

Juho Kettunen

Maintainability in cloud-native architecture

Master's Thesis in Information Technology

May 9, 2024

University of Jyväskylä

Faculty of Information Technology

Author: Juho Kettunen

Contact information: juho.kettunen@student.jyu.fi

Supervisor: Oleksiy Khriyenko

Title: Maintainability in cloud-native architecture

Työn nimi: Ylläpidettävyys pilvinatiiveissa arkkitehtuureissa

Project: Master's Thesis

Study line: Software and Telecommunication Technology

Page count: 69+6

Abstract: Goal of the thesis is to investigate how maintainability is addressed during the architectural design phase of cloud-native software development lifecycle. To this end, I conducted a survey among cloud architects at my current employer *Nordcloud*. Based on the survey, I found that maintainability is prioritized higher than scalability and performance, but lower than security and reliability. I categorized the advice given by the respondents, and saw that in order to increase maintainability, they prefer using established technologies, infrastructure-as-code, managed services, and CI/CD pipelines. To the same end, architectural considerations include loose coupling, microservices architecture, modularity, and architectural simplicity. Facilitating easy updates in the future is also seen valuable, and often even individual suggestions were relevant. Even if the comparison to literature is based on examples, the results generally appear to align with advice found in existing literature.

Keywords: maintenance, maintainability, public cloud, cloud-native, cloud architect, architecture, software architecture, Master's Theses

Suomenkielinen tiivistelmä: Tutkielman tavoitteena on selvittää kuinka ylläpidettävyys huomioidaan pilvinatiivien sovellusten arkkitehtuurisuunnitteluvaiheessa. Tämän saavuttamiseksi toteutin kyselyn nykyisen työnantajani *Nordcloudin* pilviarkkitehtien keskuudessa. Kyselyn perusteella havaitsin, että ylläpidettävyys priorisoidaan korkeammalle kuin skaalautuvuus ja suorituskyky, mutta alemmaksi kuin turvallisuus ja luotettavuus. Kategorisoin vas-

taajien antamat neuvot, jonka ansiosta näin että ylläpidettävyyden parantamiseksi he valitsevat vakiintuneita teknologioita, infrastructure-as-code -työkaluja, hallinnoituja palveluita, ja CD/CD-putkia. Yleisimmät arkkitehtuuriin liittyvät suositukset liittyivät löyhään kytkentään, mikropalveluarkkitehtuuriin, modulaarisuuteen, ja arkkitehtuuriseen yksinkertaisuuteen. Tulevaisuuden päivitysten helpottaminen nähtiin myös arvokkaaksi, ja usein yksittäisetkin ehdotukset olivat oleellisia. Vaikka vertailu kirjallisuuteen perustuu esimerkkeihin, niin tulokset vaikuttavat yleisesti ottaen noudattavan olemassaolevan kirjallisuuden ohjeistuksia.

Avainsanat: ylläpito, ylläpidettävyys julkinen pilvi, pilvinatiivi, pilviarkkitehti, arkkitehtuuri, ohjelmistoarkkitehtuuri, pro gradu -tutkielmat

Glossary

AD	Architectural Definition
BaaS	Backend-as-a-Service
CI/CD	Continuous Integration, Continuous Delivery/Deployment
CNA	Cloud-Native Application
CSP	Cloud Service Provider
FaaS	Functions-as-a-Service
IaC	Infrastructure as Code
MSA	Microservices Architecture
SBS	Service-Based System
μ SBS	Microservice-Based-System
VM	Virtual Machine

List of Tables

Table 1. Respondent experience in cloud-native architecture, $n = 15$	42
Table 2. Respondent experience in software architecture in general, $n = 15$	42
Table 3. Respondent experience in information technology in general, $n = 15$	42
Table 4. Priority options and their integer values	43
Table 5. Quality attributes, their average priorities, and other metrics, $n = 15$	43
Table 6. Experience range options and their integer values	44
Table 7. Correlations between experience in a given domain, and prioritization of maintainability.....	44
Table 8. Suggestions about platform- and technology choices, $n = 14$	46
Table 9. Suggestions about application architecture, $n = 14$	46
Table 10. Other suggestions to improve maintainability, $n = 14$	47
Table 11. Respondents' years of experience in the relevant domains.	64
Table 12. How respondents prioritize the selected quality attributes when designing cloud-native architecture.....	65
Table 13. How respondents address maintainability with platform- and technology choices.....	66
Table 14. How respondents address maintainability with application architecture.....	68

Contents

1	INTRODUCTION	1
2	BACKGROUND	3
2.1	Architecture in software development	3
2.1.1	Software architecture	3
2.1.2	Software development life cycle	7
2.1.3	General advice for achieving a good architecture	10
2.2	Maintainability	11
2.2.1	Quality Attributes	11
2.2.2	What is maintainability, and why is it important	14
2.2.3	How to measure maintainability	15
2.2.4	How to improve maintainability	19
2.3	Cloud-nativity	21
2.3.1	What is cloud-nativity	21
2.3.2	Microservices	24
2.3.3	Serverless	28
2.3.4	Cloud architecture	30
2.4	Maintainability in cloud-native architecture	32
3	METHOD	38
3.1	Research questions and choice of method	38
3.2	Survey	38
3.3	Analysis	41
4	RESULTS	42
4.1	Demographics (Q1-Q3)	42
4.2	Prioritization of maintainability (Q4)	43
4.3	Ways to improve maintainability (Q5 & Q6)	44
5	DISCUSSION	48
5.1	Research question RQ1	48
5.2	Research question RQ2	49
5.3	Research question RQ3	52
5.4	Internal and external validity	56
6	CONCLUSIONS	58
	BIBLIOGRAPHY	60
	APPENDICES	64
A	Survey responses: Q1-Q3	64
B	Survey responses: Q4	65
C	Survey responses: Q5	66
D	Survey responses: Q6	68

1 Introduction

Goal of the thesis is to investigate how maintainability should be addressed during the architectural design of a cloud-native application. Maintainability is a quality attribute, which describes how easy it is to maintain and modify a given system. Software architecture, created by the software architect, formalizes the structural and behavioral foundation upon which the rest of the software will be built. Software development life cycle can be described with different life cycle models that prescribe a way to arrange and order the different activities of software development. Cloud-nativity means using the technologies and methods best suited for the cloud computing paradigm. I will expand upon the relevant terminology in chapter 2.

I aim to answer three research questions:

RQ1 How much importance do cloud architects place on maintainability during the architectural design phase?

RQ2 How to address maintainability concerns when architecting cloud-native applications?

RQ3 Do the recommendations from literature match the views of architects working in the field?

The main contribution of this thesis are the results of a survey I conducted in order to answer the two first research question. The survey was directed at cloud architects working at my workplace *Nordcloud, an IBM company*¹. When analysing results of the survey, I will first ascertain the perceived importance of maintainability and contrast this to the respondents' years of experience in the target domain. Then I will categorize the suggested approaches for addressing maintainability. Finally, I will compare the survey results to approaches suggested in literature.

I chose this topic for three reasons. First, the maintenance phase of software development life cycle is prevalent. It takes up the majority of the system's total lifetime and costs. (Bass, Clements, and Kazman 1998). Second, choices made during the architectural design phase cascade into development and maintenance phases. Mistakes and oversights are slower and more expensive to correct later. (Bass, Clements, and Kazman 1998; Mumtaz, Singh, and

1. <https://www.nordcloud.com>

Blincoe 2021). Lastly, the cloud-native approach is highly relevant in today's technology landscape. It helps reduce time-to-market and move costs from capital expenditure to operational expenditure. Utilizing a public cloud platform also allows the company to leverage a highly scalable, reliable and secure infrastructure and a wide variety of easily integratable services. (Vettor, Pine, Jain, et al. 2022).

The thesis is structured as follows. In chapter 2 I dive into related research literature and other contemporary sources to define the necessary terminology. At the same time, I keep an eye out for best practices and other suggestions how to most effectively leverage the information in practice. Then in chapter 3 I describe how my own research is conducted in regards to the survey and analysis thereof. The results will be presented in chapter 4, and further discussed in chapter 5. I will close with conclusions in chapter 6.

2 Background

In this chapter I will review the literary findings around the thesis topic. My main database for researching existing literature was the *JYKDOK* search engine for international articles (“JYKDOK international articles search” 2023). I decided to search keywords in abstracts, instead of full-text or only in title. In my opinion the full-text search was too lax, and title-only too strict. I used the following search criteria:

- (Abstract:”cloud native” AND Abstract:maintainability)
- (Abstract:”cloud native” AND Abstract:architecture)
- (Abstract:”cloud native” AND Abstract:maintenance)
- (Abstract:architecture AND Abstract:maintainability AND Abstract:cloud)
- (Abstract:serverless AND Abstract:maintainability)
- (Abstract:serverless AND Abstract:operations)
- (Abstract:”cloud-native” AND Abstract:quality AND Abstract:attribute)
- (Abstract:software AND Abstract:”quality attribute” AND Abstract:maintainability)

With these criteria I received between 4 to 473 results per search, which was sufficient. In the end I did not find it necessary to branch out to other search engines, because I additionally employed an ad-hoc snowballing from bibliographies of promising sources. I was unable to find full-text of some articles with promising abstracts. In case of one specific article I emailed the authors ¹, and they kindly sent me a PDF of it.

2.1 Architecture in software development

2.1.1 Software architecture

In this section I will discuss the definition of software architecture and introduce related terminology. Every software system has an architecture. Bass, Clements, and Kazman (1998, 23) state that the architecture exists even if no one present knows any details about it. They provide us a definition as follows:

1. Bogner et al. (2018)

The software architecture of a program or computing system is the structure or structures of the system, which comprise software elements, the externally visible properties of those elements, and the relationships among them.

Next I will define the different parts of this definition, and adjacent terminology. A software system is a collection of software components organized to accomplish a specific function or set of functions (ISO/IEC/IEEE 42010:2007 2007, 3). These components — or elements — can be entities of varying size: objects, processes, libraries, databases, commercial software, and so on. In its most trivial case, we can consider the entire system to be a single component. This is indeed an architecture, but practically useless due to lack of necessary details (Bass, Clements, and Kazman 1998, 24). Koskimies and Mikkonen (2005, 53) define a component as an individual software unit that offers its services through well-defined interfaces. Implementation details of each component are not considered part of architecture, but their interactions are. Without describing the interactions we end up with a "bubbles and lines" diagram which is not sufficient to describe an architecture (Bass, Clements, and Kazman 1998, 24).

A software architect is a person, team, or organization responsible for systems architecture (ISO/IEC/IEEE 42010:2007 2007, 3). A system can include more than one structure (Bass, Clements, and Kazman 1998, 23). Because "structure" is an abstract concept, it can change depending on perspective. An architectural view models some aspect of the system (Koskimies and Mikkonen 2005). It is a representation of the whole system from the perspective of a related set of concerns, such as a specific stakeholder's specific concern (ISO/IEC/IEEE 42010:2007 2007, 3). A stakeholder is an individual, team or organization with interests in, or concerns relative to, a system. Examples of stakeholders include clients, users, the architect, developers, reviewers, and so on. (ISO/IEC/IEEE 42010:2007 2007, 3). Examples of architectural views include a scenario view, logical view, process view, development view, physical view, and modification view. They are partly overlapping: when a specific part of a system is changed, we often need to apply changes in multiple views. (Koskimies and Mikkonen 2005).

Architecture represents the first design decisions of a system (Bass, Clements, and Kazman 1998). It is difficult to get it right the first time, and can be laborious to alter down the road.

In worst case, there is a need for an architectural change, which might result in changes across the entire system and changes to interaction between components. If we instead can limit the change to a group of components or only a single one, the change will be easier to carry out. Out of most design artifacts, architecture has most far-reaching consequences (Bass, Clements, and Kazman 1998, 31). It will place constraints on implementation and maintenance: will the system support the required functions; how easy is it to test, maintain, change and extend; how large data masses can it process, and so on (Koskimies and Mikkonen 2005).

The architecture might even have effect on the development organization by nudging towards a specific way to divide tasks to teams by structure of its components (Bass, Clements, and Kazman 1998, 31). This is somewhat contrary to Conway's law ² which states that the technical structure of a system will reflect the social boundaries across the organizations that produced it. I suppose that in an agile organization the social structure is malleable enough to be impacted by a new system, when new teams need to be formed to facilitate the development.

According to Koskimies and Mikkonen (2005), architecture lies on a higher abstraction level than implementation. The lower abstraction level of a system manifests as source code (Bass, Clements, and Kazman 1998, 24). This might be named as "the implementation", or "design" as named in ISO/IEC/IEEE 12207:2017(E) (2017). Also the standard ISO/IEC/IEEE 42010:2007 (2007, 4) states that architecture is conceptual instead of concrete. This goes hand in hand with the used vocabulary. With architecture, we are using the vocabulary of the solution, and are further away from the problem domain than during implementation (Koskimies and Mikkonen 2005).

As we've read, architecture is of utmost importance for the implementation and future of the system. Because of this, it seems common sense to document it somehow. Bass, Clements, and Kazman (1998, 24) define an Architectural Description (**AD**) as a collection of products to document an architecture. They outscope the documentation of guiding processes from actual architecture, which means that an architecture can live independently of its description or specification. In contrast, ISO/IEC/IEEE 42010:2007 (2007, 3) mentions "... *and prin-*

2. https://en.wikipedia.org/wiki/Conway%27s_law

principles guiding its design and evolution” when defining ”architecture”. The standard states that chosen architectural concepts, as well as considered alternatives, must be justified in the description. According to the standard, the description can be utilized at multiple different levels of abstraction, at each level emphasizing the details necessary for that level. If an AD is not available, it can be reverse-engineered from an existing system. This might be needed if, for example, it was originally developed without a description, or if the original architect is no longer available for consultation. (ISO/IEC/IEEE 42010:2007 2007, 7, 67).

A software architecture doesn’t spawn out of nothing. The architecture is based on, among other things, the requirements placed on it (Bass, Clements, and Kazman 1998). Requirements codify the collective wishes and needs that various stakeholders want from the system. If requirements were the only input to software development, then an identical set of requirements would result in two pieces of identical architectures, even from two entirely unrelated organizations. (Bass, Clements, and Kazman 1998, 5–9).

