An Approach for a Cloud-based Contribution Margin Dashboard in the Field of Electricity Trading

Oliver Norkus[1], Brian D. Clark[2], Florian Merkel[3], Björn Friedrich[4], Jürgen Sauer[5], H.-Jürgen Appelrath[6]

Abstract: From our ongoing project, we present the collected requirements and business drivers from the energy industry imposed upon analytical information systems, specifically in the case of a contribution margin control as part of electricity trading. As a solution, we introduce our architecture for business intelligence in the cloud and apply these in the afore mentioned use case in the form of a research prototype. We provide the current implementation status at an early stage, on one hand to increase the transparency of the still young field of business intelligence in the cloud, and on the other hand to demonstrate the development potential of the technology for the energy sector.

Keywords: Cloud Computing, Business Intelligence, Electricity Trading, Cost Analysis, On Demand, Cloud-based Service.

1   Introduction

In our industrialized and increasingly globalized world, energy plays a major role. Without energy in the forms of uranium, coal, electricity, gas, and oil, our world and our daily life would look very different. The availability of light, heat, information, and power for the machinery locomotion is nowadays an inseparable part of our technology-society. We are in the midst of global competition for scarce resources. The energy supply is one of the central themes in science, politics, and business [AKM12].

The energy industry is undergoing changes, based on the increased production from renewable energy, the intensification of competition (especially in the context of liberalization and decentralized energy production) as well as a growing volatility, for example, in electricity trading. Electricity trading is a complex process in the contexts of market changes, legal situations, and further influence donors. In particular, the diversity

[1] University of Oldenburg, Department of Computer Science, Escherweg 2, 26121 Oldenburg, oliver.norkus@uni-oldenburg.de
[2] University of Oldenburg, brian.clark@uni-oldenburg.de
[3] University of Oldenburg, florian.merkel@uni-oldenburg.de
[4] University of Oldenburg, bjoern.friedrich@uni-oldenburg.de
[5] University of Oldenburg, Department of Computer Science, Ammerländer Heerstr. 114-118, 26121 Oldenburg, juergen.sauer@uni-oldenburg.de
[6] University of Oldenburg, Department of Computer Science, Escherweg 2, 26121 Oldenburg, appelrath@informatik.uni-oldenburg.de

of actors, the volatility of the relevant ratios, and the diversity of involved information technology (IT) systems all escalate the complexity [AKM12].

The relevant actors and roles have particularly increased by the unbundling of complexity. Through the direct marketing of electricity, remote meter reading, changes to smart grids, and (last but not least) through the trading of electricity capacity at the European Energy Exchange (EEX) raises the volatility and the velocity. Not only the minute or quarter hour accuracy, but also the volume of data results in new challenges for current IT systems. Through these factors, the challenges for business and statistical reporting as well as analysis becomes more convoluted, and as a result, the requirements on the used IT systems intensifies [NH14].

The functional requirements remain approximately the same in the areas of data procurement, processing, and analysis. The non-functional requirements increase in the areas of availability, accessibility, agility, query/analysis performance, and cost [No15].

Currently, the primarily used reporting and analysis systems are constructed monolithically. This means classical business intelligence (BI) solutions are complex, rigid, and expensive [BK10]. Therefore, traditional BI solutions are not suitable to cope with the increased demands. One approach to solving this problem is the use of cloud technology. Cloud computing has many properties that meet today's demands, especially in regards to analytical applications [No15].

Due to the novelty of the development in science and industry, there are not many field reports. Thus the acceptance of service offerings is still minor. Although first product versions from various manufactures are already available, they are in part incompatible, and cannot be combined or integrated with many existing systems, partly because fundamental issues have not been clarified in terms of business and software architecture [No15].

This is what we want to change with our approach, by utilizing a transparent, reusable cloud business intelligence system and the corresponding architecture applied in the form of a contribution margin control dashboard as part of the process of investigating the cost recovery of an energy supplier for new and existing customers in the context of direct marketing.

