Ronny Hans, David Dahlen, Sebastian Zöller, Dieter Schuller, Ulrich Lampe:
Enabling Virtual Manufacturing Enterprises with Cloud Computing – An Analysis of Criteria for the Selection of Database as a Service Offers. In: Américo Azevedo: Advances in Sustainable and Competitive Manufacturing Systems, p. 427-438, Springer, June 2013. ISBN 978-3-319-00556-0.

Enabling Virtual Manufacturing Enterprises with Cloud Computing – An Analysis of Criteria for the Selection of Database as a Service Offers

Ronny Hans[1], David Dahlen[1], Sebastian Zöller[1], Dieter Schuller[1], and Ulrich Lampe[1]

[1]Multimedia Communications Lab (KOM)
Technische Universität Darmstadt
64289 Darmstadt, Germany

Abstract

In our globalized world, small- and medium-sized enterprises in the manufacturing domain face a highly competitive environment. They are subject to various challenges, such as very short product life cycles and a strong price competition with companies from low-cost countries. To remain competitive in such an environment, new forms of collaborations, like Virtual Manufacturing Enterprises, are required. An essential part of virtual organisations is data provisioning. Thereby, data from various sources like factories' ERP systems or data provided by sensors need to be processed and stored. In this context, data storage is a crucial architectural element that influences both functional aspects and competitive aspects, especially costs, of Virtual Manufacturing Enterprises. For realizing Virtual Manufacturing Enterprises with low up-front investments, the application of new technologies, such as Cloud Computing, is required. For storage of information in databases, Database as a Service offers from the Cloud can be exploited. However, since there is a huge amount of providers acting on a non-transparent market, it is difficult to find appropriate "Database as a Service" offerings. To overcome this problem, we provide a criteria catalogue for the selection of providers and their services. Further, we show how different offers, which at the first glance look very similar, could cause very different expenses. With our work, we simplify the selection and evaluation of Cloud storage providers and provide an evaluation of current Cloud storage service offers.

1. Introduction

These days, Small and Medium Enterprises (SMEs) in the manufacturing domain face various challenges such as short product life cycles. Further, customers request a high degree of customization at low costs. To remain successful, companies need to find a way to cope with challenges effectively [9]. On the one hand, SMEs play an important role regarding economic growth, employment, and innovation. All over the world they contribute 54% to the GDP of the countries on average [3], [30]. On the other hand, SMEs do not have the financial strength like large companies to handle all the challenges individually. To remain competitive and fulfil customer expectations, SMEs need to focus on their core competencies, and thus have to collaborate with other companies. Hence, the success of an enterprise depends on the individual capability as well as efficient collaborations in well-defined value chains across enterprises [6]. To address the mentioned challenges, Agile Manufacturing (AM) is a popular concept. In this context, agility means the exploration of competitive bases through the use of reconfigurable resources in order to provide products and services regarding customer needs in a quickly changing environment [32]. Thereby, AM focuses on setting up whole organizations for production including different enterprises [16]. A basic building block to achieve AM is the concept of "Virtual Manufacturing Enterprise" (VME) [16], [32]. A VME is created by at least two companies [26] and is the enabling element for process sharing across organizational boundaries [7], [20]. The different companies within a VME share costs, skills, and competencies to produce goods which are beyond the capability of individual companies. Thus, VMEs may be equally powerful as big enterprises, but are still equipped with the ability of small enterprises to react quickly to changes of market circumstances [26].

