Building a secure and scalable cryptocurrency exchange

Florian M. Harwöck, Lukas G. Kurz

A diploma thesis
submitted to the department of Computer Science
of Higher Technical College Leonding

Principal Advisor: Franz Auernig

April 8, 2019

Declaration of Academic Honesty

I declare in lieu of an oath that I have written the present thesis independently and without external help, that I have not used sources and aids other than those indicated, and that I have identified as such the places taken from the sources used, both literally and in terms of content.

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorgelegte Diplomarbeit selbstständig und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe. Gedanken, die aus fremden Quellen direkt oder indirekt übernommen wurden, sind als solche gekennzeichnet. Die Arbeit wurde bisher in gleicher oder ähnlicher Weise keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht.

- Florian M. Harwöck

- Lukas G. Kurz

Gender Clause

For better readability, gender-specific wording is not included in this paper. It is expressly pointed out that the wording in the text refers to all genders in the same way.

Gender Klausel

Zur besseren Lesbarkeit wird in der vorliegenden Arbeit auf geschlechterspezifische Formulierungen verzichtet. Es wird ausdrücklich darauf hingewiesen, dass im Text personenbezogene Formulierungen immer gleichermaßen auf alle Geschlechter bezogen sind.

Acknowledgements

We would like to express our special thanks to our thesis supervisor *Franz Auernig*, who supported us in all phases of our work and gave us support and mentorship.
A further thank you goes to our families and friends, who helped us during the work on our diploma thesis with feedback and suggestions for improvement.

Danksagungen

Unser besonderer Dank gilt unserem Arbeitsbetreuer *Franz Auernig*, der uns in allen Phasen unserer Arbeit zur Seite stand und uns unterstützt und betreut hat.
Ein weiterer Dank geht an unsere Familien und Freunde, die uns während der Arbeit an unserer Diplomarbeit mit Feedback und Verbesserungsvorschlägen geholfen haben.

Abstract

In 2009 "Satoshi Nakamoto" published the cryptocurrency Bitcoin. Since then, around 2100[1] digital assets have been published. In total, these projects have a market capitalization of $181 billion USD[2]. In order to be able to trade all those different currencies and assets among themselves, digital exchanges are needed, which are currently often insufficiently secure or scalable. This diploma thesis shows how a cryptocurrency exchange, that fulfills modern security requirements and has sufficient scalability, can be built.

Implemented with a reasonable mix of proven software components and cloud-native products, this project shows a prototype implementation with a solid security and scalability foundation that can be used by future digital exchange projects of all kinds. It is compatible with the largest decentralized digital asset network, Ethereum's ERC20 token standard. The primary intention and objective of this thesis is to build a prototype implementation which focuses on the fundamental issues of security and scalability.

Zusammenfassung

Im Jahr 2009 hat "Satoshi Nakamoto" die Kryptowährung Bitcoin veröffentlicht. Seitdem wurden bereits über 2100[3] weitere digitale Vermögensgegenstände veröffentlicht. Insgesamt weisen diese Projekte eine Marktkapitalisierung von $181 Mrd. USD auf[4]. Um all diese verschiedenen Währungen und Vermögenswerte untereinander handeln

[1] As of https://coinmarketcap.com on the 8th of April 2019
[2] See note 1
[3] See note 1
[4] See note 1

zu können, bedarf es digitaler Börsen, welche derzeit oft unzureichend sicher oder skalierbar sind. Diese Diplomarbeit zeigt, wie ein Kryptowährungs Börse aufgebaut werden kann, welche moderne Sicherheitsanforderungen erfüllt und über eine ausreichende Skalierbarkeit verfügt.

