

The image shows a vertical stack of server racks. The top rack is dark with the EMC² logo in white. Below it is a blue horizontal band. The next rack down has the Isilon logo in blue. Below that is another dark rack with the Isilon logo in blue. The bottom two racks are also dark with the Isilon logo in blue. The overall lighting is dim, with the blue band providing a highlight.

EMC<sup>2</sup>

ISILON

ISILON

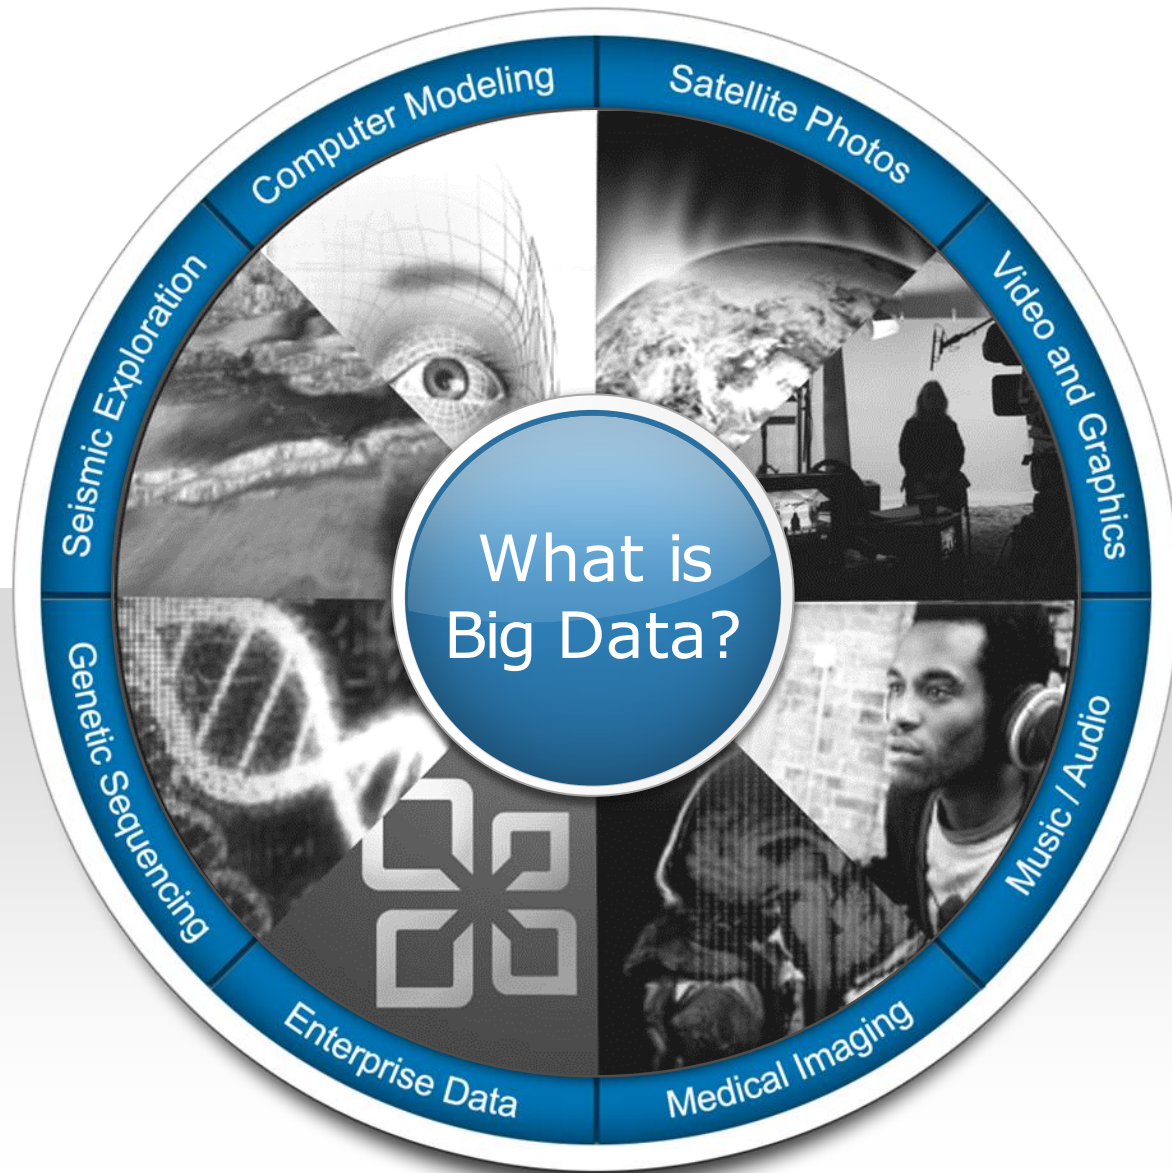
ISILON

# *The BIG Data Era has arrived...Re-invent your storage!*

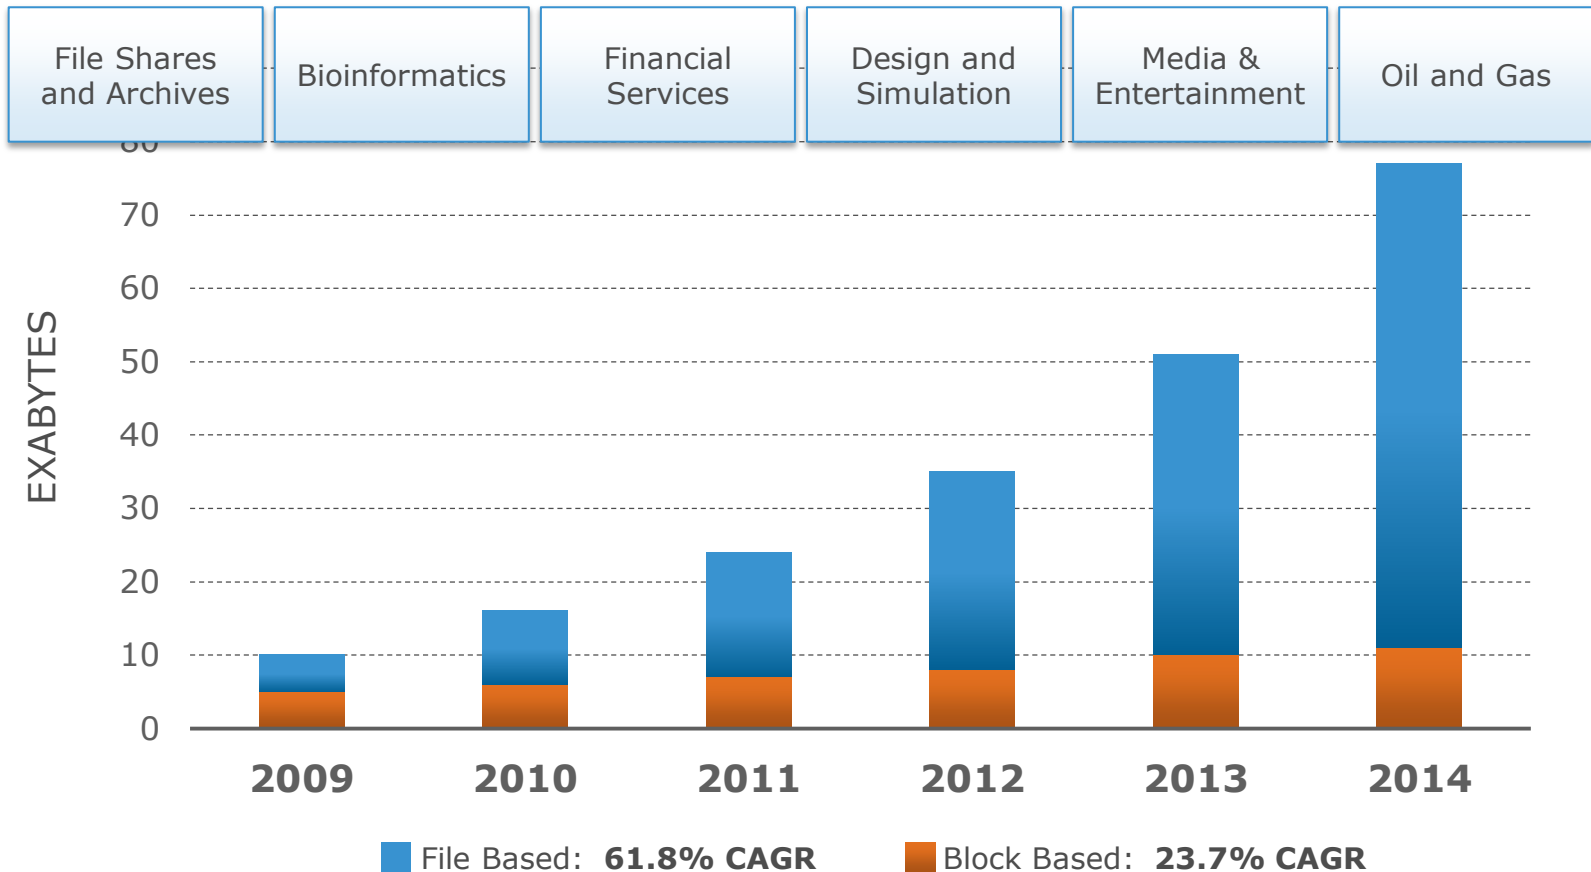
Bratislava, Slovakia, 21st March 2013

Luka Topic  
Regional Manager – East Europe  
EMC Isilon Storage Division  
[luka.topic@emc.com](mailto:luka.topic@emc.com)

EMC<sup>2</sup>



# The Big Data Challenge



**By 2013, 80% of all storage capacity sold will be for file-based data**

Source: IDC Whitepaper June 2011 "Scale Out Storage In the Content Driven Enterprise: Unleashing the Value of Information Assets" referencing IDC's 2010 Enterprise Disk Storage Consumption Model



