



**5<sup>th</sup> TF-Storage meeting**  
**Thursday - Friday, 4-5 March, 2010**  
Utrecht, the Netherlands

**Table of contents**

**Table of contents** ..... 1

**Minutes**..... 1

**1. Welcome and apologies** ..... 1

**2. Approval of agenda** ..... 1

**3. Minutes of last meeting and update of action list**..... 1

**4. Next Generation Data Centre Networks** ..... 2

**5. National updates**..... 2

**6. New TF-Storage Terms of Reference** ..... 4

**7. Participants' presentations** ..... 5

**8. Date of next meeting, aob and close** ..... 7

**Action list** ..... 7

**List of participants**..... 8

**Minutes**

**1. Welcome and apologies**

The fifth TERENA Storage Task Force meeting was held on 4-5 March, 2010 in Utrecht, the Netherlands, hosted by SURFnet. Jan Meijer (UNINETT), as the chair of TF-Storage, welcomed the 25 participants and asked for a roll call. The TERENA secretary was Peter Szegedi.

**2. Approval of agenda**

The proposed meeting schedule was agreed with the participants without any changes. The presentations are available on the TF-Storage website: <http://www.terena.org/activities/tf-storage/ws7/agenda.html>

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

There were no comments on the previous action list.

Reference	Who	Action	Status
Action Tsec(09)036-1	Jan Meijer (UNINETT)	Follow up on FileSender testing/development.	Done. FileSender v1.8 BETA release is out for testing
Action Tsec(09)036-2	Jan Meijer (UNINETT)	Draft the new ToR for TF-Storage.	Done. The new ToR has been introduced to the audience.

## 4. Next Generation Data Centre Networks

Cisco was invited to give a talk about the next generation data centre networks and technologies. Wolfgang Riedel (Corporate Consulting Engineer, Cisco Systems) kindly presented Cisco's vision to the meeting participants.

First, the Ethernet standards, protocols, and interfaces have been summarised in the context of server connectivity in data centres then, the similarities between LAN and SAN functionalities and features were pointed out. Cisco introduced a unified fabric solution where the benefits of the wire speed 10 Gigabit Ethernet, the novel Data Centre Ethernet, the Ethernet over Fibre Channel (EoFC) standard, and the virtualization optimised networking were united. The remarkable features (e.g., single, well-known framing encapsulates various data centre protocols) and practical benefits (e.g., simplified and cheaper cabling) of Cisco's Data Centre Ethernet technology were deeply discussed.

Virtualization is essential part of modern data centres. Cisco presented its solutions in VM isolation and practical use cases for virtualized adapters. VN tag (VNTAG), that Cisco and VMware jointly proposed to enable switching with Network Interface Virtualization, addresses the problem of collapsing layers of bridging within a network to provide simplified management. While the Virtual Ethernet Port Aggregator (VEPA) proposal addresses a different issue with modifying the behaviour of an embedded bridge and defining a bridge extension that provides visibility of frames that would otherwise be bridged internally. Therefore, Cisco supports both solutions.

The basic Cisco concept is to embed management, unify fabrics, optimize virtualization, and remove unnecessary components (switches, adapters, management modules) in data centres. Cisco provides a single system that encompasses network (unified fabric), compute (industry standard x86), and storage. The whole solution is virtualization optimised, provides dynamic resource provisioning, efficient scale, as well as lower costs and power consumption.

For more details, please find the presentation online: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-cisco-wolfgang.pdf>

## 5. National updates

- **CESNET**

David Antoř (CESNET) presented the status and plans of the Czech National E-infrastructure, especially focusing on its storage components. MetaCentrum, the Czech National Grid Initiative (NGI) responsible for computing and storage grids, is transforming into a national e-infrastructure. Currently it has heterogeneous resources including Shared Memory (SMP) machines and clusters. The five major sites in Czech Republic are connected via 10GE lines up to 5 ms network latency. There are local and global file systems, the overall capacity is about 200 TB in disk arrays (100 TB in Brno NFSv4) and 400 TB tape library for backups/long term archive (Brno and Pilsen).

The e-infrastructure aims to build complex coordinated e-infrastructure in the Czech Republic. Potential users, case studies, and the project administrative details were introduced. The planned storage infrastructure will contain three sites with HSM systems. The estimated capacity will be about 5 PB each, according to the current capacity/price ratio. The replications and geographic distribution will be managed by the middleware with objectives strongly inspired by Norstore.

