


Developing Microservices Architecture on Microsoft Azure with Open Source Technologies



Ovais Mehboob Ahmed Khan
Arvind Chandaka

Foreword by Robert Vettor, Principal Cloud Solution Architect, Microsoft

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



Developing Microservices Architecture on Microsoft Azure with Open Source Technologies

Ovais Mehboob Ahmed Khan
Arvind Chandaka

Developing Microservices Architecture on Microsoft Azure with Open Source Technologies

Published with the authorization of Microsoft Corporation by:
Pearson Education, Inc.

Copyright © 2021 by Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-681938-7

ISBN-10: 0-13-681938-9

Library of Congress Control Number: 2021937949

ScoutAutomatedPrintCode

TRADEMARKS

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

WARNING AND DISCLAIMER

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author, the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the programs accompanying it.

SPECIAL SALES

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

CREDITS

EDITOR-IN-CHIEF

Brett Bartow

EXECUTIVE EDITOR

Loretta Yates

DEVELOPMENT EDITOR

Rick Kughen

SPONSORING EDITOR

Charvi Arora

MANAGING EDITOR

Sandra Schroeder

SENIOR PROJECT EDITOR

Tracey Croom

COPY EDITOR

Rick Kughen

INDEXER

Cheryl Ann Lenser

PROOFREADER

Abigail Manheim

TECHNICAL EDITORS

Doug Holland

Thomas Palathra

EDITORIAL ASSISTANT

Cindy Teeters

COVER DESIGNER

Twist Creative, Seattle

COMPOSITOR

codeMantra

I would like to thank my family for supporting and encouraging me in every goal of my life.

—Ovais Mehboob Ahmed Khan

I dedicate this book to my family and friends for their everlasting encouragement and support.

—Arvind Chandaka

This page intentionally left blank

Contents

<i>Acknowledgments</i>	<i>xi</i>
<i>About the authors</i>	<i>xii</i>
<i>Foreword</i>	<i>xiii</i>
<i>Introduction</i>	<i>xv</i>

Chapter 1	Introduction to microservices	1
	Our journey with microservices	1
	Evolution of software architecture	2
	Monolithic architecture	2
	Service-oriented architecture (SOA)	3
	Comparing monolith with microservices	4
	SOA versus microservices	5
	Monolith example	6
	Microservices example	7
	Databases in a monolithic architecture	8
	Databases in microservices architecture	9
	Micro front-ends	9
	Core fundamentals of microservices	11
	Benefits	11
	Challenges	13
	Moving forward with the microservices architecture, open source, and Azure	16
	This book's goals	17
	Summary	17
Chapter 2	Modeling microservices—real-life case study	19
	Application requirements	20
	Application features	20
	Identity management and authentication scenarios	20
	Auction management	20
	Bid management	21

Payment management	22
Application flow.....	23
Decomposition principles.....	24
Decomposition strategies	24
Decomposition by business capability	25
Decompose by subdomain	25
Domain-driven design.....	27
Ubiquitous language	27
Bounded context	28
Domain categories	30
Online auction system decomposition based on DDD	31
Anti-patterns	33
Using monolith or a shared database with microservices	33
Unnecessary fine-graining of services to deeper subdomains	33
Establishing tight dependencies between code artifacts	34
Summary	34
Chapter 3 Build microservices architecture	35
Cloud-native applications.....	36
Main principles of cloud-native applications	36
Characteristics of cloud-native applications	37
Twelve-factor app methodology	40
The online auctioning system (OAS) architecture.....	41
Representation of Azure Kubernetes cluster nodes, pods, and services	42
Technologies used	43
Front-end technology	43
Technologies used for building microservices	44
Cloud technologies	45
Azure WebJobs	46
Azure Event Hubs	47
Azure API Management	47
Azure AD B2C	48

Azure Kubernetes Services	48
Azure Container Registry	49
Azure DevOps	49
Azure Application Insights	49
Azure Monitor	50
Distributed database architecture	50
Transactional data model	50
Transient data model	50
Polyglot persistent architecture	51
Patterns in distributed databases	52
Direct HTTP call	52
Aggregator pattern	53
Command query responsibility segregation	54
Summary	57
Chapter 4 Develop microservices and front-end applications	59
Developing microservices	60
Developing the auction service	60
Developing the bid service	68
Provision Cosmos DB in Azure	70
Create a bid service in the JavaSpring Boot framework	72
Developing a payment service	77
Developing an application front-end	88
Prerequisites	89
Creating a front-end application	89
Understanding the Angular project structure	90
Angular concepts	93
Developing a security module	95
Configuring environment files	101
Develop the create auction form	102
Developing an active auctions page	105
Developing a submit bid form	107
Summary	111

Chapter 5	Microservices on containers	113
	Containers Overview	113
	Docker as a container technology	114
	Install Docker	114
	Docker components	116
	Docker commands	116
	Linux versus Windows containers	117
	Build Docker images.....	118
	Containerize the auction service	118
	Containerize the bid service	120
	Containerize the payment service	123
	Deploy images to Azure Kubernetes Services	125
	Kubernetes architecture	125
	Provision Azure Kubernetes Services	127
	Provision the Azure Container Registry	129
	Push services to ACR	129
	Deploy services to AKS	130
	Create a deployment object for OAS microservices	131
	Create a service object for OAS microservices	135
	Deploy a front-end application in the Azure App Service	137
	Deploy the Kafka Listener Service as an Azure WebJob	142
	Summary	143
Chapter 6	Communication patterns	145
	Approaches to communication	145
	Synchronous versus asynchronous communication	146
	Request/response communication	147
	Pub/sub communication	148
	The best communication approach for microservices	149
	Pub/sub communication technologies.....	153
	Apache Kafka	153
	Azure Event Hubs	155
	RabbitMQ	156

	Set up Kafka to establish pub/sub communication	157
	Infrastructure setup	157
	Setting up the producer: Adding Kafka support in the Java application	159
	Setting up the Consumer: Develop Kafka Listener service with .NET Core Hosted Service	161
	Summary	173
Chapter 7	Security in microservices	175
	An overview of security and architectures	175
	IaaS and PaaS architecture security	176
	PaaS security	177
	Zero-trust architecture	179
	Authentication and authorization flows	182
	Azure Active Directory B2C	187
	End-to-end OAS security implementation	189
	User perspective	191
	Microsoft Authentication Library (MSAL)	192
	Creating a tenant	192
	Register your application	193
	Configuration	193
	User Flows	196
	Summary	200
Chapter 8	Set up Azure API Gateway	201
	Why do you need an API gateway?	201
	Azure API Management	202
	Key benefits of using Azure API Management	203
	Set up Azure API Management	203
	Configure APIs in Azure APIM	204
	Working with policies and expressions	210
	Strategies when using Azure API Management with Azure Kubernetes Services	214
	Strategies to configure APIM with Azure Kubernetes Services	218
	Summary	220

Chapter 9	Build and deploy microservices	223
	Continuous integration and continuous deployment	223
	Automating infrastructure through Infrastructure as Code.	225
	OAS Infrastructure as Code with Terraform	225
	Build a pipeline or continuous integration	229
	Deployment pipeline or continuous deployment	234
	Building CI/CD pipelines for the OAS microservices	237
	CI/CD pattern and best practices	237
	Deployment patterns	238
	The auction service build pipeline	239
	Auction service deployment pipeline	241
	A complete look at DevOps	246
	Summary	247
Chapter 10	Monitoring microservices	249
	Monitoring concepts and patterns	249
	Log information	250
	Azure Monitor	253
	Azure Application Insights	254
	Monitoring framework and best practices	257
	Azure Application Insights configuration	258
	Container Insights	265
	Dashboards	266
	Summary	268
	<i>Index</i>	269

Acknowledgments

This book involved great effort on our part. With the recent growth in microservices architecture, we saw the importance of developing a sample case study on microservices to empower our customers to build better applications. We created a use case to illustrate microservices, created a functional application, and evangelized it to our customers—all before even putting pen to paper to write this book. Our friends and family were the spectators during these laborious endeavors, and they showed nothing but encouragement and understanding during long nights and tiring weekends. We couldn't have done it without them!

A big thanks goes to our managers for their everlasting support. Barbaros Gunay, Andrew McCreary, Mekonnen Kassa—thank you! Without your support, this book could not have been done.

When we both started working at Microsoft, we were quickly taken under the OSS Enablement Team's wing. Rick Hines and Bahram Rushenas are the team leads who encouraged us to pursue the project that ultimately became this book. Whether it was discussing a design or architecture of an application, brainstorming ideas, or presenting at events, their help and support were tremendous. We want to thank you for everything you've done.

We also want to thank Sally Brennan for preparing our session outline for Microsoft Ready, our premier go-to-market conference, which led us to further refine the content we used within our book.

We also want to thank our team at Pearson. Thank you so much, Loretta Yates, Charvi Arora, Rick Kughen, Doug Holland, Thomas Palathra, Abigail Manheim, and Tracey Croom for supporting us during our journey. We would like to thank Rob Vettor for lending his time to review our book and providing a foreword. Rob is an extremely experienced individual in microservices at Microsoft and we greatly appreciate his support.

