European Union
European Social Fund

OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION & RELIGIOUS AFFAIRS
MANAGING AUTHORITY

NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Co-financed by Greece and the European Union

Using IKAROS as a data transfer and management utility within the KM3NeT computing model

Christos FILIPPIDIS (a,b), Yiannis COTRONIS (b), Christos MARKOU (a)
a:NCSR Demokritos, b: University of Athens

Large-scale scientific computations

- Large-scale scientific computations tend to stretch the limits of computational power.

- I/O has become a bottleneck in application performance.

- The most important factors affecting performance are:

1. The number of parallel processes participating in the transfers.

2. The size of the individual transfers.

3. The I/O access patterns.

I/O access patterns

1.Compulsory, consist of I/Os that must be made to read a program's initial state from the disk and write the final state back to disk when the program has finished.

2.Checkpoint/restart, are used to save the state of a computation in case of a hardware or software error which would require the simulation to be restarted.

3.Regular snapshots of the computation's progress.

4.Out-of-core read/writes for problems which do not fit to memory.

5.Continuous output of data for visualization and other post-processing.

Globally shared file systems

- Globally shared file systems have several performance limitations when used with large-scale systems, because:

1.Bandwidth does not scale economically to large-scale systems.

2.I/O traffic on the high speed network can impact on and be influenced by other unrelated jobs.

3.I/O traffic on the storage server can impact on and be influenced by other unrelated jobs.

IKAROS Framework

- Enables us to create ad-hoc nearby storage formations in order to isolate the resources being used and increase performance.

- Can use a huge number of I/O nodes in order to increase the available bandwidth (I/O and network).

- Unifies remote and local access in the overall data flow, by permitting direct access to each I/O node.

- This approach enables us to connect, at the users level, the several different computing facilities used (Grids, Clouds, HPCs, Data Centers, Local computing Clusters and personal storage devices), on-demand, based on the needs.

IKAROS Design (1/2)

