# AMADEUS MIGRATION PROCESS A SIMULATION-DRIVEN PROCESS TO ENHANCE THE MIGRATION TO A MULTI-CLOUD ENVIRONMENT

Bilel Ben Romdhanne, Mourad Boudia and Nicolas Bondoux

Artificial Intelligence Research, Amadeus SAS, Sophia Antipolis, France

***ABSTRACT***

*With the development of the cloud offers, we observe a prominent trend of applications being migrated from private infrastructure to the cloud. Depending on the application's complexity, the migration can be complex and needs to consider several dimensions, such as dependency issues, service continuity, and the service level agreement (SLA). Amadeus, the travel industry leader, had partnered with Microsoft to migrate its IT ecosystem to the Azure cloud. This work addresses the specificity of cloud-to-cloud migration and the multi-cloud constraints.*

*In this paper, we summarise the Amadeus Migration process. The process aims to drive the migration from an initial private cloud environment to a target environment that can be a public or hybrid cloud. Further, the process focuses on a prediction phase that guides the migration process.*

*This paper expects to provide an efficient decision-making process that guides managers and architects to optimise and secure their migration process while considering micro-services-oriented applications targeting an efficient deployment over multi-cloud or hybrid cloud. The prediction relies on the network simulation to predict applications' behaviour in the cloud and evaluate different scenarios and deployment topologies beforehand. The objective is to predict migrated applications' behaviour and identify any issue related to the performance, the application's dependency on other components, or the deployment in the cloud. The migration process proposed in this paper relies on SimGrid, a toolkit developed by INRIA[52] for distributed application modelling. This framework offers a generic process to model IT infrastructure and can assist cloud-to-cloud migration. Specific attention is given to predictive and reactive optimisations. The first results show predictive optimisation's impact on securing KPI and reactive optimisation to optimise the solution cost. Thus, we reach an average cost reduction of 40% in comparaison with the same deployment strategy while keeping the same SLA.*

***KEYWORDS***

*Cloud migration, SimGrid, system simulation, app modelling, decision support, cloud deployment strategy.*

## 1. INTRODUCTION

In the last decade, cloud infrastructures have become essential IT industry assets. Cloud computing emphasises flexibility, reliability, and agility while increasing efficiency and reducing operational costs. With cloud computing, it is easy to dynamically adapt the resources to the need

David C. Wyld et al. (Eds): ICDIPV, CBIoT, ICAIT, WIMO, NC, CRYPIS, ITCSE, NLCA, CAIML -2023
pp. 93-112, 2023. CS & IT - CSCP 2023 DOI: 10.5121/csit.2023.131308

instead of dedicated resources. Managing private datacentre as private clouds have become the mainstream trend. Commonly, migrating existing applications from a service-oriented design to a cloud-oriented one starts by targeting a private cloud. This step provides a secure software transformation and optimises resource usage. However, fundamental cloud concepts emphasise higher agility and efficiency features. Cloud benefits are maximised when a public cloud is involved or combined with a private one. The notion of multi-cloud deployment is even highlighted as the most scalable and efficient in the context of challenging solutions. However, migrating from private infrastructure to the cloud involves reviewing the IT strategy, the chronology, the applications' architecture, and even the business model for some cases.

Therefore, migration to the cloud is a complex and critical challenge that must be handled carefully. According to professionals' feedback, application dependability, solution redesign, and resource allocations are the most challenging issues.

Identifying and evaluating the overall dependability of a given application or product is a challenge, particularly when its architecture is based on microservices and managed services. Moreover, deploying components in a multi-cloud or hybrid cloud configuration introduces additional complexity, as the interaction between different components can impact the application's performance depending on the deployment strategy. Thus, the network becomes a critical factor in overall performance.

Resources allocation is a multi-faceted issue in the cloud. Oversizing is a strategy that uses ample computing power and high network characteristics to prevent mishaps, prioritising the quality of service over cost. This approach is considered secure during the migration phase, although it is expensive and cannot guarantee service quality or manage dependencies. Nonetheless, oversizing remains a common feature of established migration processes in both academic and industrial contexts.

Based on this observation, this work highlights how a simulation-based prediction can drive the cloud migration process, managing dependencies and optimising KPIs. The proposed migration process enables the user to move from a private cloud to a public or hybrid cloud, considering the target use case and optimising the usage cost. It relies on SimGrid [1] as a simulation framework to predict the application's behaviour. This framework offers a toolkit with core functionalities for the simulation of distributed applications running on an IT infrastructure.

The Amadeus migration process emphasises four phases: assessment, predictive optimisation, migration, and reactive optimisation. Specific attention is given to the predictive optimisation phase since it is the most valuable difference between existing migrations processes and the one highlighted in this work. The remainder of this paper is organised as follows: a state-of-the-art section reviewing related studies, the problem statement section specifies the addressed problem and in which conditions users may efficiently use the proposed process. Amadeus Migration Process section presents the global workflow; the Proof-of-Concept section emphasises the results obtained on a use case; the future work section highlights the extension axes, and the conclusion section discusses the added value of this work in the decision-making process.

## 2. STATE OF THE ART

In the last decade, IT solutions have evolved significantly through two dimensions: software architecture and deployment environment. Microservice architecture has become the architectural trend, while multi-cloud deployment is increasingly popular. This paper presents a state-of-the-art study on the multidimensional aspect of cloud migration, divided into five subsections: (1) the evolution of software architecture, (2) cloud service type, (3) cloud service model, (4) a summary of the academic literature on cloud migration processes, and (5) the use of simulation to support cloud migration.

### 2.1. The Evolution of Software Architecture

The evolution of software designed decades ago is a complex task that involves identifying internal and external dependencies and essential IT resources [1]. The first step is to move away from physical servers and use virtualisation, a process explored extensively in the literature. [2,3,4] The trend is towards industrialising software development, with Service Oriented Architecture (SOA) being a highly effective architectural approach that facilitates system enhancement without altering them. [5] This approach has been widely acknowledged as an appropriate refactoring approach. [6][7] Software industrialisation is now evolving towards the microservices pattern [8], where an application is composed of fine-grained, loosely coupled services that communicate with lightweight protocols [9]. The essence of microservice design targets that team members can develop their services independently based on precise requirements and communication features. [10]

Microservices are a great way to boost system performance and flexibly execute tasks [11]. Communication is minimised, and the advantages of microservices come with the responsibility of maintaining their decoupling [12]. Multiple interfaces or versions of the same service can be used to avoid disruption to existing users. Containerised environments are ideal for managing microservices, as they are lightweight and portable virtualisation solutions [13]. Container processes are visible to the host OS as native processes, but the container itself is not a process. All processes created by the container have access to the same set of resources and libraries [14]. Microservice design and virtualisation (VM & containers) provide high production KPIs, allowing software designers to focus on the logical flow. Moreover, software following microservice design is natively prepared to take advantage of the cloud.

### 2.2. Typologies of the Cloud Deployment

The concept of the cloud lies in pooling resources and optimising their usage. The concept of the cloud emphasises flexibility and agility while maximising performance and efficiency. The cloud computing deployment models define how resources are managed and delivered[15][16]. There are three main cloud deployment models: public, private, and hybrid. In recent years, multi-cloud deployment has emerged as a relevant model that reduces dependencies and enhances efficiency [42]. Nevertheless, we consider it as a specific case of hybrid cloud. **Public cloud**: Its name comes from the fact that its resources and services are available and provided to the general public. It is the least restrictive deployment model regarding accessibility and availability to a broad audience[15]. In the public cloud model, the cloud service provider controls the computing resources [16].

