Continuous API Evolution in Heterogenous Enterprise Software Systems

Holger Knoche
Software Engineering Group
Kiel University
Kiel, Germany
hkn@informatik.uni-kiel.de

Wilhelm Hasselbring
Software Engineering Group
Kiel University
Kiel, Germany
wha@informatik.uni-kiel.de

Abstract—The ability to independently deploy parts of a software system is one of the cornerstones of modern software development, and allows for these parts to evolve independently and at different speeds.

A major challenge of such independent deployment, however, is to ensure that despite their individual evolution, the interfaces between interacting parts remain compatible. This is especially important for enterprise software systems, which are often highly integrated and based on heterogenous IT infrastructures.

Although several approaches for interface evolution have been proposed, many of these rely on the developer to adhere to certain rules, but provide little guidance for doing so. In this paper, we present an approach for interface evolution that is easy to use for developers, and also addresses typical challenges of heterogenous enterprise software, especially legacy system integration.

I. INTRODUCTION

Many modern approaches to software engineering and delivery rely on the ability to deploy parts of a software system independently of each other. For instance, Newman [1] highlights independent deployability as one of the key advantages of microservices. Practices like Continous Delivery [2], which aim at fast feedback from production to development, would be impossible to implement without this ability.

In recent years, we have observed that also more conservative companies, such as banks and insurance companies, have begun to embrace such practices for their internal software systems [3], in particular for migrating legacy systems towards microservices [4]. There are several reasons for this development; in our case study from the insurance domain, which we will present later on in further detail, the key drivers are the desire to shorten the time-to-market, implement agile practices, and allow applications to evolve at different speeds.

However, said applications are often highly integrated, and interact with each other in complex ways. One of the key challenges is therefore to ensure that the applications remain compatible with each other at all times, even when they are independently deployed. Furthermore, some of these applications are several decades old and based on aged technologies or platforms. For such applications, libraries and utilities that are commonplace on modern platforms may not be available. Therefore, approaches are needed that can be implemented even on dated implementation platforms with acceptable effort.

In this paper, we present an approach that allows to continuously evolve the interfaces of a software application while ensuring the compatiblity with dependent applications. In particular, we aim to address the following goals:

1) The approach should be easy to use for the developers, and help to avoid common mistakes
2) The approach should have as little impact as possible on the developer's day-to-day work
3) The approach must also support aged implementation technologies and platforms
4) The runtime performance impact of the approach must be acceptable

The remainder of this paper is structured as follows. In Section II, we introduce the case study that motivated this work. An evolution example is given in Section III. Existing approaches to API evolution are discussed in Section IV. Our approach, which constitutes the main contribution of this paper, is presented in Section V, and a short evaluation is described in Section VI. Related work is discussed in Section VII, and Section VIII concludes the paper.

II. CASE STUDY

The underlying case study for this work is the core insurance software system of a German insurance company. The software system has in large parts been developed by the company itself, and consists of more than 30 individual, but highly integrated applications. Development of the first applications started in the 1970s in COBOL on mainframe computers, and even today, the vast majority of the codebase (about 14 million lines of code) is COBOL code. In the early 2000s, Java EE was adopted as a new implementation platform, and newer applications have been developed solely in Java. As of today, these applications amount to about 2 million lines of code.

In order to establish cross-platform standards, the company adopted Model-Driven Software Development (MDSD) [5] in the 2000s. A considerable amount of code is generated for both Java and COBOL, which is not included in the previous figures. In particular, code for invoking COBOL services from Java and vice-versa is fully generated. Modeling is currently performed with a graphical, UML-based modeling tool.

The majority of the COBOL programs is still run on mainframe computers, both in traditional batch processing

as well as interactive applications. The remaining COBOL programs run on Linux servers. Java EE applications are run on traditional Java application servers; some components are also deployed within an OSGi container on the mainframe.

The current release process mostly follows a traditional waterfall model with one major release every three months. As a consequence, it usually takes several months to ship a new feature into production. It should be noted that most of these applications are only used by professional users, who strongly prefer a stable environment to work in. Therefore, a high rate of change is not necessarily desirable in this setting.

However, the current process is particularly cumbersome for changes that need to be implemented at a specific point in time, such as legal requirements. Such changes are quite common in the financial services industry. Obviously, such a change needs to be part of the last release *before* the target date, and is therefore usually subject to a considerably higher delay. This can be challenging with requirements that come at short notice.

Furthermore, it has been observed that the applications evolve at different speeds. On the one hand, there are very mature applications that change seldomly, but irregularly, e.g., to keep up with changes in regulations. Other applications, on the other hand, need to evolve frequently. The fixed release schedule does not really match the needs of both these application types, and frequently causes friction.

In order to address these issues, it was decided to move towards a more flexible release model, and, in particular, to enable the individual applications to be deployed independently. However, while the underlying technology platforms provide facilities for dynamic re-deployment, the applications themselves are currently too tightly coupled to allow for independent evolution.

There are two major reasons for this tight coupling. The first is due to the fact that COBOL programs operate on "bare" memory, and therefore require all data to conform to an exact memory layout to function properly. Although this largely also applies to other languages such as C, COBOL lacks language features (such as user-defined functions) to properly abstract from this layout. Furthermore, all data fields (including character strings and tables) have static sizes in COBOL.[1] As a consequence, even changing the length of a string field breaks the memory layout and thus requires to recompile and redeploy all affected programs.

The second reason results from the way code generation is currently handled. As previously mentioned, client access code is generated automatically. This code generation is part of the provider build, and the resulting artifacts are published as libraries for use by the clients. While this approach is convenient for the developers, it leads to a significant sharing of code between applications. Such sharing can (and, often enough, does) cause problems, especially when two embedded libraries rely on incompatible versions of a third library.

---

[1] There are some exceptions, such as unbounded tables. However, to our knowledge, these features are seldomly used in practice.

III. A Simple Evolution Example

In order to better illustrate our context of evolution, consider a simple example. Note that, for the sake of brevity, we omit common attributes like IDs.

