



ROCKY HECKMAN

# DESIGNING PLATFORM INDEPENDENT MOBILE APPS AND SERVICES

  
IEEE PRESS

IEEE  
 computer  
society

WILEY

---

# *DESIGNING PLATFORM INDEPENDENT MOBILE APPS AND SERVICES*

ROCKY HECKMAN

IEEE  
computer  
society

  
IEEE PRESS

WILEY

Cover image © gettyimages.com

Copyright © 2016 by the IEEE Computer Society, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.  
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at [www.copyright.com](http://www.copyright.com). Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

**Limit of Liability/Disclaimer of Warranty:** While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at [www.wiley.com](http://www.wiley.com).

***Library of Congress Cataloging-in-Publication Data:***

Names: Heckman, Rocky, author.

Title: Designing platform independent mobile apps and services / Rocky Heckman.

Description: Hoboken, New Jersey : John Wiley & Sons, Inc., [2016] | Includes index.

Identifiers: LCCN 2016009419 | ISBN 9781119060147 (cloth) | ISBN 9781119060185 (epub) | ISBN 9781119060154 (Adobe PDF)

Subjects: LCSH: Mobile computing. | Cell phones—Programming. | Mobile apps.

Classification: LCC QA76.59 .H43 2016 | DDC 005.25—dc23 LC record available at <https://lccn.loc.gov/2016009419>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

*Thank you to all my friends who finally convinced me to write a book. Most of all, thank you to my wife Stefanie, and my two beautiful girls Elyssia and Keiralli for not only pushing me to finish, but putting up with all the time I spent doing it. I love you all very much.*

*Look Dad, I did it!*

---

# LIST OF FIGURES

1.1	Duke	3
1.2	Breadth versus Depth	9
1.3	Excel Online <a href="https://office.live.com/start/excel.aspx">https://office.live.com/start/excel.aspx</a>	10
2.1	SOAP Message Format	28
2.2	User Satisfaction versus Response Time	41
2.3	System Unreachable	56
2.4	Hash-Based Dupe Message Detection	59
3.1	High Level App Flow	65
3.2	Five Layer Architecture	66
3.3	Service Interface Layer Routing	68
5.1	Many to One Pull Message Queue	83
5.2	Point to Point Channel	84
5.3	One to Many Publish/Subscribe Message Queue	84
5.4	One to Many Push Message Queue	85
5.5	Event Hub Offsets	89
5.6	Multiple Proprietary Push Notification Services	92
5.7	Using a SaaS PNS	95
5.8	SaaS PNS Routed Through TFS	96
5.9	Message Translation Chaining	100
5.10	Content Enricher Pattern	101
5.11	Content Filter Pattern	102
5.12	Claim Check Pattern	103
5.13	Service Bus Relay with Outbound Service Connections	106
5.14	Split Service Layer for Internal and External Hosting	107
5.15	Cloud Hosted SIL as a DMZ	109
5.16	Data at Rest on Premises	110
5.17	Components of the SIL	112
6.1	Compute Nodes	115
6.2	Cost Per User Comparison	118
6.3	Node Efficiency with User Growth	119
6.4	Monolithic Service Deployment	123
6.5	Componentized Service Deployment	123
6.6	Message Routing	125
6.7	Incorrect Service Deployment	126
6.8	Clients Grouped in UI Layer, Services Behind the SIL	127
6.9	CQRS	133
6.10	Tradeoff Triangle	136
6.11	CQRS Write Side	137

6.12	Related Independent Services	144
6.13	Round Robin Load Balancing	148
6.14	Message Router Normal Configuration	149
6.15	Instance One Out of Rotation and Updated	150
6.16	Instance Two Out of Rotation and Updated	150
6.17	All Instances Updated	151
7.1	Typical Logical Interaction Flow	157
7.2	Physical Interaction Flow	158
7.3	Customer Class	162
7.4	QM Data Aggregation	165
7.5	Repository Over a Polyglot Data Layer	167
7.6	Polyglot Serialized Data	172
7.7	DAL Components	174
8.1	Basic Data Storage Mechanisms	177
8.2	CQRS/ES with a Single Relational Database	180
8.3	NameInfo and AddressInfo Families	193
8.4	Graph Database Sample	195
8.5	Data Layer Components	200
8.6	Polyglot Data Storage	201
9.1	Type One Cluster Distribution	215
9.2	Type Two Cluster Distribution	216
9.3	Type Three Cluster Distribution	217
10.1	The Five-Layer Mobile App Architecture	222