# EMC Isilon: Scale-Out Storage

EMC<sup>2</sup>

# EMC, a true Innovator, and Market Consolidator

9-Year \$26.8B Investment In Market  
Leading Technology 2003-2011.  
More than 40 acquisitions since 2003.

**M&A \$14.5B**

**R&D \$12.3B**

**2012**

**TRANSFORM**  
IT+BUSINESS+YOURSELF



**2009** Journey to the Cloud



**datadomain**

**2007** Information & Virtual Management



**2005** Information Lifecycle Management

**vmware**



**2003** (and before) Enterprise storage

**LEGATO**

# EMC ISILON Overview

**Proven, global leader** in Scale-Out Storage solutions  
(IDC)

Unrivaled product **innovation** – 7<sup>th</sup> Generation

World-class **products** and global EMC **services**

Strong ecosystem of **technology partners and  
application integration**

**~3000 worldwide customers** and growing

Acquired by EMC for \$2.5 Billion in Dec 2010

**Record Growth & Profitability**

# The EMC Isilon Advantage For The World Of Big Data

## Massive Scalability

Scales To Over 20 PB In A Single File System

## World Record Performance

1.6 Million CIFS File Operations Per Second

## Unmatched Efficiency

>80% Storage Utilization & Automated Storage Tiering

## Operational Simplicity

Single File System, Single Volume, Global Namespace



EMC<sup>2</sup>

# The EMC Isilon Advantage For The World Of Enterprise IT

## Data Security

WORM Data Retention, RBAC, Authentication Zones

## Data Protection

N+1 Through N+4 Redundancy; Efficient Snapshot Data Back-up and Restore; Rapid Data Replication for Disaster Recovery Protection

## Interoperability

Support For Industry-standard Protocols: NFS, SMB, HTTP, FTP And iSCSI; World's Only Scale-out Storage System With Native Hadoop HDFS Support

## Performance


>100GB/second Throughput





# EMC Isilon Storage Efficiency

- Single volume spans all nodes
  - Directories and files striped across the cluster
- Auto-balance capacity and performance
- Automation:
  - NO manual intervention, NO reconfiguration, NO server or client mount point or application changes
- Storage efficiency inherently optimized at 80+%



