



# SAP S/4HANA Systems in Hyperscaler Clouds

Deploying SAP S/4HANA in AWS, Google  
Cloud, and Azure

---

André Bögelsack  
Utpal Chakraborty  
Dhiraj Kumar  
Johannes Rank  
Jessica Tischbierek  
Elena Wolz

*Foreword by Alexander Zeier*



# **SAP S/4HANA Systems in Hyperscaler Clouds**

**Deploying SAP S/4HANA in AWS,  
Google Cloud, and Azure**

**André Bögelsack  
Utpal Chakraborty  
Dhiraj Kumar  
Johannes Rank  
Jessica Tischbirek  
Elena Wolz**

*Foreword by Alexander Zeier*

**Apress®**

# **SAP S/4HANA Systems in Hyperscaler Clouds: Deploying SAP S/4HANA in AWS, Google Cloud, and Azure**

André Bögelsack  
Bern, Switzerland

Utpal Chakraborty  
Kolkata, India

Dhiraj Kumar  
Kolkata, India

ISBN-13 (pbk): 978-1-4842-8157-4  
<https://doi.org/10.1007/978-1-4842-8158-1>

Johannes Rank  
Irschenberg, Germany

Jessica Tischbierek  
München, Germany

Elena Wolz  
Ismaning, Germany

ISBN-13 (electronic): 978-1-4842-8158-1

Copyright © 2022 by André Bögelsack, Utpal Chakraborty, Dhiraj Kumar,  
Johannes Rank, Jessica Tischbierek, Elena Wolz

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr  
Acquisitions Editor: Jonathan Gennick  
Development Editor: Laura Berendson  
Coordinating Editor: Gryffin Winkler

Cover image designed by Freepik ([www.freepik.com](http://www.freepik.com))

Distributed to the book trade worldwide by Springer Science+Business Media LLC, 1 New York Plaza, Suite 4600, New York, NY 10004. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit [www.springeronline.com](http://www.springeronline.com). Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail [booktranslations@springernature.com](mailto:booktranslations@springernature.com); for reprint, paperback, or audio rights, please e-mail [bookpermissions@springernature.com](mailto:bookpermissions@springernature.com).

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub.

Printed on acid-free paper

*To our families, colleagues, and friends.*

# Table of Contents

<b>About the Authors.....</b>	<b>xv</b>
<b>Foreword .....</b>	<b>xix</b>
<b>Appreciations .....</b>	<b>xxi</b>
<b>Introduction .....</b>	<b>xxv</b>
<b>Chapter 1: Introduction to Public Cloud and Hyperscalers.....</b>	<b>1</b>
Important Features of Public Cloud.....	1
Public, Private, and Hybrid Cloud .....	3
Overview of the Clouds of the Major Market Players.....	5
Amazon Web Services .....	5
Microsoft Azure .....	7
Google Cloud .....	10
Other Market Participants .....	13
SAP's Cloud Strategy .....	15
SAP RISE Program .....	15
SAP HANA Enterprise Cloud.....	20
Additional SAP Services from the Cloud.....	24
Summary.....	27
<b>Chapter 2: SAP S/4HANA Systems on Public Cloud.....</b>	<b>29</b>
Development of SAP S/4HANA.....	29
SAP S/4HANA in the Cloud .....	30
SAP S/4HANA Architecture.....	31
Overview.....	31
Data Layer .....	32
Application Layer .....	32

## TABLE OF CONTENTS

Presentation Layer.....	33
Communication Layer.....	34
Additional Components .....	34
Use Cases of SAP S/4HANA on Public Clouds.....	35
Sizing.....	35
Costs.....	43
High Availability .....	50
Disaster Recovery.....	56
Backup and Restore .....	61
Integration and Network.....	70
Automation .....	76
Horizontal and Vertical Scalability .....	80
Summary.....	84
<b>Chapter 3: SAP S/4HANA Deployment and Migration .....</b>	<b>85</b>
Contractual Basis and Support .....	85
Important Points to Be Considered.....	85
Direct Use or Use via a Partner.....	87
Needed Support.....	88
DSVGO .....	89
Support for Projects.....	89
GxP Framework and Important Certificates .....	90
GxP in General .....	90
Split of Responsibilities .....	91
Additional Certifications .....	92
Management of an SAP S/4HANA Landscape.....	94
Important Features of SAP Solution Manager .....	94
Application Operation Using SAP Solution Manager.....	96
Using SAP Cloud ALM .....	97

