NetApp®

Go further, faster®

HPC & Big Data: Trends to Watch

Dave Ellis

HPC Architect

U. S. Public Sector, NetApp, Inc.

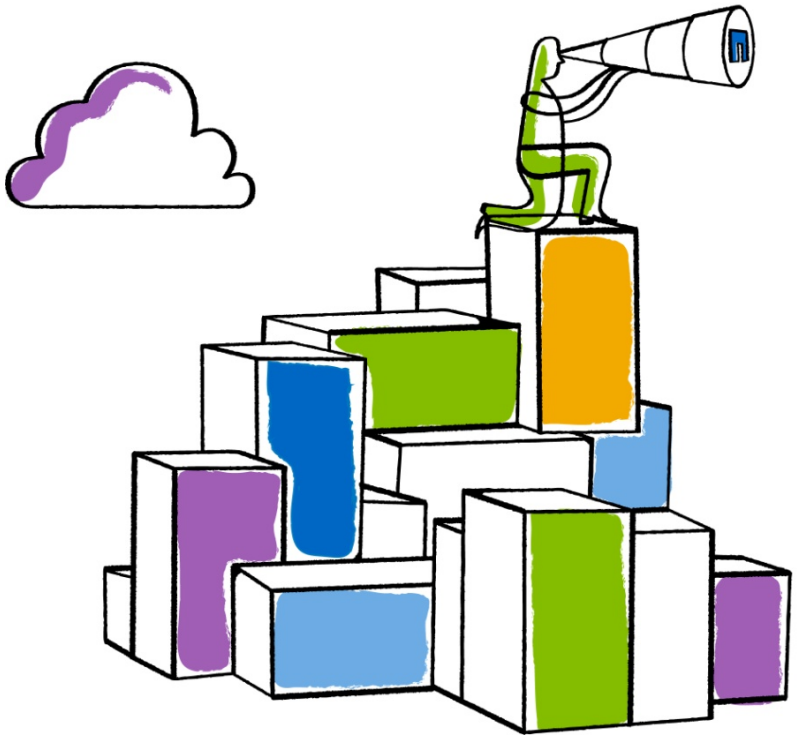

September 3, 2012

What is “Big Data”?

Coined by **Francis Diebold**, professor of economics at the University of PA in 2000, when “Big” meant Gigabytes / day[1]

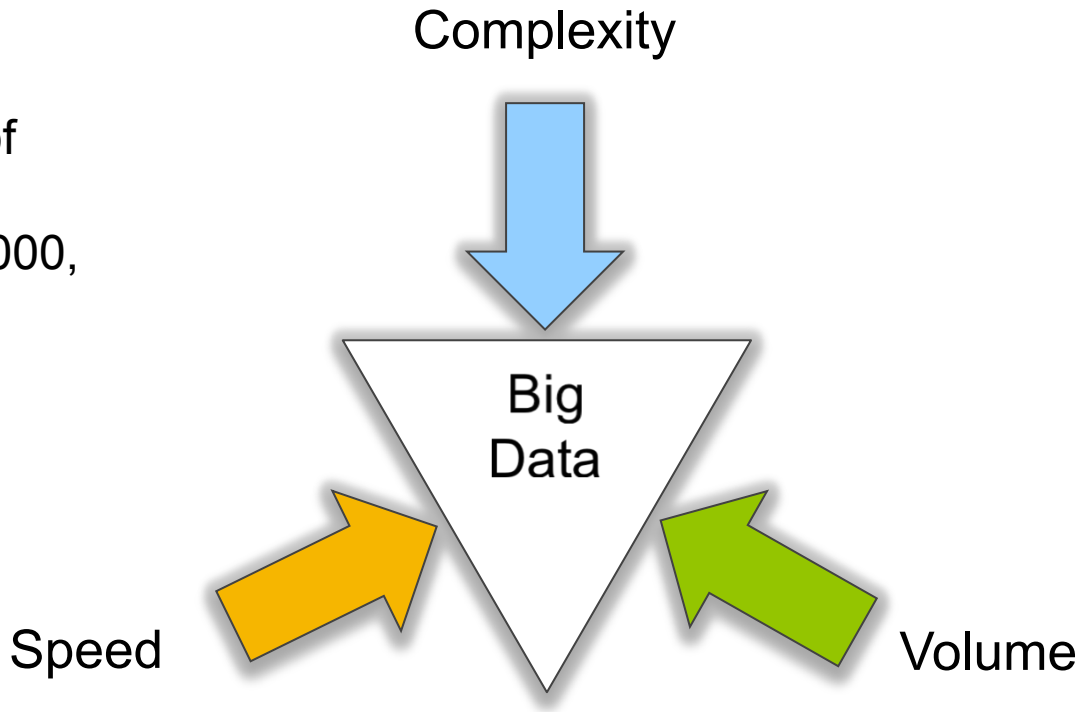

“Big Data” refers to datasets whose volume, speed and complexity is beyond the ability of typical tools to capture, store, manage and analyze.

Big Data Solution Portfolio

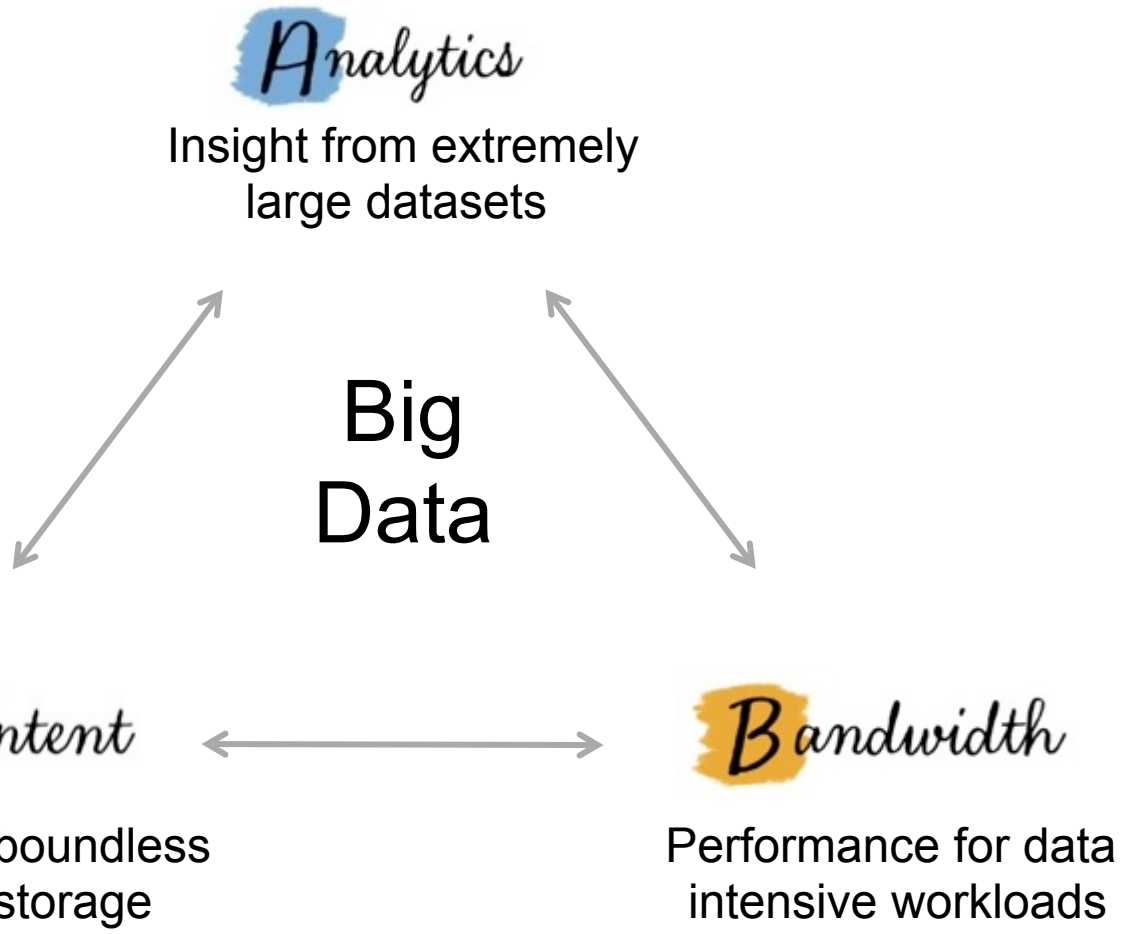

Quantifying The Big Data Challenge

60 Zettabytes

Estimated size of the digital universe in 2020

5 Billion

Smart phones

30 Billion

Pieces of new content to Facebook per month

Sensors
Video
Music
Location
Weblogs

80%
of data is unstructured

Growth Over the Next Decade:

Servers (Phys/VM):	10x
Data/Information:	50x
#Files:	75x
IT Professionals:	<1.5x

Source: *Gantz, John and Reinsel, David,* "Extracting Value from Chaos", IDC IVIEW, *June 2011, page 4.*

