COOPERATIVE COMMUNITY CLOUDS FOR SMALL AND MEDIUM ENTERPRISES

Facilitating Cloud Computing for Small and Medium Enterprises with the Cooperative Paradigm

Till Haselmann, Gottfried Vossen
European Research Center for Information Systems (ERCIS), University of Münster
Leonardo-Campus 3, 48149 Münster, Germany

Stefanie Lipsky, Theresia Theurl
Institute for Cooperative Research (IfG), University of Münster, Am Stadtgraben 9, 48143 Münster, Germany

Keywords: Cloud computing, Cloud services, Cooperative, Community cloud, Small and medium enterprises.

Abstract: Cloud computing, or more generally cloud services, ought to be particularly attractive for small and medium enterprises (SMEs). According to expert opinion, these enterprises should be able to benefit overproportionally from outsourced IT services due to a variety of reasons. As of today, however, many SMEs still refrain from adopting cloud services because they do not trust the cloud service provider sufficiently and they cannot assess the legal implications of using cloud services. In this paper, we present an argument in favor of a cooperative community cloud that applies the paradigm of a cooperative to cloud operations. We sketch our vision of a "trustworthy" cloud, argue for its viability, and identify areas for future research.

1 INTRODUCTION

Many small and medium enterprises still refrain from adopting cloud services because of a lack of trust in both the cloud service provider and the legal situation. In this paper, we present an argument in favor of a cooperative community cloud that applies the paradigm of a cooperative to cloud operations.

Cloud computing is a new variety of IT outsourcing that has been gaining much attention over the last few years. In this paradigm, a cloud service provider (CSP) offers very standardized IT services which are accessed over a network, usually the Internet. These services comprise products on different levels of abstraction, ranging from software usable by the end user (Software-as-a-Service, SaaS) over platforms that allow storing varying degrees of custom logic (Platform-as-a-Service, PaaS) to virtualized IT hardware (Infrastructure-as-a-Service, IaaS). The CSP provides a seemingly unlimited supply of resources and allows the user to quickly make use of more or less resources depending on the current demand—a concept commonly referred to as *elasticity*. Users pay only for the resources that are actually utilized (*pay-per-use*), sometimes plus a small periodic fee (Armbrust et al., 2010; Mather et al., 2009; Mell and Grance, 2009).

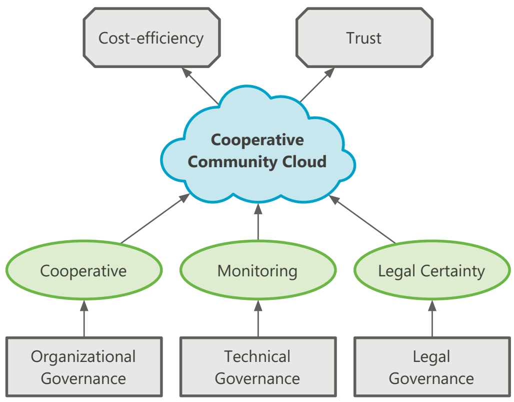

Figure 1: Schematic overview of the suggested approach.

One important trait of cloud computing is the provided abstraction that masks all the internals of the cloud to the end user. On the one hand, this is very beneficial because it allows the CSP to provide the cloud services in the most cost-efficient way possi-

Haselmann T., Vossen G., Lipsky S. and Theresia T..
COOPERATIVE COMMUNITY CLOUDS FOR SMALL AND MEDIUM ENTERPRISES - Facilitating Cloud Computing for Small and Medium Enterprises with the Cooperative Paradigm.
DOI: 10.5220/0003443001040109
In *Proceedings of the 1st International Conference on Cloud Computing and Services Science* (CLOSER-2011), pages 104-109
ISBN: 978-989-8425-52-2

ble. In addition, the CSP can perform maintainance and restructuring operations in the cloud data centers without impacting the end user's experience. On the other hand, however, this abstraction also means that the end user gives up control over most aspects of data processing and IT operations, a fact that proves to be a significant obstacle for many companies, especially smaller ones.

