NRENs' Strategic Perspective on Storage and Cloud

"Build or Buy"

Green Paper v0.5

Peter Szegedi (TERENA)

April 2011

Preface

By definition, a green paper is open ended. Green papers, also known as consultation documents, may merely propose a strategy to be implemented in the details of other conditions or they may set out proposals on which the author wishes to obtain public views and opinion.

Acknowledgements

This green paper is based on the information provided by the TERENA Storage Task Force (TF-Storage) and Management of Service Portfolios Task Force (TF-MSP) participants so far. It also summarises the outcome of the outsourcing panel discussion at TERENA Networking Conference 2010 and contains references to press releases, news items, and white papers recently issued by the NREN community.

The author gives special thank to NREN people who have provided detailed information about the storage and cloud related strategy of their NREN.

Table of Content

1. Introduction

This paper aims at providing an overview on the NREN's strategic perspective related to storage technologies and services (with special focus on cloud storage) as well as to cloud services and service outsourcing in general, as it has been initiated by the TERENA task force TF-Storage participants. The primary focus of the task force is on storage so first of all the storage service scenarios are discussed. The wide variety of NREN services related to data storage and content repository is also summarised in this paper. Since cloud services are emerging (and commercial cloud service providers are approaching the research and higher education community intensively) NRENs are started to work out strategies related to both outsourcing their commodity services to public clouds and keeping their core services and ability to innovation in-house. Universities are also facing with the question of operating campus services in-house or outsource them to their NREN or even to commercial cloud providers that are available globally. This paper gives an overview on various aspects of outsourcing as well as on the potential strategic choices related to storage, cloud, and e-Infrastructure services facilitated by NRENs.

2. Storage services

E-science infrastructures, supporting higher education and research purposes, must have wide variety of storage devices due to the data-intensive education and scientific applications that generate ever-increasing amount of digitalized information. The storage services of those infrastructures can be divided into two major categories:

- Primary or production storage serves active applications and is accessed randomly. The primary category includes most familiar direct-attached disks (DAS), storage area networks (SAN), and network-attached storage (NAS). Newcomers in the primary category include content-addressable storage (CAS) and cloud storage services.

- Secondary storage is used for data protection and is normally accessed sequentially. Tape media and optical discs were the traditional secondary storage types, but disk-based systems including virtual tape libraries (VTL) have recently become popular. CAS and cloud systems are also often used for secondary storage due to their lower cost.

The performance and capability of primary storage systems varies greatly, as does the price. Campus network designers as well as NRENs are facing with a wide variety of primary storage options. Dozens of vendors build and sell storage devices, and these leverage a variety of (sometimes proprietary) connectivity protocols. Each type of storage presents a trade-off in a number of areas, from performance to cost. There is no intrinsic reason to reject one type or adopt another. The selection process must take into account the technical and business requirements of the scientific applications that will use it [1]. TERENA TF-Storage facilitates the efficient information sharing on the aforementioned aspects among the task force participants.

2.1. Traditional non-distributed SAN/NAS storage

Traditionally, the storage market has been divided between block-based SAN and file-based NAS. The limitations of these block- and file-focused paradigms have led to the development of content-addressable storage. CAS systems discarded traditional protocols and concepts in favour of application-focused APIs and a universal naming standard for unique objects. Some scientific applications take advantage of the capabilities of the unique capabilities of these systems, especially in the document management and archiving space [1].

	DAS	SAN	NAS	CAS
Example	Seagate disk Dell PowerVault	HDS USP 3PAR InServ	NetApp Filer Windows server	EMC Centera Caringo CAStor
Protocol	SATA, SAS	FC, iSCSI	SMB, NFS	API, XAM
Access Method	block	block	file/directory	object
Connectivity	Copper cable	Fiber optic Ethernet	Ethernet	Ethernet
Throughput	1.5 Gb/s- 3.0 Gb/s	1 Gb/s- 10 Gb/s	100 Mb/s- 10 Gb/s	100 Mb/s- 1 Gb/s
Latency	5-10 ms	5-10 ms	20-50 ms	50-100 ms
Use case	Bulk storage OS/boot	Enterprise applications	Unstructured data	Archival data

Tab. 1 Large number of options is available for primary storage [1]

Traditional storage capacity has high availability and excellent performance but not scalable and can be costly due to the fact that it is closely tied to the local computing facilities in use. Many storage solutions that are currently used in the e-Science infrastructure have a limited, non-scalable approach or they use specially developed storage solutions that are only suitable in a specific research area such as nuclear physics or astronomy. Scalable, high-capacity commercial solutions could be applied in these environments, but those are expensive. Additionally, the non-distributed storage technologies often do not exploit the unique feature of high-capacity network infrastructures made available by NRENs [2].

2.2. Cloud storage

Cloud storage is basically a storage accessed over a network (internal or external) via Web Services APIs. It leverages the openness of the Internet and modern programming concepts, incorporating the Internet Protocol (IP), HTTP, SSL, REST, and SOAP. Cloud storage was developed independently from all historical storage concepts, although it might appear to be an evolution of CAS. Both are object-based, use APIs rather than traditional storage protocols, and include per-object metadata but cloud storage goes further in terms of application integration and programmability [1].

Cloud storage services, such as those offered by Amazon and Google, are examples of a fully distributed storage facility. These services have a number of attractive features, such as pay per usage and (some) guarantees that data is safely stored (through redundancy). They do not, however, provide high speed access and support for a wide range of different applications. This is also true for storage solutions applied to grid computing; in a grid infrastructure, storage facilities are focused on supporting computation at grid nodes, not on providing high-speed data access to applications at the edges of the network [2].

Most cloud systems include basic web interfaces. More interestingly, many interface solutions have been developed to bridge traditional storage protocols to the cloud. One major contributor to the success of Amazon's S3 storage offering was Jungle Disk, a consumer-oriented application that allows users to automatically back up their files to the service. Nirvanix developed CloudNAS for enterprise users, which presents cloud storage service as a Linux filesystem or Windows drive. EMC and Emulex recently revealed that they are working on a bridge between block-based SANs and cloud storage [1].

2.3. Peer-to-peer storage

Recent years have seen a strong increase in networked storage solutions for end users, especially in the area of cloud storage services. Another networked storage technology (now more and more available) relies on members of a group storing each other's data. This technology is called peer-to-peer (P2P) storage [3].

Unless storing a small quantity of data, cloud storage services are paid services, where the user pays, for instance, per amount of data kept in storage and the amount of network traffic generated to upload or download files. An opportunity exists to realize a distributed storage facility without the need for a paid subscription, but rather trade local resources (disk space, network bandwidth) to store data at peers, using P2P technologies. A large number of research projects addressed this topic and currently a number of relatively new products exist that allow users to collaboratively create a P2P storage facility. From the perspective of the user, the cost of sharing existing resources with peers may compare favourable to the repeated monetary cost for a storage subscription. Especially, when these resources are abundant and of high quality, P2P storage may be a viable and attractive alternative to both traditional storage methods as well as cloud storage services [3].

SURFnet study [4] presents an overview of existing P2P storage concepts and products (Fig. 1) that may be useful to NRNEs and campus networks.

