

# Cloud Storage as a new Storage Class: QoS Characterization and Cost Analysis (S3) \*



S. Andreozzi, A. Forti, L. Magnoni, R. Zappi

INFN-CNAF, Viale Berti Pichat 6/2, 40127 Bologna (Italy)



## Amazon S3 Object Store Model

The Amazon Simple Storage Service (S3) is a storage cloud service part of the Amazon Web Services (AWS) suite, a set of services enabling Amazon to expose resources part of its infrastructure for easy integration into external applications.

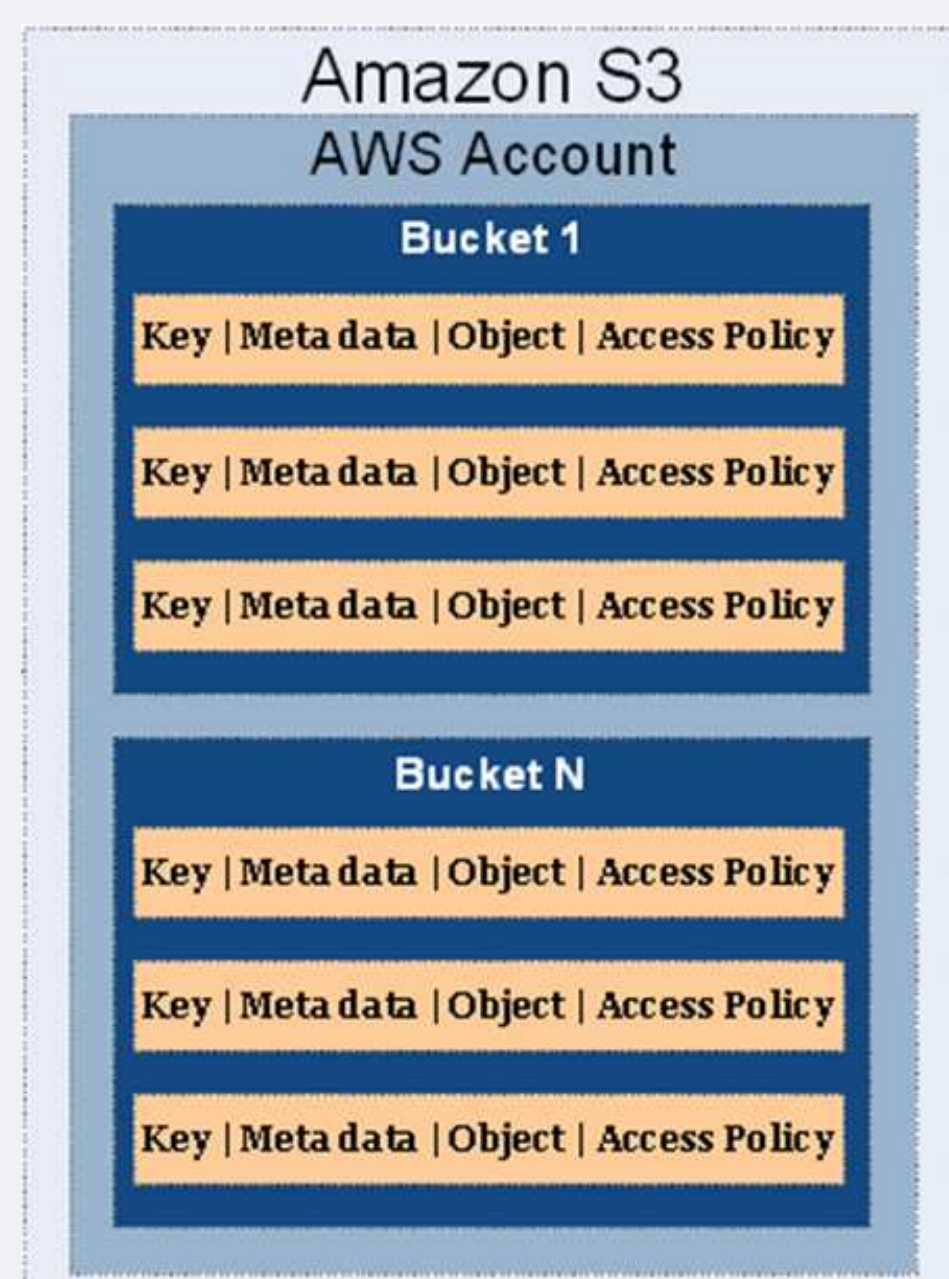


Figure 1: Amazon S3 Object Store model

- Amazon S3 provides a straightforward object store model (Fig. 1):
- Data is owned by Amazon AWS account
  - Bucket is the container for Objects
  - Object is the fundamental entity
  - Key is the unique identifier for an object in a bucket

