

Increasing the System Availability of Data Centers by Using Virtualization Technologies

Chris Ewe, Naoum Jamous, Holger Schrödl

Abstract—Like most entrepreneurs, data center operators pursue goals such as profit-maximization, improvement of the company's reputation or basically to exist on the market. Part of those aims is to guarantee a given quality of service. Quality characteristics are specified in a contract called the service level agreement. Central part of this agreement is non-functional properties of an IT service. The system availability is one of the most important properties as it will be shown in this paper. To comply with availability requirements, data center operators can use virtualization technologies. A clear model to assess the effect of virtualization functions on the parts of a data center in relation to the system availability is still missing. This paper aims to introduce a basic model that shows these connections, and consider if the identified effects are positive or negative. Thus, this work also points out possible disadvantages of the technology. In consequence, the paper shows opportunities as well as risks of data center virtualization in relation to system availability.

Keywords—Availability, cloud computing IT service, quality of service, service level agreement, virtualization.

I. INTRODUCTION

COMPANIES like Amazon or Google operate comprehensive server environments in order to provide IT services to their customers. These data centers enable for instance to operate a web page by using processing power of an infrastructure service [1], [2]. Companies like Dropbox, Airbnb or Netflix offer IT services themselves, but do not have their own infrastructure. Instead, they use an external provider like Amazon Web Services (AWS) with the Amazon Elastic Compute Cloud (EC2) [3]. AWS itself follows several goals, such as the maximization of their profit or the improvement of the company's reputation [3], [4]. In order to achieve these goals, the compliance of a given quality of service (QoS) to the customer is a key factor [5]-[7]. For this purpose, the service level agreement (SLA) determines inter alia functional- and non-functional properties, lease terms, pricing as well as regression rules [8]. Thus, the SLA contracts rights and obligations of both the service provider and the service customer. As the central non-functional property, this paper will identify the system availability because it is the prerequisite for all other properties of an IT service [1], [7], [9]-[11]. Amazon EC2, for instance, claims to guarantee an

availability of 99.95% per year [12]. If this target is not met, customers get a 10% discount [12]. Thus, the achieved service level regulates the pricing and possible damage claimed by the customer. The Amazon EC2 SLA, as an example, only specifies availability, which can therefore be considered as the most important non-functional property for Amazon [12].

The IT Infrastructure Library (ITIL) declares that the objective of the IT service area is to provide high quality service at the lowest possible price [5]. Therefore, IT service providers are faced with the challenge of guaranteeing the SLA, all the while consuming resources at a minimal cost. To meet this requirement they can use the concept of virtualization [1]. This technology provides software solutions from various vendors such as VMWare. They try to enhance inter alia the availability of a system with the facilitation of features like virtual machine (VM) migration or disaster recovery. To better understand the influence of different virtualization functions on the availability of a data center, this paper provides a first model. The aim of this model is to increase the understanding of the relation between features of virtualization and availability. Therefore, it visualizes the connections between the separate components. To illustrate this, the contribution refers to the example of cloud services. According to [1], virtualization is the engine that moves cloud computing. That is why users of the cloud obtain virtualized IT services from their provider. Regarding this point of view, data centers for cloud services, like AWS, are not able to operate without using virtualization technologies [1]. Thus, those services represent a vivid example of the potential of virtualization.

II. RESEARCH METHOD AND STRUCTURE

The aim of this paper is to find correlations between the data center virtualization and the availability of the system. Therefore, we show different functions of virtualization and their impact on the systems availability. In its first stage, the model describes the basic influences, but does not specify if the effects are positive or negative. However, we give approaches to consider the impacts of the effects. For example we point out potential problems. To fulfil this, the research approach of design science according to the principles established by [13] is used. In this context, the research aims to solve a problem by using pre-existing knowledge through the construction of an artifact. The new findings in turn expand the knowledge base and enable future research. This context is represented in Fig. 1.

C. Ewe is with the Chair of Business Informatics, Otto-von-Guericke University Magdeburg, Magdeburg, 39104 Germany (phone: +49 (0) 170 489 179 5; e-mail: chris.ewe@st.ovgu.de).

N. Jamous is with the Chair of Business Informatics, Otto-von-Guericke University Magdeburg, Magdeburg, 39104 Germany (e-mail: naoum.jamous@ovgu.de).

