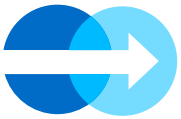

Fraunhofer SIT

Whitepaper

OmniCloud –
The Secure and Flexible Use of Cloud Storage Services

Thomas Kunz, Ruben Wolf

Fraunhofer Institute for Secure
Information Technology SIT
Rheinstrasse 75
64295 Darmstadt
Germany

2014

"'Investment into Your Future"'

The investment for this development was
co-financed by the European Union from the
European Fund for Regional Developments and
by the State of Hesse

Contents

Abstract

Despite the benefits that cloud services offer many companies still have considerable concerns to entrust their business data to a cloud storage service. With OmniCloud Fraunhofer SIT is offering a solution that allows businesses to use cloud storage services in a secure manner. The fundamental idea behind OmniCloud is to connect any kind of application and backup software with any cloud storage service. OmniCloud guarantees the stored data's confidentiality – independent from the actual security mechanisms or current security incidents at the cloud provider. OmniCloud enables businesses to use cloud storage offers securely, and thus to save costs when securing their digital data. OmniCloud also supports companies efficiently in switching over from one cloud provider to another, thus preventing an unwanted commitment to a specific provider.

1. Introduction

In recent years the popularity of cloud storage services has increased very strongly. Services such as Dropbox has more than 100 million registered users that store over 1 billion data on a daily basis (as of November 2012)[Con12]. Cloud storage services are very appealing to companies to store and secure their data externally. Besides the benefits of a professional management and the data's physical security, the available cloud storage space is nearly unlimited and extremely cost-efficient.

But using cloud storage services also harbors risks, especially with regard to the security and confidentiality of the stored data. Concerns regarding availability and loss of data control are being discussed as well. Using cloud storage services requires trusting the cloud provider. Internal security policies blow up due to the cloud providers' additional service level agreements (SLA). Oftentimes it is not easy to determine where and by whom the data is actually being processed and stored. Cloud providers partially use other cloud providers as subcontractors (e. g. Dropbox uses the Amazon S3 service). This may result in users and cloud providers being subjected to different legal conditions, they may not be able to meet compliance demands, that data may not be protected against government access (see "Patriot Act").

Furthermore, the cloud storage service market is still very heterogeneous. New services and more appealing offers emerge, less successful services disappear. On the one hand it is therefore important for enterprises not to be bound permanently to a specific cloud service ("provider lock-in"), on the other hand it should be easy for a customer to change to a different cloud provider, for example in the case of the current cloud provider's insolvency.

With OmniCloud Fraunhofer SIT has designed a solution that takes the special aspects of cloud storage use in enterprises into account and that renders cloud service use secure and flexible. OmniCloud encrypts all data locally before it leaves the company network to be transferred to the cloud storage. The company data remains confidential, independent of the security mechanisms offered by the respective cloud service provider. When using OmniCloud, actual security incidents at the cloud storage provider's will not have any impact or at least no vigorous impact on the company's stored data. At the same time OmniCloud actively supports the company in changing to another cloud provider. The OmniCloud migration service shifts the data efficiently from the old to the new cloud provider.

Currently OmniCloud is available as a concept. Prototypes for some of its functions have been realized and were presented at various events. Due to the big success and the positive feedback of potential users, Fraunhofer SIT plans to develop OmniCloud further into an applicable and marketable product. Fraunhofer SIT is looking for suitable partners from commerce for the development and distribution of OmniCloud.

1.1. Data Storage in the Cloud

Today, a substantial amount of valuable information such as contracts or business plans is available in companies often only in digitalized form. This trend is expected to increase further in the future. In a worst case scenario the irrevocable loss of this data may even result in a company's ruin, so the data has to be protected. Beyond that, companies are legally obligated to keep specific data on file for a certain time, for example tax records.

For these reasons nearly all companies generate backups of their data. To generate correct backups is complex, however: Backups have to be stored physically and spatially separate from the original data, as in the case of a theft or fire damage both the original data and the backup data may be affected, for example. Besides, regularly created backups require a large storage space. Major enterprises may be in a position to build dedicated data centers ("private cloud") for such purposes. For financial reasons this frequently does not constitute a solution for small and medium enterprises (SME). Particularly for these SMEs cloud computing is offering a solution for these problems. Specific cloud storage services promise almost unlimited storage capacity at very competitive prices. The number of such cloud storage service providers is very large. Their data centers are being operated in a professional manner, they are protected physically against theft and damage by fire, and are equipped with the respective hardware to facilitate round the clock operation, including uninterruptible power supplies in order to be protected against power outages as well.

Using cloud storage services as a backup medium thus seems to be the ideal solution, reducing costs while increasing data availability. Yet, many enterprises hesitate to entrust their data to a cloud. They fear that their data is not sufficiently protected: both the cloud provider as well as an attacker may be able to attain confidential data. Furthermore, the data is not protected against access by governments ("US Patriot Act"). When data is entrusted to another enterprise (cloud provider) legal concerns arise as well. Last but not least enterprises fear that they are bound to a specific cloud provider, because often it is not easily possible to change to another provider ("provider lock-in").

1.2. Cloud Services Security

Despite the many benefits that cloud services offer a lot of enterprises are still very concerned to entrust their business data to a cloud storage service. Many security aspects such as data protection, data integrity and availability have to be given careful consideration. Numerous independent studies reached the same conclusion. In 2011[Deu11] Deutsche Telekom interviewed decision makers from various companies on their view about the significance of IT security. The conclusion was that IT security is either extraordinarily significant (67%) or at least very (29%) significant. A CapGemini study concluded that public cloud service usage will not increase dramatically as long as the associated security issues are not being resolved [Cap12]. The Bundesamt für Sicherheit

in der Informationstechnik (BSI - Federal Office for Information Security) identified another obstacle when using cloud services: The loss of control and direct influencing with regard to implementing security mechanisms[BSI12].

In March 2012 Fraunhofer Institute for Secure Information Technology SIT carried out a study [BHH+12] taking a closer look at the security of cloud storage providers. The study analyzed the security of seven cloud storage providers primarily targeting private customers. The study showed that the providers of such cloud storage services actually do realize the significance of IT security, but none of the tested services met the most necessary security requirements even in the slightest. The typical shortcomings detected include

- missing local data encryption,
- problems when releasing files for other persons, for example because unauthorized persons are allowed to access the respective data as well,
- missing e-mail address verification during registration enabling attackers to assume a false identity and register with a service, thereby allowing them to plant malware and carry out data espionage,
- vulnerabilities during registration and login allowing to select weak passwords for the registration or to sniff out the e-mail addresses of already registered users, for example.

2. OmniCloud

With OmniCloud Fraunhofer SIT is providing a solution that allows enterprises to use cloud storage services securely. The basic idea behind OmniCloud is to connect any application or backup software with any cloud storage service. The data is stored in such a manner that its confidentiality is protected, independent of the cloud provider's actual security mechanisms. With this Fraunhofer SIT reacts to the above mentioned concerns regarding cloud storage usage and the demands concerning secure cloud storage expressed especially by small and medium enterprises. OmniCloud enables enterprises to use cloud storage services in a secure manner while saving costs when securing their digital data.