As we can see, the organization, architect or architects, and the technical environment greatly influence how software systems are built (Bass, Clements, and Kazman 1998). The organization might have other existing products, architectures or assets that can be utilized. The organization has some individual structure, goals and long-term investments, and personnel with specific skills and availability. Especially the architects’ background and experience is pivotal in deciding the course for the architecture. The architect might lean towards a familiar architectural style that they have had success with in the past, and vice versa with an unhappy experience. Their professional community can infuse into the new system a set of standard industry practices and software engineering techniques. (Bass, Clements, and Kazman 1998, 5–9).

The architecture cause feedback to its background factors, such as the organization, customers, architects and the software engineering culture in general (Bass, Clements, and Kazman 1998). Because an architecture dictates the structure of a software system, it influences the structure of the development project with respect to team formation, schedules, and budgets. If the architecture is seen as successful, it might be used in future projects, and then the team structure might become a permanent part of the organization. The software might also offer a foothold in a certain market area, and this way the organization could update its goals

to take advantage of this new possibility. Customers will likely grow fond of a good software product, which lowers the threshold for reusing a battle-tested architecture, because this is more economical than building a new one from scratch. The architects themselves gain more experience from this specific system, which means that the system's success or failure will affect architectural choices made in future projects. When looking through a wider lens, an especially novel and successful architecture might influence the software engineering culture as a whole, outside of the organization from which it originated from. Imagine, for example, the tides created by the first relational databases, spreadsheets, windowing systems, or WWW itself. (Bass, Clements, and Kazman 1998, 10–11).

2.1.2 Software development life cycle

In addition to every software system having an architecture, every software system also has a life cycle (ISO/IEC/IEEE 12207:2017(E) 2017, 17). A life cycle model is a framework that contains the processes, activities and tasks used during the life of the system, from initial requirements gathering to the termination of its use (ISO/IEC/IEEE 42010:2007 2007, 3). The specific life cycle stages depend on the chosen life cycle model, of which there are many to choose from: incremental, spiral, iterative, evolutionary, and so on. Usually there is an initial planning stage, where the need for a new software system arises. After the need is verified, requirements are gathered. Next there might be a prototyping phase where multiple architectures are rapidly mocked and evaluated. Once a rough architectural direction is chosen, an initial architectural definition is formed. After it is validated, there often follows an iterative cycle where analysis, design, implementation, verification, validation and delivery are repeated and intertwined. (ISO/IEC/IEEE 12207:2017(E) 2017, 18).

Koskimies and Mikkonen (2005) suggest an architecture-oriented software development process. It highlights architectural design and evaluation as a crucial life cycle stage before moving onto detailed design and implementation. Architecture should be devised incrementally and iteratively from the requirements most relevant to architecture. Usually it is good to start from functional requirements, after which the non-functional qualities must be pursued. (Koskimies and Mikkonen 2005).

In this thesis I don't cling to any particular software lifecycle model. Instead I rely on the general notion that the initial architecture definition and architectural design happens early on in the life cycle of a software system. If using an iterative life cycle model, the requirements, architecture or detailed design might naturally need refinement also later on during operations and maintenance (ISO/IEC/IEEE 12207:2017(E) 2017). I will focus on the early choices and design decisions that influence the activities for the rest of the software lifecycle. We can call these architectural activities as "architecting", which includes defining, documenting, maintaining, and improving an architecture, and certifying its proper implementation (ISO/IEC/IEEE 42010:2007 2007, 3). In addition to these architecting activities, Bass, Clements, and Kazman (1998, 12) mention initially creating a business case for the system, before any architecture is created.

The standard ISO/IEC/IEEE 12207:2017(E) (2017) defines processes related to architecture and design. Next I will make a distinction between two related but separate processes: the Architectural Definition process, and the Design Definition process. The purpose of the Architectural Definition process, as outlined in ISO/IEC/IEEE 12207:2017(E) (2017, 66), is threefold. First, the architect should create alternatives for a system's architecture. Then, we should choose one or multiple alternatives that properly frame the stakeholders' needs and fulfill the requirements placed on the system. Lastly, all of this should be described with a set of uniform architectural views.

At a lower abstraction level, the Design Definition process (ISO/IEC/IEEE 12207:2017(E) 2017, 71) is expected to provide enough detailed data and information about the system and its components for implementation to be in line with the architectural entities. It involves preparing for software system design definition by e.g. prioritizing design principles and characteristics, identifying and planning for necessary enabling systems and services, and obtaining access to those enabling systems and services. Then we should establish designs related to each software element. This goes deep in the technical details, such as database structures, provisions for memory and storage, software processes, and external interfaces. (ISO/IEC/IEEE 12207:2017(E) 2017, 72).

As I mentioned above, according to ISO/IEC/IEEE 12207:2017(E) (2017, 71–72) selecting the technologies for implementation and considering compatibility with technologies also

falls under the Design Definition process. In my experience with public cloud providers such as Amazon Web Services, Microsoft Azure, and Google cloud, the infrastructure technologies are decided by architects instead of software developers. The architect wants to know which cloud platform is preferred by the customer, in order to reduce the technology space to a reasonable set of alternatives provided by the platform. The architect can then work with the available services to compose an architecture that best serves the stakeholders' needs. The finer details, such as the choice of programming language or code-level design patterns, can be left to software developers, as also implied by the ISO/IEC standards.

Maintenance stage is a phase in software life cycle that is often considered less glamorous than working on greenfield projects and building the next software hit. Despite this, it is generally accepted that a substantial portion of the total cost of software is spent on maintenance. Bass, Clements, and Kazman (1998, 32) propose that this figure is close to 80%. Obviously the actual numbers might be slightly different these days 25 years after their estimate, but the implication holds: considering maintenance efforts as early as possible in the software life cycle can pay off. Mumtaz, Singh, and Blincoe (2021, 1) echo that architectural issues should be tackled in a timely manner, because the cost of fixing those kinds of problems later can be exceptionally high.

The standard ISO/IEC/IEEE 12207:2017(E) (2017) outlines the maintenance process, which takes place later in the software life cycle, after Architectural Definition and Design Definition processes. Its purpose is to sustain the capability of the system to provide a service. This can be achieved by monitoring the system's capability to deliver services; recording incidents for analysis; taking corrective, adaptive, perfective and preventive actions; and confirming restored capability. In practice, preparing for maintenance should begin already during architecting. It is necessary to consider constraints placed on the architecture by maintenance-related requirements. This can be done, for example, by emphasizing encapsulation, modularity, and scalability. Documenting the architecture and system reduces the efforts required to reverse-engineer the system and components when a fix is needed. The decisions made during architecting can also affect the need and possibilities for e.g. remote diagnostics and maintenance, roll back, backup, and recovering data. (ISO/IEC/IEEE 12207:2017(E) 2017, 95–96).

Tasks performed during maintenance activities vary between systems and organizations. There is a need to review stakeholder requirements, complaints, events, incident and problem reports in order to identify the type of maintenance effort needed (ISO/IEC/IEEE 12207:2017(E) 2017). If the organization has adopted an iterative software life cycle model, software requirements might change during maintenance. This change in requirements can be a source for adaptive and perfective maintenance activities. In that case there will be a need to make corrections, improvements or other changes to software that is already deployed (ISO/IEC/IEEE 12207:2017(E) 2017, 97). As Bass, Clements, and Kazman (1998, 32) points out, having a good grasp of the architecture will help manage changes and reason about them.

2.1.3 General advice for achieving a good architecture

Because every system has an architecture, you can create an architecture through the simple method of trial and error. Unfortunately, an architecture built this way is unlikely to fulfill the requirements and goals placed on it (Bass, Clements, and Kazman 1998, 24). They outline some helpful advice for creating an architecture from the correct principles. These rules of thumb are divided into two categories: process-related and structure-related advice.

First we will delve into the process-related topics. As a general rule, there should be a single architect or a team of architects with a designated leader (Bass, Clements, and Kazman 1998, 15). Functional requirements and a prioritized list of quality attributes should be provided, so that the architecture can be steered towards achieving them. The architecture should be formally evaluated for quality attributes, and analyzed for applicable quantitative measures, such as maximum throughput. Having an architecture support iterative implementation makes integration and testing easier: a "skeleton system" can prove that component communication works, even with otherwise minimal functionality. The architect is also responsible for defining, distributing, maintaining, and enforcing resource constraints to the responsible implementation teams. For example, there might be limits to network traffic or hardware performance that the teams must heed. A good documentation is also a must, with at least one static view and one dynamic view, written in a notation that all stakeholders can understand with minimal effort. Speaking of stakeholders, they should be actively partici-

pating in evaluating the architecture, to ensure their needs are being met. (Bass, Clements, and Kazman 1998, 15).

Next in line are the structure-related pieces of wisdom. Bass, Clements, and Kazman (1998, 16) state that modules should be well defined, with a clear separation of concerns and following the principles of information hiding. Their interfaces should also be well defined, and encapsulate any details that might change during implementation, such as choice of data structures. To improve modifiability, the modules that produce data should be kept separate from those that consume data. Between the modules, there should be only a small set of different ways of interaction, which is to say, the system should perform similar tasks in a similar fashion all around. This helps improve understandability, reliability, and modifiability, with the added benefit of shorter development time and conceptual integrity of the architecture. Each quality attribute should be strived towards with well known architectural tactics that suit it. The architecture should not depend on any specific version of a commercial product or tool, or if it does, changing the product or tool should be both straightforward and cheap. If working close to hardware, changing the processor allocated for each task or process should be easy even at runtime. In case of parallel-processing systems, special care should be given to processes or tasks that touch multiple modules. (Bass, Clements, and Kazman 1998, 16).

2.2 Maintainability

2.2.1 Quality Attributes

We of course expect that the software system fulfills all its requirements, gives us correct and accurate results, and that it interacts with other systems in the expected manner (Bass, Clements, and Kazman 1998, 76). Additionally, there usually is an implicit or even explicit goal of creating software of high quality. According to Gorla and Lin (2010, 602), this quality can be influenced by factors that are organizational, technological, or end user-related. Organizational factors can include the IT budget, number of people in system development in that organization, and quality of documentation. Level of user involvement and user trainings for the systems are end user-related factors. In contrast, technological factors include

experience and skill level of the development staff, and type and suitability of development method, programming language or database model (Gorla and Lin 2010, 602). Few factors outside the technological category are relevant in this thesis, which is why I mostly – although not exclusively – focus on factors that can be influenced by the software architect.

Quality attributes are used to describe specific aspects of quality as it relates to software. Bass, Clements, and Kazman (1998, 76) suggest that there are two categories of quality attributes that can be measured on architectural level: attributes that can be observed by running the system, and those that cannot. The first category tells us how well the system fills its behavioral requirements during runtime. Quality attributes that indicate this include (Bass, Clements, and Kazman 1998, 79–81; Li et al. 2021, 13; Arvanitou et al. 2017, 60):

- performance
- security
- availability
- functionality
- usability
- monitorability
- reliability
- safety
- fault prediction
- defect proneness

The second category of quality attributes — those that cannot be observed by running the system — tell us how easy the system is to test, integrate and modify (Bass, Clements, and Kazman 1998, 81). They are more related to technical implementation than the system’s behavior. It contains, among others, these attributes:

- modifiability
- portability
- reusability
- integrability
- testability

Other authors use related terms "functional" and "non-functional requirements". Bass, Clements, and Kazman (1998, 76) consider this dichotomy useless. They argue it suggests that all attributes can be divided strictly into those that model the system's ability to perform the actions it is expected to, and "everything else". The latter non-functional category acts as a bucket. This too easily leads to a situation where some qualities are given less consideration, leading to a diminished quality of the system as a whole. In fact, many quality attributes are intertwined to behavior of the system. For example, it might be impossible to reach lightning fast performance with a feature that processes very large images. On a higher level, though, Bass and others consider functionality to be orthogonal to quality attributes. If this wasn't the case, implementing a certain functionality would always dictate the system's security, usability, or other attributes. The program might function perfectly, but it might also be difficult to modify or costly to build. (Bass, Clements, and Kazman 1998, 77).

Instead, architectural choices made by the architect decide the level of quality possible for the system (Bass, Clements, and Kazman 2003, 72). Quality must be considered at every stage of design, implementation and deployment. Different quality attributes represent themselves differently at different stages. For example, some usability decisions, such as details of a UI, are not architectural at all. Modifiability, on the other hand, is primarily architectural. Modifiability relates to how a system is split into components, and a system is modifiable if changes do not require touching multiple components. Performance depends on both architectural and non-architectural considerations. Architect can influence how functionality is divided between each component, and how they communicate. But the finer details are left to developers, who are usually responsible for detailed design. An example of this is the choice of an algorithm and turning it into source code. Another thing to consider is that quality attributes do not exist in a vacuum, but instead influence each other in good or bad (Bass, Clements, and Kazman 1998, 78). For example, improved monitorability likely enables improved testability, which in turn leads to an increase in security (Li et al. 2021, 18). The architect must prioritize the quality attributes and decide which ones to emphasize (Bass, Clements, and Kazman 1998, 129).

2.2.2 What is maintainability, and why is it important

The quality attribute maintainability describes the efforts needed to fix errors or to make enhancements to the system (Gorla and Lin 2010, 603). Based on their mapping study on design-time quality attributes and metrics, Arvanitou et al. (2017) found that maintainability is the most often researched quality attribute in most stages of the software development life cycle; it receives interest in design, architecture, implementation, and maintenance stages. In other stages such as project management, requirements, and testing, maintainability has not been subject of much research. (Arvanitou et al. 2017, 61–62).

Like many other quality attributes, maintainability can be further subdivided into more specific quality attributes, some of which we have discussed before. The standard ISO/IEC 5055:2021(E) (2021, 233) lists modularity, reusability, analyzability, changeability, modification stability, testability, and compliance as components of maintainability. Many of these are also echoed by Arvanitou et al. (2017, 61). Some other categories, such as understandability, are sometimes related to maintainability. Arvanitou et al. (2017) consider it a separate quality attribute for two reasons. First, some other quality models make a distinction between maintainability and understandability. Second, a unique research community has already gathered around understandability as a standalone concept. (Arvanitou et al. 2017, 61).

There is no single metric "maintainability of the software system", because different components of the system likely have different purposes and implementations (Broy, Deissenboeck, and Pizka 2006, 26; Bass, Clements, and Kazman 1998, 192). As such, a phrase like "the system is highly maintainable" doesn't tell anything practical about the system, because quality attributes receive their meaning from surrounding context. For example, the software might be highly modifiable in relation to specific classes of changes. (Bass, Clements, and Kazman 1998, 192).

