

Industry Case Study: Design Antipatterns in Actual Implementations

Understanding and Correcting Common Integration Design Oversights

Mihaela Iridon

Cândeia LLC

Dallas, TX, USA

e-mail: iridon.mihaela@gmail.com

Abstract— The design of any extensible integration solution involving systems intended to communicate efficiently with one another and/or with data repositories usually begins as a proof of concept or prototype, especially if new technologies and platforms are involved. In some instances, the focus on functional features and tight deadlines lead to inadequate attention placed on non-functional system attributes, such as scalability, extensibility, performance, etc. Many design guidelines, best practices, and principles have been established, and antipatterns were identified and explained at length. Yet, it is not uncommon to encounter actual implementations suffering from deficiencies prescribed by these antipatterns. This paper discusses Leaky Abstractions Mixing Concerns, and Vendor Lock-in, as some of the more frequent offenders in case of system integration. Ensuing problems such as the lack of proper structural and behavioral abstractions are described, along with solutions aiming to avoid costly consequences due to integration instability, constrained system evolution, and poor testability. Moreover, unsuitable technology and tooling choices for database design and release management are shown to lead to a systemic incoherence of the data layer models and artifacts, and implicitly to painful database management and deployment strategies.

Keywords- *integration models; design antipatterns; leaky abstractions; database management.*

I. INTRODUCTION

Translating business needs into technical design artifacts and choosing the right technologies and tools, demands a thorough understanding of the business domain as well as solid technical skills. Proper analysis, design, and modeling of functional and non-functional system requirements is only the first step. A deep understanding of design principles and patterns, experience with a variety of technologies, and excellent skills in quick prototyping are vital. Although conceptual or high-level design is in principle technology-agnostic, ultimately specific frameworks, tools, Application Programming Interfaces (APIs), and platforms must be chosen. Together they enable the translation of the design artifacts into a well-functioning, efficient, extensible, and maintainable software system [1].

Designing a solution that targets multi-system integration increases the difficulty and complexity of the design and prototyping tasks considerably, bringing additional concerns into focus. Identifying integration boundaries and how data

and behavior should flow between different components and sub-systems, maintaining stable yet extensible integration boundaries, and ensuring system testability, are just a few of such concerns. This paper intends to outline a few design challenges that are not always properly addressed during the early stages of a project. and which can quickly lead to brittle integration implementations and substantial technical debt.

A few recognized design antipatterns and variations thereof are explained here, including concrete examples from actual integration implementations as encountered on various industry projects. Solutions to refactor and resolve these design deficiencies and issues are recommended as well.

Section II will address architectural and integration modeling concerns. The structural aspects discussed in this section range from low granularity models (i.e., data types which support the exchange of data between systems) to large-grained architectural models (i.e., system layers and components). The consequences of designing improper layers and levels of abstractions are outlined, followed by recommendations on how to avoid such pitfalls by refactoring the design accordingly.

In Section III, some antipatterns covered in Section II are extended to the design of the data models and relational databases, discussing also the ability to customize external open-sourced systems that participate in the integration. The focus is later shifted to the management and delivery of the data layer components and artifacts, as databases are an integration concern that goes beyond the data exchanged between the application tier and the data tier. This section intends to explain how the choice of tools and frameworks can have a significant impact on the overall realization, management, and delivery of the integration solution.

Finally, Section IV summarizes the integration design concerns and issues and the recommendations presented in this paper.

II. TIGHT INTEGRATION: LEAKY ABSTRACTIONS AND VENDOR LOCK-IN

A. The Problem Definition

Let's assume the specification of some business needs for building a software system to integrate with - and consume - a third-party service. The exposed data transport models, e.g., Representational State Transfer (REST) models or Simple Object Access Protocol (SOAP) data contracts, are already defined, maintained, and versioned by some external

vendor or entity (the service provider). Note that this scenario can easily be extended further, to integrations with an arbitrary system by means of some third-party APIs that expose specific behavior and data structures as containers for some meaningful payload/data.