An application that manages customer data provides an API consisting of two types, as depicted in Figure 1a: The `Customer` type consists of the customer's first and last name, its gender, and its postal address. The latter is modeled by the second type, `Address`, that represents a typical street address (street name, number, city name, and postal code). As the API was inspired by an old (i.e., pre-Java 5) Java class structure, the customer's gender is modeled as an integer instead of an enumeration. The API provides two service operations. One inserts or updates a customer with given data ("upsert"), and the other formats an address according to the rules prescribed by the local postage service.

Due to requests from client applications, the following changes are made to the API:

1) *Field Addition:* The customer's date of birth, which is already present in the existing data, is added to the `Customer` type.
2) *Field Renaming:* As customers may have secondary addresses, a new list field `secondaryAddresses` is added to the `Customer` type. As part of this change, the `address` field is renamed to `primaryAddress`.
3) *Type Change:* A `Gender` enumeration is created, and the type of the field `gender` is changed accordingly.
4) *New Enum Constant:* A third gender is added to the `Gender` enumeration.
5) *New Specialization:* Support for P.O. box addresses is added. These consist of a box number, a postal code, and a city name. Therefore, the former `Address` type is renamed to `StreetAddress`, a new type `POBoxAddress` is created, and the common fields are moved to a common a supertype `Address`.

The resulting API is shown in Figure 1b. In the following section, we will investigate how these changes could be realized using the state of the art, and highlight the shortcomings we intend to address.

IV. API Evolution

A. API Evolution Concepts and Patterns

The main goal of API evolution is to apply changes to APIs without breaking the clients, i.e., the clients can remain unchanged and will still work as intended. A fundamental pattern to allow for such changes is the *tolerant reader pattern* [6]. The core idea of this pattern is that an application (*reader*) that receives data from another application (*writer*) does not require the data to exactly match its expectation, but tolerates certain deviations. For example, additional data fields can usually be simply and safely ignored by the reader.

In order to identify such additional fields, however, the reader requires knowledge about the underlying schema of the data, such as field names. Therefore, data formats with embedded field names are commonly used, e.g., XML or

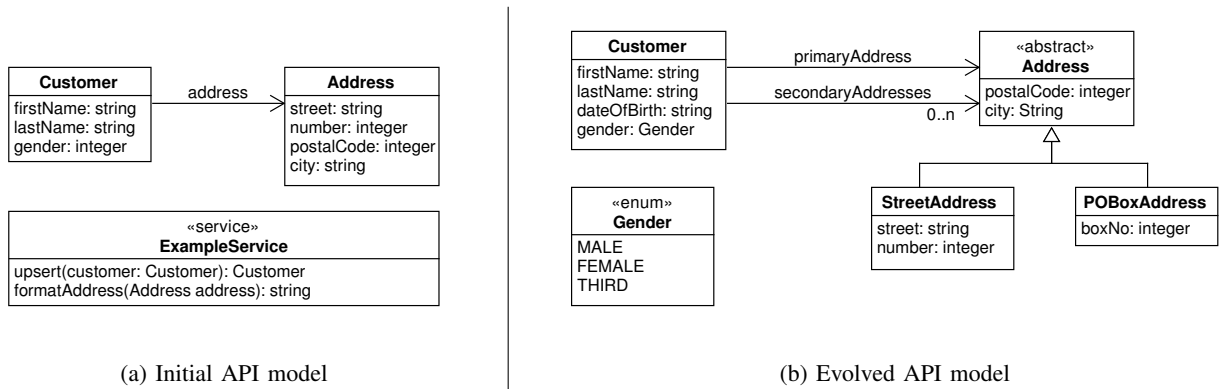(a) Initial API model   (b) Evolved API model

Fig. 1: UML model of the evolution example

JSON. And dynamic query facilities like XPath provide convenient options for creating tolerant readers. However, the field names make these data formats quite verbose. This has led to more optimized formats like Apache Avro, which is described below.

While tolerant readers allow for the safe addition of new fields by a writer, additional measures are required to support other changes, such as renaming a field. In such cases, the tolerant reader can be complemented by a *magnanimous writer*, i.e., a writer that provides more data than necessary. For instance, a renamed field can be written both under its new and its old name, and since the readers ignore unexpected fields, neither old nor new readers break.

Deletions of fields or type changes are more difficult. Mandatory fields cannot be deleted as long as there are still readers that expect them. Type changes would require to write two different values *under the same name*. As a consequence, these changes cannot be addressed by tolerant readers and magnanimous writers alone, but require more elaborate evolution patterns. A selection of common patterns is described in [7]. In the following paragraphs, we refer to these unless noted otherwise.

A common strategy to address potentially breaking changes like field deletions is to introduce *version identifiers*, often based on *semantic versioning*.[2] Semantic versioning expects that a version number is of the form *major.minor.patch*, where the major version number is incremented if and only if an incompatible (i.e., breaking) change is made.

To make use of versioning, each client must provide the identifier of the desired API version when initiating an interaction with a provider. Different approaches for accessing specific API versions are employed in practice, e.g., specific URLs or custom HTTP headers [8]. It is then the provider's responsibility to proceed according to the given version, or to refuse the interaction if it is unable to do so. Obviously, a provider must support at least two, sometimes even more (Pattern *Two in Production* or *N in Production*, respectively),

versions concurrently for a certain amount of time to avoid breaking clients.

Supporting multiple versions concurrently can be tedious. Therefore, providers try to keep the number of such versions low. A common pattern is to give a *Limited Lifetime Guarantee*, which means that already upon publication, the provider announces until when this particular version will be supported. A slightly different approach is *Aggressive Obsolescence*. Here, the provider declares an API as deprecated at a time of his own choosing (e.g., when a new version becomes available), and gives the clients a sufficient grace time to migrate to a supported version.

Such processes cannot be reliably managed without tool support. Registries keep track of published APIs and manage their life cycles [9]. Furthermore, such registries may provide additional functionality like usage monitoring for operations as well as monetarization.

