Available online at www.sciencedirect.com

ScienceDirect

Procedia Engineering 100 (2015) 574 – 583

Procedia Engineering

www.elsevier.com/locate/procedia

25th DAAAM International Symposium on Intelligent Manufacturing and Automation, DAAAM 2014

Secure Implementation of an On-Premises Cloud Storage Service for Small and Medium-Sized Enterprises

Bernd Gastermann[a]*, Markus Stopper[a], Anja Kossik[b], Branko Katalinic[c]

[a] *MKW® Austria, Industrial Research Center, Jutogasse 3, 4675, Weibern, Austria*
[b] *Promise IBC s.r.o, Slánska 20A, 080 06, Prešov, Slovakia*
[c] *Vienna University of Technology, Karlsplatz 13, 1040, Vienna, Austria*

Abstract

The purpose of this paper is to demonstrate parts of the implementation of a secure cloud storage solution for small and medium-sized enterprises (SMEs). It shows by example of an Austrian SME how secure cloud-based storage services can be applied in practice. In order to give an introduction to the topic, the paper first discusses some basic principles of cloud storage. The core of this paper eventually focuses on a cloud storage solution considered to be optimal for SME purposes and exemplifies how it could be implemented within such a company considering its particular requirements. As the primary focus of this work is on security aspects, it discusses how the system's attack surface is minimized and necessary software components are hardened in order to defy network-based attacks. The paper concludes with a discussion on various data security and encryption mechanisms.

Peer-review under responsibility of DAAAM International Vienna

Keywords: Cloud Storage; Data Security Aspects; OwnCloud; SME; Windows

1. Introduction

Nowadays, employees of an enterprise can easily work outside their office space or even from home, when applying the latest information technologies. But all enterprises that offer working opportunities outside of a controlled company environment are confronted with the difficult task to efficiently and easily provide their employees with the necessary documents. Simultaneously they need to ensure that all data are securely stored and

* Corresponding author. Tel.: +43-664-6111407
E-mail address: bernd.gastermann@mkw.at

Peer-review under responsibility of DAAAM International Vienna
doi:10.1016/j.proeng.2015.01.407

that external third parties will have no access to potentially sensitive documents. Especially in small and medium-sized enterprises (SMEs), the awareness, the resources, or even the knowledge of how to sufficiently protect these data in such situations are frequently not available. It might even be possible that the field staff itself is responsible for an adequate administration of the necessary company documents. In this case they also willingly resort to the broad range of available public cloud storage solutions. The advantage is self-evident: existing services are characterized by high availability, they are integrated into a lot of popular products and, above all, they are easy to use. In case no regulations exist on the part of the company, field staff will most likely use such services for future business applications that previously could have been adopted for private purposes.

Recently a new trend towards cloud solutions became apparent in the IT-world, which was also shown in a study on usage and distribution of cloud computing in Germany [8]. Many renowned companies promote their cloud services and integrate them deeply into their products, which causes more and more data to be transferred into the cloud. Eventually, in a worst case scenario one could completely lose control over data that have been stored there. When Edward Snowden, the former US intelligence assistant, exposed numerous top secret surveillance programs and released respective information to the public, it became evident by then that private as well as business data require special protection in the internet. Regarding these revelations it is especially critical that most of the providers offering cloud services are US-based and are therefore not subject to the same high standards for data security as European companies.

The disclosure of intelligence programs by Snowden caused an upset regarding data security and increased skepticism towards cloud services as already shown by several studies [7,8]. As the awareness of Austrian SMEs regarding data security increased because of this incident, one question is asked more and more frequently: How can the data security of cloud services be improved? This paper will focus on said quintessential question. Within this context, a company needs to evaluate in how far present dependencies of field staff on public or US-based third-party cloud services can be reduced and how consequently an equally convenient but more controllable and, above all, secure replacement can be provided. The paper will thus propose a way on how to implement a private and secure cloud storage service on existing hardware infrastructure and will eventually discuss various aspects on how to improve the general security of this service.

2. Cloud storage basics

Before this paper can address the implementation of a cloud storage solution suited for SMEs, some cloud related terminology needs to be discussed in advance. Therefore, this chapter provides the basis for the comprehension of subsequent sections and the knowledge of the following explanations will be assumed hereafter. Apart from basic definition of terms, relevant deployment and service models will be discussed.