An essential part of the VME infrastructure is Information and Communication Technology (ICT) [11]. Because of the fact that SMEs do not have the same financial power like large companies, an ICT infrastructure is required that can be established with reasonable costs. To achieve that, we suggest the use of Cloud Computing (CC). CC offers different advantages, such as cost reduction, charging in a pay-per-use manner [4], flexibility, and scalability [2]. With these characteristics, the corresponding ICT infrastructure can grow or shrink according to the business needs. Further, CC offers the access to data independently from the location of the user [13]. Within an ICT architecture for VMEs, data storage plays an outstanding role, e.g., all monitoring data, product information and data from sensors need to be stored and provided. All possible types of storage services offered by Cloud providers are referred to as the common term *Data Storage as a Service (DaaS)* [31]. Database functionality, as a specific subset of DaaS, is referred to as *Database as a Service (DBaaS)* [19]. Among the various offers in the market, an appropriate choice is difficult because of the large amount of cloud service providers, a low market transparency, and tedious access to relevant information [27]. Thus, we provide a criteria catalogue to qualitatively evaluate DBaaS offers. As the application of the derived criteria catalogue, we provide a qualitative and quantitative evaluation of three popular offers from Amazon, Microsoft and Google.

The paper proceeds as follows: In Section 2, we derive a criteria catalogue with essential functional and non-functional criteria, which puts SMEs in the position to select Cloud storage offers effectively under consideration of their specific needs. In Section 3, we apply the criteria catalogue to selected DBaaS offers from Amazon, Microsoft, and Google. Furthermore, we calculate the costs of these offers depending on different load profiles. We show that similar looking accounting models can cause extremely different costs. Related work is presented in Section 4. Finally, conclusions and an outlook on future work are provided in Section 5.

2. DEVELOPMENT OF A CRITERIA CATALOGUE

In this section, we identify key criteria in order to enable the comparison of different DBaaS offers. Based on a literature review we find, identify, and describe potential criteria for selecting appropriate offers. The references indicating the usage of such criteria are provided in Table 1.

Table 1: References for Cloud Storage Selection Criteria

Criteria	**Repschläger et al. [27]**	**Hetzenecker et al. [14]**	**Cryans et al. [8]**	**Jatana et al. [17]**
Pricing Model	X	X		
Cost Transparency	X	X		
Data Model	X		X	
Scheme			X	
Indexing				X
Scalability	X	X	X	X
Service Level Agreement		X		
Data Security and Compliance	X	X		
Encryption	X	X		
Performance	X	X		X

To begin with, the *pricing model* describes how the cost for a cloud storage offer is calculated. That is, it represents the key parameters that determine the total usage cost. In order to be able to assess the cost arising from performing IT projects, a certain degree of *cost transparency* is indispensable. If an IT project is to be realized using the own data processing center, costs for servers, electricity, maintenance staff, network, and facilities have to be considered [5]. Referring to the Cloud, different parameters come into play, depending on the pricing model applied by a cloud storage provider. A high degree of cost transparency thereby enables solid decisions with respect to selecting a certain cloud storage offer for a concrete project. Further, a comparison of these costs with the cost of providing these IT resources in-house is enabled. For assessing the cost transparency of a cloud storage offering, mainly two factors are of importance. On the one hand, the complexity of the pricing model has to be considered. If, for instance, lots of parameters are required for determining the total cost, the complexity of the pricing model is rather high, resulting in a lower cost transparency degree. On the other hand, the predictability of the mentioned parameters plays a role.

Another criterion which should be considered when a DBaaS offer has to be selected is the *data model*. In addition to the typical, relational architectures, further data models exist, commonly referred to using by the term “NoSQL” [28]. Obviously, different data models are differently appropriate for concrete tasks and problem instances, respectively. Thus, the applied data model should serve as a key criterion for the selection of a certain cloud storage offer. Depending on the concrete application scenario, different *schemes* with advantages as well as disadvantages need to be taken into account to distinguish cloud storage offerings from each other. Relational databases generally require exact definitions of data structures, while flexible schemes are allowed by NoSQL databases [10]. In order to speed up the access to datasets, *indices* are commonly used. Thereby, the access pattern of the application logic determines the selected indices. In this respect, it has to be noted that the creation and application of indices may be either limited or required by different cloud storage providers. Amazon's DynamoDB[*] and Microsoft's Azure Table Storage[†], for instance, do not support user defined indices [1], [21]. This may require for the user to adapt an existing application logic. Therefore, indices should be considered as additional cloud storage selection criterion.