B. Evolvability in Existing RPC Implementations

After establishing the basic patterns for API evolution, we now proceed to investigate how evolution is supported in two common RPC implementations. We chose Apache Thrift because its evolution capabilities are well documented, and Avro due to its dynamic schema matching, which was a major inspiration for our approach.

1) Apache Thrift: Originally created at Facebook, Apache Thrift[3] is a cross-platform RPC framework that supports numerous target languages. Thrift APIs are built from five major elements:

- *Base types*, such as binary integers or character strings,
- *Structures* composed of fields with types defined within the API,
- *Containers*, such as lists or sets,
- *Services*, which describe the operations provided by the API in terms of the defined types, and
- *Exception types*, which are effectively structures representing an error

---

[2]https://semver.org

[3]https://thrift.apache.org/

Thrift provides an own interface definition language (IDL) to specify APIs. From such an API specification, a code generator produces the required source code artifacts for serializing and deserializing data.

Thrift's evolution capabilities are discussed in [10]. They are based on field identifiers and type specifiers: At code generation time, each field receives a unique numeric id that is encoded along with its according type specifier and the actual value. Although field identifiers can be assigned automatically, it is considered good practice to specify them explicitly.

The use of field identifiers makes Thrift definitions resilient to name changes and field reordering. The type specifier ensures that the type of each encoded value is always known, even (and especially) for unexpected fields. This, together with the fact that all types in Thrift are encoded in a self-delimiting way, allows to skip unexpected fields without error.

With respect to our evolution example, Thrift supports the field addition (Change 1) and the renaming (Change 2) very well. This includes a non-obvious detail of the field addition due to the request-response communication. As the field is added to a type that is used as an input parameter, the provider must not expect this field to be supplied, as older clients are not aware of it. To account for this, Thrift fields are, by default, optional for input (opt-in), i.e., they are optional for the client, but mandatory for the provider.

The type change (Change 3) must be implemented by adding a new field with a different name. Additional enum members can be freely added; readers treat unknown values as if no value had been provided at all.

Thrift provides no immediate support for inheritance, but it can be emulated by means of unions. Unions in Thrift are special records with multiple fields, only one of which has a value. As the entire construct needs to be manually emulated, there is no particular support for Change 5.

2) Apache Avro: Apache Avro[4] is a data serialization system that was originally developed in the context of Apache Hadoop, a platform for large-volume batch data processing. Although primary a serialization system, Avro also specifies an RPC protocol.

Avro is particularly interesting for evolution as it employs dynamic schema resolution. Unlike Thrift (and many other such frameworks), Avro does not rely on generating code from a common IDL definition; it only requires that both reader and writer have compatible schemas. The Avro RPC protocol contains a handshake mechanism so that client and provider can agree on schemas before invoking a procedure.

Serializing and deserializing in Avro roughly works as follows. The writer serializes the input using its own schema, and supplies the schema definition (or a reference to it) along with the serialized output. The reader then resolves its own schema against the writer's schema, identifying potentially matching elements by name. Schema evolution is enabled by allowing certain deviations between the two schemas [11]. For instance, a reader may specify alias names for writer fields,

and certain types can be promoted (e.g., `int` can be promoted to `long`).

The result of the schema resolution is essentially a partial mapping from the writer's schema to the reader's schema. This mapping is then applied to transform the data, which is deserialized using the writer's schema, to match the reader's expectations.

This schema resolution has three important consequences. First, it removes the need to embed field identifiers in the actual payload data ("untagged data") and can therefore be applied to a wider variety of data formats. Second, it allows for provider and client schemas to be specified independently, as long as they are sufficiently compatible. Third, the explicit client schema enables the provider to deliver the data according to the needs of the client. This is crucial for contexts where rigid and predefined data formats must be used, such as COBOL programs. Therefore, our approach uses a very similar mechanism that only provides additional features to account for the static nature of COBOL.

It is apparent from the aforementioned examples that Avro's evolution capabilities primarily facilitate *backwards compatibility*, i.e., enable a newer reader to read data written by an older writer. To actually achieve backward or forward compatibility, certain rules must be actively followed by the developer, such as not adding or removing fields without default values [12].

As for our evolution example, Avro supports the field addition (Change 1) in a way similar to Thrift. However, the new field needs to be explicitly specified as optional, since Avro does not provide an equivalent concept to Thrift's opt-in. The name change (Change 2) is partially supported by aliases. But as these are only applied when reading, the (unchanged) client may not have an appropriate alias when reading the response. The type change (Change 3) also needs a new field in Avro; but Change 4 requires the new enum to specify an explicit default value that is used in case of an unexpected value. Similar to Thrift, there is no inheritance mechanism in Avro, but again, it can be emulated with union types.

Avro is extensively used in stream processing platforms such as Confluent or Apache Pulsar.[5] As such platforms distribute immutable data, they match Avro's reader-writer model very well. In order to keep readers and writers compatible, both Confluent and Pulsar provide a schema registry, which manages the life cycle of the schemas and is capable of detecting (and possibly rejecting) changes that break the desired level of compatibility.[6][7]

C. Summary

In conclusion, several of the changes of our example are already quite well supported by existing implementations. But especially more complex operations like type changes must be performed manually, which is potentially error-prone.

---

[4]http://avro.apache.org

[5]http://pulsar.apache.org/

[6]https://docs.confluent.io/platform/current/schema-registry/avro.html

[7]https://pulsar.apache.org/docs/en/schema-evolution-compatibility/

```
record Customer {
 string firstName
 string lastName
 string dateOfBirth
 Address primaryAddress replaces address
 Address* secondaryAddresses
 Gender gender as genderNew
}
```

Fig. 2: Example provider API definition with evolution

```
record Customer as Person {
 string firstName
 string lastName as familyName
 integer gender
}
```

Fig. 3: Example client API definition

Furthermore, when using semantic versioning, it is the developer's responsibility to manually identify breaking changes and adjust the version number accordingly. This is, however, far from trivial, especially in complex APIs. In practice, while semantic versioning is more and more adopted, it has been found that breaking changes are commonly introduced also in minor and patch revisions [13].