Focusing on maintainability can pay back the cost invested into it. It influences how much money and time needs to be spent on maintenance tasks, so if maintainability is not considered, cost of change can be unnecessarily inflated (Broy, Deissenboeck, and Pizka 2006, 21)[2]Vale2022. Personnel requirements have to also be taken into account, because some-

one needs to perform the necessary updates and fixes. Wiggins (2011) warns that if maintenance work is left to the one single team of software developers, the time is often deducted from development of new features.

If a system is not maintainable, it will be more difficult to address other problems, such as those related to performance or reliability (Bouwers and Deursen 2010, 46). If different services are highly coupled, one of them cannot be altered without changing others that are depending on it (Vale et al. 2022, 2). At the business level, maintainability of an important system defines the organization's ability to adapt their business processes to changing market circumstances, and to implement innovative products and services (Broy, Deissenboeck, and Pizka 2006, 21).

Even if the software itself stays unchanged for years, it can break if and when its environment changes. This phenomenon is called software rot³. In other words, software rot happens when a system's assumptions go out of date. A good example of this is the infamous "Y2K" event which resulted in uncountable bugs when two-digit year counters underwent wrap-around at the start of the new millennium. Even outside such one-off situations, software can touch the underlying operating system and network infrastructure in many places: library versions, directory paths, IP addresses and hostnames, and so on (Wiggins 2011). The Jargon File⁴ entry about software rot advises to employ robust design to deflect such problems. Even if the Jargon File is not a scientific source of wisdom, based on what we've discussed in this chapter, I consider taking care of maintainability part of robust design.

2.2.3 How to measure maintainability

As discussed before, problems discovered early in the software development life cycle are easier to address (Bass, Clements, and Kazman 1998). At that point you still have a chance to adjust requirements, specification or design with relatively low cost. Software quality cannot be "appended" to the project afterwards, but must instead be baked in deliberately. Qualities of a system can be anticipated by evaluating the architecture. This means we should evaluate

3. <http://www.catb.org/jargon/html/S/software-rot.html>.

4. The Jargon File is a "comprehensive compendium of hacker slang illuminating many aspects of hackish tradition, folklore, and humor." <http://www.catb.org/jargon/html/online-preface.html>

the architecture as soon as feasible. To evaluate a certain quality attribute, we must make sure we have the necessary architectural descriptions at hand. For example when evaluating modifiability — a subcategory of maintainability — we need at least a decomposition that shows us the different modules and how tasks are distributed between them. (Bass, Clements, and Kazman 1998, 32, 190–191).

It is necessary to decouple quality attributes from effort estimations in our daily work. Moses (2009) points out that maintainability of a software module is intended to signify possible challenges posed by code for maintenance tasks. Maintainability cannot give you an estimate of how difficult a certain maintenance change will be or how long it will take. Unless the same change has been performed in exactly the same circumstances before, duration and possibility of success of a code change are only subjective estimates. Moses suggests two methods for estimating practical maintenance tasks. It can be given on an ordinal scale of e.g. 1-5 instead of an absolute time estimate. If a concrete estimate is necessary, two experts should voice their opinions, from which we can calculate an expected duration with confidence limits. (Moses 2009, 204).

Vice versa, a prediction system for maintenance effort is not a prediction system for the quality attribute maintainability (Moses 2009, 206). Each maintenance effort includes factors that influence the duration of the activity. These factors include availability of documentation, amount of testing after the change, and skills of the maintenance team. Because it is obviously impossible to carry out and measure all possible maintenance activities for a system, the effort can only be measured from a subset of all possible maintenance activities. This means that, even in the best scenario, a system's maintainability can be estimated only for activities that are similar to ones that have already been performed before to this system. In the worst case, the measured effort can be misleading, if measured activities do not represent actions that are carried out in practice. (Moses 2009, 206).

Broy, Deissenboeck, and Pizka (2006) state that an evaluation criteria must be well-founded and checkable. Evaluating a whole software system is quite involved, and requires multiple kinds of evaluation methods. Source code can be automatically checked for syntactical properties. Other methods are semi-automatic, requiring some human actions. The rest of the methods are thoroughly manual, such as reviewing documentation, or evaluating whether

the system employs correct data structures (Broy, Deissenboeck, and Pizka 2006, 22). Our focal point of interest, maintainability, is a high-level quality attribute, because it cannot be directly calculated from metrics (Arvanitou et al. 2017, 60).

In their study about designing microservice systems using patterns, also Vale et al. (2022, 9) found that the participants do not measure quality attributes directly. Instead, the participants measured quality attributes as a sum of lower-level attributes or processes related to the quality attributes in question. Such lower-level attributes include measuring code coverage, time taken to resolve issues, and analysis of architectural debt, which all are indirect metrics for maintainability. Vale and others reason that this might be because lower level attributes are easier to measure and check directly than maintainability, which is in line with Arvanitou et al. (2017). They also postulate that subjective understanding of quality attributes is very much personal and often vague. This makes it difficult to mutually agree on a definition inside an organization. Instead, participants preferred to think of techniques and tools they used to tackle the related problems. (Vale et al. 2022, 7–10).

Automatic methods for assessing maintainability of a system usually limit themselves to source code level. Such methods cannot be used to to evaluate an architecture before implementation has taken place. Despite this, they are useful after the fact, for example when following an iterative life cycle model. The standard ISO/IEC 5055:2021(E) (2021, 1–2) points out that faults in architectural or design practice can also be detected by statically analysing design specifications. This requires that the specification is written in a design language with a formal syntax and semantics. The standard presents a long list of "Automated Source Code Quality Measures" (ASCQM) that provide strong indicators for the quality of a software system in each area. In my opinion, the following maintainability weaknesses are related to architecture (ISO/IEC 5055:2021(E) 2021, 38–47):

- invocation of a control element at an unnecessarily deep horizontal layer (layer-skipping call)
- invokable control element with excessive file or data access operations
- invokable control element with large number of outward calls (excessive coupling or fan-out)
- modules with circular dependencies

The standard also provides multiple detection patterns (ISO/IEC 5055:2021(E) 2021, 2). They help detect structural weaknesses from code. One pattern can be used to detect multiple weaknesses. I would list the following maintainability detection patterns as related to architecture (ISO/IEC 5055:2021(E) 2021, 48–54):

- limit number of data access
- limit number of outward calls
- ban circular dependencies between modules

Arvanitou et al. (2017, 63) note that there are two categories of metrics for quantifying quality attributes: those that can be evaluated through a single metric, and those evaluated with multiple. Arvanitou et al. (2017, 66) list the following metrics relating to the architectural stage:

- cyclomatic complexity
- coupling factor
- coupling between objects
- number of modules
- information flow (fan-out)
- non-functional coverage
- absolute adaptability of a service
- relative adaptability of a service
- mean of absolute adaptability of a service
- mean of relative adaptability of a service
- level of system adaptability

Most of the tools to calculate these metrics are unnamed, likely created in academia by research teams, and not available for public use. Maintainability is often connected to metrics, but is not always quantified as their function. Evaluating a quality attribute with a single function is non-trivial, which is why few research have suggested such a function for high-level attributes like maintainability. (Arvanitou et al. 2017, 67).

2.2.4 How to improve maintainability

In this section I write about ways to improve maintainability for a software system. For example, there will be less need to reverse-engineer systems and elements during maintenance work if we documenting the system is an explicit requirement during design (ISO/IEC/IEEE 12207:2017(E) 2017, 96). Utilizing the available standards and system development methodologies helps create systems that are easily understandable. This results in systems that require less effort to maintain. (Gorla and Lin 2010, 608).

Broy, Deissenboeck, and Pizka (2006) remind us to separate "facts" from "acts". An activity, or act, is something that we try to do. In context of software maintenance, example acts can include concept location, impact analysis, coding, and modification. A fact is some circumstance surrounding the system, such as readability of documentation, processes used in the organization, or availability of debugging tools. Maintainability also depends on things outside of the product itself. For example organizational facts are not part of the product, and often overlooked when assessing maintainability of a system. Skill of the maintainers plays a key role in maintainability. A high turnover rate can make a system more difficult to maintain, if the company loses a large share of its developers to a competitor. Well-defined processes, such as a configuration management process, can help create a quality product. Trustworthy tools like debuggers, and visualization and refactoring tools are critical for performing effective maintenance. (Broy, Deissenboeck, and Pizka 2006, 25–26).

Bass, Clements, and Kazman (1998) collate that the most important thing in improving non-runtime attributes like modifiability and testability is to enable and facilitate making changes to the system. The system likely suffers from poor modifiability if it relies on a shared memory solution with global variables accessible to all components, where a change to data format could require changes to all components (Bass, Clements, and Kazman 1998, 89). A change should only affect a small number of components. ISO/IEC/IEEE 12207:2017(E) (2017, 96) concurs that systems which highlight encapsulation, modularity, and scalability can be simpler to maintain. Achieving this requires foresight into what kinds of modifications will be necessary in the future. The foresight can be achieved by an experienced architect with available change histories for this and related systems. (Bass, Clements, and Kazman 1998, 118).

To improve maintainability, Vale et al. (2022, 7) suggest tracking bug reports with *Slack* and *Jira*. Static code analysis tools such as *SonarQube* and *Codacy* help on the source code level. Code reviews, automated pipelines, and domain-driven design are also beneficial. According to their research, there are design patterns one can use to gain an edge with maintainability. Strangler, Ambassador, External Configuration Store, Gateway Offloading, Backends for Frontends, Pipes and Filters, and Static Content Hosting patterns were reported by at least one interviewee to have improved maintainability of a system they have worked with. (Vale et al. 2022, 4–6). Design patterns create a common language between designers, to share an understanding of solutions that have been found to work well.

Related to design patterns, Bass, Clements, and Kazman (1998) talk about "Unit Operations", which are general design operations or techniques that allow the architect to achieve quality attributes. Unit operations have their roots in the engineering world, and are at a higher level of abstraction than design patterns. Examples of unit operations are Separation, Uniform Decomposition, Replication, Abstraction, Compression, and Resources Sharing. The unit operations Separation and Resource Sharing can have a positive effect on maintainability by improving modifiability. In Separation, functionality is detached as its own component with a well-defined interface to its surroundings. This has numerous advantages: it enables distribution, which in turn enables parallelism; it's easier to divide development work to development teams; and it improves modifiability and portability. Resource Sharing means that data or services are encapsulated and then shared between multiple independent consumers. Examples of this are databases, blackboards, and servers in a client-server-architecture. This unit operation decreases coupling between components, which leads to improved integrability, portability and modifiability. (Bass, Clements, and Kazman 1998, 123–126).

The *Heroku* cloud application platform (Wiggins 2011) utilized "Explicit Contracts" to maintain a strong separation between the platform and the applications running on it. This results in what they call "erosion-resistance", which is synonymous to improved maintainability. An example of an explicit contract is that in a *Ruby* application, the developer must use a *Gem Bundler* and a *Gemfile* to declare dependencies of their application. In a *Node.js* application this is achieved via *package.json* file instead. Other contracts dictate how environment

variables should be used, how to implement logging, and how to tell the platform how you want your application to be launched. Adhering to these explicit contracts allows Heroku to change the infrastructure and platform without breaking applications running on it. (Wiggins 2011).

Different quality control activities should be performed at different intervals (Broy, Deissenboeck, and Pizka 2006, 25). To avoid redundancy, copy-pasted sections of new code should be checked up to daily. Then again, it is enough to review documentation at certain points of the development process, and not even weekly. This is mainly because manual quality checks are laborious, and thus costly. Whenever feasible, manual review should be supported or replaced with automated tools to as high a degree as possible. This makes good quality check tooling valuable, because we can receive high-quality assessments more often. (Broy, Deissenboeck, and Pizka 2006, 25).

A system might be labeled as a "legacy system" due to factors that do not negatively affect usage or functioning of the software (Broy, Deissenboeck, and Pizka 2006, 22). Such factors can be the coding style, or the implementation language. A label of "legacy" hints at unmaintainability, and the will to replace the system with something else. Broy, Deissenboeck, and Pizka (2006, 22) argue that a legacy system should not be replaced unless the new system increases business value or maintainability. Li et al. (2021, 18) concurs: a new architecture could mitigate some problems of the legacy system, but implementing it properly will cost time and effort.

2.3 Cloud-nativity

2.3.1 What is cloud-nativity

The cloud is a distributed architecture of individual cloud-native services, providing resources as services in a tiered fashion to construct a full technology stack from hardware to middleware platforms to applications (Pahl, Jamshidi, and Zimmermann 2018, 1). Birth of the public cloud can be dated back to 2006, when *Amazon Web Services* launched their *Simple Storage Service (S3)* and *Elastic Compute Cloud (EC2)* (Kratzke and Quint 2017, 1). The so-called first wave of cloud computation meant that companies replaced their own

data centers with virtual machines (VMs) running on data centers owned by cloud platform providers (Gannon, Barga, and Sundaresan 2017, 17). While revolutionary, it was still difficult to support both scalability and security at the same time, and there was no cloud-based tooling for data, event or error correction to be used in debugging. (Gannon, Barga, and Sundaresan 2017, 18).

”Cloud-native” is a term often used, but rarely elaborated beyond ”we used the cloud instead of our own data center” (Gannon, Barga, and Sundaresan 2017, 17). Terminology and related technology have gone through many steps to get where we are now, touching Service Oriented Architecture (SOA), VMs, cloud computing, containers, and microservices (Kratzke and Quint 2017, 3). Cloud-nativity is an approach that only makes sense on top of cloud infrastructure, which means that the physical environments have been virtualized, and infrastructure is disposable. When using disposable infrastructure, units of infrastructure are created, scaled, and destroyed quickly with help of automation. There is a fitting metaphor of ”cattle vs. pets”, where disposable infrastructure is the cattle: all instances are identical, and if one of them requires an update or repairs, it will be replaced instead of fixed. (Vettor, Pine, Jain, et al. 2022).

Kratzke and Quint (2017, 8) searched ”cloud-native” from *Google Trends* for an overview of its usage. The term was widely used around 2007, around the first wave of cloud computation. Usage tapered after that, but around 2015 the term enjoyed newfound interest. The researchers admit that Google Trends is not the most reliable of metrics due to its sensitivity to industry buzzwords. These results, however, correspond to the results of their research. ”Cloud-native” was first mentioned in academic literature in 2012^{5 6}. (Kratzke and Quint 2017, 8).

Vettor, Pine, Jain, et al. (2022) offer a definition for ”cloud-native”:

Cloud-native architecture and technologies are an approach to designing, constructing, and operating workloads that are built in the cloud and take full advantage of the cloud computing model.