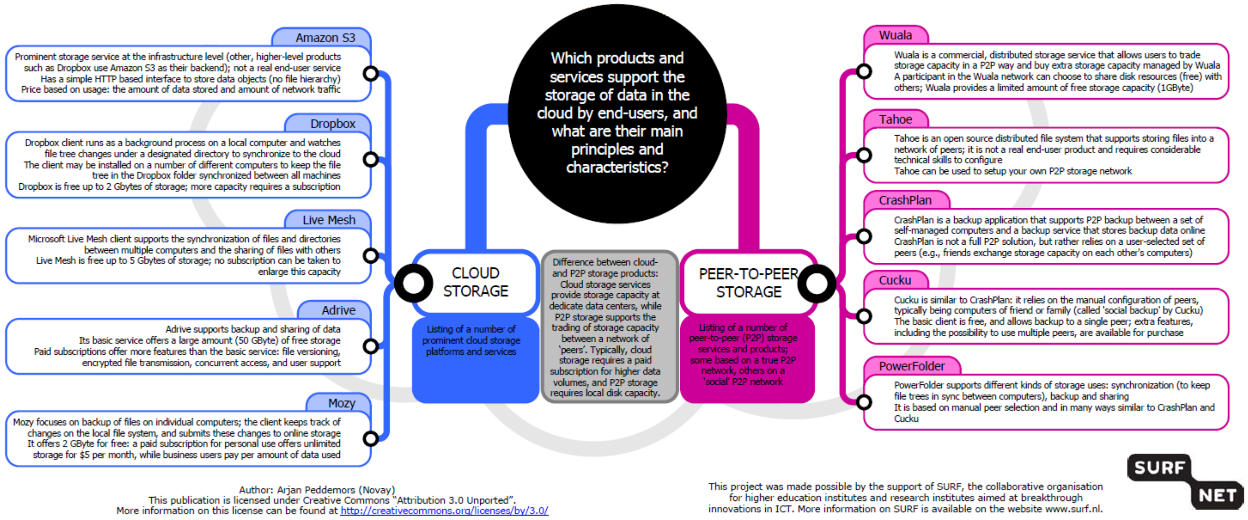

Fig. 1 Overview on cloud and peer-to-peer storage products [4]

In summary it can be said that the entire storage industry is moving toward greater levels of application awareness and integration. This is a necessary step to bring about a real storage revolution that will see a transition from bulk management of capacity to granular management of data to integrated use of information.

2.4. Storage related services of NRENs

The TERENA Compendium (2010 edition) [5] gives a snapshot on the actual housing, storage, hosting, and content delivery services provided by NRENs in Europe. NREN users want access to a range of services to support their teaching, learning and research activities. The survey focused on six areas in this category:

- Distributed storage specifically for Grid users;
- Distributed storage for any NREN users;
- Dedicated/special high-level connectivity to commercial content servers or commercial content;
- Hosting of commercial content servers or commercial content on the NREN network;
- Video servers for use by NREN sites;
- Mirroring of content from outside the NREN network.

For each of these areas, NRENs were asked to indicate whether they currently deploy the service, are planning to deploy it, or have no interest in it. The results are summarised in Fig. 2.

	Grid storage	Data storage	Peered commercial	Hosting commercial	Video service	Mirroring
EU/EFTA countries	31%	17%	21%	28%	45%	72%
Other European and Mediterranean countries	50%	8%	25%	0%	42%	33%

Fig. 2 Storage and related NREN services [5]

Mirroring (backup) is the service that seems most popular in the EU/EFTA area, having grown from 16 NRENs in 2008 to 21 NRENs in 2010. 13 EU/EFTA NRENs (45%) currently offer a content repository and distribution (i.e. video service) and 8 more are planning to introduce it. This is just one of a range of real-time and recorded multi-media services that are currently being investigated by NRENs (e.g., in TF-Media). This development warrants further investigation in future years [5].

However, the aforementioned six categories investigated by the TERENA Compendium do not represent the full set of storage related NREN services. The TERENA Task Force on Management of Service Portfolios (TF-MSP) [6] has been working on the common description of services and service categories in order to draw the map of NREN's service portfolios. According to the latest definitions [7], the storage related NREN services can be categorised as follows:

- Housing /co-location facilities
- Webhosting /hot standby
- Mail relay / back up services
- Disaster recovery/off site back up services
- Storage Area Network (SAN) infrastructure
- Netnews/Usenet server
- Academic/educational software distribution: frame agreements & clearing
- FTP & Mirroring services (proprietary & non proprietary software, Wiki, etc.)
- Hosting services/applications for research and educational community (e.g., scientific databases, Wiki, administration tools) ,
- Media storage/streaming facilities: media portals, streaming facilities (streaming server, podcasting, peer-to-peer facilities), media conversion services

It can be seen that the wide range of storage related NREN services cannot be restricted to data storage only therefore the emerging cloud services cannot be restricted to cloud computing or cloud storage only. In the next section the cloud services are discussed in general.

3. Cloud services

Cloud environment is developing very rapidly. Although it is best known for cloud computing the cloud paradigm extends to a range of services. The term "cloud computing" means different things to different people and type of institutions. It is often used to loosely describe a broad range of activities, ranging from outsourcing a specific activity to a single external provider (which many would argue is not cloud computing) to delivering a set of services from the cloud in such a way that users are not even sure where their data is being housed or where it is being processed.

Obviously, the cloud concept is still changing and it is hard to find a definition for clouds. One comprehensive definition could be:

Clouds are a large pool of easily usable and accessible virtualized resources (such as hardware, development platforms and/or services). These resources can be dynamically reconfigured to adjust to a variable load (scale), allowing also for an optimum resource utilization. This pool of resources is typically exploited by a pay-per-use model in which guarantees are offered by the Infrastructure Provider by means of customized SLAs [8].

The set of features that most closely resemble the minimum cloud definition would be ease of use (typically via a web interface), scalability, pay-per-use utility model, and virtualization.

Preliminary findings of recent JISC (Joint Information Systems Committee in the United Kingdom) studies [9] confirm confusion about the terms cloud and cloud computing and also suggest that for some in the research computing community, cloud computing is just another term for grid computing. The Next Generation Grid expert group (NGG) [10] has developed a vision which underpins the evolution of grid from a tool to solve compute and data intensive problems towards a general-purpose utility infrastructure. Grids need to accelerate the incorporation of virtualization technologies to gain some advantages that clouds natively present (e.g., migrability, hardware level scalability). Several approaches exist that combine clouds and grids together, which can also be seen as a combination of advanced networking with sophisticated virtualization. However, clouds are also said to offer a limited set of features exposed (i.e. they present a higher abstraction level to the user). For instance, the Simple Storage Service by Amazon can be regarded as a limited data grid when compared to the CERN Data Grid [8].

3.1. Cloud service classification

Cloud is typically divided into three levels of service offerings as showed in Fig. 3; Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). These levels support virtualization and management of differing levels of the solution stack [11].