V. OUR APPROACH TO API EVOLUTION

A. Overview

The central idea of our approach is not to have schemas explicitly modeled by the developer (or API designer), but to derive these as well as the internal representations from a set of supported API definitions. These definitions always exist as part of a revision history. Thus, we always have access to adjacent revisions (provided that they exist), and are able to specify evolution steps from the previous revision directly in the API definition. This way, we can – at least to a large extent – automatically ensure that the required rules are followed, and do not need to burden the developer with them.

To illustrate the general approach, consider the example shown in Figure 2; the definition language is described in more detail later on. In the example, we see an excerpt from an API definition representing the first three changes to the Customer type in our running evolution example.

First, recall that in its initial state, the Customer type had four fields, namely firstName, lastName, address, and gender. The latter was modelled as an integer. The first two fields still exist with the same name and the same type, and are therefore considered to be the same fields. The dateOfBirth field is identified as a new field, since no matching field exists in the previous revision (Change 1). Similarly, the primaryAddress and secondaryAddress fields do not have a match; however, the replaces clause reveals primaryAddress to be the same field as address from the previous revision (Change 2). The gender field does exist in the previous revision, but with a different type (note that due to the same name, no explicit replaces clause is required). Therefore, it is considered a type change, resulting in an implicit deletion of the previous field and the creation of a new one, but with the same name (Change 3).

The type change requires special treatment by the provider code. As long as not all clients have updated to the respective revision, the provider must be able to handle the old field as well as the new field. Therefore, the provider-internal representation of the type, e.g., the corresponding Java class, must contain both fields, although only one of them will be provided by or returned to a given client. To avoid name clashes in the internal representation, the as clause allows us to provide an internal name of our choice. This internal name is, of course, not part of the external API. By always assigning new internal names to new fields, we can minimize the required changes to existing code. In the given example, the old gender field will remain almost unchanged in the internal representation. It will, however, become optional, as newer clients will no longer provide it.

On the client side, we also use an explicit API definition, albeit more simple. Client API definitions are not revisioned; instead, they always refer to a specific provider API revision. Therefore, they cannot contain replaces clauses. Mapping public to internal names, however, is also possible in client definitions. This way, we can prevent name changes from trickling into the client code. An excerpt of an example client API definition is shown in Figure 3.

The process of invoking a service can be summarized as follows. First, the client encodes the request data according to a schema derived from its local API definition. The definition itself, or at least a reference to it, is passed along with the data. The provider then matches the client definition against the referenced revision; a process similar to Avro's schema resolution. It then proceeds to decode the given data using the schema derived from the client's definition, and transforms it into the appropriate internal representation. After returning from the service operation, this process is done in reverse order. By using the client's definition to encode the result, we ensure that the response can be processed by the client.

Before looking into the details of the approach, we take a quick glance at Changes 4 and 5 of our running example. We support both adding new enum members and sibling types, but these changes cannot be handled fully transparently. Because both introduce new, valid values that cannot be represented in earlier API revisions, but at existing locations. For instance, a new member in an existing enumeration (Change 4) can occur at any existing field of this type. The addition of the new address type (Change 5) is similar: Although the provider still supports the old primaryAddress field, it may be unable to determine a valid value, since a customer may now choose a POBoxAddress as its primary address. Therefore, an explicit treatment of unrepresentable values is required.
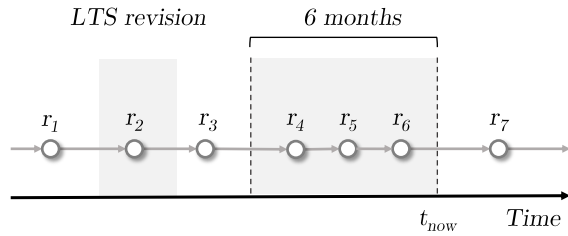
Fig. 4: Revision history with supported revisions

As discussed in the previous section, explicit API versioning is usually employed for breaking changes such as the removal of a field. We take a different approach: Each change (set) to an API definition leads to a new revision, whether breaking or not. The revisions are stored in an API registry. The provider then selects a set of supported revisions from the revision history, and our approach derives the necessary structures to support these revisions. As a consequence, removing a field in a new revision does not mean that the field is removed immediately, but represents the *intent* to remove this field. While clients aware of the new revision will never receive it, the provider is still required to support the field until the last revision containing it is removed from the supported revision set. To support this process, the registry also keeps track of client definitions. The supported set approach allows to combine the previously presented lifecycle patterns in various ways. Consider, for instance, the situation depicted in Figure 4. In this example, the provider supports all new revisions for six months (revisions r_4 to r_6), and additionally provides long-term support for revision r_2.

The central elements of this approach and their connections are summarized in Figure 5. Selected elements of the approach are described in detail in the following paragraphs.

B. The API Description Language

The first aspect of our approach that we will describe in more detail is the specification of API definitions. Similar to Thrift definitions and Avro protocols, an API definition consists of a number of service operations, together with a description of their input and output types.

An overview of the grammar of our API definition language, which we already used in the previous examples, is given in Figure 6. Note that for the sake of brevity, we do not give definitions for some non-terminals that are obvious or will become obvious by the examples below. An *API definition* consists of a number of user-defined types (i.e., record types and enumeration types) and services.

Enumeration types consist only of named values; unlike, for instance, Java enumerations, we do not support fields on members of an enumeration. *Record types* comprise named fields of arbitrary types, as described below. In order to facilitate modeling records with common fields, such as the addresses in the running example, a simple inheritance mechanism exists on record types. However, no overloading is supported. Record types can be declared `abstract` or given an optionality

modifier, which serves as the default for the fields declared within the type. The optionality modifier is inherited to the record's subtypes, but can be overridden by them.

