Object Storage

Key Role in Future of Big Data

Anil Vasudeva
President & Chief Analyst
IMEX Research.com

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA and the author.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - SNIA and author must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

Object Storage: Key Role in Future of Big Data

- This session will appeal to Data Center Managers, Development Managers, and those that are seeking a fundamental understanding of evolution of Object Storage to deal with the explosive growth in unstructured data. The session delves into the impact of Object Storage architectures on next generation storage in Private and Hybrid Clouds as well as its ability to gain meaningful business insights from Big Data given Object Storage flat hierarchical structure and associated metadata attributes governing policies applicable to Big Data. The metadata in object storage can become key in application driven Software Defined storage. The rapid adoption of Object storage by both startups and major IT vendors is a testament to their advantage as the next generation storage for Big Data. The session is targeted to bring a clear understanding of this key technology to developers as well as systems admins and large service providers providing Big Data Analytics.

Unstructured Data Explosion

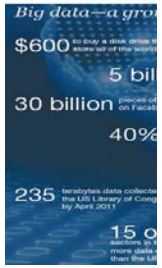

The Unstructured Data Explosion

- Growing 100X every 10 years
- Requires a new approach

New Data

Hi-Res:
50,000,000 KB / object

Video:
5,000,000 KB / movie

Audio:
5,000 KB / song

Traditional Data

Images:
1,000 KB / picture

58674322
98323456
20419335
49567053

Numbers:
5 KB / record

Text:
500 KB / record

Harnessing Big Data Infrastructure

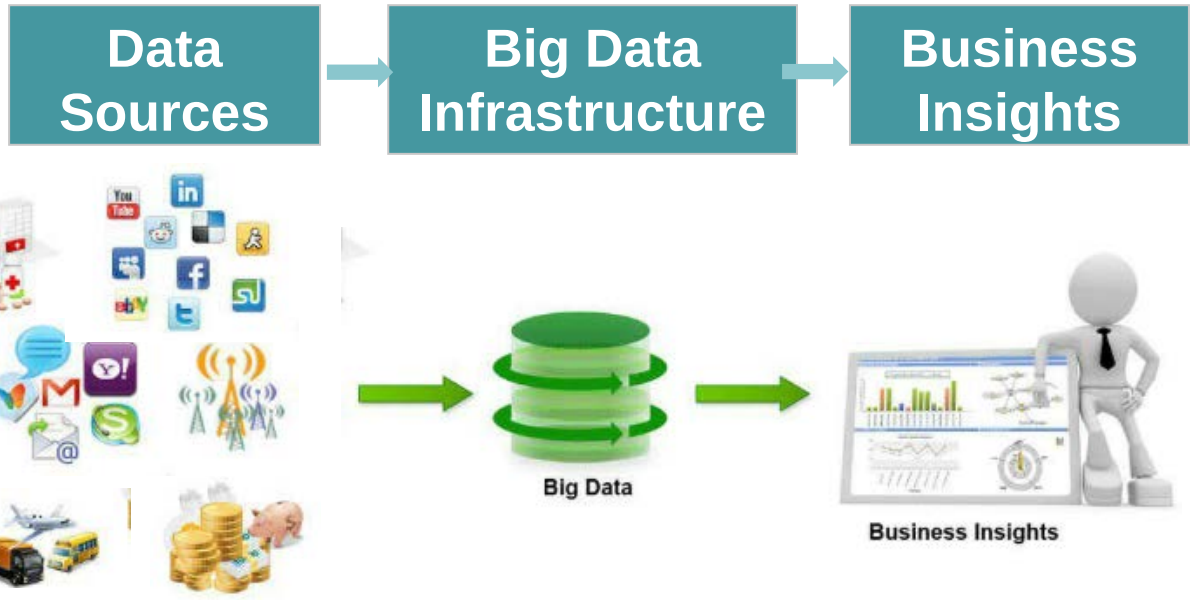

Information is at the center of New Wave of opportunity

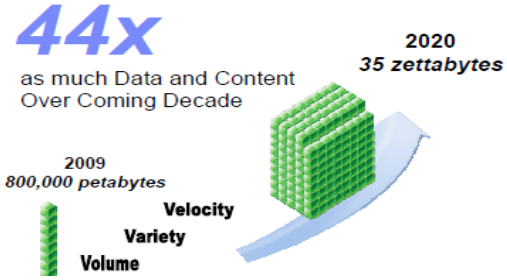

Majority of data growth is being driven by unstructured data and billions of large objects

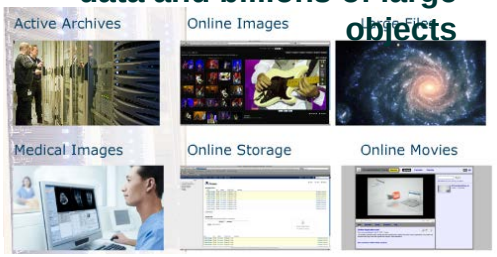

80% of world's data is unstructured driven by rise in Mobility devices, collaboration machine generated data.