5. Andrikopoulos, Fehling & Leymann, 2012. Designing for CAP - The Effect of Design Decisions on the CAP Properties of Cloud-native Applications

6. Garca-Gmez, et al., 2012. 4CaaS: Comprehensive Management of Cloud Services through a PaaS

In my opinion this definition is vague, with the main conceptual burden carried by "cloud" and "cloud computing model". Patrizio (2018) offers additional details to the definition: cloud-native is a modern way to build and run software systems, and leverages the elasticity, scalability and resilience offered by cloud computation. A still more specific definition comes from Kratzke and Quint (2017), whose systematic mapping study aimed to collate the research trends around cloud-native software development practices. They also wanted to define the Cloud-Native Application (CNA), for which the resulting definition is as follows (Kratzke and Quint 2017, 13):

A cloud-native application (CNA) is a distributed, elastic and horizontal scalable system composed of (micro)services which isolates state in a minimum of stateful components. The application and each self-contained deployment unit of that application is designed according to cloud-focused design patterns and operated on a self-service elastic platform.

From this we can conclude that a CNA is not merely a distributed system, but especially distributed, horizontally scaling, resilient, and load-adaptive (Kratzke and Quint 2017, 13). Scalability is at the heart of cloud-native software (Noval et al. 2022). A CNA must be more than "available via the internet": it must achieve global scale by serving thousands of concurrent users (Gannon, Barga, and Sundaresan 2017, 17). If you think of global, culture-penetrating companies, you might come up with *Netflix*, *Spotify*, *Uber*, *Airbnb*, *Facebook*, or *Twitter* (Gannon, Barga, and Sundaresan 2017; Patrizio 2018).

Companies such as *Microsoft*, *Google*, *Amazon*, *Oracle*, *IBM*, *Alibaba*, *Heroku*, and *Redhat* provide public cloud platforms for other companies to build cloud software on. They are called Cloud Service Providers (CSP). CSPs' own software products might utilize the CSP's own infrastructure, which makes sense for simplicity of logistics. Some services such as *AWS Kinesis*, *DynamoDB*, *SQS*, and *Amazon Redshift*; *Microsoft Azure CosmosDB* and *Data Lake*; and *Google BigQuery* are themselves cloud-native under the hood, and used to build CNAs. (Gannon, Barga, and Sundaresan 2017, 17).

CNAs share some common features. They are expected to be always running. At the same time, a CNA must expect the physical infrastructure to vary, which leads to related intermit-

tent errors (Gannon, Barga, and Sundaresan 2017, 17). This is different from single-machine applications, where we can usually assume that the platform — which consists of hardware and an operating system — stays the same during application runtime and always works. For CNAs, the reality is that once the application starts receiving a sufficient number of requests, something is always either breaking or broken. It is necessary to architect the software in such a way that updates and testing are possible without disruptions to the production instance of the software. Security should also be an integral part of the architecture, due to multiple small components, and a need to manage access on multiple levels of the software. A simple firewall is not enough. (Gannon, Barga, and Sundaresan 2017, 17).

CNAs utilize technology and techniques suitable for dynamic cloud environments. These cloud environments can be either public, private or hybrid (Noval et al. 2022). The techniques include microservices, containers, service meshes, immutable infrastructure, declarative APIs, continuous delivery technology, backing services and automation (Noval et al. 2022; Patrizio 2018; Vettor, Pine, Jain, et al. 2022). Methodologies such as DevOps and agile are often adopted when building CNAs (Patrizio 2018). Other CNA development methodologies are often based on design patterns (Kratzke and Quint 2017).

The three most common ways to build CNAs are microservices architecture, fully managed high-abstraction services, and serverless computing (Gannon, Barga, and Sundaresan 2017, 17). In fully managed services, the service handles everything outside of business logic: infrastructure, management, scaling, and so on. Microservices architecture is the most common approach. Serverless computing means utilizing CSP-provided services to build software without VMs and sometimes even without containers. (Gannon, Barga, and Sundaresan 2017, 17). I will discuss microservices and serverless in the following chapters.

2.3.2 Microservices

Microservices architecture (**MSA**) is an increasingly accepted and adopted architectural style, and is much used in modern softwares (Vale et al. 2022, 10). It was the first major style for cloud-native applications, dating back to 2013 (Gannon, Barga, and Sundaresan 2017, 18). Scalability is important when moving to a cloud-first paradigm, and MSA can

help overcome the limitations of traditional monolithic systems related to scalability (Vale et al. 2022, 10). Monolithic software is often built using layered architecture, relational databases shared across all services, and is run as a single process (Vettor, Pine, Jain, et al. 2022). MSA on the other hand encourages you to implement small-scale, independently distributed services instead of tying up all functionality into a single monolith (Li et al. 2021, 1). When using MSA, you will disassemble the application into its basic building blocks by encasing a single functionality into a single service (Gannon, Barga, and Sundaresan 2017, 18). This also allows for technological heterogeneity, when different services can use the technologies that best allow it to reach its goal (Li et al. 2021, 18). Together these services form the complete application (Vettor, Pine, Jain, et al. 2022).

A single microservice implements a specific business functionality as part of a larger domain context (Vettor, Pine, Jain, et al. 2022). This context should be limited, giving the service limited responsibility, with only a small number of dependencies on other services (Gannon, Barga, and Sundaresan 2017, 18). According to Vettor, Pine, Jain, et al. (2022), the service should contain the necessary logic, state, data, external dependencies and programming platform (Vettor, Pine, Jain, et al. 2022). Meanwhile, Gannon, Barga, and Sundaresan (2017, 18) recommend that each service should pursue statelessness, where the state is not saved into the microservice itself, but in another service, such as a database, cache or directory. In my opinion these differing views on handling state are not conflicting. Absolutes are harmful, and as prescribed in the CNA definition by Kratzke and Quint (2017, 13), the main idea is to isolate state in a *minimum* of stateful components.

The microservice should run its own process independently of other active services, and communicate with them using standard messaging protocols, such as *HTTP/HTTPS*, *gRPC*, *WebSockets* or *AMQP*. Every microservice should be able to be separately managed, replicated, scaled, updated, and deployed (Gannon, Barga, and Sundaresan 2017, 18). Vettor, Pine, Jain, et al. (2022) elaborate that independent scaling enables a more accurate control over the system. At the same time this lowers the total costs involved, because there is no need to scale services with lower usage alongside those that need it. Autonomous development and deployment of the separate services removes the need to wait for e.g. quarterly release schedule to roll out an anticipated feature or fix. Independent deployment also re-

duces the risk of total system malfunction, as the changes are isolated to smaller context. (Vettor, Pine, Jain, et al. 2022). Security must also be considered. One microservice should only have access to select few other microservices, and the target service must verify access rights of the requester service. Using role-based access control such as *OAuth*, *AWS IAM*, or *Azure RBAC* can be instrumental in achieving this. (Gannon, Barga, and Sundaresan 2017, 18).

Gannon, Barga, and Sundaresan (2017, 18) remember a time when a need arose to pack microservices into spaces smaller than an entire VM image. When using containers, the code, dependencies, and runtime are packed into a container image (Vettor, Pine, Jain, et al. 2022). The images are then stored in a container registry, where they can be fetched when deploying a service that uses them. It is possible to run multiple containers in the cloud on a single VM, and achieve startup times even lower than one second (Gannon, Barga, and Sundaresan 2017, 18). Other reasons to select containers as a basis for your application include portability, and lack of a need to configure every environment separately with the required frameworks, libraries, and runtime engines. This helps to achieve uniform environments with greater accuracy and speed. (Vettor, Pine, Jain, et al. 2022). Companies such as *Docker* popularized the container technology for use in MSAs (Gannon, Barga, and Sundaresan 2017, 18). According to Vettor, Pine, Jain, et al. (2022), *Docker* is the most popular container vendor, and it has become industry standard for packaging, deploying and running cloud-native software.

Containers are often arranged in a service mesh, using a container orchestration solution like *Kubernetes*, *Apache Mesos*, *Docker Swarm*, *Azure Service Fabric* or *IBM Blue Container Service* (Gannon, Barga, and Sundaresan 2017, 18). The orchestrator takes care of things like scheduling, affinity, health monitoring, failover, scaling, networking, service discovery, and rolling upgrades (Vettor, Pine, Jain, et al. 2022). Scheduling means automatic provisioning of the container instances. Affinity and anti-affinity describe how far away the containers are from each other, measured in respect to hardware topology. Health monitoring allows for automatic detection and recovery from failures. If a failure happens, a failover process will automatically provision the failed instance to a healthy host. As we can extrapolate from previous definitions of the word, scaling means automatically adding or removing container

instances depending on the demand. The orchestrator handles the networking overlay for container communication. Service discovery makes it possible for containers to find each other. Lastly, the orchestrator coordinates incremental updates and rollback of problematic changes to achieve rolling upgrades. (Vettor, Pine, Jain, et al. 2022).

Jumping aboard the MSA bandwagon is not all sunshine and roses. Li et al. (2021, 18) point out that if the monolith is complex to begin with, componentizing it will bring this inner chaos on display. The end result might still be complex, even if the complexity is now slightly different from previous state of affairs. Gannon, Barga, and Sundaresan (2017, 20) remind that we need to allocate efforts to manage computation resource clusters in order to run our services. As I see it, while the container orchestrator service take care of automatic operations day-to-day, the initial setup is non-trivial, and monitoring is still required.

Vettor, Pine, Jain, et al. (2022) presents some challenges for MSA systems. I'd like to reframe these challenges as things to consider when architecting such systems, rather than strictly negative aspects of the MSA approach. The listed challenges are most often related to communication, resiliency, or distributed data. Firstly, because microservices need to communicate via APIs and network protocols instead of in-process, the architect should account for network congestion, latency and intermittent errors (Vettor, Pine, Coulter, Schonning, et al. 2022). It is especially important to retry failed requests, which is made easier if all state-altering requests are idempotent. There will also be some overhead for every service from serializing, deserializing, decrypting, and encrypting messages. (Vettor, Pine, Coulter, Schonning, et al. 2022).

While the cloud platform will automatically handle most of the resiliency issues that CNAs face, the architect still needs to consider them while designing the system (Vettor, Pine, Coulter, Wenzel, et al. 2022). In my view, this means minimizing disruptions caused by A) the issues themselves, and B) the actions that the platform takes to respond to the issues. The platform might restart and migrate a failed process to a new host. A container orchestrator might be in the middle of a rolling upgrade, or move the service from one node to the next, for one reason or another. (Vettor, Pine, Coulter, Wenzel, et al. 2022).

When following the MSA model to the letter, we will end up with distributed data. Because

all data is no more in a single shared database, it will be more difficult to compose queries across multiple services (Vettor, Pine, Coulter, Victor, et al. 2022). The architect must also remember that the ACID principles do not hold across the entire system at any one point in time. This is because propagating a change of state through multiple data stores will take some time. On the plus side, each data store can be individually scaled and developed without direct effect on other services. Each of them can select a data model that best suits that specific service: relational, document, key-value, graph, or some other database paradigm. This boosts the agility, performance, and scalability of the system. (Vettor, Pine, Coulter, Victor, et al. 2022).

2.3.3 Serverless

Cloud-nativity doesn't constrain itself to MSA. Another relevant approach is striving for the so-called serverless architecture. The title "serverless" is something of an oxymoron, because there is a server, but it's not you who is responsible for it. Roberts (2018) puts it succinctly: though it might mean zero system administration work for you, the abstraction might leak at times, and then you'll become aware that a Systems Administrator is supporting your application somewhere. Broadly speaking, a serverless architecture can be achieved with two approaches: Backend-as-a-Service (**BaaS**) or Functions-as-a-Service (**FaaS**). (Roberts 2018).

With BaaS, you manage server-side logic and state by significantly or fully incorporating third-party cloud-hosted applications and services (Roberts 2018). These services handle everything else but your business logic: infrastructure, management, scaling, and so on. When such services are combined, we will end up with an application that fulfills the criteria for cloud-nativity (Gannon, Barga, and Sundaresan 2017, 20). Sometimes BaaS is used with the prefix "Mobile" BaaS, due to their popularity with application backends written for mobile devices. Examples of BaaS are *Parse* and *Firebase* databases, and *Auth0* and *AWS Cognito* authentication services. (Roberts 2018).

When going down the FaaS route instead, you use stateless, event-triggered, ephemeral, fully managed compute containers to run your server-side logic (Roberts 2018). Vendor

handles all the underlying resource provisioning and allocation by creating and destroying FaaS instances based on runtime need. FaaS instances are typically triggered as a response to inbound HTTP requests. Because you don't have control of the lifecycle of individual instances, you must store any persistent data to an external service, such as a database. While the management of the underlying compute resource is abstracted away to a very high degree, there's still monitoring, deployment, networking, support, and production debugging to handle (Roberts 2018). Examples of FaaS services are *AWS Lambda*, *Azure Functions*, *Google Cloud Functions* and many others. They are often paired with serverless services such as *AWS SQS*, *Kinesis*, *Amazon ML* and *DynamoDB*, or *Azure Event Hub*, *Stream Analytics*, *AzureML* and *CosmosDB*. (Gannon, Barga, and Sundaresan 2017, 17).

The two approaches BaaS and FaaS are similar in that neither of them concern themselves with resource management. They are also frequently used together. For example, a FaaS function can rely on a BaaS database for state management. Roberts (2018) describes the two approaches as preferring "choreography over orchestration". I understand this phrase to mean that the control flow is born from interaction of the individual components, instead of keeping control in some centralized location. This means that each component is more architecturally aware, which is a common idea also with MSA applications (Roberts 2018).

Serverless architecture does have some drawbacks, some of which are inherent to the paradigm and some which can be at least partially mitigated with careful implementation (Roberts 2018). Inherent drawbacks cannot be entirely fixed, and need to be always considered. With BaaS and FaaS, you have to give up some control of your system to a third-party vendor. As features vary between vendors, you might suffer from vendor lock-in if migrating to another service would require significant changes to code, operational tooling, or even architecture. Because it is more economic, vendors might run multiple instances of software for several different customers on the same machine, possibly even on the same hosting application. This multitenancy might be unacceptable in strictly regulated domains. More vendors results in a larger surface area for cyber attacks, and more FaaS functions means that identity- and access management issues such as IAM policy errors are compounded. (Roberts 2018).

