**: The foundation of security in the public cloud environment (identity acts as a first barrier of security to access the public cloud environment), encompassing identity and access management systems.

- 2. **Migration**: Services facilitating the smooth transfer of resources, services, and data to and from the cloud, including the importance of graceful exits from a cloud platform.
- Integration: Technical aspects enabling seamless communication between different services and environments.
- 4. **Storage**: Services essential for data storage, necessitating considerations in security, risk, technology governance, and organizational aspects.
- Infrastructure Services / Compute: Services related to virtual machines and complementary offerings such as remote access and operating systems.
- Mobile and Web Development Services: Frameworks and API interfaces supporting developers in deploying solutions in the cloud.
- 7. Artificial Intelligence (AI) and Machine Learning (ML): Services designed for developers and data scientists to improve e.g. business processes or data management.
- 8. **Database Services**: Essential for data processing, for storing structured and unstructured data.
- Analytics Services: Vital for organizations processing large volumes of data, safeguarding one of their most valuable assets.
- 10. Security and Protection Services: A key focus for companies entering the cloud, encompassing prevention of attacks and mitigation of risks.

- 11. Governance and Platform Management Services: Toolsets provided by CSPs (and sometimes third-party solutions) to manage costs, identify resources, maintain inventory, and generate reports.
- 12. Hybrid and Multi-Cloud Solutions: Options for organizations with complex requirements, allowing interoperability and utilization of the best services from different clouds.
- 13. **Development Tools**: Native automation tools offered as part of the platform or third-party solutions supporting activities like building pipelines and deployments.
- Internet of Things (IoT) Services: Enabling management of IoT solutions, secure interface publication, analysis, and reporting.
- 15. **Containers**: Managed services provided by CSPs, typically utilizing Kubernetes, to simplify the deployment and management of containerized applications.
- 16. **Network Services**: Essential components encompassing network infrastructure, IP management, network packet inspection, monitoring, firewalls, network rules, and load balancers.

Please note that not all CSPs offer every category of service listed above. This compilation serves as a reference to illustrate the breadth of cloud services available to you. Certain services are inherent elements of every cloud ecosystem, such as identity, governance services, and network infrastructure. Additionally, specific services may be included in the cloud portfolio to meet market demands, such as containerization

Chapter 3 – Navigating Cloud Cautiously

"I have not failed. I've just found $10,\!000$ ways that won't work."

~Thomas Edison

The *Chapter 2* (p.40) presented the numerous advantages of cloud computing platforms. However, as with any innovation, it is crucial to recognize that there are potential downsides.

In a program I participated in, I observed a team selecting cloud services and features for their application. The team was in the process of exploring and testing various cloud services, which was a commendable approach. However, their method of choosing these services and their tiers relied more on intuition and hypothetical future scenarios rather than real-life use cases.

Consequently, they ended up with capabilities that were unlikely to be utilized, and they lacked proper knowledge on configuring these features. This led not only to increased service costs but also raised the likelihood of misconfigurations, potentially opening avenues for misuse.

This chapter will draw your attention to the areas where caution is necessary when working with the cloud.

MONITOR YOUR EXPENSES

In the previous chapter (*Chapter 2*, p.40), you explored how each Cloud Service Provider (CSP) equips you with tools to monitor, manage, and optimize your cloud expenses. But here is the catch: while these tools provide cost-related capabilities³¹, it is ultimately your responsibility to leverage them effectively. Understanding what you are paying for and aligning with your

³¹ AWS: <u>https://docs.aws.amazon.com/cost-</u>

management/latest/userguide/billing-security-logging.html Azure: <u>https://learn.microsoft.com/en-us/azure/azure-monitor/cost-usage#view-azure-monitor-usage-and-charges</u> GCP: https://cloud.google.com/cost-management?hl=en

core requirements is key. The cloud cost tools offer insights into your environment and can provide cost forecasting based on your current usage, but they do not do all the work for you. The same applies to advisory tools. They suggest steps to reduce expenses, but these recommendations are generic and programmed by the CSP. They may not always apply to your specific situation.