H. Schrödl is with the Chair of Business Informatics, Otto-von-Guericke University Magdeburg, Magdeburg, 39104 Germany.




Fig. 1 Design Science Research principal referring to [13]

The treated problem (illustrated in Fig. 1 by arrow 1) is the lack of simple processed information on the potentials of virtualization technologies for increasing the availability of data center IT services. By using the pre-existing knowledge (arrow 2), we create an artifact in the form of a model to solve this problem (arrow 3). The findings of the contribution in turn enhance the knowledge base and provide the basis for further research (arrow 4). The research question can be formulated as: With which functions can virtualization contribute to the improvement of the system availability of a data center?

For this purpose, we use the four-stage process of analysis, design, evaluation, and communication, imposed by [14]. The analysis is based on a literature review and a case study. It concludes by describing and isolating the problem. In the course of the literature research, we used the scientific databases of ACM Digital Library, IEEE Xplore, and Lecture Notes in Computer Science as well as research documents from the University of Magdeburg. Therefore, we used the keywords 'IT Service', 'Quality of Service', 'Service Level Agreement', 'Availability', 'Data Center', 'Virtualization', 'Cloud Services' and 'Cloud Computing'. Chapter III presents the research background to IT services, quality of service, and virtualized data centers. The step of the design in Chapter IV forms the core of the work and describes the creation of the artifact in form of the model. The evaluation of the work is shown in Chapter V. It was conducted by consulting a group of experts. They were asked to their opinion about the importance of service quality as a well as the benefit of the model. In Chapter VI a conclusion is drawn and further research needs are identified.

III. RESEARCH BACKGROUND

A. IT Services and the Example of Cloud Services

Through the development to outsourcing and outtasking of computing power, there is currently a shift from the far more technically oriented IT world into the direction of service orientation [7]. Reference [5] defines a service generally as "a means of delivering value to customers by facilitating outcomes customers want to achieve without the ownership of specific costs and risks". Therefore, [15] describes an IT service: "A Service provided to one or more Customers by an IT Service Provider. An IT Service is based on the use of Information Technology and supports the Customer's Business Processes. An IT Service is made up from a combination of people, processes and technology and should be defined in a Service Level Agreement." Another important definition can be found in [16]. He defines an IT service as a service of information technology, which includes consulting, planning, and provision of services, such as hardware and software. Reference [7] says that a business arises when the focus shifts to services as an interaction between the customer and a

service provider. Thus, services in their different forms, move over to become the focus of modern IT companies.

Reference [17] says that these IT services are offered by IT service providers. Furthermore they distinguish various types, which specify their exact field. One of those types is provider of cloud services [17]. On this occasion, three service models are distinguished: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) [8], [18]. A more detailed description of these services can be found in the National Institute of Standards and Technology (NIST) by [19]. The central message of the definitions by the NIST in the context of this work is that cloud computing offers virtualized resources that can be used by customers as needed [19]. A representative of those is the already mentioned AWS. Therewith, customers use e.g. Amazon EC2 as "a web service that provides resizable compute capacity in the cloud" [12]. The infrastructure required for this purpose is presented in the form of the AWS data centers. These include resources in the form of hardware (e.g. central processing unit (CPU), storage, networking), as well as software. According to [1], the cloud provides an easy way to access data and applications, and is a vivid example for virtualized IT services of a data center.

B. Service Quality as a Part of the Service Level Agreement

The ISO 20000 defines a service level agreement (SLA) as a "written agreement between a service provider and a customer that documents services and agreed service levels" [5]. Thus, the SLA creates clear business requirements in terms of quality and quantity of IT services for the service customer and the service provider [5], [7], [10]. This relationship is illustrated in a simplified form in Fig. 2.




Fig. 2 SLA between a provider and its customer based on [7]

The basic consensus is that the provider, in this case Amazon, provides its customer, here Dropbox, infrastructure, and that Dropbox pays for this in return. The central idea of the SLA is to ensure the agreed service quality [16]. Inter alia, the document includes conditions such as prices, contract terms, and rules that determine customers' claims in case of a breach of the agreement [8]. In addition to these conditions, the document contains functional and non-functional properties that also affect the QoS [8]. Reference [20] says that SLA "captures the mutual responsibilities of the provider of a service and its clients with respect to non-functional properties". Reference [21] also says that functionality is complemented by quality. Beside the availability of the system, which is described more detailed in the following chapters, response time, performance and ease of use are among the most important non-functional properties which decide if a high QoS is given [1], [5], [7], [9], [10].

C. Foundations of Virtualization

According to the dictionary, the term virtual can be

translated with ‘thought’ or ‘apparently’ [22]. In one of the first works on this topic, [23] defined: “A virtual machine is taken to be an efficient, isolated duplicate of the real machine.” Reference [24] notes that a VM is no longer necessarily a duplicate of the real machine, but that this definition is still valid in its basic features. In their paper, they expand the original definition from [23] and define virtualization as a tool for decoupling the services from physical hardware resources [23]. As this view captures the core of virtualization in the sense of this paper, we proceed with this definition. To illustrate the concept of virtualization solutions, a three-layered structure basing on the work of [1] is shown in Fig. 3.