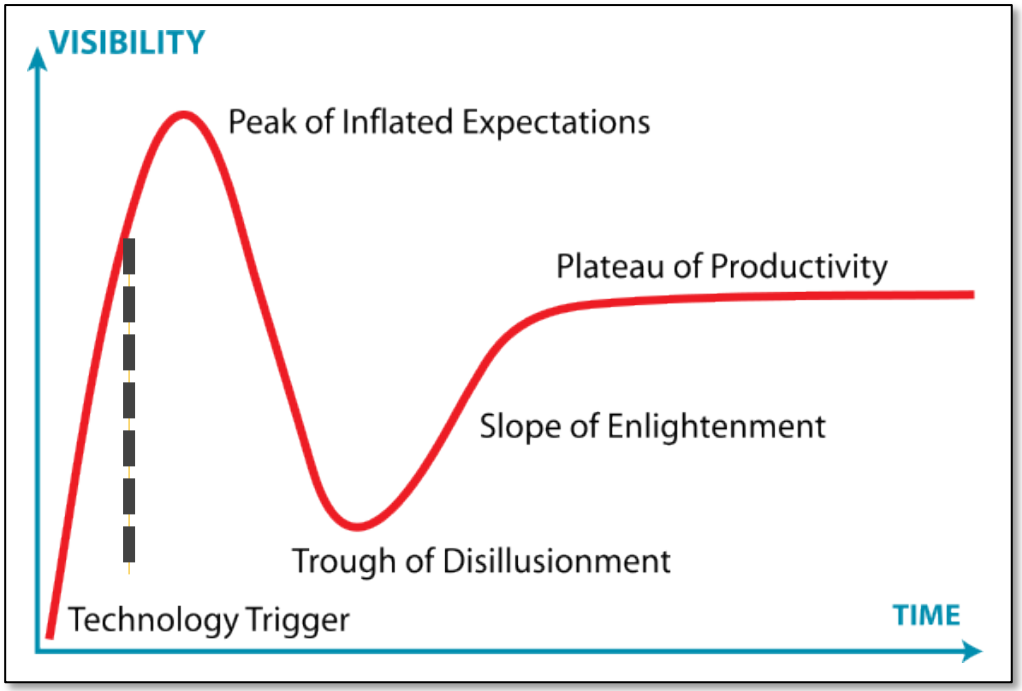

1000 PB – Some perspective…

- > 20 “Normal” Data Centers
- > 750 Million Customers
- > 48,000 Emails / Second
- > 1 Million Disks
- 24 Full Time Admins (> 40 PB/FTE)

http://wikibon.org/blog/wp-content/uploads/2012/05/big-data.v.1.0.png

Experience Managing Data at Scale

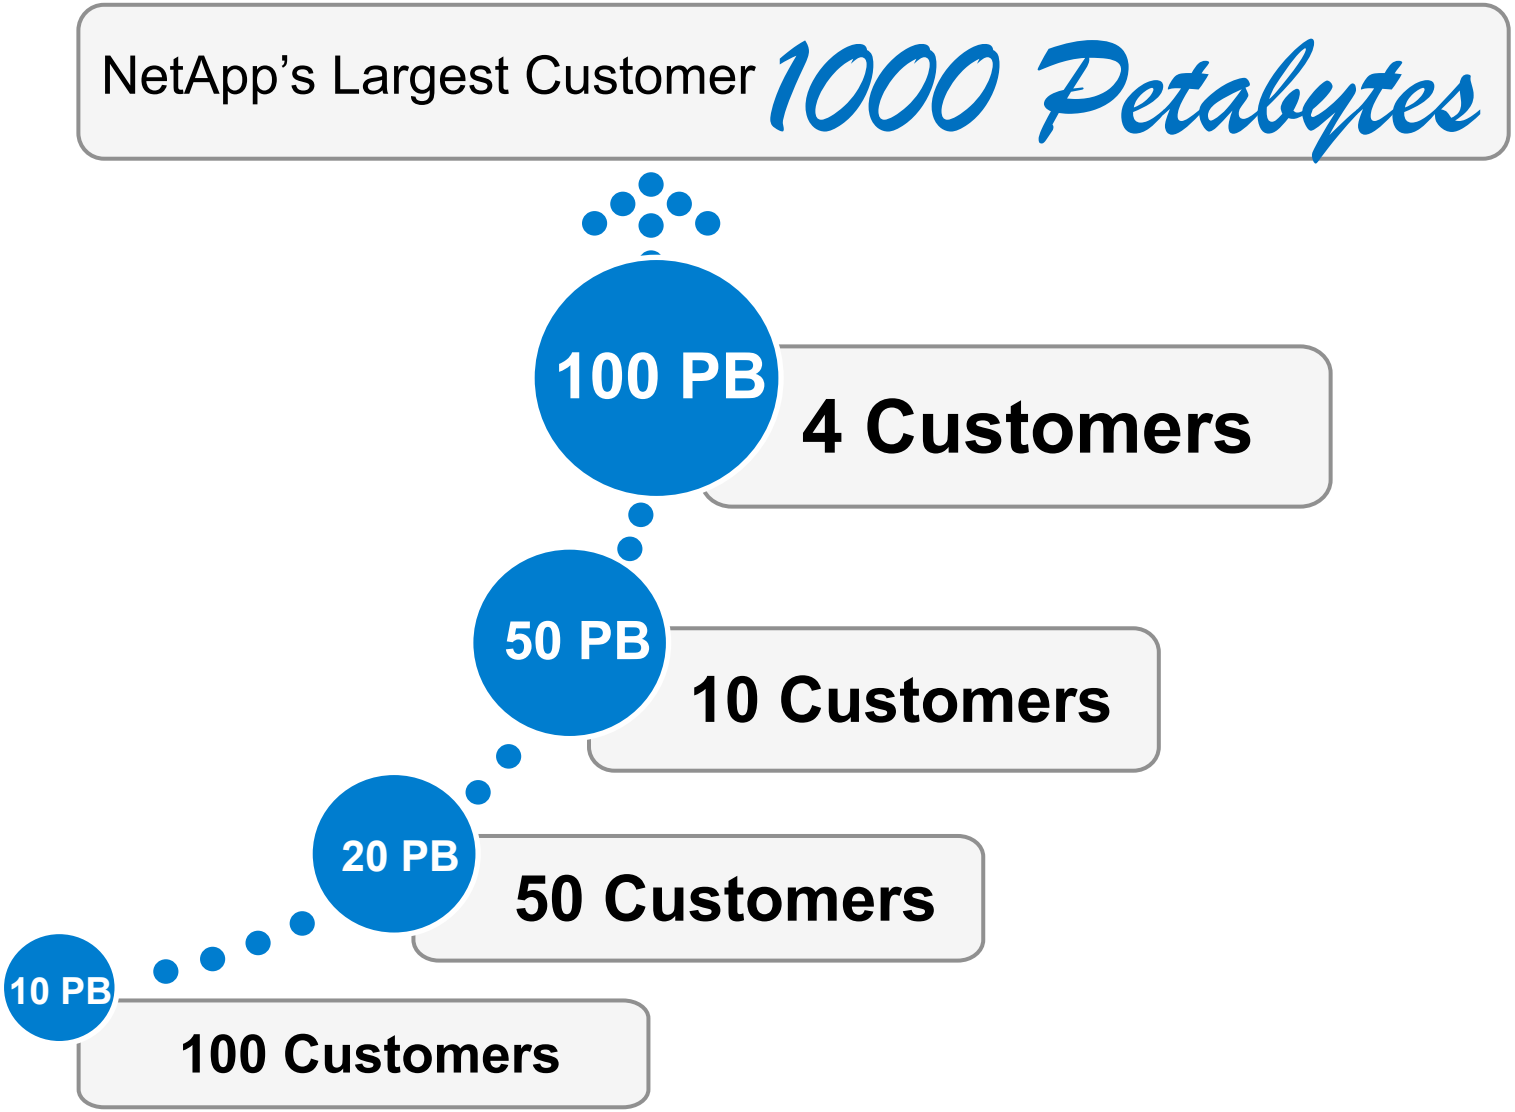

State-of-the-Science in HPC Today

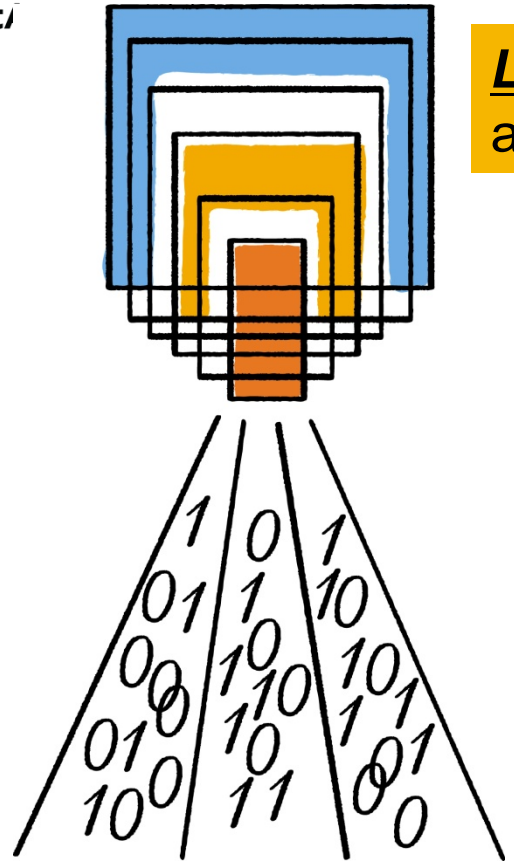

Now #1 of the Top 500 HPC Systems in the world. http://top500.org.

