Towards IT Workload Hybrid-Cloud Placement Advisory in Enterprise

André Hardt, Abdulrahman Nahhas, Hendrik Müller and Klaus Turowski
Faculty of Computer Science, Otto von Guericke University, Magdeburg, Germany
{*andre.hardt, abdulrahman.nahhas, hendrik.mueller, klaus.turowski*}@*ovgu.de*

Keywords: Enterprise Applications, SAP, Cloud Placement, BPMN.

Abstract: Placement of IT workloads in a cloud or hybrid-cloud environment is not always straightforward and requires taking into account various requirements, cloud offering capabilities, and costs. This fact has led researchers and industry practitioners to develop various automation solutions to support this decision process. However, the exact procedure for applying these solutions in practice, especially in the enterprise environment, is typically not discussed. In this work, we propose a formalized systematic business-centric process for delivering a service that relies on a data-driven automation solution, as a tool for experts, for relevant data management and placement optimization in a hybrid-cloud. We performed preliminary field testing of the proposed approach on real-world enterprise IT landscapes running SAP enterprise applications with the application of a user-friendly placement optimization automation solution. Finally, the stakeholder feedback and key takeaways from the field testing are summarized, noting the feasibility and potential usefulness of the presented formalized process.

1 INTRODUCTION

Selecting cloud services and workload placement, including Enterprise IT workloads, in a public cloud and hybrid-cloud environments can be a daunting task. And when a cloud strategy is unsuccessful, it can lead to significant negative outcomes (Venkatraman and Arend, 2022). Tackling such challenges can be difficult, which, as discussed further in section 2, has led researchers and practitioners to propose and develop various tools that can assist in relevant decision-making steps.

In this work, we aim to propose a business process model derived from empirical observations and in close collaboration with the industry experts. It relies on an automation tool designed to support decision-makers in selecting optimal placement for their enterprise IT workloads. Specifically, we focus on the placement of the standard enterprise applications (EA) in a hybrid-cloud environment. Our proposed business process model incorporates the interaction of non-technical decision-makers, domain experts, and the placement recommendation tool.

Enterprise applications placement and selecting infrastructure for these are complex tasks that require considerable expertise and insight into business processes to make a sound placement decision. Therefore, we suggest that the centralization of this expertise backed by an automated solution can lead to useful outcomes. Furthermore, we argue that formalizing a process of applying such a solution in an enterprise environment may help to bridge the gap between academic research and industry, as it has the potential to propose a clear procedure for applying algorithmic-based (i.e., heuristic, metaheuristic, machine learning) automation solutions in real-world enterprises.

2 RELATED WORK

A number of tools were proposed over the years to assist decision-makers in selecting the optimal placement for various systems in the cloud. Dubbed CloudAdvisor is such an example tool (Sahu et al., 2024) that consists of a web frontend and a backend. The rich web frontend is used to visualize and present possible placement options in the public cloud, while the backend relies on a combinatorial optimization method to compute possible placements. Additionally, the collection of user-defined requirements and constraints is done via the front as well. The combinatorial optimization is focused on selecting the placement option for the user's data in the cloud so that access latency is minimized. The stated target user base of the tool is the consultants in the area of cloud resource management.

CloudRecoMan (Mettouris et al., 2022) is another

Hardt, A., Nahhas, A., Müller, H. and Turowski, K.
Towards IT Workload Hybrid-Cloud Placement Advisory in Enterprise.
DOI: 10.5220/0013361800003929

In *Proceedings of the 27th International Conference on Enterprise Information Systems (ICEIS 2025) - Volume 1*, pages 830-839
ISBN: 978-989-758-749-8; ISSN: 2184-4992

similar tool, similarly presenting a web-frontend and a backend with a recommender system. Within the proposed methodology on which CloudRecoMan is based, the recommendation for cloud placement selection is made based on the company profile entered into the system by a decision-maker who does not necessarily possess sufficient technical background.

Another company-oriented approach was proposed earlier than the two aforementioned ones and based on a muti-criteria-decision-method. The approach is named $(MC^2)^2$ (Menzel et al., 2013), with a web-frontend and a backend prototype architecture. The user is required to supplement placement alternatives, which can either be determined by the experts and supported by integrating external data sources. The user must also supply criteria and requirements, in the specific format, that will be used for ranking the placement alternatives. Certain aspects of cloud alternatives, such as performance, are evaluated by integrating external measurement benchmarks.

Another tool named CloudAdvisor, with no relation to the aforementioned tool of a similar name, was proposed as a service-oriented solution (Jung et al., 2013) for recommending cloud placement configuration based on the application workload. Similarly to all aforementioned solutions, the proposed solution has a frontend view for the users as well as the backend where the necessary calculations are performed, including a combinatorial search of suitable alternative placement solutions. The solution is corporate-oriented and focused on VM placement. The user must specify the workload information via the frontend, as well as the performance, energy consumption, and budget information. After that, the user is presented with a set of alternatives and their costs for cloud placement for the given workload.

All of the aforementioned examples of workload placement selection methodologies and tools follow different approaches but generally present the tool that is designed to assist the decision-makers using algorithmic solutions. However, most of these require both business and technical understanding to operate. The only exception is CloudRecoMan (Mettouris et al., 2022), which is a recommendation tool based on the company's business profile, rather than details of the IT workload.

