

WP3

Zetabyte –Exascale Storage  
Virtualization

# How does it fit in?





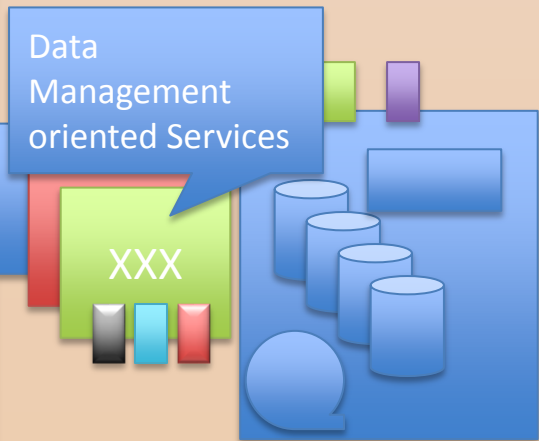
“user” oriented Services

User communities (small, large, stable, transient, funded, opportunistic )  
Storage and access needs, apps with different interfaces, Metadata needs (including Namespaces), retention policies



## Networks

Virtualization Layer (AAA, Mon, Access, Polices, Retention, locality, metadata....)



Providing: Work/Archive Storage, AAA interfaces, local namespaces, accounting, monitoring (local, global) .....

Traditional Providers    Cloud Providers    XXX Providers

# WP3 Challenges

- Defining, mapping common interfaces for functionality isn't enough → “just” ensures functionality (and is already partially solved)
- New approaches are needed that work efficiently at any scale
  - Blueprint and concrete architecture needs to be developed
    - critical elements modelled and verified by prototypes
  - Understanding who handles which functionality
    - ex. Global Namespaces (User/Infrastructure or Federation of local sub-spaces)?
    - Access control ( reading )? Community/Infrastructure ?
    - .....

# WP3 Shopping List

- Can't be complete at the moment (WP1&WP2)
- Main Domains:
- AAA
  - federating, access policies, accounting, costing...
- Locating Data
  - Namespaces, Metadata, ...
- Retention Policies
  - archive, replicas, caching, ....
- Data Management Infrastructure
  - data movers, replicators, cleaners (deletion services)
  - monitoring (performance, availability/reliability, resources, SLAs)
- Data Access (read/write)
  - Protocols
  - Access Control implementation
  - Network infrastructure interfaces .....

# Milestones and Deliverables?

- “Clean-slate approaches for multi-centric Zetabyte-Exascale Storage”
  - this mission statement will drive the milestones and deliverables

# WP3 Deliverables

- Architecture Blueprint Document for Virtualization
  - after about 1/2 of the project
  - describing the architecture and reasoning
  - supported by some modelling to verify the scalability of the concepts
  - Identifying the critical elements of the architecture that need to be verified by prototype implementations.
- Scalability and performance evaluation of prototypes
  - 2/3 of the project
  - individual building blocks ( as named in the Blueprint)
  - in collaboration with existing infrastructures (for scale)
  - these could be like:
    - Scalable Name/Object space
    - Data replication within the infrastructure
    - Metadata management, access, query, import, export
    - Storage federation mechanism
    - Identity federation.....

# Milestones?

- Initial milestone could deal with the interfaces between WP2 / WP3 / WP4 and the way policies are expressed
- **Functionality Definition Document** (early)
  - Specifying the functionality and policies from a user perspective. It should not be a classical requirement document driven by users (better horses problem)
- **Interface definition document** (PM10?)
  - Describing the protocols and interfaces between WP2 -WP3 WP3\_WP4.
  - Starting point only, needs updates as the architecture matures
- **Demonstrating scalability and performance of the integrated WP2-WP3-WP4 prototypes** (late)
  - This is a report on the stress tests carried out on the prototype services, strongly linked to the Deliverables