Objectives and Benefits

1. **OmniCloud makes cloud storage more secure.** OmniCloud encrypts all files locally in the enterprise network before they are transmitted into the cloud. A seperate encryption key is used for each file. In addition to the actual file content

all file names and directory structures are concealed as well. Users will be authenticated before they can access OmniCloud. A fine-grained access control mechanism defines which user may access which files.

2. **OmniCloud makes software cloud ready.** OmniCloud serves as an adapter that makes any type of application or backup software cloud ready. To achieve this OmniCloud has a series of widespread standard communication interfaces that are supported by all the customary operating systems and software products. Installing OmniCloud on the individual enddevices is unnecessary. OmniCloud is presented to its users like a network drive or an FTP server, to be used either directly for storage or to be assigned to a drive letter.

3. **OmniCloud prevents cloud provider lock-in.** OmniCloud prevents customers from being dependent on one specific cloud storage provider. Migrating their stored data typically represents a big obstacle for enterprises. OmniCloud allows data to be moved easily and quickly to a new cloud storage provider. OmniCloud supports the move actively by providing its own migration service. Data does not need to be downloaded into the enterprise network nor does it have to be re-encrypted.

The Typical OmniCloud User

OmniCloud was developed for business use and is directed especially at small and medium enterprises that would like to use the advantages offered by cloud storage but also have high requirements concerning the stored data's security. OmniCloud is of interest specifically for those companies that do not have a budget for implementing their own private cloud solution. OmniCloud allows them the secure use of public cloud solutions.

2.1. Use Cases

OmniCloud facilitates a multitude of use cases due to its modular structure. The actual functionality results from the chosen and configured individual OmniCloud functional modules. Thus OmniCloud offers enterprises a great degree of security while remaining highly flexible. The following lists some examples of OmniCloud use cases. Many other scenarios are possible as well.

Redundant data backup both in the cloud and locally. OmniCloud permits to connect several cloud storage services and local storages at the same time. This provides for redundant storage to be used for data backups. The data will be stored encrypted at various providers and locally in the enterprise network as well.

Combine several cloud storages to create one big one. OmniCloud is capable of combining different cloud storage offers. This allows users to combine offers (such as offered by DropBox or similar providers) and integrate them in form of a large drive into

their own business environment. OmniCloud checks the fill level of the individual offers, and when a cloud storage is full Omnicloud switches to another offer.

Data exchange over a joint project folder. Using OmniCloud enterprises can create department or project team folders to be used by the employees for exchanging files. From the user's point of view OmniCloud acts here exactly like a network drive.

Dynamic teams and responsibility adjustments. In contrast to many other encryption solutions OmniCloud is well suited for dynamic teams and allows for typical business situations such as employee absences and responsibility adjustments, for example.

2.2. OmniCloud Features - An Overview

OmniCloud offers a multitude of functions and features to organize the cloud service use securely and flexibly. The following gives an overview of the integral features.

Data encryption is OmniCloud's most important function. OmniCloud encrypts all files before they leave the company network. The keys used for the encryption are generated within the company network and are therefore not known to the cloud storage provider. By applying additional security mechanisms such as **authentication and access control** an enterprise can define exactly which users may be permitted to access which files in OmniCloud.

OmniCloud is **easy to implement** into an enterprise's existing system infrastructure. It **does not require installing** OmniCloud software on the users' end devices. OmniCloud supports various **standard communication protocols** that are supported directly by many operating systems and applications as well.

OmniCloud supports a **multitude of cloud storage services** such as Amazon S3, Dropbox or Box. New services can be integrated easily.

OmniCloud enables **saving on cloud storage space** and thus **reduces cloud storage costs**. Enterprises often store the same data at different locations. OmniCloud recognizes such duplications and ensures that only one copy ends up in the cloud storage (deduplication).

Storing strategies allow the user to determine how OmniCloud will distribute incoming data on the configured cloud storage services. OmniCloud is responsible for adapting to or mapping between the different cloud interfaces. For example, a storage strategy may define that the encrypted data is to be filed in multiple cloud storage services and local databases as well ("mirroring") to increase **redundancy**. OmniCloud also facilitates creating one large cloud storage by combining several small cloud storages. OmniCloud monitors the cloud storages' fill levels and distributes the incoming data accordingly ("striping"). Other storing strategies are easy to implement as well.

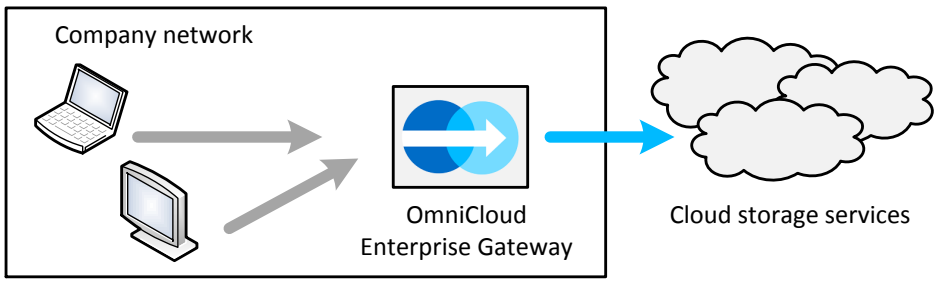

Figure 1: OmniCloud as Enterprise Gateway

OmniCloud offers a **migration service** for the data in the cloud storage. This migration service can prevent the undesired binding to the provider of one specific cloud storage service ("provider lock-in"). The migration service itself runs within the cloud and can thus profit from the fast network data transmission rates. Downloading and re-encrypting the data is not necessary.

OmniCloud is based on a **trust model** that distinguishes strictly between the local enterprise network on which OmniCloud is running and the cloud provider. The significance of this feature is demonstrated when looking at cloud storage services such as Crashplan or Mozy that offer local encryption as well. On the one hand, these services do not require a user to trust the service provider because all data is encrypted locally before they leave the user's computer. On the other hand, the service suppliers provide the users with software that has to be installed on their end devices, which then assumes key generation and encryption. The trust model on which these solutions are based is highly problematic. Though it is assumed that none of the cloud providers has any dishonest intention, there still exist possibilities that client software may not behave as desired, for example due to a software error. Being aggregators of large amounts of data, cloud storage providers represent an appealing target for espionage and infiltration.

3. Technical Overview

3.1. The Enterprise Gateway Approach

OmniCloud was designed in such a way that it runs within an enterprise's network. OmniCloud functions as an enterprise gateway in the form of a network appliance or server, as illustrated in the figure 1. With the OmniCloud gateway all users can securely access the cloud storage services used by OmniCloud. OmniCloud realizes the security mechanisms as invisible to the user (for example data encryption). Using standard software (for example file managers or backup programs) users may connect to OmniCloud and copy or exchange data. Installing a dedicated OmniCloud application on the users' end devices is not necessary.