## TABLE OF CONTENTS

Choosing a Hyperscaler .....	99
Quantitative Factors .....	99
Qualitative Factors.....	101
Deployment and Migration.....	104
Transformation Is a Must.....	105
New Approaches Using the Cloud .....	106
Greenfield Deployment.....	106
Brownfield Deployment .....	107
Migration Scenarios for SAP S/4HANA Transformations .....	110
SAP S/4HANA As SaaS Services .....	111
Overview.....	111
Differences Between the Editions .....	112
Advantages and Disadvantages .....	115
Use Cases.....	116
Summary.....	116
<b>Chapter 4: SAP S/4HANA on AWS Elastic Compute Cloud – Concepts and Architecture.....</b>	<b>119</b>
History of Amazon AWS.....	119
New Infrastructure and Innovations .....	121
Reasons for AWS .....	122
AWS Business Growth in the Last Six Years .....	124
Offered Cloud Services .....	125
Compute .....	126
Storage .....	126
Networking .....	128
Identity Management.....	130
Security and Compliance.....	131
Management Tools .....	132

## TABLE OF CONTENTS

Regions and Availability Zones/Sets .....	134
AWS Region .....	134
Availability Zones.....	134
AWS Placement Groups.....	136
The AWS Console .....	141
Supported Browsers.....	141
User for Management.....	142
IAM Role and IAM User Group.....	144
Getting Started with the AWS Console.....	144
Changing the Region .....	145
Navigating to Services.....	146
Search AWS Service .....	147
Billing and Cost Management Service.....	149
Integration into Core Services.....	150
Internal Access Architecture.....	154
Internal and Limited External Access .....	156
Active Directory .....	166
Dynamic Host Configuration Protocol (DHCP).....	181
Summary.....	186
<b>Chapter 5: SAP S/4HANA on AWS Elastic Compute Cloud – Deployment.....</b>	<b>187</b>
Exemplary Architecture.....	188
Deployment Scenarios .....	188
Single-AZ and Single-Node Architecture .....	188
Single-AZ and Multinode Architecture.....	189
Multi-AZ, Single-Node SAP S/4 High Availability Architecture.....	191
Option A: Overlay IP Using AWS Transit Gateway .....	191
Option B: Overlay IP Routing with Network Load Balancer.....	193
SAP HA Cluster Setup.....	194
SAP NetWeaver/HANA High Availability Setup on RHEL and SUSE.....	194

## TABLE OF CONTENTS

Windows Server Failover Clustering (WSFC) Using Amazon FSx File Share Solution.....	199
Windows Server Failover Clustering (WSFC) Using SIOS Protection Suite on AWS .....	200
Compute.....	200
Available EC2 Instance Family.....	201
EC2 Instance Licenses.....	201
SAP S/4 Certified EC2 Instance Types .....	202
EC2 Provisioning.....	202
Storage.....	219
Amazon EBS .....	220
Amazon EC2 Instance Store .....	221
Amazon EFS.....	221
Amazon S3 (Simple Storage Service).....	221
Storage Layout for S/4HANA System.....	223
PAYG vs. Commitment – Calculator.....	228
On-Demand Instances .....	228
Reserved Instance .....	229
Payment Option .....	230
AWS Calculator .....	230
Security on AWS.....	234
Tools .....	237
Amazon Virtual Private Cloud Security (VPC).....	240
Subnet and Routing Tables.....	242
Network Access Control List.....	242
Security Groups (SGs).....	242
Virtual Private Gateway .....	243
Internet Gateway .....	243
Backup and Restore in AWS.....	243
SAP HANA DB Backup/Recovery via AWS Backint Agent.....	244
Restore Snapshot from AWS Console .....	255

## TABLE OF CONTENTS

<b>Disaster Recovery via AWS .....</b>	<b>257</b>
Passive DR Architecture – RPO Is More Than Zero and RTO Is Higher.....	257
Semi-active DR Solution [RPO Is Near to Zero and RTO Is Medium] .....	259
Active-Active DR Solution [RPO Near to Zero and RTO Is Very Less] .....	261
<b>Summary.....</b>	<b>263</b>
Cloud Foundation Design and Implementation.....	263
SAP S/4 Architecture Design and Implementation .....	264
<b>Chapter 6: SAP S/4HANA on Microsoft Azure – Concepts and Architecture.....</b>	<b>267</b>
Historical Overview of Microsoft Azure .....	268
Azure Control and Subscriptions.....	270
Azure Organizational Structure.....	270
Subscription Models.....	272
Subscription Management .....	273
Azure Resource Management.....	277
Resource Provisioning Using Resource Manager .....	277
General Deployment Options .....	280
Computing .....	283
Storage .....	297
Network and Services .....	305
Support and Licensing.....	315
Azure Disaster Recovery Services .....	315
Azure Backup .....	316
Backup Vault.....	316
Recovery Services Vault .....	316
S/4 on Azure Architecture .....	318
Basic Reference Architecture .....	318
Virtual Private Network Reference Architecture .....	321
Summary.....	323