In consequence, cloud computing is still "terra incognita" for most small and medium enterprises (SMEs). According to a study by the University of Münster, the majority of SMEs are not yet using cloud services and have no plans of changing this in the foreseeable future (Haselmann and Vossen, 2011). Major reasons for the reluctant adoption are:

- lack of trust in the security of the cloud services and the CSPs,
- lack of control over processes in the cloud, e. g., with regard to billing, but also with regard to data protection, as well as
- lack of certainty about the legal compliance of the CSP and the cloud services with respect to contract design, accountability, and warranty.

In this paper, we argue that purely technical approaches cannot sufficiently address and, in particular, cannot eliminate these barriers to adoption. In fact, we outline why this it is at least as much an organizational problem. Thus, we suggest to employ a three-faceted approach that tackles the problem not only from the *technical* point of view but also from *organizational* and *legal* standpoints. We argue that SMEs are able to pursue cloud computing more easily by applying the paradigm of a cooperative to the cloud data center operation. This is flanked by a comprehensive monitoring approach on the technical level and specific recommendations on the legal level. In essence, SMEs can thus build trust into "their own cloud" and still benefit from the cost-efficiency of the cloud computing paradigm. Fig. 1 shows the approach in a schematic overview.

2 BACKGROUND & RELATED WORK

2.1 The Cooperative Paradigm

A cooperative is a business organization owned and operated by a group of individuals (or companies) with a common goal and for their mutual benefit (MacPherson, 1995). Traditionally, cooperatives are found in economic sectors such as agriculture, finance, and the real estate industry. In the past years, however, many newly founded cooperatives can be ascribed to expanding, future-oriented industries, such as IT service providers (Theurl and Schweinsberg, 2004).

Generally speaking, the cooperative paradigm allows for a more flexible and more diverse alignment of the supplied products with the customers' demands. It implies a common administration as well as sharing of relevant resources. Thus, it often leads to lower costs per unit. In addition, decentralized knowledge from the various companies can be combined without them losing their independence (Theurl, 2005b).

SMEs forming a cooperative typically want to realize synergies from a joint organization of elements on their value chains. This leads to economies of scale, scope and skills (Williamson, 2005). By concentrating their market power, the SMEs can also compensate competitive disadvantages towards large companies (Theurl and Meyer, 2005).

Cooperatives are characterized by a high degree of institutionalization and a standardized scope for action. The applicable "rules of the game" are defined in legislature and in the statutory regulations of the cooperative (Theurl, 2005b). The clear rules facilitate handling uncertainty and can thus foster credibility and trust among the members. The strategic orientation of the cooperative is built on the notion of creating value only for the members, which are simultaneously the owners of the cooperative. This way, anonymous vested interests are excluded from all strategic decisions. The values of the cooperative are grounded in the *MemberValue*, i. e., a special type of shareholder value (Theurl, 2005a; Theurl, 2009). The MemberValue can be interpreted as the overall business value of the cooperative.

Applied to the cloud services domain, the cooperative paradigm attempts to alleviate deficiencies in the individual IT infrastructures and to build collective competence with regard to cloud services. A working example of this is the German DATEV eG, a cooperative that manages the IT infrastructure and IT services for tax consultants in Germany. Founded around 40 years ago, DATEV has been providing data center services, some of which may nowadays be subsumed as cloud computing, mainly SaaS and IaaS in a private cloud model. The cooperative paradigm provided the foundation for continuous and sustainable improvement of both service quality and product portfolio for the currently more than 39,000 members (Gulden, 2010). Of course, DATEV was not created with a particular focus on cloud computing, but rather to satisfy the various IT demands of its members. The application of the cooperative paradigm specifically to cloud operations has not yet been researched.