- Amazon S3 security model
- Authentication is based on
    - Access Key ID (a 20-character, alphanumeric string)
    - Secret Access Key (a 40-character string)
  - S3 Requests validation is based on:
    - HMAC-SHA1 based signature attached to the request as parameter
    - Support query string authentication

Access control can be enforced relying on ACL on Buckets and Objects, specifying the Operations type: *READ*, *WRITE*, *ACL management* and the Subject of the ACL entry: *AWS users*, *AWS group* or *anonymous*.

## Storage Class on Amazon S3

WLCG collaboration adopts SRM [1] as the standard interface to manage data shared in a Grid environment. The concept of Storage Class has been added to the SRM paradigm to identify the quality of services provided by storage system. This concept defines basic entities for space management, such as Storage Area, space tokens and space alias, that have to be mapped in the S3 object model.

The main concepts are represented as follows:

- The owner of AWS account is the site
- The AWS account it is only used by StoRM service
- Buckets correspond to WLCG storage areas
- Objects are not hierarchical, but keys can be organized replicating the usual namespace adopted for le names
- Listing operation can be done on key prefix

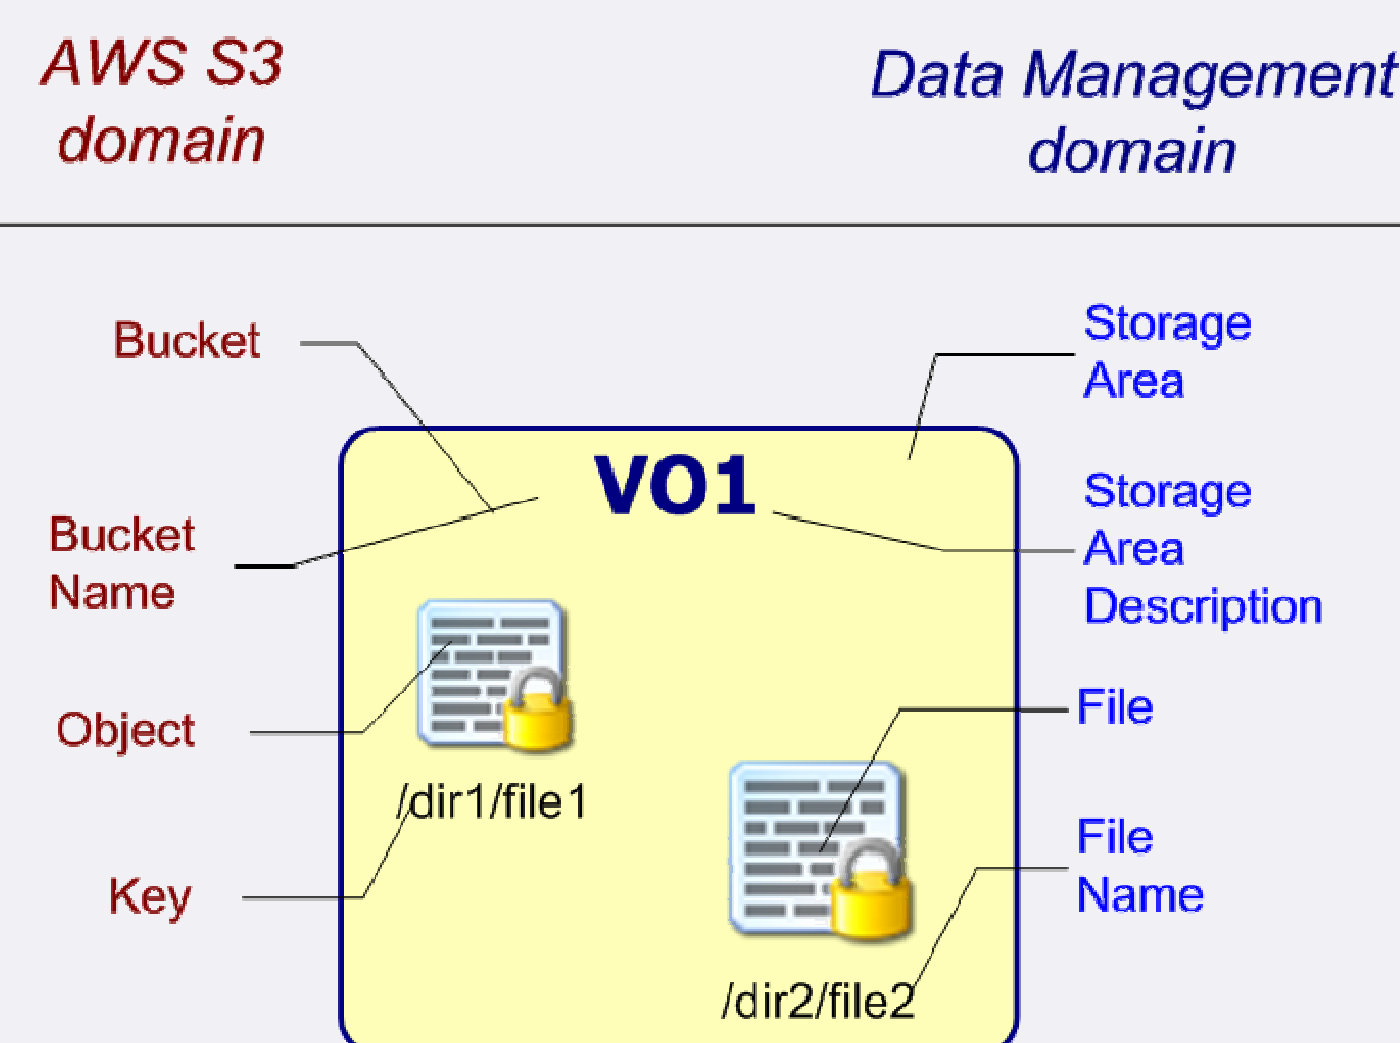


Figure 2: WLCG storage concepts in Amazon S3 domain

## Pricing in GLUE 2.0

Amazon S3 defines a clear and simple pay-per-use pricing model. A resource provider may want to expose such a pricing model in order to inform users that resource usage will be charged according to it. Hence, we have investigated how to expose pricing information associated to a Storage Service relying on the GLUE 2.0 Information Model.

GLUE 2.0 does not provide any price-specific attributes, therefore we suggest to use the extension capabilities defined in terms of the GLUE Extension class. Such a class enables to associate (key, value) pairs to existing classes. The GLUE Storage Share represents a storage extent which is made available to a certain UserDomain. Therefore, we propose to associate the pricing details to the Storage Share by using the GLUE Extension class.