Focusing on the data structures rather than behavior, once service model proxies have been generated via some automation, they tend to become part of the design artifacts for the rest of the system. Their use extends beyond the point where they are needed to exchange data with the external application. These models will percolate throughout the various layers and components of the integrating system. It is not unusual to see development efforts proceed around them, with application and business logic rapidly building on top of these data types. Development costs and tight deadlines, and sometimes the lack of design time and/or expertise, are the main reasons leading to this undesirable outcome.

Models exposed by external vendors were not designed with the actual needs of other/integrating systems in mind. External models are characterized by potentially complex shapes (width: number of exposed attributes or properties; depth: composition hierarchy). They cater to most integration needs (“one size fits all”), so they tend to be composed of an exhaustive set of elements to be utilized as needed.

Moreover, allowing these structural characteristics to seep into the application logic layer, beyond the component that constitutes the integration boundary, introduces adverse and unnecessary dependencies to external concerns. Therefore, the system is now exposed to structural instability and will require a constant need to adapt whenever these externally derived models will change. The integration boundary is no longer a crisp and well-defined layer that can isolate and absorb all changes to the external systems – speaking from a data integration perspective.

B. The Antipatterns

The lack of proper structural abstractions and allowing integration concerns to infiltrate into the integrating system is a costly design pitfall and is in fact a variation of the “Leaky Abstractions” problem – as originally defined by Joel Spolsky in 2002 [2]. Such deficient abstractions can be identified not only relative to structural models, but also to behavioral models, which could expose the underlying functional details of the software components to integrate with. This will inevitably lead to increased complexity of the current system, jeopardizing its extensibility and its ability to evolve and to be tested independently. Ultimately this results in a tightly coupled integration between the two systems (with strong dependencies on the target of the integration).

Another perspective or consequence of the problem described is an imposing reliance on vendor-specific technologies, their libraries, and even implementations. This problem is also known as the “Vendor Lock-In” antipattern [3]. External system upgrades will necessitate system-wide changes and constant adjustments on the integration side and will impact the overall stability of the system and the integration solution itself.

Examples range from adopting specialized libraries catering to cross-cutting concerns (logging, caching, etc.) to

domain-specific technologies (telecom, finance, insurance, etc.). Vendors will encourage integrators to infuse their specialized technology everywhere, leading to entire (sub)systems taking on pervasive dependencies on their technologies, thus making it impossible to separate. Such vendor-dependent architectures must and can be avoided with added effort during the design phase, as described next.

C. The Solution

To avoid such scenarios, the design must unambiguously identify the integration boundary and define custom integration models that abstract away any and all structural and behavioral details related to the system targeted for integration. This architectural approach is exemplified in the component diagram in Figure 1. The integration layer should also hide the underlying technology (REST vs. SOAP, message bus vs sockets, etc.) to avoid tight and unnecessary dependencies. An example of defining canonical models based on the “ubiquitous” integration language in case of multi-system integration is presented in “Enterprise Integration Modeling” [4].

Based on the author’s experience, designing proper model abstractions proved extremely useful in the case of building custom integrations with real-time systems. For example, Session Initiation Protocol (SIP) soft switches used in telecommunications networks, such as those from Genesys, the leader in customer experience, pertaining to contact center technology (call routing and handling, predictive dialing, multimedia interactions, etc.). In this case, an extensive array of data types, requests, events, etc., are made available to integrators as part of the Genesys Platform SDKs [5]. These facilitate communication with the Genesys application suite – which in turn enables integration with telephony systems, switches, IVR systems, etc. Most of these data types are very complex and heavy, and introduce acute dependencies on the underlying platform, exposing many implementation details as well. Employing code generation

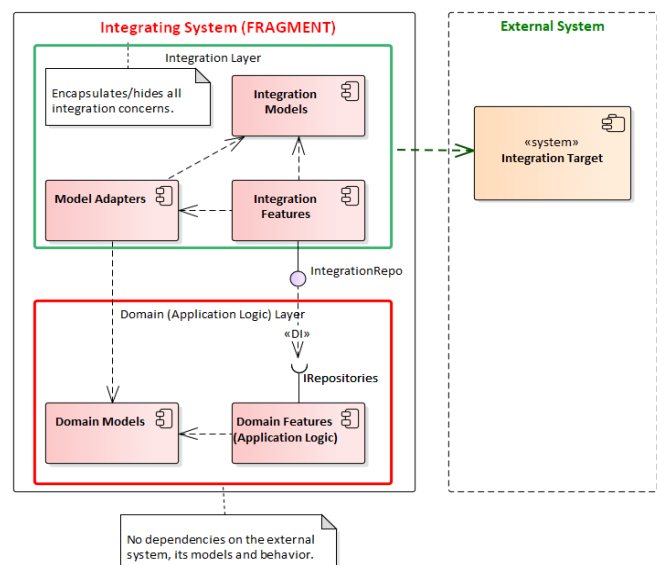


Figure 1. Integration components and the Integration Layer (Adapter) isolating the integrating system from the external system.