Another criterion is given by the *scalability* of offered services. In this respect, we aim to assess the ability of a DBaaS offer to satisfy the need for fast scalability. In general, databases may be scaled in two ways, i.e., vertical vs. horizontal scaling. While vertical scaling is performed by enhancing the performance of the according database server, horizontal scaling refers to interconnecting several, different database servers [18]. The criterion of scalability thereby should indicate which type of scaling is supported in which way. In order to assess, whether a Cloud storage offer satisfies the requirements of an IT project, dedicated information regarding the offered *performance* are needed. This criterion aims at collecting and providing observations on the performance of different Cloud storage offerings in order to enable comparisons among them. In this context, evaluation data as well as benchmarks are commonly applied. With respect to the domain of databases, the benchmarks of TPC [29] are predominant for relational databases.

Cloud providers might make use of *Service Level Agreements (SLAs)* in order to provide their customers with guarantees concerning different non-functional properties, such as availability or response time. As there is no dedicated, individual contract between a cloud user and the respective provider, the provisioning of such prescribed guarantees and may support the trustworthiness of a cloud provider form a customer's perspective. With respect to data transmission and storage of data, legal regulations as well as company internal requirements exist and have to be adhered to. For instance, the German Data Protection Act[‡] specifies and provides mentioned legal regulations. As data is handed over to the Cloud and thus, a third party, it has to be ensured that the respective cloud provider adheres to the according guidelines. For this, the criterion *data security and compliance* should indicate whether a cloud storage provider has undergone formal certification, indicating whether mentioned legal and company internal requirements are met. Whether *encryption* is used for transmitting and storing data and which encryption algorithm thereby is applied, may also be prescribed according to organizational or legal requirements. For instance, the German Data Protection Act requires taking "appropriate" actions for the protection of personal data. For this, the criterion *encryption* should indicate whether the data communication as well as the storage of data supports encryption.

3. Evaluation of Cloud Storage Offers

In the following, we exemplarily apply the criteria that were deduced in Section 2 to three selected DBaaS services. We focus on DBaaS offers from three major providers in this context, namely Amazon DynamoDB, Google App Engine Datastore[§], and Microsoft Windows Azure Table Storage. First, we provide a qualitative analysis of the DBaaS offers based on the criteria catalogue from Section 2. After that, we conduct a quantitative analysis of the selected DBaaS offers with a special focus on monetary costs associated with the offers for an example scenario, since monetary costs constitute one of the most relevant decision criteria in the very cost-sensitive manufacturing domain. Thus, we provide not only a guideline on how to employ our criteria catalogue in the selection process for a DBaaS provider, but also a comparison between the services offered by prominent providers that are currently active in this market.

3.1 Qualitative Analysis

In the following, we provide an exemplary analysis of Microsoft's Windows Azure Table Storage service according to our criteria catalogue from the previous section. We conducted this analysis as well for the Amazon DynamoDB and the Google App Engine Datastore services. An overview of the results can be found in Table 2. However, as the quantitative analysis in Section 3.2 outlines, Microsoft's service seems very promising from an economic point of view. In consequence, we present the details of its qualitative analysis in more depth here. With respect to this highly dynamic area and the fast changing offers of Cloud provides, it has to be noted that all used information and considered data have

[*] http://aws.amazon.com/de/dynamodb/
[†] http://msdn.microsoft.com/en-us/library/windowsazure/ee924681.aspx
[‡] http://dejure.org/gesetze/BDSG
[§] https://developers.google.com/appengine/docs/python/datastore/

been collected in October 2012"*Pricing Model:* Pricing depends on three parameters, namely, the storage space used in a month, the number of memory transactions, and the outgoing data transfer. The basis for the storage space accounting is the average daily storage space allocation. Depending on the arranged level of redundancy, costs between US$ 0.093 and 0.125 per month and GB accrue, with discounts for data volumes over 50 TB. Memory transactions cost US$ 0.01 per 100,000 transactions. Further, the outgoing data transfer is charged with US$ 0.12 up to US$ 0.19 per month and GB depending on region [23].