2.1. Terminology

The origin of the IT-related term “cloud” can be attributed to the cloud symbol that is frequently used in network diagrams. There, it metaphorically represents whole networks or network parts, which are either too spacious to be depicted in detail or which are of unknown structure or irrelevant to the diagram, respectively. Consequently, as an example, the internet is frequently represented by the cloud in network plans.

When talking about the cloud, this usually refers to the computing operations taking place in this “network cloud”. The term “cloud computing” subsumes multiple transparent and scalable IT aspects like processing power, storage capacity but also software applications that are provided by an abstract service situated on a few specialized locations (i.e. data centers) and are then distributed over a network to many individual customers or computers, respectively. It is the most obvious advantage of this solution that, within a cloud service, complex computing tasks are distributed to many high-performance machines, big data volumes are transferred to specialized high-capacity data centers, and software can be provided and administrated in a centralized way. All this usually happens completely transparent for the end-user. Transparency in this context means that no conclusions can be drawn about the number of computers performing a certain task or the physical location of particular data records. While cloud services aim to separate applications from subjacent infrastructure or mechanisms, customers are to be provided with improved availability and cooperation as well as a potential cost reduction by optimization and increased efficiency.

The term "cloud storage" describes a specific area of application of cloud computing and specifically addresses the storage of digital data in the cloud. A central aspect is data hosting, which is typically performed by cloud storage providers, who own and administrate the necessary physical infrastructure. The supplier and operator of a cloud storage service is responsible for keeping the data available at any time and for protecting them from unauthorized access. In turn, customers have the possibility to buy or rent storage capacity from such a provider. The acquired storage space can typically be used for all kinds of data, irrespective of their format or intended purpose [6]. The mode of access to the cloud storage is dependent on the provider. Either a standardized data transfer protocol is used or a proprietary application programming interface (API) is provided [1,6].

Cloud storage services usually employ a vast logical pool of virtualized hard drives, which – in order to increase their availability – are redundantly distributed over various server and data centers [6]. Additionally, this can result in the worldwide distribution of uploaded customer data over multiple physical hard drives and data centers in case no exceptions have previously been negotiated between customer and service provider. This physical distribution typically happens in a user-transparent way [6]. On the other hand, this virtualized and distributed architecture has many advantages for the operator as it increases the system stability and cost efficiency as well as the elasticity and scalability of the whole service [2]. Storage capacity, for example, can be added or removed upon demand, without any external evidence for the end-user. The redundancy and distribution of infrastructure and data also has a positive influence on the general fault tolerance and availability of the cloud storage service.

2.2. Cloud service models

Cloud services can provide a wide range of IT services. As previously discussed, these can comprise computing power or time, storage space or software applications. Such services can be abstracted by a layer model – the so called cloud service stack. This model allows to divide various cloud-based services into service levels in order to distinguish them technically. The resulting taxonomy basically differentiates between three levels building upon each other [4,9]. The base and thereby lowest level of this service model comprises the physical infrastructure (i.e. network and server infrastructure). Building upon this layer follows the logical platform level (operating system, database, APIs, etc.). The last layer with the highest abstraction level comprises the software (applications like document management, calendar, text processing, etc.). A cloud service can be considered more specific, the higher it is located in the layer model. Because of these three levels (software, platform, infrastructure) this model is also called the SPI model [4]. The following Fig. 1 visually depicts the typical layout of the SPI model including the elements of each respective layer as well as the administrative responsibilities of provider and customer.

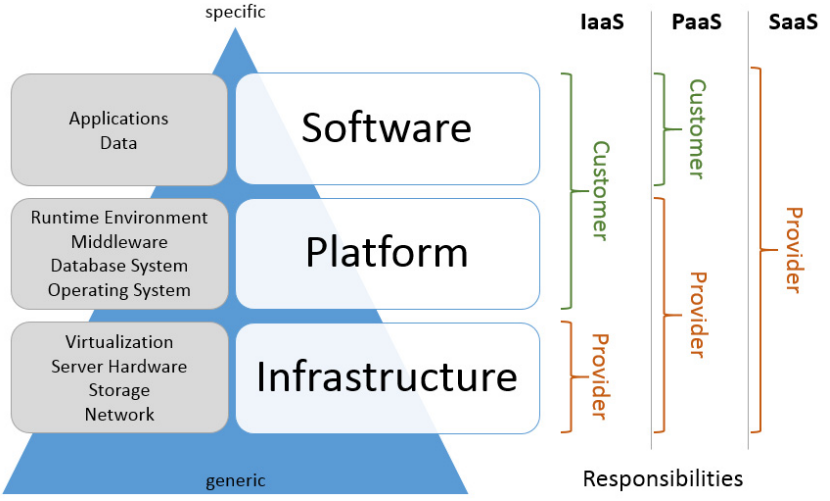

Fig. 1. Layers of the cloud service model and their respective responsibilities.