The delivery model of the solutions discussed above presumes end-user-oriented use, where the user interacts with the solution via a graphical frontend. However, the exact business process of delivery and application for these automation tools or services in an enterprise environment is not discussed in detail for any of the proposed methods. In this work, we aim to lay the foundation for bridging this gap.

3 PROPOSED PROCESS

In this work, we propose a formalization of a business process for delivering an advisory service focused on optimization of enterprise application placement in a hybrid-cloud. It is aimed at larger enterprises. We argue that automation solutions, such as discussed in section 2, can be used to support the experts by automating the tedious tasks of data management, cost estimation, placement optimization calculation. However, it is important to clearly understand how such solutions can be applied in real enterprises, where various stakeholders and domain experts must be involved to achieve the goal of hybrid-cloud adoption or sizing for existing IT infrastructure. We believe that formalization of the application might positively affect the adoption of automation solutions based on state-of-the-art algorithms to support experts in enterprise environments.

To the best of our knowledge, there is no published literature proposing a business-centric formalization of a process for providing an advisory service of enterprise IT systems placement optimization in a hybrid-cloud, while relying on a data-driven automation optimization system for supporting the experts in the field. While there are proposed tools, there is a lack of discussion on how precisely these can be applied in a real enterprise environment. Such environments typically involve various experts and stakeholders, which might benefit from being supported by a data-driven automation solution for placement optimization. The proposed formalization is the result of collaboration with field experts and empirical field testing.

The focus of the proposed process lies in the involved business activities. Therefore, we present a high-level overview of this process as a Business Process Model and Notation (BPMN) (OMG, 2010) model in Figure 1.

As seen from the aforementioned figure, the process encompasses two main actors. First is the *hybrid-cloud Placement Optimization Service* (HC-POS), which in turn encompasses the placement experts as well as the automation system for placement selection optimization. This service relies heavily on the automation provided by an Optimization Automation System (OAS), which is discussed in more detail further in subsection 3.2. The OAS is operated by the domain experts, whose tasks and roles are discussed in subsection 3.3.

The experts interact not only with OAS but also communicate with the second actor of the proposed process. This actor is the *organization*, which is a consumer of the provided service and the final recipient of the IT workload placement recommendation for the IT landscape. In the business environment, it can be seen as essentially a first actor's customer according to the proposed process. The assumptions and tasks of the organization are discussed further in subsection 3.1.

3.1 Organization

As mentioned before, under the term *organization*, in our proposed process, we understand the final recipient of the service, or in other words, the customer. As such, the stakeholders within the organization initiate the entire process. Specifically, it is expected that the organization must state a goal for its IT landscape transformation and (hybrid) cloud adoption or resizing. The expectation is that the organization involves the relevant stakeholders on the business site in the process. It is worth noting that in a sufficiently large enterprise, the *organization* (customer) might as well be a subdivision instead of an external service consumer.

The proposed process is oriented on IT landscape transformations in organizations that already have running IT landscapes. The initiation of such IT landscape transformation can be driven by a variety of business goals, which, in the end, directly influence the functional and non-functional requirements, as well as the outcome of such a transformation.

The primary motivation for designing this process specifically for the situation where the IT landscape already exists, instead of providing a "green field" [1] assessment, is that the existing infrastructure can be measured. Specifically, we rely on the data acquired by recording performance counters reflecting the system performance over a pre-determined, representative set of time. These performance counters include the capacities and consumption of the available computational resources (i.e., CPU, main memory, storage, network) as well as the application-specific measurements (e.g., number of transactions, number of users, etc.).

These measurements allow us to construct a representative workload profile. Such a profile reflects the capacity needs of the existing IT landscape within the real-world load, reflecting the real business needs of the organization. Furthermore, these measurements might reveal capacity deficiencies of the existing infrastructure, which might hinder the performance of the systems within the IT landscape and lead to negative consequences for the dependent business processes. This data and insights derived from it can serve as a basis for decision-making in an IT landscape hybrid-cloud transformation or resizing, as it reflects the existing business processes and needs in terms of computational resources.

[1] Planning and implementation of an IT infrastructure from ground up without existing legacy systems.

3.2 Optimization Automation System

One of the main assumptions of our proposed process is the existence of an *optimization automation system* (OAS). Such a system is a ready-to-use solution that is accessible to its users by means of a graphical user interface (GUI) and requires no technical or software engineering knowledge from the end users. Typically, such OAS is a solution consisting of frontend and backend components.

The backend component consists of the databases, data processing routines, and business rules implemented in the logic of a chosen programming language, as well as possibly complex computational algorithms (heuristic, metaheuristic, or machine learning). The backend can also provide various application programming interfaces (API), for integration with other information systems. However, the complexity or intertwined components that are critical to the functionality are obscured from the end users. The internal or external IT teams carry out the hosting and maintenance.

The frontend is the component that provides the user-friendly GUI and is the main layer of interaction with the OAS by the end users. Within our proposed process of overall service delivery, the GUI should be oriented toward the end users who are experts in the primarily business-oriented fields but with enough technical background to formulate inputs and interpret outputs of the OAS. Deep knowledge of the internal functionalities of the OAS and its components (e.g., internal optimization algorithms, cost models, requirements model) is not a prerequisite for the end users.