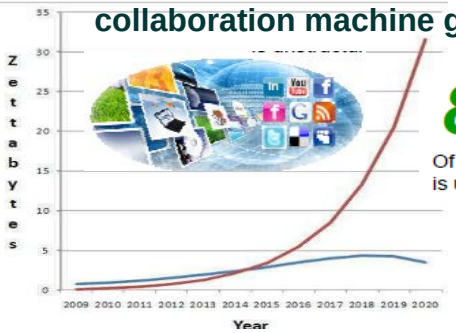

80%
Of world's data is unstructured

Object Storage Key Role in Future Data

Big Data Analytics

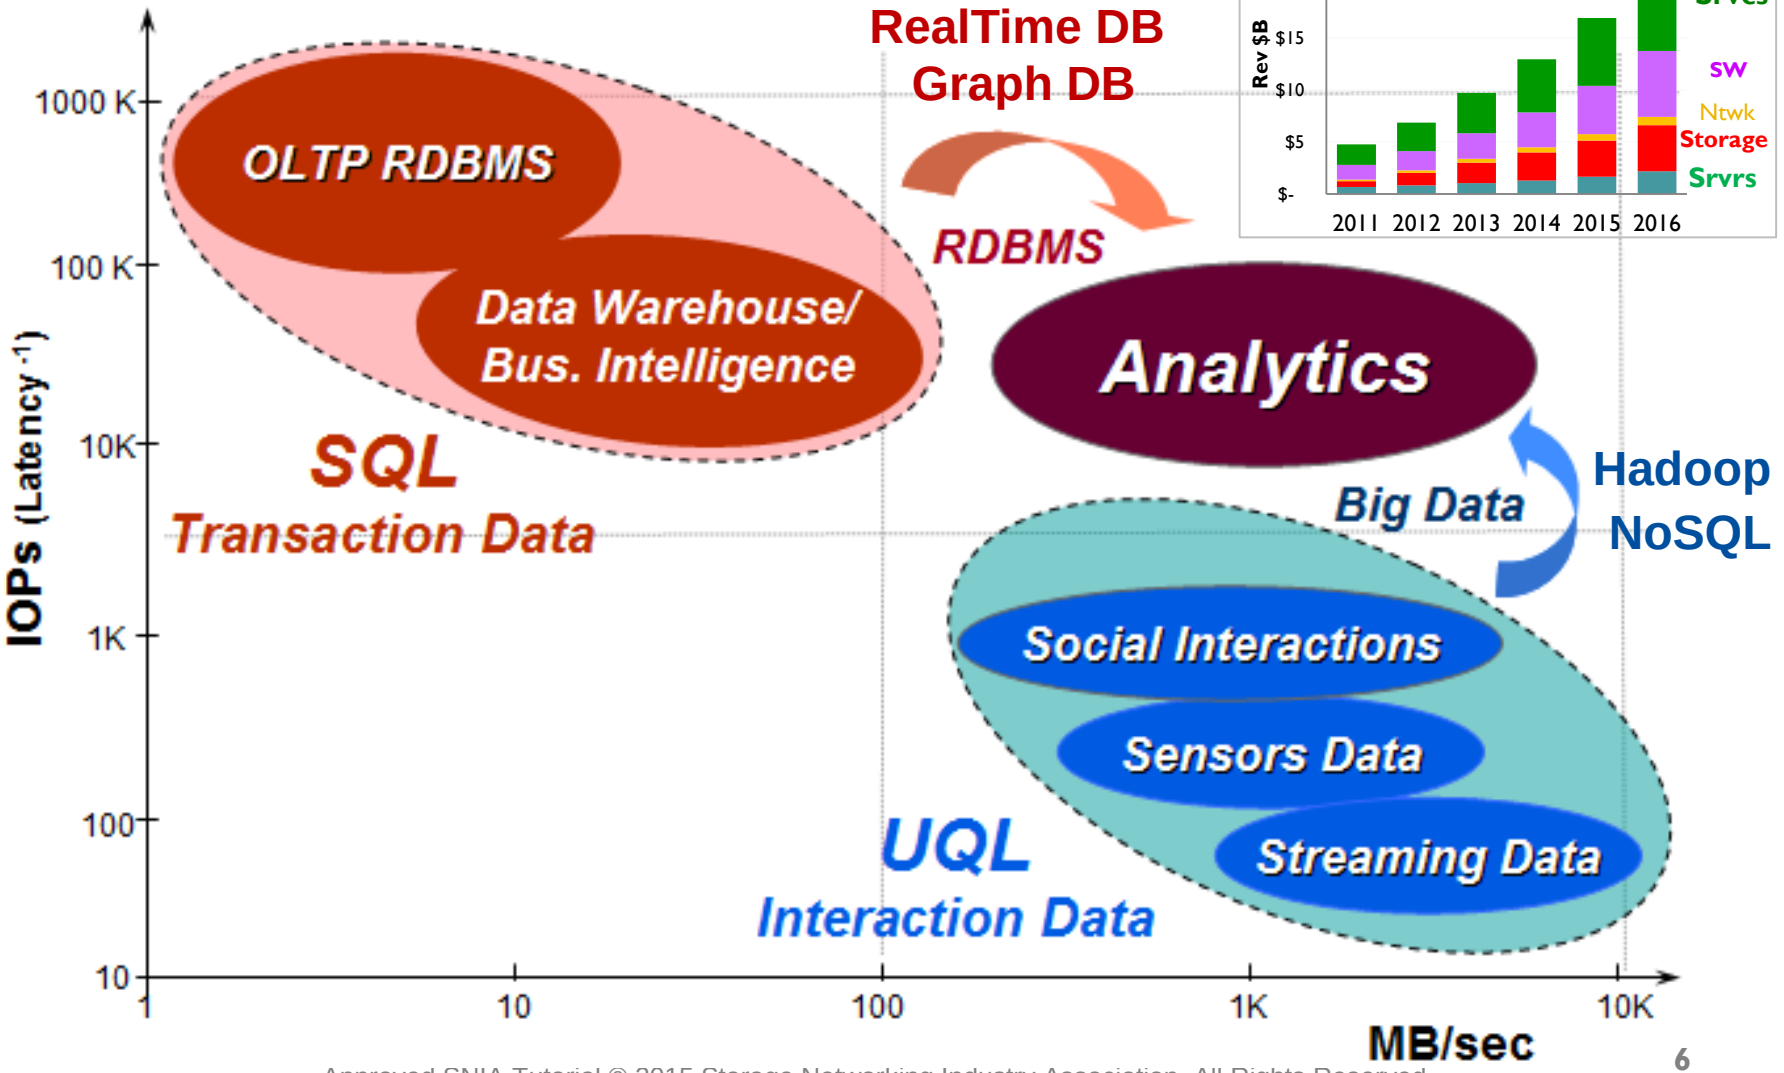

Big Data Ecosystem

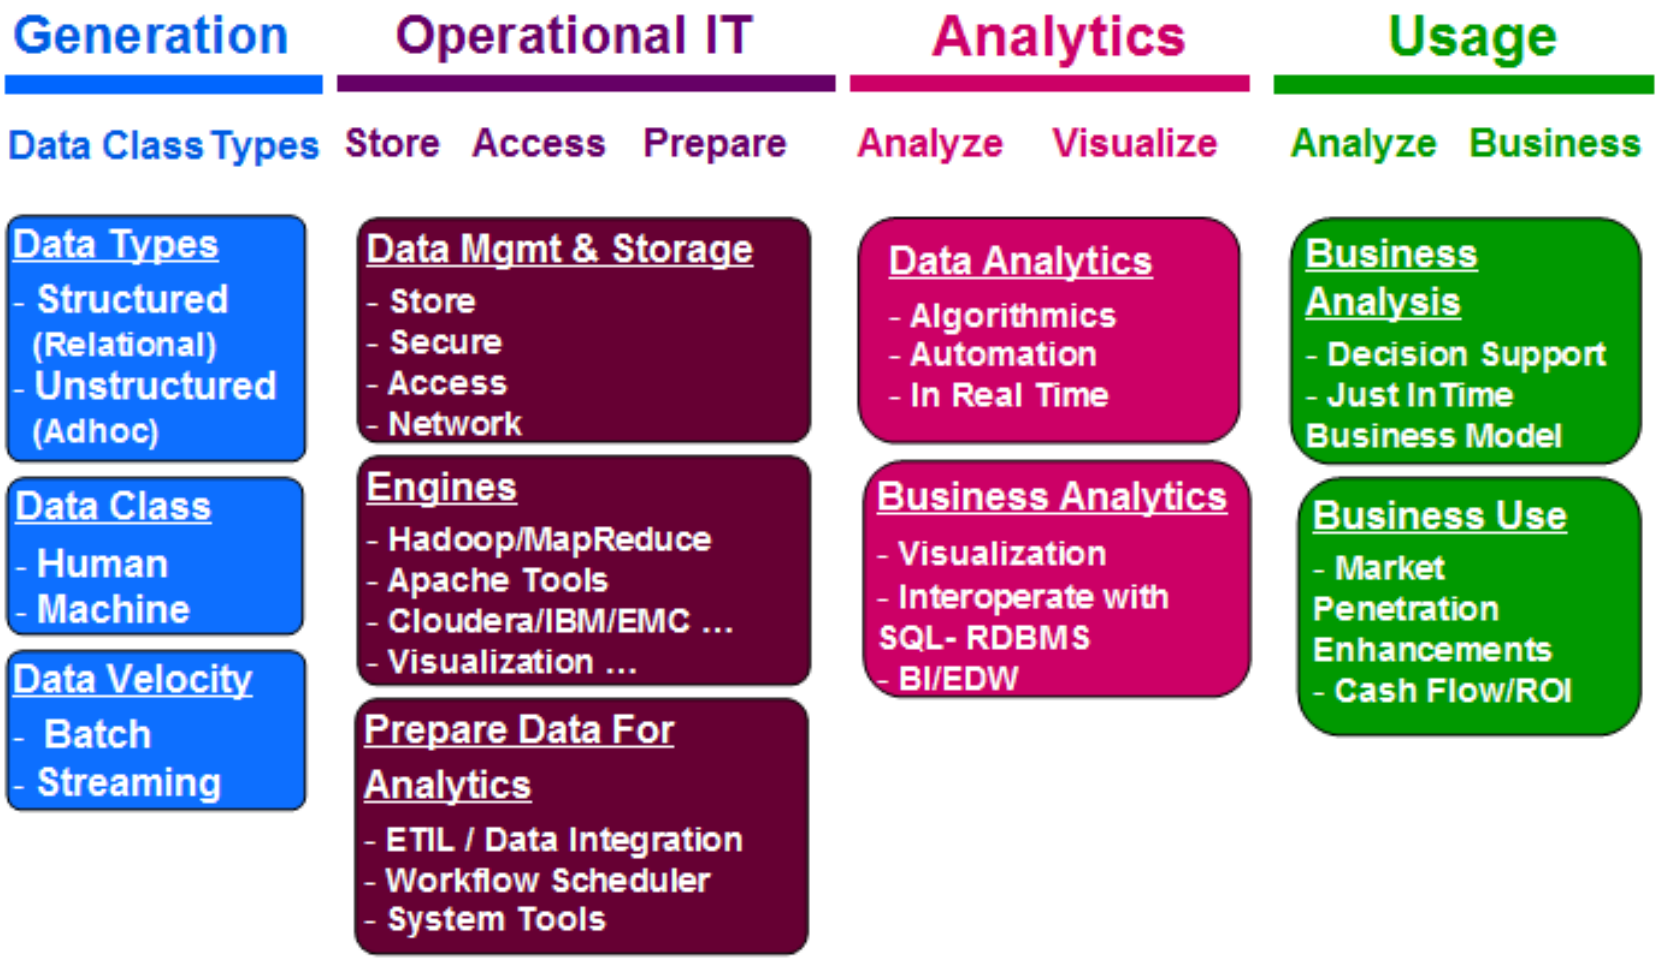

Big Data HDFS Architecture

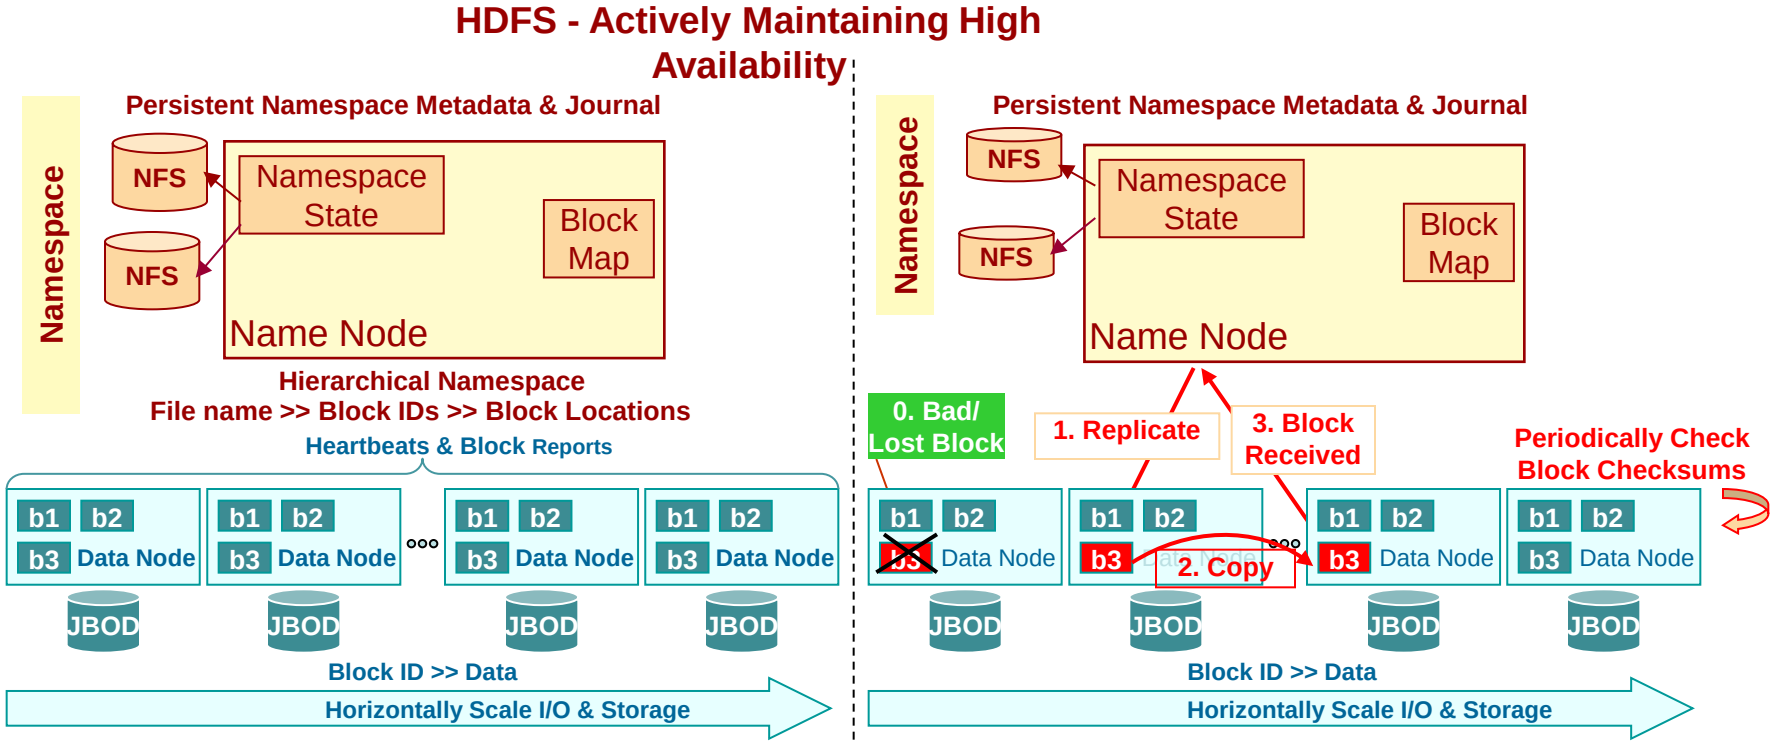

HDFS

- Immutable File System – Read, Write, Sync/Flush – No random writes
- Storage Server used for Computation – Move Computation to Data
- Fault Tolerant & Easy Management – Built In Redundancy, Tolerates Disk & Node Failure, Auto-Managing addition/removal of nodes, One operator/8K nodes
- Not a SAN but high bandwidth network access to data via Ethernet
- Used typically to Solve problems not feasible with traditional systems: Large Storage Capacity >100PB raw, Large IO/computational BW >4K node/cluster, scale by adding commodity HW, Cost ~$1.5/GB incl. MR cluster