The IT services regularly provided by a cloud service can typically be assigned to one of these three levels. In the literature the features and resources that are provided in each level are labelled with the addition „as-a-Service“, leading to the following terms for the three levels of the SPI model [9]: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS).

The lower a specific cloud service is located in the service model, the more universal the area of application and the bigger the possible influence of the customer. This is understandable insofar, as the lowest infrastructure level is only related to hardware and the design of the upper levels is therefore left to the customer. The customer can now freely implement additional platforms and applications on the provided hardware. On the higher levels this creative leeway is increasingly restricted but they address other target groups with different technical requirements instead.

2.3. Cloud deployment models

As the service models examined the more technical aspects of cloud computing, an organizational view of the concept follows. In this section possible deployment models of the cloud are introduced. They show in which way cloud services can be provided to the consumers and therefore mainly describe the relationship between customer and cloud service. In most cases it can be differentiated between two types: public and private clouds. Other types like community clouds or hybrid clouds also exist but they are less relevant in the context of this paper [9].

The public cloud is a service that is made available by a provider via internet and addresses an unlimited group of customers. This open model represents the classical and initial concept of cloud computing, in which a service can be used by everybody with corresponding demand. The customer base for such a public cloud service are private users as well as companies or academic institutions. As public cloud services are usually operated by providers that are not under the customer’s influence and are almost exclusively reachable via the internet, problems in regard to data protection and security can rapidly arise [4]. Depending on the type of service and the features the service offers, potential security issues regarding this model can be more or less pronounced.

The counterpart to the public cloud is the private cloud. Services of this model are not publicly available, but are at the exclusive disposal of a relatively small user group, like the employees of a specific company. Such private cloud services can either be run by external service providers or by the customers themselves [3]. For instance, a company could make use of the public IaaS offering of a third-party and run their own private cloud service based on the provided infrastructure. This approach allows the reduction of investment in infrastructure and increases scaling potential. The private cloud service together with the responsibility for data privacy remains under the control of the using organization, thereby increasing the service’s design options. Individual requirements regarding functionality, compatibility, and security can be met more easily, especially if the service is operated on one’s own hardware independent of any third-party providers. On the other hand, this approach is accompanied by increasing efforts, costs, as well as responsibilities, which are necessary for its smooth and secure operation.

The different types of cloud services presented so far generally do not specify under which sphere of influence they are operated. From the customer point of view, for example, private clouds can either be localized in a foreign data center or on the customer’s local infrastructure. In order to differentiate where the cloud service is hosted, the terms “on-premises” and “off-premises” are used, where “on-premises” refers to the traditional realization of a cloud service on internal infrastructure owned by the customer [5]. Contrary, the term “off-premises” denotes that the service is hosted outside of the customer’s environment, which is mostly the case for public clouds hosted directly by a third-party service provider. As the on-premises solution requires own infrastructure, this type of service is understandably related to higher investments.

3. Implementing secure corporate cloud storage

SMEs generally have two options for implementing a cloud storage solution in the business environment: They can either use one of the numerous off-premises services that are readily available on the market or establish an appropriate cloud service platform within the company on internal hardware (i.e. an on-premises solution). However, as already stated at the beginning of this paper, the use of third-party cloud services may not always be the best choice when it comes to data security and privacy protection [5]. Storing sensitive data on a cloud storage service requires it to be both controllable and secure. Therefore, companies need to evaluate in depth, which of the

available cloud storage services fit their individual requirements best. This third chapter of the paper will use an Austrian SME as an example and will demonstrate one possible way on how to implement a secure corporate cloud storage solution within this company.

Nowadays, a multitude of cloud storage solutions exist, ranging from generic IaaS-based services to more specific SaaS services. Before one of them can eventually be adopted, SMEs need to identify the most suitable cloud storage solutions based on their respective data storage requirements. However, as the definition of SMEs comprises an extensive range of different companies, their respective requirements for a cloud storage solution can vary considerably. It is therefore not feasible to give a universal recommendation for specific cloud storage solutions in the context of this paper. This is why the following assessments will be based on said exemplary Austrian SME.

