

Data-intensive Storage Services on Clouds: The VISION Cloud Project



Simona Rabinovici-Cohen, Hillel Kolodner
IBM Research - Haifa

- Introduction
- Innovations
- Use Cases
- Architecture

VISION Cloud: Virtualized Storage Services Foundation for the Future Internet

- Goal
 - Architect and implement an infrastructure for the reliable and effective delivery of data-intensive storage services, facilitating the convergence of ICT, media and telecommunications


- Innovations
 - **Raise Abstraction Level of Storage:** objects with user- and system-defined metadata
 - **Computational Storage:** technology for specifying/executing computations close to storage
 - **Content-Centric Storage:** facilitate access to data by content and its relationships
 - **Advanced Capabilities for Cloud-based Storage:** support delivery of data-intensive services securely, at the desired QoS, at competitive costs
 - **Data Mobility and Federation:** enable comprehensive data migration and interoperability across remote locations

- Use cases: Media, Telco, Healthcare, Enterprise

- Facts:
 - A 3-year project, started Oct 2010
 - www.visioncloud.eu

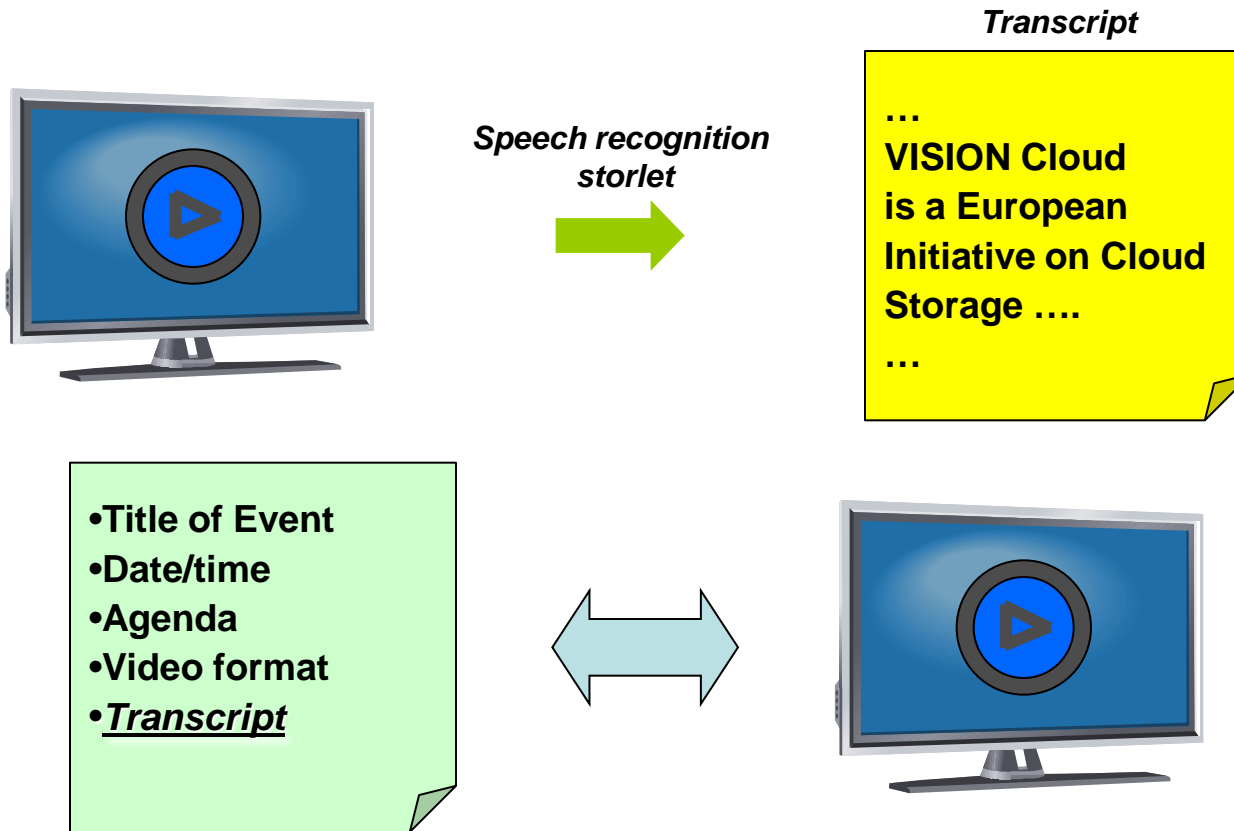


- ❑ **Store video of the summit together with rich metadata**



- ❑ **What is new:**
 - ❑ Metadata is an integral part of the storage
 - ❑ Rich metadata model describing both handling of an object and its content

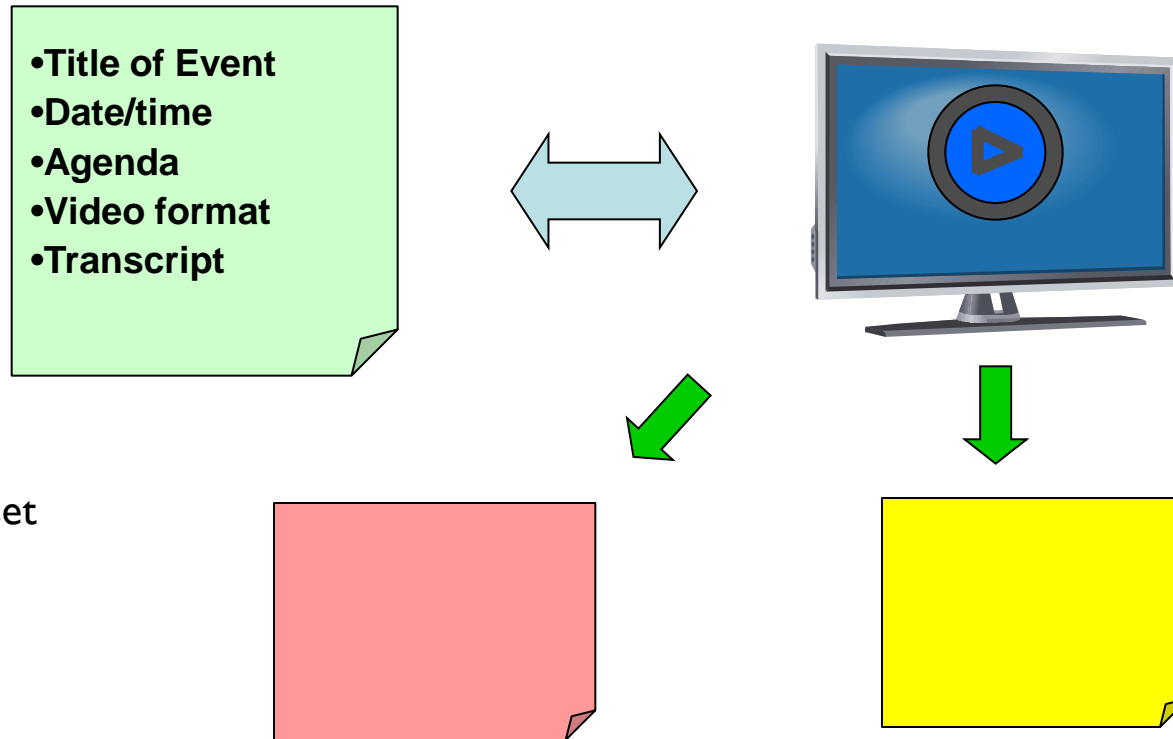
- A storlet is triggered to automatically extract metadata



- What is new:
 - Architected and safe way to run computations in the storage system

Content-Centric Storage

- ❑ Access data according to metadata values
- ❑ Build content networks

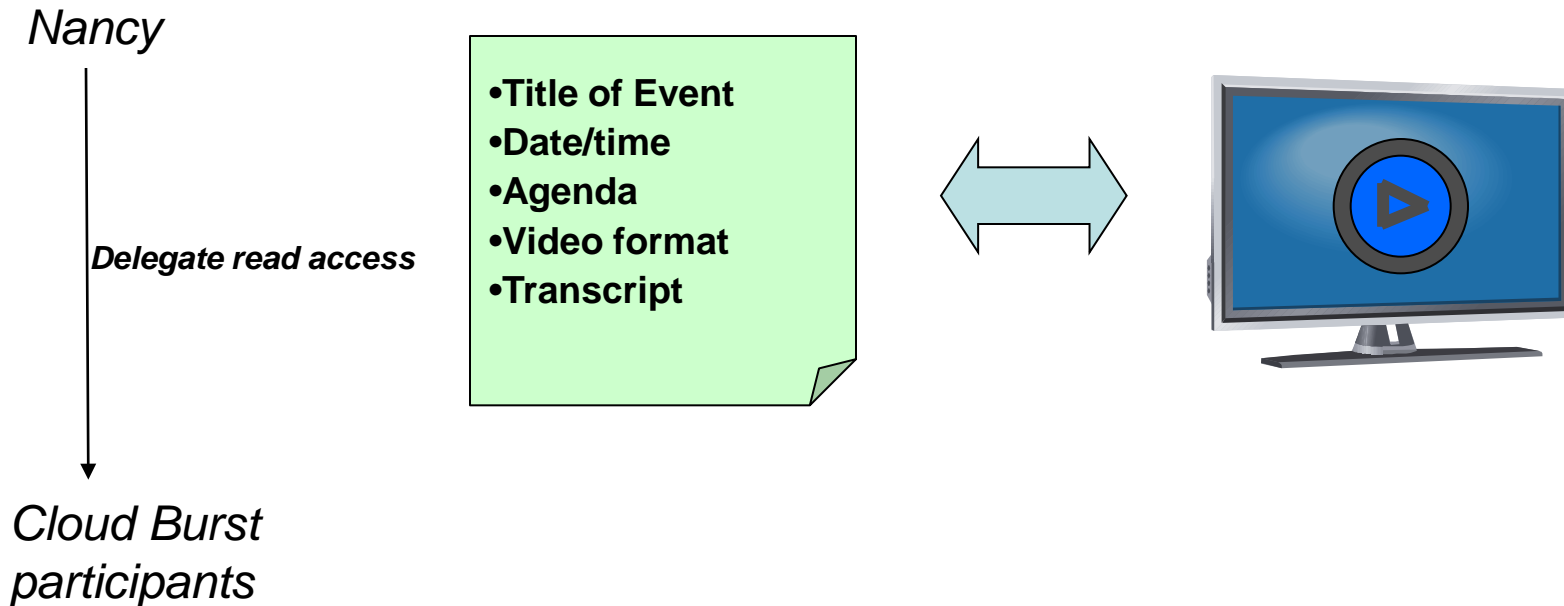


Relations:

Equivalence, list, set

- ❑ What is new:
 - ❑ Storage leverages metadata and content networks to optimize itself

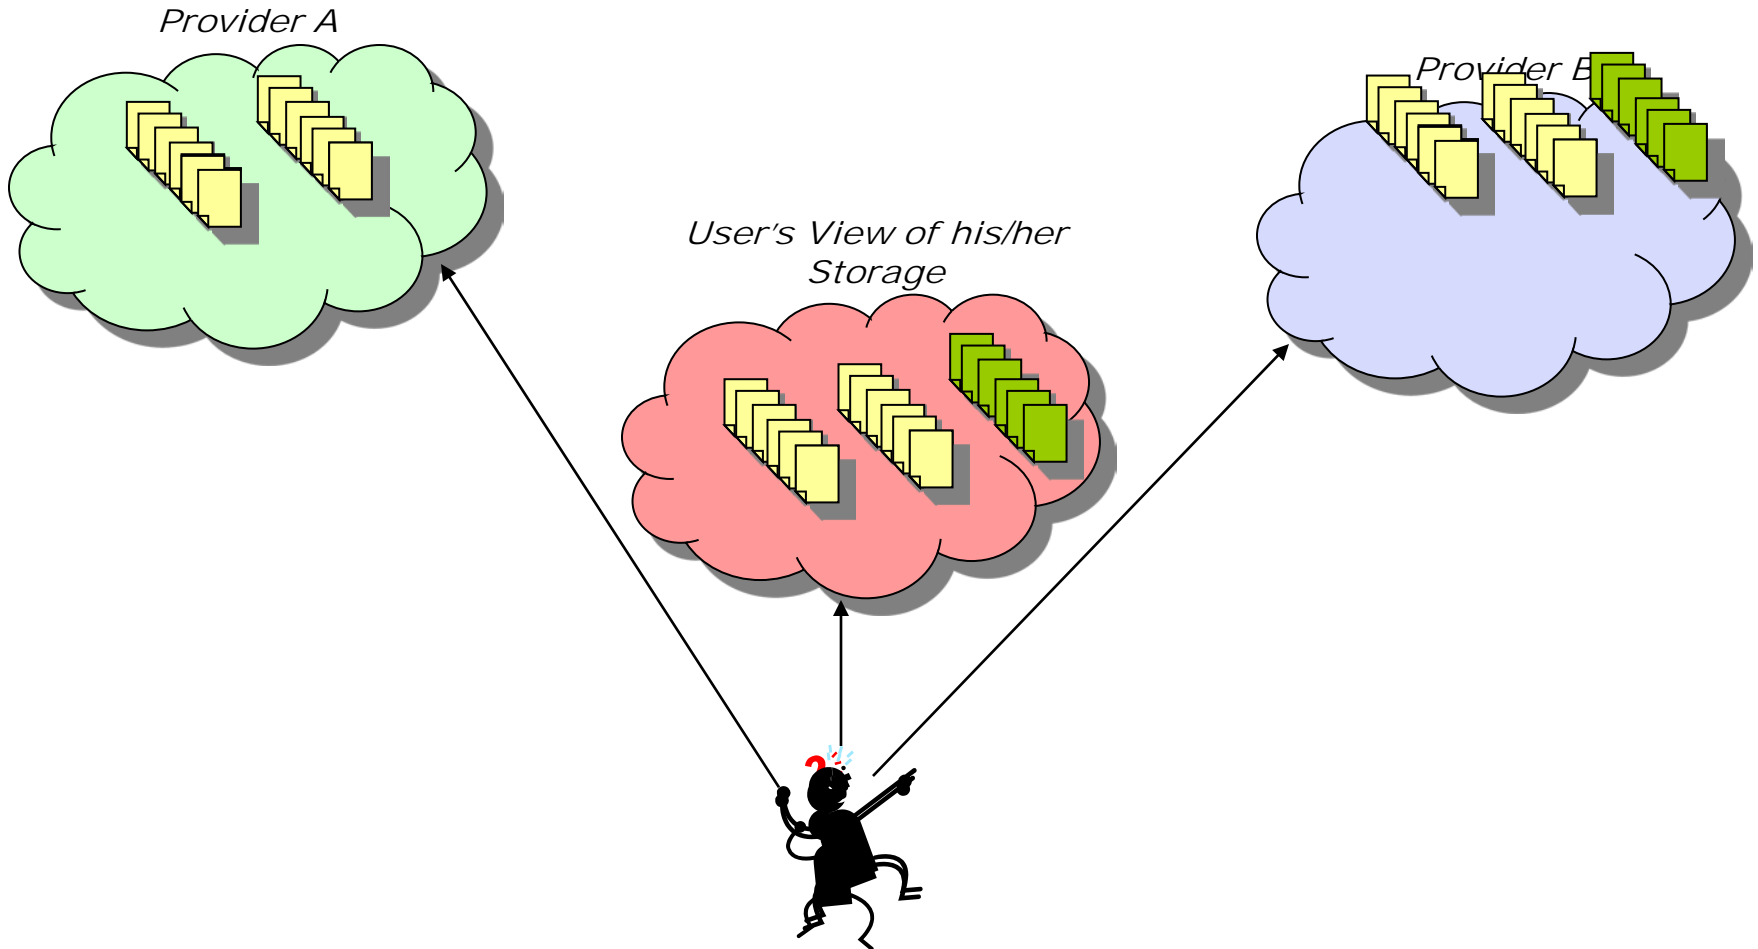
- Delegate right to access an object to people that are not known by the storage system



