CHEP 2015

Data Lake and Elastic Cloud by EMC Federation

EMC Japan
Nobuyasu Wakamatsu

EMC - Who we are

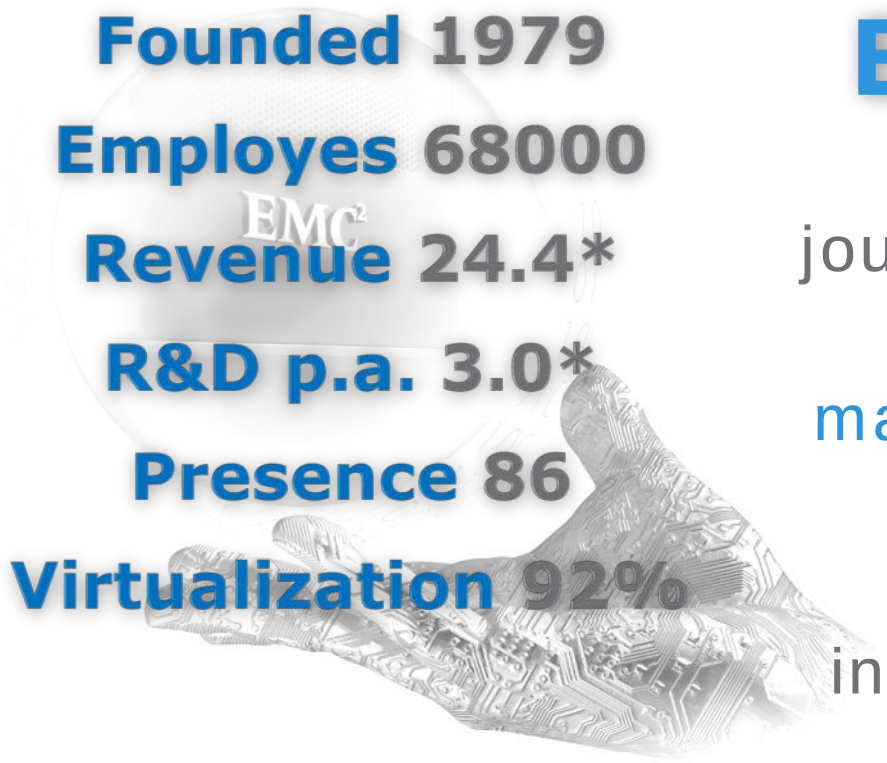

EMC's Mission

To lead customers on a journey to cloud computing by enabling them to store, manage, protect, and analyze their most valuable asset – INFORMATION – in a more agile, trusted, and cost-efficient way.

(*) Revenue and R&D spendings in 2014 in Mrd. US$

EMC ACQUISITIONS 1999 – 2014

Investment Continues

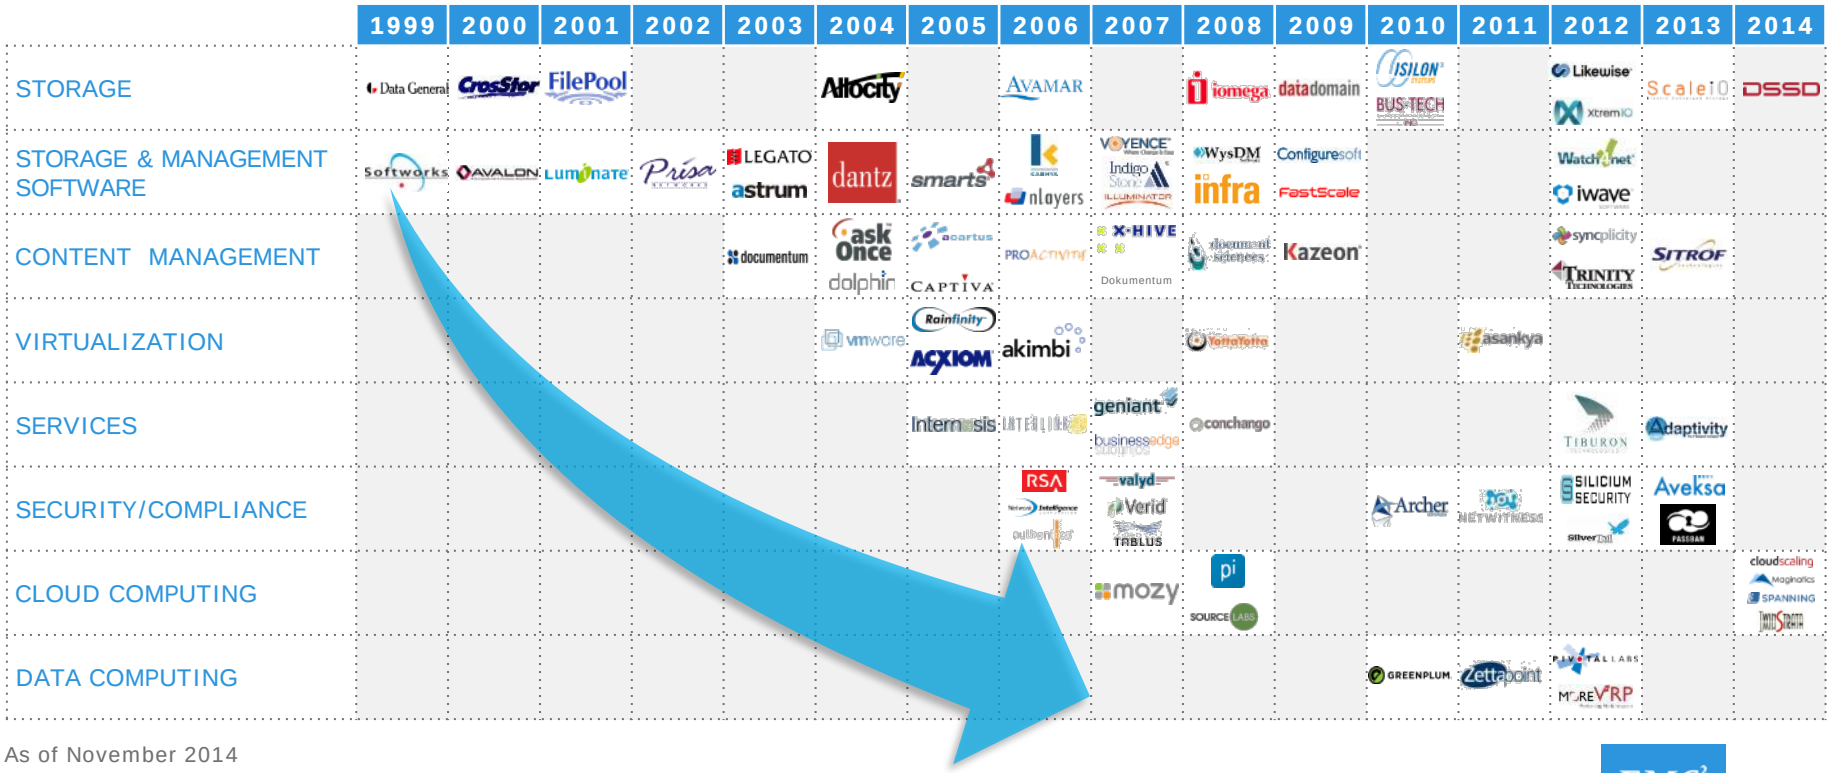

As of November 2014

A UNIQUE FEDERATION OF COMPANIES

Best Of Breed for the Software-Defined Enterprise. Architected Horizontally. Unparalleled Choice.

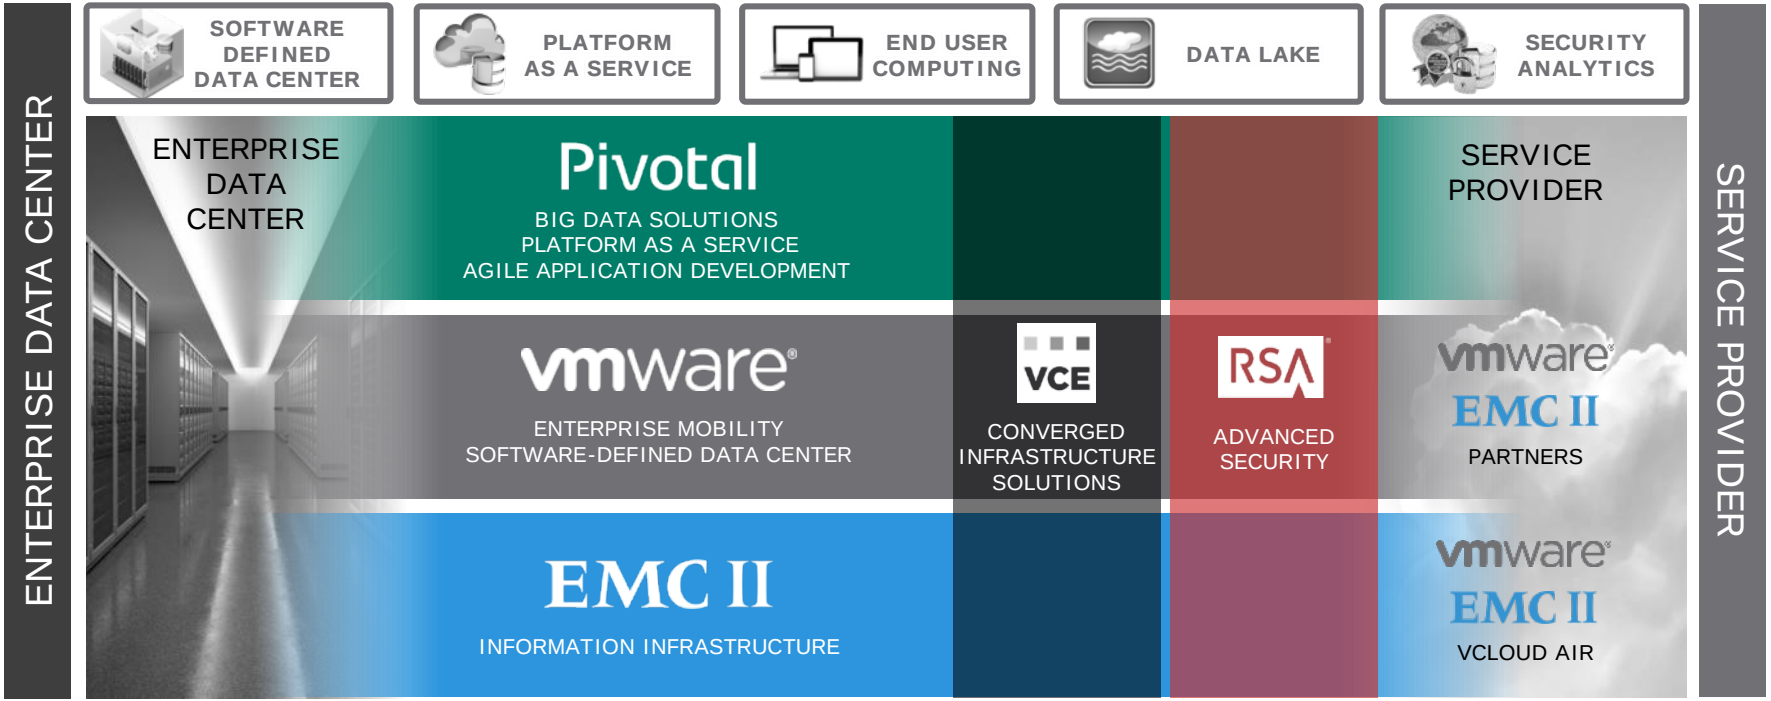

REDEFINE IT

Data Growth and Silo

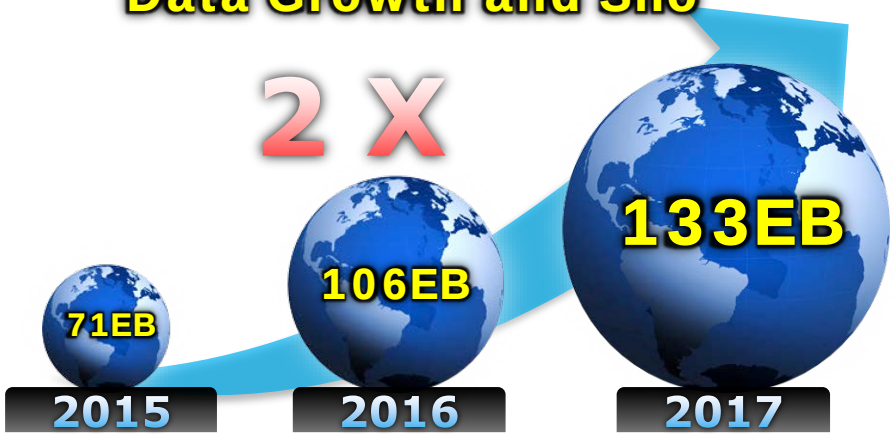

Total Capacity Shipped, WW

Source: March 2014, IDC Structured vs. Unstructured Data, The balance of power continues to shift