Big Data - Market Requirements

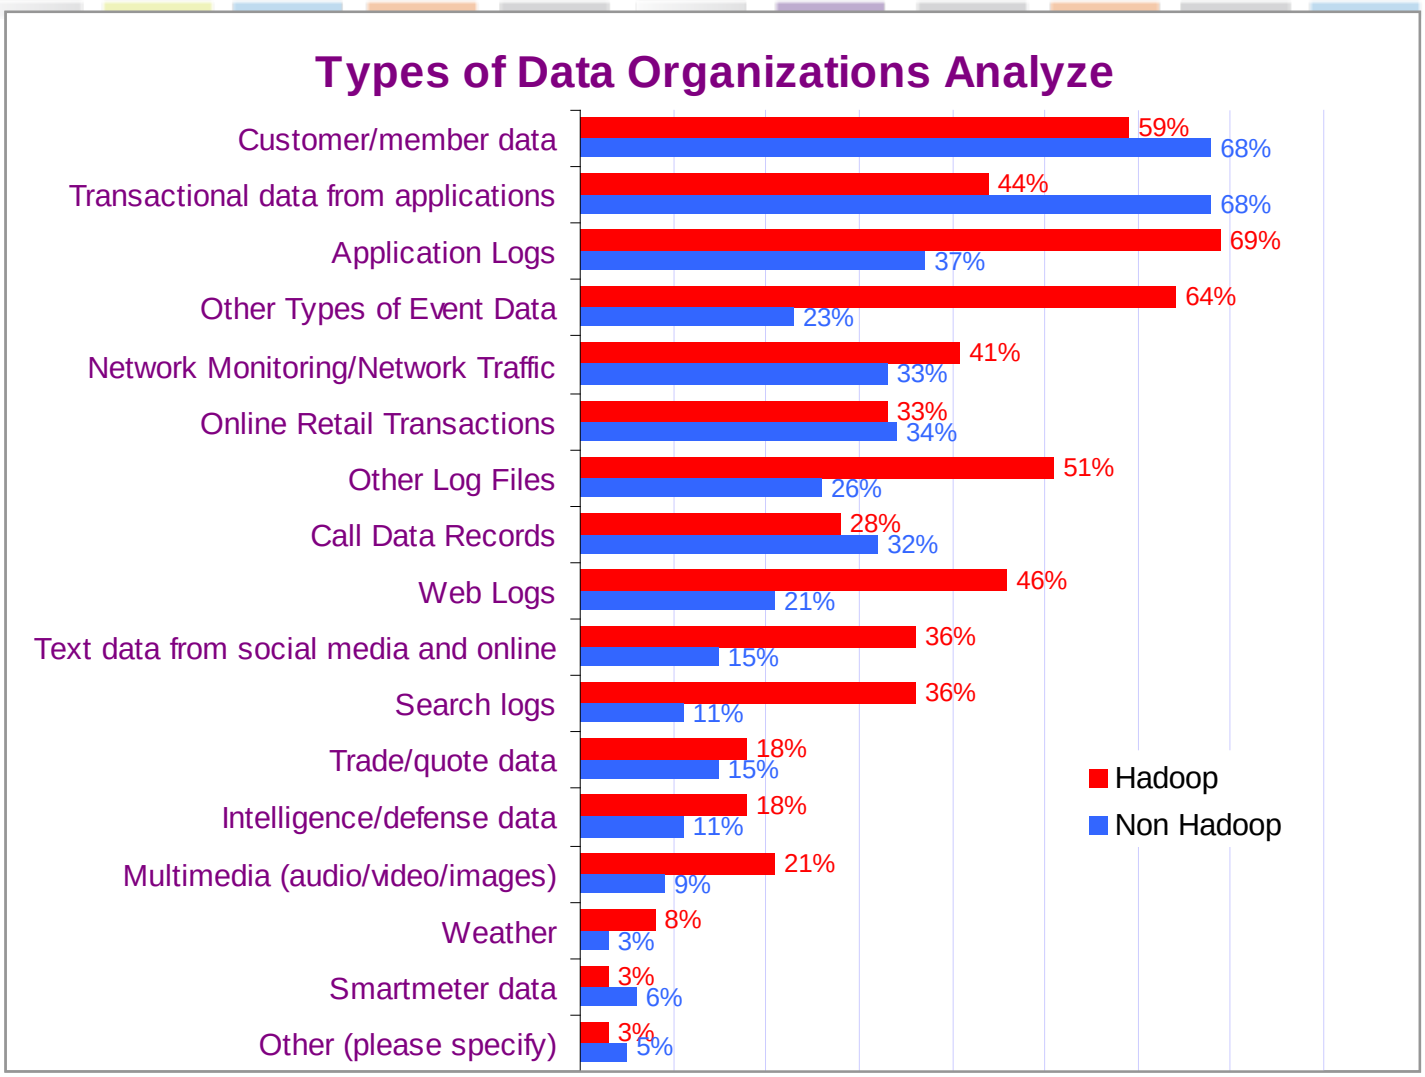

Big Data Savings with Open Source

Legacy BI vs. Open Source Big Data Analytics

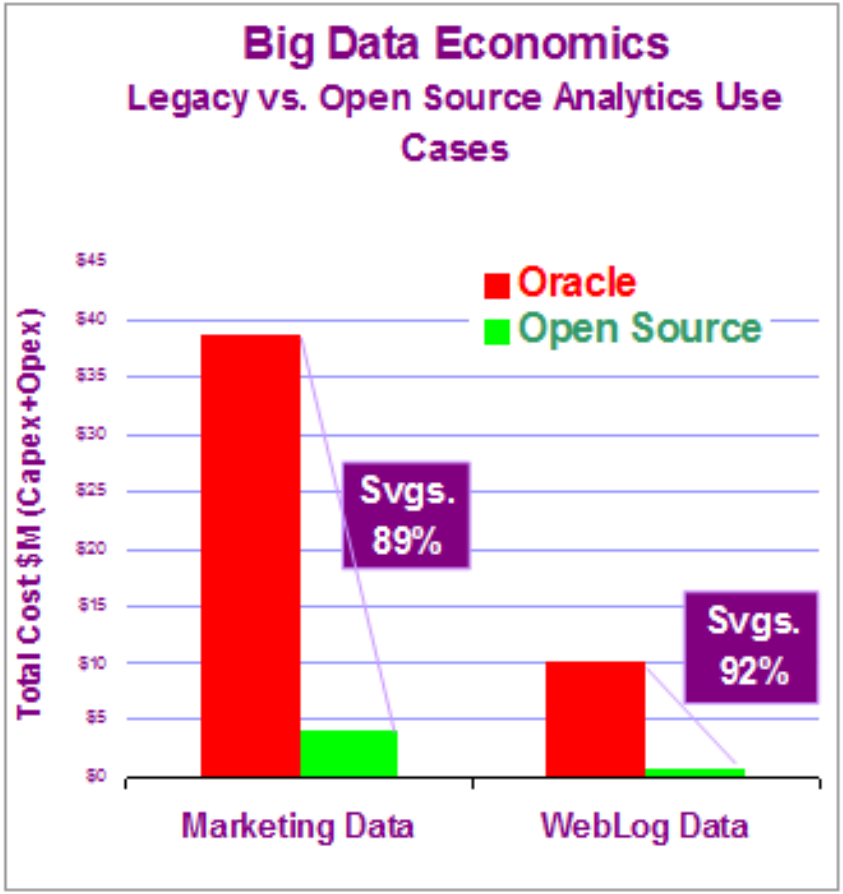

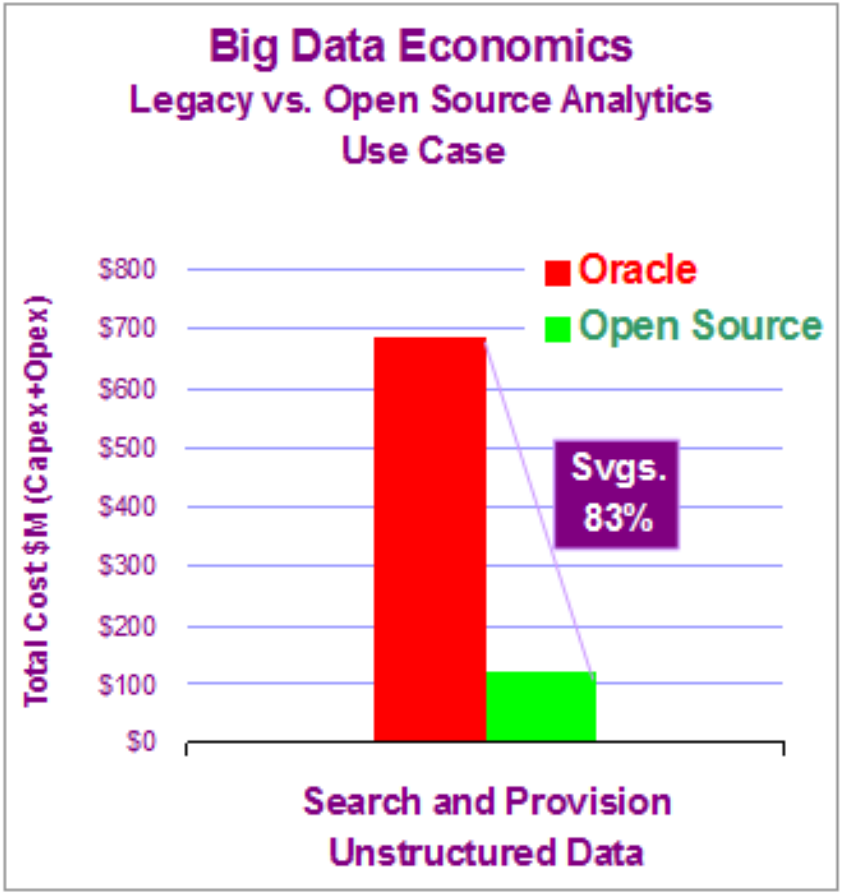

Targets for Object Storage

SNIA

Collaboration Information

Instantly available
On Boot Ups

Entertainment-Data Repository

Most Accessed Videos
Very Read Intensive

Commercial Mktg. Data

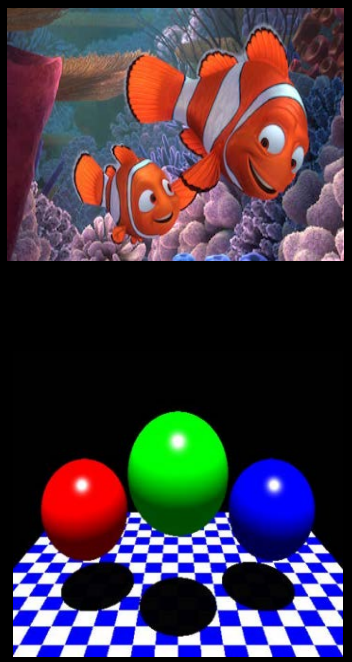

Rendered (Texture & Polygons)
Stored Data-Very Read Intensive

Medical Data Diagnostics

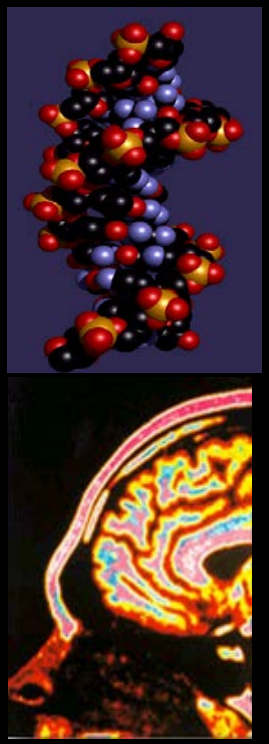

Historical Data Decision Support

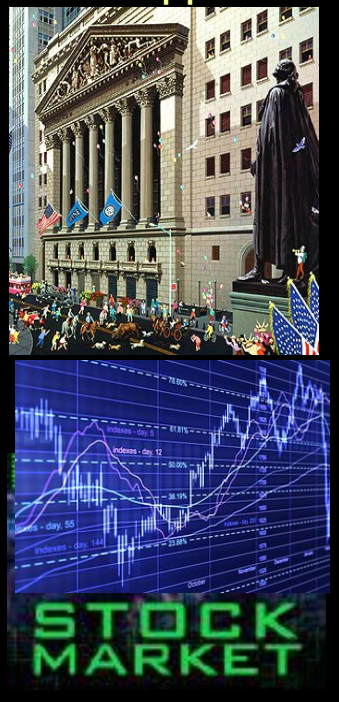

Wall St/Financials/
Federal/State/CityData
Geo-Resources Maps

Data: IMEX Research & Panasas

Object Storage - Goals

Tiered Storage

- Optimize storage costs and service levels by using cloud as a new tier of storage

Distributed Content and Collaboration

- Consolidate, classify and distribute content for re-use by employees, or end users

Offsite Data Protection

- Preservation of valuable data
- Compliance
- Legal hold

Object Storage Goals

- **Reduce cost and complexity**
  of unstructured data storage infrastructure.
- **Auto-manage unstructured data**
  documents, emails, presentations, audio files, messages..
- **WORM**
  Upload once, access from any geographic region
- **Unlimited File Size.**
  Upload files of virtually any size
- **Data Consistency:**
  Changes made to a file in one location immediately available at all locations
- **True Global Namespace:**
  Leverage auto geo-positional routing for increased performance and consistency.
- **Self-Healing Data Integrity**
  Continuous hash check , if problem detected, system repair the file automatically.
- **Self-Provisioning:**
  Provision your own storage as needed.

Why Object Storage

- **Object-based storage supports shared tenancy.**

  Objects have their own custom metadata functioning as fairly autonomous data instances - can carry access policies to support multi-tenancy - a big aspect of the relevance of object storage to public cloud infrastructures.

- **Object storage similar to low-cost, grid like architectures.**