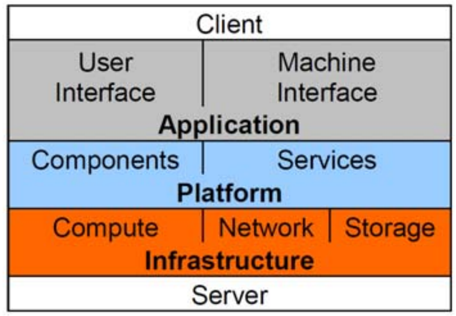

Fig. 3 Cloud service stack [11]

Depending on the type of provided capability, there are three scenarios where clouds are used:

- Application or Software as a Service

  There are services of potential interest to a wide variety of users hosted in cloud systems. This is an alternative to locally run applications. An example of this is the online alternatives of typical office applications such as word processors. This scenario is called Software as a Service (SaaS). A SaaS provider typically hosts and manages a given application in their own data centre and makes it available to multiple tenants and users over the Web. Some SaaS providers run on another cloud provider's PaaS or IaaS service offerings. Oracle CRM on Demand, Salesforce.com, and Netsuite are some of the well known SaaS examples.

- Platform as a Service

  Cloud systems can offer an additional abstraction level; they can provide the software platform where systems run on. The sizing of the hardware resources demanded by the execution of the services is made in a transparent manner. This is denoted as Platform as a Service (PaaS). It facilitates development and deployment of applications without the cost and complexity of buying and managing the underlying infrastructure, providing all of the facilities required to support the complete life cycle of building and delivering web applications and services entirely available from the Internet. This platform consists of infrastructure software, and typically includes a database, middleware and development tools. A virtualized and clustered grid computing architecture is often the basis for this infrastructure software. Some PaaS offerings have a specific programming language or API. For example, Google AppEngine is a PaaS offering where developers write in Python or Java. EngineYard is Ruby on Rails. Sometimes PaaS providers have proprietary languages like force.com from Salesforce.com and Coghead, now owned by SAP.

- Infrastructure as a Service

  Infrastructure providers manage a large set of computing resources, such as storing and processing capacity. Through virtualization, they are able to split, assign and dynamically resize these resources to build ad-hoc systems as demanded by customers, the service providers. They deploy the software stacks that run their services. This is the Infrastructure as a Service (IaaS) scenario. IaaS is the delivery of hardware (server, storage and network), and associated software (operating systems virtualization technology, file system), as a service. It is an evolution of traditional hosting that does not require any long term commitment and allows users to provision resources on demand. Unlike PaaS services, the IaaS provider does very little management other than keep the data centre operational and users must deploy and manage the software services themselves just the way they would in their own data centre. Amazon Web Services Elastic Compute Cloud (EC2) and Secure Storage Service (S3) are examples of IaaS offerings.

3.2. Question of outsourcing services to clouds

This section summarises the outcome of the "outsourcing or not" panel discussion that was held during the TERNEA Networking Conference 2010, in Vilnius, Lithuania [12]. The discussions considered two basic scenarios:

- Universities and higher education institutes outsource services to public clouds or to their NREN.
- NRENs and research organisations outsource services or sub-services to public clouds.

The general question to the panellists was that what kind of services, sub-services or functions can be outsourced by Universities and/or NRENs to public cloud service providers, how, and under what conditions and circumstances?

The participants agreed that the core services of an organisation (either it is a university or NREN) must be kept in-house! But the definition of core services is obviously different in case of university IT departments, research institutes, and NRENs.

Taking the example of e-mail service, even for larger universities, the staff e-mail and student registration services are part of the university's core business so that should not be outsourced. However, students are demanding for outsourced e-mail services (e.g., to Google Mail) simply because they can get 10G storage space there for free while the e-mail service run by the university often has no such option. It was noted that students are re-directing their mails, anyway. Here comes the data protection issue in the picture that is more relevant to university staff e-mails rather than to student e-mails. Distinction between outsourced student and in-house staff e-mail services can be an appropriate solution but may overcomplicate the whole administration.

For data intensive research institutes, the whole IT department can be considered as core (including the storage infrastructure) however, the Amazon cloud storage service provides attractive features (i.e. just type in the credit card details and get the resources immediately) to researchers. In case of NRENs some basic support services, like the calendar system, can easily be outsourced because that is not part of the NREN's core interest. But the outsourcing of specific Network Operation Centre functions should be done carefully because that is part of the core competence of an NREN. A bad example can be the outsourcing of website support of higher education institutes that makes each and every website update costly (i.e. pay per change) and can slow down or break the primary communication channel (i.e. the public face) of the university. Outsourcing cannot be done without thinking!

In conclusion it can be said that those services can be outsourced where there is a clear vision and expectation regarding the future operation of the service so the organisation can simply get rid of the commodity stuff but keep the core service parts in-house. The clever contracting means that the organisation can take the advantage of economy of scale while keep the organisational control over the outsourced service.

Taking the aforementioned two scenarios into account the panel concluded that:

- In case of universities and higher education institutes the outsourcing of commodity application services (e.g., student e-mails, document sharing) to public clouds can be done with low risk (SaaS scenario). However, joining forces at national level and let the NREN to negotiate with the commercial cloud service providers on behalf of a group of universities can have significant cost benefits and can simplify the procurement process. NRENs can play the coordination role of national "buying syndicates" here.

  The outsourcing of infrastructure related services (e.g., storage, computing) to public clouds has more risks for individual universities (IaaS scenario) concerning the service operation, data protection, authentication and access control issues. In case of these services the outsourcing to NRENs (where it is appropriate) has lower risk and universities can have natural trust in their NREN. NRENs can play the role of an infrastructure service provider or an infrastructure service proxy to public clouds here.

- In case of NRENs, the outsourcing of commodity application services (e.g., calendar system) to commercial clouds seems to be straight forward (SaaS scenario). The NRENs should get rid of the commodity services and concentrate on new service development and innovation for the benefit of their users.

  In case of infrastructure-related services (e.g., network operation, videoconferencing, storage, computing) the mixing of NRENs own infrastructure service with public clouds seems to be a value-added IaaS scenario. NRENs can hide the non-attractive features of public clouds and provide commodity cloud service to universities with tailor-made features (including federated access, data protection assurance, etc.) exploiting public cloud back-ends.

In the followings we will focus on NRENs' strategy in application services (SaaS) as well as infrastructure services (IaaS). From the storage perspective both applications as well as infrastructure can partly be outsourced to public clouds but the appropriate NREN strategy should be chosen that is not trivial.

4. NRENs' strategy in application services

This section aims at summarising the main driving forces and strategies concerning collaborative development or joint procurement of application services for scientific community. It focuses on storage related application in the first place.

4.1. Collaborative development of application services

The e-science community, including research and higher education institutes, has specific requirements on applications and services. These requirements are very well known by the NRENs have been serving the e-science community for decades. New service development and innovation is the core business of NRENs and that has been done closely together with university partners.

The decision on any new service development must be preceded by extensive research on the community requirements and comprehensive analysis of the available (commercial) services in that field. This needs the involvement of various target user communities not only in the requirement definition phase but also in the service development and testing phases. If NRENs join forces in development, an optimal mixture of specific expertise, human resource, and financial contribution can be achieved. This kind of collaboration can be facilitated by TERENA at pan-European scale.

FileSender case for joint application development

FileSender is a web based application that allows authenticated users to securely and easily send arbitrarily large files to other users. Authentication of users is provided through SAML2, LDAP and RADIUS. Users without an account can be sent an upload voucher by an authenticated user. FileSender is developed to the requirements of the higher education and research community.