and metadata inspection via reflection, for example, simpler connection-less models were designed to mimic and expose only the needed structural details and are currently used in several production systems. Furthermore, defining and realizing the proper architectural isolation layers will ultimately provide independence from vendor-specific platforms for the rest of the system.

For example, considering the integration scenario mentioned above using Genesys' Platform SDK, recently the company (Genesys) has been pushing for a new approach to integrate with their systems, specifically using the Genesys Web Services (GWS) [6], a RESTful API. From an integration viewpoint, this substitution is practically equivalent to switching to a different vendor, as the two integration facilities are based on different technologies (web calls versus direct socket connections) and using completely different models, from both a structural model perspective as well as behavioral and consumption views.

Building an explicit and clean integration layer as shown in Figure 1, when dealing with such a significant change (vendor or technology replacement) changes will be isolated to this adapter layer without any impact on the business domain layer of the integrating system (assuming similar data and functionality). This includes the specifics of the technology used to communicate between the two systems.

Finally, it is noteworthy that four out of the five SOLID design principles [7] substantiate and drive towards the proposed solution:

- Single Responsibility (SRP), from the component and layering perspective,
- Open-Closed, to avoid changing the underlying implementation every time the integration endpoints change,
- Interface Segregation, exposing only the necessary data types for consumption by the business logic layer,
- Dependency Inversion, where the Domain does not directly depend on the external system, its data and behavior, but rather on abstractions – the repository contracts realized by the integration layer.

D. Added Architectural Benefit

Proper design and isolation of the integration components and the use of interfaces and model adapters will enable adequate testing of the custom system without demanding the availability of the external system for integration testing until most defects within the custom system are resolved.

Furthermore, this design approach supports building synthetics that simulate or mock the data and behavior of the external system, providing the means to prototype and test the integration points and functional use cases. Even if only a reduced set of features is synthesized, deferring the needs for actual integration testing can be cost-effective, especially in situations where the external system is a shared resource, perhaps expensive to manage and to access in general.

Employing Dependency Injection (DI) [8], either the real or the mock implementation of the integration contracts can be injected into the Domain layer, making it easy to swap between the two implementations.

III. DATA TIER DESIGN, ACCESS, AND MANAGEMENT CONCERNS

One of the most common system integration use cases for many enterprise applications is related to data persistence and access. Integration with (relational) databases that are either part of the custom system or accessible (co-located) components of a third-party system is a pervasive requirement, whether the data tier is needed for storing configuration data, audit/logging, security-related aspects, or to support concrete operational or reporting needs.

This section focuses on several issues related to database design and management, as well as accessing the data itself.

A. 'Inverted' Leaky Abstractions in Data Integrations

1) The Problem

The previous section discussed Leaky Abstractions that result from allowing third-party concerns infiltrate custom systems when designing and implementing an integration solution. The directionality of the "leak", as described earlier, is from the external system into the current one. However, it is also possible to encounter the reverse scenario, when the integration target is open or transparent to the integrators who then take advantage of this fact to develop and apply their own customizations.

Here are two examples:

(a) An Original Equipment Manufacturer (OEM) and/or White Label license of the external system is available to integrators, including access to source code for additional customization and integration options.

(b) The external system contains database(s) accessible to the integrators, on-premise or in a cloud environment, and is open/accessible to change.

In the first example, the same issues and solutions apply, as already discussed in the previous section, only this time from the perspective of the external system. If customization design is not executed properly, software upgrades of the open-sourced third-party system will result in continuous maintenance, or worse, breaking the custom code. Both scenarios will incur high development and system integration testing costs, among other problems.