OmniCloud supports typical interaction patterns and use cases such as storing data copies in the cloud, generate backups in the cloud or working directly on remotely stored files. In OmniCloud "supports" means that the user uses appropriate software (for example a backup software) that offers standardized interfaces such as FTP[PR85], Amazon S3[Ama] or Secure Copy (SCP)[Ora]. OmniCloud offers the following storage functions:

Copy. The copy function generates an *exact copy* of the current local data in the cloud *ad-hoc*. This way the typical user wants to ensure that the data will still be available when the local hardware is failing (for example due to damage or theft). It must be noted, however, that OmniCloud is responsible for data encryption and decryption, and that the data copies in the cloud can be accessed only via OmniCloud. Data can be accessed from outside of an enterprise's network (for example from an Internet cafe or a customer's) only via a VPN connection to the OmniCloud gateway or by activating enterprise firewall ports for individual OmniCloud services.

Data Backup. The backup function allows restoring *any version* of the previously stored files or directories over a long period of time. Normally generating backups in the cloud is an automated process that *regularly* generates data copies and transports them into cloud storage, so that they can be restored in case of damage or loss.

Shared Directory. This function facilitates using OmniCloud like a network drive. Users can work directly on remotely stored data by using appropriate software and communication protocols such as FTP or WebDAV or by mapping OmniCloud as a network storage to their operating system.

3.2. Design Principles

Easy Implementation. The goal was to implement OmniCloud easily into existing IT infrastructures. To achieve this OmniCloud supports different communication protocols (e. g. FTP, WebDAV, various cloud interfaces such as Amazon S3), which allows using a variety of client applications (e. g. backup or FTP software) with OmniCloud to store files in the cloud. This means that OmniCloud does not need a dedicated OmniCloud client to be installed on the users' computers. Beyond that OmniCloud supports a number of cloud storage services enabling the user to select a suitable cloud storage service.

File and Key Separation. OmniCloud stores only encrypted files in the cloud. The keys for encrypting or decrypting the files remain within the enterprise and are managed by OmniCloud itself. Thus cloud storage services are in no position to decrypt any of the files.

Meta-Information and Key Separation. OmniCloud stores meta-information about the files (e. g. file name, file owner, location of the files in the cloud) in a local database. The keys required for encryption and decryption are stored locally as well, but separate

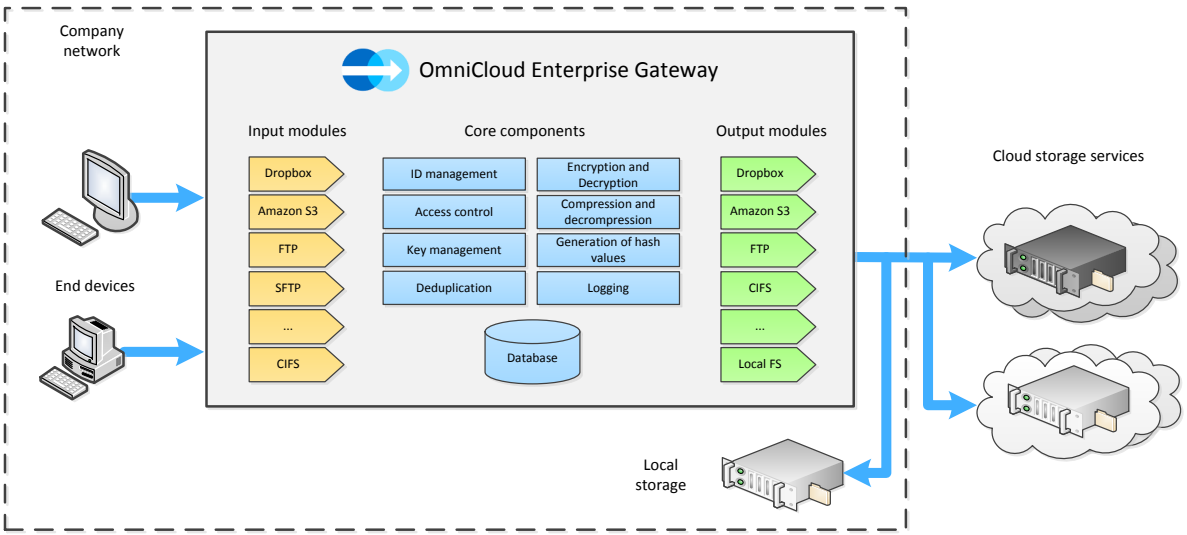

Figure 2: OmniCloud components

from the meta-information. Instead OmniCloud applies a secure key manager to store the keys. This prevents unauthorized insiders from gaining access to the keys.

Self-Replication. Once installed an OmniCloud instance can be recovered at any time in the case of a severe error (e. g. after a crash of the hard drive). In order to do this all the data managed by OmniCloud in the local database and all the configuration files are integrated into a backup process. This backup process encrypts this data in periodic intervals using a dedicated master key and stores the data encoded in such a manner at a suitable location. The cloud storage service used for storing these OmniCloud backups does not necessarily have to be the same one as used for storing the encrypted user files.

In case of data loss the OmniCloud administrator is able to download these OmniCloud backups and to recover the OmniCloud installation using the retained master key.

Joint Resources. OmniCloud may be used to grant various users access to the same resources stored in the cloud. This means that OmniCloud provides these users with a joint view at the jointly used resources. To do so, all the operations on the files are being synchronized transparently in the background. The resources may be accessed simultaneously via different protocols as well (e. g. FTP and CIFS[HLP96]).

3.3. Architectural Components

This section gives an overview of OmniCloud's fundamental architectural components (compare figure 2). The components are grouped into input modules (yellow), output modules (green) and core components (blue). On the left side are the users' clients accessing the OmniCloud gateway over the enterprise network.

Data Sinks. OmniCloud can use several external cloud storage services (on the right in figure 2) to mirror user files on several cloud storage services, for example. In addition, OmniCloud can store the files on local drives within the enterprise network to increase file availability in the case of internet problems. In the following data sinks refer to both internal and external storages.

Input Modules. The input modules are responsible for authenticating the users; they receive the client requests and forward them to the core components. Each input module supports a specific communication protocol (e. g. FTP or Amazon S3). OmniCloud is capable of operating several input modules in parallel and can thus offer a variety of communication protocols.

Output Modules. The output modules are responsible for communication with the data sinks. They accept the core component requests and translate them into service specific requests. Just like the input modules each output module supports exactly one data sink (e. g. by using a REST interface[Fie00]). The simultaneous use of several output modules facilitates redundant storage in different data sinks.

Both input and output modules were designed as lightweight components with low functionality. This immensely simplifies the implementation of additional modules. For more sophisticated functionalities such as user authentication the input and output modules can use the appropriate core components.

Core Components. The core components process the user requests. Each core component is responsible for a specific task, for example identity management (IDM), key management, file encryption or decryption, or generating the files' hash values.

OmniCloud Services. OmniCloud services represent a concept with which to group input and output modules, data sinks and varying configurations. That is from a technical point of view an OmniCloud service is an assembly consisting of a number of input modules, a number of output modules, a number of supported operations[1], and a concrete OmniCloud configuration as well.

The OmniCloud services thus represent a very flexible means to provide different functionalities. OmniCloud can define a random number of services with each service having its own configuration and thus a different functionality. For example, one service may encrypt every file because the files are stored with an external cloud storage provider, another service, however, could dispense with that because it stores the files only locally. Each OmniCloud service may be allocated its own amount of authorized users [2]. Depending on the service's respective security policy users of one service may be authorized to access the files of other users of the same service. In principle, the required access rules for this may be defined in the desired granularity. In doing so the OmniCloud service concept enables defining jointly used resources.