In the S3 pricing model, three billing items are identified: storage space, in/out bandwidth and number of requests.

For these items, prices can be given in two possible ways:

1. per quantity over time in a given range
2. per quantity.



Key	Value
space:price:rangetime	0.18\$:1GB:month:0:50TB
space:price:rangetime	0.17\$:1GB:month:50:100TB
space:price:rangetime	0.16\$:1GB:month:100:500TB
space:price:rangetime	0.15\$:1GB:month:500:MAX
transfer-in:price:range	0.10\$:1GB
transfer-out:price:rangetime	0.17\$:1GB:month:0:10TB
transfer-out:price:rangetime	0.13\$:1GB:month:10:50TB
transfer-out:price:rangetime	0.11\$:1GB:month:50:150TB
transfer-out:price:rangetime	0.10\$:1GB:month:150TB:MAX
request:price:range	0.10\$:10000 Requests

Figure 3: Extension in GLUE 2.0 for pricing

## StoRM

StoRM [2] is a flexible and scalable SRM service able to work on different storage systems. The characteristics of StoRM are suitable for a smooth integration with Amazon S3 [4].

- S3 provides Application Programming Interfaces (API) that uses standards-based REST and SOAP interfaces for operation on Buckets, Objects and S3 service

- StoRM interacts with storage system through a driver interface that decouple internal logic from storage characteristics

- A dedicated driver uses S3 functionalities to reproduce classic file system capabilities




Figure 4: StoRM architecture

## StoRM and Amazon S3: Integration Models

Amazon S3 can be integrated in StoRM with two different approach. Depending on how it is used, the resulting storage resources provide different quality of service.

This first approach is to consider S3 as any other storage resources, as well as standard file system. StoRM sets up proper buckets and objects configuration following storage area requirements, and each SRM request is mapped to the S3 domain.