**Private cloud**: One organisation owns and controls services and resources. An organisation can consist of multiple business units as users of cloud services. Ownership, management, and operations of the private cloud may be delegated to a third party, and the physical location may be on or off the premises of the cloud service customer [16].

**Hybrid cloud**: A hybrid cloud combines two or more distinct cloud deployment models, such as public, private, or a mix of both, while keeping them separate entities. Innovative technology allows data and applications to be transferred and shared between clouds [15][16]. Ownership of the hybrid cloud may be shared between the organisations that control each component. In this setup, the private cloud may be owned by one organisation, while the public cloud is provided by a service provider. Furthermore, customers can host resources on- or off-site [16]. Balancing between the clouds, or extending the private cloud to include public cloud resources and services, are further possibilities.

Multi-cloud deployments, which spread services geographically or between multiple cloud providers, are a specialised case of the hybrid cloud model. They usually consist of multiple public clouds and can also incorporate a private cloud. As public cloud adoption continues to grow, multi-cloud deployments are becoming more prevalent [24].

## 2.3. Typologies of Cloud Service Models

Most of the literature classifies the cloud service models into three categories based on the abstraction level: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS).

**Infrastructure as a Service** (IaaS) provides fundamental computing resources such as processing, storage or networking resources to deploy and run arbitrary software, including operating systems and applications [16]. In the IaaS service model, the provider manages and controls the physical cloud infrastructure. However, the customer controls operating systems, storage, applications deployment and how the virtual networking components are configured [15]. An IaaS cloud can reduce IT infrastructure capital costs and allow flexibility to meet fluctuating capacity demands while paying only for the actual allocated capacity [18]. IaaS represents 22.34% of the cloud market, reaching 91.3 $billion in 2021. The growth is a consequence of the growing (33.9% year over year growth) demand for applications and workloads requiring infrastructure that traditional data centres cannot meet. The top five IaaS vendors, Amazon (40%), Microsoft (21.9%), Alibaba (6.1%), Google (5.5%) and IBM (2.5% ), covered 76% of the global IaaS market [19].

In **the platform as a Service** (PaaS) model, the cloud service consists of a set of platforms that allow customers to deploy and run either created or acquired applications without the complications of managing and controlling the underlying infrastructure [16]. PaaS is an abstraction layer on top of the IaaS services. Instead of configuring the servers, storage and operating system, the customer configures the execution environment and manages the applications that use the platform services via programming languages, libraries and service interfaces provided by the provider [15]. PaaS represents 16.7% of the cloud market, reaching 68.2 $billion in 2021. The growth is a consequence of the growing demand for productivity compared to IaaS (39.1% year-over-year growth). The top five PaaS vendors are the same as IaaS, Amazon (34%), Microsoft (22%), Alibaba (9.1%), Google (10%) and IBM (5%) covered 80% of the global IaaS market in 2021 [19].

**Software as a Service** (SaaS) means the cloud hosts the provider's applications and makes them available to customers [15, 17]. Access to the SaaS application is typically provided through a web browser, web service or APIs [16]. The provider hosts all the required infrastructure, and the customer does not manage or control the cloud infrastructure. Customers cannot control application capabilities but may use limited user-specific configuration settings [17].

Recently, the function as a service (FaaS) service model has been separated from PaaS & SaaS [20]. FaaS operates in response to events and is used for short-running, stateless computation and event-driven applications. As with PaaS, it enables running the code without provisioning or managing servers but without the need to address the instances or pre-allocate resources [20].

## 2.4. State-of-the-Art Regarding the Migration Process

Cloud migration methodologies help plan, design, and implement migrations of chosen applications and workloads to the cloud[22][35]. Cloud actors provide their processes and tools for the migration process, which involves moving an application or computational workload into a cloud environment [21][38]. The literature explores several approaches to cloud migration. Summary of the major cloud migration processes [17] reveals more than nine migration processes from both academic and industrial sources [39].

Cloud Reference Migration Model (Cloud-RMM) [26] synthesises and consolidates existing research. Variability-based, Pattern-driven Architecture Migration (V-PAM) [24][17] enhances the academic picture. Cloud actors provide their processes operating their services and tools. Each of the four major cloud providers, Amazon [28], Google [32][22], IBM [29], and Microsoft[30], published their migration process.

This section focuses on V-PAM [31][17] and the Microsoft migration process[33][34], which inspire and influence the Amadeus migration process.

V-PAM encompasses nine activities: Define Organization Profile, Evaluate Organisational Constraints, Define Application Profile, Define Cloud Provider Profile, Evaluate Technical and Financial Constraints, Address Application Constraints, Change Cloud Provider, Define Migration Strategy and Perform Migration. It deals with multi-cloud environments and links migration patterns for cloud architectures to cloud service models. Migration graphs enable users to plan incremental optimisations. Additionally, the process model has a feedback loop[24].

Microsoft's four-phase migration process[33][34] consists of Assess, Migrate, Optimise, and Secure and Manage. During assessment, business and IT stakeholders align objectives. In the migration phase, a migration strategy is chosen based on the organisation's requirements. The optimisation phase monitors the cloud environment's performance, scalability and economics. Finally, the secure and Manage phase includes securing cloud resources, data protection, backup, failover and recovery planning.

V-PAM is generic and complete, and Microsoft is responsive, relying on feedback to optimise. Microsoft includes optimisation and security and management in the cloud migration methodology, whereas V-PAM does not.

## 2.5. State-of-the-Art, Simulation-Driven Migration.

Most of the listed processes rely on analytics or optimisation after deployment on the cloud [36][37]. However, there are several approaches in the literature addressing the prediction of an application's behaviour before cloud deployment, through system simulation or emulation [4345].

These require three models: the original environment, the application model and the target environment [47][48]. Three core simulation ecosystems have been identified: NS3[49-50], CloudSim[51-52] and SimGrid[53-56].

NS3 is a low-level engine, requiring significant effort to model an application and the target environment, limiting its usage to academic and innovation rather than production support. CloudSim, while providing a higher abstraction level, is not up to date with the latest public cloud updates [52]. SimGrid, on the other hand, has a mature software framework and a wide users' community, which makes it the most suitable option for our framework [53-54]. It offers multiple abstraction levels, allowing users to manage the simulation at the network packet level, while also allowing the use of real application code, microservice or components [56]. It bridges the gap between research and production, providing a high degree of productivity and quality. Therefore, SimGrid is our chosen simulation engine.

## 3. PROBLEM STATEMENT

The state-of-the-art regarding Cloud migration methodologies covers a wide range of industrial requirements. While these works provide a comprehensive coverage, two significant breaches remain. First, cloud-to-cloud migration needs to be addressed - a frequent occurrence in industrial settings - which generates requirements that differ from general migration cases. Refactoring software becomes inefficient here, as the target application is usually already cloud compliant. Secondly, performance evaluation and optimisation is lacking, as deployment optimisation is a post-deployment approach, relying on gathering of cloud monitoring traces into useful KPIs.