Let's consider an example. You have a Virtual Machine (VM) deployed for testing purposes. You cannot use a spot³² machine to decrease costs because the system needs to be consistently up and running to collect sufficient data for analysis. The application running on the VM has specific RAM and CPU requirements, forcing you to use a specific VM type from the cloud provider's portfolio. While the machine may be expensive and underutilized most of the time, you need to meet the test criteria set by the application's vendor. In this case, the advisory tool might recommend downsizing the VM to avoid paying for underutilized resources. However, this is not feasible given your testing requirements.

Another suggestion might be to purchase a reserved instance to significantly reduce costs. But again, in your scenario, committing to a long-lasting (1-year or 3-years) reservation does not make sense since the machine is only needed for a specific period. Hypothetically, spot machines could be proposed as an alternative, but as mentioned earlier, they are not acceptable in your case.

³² Spot Virtual Machines is a type of VMs that allows you to buy unused compute capacity at significant cost savings (run your VM cheaper). However, these resources are not permanent, so this solution is recommended for solutions running interruptible workloads.

These examples serve to illustrate potential suggestions, but they may not necessarily apply to your deployments or match your specific scenario. The point is, blindly following recommendations without evaluating their potential impact can lead to adverse effects. It is always wise to follow the simple rule of "trust but verify". Take the time to assess if the recommendations align with your unique circumstances to ensure you make informed decisions.

Use what you need

Cloud providers offer a wide range of cloud services tailored to different customer needs. As you have learned, in the cloud you pay for the resources you use. However, it is key to understand that paying for what you use, does not always align with what you are actually utilizing.

> Let's illustrate this with a simple example. Suppose you deploy a virtual machine for testing purposes. Once the test is complete, you should decommission the instance to release the compute resources. However, if you forget to do so, the machine continues to run, and technically you are still being charged for the resources, even though you're no longer utilizing them.

This example demonstrates the concept of paying for what you use. Now, let's delve into the aspect of using what you need. For

the purpose of this explanation, let's focus on the conference app from the *Chapter 2* (p.40), with some modifications.

Imagine you are organizing a regional conference where attendees will access the application from a single location. Since you do not require a highly distributed infrastructure to serve attendees from multiple regions, you can choose a single location and designate another region for backup and recovery purposes.

For this example we will use Azure services. For the conference application, you need to balance availability and security against common threats like SQL injection³³ or cross-site scripting (XSS³⁴). Two possible services are Azure Front Door and Azure Application Gateway. Both services offer similar functionalities, but the key difference is that Azure Application Gateway operates within regional boundaries, while Azure Front Door is a global service. While there are other differentiating features (based on example's requirements), I am focusing on the region. Considering this, Azure Application Gateway appears to be the better option.

³³ A SQL injection attack consists of insertion or "injection" of a SQL query via the input data from the client to the application. A successful SQL injection exploit can read sensitive data from the database, modify database data (Insert/Update/Delete), execute administration operations on the database (such as shutdown the DBMS), recover the content of a given file present on the DBMS file system and in some cases issue commands to the operating system. SQL injection attacks are a type of injection attack, in which SQL commands are injected into data-plane input in order to affect the execution of predefined SQL commands. (Source: <u>https://owasp.org/wwwcommunity/attacks/SQL_Injection</u>)

³⁴ Cross-Site Scripting (XSS) attacks are a type of injection, in which malicious scripts are injected into otherwise benign and trusted websites. XSS attacks occur when an attacker uses a web application to send malicious code, generally in the form of a browser side script, to a different end user. Flaws that allow these attacks to succeed are quite widespread and occur anywhere a web application uses input from a user within the output it generates without validating or encoding it. (Source: <u>https://owasp.org/www-community/attacks/xss/</u>)

Note, that in real-life scenarios, there are additional factors to consider beyond the one highlighted here. Always ensure that you incorporate all relevant factors based on your specific situation.

The principle of using what you need applies not only to the capabilities within each service tier but also to services that you may no longer require. This is similar to the forgotten virtual machine scenario mentioned earlier. When considering service tiers, your decision may be influenced by the features offered within each tier. Evaluate your needs carefully to avoid paying for resources that you will not utilize.

CHALLENGES WITH PROTECTING DATA

(PRIVACY AND PERSONAL INFORMATION)

Safeguarding your data is the foundation of any robust infrastructure. That is why the Zero Trust and Defense in Depth security models were introduced, and I will dive into these frameworks in later chapters (see *Implementing selective protection mechanisms*, p.181).