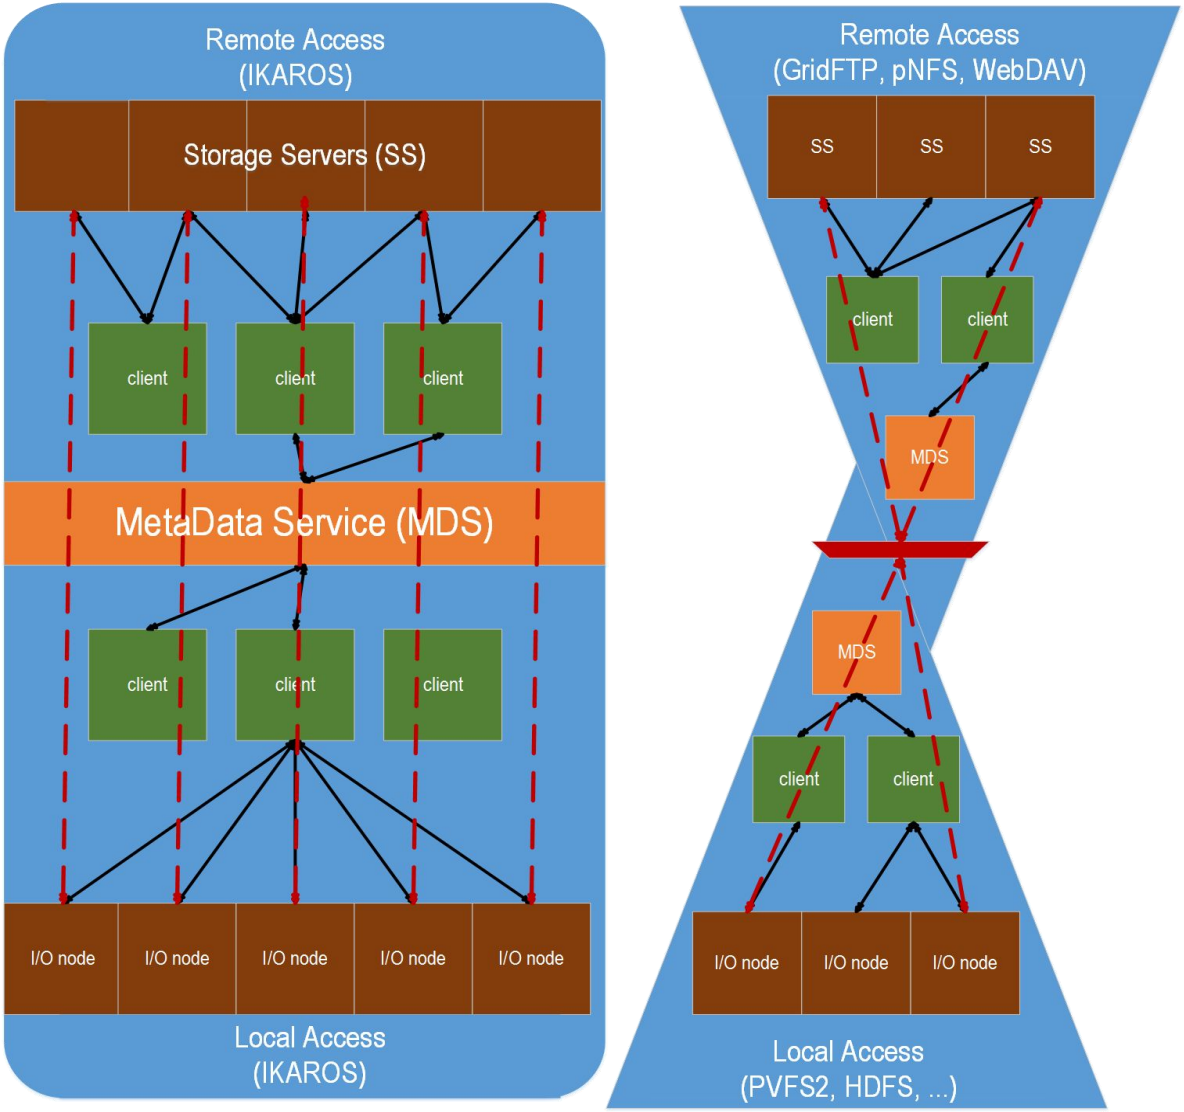

IKAROS Design (2/2)

- Allows data in a file to be striped across multiple disk volumes on multiple heterogeneous nodes.

- Provides the utility for the storage system to access and transfer a data file in parts and in parallel mode, without a specific order, according to a client request.

- Defines three types of nodes: The User Interface(UI)/Client node, the Meta-data node and the I/O node.

- Node types are peers with the ability to act in any mode.

- First version was developed as an Apache Dynamic Shared Object (DSO), the latest version is written in nodeJS.

HDFS, PVFS2, GPFS, IKAROS Features

	HDFS	PVFS2	GPFS	IKAROS
Deployment model	Co-locates compute and storage on the same node	Separate compute and storage nodes	Separate compute and storage nodes	The user/app can choose both models, on the fly
Data layout	Exposes mapping of chunks to data-nodes to Hadoop applications	Maintains stripe layout information as extended attributes but not exposed to applications	not exposed to applications	Exposes mapping of chunks to applications and users
Compatibility	Custom API and semantics for specific users	UNIX	UNIX	UNIX, WINDWOS, MAC
WAN capabilities	Can be exported through webdav	Can be exported through pNFS	Can be exported through pNFS	Build-in remote access capabilities. Supports parallel channels WAN data transfers, stripping servers, third party data transfers.