- ✓ **COST EFFICIENCY**
  FOR STORAGE AND DATA MANAGEMENT
- ✓ **FAST DATA ANALYSIS**

Workload Boost

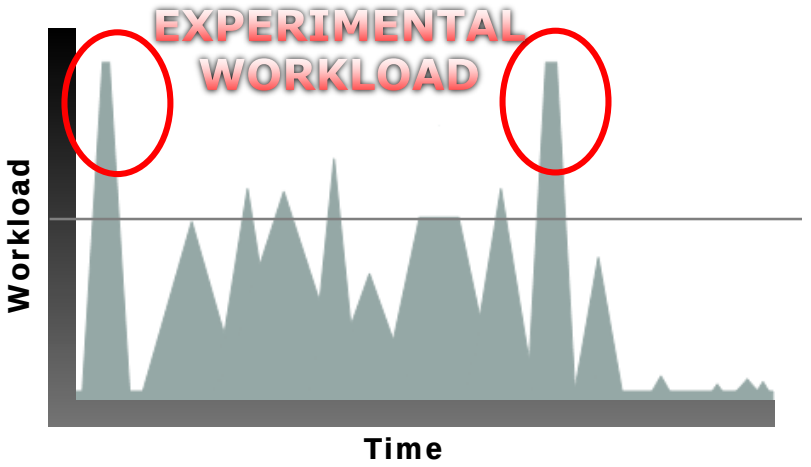

- ✓ **ELASTIC COMPUTE & STORAGE RESOURCE**

Data Growth and Silo

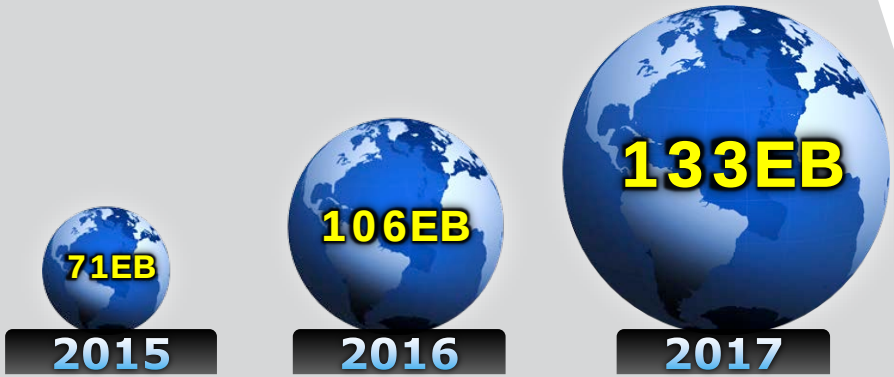

Source: March 2014, IDC Structured vs. Unstructured Data, The balance of power continues to shift

- ✓ **COST EFFICIENCY FOR STORAGE AND DATA MANAGEMENT**
- ✓ **FAST RESEARCH & ANALYSIS**

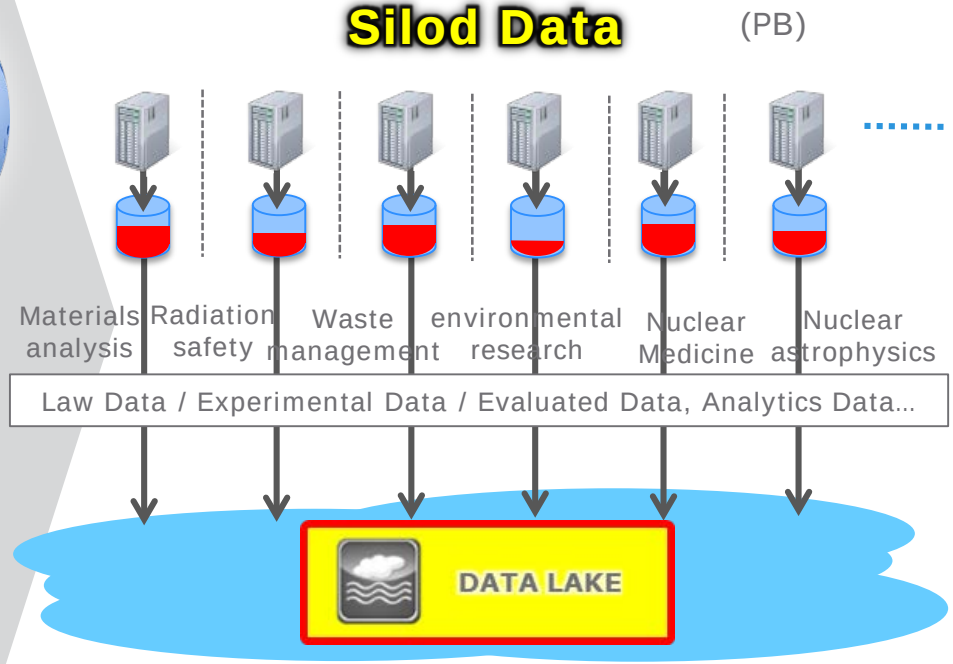

DATA LAKE TECHNOLOGY

BUT...

HADOOP

- HADOOP DISTRIBUTED FILE SYSTEM
  - Highly scalable & portable
    - Apache Open Source Specification
  - Structured and unstructured data
  - Analytics API interface standard
  - Storage hardware flexibility
  - Performance optimized for large file access

- HDFS TRADE-OFFS
  - Optimized for streaming writes; poor for random seeks
  - Write once file system
  - Hardware failure results in reduced performance
  - Specialized file system, not designed for general use