The purpose of the software is to send a large file to someone, have that file available for download for a certain number of downloads and/or a certain amount of time, and after that automatically delete the file. The software is not intended as a permanent file publishing platform.

FileSender is released under the BSD license. It is open source software and available for free. The FileSender development project was initially funded by AARnet, HEAnet and UNINETT, other NRENs such as SURFnet, BELNET joined later to the project. The coding work is outsourced to RicoShae Development Ltd.

For more details visit: http://www.filesender.org

In principle, if:

- the user requirements are clear,
- an appropriate service is not available in the commercial market that meets with the user requirements or can be tailored to the requirements cost-effectively, and

- any of the NRENs has no such a service in place that can easily be adopted or made available for the specific user community, then

a collaborative application/service development project makes sense. Perfect example of this collaborative application development activity is the FileSender project [13] that is the spinoff of the TF-Storage task force discussion of TERENA.

4.2. Joint procurement of commodity application services

As it has already been mentioned, cloud services are emerging and commercial cloud service providers are approaching the research and higher education community intensively and vice versa. Especially the commercial Software as a Service (SaaS) market is huge (Fig. 4) and rapidly growing. That is why NRENs are started to work out strategies related to both outsourcing their commodity services to public clouds and keeping their core services and ability to innovation in-house. However, not just the NRENs' but their users' position must be investigated and understand too since that can affect the service offerings of NRENs.

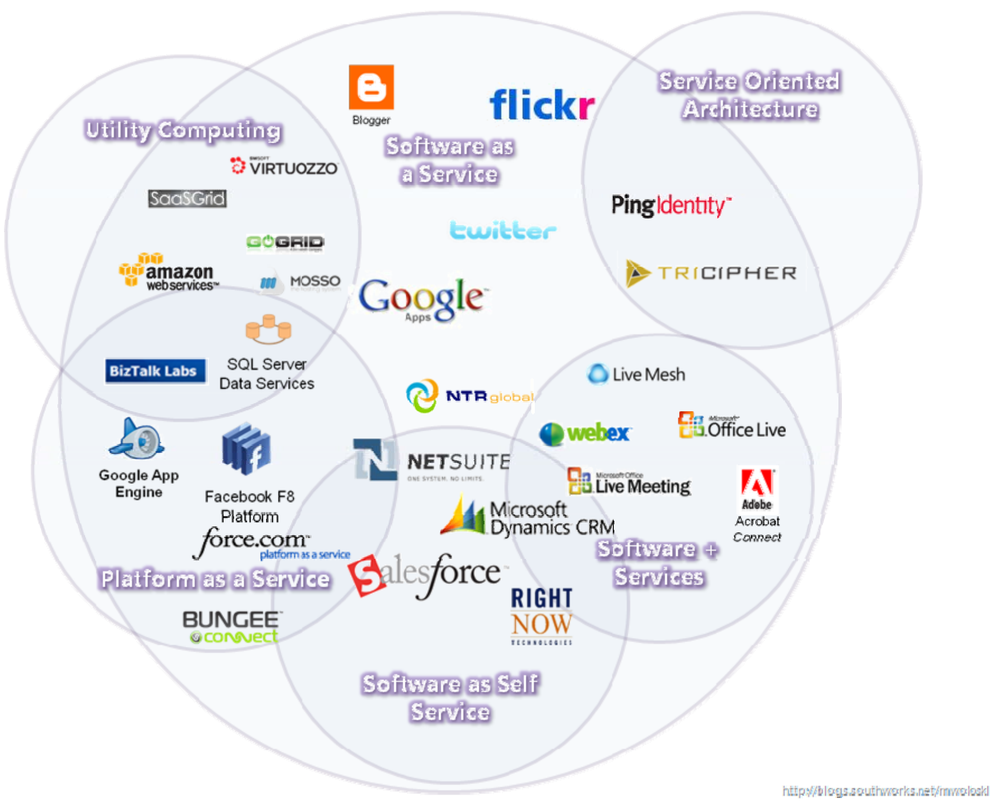

Fig.4 Taxonomy of commercial cloud services

Universities and higher education institutes have started to investigate the public cloud service offerings of commercial companies because it is highly demanded by the users (i.e. students, university staff, and researchers) itself. In terms of scalability and ease of use commodity services such as Google Mail or

Microsoft Live@edu applications can compete with services provided by University IT departments or even the NRENs. Therefore universities, in order to better serve both the students and staff, have started to outsource services to public clouds. But they have soon realised that outsourcing of services requires completely different skill set (service integration, contract management, legal sensitivity, knowledge on data protection rules, etc.) than engineering of own services. Moreover, universities cannot blindly trust in commercial providers as they always have hidden business interest. Individual universities have often no significant buying power to puss down the commercial prices as much as it would have been possible.

Analysing these aspects, NRENs have realised that there is a definite role in the middle of the service delivery chain for NRENs. This can bring both technical and non-technical benefits for the end users:

- NRENs can take the role of providing federated support for commodity services like Google Apps or Skype. Most importantly NRENs can bring trust into the service delivery procedure (universities cannot trust in commercials but have high level of trust in their NREN). NRENs can broker universities to cloud providers and provide technical expertise in handling data (e.g., sending data in and out) via research and education networks. Technically, there is the risk in commercial cloud providers to lock in proprietary APIs. NRENs' role can be to try and minimise that risk for the research and education community.

- NRENs can exploit the joint buying force by collecting an aggregating user demands (i.e. buying syndicate) towards commercial cloud service providers. NRENs can hide the pay-as-you-grow component that is simply does not work for universities who have to know the IT budget upfront. Emerging new role of NRENs is to play intermediary role between commercial cloud services and higher education needs.

SURFnet case with Microsoft and Google

SURFnet concluded agreements with Microsoft and Google in 2010 for the delivery of Live@edu and Google Apps to institutions connected to SURFnet, through its coordinating federation platform SURFfederatie.

Although SURFnet has previously been engaged primarily in the building and development of networks, it has noted the interest in cloud services communicated by the user institutions and has been focusing more and more on the available options for institutional networks; a demand that can be answered in part by cloud services.

For more details visit: http://www.terena.org/mail-archives/tf-msp/pdfsIUEA3OSlJ.pdf

In principle, if:

- the service provided by the University IT department or NREN is not efficient and scalable enough, and
- the service provided by commercials are not cheap enough or require specific contract management skills to negotiate for and deploy them

a joint application/service acquisition process coordinated by the NREN to better exploit collective buying power makes sense. A non-exclusive example of joint procurement procedure is the SURFnet case [14] in The Netherlands.

The commercial offers of public cloud services are available worldwide and big commercial companies (such as Google, Microsoft, Amazon, etc.) often have market specific sales departments in each continent. In principle, if universities and higher education institutes in Europe can join forces and approach cloud service providers with aggregated demands better overall deals can be achieved. Since the national demands can be aggregated by the NREN, as it was discussed above, it is an obvious step forward that the NREN demands can be aggregated by a pan-European organisation (e.g., TERENA) that can deal with the brokering at higher level (Fig. 5). TERENA can play the role of a European level mediator that can coordinate the joint procurements.

