



1ST TF-Storage meeting minutes

by Peter Szegedi (TERENA)

Planned agenda

- 09:30 - 09:45** Welcome, agenda bashing, round of introductions
09:45 - 10:00 Minutes and summary of last meeting, status of action items
10:00 - 10:30 USENIX FAST report (Jan Meijer)
10:30 - 11:00 Storage measurements (Stijn Eeckhaut)
11:00 - 11:15 Coffee
11:15 - 12:15 * Petabyte storage @ STC - Rutherford labs (David Corney)
* NorStore, petabyte storage in Norway (Jacko Koster)
TF-Storage work items
* Poste restante
* Storage middleware taxonomy
* Storage technology taxonomy
12:15 - 12:45 * Measuring storage performance
* Storage and AAI
* Energy consumption considerations
* Backup and/or disaster recovery services
* Overview of (national) activities and deployments
* Collection of best practices and service requirements
12:45 - 13:45 Lunch break
13:45 - 15:15 TF-Storage work items (continued)
15:15 - 15:45 Coffee
15:45 - 16:45 TF-Storage work items (continued)
16:45 - 17:30 Cisco Research on Storage (Bjørn Martinussen)

Meeting minutes

09:45: Introductions

- Jan (UNINETT) welcomed the participants. Round of introductions.
- Agreement on the agenda.
- Jan (UNINETT) explained the briefings of the last two Storage meetings, where the audience indicated willingness to participate and contribute to a joint storage activity. That is why it has been decided to formulate a TERENA Task Force.
- The technical aims of the task force and the detailed Terms of Reference was prepared by Jan Meijer and Magnus Stromdal (UNINETT), finally the ToR has been accepted by the TERENA TTC.
- There were no comments on the last meeting minutes.

- Jan (UNINETT) mentioned that 3 face-to-face TF-Storage meetings pro year would be sufficient. No comments.
- Peter (TERENA) noted that the storage mailing list remains the same (no tf-storage!), there was no objection against that. The availability of the TF-Storage web page <http://www.terena.org/activities/tf-storage/> and the TF-Storage wiki <http://wiki.terena.org/index.php/TF-Storage:Contents> was also mentioned.

10:00: USENIX FAST report

- Jan (UNINETT) provided some information about FAST findings. One of the important messages is the emerging reliability of the storage stuff.
- Jan (UNINETT) mentioned that because of the cheap storage equipment usage (hardware parts, disk controllers, etc.) there is a high probability of unexpected failures. There is a work in progress on the modelling and measurement of storage solutions.
- Pietari (CSC) noted that the expensive stuff does not mean less failure. That would be useful to have a list of trusted vendors in sense of reliability.

10:10 HEAnet - Storage Procurement

- Brian (HEAnet) held a presentation on the storage procurement at HEAnet: http://www.terena.org/activities/tf-storage/ws3/II_heanet-storage.pdf HEAnet has modest requirements. This tender is about those requirements only. That said 'we would like to see if clients can make use of the technology'.
- Niels (SURFnet) asked about the network implementation.
- Brian (HEAnet) answered that in a small country some headends connected via GE connections is sufficient. It is under the EC tender level of procurement (e-tender: suitable for projects less than EUR 220k budget).
- Jerome (RENATER) asked about the cost of the storage components.
- Brian (HEAnet) answered that they will use relatively cheap equipments.
- Kostas (GRNET) asked about the size of the storage system.
- Brian (HEAnet) answered that the minimum is 40TB storage capable of scaling towards 1PB.
- Jan (UNINETT) asked about the AAI solution.
- Brian (HEAnet) answered that no strong AAI requirement/implementation.

10:20 SURFnet – VP-X

- Niels (SURFnet) held a presentation about a webservice oriented platform for third party content delivery: http://www.terena.org/activities/tf-storage/ws3/III_VP-X.pdf VP-X is a middleware audio/video platform which allows institutions to plug in this backend into local applications. Important elements are the open source licensing and no storage lock-in.
- Cándido (Red.es) asked about its integration into the federation of Surfnet.
- Niels (SURFnet) answered that this is a back end application, no user

- interface, not integrated. But the front end applications will be integrated.
- Brian (HEAnet) asked: Do you have users developing front end applications? Is it a useful tool for them?
 - Niels (SURFnet) answered that there are some users developing that and they are satisfied.
 - Pietari (CSC) and Kaspars (SigmaNet) asked several questions about the transcoding issues: when, how, what is the resolution.
 - Niels (SURFnet) answered that they store all of them; the files will be transcoded in the same size.