Roberts (2018) points out that implementation drawbacks can be mitigated by either your actions, or by the vendor improving their service. Immature platforms might offer less

options for configuration, monitoring, deployment, packaging, versioning, and debugging (Roberts 2018). Execution duration for FaaS functions is limited, which might require you to re-architect tasks that require long running times. It is entirely possible for you to run a denial-of-service (DoS) attack on yourself unwittingly if you accidentally or intentionally launch a large number of instances, because e.g. AWS has a limit for maximum number of running instances per cloud account. (Roberts 2018).

Testing can become a pain point. Sometimes a BaaS vendor might not allow load testing at all, or will at least bill you dearly (Roberts 2018). Setting up and tearing down state might be difficult with third-party services. FaaS approach will almost certainly lead to a high number of small units, and replicating the entire cloud environment on your local machine is impossible. This makes integration testing more important than with other architectural styles. (Roberts 2018).

Despite these drawbacks, serverless architecture brings many benefits to the table. Amazon Web Services (2023) advertise their serverless offerings to provide automatic scaling, built-in high availability, pay-for-value billing model, and built-in service integrations. It eliminates infrastructure management tasks, which increases agility and optimizes costs (Amazon Web Services 2023). Reduced operational and development costs make sense due to economies of scale and commodified services for common functionality (Roberts 2018). FaaS is especially economical for occasional or inconsistent traffic, because the FaaS service spins only the necessary number of instances up or down, and you often are billed down to microsecond accuracy. This way you don't need to overprovision a VM for peak traffic. For predictable and consistent traffic however a traditional VM might be cheaper, as then it is easier and safer to keep the computing hosts at high CPU utilization. Also, time-to-market can be considerably shortened because continuous iteration is easier due to simple and fast redeployment of individual components. (Roberts 2018).

2.3.4 Cloud architecture

A typical layered architecture might include separate layers for presentation, business logic, and data management (Pahl, Jamshidi, and Zimmermann 2018, 5). Even though the devel-

oper has no direct access to the hardware of the cloud platforms, cloud software is mapped onto physical tiers and their services. A physical tier means the physical location of the application execution, and the places where software layers are deployed and where they run. Each node in a cloud architecture should handle management and recovery tasks themselves, which is why cloud services are self-managed entities. (Pahl, Jamshidi, and Zimmermann 2018, 5).

Due to these factors and the need for continuous services, Pahl, Jamshidi, and Zimmermann (2018, 5) think that the cloud requires its own architectural style. They define cloud architecture as *”an abstract model of a distributed cloud system with the appropriate elements to represent not only application components and their relationships but also the resources these components are deployed on and the respective management elements”*. This aligns with my experiences, where architectural diagrams for cloud-native applications always seem to include the names of the used platform services. For example in case of an application built on AWS, the diagram will specify each usage of *SNS, SQS, Cognito, S3, Lambda, DynamoDB, API Gateway* and so forth, instead of using generic terminology.

Cloud-native architectural style described by Pahl, Jamshidi, and Zimmermann (2018, 5) includes a set of principles and patterns. Principles are service-orientation, virtualization, adaptation, and uncertainty. Service-orientation can be considered a group of three related principles: layering, modularity, and loose coupling. Layering will be necessary, because services will likely have different lifecycles, interface granularities, and so on. This will lead to the communication containing more indirections, which might require additional communications infrastructure. Modularity means refactoring an application into services and exposing only their interfaces, encapsulating implementation details. This makes integrations and composition easier, but requires a way to locate and call these smaller services. Care must also be taken to avoid service calls creating undesired side effects for state management. Because the separate services in a modular application are in some way dependent on each other, coupling should be minimized where possible. Loose coupling can be achieved e.g. with a messaging system, which decouples time, location, and even platform. Stateless services make loose coupling easier, as service calls to stateless services are less likely to result in side effects. (Pahl, Jamshidi, and Zimmermann 2018, 3–8).

Virtualization is generally considered a necessity to manage a shared pool of cloud resources effectively, with optimised cost benefits and economies of scale (Pahl, Jamshidi, and Zimmermann 2018, 8). CSPs offer services for virtualizing infrastructure such as VMs, and platforms such as container orchestration services. Emphasizing adaptation means that the service should meet its requirements with minimal human intervention despite changing circumstances such as system resources, errors, or usage patterns. Unfortunately, in the cloud it is difficult to map requirements to underlying infrastructure. Another thing to keep in mind is that placing a burden on the infrastructure will also map directly to the cloud cost model. This means that cost should be considered in adaptation requirement choices, excluding some as too costly. (Pahl, Jamshidi, and Zimmermann 2018, 3–10). In my view it is beneficial if an architect knows the cloud service provider’s offerings almost to the infrastructure level, to the extent it is possible.

Uncertainty is always present in the cloud. It is caused by physical and logical distribution, heterogeneity, and multi-user involvement in changing contexts (Pahl, Jamshidi, and Zimmermann 2018, 3). It is also possible that monitoring data is unreliable or inadequate, if multiple distributed monitoring systems are in use. This makes reliable measurement of quality attributes of cloud software difficult. For example, creating or freeing cloud resources such as VMs is not immediate, which contributes to uncertainty during the change as the workload varies. (Pahl, Jamshidi, and Zimmermann 2018, 3–11).

2.4 Maintainability in cloud-native architecture

In this section I will tie together the topics I discussed in previous sections: architecture, maintainability, and cloud-nativity.

The Twelve-Factor App (Wiggins 2017) is a methodology for building software for the web and software-as-a-service model (**SaaS**), and as such fits the cloud-native paradigm. It is intended to increase awareness of system-level problems in modern software development, offer a shared vocabulary to discuss them, and showcase a selection of broad concept-level solutions to them. Some of these solutions will support maintainability. Solution #2 encourages to explicitly declare and isolate dependencies. This means that one should never rely

on the supposed existence of system-wide packages or tools. If you need them, you should package them along with your application. All dependencies should be explicitly declared, and the environment should be isolated from external dependencies. These can be achieved by using the preferred tools in your programming language, such as *Gemfiles* and *bundle exec* with *Ruby* programming language, or *pip* and *virtualenv* with *Python*. (Wiggins 2017). Vettor, Pine, Jain, et al. (2022) concur that packaging dependencies in the container image of a microservice belonging to an MSA application fulfills this solution. This speeds up development, as a new developer will only need the codebase, the language runtime, and a dependency manager to run the application.

A "log" is a stream of aggregated, time-ordered events collected from the output streams of all running processes and backing services (Wiggins 2017). Logs are often saved to a text file for later browsing. I consider software logs to be an inseparable part of a software maintainer's daily life. Twelve-factor app solution #11 says that logs should be treated as event streams, and forwarded unchanged into the standard output. This means relying on the platform's log management capabilities instead of trying to manage the log files yourself. Then the logs will be directed to the execution environment, which manages the logs for you. (Wiggins 2017).

In regards to software, a "smell" is considered synonymous to antipattern, flaw, or anomaly (Mumtaz, Singh, and Blincoe 2021, 1). In their paper about architectural smells detection, Mumtaz, Singh, and Blincoe (2021, 20) found that research tends to focus on dependency smells, which are usually connected to coupling issues in the architecture. Most detection tools can detect common dependency smells. Service-oriented architecture smells were covered nicely in their research, but unfortunately there was little literature about smells related to MSA or cloud architectures; the researchers identified only one research about cloud architecture smells. The researchers suggest more investigations into these architectural styles that are currently popular in the industry. The most commonly studied quality characteristic was maintainability, which can shed light into the efforts required to fix the problems resulting from the smells. There is however still room for investigating the sub-characteristics of maintainability, such as testability and reusability. Modularity and modifiability were well covered in the available research. (Mumtaz, Singh, and Blincoe 2021, 21–22).

Almost like an answer to the suggestion by Mumtaz, Singh, and Blincoe (2021), Lichtenthäler and Wirtz (2022) formulated a quality model⁷ that is focused on the design time attributes of cloud-native software. This makes it possible to evaluate architecture before the implementation, because it doesn't rely on source code. They say that the model can be used to approximate the level of "cloud-nativeness" of the application. The higher-level quality aspects in the model are mapped from the ISO/IEC 25010 standard⁸. According to the ISO/IEC 25010 standard, the quality aspect maintainability can be split into sub-characteristics modularity, reusability, analysability, modifiability, and testability. Lichtenthäler and Wirtz (2022) add simplicity to this list. The model then connects these quality aspect to the so-called "product factors". Product factors characterize an architecture, and evaluating them will give us an idea how to practically improve the quality of the architecture in relation to that aspect. Additionally, the researchers list the metrics related to each product factor, but I omit them here for sake of brevity. (Lichtenthäler and Wirtz 2022).

The quality aspect reusability can be improved by focusing on standardization on system, component, and link levels (Lichtenthäler and Wirtz 2022). The goal is to have similarity between components. When components are similar to each other, reuse becomes easier. "System" means the cloud-native application as a whole. A "component" is an abstract entity for representing a distinct part of the system that provides certain functionalities. A component can for example be a service or a certain cloud resource. A "link" then is a directed potential connection between a specific component and a specific endpoint of a different component. (Lichtenthäler and Wirtz 2022).

The quality aspect modularity is related to product factors service-orientation, isolated state, and loose coupling (Lichtenthäler and Wirtz 2022). Each of these three contains additional nested product factors. Service-orientation can be deconstructed into limited functional scope, and separation by gateways. Limited functional scope consists of nested product factors limited data scope, limited endpoint scope, and command query responsibility segregation. Isolated state means moving majority of the system state into specialized stateful services, and keeping most of the other services as stateless as possible. Loose coupling can

7. <https://r0light.github.io/cna-quality-model/>

8. <https://iso25000.com/index.php/en/iso-25000-standards/iso-25010>

be achieved by implementing asynchronous communication, and communication partner abstractions. Focusing on these factors would improve modularity as a whole. (Lichtenthäler and Wirtz 2022).

The quality aspect analysability is improved by implementing automated monitoring. This can mean consistent centralized logging and metrics, distributed tracing of invocations, and health and readiness checks (Lichtenthäler and Wirtz 2022). Monitorability is related to analysability, and is especially relevant for MSA systems, because such systems contain many dynamic structures and behaviors (Li et al. 2021, 13). It is possible to monitor many aspects of the stack separately: infrastructure, such as VM or a container metrics; software metrics such as response time; or the platform metrics, such as network latency. However with CNAs, we only need to worry about monitoring our own services, and leave the platform and infrastructure monitoring to the cloud service provider. In general, monitoring includes generating, storing, processing, and presenting the monitored data (Li et al. 2021, 13). These steps might require installing an agent and other additional tooling, but luckily cloud services often provide basic functionality out of the box.

The quality aspect simplicity can be broken down into sparsity, operation outsourcing, and usage of existing solutions for non-core capabilities. Operation outsourcing means using managed infrastructure, and managed backing services (Lichtenthäler and Wirtz 2022). It is understandable how this increases maintainability: the lower the scope of your maintenance responsibilities, the higher the maintainability. Cloud service providers offer backing services to facilitate building software. Backing services can help with monitoring, streaming, security, analytics, message brokering, storage, caching, and other needs (Vettor, Pine, Jain, et al. 2022). One should favor backing services provided by the cloud service provider in order save time and money by not having to worry about the backing service's performance, security or maintenance. Backing services are best utilized as an attached resource, that are dynamically attached to the application as an external configuration, for example via environment variables. This enables us to attach and detach a backing service without code changes. Another best practice is to implement a middle layer that sits between the business code and backing service. This helps avoid being tightly coupled to specific APIs, and changing the service will be easier if needed. (Vettor, Pine, Jain, et al. 2022).

The quality aspect modifiability can be improved by emphasizing service independence, which in turn can be achieved with low coupling, functional decentralization, limited request trace scope, logical grouping, and backing service decentralization. Implementing abstractions for addressing these services will also benefit modifiability. Automated infrastructure provisioning is important, for which the natural companion is infrastructure-as-code (**IaC**). It helps automate platform provisioning and software deployment, because you are able to utilize software development practices such as testing and version control for the infrastructure definitions (Vettor, Pine, Jain, et al. 2022). IaC technologies include *AWS Cloudformation*, *Azure Resource Manager (ARM)*, *Azure Bicep*, *Terraform*, and *Ansible*. Resource names, locations, capacities and secrets are parametrized and dynamic in IaC. IaC scripts are also designed to be idempotent, meaning that you can change one resource's definition and rerun the script without altering the other resources. This makes IaC play nicely with continuous integration and continuous deployment (CI/CD) technologies: changes to infrastructure are easy to execute. If deployments are easy, integration and deployment likely happens more often than once a quarter, which leads to better communication and mistakes that are cheaper to fix, resulting in higher quality software. (Vettor, Pine, Jain, et al. 2022).

If a system is highly testable, it is able to demonstrate its flaws through testing (Li et al. 2021, 17). Testability is important with MSA applications, because microservices go through many changes during their life, and their interactions are often complex. By maximizing testability, we aim to minimize threats for performance, availability, and security, in addition to maintainability. Testability can be increased by API documentation and management, where we manage API versions and changes over time. This management should be automated as far as possible by tools such as *Swagger* and *RAML*, to avoid "forgetting" to update documentation. (Li et al. 2021, 17).

Testing can be made easier - and in turn more effective - by implementing automatic test procedures (Li et al. 2021, 17). Increasing automation in the testing process reduces the complexity and effort required. Comprehensive manual testing is in many cases impossible, due to the complexity inherent to MSAs. Automated testing procedures align well with the DevOps approach, which emphasizes automating away as much manual work as possible, in order to avoid tedium and human mistakes. Many different testing paradigms can be au-

tomated, not only unit tests. As a few examples, *LBTest* uses machine learning to automate end-to-end testing, *EVOMASTER* uses an evolutionary genetic algorithm for service testing, *Diffy* uses a continuous delivery pipeline for regression testing, and *GMAT* uses graph-based dependency analysis for contract testing. Integrating these tools into the development process will require a learning curve, especially when the aimed at production software. (Li et al. 2021, 18).

This concludes my overview of the scientific background of maintainability in cloud-native architecture. In the following chapters I will present the basis of my own research: research method, results, and analysis. After that I will draw conclusions, and describe how they relate to the literary background I presented in this chapter 2.

3 Method

In this chapter I present the research questions, method, and the research process.

3.1 Research questions and choice of method

Through the research, I aimed to answer three research questions:

RQ1 How much importance do cloud architects place on maintainability during the design phase?

RQ2 How to address maintainability concerns when architecting cloud-native applications?