About the authors

Ovais Mehboob Ahmed Khan is a seasoned programmer and solutions architect with nearly 20 years of experience in software development, consultancy, and solution architecture. He has worked with various clients across the United States, Europe, Middle East, and Africa, and he currently works as a Senior Customer Engineer at Microsoft, based in Dubai. He specializes mainly in application development using .NET and other OSS technologies, Microsoft Azure, and DevOps.

He is a prolific writer and has published several books on enterprise application architecture, .NET Core, VS Code, and JavaScript, and he has written numerous technical articles on various sites. He likes to talk about technology and has delivered various technical sessions around the world.

Arvind Chandaka is a product manager at Microsoft and has led products across Azure and Cloud + AI Engineering. He has worked with strategy executive teams to grow Microsoft Azure as well as built IP at the company resulting in several services and products. He is a recognized SME in enterprise technology and has advised many Fortune 100 companies and international clients around the world. He specializes across infrastructure, identity, cyber security, open source, and more. A strong believer in empowerment, he works closely with startups, is a founder himself, invests in the entrepreneurial ecosystem, and serves on the board of non-profits in New York City.

Arvind earned his Bachelor of Science at Cornell University where he studied several disciplines ranging from computer science to business. In his spare time, you will likely find him at a tennis court or traveling the world for the next great foodie destination.



Foreword

For many years, the guidance for building business applications was clear-cut:

- Create a large, monolithic core that contains business logic
- Persist data to a shared relational database
- Expose functionality through front-end UIs and API endpoints
- Build, package, and deploy the application, and you're in business.

While often challenging, such monolithic applications were *straightforward* to build, test, deploy, and troubleshoot. With code executing in a single process, performance could be good. When the need arose, you could scale the application up (vertically) by adding more server resources.

This model, however, had limitations. When an application proved popular, the monolithic architecture would eventually break down. As the app grew larger, more complex, and more “coupled,” it became less agile:

- Feature updates and fixes would cause unintended and costly side effects, such as breaking existing functionality.
- One unstable component could crash the entire system.
- Scaling any component required scaling the entire application.
- New technologies and frameworks were not an option.
- Each change required a full deployment of the application.

Perhaps the biggest challenge was *agility*. Monolithic applications prevented the business from moving fast.

Fast-forward into the world of modern apps and microservices. They are all about speed and agility. By decomposing the business domain into independent services, microservices enable business systems to strategically align with business capabilities. They embrace feature releases with high confidence. They allow for updating small areas of a live application without downtime. They empower the business to adapt rapidly to evolving market and competitive forces.

While microservice adoption is rapidly increasing, not everybody is adopting it correctly. Implementing a microservices architecture requires substantial investment. For it to be successful, developers and architects must understand the principles, patterns, and best practices of distributed microservices architecture.

This book *spoon-feeds* microservices to you. From design to development and from deployment to operation, it provides a comprehensive education into the world of microservices. Moreover, the book wraps microservice construction into

the context of cloud native design and open-source technologies, all hosted in the Azure cloud. Kudos to authors Ovais Mehboob Ahmed Khan and Arvind Chandaka for making this reference available. Kudos to you, the reader, for investing the time to study and master the architecture.

—Robert Vettor
Principal Cloud Solution Architect
Microsoft

Introduction

Welcome to *Developing Microservices Architecture on Microsoft Azure with Open Source Technologies*. Today, organizations are modernizing application development by integrating open-source technologies into a holistic architecture for delivering high-quality workloads to the cloud. This book is a complete, step-by-step guide to building an application based on microservices architecture by leveraging services provided by the Microsoft Azure cloud platform and a litany of open-source technologies such as Java, Node.js, .NET Core, and Angular.

This will be a reference guide to learn key building blocks of microservices architecture by representing a real-life case study of building a trading system to auction items. This book will lead readers through a step-by-step journey to learn modern application development scenarios and leverage cloud-native capabilities in Microsoft Azure to provide distinct value to customers.

This topic is in very high demand based on statistics and our own experiences from the field. There have been books made on microservices, but they almost always cover business value proposition alone rather than technical implementation on how to establish end to end infrastructure, application development, deployment automation and to realize actual value through constructing it using Azure and open source.

We start with an introduction and modeling of microservices and then move on to build a cloud-native microservices architecture by developing services, containerizing services, and deploying them on AKS. We also cover various topics related to communication patterns, security, and API gateways, and we touch on topics related to automating builds and deployments using CI/CD and monitoring.

Who is this book for?

This book is for architects and developers. It especially targets those who have some development background to build applications or have worked in the capacity of designing and architecting applications. It will also be useful for those with some skills with Azure.

How is this book organized?

This book is organized into three parts:

- Part I: Introduction and modeling of microservices
- Part II: Designing and building microservices architecture
- Part III: Implementing patterns, security, DevOps, and monitoring

The first part of this book consists of two chapters that focus on a thorough introduction to microservices and how you can decompose or model microservices using domain-driven design (DDD).

The second part of this book consists of three chapters that focus on designing an architecture and building an application that follows this architecture with various open-source technologies such as Java, .NET Core, Node.js, and Angular. The architecture is cloud-native and leverages multiple managed services within Azure.

The third part of this book consists of five chapters that focuses on topics related to communication patterns, security within microservices, API gateways, automating builds and deployments, and monitoring.

System requirements

The requirements for running through the examples in this book should be:

- Any Linux distros or Windows operating systems
- 4 GB of RAM
- 1 GB of available disk space
- Docker with Linux Containers
- Microsoft Azure subscription
- Azure DevOps account
- Visual Studio Code or other editors of your choice

About the companion content

The companion content for this book can be downloaded from the following page: MicrosoftPressStore.com/DevMicroservices/downloads

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

MicrosoftPressStore.com/DevMicroservices/errata

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit
MicrosoftPressStore.com/Support.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

Stay in touch

Let's keep the conversation going! We're on Twitter:

<http://twitter.com/MicrosoftPress>.

This page intentionally left blank

Introduction to microservices

In this chapter, you will:

- Learn about the history and evolution of software architecture
- Understand the differences between monolithic and microservices architectures
- Learn about the core benefits and challenges of using a microservices architecture

Microservices is a buzzword that is thrown around frequently in the technology community. Many infer it as being a further dissection of modular componentization, and although that addresses part of it, it doesn't encompass the full picture. Microservices are a group of back-end services that provide business operations to form an application. When combined with other pieces such as a front-end and various communication platforms, it creates what is known as a *microservices architecture*. Although that is a simplistic way to approach this architecture, it is a vital design orientation and metaphor to learn for implementing modern application development scenarios geared to achieving better results.

This particular chapter will focus on the core fundamentals of microservices, including the transition from historical architecture to present, benefits and challenges, and a comparison of a monolithic application architecture with microservices architecture.

Our journey with microservices

In recent years, microservices architecture have become very popular with firms. Many organizations and enterprises have expressed a desire and need to move their applications to a microservices architecture.

We have worked with customers that were very concerned with increasing development velocity to ship code faster and at more frequent intervals. Other customers are focused on modularizing their codebase further to iterate on particular services for better application performance and additional support. Some are jaded by financial and time constraints that constantly plague their organization due to large scale migrations or refactoring and want to put an end to it. We have seen these pains, frustrations, complex scenarios, and wishes from a myriad of customers time and time again.

We have worked at Microsoft and in the IT space with many enterprise customers. We met with various internal open-source groups to find solutions to these problems of customers. Microservices architecture was the key behind this, and ultimately, we wanted to have a streamlined way to demonstrate the theoretical value of microservices and how to build one, too.

As a result, we created an application called the Online Auction System (OAS) to serve as a real-life case study to teach and empower customers to build this architecture for their enterprise scenarios. We accomplished this by using a variety of open-source technologies, and we used Microsoft Azure as the bedrock for the application. Throughout this book, we will use the approach to provide this background knowledge in conjunction with the parallels to this case study, which helps map theoretical and experiential learning together.

Evolution of software architecture

Before we can understand the details behind microservices architecture, let's take a brief look at how software architectures have evolved.

Monolithic architecture

About a decade ago, the monolithic application was very popular, and consequently its N-tiered architecture backbone was a norm. This consisted of an application decomposed into several layers. These layers vary depending on the complexity of the application, but traditionally, there were three main components. See Figure 1-1.

Monolithic architecture




FIGURE 1-1 The monolithic architecture model

- **Presentation layer** The presentation layer traditionally represents the front-end. Some of the artifacts include view models, view pages, user interfaces, and other artifacts that are oriented toward front-end development.

- **Business layer** The middle layer is a business layer that holds all the logical flows of the application. This includes most of what we consider as the back-end code necessary for the primary function of the overall application. This area is where languages such as .NET or even Java are used to build back-end services that contain business managers and business objects consisting of the application logic.
- **Data access layer** Finally, the data access layer is responsible for all database operations. This consists of a database and interface to perform read/write operations on data collected and used within the application as part of its overall functionality.

This layered architecture provides a basic separation of concerns and decomposition of applications into functional components. However, this architecture results in tightly coupled services that result in many dependencies.