HDFS Architecture

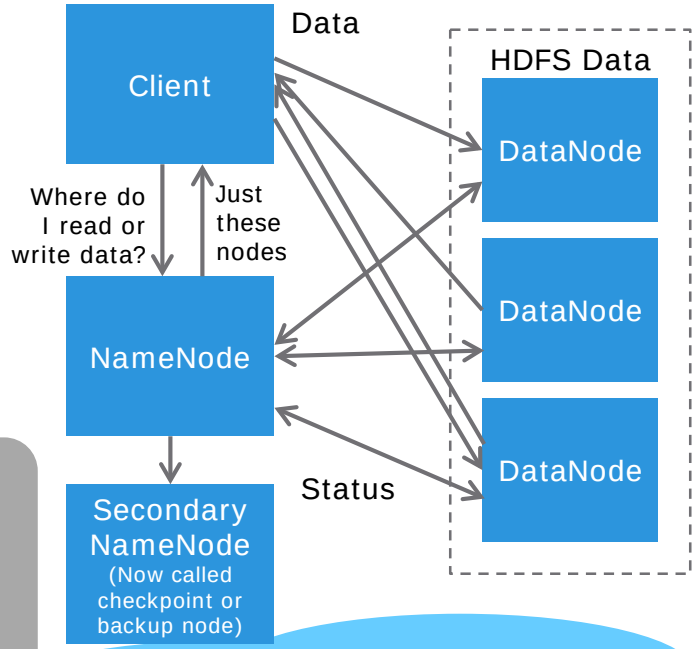

BUILDING FEDERATION DATA LAKE

EMC² Pivotal vmware

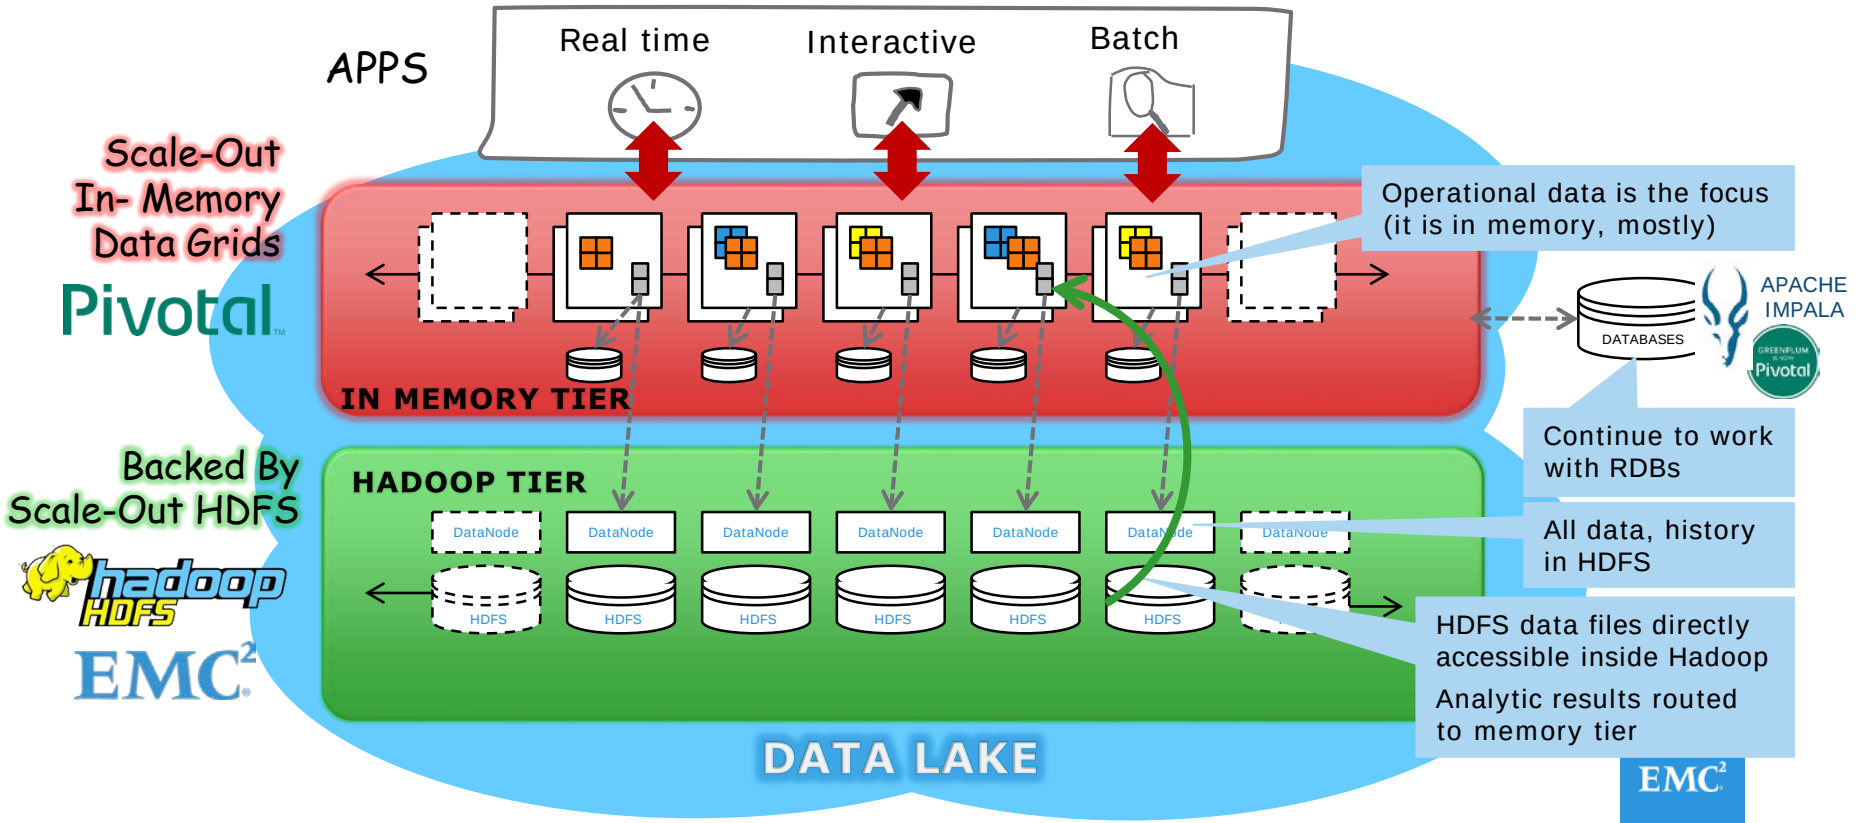

EMC²

BUILDING FEDERATION DATA LAKE

EMC² Pivotal vmware

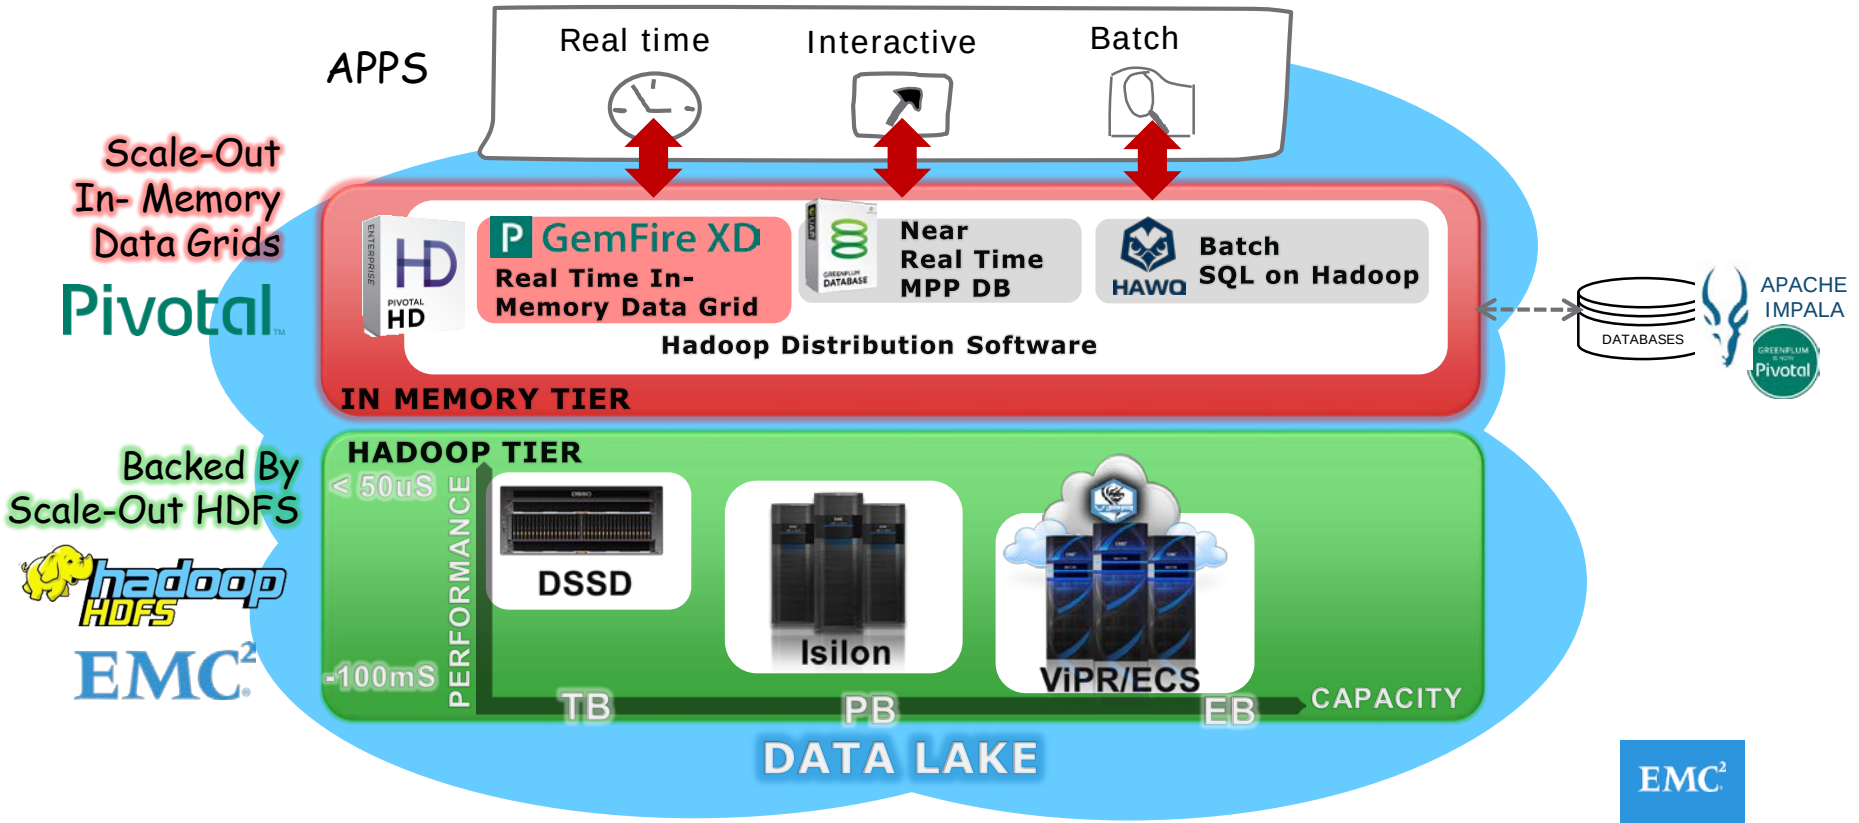

EMC Isilon

5,800+ Isilon Customers

#1 Market Leader in Hadoop Shared Storage

25% YoY Growth

600+ Analytics Customers

Gartner IDC Scale-Out NAS leader

Pivotal cloudera Ask Bigger Questions Hortonworks

HADOOP ARCHITECTURE – DAS VS ISILON

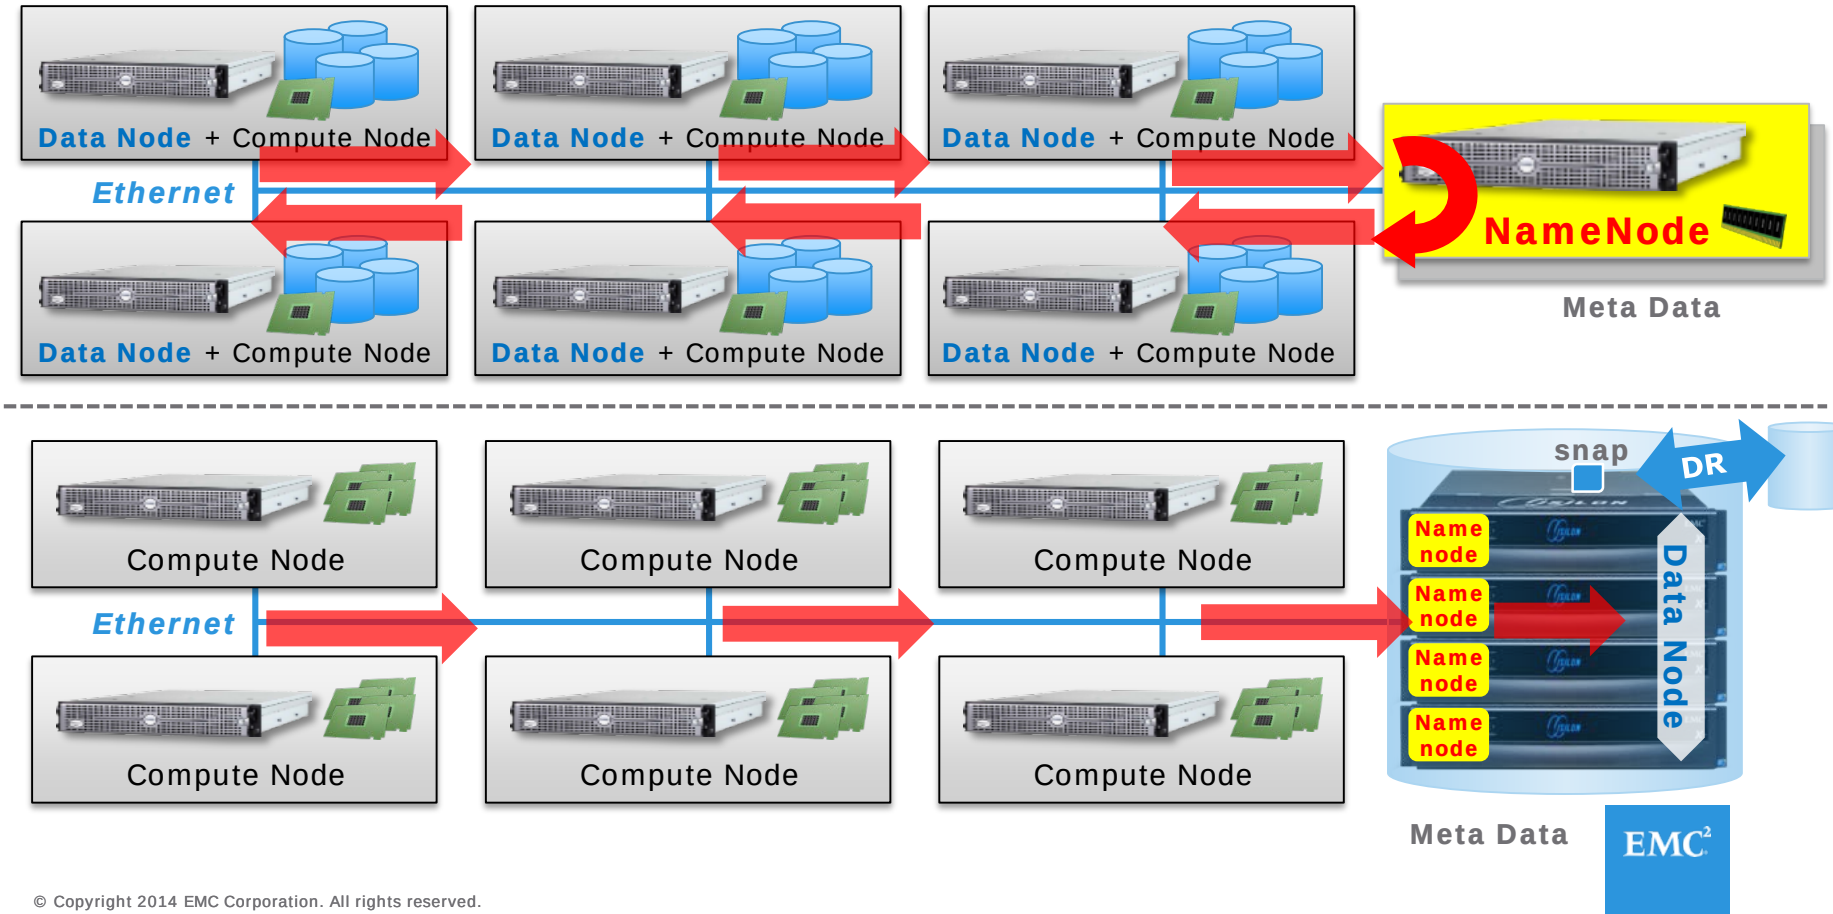

HADOOP ARCHITECTURE
– TRADITIONAL DAS

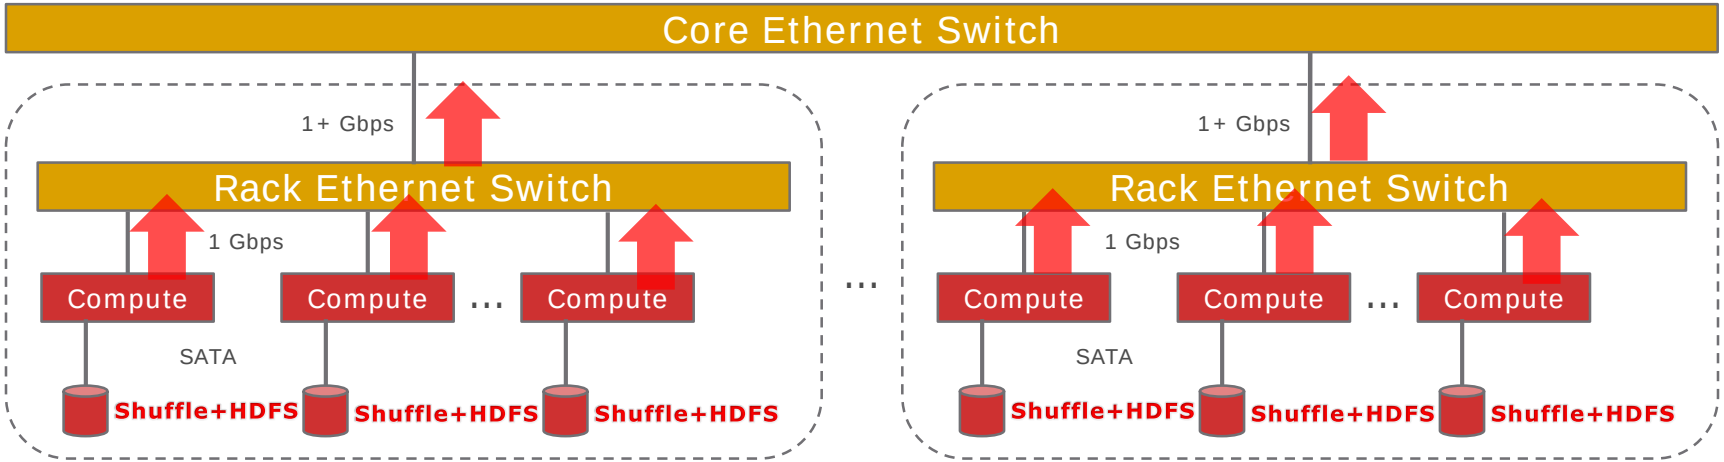

Fixed ratio of resource

20-90% of HDFS I/O go through Ethernet

Local disk usage is shared between shuffle I/O (50-75% of all I/O during terasort) and HDFS I/O