3.3 Cloud Placement Experts

As mentioned before, within our proposed process, under the term *experts*, we mean the cloud and the IT software experts on the side of the service provider. The specific focus of our service delivery process is the placement of the standard off-the-shelf IT applications in the hybrid-cloud infrastructure. Therefore, the typical cloud infrastructure knowledge of the expert users will lie in the business-process-focused domains surrounding the target enterprise application (EA). Such areas include cloud services and architectures, capacity sizing of the server solutions, licensing issues, operational costs, and interpretation of busi-

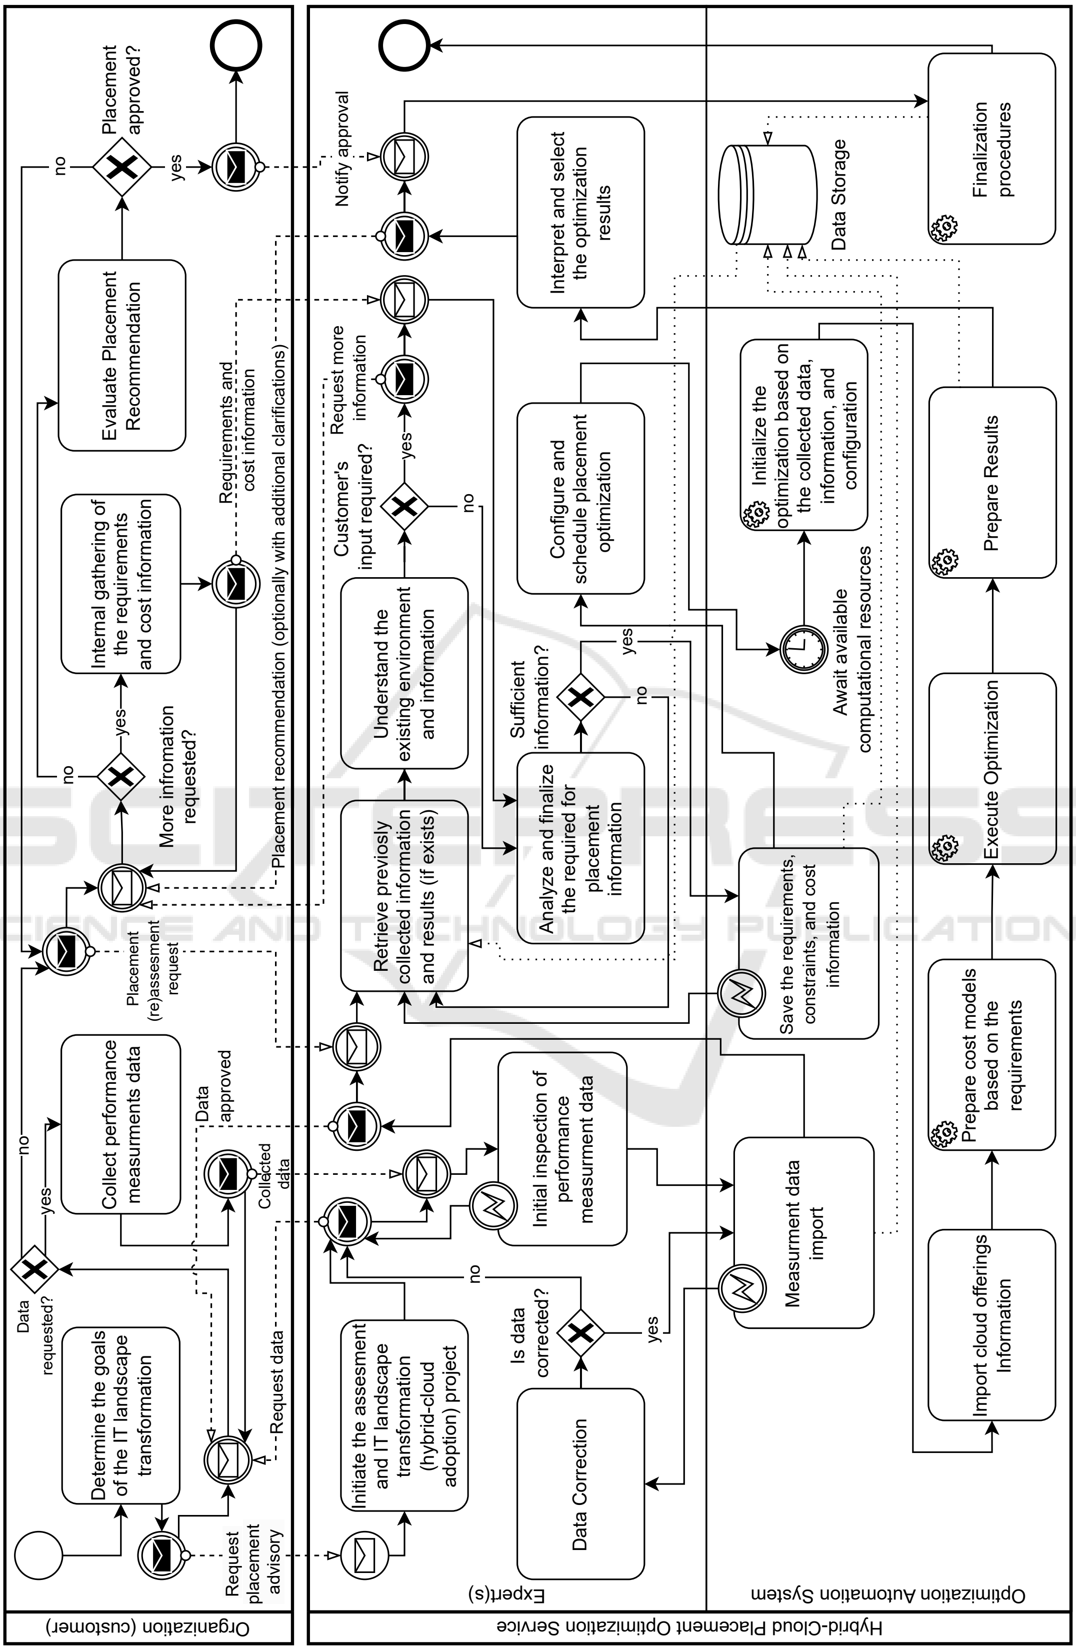

Figure 1: BPMN model of the service delivery process.

ness requirements. These experts are typically not familiar with state-of-the-art algorithmic solutions, so the goal of AOS is to provide these as tools for the experts.

3.4 Process Flow

As depicted in Figure 1, the business process proposed in our process of the IT landscape hybrid-cloud transformation starts and ends with the consumer of the service (i.e., the customer organization). The process ends with the organization's stakeholders receiving and approving a new IT landscape placement recommendation. If the approval can't be achieved, part of the entire described process is restarted, hence introducing a partially iterative aspect.

3.4.1 Assessment of the Existing IT Landscape

The organization initiates the process and determines the intent and goals for the planned IT landscape transformation project. After that, the organization engages HCPOS and hands over the stated goals. The HCPOS experts process these goals and may assist the organization in collecting the required measurement data from the running IT landscape.

Suppose the company already has collected historical data describing the IT landscape workload profile in sufficient detail and length. In that case, this data can be used directly and should be sent to the experts. If such a dataset does not yet exist, the process of collecting measurements is conducted according to the customer's requirements as well as the stated goals. It ranges from multiple weeks to multiple months. The goal is the construction of a reliable workload profile for the existing IT landscape. After the collection period has elapsed, the gathered data is sent to the HCPOS experts for initial pre-processing.

The experts then import the data into the OAS, where the collected data is further automatically processed, cleansed, and analyzed. However, at this stage, the OAS can indicate that the data is not suitable for further processing according to various rules and checks (e.g., erratic workload profile, too many missing values, errors). If that occurs, the expert can attempt to correct the data or initiate further data collection.