RQ3 Do the recommendations from literature match the views of architects working in the field?

To me, a digital survey appeared to be the best-fitting research method for answering these questions. It probably would have been possible find existing data sets through a literary review, but finding a reasonable amount of existing data seemed unlikely for RQ1 and RQ3.

3.2 Survey

I targeted the survey at cloud architects of *Nordcloud*, an IBM company, where I work. When weighing my options how to collect the necessary data, I decided to only utilize my current workplace as the target organization. The reasoning behind this was twofold: proximity through my current employment, and the fact that the company employs numerous cloud architects who had experience and, hopefully, opinions about their craft. I had access to the company's collaboration software that's used for internal communications, which made contacting potential survey respondents convenient.

I used *Google Forms* for creating the survey and collecting the results. Before the respondent could proceed to answer the actual questions, the survey briefly described the research topic. There was also a data collection disclaimer with a mandatory agree-disagree choice. If the respondent disagreed with the disclaimer, they were thanked for initial interest, and encour-

aged to submit the form nevertheless, so I would know that the terms were not preferable. The first half of the survey gathered some demographic insights:

- Q1** Years of experience in cloud-native architecture?
- Q2** Years of experience in software architecture in general?
- Q3** Years of experience in IT in general?

Each domain field mentioned in these questions was intended to be a subset of the next: in my view, IT is a superset of software architecture, which is a superset of cloud-native architecture. All three questions were multiple-choice, single answer with these ranges:

- 0-2 years
- 2-5 years
- 5-10 years
- 10+ years

Second half of the survey was dedicated to the main subject matter of maintainability and its relation to other quality attributes:

- Q4** How do you prioritize these software quality attributes when designing cloud-native architecture?
- Q5** How do you address maintainability with platform- and technology choices?
- Q6** How do you address maintainability with application architecture?

Survey question Q4 is directly linked to research question RQ1. The question listed a predefined selection of software quality attributes:

- Performance
- Reliability
- Maintainability
- Scalability
- Security

The respondent needed to set the above quality attributes in a priority order. Available options were:

- Most important
- More important
- Important
- Less important
- Least important

Survey questions Q5 and Q6 are both linked to research questions RQ2 and RQ3. The input was a free text field, so that the respondents could express their thoughts in a non-structured way. I did not want to limit the answer space with my own lack of knowledge by providing a predefined set of options. My intention was to later interpret the answers in a qualitative fashion to find common factors between the responses.

On Monday 20.03.2023 I posted an introductory message along with a link to the survey to a few of the company's internal communication channels in order to reach most of their cloud architects. Some of the channels had overlapping audience, but the total reach was over a thousand people. My goal was to get 10-20 answers, to be able to conduct any amount of meaningful statistical analysis on the data. I was prepared to repost the survey once a week up to four times until I decide I have enough answers to proceed. In the end there was no need to repost, because I received a sufficient number of responses after the initial posting. In early April I managed to convince one additional architect in-person at the company's Jyväskylä office to respond. Some stray responses were recorded up to 21.4.2023, with a final tally of 15 responses.

I could have widened the target audience to other companies, if I wanted to collect a larger sample base. I decided against it, because my preliminary plans at the time did not require more than 20 samples. I planned to do some qualitative processing on the results, in addition to basic quantitative analysis, and qualitative analysis is more labor intensive per response than quantitative methods. If I had included other companies in the research, I would have likely approached the companies via email instead.

I used my own university-provided *Google* account for all survey-related functions. Individual answers and their aggregated analysis will be stored in *Google Forms* and *Google Sheets* in this student account until the end of my study rights and subsequent removal of the

account. All answers were collected and analyzed anonymously. Email addresses or other pieces of personal information were not collected. Responding to the survey required logging in with a *Google* account, but it was only enforced to kindly remind the busy architects if they had already responded to the survey.

3.3 Analysis

My approach to the analysis of the survey results was not strictly quantitative or qualitative, but rather a mix of both depending on the source data. Questions Q1 to Q4 had a limited input space with multiple-choice fields, and as such were a natural candidate for basic quantitative processing. To answer research question RQ1, I decided to calculate the priority of maintainability relative to the other quality attributes based on answers to question Q4.

Q5 and Q6 benefitted from a more qualitative approach. Because the responses were in an unstructured format, I needed multiple steps to transform the data to a format that could be aggregated and analyzed as an answer to research questions RQ2 and RQ3:

- Interpret and list the suggestions from answers to questions Q5 and Q6.
- Normalize the suggestions to find common categories.
- Find the most popular categories and point out possible outliers.
- Compare the categories to solutions proposed in literature.

Because I gathered some demographic data with survey questions Q1 to Q3, I could formulate additional insights by combining those results with results to the domain question Q4. It is possible to calculate statistical correlation between years of experience and perceived importance of maintainability.

4 Results

4.1 Demographics (Q1-Q3)

The survey began with demographic questions about the respondents' working experience. Table 1 shows respondent experience in cloud-native architecture specifically. Table 2 shows respondent experience in software architecture in general. Table 3 shows respondent experience in information technology in general. Individual responses are available in appendix section A.

Table 1. Respondent experience in cloud-native architecture, $n = 15$

Option	Responses	Percentage	Rank
0-2 years	2	13.33%	3
2-5 years	6	40.00%	2
5-10 years	7	46.67%	1
10+ years	0	0.00%	4

Table 2. Respondent experience in software architecture in general, $n = 15$

Option	Responses	Percentage	Rank
0-2 years	3	20.00%	3
2-5 years	4	26.67%	2
5-10 years	3	20.00%	3
10+ years	5	33.33%	1

Table 3. Respondent experience in information technology in general, $n = 15$

Option	Responses	Percentage	Rank
0-2 years	0	0.00%	4
2-5 years	1	6.67%	3
5-10 years	3	20.00%	2
10+ years	11	73.33%	1

4.2 Prioritization of maintainability (Q4)

To calculate a priority score for each quality attribute, I gave the options an incremental integer value from 1 to 5. The value 1 corresponds to option *Least Important*, while value 5 corresponds to *Most important*. The exact values chosen here are not of importance as long as the difference between each consecutive option is equal, to avoid introducing additional bias. This makes our scale analogous to the *Likert scale* (Likert 1932), which is often used in research questionnaires. Table 4 lists the integer values given to each priority option.

Table 4. Priority options and their integer values

Priority	Value
Most important	5
More important	4
Important	3
Less important	2
Least important	1

Table 5 portrays relevant metrics calculated for each quality attribute.

Table 5. Quality attributes, their average priorities, and other metrics, $n = 15$

QA	Average	Δ to next highest	Min	Mode	Max
Security	4.07	-	1	5	5
Reliability	3.80	0.27	1	5	5
Maintainability	2.67	1.13	1	2	4
Scalability	2.27	0.40	1	1	5
Performance	2.20	0.07	1	1	4

Additionally, I calculated how much respondents' experience correlates with perceived importance of maintainability. To accomplish this, I calculated the correlation coefficient for each question Q1, Q2, and Q3 in relation to Q4. I mapped each experience range option in the demographic questions to a numerical value, in order to calculate the correlation. The numerical integer values for each option are listed in Table 6. I chose the values to be the midpoint of each range, with the exception of final option "10+ years", where value follows

the increasing trend across the options. The resulting correlation coefficients are shown in Table 7.

Table 6. Experience range options and their integer values

Experience range	Value
0-2 years	1
2-5 years	3.5
5-10 years	7.5
10+ years	13

Table 7. Correlations between experience in a given domain, and prioritization of maintainability.

Domain	Correlation Coefficient
Cloud-native architecture	0.23
Software architecture	0.12
Information technology	-0.10

4.3 Ways to improve maintainability (Q5 & Q6)

Responses to questions Q5 and Q6 are available in appendix section C and section D respectively. I originally intended the two questions Q5 and Q6 to be scoped to two separate areas: choices about technology, and those related to application architecture. Instead, the responses were varied and included suggestions touching multiple different areas. I think this was luck on my part, because otherwise I would not have received answers related to processes or architectural philosophies. As a consequence, I saw no need to keep the results of the two questions separate, and instead decided to pool the answers together. In my opinion, dividing results based on the question number would be technically correct, but it would add artificial complexity to the analysis. One respondent did not provide answers to either of the questions Q5 or Q6, which resulted in a pool of 14 responses to analyze.

I categorized the results to three top-level categories, two of which align with the survey question prompts Q5 and Q6:

- Platform and technology choices
- Application architecture
- Others (processes, best practices, philosophy, ...)

For each top-level category, I further interpreted subcategories from each suggestion, and aggregated them where reasonable. For example, I consider all the following responses to belong to subcategory "established technologies":

- "Future-proof technologies" (Respondent 2)
- "Use up-to-date (but established) frameworks" (Respondent 4)
- "By using commonly used technologies" (Respondent 13)

Below are the resulting categorizations, with list of respondents who had the same suggestion. Table 8 contains suggestions about platform and technology choices, Table 9 shows suggestions about application architecture, and Table 10 lists the remaining suggestions that do not strictly belong to either of the prior categories.

Table 8. Suggestions about platform- and technology choices, $n = 14$

Suggestion	Respondent IDs	Response count	Rank
CI/CD pipelines	R3, R8	2	3
Configuration-as-code	R7	1	4
Consider vendor support	R1	1	4
Disaster recovery	R15	1	4
Don't avoid vendor lock-in	R8	1	4
Established technologies	R2, R4, R9, R10, R13, R14	6	1
Infrastructure-as-code	R3, R7, R12	3	2
Limited dependencies	R2	1	4
Managed services	R11, R13	2	3
PaaS offerings	R3	1	4
Prefer a single vendor	R1	1	4
Technologies the team already knows	R1	1	4
Technologies with high observability	R1	1	4
Ticketing system	R9	1	4
Version control system	R9	1	4

Table 9. Suggestions about application architecture, $n = 14$

Suggestion	Respondent IDs	Response count	Rank
Autoscaling	R3	1	3
Centralized logging	R1	1	3
Common architectures	R14	1	3
High availability	R15	1	3
Loose coupling	R3, R10, R14	3	1
Microservices	R3, R10, R14	3	1
Modularity	R1, R8, R9	3	1
Scalability	R15	1	3
Simple architectures	R4, R6	2	2
SOLID	R2	1	3
Statelessness	R3	1	3

Table 10. Other suggestions to improve maintainability, $n = 14$

Suggestion	Respondent IDs	Response count	Rank
Architecture decision log	R9	1	2
Avoid breaking changes	R7	1	2
Avoid shortcuts / Don't cut corners	R4	1	2
Avoid time pressure in implementation phase	R9	1	2
Clear resource naming	R4	1	2
Clear responsibilities	R12	1	2
Clearly defined processes	R4	1	2
Consistency	R6	1	2
Deliver security fixes	R12	1	2
Design with goal in mind	R1	1	2
DevOps	R1	1	2
Documentation	R9	1	2
Facilitate easy updates	R7, R13	2	1
KISS	R9	1	2
Regular reviews	R12	1	2
Testing	R7	1	2
Update libraries and dependencies before EoL	R12	1	2
Versioned app releases	R12	1	2

5 Discussion

In this chapter I further discuss the results presented in previous chapter. I subdivide the chapter around my research questions RQ1, RQ2, and RQ3, which were introduced in section 3.1. At the end of the chapter are the other miscellaneous topics and thoughts about the results and related literature.

5.1 Research question RQ1

Research question RQ1 focuses on the perceived importance of maintainability from the point of view of cloud architects. Survey question Q4 was used to collect the data necessary to answer the research question, with questions Q1 through Q3 providing perspective about each respondent's work experience.

Q4 required the architect to prioritize various quality attributes on a *Likert*-like scale. From the average priorities and deltas Δ between each quality attribute in Table 5 we can draw multiple conclusions. The top two quality attributes, Security and Reliability, are unequivocally the most highly prioritized attributes in the minds of our cloud architects. Maintainability is sitting quite comfortably in the middle of the range. This means it is considered important, but not a priority if major sacrifices are needed for reliability or security. Performance and scalability are considered least important items to an almost identical degree. In my mind, performance and scalability are related, as scalability affects single- and multi-user performance. I probably could have replaced one of the two with a quality attribute that is not as highly linked to each other.

From the "Min" column in Table 5 we see that every quality attribute was considered "Least important" by at least one respondent. I find this slightly surprising, because I expected that the quality attribute *Security* is widely considered important, if not a top priority. Apparently not in every case; even this expectation is subject to consideration when building software, depending on the other drawbacks.

In a similar vein, four out of six quality attributes were considered "Most important" by at

least one respondent, as visible in the "Max" column. The exceptions were *Maintainability* and *Performance*, which were at most considered "More important", the second highest priority. This is more in line with my gut feeling of the subject. Despite maintainability being the most researched quality attribute in most software life cycle stages (Arvanitou et al. 2017, 61–62; Mumtaz, Singh, and Blincoe 2021, 21–22), it probably isn't the top one item in shareholders' and implementors' minds. No one builds software for the sake of having more software to maintain, instead of solving some business problem.

Survey questions Q1 through Q3 asked about experience in three different domains: cloud-native architecture, software architecture, and information technology. My interpretation and intention was that cloud-native architecture is a subset of software architecture, which in turn is a subset of information technology. Against this interpretation, respondents R3, R7, R10, and R15 reported less experience in software architecture than cloud-native architecture (see appendix section A). This discrepancy could probably have been avoided if I had explicitly stated my assumptions of subsets in the survey. But it does raise the idea that cloud-native architecting is separate from (traditional) software architecting, due to factors discussed in subsection 2.3.1.

From Table 7 we can see that correlation between respondent experience and importance maintainability is weak, at best. The correlation between cloud-native experience and importance of maintainability is positive and weak. Correlation between software architecture experience and importance of maintainability is also positive, but even weaker. Because the reported experience in the information technology sector is highly weighted towards longer tenures, the very weak negative correlation probably disappears into margin of error. I conclude that based on this dataset, there is no relevant correlation between years of experience and perceived importance of maintainability.

5.2 Research question RQ2

Research question RQ2 requests a list of concrete ways to improve maintainability already in the design phase of software lifecycle. Questions Q5 and Q6 allowed me to get such ideas directly from a sample of practicing experts. You can consider the results shown in

section 4.3 to be the high-level answer to research question RQ2. In this section I will bring forward additional reasonings provided by the respondents, and my own thoughts about the answers.

