Proposal for building a "TERENA Trusted Cloud Drive" facility

Date: 28 March 2011
Authors: Licia Florio, Peter Szegedi
Contributions: Maarten Koopmans (vrijeheid.net), Dick Visser, Christian Gijtenbeek

1. Introduction

Several NRENs in the TERENA Community are increasingly interested in offering cloud services for their constituency and many national pilots have already been established.

Cloud service is clearly the new paradigm that changes the traditional way how services are provided and accessed by costumers. For instance, Cloud Storage supports the demand for outsourcing scientific data preservation services to third-party, for using distributed (i.e. geographical diversity) resources whenever needed, and for offering remote space for users to store their data and access that using different devices. Privacy and ownership when outsourcing to public clouds (i.e. Amazon, Google, etc.) remain a matter of great concern though.

TERENA has been asked by several members to take some actions to support the community experimenting with cloud services. As a results of this community call, a small meeting was organised in the TERENA office in October 2011 (notes are provided in the Appendix). One of the outcomes of this meeting was that TERENA will investigate the feasibility to initiate a pilot project that provides storage capabilities exposing the cloud delivery model for the participating NRENs.

It was agreed to build the pilot using the open source software of elastic cloud storage developed by UNINETT Sigma as part of the Nordic project NEON[1]; this software was presented at that meeting by the developer (Maarten Koopmans). The latest version of the software can be found at: http://www.assembla.com/spaces/cloud-backed-storage/wiki

This document describes the aim of the proposed pilot, its technical requirements, and offers a roadmap for its delivery.

[1] UNINETT Sigma plans to offer the software as a service in 2012 with Amazon S3 as a storage back-end.

2. Aim of the pilot

The aim of this pilot is to explore possible deployment scenarios for a trusted storage service for NRENs. The pilot will be built upon a federated software platform ("the cloud broker facility") that offers the ability to easily connect different storage back-end (both private and public cloud storage back-end are supported) and store users data in a secure and privacy preserving way (thanks to the separation of storage data and metadata as well as the built-in encryption functionality) in the cloud.
The following aspects will also be addressed as part of the pilot:

(i) Longer term sustainability for a potential service;
(ii) Legal aspects and perceived trust issues related to the storage and management of the encryption keys and metadata;
(iii) Software scalability and performance;

Although the software already offers capabilities to test different front-end applications too, this aspect will not be fully explored during the pilot. However, requirements will be collected during the pilot lifetime and recommendations on how to further improve the front-end (end-users) functionalities will be provided.

3. Technical characteristics of the pilot

The pilot will installing and operating the "cloud broker" which will be based on the open software developed by UNINETT Sigma in 2010 as part of the NEON project.

This proposed software has been built with the basic idea of separating the storage data (i.e. encrypted content) from the metadata (i.e. encryption keys, filenames, size, date, etc).

It supports encryption, meaning that user's data can be encrypted before being stored in the cloud: the encryption keys can then be stored in a trusted location (that is outside the cloud) whilst the encrypted data are stored in the cloud. This particular feature makes the usage of public clouds particular appealing. A set of metadata (defining file sizes, data creation, encryption keys etc.) are linked to the user's data for search purposes; the metadata (together with the encryption keys) are meant to be stored and operated by a trusted party, which in practical terms means that the storage data and the metadata can be handled by different parties.

By keeping the metadata store "on premises" data confidentiality is guaranteed under the assumption that the premises are inside a "trusted domain" – e.g. TERENA.

The metadata are stored in a so called "metadata store" named Voldemort[2] developed and open sourced by LinkedIn; this store scales elastically and

2 http://www.project-voldemort.com/

across data centers. The data itself is encrypted using 128 bit AES (though any cipher could be used); the file names are replaced by a unique identifier (UUID) for the cloud provider. The mapping between the UUID and the filename takes place in the metadata store. In this way, a stored blob does not reveal any information that could be exploited by a malicious attack (i.e. “which blob do we need to attack”).

The data stored in the cloud is accessed using WebDAV, which is a stable and widely supported protocol (MacOS, Linux, Windows all support it); the WebDAV server connects to a web front-end, which makes the WebDAV transparent to the end-users.

Though many built-in clients differ slightly in their implementation, all their differences are handled by the custom designed WebDAV daemon – also for iOS and Android’s most popular applications such as Goodreader and WebDAVNav.

The software has features that are not essential part of the pilot such as a public folder where users can share data with the world, a web interface to the file storage, the ability to tag files and folders and search on these tags via the web interface, and the ability to store these searches as “search folders” that are automatically updates when new files get tagged (or tags get removed from existing files). The data sharing functionality is on the software development roadmap for 2012.