HADOOP ARCHITECTURE

– ISILON FOR HDFS

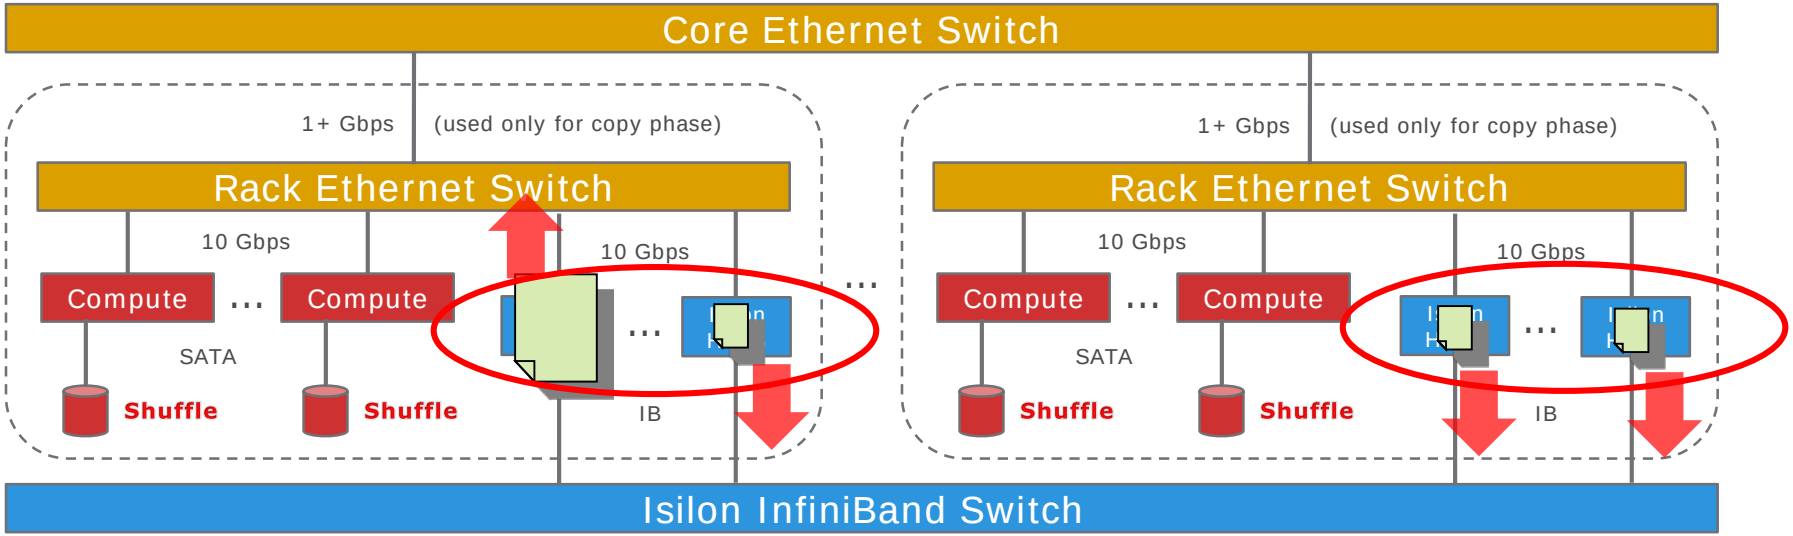

The number of compute and Isilon nodes can be adjusted independently to achieve the optimal ratio of compute and I/O bandwidth

HDFS I/O ALWAYS comes through a rack-local Isilon node which collects data blocks from all other Isilon nodes across the InfiniBand fabric

Shuffle I/O (65% of all I/O during terasort) remains on local storage. This can be flash for optimal performance.

HADOOP ARCHITECTURE

– ISILON FOR HDFS + SHUFFLE

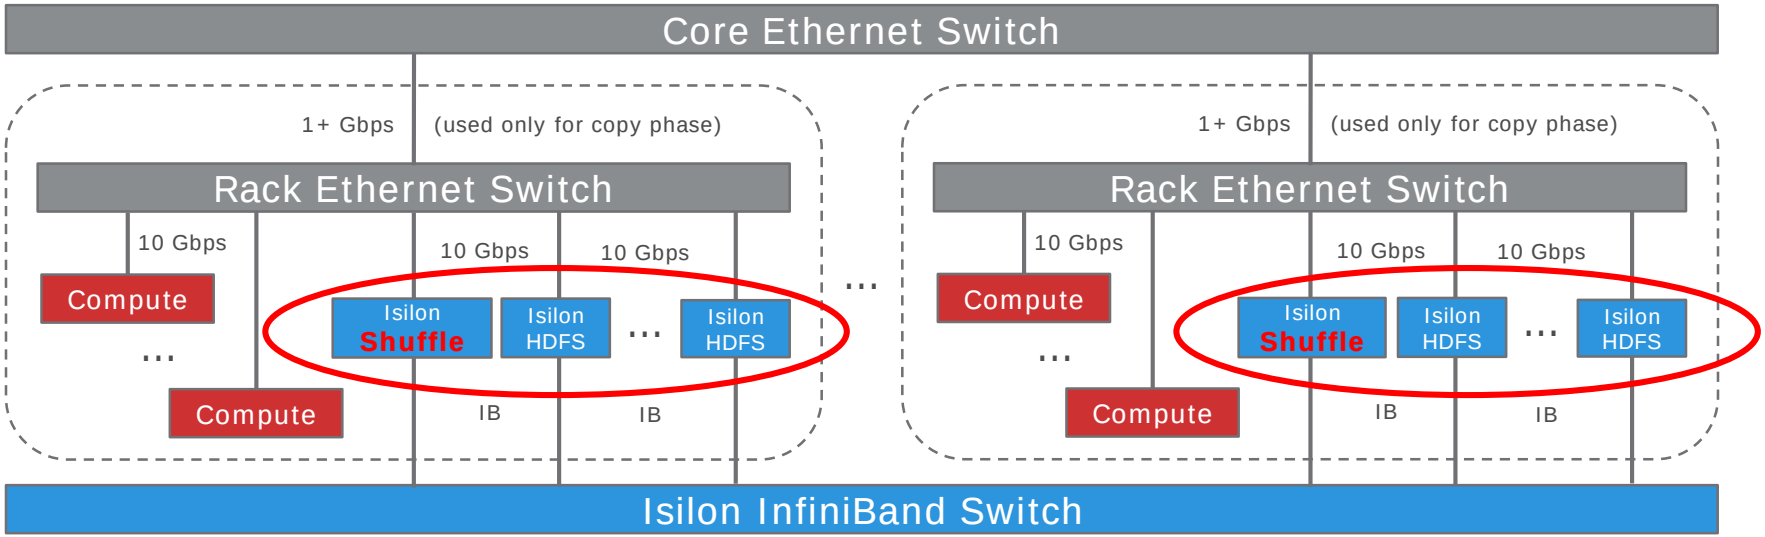

The number of compute and Isilon nodes can be adjusted independently to achieve the optimal ratio of compute and I/O bandwidth

HDFS I/O ALWAYS comes through a rack-local Isilon node which collects data blocks from all other Isilon nodes across the InfiniBand fabric

Shuffle I/O is also on an Isilon cluster. It can be a standalone Isilon cluster or tier with one node per rack. This will support the high stream count needed for optimal merge sort operations.

EMC Isilon: Massively Scalable with Simple & Ease of Use

Isilon scales from

16TB to 50PB

in a single file system, single volume cluster

- Under 60 seconds to scale with no downtime
- **NO** manual intervention
- **NO** reconfiguration
- **NO** server or client mount point or application changes
- **NO** data migrations
- **NO** RAID