Fig. 3 Three layer model of virtualization based on [1]

Fig. 3 is divided into three layers. On the bottom are the hardware resources such as CPU, storage or network. Between the individual VMs at the top level and the hardware is the hypervisor. This is the software that provides the virtual environment in which a VM operates [25]. The VMs on the main level include their respective operating system and applications. Furthermore, the ‘virtual hardware’ is shown in Fig 3. That is why we can again refer to the translation of virtual with ‘thought’ or ‘apparently’. Each VM seemingly has the possibility to access all resources in the cluster. However, in reality, the necessary resources are dynamically allocated by the hypervisor. Referring to the example of cloud services, virtualization allows the dynamic resource accesses that make up cloud computing and without which it would be even impossible [1].

IV. IMPACT OF VIRTUALIZATION ON THE AVAILABILITY OF A SYSTEM

A. General Benefits of Data Center Virtualization

To achieve their objectives, data center operators have to fulfil a number of subtasks, such as to comply with the minimization of costs or maximization of sales [4]. In many of these tasks, virtualization can be of help. Reference [26] presents server-consolidation, more flexibility and high-availability as the most important advantages. Reference [27] cites the consolidation of data center resources and the resulting effects as one of the main reasons to use virtualization. In non-virtualized data centers, the problem often arises that a plurality of server is not used, while others are congested [8], [28]. By the hypervisor of a virtualization solution, the resources of all servers are summarized and allocated dynamically as needed to VMs (known as load

balancing) [25]. Beyond the consolidation of servers, virtualization leads to other advantages such as cost reductions in the purchase of hardware, lower consumption costs (e.g. power, cooling), and less storage costs [29]. According to [30], virtualization can decrease the investment in new hardware and software by up to 70%. He also says that data center costs can decrease by up to 50% [30].

Reference [27] describes that a reduction of the physical machine and the use of sophisticated management tools allow a simplified administration, and that routine tasks can also be easily automated. Furthermore, it should be mentioned that maintainability is simplified. Since VMs consist of individual files, they can simply be moved to another server without downtime, e.g. for maintenance purposes [27]. This may also counteract human errors, one of the most common sources of errors for system failures [1]. Hence, this reduces the personnel costs because fewer employees are needed for administrative tasks [27]. This confirms that all efforts ultimately aim to meet the agreed SLA targets. Its compliance can not only save costs, but also generate revenue through customer satisfaction [7], [27]. One of these targets relates to the system availability, which is described in the next chapters.

B. Availability as the Central Non-Functional Quality of Service Property

In a series of publications, IT Infrastructure Library (ITIL) describes best practices for the implementation of IT service management [31]. For this purpose, ITIL shows how services should be provided in order to have an accepted high quality for the customer. Be considered, inter alia, incident management, availability management, or capacity management [7]. In the course of literature research in [1], [7], [9]-[11], we identified the availability of the system as the most important non-functional property. To this end, the availability management section is the focus in ITIL [7]. However, other sections, such as incident management, are critical, too. They ensure for their part the system readiness and define availability parameters in turn. Reference [1] also says that when companies examine the benefits of virtualization, a high availability is listed under their top priorities. This is justified by the fact that the availability is a prerequisite for all subsequent properties. For example, there is no response time without an available system.

ITIL defines availability as the “ability of a configuration item or IT service to perform its agreed function when required” [11]. According to ITIL [7], [11], availability results from the following four components:

- Reliability: Prevention of failures and maintaining the operability of components and services.
- Maintainability: Enable access to components and services in a normal working condition.
- Serviceability: Agreed internal and external support services.
- Security: Security measures to ensure normal operation.

Reference [16] describes that SLA allow an assessment of IT when they include the necessary information and

objectives. To determine the availability, meaningful metrics and measurement procedures have to be defined [7]. Therefore, it's crucial to know what to measure, how to measure, and how to interpret and communicate the results. To demonstrate the availability, there are various metrics that serve as a basis for measurement. According to [7], the measurement can be mainly based on:

- Mean Time Between Failure (MTBF)
- Mean Time Between System Incidents (MTBSI)
- Time to Repair (TTR)

These measures serve as proof of compliance with the target values of the SLA and are also part of them.

