



Applying Model-Driven Engineering to Stimulate the Adoption of DevOps Processes in Small and Medium-Sized Development Organizations

The Case for Microservice Architecture

Jonas Sorgalla¹ · Philip Wizenty¹ · Florian Rademacher¹ · Sabine Sachweh¹ · Albert Zündorf²

Received: 4 April 2021 / Accepted: 18 August 2021 / Published online: 16 September 2021
© The Author(s) 2021

Abstract

Microservice architecture (MSA) denotes an increasingly popular architectural style in which business capabilities are wrapped into autonomously developable and deployable software components called microservices. Microservice applications are developed by multiple DevOps teams each owning one or more services. In this article, we explore the state of how DevOps teams in small and medium-sized organizations (SMOs) cope with MSA and how they can be supported. We show through a secondary analysis of an exploratory interview study comprising six cases, that the organizational and technological complexity resulting from MSA poses particular challenges for small and medium-sized organizations (SMOs). We apply model-driven engineering to address these challenges. As results of the second analysis, we identify the challenge areas of building and maintaining a common architectural understanding, and dealing with deployment technologies. To support DevOps teams of SMOs in coping with these challenges, we present a model-driven workflow based on LEMMA—the Language Ecosystem for Modeling Microservice Architecture. To implement the workflow, we extend LEMMA with the functionality to (i) generate models from API documentation; (ii) reference remote models owned by other teams; (iii) generate deployment specifications; and (iv) generate a visual representation of the overall architecture. We validate the model-driven workflow and our extensions to LEMMA through a case study showing that the added functionality to LEMMA can bring efficiency gains for DevOps teams. To develop best practices for applying our workflow to maximize efficiency in SMOs, we plan to conduct more empirical research in the field in the future.

Keywords DevOps · Microservice architecture · Development process · Model-driven engineering

This article is part of the topical collection “New Paradigms of Software Production and Deployment” guest edited by Alfredo Capozucca, Jean-Michel Bruel, Manuel Mazzara and Bertrand Meyer.

✉ Jonas Sorgalla
jonas.sorgalla@fh-dortmund.de

Philip Wizenty
philip.wizenty@fh-dortmund.de

Florian Rademacher
florian.rademacher@fh-dortmund.de

Sabine Sachweh
sabine.sachweh@fh-dortmund.de

Introduction

Microservice architecture (MSA) is a novel architectural style for service-based software systems with a strong focus on loose functional, technical, and organizational coupling of services [54]. In a microservice architecture, services

Albert Zündorf
zuendorf@uni-kassel.de

¹ IDiAL Institute, University of Applied Sciences and Arts Dortmund, Otto-Hahn-Straße 27, 44227 Dortmund, Germany

² Department of Computer Science and Electrical Engineering, University of Kassel, Wilhelmshöher Allee 73, 34121 Kassel, Germany

are tailored to distinct business capabilities and executed as independent processes. The adoption of MSA is expected to increase an application's scalability, maintainability, and reliability [54]. It is frequently employed to decompose monolithic applications for which such quality attributes are of critical scale [9].

MSA fosters the adoption of DevOps practices, because it promotes to (i) bundle microservices in self-contained deployment units for continuous delivery; and (ii) delegate responsibility for a microservice to a single team being composed of members with heterogeneous professional backgrounds [4, 53]. Conway's Law [15] is a determining factor in DevOps-based MSA engineering. It states that the communication structure of a system reflects the structure of its development organization. Thus, to achieve loose coupling and autonomy of microservices, it is also crucial to divide the responsibility for microservices' development and deployment between autonomous DevOps teams [53]. As a result, MSA engineering leads to a distributed development process, in which several teams create coherent services of the same software system in parallel.

While various larger enterprises like Netflix,¹ Spotify,² or Zalando³ regularly report about their successful adoption of MSA, there are only a small number of experience reports (e.g., [12]) about how microservices combined with DevOps can be successfully implemented in small and medium-sized development organizations (SMOs) with less than 100 developers involved. Such SMOs typically do not have sufficient resources to directly apply large-scale process models such as Scrum at Scale [14, 77] in terms of employees, knowledge, and experience.

To support SMOs in bridging the gap between available resources and required effort for a successful adoption of DevOps-based MSA engineering, we (i) investigate the characteristics of small- and medium-scale microservice development processes; and (ii) propose means to reduce complexity and increase productivity in DevOps-based MSA engineering within SMOs. More precisely, the contributions of our article are threefold. First, we identify challenges of SMOs in DevOps-based MSA engineering by analyzing a data set of an exploratory qualitative study and linking it with existing empirical knowledge. Second, we employ model-driven engineering (MDE) [13] to introduce a workflow for coping with the previously identified challenges in DevOps-based MSA engineering for SMOs. Third, we present and validate extensions to LEMMA (Language Ecosystem for Modeling Microservice Architecture), which is a set of Eclipse-based modeling languages and model

transformations for MSA engineering [63] enabling sophisticated modeling support for the workflow.

The remainder of this article is organized as follows. In Section “[Background](#)”, we describe in detail the microservice architecture style particularly related to the design, development, and operation stages. In addition, we explain organizational aspects that result from the use of microservices. Section “[Language Ecosystem for Modeling Microservice Architecture](#)” illustrates LEMMA as a set of modeling languages and tools that address the MDE of MSA. In Section “[DevOps-Related Challenges in Microservice Architecture Engineering of SMOs](#)”, we analyze a dataset based on an exploratory interview study in SMOs to identify challenging areas in engineering MSA for DevOps teams in SMOs. Based on these challenge areas, we present a model-driven workflow in Section “[A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering](#)” and describe the extensions of LEMMA to support the workflow. In this regard, Section “[Derivation of Microservice Models from API Documentations](#)” present means to derive LEMMA models from API documentation, Section “[Assembling a Common Architecture Model from Distributed Microservice Models](#)” presents extensions to the LEMMA languages to assemble individual microservice models, Section “[Visualization of Microservice Architecture Models](#)” describes additions to create a visual representation of microservice models, Section 6.5 presents means to specify deployment infrastructure, and Section “[Generating Code From Distributed Deployment Infrastructure Models](#)” elaborates on the ability to generate infrastructure code. We validate our contributions to LEMMA in Section “[Validation](#)”. Section “[Discussion](#)” discusses the model-driven workflow and LEMMA components towards the application in DevOps teams of SMOs. We present related research in Section “[Related Work](#)”. The article ends with a conclusion and outlook on future work in Section “[Conclusion and Future Work](#)”.

Background

This section provides background on the MSA approach and its relation towards the DevOps paradigm. It details special characteristics in the design, development, operation, and organization of microservice architectures and their realization.

General

MSA is a novel approach towards the design, development, and operation of service-based software systems [54]. Therefore, MSA promotes to decompose the architecture of complex software systems into *services*, i.e., loosely coupled

¹ <https://netflix.github.io>.

² <https://labs.spotify.com>.

³ <https://opensource.zalando.com>.

software components that interact by means of predefined interfaces and are composable to synergistically realize coarse-grained business logic [25].

Compared to other approaches for architecting service-based software systems, e.g., SOA [25], MSA puts a strong emphasis on *service-specific independence*. This independence distinguishes MSA from other approaches w.r.t. the following features [53, 54, 64]:

- Each microservice in a microservice architecture focuses on the provisioning of a single distinct capability for functional or infrastructure purposes.
- A microservice is independent from all other architecture components regarding its implementation, data management, testing, deployment, and operation.
- A microservice is fully responsible for all aspects related to its interaction with other architecture components, ranging from the determination of communication protocols over data and format conversions to failure handling.
- Exactly one team is responsible for a microservice and has full accountability for its services' design, development, and deployment.

Starting from the above features, the adoption of MSA may introduce increases in quality attributes [39] such as (i) scalability, as it is possible to purposefully run new instances of microservices covering strongly demanded functionality; (ii) maintainability, as microservices are seamlessly replaceable with alternative implementations; and (iii) reliability, as it delegates responsibility for robustness and resilience to microservices [19, 20, 54]. Additionally, MSA fosters DevOps and agile development, because its single-team ownership calls for heterogeneous team composition and microservices' constrained scope fosters their evolvability [17, 80].

Despite its potential for positively impacting the aforementioned features of a software architecture and its implementation, MSA also introduces complexity both to development processes and operation [17, 73, 79]. Consequently, practitioners in SMOs perceive the successful adoption of MSA as complex [9]. Challenges that must be addressed in MSA adoption are spread across all stages in the engineering process, and thus concern the design of the architecture, its development and operation. Furthermore, MSA imposes additional demands on the organization of the engineering process.

Design Stage

A frequent design challenge in MSA engineering concerns the decomposition of an application domain into microservices, each with a suitable *functional granularity* [32, 73]. Too coarse-grained microservice capabilities neglect

the aforementioned benefits of MSA in terms of service-specific independence. Too fine-grained microservices, on the other hand, may require an inefficiently high amount of communication and thus network traffic at runtime [46]. Although there exist approaches such as Domain-driven Design (DDD) [26] to support in the systematic decomposition and granularity determination of a microservice architecture [54], their perceived complexity hampers widespread adoption in practice [9, 28].

An additional specific in microservice design stems from MSA's omission of explicit *service contracts* [57]. By contrast to SOA, MSA considers the API of a microservice its implicit contract [85], thereby delegating concerns in API management, e.g., API versioning to microservices [73]. Consequently, microservices must ensure their compatibility with possible consumers and also inform them about possible interaction requirements. Furthermore, implicit microservice contracts foster ad hoc communication, which increases runtime complexity and the occurrence of cyclic interaction relationships [78].

Development Stage

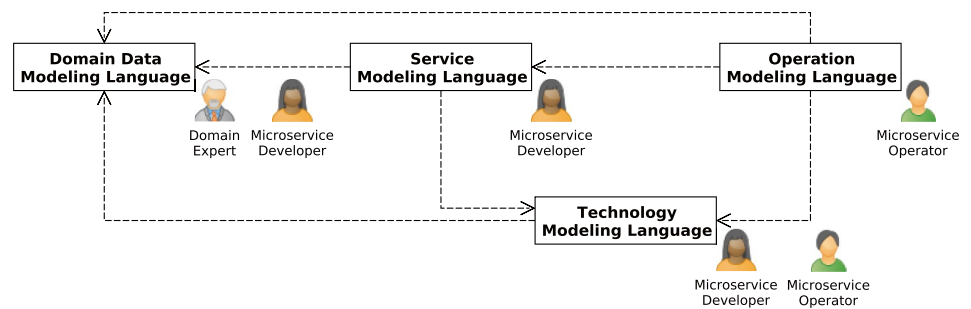
By contrast to monolithic applications, which rely on a holistic, yet vendor-dependent technology stack [19], microservice architectures foster *technology heterogeneity* [54]. Specifically, due to the increase in service-specific independence, each microservice may employ those technologies that best fit a certain capability. Typical *technology variation points* [60] comprise programming languages, databases, communication protocols, and data formats. However, technology heterogeneity imposes a greater risk for technical debt, additional maintainability costs, and steeper learning curves, particularly for new members of a microservice team [78].

Operation Stage

MSA usually requires a sophisticated deployment and operation infrastructure consisting of, e.g., continuous delivery systems, a basic container technology and orchestration platform, to cope with MSA's emphasis of maintainability and reliability [80]. In addition, microservices often rely on further infrastructure components such as service discoveries, API gateways, or monitoring solutions [5], which lead to additional administration and maintenance effort. Consequently, microservice operation involves a variety of different technical components, thereby resulting in a significant complexity increase compared to monolithic applications [73].

Furthermore, technology heterogeneity also concerns microservice operation w.r.t. technology variation points like deployment and infrastructure technologies [60]. Particularly

Fig. 1 Overview of LEMMA's modeling languages, their compositional dependencies and addressed stakeholders. Arrow semantics follow those of UML for dependency specifications [55]



the latter also involve independent decision-making by microservice teams. For example, there exist infrastructure technologies, e.g., to increase performance or resilience, which directly focus on a microservice [5]. Hence, teams are basically free to decide for suitable solutions based on different criteria such as compatibility with existing microservice implementations or available experience.

Organizational Aspects

The use of MSA requires a compatible organizational structure, i.e., following Conway's law, a structure that corresponds to the communication principle of microservices. This results in the necessity of using separate teams, each of which is fully responsible for one or more services (cf. Section "General"). The requirement that a team should cover the entire software lifecycle of its microservices automatically leads to the need for cross-functional teams. To ensure collaboration between teams, large companies such as Netflix or Spotify usually use established *large-scale agile* process models [18], e.g., the Scaled Agile Framework (SAFe) [66], Scrum at Scale [77], or the Spotify Model [71]. Establishing such a form of organization and to establish *organizational alignment* may require upfront efforts [54].

Thus, MSA fosters DevOps practices, which can result in lowered cost and accelerate the pace of product increments [53]. To this end, it is critical to foster a *collaborative culture* within and across teams to promote integration and collaboration among team members with different professional backgrounds [49].

A key enabler of a collaborative culture is the extensive *automation* of manual tasks to prevent the manifestation of inter-team and extra-team silos [49]. Specifically, it relieves people from personal accountability for a task and may thus help in reducing existing animosities of team members with different professional backgrounds [45].

Another pillar of a collaborative culture is *knowledge sharing* following established formats and guidelines [49]. It aims to mitigate the occurrence of insufficient communication, which can be an impediment in both MSA and DevOps [17, 65].

Language Ecosystem for Modeling Microservice Architecture

In our previous works we developed LEMMA [60, 63]. LEMMA is a set of Eclipse-based modeling languages and model transformations that aims to mitigate the challenges in MSA engineering (cf. Section "Background") by means of Model-driven Engineering (MDE) [13].

To this end, LEMMA refers to the notion of *architecture viewpoint* [40] to support stakeholders in MSA engineering in organizing and expressing their concerns towards a microservice architecture under development. More specifically, LEMMA clusters four viewpoints on microservice architectures. Each viewpoint targets at least one stakeholder group in MSA engineering, and comprises one or more stakeholder-oriented modeling languages.