Cost Transparency: As the costs for the usage of the service depend on the three mentioned parameters of storage space, number of memory transactions, and outgoing data transfer, their respective predictability is the main indicator to estimate the cost transparency of the service. As the costs for data storage rise proportionally to the used data storage, this parameter can be predicted quite well. The number of memory transactions correlates to the number of queries and the query types with the number and size of returned datasets. Thus, this parameter is rather hard to estimate. The same holds true for the outgoing data transfer, which is also dependent on type and number of queries and hence rather difficult to estimate.

Table 2: Qualitative Analysis of DBaaS Offerings

	Amazon DynamoDB	**Microsoft Windows Azure Table Storage**	**Google App Engine Datastore**
Pricing Model	Storage Space, reserved Read/Write-throughput, Outgoing Data Transfer	Storage Space, API calls, Outgoing Data Transfer	Storage Space, I/O-Operations, Outgoing Data Transfer
Cost Transparency	+	+	-
Data Model	Key-Tuple	Key-Tuple	Key-Tuple
Scheme	No fixed Scheme, only Primary Key fixed	No fixed Scheme, only Primary Key fixed	No fixed Scheme, only Primary Key fixed
Indexing	Only Primary Key	Only Primary Key	Any
Scalability	Horizontal	Horizontal	Horizontal
SLA	No SLA provided	99.9% availability	99.95% availability
Data Security and Compliance			
- *Datacenter Locations in EU*	Ireland	Ireland, Netherlands	Confidential
- *EU Location for Datacenter selectable*	Yes	Yes	Yes (Surcharge)
- *Safe Harbor Support*	Yes	Yes	Yes
- *Applicable Law*	USA	Ireland	USA
- *Data Owner*	User	User	User
- *Data Usage by Service Provider*	No	No	No
Encryption	No	No	No
Performance*	N/A	N/A	N/A

Data Model: The Windows Azure Table Storage service relies on a NoSQL database of the key-tupel type [24].

Scheme: Windows Azure Table Storage has no fixed scheme. Name, number, and data type are individually assigned for each dataset. However, each dataset has to possess the three mandatory attributes "*PartitionKey*", "*RowKey*", and "*Timestamp*". PartitionKey allows distributing a table over different servers to realize load balancing. RowKey allows to uniquely identifying a dataset within a partition. Timestamp is an internally used field. The service offers the possibility to specify up to 252 user-defined attributes [24].

Indexing: Data is stored as a clustered index based on a primary key, which consists of the attributes PrimaryKey and RowKey. User-defined indices are not supported [24].

* To the best of our knowledge, there is no comparable analysis of the performance available

Scalability: Windows Azure Table Storage provides automatic horizontal scaling with respect to data volume by partitioning datasets according to their PartitionKey. Datasets related to multiple PartitionKey values can be stored on one server. However, scalability directly relating to the system load is not foreseen.

SLA: A Service Level Agreement is provided with the service, which guarantees an availability of 99.9%. A refund of 10%, respectively 25%, of the monthly fee is provided in case the availability drops between 99% and 99.9%, respectively below 99%. The availability is calculated as the difference between 100% and the mean error rate in the accounting month for the memory transactions of the customer for the respective booked offer [22].

Privacy and Compliance: The locations of stored data can be specified as to be restricted to Europe. The customer remains owner of the data all the time. Microsoft takes part in Safe Harbor and the Windows Azure Table Storage service is certified with ISO27001 and ISAE3402 Type 2 [25].

Encryption: Stored data is not encrypted, but the data access via the API is realized via HTTPS [24].

