

Cloud and Serverless Computing for Scientists

Juan A. Añel • Diego P. Montes
Javier Rodeiro Iglesias

Cloud and Serverless Computing for Scientists

A Primer

 Springer

Juan A. Añel
Universidade de Vigo
Ourense, Spain

Diego P. Montes
Universidade de Vigo
Ourense, Spain

Javier Rodeiro Iglesias
Universidade de Vigo
Ourense, Spain

ISBN 978-3-030-41783-3 ISBN 978-3-030-41784-0 (eBook)
<https://doi.org/10.1007/978-3-030-41784-0>

© Springer Nature Switzerland AG 2020

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

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG.
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword

As a researcher in a number of different domains over my career, all of which have benefited from the development in the areas of distributed computing, it is clear that we will continue on this journey of discovery and development of new tools and services for some time to come. We moved away from historic paradigms, such as the mainframe, with the rise of the personal computer and from then onto the increasingly powerful workstation, generic high-performance computing, grid computing, and now “the cloud.” One disadvantage of these changes though is that proponents of these new architectures often sell the newest paradigm as being better in nearly all ways than those that have preceded it and therefore it should replace the systems following the older paradigms, moving users onto newer, bigger, better shinier systems. This is plainly not the case once hype has been ignored and as such the content of this book should be considered as aiming to educate the reader in terms of making sure that the right tool is always used for the right job.

I was struck when reading the book by the diversity of target audience within the research field, making sure that communities of practice, where the use of cloud will become ever more valuable are covered, with relevant examples is a great feature. I would though also make sure that one further cohort reads this book, the project PI or senior professor for whom this book should be required reading. Alongside more traditional responsibilities they are also now the ones who are under legal obligations to comply with requirements on open data and reproducibility, something that the migration of services and data storage to the cloud is extremely suitable for. They should also ensure that they understand the benefits and almost more importantly the challenges that a cloud-first strategy can bring.

When working with Juan and Diego prior to the development of this book, we spent a significant amount of time working through different models of the utilization of Infrastructure as a Service cloud computing for the environmental sciences, in particular climate science. Therefore, a large number of points made by this book have been developed through practical experience of the authors, not through desk-based research. Along with this with my background in the development of pan European Federated Cloud services for research, I can clearly see how lessons learnt, sometimes in very hard ways, in how researchers should be

thinking about using the cloud could have been found quicker and more easily if such a text as this existed, or that we could have shared with members of the user community. As such I also see this book being of great use for those that will in time become users not only of large-scale public clouds to support their research but also in the usage of the European Open Science Cloud as it becomes a more and more useful tool to European Research.

I would finally highlight how the authors already say that in the fast-moving and dynamic world of cloud computing that things get out of date very quickly. Examples of this of course abound, though a key highlight is where the different deployment models of cloud computing mean that it is not always the same cloud that two different researchers talk about when they describe their method of use of “cloud.” Therefore, making sure that everyone understands the underlying model of cloud computing as is done early within the book is essential. So to conclude I would say that readers should, once they think they understand how they can use cloud computing within their research, stop, read, and then think about how they want to move forward with the cloud, making sure that they follow sound scientific and IT principles in how they deploy services within the cloud, testing all the time and recognize that cloud is not a magic bullet, but in the longer term it can make things much, much easier if their deployments are well designed.

Associate Professor and Associate
Director – Innovation Oxford Research
Centre University of Oxford
Oxford, UK

Prof David Wallom

Preface

We have written this book for scientists, engineers, or anyone who wants to approach cloud computing, simply to know more about it or if they are considering or evaluating it as an alternative or complementary solution for their computing needs. Additionally, this book can be useful for IT professionals, such as Solutions Architects, who want to learn about the current and future needs of the scientific community and how they can be satisfied using cloud technologies, enabling them to offer more suitable proposals to their customers. Therefore, this text can provide a bridge between user and provider perspectives.

Moreover, this book attempts to provide a broad view of the current status of cloud computing and to serve as an introductory text for graduate, postgraduate, and Ph.D. students across every discipline, providing insight into potential uses and how it can help to develop projects and future work.

Ourense, Spain
Ourense, Spain
Ourense, Spain
November 2019

Juan A. Añel
Diego P. Montes
Javier Rodeiro Iglesias

Acknowledgements

We want to dedicate the book to our family and friends for their daily support and understanding, which makes it possible for us to devote our time to our research work, too often at the expense of our time with them.