The modeling languages enable the construction of microservice architecture models and their composition by means of an import mechanism. As a result, LEMMA allows reasoning about coherent parts of a microservice architecture [40], e.g., to assess quality attributes and technical debt of microservices [62] or perform DevOps-oriented code generation [61].

The following paragraphs summarize LEMMA's approach to microservice architecture model construction and processing.

Microservice Architecture Model Construction

Figure 1 provides an overview of LEMMA's modeling languages, their compositional dependencies and the addressed stakeholders in MSA engineering.

LEMMA's Domain Data Modeling Language [63] allows model construction in the context of the *domain viewpoint* on a microservice architecture. Therefore, it addresses the concerns of domain experts and microservice developers. First, the language aims to mitigate the complexity of DDD (cf. Section "Background") by defining a minimal set of modeling concepts for the construction of domain concepts, i.e., data structures and list types, and the assignment of DDD patterns, e.g., Entity or Value Object [26]. Additionally, it integrates validations to ensure the semantically

correct usage of the patterns. Second, the language considers underspecification in DDD-based domain model construction [59], thereby facilitating model construction for domain experts. However, microservice developers may later resolve underspecification to enable automated model processing [61]. All other LEMMA modeling languages depend on the Domain Data Modeling Language (cf. Fig. 1) because it provides them with a Java-aligned type system [63] given Java's predominance in service programming [9, 67].

LEMMA's Service Modeling Language [63] addresses the concerns of microservice developers (cf. Fig. 1) in the *service viewpoint* on a microservice architecture. One goal of the Service Modeling Language is to make the APIs of microservices explicit (cf. Section "Background") but keeping their definition as concise as possible based on built-in language primitives. That is, the language provides developers with targeted modeling concepts for the definition of microservices, their interfaces, operations and endpoints. LEMMA service models may import LEMMA domain models to identify the responsibility of a microservice for a certain portion of the application domain [54] and type operation parameters with domain concepts.

LEMMA's Technology Modeling Language [60] considers technology to constitute a dedicated architecture viewpoint [38] that frames the concerns of technology-savvy stakeholders in MSA engineering, i.e., microservice developers and operators (cf. Fig. 1). The Technology Modeling Language enables those stakeholder groups to construct and apply technology models. A LEMMA technology model modularizes information targeting a certain technology relevant to microservice development and operation, e.g., programming languages, software frameworks, or deployment technologies. Furthermore, it integrates a generic metadata mechanism based on *technology aspects* [60]. Technology aspects may, for example, cover annotations of software frameworks. LEMMA service and operation models depend on LEMMA technology models (cf. Fig. 1) and import them to apply the contained technology information to, e.g., modeled microservices and containers. In particular, LEMMA's Technology Modeling Language aims to cope with technology heterogeneity in MSA engineering (cf. Section "Background") by making technology decisions explicit [74].

LEMMA's Operation Modeling Language [63] addresses the concerns of microservice operators (cf. Fig. 1) w.r.t. the *operation viewpoint* in MSA engineering. To this end, the language integrates primitives for the concise modeling of microservice containers, infrastructure nodes, and technology-specific configuration. To model the deployment of microservices, LEMMA operation models import LEMMA service models and assign modeled microservices to containers. Additionally, it is possible to express the dependency of containers on infrastructure nodes such as service discoveries or API gateways [5]. By providing microservice

operators with a dedicated modeling language we aim to cope with operation challenges in MSA engineering (cf. Section "Background"). First, the Operation Modeling Language defines a unified syntax for the modeling of heterogeneous operation nodes of a microservice architecture. Second, it is flexibly extensible with support for operation technologies, e.g., for microservice monitoring or security, leveraging LEMMA technology models (cf. Fig. 1). Third, operation models may import other operation models, e.g., to compose the models of different microservice teams to centralize specification and maintenance of shared infrastructure components such as service discoveries and API gateways.

Microservice Architecture Model Processing

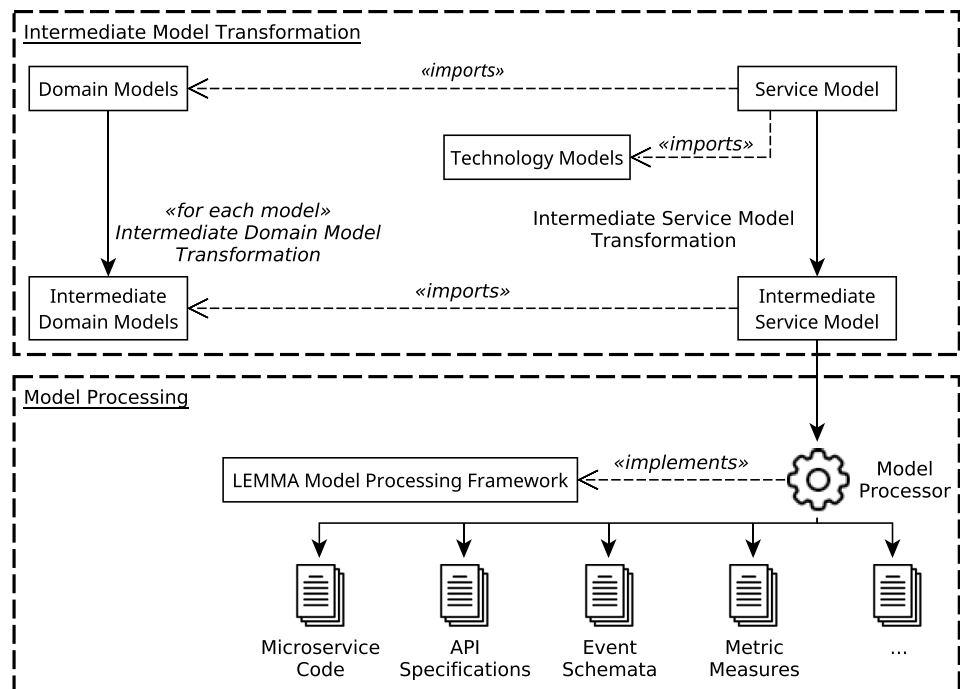
LEMMA relies on the notion of *intermediate model representation* [41] to facilitate the processing of constructed models. To this end, LEMMA integrates a set of *intermediate metamodels* and *intermediate model transformations*. The intermediate metamodels define the concepts to which the elements of an intermediate model for a LEMMA model conform. An intermediate model transformation is then responsible for the automated derivation of an intermediate model from a given input LEMMA model.

This approach to model processing yields several benefits. First, intermediate metamodels decouple modeling languages from model processors. Consequently, languages can evolve independently from processors as long as the intermediate metamodels remain stable. For example, it becomes possible to introduce syntactic sugar in the form of additional shorthand notations for language constructs. Second, intermediate metamodels enable to incorporate language semantics into intermediate models so that model processors need not anticipate them. For instance, LEMMA allows modeling of default protocols for communication types within technology models. In case a service model does not explicitly determine a protocol, e.g., for a microservice, the default protocol of the service's technology model applies implicitly. The intermediate transformation, which converts a service model into its intermediate representation, makes the default protocol explicit. Thus, model processors can directly rely on this information and need not determine the effective protocol for a microservice themselves.

Next to intermediate model representations, LEMMA also provides a model processing framework,⁴ which facilitates the implementation of Java-based model processors, e.g., for microservice developers without a strong background in MDE. To this end, the framework leverages the

⁴ https://github.com/SeelabFhdo/lemma/tree/master/de.fhdo.lemma.model_processing.

Fig. 2 Example for the interplay of LEMMA's intermediate model transformations and model processors based on its model processing framework



Inversion of Control (IoC) design approach [42], and its realization based on the Abstract Class pattern [72] and Java annotations [30]. In addition, the framework implements the Phased Construction model transformation design pattern [47]. That is, the framework consists of several phases including phases for model validation and code generation. To implement a phase as part of a model processor, developers need to provide an implementation of a corresponding abstract framework class, e.g., `AbstractCodeGenerationModule`, and augment the implementation with a phase-specific annotation, e.g., `@CodeGenerationModule`. At runtime, model processors pass control over the program flow to the framework. The framework will then (i) parse all given intermediate LEMMA models; (ii) transform them into object graphs, which abstract from a concrete modeling technology; and (iii) invoke the processor-specific phase implementations with the object graphs. As a result, the added complexities of MDE w.r.t. model parsing and the construction of Abstract Syntax Trees as instantiations of language metamodels [13] remain opaque for model processor developers. Moreover, LEMMA's model processing framework provides means to develop model processors as standalone executable Java applications. This characteristic is crucial for the integration of model processors into continuous integration pipelines [43], which constitute a component in DevOps-based MSA engineering [6, 9].

Figure 2 illustrates the interplay of intermediate model transformations, and the implementation and execution of model processors with LEMMA.

Figure 2 comprises two compartments.

The first compartment shows the structure of intermediate model transformations with LEMMA based on the example of a service model constructed with LEMMA's Service Modeling Language. The service model imports a variety of domain models and technology models constructed with LEMMA's Domain Data Modeling Language and Technology Modeling Language, respectively. As a preparatory step, the service model is transformed into its intermediate representation by means of LEMMA's Intermediate Service Model Transformation. Similarly, each imported domain model is transformed into a corresponding intermediate domain model leveraging LEMMA's Intermediate Domain Model Transformation. To this end, the transformation algorithm restores the existing import relationships between service models and domain models for their derived intermediate representations. However, the algorithm does not invoke intermediate transformations on technology models imported by service models. Instead, the applied technology information becomes part of intermediate service models so that model processors can directly access them. Therefore, LEMMA treats technology models and model processors as conceptual unities. A model processor for a certain technology must be aware of the semantics of the elements in its technology model and be capable in interpreting their application, e.g., within service models.

The second compartment of Fig. 2 concerns model processing. A LEMMA model processor constitutes an implementation conform to LEMMA's model processing framework, which thus provides the processor with capabilities for model parsing and phase-oriented model processing. Typical

Table 1 Overview of explored SMO cases [75]

C	I	Type	Domain	#Services	#Ppl	#Teams
C1	I1	Templated Greenfield	Public Administration	60	≈ 30	5
C2	I2	Migration	B2B E-Commerce	8	10	3
C3	I3	Greenfield	IoT	18	28	2 (up to 5)
C4	I4	Migration	B2B E-Commerce	34	≈ 10	2
C5		Migration	B2C E-Commerce	8	≈ 10	2
C6	I5	Templated Greenfield	Logistics	15–20	75	≈ 10

results from processing service models comprise (i) executable microservice code; (ii) shareable API specifications, e.g., based on OpenAPI;⁵ (iii) event schemata, e.g., for Apache Avro;⁶ and (iv) measures of static complexity and cohesion metrics applicable to MSA [3, 8, 24, 35, 37].

DevOps-Related Challenges in Microservice Architecture Engineering of SMOs

In this section, we present an empirical analysis of microservice development processes (cf. Section “Organizational Aspects”) in SMOs with the goal of identifying SMO-specific challenges in microservice engineering. For this purpose, we perform a *secondary analysis* [36] of transcribed qualitative interviews from one of our previous works [75]. Our analytical procedure specifically aims to identify challenges and obstacles during the development process.

Study Design

The study from which the dataset emerged is a comparative multi-case study [84]. The aim of the study was to gain exploratory insights into the development processes of SMOs. To this end, in-depth interviews were conducted on-site in 2019 with five software architects, each from a different company, and afterwards transcribed. The interviews were conducted in a semi-structured manner and covered the areas of (i) applied development process; (ii) daily routines; (iii) meeting formats; (iv) tools; (v) documentation; and (vi) knowledge management. Participants were recruited from existing contacts of our research group to SMOs. Furthermore, we constrained participant selection to the professional level or senior software architects, and SMOs that develop microservice systems with equal or less than 100 people.

⁵ <http://spec.openapis.org/oas/v3.1.0>.

⁶ <https://avro.apache.org>.

Dataset

As depicted in Table 1, the dataset includes transcripts and derived paraphrases covering six different cases (Column C) of microservice development processes in SMOs. In total, we conducted five in-depth interviews (Column I) with software architects whereby I4 covered two cases.

As shown in Table 1, we distinguish the cases into greenfield (new development from scratch), templated greenfield (new development based on legacy system), and migration (transformation of a monolithic legacy system into an MSA-based system) (Column Type). We further categorize each development process by the domain of the microservice application under development (Column Domain). The number of microservices present in the application at the time of the interview (Column #Services), number of people (Column #Ppl) and teams (Column #Teams) involved vary depending on the case. Case 3 is a special case. Although there are only two official teams, short-term teams are formed depending on the scope of customization needed per customer. This results in up to five teams working on the application simultaneously at certain points in time. In all cases the interviewees stated to apply the Scrum framework [68] for internal team organization. By contrast, the collaboration across teams was in all cases not following a particular formal methodology or model (cf. Section “Organizational Aspects”). In addition, all interviewees reported that they strive for a DevOps culture [21] in their SMOs. A detailed description of the cases can be found in our previous work [75].

Analytical Procedure

For the analysis of the dataset, we used the Constant Comparison method [70]. That is, we rescreened existing paraphrases and marked challenges and/or solutions that our interviewees told us about with corresponding codes for challenges, obstacles, and solutions. We then used the coded statements across all cases to combine similar statements to higher-level challenges.

Study Results and Challenges

Our analysis of the dataset resulted in the discovery of several common challenges across all cases. Comparable to other empirical studies, e.g., [81] or [34], our participants reported about the high technical complexity and high training effort during a microservice development process compared to a monolithic approach. Other discovered challenges in line with existing literature, e.g., [29], concern the slicing of the business domain into individual microservices and the most suitable granularity of a microservice (cf. Section “Background”).

In the following, we elaborate on two challenge areas (CA) which we found to be of particular concern for SMOs adopting a DevOps culture in more detail.

CA1: Developing, Communicating, and Stabilizing a Common Architectural Understanding

Developing a common architectural understanding of the architecture components of an application is essential for developing a software in an organization which follows the DevOps paradigm [6]. In particular, this includes an understanding of the goals and communication relationships of architecture components. The interviewees also think that the development of a general understanding of architecture among those involved in development is an important prerequisite for granting teams autonomy and trust.