- What is new:
 - Flexible yet secure access control

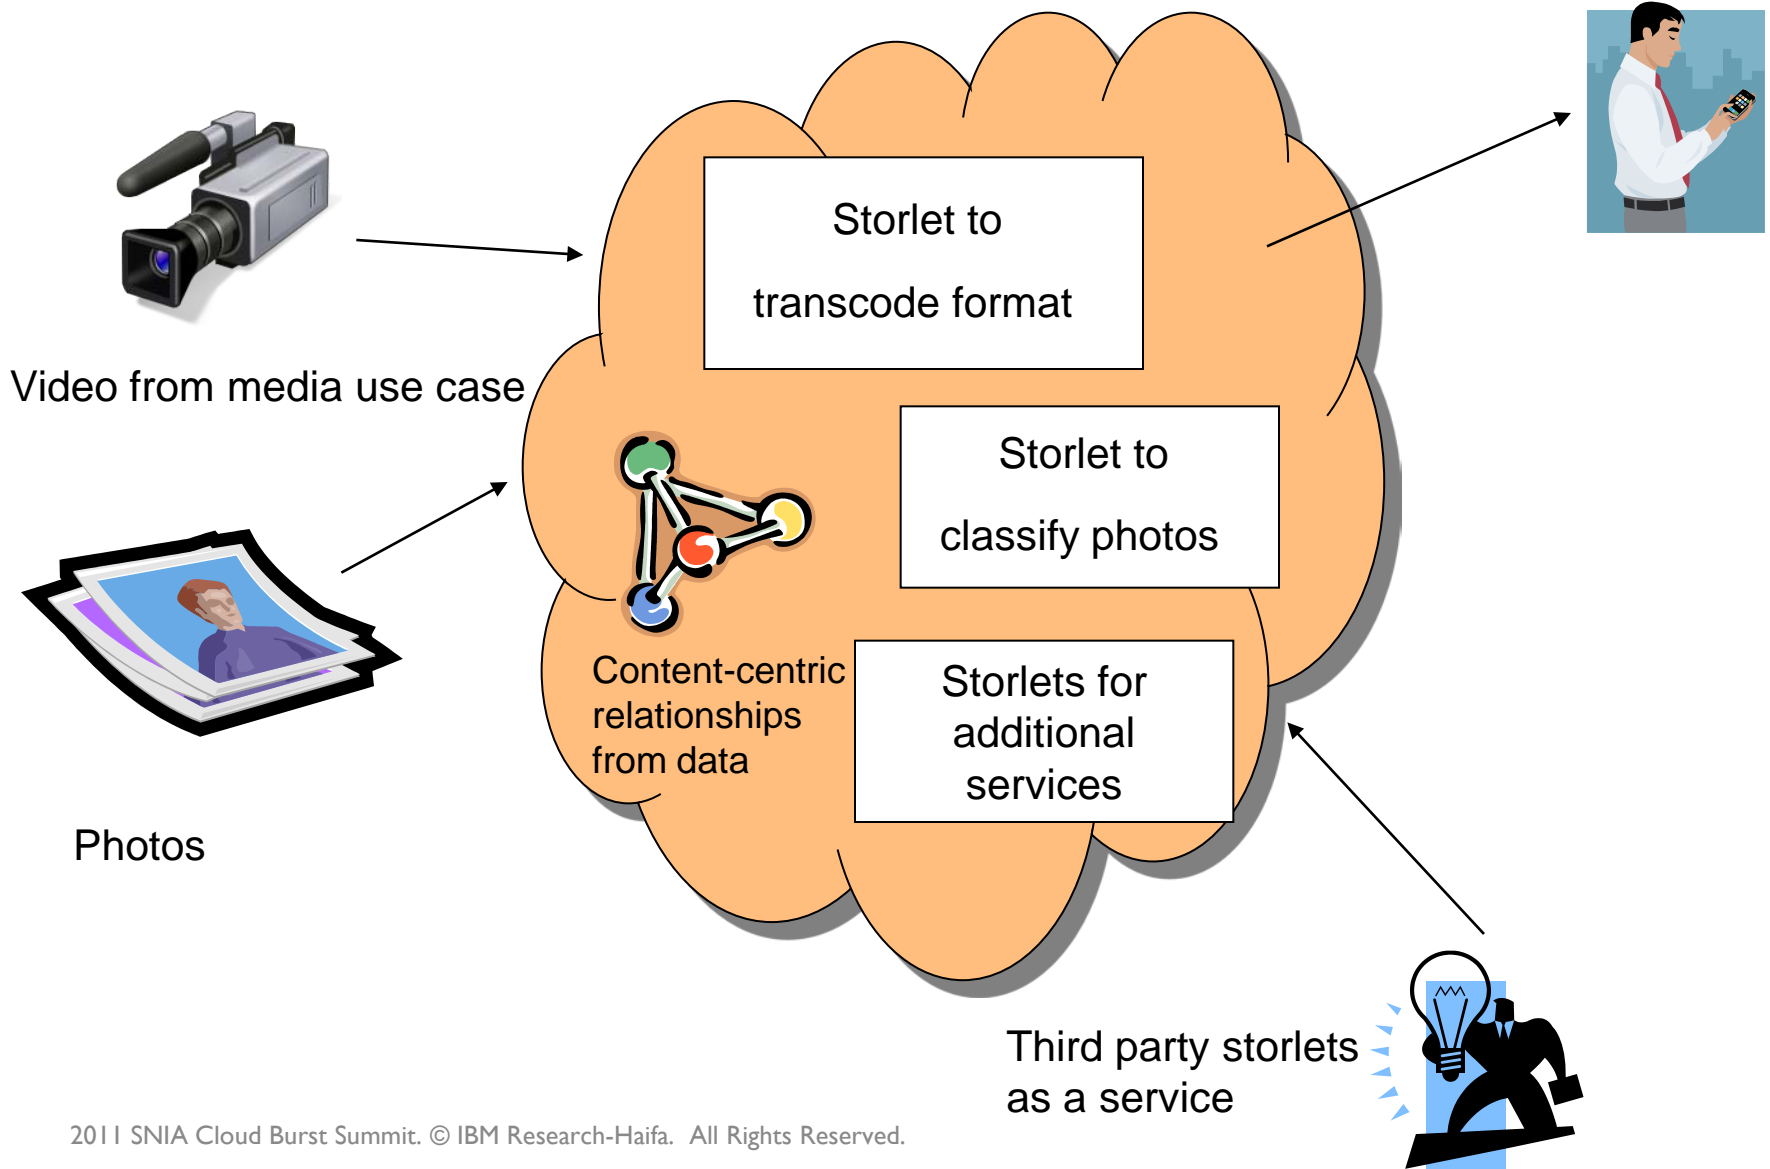
Federation and Interoperability

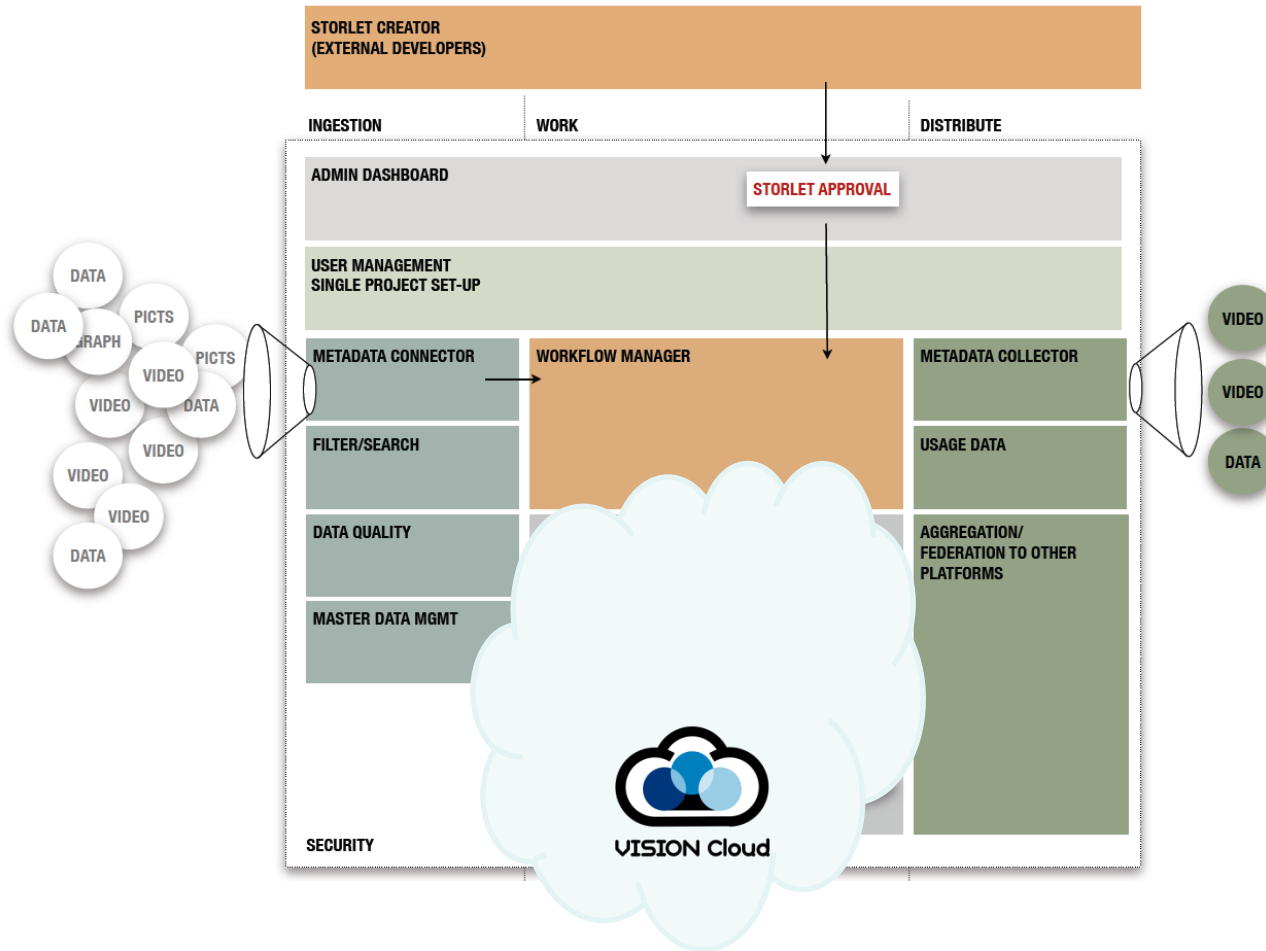
- Change storage providers without data lock-in