IKAROS Basic Usage Scenario

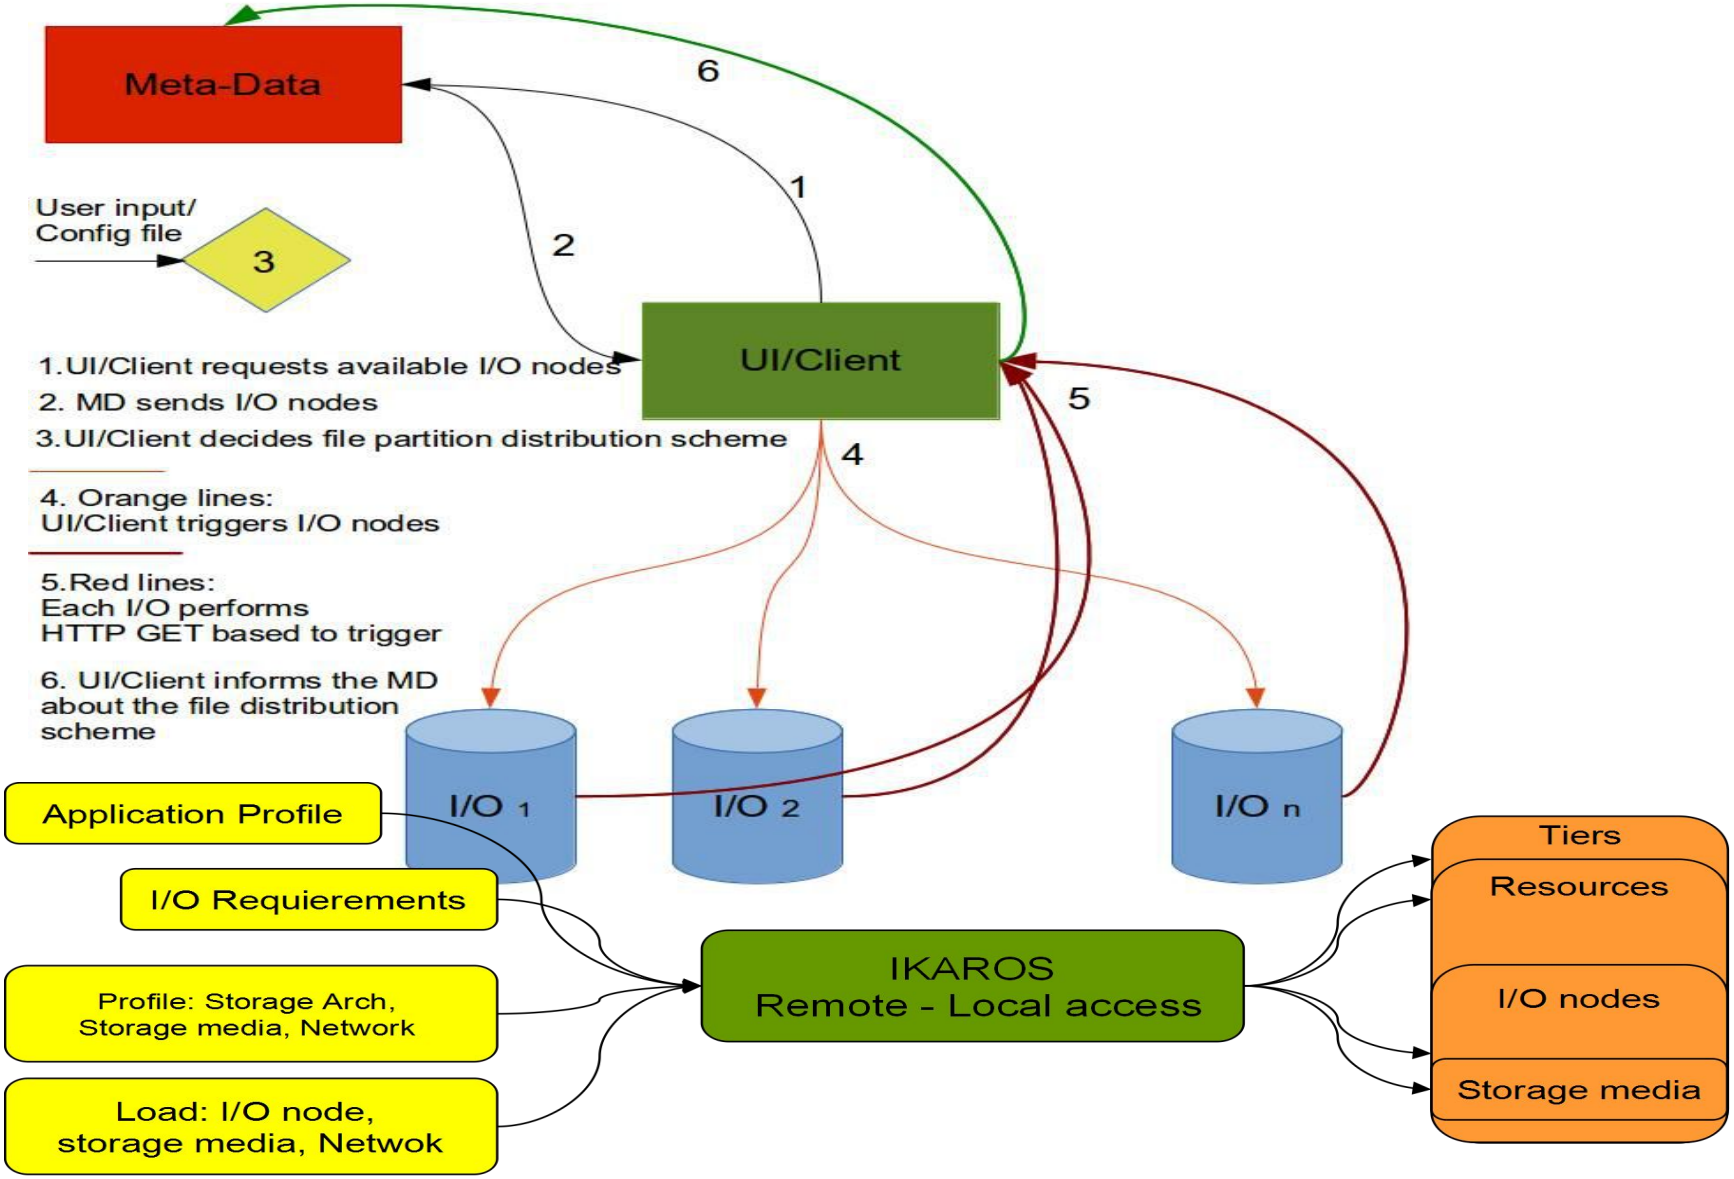

IKAROS hybrid model (1/2)

- Is using well known standards such as the HTTP and the JSON
- Creates synergies with Web 2.0 platforms.
- Tiered MDS