The Amadeus migration process goes beyond simply migrating cloud-ready, industrial, and complex applications to the target cloud, but also provides a comprehensive optimisation of those applications. Utilising the power of simulation-driven processes powered by SimGrid, Amadeus provides a high degree of accuracy in predicting the behaviour of applications on the target cloud. This predictive dimension has enabled the industrialisation of the migration process, enabling massive data processing and workflow management automation for streamlined cloud-to-cloud migrations. Furthermore, the Amadeus migration process is designed to be highly user-friendly, allowing users to quickly and easily migrate their applications with minimal disruption and downtime.

## 4. AMADEUS MIGRATION PROCESS

*Amadeus migration process* is a comprehensive, conceptually-sound approach to cloud-to-cloud migration that not only anticipates behaviours on the target cloud, but also provides assurance to the customer that their services will be delivered as promised. Through this process, business stakeholders can accurately gauge the service they will receive and the associated Service Level Agreements, while technical stakeholders can consolidate, optimise, refactor, and redesign the deployment architecture. The *Amadeus migration process* is divided into four distinct phases: Assessing, Predictive Optimisation, Migration, and Reactive Optimisation. All of these are overseen by the transversal management phase, which is not detailed in this paper. This process further enhances the Microsoft model by placing prediction at the heart of the process and making optimisation an iterative part of the migration process.

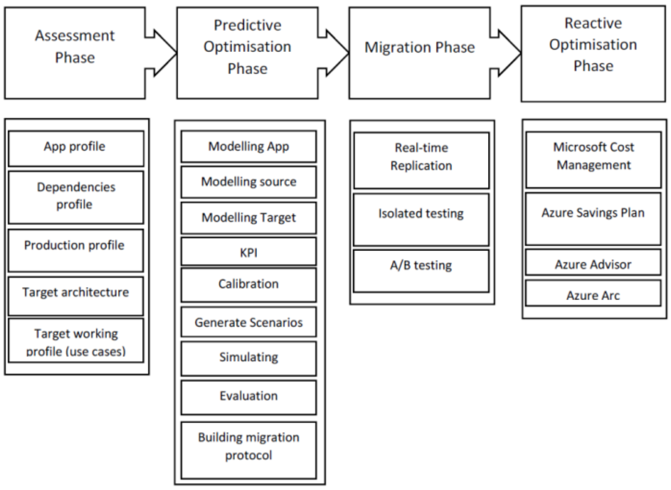


Figure 1. Amadeus Migration Process diagram: Composed of four consecutive phases. Predictive optimisation becomes the core of the process, providing a safe approach to driving the migration

## 4.1. Assessment Phase

This phase covers all necessary tasks to understand the application, the initial environment, the target environment, and the requirements. Since this process is production oriented, the assessment phase relies on a profiling approach. It is subdivided into five tasks: (1) Define the application profile, (2) Identify dependencies, (3) Define the production profile, (4) Analyse the target architecture and (5) Identify the working target profile.

### 4.1.1. Define the Application Profile

We propose a cloud-to-cloud migration process, considering applications deployed on a private cloud or compatible with a virtual environment. The aim of this approach is to identify all microservices involved in service rendering; both microservice-designed applications and containerised software are almost always considered microservices from a profiling perspective. This task requires analysing an application's design, architecture and functioning to understand its profile accurately. The internal dependency graph of the application should represent its logical design and clearly show all microservices and components, including those with low network usage. These can be obscured by active components that consume resources. Authentication components, for example, generally have low network activity yet play a vital role in the application flow. Therefore, technical stakeholders must consider hidden components when examining and validating profiling results.

### 4.1.2. Identify Dependencies

By studying application dependencies, three essential aspects must be considered: logical, hardware, and external service dependencies. Stakeholders typically supply the architecture, configuration, hardware specs, and virtualisation strategy, which is essential for comprehending the application and its environment. This analysis also takes into account the details of servers, networks, and databases on which the application relies, as well as external services dependencies. An example of an external service is Single-Sign-On.

Many tools map server dependencies yet fail to map application dependencies. To achieve a full picture of workload communication, stakeholders need a tool that can accomplish both. This will produce a visual map of applications and their respective workloads. The outcome must be an improved application graph that includes internal and external dependencies, allocated resources, and initial cloud configuration.

#### 4.1.3. Define the Production Profile

Collecting resource usage reporting (e.g. CPU, memory, and storage) is a vital step in the Assess process. It helps create a realistic calibration of the models used in prediction. Examining production logs in detail is a useful approach to comprehend application behavior and model building. These logs tend to be large, necessitating cleaning, parsing, filtering, and pairing to acquire useful data from raw logs. Sampling is recommended to manage the immense log volume while maintaining accuracy. Nonetheless, sample management is the developer's obligation. Network view plus CPU load and disk usage traces supply information on application behavior in current production settings, allowing construction of a statistical distribution of all flows and CPU load.

#### 4.1.4. Analyse the Target Architecture

Amadeus migration process expects that technical stakeholders to provide a projected target architecture. It is quite unlikely that the target architecture is a mirror deployment of the source one. Target architecture reflects the expectation of technical experts regarding resource allocation and deployment. Resources allocation addresses the sizing (VM, network and storage sizing).
The deployment addresses the geographic deployment and the management of customers.

Amadeus migration process accepts several concurrent target architectures. Each target architecture focuses on the distribution of microservices that define the application in the first place and the interaction between the application and its dependencies in the second place. For example, the first architecture candidate may define a complete migration from a private cloud to a public one. The second architecture, in contrast, may define a partial migration; the frontend is concerned with migrating to the public cloud while the backend remains in the private one.

#### 4.1.5. Identify the Working Target Profile

The proposed migration process considers that changing the deployment infrastructure may have several goals. It may target financial optimisation, but also it may target technical enhancement. For example, the migration may target reducing the response time, optimising the service or both. Accordingly, identifying the working targets to specify what the organisation expects from this process. The task helps stakeholders to identify their needs by explicitly detailing the KPI that must be reached.

### 4.2. Predictive Optimisation Phase

The predictive optimisation phase targets the estimation of the application behaviour over the target deployment infrastructure according to the expected workload. It relies on the simulation of the system to guarantee reliable and refined predictions. It is subdivided into five subtasks surrounding the simulation progression: modelling task, KPI & Models calibration, generating deployment scenarios, simulating, and evaluation and refinement.

#### 4.2.1. Modelling Task

Modelling is the process of developing a simplified algorithm to represent a complex system or component for use in a simulation engine. The abstraction level is the difference between the real component and the model, which must accurately reflect dependencies and resource utilisation. Predictive optimisation then predicts the behaviour of the studied applications when moved to the target cloud. In this purpose, three models are necessary: (i) the source cloud environment, (ii) the target cloud environment, and (iii) the application question of migration.

The Amadeus Migration process leverages the power of SimGrid, which offers an S4U API to provide abstractions and models for modelling computation, I/O, and communication activities on different hardware platforms. XML files are used to define the network topology, endpoints, compute resources, and storage resources. SimGrid simulations require two components: the simulation executable and XML configuration files. The executable includes the source code of each actor to be run in the simulation.