CESNET is looking for collaboration with institutions and teams with similar problems running comparable infrastructures. The main collaboration points could be on e.g., less latency-sensitive filesystems, federations, long-distance replicas, archiving, and long-term data preservation. For more details, see: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-cesnet-david.pdf>

Maciej Brzezniak (PSNC) expressed their interest to share information on tender evaluation and technical criteria. It was suggested to wait with tape storage tender part until LTO-5 (Linear Tape-Open) technology is released.

- **HEAnet**

Brian Boyle (HEAnet) presented HEAnet's Multimedia Sharing solution. This is basically a personal place for video files presented to clients as a Web Service. The access (i.e., to upload files) to the service is controlled by HEAnet's federation (edugate). The playback currently unrestricted. The system accepts many upload formats and transcode them to Flash and/or Windows Media formats.

The motivations behind the service development were driven by the specific user requirements such as trusted service provider (i.e., not YouTube, etc.), federated access, and Service Level Agreement (i.e., no uncontrolled takedowns). The user gets a simple web interface to upload files (it is also integrated with HEAnet's Large File Transfer application) and a custom flash player to play them back. The application developed in PHP/MySQL, there are Rhozet Carbon transcode servers and Isilon for backend storage.

Further developments will include support for Live streamed input, podcast support, and transcode to more formats. It might be possible to integrate with FileSender application and provide control access for playback. See: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-heanet-brian.pdf>

- **SURFnet**

Rogier Spoor (SURFnet) presented SURFnet's storage plans and gave a general update. He stated that storage costs are still high. While traditional SAN/NAS storage is perfect in case of highly transactional files such as database or virtual machine image, cloud storage solutions may precipitate traditional storage in case of parallel streaming writes (e.g., video surveillance), long-term storage (e.g., backup and archival files), geographically shared files (e.g., access from different geographical locations), and larger files with lots of read access (e.g., digital content, streaming, media, video, music, etc).

SURFnet's new storage project will combine cloud storage with traditional storage. The main objective is to store VM images and media content. The storage system will offer a single mount point with hierarchical/tiered storage. The plan is to have a highly scalable, vendor independent (i.e., combine different vendors) solution with low cost/petabyte, using open standards/protocols (and maybe open-source).

Beta version of the cloud storage API has been released and the new storage architecture will be ready by end of 2010. The detailed project timeline is available in the presentation: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-surfnet-rogier.pdf>

Arjan Peddemors (Novay) presented some end-user considerations and product overview on cloud and peer-to-peer storage (based on the study requested by SURFnet). Arjan explained the architectures of both the cloud and peer-to-peer storage solutions and showed many commercial products from both category. Many products have mixed cloud and peer-to-peer features but cloud storage services currently are more mature than peer-to-peer storage. For more details, see: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-surfnet-cloud-arjan.pdf>

Christiaan den Besten (Prolocation) gave a talk about Distributed Replicated Block Device (DRBD). Among the main usage scenarios the HA filesystem (for any application or service), HA filesystem with both sides writable (GFS/OCFS2), Xen guest 'backend storage' (live migration), and offsite backup (optionally re-syncing once a day) or offsite backup with onsite mirror (DRBD stacking) can be mentioned. The results of the performance test were promising. For more details, see: <http://www.terena.org/activities/tf-storage/ws7/slides/040310-storage5-drbd-chris.pdf>

There were no other updates presented by the meeting attendees.

## 6. New TF-Storage Terms of Reference

Jan Meijer (UNINETT), the chair of TF-Storage, introduced the new Terms of Reference proposal of the task force. Since the majority of the TF-Storage participants have expressed their interest in the continuation of the task force in the same format and style, another two year mandate is requested to be approved by the TERENA Technical Committee.

The aim of the renewed TF-Storage is to provide a forum for exchanging and promoting ideas, experience and knowledge and fostering collaborations among National Research and Education Networks and academic and research institutions on the topic of data storage, data management, and cloud storage.

The following work items are considered relevant to the Task Force:

- a. Forum for storage and cloud: organize interesting talks with speakers both from within and outside the TF-Storage participants to disseminate storage and cloud knowledge and expertise;
- b. Overview of (national) activities and deployments: gather information on relevant projects, activities and deployments and publish on the TF-Storage wiki on an ongoing basis, providing a technology watch function.