2.2 Cloud Monitoring

By virtue of its special governance, the cooperative is a suitable form for the organization of cloud operations that addresses current deficiencies of the cloud. However, a purely organizational approach is not comprehensive enough to establish sufficient trust with all potential cloud users among the SMEs. Therefore, the technical governance must be considered as well. Two of the most important aspects in this respect are

1. allowing the user to fully specify the permitted (and illicit) use of his or her data, and
2. elucidating the internals of the cloud in order to provide the user with a means of monitoring what processing is actually performed on that data.

An important benefit of the cloud paradigm is the abstraction from the physical resources involved, e. g., the actual computing or storage devices. Abstract interfaces mask the underlying systems from the user (Armbrust et al., 2010; Vaquero et al., 2009). However beneficial this abstraction may be in many ways, it is also the reason why many users tend not to trust cloud services. With all internal processing steps hidden from the users, it is impossible to trace the actual data flows; users can only hope that the CSPs keep their promises. In essence, they feel as though their data were "vanishing in a black hole". Using the cloud, therefore, requires a large amount of trust in the provider, an amount that most SMEs are not willing to invest (Haselmann and Vossen, 2011).

The typical approach for getting an overview of the internals of complex IT systems is the use of monitoring software, but existing approaches in the cloud-monitoring domain cannot provide a solution. The systems available, such as Hyperic CloudStatus, Nimsoft Cloud Monitoring or Rackspace's CloudKick, provide monitoring functionality on a different level: they provide the user with an overview of the used *abstract* cloud resources, analyzing and logging monetary costs. Some systems also analyze or predict usage patterns and can react to changing demands or resource shortage. Nevertheless, due to their nature, these monitoring tools do not allow any insight into the operation of the *physical* cloud systems.

The internals of a cloud must not be described in too much detail because that would deprive the CSPs of their freedom to organize the physical systems as they see fit. One possibility that we would like to explore is the concept of *spheres of control* (SOCs) (Davies, Jr., 1973; Davies, Jr., 1978). Stemming originally from a database systems background, SOCs have been used mainly in the context of workflow execution in the past. They allow grouping of operations and activities or parts of processes into semantic units that contain both transactional and non-transactional parts. These groups are nestable and can be rolled back when compensating actions are specified. In the context of cloud computing, the SOCs could be referencing parts of one or more clouds. If founded on a suitable level of abstraction, we expect to see new possibilities for the cloud monitoring.

A second promising approach is a modeling based on so-called *private enclaves* or *virtual private clouds* (VPCs) (Wood et al., 2009). In this approach, parts of the cloud are logically separated and treated as autonomous "mini-clouds in the cloud". Currently, the cloud users specify which systems should be part of a VPC. However, this concept could also be used the other way around, having the CSP identify parts of the cloud that are similar by certain characteristics. For example, the CSP could indicate systems with a very high level of security or services that are partially provided by unreliable third-party providers.

2.3 Legal Issues of Cloud Computing

The legal situation of cloud computing in the European Union (EU) is dominated by data protection and copyright legislature. In practice, the strongest reasons for not using the cloud are based on concerns about data protection and security. In Germany—as an example of a country with very strict data protection laws—, current legislature states various requirements to the cloud user that render the use of cloud services under a strict interpretation of the legal code effectively illegal. Some classic examples clarify this point:

- The cloud user is required to provide entry and access controls for the cloud data centers and the IT systems therein (Schuster and Reichl, 2010).
- The cloud user has to know the full nature and extend of the data processing steps undertaken by the CSP.
- The cloud user is responsible for ensuring that personally identifying data (PII) is not transferred to a country with an insufficient data protection legislature (Weichert, 2010).

It is obvious that the strict interpretation of these requirements is incompatible with fundamental traits of cloud services. Most SMEs therefore refrain from using cloud services because legal consequences are not assessable.

The in-depth discussion of legal issues, however important it may be in practical terms, is beyond the scope of this paper. Instead, we focus on the organizational and technical aspects of our approach.