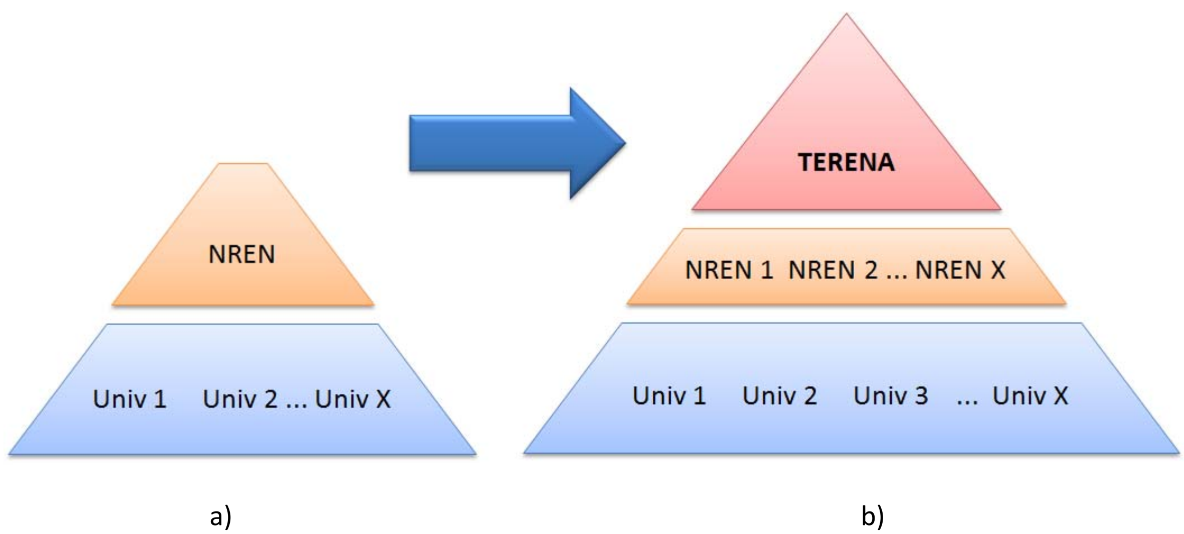

Fig.5 Brokering to cloud providers a) at national b) at European level

So, the basic question (regarding application services) is as follows:

- **Is there any chance to establish a brokering function towards commercial application service providers (in the cloud) jointly by NRENs for the benefit of the European R&E community?**

One of a kind example for the European level joint procurement is the TERENA TCS service [15] (that is not a data storage related service). TCS is a commercial certificate service provided by Comodo. TERENA is contracted with Comodo on behalf of the European NRENs and as a result NRENs may issue unlimited numbers of certificates to their users (universities, research institutes) for a flat fee. The price reduction achieved has been substantial and the technical conditions of the contract have been outstanding. It was possible because TERENA systematically analysed and aggregated all the user needs delivered by NRENs and collected experts who had deep technical knowledge, legal sensitiveness and knew what they wanted to achieve.

TERENA TCS service case for joint service procurement*

The TERENA Certificate Service (TCS) allows a variety of digital certificates to be offered to research and education institutions served by participating National Research and Education Networks (NRENs). These are provided by Comodo CA Limited, one of the largest worldwide Certification Authorities.

Comodo also offers special prices for EV (Extended Validation certificates to TCS customers. EV certificates are issued after a more thorough vetting process than for other certificates and can be used to secure online financial transactions.

TCS takes advantage of a bulk purchasing arrangement whereby NRENs may issue unlimited numbers of certificates for a flat fee. Certificates may be obtained through NRENs that are TCS Participants.

For more details visit: https://www.terena.org/activities/tcs/

** Note that TCS is not a storage related service. Only used as an example to illustrate a joint pan-European service procurement.*

The next section is about the strategy in infrastructure services. While in case of applications and application services the outsourcing decision and the potential future role of NRENs (as we discussed above) are becoming clear, in case of infrastructures and (virtual) infrastructure services the various NRENs' strategies differ (depending on the weight of the NREN) and the future direction is not so trivial. Infrastructure is definitely the core parts of an NREN that is why the basic strategies must be carefully selected.

5. NRENs' strategy in e-Infrastructure services

Providing network infrastructure for the Research and Education (R&E) community is the core business of NRENs. The continuous network development and technology deployment efforts of NRENs make them capable to provide wide range of network services (from best-effort IP to optical circuits) and to cope with the actual national demands. The GÉANT network provides the necessary services and sufficient bandwidth for NRENs at pan-European level.

The network itself is the key enabler for all kind of infrastructure services to be provided for the R&E community. For instance, high power data centres must be accessed via high-bandwidth networks in order to move huge amount of data in and out with a good throughput. Nevertheless, network access is a major cost component taking into account commercial cloud offerings.

Beyond the network, the primary cost components of building cloud storage include:

- Data centre occupancy - leased (co-location) or owned and depreciated.
- Data centre environmental - utilities, cooling, heating, etc.
- Storage hardware (leased expense or capital requirements & associated depreciation).
- File system and storage management (may be bundled in the storage hardware).
- Cloud enablement or platform (discreet or bundled with the storage system).
- Systems management and operational overhead.
- Backup and disaster recovery.

In fact, it is desired that NRENs use their own network infrastructure to access commercial cloud storage facilities (in order to eliminate that significant cost component) or implement their own data centres on top of their major network nodes (preferably hosted by local universities or research institutes in order to reduce the occupancy and environmental cost components).

Based on the user demands, networking capabilities, and funding schemes, three major national deployment strategies in e-Infrastructure service deployments are currently considered by NRENs:

- Building private (national) storage/cloud infrastructure on top of the national R&E network.
- Connecting public (commercial) storage/cloud infrastructures via the national R&E network.
- Creating hybrid storage/cloud infrastructure (i.e. mixing private and public storage/clouds infrastructures connected via the national R&E network).

The fourth potential strategy is not yet taken up by NRENs but certainly there is an emerging opportunity for creating an international community storage/cloud infrastructure either:

a) by interconnecting the national (private) infrastructures at pan-European level or

b) by aggregating national demands and go for a public (commercial) storage/cloud at pan-European scale.

5.1. Building a national (private) cloud storage infrastructure

A brief overview on the NRENs' national deployments related clouds can be found on the TERENA TF-MSP Wiki space: https://confluence.terena.org/display/msp/NREN+Cloud+Activities

Detailed information about the latest data storage related national infrastructure developments and deployments can be found on the TERENA TF-Storage website [16].

One of the NREN strategies is to focus purely on data storage services at the first place (to fulfil the actual demand of users) and do not thing about clouds (yet). CESNET's latest storage deployment is a good example for that.

CESNET case for building storage infrastructure only

The Czech national e-Infrastructure projects are to build up and to put into service a storage system of three distributed large scale repositories for saving and sharing of large volume of data including archiving. The main purpose of CESNET's storage system is to provide an easily accessible and redundant data repository for academic and scientific community. From the user (students) point of view it should provide (virtually) "unlimited storage capacities".

Technical concept is to have a HSM (Hierarchical Storage Management) system composed of disk arrays and tape libraries. The data access method and communication protocol is a combined NAS/SAN solution. Disk array (Tier1, Tier2) is about 300 TB, tape library (Tier3) is planned to be about 3 PB capacity per site (i.e. ~10PB storage capacity in total).