[1] *operation* here is to be understood as a file system operation that activates a user's client application (e. g. "store file at cloud provider", "load file from cloud provider", "delete file" and so on).

[2] but a user may be authorized to access more than one service.

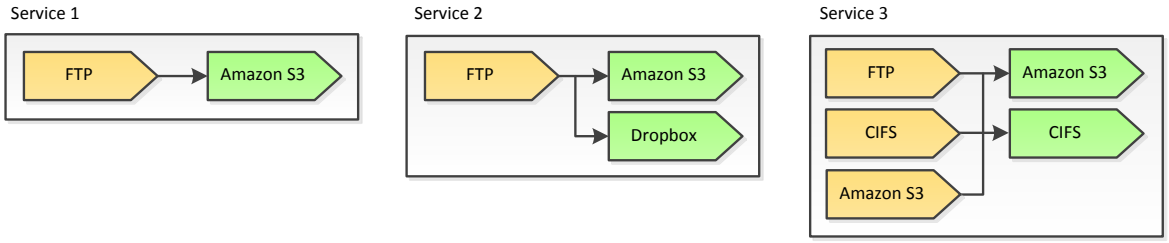

Figure 3: Communication protocol mapping

3.4. Cloud Interfaces and Storage Strategies

Interface Connection

An important concept in OmniCloud is the option to connect various input modules, which allow users to communicate with OmniCloud, to different output modules that help with storing the data in the cloud ("API mapping"). OmniCloud acts as the translator between the input and the output modules' interfaces and communication protocols. Which input and output modules are to be connected here can be defined individually for each OmniCloud service. Besides the simple *1-to-1* relations between input and output modules complex structures such as *1-to-n*, *m-to-1* or *m-to-n* relations (see figure 3) are possible as well.

Mapping between the different input and output modules gives the OmniCloud users the advantage that cloud storage offers can be used even when the client software used does not support the communication protocol offered natively by the cloud provider. OmniCloud makes software cloud ready. This allows using an already existing backup solution that may for example support only FTP, CIFS or SCP for data storage with cloud storage services such as Amazon S3 or Dropbox as well. It also helps to change cloud storage providers easily without having to make modifications in the client software. Provider lock-in can thus be prevented.

Storage Strategies

When storing data in the cloud, OmniCloud facilitates using various storage strategies. Each OmniCloud service has exactly one storage strategy. It defines the way in which OmniCloud archives the files in the cloud storages. For example, one part of the storage strategy may be to define which cloud storage services are to be used and how the data is to be distributed between said cloud storage services. The following lists storage strategies examples :

- Store all files in the same cloud storage,
- Store the files in several cloud storages ("mirroring"),

- Store the files in at least m of n cloud storages,
- Store all files on a local data storage in the enterprise network and in one of n cloud storages as well,
- Use Reed-Solomon encoding procedures to increase fault tolerance, as for example suggested in [APW10],
- Divide files into segments by applying information dispersal procedures, as for example suggested in [SGS11] and [S+11].

OmniCloud's modular approach and the existing program interfaces allow the realization of new storage strategies in a simple manner.

4. Data Security

When using cloud storage services for business purposes additional requirements must be met that result from the data's value itself for the enterprise and from the current legal framework (compliance standards) as well. These have an effect with regard to the data's security, specifically regarding effective and efficient identity, rights and key management. Existing cloud storage offers for private end users tend to consider these aspects insufficiently. OmniCloud, however, offers efficient management functions that can be integrated into existing processes for user provisioning.

OmniCloud's main objective here is to make using cloud storage services more secure from the user's point of view. For this OmniCloud offers various security mechanisms such as user authentication, access control, encryption, and related tools for identity, rights and key management. The most important OmniCloud security functions will be introduced in the following.

4.1. Security Presumptions

To provide a basis and understand OmniCloud's security mechanisms a number of security and trust presumptions must be described first.

It is presumed that OmniCloud will be installed within an enterprise's Intranet as an enterprise gateway (compare section 3.1). Only authorized employees may access the services provided by the OmniCloud gateway. To be authorized employees first have to authenticate themselves over a supported process by using the respective credentials. The credentials will be allocated to the staff within an enterprise by the accordingly designated entities. The users protect their credentials against unauthorized access and do not pass them on. The company running the OmniCloud gateway trusts its staff to

use the gateway only for the intended purpose. Additional related use regulations will be implemented both technically and organizationally.

OmniCloud services define a specific area in which service users may co-operate. All the users of one specific service have entered into a special trust relationship, even if not all users may have the same access rights to the processed files.

The OmniCloud gateway is managed by dedicated OmniCloud administrators presumed to be employees of the enterprise that is running the OmniCloud gateway. These administrators can access the administrative OmniCloud components and interfaces, set up new users, allocate credentials to the users, establish new OmniCloud services and define service access rights. They have access to all the OmniCloud databases and configuration files that store the OmniCloud relevant meta-information. OmniCloud administrators are equipped with credentials as well. It is presumed that OmniCloud administrators will protect these credentials against unauthorized access and refrain from passing them on.

OmniCloud's key management component generates and stores all the keys required for file encryption or decryption. Additional mechanisms protect the confidentiality of these keys. To initialize the key management component an OmniCloud administrator has to provide or activate the necessary credentials or keys when activating an OmniCloud gateway. This approach is comparable to starting an HTTPS webserver where the administrator is activating the private key by entering a secret combination. The credentials or keys represent the trust anchor for the OmniCloud key management components, without it new files can neither be encrypted nor existing ones be decrypted. For example, it can be implemented as a cryptographic key or as a smartcard and corresponding PIN.

It is presumed that all OmniCloud gateway users and administrators protect their credentials in such a way that potential security incidents will be minimized. If an OmniCloud user loses his credential, its further use will be disabled and a new credential will be issued. To guarantee a high degree of availability of the OmniCloud gateway, additional measures should be taken to achieve the desired degree of availability, reliability and performance. Data backups are needed as well to secure the OmniCloud internal administrative information. Suitable measures should be implemented to protect these data backups against unauthorized access and manipulation.

Concerning the cloud storage providers it is presumed that they comply with the negotiated contracts and quality of service commitments, for example with regard to the stored data's integrity and availability. The confidentiality of the stored data, however, is not presumed (even if respective quality of service commitments were made). However, it is presumed that various causes may impact the data's confidentiality, for example due to national regulations or vicious attacks. For this reason OmniCloud is offering a mechanism which ensures that the stored information remains confidential even if the files would become public at the cloud storage provider's.