3 COOPERATIVE COMMUNITY CLOUDS

3.1 Organizational Considerations

A cooperative community cloud focuses on the cooperation of SMEs from a variety of industries. The enterprises organize parts of their IT systems and processes jointly in order to become more flexible and cost-efficient. The service portfolio of the cooperative can be fine-tuned to suit the needs of both the SMEs and their customers. This reduces information asymmetries, creates transparency and reduces the potential of exploitation. It can also lead to more competition in oligopolistic markets (Harris et al., 1996). For the moment, there is a size mismatch between few large CSPs and the comparably tiny SMEs. The cooperative community cloud can create a counterbalance on the market, enabling the SMEs to negotiate with the CSPs.

Cooperatives are distinguished by their governance elements, i. e., structures for incentives, decisions, control and coordination. These elements provide stability and lead to mutual trust, as all cooperatives are designed to be long-term undertakings. This is particularly important in light of current problems of adoption of cloud services in SMEs. There are four elements of cooperative governance that are particularly important for the cloud computing domain:

- the notion of the MemberValue,
- the concept of consistent incentives,
- the systemic trust of a cooperative, and
- its size and locality.

These elements are described in detail in the following subsections.

3.1.1 Strategic Orientation: The MemberValue

The *MemberValue* describes the economic value of the cooperative for its members. As detailed in Section 2.1, only members can purchase the services of the cooperative. This aspect, the *immediate MemberValue*, is the primary motive for joining the cooperative, supported by aspects of sustainable investments and (usually a very small amount of) dividends. Since the access to the services is the primary motive, all cooperatives are not traded on financial markets and are statutorily protected from anonymous capital investors. Thus, negative market influences, such as overreactions, biased valuation, and short-term revenue orientation, are ruled out. Instead, the strategic orientation of a cooperative community cloud is based on sustainable IT operations to the benefit of its members.

3.1.2 Consistent Incentives and Mutual Control

Although the members are in charge of the fundamental strategy of the cooperative, they appoint a professional management that takes care of the operative business. Incentives of both management and members are kept in alignment by a set of mutual controls inherent in the cooperative rules and because of the multitude of roles assumed by each participant in the cooperative (e. g., the members are owners, investors and customers at the same time). The incentives of all parties are thus kept consistent, which automatically leads to more trust between them because property rights are well-defined and, thus, secure (Chaddad and Cook, 2001). With regard to a cooperative cloud, this involves, e. g., revealing inter-organizational interfaces and sensitive information about the company IT. This poses a risk for the SMEs that would usually require mitigation by elaborate legal frameworks. In case of a cooperative, the nature and extend of the cooperation is already agreed upon and regulated by the statutes.

An important particularity of a cooperative cloud is that each member has exactly one vote (*one-man-one-vote-principle*). This favors smaller enterprises which are usually not in a position to negotiate with a CSP due to the size mismatch (Harris et al., 1996).

3.1.3 Systemic Trust

It is not necessary to use artificial mechanisms to build trust into a cooperative. Instead, a functioning, economically successful cooperative will bring this about by itself due to the aforementioned incentives, the statutory regulations and the mutual controls. These mechanisms can be complemented by instruments that provide additional information. In the case of a cooperative cloud, this could be a dashboard showing the exact whereabouts of the users' data. Given this situation, it is obvious that a cooperative cloud can realize a level of trust that would otherwise require comprehensive contracts and service level agreements (SLAs), thus reducing both complexity and transaction cost between CSP and members.

3.1.4 Size and Locality

Being a member of a cooperative cloud is of high importance particularly to smaller SMEs. The combination of a variety of smaller units results in a virtual size that puts SMEs into a much more advantageous position for negotiations with other CSPs. It also allows for economies of scale with respect to bundled individual needs of the members. For example, SMEs often need to buy expensive specialized software for some tasks even though the software is used only rarely over the year. A cooperative cloud may provide a suitable means of sharing such software licenses to the benefit of all members. The realization of economies of scale, scope and skills notwithstanding, the customer-orientation can still be provided as before due to each member's local anchoring.