## TABLE OF CONTENTS

<b>Chapter 7: SAP S/4 on Microsoft Azure – Deployment .....</b>	<b>325</b>
Azure S/4HANA Reference Architecture.....	325
Deploying a Basic Network Configuration via Azure Cloud Portal.....	327
Azure Cloud Portal .....	328
Naming Conventions and the Resource Concept.....	329
Creating Resource Groups .....	330
Network Configuration.....	331
Creating Network Security Groups .....	332
Deployment of a HANA HA Cluster .....	335
HANA Cluster Architecture .....	336
Provisioning HANA Cluster Resources .....	337
Configuration of the Pacemaker Cluster.....	351
Setting Up the HANA HA Cluster .....	359
Deploying an S/4HANA HA Cluster .....	364
SAP S/4HANA Cluster Architecture .....	364
Provisioning SAP S/4HANA Cluster Resources .....	366
Highly Available NFS Storage in Azure.....	370
Installation of ASCS and ERS .....	377
Installation of PAS and AAS .....	381
Excursus: Automated SAP Provisioning .....	382
Ansible and Terraform in Azure .....	382
SAP Cloud Appliance Library .....	383
Summary.....	385
<b>Chapter 8: SAP S/4HANA on Google Cloud – Concepts and Architecture .....</b>	<b>387</b>
History of Google Cloud.....	387
Development of the Partnership Between SAP and Google.....	388
Reasons for Choosing SAP on Google Cloud .....	391
Google Cloud Organizations and Resources .....	393
Google Cloud Resource Hierarchy .....	393
Properties of Google Cloud Resources .....	396

## TABLE OF CONTENTS

<b>Google Cloud Offered Services Important for SAP S/4HANA.....</b>	<b>397</b>
<b>Google Cloud Compute Engine .....</b>	<b>397</b>
<b>Storage Considerations .....</b>	<b>401</b>
<b>Google Cloud VPC and Networking .....</b>	<b>408</b>
<b>Google Cloud Security .....</b>	<b>415</b>
<b>Google Cloud Operations .....</b>	<b>419</b>
<b>Google Cloud Customer Care.....</b>	<b>420</b>
<b>Further Google Cloud Services Relevant for SAP S/4HANA Deployments .....</b>	<b>422</b>
<b>Google Cloud Front End.....</b>	<b>424</b>
<b>Google Cloud Front-End Tools.....</b>	<b>424</b>
<b>Typical User Roles in an SAP on Google Cloud Landscape .....</b>	<b>426</b>
<b>SAP S/4HANA on Google Cloud Architecture.....</b>	<b>427</b>
<b>Licenses and Sizing.....</b>	<b>427</b>
<b>SAP S/4HANA Architecture Overview .....</b>	<b>428</b>
<b>Setup for High Availability .....</b>	<b>429</b>
<b>Disaster Recovery Setup .....</b>	<b>436</b>
<b>Primary and Secondary Storage Considerations and Recommendations .....</b>	<b>439</b>
<b>Excursus: SAP S/4HANA Scale-Out (SAP Innovation Award 2021) .....</b>	<b>448</b>
<b>Excursus: SAP HANA Fast Restart and Memory Poisoning Recovery Capabilities for SAP S/4HANA.....</b>	<b>449</b>
<b>Summary.....</b>	<b>450</b>
<b>Chapter 9: SAP S/4HANA on Google Cloud – Deployment and Setup.....</b>	<b>453</b>
<b>Example Architecture for SAP S/4HANA on Google Cloud.....</b>	<b>453</b>
<b>Planning and Deployment Checklists for SAP on Google Cloud .....</b>	<b>456</b>
<b>Google Cloud Account, Network, and Security Setup.....</b>	<b>457</b>
<b>Google Account, Cloud Billing, and Identity and Access Management Setup .....</b>	<b>458</b>
<b>Shared VPC, Subnets, Firewall Rules, Cloud NAT, and Cloud DNS Setup .....</b>	<b>459</b>
<b>Setup of Cloud DNS .....</b>	<b>463</b>