When OAS determines that a sufficient amount of data with acceptable quality is achieved, the processed data is stored in internal data storage for reuse. Data collection is a time-consuming procedure. Thus, we assume the data collection is performed only once within the frame of the same IT landscape transformation project. The following part of the proposed process is, however, iterative.

3.4.2 Requirements Collections

In the following activities, the expert can access the collected data via the frontend provided by OAS. The presentation of the data must be sufficient for the experts to efficiently understand and analyze it in order to fulfill their further activities. These activities include the collection of the function and non-functional requirements, as well as constraints, for the IT landscape and its components subject to the transformation or resizing project. The requirements are collected by interacting with the stakeholders on the organization's side in a format required by the OAS for automatic processing. Furthermore, for hybrid-cloud projects, the cost information related to the customer facilities (e.g., private data centers) must also be estimated.

Calculation of running costs for the components of the IT landscape in the private data center of the customer is difficult to estimate outside of the organization itself (Greenberg et al., 2009; Kashef and Altmann, 2012; Altmann and Kashef, 2014; Brogi et al., 2019) and, therefore, must be assisted by the organization. The role of the expert in this case is to guide the customer's representatives in this information collection process according to their own expertise in the domain.

This information collection process can be done in an iterative way according to the expert's judgment. After the expert concludes that a sufficient amount of information is collected, it is imported to the data storage of OAS for future reuse and associated with the previously collected measurement data.

3.4.3 Automated Processing

In the next activity, the expert hands over control of the process to the automation procedures and business rules encoded into the OAS. Within this activity, any additional configuration parameters might be given by the expert (e.g., select a specific preset). It is done via the GUI provided by the OAS's frontend. When it's done, a task is scheduled to process the collected measurement data with the collected requirements and customer-specific cost information.

The exact technical implementation of the automated placement recommendation processing can vary from one use case to another, however a few key activities, that are specifically relevant to our proposed service delivery process are highlighted.

Firstly, since the IT landscape transformation project involves at least one public cloud provider as a target for the placement, offerings of these cloud placements must be acquired for processing and decision-making. Depending on the provider, it can be done either by importing specific price lists or accessing provider-specific APIs. Either way, the goal of this activity is to acquire sufficient information about the cloud offerings with technical and pricing information, such that decision-making is possible by the algorithm.