Fields of a record can be of a user-defined type, a basic type (currently 32-bit binary integers, numeric strings and regular character strings), or an ordered list type. To account for languages like COBOL, all types are either bounded by construction (e.g., `int32`, records, or enums), or can be given explicit bounds, such as `string(5)` for a string of at most length 5 or `string(5)[10]` for a list of at most 10 strings of at most length 5.[8] An optionality modifier can be defined for each field, which can be either `optional` (the field is optional for both in- and output), `optin` (the field is optional for input, but mandatory for output) or `mandatory`. If no optionality is specified, the containing type's default optionality is used.

Services essentially serve as named containers for service operations. Each *service operation* has one record type as its input type, and one record type as its output type. It may furthermore throw exceptions, where an *exception type* is a record type that is restricted to be used only at this location.

Public and Internal Names: Every user-defined type, field, service, and service operation must have a *public name*, which is used for schema generation and matching. An addition, each such element can be assigned an *internal name* by means of an `as` clause, which is only used for the internal representation. While the public name must only be unique within the respective scope of the current API definition, the internal name must be unique in all supported revisions. If no explicit internal name is specified, the public name is used.

Replacement Clauses: All of the previously described elements can be used in both client and provider API definitions. An element exclusive to provider API definitions is the `replaces` clause. Such a clause may be added to any element with a public name, and denotes the element *in the previous revision* that should be considered as the predecessor of the given element. In most cases, this element is identified only by its name; in the case of fields, this name may be further qualified by a type name. If no explicit replacement is given, but a matching element of the same name exists in the previous revision, it is implicitly assumed to be the predecessor. This can be suppressed by the specifying that the given element `replaces nothing`. The semantics of replacement are discussed in detail in the following section.

C. API Evolution

The fundamental means for evolving APIs in our approach is to establish relations between the elements of consecutive API revisions using the replacement clauses of our API definition language. In total, this results in five relations, namely one on types, fields, enum members, services, and service operations. These relations define which elements are considered to be the same across revisions, and play a key role in transforming

---

[8]Note that the different types of brackets are required for syntactic unambiguity.
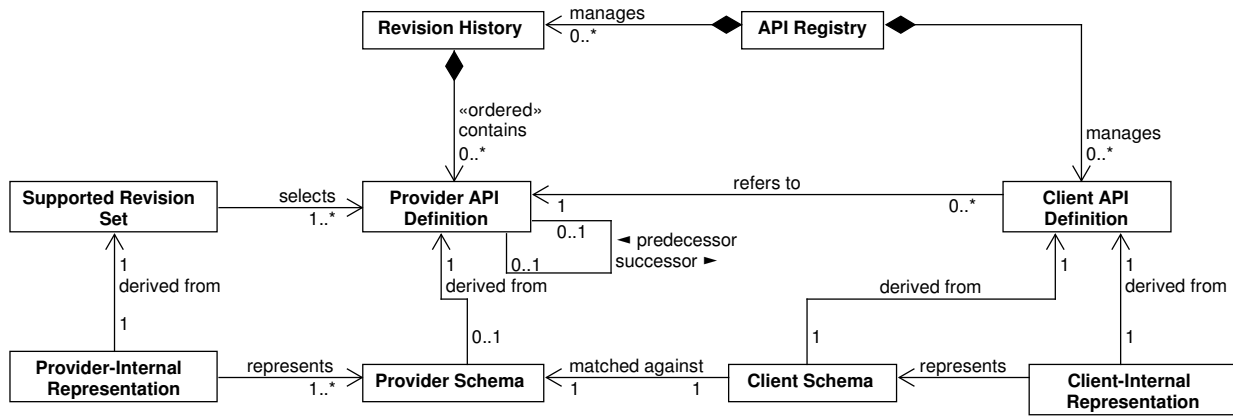
Fig. 5: Core elements of the approach

ApiDefinition: `api` *QualifiedName* {
  *(UserDefinedTypeOrService)**
}

UserDefinedTypeOrService: *EnumType* | *RecordType* | *Service*

RecordType: *(RecordModifier)** `record` *Identifier (*`extends`
Identifier)? (ReplacesClause)? (AsClause)? {
  *(Field)**
}

RecordModifier: `abstract` | **one of:** `optional`, `optin`,
`mandatory`

ReplacesClause: `replaces` *(Identifier* | `nothing`*)*

AsClause: `as` *Identifier*

Field: *(FieldModifier)? TypeRef Identifier (FieldReplacesClause)?
(AsClause)?*

FieldModifier: `optional` | `optin` | `mandatory`

FieldReplacesClause: `replaces` *(QualifiedFieldName (*`,`
*QualifiedFieldName)** | `nothing`*)*

QualifiedFieldName: *Identifier (*`.` *Identifier)?*

TypeRef: *BasicType* | *ListType* | *Identifier*

BasicType: `int32` | `numeric` `((` *IntegerLiteral* `))?` | `string` `((`
IntegerLiteral `))?`

ListType: *TypeRef* `(*` | `[` *IntegerLiteral* `])`

EnumType: `enum` *Identifier (ReplacesClause)? (AsClause)?* {
  *(EnumMember)**
}

EnumMember: *Identifier (ReplacesClause)?*

Service: `service` *Identifier (ReplacesClause)? (AsClause)?* {
  *(ServiceOperation)**
}

ServiceOperation: *RecordType Identifier* (*RecordType*)
(ReplacesClause)? (AsClause)? `throws` *Exception (*`,` *Exception)**

Exception: (`abstract`*)?* `exception` *Identifier (*`extends`
Identifier)? (ReplacesClause)? (AsClause)? {
  *(Field)**
}

Fig. 6: Grammar overview of the API definition language

representations. Unrelated elements from the "old" revision are considered deleted, and unrelated elements from the "new" revision are considered added. In order for the transformations to work, we require these relations to be injective partial functions. The reason for this requirement is given in the section on representation transformation below. Therefore, we must ensure that no two elements share the same predecessor.