## TABLE OF CONTENTS

Google Cloud Compute Setup .....	466
Deployment of the SAP HANA Database.....	466
Deployment of SAP NetWeaver Application Servers.....	478
Pricing and Billing Concepts with Best Practices .....	487
On-Demand vs. Committed Use Discounts.....	487
Reservations in Google Cloud.....	489
Google Cloud Quotas and Budgets .....	490
Pricing Calculation for the SAP S/4HANA on Google Cloud Architecture Example .....	491
Backup and Restore on Google Cloud .....	495
Backups for SAP NetWeaver Application Server .....	496
Backups for SAP HANA Database.....	498
Scripting and Automation on Google Cloud.....	500
Tools for Scripting and Automation on Google Cloud.....	500
Use Case: Automated Start and Stop of Instances .....	501
Use Case: Autoscaling for SAP NetWeaver Application Servers .....	502
Use Case: Simplification of Operations.....	503
SAP Landscape Management and Google Cloud .....	503
Disaster Recovery on Google Cloud .....	504
SAP HANA Disaster Recovery .....	504
SAP NetWeaver Application Server Disaster Recovery.....	504
Recommendations for Disaster Recovery Planning .....	504
Summary.....	505
<b>Chapter 10: Summary and Outlook.....</b>	<b>507</b>
Momentum of SAP S/4HANA.....	507
Public Cloud Is an Established Trend .....	508
Public Cloud As Innovation Driver .....	509
Outlook.....	510
<b>Index.....</b>	<b>513</b>

# About the Authors



**André Bögelsack** works as Managing Director for Accenture in Switzerland and supports customers of all industries in moving their SAP systems into the cloud and operating those systems in a cost-efficient way on Hyperscaler Cloud platforms. He received a PhD from Technische Universität München in the area of SAP.



**Utpal Chakraborty** works as Manager for Accenture in India and serves multiple clients across the industries. His main focus is migrating and operating SAP S/4 systems in the Hyperscaler Cloud. He combines a strong understanding of the technology of public clouds as well as SAP technology, which helps to mitigate common risks to such projects.

## ABOUT THE AUTHORS



**Dhiraj Kumar** works as Senior Manager for Accenture India, and, together with his team, he managed the move of SAP systems of several customers into the public cloud. He has a strong focus on the automation of SAP migrations and new ways of provisioning SAP in the cloud.



**Johannes Rank** is Head of the SAP Basis team at the SAP University Competence Center at Technische Universität München. He is a certified SAP HANA Technology Associate and oversees the operations of all the SAP S/4HANA systems with a strong focus on migration and automation. He holds a Master of Science in Business Informatics from Technische Universität München.



**Jessica Tischbirek** is Specialist Customer Engineer for SAP on Google Cloud at Google Deutschland in Munich. Before starting at Google, she was in the SAP consulting business for a long time. She holds a Bachelor of Science and a Master of Science in Business Informatics and an extensive experience in SAP technology and public clouds.



**Elena Wolz** works as SAP Basis Administrator for the SAP University Competence Center at Technische Universität München. She is the Product Owner for SAP S/4HANA and is responsible for all the deployments of S/4 systems in Hyperscaler Clouds as well as traditional hosting scenarios. She holds a Master of Science in Business Informatics from Technische Universität München.

# Foreword

The constant technological change and the compulsion to digitize have never been more evident than in the past two years when the Corona pandemic changed work habits drastically. It was at this point that Chief Information Officers (CIOs) became Chief Digitization Officers with new challenges and new demands from the wider world of work.

All those CIOs who had already consistently focused on digitization before Corona and had understood the technical change not as a challenge, but as an opportunity, could consider themselves lucky. The companies that had already focused on high flexibility and scaling before the pandemic and had already moved their IT systems to the public cloud were well prepared. They were able to pass on the full potential of the public cloud to their employees.

With the same speed that the Corona pandemic drove cloud adoption, CIOs must now address the second major IT challenge: SAP S/4HANA. The shift from traditional SAP ERP systems to a digital core based on SAP S/4HANA is a daunting task in many organizations and requires close collaboration between business and IT departments.

By combining the use of public cloud services with the transformation to SAP S/4HANA, companies open up new possibilities. Business processes can be redesigned and made more flexible, and companies create a new, highly flexible technological platform for the continued success of the company.

This book shows how companies can easily build a new technical platform for future SAP S/4HANA systems in the public cloud. However, the book does not limit itself to just one of the major public clouds but describes the three most important market players: Microsoft Azure, Amazon Web Services, and Google Cloud. The reader thus receives a comprehensive overview of the individual public clouds and learns how those can best be used for SAP S/4HANA systems.

The book is a must-read before implementation for all companies that are not yet using SAP S/4HANA and are yet to create this technical platform. It is also suitable for all those companies that are already in a public cloud but still want to gain inspiration and see how other companies are making this transition.

## FOREWORD

I am pleased that the authors have summarized the topic so comprehensively but succinctly in one book. The technical depth is remarkable and reflects the importance of the topic. The book will help you, dear reader, to answer the important questions regarding the deployment, migration, and operation of SAP S/4HANA systems in the public clouds. I hope you enjoy it!

Prof. Dr. Alexander Zeier  
1. Cloud Fellow  
CTO ASBG  
Global Managing Director  
Accenture GmbH

# Appreciations

Book projects are always a great challenge for the authors and require a lot of time and energy, which is not available for other things. Therefore, the authors would like to take this opportunity to thank their families, colleagues, as well as project partners and customers.