For cases C2, C4 and C5 (cf. Table 1), which each comprise approx. ten people and two to three teams, the practices to achieve this understanding are Scrum Dailies [68] and regular developer meetings about the current status of the architecture. However, in case of more involved people, achieving a common understanding is reported to be very challenging. For cases C1, C3, and C6, the system development initially started with fewer people, and as the software product became successful, more people and teams were added. Regarding this development and the common architectural understanding I1 states that “From one agile team to multiple agile teams is a huge leap, you have to regularly adapt and question the organization. [...] you need a common understanding of the architecture and a shared vision of where we want to go [...], we are working on that every day and I don’t think we’ll ever be done.”

A strategy that we observed to create this common architectural understanding in C1, C3, and C6 is the creation of new meeting formats. However, a contradicting key aspect of the DevOps culture is to minimize coordination across teams as much as possible [6]. The arising problem is also experienced by our interviewees. The more people and teams involved in exchanging knowledge to develop an architectural understanding, the more time-consuming the exchange becomes. In the case of C6, this has led to the

discontinuation of comprehensive knowledge exchanges due to the excessive time involved. They now only meet on the cross-team level to discuss technologies, e.g., a particular authentication framework or a new programming language. We interpret this development as a step towards the introduction of horizontal knowledge exchange formats such as Guilds in the Spotify Model [71]. As a result, C6 is currently challenged with building a common understanding of the architecture only through these technology-focused discussions. This is a problem area that is also evident in the data of other empirical studies. For example, Bogner et al. [9] report on the creation of numerous development guidelines by a large development organization to enforce a common architectural understanding. However, the development of guidelines requires that architecture decisions, technology choices, and use cases are documented [33], a practice we encountered only at C4 and C5.

In terms of technical documentation, the teams in all six cases use Swagger to document the microservices’ APIs. Other documentation, such as a wiki system or UML diagrams, either is not used or not kept up to date. In almost all cases, access to the API documentation is not regulated centrally, but is instead provided by the respective team through explicit requests, e.g., by e-mail. Only C3 has extensive and organization-wide technical documentation as it is described by I3: “Swagger is a good tool, but of course this is not completely sufficient, which is why we have an area where the entire concept of the IT platform [...] is explained. We also have a few tutorials.”

Summarizing CA1, we suspect that SMOs are particularly affected by the challenge of implementing a common architectural understanding as part of a successful DevOps culture. This may be due to a mostly volatile organization in which the number of developers and software features often grows as development progresses, as well as the reported hard transition from a single to multiple agile teams. Documenting architecture decisions, deriving appropriate guidelines, and an accessible technical documentation are key factors for an efficient development process that become more relevant with more teams and developers involved [50] and is therefore often not considered by SMOs early in the development process.

CA2: Complexity of Deployment Techniques and Tools

A recurrent challenge we identified is how to deal with the operation of microservice applications within the development process. While cross-functional teams following the DevOps paradigm are mentioned in the literature, e.g., [53], as being recommended for the implementation of microservices architectures, in each of the researched cases we found specialized units for operating microservices instead of operators included as a part of a microservice team. In

C1, C2, and C6 we encountered entire teams solely dedicated to operational aspects. In all cases, the development process included a handover of developed services to those specialized units for operating the microservice application. Although most interviewees were aware that this contradicts the ownership principle of microservices (cf. Section “General”) and they all stated to try to establish a pure DevOps without specialized teams, the effort to learn the basics of the necessary operational aspects is perceived as high. In this regard I2 comments “The complexity (*note: of cloud-based deployment platforms*) is already very, very high, you know. I would say that each of these functions in such a platform is a technology in itself that you have to learn.” In contrast to operations, the SMOs are successful in including other professions, such as UI/UX, as parts of their cross-functional teams. Our data indicates that this is due to two main reasons. First, the inherently high complexity of the operational technologies and the associated high hurdle of learning and integrating them into the microservice development process. Second, the transfer of this knowledge not only into special units but into the individual microservice teams to do justice to a DevOps approach.

Summarizing CA2, deployment and operation in the SMOs studied is not in the responsibility of the teams to which the respective microservices belong. This seems to be due to the complexity of operation technologies and the associated learning effort. This might particularly be an issue for SMOs due to the challenging environment, where there are few resources to substitute, e.g., for a colleague who needs to learn an operation technology.

Case Study

In this section, we present a case study that we will use in the following sections to illustrate and validate our model-driven workflow (cf. Section “A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering”) to address the challenges in Section “DevOps-Related Challenges in Microservice Architecture Engineering of SMOs”. We decided on the usage of a case study to show the applicability of our approach because non-disclosure agreements prevent us from presenting our approach in the context of the explored SMO cases (cf. Section “Dataset”). Therefore, we selected an open source case study microservice architecture, which maps to the design and implementation of the explored SMO cases w.r.t. the scope of our approach. More precisely, the case study (i) employs Swagger for API documentation (cf. Section “Study Results and Challenges”), (ii) uses synchronous and asynchronous communication means, (iii) is mainly based on the Java programming language, and (iv) the number of software

components matches the smaller SMOs in our qualitative analysis (cf. Section “Dataset”).

The case study is based on a fictional insurance company called Lakeside Mutual [76]. The application serves to exemplify different API patterns and DDD for MSA. The application comprises several micro-frontends [58], i.e., *semi-independent frontends* that invoke backend functionality, and microservices centered around the insurance sector, e.g., customer administration, risk management, and customer self-administration functions. The application’s source code as well as documentation is publicly available on GitHub.⁷

Figure 3 depicts the architectural design of the Lakeside Mutual application. Overall it consists of five functional backend microservice. Each microservice is aligned with a micro-frontend.

Except for the Risk Management Server, all microservices are implemented in Java⁸ using the Spring framework.⁹ A micro-frontend communicates with its aligned microservice using RESTful HTTP [27]. Additionally, the Risk Management Client and Risk Management Server communicate via gRPC. For internal service to service communication, the software system also relies on synchronous RESTful HTTP, but also on asynchronous amqp messaging over an Active MQ message broker. The Customer Management Backend and the Customer Core services also provide generated API documentations based on Swagger.¹⁰

Besides the functional microservices, the Lakeside Mutual application also uses infrastructural microservices. The Eureka Server implements a Service Registry [64] to enable loose coupling between microservices and their different instances. For monitoring purposes, the Spring Boot Admin service provides a monitoring interface for the health status of individual services and the overall application.

A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering

This section proposes a model-driven workflow based on LEMMA (cf. Section “Language Ecosystem for Modeling Microservice Architecture”) to cope with the challenges identified in Section “DevOps-Related Challenges in Microservice Architecture Engineering of SMOs”. More

⁷ <https://github.com/Microservice-API-Patterns/LakesideMutual>.

⁸ <https://www.java.com>.

⁹ <https://spring.io>.

¹⁰ <https://swagger.io/>.

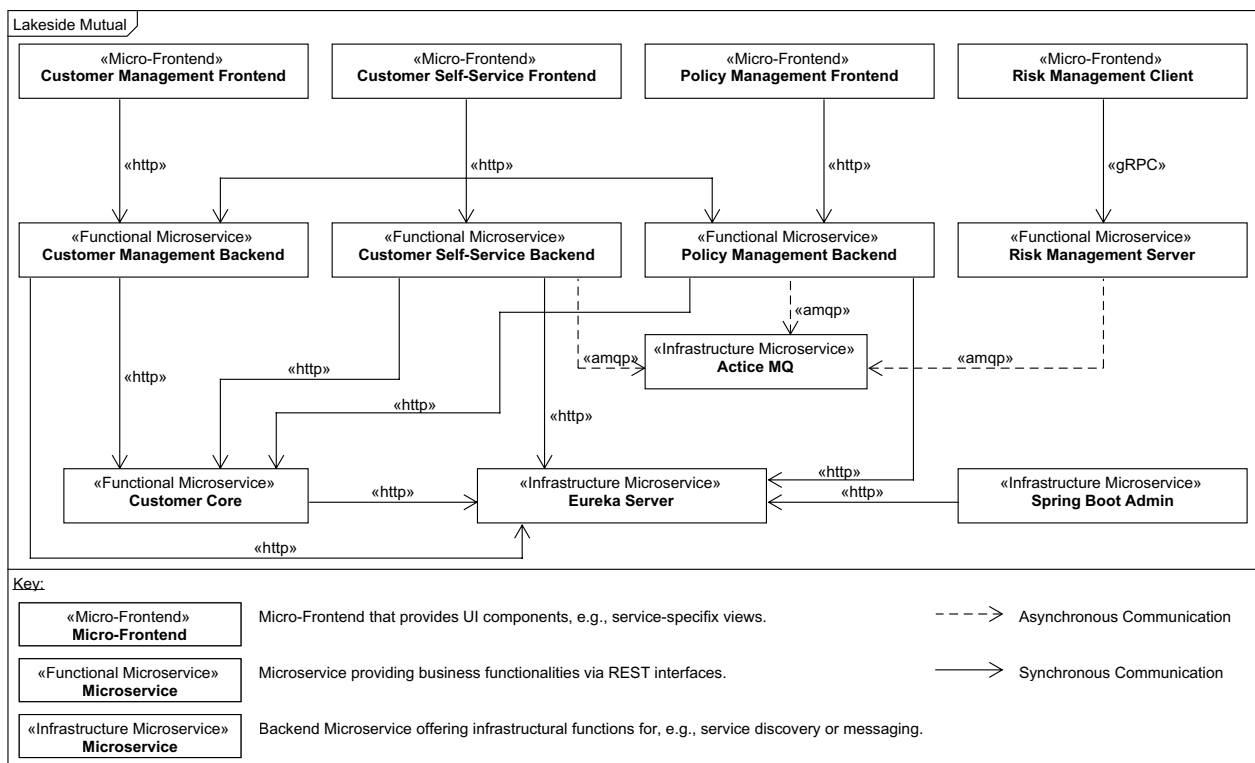


Fig. 3 Structure of the case study Lakeside Mutual microservice application

precisely, the workflow provides a common architectural understanding of a microservice application (cf. Challenge CA1 in Section “[Study Results and Challenges](#)”), and reduces the complexity in deploying and operating microservice architectures (Challenge CA2).

In the following subsections, we present the design of the workflow (cf. Section [LEMMA-Based Workflow for Coping with DevOps Challenges](#)). Next, we describe the components, which we have added to LEMMA, to support the workflow. These components include (i) interoperability bridges between OpenAPI and LEMMA models (CA1; cf. Section “[Derivation of Microservice Models from API Documentations](#)”); (ii) an extension to the Service Modeling Language to allow the import of remote models (CA1; cf. Section “[Assembling a Common Architecture Model from Distributed Microservice Models](#)”); (iii) a model processor to visualize microservice architectures (CA1; cf. “[Visualization of Microservice Architecture Models](#)”); (iv) enhancement of the Operation Modeling Language (cf. Section “[Enhancing Distributed Microservice Models with Deployment Infrastructure Models](#)”; and (v) code generators for microservice deployment and operation (CA2; cf. Section “[Generating Code from Distributed Deployment Infrastructure Models](#)”).

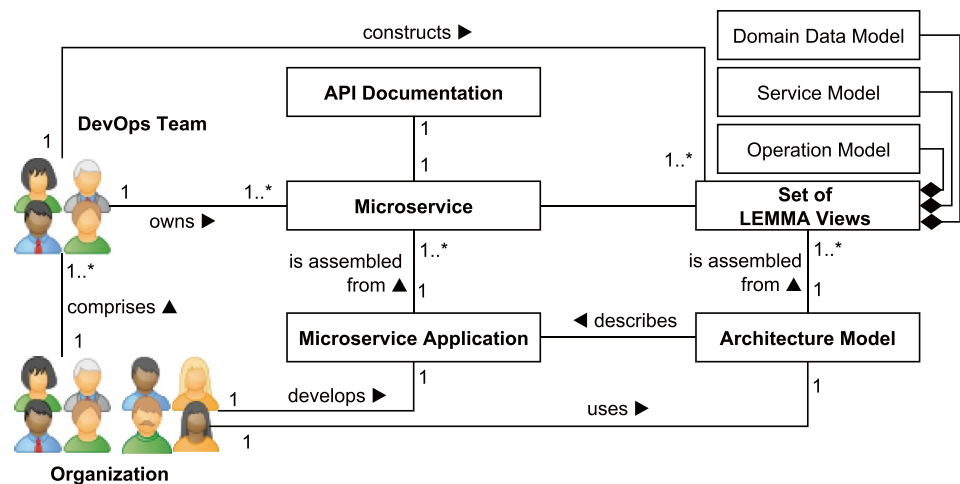
Furthermore, we present in detail prototypical components that we have added to the LEMMA ecosystem to

support the workflow. These include deriving models from API documentation (cf. Section “[Derivation of Microservice Models from API Documentations](#)”) and assembling microservice models into an architecture model (cf. Section “[Assembling a Common Architecture Model from Distributed Microservice Models](#)”) as a means to build a common architectural understanding (CA1). The presented components also comprise means to enrich microservice models with deployment infrastructure models (cf. Section “[Enhancing Distributed Microservice Models with Deployment Infrastructure Models](#)”) to more easily handle operational aspects for SMOs (CA2).

To ensure replicability of our results, we have provided a GitHub repository¹¹ which contains documentation on how to setup LEMMA and its presented extensions. It further contains all generated artifacts as well as sources and scripts to rerun the generations. Finally, it includes a manually constructed set of LEMMA models which represent all Java-based microservices of the Lakeside Mutual case study (cf. Section “[Case Study](#)”).

¹¹ <https://github.com/SeelabFhdo/SN2021>.

Fig. 4 Overview of the concepts within the workflow and their interrelationships represented as a UML class diagram [55]



LEMMA-Based Workflow for Coping with DevOps Challenges

Figure 4 shows the conceptual elements and their relationships which underlie the design of our LEMMA-based workflow for coping with the DevOps challenge areas (cf. Section “DevOps-Related Challenges in Microservice Architecture Engineering of SMOs”).