Unless explicitly suppressed, we assume that two elements from consecutive revisions are related if they have the same name and are compatible (implicit replacement). To be compatible, the elements must be of the same kind (e.g., record type, enum type, or service). For fields, we furthermore require the field's type to be related, and for service operations, both the input and the output type must be related. Anonymous types, such as lists, are considered related if their referenced types are related. Furthermore, we require bounds on bounded types to match. If an explicit replacement is specified, the element with the given name is assumed instead, and the same rules for compatiblity apply.

To allow for type changes, replacements for fields and service operations may also refer to elements that are incompatible with respect to types. These elements, however, are

not related. Therefore, this results in the removal of the old and the addition of the new element, possibly under the same name. To further illustrate these rules, several examples are given in Table I.

The inheritance mechanism of record and exception types provides additional options for evolution. Fields may be pulled up or pushed down along the inheritance hierarchy. However, we deliberately do not support changing an existing supertype of a type. While simple on the schema level, this type of change makes the internal representation and thus the programming model significantly more complex.

Before we discuss these changes in detail, we need to take a quick look at the semantics of inheritance in our approach. As previously noted, our inheritance mechanism only provides basic subtyping, and facilitates the modeling of structures with common fields. Internally, each subtype receives an individual copy of each inherited field, and the supertype is replaced by a union type consisting of all possible subtypes. If the supertype is not abstract, it is added to the union as well. This interpretation greatly facilitates handling inheritance-related evolution, as shown below.

Recall from the running example that in the last change, a new sibling type to the street address was created, and that the common attributes were moved to a new supertype. This change is best modeled by declaring `StreetAddress` as the successor to `Address`, and then pulling up the fields to the new supertype `PostalAddress`. We deliberately allow to pull up fields that are not present in all subtypes (otherwise, we would first have to publish a revision with `POBoxAddress` having them), but this requires us to explicitly specify the type from which we are pulling. We also allow to pull up multiple fields into one, as long as they have the same type.

As opposed to pulling up, a field to be pushed down is uniquely determined by its name. But in order to clearly separate a push-down from a simple replacement, we have chosen to require the field to be qualified with the respective supertype's name. A field can be pushed down to multiple subtypes, for instance, if it should remain in two of three possible types. This is the only case in which it is legal to specify two successors to the same element.

Due to our interpretation of inheritance, these changes are not in conflict with our requirement of the predecessor relations being injective functions. Especially pushing down a field to multiple subtypes appears to be a problem, but as each of the types has an individual instance of this field, these types are, in fact, unchanged. The field is only removed from fields that no longer inherit it. Conversely, pulling a field up only causes some types to receive a new field. Several examples of such evolution steps are shown in Table II.

D. Schemas and Internal Representations

The API definition serves as the foundation of the schemas and internal representations for both clients and providers.

Schema derivation works exactly the same for clients and providers. The process is very straight-forward, and is completely oblivious of the revision history. We use a schema that is very similar to Avro's, the only (but – at least for COBOL – key) difference is the support for bounded lists and base types. Conversely, if no bounded types are required, our approach can be easily adapted to Avro schemas.

Deriving the internal representation is considerably more complex, especially for the provider. This is due to the fact that the provider's internal representation must be able to represent the entire supported revision set at the same time. The client internal representation is essentially a provider representation derived from a singleton revision set.

In order to be able to represent the entire supported revision set, the internal representation must essentially be a union of all revisions in the set. Depending on the capablties of the target language and the desired programming model, the precise algorithm differs slightly.

The rough outline is as follows. First, all types and services from the latest supported revision are added using their internal names, including their respective sub-elements. For each previous supported revision, all elements are added to the internal representation that do not have a successor in the next supported revision, i.e., have been deleted afterwards.[9] Elements that do have a successor are merged, in particular:

- If the newer revision of a type does not contain a field present in the earlier revision, the field is added
- If the optionality of a field differs between revisions, the more permissive one is used
- If the exceptions of a service method differ, the union of both methods is used

For target languages that support it, it may furthermore be desirable to actually use inheritance for the internal representation. In such cases, a type must additionally be concrete if at least one of its supported revisions is concrete.

Obviously, the resulting internal representation contains every type, field, service, and service operation present in any of the supported revisions. Note that all elements are based on their latest revision, e.g., will have their latest internal name. Furthermore, the optionality of fields is chosen permissive enough for all revisions.

E. Schema Matching and Representation Transformation

As described in the overview, we employ schema matching similar to Avro's to decouple clients and providers as much as possible. All matching and converting is done by the provider, so that the client can remain unchanged.

Schema elements are matched by name and must be of the same type; type promotion is not supported. Since each client API definition is only matched against the specific provider revision it refers to, no precautions against renaming are necessary. After decoding the request data using the client schema, the resulting mapping from the schema matching is used to associate the decoded values with the appropriate types and fields of the provider schema. These, in turn, allow to identify the corresponding elements in the internal representation (e.g., the field to store the value in). As the

---

[9]Recall that the supported revision set does not need to be contiguous.

Previous revision	Current Revision	Description
`record A {`	`record B replaces A {`	Record type A is renamed to B
`  string a`	`  string d replaces a`	Field a in A is renamed to d
`  int32 b`	`  numeric(5) b`	Type change for b from `int32` to `numeric(5)`
`}`	`  int32 c replaces b`	Error, multiple replacements for b
	`  string y replaces x`	Error, x does not exist in A
	`  string z`	New field z
	`}`	
`record X {`		Type X is deleted
`  string x`		
`}`		

TABLE I: Examples for valid and invalid evolution steps

Previous revision	Current Revision	Description
`abstract record A {`	`abstract record A {`	
`  string a`	`  string a2 replaces B.b, C.c`	Pull-up of b and c into A
`}`	`  string a3 replaces B.b2, C.c2`	Error, types do not match
`record B extends A {`	`}`	
`  string b`	`record B extends A {`	
`  string b2`	`  string b3 replaces A.a`	Push-down of a into B
`}`	`}`	
`record C extends A {`	`record C extends A {`	
`  string c`	`  string c3 replaces A.a`	Push-down of a into C
`  int32 c2`	`  String c`	Error, multiple successors for c
`}`	`}`	

TABLE II: Examples for valid and invalid evolution steps with inheritance

internal representation is always based on the latest revision of an element, but the match may involve an older revision (e.g., referring to an outdated internal name), this may require to navigate the successors of the respective element until the latest revision is found.