We aim to present the first approaches from our ongoing research project, and thus make our first experiences and results transparent. Therefore, we present in this paper required foundations (see section 2) of the requirements ascertained for cloud-based reporting and analysis systems and our use case for testing and evaluation (section 3), the approach for the architecture, and the current state of the implementation (section 4). Finally, we point out next steps and further work (section 5).

2 Foundations

The demand for increased flexibility in business is ever present. One aspect of this is the ability to enable employees to locate data and perform analyses without requiring intensive training of the complicated processes involved. At the same time, the commoditization of remote access to computing resources enables the utilization of powerful hardware from lightweight portable devices. The combination of these factors gives employees the potential to initiate and monitor a variety of complex processes in the middle of dynamic business situations, such as an on-site negotiation with prospective customers.

A solution to these needs is a datacenter containing powerful hardware for running analyses, combined with an intuitive web user interface designed for mobile devices. Using virtual machines (VM), additional hardware capacity can be allocated (or vacated) to follow with the current demand on system load. For the duration that additional capacity is needed, it can either be rented or reallocated from inactive hardware in a company's private data center. This process is invisible to the user, who can dispatch the analysis process from their portable device, and expect to have results shortly thereafter.

The concepts relating to supplying normal employees (without explicit system training) to performing analyses efficiently and effectively is covered in section 2.1, while technologies covering the ability to use commoditized hardware access is covered in section 2.2. The combination of these concepts is covered in section 2.3.

2.1 Business Intelligence

BI is a broad category of applications and technologies for gathering, storing, analyzing, and visualizing information. BI is associated with supporting IT resource management in order to optimize the process of decision-making, with a focus in providing decision makers with necessary information at just the right time [Ra09].

The understanding of the term *BI* ranges from multidimensional data structures on individual information systems to complex system landscapes analyzing large quantities of data for management. Analytical information systems focus on the provisioning of information and functional support for analyses in support of professionals and managers [Wa09].

2.2 Cloud Computing

Cloud Computing (CC) is defined as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction" [MG11].

A primary aspect of CC is the notion of a cloud service, defined as the provision of virtual IT resources (i.e. logical resources that are mapped to physical hardware) demonstrating the following characteristics [MG11], [VHH12]:

- *Resource pooling*: cloud services enable the shared utilization of physical resources by means of virtualization (virtual resources instead of physical hardware) and multi-tenancy (multi-tenant data management).

- *Rapid elasticity*: services can be immediately provisioned and released as demanded and the resources available for provisioning often appear to be virtually unlimited.

- *Measured service*: service utilization is measured and commonly monetized with a pay-per-use pricing model.

- *Broad network access*: cloud services are supported by ubiquitous network access and technical standardization.

- *On-demand self-service*: consumers can unilaterally provide themselves with services as they need them, due to extensive service automation on the side of the provider.

2.3   Business Intelligence Cloud

An early definition made in 2010 is from [SB10], which describes Business Intelligence Cloud (BI Cloud) as an IT architecture with the purpose of providing analytical capabilities as a service. A distinction can be made between the focus on the application (private, public, and hybrid) and for the architectural level (infrastructure, platform, and software).

More specifically in 2014, [NA14] states that analytical applications can be deployed as a cloud service, and that the outcome is called Business Intelligence as a Service (BIaaS). Such BI Cloud services reflect on the characteristics of typical cloud services features described below [NA15].

BIaaS enables organizations to flexibly respond to changing demands. Through flexible allocation, resources can be distributed depending upon project needs, especially to enable complex queries in an efficient manner. BIaaS analyses and reports can be flexibly integrated and orchestrated into processes.

The average employee (without specialized training) is able to automatically allocate and utilize resources around the particular services they are interested in through the system, without having to specifically invoke requests through the IT department, and be blocked waiting on a response from a human. Analytical applications in the cloud promote availability and reliability. They are accessible via an Internet connection from anywhere with only a web-enabled device.