Each top-level category that I established contained suggestions that were mentioned by more than one respondent. Table 8 contains the aggregated answers related to platform and technology choices. The most popular suggestions were:

- Using established technologies (six responses)
- Infrastructure-as-code (three responses)
- CI/CD pipelines (two responses)
- Utilizing managed services (two responses)

In the context of survey results, "established technologies" means well-known programming languages and frameworks, base technologies on a mainstream platform, tech that is up to date, common, and "de facto". Infrastructure-as-code and CI/CD pipelines are widely known concepts and were explicitly mentioned by name, so there was no interpretation needed. Utilizing managed services is a boon in the cloud, and R11 writes that such services require less maintenance, presumably when compared to implementing such functionality yourself.

This does not mean that the rest of the items are less effective or relevant. For example, using a version control system (VCS) was mentioned by only one respondent, but maintaining a software system whose source code is not covered by a VCS must be horribly short-sighted and error-prone. Maybe the ability to easily track and revert changes as needed is such an obvious aspect of modern software development that the rest of the respondents did not bother to mention it explicitly.

Respondent R8 mentioned an interesting idea: don't avoid vendor lock-in. I've understood that vendor lock-in is usually considered a negative, and makes it more difficult to evolve a system in a flexible way. The respondent didn't provide more reasoning for this, but I think R1 to question Q5 is highly related: R1 advises to prefer offerings from a single vendor due to superior integrations between such services when compared to integrating third-party products. I expect that building a service across multiple CSPs will unavoidably increase the complexity of the architecture, which will make maintenance more difficult.

As discussed in subsection 2.3.4, in the cloud world the technology and platform choices are intertwined with applications architecture. The architectural suggestions are categorized in Table 9, and the most popular suggestions were:

- Loose coupling (three responses)
- Microservices architecture (three responses)
- Modularity (three responses)
- Simple architectures (two responses)

According to R14 a loosely coupled architecture, such as a microservices architecture, can be more easily maintainable than a layered architecture. R10 agrees: decoupling the components makes it easier and more flexible to maintain the components individually. Microservices architecture and loose coupling are seen as highly related, as is modularity. R1 and R8 concur that a modular architecture enables each component to be maintained or even replaced if needed. The last point however might be in opposition to the others: a highly modular architecture might be inherently complex. R4 advises to avoid convoluted architectures, and R6 reminds us to be aware that integration of each new service has a maintainability cost.

Q6 response R9 raised the very relevant point that architecture itself seems to be a lesser factor in deciding maintainability when compared to how the implementation phase is carried out. They mention time pressure as a possible risk for a maintainable implementation. This makes sense to me, as then developers are tempted or even forced to cut corners, which R4 advises against in Q5. This might result in less-than-ideal focus to be allocated for writing a robust test suite, documentation, automation, and other aspects that make the maintenance team's life easier.

The "other" response category was necessary due to high variety in answers. This category was even more varied than the two discussed prior, as seen in Table 10: only one category, "Facilitate easy updates" was mentioned by multiple respondents. Even then, the wording in those responses required much interpretation to find the common denominator:

- R7, Q5: "--, no breaking changes or clear path to update"
- R13, Q6: "--, or having updates possible to be made as easily as possible if they need to be done manually."

Many responses were left vague and without further specifications. For example in Q5, R12 suggested a review at least once a year. In Q6 the same respondent mentioned code reviews before software end-of-life, so answer to Q5 is likely related to architecture or platforms. In a similar vein, R6 called for consistency in application development. Because there was no further elaboration, this might relate to either architectural design, implementation, or both. "Designing with the goal in mind" by R1 is likely a reminder to the architect not to get stuck in the mindset of trying to fit every possible feature or technological item in the final product; "circus horse jumps only as high as it needs", as the respondent put in other words. Because especially the last top-level category I discussed contained a high number of individual answers, I will continue to the next research question RQ3, where I relate the responses to existing source literature.

5.3 Research question RQ3

The research question RQ3 concerns itself with comparing the survey results to research literature. In this section I will look back to the existing research I reviewed in chapter 2, and find common threads in the observed survey results.

Maintainability assurance (or maintainability control) is the processes, tools, and metrics used to avoid accumulation of technical debt and to control maintainability as a quality attribute during the design, development, and maintenance phase (Bogner et al. 2018, 2). Bogner et al. (2018) found that there was less scientific research about the "state of practice" than the "state of the art" around these topics. This compelled them to conduct a research into the state of practice by using a survey for software professionals. Their research questions are somewhat parallel to my thesis:

RQ1 What processes, tools, and metrics are used in the industry to measure and control maintainability?

RQ2 Is the usage of maintainability controls different for service-based systems and microservice-based systems?

RQ3 What correlations exist between applied maintainability controls and participants' overall view of maintainability in their project?

97% of respondents in their survey reported at least some symptoms of low maintainability in their project (Bogner et al. 2018). The symptoms include missing or outdated documentation, inadequate testing, time needed to add new functionality, high number of defects with new releases, architectural erosion or complexity, and low release frequency. To alleviate these, some architecture level metrics such as coupling, cohesion, and architectural violations were reported to be used by some respondents in their survey. Despite this, no respondent explicitly mentioned using tools for quality control at an architectural level. Bogner et al. (2018, 5–6) propose that this hints at neglect for architecture level maintainability or evolvability control, which is the core of my survey in this thesis.

When following cloud-native principles, the application will be loosely coupled (R3, R10, R14), resilient, manageable, and observable (Noval et al. 2022). Another staple of such software is statelessness (R3), which makes it easier to scale to multiple servers, leverage caching, use less storage space, and avoid vendor lock-in by not being reliant on any specific type of server. You can avoid much downtime because CSP backing systems (R11, R13) have inbuilt redundancy, and are geographically distributed (Patrizio 2018).

According to ISO/IEC/IEEE 12207:2017(E) (2017), maintainability can be improved in the architectural design phase by emphasizing encapsulation, modularity (R1, R8, R9), and scalability (R15). Documenting the architecture and system (R9) reduces the efforts required to reverse-engineer the system and components when a fix is needed. (ISO/IEC/IEEE 12207:2017(E) 2017, 95–96). As a sidenote, before reading this standard, I did not mentally connect scalability with maintainability, because I feel it more relates to performance and reliability. R9 had good suggestions related to documentation, including an architecture decision log. Architecture decision log is an artifact with descriptions of the chosen architectural concepts, as well as considered alternatives and related justifications. Another standard ISO/IEC/IEEE 42010:2007 (2007, 3) also states that such information must be present in the architectural description (AD).

The standard ISO/IEC/IEEE 12207:2017(E) (2017, 95–96) also notes that the decisions made during architecting can also affect the need and possibilities for e.g. remote diagnostics and maintenance, roll back, backup, and recovering data. Remote diagnostics might be supported by incorporating technologies with high observability (R1). Allocating some

effort for disaster recovery (R15) will in turn facilitate rollbacks, backups, and data recovery.

Because maintainability highly affects the time needed to add new functionality, the time can be shortened by focusing on the various subcategories of maintainability such as modularity (R1, R8, R9), reusability, analyzability/observability (R1), changeability, modification stability, testability, and compliance (ISO/IEC 5055:2021(E) 2021, 233). The quality aspect *analyzability* is improved by implementing automated monitoring. This can mean consistent centralized logging (R1) and metrics, distributed tracing of invocations, and health and readiness checks (Lichtenthäler and Wirtz 2022). Consistency (R6) between modules is also beneficial, which in the context of modularity can mean that different modules should perform similar tasks in a similar fashion (Bass, Clements, and Kazman 1998, 16).

Li et al. (2021, 18) says that breaking up a complex monolith into components will likely still be complex. It is then a worthy endeavor to take the chance to simplify the architecture (R4, R6), if a rewrite is needed anyway. As presented in section subsection 2.3.2, a highly componentized architecture such as an MSA will require additional considerations due to increased need for network calls between the components (Vettor, Pine, Coulter, Schonning, et al. 2022). Quality aspect *simplicity* contains operation outsourcing, which in the cloud computing paradigm means using managed infrastructure and managed backing services (Lichtenthäler and Wirtz 2022). This benefits maintainability, as presented in section 2.4.

Low release frequency can be countered, among other practices, by facilitating easy updates (R7, R13), adopting DevOps practices (R1) such as CI/CD pipelines (R3, R8), and using technologies that the team already knows (R1; Broy, Deissenboeck, and Pizka (2006, 25)). CNAs have high support for automation, which enables the developers to concentrate on relevant challenges instead of tedium (Patrizio 2018). Combined with the ability to deploy small individual services instead of an entire complex system every time, a CNA is easy to update. This boosts productivity when development teams can make both small-scale and sweeping changes with predictable results, and deliver new features to customers more often (Noval et al. 2022; Patrizio 2018).

Because my thesis is especially concerned with cloud-nativity, I want to raise another finding by Bogner et al. (2018) related to service-based systems (**SBS**) and microservice-based

systems (μ SBS). Based on their survey results, 67% of respondents do not treat maintainability differently for SBSs and μ SBSs versus other types of systems. Many companies are not aware of existing service-oriented approaches and metrics for maintainability, or don't see a need for special treatment despite the very different nature of SBSs and μ SBSs. On the other hand, 26% of respondents report using different strategies specifically aimed to increase maintainability with service-oriented systems. The following strategies augment and align with my survey responses:

- Additional integration and behavioral-driven testing
- Usage of reference architectures
- Standardization of interfaces
- Agile methodologies
- Different deployment process
- Different tool support due to increased technological heterogeneity
- Different metrics usage to account for additional abstraction of services
- Relaxed maintainability controls due to high trust in high base level of maintainability gained through service-orientation: flexibility, scalability, less dependency problems, separation of concerns, etc.

In my survey results, respondent R12 calls for "clear responsibilities". Without further elaboration it is unclear what the respondent exactly meant by this. I hypothesize that it has a dual purpose: clear responsibility for each software module, and clear responsibilities in the organization that produces the software. The first interpretation is linked to R2's suggestions of the SOLID design principles¹, whose first letter corresponds to the single-responsibility principle. The principles were originally written for object-oriented programming, but at least some parts can be extrapolated to architectural design: each software module should only have a single reason to change. This increases modifiability, a subcategory of maintainability, because then a change doesn't require touching multiple components (Bass, Clements, and Kazman 1998, 78). The latter interpretation – that R12 means organizational responsibilities – aligns with advice from Bass, Clements, and Kazman (1998, 15): a single architect or a team of architects with a designated leader should be responsible for the architecture.

1. <https://en.wikipedia.org/wiki/SOLID>

5.4 Internal and external validity

In this section I will discuss the internal and external validity of my thesis. For internal validity, I identify some issues with the survey. First of all, because my survey was anonymous, I don't have any way of verifying that all the respondents were, in fact, cloud architects. Instead I relied on the introductory message that I used to advertise the survey where I addressed the intended group. As a cheeky notion, I expect that everyone is so busy these days that no one is likely to volunteer for optional tasks that are directed at another group of employees, so my respondent group should be valid. As a reassurance to myself, some verified cloud architects approached me about the survey questions, so I can safely assume their responses are represented in the result set.

Second, as I wrote in the beginning of section 4.3, the suggestions to Q5 and Q6 were highly varied, and not necessarily scoped to the intended topic of each question. If this had been a series of interviews instead of a one-and-done survey, I could have honed into the intended categories with higher accuracy. But as I wrote earlier, the unexpected suggestions gave a broader look into the respondents' understanding of maintainability and how it could be improved.

Third, I also pointed out in section 5.1 that some respondents did not share my understanding of nested domains when responding to questions Q1 and Q2. This discrepancy could probably have been avoided if I had explicitly stated my assumptions of subsets in the survey. This notion also includes the obvious oversight that I didn't define *cloud-nativity* or *maintainability* at all, but instead relied on the target group having an innate shared understanding of the topics. I remember that I might originally have made a conscious decision to omit these definitions in order not to influence the results too much with a definition that might conflict with a respondent's view. In hindsight, I don't consider intentional vagueness to be good practice when doing research.

Fourth, I understand that my results for RQ1 are in relation to the set of quality attributes I chose to represent in the survey. The choice was somewhat arbitrary, as I was simply trying to include a few reasonably different quality attributes in addition to maintainability. If I had chosen to include other quality attributes such as usability, availability, or integrability, the

numerical result for prioritization of maintainability would likely have been different.

Fifth, my interpretation of responses to questions Q5 and Q6 is highly subjective and limited by my own lack of architectural experience. I probably missed many nuances and entire valid suggestions, if I didn't see how they relate to maintainability as I understand it. The difficulty was also exacerbated by short length of some answers; a single word or a sentence might leave too much for interpretation. Again, qualitative interviews would likely have given enough bandwidth for the respondents to convey their knowledge in a more intact form.

As a last point regarding internal validity, my respondents work for a consulting company. While there are longer engagements, the primary mode of software development for a given respondent might be project-based work. This means that an architect's relationship with the resulting software might be shorter than for an architect working in a company that is building its own product ². Most of my respondents will likely take part in many engagements over the years. Will project-based work cause the architects to skimp on maintainability, or will they instead double down on it, because "I will not be there to fix it later"? Variety in projects might also give the architect a better overview on the landscape of possible approaches and what works. It would be interesting to see additional research into how the company's "type" affects the architects' views on maintainability, but this is likely too niche a topic to gather enough interest.

External validity is difficult for me to judge outside of the preceding discussion available in section 5.3. If the validity should be comprehensively judged, each individual suggestion by my respondents should be searched for from the existing literature in context of maintainability in cloud-native architectures. Maybe I could have attempted to draft a more rigorous comparison to literature in a table format, for example. At the moment I don't feel that this is necessary, and I am satisfied with a blanket statement that the results generally appear to align with existing literature, as can be understood based on chapter 2 and the examples pointed out in section 5.3. I can rest assured that any advice contained in this thesis will be filtered through the expertise of the reader before being applied.

2. Thanks to my colleague Henri Nieminen for pointing out this idea.

6 Conclusions

I take this chance to sum up the results based on what I have learned. While the implementation phase has an enormous impact in determining maintainability, there are many important considerations to be accounted for already in the architectural design phase in order to lay a foundation for a highly maintainable software product. My view aligns with Bogner et al. (2018, 9): in addition to applying best practices, one should consider maintainability in specific ways depending on the chosen architectural approach, in order to avoid technical debt in long-living service-based software systems.