To encode the result from the internal representation, the process is performed the other way. In particular, the predecessor relation may be required to associate the internal revision with the appropriate elements of the schema. It is this bidirectional mapping that requires the predecessor relation to be an injective partial function. Encoding the data is again performed using the client schema to ensure that the data is returned exactly as the client expects them.

VI. EVALUATION

In order to assess the general feasibility of our approach, we conducted a short evaluation with respect to the goals stated in the introduction. To address Goals 1 and 2, we conducted a small survey with nine professional developers from our case study. We particularly selected participants that had several years of professional experience. The participants were given a presentation of the approach in small groups, and were asked to fill out a short questionnaire afterwards. Four of the respondents reported to use only Java in their everyday work, four both Java and COBOL, and one Java and JavaScript.

Regarding Goal 1, we asked the respondents to rate their agreement to three statements on a seven-point scale, ranging from "fully disagree" (encoded as -3) to "fully agree" (+3).

All participants agreed to the statement that the approach is comprehensible (fully agree: 4 / 5 / 0 / neutral: 0 / 0 / 0 / fully disagree: 0; median: 2). The statement that the approach addresses a relevant problem was also agreed to, but was apparently seen quite differently among the participants. While four participants fully agreed to this statement, three participants chose the neutral option (4 / 1 / 1 / 3 / 0 / 0 / 0; median: 2). The participants were furthermore asked whether they could envision applying the approach in their everyday work, and whether it supported the types of changes they needed. These statements were also generally agreed to (1 / 4 / 2 / 2 / 0 / 0 / 0; median: 2) and (2 / 1 / 3 / 2 / 1 / 0 / 0; median: 1), respectively.

Regarding Goal 2, we were particularly interested in identifying aspects of the approach that might negatively affect a developer's day-to-day work. Overall, handling the potentially large number of optional fields in the provider code raised the greatest concern among the participants, with 6 participants agreeing to some degree to the statement that this might impede their everyday work (0 / 1 / 4 / 3 / 0 / 1 / 0; median: 1). Creating and maintaining client schemas also led to some concern, but the respondents were mainly indifferent (0 / 0 / 1 / 4 / 1 / 2 / 0; median: 0). Maintaining the provider histories (0 / 1 / 0 / 3 / 2 / 2 / 1; median: -1) and explicitly publishing them to a repository (0 / 0 / 1 / 1 / 2 / 3 / 2; median: -2) were not considered as notable impediments.

The respondents were furthermore able to add remarks in an open text field. The most frequent remark was that a facility for highlighting semantic changes might be useful to make clients more aware of such changes.

We see the greatest threats to the validity of this survey in the size and structure of the sample, the sampling method, and potential biases. Due to the small sample size and the fact that all participants work for the same company, the results may not be generalizable. Several biases may result from the fact

that the participants were selected by the researchers, and that one of the researchers, besides his research, works for the same company. Therefore, the responses may be unduly favorable of the approach.

Despite these threats, we conclude that there is a strong indication that the approach is considered applicable in practice, although further research is necessary to support this indication. As for Goal 2, the survey revealed no show-stoppers, but suggests that future work should also focus on programming patterns and tooling to efficiently handle the additional complexity for the provider.

To address Goals 3 and 4, we created a proof-of-concept implementation of our approach. This implementation consisted of three components: One containing the API definition language and the schema resolution mechanism, and a runtime component for both Java and COBOL for performing the actual data conversions. Since published API definitions are immutable, schema resolution is a one-time operation. Therefore, the impact on runtime performance is determined by the conversion components.

In our experiments, we focused on two scenarios most relevant to our case study, namely remote invocations between Java applications and in-process calls between COBOL programs. For Java, we used reflection-based decoders and encoders to convert the serialized client data directly into the provider's internal representation and vice versa. A simple binary format was used for serialization. In our experiments, the performance of our converters was comparable to that of the JSON libraries Jackson and Yasson. For instance, converting a `Customer` object from client revision 1 to provider revision 6 and back took, on average, 1.67µs using our converters, while serializing and deserializing the object took 1.1µs with Jackson and 4.4µs with Yasson.

For our COBOL experiments, we created a converter program in C to serve as a proxy between two interacting COBOL programs and convert parameters and results as needed. The transformations themselves were defined using a simple byte-code, which can by generated dynamically, and interpreted by the C program. Although these conversions were considerably faster than the ones from our Java experiment, it must be noted that this is truly additional overhead, as COBOL parameters are usually just passed by reference.

Our proof-of-concept implementation suggests that our approach is indeed applicable to aged technologies like COBOL. However, it may have a considerable performance impact due to additional overhead. For remote communication, this is of less concern, since the conversions can be efficiently integrated into the necessary serializations and deserializations.

VII. Related Work

Evolving software as well as their data and artifacts in a compatible way is a major challenge, and a large body of research exists on various aspects of this topic. The reasons for and challenges of API evolution have been investigated in multiple studies. For instance, Dig and Johnson [14] report that the majority of breaking changes are caused inadvertently by refactorings, and conclude that tool support is required to ensure compatibility with clients. Research focused particularly on the evolution of public Web APIs reaches different conclusions. While Fokaefs et al. [15] find that mostly compatible changes such as field additions are made, other studies report that client application developers suffer considerably from non-standardized or even no procedures for API evolution [16] and poor documentation and communication of changes [17]. Wang et al. [18] observe that, in particular, the addition or removal of methods leads to many questions by client developers on Q&A sites like StackOverflow.