BIaaS provides a high level of agility: reports and analyses can be adapted to directives and the needs of stakeholders. BIaaS not only provides flexibility at the user interface level with visualization and modeling, but also at the technical level, e.g., modifying cubes, extract, transform and load (ETL) processes, or multidimensional databases, which can all be processed in a simple and flexible manner. This is especially important when there is a demand to rapidly change operational IT landscapes (e.g. due to corporate transformations).

Analytical applications in the cloud promise very strong performance levels for queries and analyses, especially rapid deployments of complex reports and a short turnaround time when changing parameters. BIaaS can be billed on a per-use basis. Payment can be calculated for some aspect of the report, for example, the number of invocations, the hardware resources used (e.g. processing time), or access to proprietary data sets.

3 Requirements and Business Drivers

We performed face-to-face expert interviews from the energy industry for the discovery of requirements and business drivers. In accordance with agile software development processes, we performed requirements engineering in multiple iterations. We evaluated each interview to identify requirements and business drivers. Afterwards we framed the requirements utilizing techniques for requirements engineering [Ru09]. In the next step, the experts evaluated the requirements. Based upon the evaluation, the requirements were overhauled. The following interview started with a short presentation of the latest version. In the first interview, only basic requirements were discussed. In each of the subsequent interviews, the level of detail increased. After all experts agreed with the collected requirements, we considered the requirements from the viewpoint of a computer scientist. The interviews revealed a few use cases. The primary use case considered here is the contribution margin control.

Energy suppliers are chiefly interested in the costs a customer needs to pay per kilowatt-hour so that they realize a profit. To date, this calculation is mostly handcrafted with the help of spreadsheet software. This calculation is called *breakeven analysis* or *contribution margin control* and the result is the *cumulative cost*. The calculation of the cumulative costs is needed when energy suppliers make offers to prospective customers. The cumulative cost per kilowatt-hour is the essential part of any offering. When a contract between a customer and an energy supplier is up for extension, the current cumulative cost must be analyzed to calculate the cost for the new offer. The process of calculation lasts up to one week and has to be completed before a sales pitch takes place. This is especially disadvantageous for the field managers. If they visit a customer site for a customer pitch, it would be impossible for them to calculate a new value of the cumulative cost when necessary. They have to return to the office and find a new appointment with the customer. This is time-consuming for both sides.

During the expert interviews, we were able to derive the following requirements for a

contribution margin dashboard. With respect to building a prototype the fulfilled and here shown requirements are focused on functional aspects.

Availability: system redundancy is required to guarantee high availability, e.g. data backups and redundant infrastructure.

Accessibility: field managers are often on site for customer pitches. Thus, the service should be available on a wide range of devices, (not only on traditional computers, but also on mobile devices), and must be available over the Internet.

Accounting: given that BI systems are expensive to set up and maintain, it is necessary for clients to pay for their usage. The charges can be calculated either with respect to the complexity of the calculations or according to the volume of reports generated.

Performance: the duration of calculation, and data transport to the user should be minimized. Poor performance leads to decreased user acceptance of BI systems [NA15]. While the quality of mobile Internet connections an unavoidable factor, keeping transmitted data small and presenting overview before detail (like loading a graphic file over a dialup connection) helps to mitigate this.

Modularity: there are multiple processes related to the calculation of analyses. There is the potential for modifications or integrations of further analyses. Modularity enables the effective management of the system.

Expandability: many data sources are required to perform the analysis. The integration of additional data sources by users must be possible.

User empowerment: field managers should not need to write complex database queries or analyze raw data. The user interface needs be usable without special training, and allow, the user to intuitively understand the interface and the analysis result. This is key to user acceptance.

Web standards: nearly every mobile device or computer supports web technologies like HTML and JavaScript. The best way to bring the service to a wide variety of devices is to use the accepted standards.

The requirements *availability* and *accessibility* lead to the two additional computer science related requirements *scalability* and *flexibility*. Many field managers from multiple energy suppliers will use the system. The hardware resources have to be provisioned as required. If one server is overburdened, another server must be started to take the load away from the overburdened one. Through resource scaling, the service's full quality is available to every user.