C. Connection Model of Virtualization and System Availability

According to [7] the focus of system availability is the prevention of downtimes. Therefore, [1] distinguishes between planned and unplanned outages. Planned outages are made for maintenance purposes or hardware/software updates [1]. Unplanned outages occur suddenly, for example, by natural disasters or errors of an employee [1]. In both cases, virtualization has features to resolve and to keep the system standby upright. Virtualization combines existing availability solutions of the physical world with new functions that are only possible with virtual machines [1]. To illustrate these functions, we describe a basic model for the relationship between virtualization and availability on three data center layers as it is presented in Fig. 4. This model has been created based on [1], [8], [27], [32]-[35].

Data center layer	Functions of virtualization – example VMWare	Availability components
Individual virtual machines	<ul style="list-style-type: none"> • Disk Mirroring • Disk Cloning • Multi-copying • VM Image • Operating System Rollout • Application Rollout 	Reliability
Group of virtual machines on a cluster of hosts	<ul style="list-style-type: none"> • Enterprise Capacity Scalability • Host Availability • Fault Tolerance • Live Migration • Secure Migration 	Maintainability Serviceability
Entire data center	<ul style="list-style-type: none"> • Provider Recovery • Host-based Replication • Storage-based Replication 	Security

Fig. 4 Model of virtualization functions to increase the system availability based on [1]

Fig. 4 is divided into three main sections/columns: 'Data center layers', 'Functions of virtualization – example of VMWare,' and 'Availability components'. The data center layers are identified by the work of [1]. He presents the division into 'Individual virtual machines', 'Group of virtual machines on a cluster of hosts,' and the 'Entire data center'. These layers determine the classification of the functions in the middle column. Therefore, the figure shows functions based on the example of VMWare. The provider was chosen due to its current leadership in the market as well as its broad

offer of functions [34], [36]. The right column shows the four availability components identified in ITIL by [7] that were initially described in Section IV B. Each function can influence each component. For instance 'Disk Mirroring' can have an effect on 'Reliability' as well as the 'Serviceability' of a data center.

To give some examples, one function of each layer will be detailed described below [1], [32], [33]:

- 'Individual virtual machine' – 'Multipathing': Multiple redundant access paths from the host to the storage array are set up so that the access path from the hard disk space to the host system is protected against failures. If, for example, two paths are set up and one fails, the other can be used. To realize load balancing, both paths can be transmitted simultaneously. This results in further performance benefits.
- 'Group of virtual machines on a cluster of hosts' – 'Fault Tolerance': Fault tolerant VMs can survive the failure of a virtualization host without them even precipitating or that their applications are affected. In addition to the primary VM, a secondary VM is installed on a different server that permanently duplicates the state of the primary. If the primary fails, the secondary takes over seamlessly. However, the disadvantage of this method is the duplicate resource requirement.
- 'Entire data center' – 'Disaster Recovery': In the event of a natural disaster, the entire data center can be destroyed. With disaster recovery, there is a second data center, either at another location of the company or at an external provider (hybrid cloud). In the case of failure of the primary data center, the backup takes over which has the identical infrastructure (replication).

When it comes to availability problems, we distinguish between two cases. On the one hand, the layers are interrelated into the form of problems on the upper level ('Entire data center') always cause problems on the lower levels ('Group of virtual machines on a cluster of hosts' and 'Individual virtual machines'). If e.g. a natural disaster destroys a data center, the single hosts with their VMs are ruined, too. On the other hand, problems with individual VMs can also influence the entire data center. The functions listed in Fig. 4 facilitate to affect these threats in a positive or a negative way. For example, the 'Fault Tolerance', described in bullet point two, brings the advantage that VMs can survive the failure of a host. However, at the same time there is the disadvantage of duplicated resources and thus higher costs of hardware. Further versions of the model have to describe these connections in a more detailed manner.

These statements show an example of how individual functions can affect the availability. Other functions can be found in [1], [8], [27], [32]-[35]. Literature and practice already show that virtualization have an effect on the system availability. The model gives a very basic overview of these impacts, but does not give a statement as to whether it is positive or negative. Therefore, the next two chapters will have a deeper look at the effect orientation. However, there are a variety of ideas how the model can be extended. For

example, functions of other vendors like Microsoft could be integrated. More on the expansion of the model is part of the outlook in Chapter VI.

D. Case Based Consideration of Virtualization Impacts

To illustrate the basic impacts of virtualization, this section presents a case in which virtualization technologies have increased system availability. The entire case is based on the study by arago Consulting GmbH [37]. In focus are Zumtobel, a company in the lighting industry, and arago, an IT service company for data management. Zumtobel provides its customers with a range of online services and portal applications to plan room lighting or view the latest products. Tom Brady, CIO of Zumtobel, said that the various user groups expect online services to be at the highest level at any time. If Zumtobel cannot offer this, the company loses sales and, in the worst case, customers.