The rest of this sub-section will focus on the second example, involving third-party databases that are accessible (i.e., open to modification) from an integration and customization perspective.

When expecting and relying on continuous upgrades and patches supplied by the vendor of the external system, it is possible that custom database artifacts (added by the integration provider) will have to be discarded and reapplied, or worse, no longer compatible with the updated system. Moreover, management of database source code targeting the customizations is more difficult if tightly dependent on the elements defined by the external entity/vendor. For example, the custom integration requirements demand two new columns on one of the third-party database tables.

Evidently, with respect to customizations of third-party components (database or otherwise), "Vendor Lock-in" is the status quo as a business-driven need and not a concern here.

2) The Solution

There are several options available and their applicability depends on concrete scenarios and business needs. Ideally, a separate, custom database could be considered, where data collected by the third party system (stored in their databases) would be extracted, transformed as needed, and loaded (ETL) [9]. Detached custom data models are easy to maintain, modify, and version-control by the integration provider. Aligning with the arguments stated in Section II, this approach enforces a well-defined data integration boundary, as shown in Figure 2 below.

Allowing for independent provisioning and evolution of both data models (one provided by the external system and one specifically designed for - and consumed by- the integrating system) will lead to improved extensibility, scalability, performance, testability, and maintainability. With this approach, upgrading the external system will potentially require updating the ETL artifacts and, if needed, some enhancements to the custom database – but both activities can be done in a detached, self-contained fashion.

Further details regarding the management of database artifacts will be discussed later, but one noteworthy benefit here is the freedom from having to maintain (a) partial custom database artifacts (divorced from their context) and/or (b) complete external database artifacts (since the database is a self-contained software system, and should not be divided further into sub-components). The reason why maintaining select/partial database artifacts is undesirable is that from a specification perspective, a database (meaning all its defining artifacts) must be valid, consistent, and complete (as it must also be from a deployment perspective).

If database customizations *must* live in the same database as the one that is part of the external system, a less optimal solution to the Inverted Leaky Abstractions (i.e., the data model), is to expend proper design effort to minimize tight dependencies and attempt to follow - as best as possible - the Open-Closed design principle at the data tier, in the context of system integration and customization.

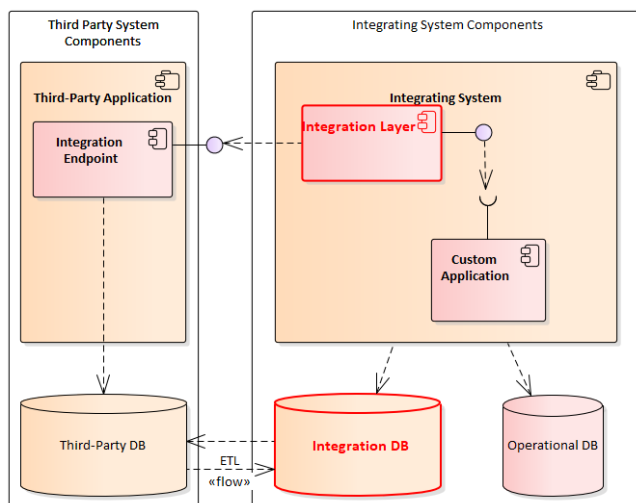


Figure 2. The integration database added to support data integration customizations and to remove direct dependencies on the third-party database.

For example, if the custom integration components require the persistence of new attributes (fields) in addition to the data captured by the external system, rather than modifying the existing third-party tables by adding new columns, association or edge tables should be considered instead, with custom data residing in new, custom tables. Custom views, parameterized or otherwise, should be designed to transform data into a ready-to-consume format (for operational, reporting, or analytical needs).

In this case, the system quality attributes mentioned earlier must however be carefully monitored, especially query performance and scalability. On the downside, database code management will become either (a) fragmented/isolated, by extracting the custom database artifacts from the rest of the database into independent scripts, or (b) more complex, by importing the entire third-party database under source control along with the custom artifacts, in order to preserve its integrity. Subsection D discusses tools that help validate the full database, warning about invalid or broken object references, binding and syntax errors, thus increasing the probability that database deployments will succeed.