  Object storage is well-suited to run on a loosely federated set of low-cost, industry-standard servers acting as nodes in a cluster. Cluster easy to expand by adding nodes - automated data layout rebalance the system to incorporate them. This drives the low acquisition cost and lower OPEX..

Why Object Storage

- **Massive scalability**

Within specified policies system can scale nearly indefinitely in number of files or objects and capacity of system..

- **Large metadata > custom control over data.**

Object Storage's expanded Metadata automatically and consistently tackles traditional storage challenges such as tiering, security, migration, redundancy, and deletion vs File Storage's only file type, creation date, and last-accessed date.

- **Data is inherently Immutable**

Saved object data is tagged with a unique ID, guaranteeing the immutability of that object. To modify an object, you create a new, and keep versioning history - a natural fit for archiving and records management guaranteeing that data has not been modified or tampered with.

Why Object Storage for Big Data

Object Storage Characteristic

- **Object is API Driven Storage**
  - Faster Development,
  - 10x reduction in Code
  - REST APIs speed Development
- **Location Transparency**
  - One Storage System/Access Point across
  - Multi-global apps
- **Self Managed Storage**
  - No LNS, Never recode when syst Change
- **No limit on number of Objects**
- **Object size can be up to 5 TB**
- **Central Data storage - All systems**
- **High Bandwidth**
- **> Five 9s of Durability**
- **Versioning Lifecycle policies**
- **Integration with industry standards**
- **Typical Suppliers:**
  Open Source/Swift, AWS/S3, Glacier, EMR,Redshift; Basho, Caringo, Cleversafe, Cloudian, Coho Data, DDN, EMC, Exablox, HDS, Huawei, IBM, NEC, NetApp, RedHat, Scality, Storiant, SwiftStack, Tarmin

Linear Scalability
Scales to billions of objects

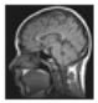

Support for large files
Object sizes are in TBs

Web friendly
Firewall friendly, http, REST accessibility

No Locking
No lock on write or create operations

Metadata and extensibility
Objects can be extended to multiple policies (Immutability, retention,

Geo-scale
Geo-replicated and distributed

Object Storage – Global Reach

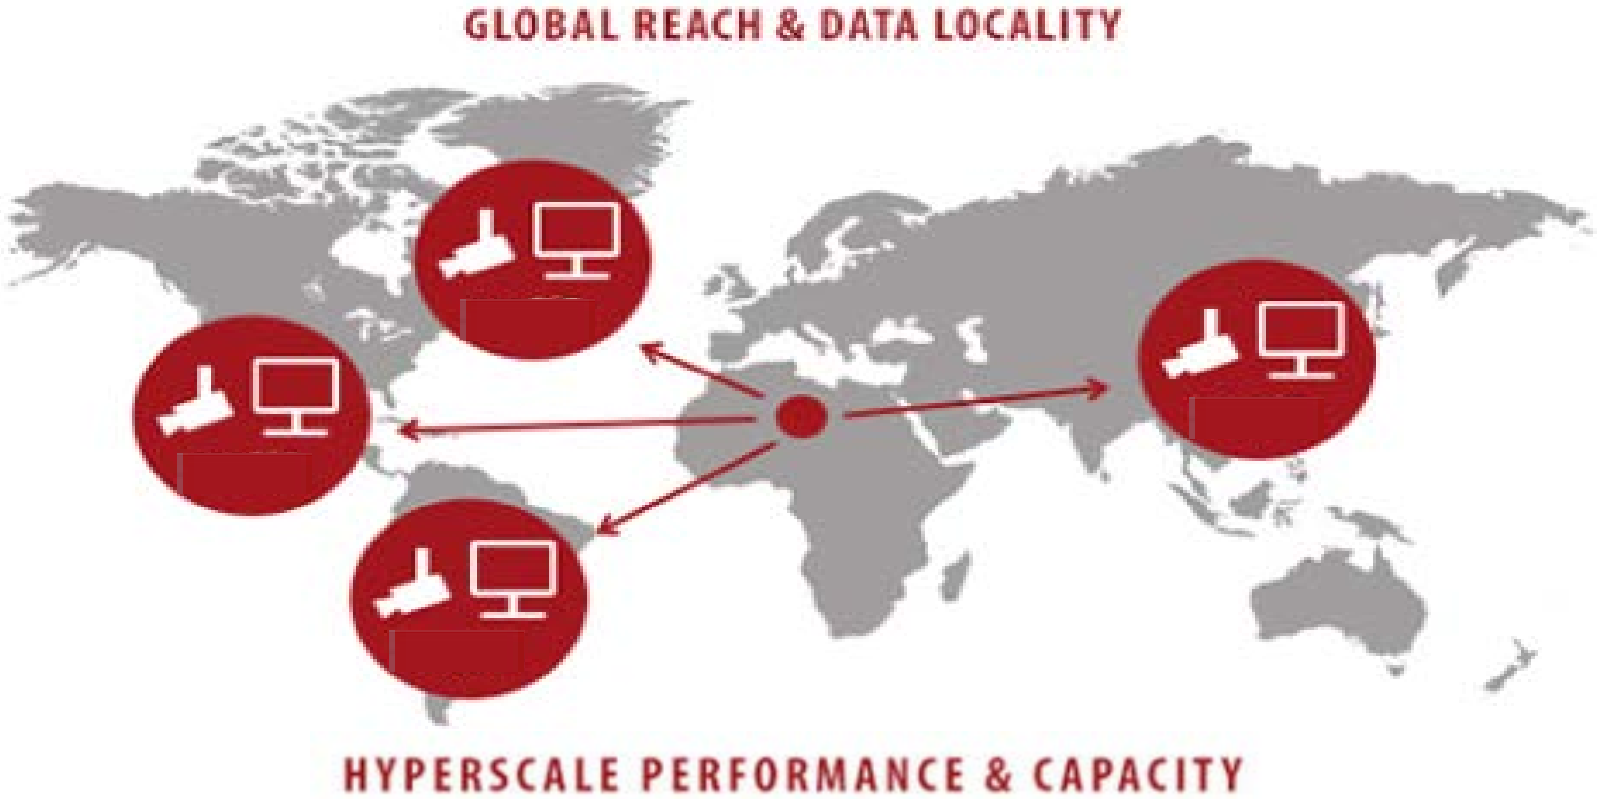

Replicate & Collaborate (Ingest, Access & Update Data at local speeds across multiple locations