---

# *LIST OF TABLES*

2.1	Serialized File Size Comparison	38
5.1	Push Notification Service Comparison	94
5.2	Translation Mapping Table	99
5.3	Message Translation Layers	99
6.1	Event Sources	135
7.1	Simple Metadata Mapping	164
7.2	Type and Property Metadata Mapping	164
8.1	Denormalized Address Table	178
8.2	Denormalized Phone Number Table	179
8.3	Normalized Address Table	179
8.4	Normalized Phone Number Table	179
8.5	Normalized User Info Table	179
8.6	Hashed Partition Key	183
8.7	KVDB Collisions	184
8.8	Single Document Links	189
8.9	Hierarchical Document Links	189
8.10	Normalized Address Table	190
8.11	Normalized Phone Number Table	190
8.12	Normalized User Info Table	190
8.13	Column Family Logical Structure	191
8.14	NameInfo Column Family	191
8.15	AddressInfo Column Family	192
8.16	PhoneInfo Column Family	192
8.17	NoSQL Choices	198
9.1	App Store Update Nuances	217

---

## PREFACE

**I WAS WONDERING** how to start this preface, and it occurred to me that writing one is backwards. After all, this is the text before the book, but you write it after the book is finished. I suppose that is analogous to how we have written software for a long time. We figure out what it is we are trying to do usually though creating systems that do what we think we want them to do; then we go back and write the documentation about what the system actually does.

This book, in large part, is aimed at helping to make that a more harmonious effort. Technology is moving so fast now that often we find ourselves trying to create mobile apps and services in a very reactionary manner. We tend to be on the back foot and playing catch up most of the time. If we could just take the time to sharpen the proverbial axe, we'd be able to get more accomplished, faster, with a lot less hair pulling.

I suppose over the past couple decades of doing this, I've seen that pattern time and time again. But there are also some good habits and patterns that I've seen along the way that in some respect were ahead of their time. Service-oriented architecture, for example, was a great idea for connecting the myriad of systems we've had inside organizations with simple, easy-to-use interfaces. In fact, this tried-and-true pattern, or collection of patterns if you will, is more relevant today in our commodity cloud computing, mobile app world than ever before.

This book is designed to provide some high level architectural guidance on how to design modern mobile apps and services so that they can adapt as technology inevitably changes. Of course, we start off with a brief history of our mobile computing explosion, and take a look at attempts to create cross platform apps and technology stacks.

Then I want to introduce what hopefully has become an obvious application stack. While we have been fairly fixated on a  $N$ -tier stack, where  $N$  usually equaled three, to truly futureproof our architectures, we really need two more clearly defined layers to provide us an abstraction boundary which insulates our code from changes in external client technology, as well as the rapidly changing data storage technologies we use today.

Once we have our layers sorted out, we'll have a look at various patterns of application development and how they apply to this layering system to create performant and resilient services for making powerful mobile applications.

I hope that you find this guidance useful. Perhaps it will make you think of things in ways you hadn't before, or validate thinking you've already implemented. In any case, I hope it prevents you from having to operate in a reactionary manner to the rapid changes of our modern computing world and lets you get on the front foot



so you can focus on creating great apps and services instead of retooling everything because a new phone hit the market.

## Target Audience

This book is for anyone who is responsible for the design, architecture, and development of modern mobile apps and the services that support them. I've written this book with futureproofing in mind. Ideally, the architectures and patterns in this book will provide you with an approach that will futureproof your designs.

By following this guidance, you should be able to create mobile app services that you can adapt, modify, update, change, or integrate without disrupting your mobile apps, or your teams. You should be able to deploy new services, change existing services, and add new client apps all without disturbing any of the running systems.

Most of all, you should be able to adapt services and apps based on this guidance to any new mobile platform that comes along. This will greatly increase your code reuse, make your teams much more efficient, and make your organization adaptable to the ever-changing mobile app landscape.

---

## *ACKNOWLEDGMENTS*

These kinds of things don't happen without a lot of people in the background pushing, pulling, helping, and sometimes simply nodding and smiling. I would like to thank Chris Bright for encouraging me and allowing me the time to put this together. I'd also like to thank Andrew Coates and Dave Glover for letting me harass them with ideas, and "what if" questions all the time.