Given the characteristics of IaaS and PaaS services, SMEs can usually refrain from adopting one of these solutions in most cases, as they are typically associated with major administrative workload as well as the development or installation of higher software levels. Therefore, the choice is mostly limited to the SaaS service model, which is based on predetermined hard- and software provided by a service operator [4].

In order to gain deeper insight into the application of one of the available cloud services, a specific product was singled out from the available SaaS services in agreement with the company where the cloud storage solution was implemented. Based on this selected product a possible approach can be demonstrated, how the secure operation of cloud storage in a company environment can be achieved. Considering the specific company requirements and conditions, a cloud software platform marketed under the name “ownCloud” was selected and implemented on already existing on-premises infrastructure. Hence, the following chapter focuses on ownCloud and sheds a light on various security-related aspects. As a software solution is only as secure as the environment in which it is applied, not only ownCloud but the entire subjacent software platform serving as the basis for its operation are subsequently considered.

3.1. Company requirements for cloud storage solutions

In order to avoid the problems outlined in the introduction of this paper, and to provide field staff with a standardized and secure option for the storage of company data, an appropriate storage system was to be implemented. This section concentrates on the specific company requirements leading to the decision to use ownCloud as software platform for the corporate cloud storage service. The following requirements were identified:

Functionality: The functional requirements of the discussed company were mainly limited to data storage, but the option for automatic synchronization of files on multiple devices for offline availability was considered important. The synchronization was supposed to happen transparently for the user and to be possible within the company network as well as via the internet. Other functionalities exceeding pure data management like calendar or contact management have not been required. An additional criterion was the support of at least 150 users and a central administration. For this purpose, the support of professional and high-capacity databases is a prerequisite.

Security: Naturally, the solution has to support basic security features like access management for single users. As security is an important issue, the encryption of data during transfer by means of a secure protocol is indispensable. Encryption of server-side data-at-rest, however, does not represent a crucial criterion.

Costs: For many companies, especially those of smaller size, product costs should not be disregarded. In general, the one-time and running expenses are supposed to remain within certain defined boundaries, which will not be subject to further discussion herein. Cloud solutions that are free of charge are preferred.

Provider Independency: Additional prerequisites for the selection of a storage solution are the independency from third-party providers and the free choice of the data storage location, respectively. In order to lower the risk of losing data control, and as the infrastructure as well as available IT resources necessary for operation are already available, the storage service should intentionally be installed and run on-premises.

Platform Support: As a major part of the existing software infrastructure is based on Microsoft technology, it is essential that the implemented cloud platform supports the Windows operating system (OS) on the client as well as on the server. Therefore, the server software needs to run at least on Windows Server 2008 R2, and client synchronization programs have to support Windows 7 or higher. Apart from that, data access has to be possible via a current web browser irrespective of the end device. Common smartphones and tablets also need to be supported.

Expandability: Finally, the general expandability of the applied solution is of major concern. This is related to the hardware as well as the software level. Accordingly, storage capacity has to be expandable at will, just like the provided basic functionality, which can be upgraded via plug-ins. Consequently, open interfaces (i.e. APIs) have to be available that can be used to add further functionality upon request. The support of the open and standardized WebDAV protocol is expected in this regard.

3.2. System architecture overview

Given these previously introduced requirements, an on-premises cloud storage solution running on the company's own hardware needs to be established. Consequently, a suitable cloud storage software is necessary that provides required functionalities. For this purpose, the freely available community edition of the open-source software called ownCloud was selected with regard to the company's requirements. The ownCloud software is web application that runs in a web server on top of existing on-premises hardware and is based on the platform independent PHP scripting language. The operation of ownCloud requires an operating system, a web server, a PHP runtime environment, as well as a suitable object-relational database management system (DBMS) used to store various application data like user accounts and settings.

The key component of the available hardware infrastructure are multiple physical servers that are combined to form a logical cluster used for virtualization. This cluster allows for several virtual machines (VM) to run in parallel and independent of the underlying physical hardware, thereby enabling efficient use of the resources. The server on which the ownCloud-based cloud storage solution is implemented is operated as such a VM, simultaneously running with other unrelated VMs on abovementioned server cluster. This cluster is designed redundantly, in order to compensate for malfunction of individual hardware components and to enable dynamic load balancing. The storage system, on which the individual VMs as well as the user data uploaded into the cloud are stored, displays a similar redundancy. Both systems are connected to an uninterruptible power supply (UPS), which can sustain the operation for a certain amount of time even after a power outage. Further technical details are not discussed here, as these specifications are irrelevant for the basic implementation of the cloud storage service.