B. Mixing Data Modeling Concerns

1) The Problem

Regardless of the targeted Database Management Systems (DBMS) technology, designing the conceptual and logical data models is a prerequisite to the implementation of the physical data models [10]. Beside ensuring that all data elements outlined by the business requirements are accurately represented, non-functional requirements, such as performance, scalability, multi-tenancy support, security (access to data), etc., will also shape the data architecture.

From an application perspective, the database is used to persist the state of the business processes supported by the application, i.e., operational needs, and to support analysis and reporting needs around the stored business data. The concept of Separation of Concerns (SoC) applies here as well but is often ignored. Operational versus reporting concerns are often mixed and data models designed specifically for operational needs are used as such for reporting or analytics purposes, although these models are usually quite different, in terms of how the data is stored and how it is accessed. Yet, it is not uncommon to find a given database used both as the operational as well as the reporting database. As a direct consequence of violating SoC with respect to data modeling (both logically and physically), stability, scalability, extensibility, and performance are the main quality attributes of the system that will be impacted.

An alternate description of this problem is known as the “One Bunch of Everything” antipattern [11], qualifying it as a performance antipattern in database-driven applications, the author aptly pointing out that “treating different types of data and queries differently can significantly improve application performance and scalability.”

2) The Solution

Following general data architecture guidelines, the solution is straightforward. In [12], Martin Fowler suggest the separation of operational and reporting databases and

outlines the benefits of having domain logic access the operational database while also massaging (pre-computing) data in preparation for reporting needs. Extract-Transform-Load (ETL) pipelines/workflows can and should be created to move operational data into the reporting database; specifically, into custom-tailored models that cater to requirements around reporting and efficient data reads.

Existing tooling and frameworks can be employed to transform and move data efficiently, on premise or in the cloud (Azure Data Factory, Amazon AWS Glue, Matillion ETL, etc.), for data mining and analytics, for historical as well as real-time reporting needs.

C. Data Access and Leaky Abstractions

1) The Problem

It has been noted [13] that Object Relational Mapping (ORM) technologies, such as Entity Framework (EF) or Hibernate, are in fact a significant cause of data architecture bleed into the application logic, representing yet another example of the Leaky Abstractions antipattern.

Although intended to ease the access to the data tier and the data it hosts, such technologies expose underlying models and behavior to the application tier. In more acute cases – depending on its usage – it also introduces strong dependencies from the domain logic to the data shapes defined in tables, views, and table-valued functions. Entity Framework, for example, while providing the ability to create custom mappings between these data models and the entity models, as designed, these object models are intended to be used as the main domain entities to build the actual domain logic around them. This forces a strong, intertwined yet inadequate dependency between two very different models, targeting different technologies, employed by very different programming paradigms (OO/functional such as C#.NET versus set-based such as SQL). This not only restricts the shape of the domain models, forcing constrained behavioral models to be implemented around them, but also causes data architecture changes to affect the domain and the application logic itself.

Not surprising, Microsoft’s EF Core framework in fact discourages against using a repository layer [14] (as prescribed by Evans’s DDD [15]) on account that EF itself implements the repository pattern/unit of work enterprise pattern [16] – alas, leading towards a rigid and potentially brittle integration. The reason is that ORM technologies push design and development towards data access logic tangled with the domain logic by encouraging multi-purpose models (domain and data access or data proxies).

2) The Solution

Just as with the integration solution presented in Figure 1, the impact of changes to database models should be constrained to one or two components – those that make up the data access layer, and prohibited from affecting the other application layers, specifically the domain and service layers. Sharing a single model across all layers of the application places unnecessary limitations on the overall design and ultimately on the extensibility and stability of the system.

Although it is uncommon to replace the database technology altogether, sometimes it may be required to

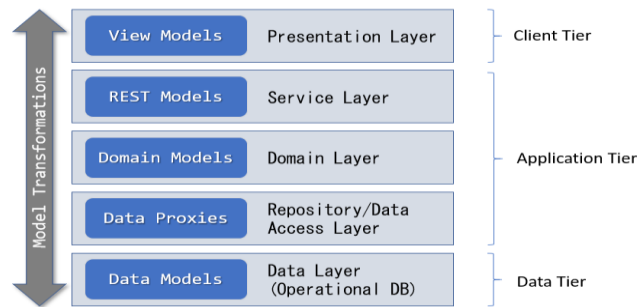


Figure 3. Layered architecture with layer-specific models and model transformations.

replace the *data access* technology due to performance and scalability concerns. Without a proper separation of data access from domain logic and models, such design changes targeting the lower layers of the system architecture are impractical without extensive refactoring of the application.