Most important of all, I need to thank my wife Stefanie Heckman, and my two girls Elyssia and Keiralli. They not only encouraged me to finish, but were patient with me, and gave me the time to keep typing away. I think, in the end, their love and enthusiasm are what really got this book over the line. So if you like it, don't forget to thank them too.

---

# TABLE OF CONTENTS

<i>LIST OF FIGURES</i>	<i>xi</i>
<i>LIST OF TABLES</i>	<i>xiii</i>
<i>PREFACE</i>	<i>xv</i>
<i>ACKNOWLEDGMENTS</i>	<i>xvii</i>
<b>CHAPTER 1</b> <i>THE MOBILE LANDSCAPE</i>	<b>1</b>
1.1    Introduction	1
1.2    Previous Attempts at Cross-Platform	2
1.2.1    Java	2
1.2.2    Early Web Apps	5
1.2.3    Multiple Codebases	7
1.3    Breadth Versus Depth	9
1.4    The Multi-Platform Targets	10
1.4.1    Traditional	10
1.4.2    Mobile	11
1.4.3    Wearables	12
1.4.4    Embedded	13
<b>CHAPTER 2</b> <i>PLATFORM-INDEPENDENT DEVELOPMENT TECHNOLOGIES</i>	<b>15</b>
The Golden Rule	15
2.1    Vendor Lock-In	16
2.2    Recommended Standards and Guidelines	18
2.2.1    Respecting the Device	18
2.2.2    Respecting the Network	19
2.2.3    Communication Protocols	21
2.2.4    Data Formats	31
2.2.5    Mobile User Experience Guidelines	40
2.2.6    Authentication	45
2.2.7    Dealing with Offline and Partially Connected Devices	47
2.3    Wrapping Up	63
<b>CHAPTER 3</b> <i>PLATFORM-INDEPENDENT DEVELOPMENT STRATEGY</i>	<b>64</b>
3.1    High-Level App Development Flow	64
3.2    Five-Layer Architecture	65
3.3    Five-Layer Architecture Detail	66
3.3.1    The User Interface Layer	66
3.3.2    The Service Interface Layer	68

- 3.3.3 The Service Layer 69
- 3.3.4 The Data Abstraction Layer 70
- 3.3.5 The Data Layer 70

**CHAPTER 4 THE USER INTERFACE LAYER**

72

- 4.1 Porting Versus Wrapping 72
- 4.2 Multi-Client Development Tools 73
  - 4.2.1 PhoneGap (<http://phonegap.com/>) 73
  - 4.2.2 Xamarin (<http://xamarin.com/>) 74
  - 4.2.3 Unity (<http://www.unity3d.com>) 75
  - 4.2.4 Visual Studio 76
- 4.3 Cross-Platform Languages 76
- 4.4 Avoid Writing for the Least Common Denominator 77
- 4.5 Wrapping Up 78

**CHAPTER 5 THE SERVICE INTERFACE LAYER**

79

- 5.1 Message Processing 79
  - 5.1.1 Push versus Pull 80
  - 5.1.2 Partially Connected Scenarios 81
- 5.2 Message Processing Patterns 82
- 5.3 High-Volume Messaging Patterns 85
  - 5.3.1 Queue Services and Microsoft Azure Event Hubs 86
  - 5.3.2 Web Sockets 89
- 5.4 High-Volume Push Notifications 91
  - 5.4.1 Third Party Notification Hubs 93
- 5.5 Message Translation and Routing 97
  - 5.5.1 Message Translation 97
  - 5.5.2 Message Routing 103
  - 5.5.3 Handling Large Amounts of Data 108
- 5.6 Wrapping Up 111

**CHAPTER 6 THE SERVICE LAYER**

114

- 6.1 Thinking in Nodes 114
  - 6.1.1 Scale Out and Scale Up 114
  - 6.1.2 Scale Out versus Scale Up 114
- 6.2 Planning for Horizontal Scaling 117
  - 6.2.1 Node Sizing 117
  - 6.2.2 Statelessness 120
- 6.3 Designing Service Layers for Mobile Computing 121
  - 6.3.1 Service Componentization 122
- 6.4 Implementation Abstraction 124
  - 6.4.1 Service Interface Abstraction 124
- 6.5 Using CQRS/ES for Service Implementation 127
  - 6.5.1 CQRS Overview 127
  - 6.5.2 Why CQRS 129
  - 6.5.3 Being Able to Separate Data Models 129
  - 6.5.4 Aggregates and Bounded Contexts 131