One of the central aspects of our approach, namely deriving schemas and representations from modifications of API definitions, is closely related to *co-evolution* or *coupled evolution*. A co-evolution approach that also uses traceability information from models to derive transformations on derived data is describe in [19]. The authors present an approach for specifying changes to an entity model, and use the traceability information to derive database migrations reflecting these changes. Wagelaar et al. [20] propose a virtual machine for executing transformation bytecode, which is a similar notion to the generic transformer for COBOL.

Compatible evolution has also been studied for binary artifacts. The primary goal of binary compatibility is to allow that a newer version of, say, a library can be dynamically linked with a program that was compiled against an older version of this library. For instance, a compiler may statically embed indexes of the virtual method table within the object code of a derived class. If a new virtual method is added to the superclass, these stored indexes may become invalid, therefore breaking binary compatibility [21].

Most modern programming languages provide information about the changes that do and do not preserve binary compatibility. The Java Language Specification [22], for instance, contains an entire chapter on this topic.

VIII. Conclusions and Future Work

In this paper, we have presented an approach for continuous API evolution for enterprise software systems in order to support the independent deployability of single applications. As shown by the proof-of-concept implementation, this approach is indeed viable and satisfies the goals described at the beginning of this paper. Although it would certainly be possible to add more powerful evolution mechanisms, we decided to aim for an approach that is both easy-to-use and easy-to-understand for a developer.

In our future work, we plan to extend this approach in several ways. First, we intend to address the findings and limitations of our evaluation. Besides more sophisticated implementations, this includes tooling for managing, querying, and editing revision histories, and a more in-depth exploration of the resulting programming model. Further opportunities for research include a more formal treatment of the API definition language and their underlying evolution semantics as well as the introduction of such concepts into existing API specifications like OpenAPI.

REFERENCES

[1] S. Newman, *Monolith to Microservices – Evolutionary Patterns to Transform Your Monolith*. Sebastopol, CA: O'Reilly, 2020.

[2] J. Humble and D. Farley, *Continuous Delivery*. Upper Saddle River, NJ: Addison-Wesley, 2011.

[3] H. Knoche and W. Hasselbring, "Drivers and Barriers for Microservice Adoption – A Survey among Professionals in Germany," *Enterprise Modelling and Information Systems Architectures (EMISAJ) – International Journal of Conceptual Modeling*, vol. 14, no. 1, pp. 1–35, 2019.

[4] ——, "Using Microservices for Legacy Software Modernization," *IEEE Software*, vol. 35, no. 3, pp. 44–49, May 2018.

[5] T. Stahl and M. Völter, *Model-Driven Software Development – Technology, Engineering, Management*. Chichester: Wiley, 2006.

[6] R. Daigneau, *Service Design Patterns – Fundamental Design Solutions for SOAP/WSDL and RESTful Web Services*. Upper Saddle River, NJ: Addison-Wesley, 2012.

[7] D. Lübke, O. Zimmermann, C. Pautasso, U. Zdun, and M. Stocker, "Interface evolution patterns: Balancing compatibility and extensibility across service life cycles," in *Proceedings of the 24th European Conference on Pattern Languages of Programs*. New York, NY, USA: Association for Computing Machinery, 2019.

[8] P. Sturgeon, *Build APIs You Won't Hate*. Leanpub, 2016.

[9] M. Medjaoui, E. Wilde, R. Mitra, and M. Amundsen, *Continuous API Management*. Sebastopol, CA: O'Reilly, 2019.

[10] M. Slee, A. Agarwal, and M. Kwiatkowski, "Thrift: Scalable Cross-Language Services Implementation," Facebook, Tech. Rep., 2007.

[11] R. Blue and D. Cutting. (2020) Apache Avro 1.10.1 Specification. [Online]. Available: https://avro.apache.org/docs/1.10.1/spec.html

[12] M. Kleppmann, *Designing Data-Intensive Applications*. Sebastopol, CA: O'Reilly, 2017.

[13] S. Raemaekers, A. van Deursen, and J. Visser, "Semantic Versioning versus Breaking Changes: A Study of the Maven Repository," in *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*, 2014.

[14] D. Dig and R. Johnson, "How do APIs evolve? A Story of Refactoring," *Journal of Software Maintenance and Evolution: Research and Practice*, vol. 18, no. 2, 2006.

[15] M. Fokaefs, R. Mikhaiel, N. Tsantalis, E. Stroulia, and A. Lau, "An Empirical Study on Web Service Evolution," in *2011 IEEE International Conference on Web Services*, 2011.

[16] T. Espinha, A. Zaidman, and H. Gross, "Web API growing pains: Stories from client developers and their code," in *2014 Software Evolution Week - IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE)*, 2014.

[17] S. M. Sohan, C. Anslow, and F. Maurer, "A Case Study of Web API Evolution," in *2015 IEEE World Congress on Services*, 2015.

[18] S. Wang, I. Keivanloo, and Y. Zou, "How Do Developers React to RESTful API Evolution?" in *Service-Oriented Computing*. Springer, 2014.

[19] S. Vermolen and E. Visser, "Heterogeneous Coupled Evolution of Software Languages," in *Model Driven Engineering Languages and Systems"*. Springer, 2008.

[20] D. Wagelaar, L. Iovino, D. Di Ruscio, and A. Pierantonio, "Translational Semantics of a Co-evolution Specific Language with the EMF Transformation Virtual Machine," in *Theory and Practice of Model Transformations*. Springer, 2012.

[21] I. R. Forman, M. H. Conner, S. H. Danforth, and L. K. Raper, "Release-to-Release Binary Compatibility in SOM," in *Proceedings of the Tenth Annual Conference on Object-Oriented Programming Systems, Languages, and Applications*, 1995.

[22] J. Gosling, B. Joy, G. Steele, G. Bracha, A. Buckley, D. Smith, and G. Bierman. (2020) The Java Language Specification, Java SE 15 Edition. [Online]. Available: https://docs.oracle.com/javase/specs/jls/se15/jls15.pdf