- Namespace and Authorization operations are managed by StoRM interacting with S3

- Access to data is performed by user through specific clients, built on S3 API

- Authentication is managed via query authentication string

- URLs point to object resources

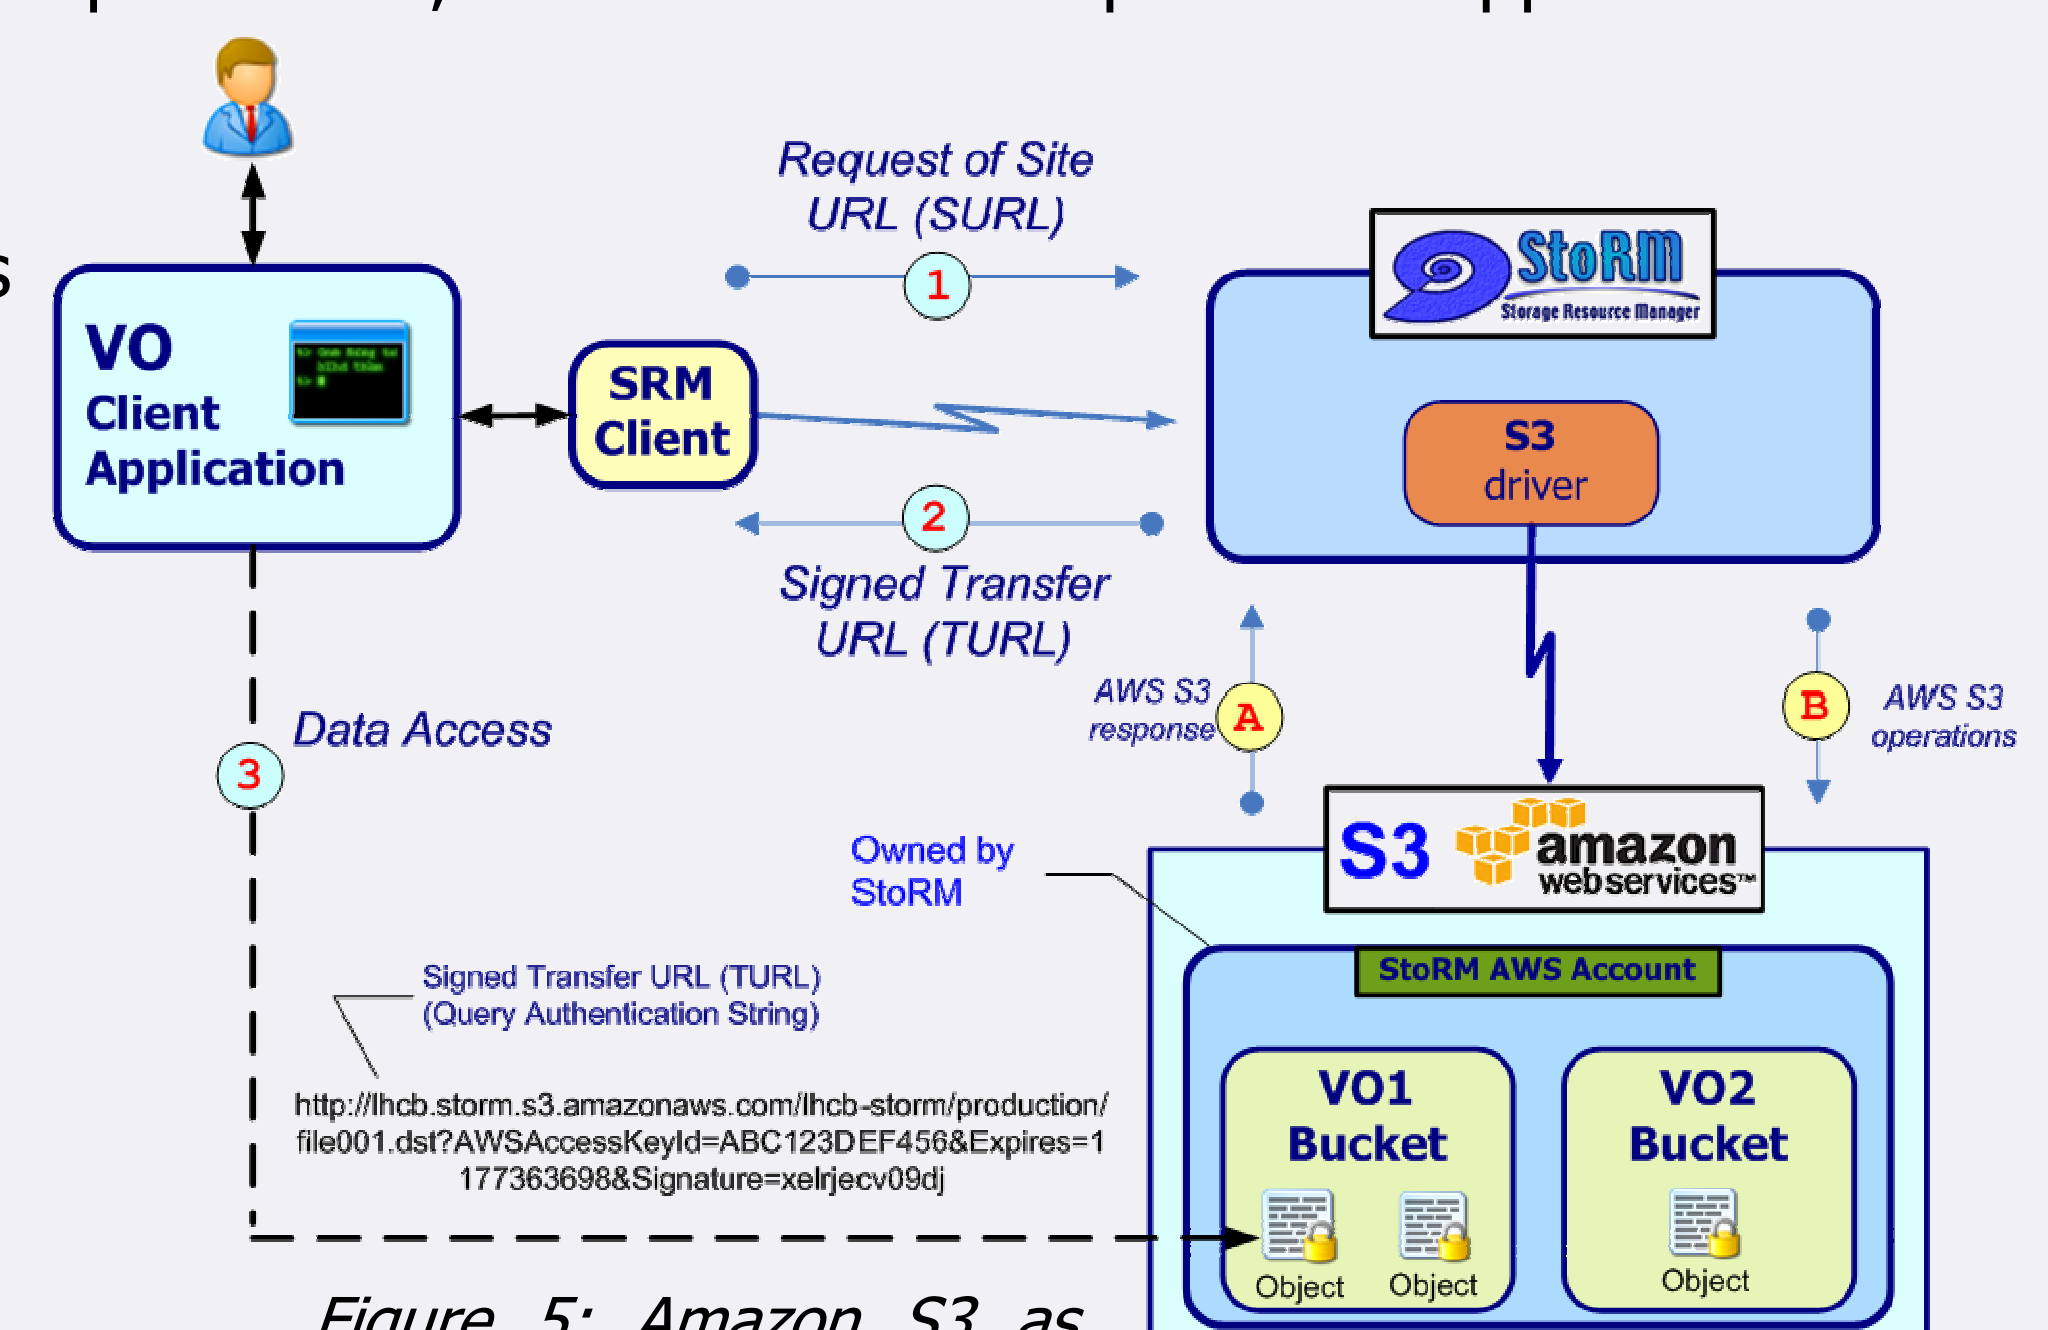


Figure 5: Amazon S3 as standard storage resource

This integration model provides sites with the capability to build up storage area for experiments relying completely on storage clouds. It can be useful in case of unplanned storage requirements or extreme storage failures.