In order to ensure the accuracy of the SimGrid model, the actor abstraction level must be tailored to the scope of the study. Thus, the micro-service is the most suitable abstraction level as it allows for accurate isolation of components from a network perspective. Consequently, each micro-service of the application is modelled using one actor, while the entire application's model is defined as the interconnection of all the actors. It is of paramount importance that the SimGrid model replicates the behaviour of the application running in the production settings with similar results, thus allowing for the comparison of the defined KPIs across the two environments. Any divergence between the KPIs metrics between the two environments is referred to as the error rate.

#### 4.2.2. KPI and Models Calibration

Models' accuracy is a highly sought-after metric in the literature, as it is essential to assess the performance of any simulation properly. As such, the Amadeus migration process puts a special focus on resource usage and network performance, as evidenced by its use of Key Performance Indicators (KPI). To guarantee the accuracy of the application model and the source cloud infrastructure, the proposed process proposes to reproduce the same behaviour as it is under production conditions since it is believed that this is sufficient to verify the accuracy of the model from a practical standpoint. To that end, the calibration methodology proposed requires two steps: Key Performance Indicator pattern definition and model calibration.

*KPI pattern definition* is an essential tool for analysing production traces in order to accurately identify the performance of any given KPI. Response-time is a common operational KPI used in the industry and is indicative of the performance of the underlying infrastructure. Technical KPIs such as CPU/GPU usage can also provide valuable insights into how a given microservice is using computing resources to complete its tasks. Ultimately, the KPI pattern definition will create a set of KPI patterns that aggregate the requirements of stakeholders, providing a goal for the model calibration subtask.

*Model calibration* consists of tuning the model’s parameters until reaching a model whose behaviour comes as close as possible to that of the modelled system. Using SimGrid, it is possible to replay production traces facing the application and infrastructure models and analyse the divergence between the original and simulated system behaviour. We combine the divergence between application and model metrics into a global fidelity score to define a measurable accuracy evaluation.

### 4.2.3. Generating Deployment Scenarios

The Amadeus migration process gathers two components under the “deployment scenario”: the deployment architecture and the simulation scenario. The deployment architecture outlines the exact manner in which an application is deployed on the execution infrastructure, such as a micro-service-oriented architecture. This includes pinpointing where each microservice is deployed (VM, container) and how the microservices are interconnected (network specifications and geographical deployment). The simulation scenario provided by the SimGrid features adds even more efficiency and productivity by enabling the separation between deployment and environment descriptions. This helps to meet the demands of industrialisation and streamline the entire process.

The Amadeus migration process supports concurrent deployment architectures to optimise performance. Cloud architects consider several options, each tailored to a specific set of Key Performance Indicators (KPIs). In order to accurately simulate a use case, real-world production traces or anticipated future use cases can be used as inputs for the simulator. For example, businesses can anticipate a load increase by a factor of 3-5 for the post-Covid recovery and use scripting language and trace files to simulate the expected use case.

Simulation scenarios may also rely on a trace generator that simulates the expected use case. In such conditions, a fourth model is also necessary to accurately represent the behaviour of the external users. The simulation scenario can be thought of as a functional representation of a use case; it is written as a software application, known as the application model. Generally, the programming language used to construct the model is C++. The result of this task is a pool of deployment scenarios, referred to as deployments candidates. These deployments candidates can then be used to test and validate the use case, ensuring it meets all the necessary requirements.

### 4.2.4. Simulating

Amadeus migration process is a specified process that combines three targets: (i) cloud-to-cloud migration, (ii) microservices-oriented application and (iii) simulation-based prediction. The simulation task summarises the activities of simulating each deployment candidate iteratively, collecting simulation traces and providing pre-processed data. Pre-processed data is the necessary input for the evaluation task. Each deployment candidate regroups all required models and data to lunch a simulation. 1A migration process acclaim the SimGrid simulation engine to ensure the simulation task.

A SimGrid simulation consists of a series of sequential threads of control that consume hardware resources. The threads of control can execute arbitrary code, exchange messages over a simulated network, run computations on simulated (multicore) hosts, and perform I/O on simulated devices. Further, SimGrid provides a VM abstraction that incorporates migration functionality. In other words, it provides all the necessary base abstractions to simulate the migration from a given cloud infrastructure to another cloud one.

SimGrid relies on a packet tracing paradigm. The raw simulation result is a packet trace log containing information about the source, destination, timestamp, size, flow-id, and transactionId. Due to the large size of the simulation traces, it is recommended to pre-process them into valuable and easy-to-use data.

The pre-processing sub-tasks gather packet traces into network flows, network throughput information, average CPU usage, disk throughput, components response time, and network response time—the relevant pre-processing approach depends on the required KPI to evaluate the performance of a given deployment candidate.

#### 4.2.5. Evaluation and Refinement

Evaluation and refinement tasks target to provide a helpful evaluation of each deployment candidate. For each candidate, the evaluation consists of processing the simulation output to generate a user-friendly KPI, helping the decision-makers. The details and abstraction level of the KPI depend on the target audience. The refinement subtask, in contrast, relies on the evaluation output to optimise the specification of the deployment candidate. Accordingly, it is an iterative process that may be repeated until reaching the expected performance. A widely considered KPI includes CPU usage, disk usage, network load, response time and energy consumption. SimGrid engine supports the monitoring and analysis of These KPIs. Optimising and reviewing the architecture of the deployment candidate may also result in resizing the resources, modification of the network topology or improvement of the deployment architecture. The output of the evaluation and refinement task is a deployment candidate defined as the most suitable solution that combines all requirements, or at least, the solution that mitigates risks.

### 4.3. Migration Phase

In the Migrate phase, migration managers will determine the strategy that best meets stakeholders' requirements—and this is usually best addressed on a per-application basis. In this step, the target is to physically move given workloads and applications (including their data) from one cloud to another. Every organisation will have a different approach, and application architects' might mix rehosting, refactoring, rearchitecting, rebuilding, and replacing your applications. As in the case of the Microsoft migration process, the Amadeus migration process recommends the usage of real-time replication and intensive isolated testing to validate the migration. In contrast, the target is not a full migration but a possible cohabitation and hybrid architecture. Thus, the Amadeus approach introduces an extended A/B testing task that provides a real-world load-balancing between replicates. Accordingly, the migration phase is subdivided into three sub-tasks: (1) real-time replication, (2) isolated testing, and (3)

#### 4.3.1. Real-Time Replication

Modern tools enable the system to cleanly migrate real-time data even when the system is actively being used. Real-time replication allows the workload to remain online and accessible during the migration. Real-time replication involves setting up a copy of the workload in the target cloud and allowing asynchronous replication to keep the copy and the original in sync. This means that data or server updates are synchronised between the copies while the migration managers are building and executing the migration plans. This model also enables groups of VMs to be connected, for example, for a multitiered application or workload. When the system has a map of the connections and dependencies among VMs, it becomes possible to create plans to ensure the VMs are bought up in the correct order when starting.

Most migration tools have a mechanism to define the replication timeframe (usually between 30 seconds and 15 minutes). This timeframe is based on the network capability and latency. Many tools also support application-aware replication automatically. Initial replication is also bandwidth intensive. In the specific case of migration to Microsoft cloud, this task can rely on dedicated mechanisms like ExpressRoute and Data Box. It's something to consider when planning the migration timeline.