In regard to the network environment the ownCloud VM is located in a so called demilitarized zone (DMZ), which, due to security reasons, separates publicly available IT services from the remaining and more sensitive network areas. A central firewall isolates the DMZ from these underlying company networks.

3.3. Reducing the system's attack surface during installation and configuration

A prerequisite for the selection of an appropriate cloud software was its ability to run under the Windows OS. This requirement was defined in order to keep the company's current system landscape mostly homogenous, and to keep administrative workload to a minimum. Therefore, a Windows OS has to be set up within the cloud server VM. The OS provides the basis for the further installation and configuration of the cloud system. Already with this first installation step it is recommended to deliberately consider how the potential attack surface of the system can be minimized. The IT security term „attack surface" denotes the entirety of all services that are exposed to the outside of a system and are therefore potential targets for an attack. The more services are targetable from the outside, the higher the attack surface of a system and the higher the probability of success for such an attack. Therefore, this surface is to be minimized by removing non-essential services, which not only leads to an increase of security, but also reduces the system's overall complexity and its consumption of resources [4].

Considering this security aspect, a minimal installation of the software components required for the operation of ownCloud is reasonable. Consequently, the so called core version of Windows Web Server 2008 R2 is applied. In Windows Server Core many features that aren't typically needed for server systems, like end-user applications and services as well as most of the graphical user interface (GUI) have been removed. Thus, administration of the server is primarily performed via the internal command-line interface (CLI) or by various tools that can remotely connect to the server via the internal network, respectively. The reduction of the attack surface is not only crucial for the OS but for all installed software components. All components that are essential for the operation of ownCloud have to be securely configured in this regard. These components comprise a web server, the PHP application library and a DBMS. The Windows-integrated Internet Information Services (IIS) has been installed as a web server and MySQL

as a database system. Each of these components must be hardened against network-based attacks through diligent configuration. Due to the complexity and scope of this topic, the installation and configuration steps of the individual components will not be discussed in this paper. Further information regarding the secure configuration of each of these components can be found in their respective product manuals or other technical literature. As soon as the OS, IIS, PHP and MySQL have been installed and secured, ownCloud can be deployed to the VM.

4. Protecting data on client, server, and during transport

This main part of the paper provides an in depth discussion of three security relevant aspects of the presented cloud storage service. It will give detailed insight into additional technical measures that can improve data security. Three areas can be identified, where such measures can be applied: the protection of data-at-rest on the server, the encryption of data during transfer, and the client-side encryption of data directly on the computer of the user before they are eventually transferred to the server.

4.1. Server-side data protection

The first data protection measure solely and directly addresses data on the cloud server itself. The server represents the central component of the system and is responsible for the storage of all uploaded user data. In this part of the paper, server-side data protection primarily refers to ensure confidentiality. Other aspects of data protection on the server comprise hardware redundancy and the physical protection of IT infrastructure but these areas are not taken into consideration in the scope of this paper.

As data are usually filed on the server as plain text, they can – under certain circumstances – theoretically be read out and manipulated by an attacker. However, server-side encryption of stored user data represents one possible countermeasure. If data are stored in cipher text, this will increase the effort for any entity with malicious intent to remain unnoticed when reading out or modifying these data. Server-side encryption is typically transparent for the user, which means that the user will not notice any difference to unencrypted data storage while using the service. In case documents are uploaded, the system takes care that they are automatically encrypted and stored as such on the storage medium. In the opposite direction, when documents are retrieved, the system reads out the encrypted data from the hard drive, decrypts them with the respective key, and then sends the document back to the user. This example also demonstrates a potential weakness of the system: key and encrypted data are both accessible by the server. Users typically do not have any control over the key or the applied cryptographic method and are therefore dependent on the confidentiality and reliability of the service and its administrative environment. From the perspective of the user, this is related to the potential risk that administrators or other persons with respective authorization can retrieve data from the server before they are even encrypted or are able to decode them by use of a system-wide key. Furthermore, insufficiently secure cryptographic methods could be applied without the knowledge of the user. With regard to server-side encryption, the difference has to be made whether it was performed with an individual user-specific key or a universal key valid for all filed data, respectively. Due to several security related reasons, the former would be preferable because many individual keys increase the effort for the attacker and could also prevent compromise of the entire system in case a key is falling into the wrong hands.