## *Individual Appreciations*

### *André Bögelsack*

I would like to thank my wife Kathrin and my two children for their understanding when I sat at the laptop on weekends or in the evenings rather than spending time with them. A special thanks goes to **Niklas Feil**, who provided a lot of content for the book. Many thanks for this! Another special thanks goes to Samiksha Munjal for keeping my back free, when I was too much focused on writing the book. For inspiring me to write another book, I would like to thank Debolina Banerjee – thank you!

My thanks also go to all the co-authors, without whose expert knowledge and time commitment we would never have written the book. I am very happy that we were able to do this together.

### *Utpal Chakraborty*

First of all, I would like to express my gratitude to Dr. André Bögelsack, who gave me the opportunity to share my knowledge and experience with a wider audience through this book, which would not have been possible without him. I would also like to thank my family who supported me in this pandemic situation, and I dedicate my achievement to my daughter Sreejani.

Last but not least, a big thank you to my co-authors, without whose support this would not have been possible.

### *Dhiraj Kumar*

My thanks to my wife Shikha for her constant motivation and understanding and to my son Agastya for giving me the time to write this book. My special thanks to my friends Ravi Kumar and Ravi Surana, Kunal, Misal, and Sushobhit for keeping me alive (and grounded) with their weekend calls.

## APPRECIATIONS

A special thanks to Dr. André Bögelsack, who believed in me more than I believed in myself. Thank you André for always being there for me. This book has taken its shape through the great collaboration with my co-authors. To all of them, my very special thanks for their time and expertise.

Last but not least, thank you to my brother Manoj Kumar, who always inspires and encourages me in whatever I do. Thank you Bhaiya.

### *Elena Wolz*

I would like to thank my family and friends for their support and understanding when I spent evenings in the company with load balancing modules and availability groups instead of with them. I would also like to thank the SAP UCC TU Munich for their support during the project, as well as my colleagues at the SAP UCC TU Munich for their everlasting encouragement. I would also like to thank my co-authors, whose contributions and collaborative exchanges helped to provide new perspectives. A special thanks goes to Johannes Rank; I appreciated the close collaboration very much.

### *Johannes Rank*

A big thank you goes to my co-authors of this book, especially to André who made it possible for me to contribute my expertise in the first place and without whose help this book would probably not have been possible. I would also like to say a special thank you to my family. To my daughter, who had to do without her dad many weekends and evenings, and to my wife, who always had my back.

### *Jessica Tischbirek*

First, I would like to thank my co-authors for the exchange of experiences and especially Dr. André Bögelsack for the request, ideas, and opportunity to contribute to this book project.

A huge thank you goes to all my colleagues who supported this book with their ideas, motivation, and especially their review of the content: Hasmig Samurkashian, Thorsten Staerk, Benjamin Schuler, Alexander Schmid, Abdelkader Sellami, Alison Hettrick, and to my manager Ian Moria.

Finally, a very special thank you goes to my partner Mike for his ongoing support, critical questions, as well as his understanding of the intense book project.

### *Joint Acknowledgments*

#### *Accenture*

The author team would like to thank all those who contributed to the success of this book and who supported the book during its creation. In particular, this includes Tobias

Regenfuss, Alan Mohr, Veronica Wolters, Bernhard Schulzki, and the entire Accenture Cloud First leadership team. Without their dedication and support, this book would not have been written.

#### *SAP*

We would also like to thank SAP and the SAP University Alliances (SAP UA) program, which has supported universities worldwide for many years and helped make book publications possible in this way.

#### *Partners and customers*

Writing a book about multiple Hyperscalers can only succeed with the support of Microsoft, Amazon, and Google. Without this support on technical issues, the book would not have come to fruition. For this, the entire author team would like to thank – especially the Alliances teams of the vendors.

To make the book more exciting with the many examples from the customers, we also need the support of the respective companies. We would like to thank LANXESS and especially Tobias Greskamp for their support.

#### *Technische Universität München*

In addition, we would like to thank the Technische Universität München, the Chair of Information Systems (i17), and the SAP University Competence Center (SAP UCC), which provide an excellent basis for research and development in the field of SAP and cloud computing. Special thanks go to the Head of the SAP UCC, Dr. Holger Wittges, and to the Chair Holder Prof. Dr. Helmut Krcmar.

#### *Google*

We would like to thank Google Cloud and all teams working around SAP on Google Cloud: Solution and Sales Specialists, Product Engineering, the Center of Excellence, SAP Management, the SAP Partnership Team, and many more. A very special thank you to the global SAP Marketing team and the Accenture Partner Alliances team.