Implementiert mit einer sinnvollen Mischung aus bewährten Softwarekomponenten und Cloud-basierten Produkten, zeigt dieses Projekt eine Prototypen Implementierung mit einer soliden Sicherheits- und Skalierbarkeitsbasis, die von zukünftigen Projekten über digitale Börsen aller Art genutzt werden kann. Der Prototyp ist kompatibel mit dem größten dezentralen Digital Asset Network, dem ERC20-Token-Standard von Ethereum. Die primäre Absicht und Ziel dieser Arbeit ist es, einen Prototypen zu implementieren mit Fokus auf Sicherheit und Skalierbarkeit.

Contents

Chapter 1

Introduction

This chapter delivers a summary of the motivation of this thesis. The main topic of cryptocurrencies and their respective digital exchanges are introduced. Furthermore, the common problems faced in the industry are briefly addressed. In addition, the defined objectives are described, and their reasoning is explained.

1.1 Motivation

In late 2017, most significant cryptocurrencies, such as Bitcoin, Ethereum, Litecoin, and Ripple, saw a drastic increase in value. Bitcoin nearly had a threefold increase in price from November to the end of December. The rise in price lead to a sort of "hype" in the community, resulting in more publicity and growth of the cryptocurrency industry. The writers of this paper, Florian Harwöck, and Lukas Kurz, also gained interest in these cryptocurrencies, but more from a technical viewpoint. They were interested in the inner workings of cryptocurrencies and cryptocurrency exchanges. There are many different exchanges, all trying to appeal to users, but upon closer inspection, many of these exchanges lack essential features. While some of them are not able to cope with massive amounts of traffic, which negatively impacts the user experience on the exchanges, others have **severe security problems**. There was daily news about exchanges being hacked, using all kinds of methods, ranging from the simplest forms of server exploitation and Social Engineering to brute force attacks.

1.2 Objectives

Both of the writers have a background in software engineering and knowledge of software architecture. To improve the situation, the writers decided to **research basic techniques for security and scalability, used in modern applications, and apply some of them to a cryptocurrency exchange**. The aim was not to devise new means for scalability and security, but instead, use and implement already well known and established concepts, from other areas. The goal was to **display ways of making exchanges more secure and scalable**, but also to **spread awareness** about specific security techniques, that can be used to protect exchanges and their customers.

Chapter 2

Thesis Results

This chapter delivers a short overview of the final achievements for this diploma thesis, without going into technical details. Short explanations of functionalities and relevant information to understand them is given. All theoretical research, thoughts, trade-offs, comparisons, etc. and the practical implementations are discussed and shown in later chapters.

2.1 The thesis' final result

This thesis reached its objectives and goals in two different areas: Firstly different concepts for improving scalability and security were designed. In conjunction with other industry proven best-practice-standards and software products, significant improvements compared to standard solutions were achieved. Secondly, most of the concepts could be directly implemented in the developed prototype and therefore validated for practicability and feasibility.

Design decisions, system architectures, security concepts, and learnings in this thesis can be used by future projects in similar directions as a solid foundation for security and scalability. The first section will introduce the reader into the more theoretical concepts behind this thesis' work. The second section will further elaborate on the prototype, and all the challenges and curiosities faced while developing a cryptocurrency exchange.

2.1.1 Theoretical achievements

Blockchain

One of the most critical aspects of a cryptocurrency exchange is the interaction with the decentralized network powering these systems. Even though the exchange software is a standard central software solution (in the meaning of a centrally controlled backend; the system itself is built on distributed microservices), one needs a deep understanding of blockchains, smart contracts, cryptography, and a few other topics.

2.1.2 Practical achievements

Order book and the trade matching algorithm

Processing orders as fast as possible is one of the tasks of an exchange. This task might seem rather trivial but requires at least a fast algorithm, so the exchange is not slowed down when more customers create trade offers. Creating such an algorithm also requires an understanding of the different kinds of modes in which orders can operate. Not only has the algorithm to support multiple types of trading offers,

but it also needs to be able to make different kinds of order modes work with each other. Otherwise the concept of multiple modes would be useless. Our algorithm is capable of handling market orders and limit orders. It was designed with three main processes in mind, validating and enqueueing orders, matching and persisting in the order book, and processing asset transactions.