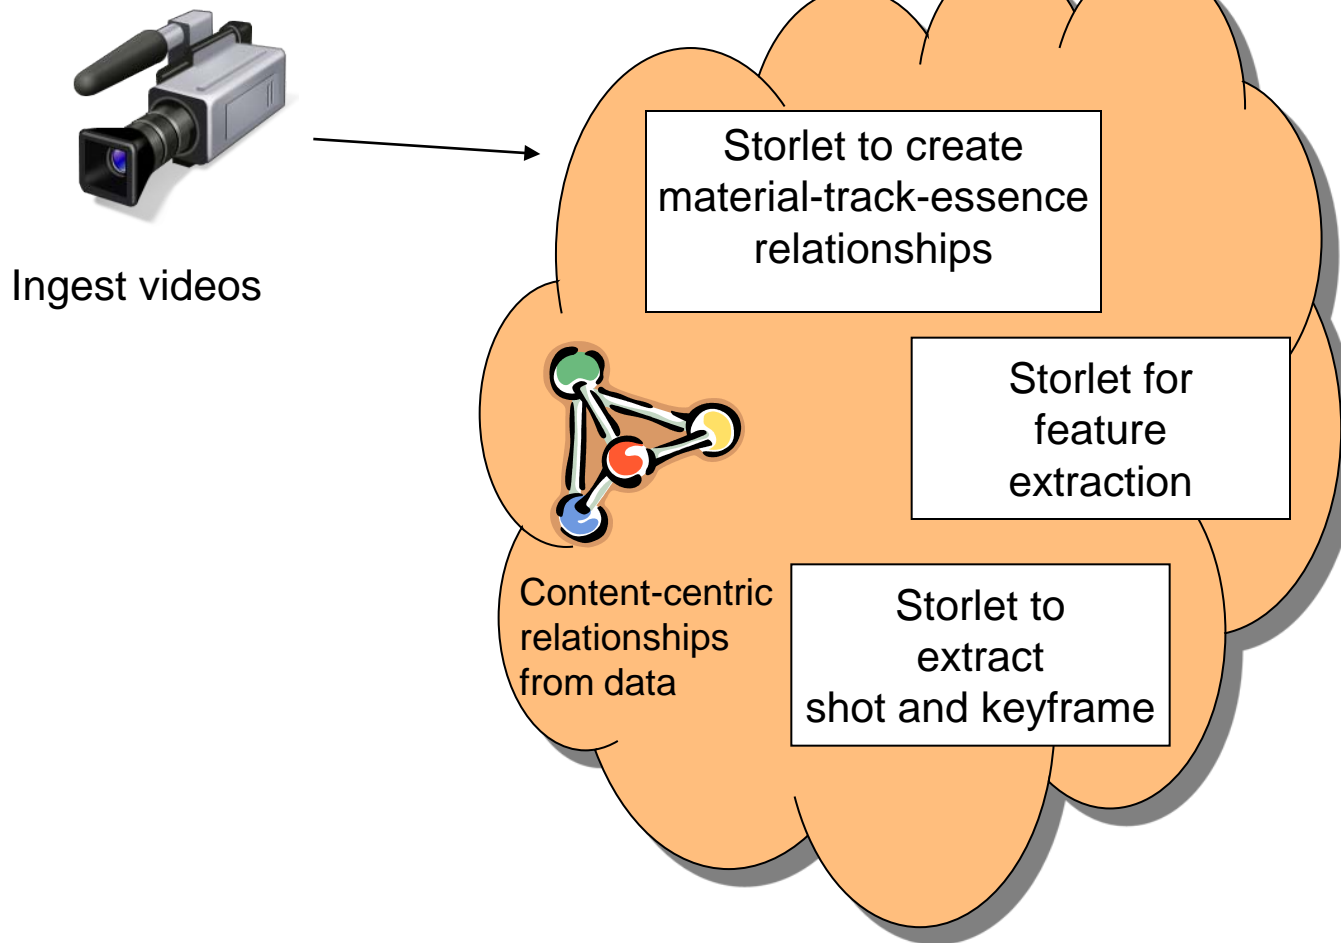
- Introduction
- Innovations
- Use Cases
- Architecture