6.5.5	The Read and Write Sides	132
6.5.6	CQRS Communications	132
6.6	Side by Side Multi-Versioning	140
6.7	Service Agility	141
6.8	Consumer, Business, and Partner Services	141
6.9	Portable and Modular Service Architectures	142
6.9.1	Designing Pluggable Services	145
6.9.2	Swapping Services	147
6.9.3	Deployment and Hosting Strategies	151
6.10	Wrapping up	152

---

## CHAPTER 7 THE DATA ABSTRACTION LAYER

154

7.1	Objects to Data	154
7.2	Using the DAL with External Services	157
7.3	Components of a DAL	159
7.3.1	Data Mapper	160
7.3.2	Query Mapper	161
7.3.3	Repository	166
7.3.4	Serializers	168
7.3.5	Storage Consideration	169
7.3.6	Cache	172
7.4	Wrapping Up	174

---

## CHAPTER 8 THE DATA LAYER

176

8.1	Overview	177
8.2	Business Rules in the Data Layer	178
8.3	Relational Databases	178
8.4	NoSQL Databases	181
8.4.1	Key Value Database	183
8.4.2	Document Database	186
8.4.3	Column Family Databases	189
8.4.4	Graph Database	194
8.4.5	How to Choose?	197
8.5	File Storage	197
8.6	Blended Approach	200
8.6.1	The Polyglot Data Layer	201
8.7	Wrapping up	203

---

## CHAPTER 9 STRATEGIES FOR ONGOING IMPROVEMENT

204

9.1	Feature Expansion	204
9.1.1	User Interface	206
9.1.2	Service Interface Layer	206
9.1.3	Service Layer	206
9.1.4	Data Abstraction Layer	206
9.1.5	Data Layer	207
9.2	Data Collection Matters	207
9.3	Multi-Versioning	209

**X**    **CONTENTS**

9.4    Version Retirement    **212**  
      9.4.1    Scale Back    **214**  
9.5    Client Upgrades    **216**  
9.6    Wrapping Up    **220**

**CHAPTER 10    CONCLUSION** **221**

---

*REFERENCES* *225*

*INDEX* *229*

---

# THE MOBILE LANDSCAPE

---

## 1.1 INTRODUCTION

---

When the idea of reaching people first struck home in the dark ages, we wanted to find ways for people to use and pay for our services. We had to find a way to let people know these services existed. In early days there were town criers, then during the industrial revolution when we could reproduce and distribute text to a largely literate audience, we had broadsheets. Then came the catalogue where mercantile companies would list their wares for sale. Once we heard the first radio waves, one of the first things they did was to sell advertising on the radio. This graduated to television advertising. Then along came the Internet. Everyone had to get themselves a website and would put their website address in their print, radio, and TV ads. Along came Facebook and everyone created a Facebook page for their companies.

Now, everyone wants to have an app for their customers to download. These apps go with customers wherever they are and provide instant interaction between consumer and supplier. We can push advertising into them, take orders through them, keep in touch with friends and relatives, and of course play games, listen to music and watch videos all in the palms of our hands. These experiences require devices, operating systems, and apps, all of which require software companies, architects, and developers to produce them. Unfortunately, these devices and operating systems often change.

In today's computing world, there is one thing you can be sure of; the leading operating system (OS) platform will change. As recently as 7 years ago, Microsoft Windows Mobile was the leading smartphone platform and tablet computers, while mobile, were large and clunky and ran full versions of the Windows XP and Windows 7 OS. Then came Blackberry which took a lot of market share from Windows Mobile. But that only lasted until the iPhone came along in 2007 and we went from a feature phone dominated world to a smartphone dominated one. This set a new benchmark and became the leading mobile computing platform. In the same year, the Open Handset Alliance re-released Linux-based Android-powered smartphones. Then in 2010, Google launched its Android-based Nexus series of devices. By 2011, Android-powered smartphones made up the majority of mobile OS-powered smartphones shooting past the iPhone.