#### *Apress*

Last but not least, we would like to thank our lectors, who helped us and guided us during the whole time. Many thanks to Jonathan and Jill for their guidance and patience.

**André Bögelsack, Utpal Chakraborty, Dhiraj Kumar, Elena Wolz,  
Johannes Rank, and Jessica Tischbierek**

# Introduction

The public cloud has existed for more than ten years, and it is impossible to imagine the current IT landscape without it. Many well-known companies rely on services from the public cloud to strengthen their IT, and Hyperscalers now form the backbone of many companies from a wide range of industries. Millions of customers and consumers of the companies use websites or order portals, which are operated on one of the large public clouds, in their daily operations.

SAP systems represent the neuralgic points in the IT operations of companies. Production lines come to a standstill or trucks can no longer leave the warehouses if SAP systems fail. Hardly any other IT system has this importance in today's world. Nevertheless, more and more companies are moving SAP systems to the public clouds. This is done due to cost pressure, pressure to innovate, or simply to expand capacity.

## Goal of the Book

With the current pressure among all SAP customers to jump to the new generation of SAP S/4HANA systems, the topic of public clouds is also gaining momentum. Companies are taking advantage of this S/4 momentum to migrate to a public cloud in addition to transforming to SAP S/4HANA.

This book will introduce you to several points of public clouds and SAP S/4HANA systems. Here, in addition to the theoretical constructs, the practical implementations are also shown. The book enables you to directly implement and provision a new SAP S/4HANA system on one of the three big public clouds: Azure, AWS, and Google. The chapters are structured identically, and thus the necessary steps for setting up and operating SAP S/4HANA systems can be compared between the public clouds.

In all chapters, the authors provide real-world examples to make the content tangible and applicable. The book will detail the following aspects:

- The emergence of public clouds and the three largest Hyperscalers:  
Microsoft Azure, Amazon Web Services, and Google Cloud
- The offerings available from SAP for the use of cloud services

## INTRODUCTION

- The architecture of SAP S/4HANA systems
- The key drivers and use cases for SAP S/4HANA systems on a public cloud
- The provisioning and operations of SAP S/4HANA systems on the three main public clouds: Microsoft Azure, Amazon Web Services, and Google Cloud

After completing all chapters, you will know what to look for when deploying public clouds and how to efficiently run SAP S/4HANA systems on public clouds without getting out of line with the effort required to run them or, alternatively, the cost of running them on public clouds.

## Organization and Structure

The book guides you through the typical life cycle of an SAP S/4HANA system, highlighting the important aspects of selecting a public cloud, provisioning SAP S/4HANA systems, and operating them on the public clouds:

- Chapter 1 provides an overview of the most important terms, such as private, public, and hybrid clouds. It also presents the three most important public clouds in detail and outlines the current market of cloud providers with, for example, the IBM Cloud. Chapter 1 then sheds light on SAP's cloud strategy, the SAP RISE program, as well as the SAP HANA Enterprise Cloud, which is present in the market as a private cloud.
- Chapter 2 describes the architecture of the SAP S/4HANA systems, which will later be taken up again and again in the following chapters. It also contains the important aspects for operating SAP S/4HANA systems on the public cloud, such as sizing, backup, high availability, and automation. The following chapters on the respective public clouds show the implementations of these points.
- Chapter 3 describes the general procedures for deploying and migrating SAP S/4HANA systems and explains, for example, points such as greenfield and brownfield deployments, but also shows possible migration scenarios for the systems. It also discusses the most important factors when selecting a new Hyperscaler.

- Chapter 4 is the first part of the book to describe the deployment of Amazon Web Services. To this end, the chapter first explains the most important terminology, which is AWS specific. It goes into the types of machines available, but also explains the basic concepts of networking and storage.
- Chapter 5 is the practical part on AWS and shows how to create and control an SAP S/4HANA system on AWS. For this purpose, the chapter shows the system to be implemented and then describes the steps that can be taken to build the system on AWS. The important typical use cases, such as recovery after a disaster case, are described successively.
- Chapter 6 is the first chapter on Microsoft Azure and describes how Azure can be used as a public cloud solution and which important basic concepts exist for the operation and provisioning of SAP S/4HANA systems.
- Chapter 7 builds on the foundation previously established in Chapter 6 on Microsoft Azure and illustrates the provisioning of a new SAP S/4HANA system and shows the implementation of the important use cases from Chapter 2.
- Chapter 8 highlights the important concepts and architectures of Google Cloud and describes how SAP S/4HANA systems are scheduled in Google Cloud.
- Chapter 9 uses the previously described concepts of the Google Cloud to show the concrete implementation of an SAP S/4HANA system in the Google Cloud.
- In the concluding chapter, Chapter 10, all the important points from the previous chapters are summarized, and an outlook is given on developments in the coming years.