Summarized, the service has to be available over many devices; field managers are often at customer sites and not at their office. The first aspect leads to flexibility according to the used devices and the second one to flexibility according to the access to the service. Both can be managed by providing the service over the Internet. Nearly every device in

the information technology domain supports web technologies. Therefore, providing the service over the Internet is preferable. Additionally the service could be used over the mobile web. This makes the service available over a great territory.

Scalability and flexibility are most relevant to the architectural. We show how we approach them in section 4 and how we intend to fulfill them in section 5.

4 Architectural Approach

In order to satisfy the identified requirements we envisioned the *Business Intelligence in the Cloud for Energy* (BICE) as a cloud service system, providing contribution margin analysis for electric utilities. As such, it needs to scale for large numbers of users while still responding quickly to requests. Additionally, the system has to enable users to add data sources and customize the generated reports. The analytical functionality is restricted to using preexisting data and does not incorporate any production planning.

In this section, we present our approach for a cloud-based contribution margin dashboard. In section 4.1, we describe the server portion and necessary system components that were identified in order to provide the required functionality. Following in section 4.2 is the description of our prototype, including the technology platform, data structures, and exemplary interaction between the components of the BICE system.

4.1 Components

The BICE system comprises of several components (seen in fig. 1) modelled with Unified Model Language (UML).

The user interacts directly with the Local Client component. It represents the user interface and control logic, as well as display data to the user. It uses the model-view-controller design pattern to manage complexity by separating the view from business logic. This component sends requests to the Load Balancer component, a proxy server that forwards the request to an available Dispatcher component running inside a VM. A Resource Monitor measures and reports the load on each VM, so that the Resource Management component can start up or shut down VM instances as needed. This response to the incoming load maintains Quality of Service for availability, accessibility, and query performance.

Depending on the type of the request, the Dispatcher communicates with the various other server components in order to answer the user request.

The User Database component stores relevant user data, such as username, password, permissions, the user's company, and usage statistics, and is encapsulated in the User Manager component. The Authorization and Rights Management component uses permissions to determine if a user may perform actions.
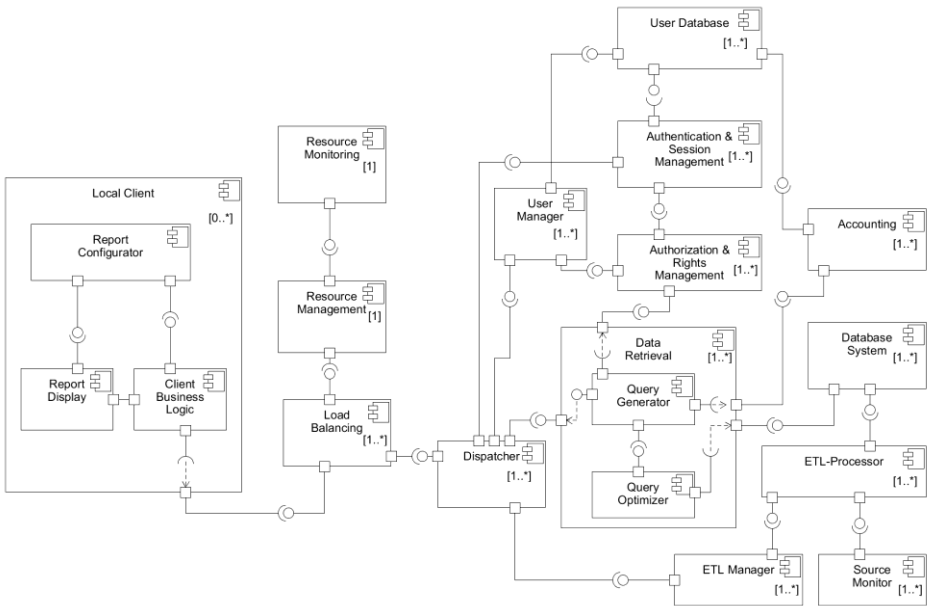
Fig. 1: Component diagram

Usage based billing is made possible by tracking the usage of individual users through the Accounting component and aggregating the data per electrical utility company. The user management related components will be deployed on a separate machine than the dispatcher.