#### 4.3.2. Isolated Testing

Testing is integral to ensuring system health before the final cutover. Many migration tools have options that let the user start up a set of VMs in an isolated environment. This means it is possible to thoroughly test the application without affecting the source cloud production versions. Once replication is complete, it is mandatory to intensively test and evaluate the deployed application or workloads using the isolated environment. Further, the testing task enables a retro-evaluation of the simulated models and the quality of the previous predictive optimisation. Migration managers can validate the candidate if the simulation results are consistent with the deployment test measurements. When the stakeholders are satisfied with the application's behaviour in the isolated environment, it is possible to throw the final task: the enhanced A/B testing.

#### 4.3.3. A/B Testing

A/B testing is a user experience research methodology that consists of a randomised experiment that usually involves two variants of a given variable (A and B). However, the concept can also be extended to multiple variants of the same variable. A/B testing is a way to compare various application deployment architectures, for example, by testing the response behaviour of variant A against variant B and determining which of the variants is more effective. In the Amadeus Migration Process context, A/B testing will be the last line to validate the migration hypotheses. A/B testing is adopted for each deployment configuration with a debut between concurrent scenarios. In this work, we identify three typical use cases that cope with the A/B testing as a last validation method: (i)The new cloud copy challenges the original cloud version, (ii) The referential deployment scenario is challenged by an innovative one, (iii) A KPI-optimized scenario challenges the referential deployment one.

If the challenging deployment configuration achieves the expected performance, feedback is provided to the stakeholders to guide their decision. In such a case, the recommended migration approach uses a load-balancing interface that distributes a low percentage of the production traffic from the reference infrastructure to the challenging one. Typically, this percentage is increased progressively until reaching a production threshold (between 20% and 30%).

### 4.4. Reactive Optimization Phase

For the purpose of continuous optimisation, this phase takes place during production to ensure continual efficiency and adapt the resources usages to the real needs and eventually update the architecture if needed. Whatever the performance of the prediction phase, there would be a difference between the prediction results and the production monitoring. Moreover, the prediction phase relies on "assumptions." regarding the production load. In contrast, the reactive optimisation phase depends on the production monitoring traces. Cloud providers promote several tools that emphasise reactive optimisation based on their ecosystem. This work relies on the Azure ecosystem in the scope of migration to the Microsoft Azure cloud. The Microsoft Azure migration methodology claims that Azure optimisation tools realise savings and improve agility by increasing the efficiency of cloud investments. Accordingly, it aims to enhance the usage rate of cloud resources and, consequently, the cloud cost.

From a perspective of academic and industrial rigour, the Amadeus migration process emphasises four subtasks: (1) optimising through Microsoft Cost Management, (2) Cost reduction using Azure Savings plan for compute, (3) Cost reduction using Azure Advisor and (4) hybrid resources management using Azure Arc.

#### 4.4.1. Optimisation through Microsoft Cost Management (MCM)

MCM provides a complete set of cloud cost management capabilities, enabling Azure customers to optimise their cloud spending and improve financial governance for that spending. It is integrated into the Azure portal, always on by default, and available at no additional cost. All usage statistics and cost data are displayed in intuitive dashboards and report. It includes a detailed monitoring dashboard that can be customised as needed. Further, it provides a precise and explicit estimation of the cloud costs that may be used. The granularity is also subject to customisation, allowing multi-level optimisation.

#### 4.4.2. Optimisation through Azure Saving Plan (ASP)

Azure savings plan for compute is a flexible pricing model that provides savings feature compared to pay-as-you-go pricing when the user commits to spending a fixed hourly amount on compute services for one or three years. This model is competitive in the case of predictable computing needs. It can reach up to 65% savings of pay-as-you-go pricing. This possibility combines perfectly with the prediction phase that we have introduced with our methodology and increases the financial efficiency of the migration plan put into action.

#### 4.4.3. Optimisation through Microsoft Azure Advisor (MAA)

MAA helps users optimise their Azure cloud deployments through personalised cloud consulting. The Advisor provides proactive, actionable, and personalised best practice recommendations. This software analyses the configuration of resources and the telemetry of their usage. As a result, it recommends solutions to improve Azure resources' cost-effectiveness, performance, reliability, and security. In addition, advisor recommendations aim to enhance cloud-deployed resources. Additionally, it identifies areas for cost reduction in Azure.

#### 4.4.4. Optimisation through Microsoft Azure ARC (MAAc )

Azure Arc is a bridge that extends the Azure platform to drive building applications and services with the flexibility to run across data centres, at the edge, and in multi-cloud environments. Azure Arc extends management and services from Azure to any infrastructure. Moreover, Azure's core control plane is accessible to customers, while all services and management are derived from Azure.

## 5. POC: Simulating a Middle-Range Application

Our first Proof of Concept was a mid-size Amadeus application. It is a critical application with discernible reliance on external applications and data sources. Further, it is a fully virtualised cloud-candidate application and industrialised under private cloud infrastructure. Finally, it is an efficient application optimised to reach an aggressive SLA. Thus, Applying the Amadeus migration process successfully to this Application while predating its behaviour correctly during the predictive optimisation phase can support taking the right decisions and avoid possible negative impacts on the SLA. Further, it will lead to an easy internal adaptation of the migration process when targeting larger applications. To consolidate this paper with industrial material while protecting Amadeus's confidentiality, we assume that the application is called FLY. In this section, we detailed the flow of the Amadeus migration process through the FLY application.

## Assessment Phase

During the assessment phase, we identify the application profile, its internal and external dependencies, the production profile, and the target architecture. Fly simplified design involves nine known microservices that ensure around 95% of the service. Fly requires the input of more than thirty-five external services; however, four services cover around 90% of the requirement. FLY interacts mainly with two productions and one log database. Figure 2 illustrates the highlevel design of FLY with its nine microservices, external dependencies and related database. Defining the production profile determines the application's resource usage and traffic load needs. It includes the identification of KPIs that must be supervised and identified. Production logs were sampled from both frontend and backend tracing frameworks. This step is the most rigid and context-dependent one. It requires a customised investment for each ecosystem and specific processing for each application. Internal traffic is approximated using mathematical models.

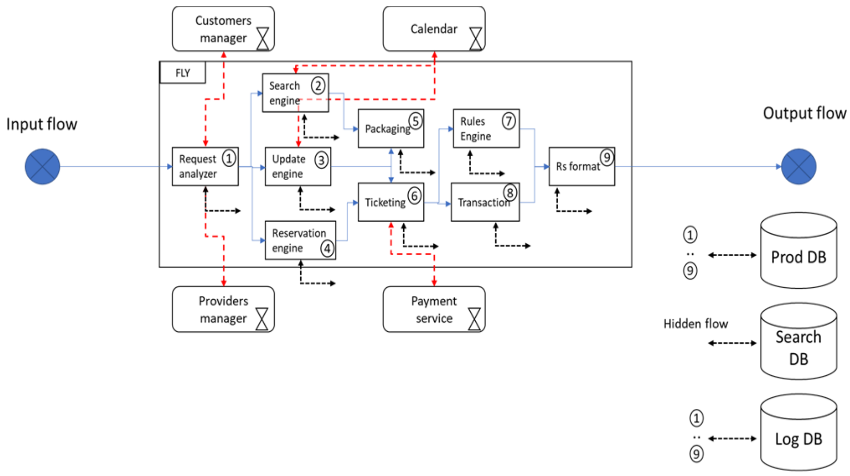


Figure 1: simplified Microservice diagram with external dependencies. One of the main outputs of the assessment phase.