In 2010, the company completed a restructuring of its data center landscape. The aim was to create a more flexible and efficient environment. The problem was that the existing infrastructure was not able to flexibly scale and effectively respond to the increased inquiries. An increased workload could lead to failures and in consequence to customer complaints. In addition, a total of 62 servers caused high costs for power, cooling, and space requirements. Brady said that in about 60 servers and dozens of applications there just have to be a potential for improvement.

The solution was to implement a fully virtualized hybrid cloud. This has only nine physical servers locally at Zumtobel and, if necessary, adds resources of the data center operator Arango (hybrid cloud). This means that both Arango and Zumtobel act as an operator of a data center since both have infrastructure. The NIST defines a hybrid cloud as “[...] a composition of two or more distinct cloud infrastructures [...] that remain unique entities, but are bound together by standardized or proprietary technology [...]” [19]. Thus, both Zumtobel and arago benefit from the virtualized infrastructure. On the one hand, it helps arago to comply SLAs with Zumtobel, and on the other hand Zumtobel can provide a high quality of service to its customers.

Outcome of the project was a significantly increased availability and performance of the system. In addition, as a result of server consolidation, hardware and energy costs could be significantly reduced. Brady says that in fast moving markets, IT should react just as dynamically as business. The virtualized infrastructure enables exactly that. Thus, the example shows the potential that virtualization offers and the positive effects a virtualization project may have on the system availability. However, negative aspects are not mentioned in this study, so the next chapter will have the focus on possibly occurring disadvantages.

E. Critical Examination of Virtualization

Besides the positive statement in the chapter above, and all positive voices from software vendors and literature, this contribution will also raise awareness that virtualization does not always lead to success and can also bring disadvantages.

According to [26] main disadvantages include partially increased complexity, the required expert know-how, and an initial performance loss. Reference [38] also mentioned that through the implementation of access by the hypervisor, VMs have less power than real hardware. The additional software layer, the hypervisor, can explain this. However, because of the very mature technology, these disadvantages lie only between 5–10% [38].

Another possible disadvantage can be found at the live migration of virtual machines [8]. In addition to increasing the availability, it also promises to improve the system performance, e.g., by using load balancing. This is not shown in every case, but strongly depends on which algorithm is used. Frequent migration operations can overload the network and have a negative effect with an increased input-output traffic to the response time and availability of the system. It is then possible that, in a completely homogeneous data center, virtualization has no positive effect at all. In this case it's even possible that factors such as the availability are influenced negatively. Therefore, virtualization is the core technology in cloud computing, since in this case there is a heterogeneous infrastructure in which virtualization can play to their strengths [1].

The points above are only some of the disadvantages that may arise. Other more general points that can be named are organizational barriers or the overall problem of such a comprehensive IT project. If we consider the model provided in Section IV C, we can see the effects. Those are either positive or negative or even a both. For virtualization, it is favourable that this offers many opportunities that cannot be offered by any other technology [1]. We notice that as long as providers are aware of the possible disadvantages, e.g. a lack of migration processes, virtualization can have the desired outcome, inter alia, on minimal downtime.

V. EXPERT INTERVIEWS ON THE RELEVANCE OF THE CONTRIBUTION

The evaluation of the work is to once again question the established hypotheses in the article. Furthermore, it will be clarified how the generated model can be used to solve the problem initially described. For this purpose, a group of experts were asked separately to give their opinion. The experts obtain their knowledge from their years of practice as well as academic work in the field of virtualization. Additionally, they show a variety of own publications on the topic.

The interviews were conducted in two different alternatives. First, we sent all chosen experts an email with a short version of the contribution. This included the abstract with the research question as well as the resulting model and its description. The email also referred to the possibility to get the entire paper for an even deeper insight. Approximately 50% of the experts took this opportunity to give their opinion. Next, we offered the possibility to have a personal interview. Two of the experts took this opportunity. Therefore, the expert interviews are documented as the replies to our emails and as the notes we took from the personal interviews. The

evaluation was carried out on the basis of the documented statements in comparison with the previous findings of the contribution. The following key messages from the interviews are analogously reproduced:

1) How Important Is a Recognized Service Quality for a Data Center in Terms of Adherence to the SLA?