Performance: In several tests, a latency of 18 ms per read and 24 ms per write operation for a dataset has been reported [12]. Hill et al. [15] could register in one of their test setups an average performance of ca. 1,150 datasets per second for CRUD-operations* reaching a maximum of 1,525 datasets per second.

3.2 Quantitative Analysis

In order to evaluate the monetary efficiency of the three mentioned DBaaS offers from Amazon, Microsoft, and Google, we devised an exemplary application scenario and calculated the monetary costs arising within this scenario. Furthermore, the gathered monitoring data shall be stored in the data store, as well. We assume a maximum length of the monitoring data of 140 characters, with an average length of 70 characters.

Table 3: Load Levels Employed in the Example Scenario

Load Level	Average Number of Monitoring Datasets per Day	Average Number of Views of the Data Source Websites per Day	Peak Load of Number of Monitoring Datasets per Second	Peak Load of Number of Views of the Data Source Websites per Second	Initially Registered Data Sources	Initially Stored Monitoring Datasets
1 (Low)	500	1000	1	5	100	182,500
2 (High)	50,000	100,000	100	500	10,000	18,250,000

Table 4: Monthly Resource Consumption in the Example Scenario

DBaaS Offers	Load Profile	Employed Storage	Employed Writes	Employed Reads	Outgoing Data Transfer
Amazon DynamoDB	1	31.7 MB	1	8	65.7 MB
	2	3.1 GB	100	750	6.4 GB
Windows Azure Table Storage	1	20.6 MB	75,000†		93.7 MB
	2	2.0 GB	7,500,000		9.1 GB
Google App Engine Datastore	1	31.2 MB	60,000	630,000	73.4 MB
	2	3.0 GB	6,000,000	63,000,000	7.2 GB

The mapping of the monitoring data to its source shall be supported. It shall be possible to create new monitoring datasets, read datasets, and delete datasets. Furthermore, all datasets associated with a certain source shall be retrievable arranged according to their timestamps, with the newest on top. The data storage solution should provide an availability level of 99.9% and a Web interface with a dedicated website for each data source, along with its identifier

*CRUD: Basic operations on data storage; acronym for: create, read, update, delete
† With Windows Azure Table Storage, write and read transactions are not differentiated. Thus, the values provided here comprise the number of consolidated memory transactions.

and description as well as its last 20 datasets. Regarding the Microsoft Developer Network[*], we assume an average size of 514 Byte for a source dataset, 188 Byte for a monitor dataset and 200 Byte for the protocol overhead per transaction. Regarding the prices for the western region we calculated US$ 0.01 per 100,000 transactions, US$ 0.125 per GB stored data per month and US$ 0.12 per transferred GB.

Table 5: Monthly Costs in the Example Scenario

DBaaS Offers	Load Profile	Costs for Employed Storage	Costs for Employed Writes	Costs for Employed Reads	Costs for Outgoing Data Transfer	Sum
Amazon DynamoDB	1	US$ 0.00	US$ 0.00	US$ 0.00	US$ 0.00	US$ 0.00
	2	US$ 3.39	US$ 77.29	US$ 120.41	US$ 0.65	US$ 201.74
Windows Azure Table Storage	1	US$ 0.13	US$ 0.01[†]		US$ 0.00	US$ 0.14
	2	US$ 0.25	US$ 0.75		US$ 1.10	US$ 2.10
Google App Engine Datastore	1	US$ 0.00	US$ 0.00	US$ 0.00	US$ 0.01	US$ 0.01
	2	US$ 0.49	US$ 5.25	US$ 43.58	US$ 0.86	US$ 50.18

Based on two different load levels (cf. Table 3), we calculate the total price of each mentioned DBaaS offer for one month of operation. This allows drawing conclusions on the cost efficiency with regard to scalability and the impact of different usage parameters. The individual results of our calculations can be found in Table 5. Figure 1 provides a comparison with respect to the monetary costs arising in the example scenario with the different load levels for the three analyzed DBaaS offers Amazon DynamoDB, Google App Engine Datastore, and Microsoft Windows Azure Table Storage.