While phones were taking off, in 2010 Apple released the iPad. Tablet computing was not new and in fact Microsoft and its Original Equipment Manufacturers

OEM partners had been trying to sell tablet computers since 2003. However, the iPad's sleek design brought tablet computing to the masses despite the clumsy and restrictive iOS operating system. This opened up the tablet computing market which Android was well suited for. After the iPad's initial success, by 2013 Google's Android-powered tablets had overtaken iPads as the tablet of choice. Additionally, although lagging considerably behind, Microsoft has re-created itself to make a run in the mobile and tablet computer markets as well. With Microsoft's massive install base, and very large developer ecosystem, they are likely to challenge Apple and Google in the mobile and tablet space eventually. With Windows 10 released as a free upgrade for over 1.5 billion eligible devices [1], it is likely to be the most common cross platform OS. That is, over half of Gartner's predicted 2,419,864,000 devices shipped into the market in 2014 [2]. Overnight the app ecosystem market leader could change again.

What this means for software developers, independent software vendors (ISVs), hobbyist app developers, and online service providers is that every few years they will have to retarget their efforts for a new platform, new development languages, new development tools, new skills, and new ways of thinking. This is not an attractive proposition for anyone. However, due to the success of the iPhone and iPad, software developers were willing to re-skill and even purchase proprietary hardware just to be able to develop applications for the new platforms. Then when Android devices surpassed the Apple devices, these same developers painfully went through the whole process again. Developers were forced to maintain three or more separate and complete codebases. This is the problem that Platform-Independent Delivery of Mobile Apps and Services solves.

If you are not planning a platform-independent strategy, you will likely be an ex-company in 3–5 years. Due to the rapid change of the consumer and enterprise mobile computing landscape, software developers must be able to adapt to new platforms, devices, and services before their competition. While cross-platform goes a long way toward this goal, it is still cumbersome and tends to lag behind a more platform-independent strategy. While it is not practical to get completely away from device-specific app code, the more you can move off of the device and put into a reusable back-end service, the less code you have to write and maintain when a new OS version or a new platform comes along. In this book we will examine strategies to do this, and provide future proof foundations to support changes in the computing landscape down the road.

**Disclaimer:** This book was written in late 2014. Everything in it was accurate at that time. If you are reading this in 2025, expect that a few things have changed. Just keep this in mind as we go through this so I don't have to keep writing "At the time of this writing...."

## 1.2 PREVIOUS ATTEMPTS AT CROSS-PLATFORM

---

### 1.2.1 Java

"Write once, run everywhere" was a slogan developed by Sun Microsystems which promised cross-platform capability for Java applications supported by Duke, Java's Mascot shown in Figure 1.1. This gained significant traction in the mid to late 1990s.





Figure 1.1 Duke

In the beginning of this era, the promise seemed legitimate. You could write the Java code once, package up your Java byte code in .jar files and run them on any system that had the Java Virtual Machine (JVM) interpreter. It worked so well that there are even C to Java compilers so your C applications can run with the same cross-platform reach that Java had.

The promise was that you could write code like this:

```
class CrossPlatformApp {  
    public static void main(String[] args) {  
        System.out.println("I am omnipresent!");  
        // Display the string.  
    }  
}
```

And it would run on every computer and device that ran Java without compiling multiple versions for each target device. All you had to do was make sure that the target device had the correct version of the JVM installed on it.

This worked fine until various vendors started creating their own versions of the JVM to run on their platforms. By 2014, more than 70 different JVMs [3] had been created that could run Java applications, for the most part. The catch was that they were each slightly different.

If we take the Sun JVM to be the standard, some of these other JVMs were better, and most were worse, at interpreting Java byte code. Some of them such as the IBM J9 ([http://en.wikipedia.org/wiki/IBM\\_J9](http://en.wikipedia.org/wiki/IBM_J9)), the Azul Zing JVM ([http://en.wikipedia.org/wiki/Azul\\_Systems](http://en.wikipedia.org/wiki/Azul_Systems)), and the Microsoft JVM ([http://en.wikipedia.org/wiki/Microsoft\\_Java\\_Virtual\\_Machine](http://en.wikipedia.org/wiki/Microsoft_Java_Virtual_Machine)) were better