Thinking of some of the scenarios from earlier, we can see why tightly coupled services and dependencies cause problems. Dependencies present drawbacks where it is difficult to make any changes to the application easily. This includes complexity in the form of large migrations and refactoring projects where many teams are involved and many times agility is replaced with a potential for bureaucracy. This introduces longer timelines because simple quick deployments cannot be made and ultimately all these problems cascade into a single point where management is then required to make massive expenditures.

Service-oriented architecture (SOA)

The transition away from monolithic architecture was driven by these reasons. The heavyweight nature of this architecture was slowly decomposed and eventually led to the formation of a new architecture known as service-oriented architecture (SOA), which is illustrated in Figure 1-2.




FIGURE 1-2 SOA architecture

This model was geared toward decomposing an application into even more granular modules to account for the drawbacks in monolithic architectures. Essentially, SOA was a way to connect different monoliths across a consistent messaging channel. Thus, as you can see in Figure 1-2, with SOA, you can have a set of services that are serving various functions connect and communicate with each other over a service bus. This service bus essentially is the communication tool that assists with message routing, transformations, security improvement, logging, and much more.

Because the services are broken down by function, it is slightly easier to deploy changes that don't affect the whole application. So, if you were building an application with a monolithic architecture, you could refactor it by using SOA architecture and addressing the litany of cascading problems we had with the monolithic architecture. Although SOA addresses some of these problems, we haven't solved a primary pain point in either architecture model—assuaging application scalability concerns.

Application scalability is an important design consideration when thinking about the growth of the application over time. With a larger user base, you will likely have more complex requirements, and thus considerations must be in place for handling additional requests and traffic. The clunky nature of the monolithic architecture doesn't address this and neither does SOA. Ultimately, with more scale, SOA breaks down internally because the single enterprise bus is overwhelmed, which slows (or throttles) during periods of increased requests and traffic.

That being said, the SOA concept ultimately introduced an important integration solution that helped applications talk to each other, and the additional consideration of scalability gave way to the expansion of the microservices architecture.

Comparing monolith with microservices

A microservices architecture is a series of services that communicate with one another to achieve complete application functionality. Each microservice has a similar pattern to a monolithic architecture but only covers one function within the entire application functionality.

This means that each service can have a data access layer that is customized to use several kinds of database technologies and a business layer that is customized to use wide varieties of technology to perform the functional aspects of the microservice. To elaborate on the business layer using our case study detailed in future chapters, this means that instead of having all the OAS back-end code in a single layer, we have individual services for setting up an auction or placing a bid.

The microservices architecture iterates on previous models with more communication and an all-encompassing front-end functionality. All the back-end services now communicate with each other to represent actions that users can experience through a single point of entry within the front-end, as illustrated in Figure 1-3.




FIGURE 1-3 Microservices architecture

SOA versus microservices

You might also wonder how SOA differs from microservices. Microservices are fine-grained services where each service is modeled for a business domain or capability. SOA services are usually coarse-grained services that are not completely decomposed based on the business domain. Microservices are more abstracted and segmented by the individual business functionality. We will cover this in detail in future chapters, but the use of domain-driven design is a major design factor in microservices that is distinct from other architectures. Domain-driven design specifically aligns services to individual business functions in an application.

Also, microservices have more complex patterns when it comes to communication and messaging. Where we once had an enterprise bus as one of the innovative features of an SOA, we now have a broader set of technologies and mechanisms, such as request-response and publish-subscribe communications. These patterns will help distribute communication channels for better performance in an application, rather than making the service bus the bottleneck of the system.

Another advantage of microservices is rooted in using lightweight protocols, such as HTTP, whereas with SOA, we use other protocols such as TCP or MSMQ. There are many examples in which microservices are the presumptive choice for your application architectures. To help illustrate this better, we will dive deeper into some of the issues with legacy architectures.

Monolith example

Now that we have talked about the high-level differences, let's run through an exercise to understand the minute differentiations between monolithic and microservices architectures. To help illustrate this, see Figure 1-4, which illustrates how the monolithic architecture requires a full application deployed in each host, making it difficult to scale.

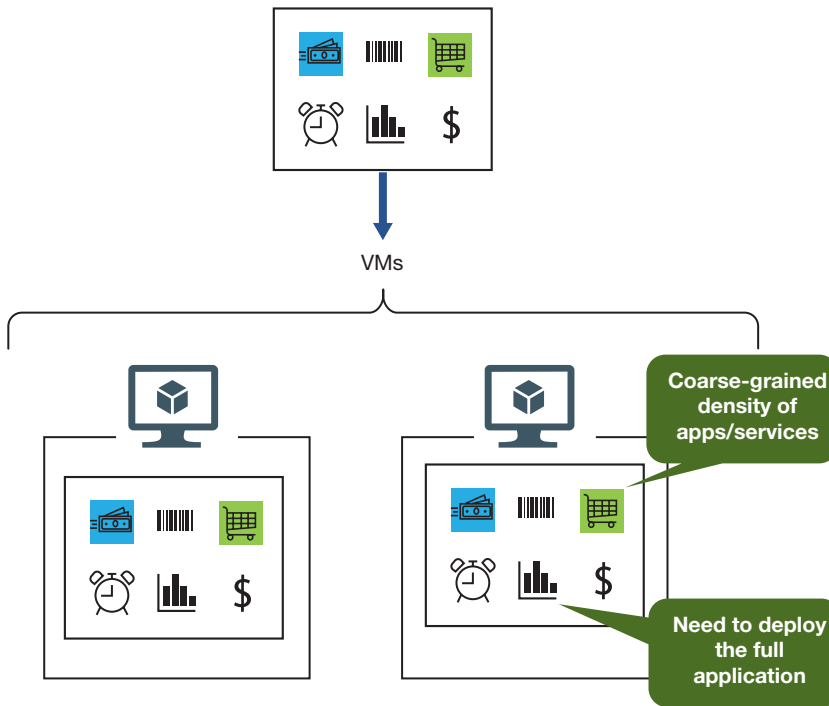


FIGURE 1-4 The monolithic architecture is difficult to scale

Figure 1-4 helps clarify the nuances of a monolithic application. Here, we have a collection of services that are ultimately deployed on two virtual machines (VMs)—all containing the same services within each host. As you can see, the scaling solution in place essentially deploys the full application on multiple VMs.

NOTE We know the monolithic architecture is scalable, however, it is expensive. We need to have the prerequisite specifications for these VMs with the same memory, compute power, and so on. Even though the deployment of the monolithic architecture might be simple because it is consistent, this can be expensive, whether you are utilizing your on-premises datacenter for the infrastructure or you are using a cloud provider with a pay-as-you-go subscription.

Microservices example

Now let's compare the monolithic architecture with a microservices approach. Figure 1-5 illustrates how the microservices architecture can be deployed independently in several different hosts and is more capable of scaling.

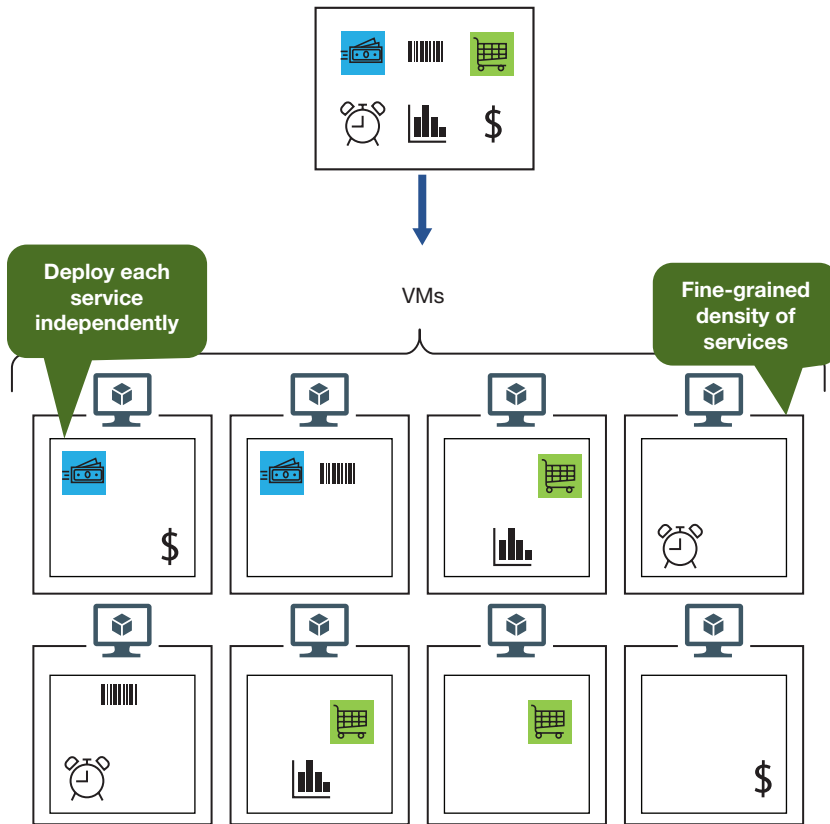


FIGURE 1-5 Microservices architecture

Figure 1-5 shows how microservices architecture differs from the monolithic approach. Virtual machines can be deployed in the microservices architecture, but instead of deploying all components of the application in a single VM, we see that each service is deployed independently on several VMs. Furthermore, with a monolithic application, any deployment that might have issues or bugs would have to be rolled back and analyzed carefully for mistakes. Not only is this tedious, but it provides a bad user experience for your app's consumers because of bugs and unexpected downtime. This allows us to scale based on functional need, doesn't result in overexpenditure of your resources, and enables a more fine-grained control of your services. If you experience deployment issues, you can easily find errors and redeploy with flexibility and velocity.