- c. File sharing services: coordinate the development and deployment of the (web based) file sharing software such as FileSender, the Swiss DocExchange and NorduShare;
- d. Measuring storage performance: maintain the storage performance measurement cookbook and investigate a cookbook for doing solid measurements on storage performance;
- e. Storage and AAI: Liaise with the TF-EMC2 on the ongoing work to use NREN AAI federations with storage services. The TF-Storage will not work on solving the problem but will assist the TF-EMC2 in its efforts.

The proposal was accepted by the meeting participants without any comments. The new Terms of Reference (proposal) is available on-line: <http://www.terena.org/activities/tf-storage/tf-storage-tor2.pdf>

***ACTION Tsec(10)006-1 on everyone to comment the new TF-Storage Terms of Reference by the 22<sup>th</sup> of March, 2010, as the latest.***

The new Terms of Reference will be approved by the TERENA Technical Committee later on.

## **7. Participants' presentations**

The second day of the meeting was started by Jan Meijer (UNINETT) who presented the history and current status of the FileSender project (<http://www.filesender.org/>). It is an open source project under BSD licence initially funded by AARNet, HEAnet, and UNINETT.

FileSender is capable to share large files in a secure and reliable way. It is a centralised service with federated access and vouchers can be sent to non-federated users. Simple web interface is sufficient under 2GB file size for easy upload, for larger files Adobe Air and Google Gears are used. In the future there are plans to move away from Gears to HTML5.

An operational beta version is available for testers in Australia and Europe. The first feedback from the 170 beta testers is very promising, they were able to send 50GB data properly and in the lab 100GB data was sent in 10 hours without any problem. The version 1.0 RC1 is expected in April, 2010. The project is open for testers, coders and release managers. You can contribute to the project with human efforts and money, as well.

More information about FileSender application functions and planned features are available here: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-filesender-jan.pdf>

Peter Szegedi (TERENA), as the leader of the Joint Research Activity and Networking Activity in FEDERICA, presented the FP7 project FEDERICA and its potential relations to clouds.

FEDERICA, an e-Infrastructure based on virtualization in both computers and network, is a fundamental tool for researchers on Future (and current) Internet. The facility allows researchers a complete control of their set of resources in a 'slice', enabling disruptive experiments at all communication layers. Particular care has been placed in reproducibility of the experiments and in the avoidance of complexity.

The physical FEDERICA infrastructure (including network and computing elements) and the applied virtualization layer were presented. On top of the physical infrastructure users are able to obtain their own virtual slices to perform wide variety of research studies. The unique

feature of FEDERICA is the real network substrate that allows user access and control down to the Ethernet layer and ensures reproducibility of the experiments with Quality of Service assurance.

Although FEDERICA has various type of resource (switches, routers, computing nodes, disks, CPU, etc.), has a virtualization layer based on VMware ESXi, and provide e.g., root access to (virtualised) resources, FEDERICA is still not a 'cloud'. What currently missing (but it is in the development phase) is all the automated functions in virtual resource/infrastructure provisioning.

FEDERICA mandate will expire in mid 2010 but the proposal for the continuation has been submitted to the European Commission's infrastructure call. In case of success, FEDERICA 2 will expand the infrastructure installing massive storage devices, connect photonic islands, and interfacing with public/private clouds. The overall aim is to provide a rich platform for multi-disciplinary researchers. More about the project: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-federica-peter.pdf>

Maarten Koopmans (Vrijheid.net) gave an update on the NEON project. The project addresses the state-of-the-art of cloud computing, exploring the cost of moving and running non-HPC jobs on a cloud computing environment, and trying to identify its risks/benefits on a short/long perspective. Focusing to the storage part of the project, the main requirements are to have: no client to access the data, transparent versioning and encryption (both transport and storage), AAI integration, and finally allow sharing of resources.

The current cloud storage is too complex for end users and there is no AAI integration at all. To improve/implement these features, NEON 'stands on the shoulders of giants' and use available solutions. The initial lessons have been learnt are summarised in the presentation: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-neon-maarten.pdf>

Marian Hellema (Koninklijke Bibliotheek, National Library of the Netherlands) had been invited to give a presentation about the digitisation of historical newspapers and its IT/storage infrastructure. KB's project founded to digitalise 60 million pages of historical text and 8 million pages of old newspapers aiming both the conservation of original documents and on-line access to cultural heritage. The project steps include the preparation phase, the digitalisation itself (done by manually scanning the pages), the data processing (metadata creation), search and retrieval, and finally the presentation (currently .pdf but can be anything else later on).