Telco use case

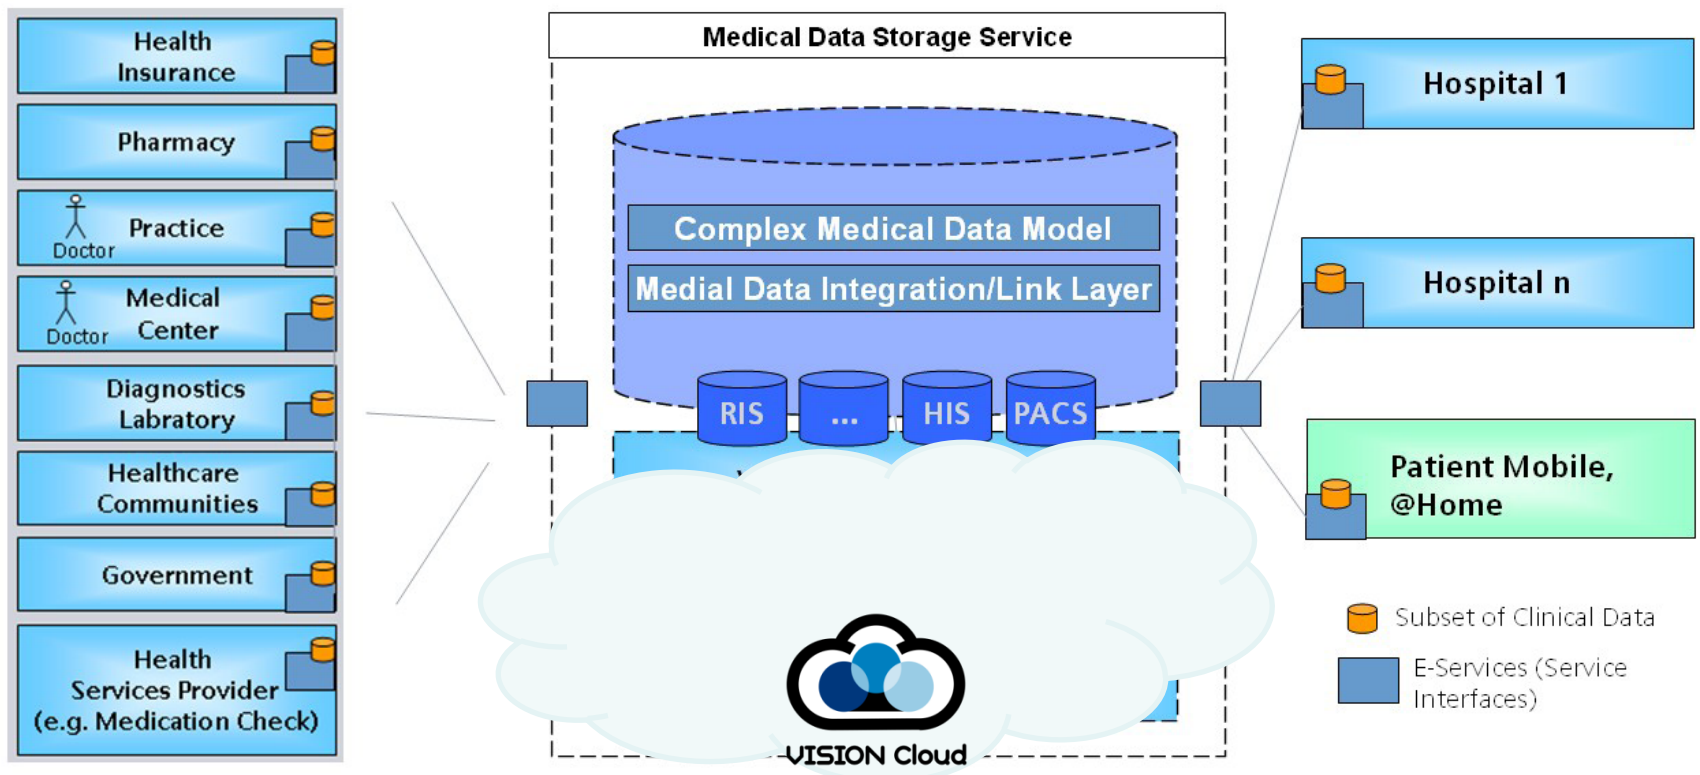




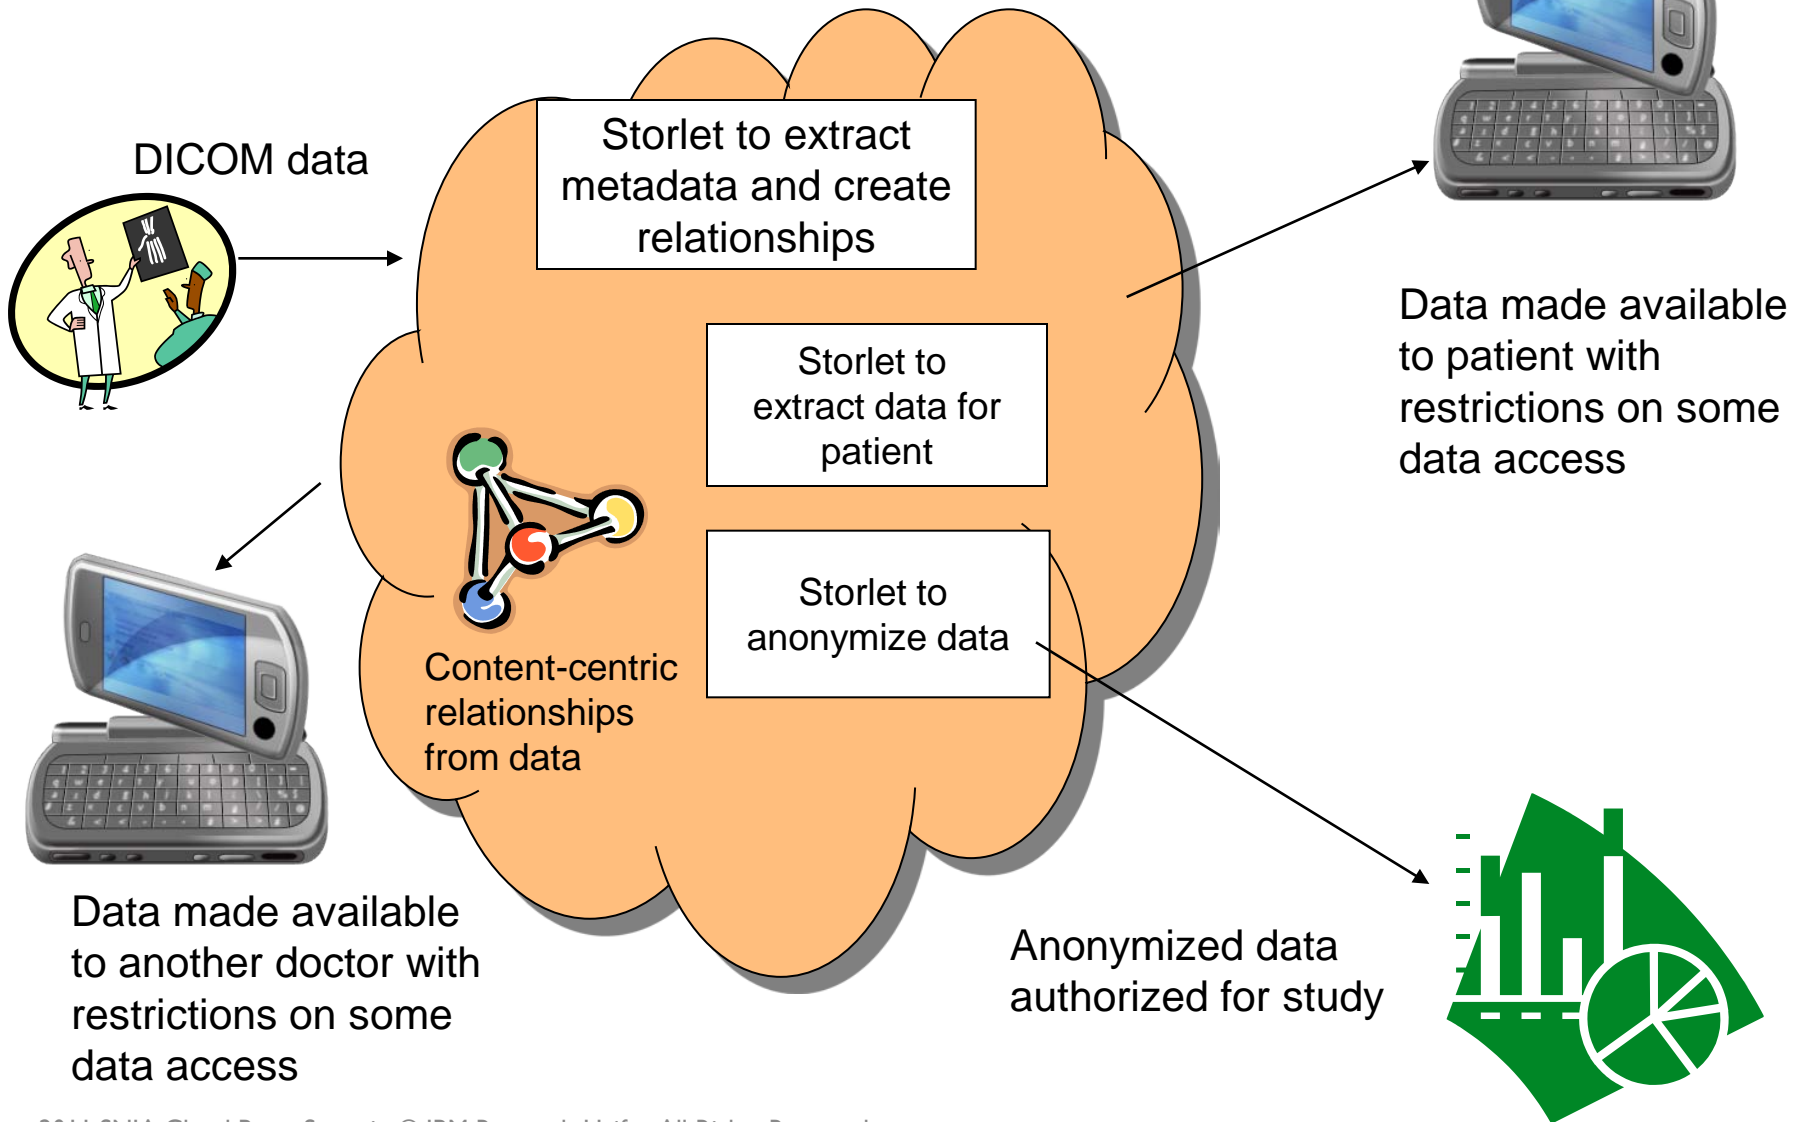
Media use case



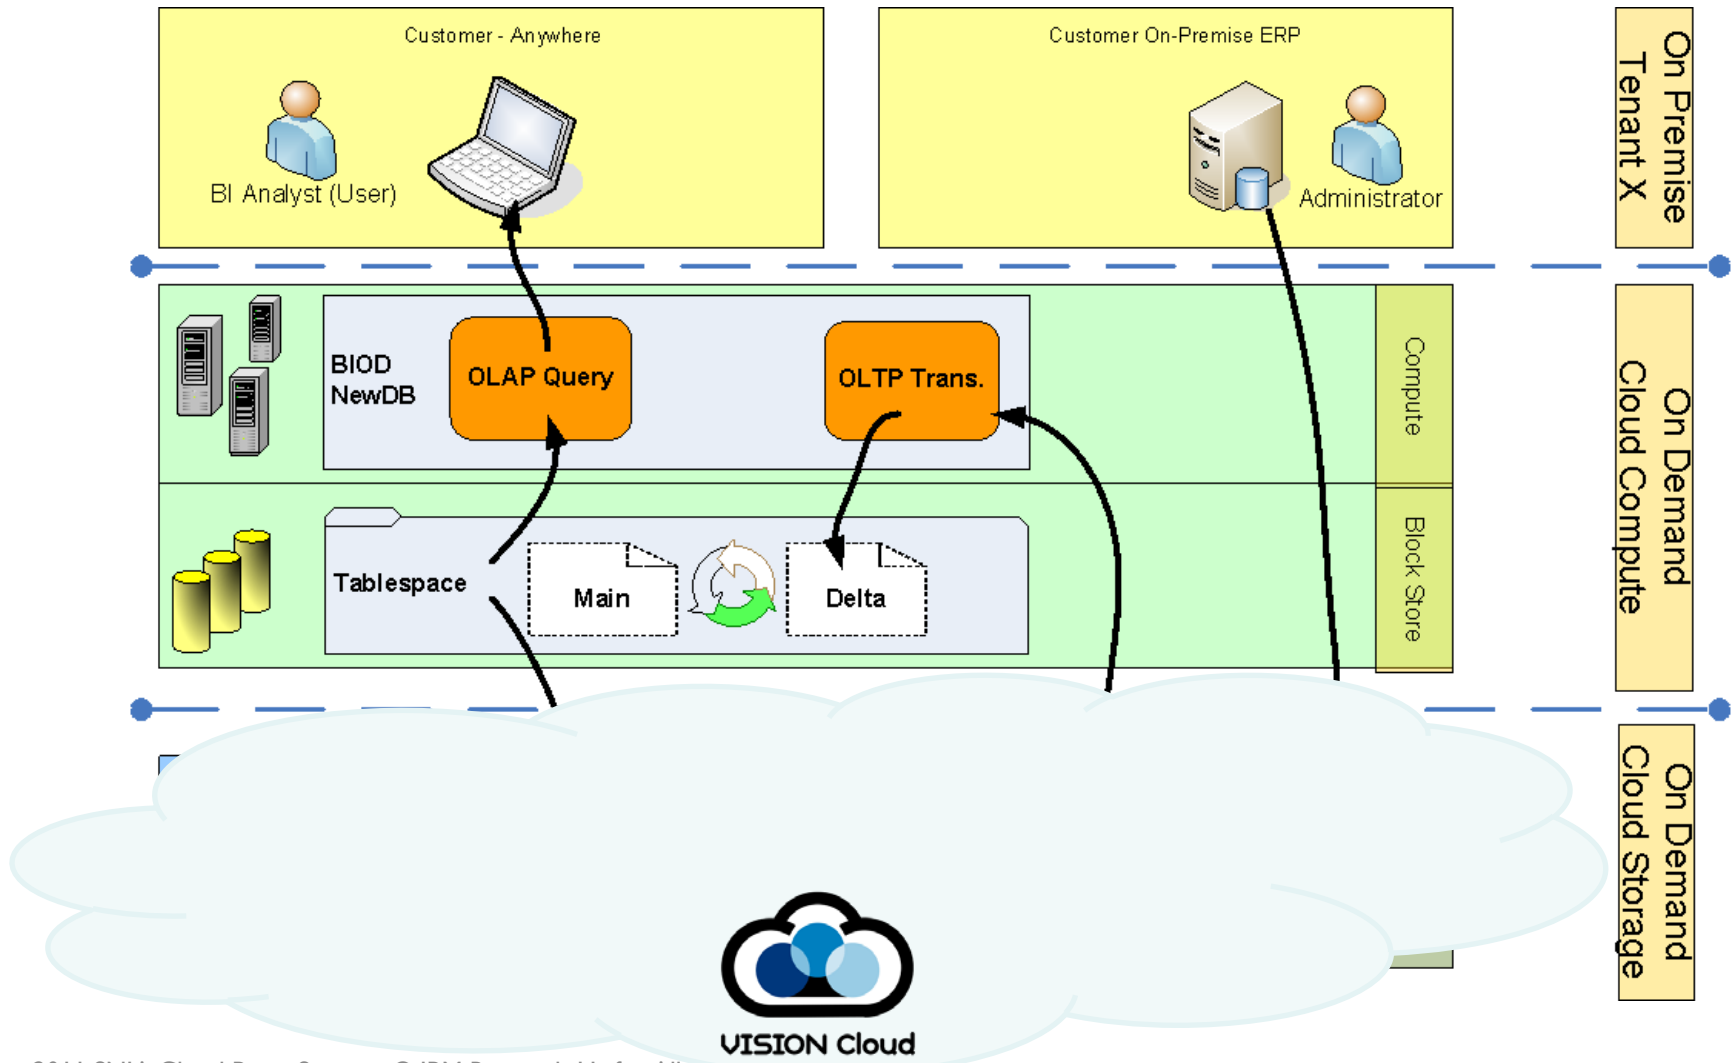
Healthcare



Healthcare use case

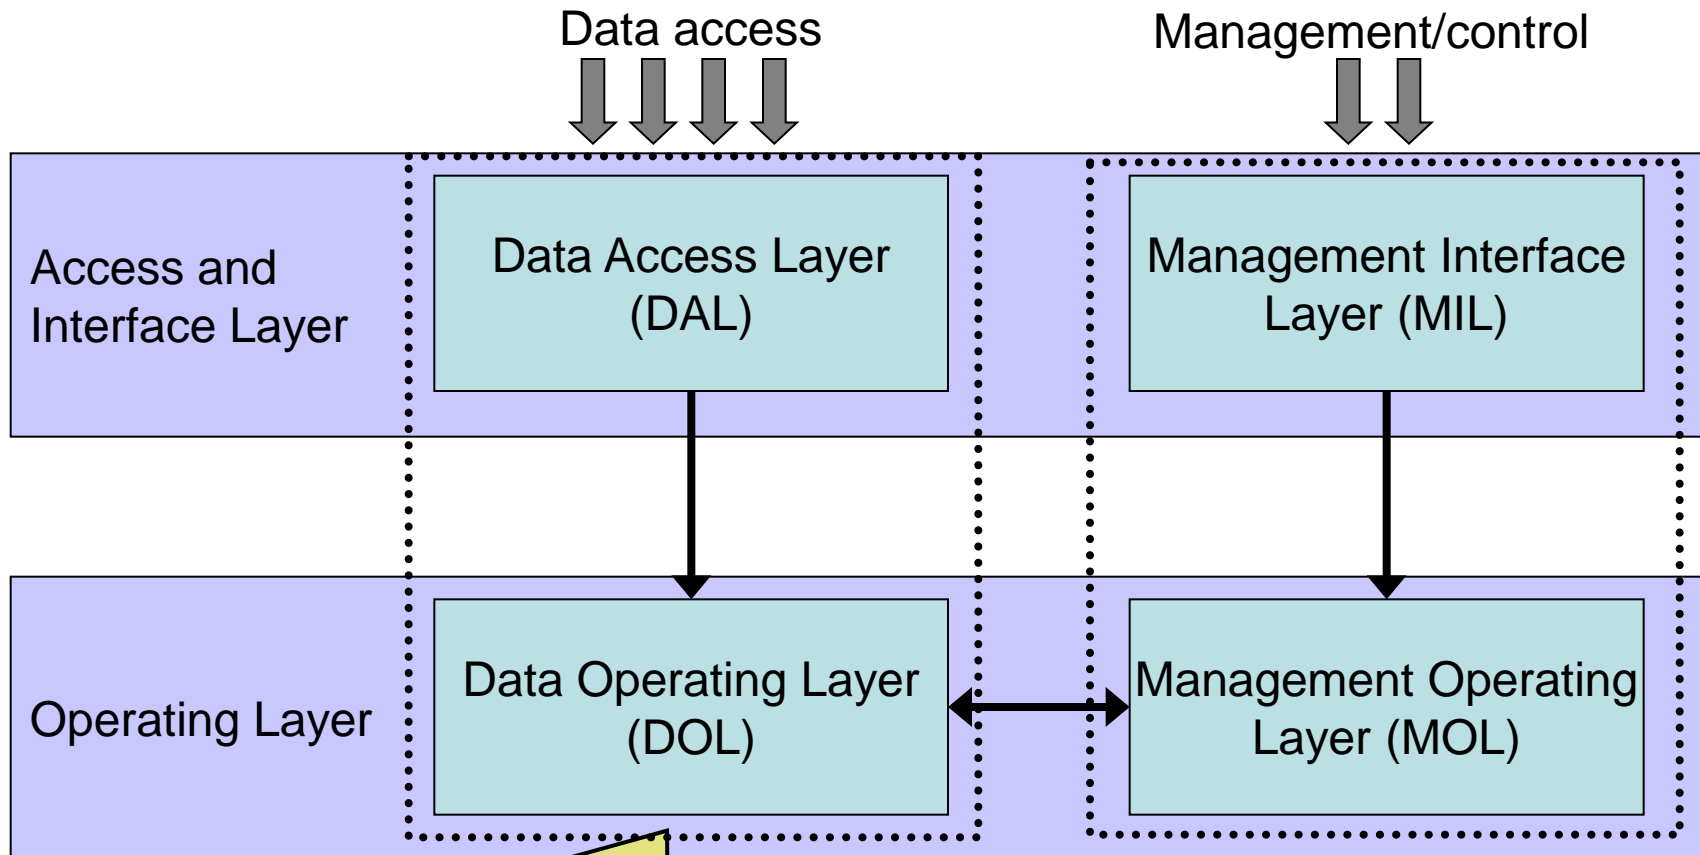


Enterprise Software



- Introduction
- Innovations
- Use Cases
- Architecture