# EMC Isilon Business Benefits

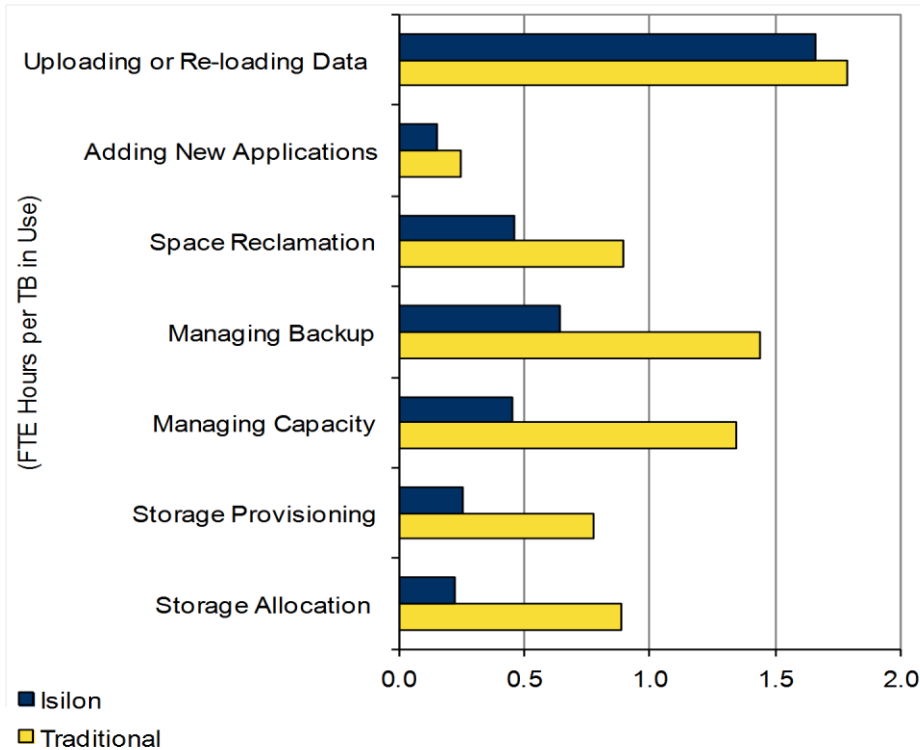
- Ease of Management
  - Reduces OPEX
- Linearly Scalable Performance and Capacity
  - Predictable SLAs
- Inherent 80%+ Storage Efficiency
  - Reduces CAPEX
- Unmatched Reliability and Availability
  - N+4 Protection



# What IDC is saying...

# EMC Isilon

## Ease of Use And Management



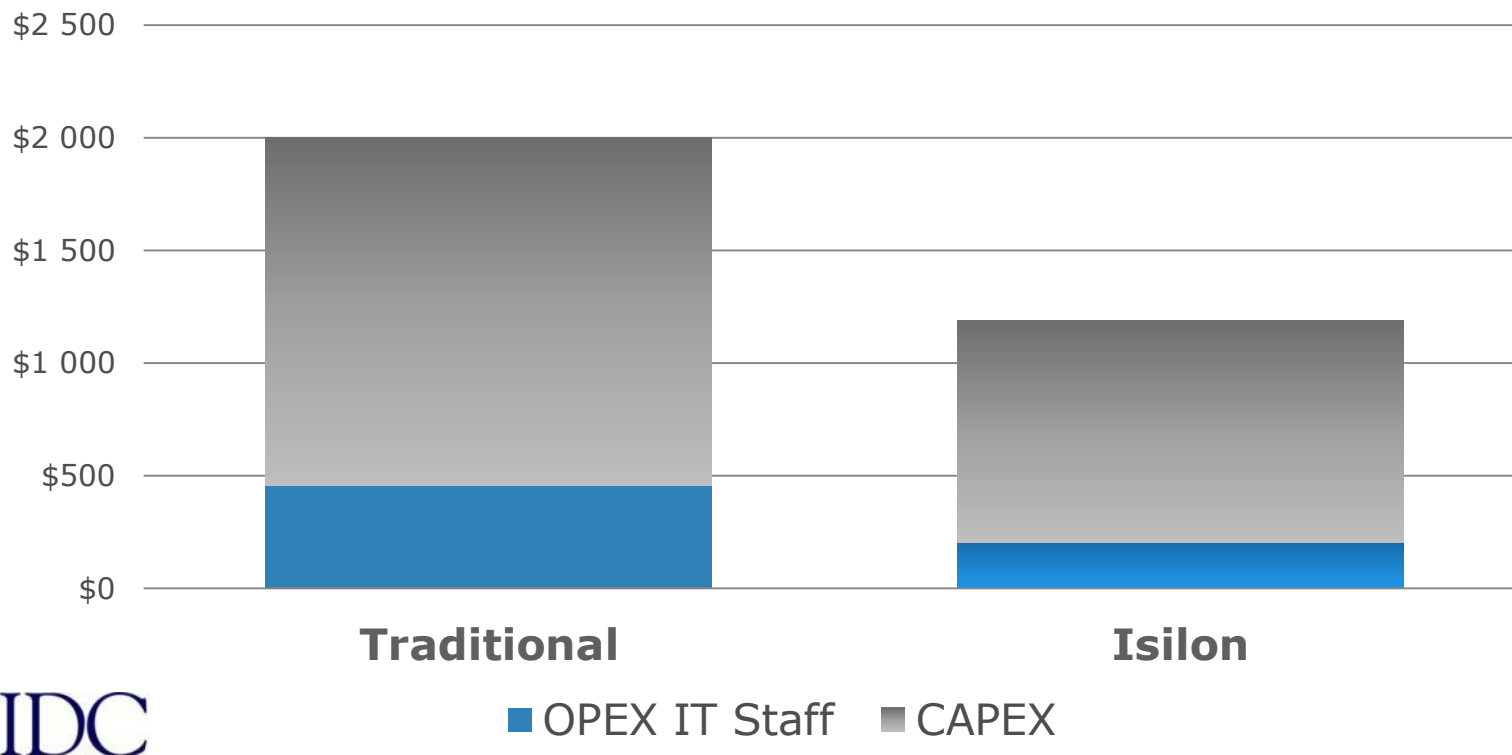
## IDC: Isilon Improves IT Productivity by 48%, Reduces OPEX\*

