



**4<sup>th</sup> TF-Storage meeting**  
**Tuesday, 15 September, 2009**  
Copenhagen, Denmark

**Table of contents**

<b>1. Welcome and apologies</b> .....	<b>1</b>
<b>2. Approval of agenda</b> .....	<b>1</b>
<b>3. Minutes of last meeting and update of action list</b> .....	<b>1</b>
<b>4. Participants' presentations</b> .....	<b>2</b>
• Mixing public and private clouds, Maarten Koopmans (Vrijheid.net) .....	2
• BalticCloud, Aake Edlund (BalticCloud) .....	3
• National Data Storage in the PIONIER network, Maciej Brzezniak (PSNC) .....	4
• Pithos: experience and lessons, Panos Louridas (GRNET) .....	4
• SURFnet storage pilot plans, Rogier Spoor (SURFnet) .....	4
• Personal data traversing organisational borders: legal issues, Walter Tvetter (University of Oslo) .....	5
• Scalable Storage for Streaming Applications, Kostas Magoutis (FORTH) .....	6
<b>5. Update on FileSender development</b> .....	<b>6</b>
<b>6. Future of TF-Storage</b> .....	<b>7</b>
<b>7. Date of next meeting, aob and close</b> .....	<b>7</b>

**Minutes**

**1. Welcome and apologies**

The fourth TERENA Storage Task Force meeting was held on 15 September, 2009, in Copenhagen, Denmark hosted by UNI-C and NORDUnet. Jan Meijer (UNINETT), as the chair of TF-Storage, welcomed the participants and asked for a roll call. Valentino Cavalli substituted Peter Szegedi as the secretary of the task force.

**2. Approval of agenda**

The proposed meeting schedule was agreed with the participants without any changes although there were some minor changes in the presentation titles compared to the planned agenda. The presentations are available on the TF-Storage website: <http://www.terena.org/activities/tf-storage/ws6/agenda.html>

**3. Minutes of last meeting and update of action list**

There were some updates on the action list defined during the last TF-Storage meeting in Dublin. The comments are shown in the table below:

Reference	Who	Action	Status
Action Tsec(09)007-1	Jan Meijer (UNINETT)	Discuss with CESNET, Belnet and any others that are interested, about them joining the Flash based, open source, Poste Restante software.	<i>Done. Reported during the meeting.</i>
Action Tsec(09)007-2	Jan Meijer (UNINETT)	Update the Storage Wiki with the latest information.	<i>Still open.</i>
Action Tsec(09)007-3	David Corney (STFC)	Identify organisations and people who worry about the (legal) aspects of cross-domain data sharing.	<i>Interested but not available – action still pending, but a legal expert from Norway will talk at the meeting later on.</i>
Action Tsec(09)007-4	All	Think about how to establish a maintainable storage taxonomy document on the Storage Wiki.	<i>Pending.</i>
Action Tsec(09)007-5	David Corney (STFC)	Organise a meeting with IBM Cloud Research Centre and interested participants to discuss collaboration possibilities.	<i>Done. IBM solution requires a very specific software environment; TF-Storage people would prefer to work with open tools.</i>

#### 4. Participants' presentations

Round of presentations follows according to the agenda.

- **Mixing public and private clouds, Maarten Koopmans (Vrijheid.net)**

Maarten presented a practical perspective on mixing public and private clouds. He mentioned a paper from Barkley: "Above the clouds" that is a good introduction to cloud computing. In summary, a cloud can be seen as a pay-as-you-go public environment following the Software as a Service and Utility Computing principles.

The cost model is very promising with small upfront investment and an additional pay-per-resource component. It was suggested that the public cloud model could be an example for building private clouds, cross-NREN clouds and public-NREN clouds.

Jan Meijer was wondering why NRENs should build something themselves? Maarten replied that this would be more cost effective; Klaas Wierenga added that one benefit could be the possibility to swap among different providers as needed while still exposing the same interface to customers.

Maarten presented the results and the lessons learned from experimentation with Google App Engine and Eucalyptus carried out by university students during the summer. Google App Engine is really well suited for certain web applications. Eucalyptus works (and was effectively used for processing images in parallel) but is not mature enough. The amount of time to get an infrastructure working on par with Amazon Web Services is too much. It is expected to change within a year. The essence of clouds is resource sharing and credential management is key.