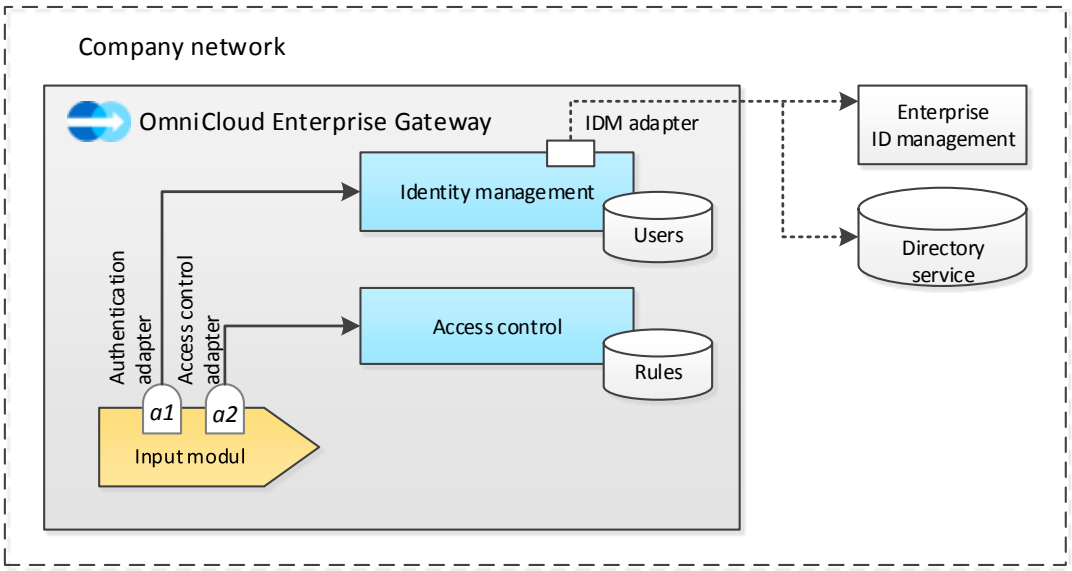

Figure 4: Identity management and access control component connection

4.2. Authentication

The authentication of OmniCloud users is the underlying security mechanism on which the subsequent security mechanisms build. Users must authenticate themselves to an OmniCloud service's input module, for example to an OmniCloud service's FTP server. Since users may have access to different OmniCloud services, they may in turn have different credentials for the various authentication types as well. There may be different authentication procedures that may be employed in a service's individual input modules. For example, the authentication procedure of an SFTP server[Fil] differs from that of a CIFS server.

Input modules were developed as light-weight components, i. e. they themselves may not provide for the overall functionality with regard to the ID management. Instead they can use the OmniCloud gateway's core components such as the identity management component, which contains all the reference values for verifying credentials. This identity management component is capable of handling the credentials for the different authentication protocols (for example passwords, hash values of passwords, or public keys for public key based procedures).

The work between the input modules and the identity management component is divided as follows: The input modules realize the respective authentication protocols while the identity management component verifies the submitted credentials. This is achieved by appropriate adapters in the input modules. Figure 4 shows a sample input module with the authentication adapter *a1* connecting to the identity management component.

The identity management component is also responsible for mapping the input module specific identities to OmniCloud-wide identities. The identity management component

can be connected to the enterprise's central identity management component via an IDM adapter (see figure 4) for simplifying the OmniCloud identity administration.

Credentials for the OmniCloud administration are managed separately from those of the regular OmniCloud users. A more detailed description of the underlying mechanisms will be dispensed with herein. Besides the user authentication OmniCloud components authenticate each other at certain points as well. The respective details are not the subject of this documentation either.

4.3. Access Control

OmniCloud offers a mechanism for role-based access control. For each OmniCloud service it can be defined what user is allowed to carry out which operations on what files. To simplify rights management the roles can be used for grouping the users. In order to achieve a use case adequate compact amount of rules the access rules may differ from service to service with regards to their level of granularity. Access rules can be defined based on the user, the assigned role, directories, files or file characteristics. Even more complex access rules can be implemented, for examples rules that take dynamic context information into account such as date, time or IP address. This can be done by defining context sensors that deliver the current context value at the time when the access control is being decided upon (compare [WKS03, NS04, AW09]). To facilitate access rule management while increasing usability and efficiency at the same time, concurrent access rules may be defined in varying degrees of granularity.

Access control is integrated into the input modules similarly to authentication. Access control decisions are enforced in the input module by means of an access control adapter and with the support of the core component for the access control (see figure 4). The access rules are defined using the ORKA Policy Language (OPL) [AWP09]. The overall syntax and related formal semantics were described in [AW09].

Expressiveness. The following access rules represent examples for OPL's strong expressiveness: (1) All users of a service can access all the files of the service. (2) Only user A, B and C may overwrite files in directory X. (3) All users may read the files but only the file owners may modify them. (4) Files ending in *.docx* may be overwritten only if the operation is not taking place on a weekend and the temperature in Darmstadt is more than 20 degrees Celsius.

Configuring OmniCloud services for certain use cases (projects, departments, etc.) and defining the associated access rules is an effective instrument to control access to the files stored in the cloud storages.

4.4. Local File Encryption and Key Management

As aforementioned, OmniCloud does not rely on the cloud storage provider's measures to ensure the stored data's confidentiality. All files will be encrypted locally at the OmniCloud gateway before they are stored externally. All files will be encrypted with different keys. OmniCloud supports different symmetrical encryption and decryption algorithms and modes.

To increase the security level the encryption keys will not be stored together with the meta-information related to the files. Instead a key management component is being used, the main purpose of which is to generate new keys, store them with a corresponding ID, and to supply these keys upon request or presented ID to the encryption or decryption component respectively. The key management component's interface has been designed in such a way that it is possible to integrate other implementations as well (including dedicated hardware components with the respective program interfaces). For example, this would facilitate using an especially secure entropy source for key generation. The data of the key management component is protected by specific credentials or security tokens that have to be presented for the initialization when the OmniCloud gateway is started. Encryption keys will not be stored unencrypted on the gateway nor will they be stored with the encrypted files at the cloud storage provider's.

The keys used for encryption are generated pseudo-randomly (or if a suitable entropy source is available then randomly as well). In contrast to other cloud encryption solutions OmniCloud does not employ password-based encryption ("PBE"). Background: In password-based encryption procedures the cryptographic keys are derived from passwords. In practice, the range of keys resulting from PBE is tiny in comparison to the range of keys from the encryption process, even when the passwords are very long and have a lot of different characters. It is therefore an erroneous inference to believe that cryptographic keys generated by password-based derivation functions are random and comparable to the security provided by random keys. Such procedures can be targeted by brute force or dictionary attacks, even if well established and common derivation functions such as PBKDF2 from the PKCS#5 standard [RSA99] are applied. This can lead to an attacker being able to decrypt confidential data without breaking the associated encryption algorithm (for example AES). In this context 17 password managers on mobile end devices were analyzed in a recent study [BS12] that encrypt their data with PBE. All 17 password managers were broken open after one day at the very latest. Furthermore, OmniCloud does not derive the encryption keys from the file content. The disadvantages of such an approach were described in the Fraunhofer SIT cloud study [BHH+12].

Encryption. When a file is being encrypted the following steps will be carried out: (1) The encryption component requests a new encryption key from the key management component, (2) the key management component returns the generated key together with a key ID, (3) the encryption component encrypts the file with the key while applying

the configured encryption procedure, (4) the encryption component stores the key ID as the file's internal meta-information for later decryption.

Decryption. Decrypting a file is done according to the following steps: (1) The decryption component establishes the file's key ID from the internal meta-information, (2) the decryption component requests the key associated with this key ID from the key management component, (3) the key management component delivers the respective key, (4) the decryption component decrypts the file applying the rendered key.