The experts highlighted that a high service quality is crucial. The QoS is an integral part of the SLA and therefore very important regarding their adherence. The quality of the company is associated with the quality of the delivered product. Infringements will result in sanctions and in addition anything from directly measurable costs to reputation loss. Thus, a loss in revenue incurred in the form of migrated customers. One expert also pointed to the trade-off for each service provider. On one side, the service quality should be kept as high as possible and on the other side, the costs should be kept as low as possible. Each provider has to face this challenge. The experts also pointed out that when we talk about "recognized" or "high" quality we have to strictly define these words. Therefore the paper gave first impressions of what we understand by these terms. However there is still work to do in defining these in compromise of the broad views that exist at the moment.

2) Which Non-Functional Properties Do You Know that Influence the Service Quality, and Which Do You Consider as the Most Important?

The experts gave a broad number of non-functional properties, which were: availability, performance, maintainability, complexity, security, privacy, response time, reliability, usability, fault tolerance, scalability, accessibility, and adaptability. The interviews made clear that it is hard to give a general statement as to which of the properties is the most important. One expert gave a good explanation for why it is a difficult problem. He said, that first, this depends on the specific SLA of a provider. For example, the only non-functional property mentioned in AWSs SLA for the EC2 is the system availability. Therefore, it is likely that for AWS, availability is more important than e.g. performance, as we also mentioned earlier. However, this might be different for other vendors. Second, not all non-functional properties can be considered in isolation. For example, a very low response time might have an impact on the availability. Other experts were also aware the problem, but instead clearly stated that availability is the key criterion, since it is a pre-requisite for all other properties. They pointed out that without an available system the services are not working at all and so there is e.g. no performance. In addition, we say that we are convinced that it is hard to give a general statement that availability is the most important non-functional property. However, the expert interviews and our own experiences made us confident about the high relevance of availability.

3) How Can Virtualization Increase the Availability of a Data Center?

Our group of experts provided a broad knowledge about the possibilities of virtualization. In their experience,

virtualization provides various features that can be used in order to increase availability. Examples are using spare virtual images for a fast failover, elimination of single points of failure, or a fault tolerant multipath storage system that can be used without having to buy dedicated hardware. Therefore, the experts especially highlighted the live migration and the possibilities for disaster recovery. They stressed again that virtualization offers everything what a physical server environment offers, extended by new features and improvements to existing features.

4) Do You See Cases in Which Virtualization Cannot Improve the Service Quality, Particularly in Terms of Availability?

This was a crucial question for most experts. Virtualization is just one part of a technology stack that can be used in order to ensure specific service levels and therefore also availability. Hence, a high availability cannot be guaranteed only by virtualization technologies. For example, a service will still become unavailable if the underlying hardware fails. Therefore, virtualization cannot improve availability if it is the only part of the system that is optimized. One expert also stressed that, some features can have advantages as well as disadvantages, like we mentioned with fault tolerance as an example. He also said that in certain cases, no benefits at all arise. That we mentioned by describing the case of a totally homogeneous data center. Thus, each data center operator has to weigh the pros and cons of virtualization in its specific situation to get aware of the case based effects that will appear.

5) Do You Agree that the Provided Model Can Facilitate the Assessment of the Effect of Virtualization Functions on the Parts of a Data Center in Relation to the System Availability?

The experts agreed that the provided model would help in understanding the impact of virtualization on the availability of the system. The model delivers an overview and an added value for both IT service providers and their customers. However, one expert pointed out that there is still work to do. In his opinion further versions of the model inter alia should more clearly point out why specific functions like fault tolerance are part of the model. He mentioned the huge potential of the artefact, since the extension of the model to an additional dimension would create an even better overview.

Thus, the expert interviews can confirm the key assumptions of the work. Furthermore, it was demonstrated that the established model has an added value for providers as well as customers of IT services. They can gain a better understanding of the impact of individual functions of virtualization on the availability. Moreover, the future work that can be done on the basis of this paper was highlighted once again.

VI. CONCLUSION AND FUTURE RESEARCH

According to [7] there is currently a shift from the more technically oriented IT world into the direction of service orientation. That he leads back to the development towards

outsourcing and outtasking of computing power. In this context provider like Amazon Web Services offer their customers individually adaptable data center IT services inter alia in the form of infrastructure. Companies like Dropbox use these in turn to operate their own IT services. To secure a recognized service quality, the provider and the customer settle on a so called service level agreement (SLA). This contract determines functions of the system itself, frameworks such as contract terms as well as non-functional properties. As this paper showed, the system availability is a crucial non-functional property to guarantee the defined service quality and therefore to satisfy the customer. To guarantee a generally accepted quality of availability, virtualization solutions offer a wide range of functions. To foster progress in this area, a common understanding of the individual influencing factors is central. In the core of the work we introduced a model which provides a basic overview of the effects of virtualization functions on the parts of a data center in relation to the system availability. In this context, we also describe first impressions to consider the impact of the shown effects. At last we questioned our hypotheses and results by asking a group of experts on the topic. In conclusion the experts agreed with our findings and emphasized the huge potential of the prototype.