Meeting attendees were wondering about the benefit of combining the grid and the cloud models. It was remarked that support for Grid users could be combined with support to clouds users, but the user applications in the case of clouds are less well known. NRENs trust in more collaborative environment than competitive commercial offerings. NRENs should now look at the existing clouds experiences and start thinking, learning lessons, getting ready to make the right decision when the moment for that will come (expect 1-2 years from now). Klaas noted that it should be possible to join forces and take advantage of high capacity networking.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/3-MixingClouds-Koopmans.pdf>>

- **BalticCloud, Aake Edlund (BalticCloud)**

Aake summarised the facts of BalticGrid which BalticCloud is built on. The initial aim of BalticGrid Innovation Lab is to help early stage high-tech Internet based companies to try their services on new platforms, resulting in early proof of concepts and later exploitation of grid and cloud in the region.

BalticGrid asked for one centre per country to install a cloud instance. SNIC (Swedish National Infrastructure for Computing) supported the coordination of the work and took early contact with open source cloud vendors. It was decided to focus on Eucalyptus integrated with Ubuntu Linux and RightScale as an open source "cloud in a box" solution. The first two use cases were the virtual world installations (Wonderland) for the Immersive Education organization and the film rendering for Estonian Film Company.

The major concerns were the inefficient support of the open source alternatives, the missing standards for interoperability, and the various lock-in issues. The BalticCloud status, bug reports, etc. are available at <http://cloud.balticgrid.eu> (see the slides for report on individual countries' experiences).

The expectations by the end of 2009 include the BalticCloud interconnection with the Northern Europe Cloud (Sweden, Norway, Denmark, Estonia, Latvia, Lithuania, Belarus, Poland and probably Iceland and others) as well as connecting with other clouds (e.g., in The Netherlands) supporting availability zones. There are low expectations on standard APIs. In summary, it can be said that BalticCloud have learned a lot about the current open source cloud technologies. They realised that the quality assurance is missing some points but the user side is evolving. The project, at least, has created some basic courses and a cloud to play with. The next level, beyond this pilot project, is to be part of the Northern Europe Cloud.

The Northern European cloud focuses on what the cloud could provide for science, how clouds could impact on research infrastructure, etc. There is a room also for collaboration with IT

departments of universities. One of the open issues is how to get resources, because it is very hard to work with Eucalyptus due to lack of management tools, support, etc. The billing issues (specifically for Cross-NREN resource sharing) have to be considered, as well.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/4-BalticCloud-Edlund.pdf>>

- **National Data Storage in the PIONIER network, Maciej Brzezniak (PSNC)**

Maciej explained the National Data Storage project in details (see slides). The project's main objective is to provide backup/archive storage service for edu.pl, gov.pl, etc. users. It could be a secondary storage and not a primary one (it is assumed that the primary storage systems exist). The potential users are: universities, Polish Academy of Science institutes Supercomputing Centres, Polish NREN, MANs, local government units, e.g. city hall, police department, clinical hospitals, virtual laboratories, and digital libraries. The user requirements are detailed in the slides. The main design decision was to give up some extra (unnecessary) features but keep reasonable robustness level. Check the slides for the details of the architecture and the workflow.

There was a question on the wide range of potential user groups. Maciej noted that they are real contacts and have some commitment to use the system. He could provide detailed requirements if people are interested in them.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/5-PSNC-Brzezniak.pdf>>

- **Pithos: experience and lessons, Panos Louridas (GRNET)**

Panos described the Pithos system implementation both on the server (App Server: Java/JBoss, RDBMS: Postgres, Index: Solr) and the client (Web client: Google Web Toolkit (GWT), Firefox extension: Javascript, CLI: Java, WebDAV: AppleScript) sides. Note that GWT would not be used for developing the web client anymore. Pithos is supported by GRNET and has also a user forum for people to post.

It is available to all customers of GRNET (currently 22 universities and research centres) as long as they are part of Shibboleth federation and agree to pass a few attributes to GRNET when they log in. The first usability tests were positive - want to expand to cloud service in the future.