We would also like to thank Springer for the opportunity to publish this book and their interest in it, especially Zachary Romano, our publishing editor here, and the AJE language editing service. Special thanks to Paulo Rodríguez from Dropbox for facilitating part of the work that has contributed to this book and useful discussions. Additionally, we would like to thank Google Enterprise, Microsoft Research, and Amazon Web Services for granting us computational resources and facilitating technical support to perform experiments using their cloud computing platforms. We are deeply grateful to Tom Grey, Barak Regev, and Shailesh Rao from Google and Kenji Takeda from Microsoft Research. We are also grateful to our colleagues at ClimatePrediction.net at the University of Oxford, in particular David Wallom, all the staff from the Supercomputing Centre of Galicia (CESGA), in particular Carlos Fernández, and Tomás Fernández Pena of the Universidade de Santiago de Compostela.

We thank Dr. Sulay Tovar from the CIMUS (Universidade de Santiago de Compostela) and Prof. José Agustín García from the Universidad de Extremadura for providing feedback on draft versions of this book. Also, we thank Alexandra Iglesias for her contribution to the design of the cover of the book.

Part of the research that we have developed for years, therefore enabling us to accumulate the knowledge to write this book, has been funded by different agencies and bodies, among them is the contribution by the European Regional Development Fund (ERDF). Additionally, Juan A. Añel has been supported in recent years with a Ramón y Cajal Grant from the Government of Spain (RYC-2013-14560).

Contents

- 1 Why This Book?**..... 1
 - References 3
- 2 Why the Cloud?** 5
 - 2.1 Introduction 5
 - 2.2 Requirements and Needs 6
 - 2.2.1 Sasha in Their House or with a Self-employed Business..... 7
 - 2.2.2 Sasha as Entrepreneur 7
 - 2.2.3 Sasha in Their Laboratory 10
 - 2.2.4 Sasha Has an Idea 13
 - 2.2.5 Sasha Needs Mobility 14
 - 2.3 Everyone Is Sasha 14
 - References 15
- 3 From the Beginning to the Future**..... 17
 - 3.1 What Is Cloud Computing?..... 17
 - 3.2 Types of Cloud 18
 - 3.2.1 Commercial and Non-commercial..... 18
 - 3.2.2 Cloud Models 20
 - 3.2.3 Types of Service 21
 - 3.3 Features of the Cloud 23
 - 3.3.1 Elasticity and Scalability 23
 - 3.3.2 Self-service Provisioning 23
 - 3.3.3 Application Programming Interfaces 24
 - 3.3.4 Performance Monitoring and Measuring 24
 - 3.4 Cloud Native Concepts 24
 - 3.4.1 Cloud Native 25
 - 3.4.2 Containers 25
 - 3.4.3 Microservices 26
 - 3.4.4 Kubernetes 26
 - 3.5 Environmental Issues 26

| | | |
|----------|--|----|
| 3.6 | Success Stories of Cloud and Serverless Computing in Science | 27 |
| 3.6.1 | Archeology | 27 |
| 3.6.2 | Computer Sciences | 28 |
| 3.6.3 | Energy | 28 |
| 3.6.4 | Geosciences | 28 |
| 3.6.5 | Library Science | 28 |
| 3.6.6 | Mathematics | 29 |
| 3.6.7 | Medical Sciences | 29 |
| 3.6.8 | Physics | 30 |
| 3.6.9 | Social Sciences | 30 |
| 3.6.10 | Space Research | 30 |
| | References | 30 |
| 4 | Show Me the Money | 33 |
| 4.1 | Money for Research Using Cloud Computing | 33 |
| 4.2 | Sharing Costs | 34 |
| 4.3 | Selling Research | 35 |
| 4.4 | What Happens If I Do Not Pay? | 35 |
| 4.5 | What Happens If My Provider Does Not Deliver | 36 |
| 4.6 | Assessing Migration | 36 |
| 4.6.1 | Selecting a Cloud Computing Provider | 37 |
| 4.6.2 | Problem Definition | 37 |
| 4.6.3 | Gather Requirements | 37 |
| 4.6.4 | Obtain the Data | 39 |
| 4.6.5 | Review and Decide | 39 |
| 4.7 | Conclusions | 40 |
| | References | 40 |
| 5 | Tools in the Cloud | 41 |
| 5.1 | Sasha Wants Better Tooling for Their Organization | 41 |
| 5.2 | General Productivity | 41 |
| 5.2.1 | Office | 42 |
| 5.2.2 | Communication | 42 |
| 5.3 | Scientific and Engineering | 43 |
| 5.3.1 | Document and Paper Edition and Sharing (LaTeX) | 43 |
| 5.3.2 | Project and Issue Management | 43 |
| 5.3.3 | Version Control | 44 |
| 5.3.4 | Other Relevant SaaS | 45 |
| 5.4 | Sasha Decision | 45 |
| | References | 46 |
| 6 | Experiments in the Cloud | 47 |
| 6.1 | Sasha Is Ready to Start a New Experiment | 47 |
| 6.2 | Running Code on the Cloud | 48 |
| 6.2.1 | Understanding Software Development on the Cloud | 48 |
| 6.2.2 | Application Programming Interfaces (APIs) | 49 |