LLNL Sequoia – Currently the fastest supercomputer and storage combination on the planet

- 20 Peak PetaFLOPS (20 x 10^{15})
- 1.6 Million Cores
- 1.6 PB main memory
- **55 PB (usable) NetApp E-Series storage**
  - Lustre Cluster Parallel File System
  - QDR (40 Gigabit) Infiniband
  - > 1 Terabyte/sec aggregate bandwidth
- ~8 MW of Power
- Simulations for
  - Nuclear Weapons Viability
  - Counter Terrorism
  - Energy Security
  - Climate Change

Press Release: http://www.netapp.com/us/company/news/news-rel-20110928-990734.html

NetApp

Dispelling the Misconceptions About Big Data

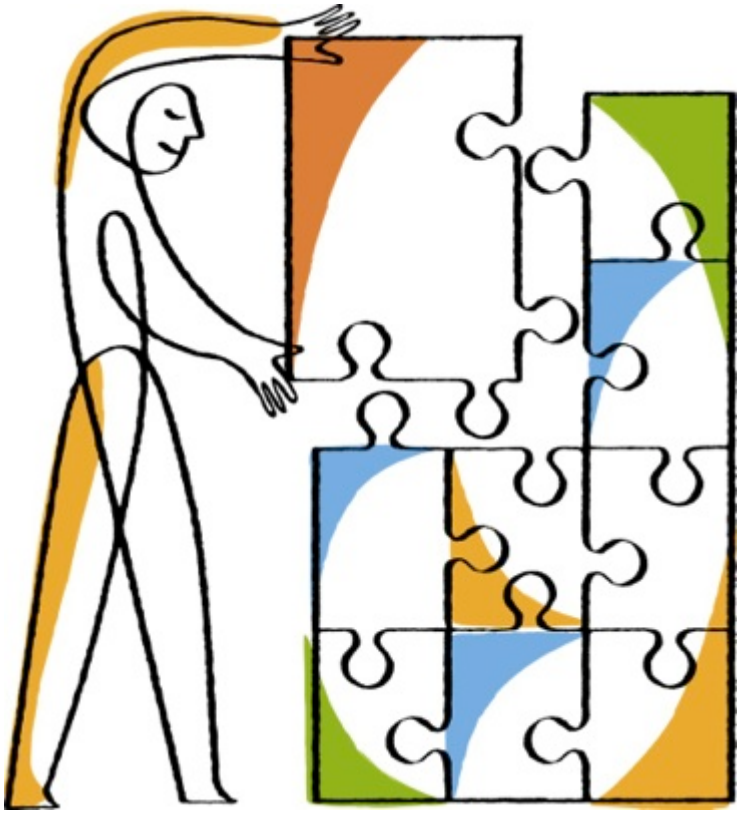

Big Data Is NOT New

National Oceanographic and Atmospheric Administration

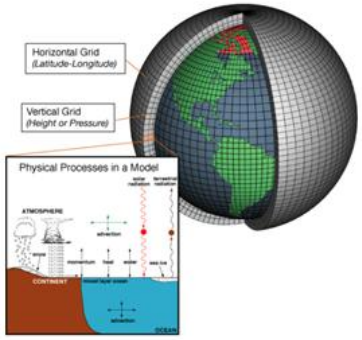

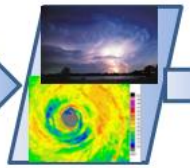

"The Tank"
24-hour
Ring Buffer

Data Assimilation System

Initial Conditions

Global Atmospheric Model (HIRAM)

HPC Models

Other High Res Specialized Models:

- Hurricane (3 Km)
- Thunderstorms
- Tornados
- Fire Weather
- Ocean Models
- Volcanic Ash
- Etc.

3.5 billion observations per day from NOAA sensors

6-Hour Update Cycle

Forecast & Warning Guidance to Public and Private Sector Forecasters

15 million information products per day


Thunderstorm Warnings

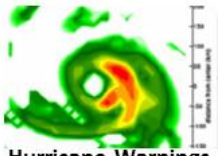
Hurricane Warnings


Flood Warnings


Fire Weather

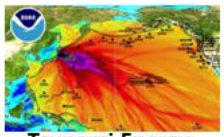
Tsunami Energy Vectors)

30 PB of New Data Annually

National Weather Service Data Stream

Big Data = Big Analytics = Hadoop?

- That's What The Media Hype Implies, but it is NOT true!
- Traditional analytics (BI/DSS/DW) dominates the analytics market
- Like other technologies vying to gain broad adoption in Enterprise IT (e.g., Traditional Analytics, HPC & Cloud), it shows promise

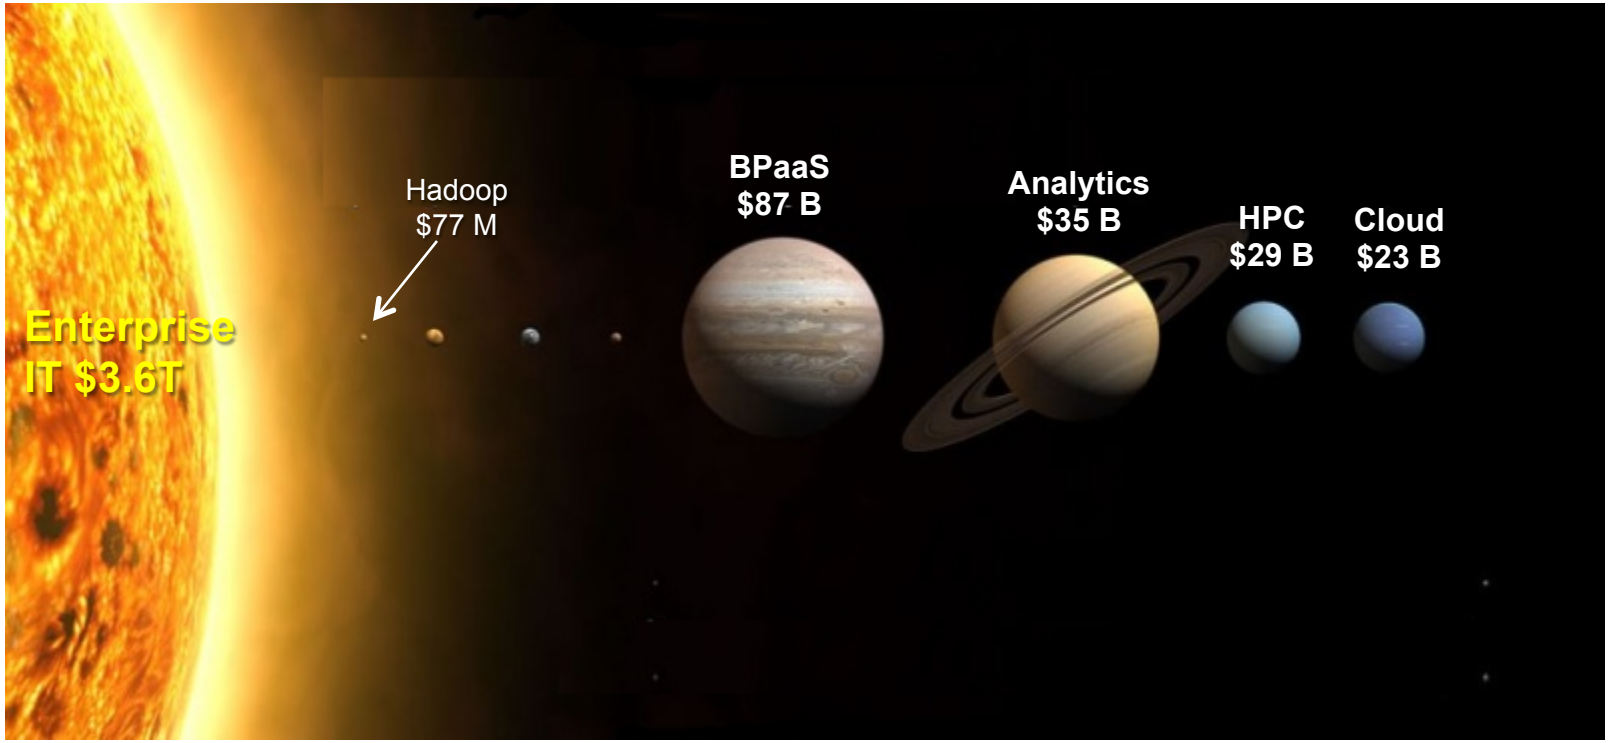

Worldwide Market Numbers

	2011	2012	2013	2014	2015	2016	CAGR	Comments
Hadoop	0.077	0.123	0.198	0.317	0.507	0.812	60%	Source: IDC - http://bit.ly/S97FcP
All Analytics	31.8	34.9	38.3	42.1	46.2	50.7	9.8%	Source: IDC - http://bit.ly/S97FcP
HPC	27.5	29.4	31.4	33.6	36.0	38.5	7.0%	Source: Intersect360 - http://bit.ly/S9bQWf
Public Cloud Services	17.0	22.5	29.4	36.6	45.0	52.9	25.5%	w/o BPaaS. Source: Gartner - http://bit.ly/S9HpPK
Public Cloud Services	91	109	130	154	180	207	17.9%	w/ BPaaS. Source: Gartner - http://bit.ly/S9LLGz
Overall IT	3500	3640	3786	3937	4095	4258	4.0%	Source: Gartner - http://bit.ly/S9LLGz
	In Billions US$							
	These are Worldwide Numbers							

NetApp

Analytics

Hadoop

- Runs on a collection of cheap, commodity servers, in a distributed, shared nothing architecture
- Two key components
  - HDFS
    - Hadoop Distributed File System
  - MapReduce
    - Programming model for processing and generating large datasets