10:40 IBBT Univ. Ghent - Storage benchmarking cookbook

- Stijn (IBBT) held a presentation about how to perform solid storage performance measurements: http://www.terena.org/activities/tf-storage/ws3/IV_StorageBenchmarkingCookbook-StijnEeckhaut-final.pdf The storage systems can be very complex, complexity can impede correct measurement of the storage system. The bottom lines are: the reproducible measurement, the applied test load reflects the real load of the system and to measure the right bottleneck.
- Pietari (CSC) mentioned that usually there is no access to the benchmark, closed by the vendors.
- Stijn (IBBT) noted that it has a cost; you have to pay for the access.
- Jan (UNINETT) asked that the academic use is free?
- Stijn (IBBT) answered that it is not sure. Benchmark is a comely stuff we have to apply for.

11:00 Coffee break

11:30 STFC - Large scale Data Storage Services for Science at STFC

- David (STFC) held a presentation about the Science and Technology Facilities Council (UK), user communities, major trends and Data Storage and Management Services: http://www.terena.org/activities/tf-storage/ws3/V_TERENA-Apr08-DCorney.pdf
- Kostas (GRNET) asked about the legacy dCache service.
- David (STFC) answered that: we are running dCache (phase out soon) but CASTOR is also good enough. For us it is fine.
- Jan (UNINETT) asked about the planned man-months spent on iRODS.
- David (STFC) answered that there is no fixed budget. It is not clear how many efforts will be on it. Do not drop it!
- Andreas (T-Systems) mentioned that there will be a conference on iRODS:
Building data grids with iRODS
27 May, 08 09:30 AM - 30 May, 08 04:00 PM
e-Science Institute, 15 South College Street, Edinburgh
<http://www.nesc.ac.uk/esi/events/866/>
- That could be interesting for all of us.

12:05 UNINETT – NoreStore

- Jacko (UNINETT) held a presentation on national storage infrastructure for Norwegian research: http://www.terena.org/activities/tf-storage/ws3/VI_norstore_janmeijer-apan.pdf The presentation was about the trends in Norwegian e-Infrastructure and current investments. The main objective of NoreStore: Establish and maintain a broad and sustainable infrastructure for the curation, archiving and preservation of data from computational science and the natural sciences.
- Niels (SURFnet) asked about how data remains accessible in case of failures using the NoreStore infrastructure.
- Jacko (UNINETT) answered that it is not just storage but services on top of that (metadata services).
- Jerome (RENATER) asked that: Are 10G links dedicated?
- Jacko (UNINETT) answered that only two links are dedicated 10G, but more dedicated links are needed in the near future because of the data growth.
- Jan (UNINETT) asked that is there any European collaboration like this?
- Jacko (UNINETT) answered yes, e.g. in Sweden (same storage policy, same backup technology, etc.)

12:30 Discussions on the TF-Storage work items

E - Taxonomy of storage (virtualization) middleware

- Lars (NORDUnet) presented some slides on storage systems taxonomy: http://www.terena.org/activities/tf-storage/ws3/VII_TF%20Storage%20-%208%20April%202008.pdf
- The main goal is to create an accessible overview of existing systems. Strong experiences are needed. As many contributors as many platforms! It could be hard to deal with the similarities/differences. The main question is: what stuff for which purpose?
- To do list:
 - o Select platforms to include
 - o Find experts to describe platforms
 - o Collect platform synopsis'
 - o Produce overall synopsis
 - o Pick classification attributes
 - o Produce taxonomy
 - o *Iteration may be needed*
- Jan (UNINETT) mentioned that the list of partners has been started to collect.
- Brian (HEAnet) was volunteered to lead this work item.
- Andreas (T-Systems) will contribute to this work item.
- Kostas (GRNET) noted that the most efficient way to get information about the platforms is to invite vendors (SUN, IBM, etc.) for the next TF-Storage

meeting.

- It was agreed that a half-day meeting with vendors would be useful/sufficient during the next meeting.

12:45 Lunch break

13:50 Discussions on the TF-Storage work items ctd.

F - Taxonomy of storage technologies

- Review on storage technologies is needed (draw a tree of available technologies).
- Compile experiences (independently from vendors).
- Product evaluations vs. Believe in what the vendor says?
- Andreas (T-Systems) noted that during the product evaluations only list prices can be comparable. But prices are not really interesting (just roughly cheap or expensive).
- Pietari (CSC) mentioned that the total cost depends on the support, as well.
- Brian (HEAnet) was volunteered to contribute (not lead) to this work item.