In order to encapsulate and optimize database access, all analysis requests go through the Data Retrieval component running on separate VMs, so that resource intensive operations do not affect the overall responsiveness. To avoid unnecessary database queries and to enforce security, the Data Retrieval component verifies that the user who initiated a request has sufficient permissions to execute it.

The ETL-Manager, ETL-Processor, Source Monitor, and Database System components provide data extraction, transformation, storage, and analysis. The first three components play an important role in realizing flexible operation by enabling the user to incorporate data sources through the ETL-Manager component, which will load and transform data using an ETL-Processor. The Source Monitor component permits the monitoring of external data sources for changes, and signaling the corresponding ETL-Processor to start its work. During this process, heterogeneous data needs to be transformed into a format that facilitates the required data analyses provided by the system.

4.2 Prototype

We are developing a prototype of our architecture based on state of the art technologies provided by the Hasso-Plattner-Institute (HPI) to assess the feasibility of the conceptualization. The VMs required for running the components will be run and managed using HP Converged Cloud™, which adheres to the OpenStack™ architecture. The in-memory database system is a shared SAP HANA™ instance. The prototype will be evaluated in the scenario of analyzing the contribution margin of an electric utility companies for new and existing customers in the context of over the counter energy trading.

During the expert interviews, four key performance indicators were identified: total cost for supplying electricity to a customer, production cost, labor cost, and profit. These figures can be aggregated over time, the customer, the generators, and the distribution grids.

The data structure shown in fig. 2 was devised according to these dependencies. Despite the analytical focus of the system, the data is not structured as a star or snowflake scheme in order to avoid redundancy in an in-memory database system.
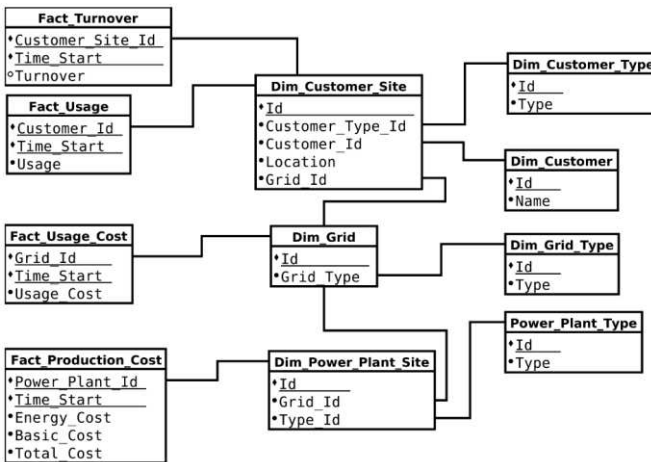

Fig. 2: Data fact table

The left direction in fig. 2 indicates the data fact tables that hold the data used to compute the indicators. Each record is linked to the entry in the dimension table it depends on, which in turn can reference additional records that provide information orthogonal to the dimension (e.g. Dim_Customer_Type of Dim_Customer_Site) and records that represent coarser representations along the dimension (e.g. Dim_Customer of Dim_Customer_site). The Time_Start field incorporates the time dimension in the fact tables. It describes the start of a 15-minute time interval for which the record is valid.