The key management component supports the use of differing encryption procedures and modes for different files. This facilitates modifying the standard encryption procedure used in an OmniCloud service during a service's life cycle. In such a case the existing files remain encrypted with the previous procedure and key while new or modified files apply the new procedure. From a company's perspective it is thus very easy to replace older but yet secure encryption procedures successively with newer procedures without having to re-encrypt all the files at once.

In contrast to a lot of other encryption solutions OmniCloud is suitable for dynamic teams and allows for typical business situations such as employee absences and changes in responsibilities. This is possible due to identity management and key management being separate. This facilitates the easy implementation of stand-in regulations – without having to pass on passwords or do extensive data re-encryption.

4.5. Obfuscating File and Directory Names

Besides the actual file content the associated meta-information such as the file name or the directory name may contain confidential information as well. The potential content of a file named *"john-doe-dismissal-without-notice.docx"* can be guessed even without being able to look inside the file. A folder carrying the name *"takeover-measures-company-xyz/"* already reveals the enterprise's strategic objectives. This is why OmniCloud not only encrypts a file's content but renames all the files before they are stored in the cloud. The corresponding file names are generated randomly, file endings such as *".pdf"* will be removed. Before cloud storage the complete directory structure will be deleted to be managed internally by OmniCloud. Thus neither an attacker nor a provider can draw conclusions to the data stored in the cloud. Figure 5 shows a folder content example of a cloud storage used by OmniCloud.

4.6. Other Data Security Aspects

Besides the previously described security mechanisms OmniCloud protects all internal management information against unauthorized access as well by employing suitable security mechanisms (for example encrypted storage of all encryption keys). Information

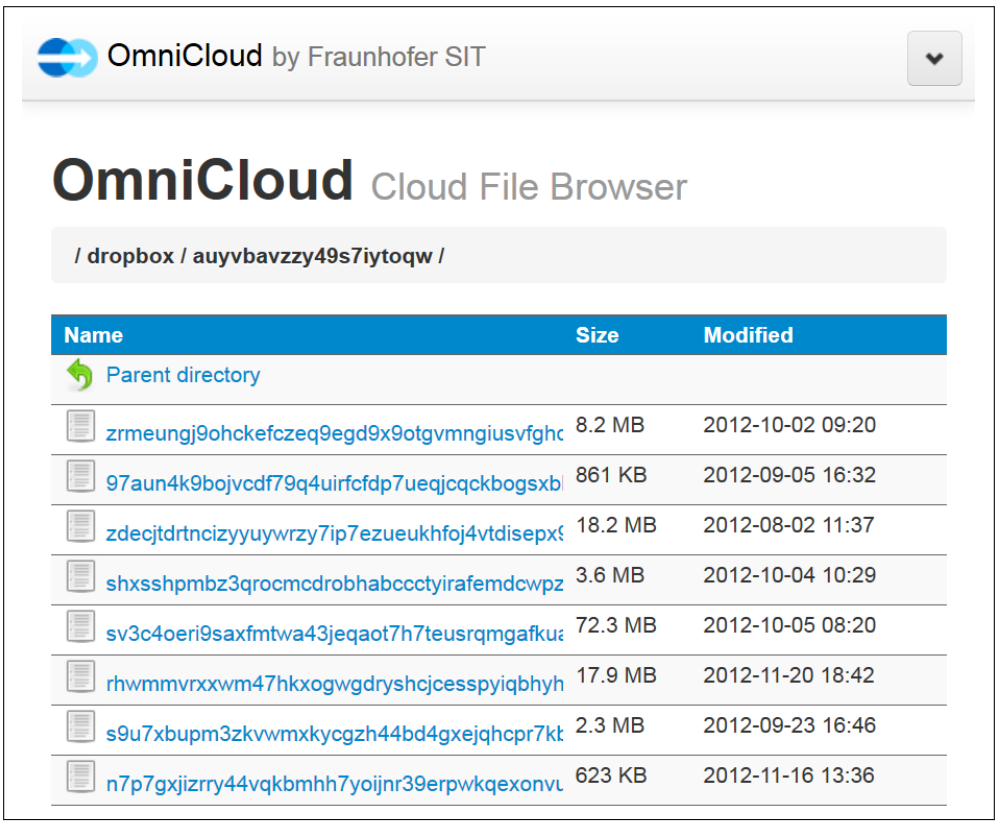

Figure 5: Obfuscating file and directory names

exchanged between the user end device and the OmniCloud gateway or between the gateway and the cloud storage provider is protected by appropriate measures (as far as this is possible with the communication protocols used or any other means). Beyond that, all data protections of OmniCloud internal information is safeguarded against unauthorized access and unnoticed modification.

5. Duplication Prevention / Deduplication

The term *deduplication* describes a popular technique that enables cloud storage providers to reduce the amount of required storage space significantly. The basic principle of deduplication is that only one copy of any file will be stored. If a user would like to store a file, which the cloud storage provider had already stored previously, the service will simply generate a cross reference to the existing file instead of storing another copy. Deduplication distinguishes between different versions:

1. **File-based deduplication versus block-based deduplication.** File-based deduplication means that only one copy of each file will be stored. In the block-based deduplication each file is split into blocks and only one copy of each block

will be stored. Comparing the hash values of files and blocks with a list of all known files and blocks allows identical file or block detection.

2. **Server-side deduplication versus local deduplication.** In the case of server-side deduplication each file a user would like to store is being sent to the cloud storage service first. The service reviews each file, whether it needs to be stored or if only a cross reference needs to be generated to an already existing file. In this case the user cannot determine whether the cloud storage service supports deduplication. In the case of local deduplication the client application first only sends the file's hash value to the service. Only if the service does not yet store the file the client application will be requested to send the complete file. This type of deduplication has the effect that besides cutting down on storage space it saves bandwidth as well. A log file analysis or observation of the transmitted data makes it easy to determine whether a service uses this type of deduplication.

3. **Deduplication per user versus user-shared deduplication.** Deduplication per user means that deduplication will be carried out separately for each individual user: If user A would like to store a file that he had previously stored (possibly in a different directory) the cloud storage service will merely generate a cross reference to the file. However, if the file had been stored previously by just one other user another copy of the file will be produced. This is not the case in user-shared deduplication. Deduplication will be carried out across all the users: The cloud storage service will produce a cross reference to an already existing file, if user A wants to store a file that another user B has already stored.

With OmniCloud users can profit from the advantages deduplication provides even if the cloud storage service itself does not support deduplication. Since OmniCloud stores the hash values of all the files, OmniCloud is in a position to decide whether a file has already been transmitted to a cloud storage service or not. If a file has already been sent to a cloud storage service OmniCloud generates a reference in its own database and marks the file as de-duplicated. OmniCloud carries out deduplications separately for each OmniCloud service. This means that both deduplication per user and user-shared deduplication are possible, depending on the group of authorized service users.

Some data protection issues still exist that need to be taken into account when using deduplication [HPSP10]. These data protection issues, however, can only arise if the cloud storage service employs both local and user-shared deduplication. If this is the case some side channel attacks are possible. These attacks work according to the following principle: An attacker, who has access to a cloud storage service, may use deduplication to find out which files have already been stored with this service: He sends a file to the service and observes what will be happening. If his client application does not transmit the file the attacker will know that the file is already present at the cloud storage service. More exactly, he knows that another user has the same file, but he does not know which one of the users. More sophisticated variations of this attack use this principle to obtain information about a specific user.