An *Organization* includes multiple *DevOps Teams*, each responsible for one or more *Microservices* (cf. Section “Organizational Aspects”). The sum of all microservices forms the *Microservice Application* that is developed by the organization. Associated with a microservice is a corresponding documentation of its interfaces (*API Documentation*). For each microservice owned by it, the team constructs a *Set of LEMMA Views* as a model representation (cf. Section “Language Ecosystem for Modeling Microservice Architecture”). The sum of all LEMMA models forms an *Architecture Model* which describes the system’s architecture. This model can be used by the organization, e.g., to gain insight into existing dependencies between the microservices involved.

Based on the conceptual elements and their relationships, Fig. 5 shows our model-driven workflow for DevOps-based microservices development in SMOs as a UML activity diagram [55].

We depict the workflow from the perspective of a single DevOps team including all steps required for the development of a new microservice. When incremental changes are made to individual aspects of a microservice, only the steps affected by the changes need to be performed.

The process starts with the planning of the development. The team decides whether to follow a *code-first* or *model-first* approach. We support both variants to allow the teams autonomy according to the DevOps paradigm [6].

Code-First Approach Here, the team first implements the microservice consisting of structure and behavior. Based on the finished implementation, the team creates an *API Documentation*, which can be done manually or automatically with tools such as Swagger.¹² Using the API documentation, a LEMMA domain model and a LEMMA service model are automatically derived (cf. Section “Derivation of Microservice Models from APIDocumentations”) and, if necessary, refined by the team. In parallel, the team creates a LEMMA operation model, since the information required for this kind of model cannot be derived from the API documentation (cf. Section “Derivation of Microservice Models from APIDocumentations”).

Model-First Approach Alternatively, the team can decide to first model the structure and operation of the microservice using LEMMA. In the subsequent implementation activity, the structural aspects can be generated based on the previously constructed models and only the manual implementation of the behavior is necessary (cf. Section “Language Ecosystem for Modeling Microservice Architecture”).

Regardless of which of the two approaches was chosen, in the end LEMMA domain, service, and operation models are available and describe the *Dev* and *Ops* aspects of the microservice under development.

The operation model is then used to *Generate* a *Deployment Specification* for a container-based environment which mitigates the complexity of the operation (cf. Section “Generating Code from Distributed Deployment Infrastructure Models”). The team refines this specification as needed and then deploys the microservice. In parallel, the models generated during the workflow are sent to a central model repository and made available to the entire organization where they can be used by other teams to gain insight

¹² <https://swagger.io/>.

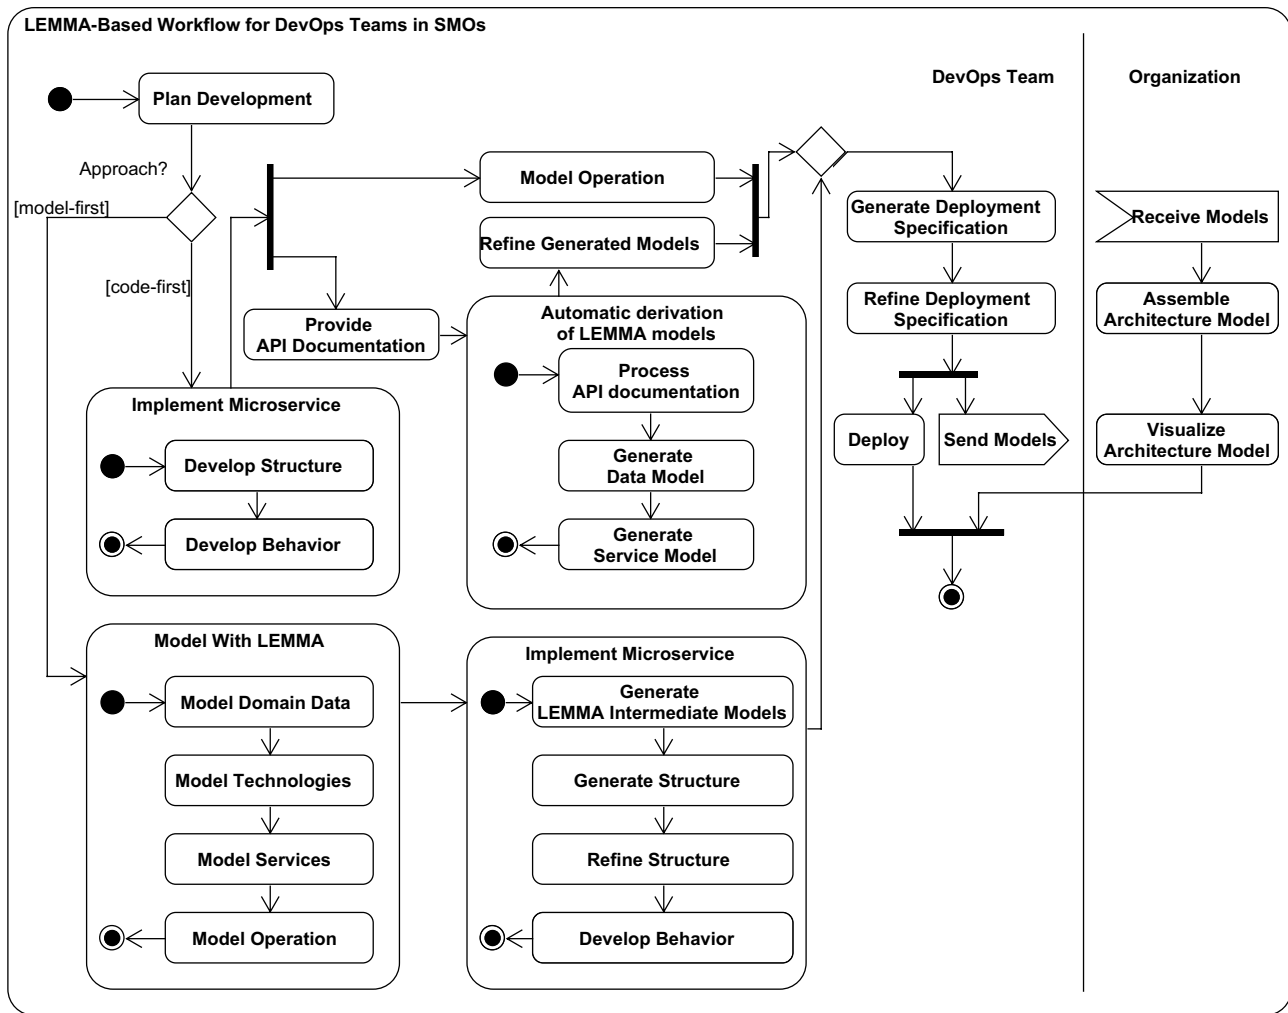


Fig. 5 Proposed workflow for DevOps teams for model-driven microservices development represented as a UML activity diagram [55]

and a common understanding of the application’s architecture, e.g., by visualizing its structure.

Based on the use of model transformations and code generation steps, we argue that the application of the workflow is possible with almost the same resources as the current development processes in the individual DevOps teams that we were able to explore as cases in the empirical study (cf. Section “DevOps-Related Challenges in Microservice Architecture Engineering of SMOs”). This applies to both the code-first and model-first approaches. At its core, the code-first approach relies on the same development steps, i.e., implementing structure and behavior of a microservice, as non-model-based processes in the individual teams, so that even teams without experience in MDE can adapt the flow in a non-invasive way. Besides the actual implementation, the workflow provides a service’s description in the form of LEMMA viewpoint models, which can be used as a communication basis and for knowledge transfer to create a common architecture understanding (CA1;

cf. Section “Study Results and Challenges”) in the organization. This can be used to, e.g., accelerate verbal coordination processes between teams, improve the documentation, or identify microservice bad smells [78]. In addition, by using LEMMA operation models and generating deployment specifications, it is easier for teams of an SMO to address the *Ops* aspects themselves without passing on the responsibility for deployment to another unit (CA2; cf. Section “Study Results and Challenges”). This enables teams to foster the ownership principle of MSA (cf. Section “General”).

Derivation of Microservice Models from API Documentations

To enable the model-driven workflow with sophisticated modeling support by LEMMA, we extended the ecosystem with the ability to derive data and service models from API documentation into LEMMA modeling files. In particular, our extension targets API documentation that conforms to

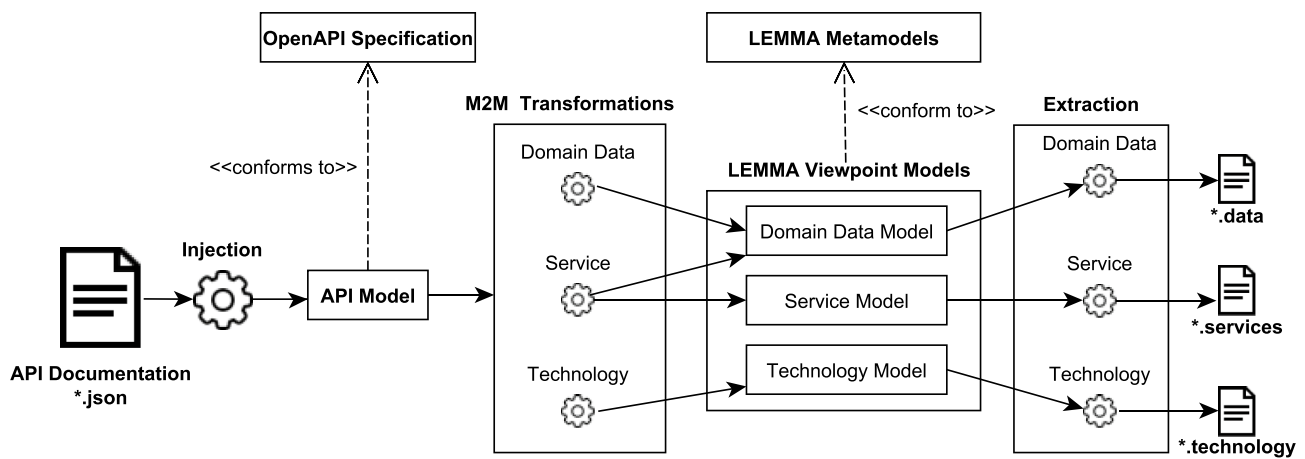


Fig. 6 Process to transform OAS conform API models into LEMMA models

Table 2 Applied mappings between OAS and LEMMA concepts

OAS Objects	LEMMA model	Description of mapping
Info	Domain Data	Name and version of the Domain and its Context
Schemas	Domain Data	DataStructures of the Context
Paths	Domain Data	Inline schemas are transformed to additional DataStructures and arrays are mapped to ListTypes
Info	Service	Name and commentary of the Microservice
Tag	Service	Used to derive a service’s Interfaces which cluster Operations
Paths	Service	Operations of an Interface comprising its Endpoints, Parameters, HTTP request method as Aspect, and Commentary
Media Types	Technology	All mentioned media types in a OAS model are transformed to DataFormats of the RESTful HTTP Protocol.
Data Types	Technology	Mapped to TechnologySpecificPrimitiveTypes.

the OpenAPI Specification¹³ (OAS) [56]. OAS defines a standardized interface to describe RESTful APIs. One of the most popular tools implementing OAS is Swagger, which was used by all SMOs in the qualitative study (cf. Section “DevOps-Related Challenges in MicroserviceArchitecture Engineering of SMOs”).

The transformation of OAS files into LEMMA files can be classified as an interoperability issue in which OAS models are to be converted into LEMMA models. We therefore applied the *interoperability bridge* process proposed by Brambilla et al. [10]. Figure 6 shows the applied interoperability bridge process.

First, using the Swagger parsing framework,¹⁴ an OAS conform API model in the YAML [7] or JSON [22] format

is converted into an in-memory API Model. We then perform three model transformations where the information from the API Model is transformed into a Domain Data Model, Service Model, and a Technology Model which each correspond to their respective LEMMA metamodel (cf. Section “Language Ecosystem for Modeling-Microservice Architecture”). Table 2 describes which OAS objects¹⁵ are mapped to which LEMMA model kind.

To be able to transform the in-memory LEMMA models as files, we extended LEMMA with *extractors* [10] for technology, service, and data models.

Listing 1 and Listing 2 illustrate the application of the process.

¹³ Version 3.0.3.

¹⁴ <https://github.com/swagger-api/swagger-parser>.

¹⁵ <https://spec.openapis.org/oas/v3.0.3#openapi-object>.

```

"paths":{
  "/cities/{postalCode}":{
    "get":{
      "tags":["city-reference-data-holder"],
      "summary":"Get the cities for a postal code.",
      "operationId":"getCitiesForPostalCodeUsingGET",
      "produces":["*/*"],
      "parameters":{
        "name":"postalCode",
        "in":"path",
        "description":"the postal code",
        "required":true,
        "type":"string"
      },
      "responses":{
        "200":{
          "description":"OK",
          "schema":{"$ref":"#/definitions/CitiesResponseDto"}
        },
        /*...*/
      }
    }
  }
}

```

Listing 1 Excerpt of the OAS conform OpenAPICustomerCoreBackendService.json describing the CustomerCore microservice.

Listing 1 shows an excerpt of the API documentation file of the Customer Core microservice from the case study (cf. Section “Case Study”). In detail, the listing presents the OAS description for an HTTP GET request on the path `cities/{postalCode}` (Lines 2 and 3). This includes, e.g., the unique id `getCitiesForPostalCodeUsingGET` (Line 6) of the operation, the incoming parameters (Lines 8–14), and the information that a response returns an object based on the `CitiesResponseDto` schema (Lines 15–20). The excerpt shows only the response for HTTP status code 200 (Line 16). OAS also offers the possibility to define responses for other status codes, e.g., HTTP status code 404, but these are currently not considered in the transformation to LEMMA in our prototypical implementation.

```

import datatypes from "customerCore.data"
  as CustomerCoreAPIData
import technology from "OpenApi.technology" as OpenApi

@technology(OpenApi)
public functional microservice
  com.lakesidemutual.customercore.CustomerCore {
    interface cityReferenceDataHolder {
      ---
      **Type** GET Operation for path /cities/{postalCode}
      **Summary** Get the cities for a postal code.
      **Description**
      @required returnValue [INSERT PARAMETER DESC HERE]
      ---
      @endpoints(OpenApi::_protocols.rest :
        "/cities/{postalCode}");
      @OpenApi::_aspects.GetMapping
      getCitiesForPostalCodeUsingGET(
        sync in postalcode : string,
        sync out returnValue : CustomerCoreAPIData::v100.
          CustomerCoreAPI.CitiesResponseDto);
    }
  }
}

```

Listing 2 Excerpt of `customerCore.services` resulting from transforming the CustomerCore API documentation.