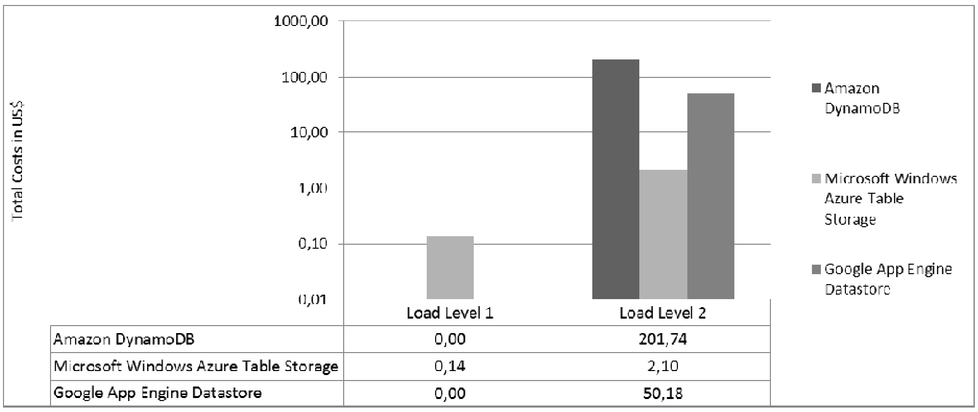

Figure 1: Cost comparison for DBaaS offerings for the introduced example application scenario

Our results show that for the different load levels in our sample application scenario, the overall costs highly differ. Notably, for load level 1, the Amazon and the Google DBaaS service can be used for free, wheras the Microsoft service incurs minimal costs. A signifcant difference between the offered services is found by analyzing load level 2. In this scenario, the costs when employing the Microsoft solution are nearly negligible, wheras the usage of the Google service would incur costs of roughly US$ 50 and the Amazon service would even incur costs of roughly US$ 200 per month.

In consequence, it is essential to spend sufficient time on a thorough analysis of the application scenario and the respective current and estimated future requirements for the data storage solution to be employed. Due to very different pricing schemes, one solution can be cost-efficient in a certain application scenario, but in another scenario, its employment can be disadvantageous and result in significant additional spendings. Nevertheless, a selected service

[*] http://blogs.msdn.com/b/windowsazurestorage/archive/2010/07/09/understanding-windows-azure-storage-billing-bandwidth-transactions-and-capacity.aspx.

[†] With Windows Azure Table Storage, write and read transactions are not differentiated. Thus, the values provided here comprise the number of consolidated memory transactions.

needs to fulfil both qualitative and qunatitive requirements. For exampe, DBaaS offers that are not in accordance with data security and compliance requirements are not of interest for enterprises, even if they are the cheapest solution.

4. RELATED WORK

Repschlaeger et al. [27] developed a classification framework for comparison and selection of Infrastructure as a Service (IaaS) providers. Their findings are based on a literature analysis, on an analysis of current cloud service offers, and on expert interviews. The authors aim to simplify the cloud provider selection process and to increase the market transparency. Hetzenecker et al. [14] developed a model of requirements to asses CC providers. Their findings based on a literature analysis, on expert interviews, and on an online survey. The provided model should help CC users to identify the individual needs and evaluate providers regarding their requirements. Cryans et al. [8] present new database technologies based on Hadoop, a scalable distributed file system. They give an overview of HBase and the corresponding infrastructure. Further, they propose comparison elements to contrast relational databases with distributed databases like HBase. Jatana et al. [17] examine essential characteristics of relational and non-relational databases. They give a basic overview of both database types and conclude with a comparison of both database types by means of ten characteristics. In summary, all related work focuses either on Cloud-specific or on database-specific criteria. None of them provides a comprehensive criteria list for DBaaS offers. Our work comprises essential elements from both worlds to take the specific characteristics of Cloud-based database offers into account.