For more details visit: http://www.terena.org/activities/tf-storage/ws9/slides/10092010-7tf-storage-vercimak.pdf

Other strategy is to consider raw storage infrastructure as an enabler for future cloud services. NIIF/Hungarnet example clearly shows that how raw data storage infrastructure can be used to develop sophisticated cloud services on top. The Hungarian cost efficient storage development had entered into a new stage by pushing national scale storage improvement into the forefront of NIIF's infrastructure developments plans.

NIIF/HUNGARNET case for providing raw data storage as enabler for clouds

The national storage network (part of the large-scale project initiatives called HBONE+) has three geographically separate sites with 0.5 PB raw data capacity in total. The plan is to upgrade to 1.2 PB soon. The device offers raw disks, and disk related services, such as RAID structures, volumes within RAID groups via iSCSI protocol. The devices are currently managed manually, based on established procedures, yet there is a new storage management utility being developed that will allow the users to configure the storage devices through a straightforward graphical user interface. The storage management utility runs on two dedicated servers to offer redundancy.

An on-going development is to build an IaaS cloud on top of that raw data storage facility. The service is going to be public for NIIF users and member institutions. KVM (Kernel-based Virtual Machine) will be used as virtualization technology and OpenNebula will be used as cloud management system. The initial configuration will cover ~60 nodes (4 CPU cores and 12 GB RAM each) distributed to ~8 sites in Hungary. Its features will include self-service management, live migration of virtual machines between sites, and QoS provisioning (guaranteed minimum CPU performance for virtual machines).

For more details visit: http://www.terena.org/activities/tf-storage/ws9/slides/10092010-7tf-

The overall strategy of PSNC, the Polish NREN, is to provide storage services at the national level as it allows the robust, scalable, efficient deployment and operation. Currently, some local and regional solutions and approaches exist, however the aim is to run a common data storage and backup/archive services covering the country. PSNC's general approach is to deliver a service rather than "pure" storage resources. PSNC is not going to provide the direct, block-level access to the storage resources such as the access to disk volumes through iSCSI protocol in a predictable future (opposite e.g., to NIIF strategy).

PSNC case for providing storage as a service

In Poland there is no on-line storage service which could be treated as and replace the users' primary storage in the Polish National Research and Educational Network at the moment. However, PSNC and the PIONIER network consortium members are deploying remote, networked storage service for secondary storage and backup/archive purposes. It is called Popular Archive Service, run as one of the services of the PLATON project (PLATON – Service Platform for e-Science), funded under EU structural programme. The service will go into production in 1Q2011.

The backup/archive service is accessible through a set of standard protocols including SCP, SFTP, WebDAV/HTTPs and GridFTP. Using these protocols, the users can store and retrieve files from a remote, virtual filesystem. The geographical replication of the data and meta-data handling is made by the system transparently to the users.

The National Data Storage (NDS) system and services are deployed in the infrastructure composed of 12.5 PB of tape storage and 2 PB of disks as well as 70 servers and 50 Ethernet and SAN network switches located in 5 academic HPC computing centres and 5 MANs and universities across the country. The infrastructure purchase and maintenance as well as the deployment of the NDS software stack in this infrastructure are funded under PLATON project.

For more details visit: http://www.man.poznan.pl/online/en/projects/50/PLATON.html

The aforementioned basic deployment strategies aim at data storage infrastructures. What makes the difference to call them as a cloud is the list of features below [17]:

- On-demand self-service. A consumer can unilaterally provision system capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.
- Resource pooling. The provider's resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.
- Rapid elasticity. Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out, and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

- Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Following these features the most visionary strategy is to develop and deploy a cloud infrastructure (including storage, computing, network, etc.) in a fully integrated way. GRNET's cloud service offering integrates all the necessary features in one solid infrastructure.

GRNET case for building a cloud

GRNET's strategy aims at enabling the future service-oriented Internet for the Greek R&E community. It can be achieved by a centralised services provisioning system that can provision virtual working environment services.

The cloud infrastructure integrates the existing virtualised/cloud services of GRNET (i.e. GRNET Virtual Machine (ViMa) Service, GRNET AAI Federation, HellasGrid infrastructure, Pithos online storage service) taking the fact that the actual cluster management, a firm hypervisor, a Web frontend for OOB access / power control and more importantly an open architecture concept are given.

The necessary deployments include: KVM, Google's Ganeti project results (with GRNET NOC contributions), and a custom web frontend in Django (under development).

For more details visit: http://www.terena.org/about/ga/ga34/20101021GRNETgaClouds.pdf

HEAnet is actively thinking of providing cloud service including the network side (participating in the FP7 project Mantychore that is about IP Networks as a Service/Infrastructure as a Service scenario). HEAnet is also looking at such cloud service related to storage and virtual machines (which includes help to their clients to determine the technical specifications for such clouds; like network connectivity and resilience). An integration of these components (i.e. network and higher layer services) is seen important.

HEAnet case for dynamic networking as a service for e-infrastructure users

EC funded FP7 project Mantychore aims to provide dynamic networking as a service for e-infrastructure users. The e-science community has expressed a need for flexible and application-driven networks. Traditional methods for configuring networks for research users are impractical when the scale is international and involves several networks. Mantychore takes an "infrastructure as a service" (IaaS) approach to networking to enable National Research and Education Networks (NRENs) and other e-Infrastructure providers to deploy a configurable, automatable network that allows virtual research communities to control routers, switches, optical devices, and IP networks as necessary to meet the needs of their applications. Grid and cloud computing will provide some of the first real uses for this technology.

For more details visit: http://www.mantychore.eu/

However, as there might be some cloud providers in Ireland, HEAnet is also interested in how to make a framework to connect such providers to NRENs' networks (peering with NRENs, general Internet and/or GEANT IP/Plus).

The next section discusses some basic strategies towards public (commercial) cloud inclusion and infrastructure service outsourcing in general.

5.2. Connecting to public (commercial) cloud infrastructure

Given the NRENs' different funding models, income streams and ability for capital investment it is unlikely that an Amazon-like cloud offer (i.e. 1 hour of compute time on 1000 machines costs you as much as 1000 hours on 1 machine) can soon be achieved by building national cloud storage infrastructures of individual NRENs. In contrary, taking the fact that wide range of commercial cloud offerings are currently available on the market, a reasonable national strategy would be to try and negotiate with commercial cloud providers and connect NREN users (via NREN network) to public clouds.

Aggregating infrastructure service demands at national level can be managed by NRENs (just like in case of application services) since NRENs have the network to connect their users to commercial clouds. JANET(UK) 's example shows that how commercials can natively extending NRENs into their data centre and cloud environment.

JANET(UK) case for connecting a cloud

International solutions provider Logicalis has announced that its high-density data centre (HDDC) and range of Made to Measure cloud services are now natively connected into JANET. Logicalis' connection will provide a range of data centre and cloud services, including Infrastructure as a Service, application hosting for Independent Software Vendors (ISVs) servicing the education community, through to high density hosting of virtualisation and supercomputing environments.