The amount of cloud-specific information and composition of the offerings can vary depending on the IT landscape, as well as requirements and constraints. For example, if one of the requirements is that one of the systems in the IT landscape requires a particularly high availability level, this would require the retrieval of a set of high-availability offerings, or composition of such offerings (Salapura and Mahindru, 2016), suitable for the specific enterprise system. The decision of which offerings are required according to the specific requirements is derived from the business rules encoded into the automation processes of OAS and will differ from use case to use case, from one enterprise IT system vendor to another.

In the next activity, the OAS must prepare sufficient cost models for the comparison purposes of various placement compositions of the IT landscape in a hybrid-cloud environment. Expenses associated with running specific components of the IT landscape in the private data centers of the organization are reused from the information collected by the expect from the organization stakeholders. However, the cost estimation for the public cloud can also be daunting because of the wide variety of pricing models available, even for the same cloud service (Wu et al., 2019). Therefore, it is imperative that the encoded business rules of cost estimation for all cloud providers are use-case-specific and validated prior to the rollout of the OAS.

The next activity executed by OAS is actual decision-making by employing a state-of-the-art algorithm (e.g., metaheuristic, machine learning). The algorithm and its implementation will strongly depend on the target use case. The central goal of this activity is to execute an automation solution that seeks viable IT landscape placement solutions and selects the most suitable one according to the encoded business rules, technical specifications, data, customer requirements, and constraints. It's a typical case of combinatorial optimization or multi-criteria-decision-methods applications, depending on the complexity of the problem, similar to many solutions discussed in section 2.

The next and final activity of OAS within a single iteration is responsible for the processing of the output generated by the selected algorithm in the preceding activity. This step is critical because, depending on the type, the optimization algorithm might generate a large volume of data, which must be translated into human-readable form. In this case, it has to be an interpretation of the algorithm output to specific placement solutions of the given IT landscape according to the requirements, constraints, and measured data. Therefore, it is imperative that the output contains the exact specifics of the placement compositions for the cloud solutions with the pricing information.

Furthermore, the prepared results contain not only a sufficient amount of information but are also presented in a form that is understandable not only to the expert but also to the stakeholders in the customer organization. This typically means providing the expert with the interactive view in the frontend of the OAS. Additionally, it is prudent to automate the generation of the reporting material for dispatching to the stakeholders in a format that the organizations in the given domain typically accept. Automation of the executive-level visualizations, specifically with a concise explanation of the decision-making process made by the employed algorithm, can have a positive effect on the final approval stage (Dimara et al., 2022). The automation of generating such materials can considerably reduce the efforts and time required by the experts in the following activity focused on preparing these for dispatching to the customer organization for approval.

3.4.4 Iteration Finalization

Finally, the prepared results and recommendations for placement of the IT landscape are sent back to the organization for approval, optionally supplemented with additional proposals or clarifications from the experts. If the results of the process and the placement recommendation are found to fit the requirements of the stakeholders, the process of our proposed process ends. If the approved recommendation is accepted, the experts are notified, and the project is finalized appropriately. This might include the archival or even removal of the customer's data from the OAS to ensure the security of the customer's infrastructure. At this point, it is assumed that the organization can enact the proposed IT transformation plan.

However, if the final results are not approved, the process is returned to the expert, who attempts to adjust the requirements, constraints, and optimization settings while taking into account the feedback from the previous interaction of the process execution. The expert can also choose to request further input from the customer. This additional information from the company stakeholders is further supplemented to the data storage for OAS. This iterative process continues until the final refinements produce the results that are approved by the customer organization.

4 PROOF OF CONCEPT

We empirically test the overall feasibility of the proposed process in a real-world environment by providing advisory support to stakeholders in the organizations operating standard IT enterprise applications. The goal is to verify the proposed process's overall usefulness and applicability in the real-world enterprise environment. The field testing was based on production SAP enterprise landscapes with the goal of assessing the viability of the hybrid-cloud transformation of the existing IT landscape.

Selection of the placement configuration in a hybrid-cloud for SAP landscapes can be a complex task and subject to a variety of considerations and constraints (Berhorst et al., 2021). In our field test, we focused on the Infrastructure-as-a-Service (IaaS) (Badger et al., 2012) options of the public cloud providers as placement options, as well as private (on-premise) data centers of the IT landscape owner. For the cloud placement options in the hybrid environment in our field test we relied on Azure (Bögelsack et al., 2022), their SAP-certified VMs, required networking and storage infrastructure, and APIs to automate offering selection and costs estimation.

It is shown in previously published work that this task can be solved by using metaheuristic algorithms (Kharitonov et al., 2023), which also served as a basis for the field testing discussed in this work. However, in principle, any suitable optimization mechanism can be applied. Therefore, the specifics of optimization are beyond the scope of this work, as we concentrate on how these can be actually used in a business environment. For the purposes of this proof of concept, the aforementioned metaheuristic-based approach was wrapped into a specially developed and hosted containerized prototype solution, consisting of frontend (web-UI) and backend (relational database, queues, various workers), as discussed in section 3.2.