In the Amadeus migration process context, we schematise the production behaviour as a weighted messages graph, as summarised in Figure 3, where we find the probability of generating a message in each link as a response to a random input request. This model is the conclusion of the average behaviour during an observation window. Depending on the required granularity, the modelling process may consider additional parameters such as the probability function of the time and the input request's nature. Both parameters were evaluated during our study. Nevertheless, for confidentiality requirements, we consider the average value in this publication. During the assessment phase, it is required to identify and build a schematic presentation of the production infrastructure, including the involved computing power and the network capabilities. The abstraction level depends on the deployment target capabilities. Thus, if users cannot choose the network hardware, the modelling can focus on network characteristics. In this PoC, we consider the link latency and the bandwidth for each link.

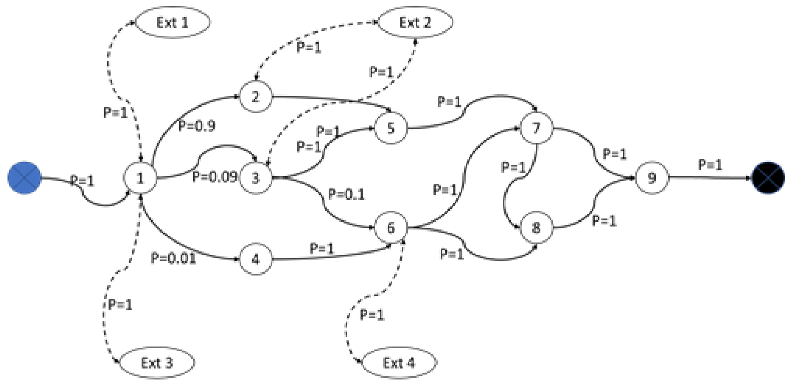


Figure 2: Message diagram with generation probabilities; it is necessary to drive the model created during the optimisation phases.

## Predictive Optimisation Phase

In this work, we highlighted a comparison of evolving four deployment candidates: (A) Full migration to the cloud (FLY+ external services), (B) Full application migration (FLY), (C) hybrid deployment where only search flow is offloaded to the public cloud, and (D) hybrid deployment where FLY except the search flow is offloaded. We expose to summary results: the response time function of the input flow as shown in Figure 4 and the response time function of the VM size, shown in figure 5. For confidentiality motives, the response time is expressed as a percentage of the target value (the SLA); a given point of each curve means that the response time of X% of the transaction is lower or equal to that percentage of the SLA. Simulations highlighted in Figure 4 use a VM model sized to cope with the nominal flow and referred to as V. The nominal flow is referred N, 2N means that we simulate the double of the nominal flow. Figure 5 summarises the simulation that targets the nominal flow where we change the VM's size; 0.5v implies that the used VM have half core compared to V.

Simulation results confirm that candidate D cannot cope with the expected response time and can be eliminated. Candidates A, B, and C perform in the acceptable zone, up to two times the nominal input load. Candidate A will ensure acceptable performance up to three times the nominal load. In contrast, candidates that perform lower than the acceptable zone also present a higher cloud cost. In addition, the simulation summarised in Figure 5 allows identifying the intrinsic network inertia when VMs are oversized. Further, we can identify the smallest VM size that can cope with the expected performance for each candidate.

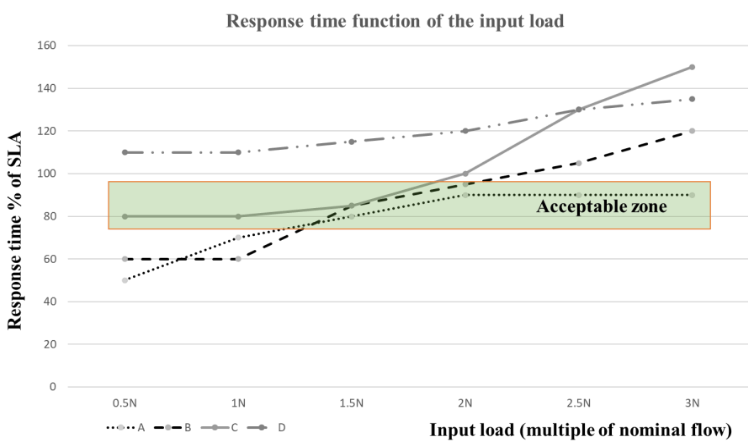

Figure 4. FLY response time of each deployment candidate function of the input load.

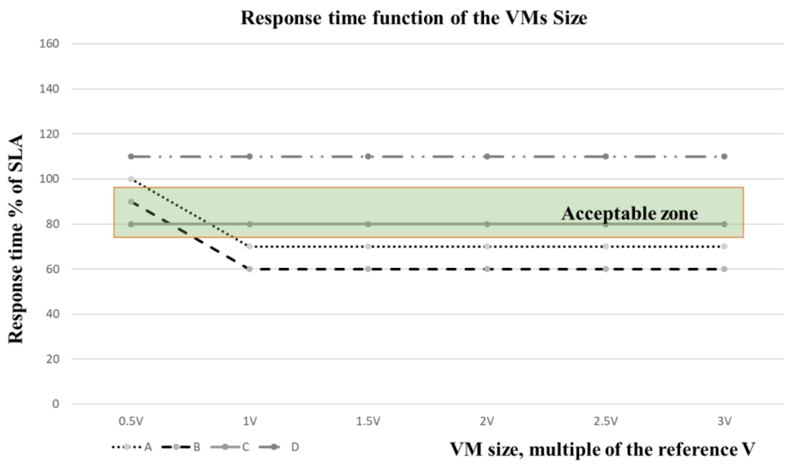


Figure 5. FLY response time of each deployment candidate function of the VM size.

## 6. Future Work

Amadeus' migration process is a simulation-driven approach that aims to improve the transition of a cloud-ready application from a local-cloud infrastructure to a multi-cloud one. The process structures the migration to ensure the expected performance of the deployed system by simulating it. The user needs to provide detailed inputs such as the application, source infrastructure, and target environment, as well as target deployment scenarios and use cases. The process provides a recommendation on how to proceed, but the user needs to have an understanding of the target cloud(s) and know how to optimize the solution. The process allows for comparison between several scenarios yet does not guarantee the best results.

Future work is focused on making the process more user-friendly. This includes developing a deployment scenario generator, an accuracy rate estimation module, and an intuitive interface. Moreover, an integrated framework is in the works, where users can automate all steps in one interface for optimal productivity.

The deployment scenario generator module utilizes machine-learning techniques to generate deployment scenarios that meet user expectations and the target cloud's specifications. This approach requires a large number of existing, evaluated deployment scenarios that have been classified by Amadeus experts. Therefore, hundreds of scenarios must be evaluated to train the module and produce an efficient solution. The accuracy rate estimation module will provide a fast estimation of the process' accuracy, taking into account the used models' granularity, the deployment scenario's clarity, and the existing knowledge about the simulation accuracy.