Data Integrity of Big Data Storage

Issues

Preventing Data Loss
Maintaining ‘Always On’ System
Scaling Storage Capacity while Controlling Costs
Preventing Data from unauthorized Access

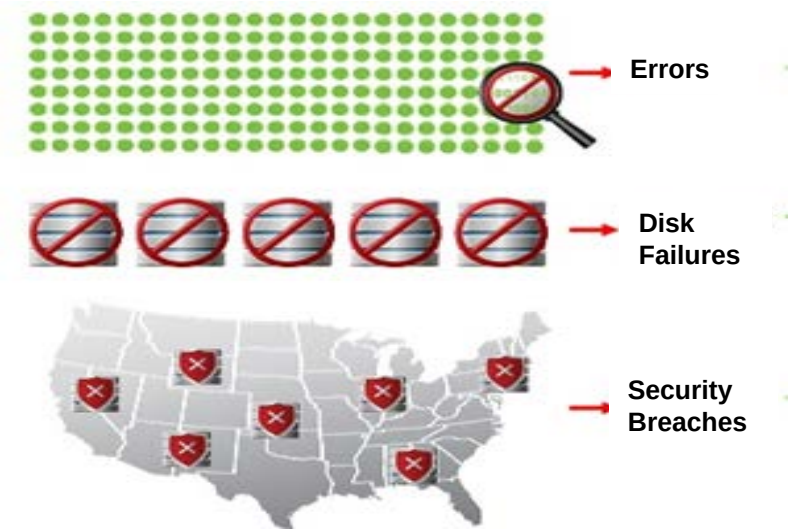

Solutions

Choose Object Based storage vs. File Based Systems
  To manage billions of objects
Information Dispersal Algorithm
  Using Error Codes
Preventing Data from unauthorized Access
Scaling Storage Capacity while Controlling Costs