The Pithos code is available at the following URLs:

- API: <http://code.google.com/p/gss/wiki/API>
- Server code and web client code: <http://code.google.com/p/gss/source/browse/>
- Other client code: <http://code.grnet.gr/projects/pithos>

The business case behind the Pithos project was to increase the usability of the network; as GRNET is expanding in the service business this was considered as a good example to start getting experience.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/6-GRNET-Louridas.pdf>>

- **SURFnet storage pilot plans, Rogier Spoor (SURFnet)**

Rogier started with the basic motivations for looking into storage services. There are two ways to innovate; go up or go down in the OSI/ISO protocol stack. Innovation at the lower layers is not a logical choice while innovation at the higher layers can only be done at the application layer. The most logical way is to go for application based services. The major challenges are to have computer power, storage and network facilities to be fully integrated (these are not integrated yet at SURFnet) following a single, reliable and scalable architecture principle.

A remote storage project was started in 2007 with two phases. The first phase has resulted a strategy paper the second phase was a practical test how to outsource storage. Ethernet over lightpaths was used instead of IP and FiberChannel. A new storage project will start in Q4 2009. It will combine different storage techniques, store VM images and media content and offer a single mount point with easy scalability and low cost/petabyte ratio. The details of the proposed design (with the components of: DRBD (distributed replicated block device) as SAN/NAS, Mogile/Castor as cloud storage, Fuse as switching layer) can be found on the slides.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/7-SURFnet-Spoor.pdf>>

- **Personal data traversing organisational borders: legal issues, Walter Tvester (University of Oslo)**

Walter introduced the key terms like personal data, sensitive data and “the controller” as defined by the directive 94/46/EC. Sending data from one entity to another will have different implications depending on whether or not the controller is the same. The use of subcontractors (data processors) is more or less unproblematic within the scope of the directive. However, two important points must be emphasised:

- There must be a written agreement
- The controller bears the responsibility for the data processors actions (you can reduce liability in civil matters by contract-based indemnification clauses, but you cannot escape criminal liability).

If the processing is to be made by a new controller, additional rules come into play. This will be a new processing. The intended purpose of the original processing must be taken into account. Data from one processing cannot be used for another processing if the purpose of this new processing is incompatible with the purpose of the first one.

Data transfer across border could be problematic in case of third country (not EU/US). Third country must have adequate level of protection. General rule (from the directive) is that one has to abide by the rules of the country where the controller is located, as well as the rules in the country of the controller (if from EU countries) has an established presence. If no EU controller exists, he/she has to have a representative in the country where the processing will take place.

This directive is the first generation of EU directive on data protection. It does not deal with the special issues of clouds (e.g., encryption is not sufficient protection, issue who owns the encryption key, etc.). There should be a new version of this directive in less than 5 years.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/8-LegalIssues-Tvester.pdf>>

- **Scalable Storage for Streaming Applications, Kostas Magoutis (FORTH)**

Kostas introduced the EC funded STREAM project: [www.streamproject.eu](http://www.streamproject.eu) The overall architecture and the details of various functional modules were described. The aim is to provide high performance networking substrate to data streaming middleware Early results indicate significant boost over current performance, especially at speeds above 10Gbps.

The conclusion is that scalable streaming I/O is achievable. The key challenges to solve are:

- Network overhead on the I/O path
- Global timestamp-ordered atomic-append semantics
- Seamless switchover between network, storage paths

Currently there are some pilots with each of the three industrial partners of the project.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/11-ScalableStorage-Magoutis.pdf>>

## 5. Update on FileSender development

Jan did recap the general idea: easy exchange of big files. It was initially presented by CSC and other implementations were also discussed during the past meetings. The group found the SWITCH's DocExchange implementation as future proof but there were some licensing issues with the code. It was then decided to go for an open source FileSender development project.

The project has now produced the alpha1 and beta1 releases and is looking for testers. The beta2 and 1.0 releases are due on 15 November and 20 December 2009 respectively. CESNET and SURFnet will join the initial project group (AARnet, HEAnet, UNINETT) for 2.0 release.

A demo installation is available here (latency issues because it is hosted in Australia):

<http://apps2.vetvirtual.org/simplesaml/www/bin-debug/Default.htm>

There is an issue with 'big' file upload, waiting for the bug fix in Adobe Air implemented with GoogleGears.