\*Results of an IDC Business Value Strategy Analysis - 28 November 2011

# EMC Isilon Customers

Reduced Big Data Storage Costs by 40%

## Average Annual Cost Per TB in Use

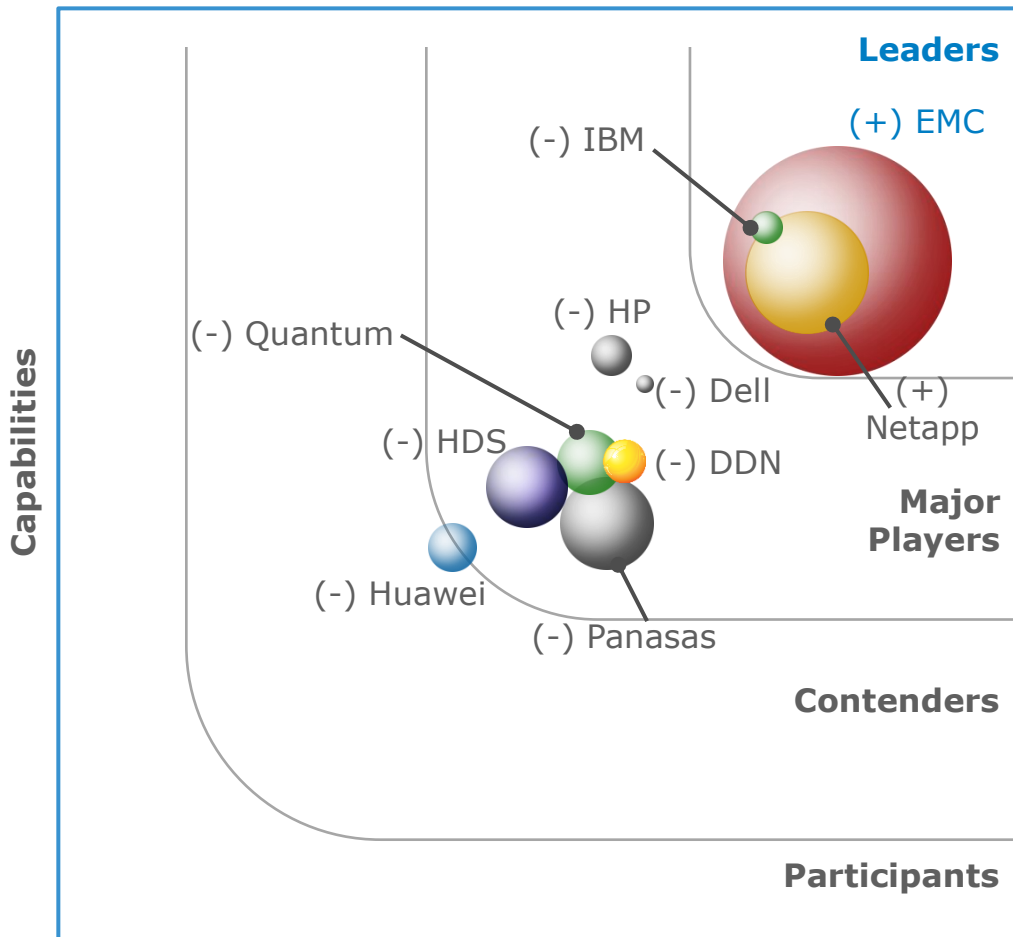


■ OPEX IT Staff ■ CAPEX

Source: IDC Whitepaper Quantifying the Business Benefits of Scale-Out. October 2011



# Isilon: Leader in IDC's Scale-Out Storage MarketScape Report



IDC MarketScape - Scale-Out File-Based Storage Market  
January, 2013



“EMC’s backbone for the scale-out file-based storage segment is Isilon...”

“EMC acquired Isilon in late 2010 and positions Isilon's unique clustered architecture for customers that face rapid growth of data, high bandwidth needs, and high availability demands.”