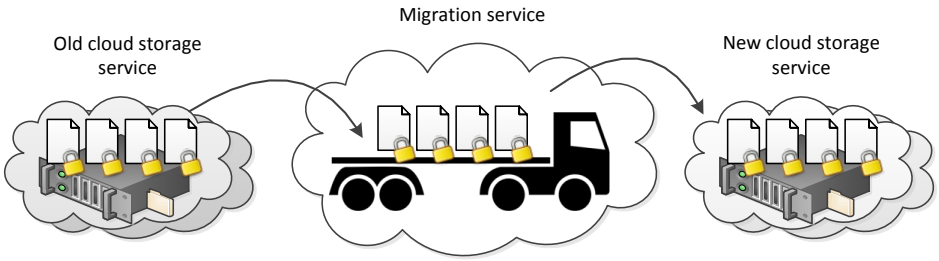

Figure 6: OmniCloud Migration Service

OmniCloud averts such types of attacks. From the perspective of the client application OmniCloud is carrying out a server-side deduplication, because the client has to transmit every file to OmniCloud regardless whether this file has been transmitted previously or not. Furthermore, deduplication is performed only within an OmniCloud service. But only those users can access an OmniCloud service that have a high degree of trust in each other, which is why side channel information does not need to be considered as being particularly crucial. Implementing customized OmniCloud services for each individual user or deactivating deduplication for individual OmniCloud services prevents such an information flow.

6. Migration Service Makes Provider Change Easy

One major reason why many companies hesitate to use cloud storage services is the fear of being locked in to one specific cloud provider. Since there is no efficient process available for migrating the data between the various cloud providers, users tend to be bound to the cloud provider they have chosen originally. To counteract this OmniCloud offers a migration service that helps users to switch easily over to another cloud storage service.

One special feature of the migration service is that it is not part of the local OmniCloud gateway. Instead it operates as a virtual machine at the cloud provider's location, which is outside of the enterprise's Intranet. The local OmniCloud gateway merely manages the migration service. When data migration is carried out the migration service takes all the encrypted files from the old cloud storage service and transfers them to the new service (compare figure 6). It is unnecessary to decrypt the files and subsequently encrypt them again.

This is an immense advantage as opposed to a migration service that runs within the enterprise's Intranet, because this solution profits highly from the high speed Internet connections between the cloud services. For example, if a local migration service would use a rather slow DSL connection to download all the files with an overall size of 500 GB from cloud storage service A and then upload it to service B the data transfer time would

take more than seven weeks [3]. Besides, the local internet would have to carry a very high load during the data transfer, which in turn could restrict OmniCloud gateway usage, which is running on the same Intranet. On the other hand, the transfer rate increases dramatically if the data is copied directly from cloud to cloud. An analysis by Nasuni concluded that copying 12 TB of data from one cloud to another will take only between 4 hours and one week [Nas12].

[3] computation basis: DSL 16.000 with a net download rate of 1 MB/sec. and a net upload rate of 128 KB/sec.

References

[Ama] Amazon Web Services. Amazon Simple Storage Service (Amazon S3). http://aws.amazon.com/de/s3/ (Last visited on 12/09/2014).

[APW10] H. Abu-Libdeh, L. Princehouse, and H. Weatherspoon. RACS: A Case for Cloud Storage Diversity. In *Proc. of the 1st ACM Symp. on Cloud Comp.*, pages 229–240. ACM, 2010.

[AW09] C. Alm and R. Wolf. The Definition of the OPL Access Control Policy Language. Technical Report MIP-0902, University of Passau, Germany, 2009.

[AWP09] C. Alm, R. Wolf, and J. Posegga. The OPL Access Control Policy Language. In *TRUSTBUS*, volume 5695 of *LNCS*, pages 138–148. Springer, 2009.

[BHH+12] M. Borgmann, T. Hahn, M. Herfert, T. Kunz, M. Richter, U. Viebeg, and S. Vowé. On the Security of Cloud Storage Services. SIT Technical Reports SIT-TR-2012-001, https://www.sit.fraunhofer.de/en/cloud-security/ (Last visited on 12/09/2014), March 2012.

[BS12] A. Belenko and D. Sklyarov. "Secure Password Managers" and "Military-Grade Encryption" on Smartphones: Oh, Really? http://www.elcomsoft.com/WP/BH-EU-2012-WP.pdf (Last visited on 12/09/2014), March 2012.

[BSI12] BSI. Studie zur IT-Sicherheit in kleinen und mittleren Unternehmen. https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/Studien/KMU/Studie_IT-Sicherheit_KMU.pdf?__blob=publicationFile (Last visited on 12/09/2014), January 2012.

[Cap12] Capgemini. Studie IT-Trends 2012. Business-IT-Alignment sichert die Zukunft. http://www.de.capgemini.com/ressourcen/it-trends-studie-2012 (Last visited on 12/09/2014), January 2012.

[Con12] Josh Constine. Dropbox Is Now The Data Fabric Tying Together Devices For 100M Registered Users Who Save 1B Files A Day. TechCrunch, Available from http://techcrunch.com/2012/11/13/dropbox-100-million/ (Last visited on 12/09/2014), November 2012.

[Deu11] Deutsche Telekom / T-Systems. Sicherheitsreport 2011. http://www.telekom.com/static/-/22846/1/110916-sicherheitsreport-si (Last visited on 12/09/2014), September 2011.

[Fie00] Roy Thomas Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, Irvine, 2000.

[Fil] FileZilla Wiki. SFTP specifications. Available from http://wiki.filezilla-project.org/SFTP_specifications (Last visited on 12/09/2014).

[HLP96] I. Heizer, P. Leach, and D. Perry. Common Internet File System Protocol (CIFS/1.0). Internet Draft. Available from http://tools.ietf.org/html/draft-heizer-cifs-v1-spec-00 (Last visited on 12/09/2014), June 1996.

[HPSP10] D. Harnik, B. Pinkas, and A. Shulman-Peleg. Side Channels in Cloud Services: Deduplication in Cloud Storage. *IEEE Sec. & Priv.*, 8(6):40–47, November 2010.

[Nas12] Nasuni. Bulk Data Migration in the Cloud. Whitepaper, http://www.nasuni.com/cloud-migration (Last visited on 12/09/2014), March 2012.

[NS04] G. Neumann and M. Strembeck. An Integrated Approach to Engineer and Enforce Context Constraints in RBAC Environments. *ACM TISSEC*, 7(3):392–427, 2004.

[Ora] Oracle. Jan Pechanec's weblog: How the SCP protocol works. https://blogs.oracle.com/janp/entry/how_the_scp_protocol_works (Last visited on 12/09/2014).

[PR85] J. Postel and J. K. Reynolds. RFC 959: File transfer protocol, October 1985.

[RSA99] RSA. PKCS #5 v2.0: Password-Based Cryptography Standard, March 1999.

[S+11] J. Spillner et al. Information dispersion over redundant arrays of optimal cloud storage for desktop users. In *Proc. of Utility and Cloud Computing (UCC) 2011*, pages 1–8. IEEE, 2011.