Users will access the system by logging into a website (federated access). At a later stage it could be desirable to provide (beside the web access) a dedicated application (WebDAV client) which could be accessed in a federated fashion using Moonshot.

The picture below provides a high-level overview of the pilot to be implemented.

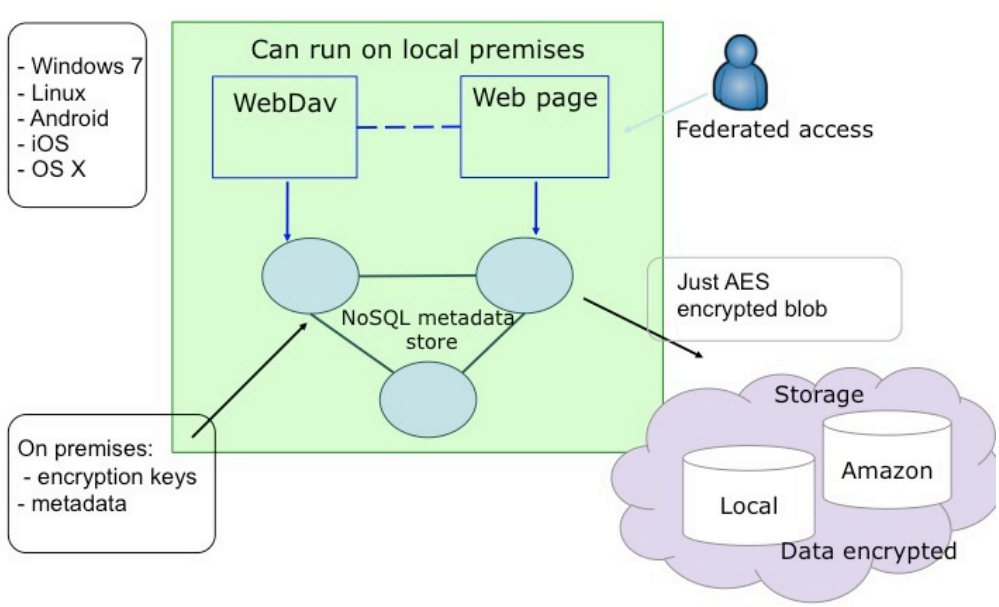

In the scenario depicted above a user logs into the system website via their federation; the storage is envisaged as a combination of a public storage (e.g. Amazon) and a local storage; the local storage will be offered by TERENA and the participating NRENs. There are no limitations on the type of storage that can be used, as the software is able to support different cloud storages such as Amazon S3, RackSpace CloudFiles etc., which makes it particularly flexible. Users will not be aware of the details of the cloud storage back-end, from their perspective upon logging in via their federations they will be able to store and retrieve their data as desired.

The frontends are elastically scalable as they are stateless – all data resides in the metadata store. Stable operations of the metadata store and fast connections (low latency, low round trip time) will improve the end user experience when keeping the NoSQL metadata store on premises.

4. Delivering the pilot

The technical part of the pilot will consist in installing all the components depicted in the picture above: namely a centralised cloud broker for the TERENA's community (the green box depicted in the picture above), the web portal to access the system (front-end) and the storage back-end. The pilot will be carried out in two phases:

i. Phase i - Local installation of the platform at the TERENA office. During this phase the cloud broker (the elements in the green box above) will be installed and connected to a limited storage backend offered by TERENA. A simple web portal and the necessary support for the federated access will also be developed. For this phase TERENA will sub-contract the software developer, Maarten Koopmans who will provide the necessary support for the installation. The platform will be evaluated and tested by a limited number of NRENs' experts coordinated by TERENA.

ii. Phase ii - Upon successful test of Phase i, NRENs will be invited to participate in the pilot (NRENs that have already expressed their interest in participating are HEAnet, NIIF, BELNET, PSNC, and CARNet/Srce) either adding their own cloud storage back-end and/or developing new front-end applications to the cloud broker. An additional public storage back-end will also be added. During this phase it is envisage that NRENs will offer a limited number of end-users to provide feedback on the usability of the system. Although most of the user requirements will not implemented during the pilot phase, they will help shape and understand the type of service users would be looking for.

The pilot Phase ii will be operated for a 9-month period after which an evaluation will follow to assess the success of the pilot and to agree on the following steps.

4.1 How to participate

There are three ways for NRENs to participate in the pilot:

- At the end of Phase i, NRENs can participate in the testing of the pilot service implemented at TERENA. This phase will be limited a selected number of NRENs' experts and it is meant to verify if all the components work as expected.
- In Phase ii, NRENs can participate in hooking up their storage back-end with the cloud broker and experiment with possible ways to deliver a cloud storage service for their constituency.
- In the last period of Phase ii, NRENs can participate in developing custom front-end applications to the cloud broker.