5. SUMMARY AND OUTLOOK

Small and Medium Enterprises (SMEs) in the manufacturing domain act in a highly competitive environment, which is hallmarked by short product life cycles, a high level of customization, and a cheap production in low-cost countries. To handle these challenges efficient new collaboration forms like Virtual Manufacturing Enterprises (VMEs) are required. A major enabler for such agile organizations is Information and Communication Technology (ICT). Traditionally, the deployment of new ICT causes high up-front investments which could be a hurdle for SMEs. To overcome this obstacle we propose the utilization of Cloud Computing (CC). By doing so, the ICT infrastructure can dynamically grow or shrink according the business needs. Besides a wide range of advantages, CC causes challenges regarding the selection of appropriate services. In this paper, we present a catalogue of criteria to assess, compare, and select *Database as a Service (DBaaS)* offerings regarding the business needs. Further, we conduct a qualitative and quantitative analysis of DBaaS offers from Microsoft, Amazon and Google. Our results show that for different load levels different offers are more or less appropriate. In our future work, we plan to implement a Cloud Storage solution as a part of a complete architecture for VMEs and thus to enhance our calculation with results from real world deployments.

ACKNOWLEDGEMENTS

This work is supported in part by the Commission of the European Union within the ADVENTURE FP7-ICT project (Grant agreement no. 285220), within the SIMPLI-CITY FP7-ICT project (Grant agreement no. 318201), and by E-Finance Lab e. V., Frankfurt am Main, Germany (http://www.efinancelab.com).

REFERENCES

[1] Amazon: "Amazon DynamoDB Developer Guide", http://docs.amazonwebservices.com/amazondynamodb/latest/developerguide, 2012, (accessed January 19, 2013).

[2] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia: "A View of Cloud Computing", Communications of the ACM, Vol. 53, No 4, pp.50–58, 2010.

[3] M. Ayyagari, T. Beck, and A. Demirguc-Kunt: "Small and Medium Enterprises Across the Globe", Small Business Economics, Vol. 29, No. 4, pp. 415–434, 2007.

[4] R. Buyya, C. Yeo, and S. Venugopal: "Market Oriented Cloud Computing: Vision, Hype and Reality for Delivering IT Services as Computing Utilities", In 10th IEEE International Conference on High Performance Computing and Communications, 2008.

[5] Y. Chen and R. Sion: "To Cloud or not to Cloud? Musings on Costs and Viability", In 2nd ACM Symposium on Cloud

Computing, 2011.

[6] *Y. Choi, D. Kang, H. Chae, and K. Kim: "An Enterprise Architecture Framework for Collaboration of Virtual Enterprise Chains", The International Journal of Advanced Manufacturing Technology, Vol. 35, No. 11, pp. 1065–1078, 2008.*

[7] *V. Corvello and P. Migliarese: "Virtual Forms for the Organization of Production: A Comparative Analysis", International Journal of Production Economics, Vol. 110, No. 1–2, pp. 5–15, 2007.*

[8] *J. Cryans, A. April, and A. Abran: "Criteria to Compare Cloud Computing with Current Database Technology", In Proceedings of the International Conferences on Software Process and Product Measurement, 2008.*

[9] *S. Dowlatshahi and Q. Cao: "The Relationships Among Virtual Enterprise, Information Technology, and Business Performance in Agile Manufacturing: An Industry Perspective", European Journal of Operational Research, Vol. 174, pp. 835–860, 2006.*

[10] *S. Edlich, A. Friedland, J. Hampe, B. Brauer, and M. Braeckner: "NoSQL: Einstieg in die Welt nichtrelationaler Web-2.0-Datenbanken", Hanser, 2011.*

[11] *W. Faisst: "Information Technology as an Enabler of Virtual Enterprises: A Life-cycle-oriented Description", In Proceedings of the European Conference on Virtual Enterprises and Networked Solutions, 1997.*