In a layered component-based architecture – as shown in Figure 3 above, it is easy and natural to allow each layer to define its own models (darker boxes) and provide adapters to translate from one model to another as data flows through the layers of the application. Although this would seem wasteful at first sight, especially if some models hardly vary from one layer to the next, this approach offers two core benefits. It allows for independent evolution of the models, customizing them to serve very specific needs of the layer they belong to, and keeps the propagation of model changes confined to the corresponding adapter (translation) components.

In case of ORM technologies, the data access layer overlaps with the domain layer, while entity models (shown as data proxies in Figure 3) represent the actual domain models. Interestingly enough, even as ORM is recognized as a Leaky Abstraction, its use is nevertheless encouraged [17], most likely because in unsophisticated implementations, it may be able to deliver acceptable results.

However, as [13] points out, ORM tools can be successfully used “if there is proper separation of concerns, proper data access layer, and competent developers who know what they are doing and really, really understand how relational databases work.” Sooner or later, the inherent deficiencies of such technologies, compounded by inadequate implementations due to the lack of understanding of how the underlying technology works, will surface, in most cases under system load and/or when new features are added.

D. Improper Management of Database Artifacts

1) The Problem

Source code, regardless of the language it is written in, is “a precious asset whose value must be protected”, as Atlassian’s Bitbucket web site states on their “What is version control” online tutorial [18]. All software-producing companies will employ one tool or another for version control. This allows software developers to collaborate, store (or restore/rollback) versions of the software components they build and perform code reviews, and providing a single

stable “source of truth” of the software artifacts they create and release/deploy. As advocated in [19], “source files that make up the software system aren’t on shared servers, hidden in folders on a laptop, or embedded in a non-versioned database.” Yet, it is rather commonplace to find database implementations that are improperly managed, leading to frustration, bad deployments, making the data tier integration and overall solution delivery unreliable and difficult. There are many online articles and blogs describing such cases.

As encountered by the author, while being engaged as a solution architect and consultant on several projects at various clients, the actual data models and database artifacts were often created and delivered as *ad-hoc implementations* in some arbitrary database, hosted under some arbitrary Microsoft SQL Server instance. Several teams needed these database artifacts: Development for implementation and integration, Quality Assurance for testing, DevOps for deployment. The most common process for deploying this database (fresh install or incremental) to some other environment was to generate and pass around SQL scripts when needed. In somewhat more fortunate situations, these scripts were maintained in some form of source control as SQL/text files, but lacking the ability to validate them or trace the source back to the developer responsible for the actual implementation (in the original database).

So then, where does the “source of truth” for the database definition reside? How can multiple developers work on the database code without overwriting each other’s changes and without being aware of the latest updates? How does the organization deliver incremental deployments to any number of target environments? When onboarding new team members, what database code should they be pointed to?

The problems derived from not having a stable, accurate, up-to-date, and complete definition of the database source code, one that is under version control and that can be validated before a deployment, are numerous, acute, and rather obvious. Just as one maintains all other application code under source control, entire solutions composed of many components, why should database implementations not follow the same standards and take advantage of the same acclaimed benefits of code well-managed?

Furthermore, when the database (source) code resides in some database, invalid object references (because someone dropped a column on a table or deleted a stored procedure) will surface only at runtime. Often, changes are made to the database post deployment, even in Production environments, changes that could potentially break the code, or which are at best confined to that environment alone, but without being retrofitted/updated back into the “source code database”.