The performance metrics data collection from a real-world SAP IT landscape was carried out using a software solution that integrated with the SAP systems via the standard APIs and did not impact the overall performance of the system. The data collected contains the performance counters collected as time series over a number of weeks. The performance counters include both the utilization and capacity standard metrics (e.g., CPU, RAM, network, storage), as well as the SAP-specific metrics (e.g., number of transactions, SAPS(Marquard and Götz, 2008)).

The goal is to take advantage of the data and prevent overprovisioning by selecting the placement options according to the real workload profile. Choosing the right size of the target cloud environment can significantly reduce future operational costs (Aloysius et al., 2023) of an SAP system.

The collection of the requirements and constraints was carried out by a team of domain experts with close cooperation with the owners (organization) of the SAP IT landscape. The collection and processing of the requirements are supported by the frontend of our prototype OAS, thus ensuring the correctness and consistency of the collected values.

4.1 Feedback

At the finalization phase of the field test, we have collected feedback from the participating stakeholders. It is important to note that in this work, we focus on the feedback related to the overall business process validity and usefulness, while the results of the placement optimization itself are consistent with the work we relied upon to develop the proof of concept prototype (Kharitonov et al., 2023).

Two real-world companies participated in the field test. Company A is a manufacturing business that maintains its own SAP IT landscape for internal use. Company B is a service provider with a higher degree of internal competence in the field of public cloud.

The management-level stakeholder, in the role of the head of the SAP infrastructure department in company A, declared that the primary objective of the proof of concept is to acquire fact-based executive-level answers about cloud adoption within the given business realities. The aforementioned stakeholder of company A remained satisfied with the assessment results achieved in the proposed business process section 3.

The technical-level personnel in the same department demonstrated initial skepticism regarding the application of the proposed process and voiced doubts it would provide useful insights beyond what is known already. However, the final feedback was positive and concentrated on the number of relevant insights presented as a result, underlying the usefulness of the proposed formalized process for integrating state-of-the-art algorithmic approaches to real-world enterprise processes. Specifically, it was noted that the presented cost-driving factors are hard to calculate just within the day-to-day department activity.

Within Company B, two groups of stakeholders can be viewed according to their expectations. First are the stakeholders in the field of strategy and technological transformation, including strategic management in the fields of data centers, cloud, and infrastructure. These initially expressed an interest in the validation of their own efficiency, costs, and innovation advantages in the field of cloud adoption.

Second is the technology-oriented executive and a technology expert. They expressed an interest in using the tool within our proposed process to acquire additional insights and expertise. This expertise would allow them to decide on offering cloud solutions to their own customers if such makes sense in the particular business environment.

As the feedback from the stakeholders of company B, it was noted that the use of an automated solution within the proposed process provides valuable insights and information for the management-level discussions. It was noted that the final approved results of the executed process are detailed and correspond to the expectations of the stakeholders. The transparency of the analysis done using the proposed process was also noted positively. The interaction between the stakeholders and the placement experts within the meetings is also positively noted as an opportunity for the experts to provide detailed explanations for the results and the decision process provided by AOS.

The last point is specifically notable as a positive aspect of relying on the proposed process instead of simply providing the AOS to the stakeholders as a self-service. The value of the data and information provided by the owners of the IT landscape can be truly put to use by experts in the field of IT workload placement. As we initially hypothesized, this experience is complemented by the specific business-domain knowledge of the stakeholders within a controlled process of gathering, processing, storing the relevant data and information, and finally providing a relevant, useful result.

In the summary of the feedback, it can be said that an application of an automation tool for the assessment of real-world IT landscapes and evaluation of the cloud adoption possibilities is useful when done in a systematic manner. The systematic collection of the functional and non-functional requirements and the use of data-driven automation can result in providing the stakeholders with the expected information that can further serve as a basis for management-level decision-making.

Delivery of the information relevant for the specific stakeholders, instead of simply relying on the ability of the automation solution to deliver some data, is the main advantage of relying on a well-defined expert and automation-supported process of collecting and refining the requirements, as well as evaluating the results. This is the main advantage of relying on a systematic process for collecting relevant information from the customer as part of the consultancy service, instead of simply relying on the data and the automation as a self-service.

4.2 Lessons Learned

During the field testing of the proposed process, a few key takeaways can be formulated based on the feedback received from the owners of the IT landscape (decision-makers) and the experts. The takeaways concern both the functionality of the prototype OAS, as well as the formulation and presentation of the placement recommendation results.

4.2.1 Intelligent Results Presentation

Within our field testing, we relied on evolutionary combinatorial optimization to achieve the optimal recommendation within the OAS. This type of optimization evaluates a vast number of possible placement combinations for the IT landscape and, in the end, presents the best-selected combinations according to the given requirements. While it might be tempting to simply present only the best solution to the customer, this turned out to be an insufficient approach as the decision-makers displayed interest in comparisons against suboptimal solutions to better understand the influence of their requirements and constraints on the recommendation.

However, providing too many such options for examination via the visualization and data generated by OAS was noted to be overwhelming for the decision-makers. At the same time, a simple application of the so-called Hall of Fame (HoF) strategy, where only a certain number of the top best solutions are selected, is displayed, resulting in the critique of the lack of sufficient diversity of the solutions. The solutions that approach the optimum tend to have only slight differences between them.