After completing all chapters, you, dear reader, will have a comprehensive overview of all aspects of the topic “SAP S/4HANA systems on Hyperscaler clouds” and will know how to plan, conceptualize, implement, and control these systems on public clouds.

## CHAPTER 1

# Introduction to Public Cloud and Hyperscalers

Before entering the world of SAP S/4HANA systems on the public cloud, this chapter first describes the basic features of cloud computing and introduces the most important terms. This makes it easier to understand the rest of the book.

## Important Features of Public Cloud

Cloud computing emerged in the mid-2000s and was developed on the basis of excess capacity in data centers. Originally, the approach was to allocate excess computing capacity to other customers. The goal was to achieve consistently high utilization throughout the year.

Cloud computing is, on the one hand, about a new way of delivering services to customers and businesses. For example, cloud services and services from the cloud can be used *without major prior contract negotiations*. Only the general terms and conditions (GTCs) have to be accepted and a valid payment option has to be deposited, and then the services can already be used.

On the other hand, cloud computing is an important factor in *innovations* and the introduction of completely new services. In the meantime, not very much has remained of the original idea of making surplus capacities available. Cloud providers are constantly making new services available and regularly also new versions of the services. Customers benefit from this because new services do not have to be explicitly requested but come automatically with the use of the cloud.

In addition to the *simplicity* of using the services and the very high level of innovation, another point is in focus. In principle, the cloud knows no limits, and so customers and companies can assume limitless capacities in a cloud. This distinguishes clouds from conventional providers, where resources are usually limited.

## CHAPTER 1 INTRODUCTION TO PUBLIC CLOUD AND HYPERSCALERS

The *high scaling* of the clouds, that is, the large number of customers, means that the economy of scales applies in full. Here, customers benefit from providing the same services to other customers and achieving a very low price.

If the most important features of a cloud are summarized, the following characteristics are the key points of clouds:

- Simplicity of use and access
- High level of innovation through new services
- Very large capacities/infinite capacities
- Very high price attractiveness

In the context of Hyperscalers, such as Amazon, Microsoft, and Google, the preceding criteria apply, and companies can take advantage of the benefits. However, the many advantages are also paired with some disadvantages when using the clouds. Here, it is not so much technical, but rather organizational and process-related points that can be interpreted as disadvantages.

Although the use of clouds is very simple, the use must also be constantly and continuously controlled. Very often, a “shadow IT” exists alongside the actual IT, in which specialist departments have provided themselves with IT services from the cloud. This contradicts central IT and efficient control of the entire IT.

Although the prices of the cloud are very good, the use of the cloud in a first step can become more expensive than if a company had continued to work without a cloud. So, to a certain extent, the use of the cloud must be accompanied by a transformation of IT. Simply transferring IT services to the cloud without transformation usually does not bring the hoped-for effects.

Unlimited resources in a cloud were a reality for a very long time. There were never any reports of bottlenecks or unavailability of resources. Since the Corona crisis, however, it has become clear that clouds are also “just” data centers with conventional capacities. At the latest since the Microsoft Azure Cloud ran out of spare capacity for certain SAP workloads in March 2020, the myth of infinite resources has been disproved.

Companies have to do a lot themselves when using clouds. This is very different from using services from a service provider. Here, the provider does a lot of tasks and takes responsibility for the tasks and results. A cloud provider will only do this for the services that are their responsibility. In contrast to a service provider, however, these are much fewer. Therefore, the customers have to take over tasks or commission another provider with them.

When operating SAP systems on a cloud, the difference to the service providers must be taken into account. This is illustrated in Table 1-1.

**Table 1-1.** *Responsibilities in the Public Cloud*

Component	Responsibility with Customer	Responsibility with Cloud Provider	Responsibility with Service Provider
Data center			
Compute/storage/network			
Security and compliance			
Operations of infrastructure			
Stack design			
Operating system			
Operations of SAP basis			
Operations of non-SAP applications			

All dark fields in the table show the responsibilities of the respective parties. It is clear that the typical service provider can map the complete stack and provide all services in a traditional hosting. When using the cloud, customers must provide services themselves, as the cloud provider does not.

## Public, Private, and Hybrid Cloud

The cloud computing market is very heterogeneous and permeated by many trends. Three main clouds are used for SAP workloads and operations in a cloud: the public cloud, the private cloud, and the hybrid cloud.

## CHAPTER 1 INTRODUCTION TO PUBLIC CLOUD AND HYPERSCALERS

The **public cloud** is a cloud that most companies are familiar with. It is the type of cloud which is offered by all Hyperscalers (i.e., Amazon, Microsoft, and Google). The name public already implies the most important feature of the cloud – it is public and equally accessible to every customer. However, this does not mean that all companies can see all workloads and all data of all other companies. It simply means that the cloud is available to all companies equally and that the cloud has not been built dedicatedly for just one company. The providers of the public cloud have implemented mechanisms for this, which also make access to the resources of a company possible only for the actual company.