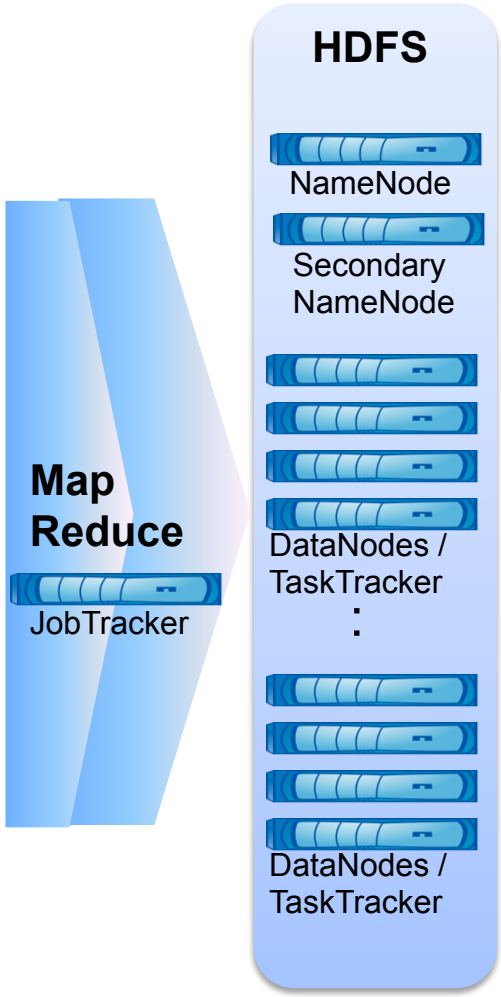

More than just Map/Reduce

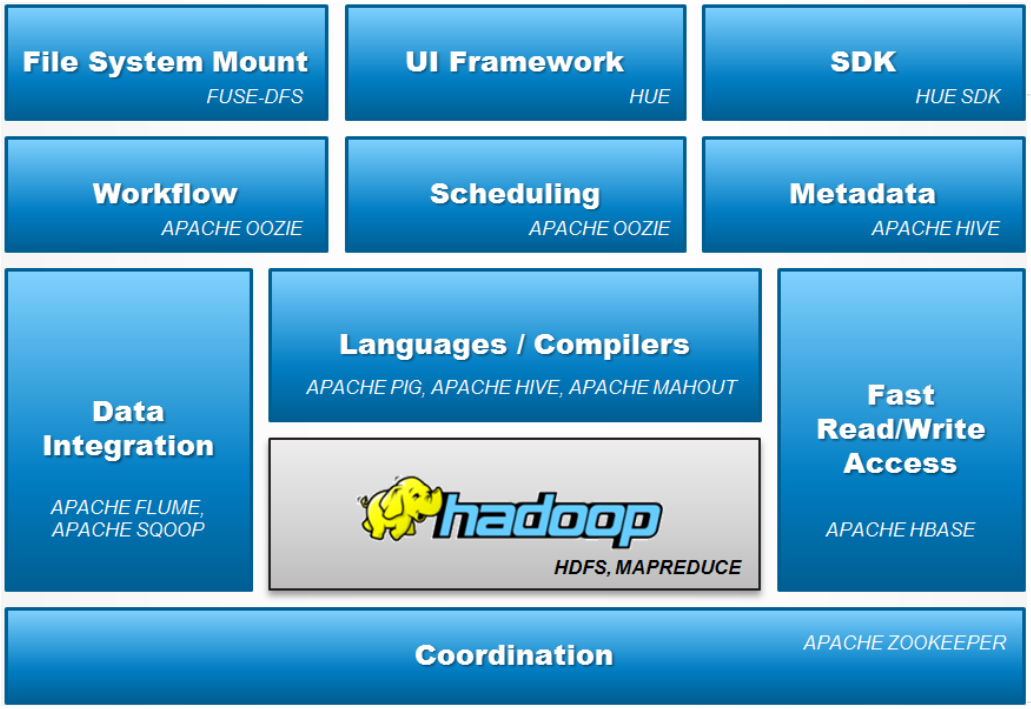

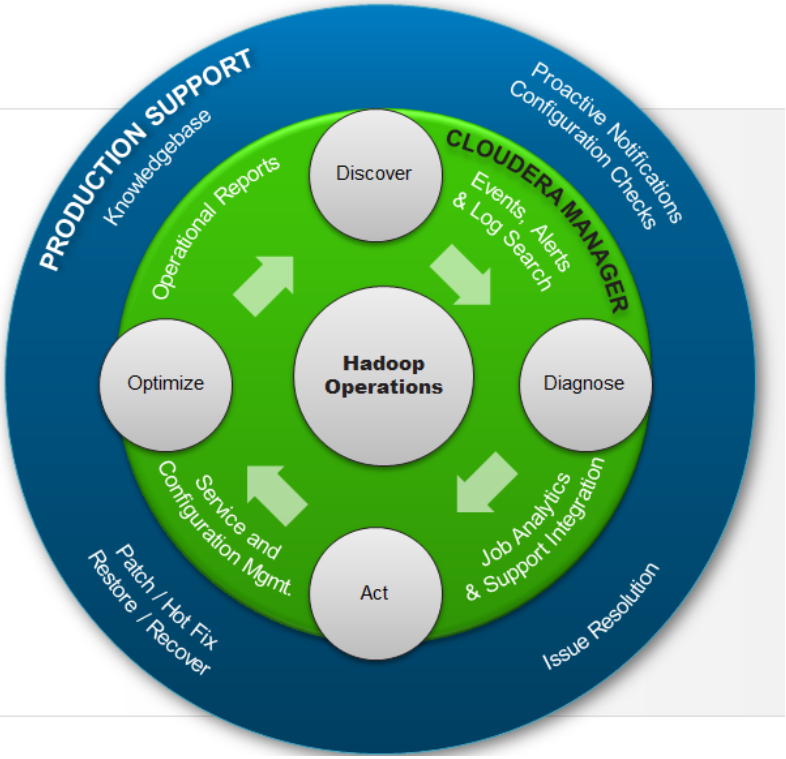

Cloudera's Distribution:

- **Enterprise Support** that gives you the expertise and responsiveness you need to run Apache Hadoop in Production
- **Management Suite** that improves quality of service, increases compliance and reduces administrative costs
- **Customer Portal** access to submit and manage cases with Subject Matter Experts
- **Knowledge Base** that your team can use to expand their Hadoop knowledge with hundreds of articles and tech notes
- **Optimized Connectors** that integrate your Hadoop clusters with existing data systems