[SGS11] R. Seiger, S. Groß, and A. Schill. Seccsie: A secure cloud storage integrator for enterprises. In *Commerce and Enterprise Computing (CEC), 2011 IEEE 13th Conference on*, pages 252–255. IEEE, 2011.

[WKS03] R. Wolf, T. Keinz, and M. Schneider. A Model for Context-dependent Access Control for Web-based Services with Role-based Approach. In *DEXA Intern. Workshop on Network-Based Information Systems*, pages 209–214. IEEE Computer Society Press, 9 2003.

A. Frequently Asked Questions (FAQ)

General Questions

1. **What is OmniCloud?**

   OmniCloud is a solution that facilitates enterprises to be able to use any kind of cloud storage service securely and flexibly.

2. **Who developed OmniCloud?**

   OmniCloud is being developed by Fraunhofer Institute for Secure Information Technology SIT in Darmstadt.

3. **Is OmniCloud a cloud storage service?**

   No, OmniCloud is not competing with other cloud storage services. OmniCloud itself does not store any files, it encrypts the files and transfers these encrypted files to the respective cloud storage service.

4. **Which cloud storage providers does OmniCloud cooperate with?**

   OmniCloud currently supports Amazon S3, Dropbox, Box and any FTP server. Other providers will follow shortly.

5. **Can I use OmniCloud with local storage space as well?**

   OmniCloud supports both local and cloud storage. Local storage can be linked via FTP or directly via a file system interface, for example.

Licensing, Purchase and Distribution

6. **What is OmniCloud's current status?**

   OmniCloud is still being developed. The concept and the OmniCloud prototype were developed by Fraunhofer SIT and presented at various events. Fraunhofer SIT intends to develop OmniCloud further into a truly utilizable and marketable product.

7. **How can I purchase or license OmniCloud?**

   OmniCloud software development and productization is in progress. Currently there are talks with potential distribution partners. The licensing model has not yet been completed.

8. **How can I myself, a system vendor or system integrator, become an OmniCloud distributor?**

   Please send us your application. Our contact information can be found at the end of this document.

9. **I am interested in using OmniCloud in our company. How can I get the information that OmniCloud has been completed?**

   We will be happy to keep you informed about OmniCloud. Please send us a short message to this regard. Our contact information can be found at the end of this document.

10. **I would like to write an article about OmniCloud in a magazine or a blog. Where can I get more information on OmniCloud?**

    Many details about OmniCloud are in this whitepaper or can be found on our OmniCloud website at http://www.sit.fraunhofer.de/omnicloud.
    If you require further information please contact us. Our contact information can be found at the end of this document.

Installation and Use

11. **Where will OmniCloud be installed?**

    OmniCloud was designed for business use. The OmniCloud software has to be installed in the company's network on a server, the so-called OmniCloud gateway. All OmniCloud users can access the secure cloud storage via this gateway.

12. **What operating system do I need on my end device in order to use OmniCloud?**

    OmniCloud can be used independently from any specific operating system. Applications, which support standard protocols such as CIFS, FTP, SFTP, SCP and Amazon S3, and which can be used for communicating with OmniCloud, exist for all the typical operating systems.

13. **Is it necessary to install the OmniCloud application on each individual user's end device? How do the users access OmniCloud?**

    Installing the OmniCloud application is not required in order to use it. Instead use standard software such as a file manager or backup software to communicate with OmniCloud via standard protocols such as CIFS, FTP, SFTP, SCP and Amazon S3. Many operating systems allow opening up network resources for all local applications.

14. Can the data stored with OmniCloud be used conjointly?

Yes. To the user OmniCloud resembles a shared directory, which may be accessed by various users.

15. For example, can field representatives access the data stored via OmniCloud in the cloud as well?

Yes. Just like with using any other Intranet service the employee first has to connect to the enterprise's network via a VPN tunnel. Alternatively, individual OmniCloud services may be made available by activating a firewall port for outside access as well. In any case, it is required that the user is authenticated.

16. Can OmniCloud be used to release files for external business partners?

Yes. If business partners are being granted access to OmniCloud they can access the individual OmniCloud services. The business partners need to authenticate themselves with OmniCloud. Access rules allow defining in a fine granular manner which business partners may access which files.

17. How can I access my files in case that my end device is defective?

No problem. Use another end device and log into the OmniCloud gateway using your login details. You will immediately have access to your files again.

18. Can OmniCloud manage multiple versions of one file?

No, not yet.

19. What are the running expenses for the enterprise during the operation of OmniCloud?

In addition to the OmniCloud license costs there may possibly be charges for storing, recalling and transfering the files at the cloud storage providers used. When using the OmniCloud migration service there may be additional cost for transfering the files from the old to the new cloud storage provider.

Security

20. Can a cloud provider discover what kind of data are stored?

No. OmniCloud encrypts the file's content. File name and file endings are replaced with random designator and directory structures removed.

21. Which encryption procedure does OmniCloud use?

OmniCloud supports different symmetrical encryption procedures. AES-256 is used by default.

22. Where will the data be encrypted?

The data will be encrypted in the enterprise network on the OmniCloud gateway before they are transferred to the cloud storage provider.

23. Will the same key be used for all the files?

No. OmniCloud generates for each file its own key. Even when a file is being updated a new key will be generated.

24. Will the encryption keys be derived from the name or the content of the file to be encrypted?

No. OmniCloud generates all keys (pseudo) randomly. The generated keys do not refer in any way to the file's name or content. If such a reference existed it could have a negative impact on the security and privacy.

25. Will I, as the OmnCloud user, have to manage all keys myself?

No, OmniCloud takes care of the overall key management.

26. If my end device breaks, are all the keys lost?

No, the keys are not stored on the user's end device but on the OmniCloud gateway.

27. Can all the OmniCloud users access my stored data?

No. Implemented access regulations will define exactly what user may access what data.

28. How do I log into OmniCloud?

The actual login procedure depends on the configured input modules (for example with user name and password at the FTP input module). Users will receive their own access authorizations for the individual OmniCloud services.

About Fraunhofer SIT

Fraunhofer Institute for Secure Information Technology SIT is one of the oldest and most respected research institutions on IT security worldwide. A staff of over 165 employees supports enterprises and government agencies in protecting their data, services, infrastructures and end devices.

Fraunhofer SIT carries out applied research with the objective to make new technology market ready in such a way that its potential can be used securely and completely. Together with its partners the Institute designs new methods and procedures, generates prototypes, develops individual IT solutions and tests existing products and systems.

Contact

Fraunhofer Institute for Secure
Information Technology SIT
Rheinstrasse 75
64295 Darmstadt
Germany
Phone: +49 6151 869-213
Fax: +49 6151 869-224
info@sit.fraunhofer.de
www.sit.fraunhofer.de

Contact person for OmniCloud

Thomas Kunz
Phone: +49 6151 869-164
Fax: +49 6151 869-224
omnicloud-info@sit.fraunhofer.de

Ruben Wolf
Phone: +49 6151 869-178
Fax: +49 6151 869-224
omnicloud-info@sit.fraunhofer.de

OmniCloud web page

www.sit.fraunhofer.de/omnicloud