3.2 Technical Considerations

As stated above, the abstraction from the physical implementation of a service is one important advantage of the cloud paradigm. However, in certain cases—particularly in order to verify what processing and transmission steps are actually taking place—it is necessary to have an "unclouded" view on things. We argue that it is necessary to strike a balance between abstraction when it's beneficial and insight when it's required.

In order to provide such insight, it is necessary to label one's data in some manner in order to be able to identify and track it. The solution to this problem is divided in two parts:

1. The design of a means of labeling, identifying and tracking data, along with a system that allows the user to see where the data resides, in what state it is (e. g., encrypted), and what processing is being performed on it.
2. The design of a means of specifying what processing steps are allowed and disallowed for certain subsets of the data, along with a system run by the CSP that ensures these specifications are followed while the data resides in the cloud.

3.2.1 Identifying and Tracking Data

The first step is to identify the correct level of abstraction that allows both the customer to gain meaningful insights into the cloud internals and the provider to keep the infrastructure flexible. At first glance, this seems to be contradictory to the basic notion of cloud services. This is true most of the time. However, we argue that in certain "exceptional" situations, the user will want to dig into the details, in which case the provider may be able to generate the required visualization for the specific situation despite the abstraction usually enforced.

For a meaningful investigation, the cooperative cloud needs to provide a means of visualizing the relevant information in an aggregated way, e. g., a dashboard. The whereabouts of one's data are shown on a "data map" which offers the ability to drill-down into a more detailed view of the information. The data map also manifests additional meta-data that the CSP generated, such as billing information, usage statistics, or indexes. Of course, these meta-data must be treated with the same care (and under similar policies) as the original data because they, too, can contain very sensitive information (e. g., information giving clues about the current business situation of the cloud user). Even more problems arise when data is forwarded to third-party providers outside of the scope of the cooperative cloud.

When designing the approach, the overall security concept must already be taken into consideration. As the data and meta-data that are handled and visualized in the cloud dashboard and data map may be highly sensitive, there must be a fine-grained access control mechanism to protect them.

3.2.2 Data Processing Policies

In order to have a viable solution for data processing policies, the CSP needs to be able to describe the physical systems in a way that allows varying levels of abstraction while still retaining freedom over the design of the IT landscape. One approach might be designed based the spheres of control (SOCs) as outlined in Section 2.2. The CSP could then provide a somewhat abstract description of the cloud systems. The cloud user, in turn, specifies processing policies based on this description. Combined with the aforementioned data tracking, the CSP can design the cooperative cloud such that it automatically complies with the policies.

4 CONCLUSIONS & OUTLOOK

One major show-stopper for many SMEs that consider using cloud services is the lack of trust. On the one hand, this concerns data protection and data security issues. On the other hand, this also concerns doubts about the legal situation surrounding cloud computing. The lack of trust is at least as much a social and organizational problem as it is a technical one. Therefore, we suggest to tackle this approach on all three facets: organizational, technical, and legal. A cooperative community cloud seems to be a promising approach to establish some of the required trust on an organizational level, also contributing a good deal of the legal certainty.

The general notion of a cooperative has been around for many years and is well understood. However, the specific application of the concept to the domain of cloud computing has not yet been researched. First, organizing a community cloud as a cooperative requires the careful design of the "basics", e. g., suitable statutes, optimal management structures, and a viable set of checks and balances. More importantly, however, the processes within the cooperative and among the participating SMEs need to be designed and optimized. In this regard, there is much to be researched, such as possible synergies stemming from shared use of resources (e. g., software) and common processes (e. g., joint IT procurement).

From a technical point of view, the question is how to provide sufficient means of control to the cloud users. While some trust in the CSP certainly stems from the fact that each SME is also owner of the cooperative, technical measures are still required to provide a more comprehensive overview of the cloud internals. A cloud dashboard with a data map seems to be an appealing solution. This, however, would require an approach of tracking one's data in the cloud, particularly also with respect to associated meta-data. Such an approach is not yet available and it is not obvious how one might tackle this problem.