Listing 2 shows the LEMMA service model automatically transformed from the CustomerCore OAS model in


Listing 1. First, the results of the other transformations are imported into the service model. This includes the previously transformed LEMMA domain data model `customerCore.data` resulting from the OAS schemas (Lines 1 and 2), which contains all data structures such as `CitiesResponseDto`, and the technology model `OpenApi.technology` (Line 3), which contains, e.g., the OpenAPI-specific primitive data types and the media types used in the CustomerCore OAS model. Line 5 enables the OpenApi technology for the `com.lakesidemutual.customercore.CustomerCore` microservice, whose definition starts in Lines 6 and 7. The microservice comprises an interface named `cityReferenceDataHolder` which was derived by the associated tags in the OAS model (Line 8). The interface consists of the operation `getCitiesForPostalCodeUsingGET` named after the OAS `operationId` (Lines 18–22). The operation commentary section (Lines 9–14) is populated using the summary information from OAS. The OAS path is added as an endpoint (Line 15) and the operation classified as an HTTP GET request (Line 17). The OAS response associated with the HTTP status code 200 is modeled as an OUT parameter and named `returnValue` (Lines 20 and 21).

Assembling a Common Architecture Model from Distributed Microservice Models

Microservices can interact with and depend on each other to realize coarse-grained functionality [54]. In the case study (cf. Section “Case Study”), such a relationship is found between the microservices `Customer Management Backend` and `Customer Core`. Such dependencies cannot be derived from an API documentation, since its purpose is to describe the provided interface of a service and not the invocation of functionality provided by other architecture components. However, these dependencies are essential to be able to assemble and assess an architecture model and to raise a common architectural understanding across the whole organization (cf. Section “Study Results and Challenges”). Therefore, within the workflow (cf. Section “LEMMA-Based Workflow for Coping with DevOps Challenges”), the dependencies should be added manually by the teams in the LEMMA models. This can be done during the `Model Services` activity when using the model-first approach and during the `Refine Generated Models` activity when using the code-first approach.

However, LEMMA service models originally were only able to depend on other LEMMA service models if they are accessible in the local file system. Therefore, to allow teams the expression of interaction dependencies with the microservices of other teams, we have extended LEMMA to allow external service imports. Listing 3 shows the service model of the

Fig. 7 Model-to-Text transformation to generate a visual representation of a microservice application's architecture



Customer Management Backend microservice from the case study. The microservice imports show the two alternatives. The syntax for importing locally accessible service files is shown in Lines 3 and 4. Alternatively, the import in Lines 6–8 exemplifies the mechanism for external imports.

```

import datatypes from "customerManagementBackend.data" 1
as customerManagementBackend                          2
import microservices from "../customer-core/customerCore 3
.services" as customerCoreServices                     4
//External import alternative to previous local import 5
import microservices from "https://repo.lakeside.com/ 6
teamB/customercore.json" to "../customer-core/ 7
customerCore.services" as customerCoreServices        8
...                                                    9
public functional microservice                         10
com.lakesidemutual.customerManagementBackend         11
.CustomerManagementBackend {                          12
  required microservices {customerCoreServices::com. 13
    lakesidemutual.customercore.CustomerCore}        14
  ...                                                  15

```

Listing 3 Excerpt of `customerManagementBackend.services` illustrating the external import feature.

As soon as the Eclipse IDE detects such an external import in the model, it offers a quickfix that automatically downloads the referenced file and, if it is OAS-compliant API documentation, starts a corresponding transformation to LEMMA (see Section “[Microservice Architecture Model Processing](#)”). This also makes it possible to model a dependency to a service of another team, even if this team does not yet provide its own model but only API documentation.

Since LEMMA models are textual [63] and with the extension it is possible to import external sources, the model files of the different teams can be managed centrally as an architecture model by a version management system such as Git and thus integrated into CI/CD pipelines, e.g., by a Git hook¹⁶ that copies the models to a central model repository with each release of the microservice.

Visualization of Microservice Architecture Models

To enable visualization of the architecture using LEMMA (cf. Section [LEMMA-Based Workflow for Coping with DevOps Challenges](#)), we have developed the LEMMA Visualizer.¹⁷ It is able to transform several LEMMA

¹⁶ <https://git-scm.com/book/en/v2/Customizing-Git-Git-Hooks>.

¹⁷ <https://github.com/SeeLabFhdo/SN2021/blob/master/de.fhdo.lemma.visualizer-0.8.0-SNAPSHOT-standalone.jar>.

intermediate service models (cf. Section “[Microservice Architecture Model Processing](#)”) into a single graphical representation using a *model-to-text transformation* [13]. In its current form, the visualizer is a standalone executable Jar file. The Jar file can be passed the paths to several service models, which serve as input, and a target path for generating the visualization as arguments through the command line. The steps of the transformation are depicted in Fig. 7.

First, one or more intermediate service model files are passed to the LEMMA Model Processor (cf. Section “[Microservice Architecture Model Processing](#)”). These are converted to their in-memory representation and then processed. Using the JGraphT framework [51], we create a directed graph that we populate with microservices found during the processing of the intermediate models as vertices and any existing imports from other services as edges. Microservices that are neither imported nor import another service and are thus without an edge are added as isolated vertices. Then we use JGraphT’s DOTExporter to convert the graph into a textual representation of the graph based on the DOT language.¹⁸ During the export, we enrich the DOT representation with attributes which describe the appearance of the vertices and edges for the later visualization, e.g., we add coloring and describe vertex shapes. Finally, we use GraphViz [23] to generate an image of the graph’s DOT representation that represents the system architecture in the form of a box-and-line diagram as a Portable Network Graphics (.png) file. The visualizer supports the setting of different levels of detail of the display in relation to the attributes of microservice vertices. A resulting architecture image which shows the functional microservices from the case study is shown in Fig. 8. To generate the image, the visualizer was configured to render with a detail level which shows interfaces but not operations.

Enhancing Distributed Microservice Models with Deployment Infrastructure Models

In this subsection, we elaborate on the creation of an operation model (cf. Section “[Case Study](#)”) that complements the previously created data and service models to form a complete set of LEMMA views describing a microservice

¹⁸ <https://graphviz.org/doc/info/lang.html>.

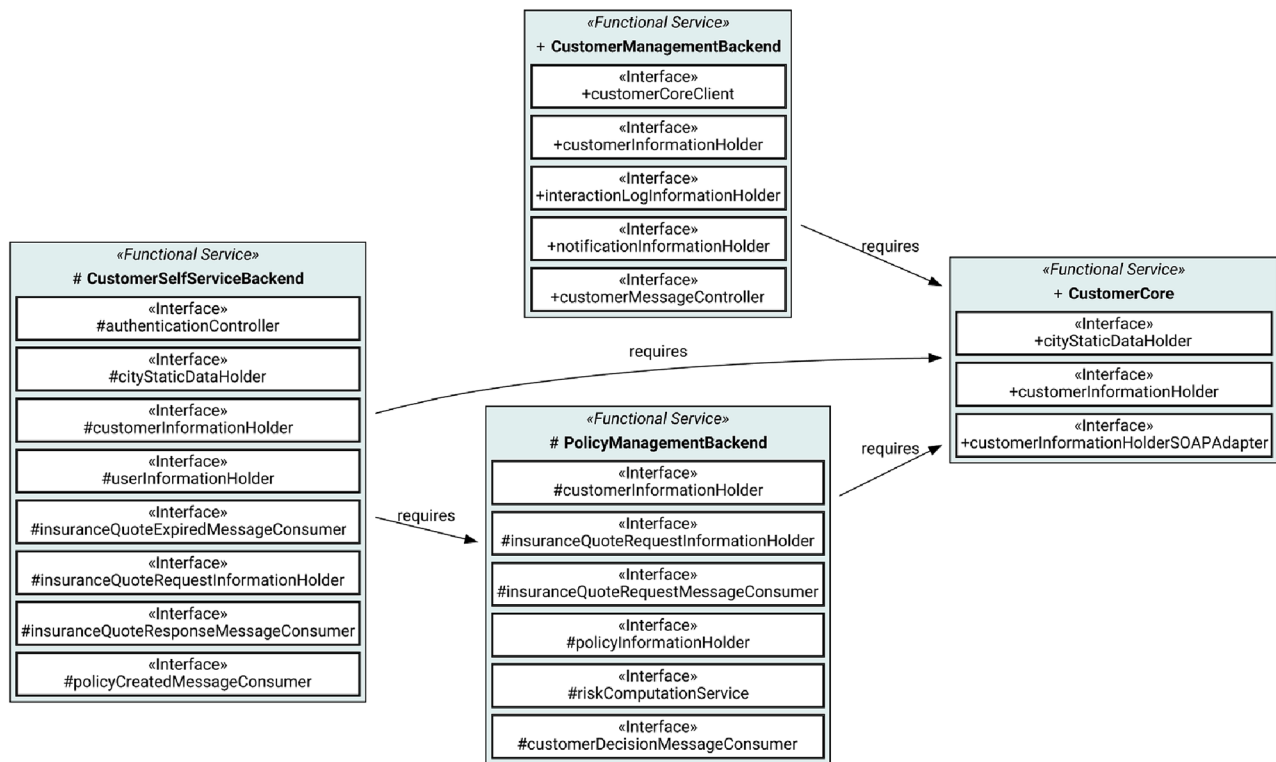


Fig. 8 Generated visualization of the architecture of the case study (cf. Section 5)

(cf. Section “[LEMMAs-Based Workflow for Coping with DevOps Challenges](#)”). The operation model is constructed using LEMMA’s Operation Modeling Language (OML) Section “[Microservice Architecture Model Construction](#)”). This approach specifically addresses the Operation and Deployment Stages of MSA (cf. Sections “[Operation Stage](#)” and “[Development Stage](#)”) and therefore addresses CA2 (cf. Section “[Study Results and Challenges](#)”) by providing functionalities for describing the microservice models’ deployment, including their dependencies to infrastructural services, e.g., API gateways, services discoveries, and databases. Additionally, OML abstracts from concrete technology-specific deployment configurations and reduces the overall complexity of deploying a microservice application.

To enable DevOps teams in SMOs to take full ownership of their respective services, which mitigates the need to apply specialized teams dedicated to operating the whole microservice application (cf. Section “[DevOps-Related Challenges in Microservice Architecture Engineering of SMOs](#)”), we have extended the OML with means to import other operation models as *nodes* and, therefore, nest operation specifications with each other. I.e. teams do not have to maintain individual models for infrastructure microservices, but can use the new mechanism to import the operation model, e.g., for a Eureka service discovery

(cf. Section “[Case Study](#)”), from a central model repository (cf. Section “[Assembling a Common Architecture Model from Distributed Microservice Models](#)”).

OML now enables the DevOps Team to describe the deployment of a microservice and all necessary dependencies. Listing 4 shows an excerpt of the operation model for the deployment of the `CustomerCore` microservice. Lines 1 and 2 of the listing import the `customerCore.services` model derived from the services’ Open API specification (cf. Section “[Derivation of Microservice Models from API Documentations](#)”). The following Lines 3–6 deal with the import of the technology for service deployment. The `Container_base` technology model uses Docker¹⁹ and Kubernetes²⁰ for service deployment. Lines 7 and 8 illustrate the new possibility to import other operation models as nodes by importing the `eureka.operation` model that describes the deployment of a service discovery by the Eureka²¹ technology.

¹⁹ <https://www.docker.com>.

²⁰ <https://kubernetes.io/>.

²¹ <https://github.com/Netflix/eureka>.


```

1 import microservices from "customerCore.services"
2   as customerCoreServices
3 import technology from "../technology/"
4   container_base.technology" as container_base
5 import technology from "../technology/"
6   javaWithSpring.technology" as protocolTechnology
7 import nodes from "../eureka-server/"
8   eurekaServer.operation" as eureka
9
10 @technology(container_base)
11 @technology(protocolTechnology)
12 container CustomerCoreContainer
13   deployment technology
14     container_base::_deployment.Kubernetes
15   deploys customerCoreServices::com.lakesidemutual.
16     customercore.CustomerCore
17   depends on nodes eureka::ServiceDiscovery {
18     customerCoreServices::com.lakesidemutual
19     .customercore.CustomerCore {
20       eurekaUri = "http://localhost:8761"
21       basic endpoints {
22         protocolTechnology::_protocols.rest :
23           "http://localhost:8110"; }
24     }
25 }

```

Listing 4 Excerpt of the `customerCore.operation` model from our case study.

Lines 10 and 11 assign the technology to the `CustomerCoreContainer` (Line 12). The container runs the deployed microservices and clusters deployment-relevant information, e.g., dependencies to infrastructural components such as databases, service-specific configurations, and protocol-specific endpoints. For this purpose, Lines 12–14 create the `CustomerCoreContainer` and assign the Kubernetes deployment technology which is imported from the `container_base.technology` model. The deployment of the `CustomerCore` microservice via the container is shown in Lines 15 and 16. The following Lines 17–25, show the dependency to the `ServiceDiscovery` imported from the `eureka.operation` model. In detail, Line 19 starts the service-specific configuration of the `CustomerCore` microservice by specifying the `eurekaUri` responsible for configuring the dependency to the `ServiceDiscovery`. The `CustomerCore` microservice exposes its functionality via a `rest` endpoint (Lines 21–23).