Defense in Depth follows a layered approach, where each layer acts as a barrier, preventing attackers from gaining access to and taking control of your valuable data. It is a comprehensive strategy that emphasizes multiple lines of defense. **Zero Trust** is a security framework and concept that challenges the traditional perimeter-based approach to network security. In a Zero Trust model, the default stance is to distrust all entities, both internal and external, and verify every user and device trying to access resources within the network, regardless of their location or network boundaries.

When it comes to protecting your data in the cloud, there are several factors that present challenges. However, that does not mean you are left without mechanisms to mitigate risks or adversarial situations. In fact, cloud environments offer a wide array of security measures to safeguard your data.

Encryption is one such powerful tool that shields your data in transit, at rest, and even during usage, ensuring protection against hostile takeover attempts. Firewalls play a crucial role in inspecting your network traffic and implementing rules that prevent unauthorized lateral movement within your infrastructure. Additionally, cloud platforms often provide built-in DDoS protection services, fortifying your defenses against malicious attacks.

The list goes on, showcasing the extensive range of protective measures that can be employed within the cloud environment.

In this book, I will shed light on the potential hurdles you might encounter while safeguarding your data in the cloud. These challenges can make data protection more demanding, necessitating adjustments to your security strategies to account for the unique characteristics of cloud computing.

Multitenancy

This is a characteristic commonly found in public cloud environments, where the physical infrastructure is shared among multiple customers. While it is possible to purchase dedicated hardware instances for your workloads, this option tends to be more expensive and may not be universally offered by all Cloud Service Providers (CSPs).

CSPs take measures to logically isolate customers' environments, but note that there can never be a complete guarantee against potential guest escape scenarios. To mitigate such risks, cloud providers strongly recommend the use of encryption.

Encryption plays a crucial role in safeguarding data both during transit and at rest. Many services already have encryption implemented by default, and some even prevent users from disabling it. Encryption acts as a protective mechanism, preventing unauthorized access to data during communication interception or in the event of physical theft of disks from the provider's datacenter facilities.

However, understand that encryption alone does not resolve all security issues. Its effectiveness depends on the strength of the encryption algorithm employed. Therefore, it is worth to assess the encryption algorithms utilized by your provider and compare them against your internal security guidelines.

In certain situations, you may have specific requirements that prohibit sharing infrastructure with other customers or certain competitors. These requirements could come from contractual agreements or government regulations. However, in the public cloud environment, identifying which exact customers share the underlying infrastructure with you can be challenging, if not impossible. Therefore, it is necessary to thoroughly understand the specific demands you need to fulfill and develop a comprehensive plan accordingly.

Not your infrastructure

When discussing the concept of multitenancy, I would like to emphasize the aspect of infrastructure ownership. The challenges mentioned earlier arise due to the nature of the public cloud model. In *Chapter 1* (p.20), I highlighted the differences between private and public clouds, with one of the challenges being multitenancy, as discussed in the *Multitenancy* subchapter (p.74). Another aspect relates to the disparities between on-premises and public cloud environments, which were outlined in the *Responsibilities in a public cloud* subchapter (p.28). Cloud Service Providers (CSPs) manage the physical layer, as well as aspects like virtualization, networking, and storage (with the extent of responsibility depending on the chosen service model). As a result, you have limited control over your environment, and there are certain boundaries set by the CSP to ensure consistency and security across the cloud platform.

Let's explore a few examples of the challenges you may encounter in this context and how to minimize or mitigate them:

 Uncertainty regarding cloud infrastructure patching: In a public cloud, the schedule for infrastructure patching is not publicly shared knowledge. Although CSPs are bound to SLAs (Service Level Agreements), it is essential to ensure that any service interruptions do not adversely affect your obligations. While you cannot completely eliminate issues that impact the entire cloud platform's services, reputable cloud providers like Amazon, Microsoft, or Google experience such incidents infrequently due to their high availability setups. To address this challenge, design your solution with considerations for high availability, redundancy, and regional or global service dispersion.

- 2. Limited access to the infrastructure: In the public cloud model, excluding strict regulatory obligations, you do not have direct access to the infrastructure. This shift in responsibilities requires a change in mindset and operational models for cloud customers. Establishing trust in the CSP's ability to deliver the service securely is crucial, and cloud providers undergo regular audits by external parties, with the results or reports often published.
- 3. Limited influence on the service offerings: Each cloud provider has its own roadmap and agenda based on market trends and demands. While it may be challenging for individual customers to influence portfolio changes, selecting a cloud provider whose portfolio and roadmap align closely with your needs can help. Some CSPs engage in open communication with customers to identify desired features or services, shaping their portfolio development accordingly.