Nevertheless, we are convinced that a cooperative community cloud is an appropriate way to organize the IT for SMEs. As it does not set any focus on a particular business area nor does it require a certain company size, SMEs from a variety of industries and of diverse sizes can cooperate to form such a cloud. The cooperative approach fosters trust among the members and allows them to realize the cost-benefits of cloud services that are so far not available to them. We are currently working on a prototypical development of a community cloud in order to be able to substantiate our claims even further.

REFERENCES

Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., and Zaharia, M. (2010). A view of cloud computing. *Commun. ACM*, 53:50–58.

Chaddad, F. R. and Cook, M. L. (2001). Understanding new cooperative models: An ownership-control rights typology. *Review of Agricultural Economics*, 26:348–360.

Davies, Jr., C. T. (1973). Recovery semantics for a db/dc system. In *Proceedings of the ACM annual conference*, pages 136–141, New York, NY, USA. ACM.

Davies, Jr., C. T. (1978). Data processing spheres of control. *IBM Systems Journal*, 17(2):179–198.

Gulden, H. (2010). DATEV eG – Designing Future (German). Gemeinsam, IfG-Newsletter, Münster.

Harris, A., Stefanson, B., and Fulton, M. (1996). New generation cooperatives and cooperative theory. *Journal of Cooperatives*, 11:15–28.

Haselmann, T. and Vossen, G. (2011). Empirical assessment of the attitude towards software-as-a-service in small and medium enterprises. Working Paper (to appear), European Research Center for Information Systems, Münster.

MacPherson, I. (1995). *Co-operative Principles for the 21st Century*. International Co-operative Alliance, Geneva.

Mather, T., Kumaraswamy, S., and Latif, S. (2009). *Cloud Security and Privacy*. O'Reilly Media.

Mell, P. and Grance, T. (2009). The NIST Definition of Cloud Computing V15. Technical report, National Institute of Standards and Technology (NIST). online.

Schuster, F. and Reichl, W. (2010). Cloud Computing & SaaS: What are the actually new questions? (German). *Computer und Recht*, (1):38–43.

Theurl, T. (2005a). Cooperative Membership and MemberValue (German). *Zeitschrift für das gesamte Genossenschaftswesen (zfgG)*, 55:136–145.

Theurl, T. (2005b). *Economics of Interfirm Networks*, chapter From Corporate to Cooperative Governance, pages 149–192. Ökonomik der Kooperation. Mohr Siebeck, Tübingen.

Theurl, T. (2009). *Genossenschaften zwischen Innovation und Tradition – Festschrift für Verbandspräsident Erwin Kuhn*, chapter Cooperative Governance and MemberValue: Contents and Communication (German), pages 95–116. Forschungsstelle für Genossenschaftswesen, Stuttgart-Hohenheim.

Theurl, T. and Meyer, E. C., editors (2005). *Strategies for Cooperation*. Shaker, Aachen.

Theurl, T. and Schweinsberg, A. (2004). *New cooperative economy – Modern structures for cooperative governance (German)*. Ökonomik der Kooperation. Mohr Siebeck, Tübingen.

Vaquero, L. M., Rodero-Merino, L., Caceres, J., and Lindner, M. (2009). A break in the clouds: Towards a cloud definition. *Computer Communication Review*, 39(1):50–55.

Weichert, T. (2010). Cloud Computing and Data Protection (German). *Zeitschrift für Datenschutz und Datensicherheit*, 34(10):679–687.

Williamson, O. (2005). *Economics of Interfirm Networks*, chapter Networks - Organizational Solutions to future challenges, page 3–28. Ökonomik der Kooperation. Mohr Siebeck, Tübingen.

Wood, T., Gerber, A., Ramakrishnan, K. K., Shenoy, P., and der Merwe, J. V. (2009). The case for enterprise-ready virtual private clouds. In *Proceedings of the 2009 conference on Hot topics in cloud computing*, HotCloud'09, Berkeley, CA, USA. USENIX Association.