Besides modeling the deployment of microservice-specific configurations, OML also enables the DevOps team to specify infrastructural components' deployment, e.g., service discoveries and databases. Listing 5 describes the deployment of the `ServiceDiscovery`. Lines 1–3 import the `containerbase.technology` and `eureka.technology` models. The models include the specification of the technology used for the deployment of the `ServiceDiscovery`. Lines 4 and 5 import the `CustomerCore.operation` model (cf. Listing 4) because the `CustomerCore` microservice uses the service discovery. Lines 7 and 8 assign the imported technology to the `ServiceDiscovery`.

```

1 import technology from "docker.technology"
2   as containerTechnology
3 import technology from "eureka.technology" as Eureka
4 import nodes from "customerCore.operation"
5   as customerCore
6
7 @technology(containerTechnology)
8 @technology(Eureka)
9 ServiceDiscovery is Eureka::_infrastructure.Eureka
10   used by nodes customerCore::_CustomerCoreContainer {
11   ...
12   default values {
13     hostname="DiscoveryService"
14     port = 8761
15     ...
16   }
17 }

```

Listing 5 Excerpt of the service discovery operation model from our case study.

```

technology container_base {
  deployment technologies {
    Kubernetes {
      operation environments="openjdk:11-jdk-slim"
      default, "golang", "python3";}
    ...
  }
  operation aspects {
    //Aspect for configuring custom dockerfiles
    aspect Dockerfile<singleval> for containers,
    infrastructure {string content <mandatory>;}
    ...
  }
}

```

Listing 6 Excerpt of the `container_base.technology` model (`container_base.technology`) for generating deployment related artifacts.

Line 9 starts the actual specification of the `ServiceDiscovery`, which uses the imported `Eureka` technology. The following Line 10 contains the dependency to the `CustomerCoreContainer`, specified in Listing 4. The service-specific configuration of the `ServiceDiscovery` is set via the assignment of `default values` in Lines 12–16. Lines 13 and 14 set the actual `hostname` and `port` of the service.

Overall, LEMMA's OML enables the DevOps team to construct operation models which specify the deployment of microservices and their dependencies on the microservice application's infrastructural components. The operation models consist of the concepts of containers and infrastructure nodes. Containers (cf. Listing 4) specify the deployment of microservice, whereby infrastructure nodes contain the configuration for infrastructural components, e.g., API gateways, databases, and service discoveries (cf. Listing 5).




Fig. 9 Code generation pipeline for creating deployment relevant artifacts from LEMMA's operation models

Generating Code from Distributed Deployment Infrastructure Models

In Section “Enhancing Distributed Microservice Models with Deployment Infrastructure Models” we introduced OML as a methodology to describe the deployment of a service-based software system. In this subsection, we contribute a code generation pipeline for creating deployment-related artifacts based on the operation models using LEMMA's Model Processor (cf. Section “Microservice Architecture Model Processing”). As depicted in Fig. 9, the code generation pipeline consists of two consecutive stages.

The first stage of the code generation pipeline consists of a *model-to-model transformation* [13] transforming an operation model into an intermediate operation model in the sense of LEMMA's intermediate model processing (cf. Section “Microservice Architecture Model Processing”).

The second stage of the code generation pipeline deals with the creation of the deployment-relevant artifacts. Based on an intermediate operation model, the code generators already included in LEMMA (cf. Section “Language Ecosystem for Modeling Microservice Architecture”) provide a variety of different functionalities that are usually bound to a specific technology model. As already shown in Listing 4 and Listing 5, the described operation models both use the `container_base` technology model.

The `container_base` model clusters a technology stack suited for a service-based software system with focusing on container technologies [44] such as Docker, Docker-Compose, and Kubernetes. Listing 6 shows an excerpt of this specific technology model. Line 1 specifies the actual name of the model. Lines 2–7 describe the deployment technologies of the model, in this particular case Kubernetes. Additionally, Kubernetes supports the operation environments `golang`, `python3`, and `openjdk` as its default.

```

1 # syntax=docker/dockerfile:experimental
2 FROM maven:3.6.1-jdk-8-alpine AS build
3 ARG BASE=/usr/src/app
4 COPY pom.xml ${BASE}/
5 COPY src ${BASE}/src
6 RUN --mount=type=cache,target=/root/.m2
7     mvn -f ${BASE}/pom.xml install -DskipTests
8
9 FROM openjdk:8-jdk-alpine
10 COPY --from=build /usr/src/app/target
11     /dependency/BOOT-INF/lib/* /app/lib/
12 COPY --from=build /usr/src/app/target
13     /dependency/META-INF /app/META-INF
14 COPY --from=build /usr/src/app/target
15     /dependency/BOOT-INF/classes /app
16
17 EXPOSE 8110
18 ENTRYPOINT ["java", "-cp", "app:app/lib/*",
19     "com.lakesidemutual.customercore.
20     CustomerCoreApplication"]
  
```

Listing 7 Dockerfile for the customer core service created by the code generation pipeline.

The second part of Listing 6 contains the definition of operation aspects for further service deployment specification from Lines 8 to 13. Lines 10 and 11 define the `Dockerfile` aspect, which can be applied to containers in operation models. The aspect consists of a single attribute named `content` containing the actual content of the `Dockerfile`. Furthermore, the `content` attribute has the property `mandatory` to it, so it can only be configured a single time per container.

The Container Base Code Generator (CBCG)²² supports the `container_base` technology model (cf. Listing 6) to derive a set of deployment-related artifacts such as build scripts, Dockerfiles, and Kubernetes configurations. Furthermore, the CBCG integrates with existing microservice configuration files by extending them with operation-specific entries, e.g., service discovery addresses, as necessary.

Listing 7 shows a Dockerfile generated by the CBCG from the operation model in Listing 4. The Dockerfile contains a basic configuration consisting of a Docker image deduced from the modeled operation environment. In Lines

²² https://github.com/SeelabFhdo/lemma/tree/master/code%20generators/de.fhdo.lemma.model_processing.code_generation.container_base.

2–15 several artifacts are copied into the image. Line 17 configures the port 8110 on which the microservice is started. Finally, Lines 18–20 define the `entrypoint` of the Docker image to compile and run the microservice.

```

1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4  ...
5  labels:
6    app: customercorecontainer
7    name: customercorecontainer
8  spec:
9    replicas: 1
10   selector:
11     matchLabels:
12       app: customercorecontainer
13   ...

```

Listing 8 Excerpt of the deployment part of the Kubernetes file for the CustomerCore service.

While the CBCG is able to create a basic `Dockerfile`, the OML also integrates with LEMMA’s aspect mechanism so that it becomes possible to leverage the `Dockerfile` aspect from the `container_base` technology model (cf. Listing 6) to customize `Dockerfile` generation.

In addition to `Dockerfiles`, the CBCG may also generate Kubernetes deployment files. In general, a derived Kubernetes file consists of parts that concern deployment and service configuration. The deployment part shown in Listing 8 clusters the configuration of the Kubernetes pod²³ to which a microservice gets deployed. Line 1 defines the `apiVersion` the Kubernetes file uses. Line 2 contains the definition of the configuration `kind`: `deployment`. Lines 5–7 assign a name to the deployment, i.e., `customercorecontainer`. Line 8 determines the configuration of the Kubernetes deployment, specifically the set of `replicas` that shall be created for the deployment.

```

1  apiVersion: v1
2  kind: Service
3  metadata:
4  labels:
5    app: customercorecontainer
6    name: customercorecontainer
7  spec:
8    ports:
9      - name: "8110"
10        port: 8110
11        targetPort: 8110
12   ...

```

Listing 9 Excerpt of the service part of the Kubernetes file for the CustomerCore service.

Listing 9 contains the service part of the Kubernetes deployment and the configuration on how the microservices

application is exposed. Lines 1 and 2 cluster information about the `apiVersion` and configuration type of the Kubernetes file, i.e., `kind: Service`, followed by the name assignment in Lines 4–6. Additionally, the listing defines the exposure of the microservice via port 8110 in Lines 9–11.

In addition to the deployment-related generated artifacts, LEMMA’s code generation pipeline also supports the extension of existing service configurations. For this purpose, we implemented additional code generators for technologies like MongoDB,²⁴ MariaDB,²⁵ Zuul, and Eureka.

Listing 10 contains a variety of configuration options for Spring-based microservice implementations. The `spring.application.name` and `server.port` options in Lines 1 and 2 are derived from the modeled microservice’s name and its specified endpoint in the LEMMA models. Lines 3–7 result from the Eureka configuration shown in Listing 4. They configure the endpoints for connecting to the eureka service discovery.

```

1  spring.application.name=CustomerCoreService
2  server.port=8091
3  eureka.instance.preferIpAddress=true
4  eureka.client.registerWithEureka=true
5  eureka.client.fetchRegistry=true
6  eureka.client.serviceUrl.defaultZone=
7  ${EUREKA_URI:http://discovery-service:8761/eureka}

```

Listing 10 Excerpt of the CustomerCoreService configuration file.

Validation

In this section, we validate the present LEMMA extensions that implement the workflow (cf. Section “A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering”). To enable replicability of our results, we provide a validation package on GitHub.²⁶ To make the validation feasible, we first manually reconstructed the functional backend and infrastructure microservices of Lakeside Mutual (cf. Section “Case Study”) using a systematic process [62]. This step was necessary because the backend and infrastructure microservices of Lakeside Mutual are implemented in Java and not modeled with LEMMA. In detail, our reconstruction includes all four Java-based functional microservices and the infrastructural microservices Eureka Server and Spring Boot Admin (cf. Section “Case Study”).

²⁴ <https://www.mongodb.com/>.

²⁵ <https://mariadb.org/>.

²⁶ <https://github.com/SeelabFhdo/SN2021>.

²³ <https://kubernetes.io/docs/concepts/workloads/pods/>.

Table 3 Overview of the number of LoC of the different model artifacts involved in the LEMMA-based workflow

Service	Type	Viewpoint	LoC
Manually built models			
All Services	LEMMA	All	3702
All Services	LEMMA	Operation	311
Customer Core	LEMMA	All	588
Customer Core	LEMMA	Operation	50
Customer Management Backend	LEMMA	All	458
Customer Management Backend	LEMMA	Operation	49
API documentation			
Customer Core	OAS/JSON	Domain Data & Service	534
Customer Core	LEMMA	Domain Data & Service	171
Customer Management Backend	OAS/JSON	Domain Data & Service	496
Customer Management Backend	LEMMA	Domain Data & Service	174
Deployment specification			
All Services	Docker Kubernetes	Operation	285

In addition, we retrieved the current API documentation of Lakeside Mutual by putting the architecture into operation and triggering the generation of the documentation using prepared REST requests. At the end of this process, we could refer to the current API documentations of Lakeside Mutual's `Customer Core` and `Customer Management Backend` (cf. Fig. 3 in Section “[Case Study](#)”), which are the two components for which the application provides API documentation.

We then performed the individual generation steps of our workflow (cf. Section “[A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering](#)”) based on our reconstructed LEMMA models and the case study's API documentation. We illustrate the results of the application of our workflow as shown in Table 3 using the Lines of Code (LoC) metric.

As Table 3 shows, using the OAS-conform API documentation, we were able to generate 171 and 174 LoC of LEMMA Domain Data and Service files for the Customer Core and the Customer Management Backend microservices, respectively. Although the same operations and parameters for interfaces are present in the models generated by our workflow and the reconstructed LEMMA models, the LoC are higher in our reconstructed models. This is due to the fact that, e.g., the operation-related portion of LoC or technology-related annotations for databases are present in the manual models, but not in the generated ones, since no information on this is available from the API documentation.

Regarding the generation of deployment specifications, we were able to generate 285 lines of infrastructure code for Docker and Kubernetes from the reconstructed operation models of the functional microservices. Teams can abstract from technology-specific infrastructure code and, in combination with LEMMA's source code generators such as the

Java Base Generator [61], generate directly executable and deployable stubs of their services.

Discussion

The model-based workflow presented in Section “[A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering](#)” addresses the previously identified challenge areas (cf. Section “[DevOps-Related Challenges in Microservice Architecture Engineering of SMOs](#)”). In detail, the workflow provides means to establish a common understanding of architecture in an organization scaling to the level of multiple teams for the first time (CA1) and the complexity of operational aspects in microservice engineering (CA2).

We argue that by documenting the architecture in a centralized manner (cf. Sections “[Derivation of Microservice Models from API Documentations](#)” and “[Assembling a Common Architecture Model from Distributed Microservice Models](#)”), combined with the ability to visualize it (cf. Section “[Visualization of Microservice Architecture Models](#)”), teams and higher-level stakeholders, such as project sponsors, have a good basis for sharing knowledge and gaining insight into each other's development artifacts through the inherent abstraction property of the models [48]. *Box-and-line* diagrams, in particular, have the advantage that people can more easily grasp relations between concepts [13].

Another added value of our approach is the ability to seamlessly integrate deployment specifications into architecture models as a LEMMA operation model with the possibility to derive deployment configurations for heterogeneous deployment technologies, i.e., to generate them for Docker and Kubernetes (cf. Section “[Generating Code from Distributed Deployment Infrastructure Models](#)”). To this

regard, Combemale et al. [13] underline the added value of models to abstract complexity in the deployment process making the process more manageable. However, deployment technologies supported by our workflow constitute de-facto standards [69], LEMMA does limited justice to the heterogeneous technology landscape concerning cloud providers. In particular, we do not specifically address cloud-based deployment platforms such as AWS²⁷ or Azure.²⁸ Presumably, LEMMA is able to support such technologies through specific technology models (cf. Section “[Microservice Architecture Model Construction](#)”). In the future, we plan to address this limitation by providing LEMMA technology models and code generators for languages targeting the *Infrastructure as Code* [52] paradigm, e.g., Terraform [11]. As a result, LEMMA would support model-based deployment to a variety of cloud-based deployment platforms.

To implement and take advantage of the LEMMA-based workflow, team members need to learn and use a new technology with LEMMA. As the validation (cf. Section “[Validation](#)”) shows, teams can significantly increase efficiency through the available generation facilities of LEMMA. However, we need further empirical evaluation in practice (cf. Section “[Conclusion and Future Work](#)”) to more accurately assess in which cases the efficiency gains from better documentation, accessible architectural understanding, and generation of deployment specifications outweigh the effort required to learn LEMMA and in which cases they do not.

An important aspect on which the efficiency of the workflow depends is the organization-wide agreement on the level of detail of the models shared between teams. For example, if a very high level of detail is agreed upon, i.e., including as much information as possible from the source code in the models, as we applied to the reconstruction of the case study (cf. Table 3), generated artifacts must be more refined by the DevOps teams. This results in a higher effort. This can exemplified by the Customer Core Service (cf. “[Validation](#)”). The reconstructed model contains considerably more LoC, e.g., regarding technologies, than the generated model. In contrast, if the organization agrees on a low level of detail that, e.g., only considers technology-agnostic domain, service, and operation models (cf. Section “[Language Ecosystem for Modeling Microservice Architecture](#)”), very few adjustments to the generated models are necessary. This fact is also evident from our experience in performing the validation (cf. “[Validation](#)”). Although we were familiar with both LEMMA and the source code from the case study, the manual reconstruction was a tedious and time-consuming task of several hours compared to the automatic generation of the models.