A particularly curious approach to database code management and deployment was encountered on a project that used the Fluent Migrations Framework for .NET [20], self-proclaimed as a “structured way to alter your database schema [...] and an alternative to creating lots of sql scripts that have to be run manually by every developer involved.” In a nutshell, the tool calls for creating a C#.NET class every time the database schema would change (one class per “migration”). These code files (admittedly, version-controlled) attributed with metadata to identify a specific

database update, encapsulate two operations that describe the schema changes: one for a forward deployment (“Up”) and one for rollback (“Down”). With a large database, one that evolved considerably over time, with hundreds of artifacts, the number of C# migration files was astounding (thousands). Database changes were published to the target database as part of the application deployment process. Installing the database from scratch would incrementally apply every single “Up” migration specification it finds in these files, following the prescribed update. To maintain sanity, these source code files needed to be named such that the chronological order would be preserved when browsing in the development tool.

However, other more serious problems arise from using this framework, two of them being briefly discussed next.

a) SQL code as C#.NET strings??

Say a new stored procedure must be added; the code is developed and tested from SQL Server Management Studio (SSMS) in some local deployment of the database (assuming the objects the stored procedure is referencing do not change in the interim). Next, a migration file is created, with the “Up” method containing the full (CREATE) stored procedure script, as a C# string passed as input argument to the “Execute.Sql” method call.

The major and obvious problem here is the inability to validate SQL syntax and semantics and SQL object references when represented as indiscriminate plain strings, subject to typing errors.

b) No database source code??

Unless deployed on some SQL Server instance, it is impossible to even begin to understand the structure of the database, even the structure of individual objects. The data models and data logic are scattered, fragmented (across many C# files), impossible to validate (syntactically or otherwise) from where the database “source code” is stored.

Moreover, a given database object, say a table for example, can change any number of times, each change being captured in a different source file, with no unified, single view of what that table looks like, what the shape of the data is, with all its columns and corresponding types, with its keys and indexes, constraints and triggers, if any. This problem extends to all database objects, not just tables.

The data models (the source code artifacts) are practically non-existent, difficult to comprehend, and cannot be validated until they are deployed. The result is a total and indefensible representational incoherence afflicting the most important component of a data-dependent enterprise system.

2) The Solution

There are various software tools available to address this problem. Both Microsoft and Redgate, for example, provide excellent tooling for developing relational databases, managing database artifacts under source control, facilitating change management and incremental deployment, generating manual update scripts (when automated deployment is constrained), and more.

Microsoft’s SQL Server Data Tools (SSDT) [21] is a development tool, available since 2013, using the Data-Tier Application Framework (DacFx). It facilitates the design and

implementation of SQL Server and Azure SQL databases, as well as database source control and incremental deployment, all integrated under the Microsoft Visual Studio development environment.

A version-controlled database project contains all distinct database objects as individual files, and it *must* compile – targeting a specific SQL Server (or Azure) database version – before it can be deployed anywhere. Developers can check out individual objects (files) to change as needed or can add new objects using the provided templates. Just as one can see the entire schema of a database in SSMS, similarly, anyone can see and browse these objects in Visual Studio.

Tools like SSDT are also capable of identifying the changes (delta) between the source and the destination database in order to create the appropriate deployment scripts, and ultimately allowing rapid and valid delivery of database changes to any environment.

It is questionable to store Java or C# code in SQL scripts, with artifacts/classes shredded and reduced to SQL NVARCHARS, scattered in an arbitrary number of stored procedures (equal to the number of updates effected upon that class), and passed around to call other stored procedures (via EXEC statements). The reverse scenarios should be equally unacceptable. Treating the database as a proper software implementation artifact is imperative.

IV. CONCLUSION

This paper aimed to raise awareness about certain design challenges that, when not addressed early and properly, will lead to deficient architectures and rigid solutions concerning various aspects of system integration, as often encountered in practice.

When the design of software systems follows some basic guidelines and principles (SOLID), the resulting architecture will allow the system to be easily built, modified, and extended. In case of system integrations and customizations, violating these principles and particularly the multi-faceted Separation of Concerns design rule, leads to unmanageable and highly complex systems that do not scale well, cannot be extended or modified easily, with tight dependencies on external components and overall brittle integration solutions.

Many design antipatterns have been catalogued and well documented; yet deficient architectures are encountered quite frequently, leading to high technical debt and unhappy stakeholders. This paper discussed “Leaky Abstractions”, “Mixing Concerns”, and “Vendor Lock-in” antipatterns – from the perspective of concrete industry examples, as encountered and worked on by the author.