...customers  
can confirm...

# EMC Isilon Growing Momentum 3000+ Customers



EMC<sup>2</sup>



# Scale-Up vs. Scale-Out

# CIO's Turning To Scale-Out To Deal With Massive File-Data Growth

The #1 and #2 concerns of CIOs... ...are driving adoption of scale-out!

Data Growth

47%

System Performance & Scalability

37%

**Gartner** CIO Survey

84%

Already using or planning to use scale-out within the next 24 months



Scale-out NAS Market Survey

EMC<sup>2</sup>

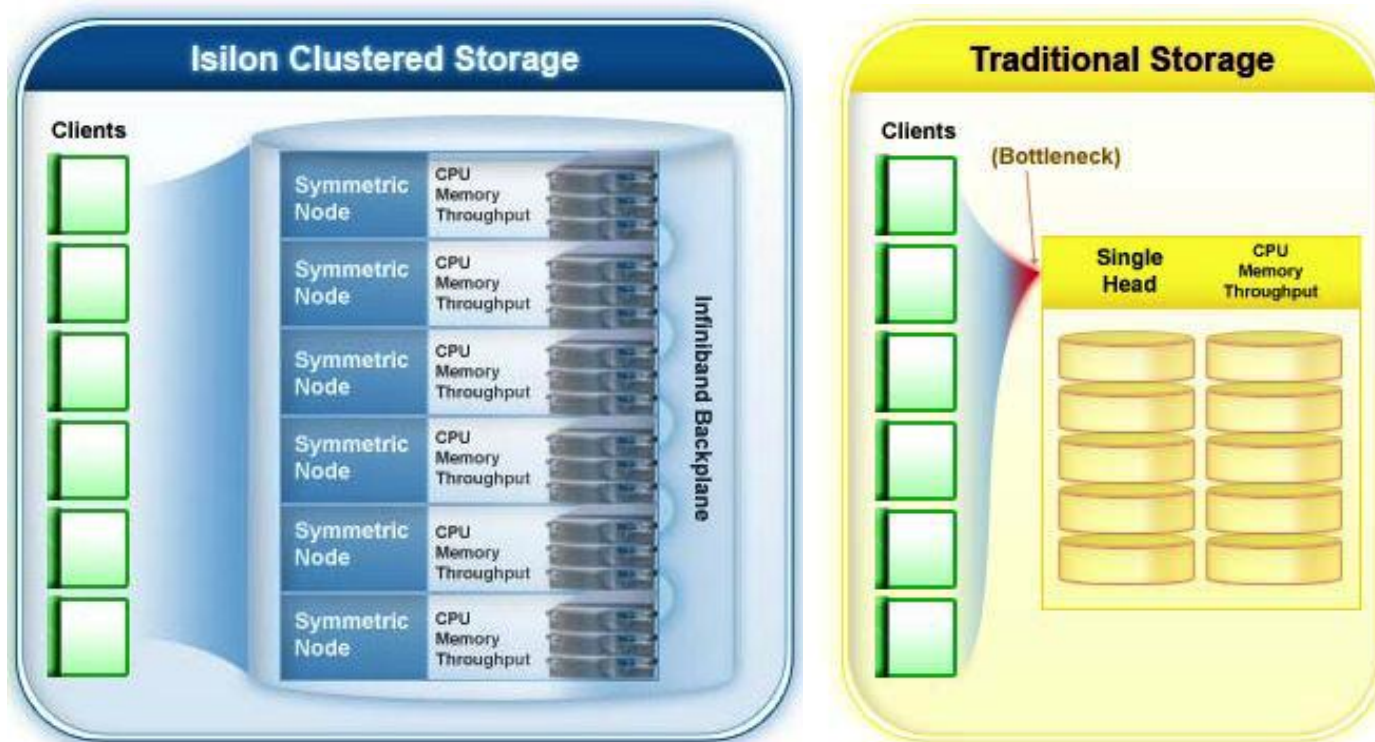
# Common terms you may know....