## 7. Conclusion

Amadeus Migration Process aims to enhance the industrialisation of the infrastructure transition to the cloud while considering multi-factor constraints. It focuses on modern applications that respect the microservice design pattern and target hybrid cloud deployment.

Amadeus Migration Process emphasises continuous optimisation that surrounds the migration. The productive optimisation upstream of the migration relies on system simulation using the SimGrid engine. It enhances the prediction of the application behaviour across the target cloud. It supports complex deployment models that involve hybrid and multi-cloud solutions. Further, the simulation-based prediction can detect architectural and design weaknesses while adjusting the overall costs.

Reactive optimisation, in contrast, ensures continuous and financial optimisation during production. It remains a provider-dependent task since it relies on provider tools. Amadeus optimisation process, in its earlier version, targeted Microsoft Azure Cloud and, thus, depends on Azure tools to ensure the reactive optimisation.

Reactive optimisation evaluation is still in progress. The first proof of concept demonstrates the possibility of identifying architectural weakness quickly and provides an empirical review of different deployment candidates. Further details about the results will be provided in the extended version of this work. The next version of the Amadeus Migration process would consider further cloud providers and the transversal management phase.

## REFERENCES

[1] Arena, L., Mola, L., Remond, N., & Rowe, F. (2020). How do enterprise software providers adapt their strategies to the cloud? An analysis through SAP Hana journey based on the evolution of SAP's discourse (2010-2018).

[2] Inverardi, P. (2006, November). Software of the future is the future of Software?. In International Symposium on Trustworthy Global Computing (pp. 69-85). Springer, Berlin, Heidelberg.

[3] Cheng, S. W., Garlan, D., Schmerl, B., Steenkiste, P., & Hu, N. (2002, July). Software architecture-based adaptation for grid computing. In *Proceedings 11th IEEE International Symposium on High Performance Distributed Computing* (pp. 389-398). IEEE.

[4] Pahl, C., Jamshidi, P., & Weyns, D. (2017). Cloud architecture continuity: Change models and change rules for sustainable cloud software architectures. *Journal of Software: Evolution and Process*, *29*(2), e1849.

[5] Valipour, M. H., AmirZafari, B., Maleki, K. N., & Daneshpour, N. (2009, August). A brief survey of software architecture concepts and service oriented architecture. In *2009 2nd IEEE International Conference on Computer Science and Information Technology* (pp. 34-38). IEEE.

[6] Stocker, M., & Zimmermann, O. (2021, September). From Code Refactoring to API Refactoring: Agile Service Design and Evolution. In *Symposium and Summer School on Service-Oriented Computing* (pp. 174-193). Springer, Cham.

[7] Krafzig, D., Banke, K., & Slama, D. (2005). *Enterprise SOA: service-oriented architecture best practices*. Prentice Hall Professional.

[8] Pacheco, V. F. (2018). *Microservice Patterns and Best Practices: Explore patterns like CQRS and event sourcing to create scalable, maintainable, and testable microservices*. Packt Publishing Ltd.

[9] Balalaie, A., Heydarnoori, A., & Jamshidi, P. (2016). Microservices architecture enables devops: Migration to a cloudnative architecture. *Ieee Software*, *33*(3), 42-52.

[10] Carnell, J., & Sánchez, I. H. (2021). *Spring microservices in action*. Simon and Schuster.

[11] Yang, T., Li, B., Shen, J., Su, Y., Yang, Y., & Lyu, M. R. (2022). Managing Service Dependency for Cloud Reliability: The Industrial Practice. arXiv preprint arXiv:2210.06249.

[12] Habibullah, S. (2021). *Evolving Legacy Enterprise Systems with Microservices-based Architecture in Cloud Environments* (Doctoral dissertation, Edinburgh Napier University).

[13] Goldschmidt, T., Hauck-Stattelmann, S., Malakuti, S., & Grüner, S. (2018). Container-based architecture for flexible industrial control applications. *Journal of Systems Architecture*, *84*, 28-36.

[14] Cinque, M., Cotroneo, D., De Simone, L., & Rosiello, S. (2021). Virtualizing mixed-criticality systems: A survey on industrial trends and issues. *Future Generation Computer Systems*.

[15] Polash, F., Abuhussein, A., & Shiva, S. (2014, December). A survey of cloud computing taxonomies: Rationale and overview. In The 9th International Conference for Internet Technology and Secured Transactions (ICITST-2014) (pp. 459465). IEEE.
[16] Mell, P. and Grance, T. The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technology.
[17] Pulkkinen, M. (2020). Cloud migration strategy factors and migration processes.D. C. Chou. “Cloud computing: A value creation model”. In: Computer Standards & Interfaces 38 (2015), pp. 72–77. issn: 0920-5489. doi: https://doi.org/10.1016/ j.csi.2014.10.001.
[18] B. Sotomayor, R. S. Montero, I. M. Llorente, and I. Foster. “Virtual infrastructure management in private and hybrid clouds”. In: IEEE Internet computing 13.5 (2009), pp. 14–22.
[19] https://www.wpoven.com/blog/cloud-market-share
[20] Fox, G. C., Ishakian, V., Muthusamy, V., & Slominski, A. (2017). Status of serverless computing and function-as-a-service (faas) in industry and research. CoRR, vol. abs/1708, 8028.
[21] Jamshidi, P., Ahmad, A., & Pahl, C. (2013). Cloud migration research: a systematic review. IEEE transactions on cloud computing, 1(2), 142-157
[22] N. Ahmad, Q. N. Naveed and N. Hoda, "Strategy and procedures for Migration to the Cloud Computing," *2018 IEEE 5th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*, 2018, pp. 1-5, doi: 10.1109/ICETAS.2018.8629101.
[23] Zhao, JF., Zhou, JT. Strategies and Methods for Cloud Migration. Int. J. Autom. Comput.11, 143–152 (2014).
[24] Jamshidi, P., Pahl, C., & Mendonça, N. C. (2017). Pattern-based multi-cloud architecture migration. Software: Practice and Experience, 47(9), 1159-1184.
[25] Ke Wang, Migration strategy of cloud collaborative computing for delay-sensitive industrial IoT applications in the context of intelligent manufacturing, Computer Communications, Volume 150, 2020, Pages 413-420, ISSN 0140-3664, https://doi.org/10.1016/j.comcom.2019.12.014.
[26] Pahl, C., Xiong, H., & Walshe, R. (2013, September). A comparison of on-premise to cloud migration approaches. In European Conference on Service-Oriented and Cloud Computing (pp. 212-226). Springer, Berlin, Heidelberg.
[27] Casanova, H., Pandey, S., Oeth, J., Tanaka, R., “Wrench: A framework for simulating workflow management systems,” 2018 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS) (pp. 74-85).
[28] Varia, J. (2010). Migrating your existing applications to the aws cloud. A Phase-driven Approach to Cloud Migration, 1-23.
[29] Boss, G., Malladi, P., Quan, D., Legregni, L., & Hall, H. (2007). Cloud computing. *IBM white paper*, *321*, 224-231.
[30] Warner, T. L. (2020). *Microsoft Azure For Dummies*. John Wiley & Sons.
[31] Jamshidi, P., Pahl, C., & Mendonça, N. C. (2017). Pattern-based multi-cloud architecture migration. Software: Practice and Experience, 47(9), 1159-1184.
[32] Google Cloud. CIO’s Guide to Application Migration, Google Cloud Whitepaper, https://services.google.com/fh/files/misc/cio_guide_to_ application_migraton.pdf.
[33] Microsoft. Cloud Adoption Framework. https://docs.microsoft.com/enus/azure/cloud-adoptionframework/opbuildpdf/toc.pdf.
[34] Microsoft. Cloud Migration Essentials. https : / / azure . microsoft . com / en-us/resources/cloud-migration- essentials- e- book/.
[35] Mirsaeid Hosseini Shirvani, Amir Masoud Rahmani, Amir Sahafi, “A survey study on virtual machine migration and server consolidation techniques in DVFS-enabled cloud datacenter: Taxonomy and challenges,” Journal of King Saud University - Computer and Information Sciences, Volume 32, Issue 3, 2020, Pages 267-286,
[36] Khattar, N., Singh, J. & Sidhu, J., “Multi-criteria-Based EnergyEfficient Framework for VM Placement in Cloud Data Centers. ,” Arab J Sci Eng 44, 9455–9469 (2019). J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.
[37] Shuaib, M. Samad, Abdus Jacobs, and Alam, Shadab, “Why Adopting Cloud Is Still a Challenge?— A Review on Issues and Challenges for Cloud Migration in Organizations,” Ambient Communications and Computer Systems, Springer Singapore, 2019, pp. 387-399.
[38] Amin, R., Vadlamudi, S., & Rahaman, M. M. “Opportunities and Challenges of Data Migration in Cloud. Engineering International, ” 9(1), 41-50. R. (2021).