²⁷ <https://aws.amazon.com/>.

²⁸ <https://azure.microsoft.com/>.

A technical limitation within the LEMMA-based workflow is the unidirectional artifact creation. Changes to the models currently have to be made by the team owning the corresponding microservice. However, to further extend a shared understanding of the architecture as well as to follow DevOps’ *minimize communication efforts* characteristic [21], it would be beneficial if other teams or stakeholders could request editing of services of other teams directly using the shared models, e.g., to add an attribute to an interface operation.

Related Work

In the following, we describe related work from the areas of service and operation modeling, comparable qualitative studies, and workflows for DevOps-oriented development of microservice architectures in the context of model-driven software engineering.

MSA Service Modeling Terzić et al. [82] present MicroBuilder, a tool that enables the modeling and generation of microservices. At its core, MicroBuilder comprises the MicroDSL modeling language. Like LEMMA, MicroDSL is based on the Eclipse Modeling Framework. Unlike LEMMA, however, MicroBuilder is closely linked to Java and Spring as specific technologies, so that the MicroDSL metamodel would have to be adapted for new technologies. MicroBuilder also addresses only the role of the developer and neglects stakeholders such as domain experts or operators. In addition, MicroBuilder does not address MSA’s characteristic of having multiple teams involved in the development process. Another model-based approach called MicroART [31] is provided by Granchelli et al. MicroART contains a DSL called MicroARTDSL which aims to capture architecture information. The purpose of MicroART is to recover microservice architectures through static and dynamic analysis. As such, MicroART can support organizations in raising a common architectural understanding similar to the visualization we proposed in Section “[Visualization of Microservice Architecture Models](#)”. However, MicroART does not provide a model-based workflow for the teams and lacks the rich ecosystem of LEMMA comprising means to also model and generate domain data, operational aspects, and different technologies. Alshuqayran et al. [1] present MiSAR an empirically derived approach for generating architectural models of microservice applications. Like LEMMA, MiSAR also leverages MDE and provides a metamodel centered around the microservice concept. It includes concepts for, e.g., interfaces, operations, and service types. However, MiSAR primarily focuses on reconstructing models from existing source code and does not provide a concrete syntax aligned with the metamodel like LEMMA’s textual notation. In addition, MiSAR has a cohesive metamodel for operation and service modeling. LEMMA,

on the other hand, supports separate model kinds derived from the roles in the MSA engineering process and also includes means for modeling domain data.

Qualitative Study Bogner et al. [9] describe a study related to our qualitative empirical analysis (cf. Section “[DevOps-Related Challenges in Microservice Architecture Engineering of SMOs](#)”) that includes 14 interviews with software architects. In contrast to our analysis, Bogner et al. do not focus on the challenges in the workflow of the organizations, but on the technologies used and software quality aspects. Another interview study was conducted by Haselböck et al. [34] focusing on software design aspects such as the sizing of microservices. A questionnaire based study on Bad Smells in MSA was conducted by Taibi et al. [78]. The study touches on organizational aspects and is included in our argumentation of the challenges (cf. Section “[Study Results and Challenges](#)”, but due to the study design as a questionnaire, the development process as a whole was not considered.

Development Workflows In the context of our proposed workflow (cf. Section “[LEMMA-Based Workflow for Coping with DevOps Challenges](#)”), there are several *large-scale agile* process models or methodologies that can foster the development of MSA by multiple DevOps teams. Examples include Scrum at Scale [77], the Spotify Model [71], or SAFe [66] (cf. Section [Organizational Aspects](#)). However, these approaches generally only become viable when an organization has at least 50 or more developers involved [18], and are therefore not suitable for SMOs facing the challenge of initially scaling their small organization from one to two or three teams. In addition, the aforementioned approaches address development at an organizational level and do not address development practices. Therefore, we expect our proposed workflow (cf. Section “[LEMMA-Based Workflow for Coping with DevOps Challenges](#)”) to integrate well with the stated large-scale approaches.

MSA Operation Modeling The essential deployment metamodel (EDMM) [83] is an approach that combines existential components of the deployment of a software system in a metamodel, taking into account concepts such as *configuration management* [16] and *infrastructure as code* [52]. EDMM makes a specific mapping concerning the technology used for the software system’s provisioning process based on the metamodel. For deploying the microservice application, EDMM supports technologies like Puppet,²⁹ Terraform,³⁰ AWS Cloud Formation,³¹ and Cloudify.³² Unlike EDMM, LEMMA addresses the deployment of

service-based systems and their data structures and service composition. Besides, EDMM provides mapping concerning specific cloud providers. On the other hand, LEMMA provides technology-specific provisioning artifacts that can be used with different cloud providers. DICER [2] represents an approach based on technology-independent models for the generation of infrastructure as code and is used to deploy the software system. DICER models encapsulate monitoring, self-adaptation, configuration management, server deployment, and software system deployment. Also, DICER fosters the transformation of models into artifacts for service deployment using TOSCA³³ and other technologies. The functional scope of DICER relates exclusively to the provisioning or operation of the software system. Furthermore, DICER does not support the modeling of data structures or service composition. Like LEMMA, DICER also provides technology-specific artifacts that can be used for the deployment process. Additionally, it also provides a graphical representation in the form of UML deployment diagrams, which LEMMA does not provide on an operational view.

Conclusion and Future Work

In this paper, we have identified two key challenge areas for SMOs through an empirical analysis of an interview study (cf. Section 4). First, it is challenging for SMOs to develop and maintain a common understanding of architecture in an organization that is scaling to multiple teams for the first time through the application of MSA. Second, deployment in particular seems challenging due to its complexity, so SMOs tend to constitute special operation teams contrary to the microservice ownership principle (cf. Section “[Organizational Aspects](#)”). This is detrimental to the implementation of DevOps practices and the benefits hoped for within the teams.

To address these two challenge areas, we have presented a model-driven workflow based on LEMMA (cf. Section “[Language Ecosystem for Modeling Microservice Architecture](#)”) for developing microservice architectures (cf. Section “[A Model-Driven Workflow for Coping with DevOps-Related Challenges in Microservice Architecture Engineering](#)”) and elaborated on the components we have added to LEMMA to support this workflow. The components comprise (i) interoperability bridges between OpenAPI and LEMMA models (cf. Section “[Derivation of Microservice Models from API Documentations](#)”); (ii) an extension to the Service Modeling Language to allow the import of remote models (cf. Section “[Assembling a Common Architecture Model from Distributed Microservice Models](#)”); (iii) a model processor to visualize microservice

²⁹ <https://puppet.com/>.

³⁰ <https://www.terraform.io/>.

³¹ <https://aws.amazon.com/>.

³² <https://cloudify.co/>.

³³ <https://cloudify.co/tosca/>.

architectures (cf. Section “[Visualization of Microservice Architecture Models](#)”); (iv) enhancing the Operation Modeling Language through the ability to import infrastructural nodes (cf. Section “[Enhancing Distributed Microservice Models with Deployment Infrastructure Models](#)”); and (v) code generators for microservice deployment and operation (cf. Section “[Generating Code from Distributed Deployment Infrastructure Models](#)”).

For future work we plan to conduct a qualitative observation and interview study which aims to evaluate the proposed workflow in practice. In particular, we plan to evaluate the integrability into CI/CD pipelines as well as the adoption of our proposed mixture of manual and automatic steps within the workflow. Furthermore, we are also currently exploring the possibility of automating our reconstruction process for LEMMA models [61]. Finally, we would like to investigate whether our proposed approach can be applied to other software architectures.

Regarding the presented LEMMA extensions, we are going to mature the prototypical development and improve accessibility for users, e.g., by providing a web-based dashboard to trigger the visualizer and access the generated images. Furthermore, we would like to develop LEMMA’s means to support a common architectural understanding in an organization not only through the presented visualization but also through analytical means such as code metrics.

Author Contributions Not applicable.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability Statement We strive to make our research as replicable as possible. To this end we have provided all modeling artifacts, images, scripts, and a documentation in a publicly available GitHub repository, which is accessible at <https://github.com/SeelabFhdo/SN2021>. Due to confidentiality agreements, we are not allowed to provide the transcribed interviews of our qualitative study. On request we are able to provide anonymized text passages should questions arise by the reviewers or editors.

Declarations

Funding Not applicable.

Conflict of interest/Competing interests On behalf of all authors, the corresponding author states that there is no conflict of interest.

Code availability Within the article, we use and extend the modeling ecosystem LEMMA developed by our research group. LEMMA is under the MIT license available on GitHub through this link <https://github.com/SeelabFhdo/lemma>. The specific source code contributions for this article are also ready together with the data and other material in this prepared GitHub repository <https://github.com/SeelabFhdo/SN2021>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alshuqayran N, Ali N, Evans R. Towards micro service architecture recovery: An empirical study. In: 2018 IEEE International Conference on software architecture (ICSA). IEEE. 2018. p. 47–56. <https://doi.org/10.1109/ICSA.2018.00014>.
2. Artač M, Borovšak T, Nitto ED, Guerriero M, Perez-Palacin D, Tamburri DA. Infrastructure-as-Code for data-intensive architectures: a model-driven development approach. In: 2018 IEEE International conference on software architecture (ICSA). IEEE. 2018. p. 156–65. <https://doi.org/10.1109/ICSA.2018.00025>.
3. Athanasopoulos D, Zarras AV, Miskos G, Issarny V, Vassiliadis P. Cohesion-driven decomposition of service interfaces without access to source code. *IEEE Trans Serv Comput.* 2015;8(4):550–62.
4. Balalaie A, Heydarnoori A, Jamshidi P. Microservices architecture enables DevOps: migration to a cloud-native architecture. *IEEE Softw.* 2016;33(3):42–52. <https://doi.org/10.1109/MS.2016.64>.
5. Balalaie A, Heydarnoori A, Jamshidi P. Migrating to cloud-native architectures using microservices: an experience report. In: Celesti A, Leitner P, editors. *Advances in service-oriented and cloud computing*. Cham: Springer; 2016. p. 201–15.
6. Bass L, Weber I, Zhu L. *DevOps: a software architect's perspective*. 1st ed. Boston: Addison-Wesley Educational; 2015.
7. Ben-Kiki O, Evans C. *Ingy dot Net: YAML ain't markup language (YAML) Version 1.2. 3rd Edition, Patched at 2009-10-01*. Tech. rep., <https://yaml.org>, 2009. Accessed 2 Apr 2021.
8. Bogner J. *On the evolvability assurance of microservices: metrics, scenarios, and patterns*. Ph.D. thesis. Stuttgart: University of Stuttgart; 2020.
9. Bogner J, Fritzsche J, Wagner S, Zimmermann A. Microservices in industry: Insights into technologies, characteristics, and software quality. In: 2019 IEEE International Conference on software architecture companion (ICSA-C). IEEE, 2019; p. 187–95. <https://doi.org/10.1109/ICSA-C.2019.00041>.
10. Brambilla M, Cabot J, Wimmer M. *Model-driven software engineering in practice*. 2nd ed. San Rafael: Morgan & Claypool; 2017.
11. Brikman Y. *Terraform: up & running: writing infrastructure as code*. Sebastopol: O'Reilly Media; 2019.
12. Buchgeher G, Winterer M, Weinreich R, Luger J, Wingelhofer R, Aistleitner M. Microservices in a small development organization. In: Lopes A, de Lemos R, editors. *Software architecture*. Cham: Springer International Publishing; 2017. p. 208–15.
13. Combemale B, France RB, Jézéquel JM, Rumpe B, Steel J, Vojtisek D. *Engineering modeling languages: turning domain knowledge into tools*. 1st ed. Philadelphia: Chapman & Hall/CRC; 2017.
14. Conboy K, Carroll N. Implementing large-scale agile frameworks: challenges and recommendations. *IEEE Softw.* 2019;36(2):44–50.