Faster Performance on Shared Storage

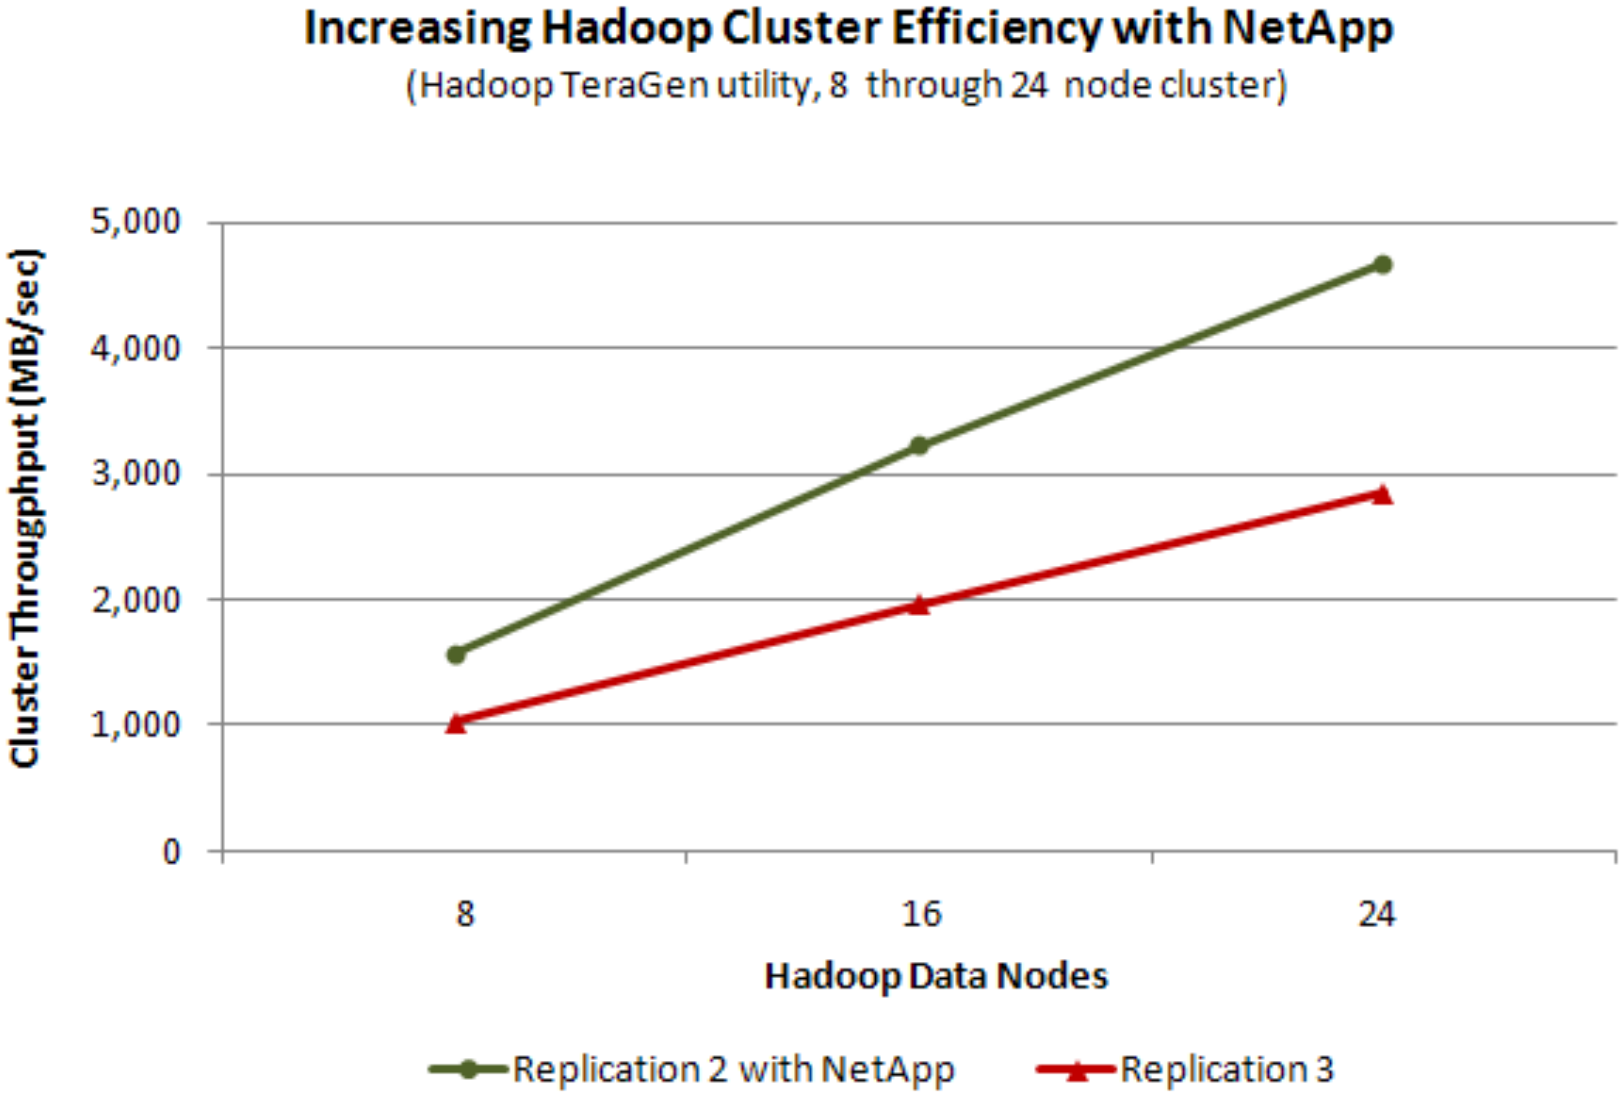

Faster Job Completion During Disk Failure

- Compared to a Hadoop cluster with the same number of nodes, but using only internal disk (no RAID) and triple copies, the NetApp Open Solution for Hadoop shows only a slight performance decrease during a disk failure with less impact to job completion.

Table 4. Drive Failure Recovery Results

Test Scenario	Job Completion Time (hh:mm:ss)	Throughput (MB/sec)	Delta (vs. Healthy Cluster)
Healthy cluster	00:30:21	1,821	N/A
NetApp E2660 drive failure	00:32:06	1,486	-5.7%
Internal data node drive failure	01:12:13	660	-237.9%

Source: Garrett, Brian and Lockner, Julie, "NetApp Open Solution for Hadoop", ESG Report, May 2012, http://bit.ly/LyYG0t

AutoSupport: Hadoop Use Case at NetApp

- "Call-home" service for all NetApp® systems
- Foundation of NetApp proactive support strategies
- Machine-generated data doubles every 16 months

CHALLENGE	NETAPP SOLUTION	BENEFITS
4 weeks to run a query on 24 billion unstructured records		**Time reduced from 4 weeks to 10.5 hours**
Impossible to run a query: 240 billion unstructured records	**10-node Hadoop Cluster w/ shared Storage**	**Previously impossible, now achievable in just 18 hours**

"NetApp ASUP is a mission-critical application"

Analytics Roadmap

- Traditional & Big Analytics side-by-side for years to come
- Hadoop moves to shared, virtualized infrastructure, for better efficiency and ease of management:
  - Hadoop remains logically distributed, shared nothing, but runs on a virtualized shared everything architecture (e.g., FlexPod for Vmware + eSeries)
  - Same as above, except Hadoop becomes logically shared everything, as HDFS is replaced by a parallel file system (e.g., Lustre Cluster, StorNext or GPFS)
- Enterprise class resiliency (no SPoF) and reliability with HPC-like performance (no need for triplicas)
- Use of a single copy of data for the map phase (higher storage utilization)
- Natural intersection with Cloud (Analytics as a Service)

NetApp

Bandwidth

Big Bandwidth Use Cases

Full Motion Video/ISR

Scalable density and performance to ingest and simultaneously analyze UAV, satellite and other data

Video Surveillance

High bandwidth & density supporting hundreds or thousands of HD cameras

Media Content Management

High ingest & play-out rates with support for media and entertainment workflows

HPC & Seismic

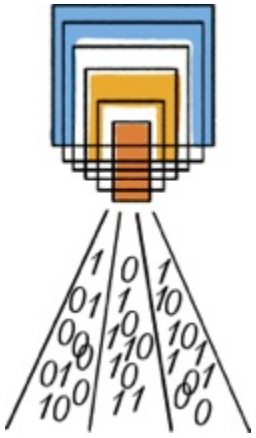

Massively parallel distributed file system for large scale cluster computing and O&G Seismic Processing

Bandwidth Driven by HPC

- THE driver for innovations in bandwidth are clearly high performance computing (HPC) applications.
- HPC falls into two camps:
  - High Performance Technical Computing (HPTC),
    - Supports scientific and engineering modeling and simulation.
    - HPTC makes up about 70% of the HPC market
    - Checkpoint/Restart
  - High Performance Business Computing (HPBC)
    - Financial Services, Media & Entertainment, Complex Event Processing, business process optimization and ultrascale business computing.

HPC Roadmap

- National Labs are already hard at work architecting the next system, pushing technologies to their limits
- Next stop: Exascale (10^{18} FLOPS) in 2016-2018
- This will require 50 – 100X Sequoia
  - 80-160 million CPU cores
  - 80-160 PB's of main memory
  - 266 – 533 TB/sec burst of sequential I/O from main memory during 5 minute checkpoint
- > 50 MW of power
- Existing storage technologies won't scale within bandwidth, power and budget constraints

NetApp

Content

Trends Shaping the Big Content Market

- Data at scale
  - Objects replace files to overcome filesystem limitations
  - Containers 10's PB in size, 10's billions of objects, Millions of users
  - "Life time retention" of enterprise and consumer data
- Policy driven infrastructure - "Set it and forget it"
  - Object/VM granular automated data mobility/management
- Simpler, efficient, self-managing systems; low price
  - Systems with lower performance & resiliency for less $/GB
  - Space, power, density @ lower infrastructure cost (~2 PB/rack)
- Centralized → distributed, cloud-friendly "system"
  - Multi-location repositories, data mobility, local view/access