The VISION Cloud Architecture



DATA SERVICE

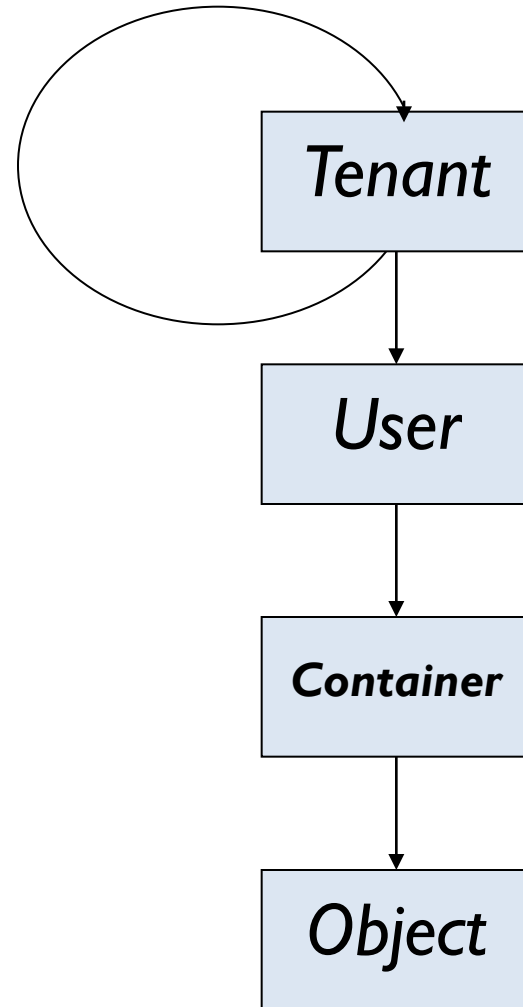
Content networks/objects,
Computation on storage,
Mobility, availability, reliability, security

MANAGEMENT SERVICE

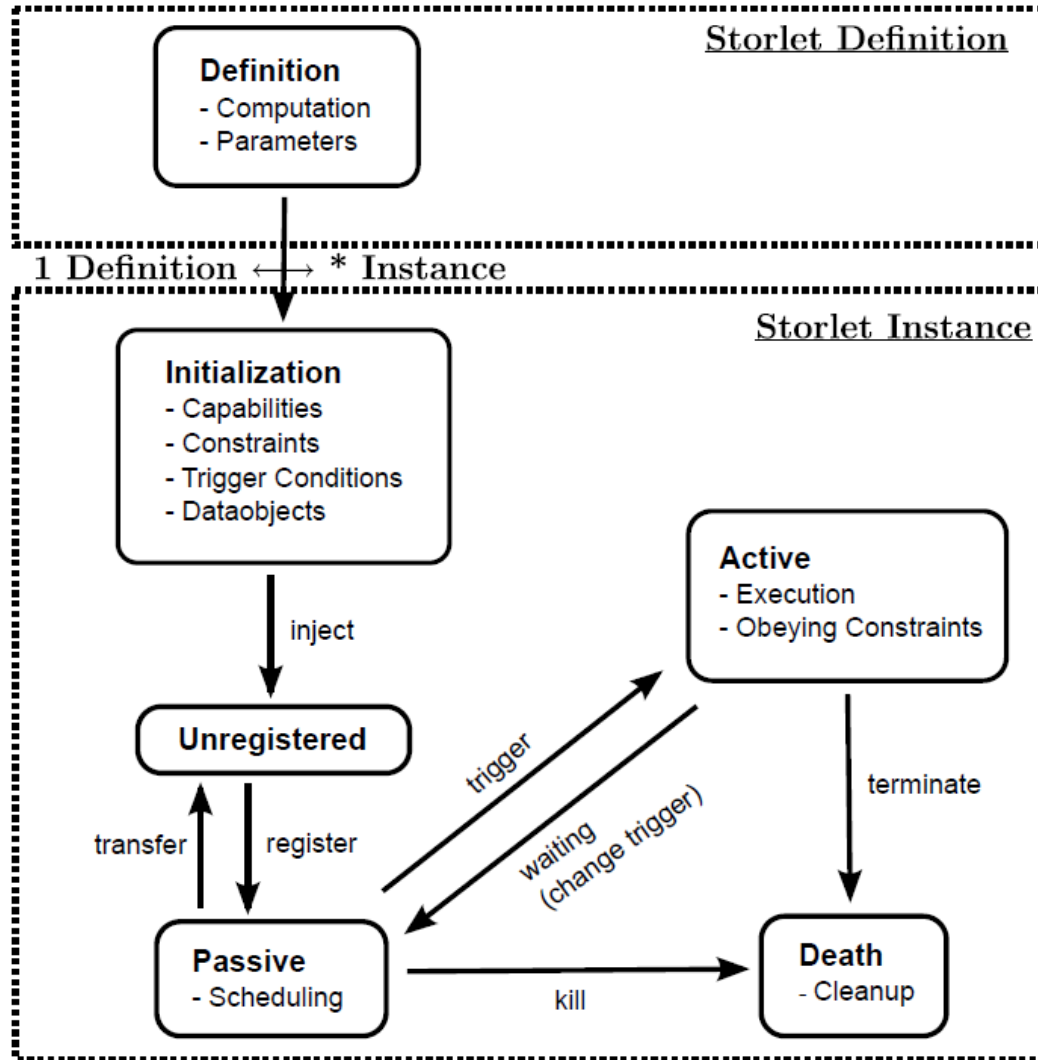
Monitoring, Metering, Billing,
Security management, Tenant/User
management, SLA