4.2 Timelines and Costs

The pilot will be delivered as follows:

Phase i) – 2 months

- Mid of March, 2012: approve the project and start Phase i
- End of April, 2012: complete Phase i and start testing.
- Mid of May, 2012: (possibly at TNC) give a short demo of Phase i and prepare to start of Phase ii

Phase ii) – 9 months

- Beginning of June, 2012: start Phase ii with the participation of interested NRENs
- Beginning of March, 2013: phase ii ends.
- Decide how to continue the pilot

An estimate of the costs for the pilot is described below.

i. In Phase i the costs will include:
   - the sub-contractor; (approx. 40h) about **3,000 EUR**
   - Support time from TERENA; (approx.5 days ITS and 5 days PDO) about **6,750 EUR**

   In the first instance TERENA will provide the necessary hardware/software and the whole system will be installed at the TERENA office. At the end of this phase, any cost related to additional equipment needed to enter phase ii will become clearer.

ii. In Phase ii the costs will be limited to support new NRENs to connect their cloud storage, enabling access to them, and ensure that the system is running smoothly; (approx. 10 days PDO and 25 days ITS) about **22,000 EUR.**

An additional **10,000 EUR** should be parked for equipment (if needed).

The estimated total cost:

- Phase i: 9,750 EUR
- Phase ii: 22,000 EUR
- Equipment: 10,000 EUR (optional)

41,750 EUR

4.3 Deliverables

There will be three deliverables produced as part of the pilot:

1. **May 2012 – Kick-off report**: System installation and technical documentation concerning the installation process (phase i).
2. **Jan 2013 – Describe possible service models**: This document will describe what service(s) can be deployed and how and will detail the service scenario recommended to phase ii and the related metrics to asses the pilot. The scenario of TERENA offering this as a (sharing) service will be considered.
3. **March 2013 – Final report**: Provide an evaluation of the pilot and recommendations for the next steps, based on the success of the pilot. Technical recommendations for NRENs that wish to run a local instance of the software will also be provided.

5. Added value for the participating NRENs

The pilot should be seen as an opportunity for NRENs to experiment with the ways to deliver cloud services and possibly to better understand the community requirements for cloud services. Because TERENA (a trusted entity) will take care of the instantiation of the pilot, NRENs' participation will require very limited effort, whilst NRENs can still enjoy the benefits.

One of the strong points of the proposed software relates to its security. As it has been pointed out in different occasions there are concerns on storing users' data particularly in public clouds. The proposed solution supports encryption so that user's data can be encrypted before being stored in the cloud.

During the lifetime of the pilot, the encryption keys and the metadata will be managed by TERENA, which will act as a secure trusted party. Particularly during Phase ii, when the storage back-end will be a combination of both private storage (provided by the participating NRENs) and public cloud storage (e.g., Amazon) the feature to encrypt user data before storing them into the cloud will result in an increased security.

The access to the system will be federated since phase i; this means that

especially in the case of which a centralized installation is offered, NRENs will be able to add an additional service to their federated service portfolio. Lastly, if the development of a Moonshot-enabled WedDAV client were agreed, interested NRENs would be able to test Moonshot technology features with limited burden on their side (the Moonshot-enabled WebDAV client could be developed by the Moonshot and reused by the interested NRENs).

At the end of the pilot TERENA will provide a short document summarizing the lesson learnt and how the pilot supported. Some of the identified use cases and guidelines will also be provided for those NRENs that wish to run a local instance of the software.

This pilot will offer a way to gain experiences in offering a cloud service and to explore possible scenarios for delivering the service; NRENs (whether those participating or not) will benefit from the results whether by using the pilot directly, by developing additional capabilities or simply by understanding the issues.

6. Risks

The risks related to this project are rather limited, considering its limited duration.

The main risk relates to the fact that to date the software used for the pilot is developed and maintained by one person (Maarten Koopmans); the risk can be mitigated by building a community behind it, especially if NRENs plan to use the software, however this process will require some time and commitments.

At this stage it is rather difficult to foresee the success of the pilot, both in terms of NRENs participating in the pilot and in terms of users' satisfaction. Considering that the initial installation will be operated by TERENA, the 'success' of the pilot might have implications on the performance. This risk can be mitigated by investigating scalability issues and by initially running the pilot on a limited scale.

Lastly, NRENs would have to take into account that from the users perspective any cloud service NRENs wish to provide will be compared to existing services (i.e. Dropbox or equivalent) and therefore users would expect similar functionalities. Although these functionalities can be provided, they will make the service's long term sustainability more challenging.