Therefore, it was noted that a more intelligent approach to selecting the final recommendations and comparison solutions is needed. In our specific field test, such an approach, that was deemed successful by the decision-makers, was to select the best possible solution, and, as a comparison, data points provide the evaluated solutions that place the entire IT landscape within the same location or provider. Such presentation provides an opportunity to clearly demonstrate various trade-offs (e.g., degree of requirement satisfaction, costs, constraints, etc.) between simpler strategies, where the entire landscape is concentrated within the same location, and more complex hybrid-placement solutions. Further sorting and filtering capabilities for navigating discovered solutions are also deemed beneficial for presentation and exploration, including for future requirements or constraints tuning.

4.2.2 Complexity of Requirements Processing

Within our field test, we have concentrated on the specific use case of SAP systems placement within a hybrid-cloud environment. While there are a number of well-understood best practices and specific rules presented by the SAP developers, direct conversion of the business-specific requirements into technical details of cloud deployment architectures is far from trivial and requires a sufficient amount of initial investment from the experts to construct such rules for translating requirements and constraints into the selection of specific combinations of cloud products.

While this initial investment pays off in reducing the workload of the experts working with OAS, the participating domain experts noted that a degree of flexibility is required in the encoded rules. The reason is that public cloud offerings do not stay static. New cloud products are introduced, while older ones might get deprecated. Therefore, changes in the placement recommendation selection rules might require frequent updates.

4.2.3 Customer-Specific Pricing

When selecting a public cloud location for placement of the IT landscape, relying solely on the list prices for all cloud products might not be sufficient. Specifically, corporate customers of the public cloud providers can negotiate discounts on certain cloud offerings (Deb and Choudhury, 2021).

However, details of such discounted contracts with the public cloud providers must be provided by the organization requiring the placement optimization recommendation and shared with the experts. This data can be sensitive and must be guaranteed to be protected from disclosure. Access to this information should be strictly controlled, even among the experts operating OAS. There must also be a mechanism to guarantee the removal of this information from OAS after the recommendation process is finished.

Furthermore, encoding the details of such discounted pricing can be tedious and require care. Any errors will result in incorrect price models and, consequently, placement recommendations. Therefore, this stage directly benefits from a specially made frontend GUI supporting the experts in this task.

The complexity of the customer-specific cost estimation extends to the on-premise costs as well. Within our field tests, we discovered that an accurate estimation of the private data center's running costs was not possible within the timeframe of the evaluation field test in one of the cases mentioned earlier. This led to potentially inaccurate placement recommendations generated by AOS. In such cases, the results have to be carefully processed by the placement experts before finally being sent to the stakeholders. The complexity of gathering customer-specific costs and estimating them should not be underestimated.

5 FUTURE WORK

Relying on SAP-based solutions in our field test allowed us to benefit from a significant degree of standardization and support from cloud providers (e.g., Azure (Bögelsack et al., 2022)), as well as an abundance of domain experts, who are vital for our proposed approach. That is typical for large enterprise solutions. However, that brings to question the generalization of such an approach. We intend to validate this approach with a similar solution aimed at smaller organizations (e.g., Odoo (Wu and Chen, 2020)).

Furthermore, we intend to investigate the feasibility and how beneficial it is to apply the proposed approach with non-IaaS placement models (e.g., RISE with SAP (Subrahmanyam, 2022)). This investigation might prompt a further adaptation of the proposed formalized process.

6 CONCLUSION

In this work, we propose a formalization based on BPMN of a systematic process for applying automation solutions to decision-making advisory in hybrid-cloud placement selection for enterprise IT workloads in a data-driven manner. The proposed process is business-centric and includes a systematic collection of the data, as well as the required information needed for the decision-making of the IT workload placement by experts collaborating with stakeholders. We aim to demonstrate how algorithmic-based automation can fit into real-world business processes.

The preliminary feasibility testing, in the form of an overall proof-of-concept evaluation, of the proposed approach is performed within controlled, real-world enterprise environments and involves stakeholders from two distinct companies. It is done as a field test with a specific use case of SAP enterprise applications placement in a hybrid-cloud infrastructure. For this purpose, an automated solution for placement optimization was used within the constraints of the proposed process. The evaluation results include the feedback received from the stakeholders. The feedback confirms the validity of the proposed approach for an enterprise environment, where trained experts rely on the automated solution and communicate with the stakeholders who own the specific IT landscape.

Several important takeaways from the field test are noted. These include the importance of presenting the results generated by the automation solution, which should be done without overwhelming the stakeholders. We also note the complexity of the initial formalization of use-case-specific requirements, such that these requirements could be supplied to the automation system. Additionally we note the particular complexity of estimating and accounting for the costs that are specific to the specific companies (e.g., cloud discount contracts, running costs of the private data center placements). This complexity can be detrimental to the validity of the final results. Finally, we note the importance of the appropriate visualization used in the final results presentations.

REFERENCES

Aloysius, G. C., Bhatia, A., and Tiwari, K. (2023). Stability and Availability Optimization of Distributed ERP Systems During Cloud Migration. *Lecture Notes in Networks and Systems*, 654 LNNS:343–354.