High performance charts

What proved to be more difficult than expected, was making charts and graphs for the prices of the currencies. These charts have to display a large amount of data, while still performing reasonably fast and also look good, as to fit in with the rest of the user interface. Finding a solution to this problem, required much research, as well as testing different frameworks and libraries for creating such charts. These libraries had to be tested for performance, but also needed to be good looking, or atleast have the option of customizing them, to fit them to our needs. Highcharts was the library that best met those needs, so it was chosen as the library to develop the charts for the prices of the currencies.

Cryptocurrency wallets

One challenging part during the development of a digital asset exchange is the storage solution for the digital assets itself. A bank has big physical vaults to safeguard assets, whereas in the digital world this isn't possible. To solve this issue as efficient and most importantly as secure as possible, a specialized service was developed. This service is the single authority for all assets stored and is also the only entity, that is allowed to make changes in the balance records of a user. This system together with the extensive use of cryptography lets the service store digital assets securely.

Microservice design

One of the goals for designing this exchange was creating a backend that would be capable of scaling to fit the varying demands due to traffic. To achieve this, a microservice architecture for the backend was used. Microservices have benefits regarding the scaling issue, but also bring their own unique set of problems with

them. A solution for letting the microservices communicate with each other had been chosen. To make development easier and communication more uniform across the microservices, we chose GRPC for our API.

Debugging the whole system is always a difficult task since the code and work are split up into different services. To handle this, we used a modular approach, making what looks like a monolith, but its inner workings function like different modules communicating with each other, just like microservices. This way the whole microservice system can be split up in the production environment but also put back together into a monolith, when the system needs debugging.

Secure storage systems

Financial platforms, like exchanges, have a lot of interesting and valuable data. This data ranges from exact personal identifiable information (PII) and passwords to metrics about financial credibility and liquidity. To keep these data points as secure as possible content encryption is performed in transit and storage. Furthermore, pseudonyms are used throughout the system to remove links between the asset-balances, trading offers, and the end user identities.

2.2 Overview of the developed prototype

As with most services nowadays the user journey starts with a registration process of a new user account. In the case of this prototype, the user is obligated to enter a valid email address and password that fulfills the minimum expected requirements. During typing in the password, the user gets a dynamic overview with red and green dots for each condition. (see figure 2.1 on page 8)

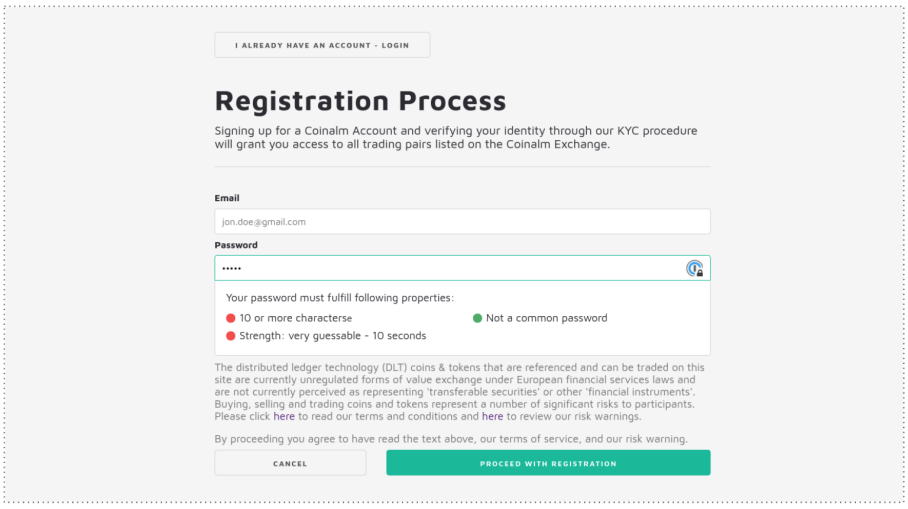

Figure 2.1: The prototype's registration page

After finishing the necessary registration process, the user is required to enter his details. That is done to comply with KYC (Know-Your-Customer) regulations from the European Union. The information entered includes name, gender, date of birth, phone number, residential address, passport number, employment status, self annual income, and expected investment. (see figure 2.2 on page 8) [1] [2]