Many service providers peer with the JANET network, but Logicalis has gone much further by natively extending JANET into their data centre and cloud environment. This gives JANET's users and our existing education customers the ability to extend their campus into our environment and benefit from highly virtualised cloud services and a high-density infrastructure, without the constraints of the Internet or the capital costs incurred through building a shared private network.

This connection supports the activities outlined in JANET's corporate plan, specifically 'supporting software as a service, cloud computing, and cost effective data storage'. JANET expectation is that their customers can benefit from an improved cloud/data centre proposition from the commercial sector, in turn enabling them to meet their needs more effectively.

For more details visit: http://www.ja.net/company/news-2011/logicalis.html

Connecting public cloud infrastructures to NRENs means that even data intensive infrastructure services (e.g., storage/backup) can be outsourced to commercials. UNINETT's vision is, for instance, to end up with outsourcing higher education IT services to the cloud at national level.

UNINETT case for outsourcing IT

UNINETT's strategic vision is to outsource IT operations of national higher education institutes. An evolutionary path would be to start with data backup service offsite (storage), then move live data and server images offsite (disaster recovery site) and finally move all IT operations offsite (i.e. cloudify Norwegian higher education sector into a nation-wide virtual data centre).

The business case behind that could be to provide a cheaper, better (in terms of reliability, scalability, functions, etc.) and future proof solution. E.g., instead of having each college to establish their own offsite backup facility with hardware to run virtual machines on, you'd have one or more centralized platforms to do this.

From the network connections' perspective, Norway can use any network services via NORDUnet and GEANT. In addition, with potential direct peering contracts, bandwidth cost could be cut when connecting to commercial cloud providers.

For more details visit: http://www.terena.org/activities/tf-storage/ws5/slides/d2-1-backup.pdf

5.3. Obtaining a pan-European community cloud infrastructure

In the previous sections examples of various national efforts either to build private infrastructure clouds or buy in (i.e. connect to) commercial infrastructure clouds can be found. There is a way, of course, in the middle to mix public and private clouds together in order to achieve the necessary flexibility and scalability of the infrastructure on the technical side as well as to obtain a good economy of scale on the business side.

Concerning the economy of scale (again, just like in case of joint procurement of application services) NRENs may team up to aggregate demands at regional or even pan-European scale and

a) do large scale purchasing of resources (i.e. build a massive data centre) from a solid technical understanding of what is really needed by the community, or

b) establish a common brokering function (i.e. connect to a "big" public cloud) towards commercial cloud service providers.

It is clearly seen that the first steps have been taken by some of the NRENs that are building or buying in national storage/cloud infrastructures or contracting with commercials for application services provided in the cloud. The obvious next step is to communicate and discuss these national strategies (facilitated by this green paper) and see if there is a chance for expansion to pan-European scale.

So, the basic questions (regarding infrastructure services) are as follows:

- **Is there any chance to build a large scale cloud infrastructure jointly by NRENs for the benefit of the European R&E community?**
- **Is there any chance to establish a demand aggregation function towards commercial cloud infrastructure providers jointly by NRENs for the benefit of the European R&E community?**

The following non-exclusive examples are clearly show that some of the NRENs are willing to team up and join forces behind this potential new strategy.

PSNC's call for joining forces

PSNC has started a NDS2 project, funded by Polish Ministry of Science and Higher Education, which develops the extended architecture of NDS and the new system features.

PSNC plans to join national and international initiatives that deal with content management issues (such as digital libraries projects), in order to (potentially) participate in their development and support their operation.

PSNC plans to participate in building Europe-wide, pervasive data storage infrastructure. For that purpose PSNC approached the recent "Data infrastructures for e-Science" published by the European Commission as the member of the EUDat consortium. However, PSNC is constantly seeking for other collaboration opportunities focused on building a European distributed data storage system, which PSNC thinks is of utmost importance for the scientific community.

"TERENA's TF-Storage seems to be a suitable place for discussing and brainstorming about such an initiative."

NIIF's call for a potential project consortium

It is still in a very early brainstorming phase, but NIIF's intention would be to go for a European project that focuses around the following topic items:

- The key aim would be to build up a scientific cloud, i.e. an NREN centered cloud service in parallel to the commercial ones (mostly on IaaS and no upper layers).
- The distinguishing features would be as follows: added value, improved environment for scientific job sandboxing (e.g., HPC-specific software components or even OS images, grids clusters)
- Knowledge transfer among NRENs in this field is essential.
- Improved standardization, interoperation with other public and private clouds are considered to be important.

"TERENA's TF-Storage seems to be a suitable place for discussing and brainstorming about such an initiative."

6. Conclusions

This paper aims at providing an overview on the NREN's strategic perspective related to storage technologies and services (with special focus on cloud storage) as well as to cloud services and service outsourcing issues in a broad spectrum.

The service offering are divided into two groups; application services and infrastructure services. In case of application services the strategy options for NRENs and their customers (e.g., universities) are also taken into consideration. A particular NREN's strategy must be in line with their users' business strategy in this field, of course, that requires continues communication, national consultations within the R&E community.

A simplified strategy decision tree for NRENs is depicted in Fig.6 (including some well-known examples). Regarding application services, NRENs can develop and provide their own services or can exploit the joint buying force of their users and brokering towards commercial cloud service providers. Regarding infrastructure services, NRENs can build their own cloud infrastructure or can aggregate user demands and channel them in to commercial cloud infrastructures.