Distributed, Object-Based Repository

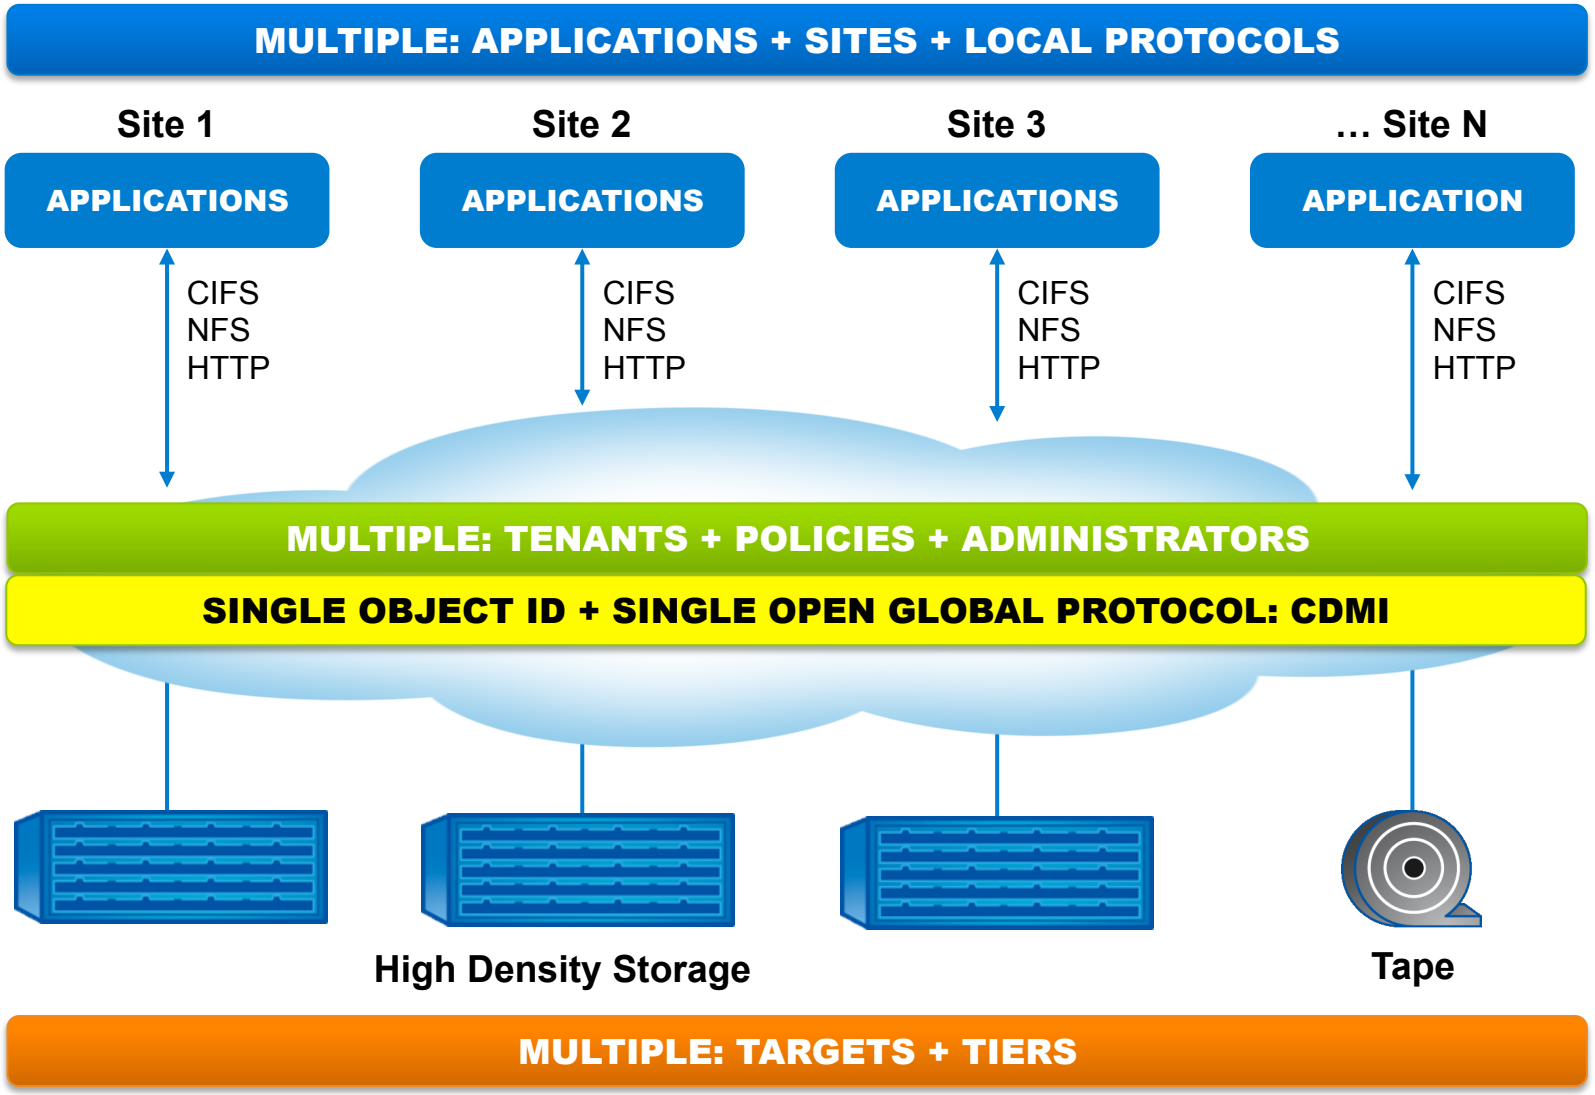

Distributed Content Repository Roadmap

- Federated CDMI (Cloud to Cloud)
- Strong ecosystem of ISVs supporting CDMI
- Exabyte size containers
- Trillions of objects
- Rack densities > 5 PB/rack
- Power densities < 1 kW / PB
- Costs < $150,000 / PB ($0.15 / GB)

Pre-configured, Certified Solutions

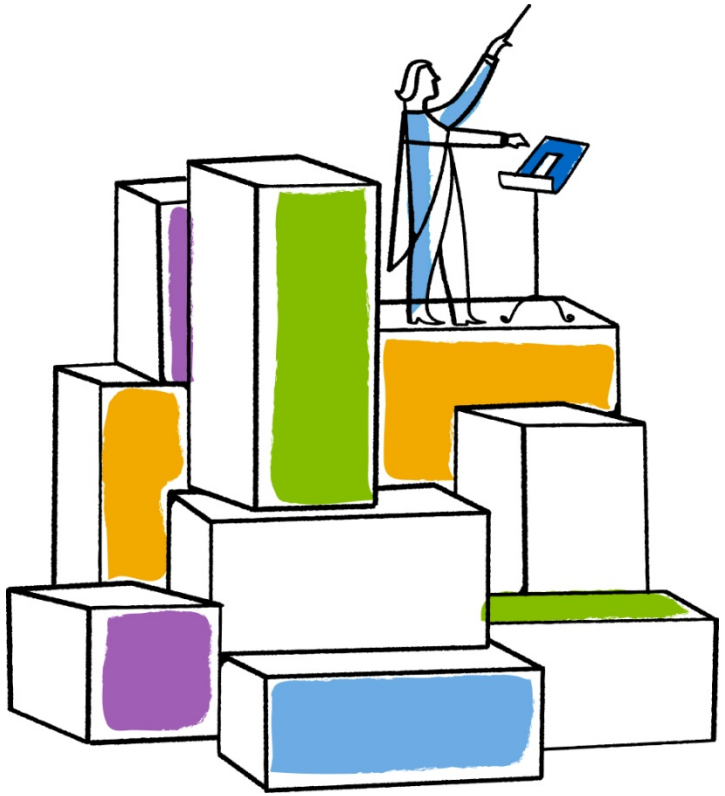

NetApp E-Series Storage Solutions

Dense and Modular Storage to Meet Any Scale

- Pre-Configured, Tested, Certified Solutions
- NetApp E-Series Hardware
- NetApp Professional Services
- NetApp Support Services

60 drives	120 drives	180 drives	240 drives	300 drives	360 drives
180TB	360TB	540TB	720TB	900TB	1080TB

E-Series Modular Design

NetApp

Disk Shelves

DE6600

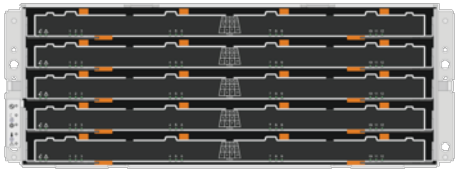

- 4U/60 x 3.5” drives
- Highest throughput
- Largest capacity/density
- 3TB NL-SAS drives

- Modular packaging architecture for maximum configuration flexibility
- Two optimized controller families

Controllers

- 8Gb/s FC, Infiniband
- 12/24GB Cache per system
- 24.8Gb/s writes actual performance
- High Bandwidth / High IOPs

5400

2600

- FC or iSCSI
- 4GB Cache per system
- 11.2Gb/s writes actual performance
- Cost effective performance

Systems

E-5460

E-2660

Two Sample Configurations:
E-Series 2660 and 5460

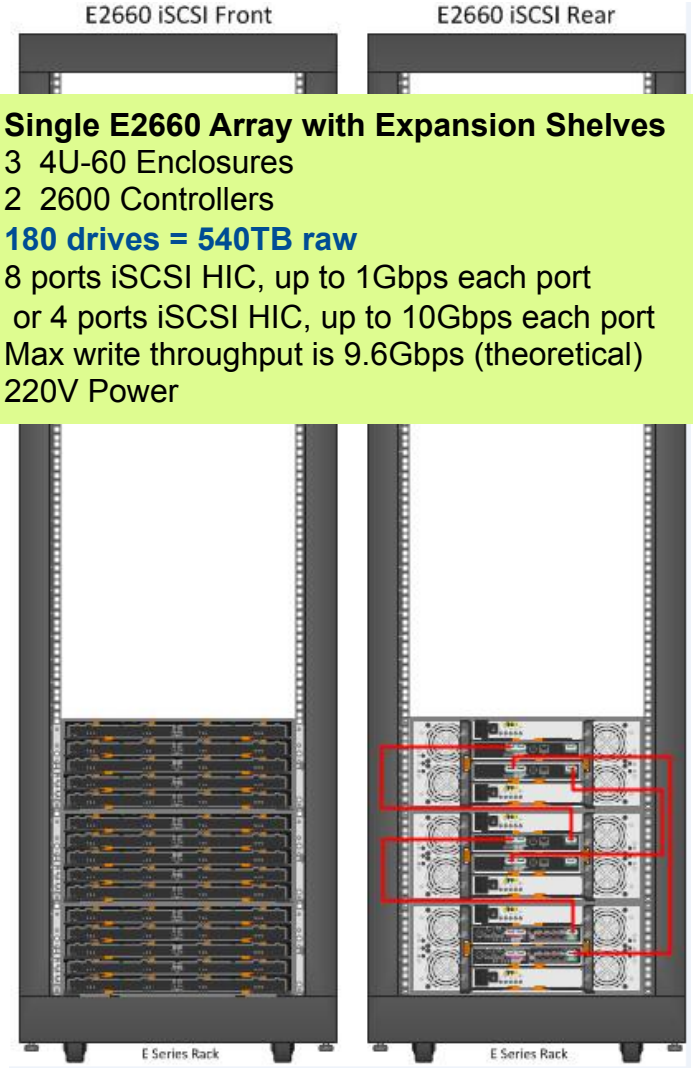

Single E2660 Array with Expansion Shelves
3 4U-60 Enclosures
2 2600 Controllers
180 drives = 540TB raw
8 ports iSCSI HIC, up to 1Gbps each port
or 4 ports iSCSI HIC, up to 10Gbps each port
Max write throughput is 9.6Gbps (theoretical)
220V Power