Several methods for the implementation of server-side encryption exist. One option could be the use of third-party software or of features provided by the OS, such as “BitLocker” or the built-in encryption mechanism of the NTFS file system. Both techniques address a level close to the hardware so that the protection is directed primarily against the readout by attackers with direct hardware access or other unauthorized user accounts at OS level. However, this protection is turned ineffective in case the attacker is able to log in via the system’s administrator account. Thus, a different approach is the use of encryption on the application level. An optional ownCloud extension is specifically intended for this purpose, allowing server-side encryption on application level. There are no functional limitations related to the activation of this feature, but the encryption is restricted to file contents only, while file names remain universally readable. Depending on the activation of other features that could use various caching techniques, it can occur that not all information gets properly encrypted. A search index used to accelerate file content look-ups, for example, could still contain plaintext contents of encrypted files.

It must be noted, however, that the application of server-side encryption inside a controllable company environment does not always provide a clear benefit. In case data (like private data) deserve special protection and should remain inaccessible even to the system administrator, it can be feasible to make use of this approach. If the use is restricted to business documents that are potentially freely accessible for the administrator or other users on other network drives, server-side encryption is of no benefit as it also increases the overall administration effort and represents an additional source of error. In a worst case scenario, where no documented error recovery method for encrypted data was defined, data can potentially get lost or the import of an earlier backup can be complicated.

4.2. Transport encryption

In this regard, transport encryption is a completely different issue. The protection of data during transfer is absolutely indispensable in any case. This is caused by the fact that the cloud storage service usually has to be accessible not only via the internal company network but also via public networks like the internet. Access via public networks is always associated with the risk that data are intercepted or read out during transfer. This issue is mitigated by Hypertext Transfer Protocol Secure (HTTPS), which is very similar to conventional HTTP but uses Secure Socket Layer (SSL) or Transport Layer Security (TLS) for authentication and encryption of communication. Thereby it is of significance, which of the two encryption protocols has been applied. For security reasons the use of TLS in version 1.1 or 1.2 in combination with modern browsers is advisable. Older protocol versions like TLS 1.0 and especially its antecessor SSL 3.0 should only be used for compatibility reasons. HTTPS is thought to ensure that the counterpart is really the computer it is supposed to be, and that sensitive data are not easily read along by others.

The configuration of the web server to use HTTPS is trivial, but it takes a valid digital certificate issued by a certificate authority (CA). Provided that such a certificate is available, the web server is configured to use the HTTPS protocol in just a few steps. Even though communication is already basically protected by this measure, further adaptations like the prioritization of the system's preferred cipher suites are still recommended. Cipher suites describe a standardized set of cryptographic algorithms, which are used for different purposes during the communication by SSL/TLS, respectively. A suite comprises algorithms related to four areas: key exchange, digital signature, hashing functions, and data encryption.

A multitude of such cipher suites with different characteristics do exist. The broad range of suites is mainly due to the fact that some of them ensure compatibility with older browser versions, and that others apply algorithms, which are more secure, faster, and less prone to certain attacks than the rest. A feature that gets more and more attention from cloud services is "perfect forward secrecy" (PFS). It describes a quality of key exchange protocols, which are able to guarantee that individual session keys are not affected even if the private long-term keys, from which individual session keys are derived, have been compromised. In other words: Even compromised long-term keys do not provide attackers with enough information regarding the disclosure of individual session keys – these have to be attacked directly and will, at worst, only disclose the respective session.

As previously stated, cipher suites can be distinguished by specific features like speed, resource consumption, and security, but not all of them offer PFS. On the other hand, not all available cipher suites are supported by all OS or web servers. Therefore, they provide a list of supported cipher suites as well as a default priority list that defines, which of them should be applied for client communication. Client and server negotiate the optimum set based on their respective cipher suite preferences, if a match can be found at all. The problem is to find a balance between the support of older browsers and systems on one side, and potentially insecure algorithms on the other. Windows Server 2008 R2 defines a set of default cipher suites, which are active right after the activation of HTTPS. This default list is associated with several problems, though. For example, cipher suites without support for PFS are preferred. Additionally, algorithms are used that are not considered to be sufficiently secure and state-of-the-art by today's standards. Even cipher suites that refrain from using encryption algorithms at all can be found.

In controlled company environments and in the case of a specific application situation like the one discussed in this paper, it is easy to judge, which platform and browser need to be supported by the cloud storage service. The list of cipher suites supported by the ownCloud server can therefore be adjusted accordingly. Knowledge of potential weaknesses or discouraged algorithms can also be considered during this configuration process.