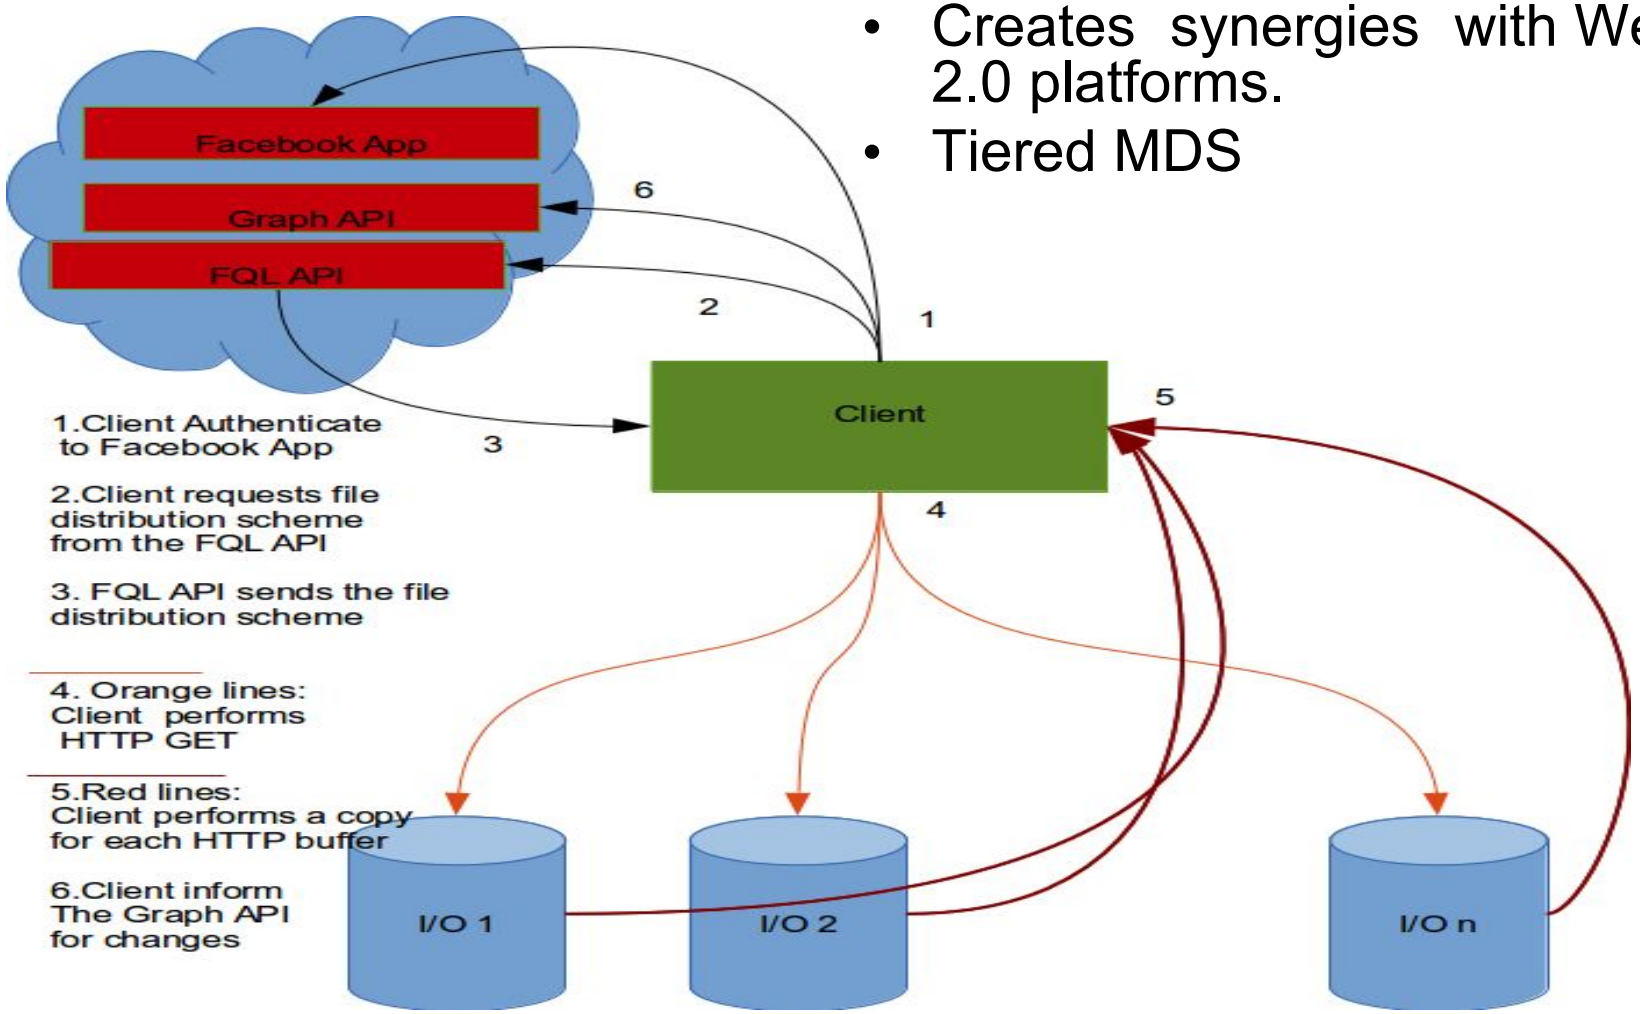

IKAROS hybrid model (2/2)

Experiments-Cytera Machine

	#	Specs
Compute nodes	96	12 Intel Xeon CPU cores, 48 GBs of RAM and 15K rpm local HDD
Storage Nodes	4	360 TBs raw disk space in 18 Raid 6 arrays each with 10 7200 rpms HDDs
GPFS-Meta data System	4, hosted at the storage nodes	Raid 10 arrays (one associated at each server)
Network Connectivity	-	QDR (40Gbit/s) infiniband

Comparing IKAROS with GPFS (Cytera-Machine)