As I mentioned before, I think that my research would have benefitted from interviews with the cloud architects. That way I could have dived deeper into the reasoning behind the suggestions gathered via the survey. While there is good variety in the discovered advice, my results are mostly a list of things to consider during architecting. Having more explanations and context around each response would allow the reader to better align the advice with the reality of their current project. In this light I echo the suggestion by Bogner et al. (2018, 9), and request future research to utilize qualitative interviews with experts in the field of μ SBSs and CNAs.

For more future research topics, I would be interested in hearing if there are differences in how a cloud architect views maintainability if they work in a consulting company, versus if they work in a product company instead. I hypothesize that while the different nature of the architect-product relationships in these two types of companies influences the prioritization of maintainability to some degree, the architect's personality and overarching professional views play a larger role.

Based on the discussion in section 5.3 we can conclude that there are many connection points between my survey, existing literature, and survey by Bogner et al. (2018). The discussion in that section is only a subset of all possible inquiries that can be derived from my survey results. Fully connecting the dots in a formal mapping would require a systematic mapping study or a systematic literature review, both of which are out of scope in this thesis. I raise this item as another topic for future research.

As closing thoughts, this thesis project was my humble attempt at passable science¹. Much could have been improved in the process, but I will leave that for my next thesis. I did reach my unspoken goal, however: to unearth as much of my chosen subject as I can within the framework of a thesis, which has a limited scope by nature. The topic was and still is highly relevant to my professional interests, and I aim to carry and enhance these learnings in my future work engagements.

1. In addition to my supportive wife Sanna, I extend my thanks to Nathan W. Pyle for his comic strip "I have attempted science" which loaned me inspiration to power through. View the comic here before link rot sets in: <https://nathanwpyle.threadless.com/designs/strange-planet-i-have-attempted-science>

Bibliography

Amazon Web Services. 2023. “Serverless Computing”. Visited on September 2, 2023. <https://aws.amazon.com/serverless/>.

Arvanitou, Elvira Maria, Apostolos Ampatzoglou, Alexander Chatzigeorgiou, Matthias Galster, and Paris Avgeriou. 2017. “A mapping study on design-time quality attributes and metrics”. *Journal of Systems and Software* 127:52–77. ISSN: 0164-1212. <https://doi.org/https://doi.org/10.1016/j.jss.2017.01.026>.

Bass, Len, Paul Clements, and Rick Kazman. 1998. *Software Architecture in Practice*. USA: Addison-Wesley Longman Publishing Co., Inc. ISBN: 0201199300.

———. 2003. *Software Architecture in Practice*. SEI series in software engineering. Addison-Wesley. ISBN: 9780321154958. <http://books.google.fi/books?id=mdiIu8Kk1WMC>.

Bogner, Justus, Jonas Fritzsich, Stefan Wagner, and Alfred Zimmermann. 2018. “Limiting Technical Debt with Maintainability Assurance - An Industry Survey on Used Techniques and Differences with Service-and Microservice-Based Systems”. In *1st International Conference on Technical Debt (TechDebt'18)*. <https://doi.org/10.1145/3194164.3194166>.

Bouwers, Eric, and Arie van Deursen. 2010. “A Lightweight Sanity Check for Implemented Architectures”. *IEEE Software* 27 (4): 44–50. ISSN: 1937-4194. <https://doi.org/10.1109/MS.2010.60>.

Broy, Manfred, Florian Deissenboeck, and Markus Pizka. 2006. “Demystifying Maintainability”. In *Proceedings of the 2006 International Workshop on Software Quality*, 21–26. WoSQ '06. Shanghai, China: Association for Computing Machinery. ISBN: 1595933999. <https://doi.org/10.1145/1137702.1137708>.

Gannon, Dennis, Roger Barga, and Neel Sundaresan. 2017. “Cloud-Native Applications”. *IEEE Cloud Computing* 4 (5): 16–21. ISSN: 2325-6095. <https://doi.org/10.1109/MCC.2017.4250939>.

Gorla, Narasimhaiah, and Shang-Che Lin. 2010. “Determinants of software quality: A survey of information systems project managers”. *Information and Software Technology* 52 (6): 602–610. ISSN: 0950-5849. <https://doi.org/https://doi.org/10.1016/j.infsof.2009.11.012>.

ISO/IEC 5055:2021(E). 2021. *Information technology - Software measurement - Software quality measurement - Automated source code quality measures*. Standard. Geneva, CH: International Organization for Standardization.

ISO/IEC/IEEE 12207:2017(E). 2017. *Systems and software engineering – Software life cycle processes*. Standard. International Organization for Standardization. <https://doi.org/10.1109/IEEESTD.2017.8100771>.

ISO/IEC/IEEE 42010:2007. 2007. *Recommended Practice for Architectural Description of Software-Intensive Systems*. Standard. International Organization for Standardization. <https://doi.org/10.1109/IEEESTD.2007.386501>.

“JYKDOK international articles search”. 2023. Visited on February 26, 2023. <https://jyu.finna.fi/primo>.

Koskimies, K., and T. Mikkonen. 2005. *Ohjelmistoarkkitehtuurit*. Valikko-sarja. Talentum. ISBN: 952-14-0862-6.

Kratzke, Nane, and Peter-Christian Quint. 2017. “Understanding cloud-native applications after 10 years of cloud computing - A systematic mapping study”. *Journal of Systems and Software* 126:1–16. ISSN: 0164-1212. <https://doi.org/https://doi.org/10.1016/j.jss.2017.01.001>.

Li, Shanshan, He Zhang, Zijia Jia, Chenxing Zhong, Cheng Zhang, Zhihao Shan, Jinfeng Shen, and Muhammad Ali Babar. 2021. “Understanding and addressing quality attributes of microservices architecture: A Systematic literature review”. *Information and Software Technology* 131:106449. ISSN: 0950-5849. <https://doi.org/https://doi.org/10.1016/j.infsof.2020.106449>.

- Lichtenthaler, Robin, and Guido Wirtz. 2022. “Towards a Quality Model for Cloud-native Applications”. In *Service-Oriented and Cloud Computing*, edited by Fabrizio Montesi, George Angelos Papadopoulos, and Wolf Zimmermann, 109–117. Cham: Springer International Publishing.
- Likert, Rensis. 1932. “A Technique for the Measurement of Attitudes”. *Archives of Psychology* 140:1–55.
- Moses, John. 2009. “Should we try to measure software quality attributes directly?” *Software Quality Journal* 17 (2): 203–213. ISSN: 1573-1367. <https://doi.org/https://doi.org/10.1007/s11219-008-9071-6>.
- Mumtaz, Haris, Paramvir Singh, and Kelly Blincoe. 2021. “A systematic mapping study on architectural smells detection”. *Journal of Systems and Software* 173:110885. ISSN: 0164-1212. <https://doi.org/https://doi.org/10.1016/j.jss.2020.110885>.
- Noval, Ahmad, Chris Aniszczyk, Vincent Rabah, Orlin Vasilev, Michal Jakobczyk, Mario Cisterna, Dominik Liebler, et al. 2022. “CNCF Cloud Native Definition v1.0 · cncf/toc”. Visited on April 10, 2023.
- Pahl, Claus, Pooyan Jamshidi, and Olaf Zimmermann. 2018. “Architectural Principles for Cloud Software”. *ACM Trans. Internet Technol.* (New York, NY, USA) 18 (2). ISSN: 1533-5399.
- Patrizio, Andy. 2018. “What is cloud-native? The modern way to develop software”. *InfoWorld.com* (San Mateo), <https://www.proquest.com/trade-journals/what-is-cloud-native-modern-way-develop-software/docview/2055357601/se-2?accountid=11774>.
- Roberts, Mike. 2018. “Serverless Architectures”. Visited on September 2, 2023. <https://martinfowler.com/articles/serverless.html>.
- Vale, Guilherme, Filipe Figueiredo Correia, Eduardo Martins Guerra, Thatiane de Oliveira Rosa, Jonas Fritsch, and Justus Bogner. 2022. “Designing Microservice Systems Using Patterns: An Empirical Study on Quality Trade-Offs”. (Ithaca), <https://www.proquest.com/working-papers/designing-microservice-systems-using-patterns/docview/2619059217/se-2>.

Vettor, Rob, David Pine, David Coulter, Nick Schonning, and Maira Wenzel. 2022. “Cloud-native communication patterns”. Visited on March 25, 2023.

Vettor, Rob, David Pine, David Coulter, Youssef Victor, Nick Schonning, and Maira Wenzel. 2022. “Cloud-native communication patterns”. Visited on March 25, 2023.

Vettor, Rob, David Pine, Nick Coulter, Maira Wenzel, and Steve Smith. 2022. “Cloud-native communication patterns”. Visited on March 25, 2023.

Vettor, Rob, David Pine, Tarun Jain, Szanto Peter, Omair Majid, Genevieve Warren, Kent Sharkey, et al. 2022. “What is Cloud Native?” Visited on March 25, 2023.

Wiggins, Adam. 2011. “The New Heroku (Part 4 of 4): Erosion-resistance & Explicit Contracts”. Visited on March 26, 2023. https://blog.heroku.com/the_new_heroku_4_erosion_resistance_explicit_contracts.

———. 2017. “The Twelve-Factor App”. Visited on March 26, 2023. <https://12factor.net/>.

Appendices

A Survey responses: Q1-Q3

Table 11. Respondents' years of experience in the relevant domains.

Respondent	Cloud-native architecture	Software architecture	Information technology
R1	2-5 years	10+ years	10+ years
R2	2-5 years	2-5 years	5-10 years
R3	2-5 years	0-2 years	10+ years
R4	2-5 years	10+ years	10+ years
R5	5-10 years	5-10 years	10+ years
R6	5-10 years	10+ years	10+ years
R7	5-10 years	2-5 years	10+ years
R8	5-10 years	5-10 years	10+ years
R9	2-5 years	5-10 years	5-10 years
R10	5-10 years	0-2 years	10+ years
R11	0-2 years	2-5 years	2-5 years
R12	5-10 years	10+ years	10+ years
R13	0-2 years	2-5 years	5-10 years
R14	2-5 years	10+ years	10+ years
R15	5-10 years	0-2 years	10+ years

B Survey responses: Q4

Table 12. How respondents prioritize the selected quality attributes when designing cloud-native architecture.

Respondent	Performance	Reliability	Maintainability	Scalability	Security
R1	More important	Important	Less important	Least important	Most important
R2	Least important	Important	More important	Less important	Most important
R3	Important	Most important	Less important	Least important	More important
R4	Less important	More important	Important	Least important	Most important
R5	Important	Less important	Least important	Most important	More important
R6	Less important	Least important	More important	Important	Most important
R7	Least important	More important	Important	Less important	Most important
R8	Least important	Most important	Important	Less important	More important
R9	Least important	Most important	More important	Important	Less important
R10	Least important	More important	Less important	Important	Most important
R11	More important	Most important	Less important	Least important	Important
R12	Least important	Important	More important	Less important	Most important
R13	Important	More important	Less important	Most important	Least important
R14	More important	Most important	Least important	Less important	Important
R15	Less important	More important	Important	Least important	Most important

C Survey responses: Q5

Table 13: How respondents address maintainability with platform- and technology choices.

Respondent	Answer
R1	<p>Usually I would stick with offerings from the same provider... i.e. if hosting on Azure and I'm looking for a security suite, I'd highly prefer using Microsoft Defender and Sentinel, rather than a third-party product. Integrations between them are usually superior.</p> <p>Another factor is the capability of the team. Higher preference will be given to technologies that the team already knows. The availability to the inner workings (i.e. exposing debug logs to aid maintaining) is also a factor.</p> <p>Lesser factors to consider are support offered by the provider (which obviously affects the choice of purely open-source products).</p>
R2	Future-proof technologies, limiting number of technological dependencies, proper DevOps practices
R3	PaaS where possible, autoscaling + stateless where possible, Infrastructure as code, automated CI/CD pipelines.
R4	Use up-to-date (but established) technology, avoid shortcuts/cutting corners, CLEAR RESOURCE NAMING, clearly defined processes
R5	-
R6	Consistency is what is heavily missing during evolving application development
R7	IaC, tests when updates are planned, no breaking changes or clear path to update
R8	Don't avoid vendor lock-ins, but make sure it's split up in small components, so each can be maintained and replaced if needed.

Continued on the next page...

Table 13: Continued from previous page

Respondent	Answer
R9	VCS, ticket system, tests, use well known language and framework, write docs, architecture decision log, stick to base technologies in (for example) AWS
R10	picking mainstream platform if possible, using the more widely used technologies makes it easier to find people with skills to maintain it in the future
R11	Most of the cloud-native service provides the managed service required less maintenance
R12	IaC, clear responsibilities, at least once-a-year review + changes as needed
R13	By using commonly used technologies
R14	I prefer de facto solutions to assure platform and technology maintenance by the vendors.
R15	Scalable (use case to use case, VMSS, HA) solution with resiliency (AV set, AV zone) and highly secure (WAF, App gateway(depends on external/internal app), Network rules (NSG/FW), right sizing (performance), cost (depends on environment of the application (dev/test, UAT, prod)

D Survey responses: Q6

Table 14: How respondents address maintainability with application architecture.

Respondent	Answer
R1	Modularise the architecture, and enable each component to be maintained/swopped out individually. However, centralise logging for architecture wide debugging.
R2	limiting number of technological dependencies, proper application architecture (SOLID), designing with the goal in mind, "circus horse jumps only as high as it needs"
R3	loosely coupling, stateless components, microservices
R4	Use up-to-date (but established) frameworks, avoid convoluted architectures, avoid shortcuts/cutting corners, CLEAR NAMING
R5	-
R6	Once more consistency - instead of integration of many services you should be aware of each new cloud service costs maintainability
R7	IaC, configuration as a code
R8	A great CI/CD goes a long way
R9	Modularity and KISS, feel like application architecture is less of an issue than HOW the implementation phase is going (time pressure,)
R10	if possible, microservices, decoupled components make it easier to maintain in parts and more flexibly
R11	Use managed cloud service
R12	IaC, clear responsibilities, app releases versioned, code reviewed and replaced well before EOL, security fixes delivered
R13	By using services that update automatically, or having updates possible to be made as easily as possible if they need to be done manually.

Continued on the next page...

Table 14: Continued from previous page

Respondent	Answer
R14	<p>Microservices architecture (loosely coupled) can be more easily maintainable than a layered architecture. Just need to avoid too granular/small microservices to prevent complexity in service relations and interactions. Every new microservice in transactions call path/sequence brings extra delay and complexity.</p>
R15	HA (with DR for prod) with secure and scalable