Altmann, J. and Kashef, M. M. (2014). Cost model based service placement in federated hybrid clouds. *Future Generation Computer Systems*, 41:79–90.

Badger, M., Grance, T., Patt-Corner, R., and Voas, J. (2012). Cloud Computing Synopsis and Recommendations. Technical Report NIST Special Publication (SP) 800-146, National Institute of Standards and Technology.

Berhorst, R., Dürk, M.-O., Hauzeneder, C., Sperzel, S., Wagner, P., and Zimmermann, R. (2021). DSAG-Leitfaden: Hybrid Operations. Technical report, Deutschsprachige SAP Anwendergruppe e.V.

Bögelsack, A., Chakraborty, U., Kumar, D., Rank, J., Tischbierek, J., and Wolz, E. (2022). SAP S/4HANA on Microsoft Azure – Concepts and Architecture. In Bögelsack, A., Chakraborty, U., Kumar, D., Rank, J., Tischbierek, J., and Wolz, E., editors, *SAP S/4HANA Systems in Hyperscaler Clouds: Deploying SAP S/4HANA in AWS, Google Cloud, and Azure*, pages 267–323. Apress, Berkeley, CA.

Brogi, A., Corradini, A., and Soldani, J. (2019). Estimating costs of multi-component enterprise applications. *Formal Aspects of Computing*, 31(4):421–451.

Deb, M. and Choudhury, A. (2021). Hybrid Cloud: A New Paradigm in Cloud Computing. In *Machine Learning Techniques and Analytics for Cloud Security*, chapter 1, pages 1–23. John Wiley & Sons, Ltd.

Dimara, E., Zhang, H., Tory, M., and Franconeri, S. (2022). The Unmet Data Visualization Needs of Decision Makers Within Organizations. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4101–4112.

Greenberg, A., Hamilton, J., Maltz, D. A., and Patel, P. (2009). The cost of a cloud: Research problems in data center networks. *SIGCOMM Comput. Commun. Rev.*, 39(1):68–73.

Jung, G., Mukherjee, T., Kunde, S., Kim, H., Sharma, N., and Goetz, F. (2013). CloudAdvisor: A recommendation-as-a-service platform for cloud configuration and pricing. *Proceedings - 2013 IEEE 9th World Congress on Services, SERVICES 2013*, pages 456–463.

Kashef, M. M. and Altmann, J. (2012). A Cost Model for Hybrid Clouds. In Vanmechelen, K., Altmann, J., and Rana, O. F., editors, *Economics of Grids, Clouds, Systems, and Services*, pages 46–60, Berlin, Heidelberg. Springer.

Kharitonov, A., Nahhas, A., Müller, H., and Turowski, K. (2023). Data driven meta-heuristic-assisted approach for placement of standard IT enterprise systems in hybrid-cloud. In *Proceedings of the 13th International Conference on Cloud Computing and Services Science - CLOSER*, pages 139–146. SciTePress / INSTICC.

Marquard, U. and Götz, C. (2008). SAP Standard Application Benchmarks - IT Benchmarks with a Business Focus. In Kounev, S., Gorton, I., and Sachs, K., editors, *Performance Evaluation: Metrics, Models and Benchmarks*, pages 4–8, Berlin, Heidelberg. Springer.

Menzel, M., Schönherr, M., and Tai, S. (2013). (MC2)2: Criteria, requirements and a software prototype for Cloud infrastructure decisions. *Software - Practice and Experience*, 43(11):1283–1297.

Mettouris, C., Vanezi, E., Zampas, G., Argyrou, S., Achilleos, A., Constantinides, A., and Papadopoulos, G. A. (2022). CloudRecoMan: Cloud adoption made easy. *ACM International Conference Proceeding Series*, pages 110–117.

OMG (2010). Business process model and notation (BPMN) version 2.0. Standard, Object Management Group (OMG).

Sahu, P., Roy, S., and Gharote, M. (2024). CloudAdvisor for Sustainable and Data Residency Compliant Data Placement in Multi-Cloud. In *2024 16th International Conference on COMmunication Systems & NETworkS (COMSNETS)*, pages 285–287.

Salapura, V. and Mahindru, R. (2016). Enabling enterprise-level workloads in the enterprise-class cloud. *IBM Journal of Research and Development*, 60(2-3):3:1–3:8.

Subrahmanyam, R. S. (2022). *An Overview of SAP S/4HANA*, pages 1–23. Apress, Berkeley, CA.

Venkatraman, A. and Arend, C. (2022). A Resilient, Efficient, and Adaptive Hybrid Cloud Fit for a Dynamic Digital Business. International Data Corporation (IDC). Report No.: IDC #EUR149741222.

Wu, C., Buyya, R., and Ramamohanarao, K. (2019). Cloud Pricing Models: Taxonomy, Survey, and Interdisciplinary Challenges. *ACM Comput. Surv.*, 52(6):108:1–108:36.

Wu, J. Y. and Chen, L. T. (2020). Odoo erp with business intelligence tool for a small-medium enterprise: A scenario case study. In *Proceedings of the 2020 11th International Conference on E-Education, E-Business, E-Management, and E-Learning*, IC4E ’20, page 323–327, New York, NY, USA. Association for Computing Machinery.