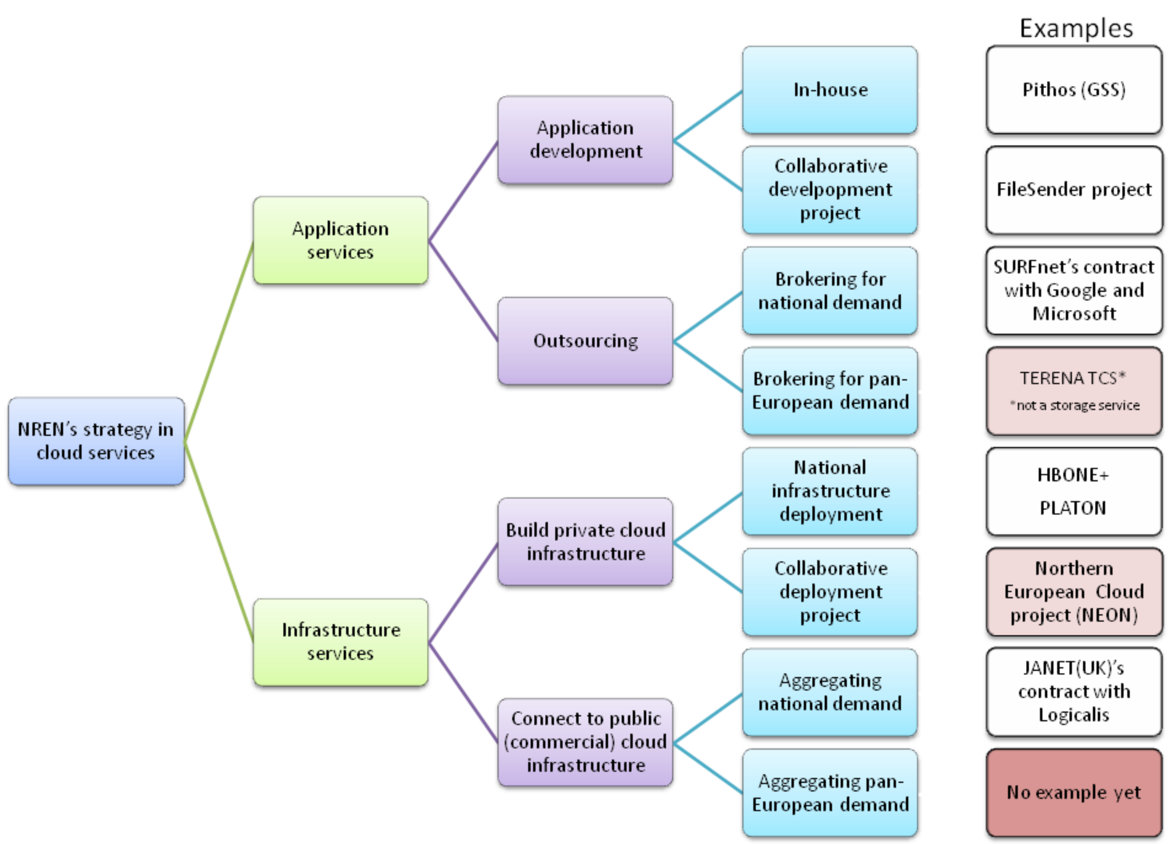

Fig.6 NREN's basic strategy decision tree with known examples

In summary, the open questions of this green paper address those boxes in Fig.6 (marked in red) where there is no known example of a particular NRENs' strategy, at least in the specific field of data storage and cloud storage services at pan-European scale.

The basic question (regarding application services in the cloud) is as follows:

- **Is there any chance to establish a brokering function towards commercial application service providers (in the cloud) jointly by NRENs for the benefit of the European R&E community?**

The basic questions (regarding cloud infrastructure services) are as follows:

- **Is there any chance to build a large scale cloud infrastructure jointly by NRENs for the benefit of the European R&E community?**
- **Is there any chance to establish a demand aggregation function towards commercial cloud infrastructure providers jointly by NRENs for the benefit of the European R&E community?**

Epilogue

The Royal Society which is fellowship of some of the world's most eminent scientists and this report strongly endorses the importance of international collaborative research.
http://royalsociety.org/policy/reports/knowledge-networks-nations/

"With human society facing a number of wide-ranging and interlinked 'global challenges' such as climate change, food security, energy security and infectious disease, international scientific collaboration is essential if we are to have any chance of addressing the causes, or dealing with the impacts, of these problems… the primary driver of most collaboration is individual scientists. In seeking to work with the best of their peers and to gain access to complementary resources, equipment and knowledge, researchers fundamentally enhance the quality and improve the efficiency of their work." Many countries simply don't have the financial resources to carry out world class research in significant number of fields of research and so international collaboration becomes increasingly important. Unfortunately, in many of these countries, funding councils still largely focus on small scale domestic research which often results in underfunded research and second rate results. This is particularly critical as the science community moves ever more so to e-Research and needs access to vast international computing resources, datasets and models. Stand alone, small scale provincial computing resources, data facilities, clouds, grids or networks will not enable world class science [18].

References

[1] Nirvanix Developer Center, "**What Makes Cloud Storage Different from Traditional SAN and NAS?**", June 2009
http://developer.nirvanix.com/blogs/strategies/archive/2009/06/19/what-makes-cloud-storage-different-from-traditional-san-and-nas.aspx

[2] Arjan Peddemors, Christiaan Kuun, Rogier Spoor, Paul Dekkers, Christiaan den Besten, "**Survey of Technologies for Wide Area Distributed Storage**", June 2010
http://www.surfnet.nl/nl/Innovatieprogramma's/gigaport3/Documents/EDS-3R%20open-storage-scouting-v1.0.pdf

[3] Rogier Spoor, Arjan Peddemors, "**Cloud Storage and Peer-to-Peer Storage - End-user considerations and product overview**", June 2010
http://www.surfnet.nl/nl/Innovatieprogramma's/gigaport3/Documents/EDS-3R%20Cloud%20and%20p2p%20storage-v1.1.pdf

[4] SURFnet GigPort3 project, "**Cloud Storage and Peer-to-Peer (P2P) Storage**",
http://www.surfnet.nl/nl/Innovatieprogramma's/gigaport3/Documents/conceptmap-v1.0.pdf

[5] **TERENA Compendium of National Research and Education Networks in Europe**, 2010 edition, Section 5.3, pp. 59, http://www.terena.org/activities/compendium/2010/pdf/TER-compendium2010-complete-web.pdf

[6] TERENA Task Force on Management of Service Portfolios, TF-MSP,
http://www.terena.org/activities/tf-msp/

[7] Koen Schelkens, "**Work area NREN Service Portfolios**", TERENA TF-LCPM: action 20060309-03
https://portal.nordu.net/download/attachments/15699252/nren-service-cats.pdf?version=1&modificationDate=1263565062000

[8] Luis M. Vaquero, Luis Rodero-Merino, Juan Caceres, Maik Lindner, "**A Break in the Clouds: Towards a Cloud Definition**" ACM SIGCOMM Computer Communication Review, 55 Volume 39, Number 1, January 2009
http://ccr.sigcomm.org/drupal/files/p50-v39n1l-vaqueroA.pdf

[9] Dr Max Hammond, Dr Rob Hawtin, Dr Lee Gillam, Prof Charles Oppenheim, "**Cloud computing for research**", June 2010
http://www.jisc.ac.uk/media/documents/programmes/research_infrastructure/cc421d008-1.1_cloud_computing_for_research_final_report_annexes.pdf

[10] Next Generation Grids (NGG) Experts Group, "**Requirements and options for European grids research 2005-2010 and beyond**", Technical report, European Commission, 2004
http://www.semanticgrid.org/docs/ngg2_eg_final.pdf

[11] Sushil Bhardwaj, Leena Jain, Sandeep Jain, "**Cloud Computing: A Study of Infrastructure as a Service (IaaS)**", International Journal of Engineering and Information Technology Vol 2 , No. 1, Feb 2010 http://ijeit.org/index_files/vol2no1/CLOUD%20COMPUTING%20A%20STUDY%20OF.pdf

[12] TERENA Networking Conference 2010, Vilnius, Lithuania
http://tnc2010.terena.org/schedule/sessions/show.php?sess_id=21

[13] FileSender project http://www.filesender.org
[14] SURFnet, "**Connecting to cloud services - A guideline for NRENs**", November 2010 http://www.terena.org/mail-archives/tf-msp/pdfsIUEA3OSlJ.pdf
[15] TERENA TCS Service https://www.terena.org/activities/tcs/
[16] TERENA Task Force on Storage and Cloud, TF-Storage http://www.terena.org/activities/tf-storage/
[17] Peter Mell, Timothy Grance, **"The NIST Definition of Cloud Computing (Draft)"**, January 2011 http://csrc.nist.gov/publications/drafts/800-145/Draft-SP-800-145_cloud-definition.pdf
[18] Bill St. Arnaud, **Royal Society - Knowledge, Networks and Nations - open exchanges and future of R&E networks**, April 4, 2011 http://billstarnaud.blogspot.com/2011/04/royal-society-knowledge-networks-and.html