- RAID
- RAID sets/stripes
- LUN
- Aggregate
- Filer
- Disk Tray
- HBA
- Metadata Server
- Hot spare
- Hot spots
- LUN thrashing
- Forklift upgrade
- 16TB LUN limit
- File System defrag
- FSCK
- Discretionary reserve
- 70% performance cliff
- Avg. 50-60% utilization
- Snapshot limits
- Snapshot reserve
- Multiple volumes
- Client side drivers

....that don't apply to EMC Isilon



# Architecture

## Distributed Cluster versus Single-Server



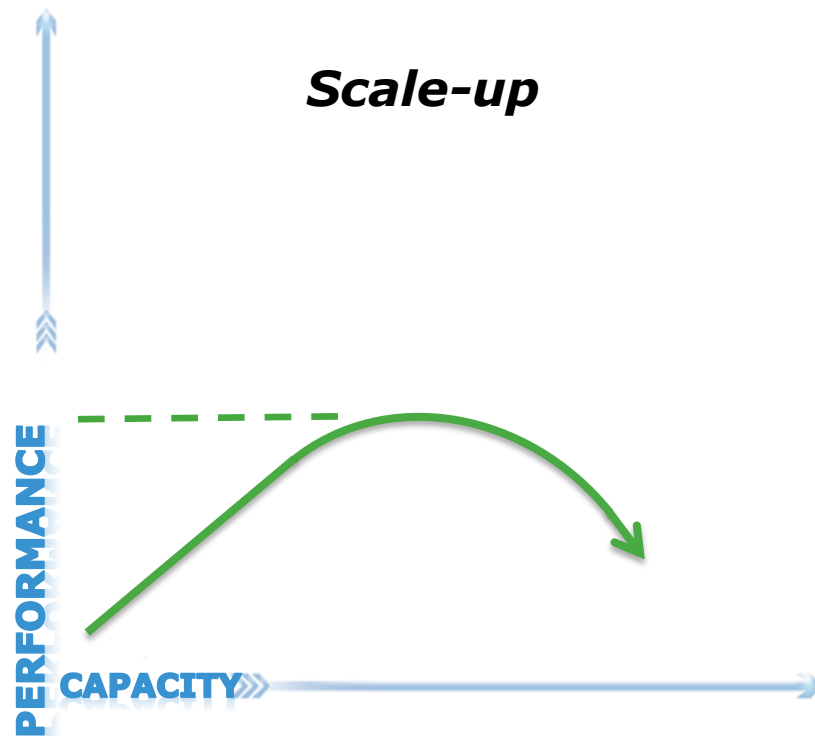
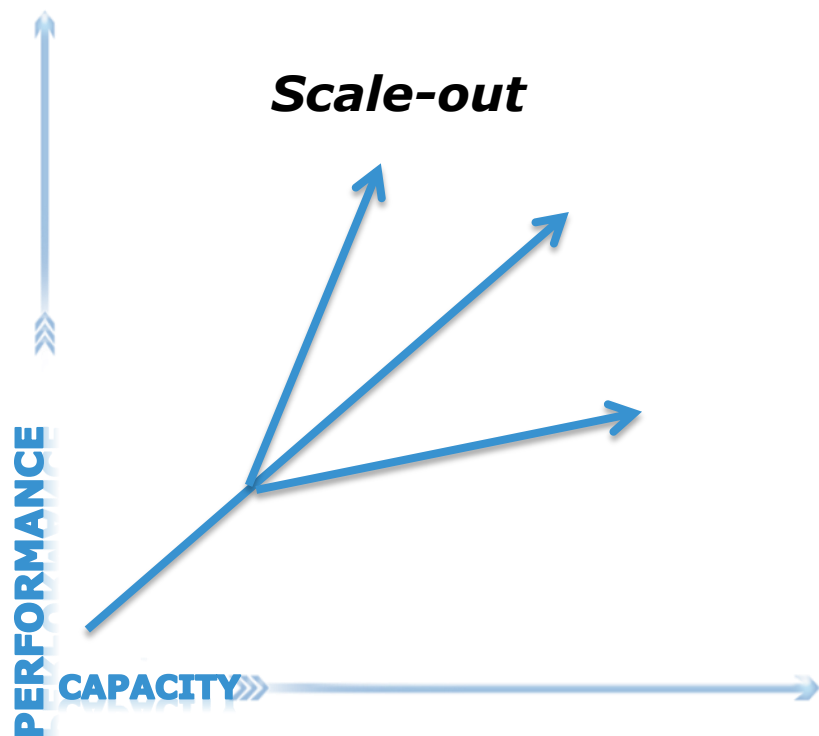
	<i><b>Isilon OneFS</b></i>	<i><b>Traditional</b></i>
<b>Architecture</b>	<ul style="list-style-type: none"> <li>• Symmetric, shared</li> <li>• Fully distributed OS</li> </ul>	<ul style="list-style-type: none"> <li>• Controller with shelves</li> <li>• Legacy single server OS</li> </ul>