At this stage there are two ways to support the FileSender initiative:

- The project needs more money and/or developer for further improvements. Jan invited attendees who want to join in to contact him.
- The project needs alpha/beta testers (do not need to be part of a federation).

During the meeting RedIRIS expressed their interest in testing and other participants are also excited. Jan (UNINETT) will send out a call for testers on the Storage/FileSender mailing lists.

***ACTION Tsec(09)036-1 on Jan (UNINETT) to follow up on FileSender testing/development.***

The further development of FileSender could include the possibility to connect different storage backend as well as PHP + SimpleSAMLPHP implementation for the server side.

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/10-FileSender-Meijer.pdf>>

## 6. Future of TF-Storage

The current TF-Storage mandate will expire in February 2010. The obvious question is do we want to continue?

Rogier thinks that a lot of development is still going on and TF-Storage is a very valuable platform that we should continue. In general, NRENs are being more and more interested in providing services and storage is one of them. There are common issues on storage and clouds; lot of storage initiatives in the future will have more of cloud components. Majority of the participants agreed with it and would come again to discuss such things.

Basically, the continuation and the format of the meetings are fine with the participants. Jan (UNINETT) mentioned that the task force goals are still valid and lot has been achieved. There were good talks, interesting presentations (about national activities, storage performance evaluations, etc.) and finally popped up a collaborative project (FileSender) but not much in working together so far. Maybe the time is right now to collaborate more on the storage related issues. For instance, it might be considered to extend the Northern European Cloud initiative to more countries. The point is to facilitate the way of identifying collaboration opportunities before they are actually presented at meetings and/or talk to other task forces getting brief national updates/plans like in TF-Mobility and TF-EMC2.

Valentino remarked that collaboration is not a must in a task force provided that there is sufficient motivation for exchange of information and discussion on specific technical matters.

The renewed task force will focus on storage and clouds. It will investigate new projects making links with them. Identity management should be part of the picture at this stage. Products and/or technology evaluation could be important for the participants.

TF-Storage will meet at least 2 times per year for one- or two-day events (depends on how difficult to put the agenda together). A third virtual meeting might happen. It is important to exchange information and/or produce information before planning the local activities.

***ACTION Tsec(09)036-2 on Jan (UNINETT) to draft new Terms of Reference for TF-Storage.***

The new Terms of Reference should be approved before the expiration of the current one (February 2010).

<Slides: <http://www.terena.org/activities/tf-storage/ws6/slides/9-tf-storage-future-Meijer.pdf>>

## 7. Date of next meeting, aob and close

The next TF-Storage meeting will be held on 4-5 March, 2010, in Utrecht, The Netherlands, hosted by SURFnet. It is going to be right after the start of the renewed task force mandate.

**Action list**

Reference	Who	Action	Deadline
Action Tsec(09)036-1	Jan Meijer (UNINETT)	Follow up on FileSender testing/development.	Next meeting
Action Tsec(09)036-2	Jan Meijer (UNINETT)	Draft the new ToR for TF-Storage.	By the next TTC meeting: 26 Nov. 2009.

**List of participants**

Name	Affiliation
Aake Edlund	SNIC and BalticGrid
Andreas Landhäußer	T-Systems Solutions for Research GmbH
Bill St. Arnaud	CANARIE
Brian Boyle	HEAnet
Candido Rodriguez Montes	RedIRIS
Chris Ariyo	CSC, IT Center for Science
Christoph Witzig	SWITCH
Hans Wallberg	SUNET
Harold Teunissen	SURFnet
Jacko Koster	UNINETT
Jan Meijer	UNINETT
Jean-Christophe Real	BELNET
Keith Rochford	Dublin Institute for Advanced Studies
Kimmo Koivunen	CSC - IT Center for Science
Klaas Wierenga	Cisco Systems
Kostas Magoutis	ICS-FORTH
Maarten Koopmans	Vrijheid.net
Maciej Brzezniak	PSNC, Poznan
Panos Louridas	GRNET
Paul Dekkers	SURFnet
Peter Szegedi (remote)	TERENA
Rogier Spoor	SURFnet
Rosend Llurba	NCF
Sæþór L. Jónsson	University of Iceland / RHnet
Valentino Cavalli	TERENA
Walter Tveter	USIT, University of Oslo
Zsombor Nagy	NIIF/HUNGARNET