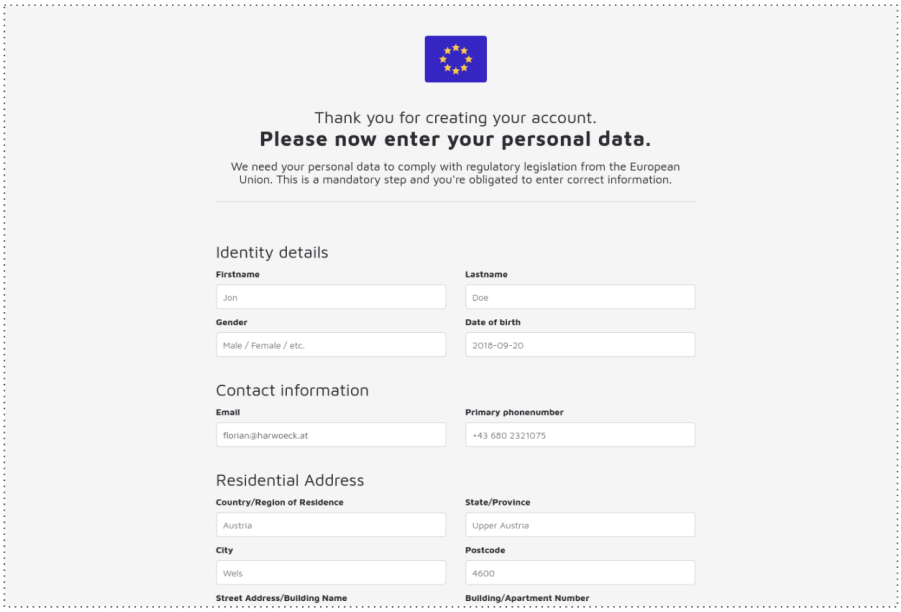

Figure 2.2: The personal data collection part of the registration process

After finishing the registration process, users are redirected to the dashboard. In this main application view, the user can see his estimated account balance (only if the user already deposited digital assets) and other information. Furthermore, the user can log out of the application. (see figure 2.3 on page 9)

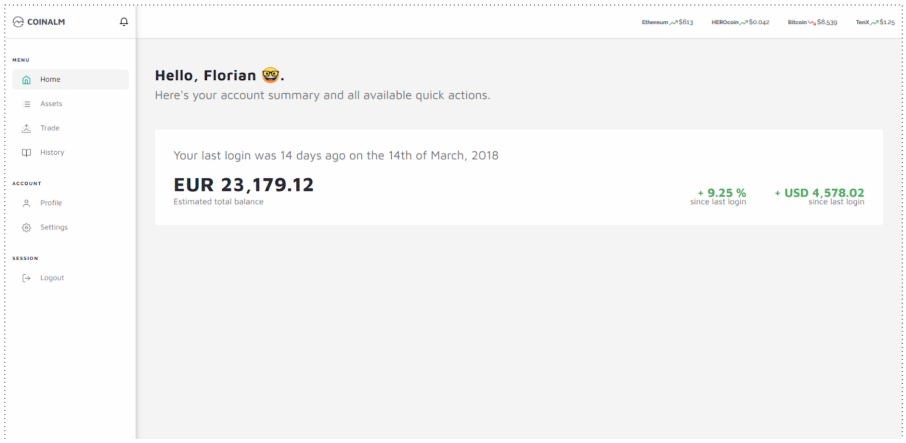

Figure 2.3: The dashboard of the implemented prototype

The user can display his details, entered previously in the KYC part of the registration, by visiting the profile page. In an arranged overview the user can verify the accurateness of his information. (see figure 2.4 on page 9)