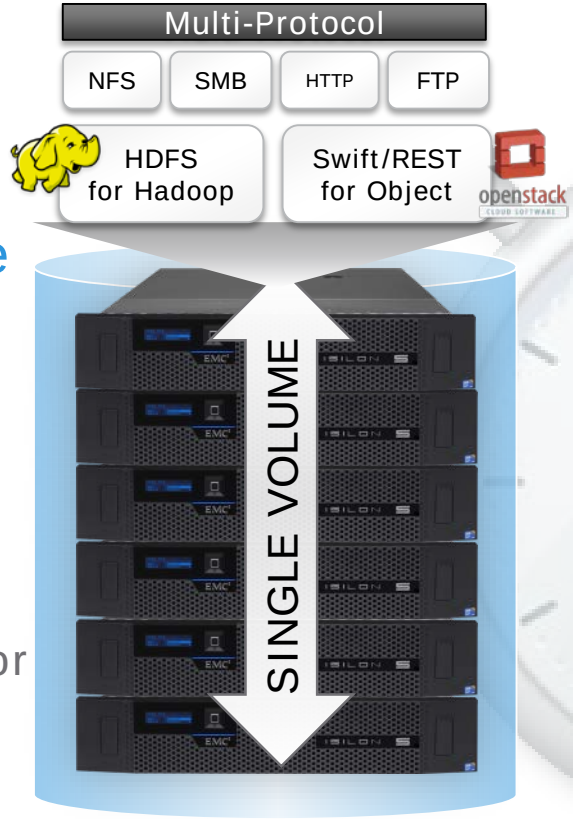

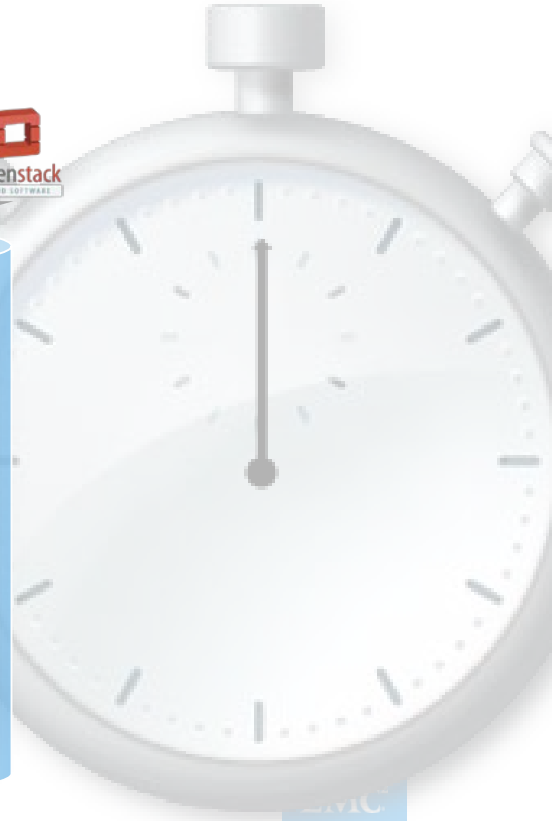

EMC Isilon Scale-Out NAS Product Family

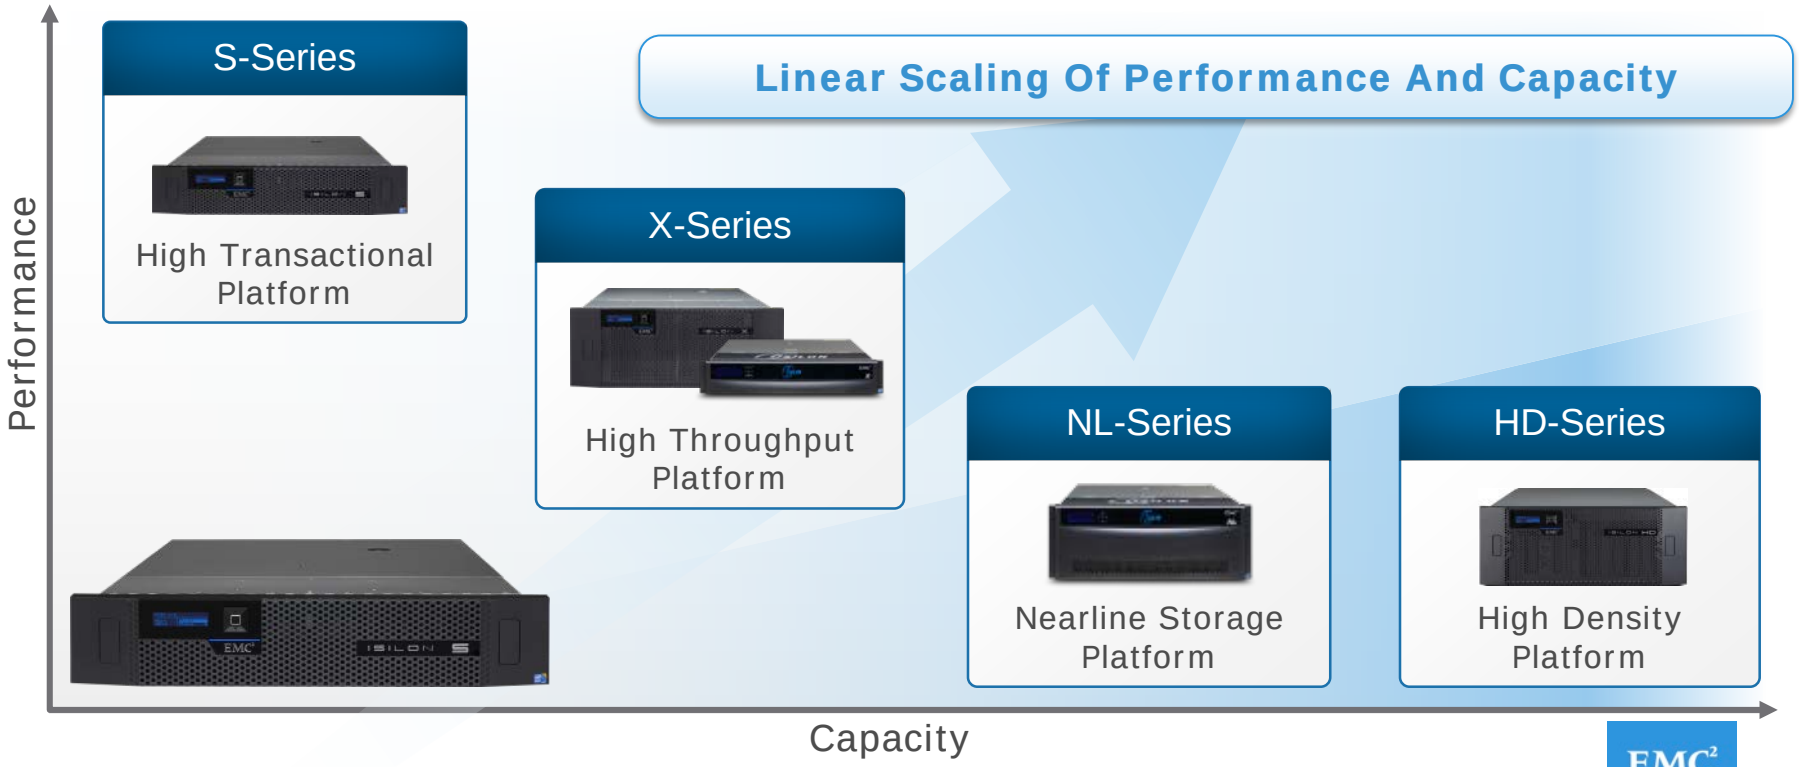

Computing Centre of the National Institute of Nuclear and Particle Physics

Using an EMC Isilon cluster to support cutting-edge research

Challenge

- Highly scalable data storage infrastructure
- On-demand cloud computing platform

Solution

- EMC Isilon X200

Applications

- Cloud Computing
- Infrastructure as a service

BENOIT DELAUNAY

Leader Network Infrastructure, system, storage team

"The EMC Isilon cluster has the capacity to fully accommodate the evolution requirements of our cloud computing platform in terms of storage capacity as well as bandwidth, as the 'hot' addition of modules is enabled without the need for any system downtime. Its rapid implementation and short learning curve translated into significant saving of time for our teams."

Results

- Infrastructure-as-a-service (IaaS) solution
- Implementation of a new "on-demand" computing platform
- Fully scalable in capacity and bandwidth

EMC ECS Appliance

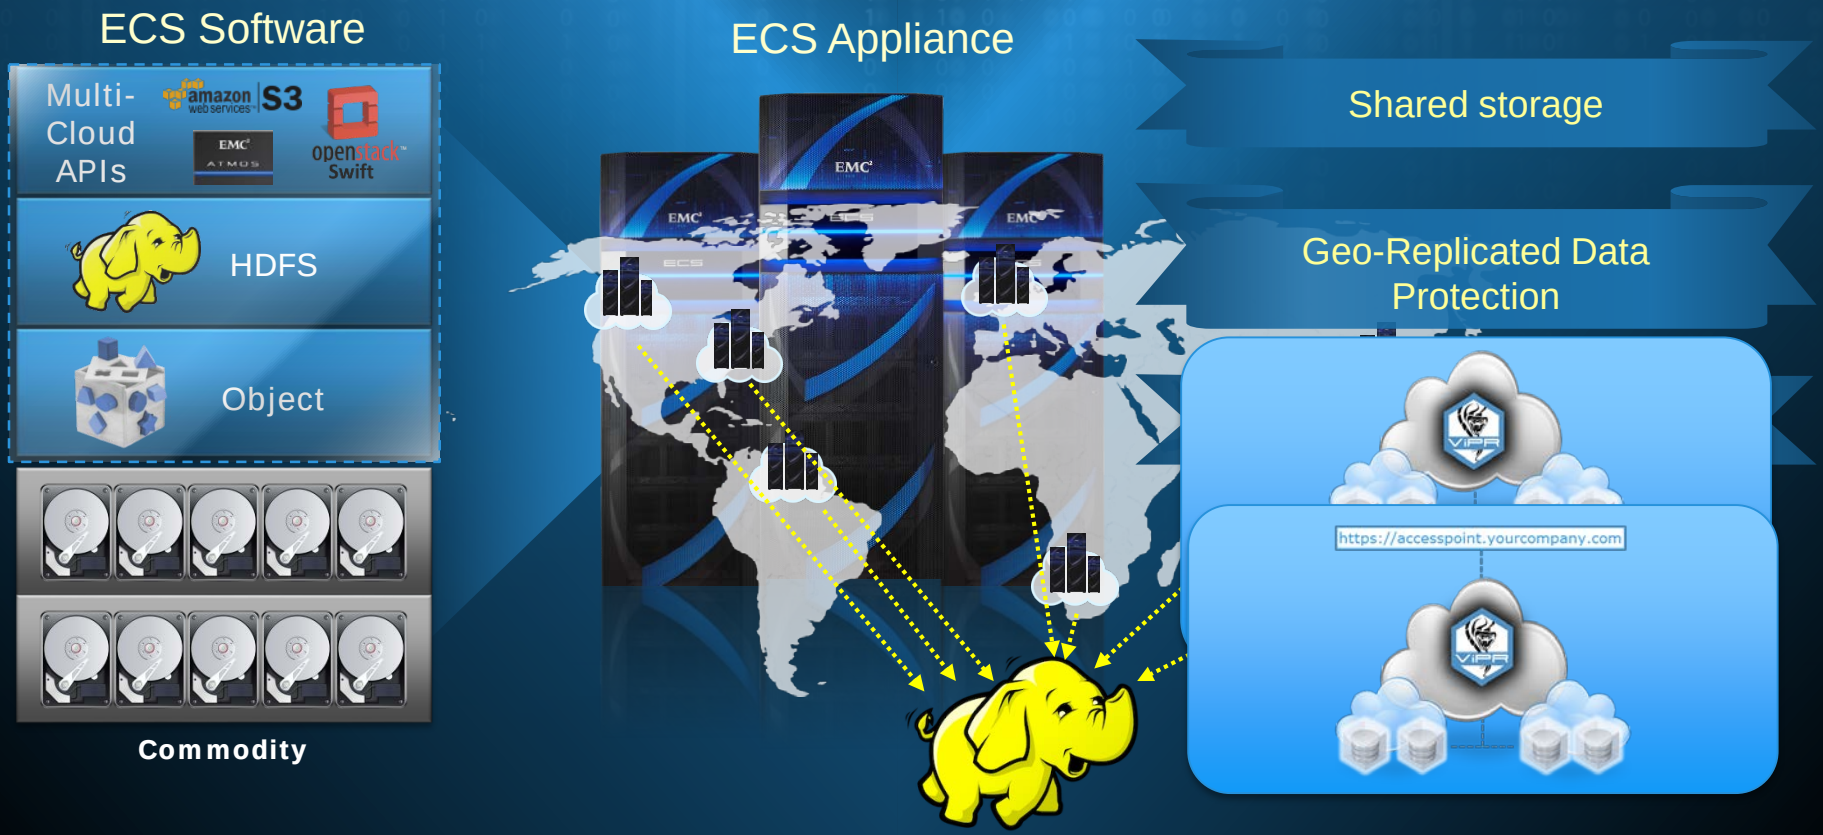

DSSD: RACK SCALE Flash STORAGE

- Wicked fast, low latency persistent storage!
  - Up to 144 TB in 5u
  - ViPR controller
- 10-20 million IOPS
- 100 GB/s bandwidth
- <100µs latency
- New data access - flood datapath
- Direct app API interface

nvm EXPRESS

REDEFINE IT

Workload Boost

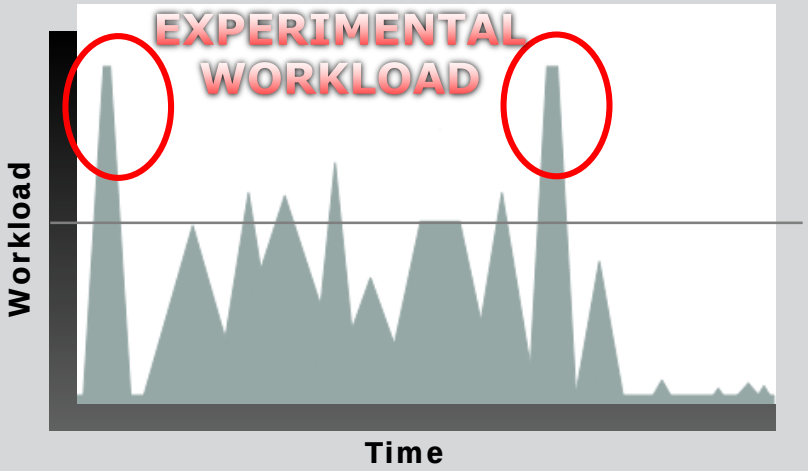

✓ **ELASTIC COMPUTE & STORAGE RESOURCE**

Elastic Cloud

On-Premise Private Cloud

Software-Defined
Elastic Storage

ScaleIO
Elastic Converged Storage

SOFTWARE DEFINED DATA CENTER

Hybrid Cloud

Enterprise Hybrid Cloud

EMC ENTERPRISE HYBRID CLOUD

PLATFORM AS A SERVICE

EMC²

BLOCK

Software-Defined Elastic Storage

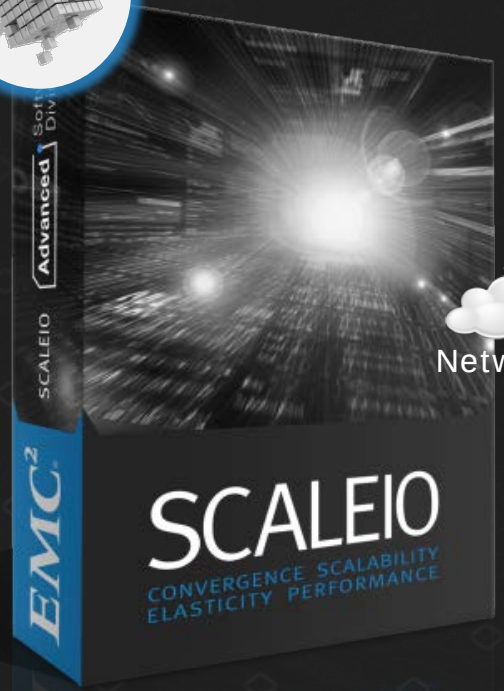

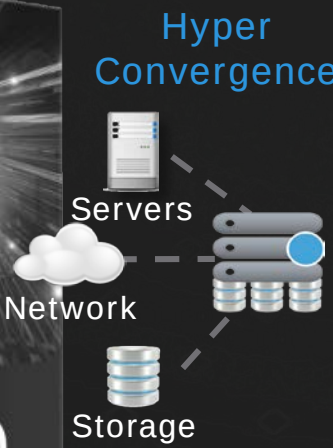

- **Media/Server/OS** – agnostic
- **Scale-Out** of thousands of Servers
- **Elastic** – add/remove on the fly
- **Performance** – Massive I/O parallelism
- **Integration** – OpenStack Cinder