To understand how microservices promote scalability, suppose we are running an order management system in which users can place orders. On weekends, the system load is increased by a high number of user orders, which causes response timeouts and improper handling of resources. With the appropriate monitoring framework in place, we can identify when the order service is not handling a larger amount of requests. Based on this, we can just scale out the order service, rather than scaling out the whole system. This allows us to save our resources inexpensively.

NOTE Although we are demonstrating some of these basic concepts with microservices spread across VMs, a more regimented enterprise approach will use infrastructure such as Kubernetes, containers, and continuous integration and continuous development (CI/CD) as components to bolster some of the microservices benefits.

Databases in a monolithic architecture

Let's focus more on the data tier of the architecture. Figure 1-6 shows that a monolithic database might not be a good design choice when it is shared across many services.

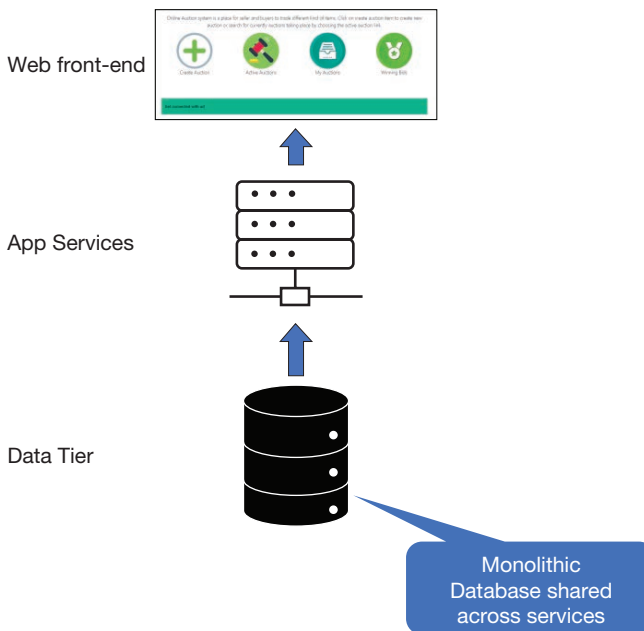


FIGURE 1-6 Monolithic database when shared across many APIs

As you can see, the data access layer in Figure 1-6 utilizes a single monolithic database that is shared across multiple services. This database also tends to be a relational database to store information, and all the app services that are operating in the middle tier perform actions directly on this single data tier.

If your application continues to grow, you might start noting issues with the scaling of your database and with maintaining this data. All the different services are recording particular information relevant to their service's operations, and it becomes very difficult to segregate this within your database.

Databases in microservices architecture

The microservices perspective is shown in Figure 1-7.




FIGURE 1-7 The microservices architecture

In microservices, you design with the domain or business functionality in mind. This means that you can segregate the services by business functionality, so each service can have its own particular database. Furthermore, you can specifically pick the technology behind the database of your choosing based on the need of the service. This provides even more customization, and it is a great benefit to leverage. Granted, when we abstract these services even more, we face difficulties with database consistency. Thankfully, as we mentioned earlier, there are many patterns in place to address that concern, which shows how microservices can truly make a difference in your application architecture.

Micro front-ends

When discussing the architectural features of microservices, we should also be aware of new movements and trends that continue to positively affect this paradigm, such as micro front-ends.

A micro front-end is similar to having various back-end microservices because the front-end is segmented along the same verticals. In the future, we will see more micro front-ends as we continue to evolve software architecture (see Figure 1-8).




FIGURE 1-8 Microservices architecture versus micro front-end architecture

Figure 1-8 shows the current evolution of microservices architecture. When it comes to micro front-ends, we see that the front-end is further abstracted and connected to the back-end microservices over the same API gateway. This means that the UI/UX components of each of the services are segmented to those applications only. Then, all these separate components are placed together in a container app that serves as a placeholder for these HTML components with embedded data. When deployed, the components are rendered into a single front-end that is seen when the user accesses the overall application, say, through a web browser (assuming it was a web app). In traditional microservices, data components are received as JSON payloads, but in micro front-ends, data components are received as HTML components.

This componentization has similar benefits to those discussed in the next section, but this example helps illustrate its usefulness.

Let's take three teams—named Service 1, Service 2, and Service 3—operating in distinct service verticals. This means the Service 1 team can build and deploy the Service 1 full stack microservice, meaning they can build and deploy its respective data layer, business layer, and front-end. Likewise, the Service 2 and Service 3 teams can do the same, without any dependency on the other verticals.

Many of the benefits and challenges of micro front-end are strongly correlated to their respective back-end microservices. We will cover this in more detail in the next section. In this book, we use the OAS as a way to illustrate many of these concepts in practice, though we did

not use this particular feature in our case study. We still want you to be aware of micro front-ends because you will likely see their underlying patterns time and time again. Different layers and components of application architecture will always be further abstracted in order to create simplified, yet richer, scenarios for developers, engineers, and more. This is something that will always grow in popularity, and is a trend to observe and track as we continue evolving our application design strategies.

Core fundamentals of microservices

A fundamental understanding of microservices is vital to understanding the examples used throughout this book, as well as the real-life case study.

First, it's important to know that microservices are autonomous, independent, and loosely coupled services that cover business scenarios and collaborate with each other to form an application. As we know, each service will have the back-end code needed to perform a particular operation. In addition, each service will most likely have an endpoint to communicate with other services for entire application functionality. Finally, these services will also have a data access layer for maintaining service-specific data in individual databases.

Benefits

We have seen some of the scenarios and understood some of the pains that enterprise architects, developers, and consultants face when it comes to application architecture. Based on our experiences, we divided these benefits into three different categories:

- Agility
- Scalability
- Technology diversity

Agility

As shown in Figure 1-9, microservices are representative of individual functions within an overall application, and as such, business functionality being encapsulated into small targeted services is one of the major points of agility. In a traditionally monolithic application, there are three layers: presentation layer, business layer, and the data access layer. In this form of architecture, all the business logic is packed together in a single layer. With microservices, this is abstracted so that each microservice is representative of one component of the overall business logic, and a single team can be responsible for building that out. In fact, this attribute alone is considered one of the major hallmarks of microservices.

This is a method to ensure logical separation of concerns. When individual teams are aligned to microservices, this tends to lend itself to more agile development methodologies and ultimately, sections of the app can be changed without regard to issues stemming from tightly coupled functions.

The second benefit in agility is that each service can evolve independently and can be deployed frequently. This is of course with the assumption that the service endpoints are well defined and agreed upon by the teams building and consuming the services. In the team dynamic setting mentioned earlier in this chapter, we noted that each service can be built out by different teams because the overall business logic of the application is further abstracted. Additionally, there isn't dependency on the deployment side either, and separate continuous integration and deployment pipelines can be set up to make changes to each team's respective service. Ultimately, one team's work on a particular service will not affect another team's work. This independence ultimately reduces the mean time to delivery from a time-to-market perspective.

This results in development teams being more agile and being able to produce at a high velocity. Ultimately, they can focus better on their individual services. Development teams are now abstracted from the internal intricacies present when working with other teams from both a technological and process standpoint, which means they can be more productive overall.

Scalability

Microservices are very scalable in relation to an application. They can scale on respective demands without affecting overall performance. This occurs because of the abstraction of business functions combined with components of the overall architecture that build out the ability to handle more traffic and requests, and it has high availability. In fact, scalability was the feature that drove the transition from SOA to microservices.

We touched on some of these components earlier. Granular services mean that requests are oriented to particular areas within the architecture. However, when we have services communicating with others for an end-to-end scenario, communication patterns and messaging channels between individual services help streamline logic to complete scenarios. For microservices architecture, we have individual communications and patterns such as publish-subscribe which makes communication easier and doesn't throttle the performance of the application.

Also, microservices tend to be containerized in enterprise scenarios and thus can be deployed efficiently and frequently with the use of infrastructure to containerize the services (like Docker), deploy with scale with Kubernetes, and be able to deploy quickly with pipelining and continuous integration and deployment.

Technology Diversity

Often, balance and flexibility are overlooked in developer environments, which gives way to some rigid architecture that has specific problems and thus requires specific answers.

Developers who work on applications that utilize microservices architectures can mix and match different programming platforms and storage technologies. This offers great flexibility for developers, architects, and consultants to design with almost limitless permutations for a final business solution. This flexibility to leverage numerous technologies is an added benefit that we see time and time again in enterprise environments.