Appendix – TERENA Cloud Panel minutes

TERENA Cloud Panel Notes by: Licia Florio
17 October 2011 John Dyer

TERENA Offices, Amsterdam 17 October 2011

Attendees:

Brian Boyle, HEANET
Maarten Koopmans, representing UNINETT
Szekelyi Szabolcs, NIIF
Valentino Cavalli, TERENA
John Dyer, TERENA
Licia Florio, TERENA
Peter Szegedi, TERENA

John Dyer provided an introduction to the meeting expressing the wish that the meeting should reach some agreement on how to implement the TERENA strategy on Clouds.

In short TERENA intends to:

- Support NRENs in submitting a proposal to the EC FP7-ICT-2011-8/1.2 call sue January 2012
- As part of APSIRE contribute to a study that will make recommendations on cloud provision for the NREN community
- During 2012 develop a simple pilot cloud service to gain some practical community experience in operating a multi-domain federated pilot service
- Provide a single entry point on its website for all TERENA cloud related activities

Presentation by Peter Szegedi, TERENA

Peter reported on the initiatives in the TF-Storage community; some NRENs (PSNC, NIIF, TERENA, others?) to prepare for an EC project for the January 2012 EC call.

Peter presented his idea on the scope for an EC project; the idea is to build on current NRENs experiences to integrated NRENs private clouds with a public cloud by using a platform as a service.

Maarten asked if TERENA would aim to provide guidelines only or to also procure a service.
John Dyer answered that TERENA would like to do both. We need to identify the requirements for the tender. Maarten noted that due to the variety of groups, TERENA might have to select one group and tender for them.

The Irish experience provided good learning point. HEANET had tendered for storage space the tender was very open which was a bad thing and a good thing maybe. Amazon did not respond to tenders, maybe due to:

- Amazon not willing to change their standard business model for a small community such as the NRENs community
- The main cost for amazon is not related to the costs of transiting traffic over the wide area, but it is more related to their high quality infrastructures.

Lesson learnt:

- Tender documents should not be too long
- Make it clear what you have in terms of resources and infrastructures

Updates from Maarten Koopmans on behalf of Sigma and UNINETT

Maarten presented that finding of a study (NEON) done by SIGMA in 2009. (See NEON report.)

Maarten also mentioned the "Cloud Back-Storage" service being built by SIGMA. This service uses WebDAV to talk to the storage system and to connect to the web interface offering offer federated access. Local storage and Amazon S3 will be the first storage systems available but in principle other cloud systems such as Rackspace for instance can be added in parallel.

This service is mostly driven by the HPC community who have developed a process of working together across borders, but there is no reason why such a system would not be able to support NREN/Campus users. Maarten reported that SIGMA would like to collaborate with NRENs and other institutions and TERENA could provide a focus for this cooperation. The software is open source.

NIIF activities, Szekelyi Szabolcs

Szabolcs reported that NIIF are building federated clouds that can be used in much the same way as they are using Grid Services. They are currently using an old bit stable version of the OpenNebula stack (version 1.4) during development. The intention is to develop their own private stack to replace OpenNebula. This will be Open Source software.

Next Steps

Peter Szegedi is coordinating an effort initially with PSNC and NIIF in response to the FP7 call 8 from the European Commission. The proposal will also include SMEs such as a commercial cloud provider and further NRENs that are able to contribute to the project.

John Dyer will be coordinating the ASPIRE foresight study on the provision of Cloud Services in the NREN and R&E community. It is expected that this will contain much of the report information that is required for the NREN community. The study will start in November 2011 and most of the material wil be available in 1Q2012.

We need to explore what infrastructure resources could be made available for the pilot for a 6-9 month period in 2012. Szabolcs thought NIIF might be able to offer some resources. Other NRENs would also be asked if they could contribute.

Licia Florio will coordinate a small group of NRENs to define the requirements for the pilot. NIIF and HEANET expressed an interest. Licia will contact SURFnet to see if they are also interested.

Gareth Eason of HEANET had been investigating the use of SparkleShare. Brian Boyle offered to try and talk with Gareth to see if that might also provide a starting point for the pilot or a longer term project for the community. It is not clear at this moment how much development would be required and whether SparkleShare would indeed be a viable option for a pilot in the short term.

John Dyer will create a "cloud" activity index and information page linking all the TERENA cloud-related activities in one place.

We have two email distribution lists related to these activities:

cloud-panel@terena.org for this core group discussion (15 members including 5 TERENA staff)
nren-clouds@terena.org general cloud discussion list (currently 84 members)