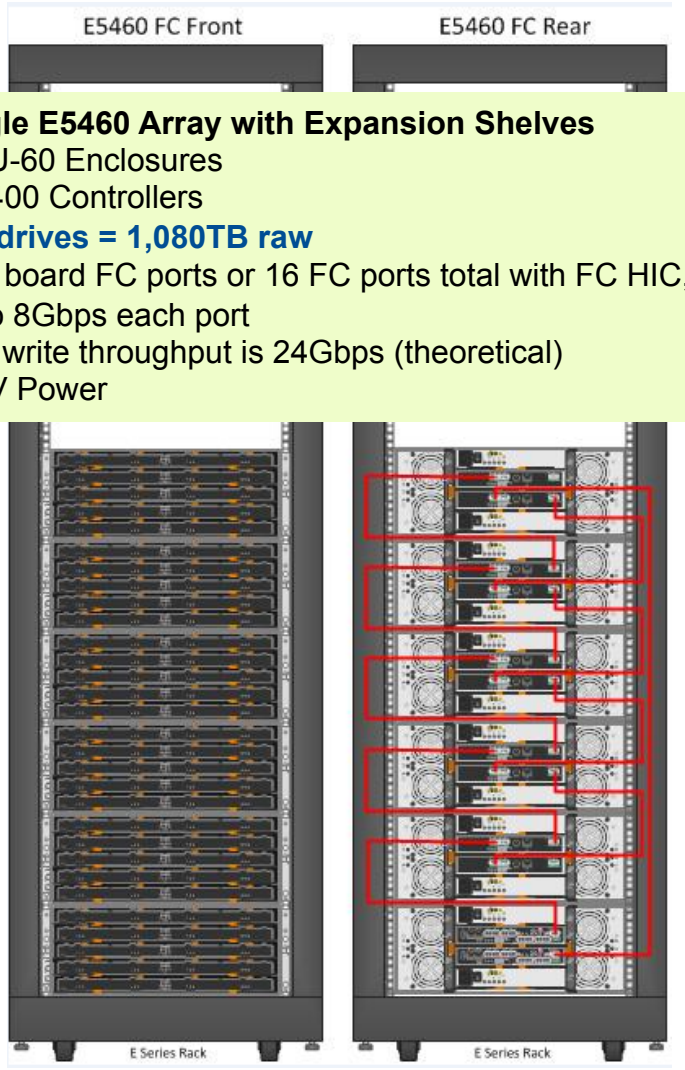

Single E5460 Array with Expansion Shelves
6 4U-60 Enclosures
2 5400 Controllers
360 drives = 1,080TB raw
8 on board FC ports or 16 FC ports total with FC HIC, up to 8Gbps each port
Max write throughput is 24Gbps (theoretical)
220V Power

Summary

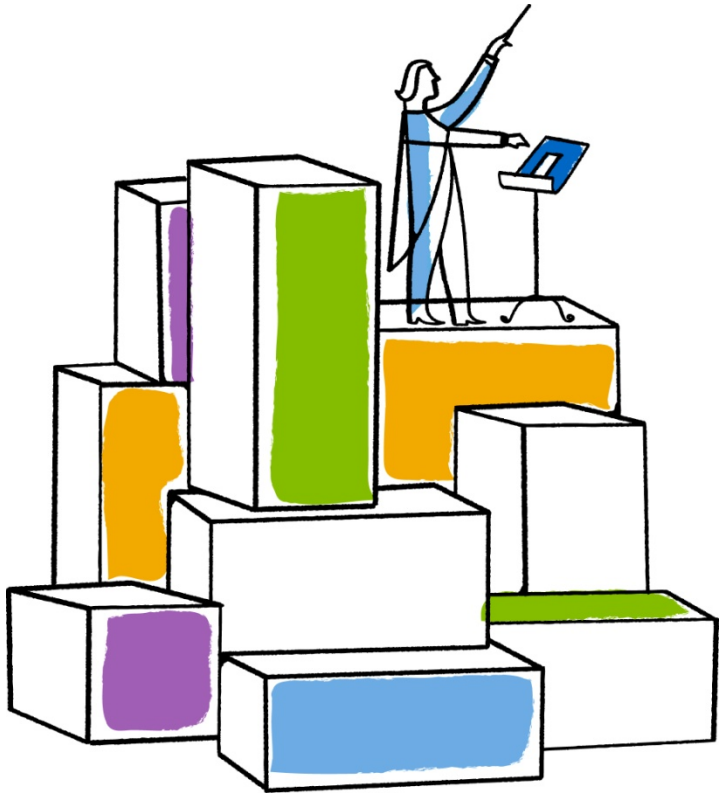

So, What Does All This Mean to You?

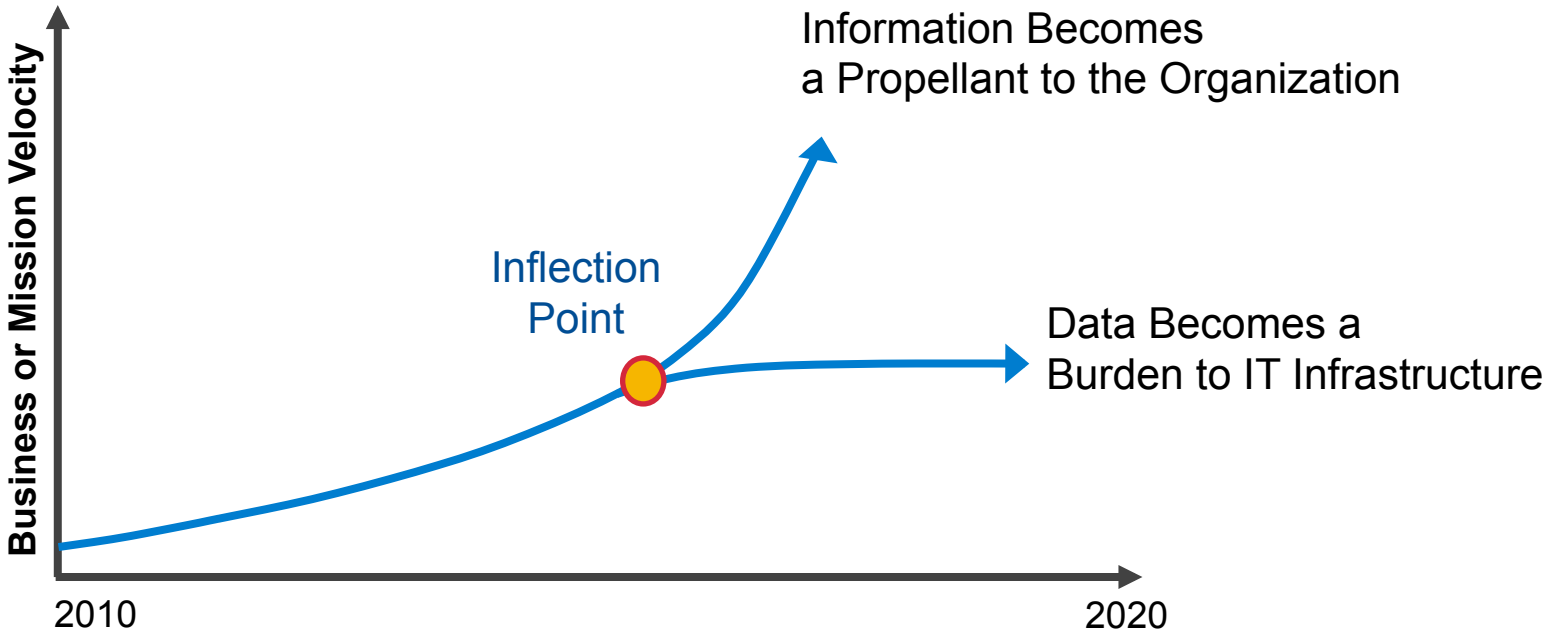

You are also at an Inflection Point: You also have a decision to make, as “business as usual” may not cut it!

Summary

- Despite the hype, **Big Data is not new and is more than just analytics! (**Many agencies and private companies have struggled with Big Data for decades)
- **Analytics:** Traditional BI/DSS analytics still dominate. Importance of newer NoSQL & Columnar DB applications, enabled by MapReduce will grow with the growth of multi-structured data
- **Data bandwidth** will be driven by HPC Exascale initiatives
- Continued data growth that outstrips network bandwidth growth and file system scalability will drive low cost, geographically distributed, object-based, **Content Repositories** with federated access through CDMI
- Big Data applications, such as Hadoop, will need to adopt shared, virtualized infrastructure (and its management benefits) if they are to be widely adopted by Enterprise IT