The last two challenges are indirectly related to data privacy but are still essential considerations:

5. Limited ability to influence the underlying infrastructure: As the physical layer is beyond the reach of cloud customers for security reasons, it may be difficult to obtain detailed information about the hardware used in CSP facilities. However, in some cases, when selecting specific types of virtual machines, you may learn more about the computing components, such as the type of processor. This information can aid in determining the most suitable tier or type of service for your solutions.

6. No access to datacenter facilities: In the public cloud, customers do not have access to the underlying infrastructure or datacenter facilities. For activities requiring direct access, such as transferring large amounts of data quickly, providers offer high-speed storage options. These allow you to upload your data, which is then transferred to the CSP's datacenter by their personnel. While physical visits to datacenters are not typically accessible, some providers offer virtual datacenter tours, providing insights into exemplary facilities.

Legal, regulatory, jurisdictional and contractual obligations

Be aware that compliance with regulations and legal frameworks is necessary to protect sensitive data, especially when dealing with globally dispersed solutions and a worldwide audience. Meeting various legal obligations can lead to controversies, challenges, and uncertainties. In a public cloud environment, you have numerous opportunities, but it is essential to approach legal, regulatory, jurisdictional, and contractual obligations with diligence. Here are some key aspects to consider:

- Location of data: Determining the location of data can be challenging, particularly with SaaS applications. Considerations include:
 - a. When your application is globally dispersed, you may need different security controls based on regional legal

requirements, which can impact your application unpredictably.

- b. With globally dispersed infrastructure and the use of CDNs, identifying the precise location of data storage and understanding applicable regulations can be demanding.
- c. In some cases, sensitive data may be stored and processed in two locations with conflicting legal rules, making it difficult to apply proper security controls.
- Data sovereignty: Data sovereignty means that the customer has control over the data through their security policies, controls, and mechanisms, and it is governed by local law to ensure compliance with regulations.
- 3. Data residency: Data residency refers to the obligation of storing data within a specific geographic location, such as within country boundaries. This requirement can originate from governmental or non-governmental institutions, including industry standards, and may be enforced through contractual agreements.
- 4. Compliance: Compliance relates to industry, national, or global standards and obligations that must be met. IT standards and regulations provide guidance on governance, policies, data protection, and other aspects. Compliance is a mandatory element for any deployment in the cloud, and violations can have severe consequences.

When considering compliance³⁵, be aware of respected security standards, formal certifications, and third-party audits that demonstrate adherence to IT security, privacy, and data management requirements. Some noteworthy standards include:

- ISO/IEC 27001: Information Security Management Systems (ISMS) provides a standardized model for developing and implementing security policies, procedures, and standards.
- 2. **ISO/IEC 27017**: Standard security controls for the use of cloud services.
- NIST 800-122: A resource for ensuring contractual and regulated Personally Identifiable Information (PII) requirements are met.
- ISO 27018: Code of practice and security techniques for processing Personally Identifiable Information (PII) in the cloud, emphasizing principles like consent, control, transparency, communication, and independent yearly audits.
- PCI DSS: Payment Card Industry Data Security Standard sets requirements for entities handling credit or debit card transactions, including security policies, controls, and monitoring techniques.
- SOC (System and Organizational Controls): SOC reports, including SOC 1, SOC 2, and SOC 3, provide insights into the safety and robustness of an organization's control

³⁵AWS: <u>https://docs.aws.amazon.com/whitepapers/latest/aws-risk-and-</u> <u>compliance/customer-cloud-compliance-governance.html</u> Azure: <u>https://learn.microsoft.com/en-us/azure/cloud-adoption-</u> <u>framework/govern/policy-compliance/regulatory-compliance</u> GCP: <u>https://cloud.google.com/compliance?hl=en</u> Want to read the full book?

Amazon

https://www.amazon.com/dp/B0CKQ3DLKH

Gumroad

<u>https://mwojnarowskapietrzak.gumroad.com/l/mi</u> <u>ndthegap</u>

For more information about book visit <u>https://mindtheqapcloudmistakes.com/</u>