Based on our model, there is plenty of space for further research. A first step would be to extend the model by adding more functions from other vendors on the three layers. Furthermore, a detailed descriptions of the functions, interdependency between them as well as a description which function influences which availability component, could be added. Moreover the model could be extended to include additional dimensions. For instance various types of failures, maybe divided into planned and unplanned outages, could be used. Additionally, measuring methods for availability like the Mean Time Between Failure (MTBF) could be a part of the model. Another distinction could be made in the middle column, where the functions could be assigned to different virtualization types, such as full- or partial-virtualization.

This model respectively its successors could also be used for other non-functional properties. For example performance, maintainability, complexity, security, response time, reliability or usability could be detailed defined in the right column. Afterwards the illustration could show the influences of the virtualization functions on the components of these non-functional properties. In an advanced version a combination of all resulting models could show the entire influence virtualization functions have on non-functional properties. This construct would provide statements on the influence of specific virtualization functions on properties for specific cases. Although our model is only the prototype, all parties such as data center operators, IT service customers as well as researchers, can already benefit by these single point of easy understandable knowledge.

REFERENCES

- [1] M. Portnoy, *Virtualisierung für Einsteiger*. Weinheim: Wiley VCH, 2012, pp. 21, 29–40, 203–210, 277–294, 316–319.
- [2] M. Teuffel, “Service-Qualität im Rechenzentrum,” in *Das Rechenzentrum*, vol. 4, no. 2, 1981, pp. 119–126.
- [3] R. Narendula, “Amazon Web Services: a Case Study. Course: Business Process for IT Services 2012, EPFL”, (online), Retrieved 12 Nov, 2014 from <http://infoscience.epfl.ch/record/181192/files/AWS-Case%20Study-RammohanNarendula.pdf>, 2012.
- [4] P. Haric, “Gewinnmaximierung”, (online), Retrieved 12 Nov, 2014 from <http://wirtschaftslexikon.gabler.de/Archiv/54767/gewinnmaximierung-v5.html>, 2014.
- [5] R. Buchsein, F. Victor, H. Günther, and V. Machmeier, *IT-Management mit ITIL V3: Strategien, Kennzahlen, Umsetzung*. Wiesbaden: 1st ed., Vieweg + Teubner, 2007, pp. 12–43, 70–94, 272.
- [6] G. Disterer, “Zertifizierung der IT nach ISO 20000,” in *Wirtschaftsinformatik*, vol. 51, no. 6, 2009, pp. 530–534.
- [7] A. Olbrich, *ITIL kompakt und verständlich: Effizientes IT-Service-Management - den Standard für IT-Prozesse kennenlernen, verstehen und erfolgreich in der Praxis umsetzen*. 4th ed. Wiesbaden: Vieweg + Teubner, 2008, pp. 3, 16–42, 84–122, 260.
- [8] K. Turowski, *System Landscape Engineering: Very Large Business Applications I*. Otto-von-Guericke-Universität Magdeburg, 2013/2014, pp. 119, 128, 169–207, 213.
- [9] C. B. Border, “Cloud computing in the curriculum,” in *Proceeding of the 44th ACM technical symposium on Computer science education*, 2013, pp. 147–152.
- [10] W. Elsässer, *ITIL einführen und umsetzen: Leitfaden für effizientes IT-Management durch Prozessorientierung*. 2nd ed., München: Hanser, 2006, pp. 57, 70 p. 59 p. 58
- [11] A. van der Veen, and J. van Bon, *Foundations of ITIL®*. Haren Publishing: Van, 2011, p. 326.
- [12] Amazon Web Services, “Amazon Elastic Compute Cloud (EC2) Service Level Agreement (SLA)”, (online), Retrieved 17 Nov, 2014 <http://aws.amazon.com/ec2/sla/>, 2014.
- [13] A. Hevner, S. March, J. Park, and S. Ram, “Design science in information systems research,” in *MIS Quarterly*, 28 (1), 2004, pp. 75–105.
- [14] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, “A Design Science Research Methodology for Information Systems Research,” in *Journal of Management Information Systems*, vol. 24, no. 3, 2007, pp. 45–77.
- [15] Office of Government Commerce (OGC), *Service strategy*. London: TSO. The Stationery Office Ltd, 2007.
- [16] C. Stych, *ITIL, Informatik im Fokus*. Berlin, Heidelberg: Springer, 2008, pp. 6, 28, 86.
- [17] L. M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, “A break in the clouds,” in *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, 2008, p. 50.
- [18] R. Dukaric, and M. B. Juric, “Towards a unified taxonomy and architecture of cloud frameworks,” in *Future Generation Computer Systems*, vol. 29, no. 5, 2013, pp. 1196–1210.
- [19] P. Mell, and T. Grance, “The NIST Definition of Cloud Computing. Recommendations of the National Institute of Standards and Technology”. NIST SP 800–145, 2011.
- [20] D. D. Lamanna, J. Skene, and W. Emmerich, “SLAng: A language for defining service level agreements,” in *Distributed Computing Systems*, 2003, pp. 100–106.
- [21] K. Turowski, *System Architecture: Very Large Business Applications II*. Otto-von-Guericke Universität Magdeburg, 2014, p. 169.
- [22] Duden, “virtuell”, (online), Retrieved 17 Nov, 2014 from <http://www.duden.de/suchen/dudenonline/virtuell>, 2014.
- [23] G. J. Popek, and R. P. Goldberg, “Formal requirements for virtualizable third generation architectures,” in *Communications of the ACM*, vol. 17, no. 7, 1974, pp. 412–421.
- [24] H. Jehle, H. Wittges, A. Bögelsack, and H. Kremer, “Virtualisierungsarchitekturen für den Betrieb von Very Large Business Applications,” in *Multikonferenz Wirtschaftsinformatik (MKWI)*, 2008, pp. 1901–1912.
- [25] T. C. Bressoud, and F. B. Schneider, “Hypervisor-based fault tolerance,” in *ACM Transactions on Computer Systems*, vol. 14, no. 1, 1996, pp. 80–107.
- [26] F. Lampe, *Green-IT, Virtualisierung und Thin Clients: Mit neuen IT-Technologien Energieeffizienz erreichen, die Umwelt schonen und Kosten sparen*. 1st ed, Wiesbaden: Vieweg + Teubner, 2010, pp. 74, 103–104.
- [27] C. Baun, M. Kunze, and T. Ludwig, “Servervirtualisierung,” in *Informatik-Spektrum*, vol. 32, no. 3, 2009, pp. 197–205.