Consequently, applications can be refactored, modernized, or experimented on with new languages and practices, one service at a time. However, the introduction of new technology into longstanding enterprise environments tends to create new problems and headaches for incumbent IT professionals. This signals the start of large-scale refactoring and modernization projects, which equate to massive expenditures in time and money. Despite this, microservices change the way we think about these initiatives and instead of generating more problems, provides solutions for IT managers.

To understand this better, suppose we started with a monolithic application a decade ago to develop a product. In the last 10 years, we find that the application code base has become huge, and the database has become monolithic, too. Today, we realized that the technology we used earlier is getting outdated and making modifications to the existing application is difficult as one change will potentially break other pieces of the monolithic application. We want to provide a modern user experience to the consumer with new technologies and applications. This is where microservices solves these problems. With microservices, because the application is segregated into set of different services, the front-end could be a micro front-end or even a basic terminal that contains some HTML pages and consumes those services to represent application features. The segregation allows teams to independently make improvements to their service with minimal impact to the rest of the application. Adding new functionality is not only easy, but the adoption of cutting-edge technologies are faster and complementary.

Challenges

We've seen some of the fantastic benefits of using microservices architecture, many of which have themes related to componentization and flexibility. On the other hand, many of these same benefits pose new challenges for your IT professionals as well.

These challenges include:

- Learning curve
- Deployment
- Interaction
- Monitoring

Learning curve

In the modern enterprise environment, we have seen the shift from a traditional team setting to a DevOps-oriented team. A traditional company might have individualized teams of developers, QA engineers, an infrastructure team, an information security team, an operations team, and database admins—all in individual silos that have to work together in a waterfall-oriented project.

Today, agile projects are the norm and teams in a DevOps setting are almost mandatory, which means that teams are composed of a smattering of people from different traditional teams to make a completely iterative team that is capable of building out and maintaining some function.

However, this means that everyone on that agile team that is assigned to building out a particular microservice must ramp up based on the different kinds of technology that are selected for the final design.

The number of technologies here poses a problem, where more options cause decision fatigue and might exacerbate the learning curve. From the developer perspective, a design choice could be made to use Java for the business logic in a particular microservice. If this isn't the primary language that developers are used to building on, then ramping up to learn the language from a syntax, semantic, and even a knowledge perspective (considering libraries available) could be a significant barrier, which can pose both time and fiscal challenges when training time and resources are taken into consideration.

This isn't a situation that is unique to developers; it can happen to any role within your agile team. Let's take an infrastructure team who is responsible for deploying the underlying components needed to build out the application. The decision on what tool to use for infrastructure-as-code (IaC)—such as Terraform, ARM Templates, and so on—and for pipelines (build and deployments, i.e. Jenkin, Azure DevOps etc.) could cause delays and expenditures out of budget for training.

Building expertise on many different kinds of technology is difficult for teams in general, and it is even more dynamic when working with the microservices architecture. That being said, good planning and an understanding of your team's strengths and weaknesses can enable you to make the right decisions to balance the time and money to make the team productive when using microservices, despite these challenges.

Deployment

Agility is one of the big benefits of microservices, particularly because builds and deployments are faster and more frequent because teams own smaller parts of an overall application. However, this bonus comes with difficulty, where one needs a strong focus and expertise in automation for faster and frequent deployments.

A manual process to address builds and deployments would be cumbersome and painstaking. Thus, automation is required to reduce this overhead. The problem is deeper here because automation of pipelines in continuous integration and continuous deployment requires specialized technologies and skills, and automated pipelines are complex to set up based on how customized your environment is. Again, this could hinder the potential pros of using microservices.

Let's understand this better by walking through a scenario. With the assumption that the code base for your service is already complete, we need to ascertain where we will be storing the code. Traditionally, tools such as GitHub provide repos and source code version control, which is useful for developers to push code, generate builds, or make improvements to the service. Primarily, we want to focus on the integration of these pushes to the repo in kicking off the workflow for a build. There are many ways to do this with your tools, but then there is a need to create the build using the artifacts created from the code and running through tests

to ensure the build is working before moving it over for automatic deployment to your various environments. Each layer and step of this is complex and requires a deep understanding of multiple roles to work. This is why we see there is a high demand for DevOps and automation engineers. Furthermore, the tools that you could use for this end-to-end workflow can all vary in capability and have their own associated learning curve.

That being said, many enterprise IT professionals are observing and reacting to this need by seeking out people to help accomplish this in their environments based on incumbent talent, hired talent, and consultants they might be bringing in for these particular tasks. Microservices aside, to be successful in an agile IT world, we see this as being a scenario you cannot ignore.

Interaction

Adding on to some of the knowledge base needed to work in microservices are patterns and messaging channels. In a microservices architecture, although services are disparate and independent, the overall application functionality requires services to talk and communicate with one another for triggering of actions and functions to conceive an end-to-end cohesive workflow. Therefore, knowing messaging patterns and related technology is vital. In a traditional monolithic application, all the layers are stacked on top of each other, allowing interfacing between all parts of the application. In microservices, each function is a miniaturized monolithic application, but now it needs to communicate with several others to complete an end-to-end scenario and provide a rich user scenario/experience. Ultimately, seamless communication and independent scaling can only be achieved if services can asynchronously communicate with one another. This gave rise to technologies such as service buses and queues to help delineate communication between services.

Let's walk through an example that will reappear in detail in Chapter 6, which is how we implemented interaction with our OAS case study. The goal of this application is to provide a one-stop shop for people to auction off items, bid for items, pay once a bid is won, and subsequently receive the items they've purchased. This is a simplification of the scenario, but as you can see, with the moving parts of an application in place, several systems must talk to each other. We can see that the auction and bid services must communicate with each other, the front-end would have to communicate to all the back-end services, and the payment service must communicate with the bid services.

A service-oriented architecture has a single enterprise bus and doesn't account for all of these scenarios. Thus, a more loosely coupled architecture will require you to learn more patterns (such as request/response and pub/sub) and messaging technologies (such as Apache Kafka and RabbitMQ) to be productive. Again, this is a burden on the developer and infrastructure teams.

However, there are standard ways of learning this information, so although there is a learning curve, the information you gather will be useful in understanding how to design further architecture to handle various interactions.

Monitoring

Monitoring is a core component of many developers' day-to-day jobs. Monitoring is vital for understanding the performance of your particular application and helping identify bottlenecks, and it is used for overall troubleshooting. This is something already ingrained in the mindset of most developers. However, monitoring gets more difficult and is a strict requirement when it comes to microservices.

NOTE It is vital to log each and every detail relevant to these services, and you must be able to visualize these results. This involves the use of potential log ingestion/collection engines, visualization tools, alerting systems, and much more. Again, this overhead needs to be considered when shifting to this paradigm.

There are several choices when it comes to monitoring, but we have found certain technologies and frameworks that work particularly well for microservices. One of the most useful things to do would be to understand the native features available from the cloud provider you are using. Today, it is common for many businesses to build modern, cloud-native applications. Because of this, there are likely cloud architects on teams that are experts on the monitoring capabilities native to your cloud platform. It is extremely valuable to leverage their insights and skills to building this piece of the puzzle.

In the OAS, we leveraged Azure and so our application metrics and performance are recorded by tools such as Azure Monitor, Log Analytics, Application Insights, and much more. Whatever your platform might be, the key to addressing a potential challenge here is being informed and intelligent about the different options available to you, whether it be cloud-native, third-party tools, or open-source tools, to make good decisions for your monitoring frameworks.

Moving forward with the microservices architecture, open source, and Azure

We hope that this provides a level set for readers to understand microservices architecture as we go deeper into our case study and the decisions we've had to make with Open Source Software (OSS) tooling and using Microsoft Azure.

Moving forward, we will go into greater depth in this book, describing different concepts in detail that we've touched on in this chapter. Our case study—the OAS shown in Figure 1-9—will be the context for understanding and using these concepts as we go through this journey together.

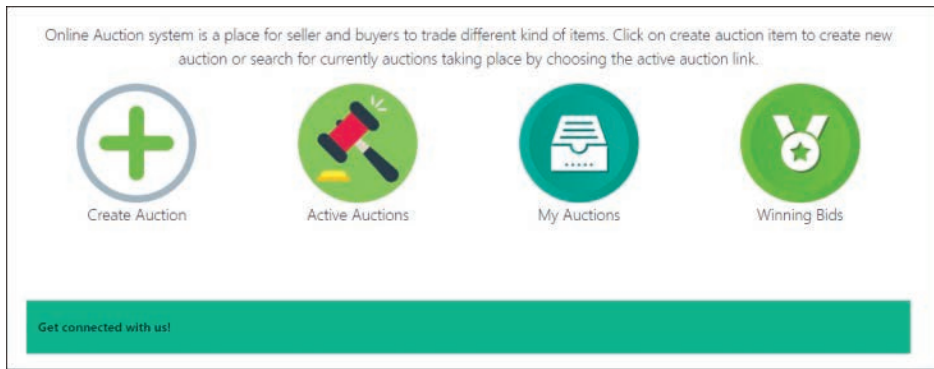


FIGURE 1-9 The Online Auction System is our real-life case study to help you understand how to use microservices architecture to build out your own modern, cloud-native applications

This book's goals

Ultimately, we hope this book will arm you—whether you are a developer, architect, consultant, or simply someone who wants to dive deeper into this topic—with the information you need to better understand the microservices architecture.