The basic storage requirements are the fast access for web services and the storing capacity of 40 TB for newspaper project and 100 TB for other digitised documents. The most important thing is to do the metadata creation right. Open Archives Initiative Protocol for metadata harvesting is applied at KB. The project uses Dublin Core, MPEG21-DIDL and optionally more formats, Oracle database, and Verity K2 search engine. More about the digitalisation project can be found here: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-bibliotheek-marian.pdf>

Leif Johansson (NORDUNet/SUNET) spoke about NORDUshare. NORDUshare is a data distribution platform based on bittorrent technology. Superficially, NORDUshare looks like most torrent sites but with an authentication and authorization infrastructure on top of that. The use cases have been considered during the design phase are: sensor networks, provisioning of bulk content (e.g., video, SCORM objects) for shared LMS or media distribution systems, distributed "FTP" service, and VO data sharing. The design objectives include programmable API, basic client for entry-level users (java-based), compatible with current thinking around federation

VOs (entitlements), scalable (no single point of storage), and behave well in the face of mobile, intermittent and varying quality of Internet connectivity. The details of the service architecture can be found in the presentation: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-nordushare-lief.pdf>

The first public beta version of the software (working name Lobber) will be available by May, 2010 using Kalmar Union (eduGAIN is also a possibility).

The last speaker, Zsombor Nagy (NIIF/Hungarnet), presented the Chelonia storage system, which is a lightweight, self-healing, distributed storage. The motivations behind the software development were to create a storage resource from available disk space on any computer, to connect storage resources through the Internet, to provide an easy-to-understand way for users to upload and share their files, and to use these files in a Grid environment. The system should be self-healing, easy to deploy, easy to maintain, lightweight and with no single point of failure.

Chelonia has been developed by the KnowARC project and will be maintained by the NorduGrid Collaboration. The architectural details can be found here: <http://www.terena.org/activities/tf-storage/ws7/slides/050310-storage5-chelonia-zsombor.pdf>

Finally, Zsombor briefly presented the ClusterGrid project in Hungary, maintained by NIIF, the Hungarian NREN. Some new storage equipment will soon be purchased and deployed at selected Hungarian universities with Chelonia software on top. The ClusterGrid users will be able to upload and download (input/output) files to the Chelonia cloud and refer to these files in the job descriptions.

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

The next coming half-day TF-Storage meeting will be held on Thursday afternoon, 3 June, 2010, co-located with the TNC 2010 conference in Vilnius, Lithuania. The plan is to exploit the unique opportunity and invite large cloud vendors to give technical presentations.

***ACTION Tsec(10)006-2 on Jan and Peter to invite large cloud vendors/providers to the Vilnius meeting.***

The autumn meeting will be held on the 9-10<sup>th</sup> of September, 2010. PSNC kindly offered to host the meeting in Poznan, Poland.

## Action list

Reference	Who	Action	Deadline
Action Tsec(10)006-1	all	Final comments on the new TF-Storage Terms of Reference	22 March 2010
Action Tsec(10)006-2	Jan, Peter	Invite cloud vendors/providers to the Vilnius meeting	end of April, 2010

## List of participants

Name	Affiliation
Arjan Peddemors	Novay
Auke Folkerts	Universiteit van Amsterdam
Brian Boyle	HEAnet
Christiaan den Besten	Prolocation
Damir Danijel Zagar	Srce
David Antos	CESNET
David Corney	Rutherford Lab UK
Faidon Liambotis	GRNET
Harri Salminen	CSC
Henny Bekker	SURFnet
Jan Meijer	UNINETT
Jeroen Roodhart	Universiteit van Amsterdam
Klaas Wierenga	Cisco Systems
Leif Johansson	NORDUnet/SUNET
Maarten Koopmans	vrijheid.net
Maciej Brzezniak	PSNC, Poznan
Marian Hellema	National Library of the Netherlands
Mario Vandaele	Belnet
Martin Kämpf	SWITCH
Paul Dekkers	SURFnet
Peter Szegedi	TERENA
Rogier Spoor	SURFnet
Rossend Llurba	NCF
Wolfgang Riedel	CISCO Systems
Zsombor Nagy	NIIF/HUNGARNET