ScaleIO
Elastic Converged Storage

Before ScaleIO
With ScaleIO

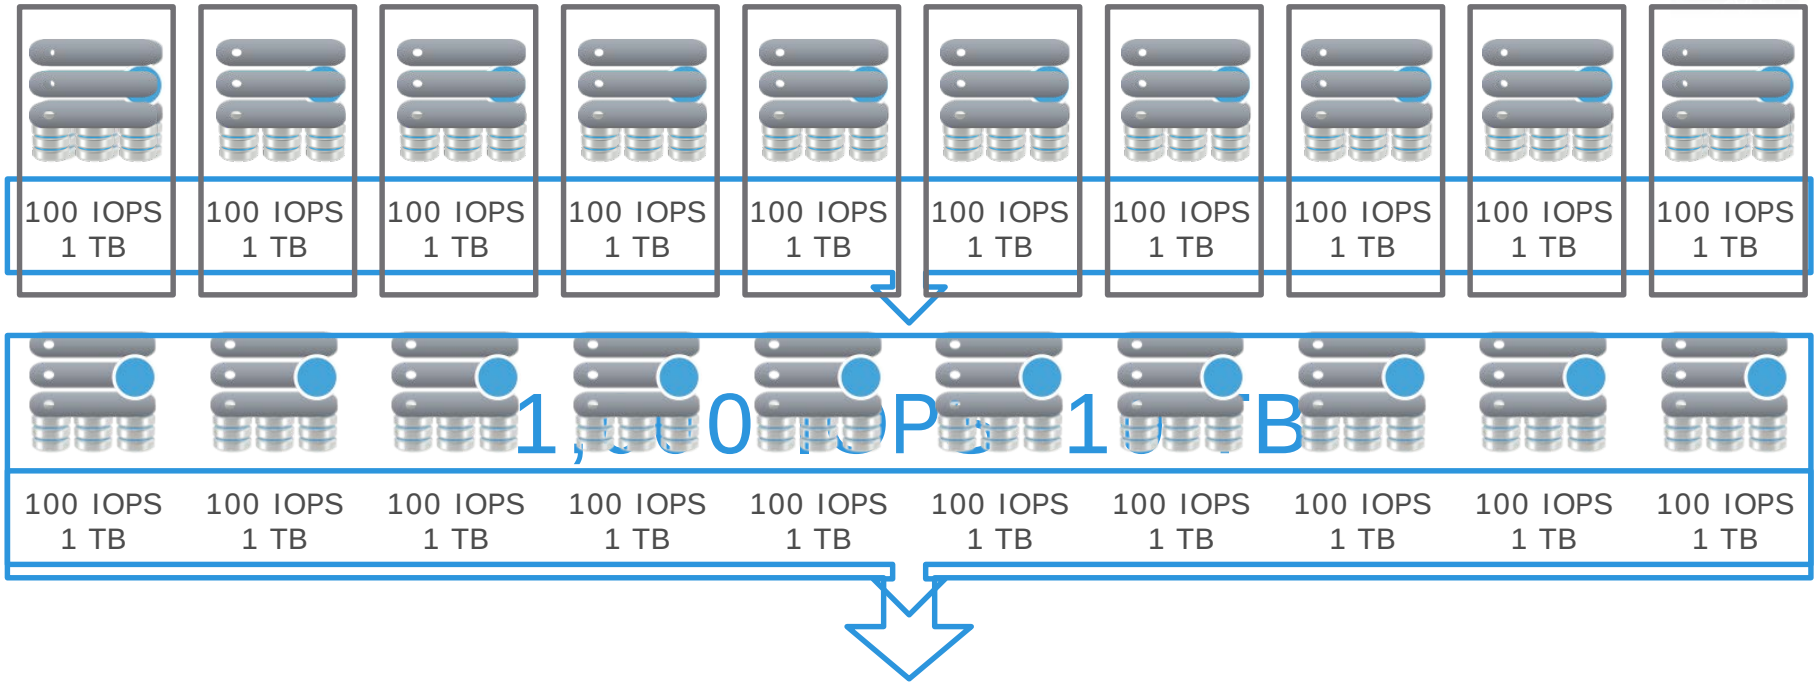

ScaleIO: Auto-Rebalance

Add Nodes or Disks Dynamically

—System Automatically Migrates and Rebalances Storage

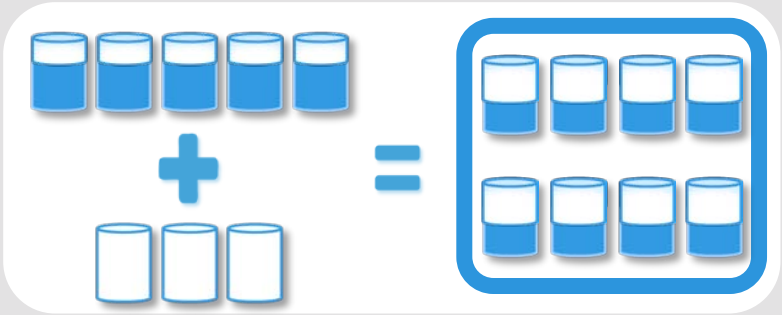

Remove Nodes or Disks Dynamically

—System Automatically Migrates and Rebalances Storage

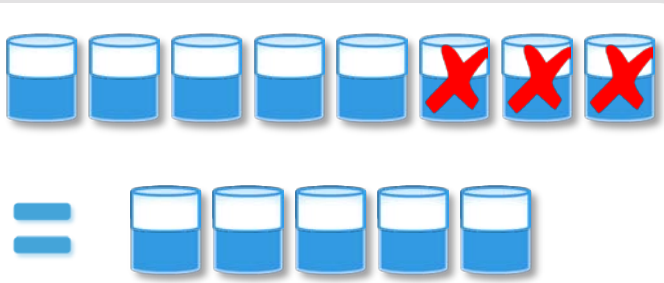

ScaleIO : Configuration

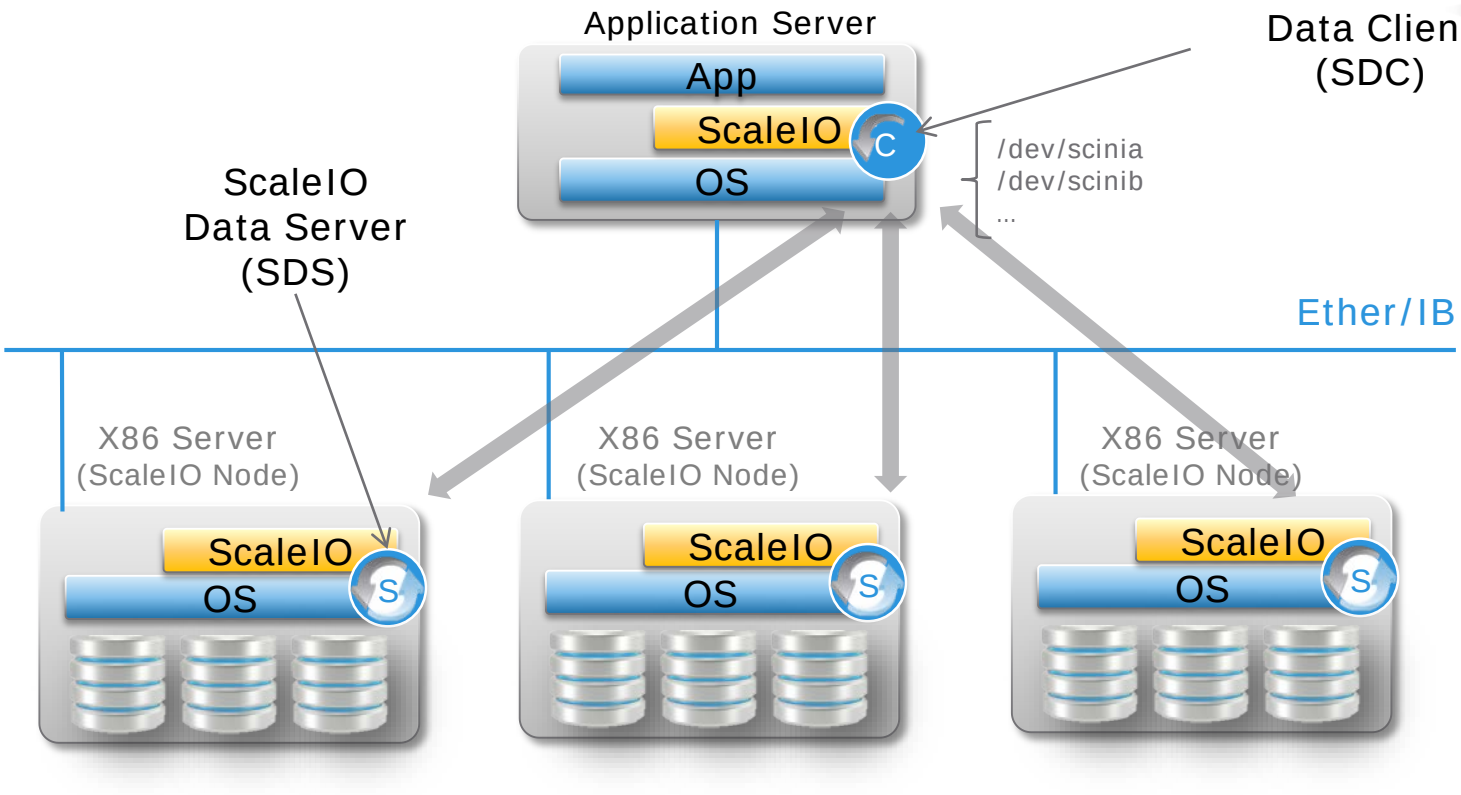

ScaleIO : Architecture

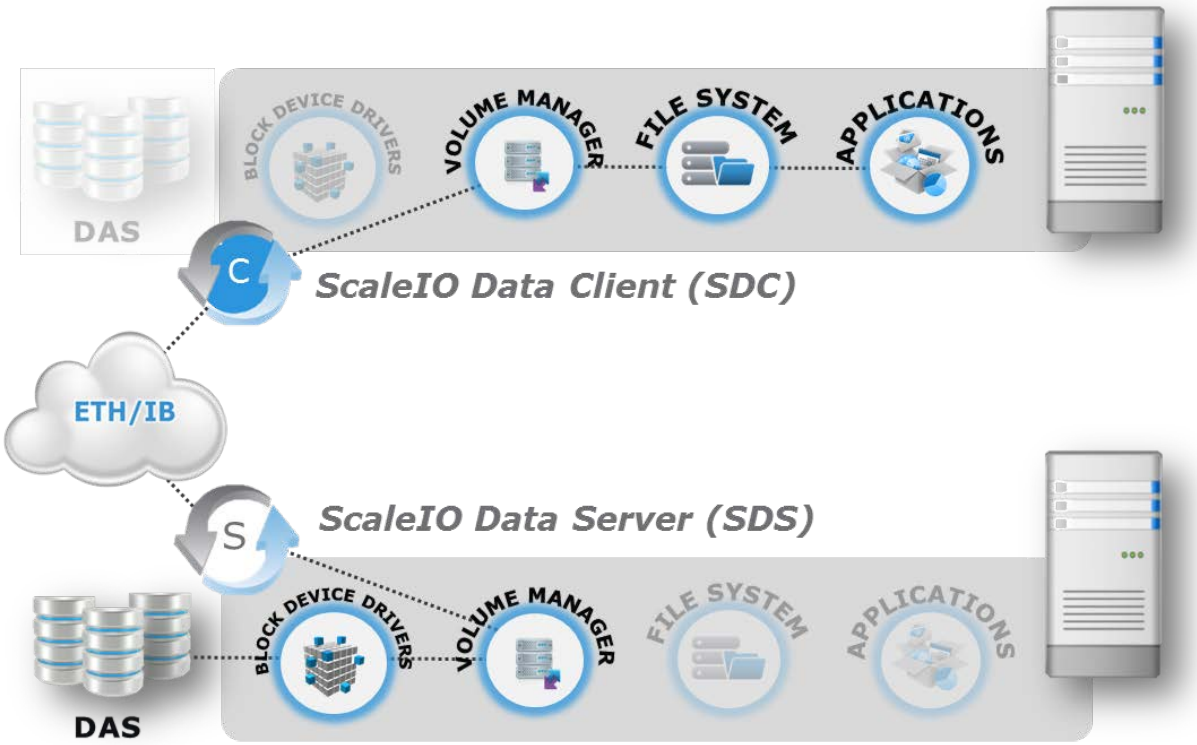

ScaleIO Data Client (SDC)