[12] *J. Giardino, J. Haridas, and B. Calder: "How to Get Most out of Windows Azure Tables", http://blogs.msdn.com/b/windowsazurestorage/archive/2010/11/06/how-to-get-most-out-of-windows-azure-tables.aspx, 2010. (accessed January 20, 2013).*

[13] *B. Hayes: "Cloud Computing", Communications of the ACM, Vol.51, No. 7, pp. 9–11, 2008.*

[14] *J. Hetzenecker, S. Kammerer, M. Amberg, and V. Zeiler: "Anforderungen an Cloud Computing Anbieter", In Tagungsband der Multikonferenz Wirtschaftsinformatik, 2012.*

[15] *Z. Hill, J. Li, M. Mao, A. Ruiz-Alvarez, and M. Humphrey: "Early Observations on the Performance of Windows Azure", In 19th ACM International Symposium on High Performance Distributed Computing, 2010.*

[16] *S. Jain: "Virtual Factory Framework: A Key Enabler for Agile Manufacturing", In Proceedings of Symposium on Emerging Technologies and Factory Automation, 1995.*

[17] *N. Jatana, S. Puri, M. Ahuja, I. Kathuria, and D. Gosain: "A Survey and Comparison of Relational and Non-Relational Database", International Journal of Engineering Research & Technology, Vol. 1, No. 6, pp. 1-5, 2012.*

[18] *D. Kossmann, T. Kraska, and S. Loesing: "An Evaluation of Alternative Architectures for Transaction Processing in the Cloud", In ACM SIGMOD International Conference on Management of Data, pages, 2010.*

[19] *D. Linthicum: "Cloud Computing and SOA Convergence in Your Enterprise: A Step-by-Step Guide", Addison-Wesley, 2009.*

[20] *M. Martinez, P. Fouletier, K. Park, and J. Favre: "Virtual Enterprise - Organization, Evolution and Control", International Journal of Production Economics, Vol. 74, No. 1–3, pp. 225–238, 2001.*

[21] *Microsoft: "Developer Center", http://www.windowsazure.com/en-us/develop/overview/, 2012, (accessed January 20, 2013).*

[22] *Microsoft: "Windows Azure Support: Service Level Agreement", http://www.windowsazure.com/en-us/support/legal/sla/, 2012, (accessed January 20, 2013).*

[23] *Microsoft: "Pricing details", http://www.windowsazure.com/de-de/pricing/details/, 2012, (accessed January 21, 2013)*

[24] *Microsoft: "Table Service REST API", http://msdn.microsoft.com/en-us/library/dd179423.aspx, 2012. (accessed January 20, 2013).*

[25] *Microsoft: "Windows Azure Trust Center", https://www.windowsazure.com/en-us/support/trust-center/, 2012, (accessed January 20, 2013).*

[26] *K. Park and J. Favrel: "Virtual Enterprise - Information System and Networking Solution", Computers & Industrial Engineering, Vol. 37, No. 1-2, pp. 441–444, 1999.*

[27] *J. Repschlaeger, S. Wind, and R. Zarnekow: „Klassifikationsrahmen für die Anbieterauswahl in der Cloud", In 41. Jahrestagung der Gesellschaft für Informatik, 2011.*

[28] *C. Strauch: "Nosql Databases", http://www.christof-strauch.de/nosqldbs.pdf, 2011, (accessed January 20, 2013).*

[29] *Transaction Processing Performance Council: "TPC Benchmarks", http://www.tpc.org/information/benchmarks.asp, 2012, (accessed January 20, 2013).*

[30] *World Bank Group: "World Bank Review of Small Business Activities", 2004.*

[31] *L. Youseff, M. Butrico, and D. Silva: "Toward a Unified Ontology of Cloud Computing", In Grid Computing Environments Workshop, 2008.*

[32] *Y Yusuf, M. Sarhadi, and A. Gunasekaran: "Agile Manufacturing: The Drivers, Concepts and Attributes", International Journal of Production Economics, Vol. 62, No. 1–2, pp. 33–43, 1999.*