After reading this book, we hope you are sufficiently prepared to take on the challenge of building this out in a cloud-native setting such as Azure to solve your business need—whether it be through modernizing your application, migrating your application workloads, changing management, and so on. We hope that the world of microservices and open source can be the catalyst for making this a reality for you. Let's dive in!

Summary

In this chapter we covered:

- The history and evolution of software architecture from monolith all the way to microservices
- We compared different architecture models to help you understand the differences and basic inner workings of each
- We reviewed the benefits and challenges of microservices to understand the trade-offs when it comes to using this architecture

In the next chapter, we will elaborate on our real-life online auction system (OAS) case study. We will start by describing the base scenario and business requirements of the OAS and will then transition to the feature breakdown. We will discuss some decomposition principles and strategies with an emphasis on domain-driven design to demonstrate how we model our microservices.

This page intentionally left blank

Index

A

- ACID (atomicity, consistency, integrity, durability), 50
- ACR (Azure Container Registry)
 - provisioning, 129
 - pushing services to, 129–130
- ADO (Azure DevOps), 49
- aggregator pattern, 53–54
- agile, 13
- AKS (Azure Kubernetes Services), 1, 48–49
 - clusters, 42–43
 - configuring with APIM (Azure API Management), 218–220
 - creating a deployment object for OAS microservices, 131–135
 - creating a service object for OAS microservices, 135–137
 - creating a YAML file, 131–133
 - deploying images to, 125
 - deploying services to, 130–131
 - ingress controller-accessible services, 220
 - Internet-accessible services, 218–219
 - intranet-accessible services, 219
 - provisioning, 127–128
 - using with APIM, 214–215
 - ClusterIP configuration, 216–217
 - External Name configuration, 217–218
 - LoadBalancer configuration, 216
 - NodePort configuration, 215
- AMQP (Advanced Messaging Queuing Protocol), 146
- Angular, 43, 89
 - components, 94
 - creating front-end applications, 89–90
 - environment files, configuring, 101–102
 - modules, 94
 - onlineauctionwebapp, 262–264
 - project structure, 90–92
 - routing, 95
 - services, 95
 - templates, 95
- anti-patterns, 29, 33. *See also* patterns
 - Direct HTTP call, 52–53
 - establishing tight dependencies between code artifacts, 34
 - unnecessary fine graining of services to deeper subdomains, 33–337
 - using monolith or a shared database with microservices, 33
- Apache Kafka, 153–154
 - adding Kafka support in the Java application, 159–161
 - BidHostedService class, 164–165
 - building the HTTP client, 171–172
 - establishing pub/sub communication, 157
 - Event Hubs, 155–156
 - infrastructure setup, 157–159
 - KafkaHostedService, 264–265
 - listener service with .NET Core hosted service, 161–164, 165–166
 - resiliency, 165–171
- APIs, 47–48, 82, 138
 - auction service, 66–68
 - configuring with APIM (Azure API Management), 204–210
 - creating a bid service using JavaSpring Boot framework, 72–74
 - gateways, need for, 201–202
 - methods used in bid service, 74
- application development, DDD (domain-driven design), 27
 - bounded context, 28, 29–30
 - core domains, 30
 - and decomposition, 31–32
 - supporting domains, 30–31
 - ubiquitous language, 27
- Application Insights, 254–257
 - auction service configuration, 259–260
 - bid service configuration, 260–261

- configuring, 258–259
- KafkaHostedService
 - instance, 264–265
- onlineauctionwebapp
 - instance, 262–264
- payment service
 - configuration, 261–262
- applications. *See also* back-end services; cloud-native applications; containers; front-end applications
 - backend, 59
 - cloud-native, 36
 - front-end, 59
 - monitoring, 249–250
 - scalability, 4, 6
 - twelve-factor methodology, 40–41
- architecture
 - Kubernetes, 125–127
 - polyglot persistence, 51–52
 - and security, 175–176
 - zero-trust, 179–180
- ARM (Azure Resource Manager), 177
- asynchronous communication, 146–147. *See also* Apache Kafka
 - pub/sub, 148
 - Apache Kafka, 153–154, 157–164
 - Azure Event Hubs, 155
 - RabbitMQ, 156
- auction form, developing, 102–105
- auction service, 60–61
 - auction table schema, 64–65
 - configuring Application Insights, 259–260
 - containerizing, 118–120
 - creating a database, 64–66
 - creating in Node.JS, 66
 - Dockerfile*, 118–119
 - pipelines
 - build, 239–241

- deployment, 241–246
 - provisioning MySQL
 - database in Azure, 61–64
 - writing API methods, 66–68
- authentication. *See also* authorization flows; security
 - Azure Active Directory, 180–181
 - identity management, 180
 - MSAL (Microsoft Authentication Library), 192
 - multifactor, 182
 - OAS (Online Auction Service), 20
 - OAuth 2.0 authorization framework, 182–183
 - security module, developing, 95–101
- authorization flows
 - client credentials grant, 185–186
 - code grant, 183–184
 - device code, 186–187
 - implicit grant, 184–185
 - OIDC hybrid, 185
- automating infrastructure through IaC (Infrastructure as Code), 225
- availability, cloud-native applications, 38
- az login command, 129
- Azure, 2, 16, 41. *See also* ADO (Azure DevOps); AKS (Azure Kubernetes Services) clusters
 - Active Directory, 180–181
 - B2B user flow, 188
 - B2E user flow, 188
 - AD B2C, 48
 - developing a security module, 95–101
 - user flow, 189
 - APIM (API Management), 47–48, 202–203
 - benefits of using, 203

- configuring APIs, 204–210
- configuring with AKS, 218–220
- policies, 210–214
- setting up, 203–204
- using with AKS, 214–215
- App Service, 45–46, 179
 - deploying front-end applications, 137–142
- Application Insights, 49
- Container Registry, 49
- deploying services to AKS. *See also* microservices
- DevOps, 224–225
 - building a pipeline, 229–234
 - CI/CD pipeline, 246–247
 - projects, 224
- Event Hubs, 47, 155. *See also* Apache Kafka
- Key Vault, 178
- Log Analytics, 250–252
 - KQL (Kusto Query Language), 252–253
- Monitor, 50, 253–254. *See also* Application Insights
- Application Insights, 254–257
- provisioning a MySQL database, 61–64
- provisioning ACR, 129
- provisioning Cosmos DB, 70–72
- pushing services to ACR, 129–130
- Security Center, 176
- SQL, creating a payment service database, 79–81
- Storage Accounts, 178
- WebJobs, 46–47
 - deploying Kafka Listener Service as, 142
 - listener service, 171

B

back-end services, 40. *See also* auction service; back-end services; OAS (Online Auction Service); payment service

- auction service, 60–61
 - auction table schema, 64–65
 - creating a database, 64–66
 - creating in Node.JS, 66
 - provisioning MySQL database in Azure, 61–64
 - writing API methods, 66–68
- bid service, 68–70
 - API methods, 74
 - creating a bid controller, 75
 - creating a bid model, 74–75
 - creating in JavaSpring Boot framework, 72–74
 - creating the bid method, 76–77
 - database schema, 72
 - get bids using the auctionID method, 75–76
 - provisioning Cosmos DB in Azure, 70–72
- OAS (Online Auction Service), 60
- payment service, 77–79
 - create a code-first model using EF Core, 83–84
 - creating a database, 79–81
 - creating a payment service context, 85–87
 - creating an auction payment controller, 87–88
 - creating in ASP.NET Core, 81–82

- bid service, 68–70
 - API methods, 74
 - configuring Application Insights, 260–261
 - containerizing, 120–123
 - creating a bid controller, 75
 - creating a bid model, 74–75
 - creating in JavaSpring Boot framework, 72–74
 - creating the bid method, 76–77
 - database schema, 72
 - Dockerfile*, 120–121
 - get bids using the auctionID method, 75–76
 - provisioning Cosmos DB in Azure, 70–72
- bids, OAS (Online Auction Service), 21–22
- binding, 95
- bounded context, 28
 - boundaries, 29
 - communication between contexts, 29
 - reusing code between contexts, 29–30
- business layer, 10

C

- characteristics, cloud-native applications, 37
 - availability and scalability, 38
 - cost-effective solutions, 38
 - easy provisioning and maintenance, 38
 - resiliency and reliability, 38
- CI/CD (continuous integration/continuous deployment), 224
 - best practices, 237
 - pipelines, 224–225, 246–247
 - auction service build, 239–241

- auction service
 - deployment, 241–246
 - creating, 229–234
 - deploying, 234–237
- circuit breaker pattern, 166–167
- client credentials grant flow, 185–186
- cloud technologies
 - ADO (Azure DevOps), 49
 - AKS (Azure Kubernetes Services), 48–49
 - Azure AD B2C, 48
 - Azure API management, 47–48
 - Azure App Service, 45–46
 - Azure Application Insights, 49
 - Azure Container Registry, 49
 - Azure Event Hubs, 47
 - Azure Monitors, 50
 - Azure WebJobs, 46–47
- cloud-native applications, 36, 113. *See also* OAS (Online Auction Service)
 - characteristics, 37
 - availability and scalability, 38
 - cost-effective solutions, 38
 - easy provisioning and maintenance, 38
 - resiliency and reliability, 38
- IaaS (Infrastructure as a Service), 38
 - security, 176
- PaaS (Platform as a Service), 38, 137
 - app service security, 179
 - Azure Storage Accounts and key management, 178
 - database security, 179
 - security, 177–178
- principles, 36–37