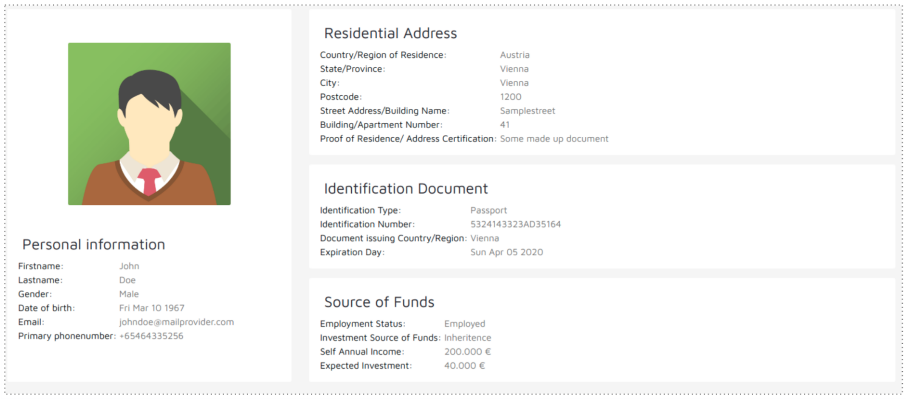

Figure 2.4: The overview of a user's entered details

To get to the core feature of the prototype, the trading page, the user first needs to select the trading pair he wants to trade. All available trading pairs are displayed in a list. (see figure 2.5 on page 10)

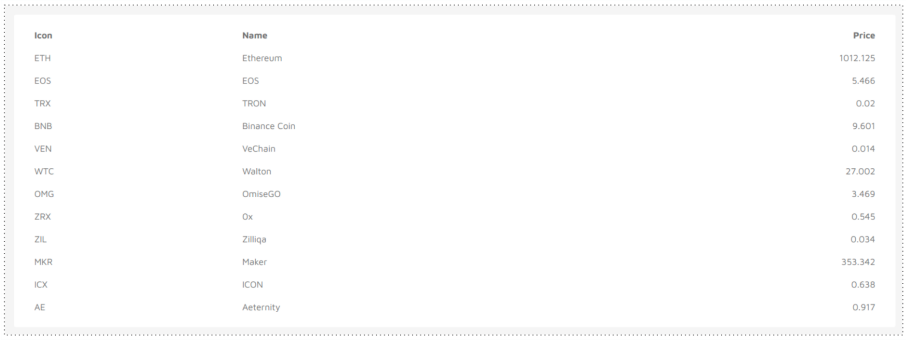

Icon	Name	Price
ETH	Ethereum	1012.125
EOS	EOS	5.466
TRX	TRON	0.02
BNB	Binance Coin	9.601
VEN	VeChain	0.014
WTC	Walton	27.002
OMG	OmiseGO	3.469
ZRX	0x	0.545
ZIL	Zilliqa	0.034
MKR	Maker	353.342
ICX	ICON	0.638
AE	Aeternity	0.917

Figure 2.5: The list of all trading pairs available

After the user has selected one of the trading pairs, he will be redirected to the trading details page. In this view, a history price graph, information resources, and the buy- and sell-offer form is found. The resources contain a small introduction text to the digital asset and links to their white-paper and websites. Users can use this information to educate them further about specific assets. To create a trading offer the user needs to enter the price at which he wants to buy or sell and the amount of the assets in question. (see figure 2.6 on page 10)