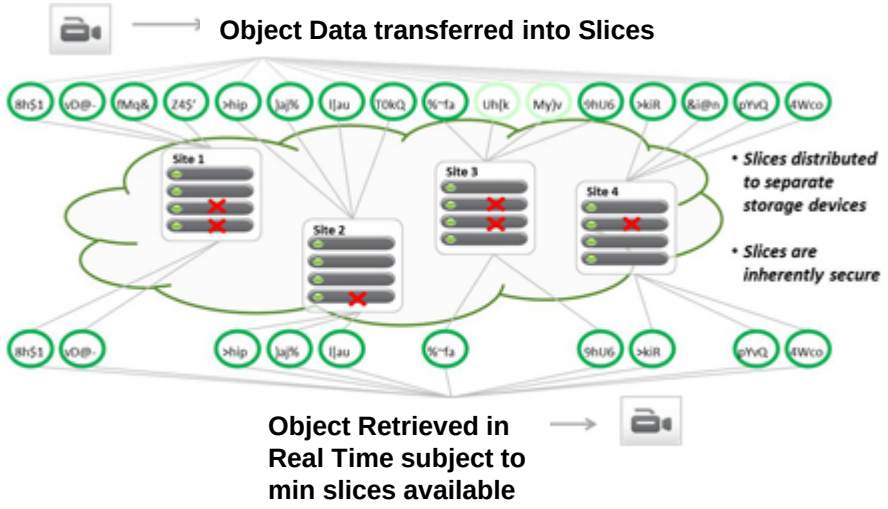

Three Types of Storage

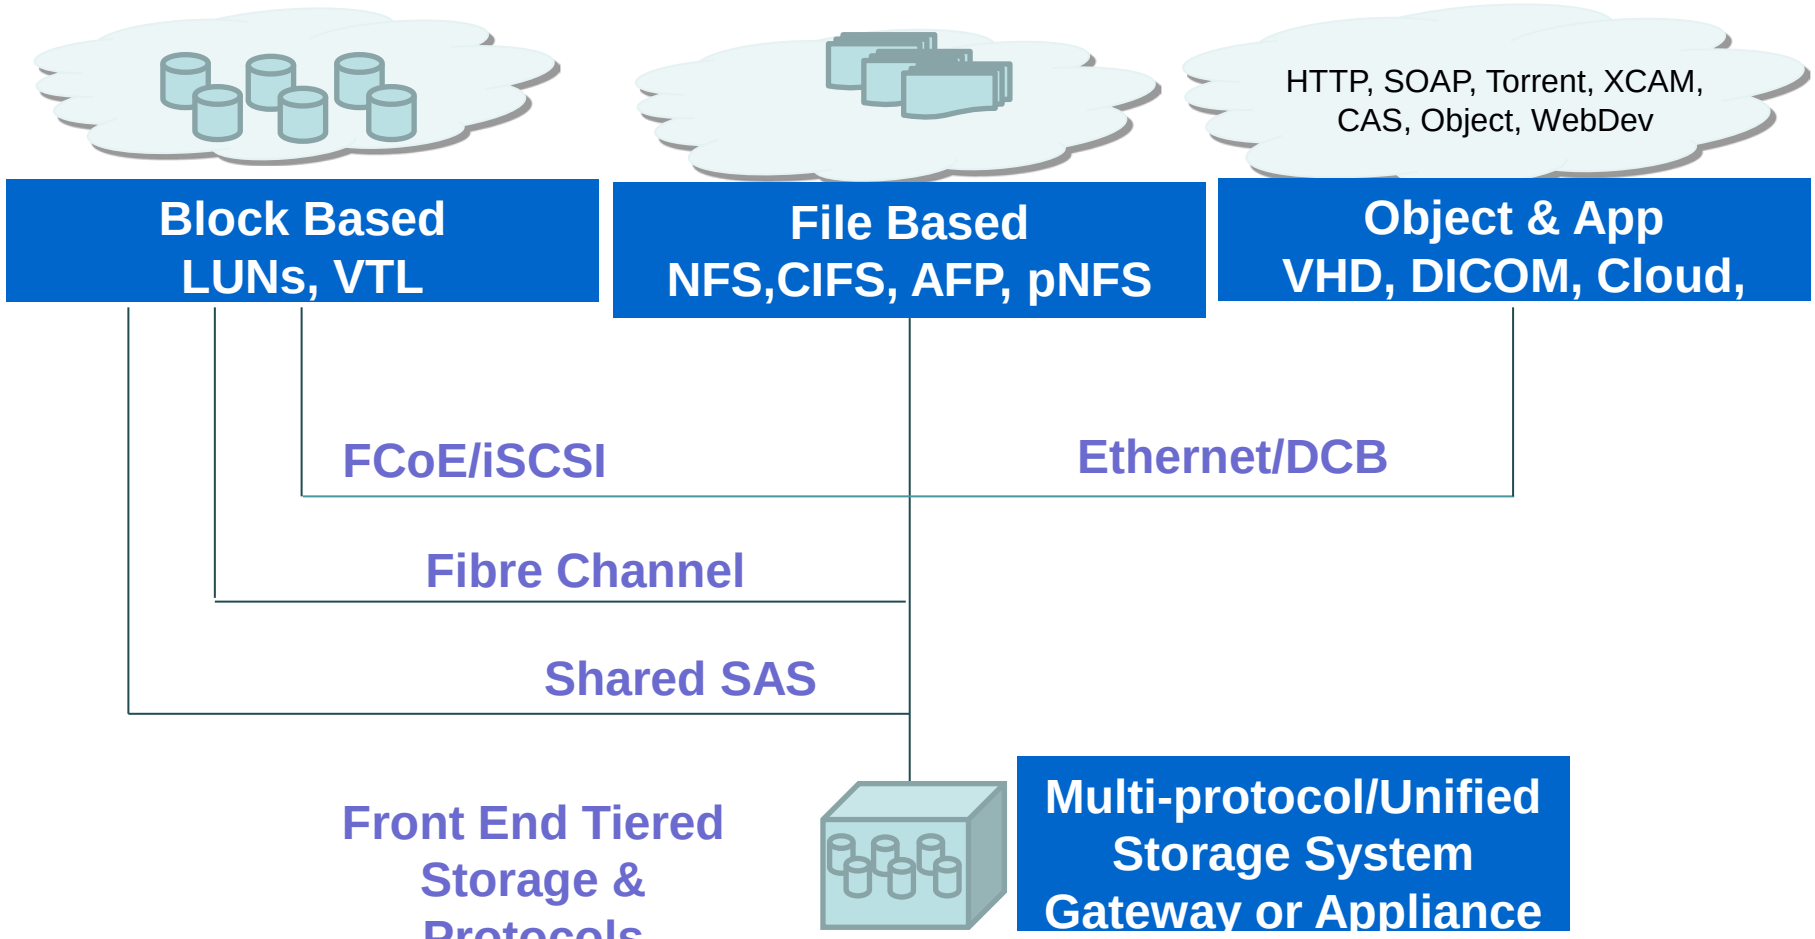

Data Storage Tiers Evolution

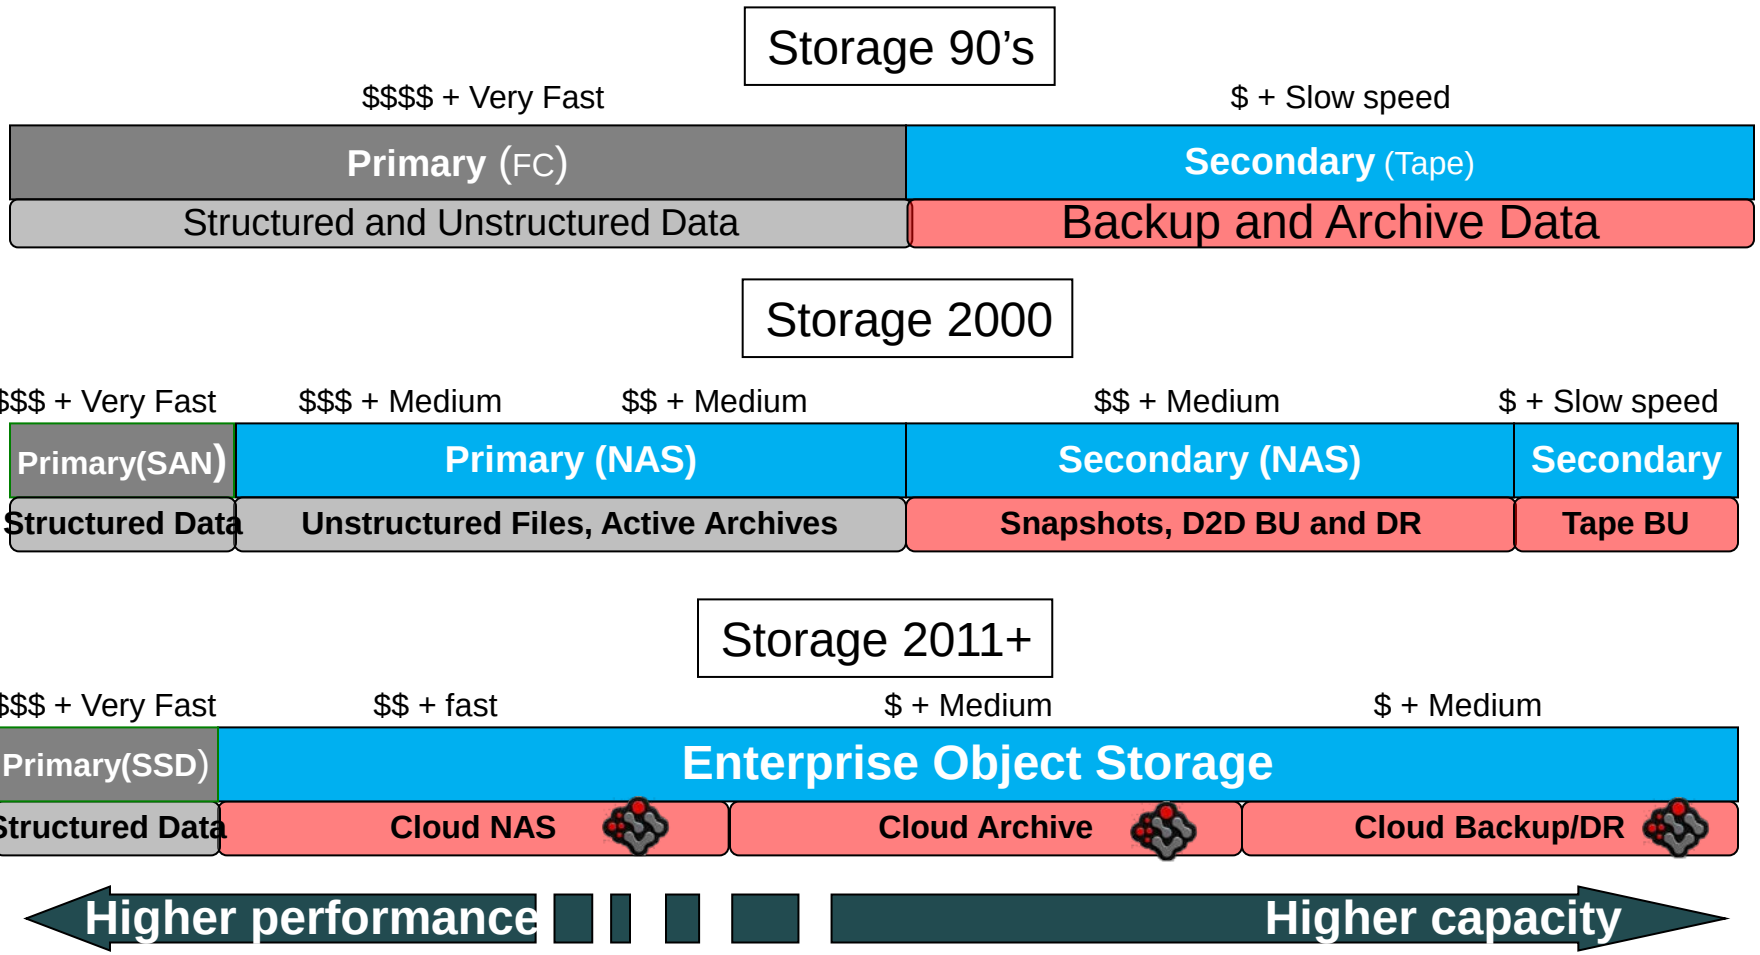

File vs Object Storage Characteristics

2012-2016 Workload Growth

Enterprise IT must adapt to support mixed workloads

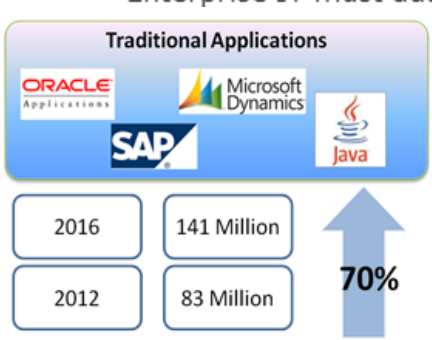

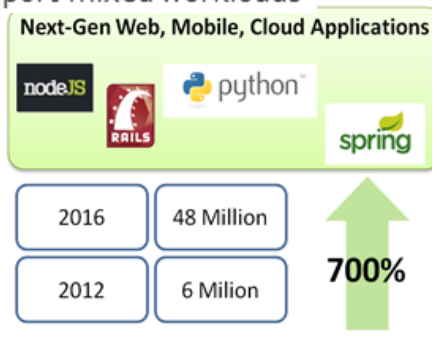

File-based Storage

- Optimized for structured/semi-structured
- CRM, ERP, Database apps, email, etc.
- Built for performance (10s of ms latency)
- Directory structure, location-dependent
- Manual IT administration & provisioning

- Writes to file require exclusive lock
- Limit on number of files in a directory
- No user meta-data
- Large files hard to seek
- File create require directory exclusivity lock
- Not Web or firewall-friendly

Object/Cloud Storage

- Optimized for unstructured
- Archiving, Web, mobile, cloud apps
- Built for scale/efficiency (100s of ms latency)
- Location-independent - Metadata & objects stored together
- Multi-tenancy - self-service access

- Object supports multiple writes, no locking
- Objects are limitless in size, 1 MB to TBs
- Objects support extensible meta-data
- Objects can be viewed with no limitation
- No locking required to create files
- HTTP, REST-based access, Web & firewall-friendly

Hierarchical vs Flat Storage Schema

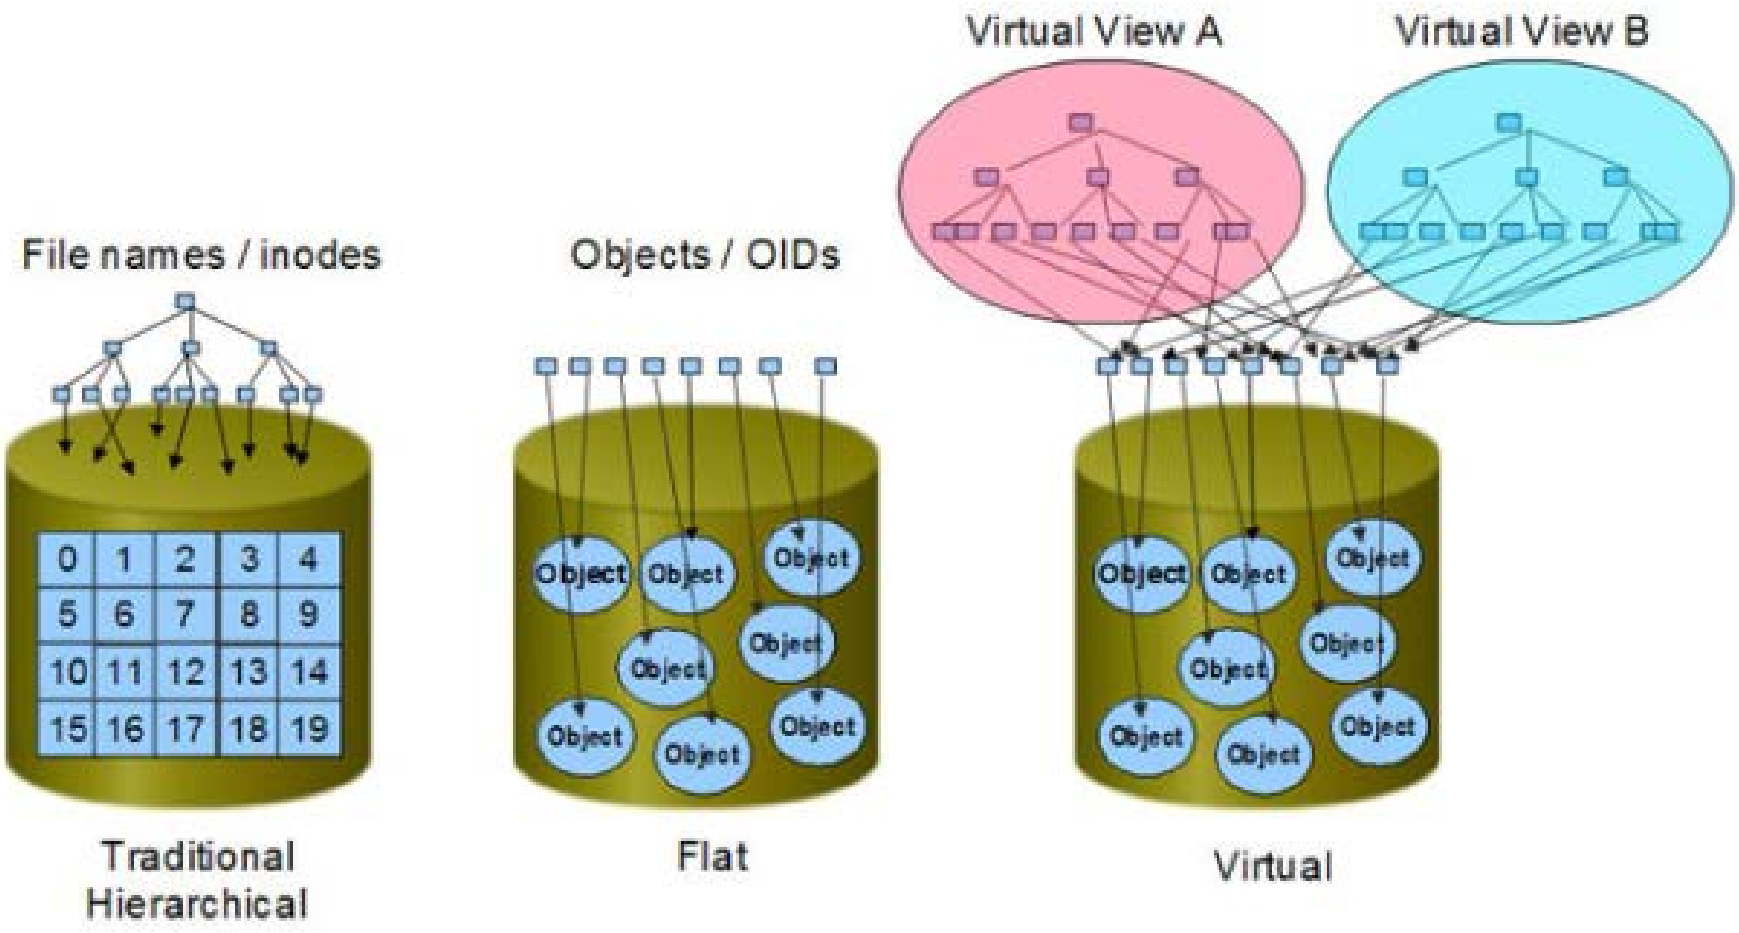

Metadata: Key to Object Storage Adoption

File

File Name: CATSCANJQSMITH
Created By: Technician 1
Created On: 01-01-2001
File Type: .DICOM

VS.