- ✓ Block Devise Driver
- ✓ Exposes volumes to applications
- ✓ Service must run to provide access to volumes

ScaleIO Data Server (SDS)

- ✓ Abstracts storage media
- ✓ Contributes to storage pools
- ✓ Performs I/O operations

ScaleIO :
Fully Converged Configuration

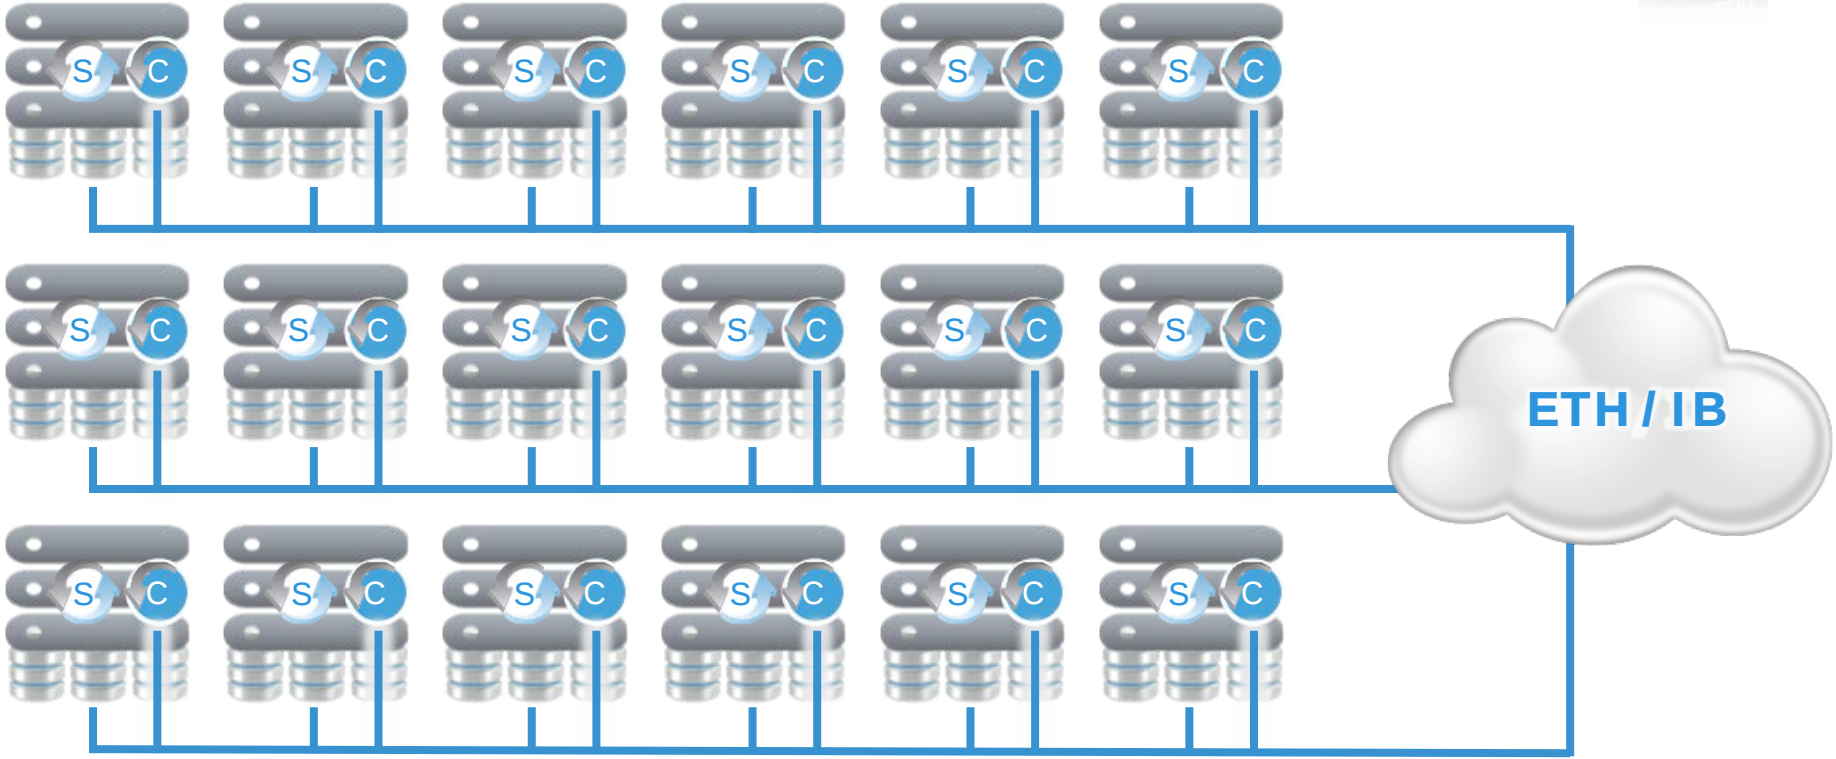

ScaleIO :
Two-Layer Configuration

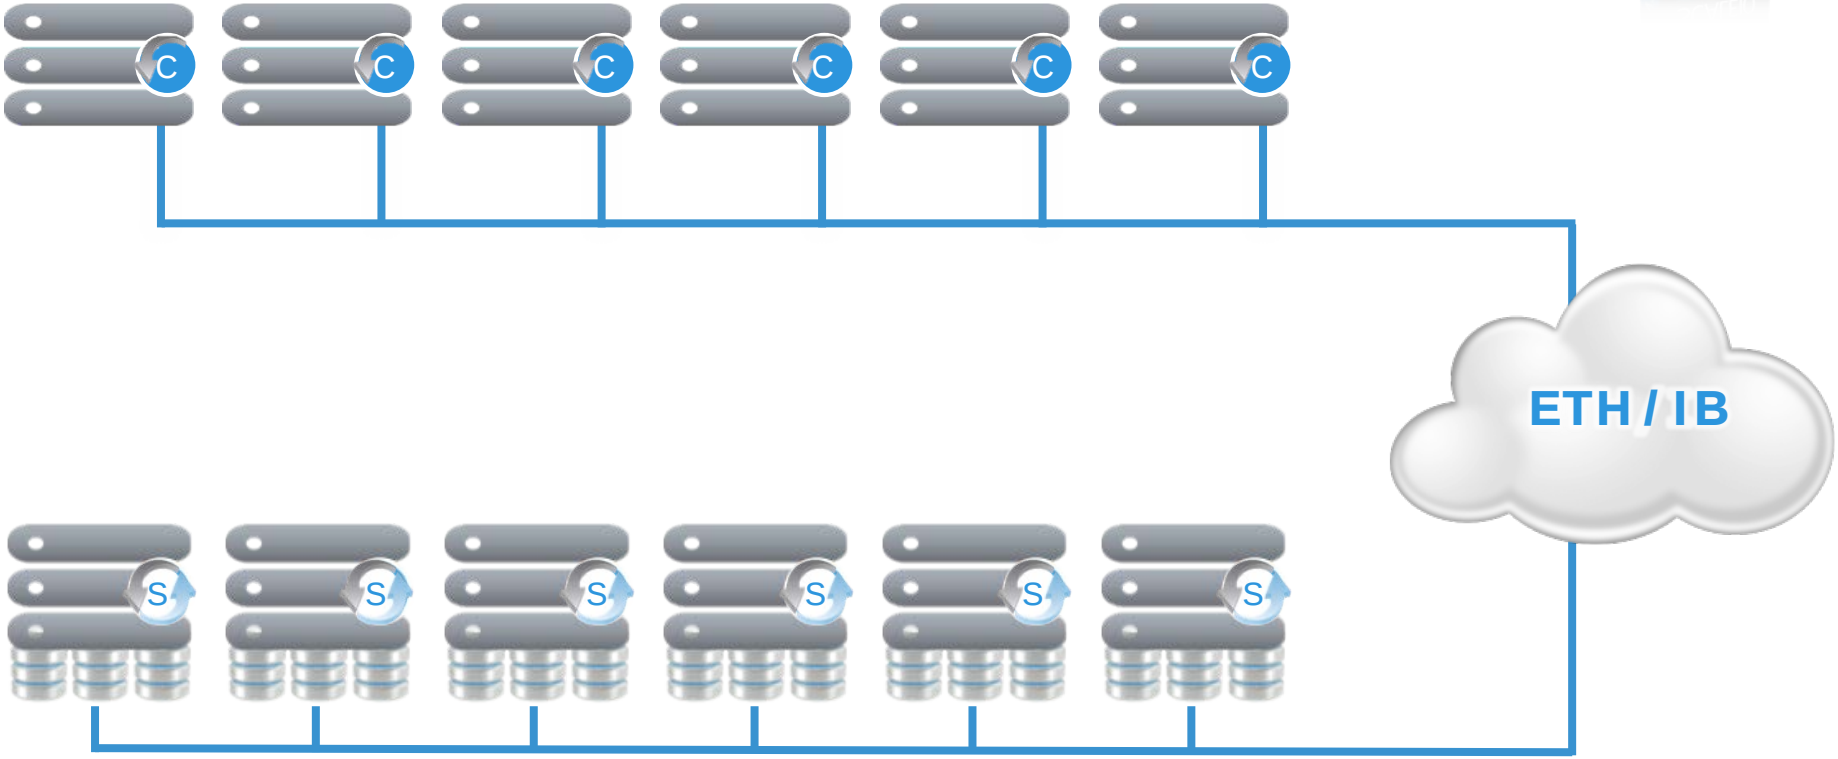

ScaleIO :
Two-Layer Configuration

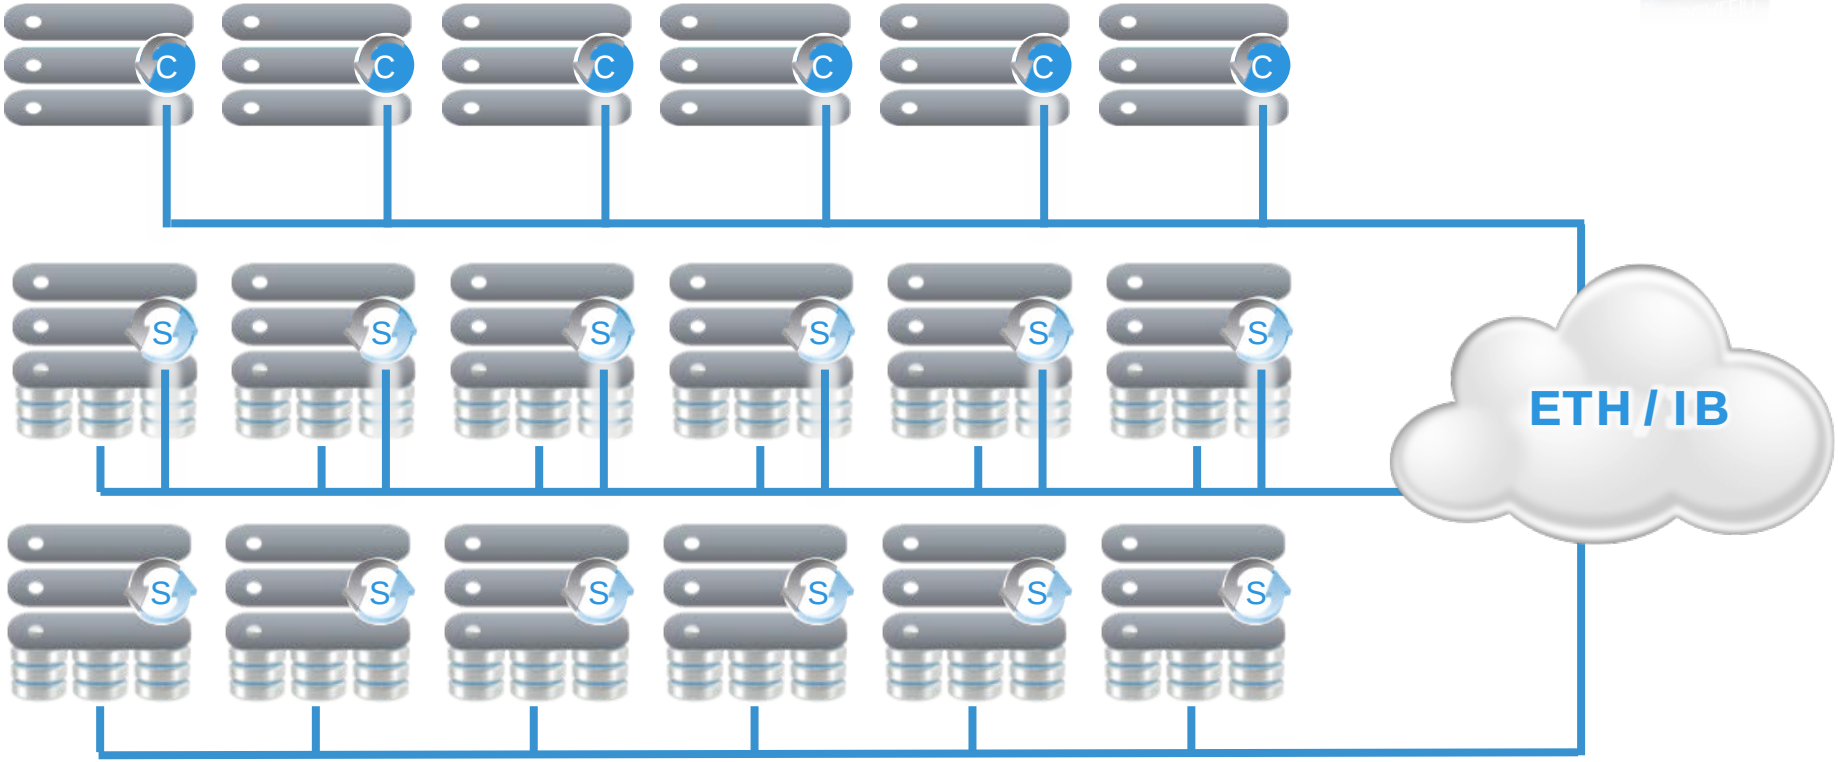

ScaleIO – Massive scalable Performance

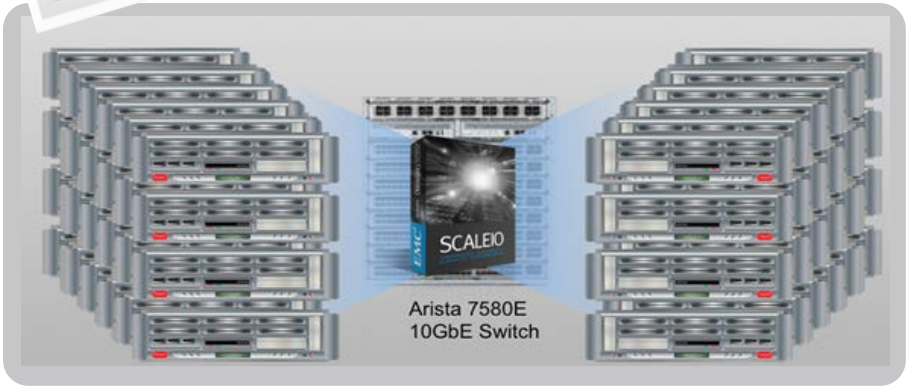

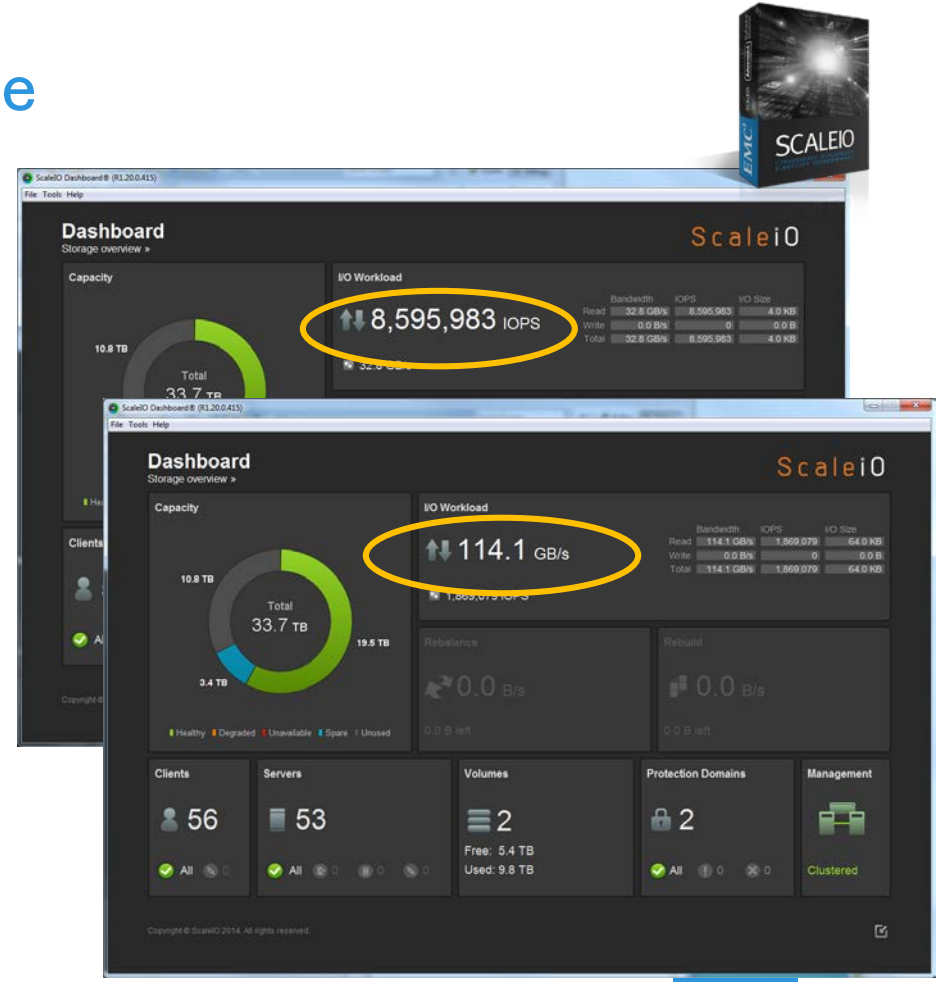

Configuration (per Node):
Cisco UCS C240M# Servers
2 x E5-2680 (2.8GHz) w/64GB RAM
Intel X520 Dual Port 10GbE Adapter
1 x 700GB SLC PCIe Flash Card

EMC Enterprise Hybrid Cloud (EHC)

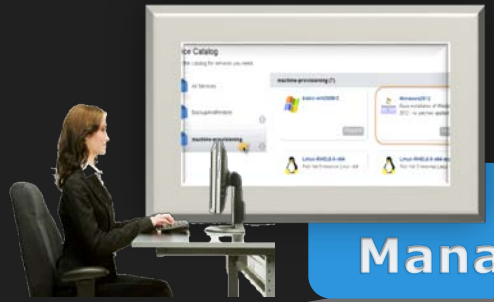

Self Service Portal

Integrated Management & workflow

Private Cloud

Public Cloud

Pivotal CF (Cloud Foundry) : Platform provisioning

Software-Defined Data Center

VMware vCloud : Integrated Cloud Management & Automation

VMware SDC&SDN

EMC Software-Defined Storage

Compute | **Network** | **Storage**

EMC, 3rd Party, Commodity

EMC BUSINESS PARTNER CLOUD SERVICE PROVIDER

CLOUD SERVICE PROVIDER

VMware vCloud Air

Pivotal

Application Mobility

vmware

VM Mobility

EMC²

Data Mobility

EMC²

VM Mobility by EHC

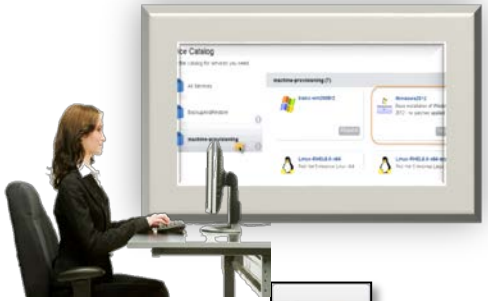