High Level Concepts and Data Model

- ❑ Objects are write all-at-once
- ❑ Metadata
 - ❑ System
 - ❑ Management directives
 - ❑ User
 - ❑ Key-value pairs
 - ❑ Schema
- ❑ Metadata can be updated
- ❑ Versioning for logical protection
- ❑ Symmetric replication for resiliency
- ❑ Eventual consistency

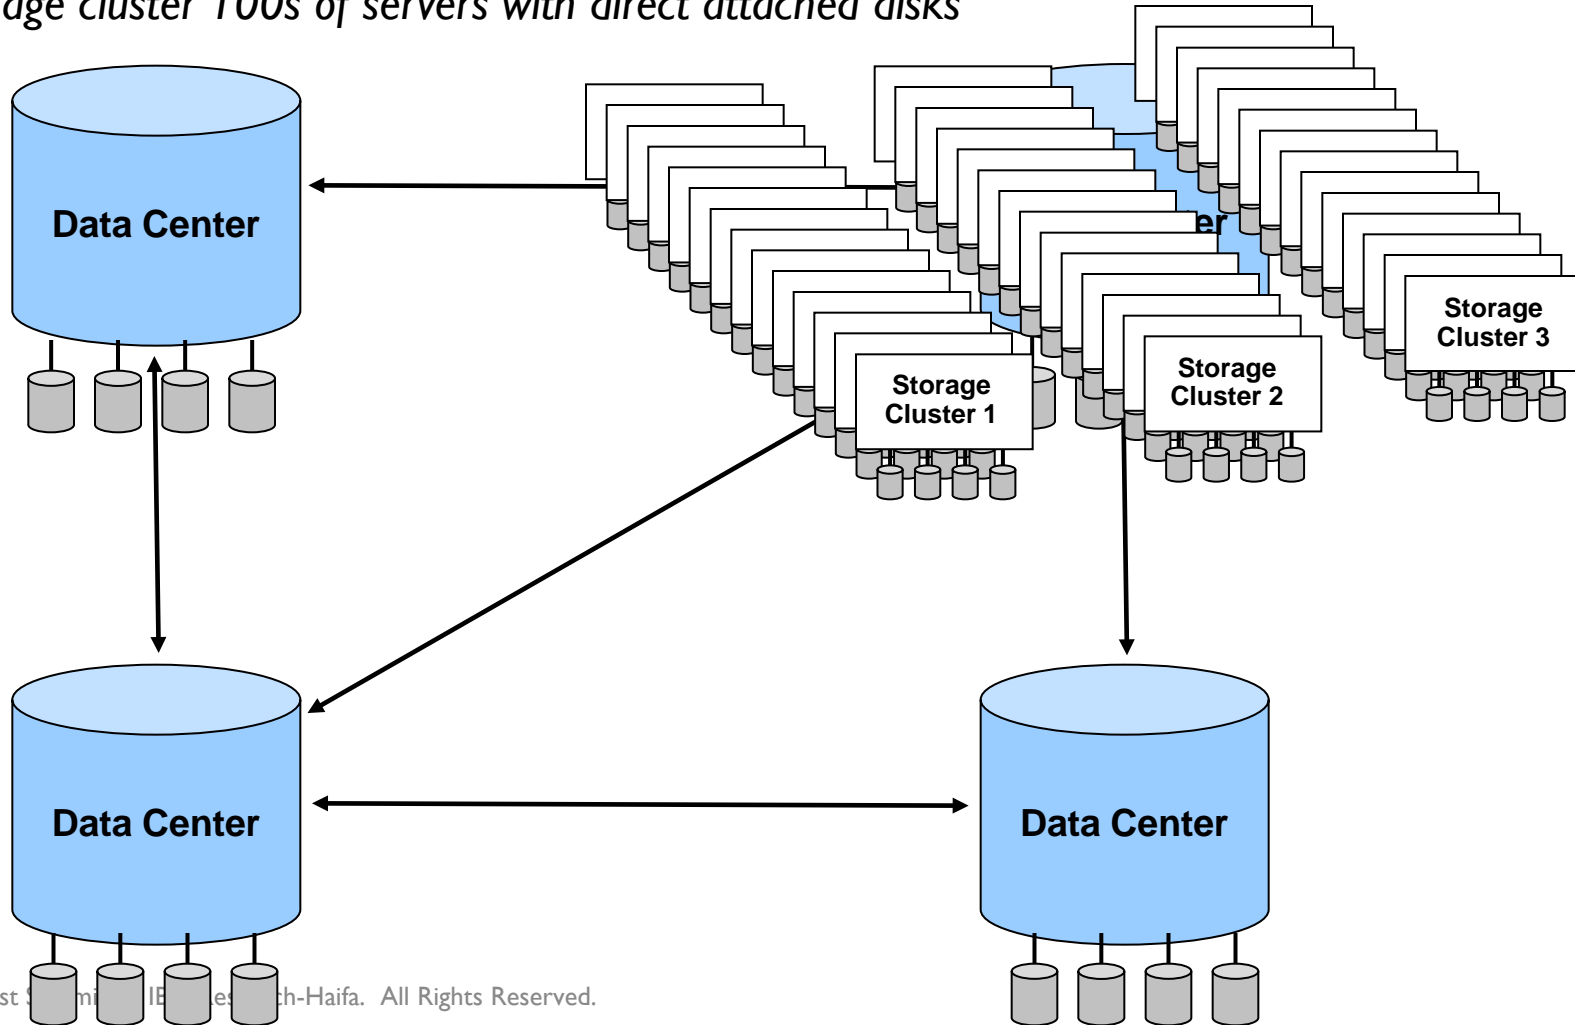


Storlet Life Cycle and States




Physical Model

- ❑ 100s of Data Centers (DC)
- ❑ Each DC 10s of storage clusters
- ❑ Each storage cluster 100s of servers with direct attached disks



Data Access Flow



Logical model

- ❑ External interface is REST
- ❑ Every server runs same software stack
 - ❑ Basic stack – Apache, Cassandra, file system
 - ❑ VISION Cloud components of DAL, DOL, MIL, MOL
- ❑ Many independent requests processed in parallel by each server
- ❑ Some servers in each cluster also run global view
- ❑ An object can be placed on a specific server
- ❑ Shared state at the cluster level belongs in the catalog
- ❑ Shared state at the cloud level belongs in the global view