Object

Object ID: 12345
File Type: .DICOM
Patient Name: John Q. Smith
Patient ID: 555-55-5555
Procedure Date: 01—1-2001
Physician Name: Dr. Organ
Physician Notes: .WAV File
Prior 1: XYZ.DICOM
Modality: XYZ
Manufacturer: XYZ
Diagnosis: XYZ
Description: XYZ
Custom Metadata: XYZ
Custom Metadata: XYZ

Source: Dell

Object Storage Positioning

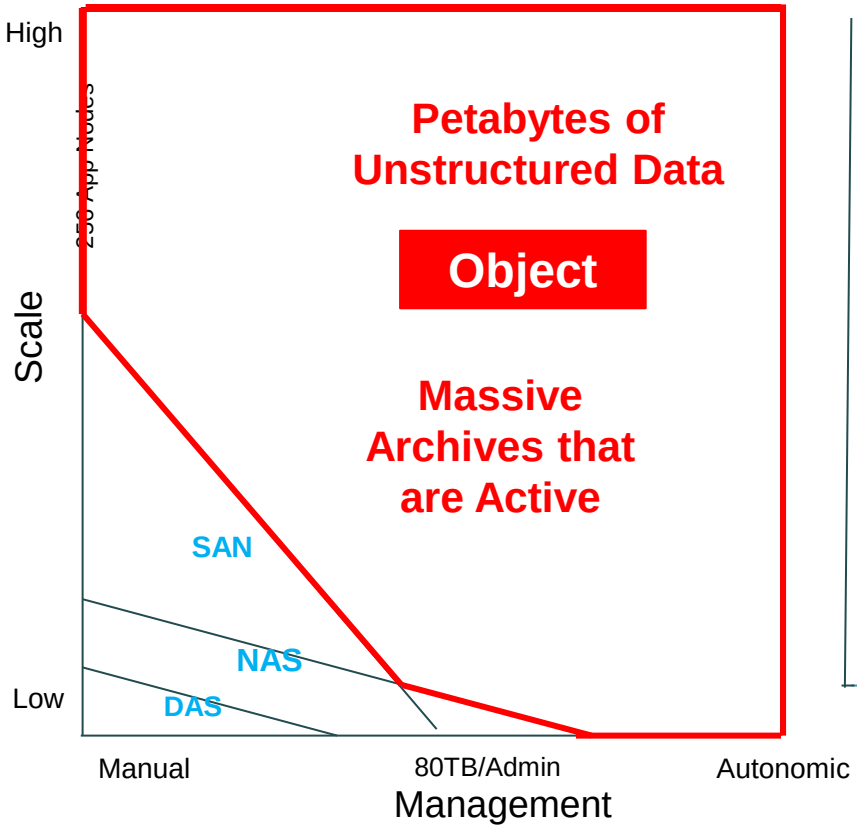

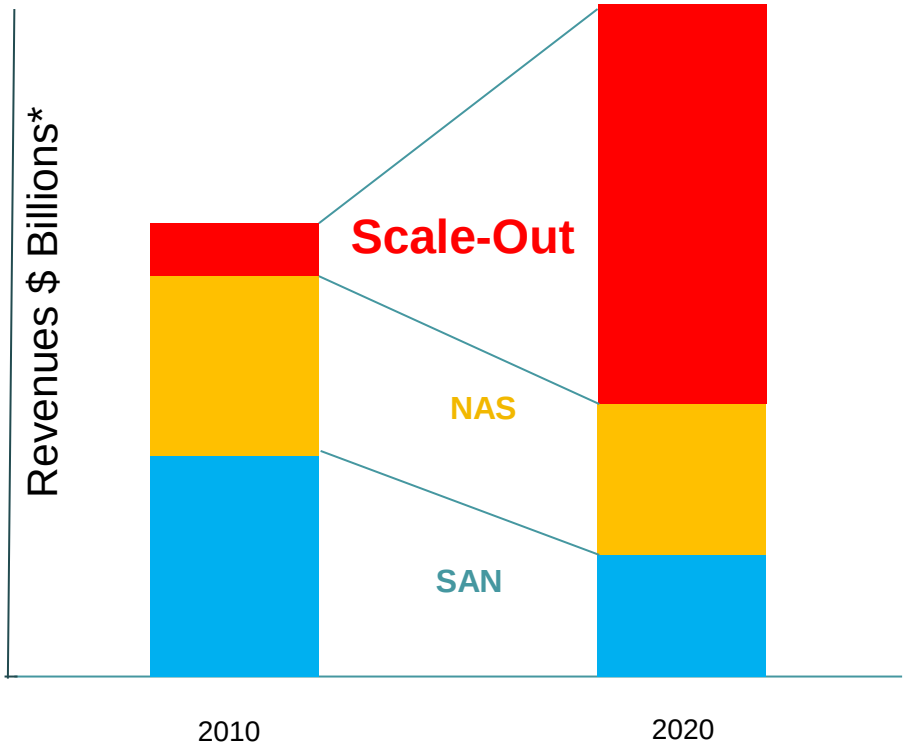

*Includes HW, SW, File System, Backup, High Availability & Management

Object Storage: Key Features

- **Object Storage:**
  - Data repository of billions of Objects
- **Key Architecture:**
  - Separation of Metadata and Data
- **Features :**
  - Flat Namespace
  - Infinite Scalability
  - Elasticity
  - Cost-efficiency
  - Data durability
  - Distributed System
  - No Single point of failure

App Driven SDS = Metadata Driven Object Storage

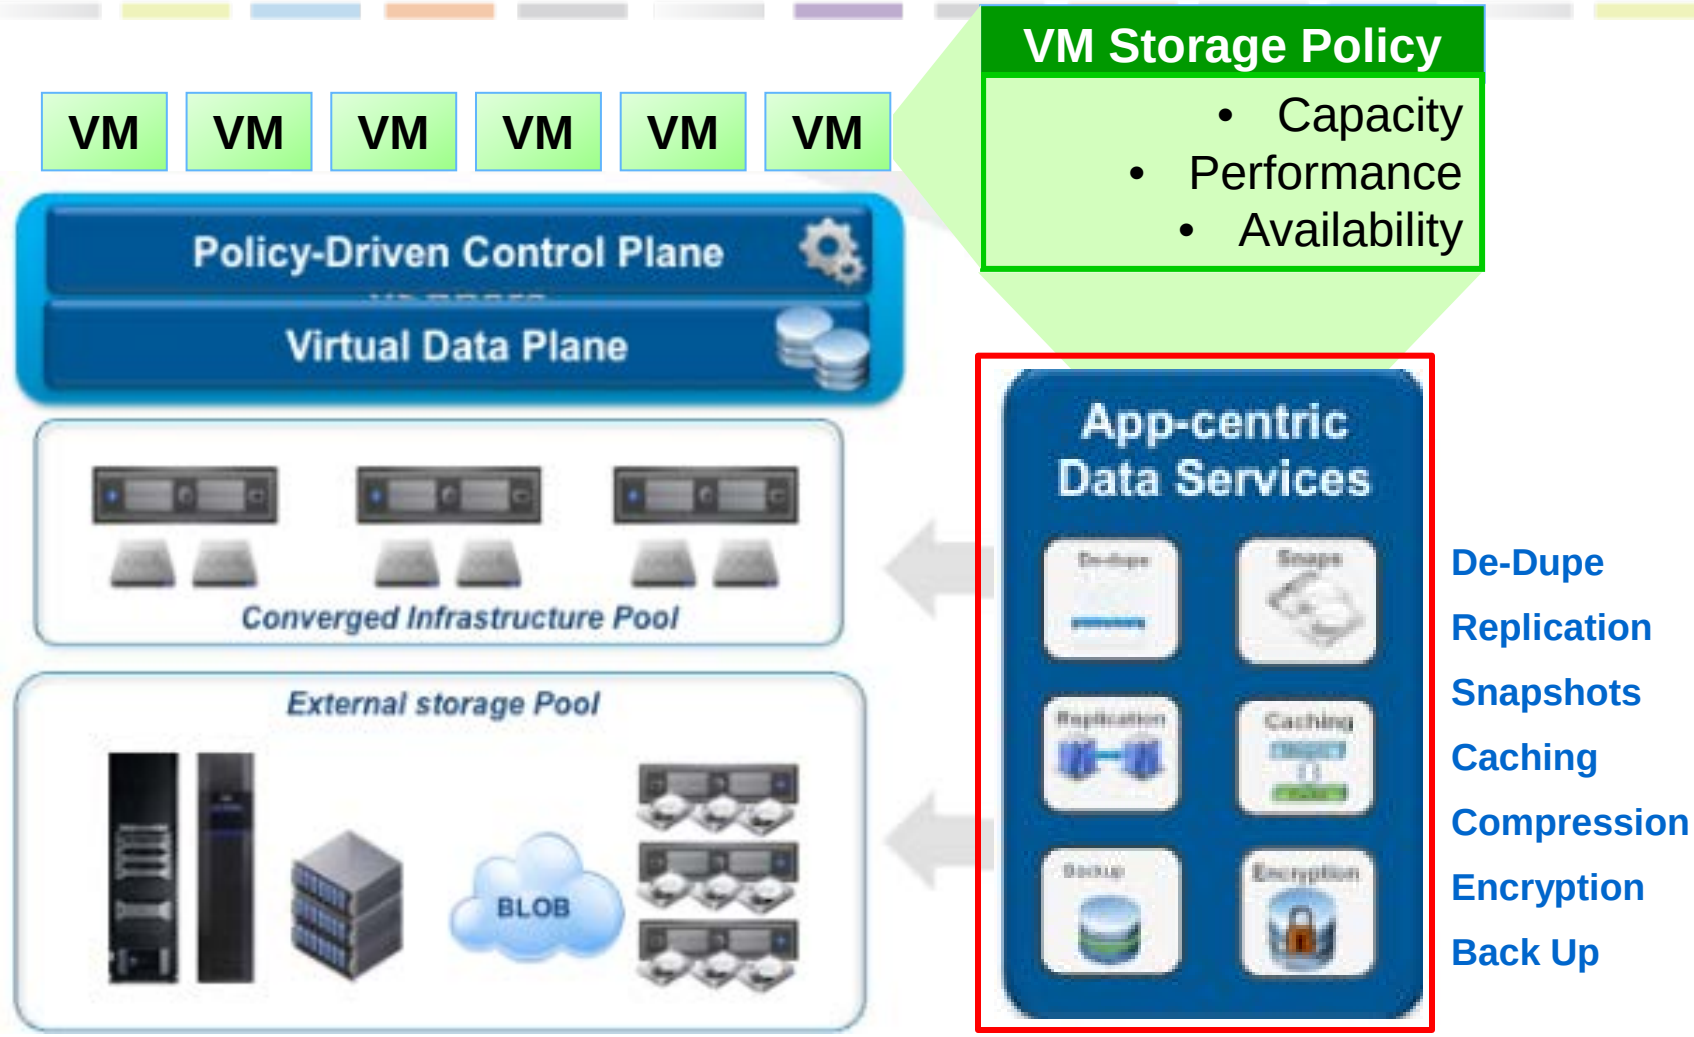

Object Storage - Implementation

In implementing object storage:

- A **unique identifier** assigned to each piece of data
- Identifier allows servers or users to get access to Object **without any knowledge of where data is physically stored**.
- Data can be broken into chunks, stored anywhere worldwide, can move within a system and have multiple copies.
- **Object tag** ensures that **right piece of data is returned** and **verifies integrity of that data**.

Open Stack Object Storage

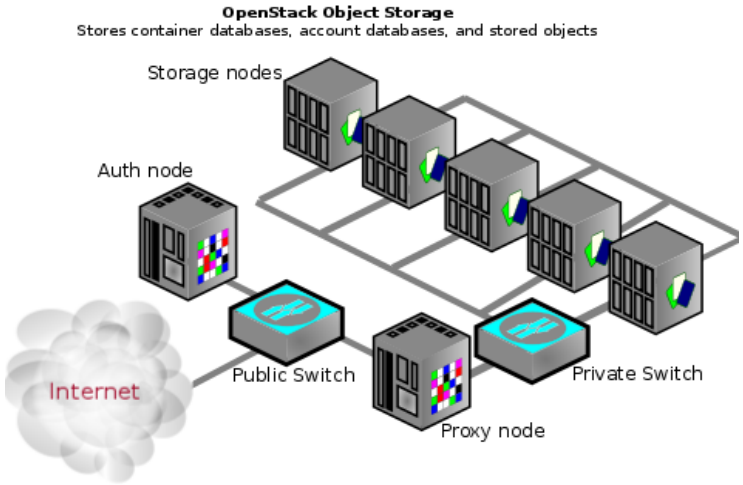

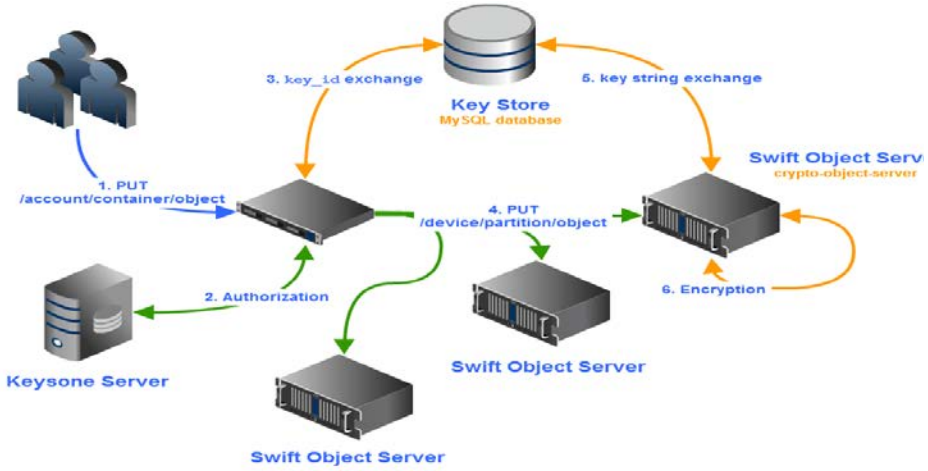

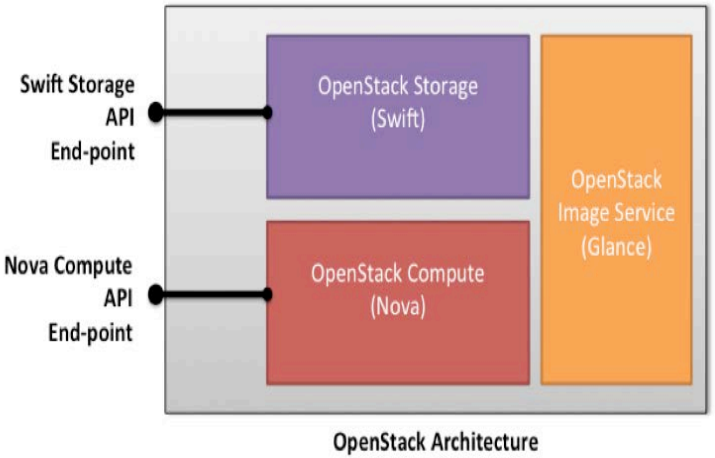

OpenStack Architecture

Centralized Storage using Public/Private Cloud

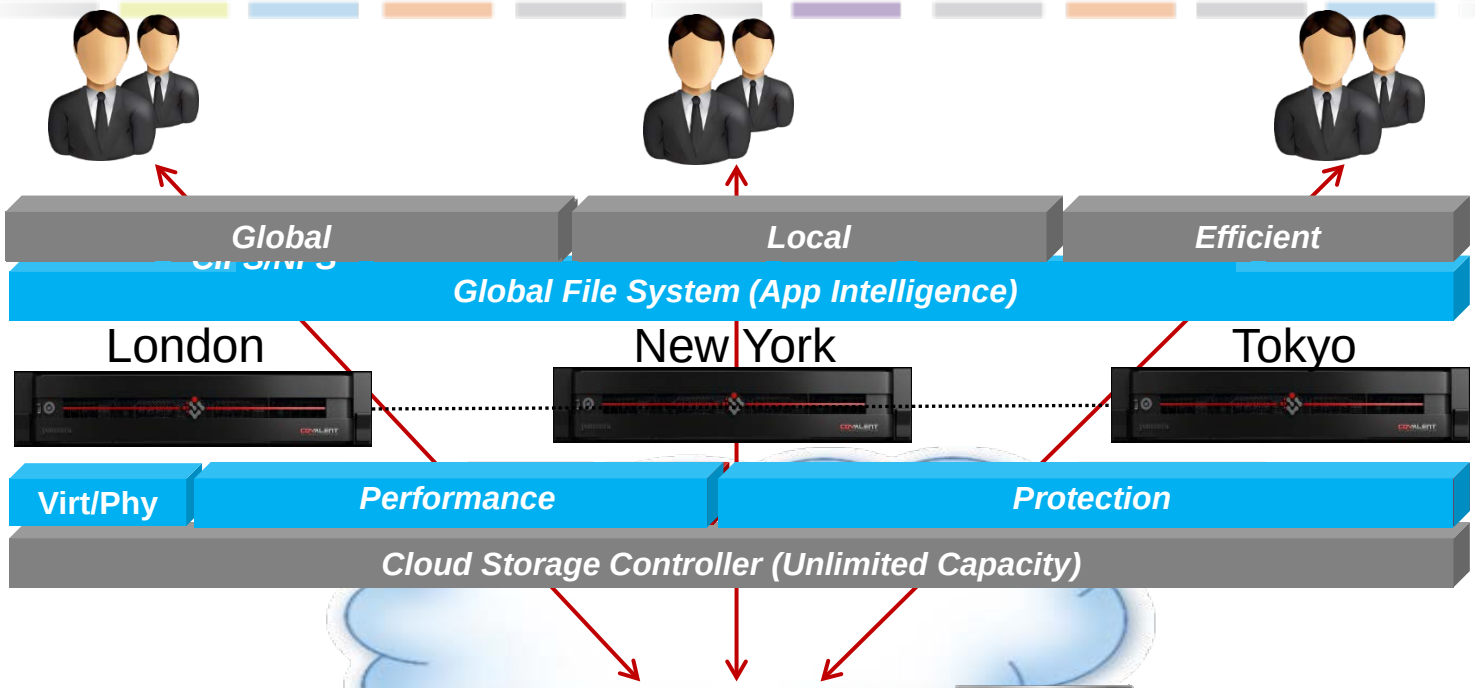

Google
amazon web services
Rackspace
F5 Networks
DreamHost
Riverbed
IBM Cloud
HDS
Gluster
AMPLIDATA
openstack
panzura
SCALITY
Cleversafe
StorSimple
Nasuni
TwinStrata
nirvanix
Data Direct Networks
Cloud Services
EMC²
Dell
NetApp/ByCast

Object Storage: Key for Big Data

Object Storage Characteristic

- **Object is API Driven Storage**
  - Faster Development,REST APIs speeds Dev.,
  - 10x reduction in Code,
- **Location Transparency**
  - Access Point across Multi-global apps
- **Self Managed Storage**
  - No LNS, Never recode when syst Change
- **No limit on number of Objects**
- **Object size can be upto 5 TB**
- **Central Data storage - All systems**
- **High Bandwidth**
- **> Five 9s of Durability**
- **Versioning Lifecycle policies**
- **Integration with industry standards**

Typical Suppliers:

Open Source (Swift), AWS (S3, Glacier, EMR, Redshift)
Basho, Caringo, Cleversafe, Cloudian, Coho Data, DDN, EMC, Exablox, HDS, Huawei, IBM, NEC, NetApp, RedHat, Scality, Storiant, SwiftStack, Tarmin

Object Storage – In Perspective

- Object Storage is rising to meet the needs of next generation storage bringing in infinite scalability, ease of management and security in Cloud Computing and opening greater insights into Big Data by leveraging extensive attributes in Metadata associated with Object Storage while significantly lowering cost of storage using volume x86 hardware to manage both servers and storage lowering CapEx and OpEx for enterprises and cloud service providers.
- Several start-ups and major IT vendors are now endorsing it.

Attribution & Feedback

The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

Authorship History

Original Author :

Anil Vasudeva, IMEX Research.com

Additional Contributors

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***