C - Poste Restante service

- Ari (CSC) presented a brief status report on the ePoste Restante and ePort Payé services: http://www.terena.org/activities/tf-storage/ws3/VIII_ePosteRestante_2008-04-08.pdf
- Niels (SURFnet) asked that the code is available for testing?
- Ari (CSC) answered, not yet. It will be ready to release in GPL license.
- Brian (HEAnet) mentioned that they do the same service for large file upload.
- Christoph (SWITCH) noted that there is a similar (internal prototype) university project in Switzerland. Update about it on the next meeting!
- Red.es, SURFnet, SWITCH, HEAnet are willing to contribute to this work item.

I - Storage and AAI

- Cándido (Red.es) presented some slides about Integrating AAI federations with storage access methods: http://www.terena.org/activities/tf-storage/ws3/IX_Storage%20and%20AAI.pdf
- Which federation technologies: Shibboleth/eduGAIN? They are almost equals.
- The problem is all the federations are web based (only browser). Issue to bridge this gap!
- Possible action items:
 - o Integrating present storage access methods with AAI
 - o Build better storage access methods integrated with AAI
 - o Integrate ssh/scp/sftp with AAI federations
 - o Check if the Grid community has integrated with AAI the grid-enabled
 - o ssh/scp/sftp

- Christoph (SWITCH) presented some slides on the Data Management Related Activities in Switzerland, project description and use cases:
http://www.terena.org/activities/tf-storage/ws3/X_080408_storage_ws_chw.pdf
- The profile is available on the SWITCH website.
- The main message is: AAI people are willing to do something in this area, we have to drive them!

15:30 Coffee break

16:10 Next meeting

- Jan (UNINETT) announced that the next TF-Storage meeting will be held in Riga, Latvia, hosted by SigmaNet.
- There was an agreement on the date of the next meeting: September 12, 2008.

16:15 Discussions on the TF-Storage work items ctd.

J - Energy consumption considerations

- Klaas (Cisco) mentioned that Cisco is interested in this work item.
- Niels (SURFnet) mentioned that SURFnet is interested a bit. Currently there is an ongoing 'Green office automation' project.
- T-Systems and RENATER are also interested in this work item.

G - Measuring storage performance

- The starting point is the presentation of Stijn (IBBT) see above.
- Univ. Ghent is interested in this work item, maybe PSNC and HEAnet?
- Pietari (CSC) noted: Not defined yet what the benchmark is.

H - Backup and/or disaster recovery services

- There was not too much discussion.
- SURFnet is interested in this work item.
- Andreas (T-Systems) will contribute to this work item.

16:40 Cisco – Research on storage

- Bjorn (Cisco) held a presentation about the Cisco's research on storage. Cisco has lot of investment in storage. Note that: Cisco is not storage company but interconnect.

17:45 Closing the meeting

- Jan (UNINETT) asked the participants: What would you like to see in the next meeting?
 - o Vendor presentations on storage platforms.

- o Presentations on how can be ensured that the data is still ok.
- Jan (UNINETT) closed the meeting.

PARTICIPANTS:

Brian Boyle	HEAnet
Maciej Brzezniak	PSNC
David Corney	Rutherford Appleton Lab
Paul Dekkers	SURFnet
Jerome Durand	RENATER
Stijn Eeckhaut	IBBT - Ghent University
Lars Fischer	NORDUnet
Pietari Hyvärinen	CSC
Sæþór L. Jónsson	RHnet
Jacko Koster	UNINETT Sigma
Kostas Koumantaros	GRNET
Kaspars Krampis	IMCS UL, SigmaNet
Andreas Landhäußer	T-Systems Solutions for Research GmbH
Rossend Llurba	Netherlands National Computing Facilities Foundation (NCF)
Dana Ludviga	IMCS UL, SigmaNet
Ari Lukkarinen	CSC Scientific Computing
Bjørn Martinussen	Cisco
Jan Meijer	UNINETT
Zsombor Nagy	NIIF (HUNGARNET)
Peter Szegedi	TERENA
Cándido Rodríguez	RedIRIS
Miroslav Ruda	Cesnet
Harri Salminen	CSC / Funet
Olaf Schjelderup	UNINETT
Magnus Strømdal	UNINETT
Niels van Dijk	SURFnet
Klaas Wierenga	Cisco
Christoph Witzig	SWITCH