The second approach is a hierarchical integration model. Leveraging on the SLA provided by Amazon [3], S3 is used to maintains a reliable copy of data, like a mass storage system. StoRM interacts with S3 through the API driving the stage and recall operations of files, following user's requirements.

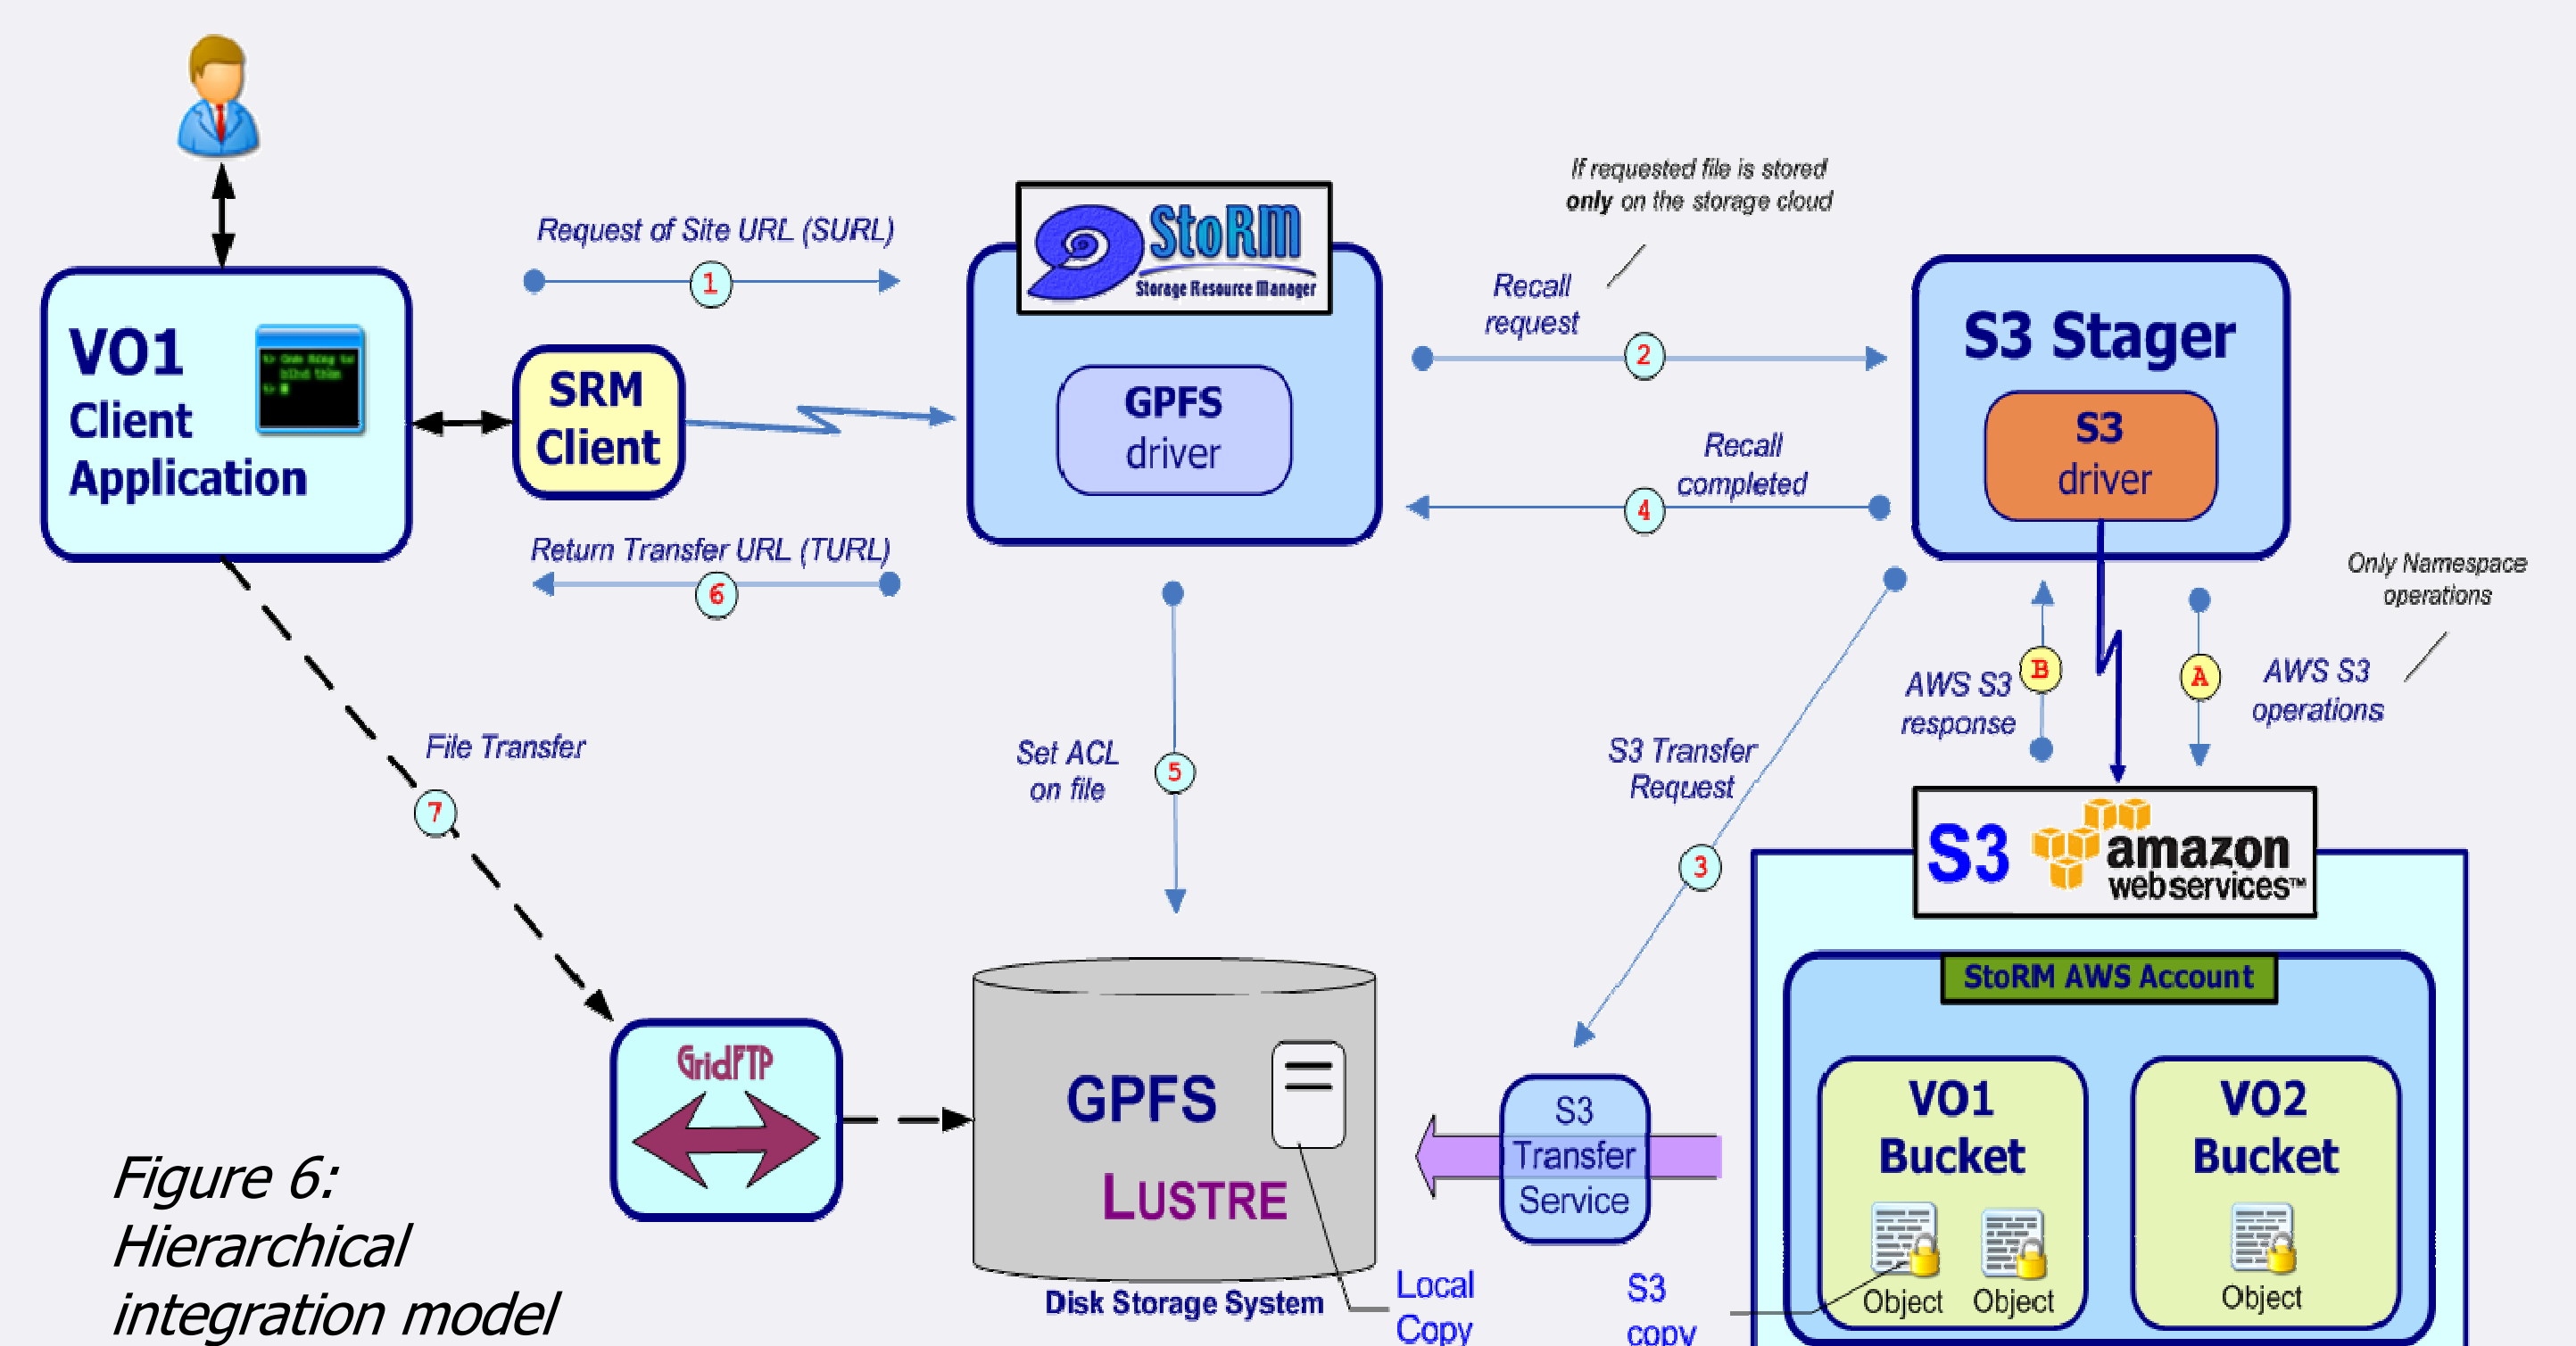


Figure 6: Hierarchical integration model

Data can be accessed directly via GPFS/Lustre/POSIX file system through standard Grid protocols and services as usual. To maximize transfer throughput the S3 stager can use multiple connections.

## References

- [1] The Storage Resource Management Working Group: <http://sdm.lbl.gov/srm-wg>, 2007
- [2] L. Magnoni, R. Zappi, A. Ghiselli, "StoRM: a Flexible Solution for Storage Resource Manager in Grid", In Proceedings of the IEEE 2008 Nuclear Science Symposium (NSS-MIC 2008), 19 - 25 October 2008, Dresden, Germany. IEEE Computer Society 2008.
- [3] Amazon Simple Storage Service (S3). <http://aws.amazon.com/s3>
- [4] S. Andreozzi, L. Magnoni and R. Zappi. Towards the Integration of StoRM on Amazon Simple Storage Service (S3). In Proceedings of the International Conference on Computing in High Energy and Nuclear Physics
- [5] A. Ghiselli, L. Magnoni, V. Vagnoni, R. Zappi, et al. "Performance studies of the StoRM Storage Resource Manager", In Proceedings of the 3rd IEEE International Conference on e-Science and Grid computing (eScience2007)