[39] N. Ahmad, Q. N. Naveed, and N. Hoda, "Strategy and procedures for Migration to the Cloud Computing, " 2018 IEEE 5th International Conference on Engineering Technologies and Applied Sciences (ICES), 2018, pp. 1-5, doi: 10.1109/ICETAS.2018.8629101.

[40] Rosati, P., Fowley, F., Pahl, C., Taibi, D., & Lynn, T. (2018, January). "Making the Cloud Work for Software Producers: Linking Architecture, Operating Cost and Revenue. " In CLOSER (pp. 364375).

[41] Rosati, P., Fowley, F., Pahl, C., Taibi, D., & Lynn, T. (2018, March). "Right scaling for right pricing: A case study on the total cost of ownership measurement for cloud migration. " In International Conference on Cloud Computing and Services Science (pp. 190-214). Springer, Cham.

[42] Khajeh-Hosseini, A., Sommerville, I., Bogaerts, J., & Teregowda, P. "Decision support tools for cloud migration in the enterprise. " In 2011 IEEE 4th International Conference on Cloud Computing (pp. 541-548). IEEE. (2011, July).

[43] F. Fittkau, S. Frey and W. Hasselbring, "CDOSim: Simulating cloud deployment options for software migration support," *2012 IEEE 6th International Workshop on the Maintenance and Evolution of Service-Oriented and Cloud-Based Systems (MESOCA)*, 2012, pp. 37-46, doi: 10.1109/MESOCA.2012.6392599.

[44] Bergmayr, A., Breitenbücher, U., Ferry, N., Rossini, A., Solberg, A., Wimmer, M., ... & Leymann, F. (2018). A systematic review of cloud modeling languages. *ACM Computing Surveys (CSUR)*, *51*(1), 1-38.

[45] Zhao, W., Peng, Y., Xie, F., & Dai, Z. (2012, November). Modeling and simulation of cloud computing: A review. In *2012 IEEE Asia Pacific cloud computing congress (APCloudCC)* (pp. 20-24). IEEE.

[46] Bilel, B. R., Navid, N., Raymond, K., & Christian, B. (2011, October). OpenAirInterface large-scale wireless emulation platform and methodology. In *Proceedings of the 6th ACM workshop on Performance monitoring and measurement of heterogeneous wireless and wired networks* (pp. 109-112).

[47] Gupta, S. G., Ghonge, M. M., Thakare, P. D., & Jawandhiya, P. M. (2013). Open-source network simulation tools: An overview. *International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)*, *2*(4), 16291635.

[48] Siraj, S., Gupta, A., & Badgujar, R. (2012). Network simulation tools survey. *International Journal of Advanced Research in Computer and Communication Engineering*, *1*(4), 199-206.

[49] Riley, G. F., & Henderson, T. R. (2010). The ns-3 network simulator. In *Modeling and tools for network simulation* (pp. 15-34). Springer, Berlin, Heidelberg.

[50] Kodali, R. K., & Kirti, B. (2020, October). NS-3 Model of an IoT network. In *2020 IEEE 5th International Conference on Computing Communication and Automation (ICCCA)* (pp. 699-702). IEEE.

[51] Calheiros, R. N., Ranjan, R., De Rose, C. A., & Buyya, R. (2009). Cloudsim: A novel framework for modeling and simulation of cloud computing infrastructures and services. arXiv preprint arXiv:0903.2525.

[52] Silva Filho, M. C., Oliveira, R. L., Monteiro, C. C., Inácio, P. R., & Freire, M. M. (2017, May). CloudSim Plus: A cloud computing simulation framework pursuing software engineering principles for improved modularity, extensibility and correctness. In *2017 IFIP/IEEE symposium on integrated network and service management (IM)* (pp. 400-406). IEEE.

[53] Casanova, H. (2001, May). Simgrid: A toolkit for the simulation of application scheduling. In *Proceedings First IEEE/ACM International Symposium on Cluster Computing and the Grid* (pp. 430-437). IEEE.

[54] Casanova, H., Giersch, A., Legrand, A., Quinson, M., & Suter, F. (2013). Simgrid: a sustained effort for the versatile simulation of large scale distributed systems. *arXiv preprint arXiv:1309.1630*.

[55] Suter, F. (2022, June). Performance Evaluation Through Simulation with SimGrid. In *Proceedings of the 2nd Workshop on Performance EngineeRing, Modelling, Analysis, and VisualizatiOn Strategy* (pp. 1-1).

[56] Camus, B., Dufossé, F., Blavette, A., Quinson, M., & Orgerie, A. C. (2018, September). Network-aware energy-efficient virtual machine management in distributed Cloud infrastructures with on-site photovoltaic production. In *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)* (pp. 86-92). IEEE.

## AUTHORS

**Bilel Ben Romdhanne** is principal research scientist at Amadeus since 2018. His research interests focus on network and HPC emerging technologies. It includes blockchain, cloud computing and distributed computing. He holds a Master's degree in Networks and a PhD in science computer.

**Mourad Boudia** is head of AI and emerging technologies department at Amadeus. His research has focused on AI, operations research and emerging technologies applications for travel and airline industry. He holds a Master's degree in Industrial Engineering and a PhD in operations research and applied mathematics.

**Nicolas Bondoux** is a lead research scientist at Amadeus in the AI Research team. He graduated from ENSIIE engineer school in Évry, France and holds a MSc in mathematics and computer sciences from the University of Manchester. His current main research interests are artificial intelligence and hardware acceleration.