# Management and Economics



	<i><b>OneFS – Scale-Out</b></i>	<i><b>Scale-Up</b></i>
<b>Management</b>	<ul style="list-style-type: none"> <li>• 1 volume, 1 file system</li> <li>• Single volume up to 15PB</li> </ul>	<ul style="list-style-type: none"> <li>• Multiple RAID, volumes, file systems</li> <li>• Single volume up 100TB (16TB typical)</li> </ul>
<b>Economics</b>	<ul style="list-style-type: none"> <li>• 80%+ utilization, &lt;1FTE</li> <li>• Business agility</li> <li>• Investment protection</li> </ul>	<ul style="list-style-type: none"> <li>• 57% utilization, Multiple FTE's</li> <li>• Silos, limitations</li> <li>• Forklift upgrades</li> </ul>

# The Scale-Out / Scale-Up Dilemma



	<i>Isilon OneFS</i>	<i>Scale Up -Storage</i>
<b>Scalability</b>	<ul style="list-style-type: none"> <li>• Scale-out</li> <li>• Performance, Capacity, Both</li> </ul>	<ul style="list-style-type: none"> <li>• Scale-up</li> <li>• Capacity only, limited performance options</li> </ul>
<b>Performance</b>	<ul style="list-style-type: none"> <li>• True linear predictability</li> </ul>	<ul style="list-style-type: none"> <li>• Degradation of performance &amp; capacity at scale</li> </ul>

# Scale Capacity

## Isilonmeter



3 Nodes

144 Nodes



Each node adds:

- Throughput
- IOPS
- Capacity
- Connectivity



# EMC Greenplum with Isilon

- Packaged Hadoop distribution
  - Innovation from Greenplum HD
  - Latest stable release of Apache Hadoop
- Enterprise storage from Isilon
  - Better storage utilization without multiple mirrored copies
  - Multi-protocol parallel access from any client (HDFS, NFS, CIFS)
  - Metadata protection (by copying the NameNode information throughout the cluster)







**SO WHAT IS A PETABYTE ANYWAY?**

**WHAT IS A PETABYTE?**

TO UNDERSTAND A PETABYTE WE MUST FIRST UNDERSTAND A GIGABYTE.

- 1 PETABYTE = 20 MILLION FOUR-DRAWER FILMING CABINETS FILLED WITH TEXT
- 1 PETABYTE = 13.3 YEARS OF HD-TV VIDEO
- 1 PETABYTE = SIZE OF THE 10 BILLION PAGES OF FACEBOOK
- 1.5 PETABYTE = INTERNET USERS DATA BACKED UP ON MOZY.COM
- 15+ PETABYTE = THE AMOUNT OF DATA PER DAY PROCESSED BY GOOGLE
- 20 PETABYTE = TOTAL HARD DRIVE SPACE MANUFACTURED IN 1995
- 50 PETABYTE = THE ENTIRE WRITTEN WORDS OF ALL CIVILIZED MAN SINCE THE BEGINNING OF TIME

**A ZETABYTE IS ONE MILLION PETABYTES:**

Facebook: More than 200 million users. More than 200 million photos. More than 200 million posts. More than 200 million messages.

DRIVING MARKETING EFFECTIVENESS BY MANAGING

# THE FLOOD OF BIG DATA

100 PETABYTE OF DATA PER DAY

US\$2.1 PER PETABYTE


4.8 PETABYTE PER DAY

US\$83.2 PER PETABYTE

294 PETABYTE PER DAY

230 PETABYTE PER DAY

**BIG DATA = BIG OPPORTUNITY**



# The Human Face of Big Data

Project which represents the most engaging attempt ever to tap into and put a human face on, the societal impact of Big Data

<http://humanfaceofbigdata.com/>



# THE HUMAN FACE OF BIG DATA

Created by Rick Smolan and Jennifer Erwitz

AGAINST ALL ODDS PRODUCTIONS



© 2012 Against All Odds Productions  
P.O. Box 1180  
Sausalito, California 94965  
thehumanfaceofbigdata.com

All rights reserved. No part of this book may  
be reproduced, stored in a retrieval system,  
or transmitted, in any form or by any means,  
electronic, mechanical, photocopying,  
recording, or otherwise, without prior written  
permission from the publisher.

ISBN 978-1-5240-0071-0  
Manufactured in Canada

1 2 3 4 5 6 7 8 9 10 11

IMAGE: MICHAEL TOFFERT

EMC<sup>2</sup>®