Generally, the application of key exchange protocols based on Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE) algorithms are recommended. Both support PFS but in regard to

performance, ECDHE should be preferred [10]. Concerning compatibility, it is advisable to use cipher suites based on the fast RSA procedure for key exchange, which does not support PFS, however. Regarding the signature algorithm, either RSA or the Elliptic Curve Digital Signature Algorithm (ECDSA) can be applied. Considering hash functions, the newer SHA-2 algorithms should be preferred as SHA-1, according to current technical knowledge, does no longer provide sufficient protection and cannot be judged resistant enough against collision attacks.

4.3. Client-side encryption

Client-side encryption offers a very high level of protection as user data are encrypted directly on the computer of the user, before they are uploaded to the server. Ideally, it allows countering the potential security issues of the two aforementioned approaches. The confidentiality of user data is therefore solely related to the confidentiality of the client computer and to the security of the used encryption algorithm, because the actual data remain encrypted even when the transport encryption is circumvented. The same is valid for data-at-rest, as neither the administrator nor other persons usually have the possibility to decode user data as long as the applied key is held private. As such a distinctive feature of client-side encryption can be pointed out: the administration of the secret key is the exclusive responsibility of the user, who has to ensure that the key is not getting compromised. This approach therefore provides data confidentiality and privacy, which is basically independent of the encryption applied during transport or on the server.

Despite these advantages, client-side encryption is considered less suited for company environments. Considering the exemplary company discussed in this paper, no sensitive or private documents are stored on the cloud server that have to be shielded from the view of the administrator. Additionally, this approach complicates assistance in case of problems or support requests. Apart from that, client-side encryption leads to extensive functional consequences in connection with ownCloud: as the encryption is mostly performed with symmetric encryption algorithms – due to reasons like performance and practicability – the ownCloud feature for document sharing loses its practical function. If encrypted data were shared with others, the private symmetric key would have to be passed down, which in turn would inevitably lead to other security problems. Data access via the web interface of ownCloud would also be complicated, as data would have to be manually decrypted after download and manually encrypted before upload using an appropriate tool. In a worst case scenario, this approach could also lead to complete data loss, if users lose their private key or forget their password, as the administrator would be unable to restore access. In total, all these reasons make client-side encryption unattractive for most business environments.

5. Conclusion and outlook

This work is intended to give an overview over a multitude of different security aspects that have to be considered when implementing a secure cloud storage service for SMEs. In the course of this paper, various security related recommendations were provided. SMEs planning to implement a cloud storage solution within their business environment first need to consider which type of service fits their requirements best. As there are a multitude of requirements for SMEs, no universally applicable recommendation can be given here. However, based on the requirements of the company discussed in this paper, SaaS-based cloud solutions seem generally more appropriate for SMEs regarding costs and administrative effort than platform or infrastructure-based cloud storage services. Another important aspect is the type of deployment. The exemplary company presented here required to establish an on-premises cloud storage service, for which ownCloud was chosen as cloud software platform.

The paper introduced the company requirements that led to the selection of ownCloud and briefly explained how it was implemented on readily available hardware infrastructure. However, the primary focus of the presented cloud storage solution and this paper is on the security of the system. Hence, the overall security of all the applied software components is of utmost importance, as a system is only as secure as the weakest link in the chain of its components. Therefore, reducing the attack surface of the cloud storage service discussed in this paper was essential. According to this intention, the minimalistic and more secure core version of Windows Web Server 2008 R2 was used as the basis for the implemented system. Additionally, it was attempted to harden all other components that are indispensable for the operation of ownCloud against potential network-based attacks.

In the course of the implementation it was also evaluated, which other measures are capable of improving the overall data security. This comprises client and server-side encryption of user data but also the protection of data during transfer. Even though the former seem to be less useful for an on-premises service within a company, the application of HTTPS is an indispensable measure nowadays. Especially considering transport encryption, a major potential for optimization was detected in regard to PFS and potential attack vectors against security protocols.

The cloud storage solution presented herein is not perfect by far and this discussion raises no claims to completeness. In regard to security the whole topic is much too complex, so that only a few important aspects could be outlined. For example, one major limitation of this work is the fact that only software-related security aspects could be covered. However, in order to build a secure system, the whole environment needs to be considered, which includes employees, company guidelines, and other IT systems as well as the physical security of the infrastructure.