Concrete approaches that address these problems to help refactor and realign the design according to best practices and principles were elaborated, explaining how they lead to scalable, extensible, testable, efficient and robust integration solutions.

Relational database design and management concerns were also presented, with focus on data model design, data access practices, and management of database artifacts. The consequences of improper tooling and frameworks were briefly covered, and a solution discussed.

ACKNOWLEDGEMENT

I would like to thank my husband and long-time mentor, Chris Moore, for his indefatigable guidance and for sharing the extensive technical knowledge and experience he possesses and masters so adeptly.

REFERENCES

- [1] R. Martin, “Clean Architecture,” Prentice Hall, 2018, ISBN-13: 978-0-13-449416-6.
- [2] J. Spolsky, “The Law of Leaky Abstractions,” [Online]. Available from <https://www.joelonsoftware.com/2002/11/11/the-law-of-leaky-abstractions/> [retrieved: June, 2019].
- [3] SourceMaking, Software Architecture AntiPatterns, [Online]. Available from <https://sourcemaking.com/antipatterns> [retrieved: June 2019].
- [4] M. Iridon, “Enterprise Integration Modeling,” International Journal of Advances in Software, vol 9 no 1 & 2, 2016, pp. 116-127.
- [5] Genesys, “Platform SDK,” [Online]. Available from <https://docs.genesys.com/Documentation/PSDK> [retrieved: June, 2019].
- [6] Genesys, “Web Services and Applications,” [Online]. Available from <https://docs.genesys.com/Documentation/HTCC> [retrieved: June, 2019].
- [7] G. M. Hall, “Adaptive Code via C#: Agile coding with design patterns and SOLID principles (Developer Reference),” Microsoft Press, 1st Edition, 2014, ISBN-13: 978-0735683204.
- [8] M. Seemann, “Dependency Injection in .NET,” Manning Publications, 1st Edition., 2011, ISBN-13: 978-1935182504.
- [9] Microsoft, “Extract, Transform, and Load (ETL),” [Online]. Available from <https://docs.microsoft.com/en-us/azure/architecture/data-guide/relational-data/etl> [retrieved: June, 2019].
- [10] G. Simson and G. Witt, “Data Modeling Essentials,” Morgan Kaufmann: 3rd edition, 2004, ISBN-13: 978-0126445510.
- [11] A. Reitbauer, “Performance Anti-Patterns in Database-Driven Applications,” [Online]. Available from <https://www.infoq.com/articles/Anti-Patterns-Alois-Reitbauer/> [retrieved: September, 2019].
- [12] M. Fowler, “Reporting Database,” [Online]. Available from <https://martinfowler.com/bliki/ReportingDatabase.html> [retrieved: September, 2019].
- [13] V. Bilopavlovic, “Can we talk about ORM Crisis?,” [Online]. Available from <https://www.linkedin.com/pulse/can-we-talk-orm-crisis-vedran-bilopavlovic/> [retrieved: June, 2019].
- [14] Jon P. Smith, “Entity Framework Core in Action,” Manning Publications, 2018, ISBN-13: 978-1617294563.
- [15] E. Evans, “Domain-Driven Design: Tackling Complexity in the Heart of Software,” 1st Edition, Prentice Hall, 2003, ISBN-13: 978-0321125217.
- [16] M. Fowler, “Patterns of Enterprise Application Architecture,” Addison-Wesley Professional, 2002.
- [17] M. Fowler, “OrmHate,” [Online]. Available from <https://martinfowler.com/bliki/OrmHate.html> [retrieved: June, 2019].
- [18] Atlassian, “What is version control,” [Online]. Available from <https://www.atlassian.com/git/tutorials/what-is-version-control> [retrieved: September, 2019].
- [19] P. Duvall, “Version everything,” [Online]. Available from <https://www.ibm.com/developerworks/library/a-devops6/> [retrieved: September, 2019].
- [20] “Fluent Migrations Framework for .NET,” [Online]. Available from <https://fluentmigrator.github.io/> [retrieved: September, 2019].
- [21] Microsoft, “SQL Server Data Tools,” [Online]. Available from <https://docs.microsoft.com/en-us/sql/ssdt/sql-server-data-tools> [retrieved: September, 2019].