Global View

Global View

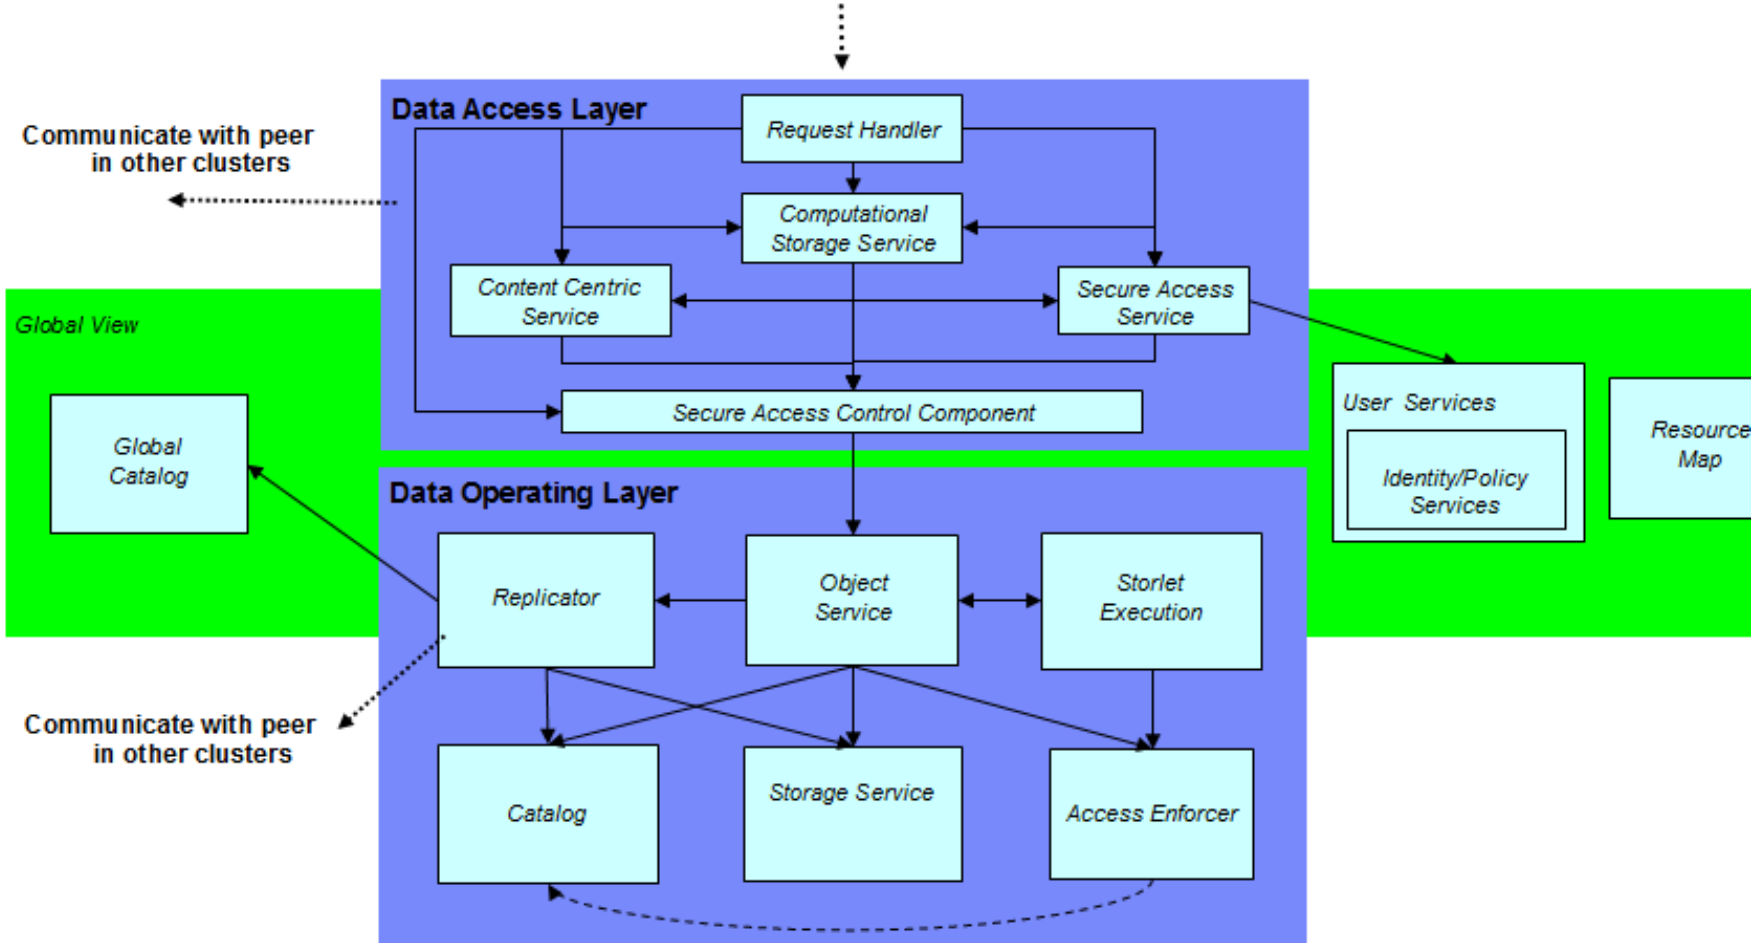
*Global
Catalog*

User Services
*Identity/Policy
Services*

*Resource
Map*

Data Access and Operating Layers

Application requests

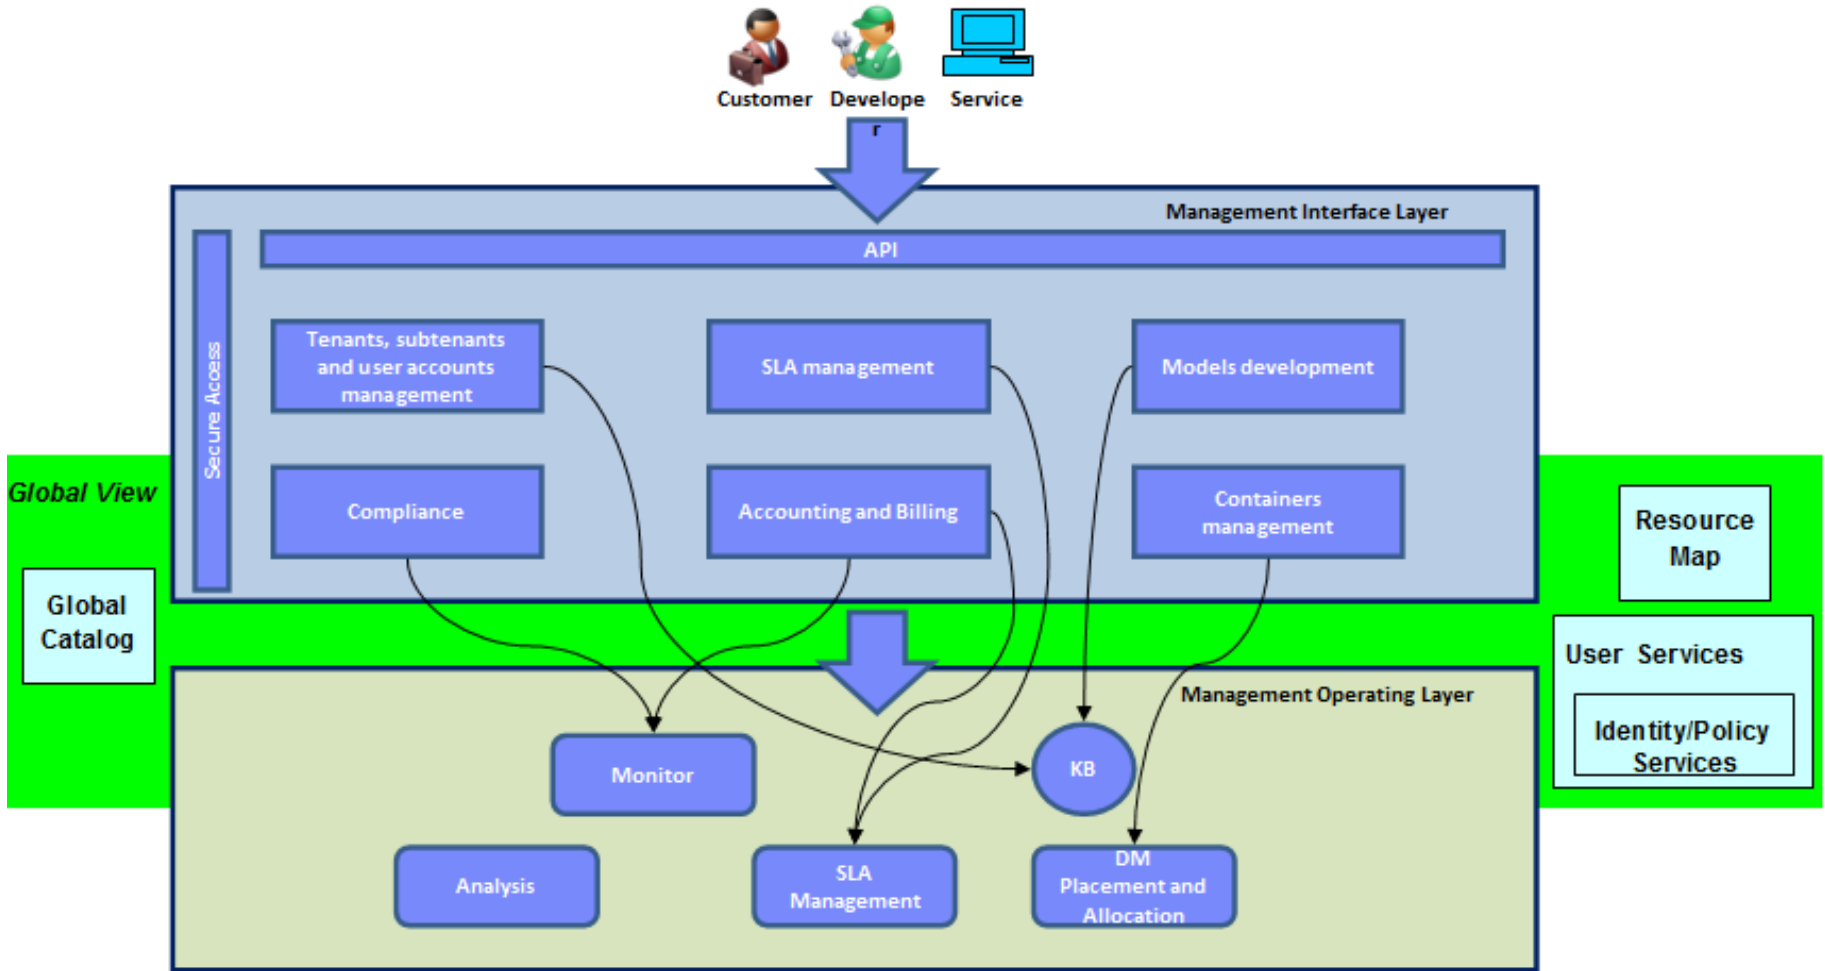


Logically global components

Components that have an independent instance on each cluster node

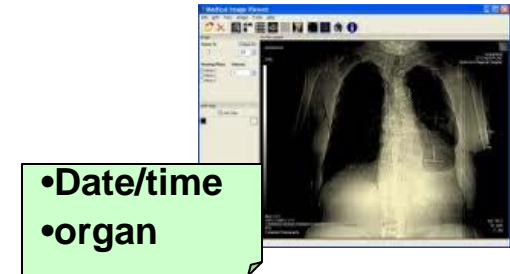
Logically global components

Management Architecture




Suggested Enhancements to CDMI

- ❑ Large binary objects with metadata
 - ❑ eliminate conversions in the JSON payload



- ❑ Advanced queries
 - ❑ Support range queries, list container, cursors, etc.
 - ❑ Not just query queues



- ❑ Computational storage
 - ❑ Add interface for managing and triggering computation in the storage