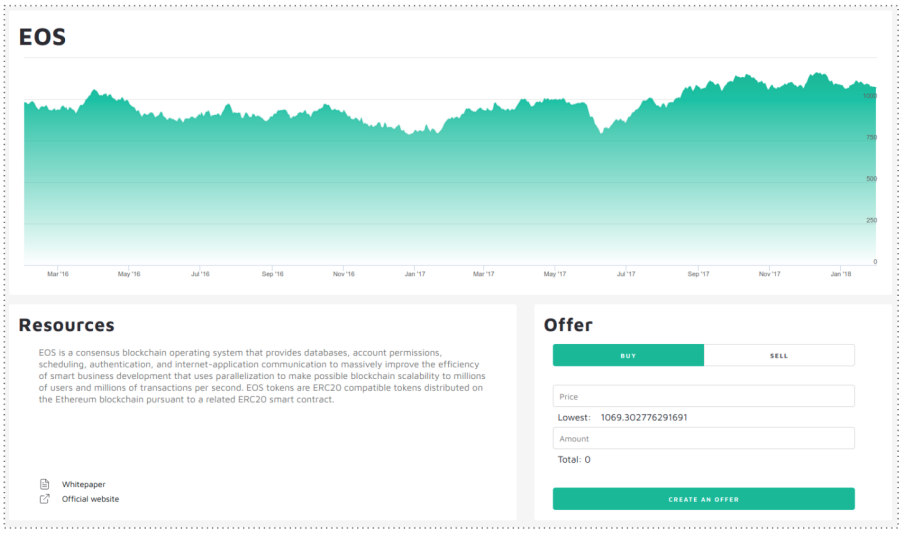

Figure 2.6: The trading page of a specific asset

After a trading offer was created and submitted, it is shown in the trading history list. A user can see all previously happened trades and their respective conditions. (see figure 2.7 on page 11)

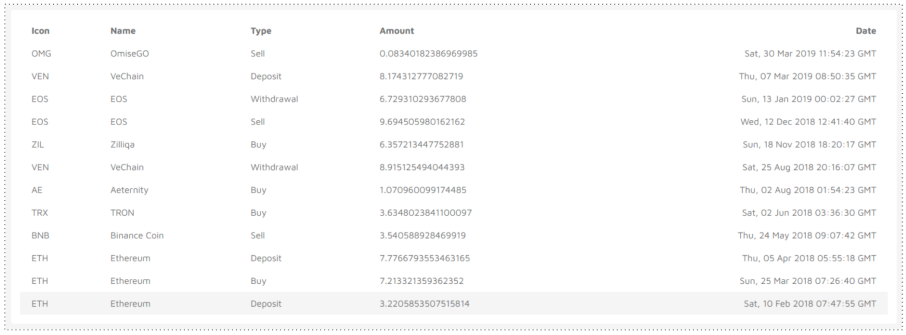

Icon	Name	Type	Amount	Date
OMG	OmiseGO	Sell	0.08340182386969985	Sat, 30 Mar 2019 11:54:23 GMT
VEN	VeChain	Deposit	8.174312777082719	Thu, 07 Mar 2019 08:50:35 GMT
EOS	EOS	Withdrawal	6.729310293677808	Sun, 13 Jan 2019 00:02:27 GMT
EOS	EOS	Sell	9.694505980162162	Wed, 12 Dec 2018 12:41:40 GMT
ZIL	Zilliqa	Buy	6.3572134477752881	Sun, 18 Nov 2018 18:20:17 GMT
VEN	VeChain	Withdrawal	8.915125494044393	Sat, 25 Aug 2018 20:16:07 GMT
AE	Aeternity	Buy	1.070960099174485	Thu, 02 Aug 2018 01:54:23 GMT
TRX	TRON	Buy	3.6348023841100097	Sat, 02 Jun 2018 03:36:30 GMT
BNB	Binance Coin	Sell	3.540588928469919	Thu, 24 May 2018 09:07:42 GMT
ETH	Ethereum	Deposit	7.7766793553463165	Thu, 05 Apr 2018 05:55:18 GMT
ETH	Ethereum	Buy	7.213321359362352	Sun, 25 Mar 2018 07:26:40 GMT
ETH	Ethereum	Deposit	3.2205853507515814	Sat, 10 Feb 2018 07:47:55 GMT

Figure 2.7: The history list of all trading offers happened for this user account