| | | |
|----------|--|-----------|
| 6.2.3 | Cloud Environment and Runtimes | 50 |
| 6.2.4 | Code in the Cloud: Considerations | 53 |
| 6.3 | Starting an Experiment in the Cloud | 55 |
| 6.3.1 | Example: Month Classifying Experiment | 55 |
| | References | 61 |
| 7 | Serverless Experiments in the Cloud | 63 |
| 7.1 | Sasha Revisits the Experiment | 63 |
| 7.2 | Thinking Serverless | 64 |
| 7.2.1 | Pros and Cons | 64 |
| 7.3 | Example: Month Classifier, Serverless Version | 65 |
| 7.4 | Example: Serverless Sensor Data Collector | 70 |
| | References | 72 |
| 8 | Ethical and Legal Considerations of Cloud Computing | 73 |
| | References | 75 |
| 9 | You Are Outdated, We Are Already Updating This Book | 77 |
| | Glossary | 79 |
| | Index | 83 |

Acronyms

| | |
|--------|---|
| AI | Artificial Intelligence |
| API | Application programming interface |
| AWS | Amazon Web Services |
| BOINC | Berkeley Open Infrastructure for Network Computing |
| CDC | USA Centers for Disease Control and Prevention |
| CERN | European Organization for Nuclear Research— <i>Conseil Européen pour la Recherche Nucléaire</i> |
| CLI | Command Line Interface |
| CNCF | Cloud Native Computing Foundation |
| DoD | US Department of Defense |
| EC2 | (AWS) Elastic Cloud Compute |
| ERP | Enterprise Resource Planning |
| ESA | European Space Agency |
| EU | European Union |
| FaaS | Function as a Service |
| FDA | U.S. Food and Drug Administration |
| FMI | Finnish Meteorological Institute |
| GCP | Google Cloud Platform |
| GCS | Google Cloud Storage |
| GPL | GNU Public License |
| HIPAA | Health Insurance Portability and Accountability Act |
| HIV | Human Immunodeficiency Virus |
| HPC | High performance computing |
| HPCaaS | High performance computing as a Service |
| IaaS | Infrastructure as a Service |
| IOPS | Input/output operations per second |
| IoT | Internet of Things |
| IRC | Internet relay chat |
| IT | Information Technologies |
| JPL | Jet Propulsion Laboratory |
| JSON | JavaScript object notation |

| | |
|------------|---|
| Met Office | United Kingdom's National Weather Service |
| MPI | Message Passing Interface |
| NAS | Network Attached Storage |
| NASA | National Aeronautics and Space Administration |
| NOAA | USA National Oceanic and Atmospheric Administration |
| NREL | National Renewable Energy Laboratory |
| OS | Operating system |
| PaaS | Platform as a Service |
| POSIX | Portable Operating System Interface |
| PUE | Power Usage Effectiveness |
| RAM | Random-Access Memory |
| S3 | AWS Simple Storage Service |
| SaaS | Software as a Service |
| SDK | Software Development Kit |
| SEO | Search Engine Optimization |
| SSD | Solid-State Drive |
| SSL | Secure Sockets Layer |
| TLS | Transport layer security |
| UI | User Interface |
| VPC | Virtual private cloud |
| VPN | Virtual private network |
| VPS | Virtual private server |