In contrast to a public cloud, the **private cloud** is set up exclusively for one company. All resources of this private cloud are only available to one company and can also only be used by the company. However, the private cloud has a decisive disadvantage. The built-up resources of the private cloud can be used until the capacities are reached and must then be expanded. Private cloud operators implement such expansions in a similar way to normal service providers: hardware resources are acquired exclusively for the customer and therefore also invoiced. Another disadvantage of private clouds is the strong focus on the intended use. The providers of private clouds do not enrich the clouds with new services, but leave them as they are.

In addition to the public cloud and the private cloud, there is also the *hybrid cloud*. This is not so much a cloud but a merger of a public cloud with a private cloud or with an existing data center of a company. This already shows that the majority of companies are directly in a hybrid cloud scenario as soon as the first services from the public cloud are consumed.

Many companies still operate SAP systems in their own data centers. These data centers often reach their capacity limits or are heading to end of support because they have become too old. This also happened at a large German manufacturing company. The data centers on the same campus were to be demolished to make room for new office buildings.

The company had several options open to it. It could use an external service provider with data center capacity, it could buy temporary data center capacity and operate it itself, or it could use the public cloud. The company strategically opted for the public cloud. Due to many projects and project activities, various numbers of SAP systems were required. When the company started provisioning these systems in the public cloud, it moved into the hybrid cloud scenario. Many systems still exist in the two data centers to be demolished, but the new systems are in the public cloud. This scenario is a hybrid one.

The preceding example shows that by simply using the first service from the public cloud, a company goes into the hybrid cloud scenario. This is what happens to all companies that go to the cloud with the first service.

Companies that combine two public clouds, for example, Google and Amazon, can also create a hybrid cloud. For example, some companies decide to move SAP workloads to one public cloud and the non-SAP workloads to another public cloud. In this way, two public clouds are combined with each other, and a large hybrid cloud is created.

## Overview of the Clouds of the Major Market Players

Currently, the public cloud market is dominated by three Hyperscalers. These are Amazon Web Services, Google Cloud, and Microsoft Azure. All three Hyperscalers are presented in this chapter. In addition, the other cloud providers, which currently have small footprints such as IBM or Oracle, are also described.

### Amazon Web Services

#### History

Amazon Web Services was the first real cloud provider on the market and is still one of the leading providers of cloud computing services. There are several versions to the reasons why Amazon Web Services started.

One version is based on a situation of excess capacity. In the early 2000s, Amazon's business was developing steadily, and due to the strong sales activities on Amazon's platform, the demands for computing power to provide the platform and its services became enormous. Particularly due to seasonal effects, such as Black Friday in the United States, demand increased dramatically and then dropped rapidly back to normal levels. Amazon reacted to this and built the platform accordingly in such a way that the load peaks were cushioned by the seasonal events and the platform ran in a performant and stable manner. This focus on maximum performance led to a steady overcapacity in Amazon's data centers during the rest of the year – outside of the seasonal effects. This excess capacity was made available as resources for customers. Then, in 2006, Amazon launched its Amazon Web Services subsidiary, which made capacity from Amazon's data center available externally through an open access portal.

Another version of the emergence of cloud computing at AWS is based on the internal processes and workflows within Amazon. Here, Amazon faced recurring issues due to the steady growth of its platform, where web developers had to repeatedly coordinate with data center teams on network, capacity, and availability. To eliminate this dependency, a type of commodity IT was introduced (AWS) that allowed web developers to assemble the necessary resources themselves. In other words, they became independent of data center teams. Amazon realized that any web developer would appreciate such capabilities and decided to make these new services available as Amazon Web Services.

Regardless of the version to emerge, the launch of Amazon Web Services was a resounding success. Within a few years, AWS managed to create a large customer base and thus achieve high usage of the platform – even if it was initially relatively complicated to get your own SAP systems up and running on AWS.

## Current Market Position

AWS is the market leader and serves a very large number of customers from global corporations, as well as small- and medium-sized businesses and home users. AWS serves this broad abundance of customers by offering a very wide range of products.

AWS stands out among customers because of the following:

- High density of AWS data centers in many different countries with availability zones
- Continuous updates of the cloud services with new features on a regular basis
- High availability of services with no major impact on a large number of customers so far
- Very large ecosystem with many third-party providers offering additional services
- Ease of adoption

However, AWS' clientele also sees some negative points:

- All services are paid, and in some cases the prices for cloud services are very high. This can lead to very high bills for using AWS.
- Strict contracts between customers and AWS with no flexibility to customize for businesses.