VM

EHC
PRIVATE CLOUD

VM
VM
VM

- Migration per VM
- Migration VM Template

Powered by VMware
vCloud Connector

VM
VM

vCloud Air

VM

EHC
PUBLIC CLOUD

VM

Expand Computing Resource

Expand Data Center

EMC Cloud Service Provider Partners

Application Mobility by EHC

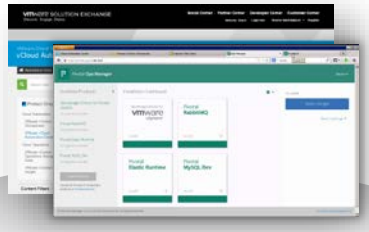

Application

VM VM VM VM

PaaS

CF Pivotal CF

EHC
PRIVATE CLOUD

Application

VM VM VM VM

PaaS

CF Pivot CF

vCloud Air
Pablic Cloud

Deploy PaaS & Application

Expand Computing Resource

Expand Data Center

EHC and the beyond

vmware®

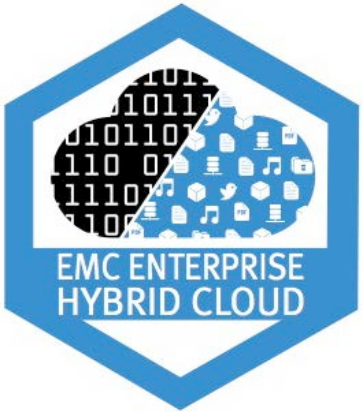

FEDERATION
SDDC EDITION

AVAILABLE NOW

Microsoft

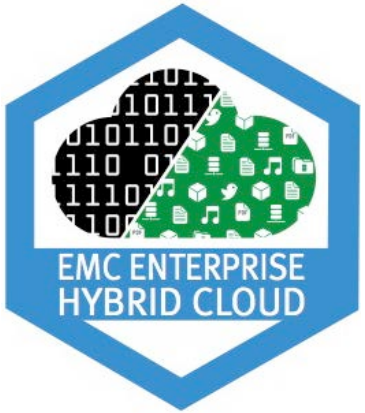

MICROSOFT
EDITION

2015

openstack™
CLOUD SOFTWARE

cloudscaling

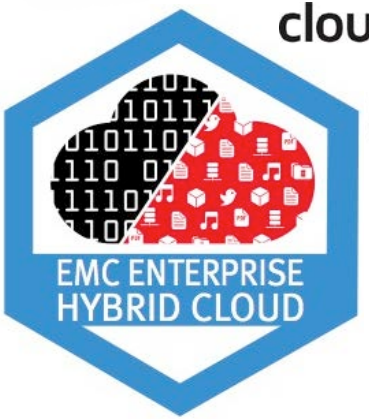

OPENSTACK
EDITION

2015

EMC²

Redefine IT by EMC Fereration

- ✓ **COST EFFICIENCY**
  FOR STORAGE AND DATA MANAGEMENT
- ✓ **FAST DATA ANALYSIS**

- ✓ **ELASTIC COMPUTE & STORAGE RESOURCE**

Federation Data Lake

Pivotal In-Memory Data Grid

GemFire XD

hadoop HDFS

EMC²

HDFS Storage

To accelerate data analytics in fast and cost efficient way

EMC ScaleIO

To utilize Elastic storage resource by commodity hardware

EMC Enterprise Hybrid Cloud

To Expand resource to hybrid cloud seamlessly

EMC²

THANK YOU.

EMC² Pivotal RSA VCE vmware