CNCF (Cloud Native Computing Foundation), 36
 code grant flow, 183–184
 collecting log information, 250
 commands
 az login, 129
 Docker, 116–117
 docker build, 119–120
 docker run, 120
 kubectrl, 125, 130–131, 135
 kubectrl apply, 125
 mwnw, 122
 communication, 5, 145.
 See also asynchronous
 communication; synchronous
 communication
 approaches to, 145
 best approach for
 microservices, 149–153
 listener services, 165
 in monolithic architecture,
 145
 pub/sub, 148
 Azure Event Hubs, 155
 RabbitMQ, 156
 request/response, 147, 149
 synchronous vs.
 asynchronous, 146–147
 components
 Angular, 94
 Docker, 116
 configuration
 APIs, 204–210
 Application Insights, 258–259
 auction service, 259–260
 bid service, 260–261
 payment service, 261–262
 containerizing
 auction service, 118–120
 bid service, 120–123
 payment service, 123–124
 containers, 37, 49, 114
 Docker, 114
 commands, 116–117
 components, 116

Hyper-V, 118
 images, 113
 deploying to AKS, 125
 Linux vs. Windows, 117–118
 monitoring, 265–266
 and VMs, 114
 Continuous Delivery, 36
 Conway's rule, 25
 core domains, 30
 Cosmos DB, provisioning in
 Azure, 70–72
 CQRS (Command Query
 Responsibility Segregation),
 64, 152
 CQRS (Command Query
 Responsibility Segregation)
 pattern, 54–55
 create auction form, 102–105
 CRON expressions, 46

D

dashboards, 266–268
 data access layer, 4, 10, 11
 databases
 Cosmos DB, provisioning in
 Azure, 70–72
 creating for auction service,
 64–66
 creating for payment service,
 79–81
 distributed architecture, 50
 in microservices architecture,
 9
 in monolithic architecture,
 8–9
 MySQL
 creating an auction table,
 65–66
 provisioning, 61–64
 polyglot persistence, 51–52
 security, 179
 transactional data model, 50
 transient data model, 50–51

DDD (domain-driven design), 5,
 26, 27
 bounded context, 28
 boundaries, 29
 communication between
 contexts, 29
 reusing code between
 contexts, 29–30
 core domains, 30
 and decomposition, 31–32
 supporting domains, 30–31
 ubiquitous language, 27
 decomposition
 based on DDD, 31–32
 by business capability, 25
 Conway's rule, 25
 Inverse Conway Maneuver
 Law, 25–26
 principles, 24
 by subdomain, 25–26
 dependencies, 40
 deployment
 challenges in microservices
 architecture, 14–15
 patterns, 238–239
 device code flow, 186–187
 DevOps, 13, 14, 36, 224
 Direct HTTP call, 52–53
 disposability, 40
 distributed database
 architecture, 50
 patterns
 aggregator, 53–54
 CQRS (Command
 Query Responsibility
 Segregation), 54–55
 Direct HTTP call, 52–53
 saga, 56
 Docker, 114
 auction service,
 containerizing, 118–120
 bid service, containerizing,
 120–123
 commands, 116–117
 components, 116

- installing, 114–116
- payment service,
 - containerizing, 123–124
- docker build command, 119–120
- docker run command, 120
- Dockerfile*
 - auction service, 118–119
 - bid service, 120–121
 - payment service, 123–124

E

- environment files, configuration, 101–102
- event binding, 95
- Express.JS, 66. *See also* Node.JS

F

- forms
 - create auction, 102–105
 - submit bid, 107–110
- frameworks, 249–250
 - monitoring, 253, 257–258
- front-end applications, 43, 88–89. *See also* OAS (Online Auction Service)
- Angular
 - modules, 94
 - project structure, 90–92
- auction form, developing, 102–105
- configuring environment files, 101–102
- creating, 89–90
- deploying in the Azure App Service, 137–142
- developing an active auctions page, 105–107
- onlineauctionwebapp, 262–264
- prerequisites for building, 89
- security module, developing, 95–101
- submit bid form, 107–110

G

- generic domains, 30–31
- GitHub, 14
- granular services, 12

H-I

- Hyper-V, 118
- IaaS (Infrastructure as a Service), 38
 - security, 176
 - Terraforms, 61
- IaC (Infrastructure as Code)
 - automating infrastructure, 225
 - OAS (Online Auction Service), 225–229
 - security principles, 177
- IaS (infrastructure-as-code), 14
- identity management, 180, 181
 - Azure Active Directory, 180–181
 - OAS (Online Auction Service), 20
 - principle of least privilege, 181
 - RBAC (role-based access control), 181
- images, deploying to AKS, 125
- implicit grant flow, 184–185
- installing, Docker, 114–116
- Inverse Conway Maneuver Law*, 25–26

J-K

- JavaSpring Boot framework,
 - creating a bid service, 72–74
- Kafka protocol, 47
 - adding Kafka support in the Java application, 159–161
 - building the HTTP client, 171–172

- deploying Listener Service as an Azure WebJob, 142
- establishing pub/sub communication, 157
 - infrastructure setup, 157–159
- key management, 178
- KQL (Kusto Query Language), 252–253
- kubectly apply command, 125
- kubectly commands, 125, 130–131, 135
- Kubernetes. *See also* AKS (Azure Kubernetes Services) clusters; VMs (virtual machines)
 - architecture, 125–127
 - creating a deployment object for OAS microservices, 131–135

L

- lightweight protocols, 5, 36
- Linux, containers, 117–118
- listener services, 165, 171. *See also* communication
 - resiliency, 165–166
- logs, 16, 40, 250
 - Azure Log Analytics, 250–252
 - collecting, 250
 - querying, 252–253
 - retention, 252

M

- MFA (multifactor authentication), 182
- micro frontends, 9–10
- microservices, 1, 4, 36, 37, 59, 145. *See also* auction service; bid service; containers; payment service
 - communication, 5
 - best approach, 149–153
 - data access layer, 4

- distributed database
 - architecture, 50
 - lightweight protocols, 5
 - monitoring, 249
- microservices architecture, 113–114. *See also*
- communication; DDD (domain-driven design); security
- anti-patterns
 - establishing tight
 - dependencies between code artifacts, 34
 - unnecessary fine graining
 - of services to deeper subdomains, 33–337
 - using monolith or a shared database with microservices, 33
- benefits
 - agility, 11–12
 - scalability, 12
 - technology diversity, 12–13
- challenges
 - deployment, 14–15
 - interaction, 15
 - learning curve, 13–14
 - monitoring, 16
- communication
 - pub/sub, 148
 - request/response, 147, 149
 - synchronous vs.
 - asynchronous, 146–147
- comparison with monolithic architecture, 4–5, 6, 7–8
- comparison with SOA, 5
- data access layer, 11
- databases, 9
- decomposition, principles, 24
- developing, 60
- laC security principles, 177
- micro frontends, 9–10
- patterns
 - aggregator, 53–54
 - chassis, 44–45
 - circuit breaker, 166–167

- CQRS (Command Query Responsibility Segregation), 54–55
- deployment, 238–239
- Direct HTTP call, 52–53
- retry, 166, 167
- saga, 56
- scalability, 7
- VMs (virtual machines), 7
- Microsoft Azure. *See* Azure
- modules, Angular, 94
- monitoring
 - Application Insights, 254–257
 - configuring, 258–259
 - Azure Log Analytics, 250–252
 - Azure Monitor, 253–254
 - best practices, 257–258
 - containers, 265–266
 - frameworks, 253, 257–258
 - logs, 16, 40, 250
 - collecting, 250
 - querying, 252–253
 - retention, 252
 - microservices, 16
 - microservices architecture, 249
 - OAS (Online Auction Service), 42
 - telemetry
 - KafkaHostedService, 264–265
 - onlineauctionwebapp, 262–264
 - visualization tools, 250
 - dashboards, 266–268
- monolithic architecture, 2–3
- application scalability, 6
- business layer, 3
- communication, 145
- comparison with
 - microservices architecture, 4–5, 6, 7–8
- data access layer, 3
- databases, 8–9
- presentation layer, 2
- VMs (virtual machines), 6

- MSAL (Microsoft Authentication Library), 192
- mvnw command, 122
- MySQL databases
 - creating an auction table, 65–66
 - provisioning, 61–64

N

- .NET Core
 - BidHostedService class, 164–165
 - creating a payment service, 81–82
 - create a code-first model using EF Core, 83–84
 - Kafka listener service, 161–164
 - Polly framework, 167
- Node.JS, creating an auction service, 66