The Database System component stores data with the described data structure. The next section shows how, the system passes requests and responses back and forth.
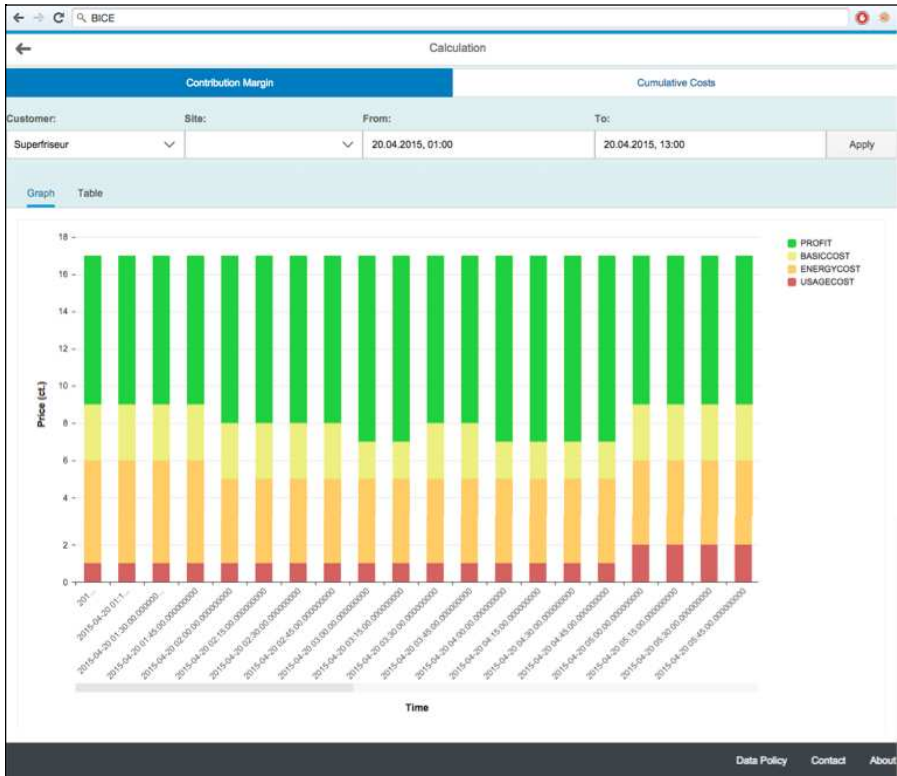

Fig. 3: Web interface for the dashboard

The user interacts with a web graphical user interface (GUI) like the one shown in fig. 3 to request and view data. The GUI targets mobile devices, specifically tablets. The choice of designing the GUI primarily for tablets is due to the larger screen size compared to smart phones, which can display more information at once, aiding in providing an overview of the data.

When the user is interacting with the system, they first need to log in with a username and password (not shown here). After setting the input fields to the desired values reflecting the user's request the *Apply* button sends the request to the server, which handles the necessary backend tasks and returns the response data. The client provides two views for displaying the requested information: a graph view, which shows the relevant indicators in a chart, and a view presenting information in table form.

In addition, the GUI also provides functionality for adding data sources and managing the user's own account.
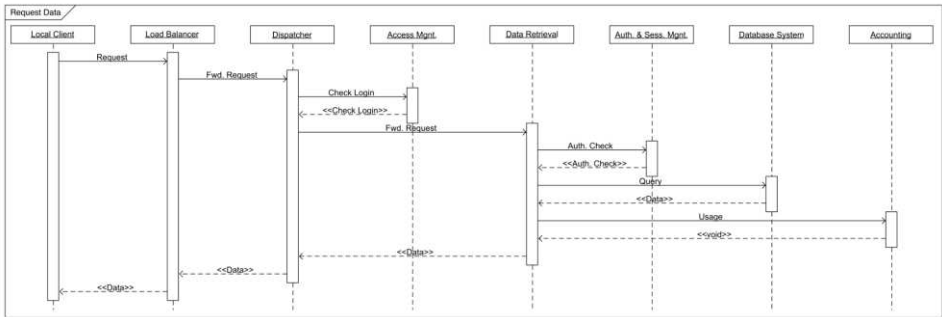
Fig. 4: Sequence diagram

The diagram in fig. 4 shows a typical sequence of interactions between the components as envisioned in the current architecture. It assumes the Local Client component is authenticated with the system when it sends a request for data. The request is received by the Load Balancer and passed to a Dispatcher component, which verifies that an authenticated client sent the request. The Dispatcher then asks for the respective data from the Data Retrieval component. After confirming that the Local Client instance has the proper access rights, it sends the actual database query. Additionally, resource usage is logged: possible measurements include the processing load on the database and/or the Data Retrieval component, or the time the request takes to be completed. The appropriate value is passed to the Accounting component, which tracks usage for individual users and provides functionality to aggregate the usage over all users of one electric utility company. After the database query returns, the requested data is passed back all the way to the client component while being transformed and converted as necessary.