NetApp

Q&A

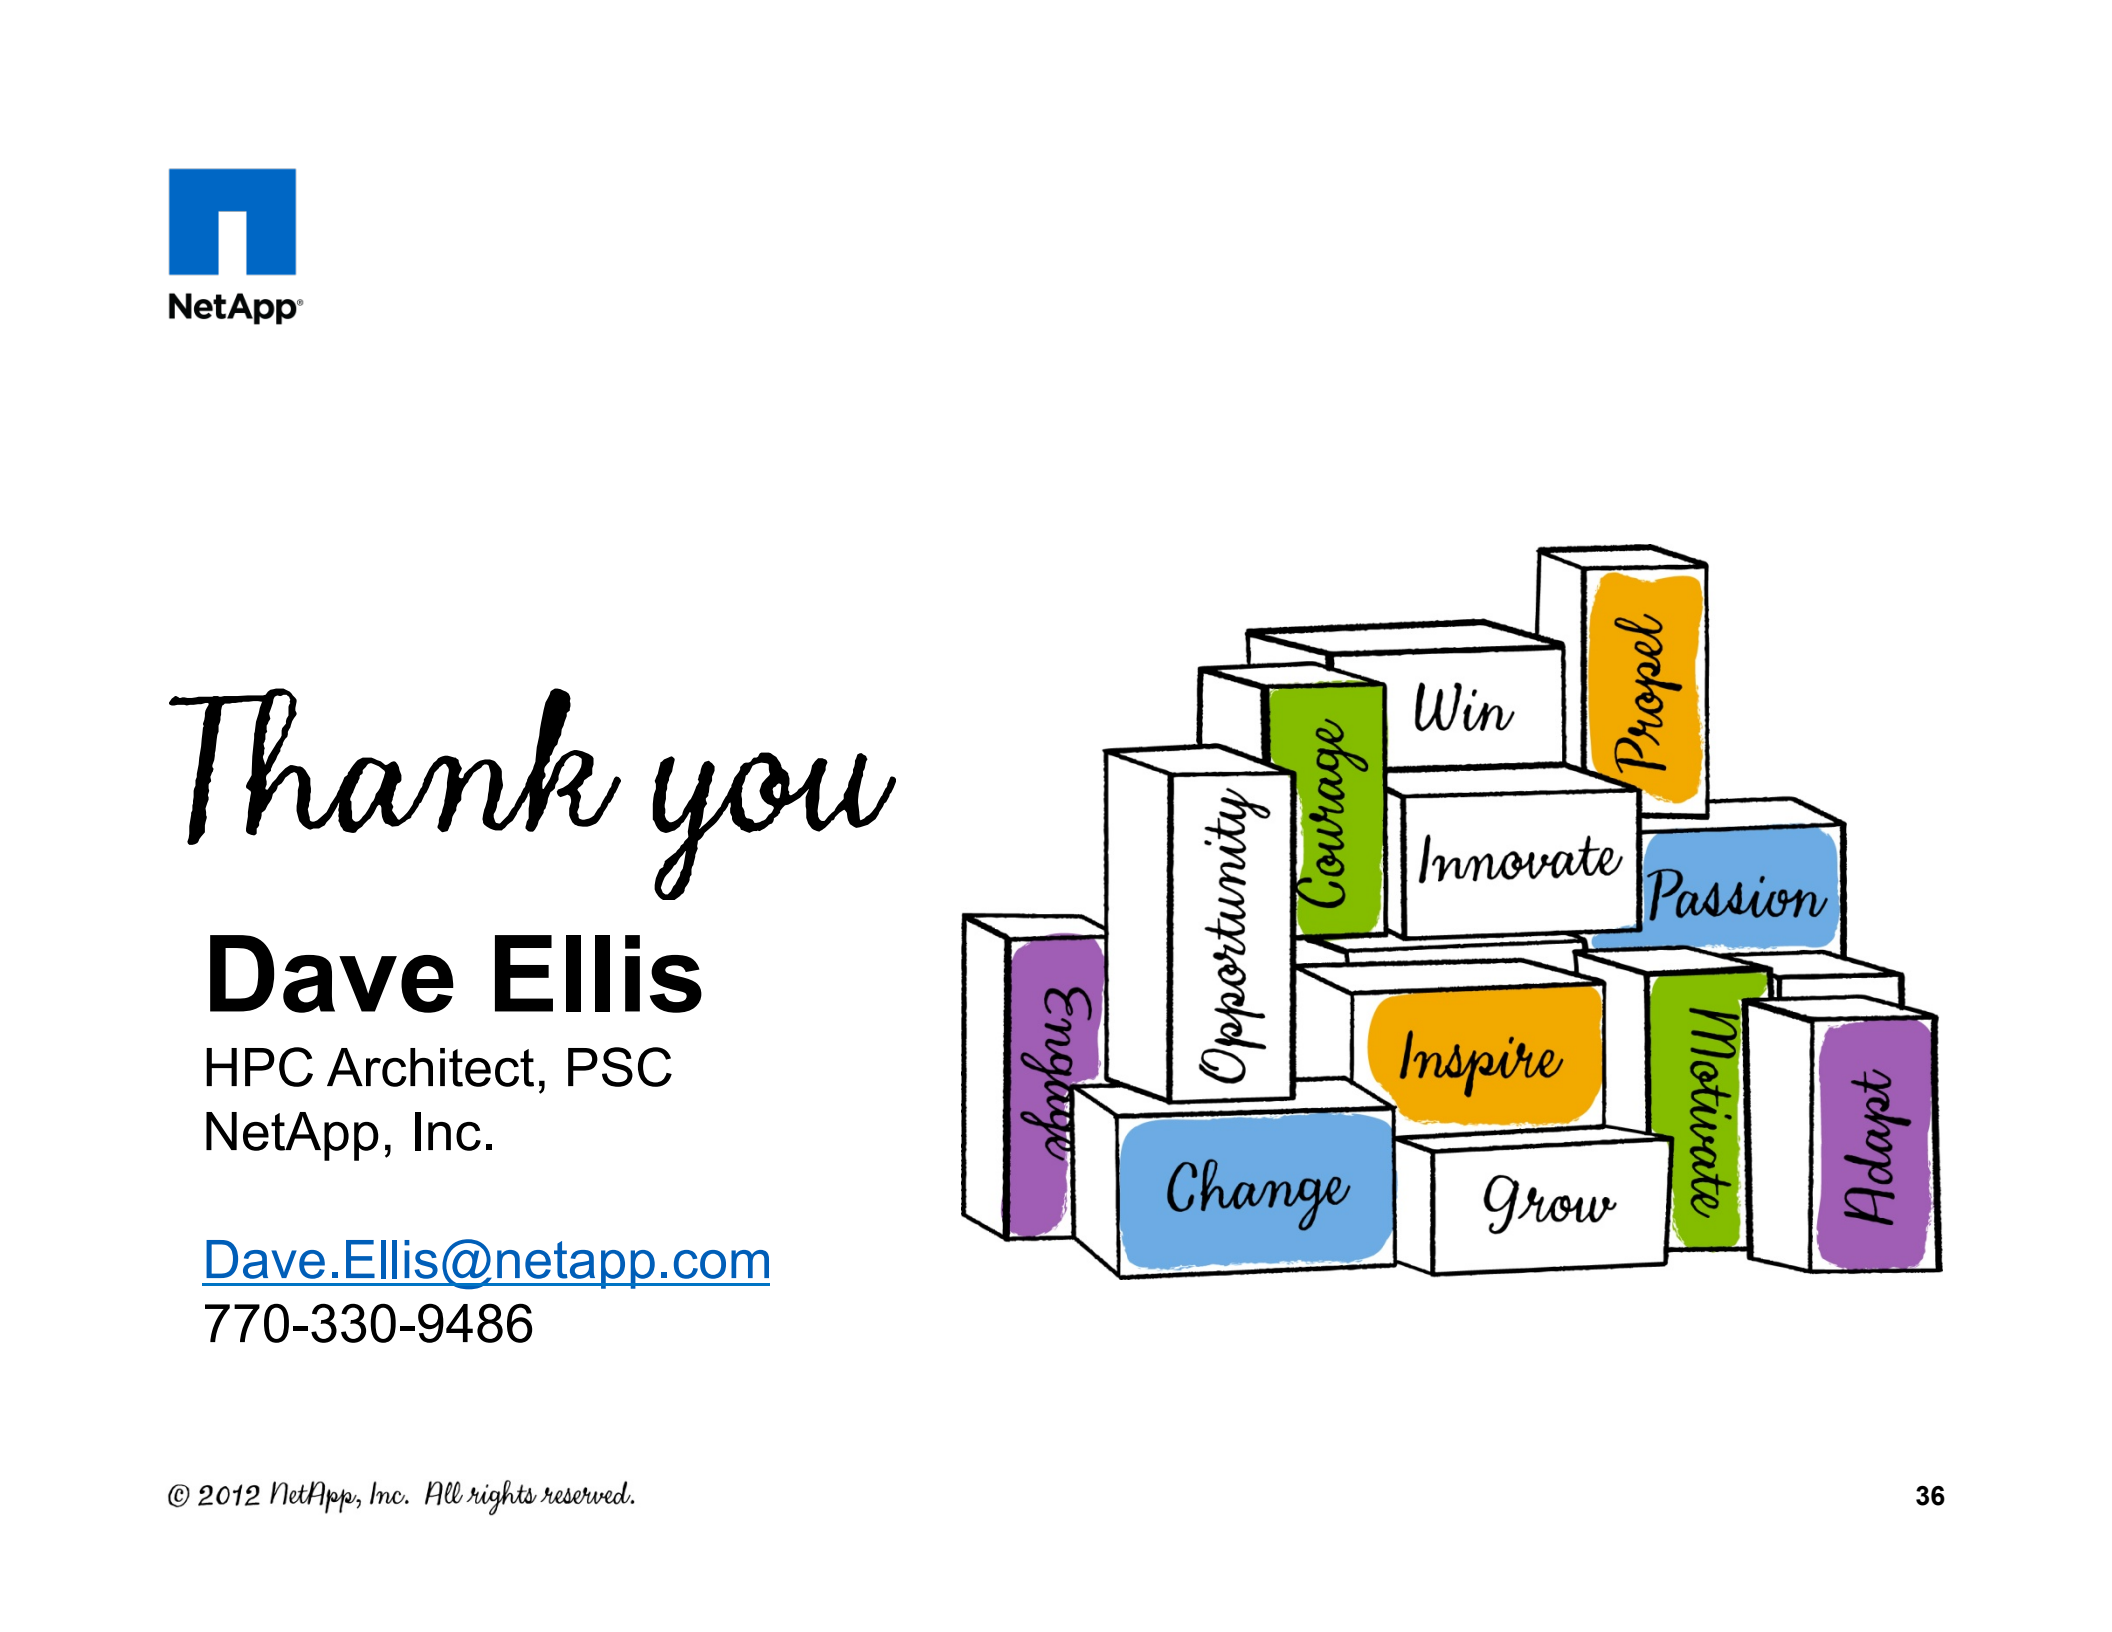

Thank you

Dave Ellis

HPC Architect, PSC
NetApp, Inc.

Dave.Ellis@netapp.com
770-330-9486