- [28] M. Bolz, "Skalierbare Lastverteilung für verteilte virtuelle Umgebungen". University of Paderborn, 2003.
- [29] A. Singh, "An Introduction to Virtualization", (online), Retrieved 17 Nov, 2014 from <http://www.kernelthread.com/publications/virtualization/>, 2014.
- [30] F. Hantelmann, "Mit Virtualisierung RZ-Kosten halbieren," in *iX*, vol. 12/2008, 2008, pp. S. 88.
- [31] V. Arraj, (2013), "ITIL_The_Basics", (online), Retrieved 17 Nov, 2014 from http://www.best-management-practice.com/gempdf/ITIL_The_Basics.pdf, 2013.
- [32] M. Langes, *Das Managementmodell zur Datenverfügbarkeit und Datensicherheit: Von der Anwendung bis zum Rechenzentrum*. 1st ed. Books On Demand, 2012, pp. 8, 100
- [33] A. S. Tanenbaum, J. Muhr, and M. van Steen, *Verteilte Systeme: Grundlagen und Paradigmen*. München: Pearson Studium Informatik, 2003, p. 411.
- [34] VMWare, "VMware in Leaders Quadrant of 2014 Gartner Magic Quadrant for x86 Server Virtualization Infrastructure", (online), Retrieved 14 Nov, 2014 from <http://blogs.vmware.com/tribalknowledge/2014/07/gartner-x86-magic-quadrant.html>, 2014.
- [35] D. Zimmer et al., *VMware vSphere 5: Das umfassende Handbuch*. 2nd ed., Galileo Press: Bonn, 2012, pp. 36, 44, 945.
- [36] T. J. Bittman, M. A. Margevicius, and P. Dawson, "Magic Quadrant for x86 Server Virtualization Infrastructure", (online), Retrieved 5 Nov, 2014 from <https://www.gartner.com/doc/2788024/magic-quadrant-x-server-virtualization>, 2014.
- [37] arago Consulting GmbH, "Fallstudie: Autopilot-gesteuerter Cloud-Betrieb von arago ermöglicht Zumtobel Group Downsizing von 62 auf neun Server bei gleichzeitiger Leistungssteigerung", (online), Retrieved 16 Nov, 2014 from https://www.arago.de/wpcontent/uploads/2012/02/Fallstudie_Zumtobel_Maerz11_DE-130212.pdf, 2011.
- [38] M. Hardt, "Virtualisation for Grid-Computing," in *Cracow Grid Workshop*, vol. 20, 2005, pp. 23–28.