5   Conclusion and Further Challenges

New challenges for BI arise with the velocity, volatility, and complexity of the markets. Old-fashioned BI systems cannot satisfy the new requirements. We introduced a flexible BI cloud system, which fulfills the new requirements.

In the use case, the BI cloud system performs the breakeven analysis faster and is (especially the field managers) more flexible. That must improve their negotiations and hence their business. However, there are many more processes in the energy industry. The breakeven analysis is only a single aspect. To date, many systems are used for providing a single process. We strive for a future of a single integrated system providing all processes in the energy business sector.

In this paper, we have shown the demands energy suppliers have for BICE, and how we gathered them. Based on the requirements, we developed the architecture for a prototype.

We finished developing the architecture for the BI cloud system and the HPI provides the hardware resources for the prototype. We are implementing based on the HP Converged Cloud™ and the SAP HANA™ in-memory database system. In respect to agile development strategies, the progress is will be presented to the experts throughout the development period. With their feedback, we are able to adapt the prototype closer to their needs. Moreover, the development and the improvement of existing features and components continues.

The next steps are to implement the further components shown in the architecture and deploy them to the mentioned hardware. The prototype is going to be evaluated with regards to the use case. We will present the results of our evaluation and further findings by means of a report and a poster as part of the HPI Future SOC Lab Days on the 4[th] November 2015.

Literaturverzeichnis

[AKM12] Appelrath, H.-J., Kagermann, H., Mayer, C. (Eds.). Future Energy Grid: Migrationspfade ins Internet der Energie. Springer-Verlag, 2012.

[NH15]   Norkus, O.: Haas, M.: Stromhandel in der Cloud. ew - Magazin für die Energiewirtschaft, EW Medien und Kongresse, Frankfurt, S.44-46, 4/2015

[No15]   Norkus, O.: Ein Ansatz zur Standardisierung von BI in der Cloud. In (Uwe Aßmann, Birgit Demuth, Thorsten Spitta, Georg Püschel, Ronny Kaiser): Proceedings of Software Engineering & Management 2015, Dresden, Gesellschaft für Informatik, Bonn, Bonner Köllen Verlag, S.299-303, LNI, P-239, 3/2015

[Ru09]   Rupp, C.: Requirementsengineering und -Management. Hanser Verlag, 5. Auflage, 2009.

[BK10]   Baars, H.: Kemper, H-G.: Business intelligence in the cloud? PACIS 2010 Proceedings. Paper 145. 2010.

[Ra09]   Ranjann, J.: Business Intelligence: Concepts, Components, Techniques and Benefits. Journal of Theoretical and Applied Information Technology.Vol 9. No 1. (pp 060-070)

[Wa09]   Watson, H. J.: Tutorial Business Intelligence - Past, Present, and Future. Communications of the AIS 25 (2009, 1), pp. 487-510.

[MG11]   Mell, P.; Grance, T.: The NIST definition of cloud computing. 2011. National Institute of Standards and Technology, 2011

[VHH12] Vossen, G.; Haselmann, T.; Hoeren, T.: Cloud-Computing für Unternehmen. Technische, wirtschaftliche, rechtliche und organisatorische Aspekte. dpunkt, Heidelberg (2012).

[SB10]   Seufert, A.; Bernhardt, N:. Business Intelligence und Cloud Computing. HMD Praxis der Wirtschaftsinformatik, 47(5):34–41, 2010.