O

- OAS (Online Auction Service), 2, 10, 15, 16, 17, 19, 36, 48, 59, 113. *See also* DDD (domain-driven design)
- active auctions page,
 - developing, 105–107
- AKS clusters, 42–43
- application flow, 23–24
- application requirements, 20
- architecture, 41–42
- auction service, 60–61
 - auction table schema, 64–65
- build pipeline, 239–241
- configuring Application Insights, 259–260
- containerizing, 118–120
- creating a database, 64–66

- creating in Node.JS, 66
- deployment pipeline, 241–246
- Dockerfile*, 118–119
- provisioning MySQL
 - database in Azure, 61–64
- writing API methods, 66–68
- back-end services, 60
- bid service, 68–70, 161
 - API methods, 74
 - configuring Application Insights, 260–261
 - containerizing, 120–123
 - creating a bid model, 74–75
 - creating in JavaSpring
 - Boot framework, 72–74
 - creating the bid method, 76–77
 - database schema, 72
 - Dockerfile*, 120–121
 - get bids using the
 - auctionID method, 75–76
 - provisioning Cosmos DB
 - in Azure, 70–72
- creating a deployment object, 131–135
- creating service objects, 135–137
- creating the auction form, 102–105
- dashboards, 266–268
- decomposition
 - based on DDD, 31–32
 - by business capability, 25
 - Conway's rule, 25
 - Inverse Conway Maneuver Law*, 25–26
 - principles, 24
 - by subdomain, 25–26
- deployment patterns, 238–239

- end-to-end security, 189–191
- environment files,
 - configuring, 101–102
- features
 - auction management, 20–21
 - authentication, 20
 - bid management, 21–22
 - identity management, 20
 - payment management, 22–23
- front-end, prerequisites, 89
- KafkaHostedService, 264–265
- Listener service, 42
- monitoring, 42
- onlineauctionwebapp, 262–264
- payment service, 77–79
 - configuring Application Insights, 261–262
 - containerizing, 123–124
 - creating a database, 79–81
 - creating a payment
 - service context, 85–87
 - creating an auction
 - payment controller, 87–88
 - creating in ASP.NET Core, 81–82
 - Dockerfile*, 123–124
 - UI, 152
- pub/sub communication, 148
- request/response
 - communication, 147
- security
 - authentication, 193–196
 - creating a tenant, 192
 - user flows, 196–200
- security module, developing, 95–101
- services, 41–42
- submit bid form, 107–110

- Terraform script for
 - infrastructure components, 225–229
- Winning Bids table, 151–152
- workflows
 - Active Auctions*, 150–151
 - Create Auction*, 149
 - Make a Bid*, 151
- OAuth 2.0 authorization
 - framework, 182–183
- OIDC (OpenID Connect), 182
 - hybrid flow, 185
- OSS (Open Source Software), 16

P

- PaaS (Platform as a Service), 38, 137
 - security, 177–178
 - app services, 179
 - Azure Storage Accounts
 - and key management, 178
 - databases, 179
- patterns
 - aggregator, 53–54
 - circuit breaker, 166–167
 - CQRS (Command Query Responsibility Segregation), 54–55
 - deployment, 238–239
 - Direct HTTP call, 52–53
 - retry, 166, 167
 - saga, 56
- payment service, 77–79
 - configuring Application Insights, 261–262
 - containerizing, 123–124
 - creating a database, 79–81
 - creating a payment service
 - context, 85–87
 - creating an auction payment
 - controller, 87–88

- creating in ASP.NET Core, 81–82
 - code-first model, 83–84
- Dockerfile*, 123–124
- pipelines, 224–225
 - auction service
 - build, 239–241
 - deployment, 241–246
 - creating, 229–234
 - deploying, 234–237
- polyglot persistence, 51–52
- port binding, 40
- principles
 - of cloud-native applications, 36–37
- IaC security, 177
- least privilege, 181
- property binding, 95
- provisioning
 - ACR (Azure Container Registry), 129
 - AKS (Azure Kubernetes Services) clusters, 127–128
- pub/sub communication, 148
- Apache Kafka, 153–154
 - adding Kafka support in the Java application, 159–161
- BidHostedService class, 164–165
- building the HTTP client, 171–172
- infrastructure setup, 157–159
- listener service with .NET Core hosted service, 161–164, 166–171
- Azure Event Hubs, 155
- RabbitMQ, 156
- pushing services to ACR, 129–130

Q-R

- querying logs, 252–253
- RabbitMQ, 156

- RBAC (role-based access control), 181
- request/response communication, 147, 149
- resiliency, 165–171
 - circuit breaker pattern, 166–167
 - cloud-native applications, 38
 - retry pattern, 166, 167
- retry pattern, 166, 167
- Route Guard, 100–101
- routing, Angular, 95

S

- SaaS (Software as a Service), 176
- saga pattern, 56
- scalability
 - application, 4, 6
 - cloud-native applications, 38
 - elastic, 71
 - microservices architecture, 7, 12
- security, 175, 176, 178
 - and architectures, 175–176
 - authentication, 182–183
 - and authorization flows, 182–183
 - Azure Active Directory, 180–181
 - identity management, 180
 - multifactor, 182
 - OAS (Online Auction Service), 20
 - security module, developing, 95–101
- authorization flows
 - client credentials grant, 185–186
 - code grant, 183–184
 - device code, 186–187
 - implicit grant, 184–185
 - OIDC hybrid, 185
- Azure Active Directory, 180–181
- Azure Storage Accounts, 178
- IaaS (Infrastructure as a Service), 176
- IaC (Infrastructure as Code), 177
- identity management, 181
 - Azure Active Directory, 180–181
 - MFA (multifactor authentication), 182
 - OAS (Online Auction Service), 20
 - principle of least privilege, 181
 - RBAC (role-based access control), 181
 - key management, 178
 - OAS (Online Auction Service) authentication, 193–196
 - creating a tenant, 192
 - user flows, 196–200
- PaaS (Platform as a Service), 177–178
 - app services, 179
 - Azure Storage Accounts and key management, 178
 - databases, 179
 - “perimeter”, 180
 - zero-trust architecture, 179–180
 - identity management, 180
- security module, developing, 95–101
- service bus, 4, 5
- Service object, creating for OAS microservices, 135–137
- services. *See also* microservices
 - Angular, 95
 - app, security, 179
 - deploying to AKS, 130–131
- SOA (service-oriented architecture), 3–4
 - comparison with microservices architecture, 5
 - service bus, 4

software architecture. *See also*
 microservices architecture;
 monolithic architecture; SOA
 (service-oriented architecture)
 application scalability, 4
 dependencies, 3, 34, 40
 domain-driven design, 5
 microservices
 agility, 11–12
 comparison with
 monolithic architecture,
 4–5, 6, 7–8
 comparison with SOA, 5
 data access layer, 4, 11
 databases, 9
 decomposition by
 business capability, 25
 decomposition principles,
 24
 deployment challenges,
 14–15
 and interaction, 15
 learning curve, 13–14
 lightweight protocols, 5
 micro frontends, 9–10
 and monitoring, 16
 scalability, 12
 technology diversity, 12–13
 VMs (virtual machines), 7
 monolithic, 2–3
 application scalability, 6
 business layer, 3
 comparison with
 microservices
 architecture, 4–5, 6, 7–8
 data access layer, 3
 databases, 8–9
 presentation layer, 2
 monolithic architecture, VMs
 (virtual machines), 6
 SOA (service-oriented
 architecture), 3–4
 comparison with
 microservices, 5

software development
 anti-patterns, 33
 establishing tight
 dependencies between
 code artifacts, 34
 unnecessary fine graining
 of services to deeper
 subdomains, 33–337
 using monolith or a
 shared database with
 microservices, 33
 SPA (Single Page Application), 59
 submit bid form, 107–110
 supporting domains, 30
 synchronous communication,
 146–147
 request/response, 147, 149

T

technologies
 cloud, Azure App Service,
 45–46
 front-end, 43
 microservices chassis pattern,
 44–45
 templates, 231
 Angular, 95
 ARM (Azure Resource
 Manager), 177
 deployment object, 131–135
 pod object, 133–135
 Terraforms, 62–63
 Terraform(s), 61, 63–64, 71, 80
 configuring Azure App
 Service, 138–142
 installing on Windows, 62–63
 nodes, 125
 provisioning ACR, 129
 provisioning AKS, 127–128
 script for OAS infrastructure
 components, 225–229
 templates, 62–63

transactional data model, 50
 transient data model, 50–51
 twelve-factor app methodology,
 40–41

U-V

ubiquitous language, 27
 visualization tools, 250
 dashboards, 266–268
 VMs (virtual machines), 6, 37, 125
 Apache Kafka, 153–154
 and containers, 114
 in microservices architecture,
 7
 in monolithic architecture, 6
 security, 176

W

web apps, 59
 progressive, C01.1032
 WebJobs
 deploying Kafka Listener
 Service as, 142
 listener service, 171
 Windows, containers, 117–118
 workspaces, Azure Log
 Analytics, 252, 265

X-Y-Z

YAML, 131–133, 224–225, 230
 zero-trust architecture
 Azure Active Directory,
 180–181
 identity management, 180
 “pesrimeter”, 180