In this regard, some weaknesses of the presented solution also need to be mentioned: First of all, the applied version of the operating system is already slightly outdated. Although older software is not insecure as long as it is serviced and supplied with security updates, problems could arise in the medium-term when support for this product is discontinued and security flaws remain untended. Furthermore, the intrinsic problems of the selected on-premises operation are obvious: all deficits discussed so far demonstrate that, apart from the investment in required infrastructure, the operation of one's own cloud storage service is related to continuous administration effort.

Overall, the implemented cloud storage service provides important benefit for the discussed Austrian SME. It provides field staff with an appropriate and secure tool to access and manage business documents from within the company network as well as from any other location with access to the internet. At the same time, the company retains control of all sensitive data as they are stored on hardware within the company on Austrian territory.

Although the implementation of a secure cloud storage system for this company was considered essential and is completed by now, further development of this cloud service has to continue. Routine service is necessary in order to keep system security up to date with the latest attacks. In order to validate the security of the implemented system an extensive penetration test has to be performed by an independent and professional entity. Such a test alone can verify system security on a deeper level and can detect potential vulnerabilities of the software components. Based on the outcome of the penetration test, further actions or research may be required to improve the overall security of the implemented cloud storage system and its IT environment.

References

[1] R. Seiger, S. Groß, A. Schill; *SecCSIE: A Secure Cloud Storage Integrator for Enterprises*; Proceedings of the 13th IEEE Conference on Commerce and Enterprise Computing (CEC); pp. 252-255; ISBN 978-1457715426; DOI 10.1109/CEC.2011.45; Institute of Electrical and Electronics Engineers (IEEE); 2011.

[2] T. Sasidhar, P. K. Illa, S. Kodukula; *A Generalized Cloud Storage Architecture with Backup Technology for any Cloud Storage Providers*; International Journal of Computer Application (IJCA); Volume 2, Issue 2; ISSN 2250-1797; RS Publication; 2012.

[3] G. Kulkarni, R. Waghmare, R. Palwe, V. Waykule, H. Bankar, K. Koli; *Cloud Storage Architecture*; Proceedings of the 7th International Conference on Telecommunication Systems, Services, and Applications (TSSA); pp. 76-81; ISBN 978-1467345491; DOI 10.1109/TSSA.2012.6366026; Institute of Electrical and Electronics Engineers (IEEE); 2012.

[4] R. L. Krutz, R. D. Vines; *Cloud Security: A Comprehensive Guide to Secure Cloud Computing*; ISBN 978-0470589878; Wiley Publishing; Indianapolis, IN, USA; 2010.

[5] T. Erl, Z. Mahmood, R. Puttini; *Cloud Computing: Concepts, Technology & Architecture*; ISBN 978-0133387520; Prentice Hall; Upper Saddle River, NJ, USA, 2013.

[6] M. Borgmann, T. Hahn, M. Herfert, T. Kunz, M. Richter, U. Viebeg, S. Vowé; *On the Security of Cloud Storage Services*; Fraunhofer Institute for Secure Information Technology (SIT); SIT Technical Reports, SIT-TR-2012-001; ISBN 978-3839603918; ISSN 2192-8169; Fraunhofer Verlag; Stuttgart, Germany; 2012.

[7] Microsoft; *Press Conference "New Deal for the Digital World": Location Factor Digital Trust - Solutions for Austria to steer clear from the current confidence crisis*; Web Address: http://www.microsoft.com/de-at/news/Press/2014/Apr14/New-Deal-fur-die-digitale-Welt.aspx; Last Access on 29th September 2014.

[8] KPMG AG, BITKOM; *Cloud-Monitor 2013: Cloud Computing in Germany*; Web Address: http://www.bitkom.org/files/documents/Studie_Cloud_Monitor_sec.pdf; Last Access on 29th September 2014.

[9] S. Celar, Z. Seremet, M. Turic; *Cloud Computing: Definition, Characteristics, Services and Models*; Annals of DAAAM for 2011 & Proceedings of the 22nd International DAAAM Symposium; Volume 22; No. 1; pp. 1-2; ISBN 978-3901509834; ISSN 1726-9679; Editor B. Katalinic; Published by DAAAM International, Vienna, Austria, EU, 2011.

[10] V. Gupta, S. Gupta, S. Chang, D. Stebila; Performance Analysis of Elliptic Curve Cryptography for SSL; Proceedings of the 1st ACM Workshop on Wireless Security (WiSE '02); pp. 87-94; ISBN 1581135858; DOI 10.1145/570681.570691; ACM; New York, NY, USA 2002.