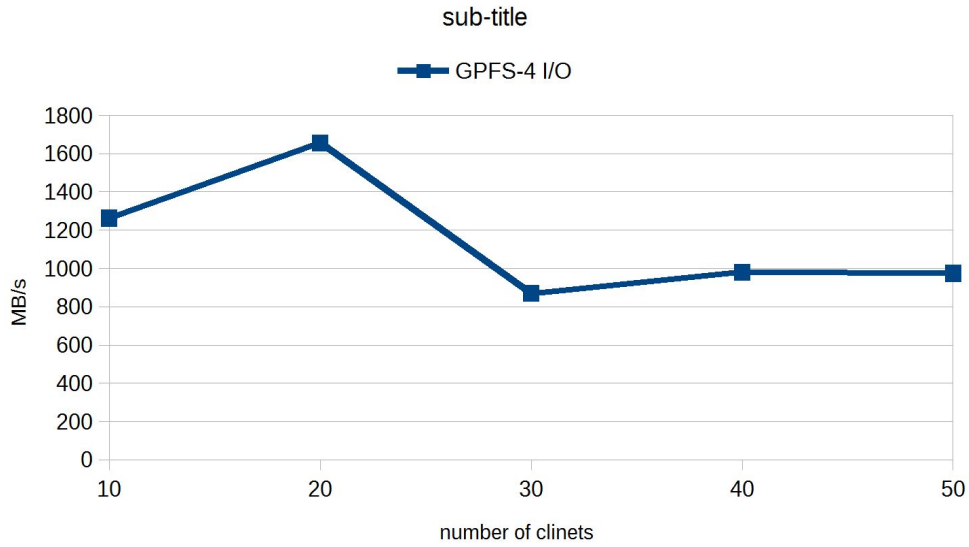

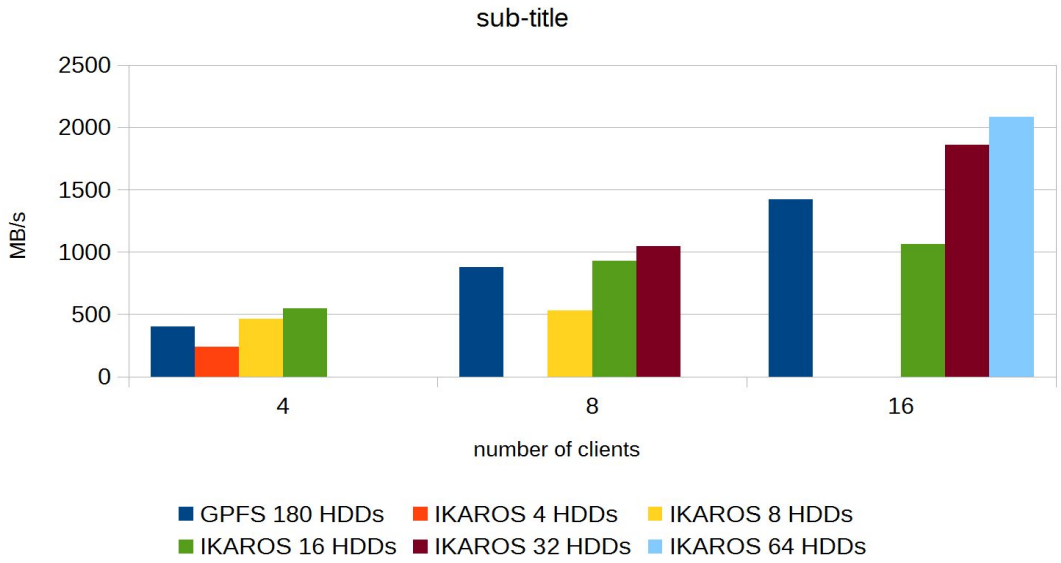

IKAROS-KM3NeT

Corsika Grid job test submission:

- 200 jobs submitted (5000 events). Output ~ 14GB send directly at Demokritos & Lyon (single step: WN-> Lyon).

- 50 jobs submitted (50000 events). Output ~74GB send directly at Demokritos & Lyon (single step: WN-> Lyon).

- Default procedure: grid WN -> Storage Element (SE) -> User Interface (UI) -> pc/cluster

Conclusions

- IKAROS helps us address several limitations which current file systems facing with large scale infrastructures.

- IKAROS approach enables us to create more user-driven computing facilities with application users and owners playing a decisive role in governance and focusing on placing computer science and the harvesting of ‘big data’ at the center of scientific discovery.