15. Conway ME. How do committees invent? *Datamation*. 1968;14(4):28–31.
16. Delaet T, Joosen W, Vanbrabant B. A survey of system configuration tools. In: *Proceedings of the 23rd Large Installations Systems Administration (LISA) Conference*. Usenix association 2010; p. 1–14. <https://doi.org/10.5555/1924976.1924977>.
17. Di Francesco P, Malavolta I, Lago P. Research on architecting microservices: trends, focus, and potential for industrial adoption. In: *2017 IEEE International Conference on software architecture (ICSA)*. IEEE, 2017; p. 21–30. <https://doi.org/10.1109/ICSA.2017.24>.
18. Dikert K, Paasivaara M, Lassenius C. Challenges and success factors for large-scale agile transformations: a systematic literature review. *J Syst Softw*. 2016;119:87–108.
19. Dragoni N, Giallorenzo S, Lafuente AL, Mazzara M, Montesi F, Mustafin R, Safina L. *Microservices: yesterday, today, and tomorrow*. In: Mazzara M, Meyer B, editors. *Present and ulterior software engineering*. Cham: Springer International Publishing; 2017. p. 195–216.
20. Dragoni N, Lanese I, Larsen ST, Mazzara M, Mustafin R, Safina L. *Microservices: how to make your application scale*. In: Petrenko AK, Voronkov A, editors. *Perspectives of system informatics*. Cham: Springer; 2018. p. 95–104.
21. Ebert C, Gallardo G, Hernantes J, Serrano N. *Devops*. IEEE Softw. 2016;33(3):94–100. <https://doi.org/10.1109/MS.2016.68>.
22. Ecma International. *The JSON data interchange syntax*. Standard ECMA-404. Ecma International; 2017.
23. Ellson J, Gansner E, Koutsofios L, North SC, Woodhull G. *Graphviz—open source graph drawing tools*. In: Mutzel P, Jünger M, Leipert S, editors. *Graph drawing*. Berlin: Springer Berlin Heidelberg; 2002. p. 483–4.
24. Engel T, Langermeier M, Bauer B, Hofmann A. Evaluation of microservice architectures: a metric and tool-based approach. In: Mendling J, Mouratidis H, editors. *Information systems in the big data era*. Cham: Springer; 2018. p. 74–89.
25. Erl T. *Service-oriented architecture (SOA): concepts, technology and design*. 1st ed. Philadelphia: Prentice Hall; 2005.
26. Evans E. *Domain-driven design*. 1st ed. Boston: Addison-Wesley; 2004.
27. Fielding RT. *Architectural styles and the design of network-based software architectures*. Ph.D. thesis (2000)
28. Francesco PD, Lago P, Malavolta I. Migrating towards microservice architectures: an industrial survey. In: *2018 IEEE International Conference on software architecture (ICSA)*. IEEE, 2018; p. 29–38. <https://doi.org/10.1109/ICSA.2018.00012>.
29. Fritzsche J, Bogner J, Zimmermann A, Wagner S. From monolith to microservices: a classification of refactoring approaches. In: Bruel JM, Mazzara M, Meyer B, editors. *Software engineering aspects of continuous development and new paradigms of software production and deployment*. Cham: Springer International Publishing; 2019. p. 128–41.
30. Gosling J, Joy B, Steele G, Bracha G, Buckley A, Smith D. *The Java language specification: Java se 13 edition*. Specification JSR-388 Java SE 13, Oracle America, Inc. (2019)
31. Granchelli G, Cardarelli M, Francesco PD, Malavolta I, Iovino L, Salle AD. Towards recovering the software architecture of microservice-based systems. In: *2017 IEEE International Conference on software architecture workshops (ICSAW)*. IEEE 2017; p. 46–53. <https://doi.org/10.1109/ICSAW.2017.48>.
32. Haesen R, Snoeck M, Lemahieu W, Poelmans S. On the definition of service granularity and its architectural impact. In: Bellahsene Z, Léonard M, editors. *Advanced information systems engineering*. Berlin: Springer; 2008. p. 375–89.
33. Haselböck S, Weinreich R, Buchgeher G. Decision models for microservices: design areas, stakeholders, use cases, and requirements. In: Lopes A, de Lemos R, editors. *Software architecture*. Cham: Springer; 2017. p. 155–70.
34. Haselböck S, Weinreich R, Buchgeher G. An expert interview study on areas of microservice design. In: *2018 IEEE 11th Conference on service-oriented computing and applications (SOCA)*, 2018; p. 137–44. <https://doi.org/10.1109/SOCA.2018.00028>.
35. Haupt F, Leymann F, Scherer A, Vukojevic-Haupt K. A framework for the structural analysis of REST APIs. In: *2017 IEEE international conference on software architecture (ICSA)*. Springer; 2017. p. 55–8. <https://doi.org/10.1109/ICSA.2017.40>.
36. Heaton J. *Secondary analysis of qualitative data*. In *The SAGE handbook of social research methods*. Thousand Oaks: SAGE Publications Ltd; 2008. p. 506–519. <https://doi.org/10.4135/9781446212165>.
37. Hirzalla M, Cleland-Huang J, Arsanjani A. A metrics suite for evaluating flexibility and complexity in service oriented architectures. In: Feuerlicht G, Lamersdorf W, editors. *Service-oriented computing—ICSOC 2008 workshops*. Berlin: Springer; 2009. p. 41–52.
38. Hofmeister C, Kruchten P, Nord RL, Obbink H, Ran A, America P. A general model of software architecture design derived from five industrial approaches. *J Syst Softw*. 2007;80(1):106–26.
39. ISO/IEC. *Systems and software engineering—systems and software quality requirements and evaluation (SQuARE)— system and software quality models*. Standard ISO/IEC 25010:2011(E), International Organization for Standardization/International Electrotechnical Commission. 2011.
40. ISO/IEC/IEEE. *Systems and software engineering—architecture description*. Standard ISO/IEC/IEEE 42010:2011(E), International Organization for Standardization/ International Electrotechnical Commission/Institute of Electrical and Electronics Engineers. 2011.
41. Jézéquel JM, Combemale B, Derrien S, Guy C, Rajopadhye S. Bridging the chasm between MDE and the world of compilation. *Softw Syst Model*. 2012;11(4):581–97. <https://doi.org/10.1007/s10270-012-0266-8>.
42. Johnson RE, Foote B. Designing reusable classes. *J Object-Orient Program*. 1988;1(2):22–35.
43. Jongeling R, Carlson J, Cicchetti A. In: *Impediments to introducing continuous integration for model-based development in industry*. In: *2019 45th Euromicro Conference on software engineering and advanced applications (SEAA)*, IEEE, 2019; p. 434–41. <https://doi.org/10.1109/SEAA.2019.00071>.
44. Kang H, Le M, Tao S. Container and microservice driven design for cloud infrastructure DevOps. In: *2016 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2016; p. 202–11. <https://doi.org/10.1109/IC2E.2016.26>.
45. Knoche H, Hasselbring W. Drivers and barriers for microservice adoption—a survey among professionals in Germany. *Enterp Model Inf Syst Arch*. 2019;14(1):1–35.
46. Kratzke N, Quint PC. Investigation of impacts on network performance in the advance of a microservice design. In: Helfert M, Ferguson D, Méndez Muñoz V, Cardoso J, editors. *Cloud computing and services science*. Cham: Springer; 2017. p. 187–208.
47. Lano K, Kolahdouz-Rahimi S. Model-transformation design patterns. *IEEE Trans Softw Eng*. 2014;40(12):1224–59.
48. Ludewig J. *Models in software engineering—an introduction*. *Softw Syst Model*. 2003;2(1):5–14. <https://doi.org/10.1007/s10270-003-0020-3>.
49. Luz WP, Pinto G, Bonifácio R. In: *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM '18 Building a collaborative culture: a grounded theory of well succeeded DevOps adoption in practice*, ACM, New York, NY, USA, 2018; p. 6:1–6:10. <https://doi.org/10.1145/3239235.3240299>.
50. Martini A, Pareto L, Bosch J. Communication factors for speed and reuse in large-scale agile software development. In: *Proceedings of the 17th International Software Product Line Conference*,

- SPLC '13. Association for Computing Machinery, New York, NY, USA, 2013; p. 42–51. <https://doi.org/10.1145/2491627.2491642>.
51. Michail D, Kinable J, Naveh B, Sichi JV. JGraphT—a java library for graph data structures and algorithms. *ACM Trans Math Softw.* 2020;46(2):1–29. <https://doi.org/10.1145/3381449>.
 52. Morris K. *Infrastructure as code: managing servers in the cloud*. 1st ed. Sebastopol: O'Reilly; 2016.
 53. Nadareishvili I, Mitra R, McLarty M, Amundsen M. *Microservice architecture: aligning principles, practices, and culture*. 1st ed. Sebastopol: O'Reilly; 2016.
 54. Newman S. *Building microservices: designing fine-grained systems*. 1st ed. Sebastopol: O'Reilly; 2015.
 55. OMG. *OMG Unified Modeling Language (OMG UML) version 2.5.1. Standard formal/17-12-05, Object Management Group*. 2017.
 56. OpenAPI Initiative. *OpenAPI specification 3.0.3*. 2020.
 57. Papazoglou MP. *Web services: principles and technology*. 1st ed. London: Pearson Education; 2008.
 58. Peltonen S, Mezzalana L, Taibi D. Motivations, benefits, and issues for adopting micro-frontends: a multivocal literature review. *Inf Softw Technol.* 2021;136:106571. <https://doi.org/10.1016/j.infsof.2021.106571>. <https://www.sciencedirect.com/science/article/pii/S0950584921000549>.
 59. Rademacher F, Sachweh S, Zündorf A. Towards a UML profile for domain-driven design of microservice architectures. In: Cerone A, Roveri M, editors. *Software engineering and formal methods*. Cham: Springer; 2018. p. 230–45.
 60. Rademacher F, Sachweh S, Zündorf A. Aspect-oriented modeling of technology heterogeneity in microservice architecture. In: 2019 IEEE International Conference on software architecture (ICSA). IEEE, 2019; p. 21–30. <https://doi.org/10.1109/ICSA.2019.00011>.
 61. Rademacher F, Sachweh S, Zündorf A. Deriving microservice code from underspecified domain models using DevOps-enabled modeling languages and model transformations. In: 2020 46th Euromicro Conference on software engineering and advanced applications (SEAA). IEEE, 2020; p. 229–36. <https://doi.org/10.1109/SEAA51224.2020.00047>.
 62. Rademacher F, Sachweh S, Zündorf A. A modeling method for systematic architecture reconstruction of microservice-based software systems. In: Nurcan S, Reinhartz-Berger I, Soffer P, Zdravkovic J, editors. *Enterprise: business-process and information systems modeling*. Berlin: Springer International Publishing; 2020. p. 311–26.
 63. Rademacher F, Sorgalla J, Wizenty P, Sachweh S, Zündorf A. Graphical and textual model-driven microservice development. In: Bucchiarone A, Dragoni N, Dustdar S, Lago P, Mazzara M, Rivera V, Sadovykh A, editors. *Microservices: science and engineering*. Berlin: Springer; 2020. p. 147–79.
 64. Richards M. *Microservices vs. service-oriented architecture*. 1st ed. Newton: O'Reilly; 2016.
 65. Riungu-Kalliosaari L, Mäkinen S, Lwakatara LE, Tiihonen J, Männistö T. DevOps adoption benefits and challenges in practice: a case study. In: Abrahamsson P, Jedlitschka A, Nguyen Duc A, Felderer M, Amasaki S, Mikkonen T, editors. *Product-focused software process improvement*. Cham: Springer; 2016. p. 590–7.
 66. Scaled Agile, Inc. *Achieving business agility with safe® 5.0*. 2019.
 67. Schermann G, Cito J, Leitner P. All the services large and micro: revisiting industrial practice in services computing. In: Norta A, Gaaloul W, Gangadharan GR, Dam HK, editors. *Service-oriented computing—ICSOC 2015 workshops*. Berlin: Springer; 2016. p. 36–47.
 68. Schwaber K, Beedle M. *Agile software development with scrum*, vol. 1. Upper Saddle River: Prentice Hall; 2002.
 69. Shah J, Dubaria D. Building modern clouds: using docker, kubernetes & google cloud platform. In: 2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC). IEEE, 2019; p. 0184–0189. <https://doi.org/10.1109/CCWC.2019.8666479>.
 70. Shull F, Singer J, Sjøberg DIK, editors. *Guide to advanced empirical software engineering*. London: Springer; 2008. <https://doi.org/10.1007/978-1-84800-044-5>.
 71. Smite D, Moe NB, Levinta G, Floryan M. Spotify guilds: how to succeed with knowledge sharing in large-scale agile organizations. *IEEE Softw.* 2019;36(2):51–7.
 72. Sobernig S, Zdun U. Inversion-of-control layer. In: *Proceedings of the 15th European Conference on Pattern Languages of Programs, EuroPLoP '10*. ACM, New York, NY, USA; 2010. p. 1–22. <https://doi.org/10.1145/2328909.2328935>.
 73. Soldani J, Tamburri DA, Heuvel WJVD. The pains and gains of microservices: a systematic grey literature review. *J Syst Sofw.* 2018;146:215–32. <https://doi.org/10.1016/j.jss.2018.09.082>. <http://www.sciencedirect.com/science/article/pii/S0164121218302139>.
 74. Soliman M, Riebisch M, Zdun U. Enriching architecture knowledge with technology design decisions. In: 2015 12th Working IEEE/IFIP Conference on software architecture. IEEE, 2015; p. 135–144. <https://doi.org/10.1109/WICSA.2015.14>.
 75. Sorgalla J, Sachweh S, Zündorf A. Exploring the microservice development process in small and medium-sized organizations. In: Morisio M, Torchiano M, Jedlitschka A, editors. *Product-focused software process improvement*. Cham: Springer International Publishing; 2020. p. 453–60.
 76. Stocker M, Zimmermann O, Lübke D, Zdun U, Pautasso C. Interface quality patterns—communicating and improving the quality of microservices apis. In: *23rd European Conference on pattern languages of programs 2018 (2018)*. <https://doi.org/10.1145/3282308.3282319>. <http://eprints.cs.univie.ac.at/5661/>.
 77. Sutherland J. *The scrum@scale guide version, vol. 2*. 2020.
 78. Taibi D, Lenarduzzi V. On the definition of microservice bad smells. *IEEE Softw.* 2018;35(3):56–62. <https://doi.org/10.1109/MS.2018.2141031>.
 79. Taibi D, Lenarduzzi V, Pahl C. Processes, motivations, and issues for migrating to microservices architectures: an empirical investigation. *IEEE Cloud Computing.* 2017;4(5):22–32. <https://doi.org/10.1109/MCC.2017.4250931>.
 80. Taibi D, Lenarduzzi V, Pahl C. Continuous architecting with microservices and DevOps: a systematic mapping study. In: Muñoz VM, Ferguson D, Helfert M, Pahl C, editors. *Cloud computing and services science*. Cham: Springer; 2019. p. 126–51.
 81. Taibi D, Lenarduzzi V, Pahl C. *Microservices anti-patterns: a taxonomy*. Cham: Springer; 2020. p. 111–28.
 82. Terzić B, Dimitrieski V, Kordić S, Milosavljević G, Luković I. Development and evaluation of MicroBuilder: a model-driven tool for the specification of REST microservice software architectures. *Enterp Inf Syst.* 2018;12(8–9):1034–57. <https://doi.org/10.1080/17517575.2018.1460766>.
 83. Wurster M, Breitenbücher U, Falkenthal M, Krieger C, Leymann F, Saatkamp K, Soldani J. The essential deployment metamodel: a systematic review of deployment automation technologies. *SICS Softw-Intensive Cyber-Phys Syst.* 2020;35(1):63–75.
 84. Yin RK. *Case study research and applications: design and methods*. 6th ed. Thousand Oaks: SAGE Publications; 2017.
 85. Zimmermann O. *Microservices tenets*. *Comput Sci Res Dev.* 2017;32(3–4):301–10. <https://doi.org/10.1007/s00450-016-0337-0>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.