IDC SC13 Breakfast Briefing

November 2013

Earl Joseph, IDC
ejoseph@idc.com

Steve Conway
sconway@idc.com

Chirag Dekate
cdekate@idc.co

IDC Has >1,000 Analysts In 52 Countries

IDC's HPC Team

Earl Joseph

IDC HPC research studies, HPC User Forum and strategic consulting

Steve Conway

Strategic consulting, HPC User Forum, market trends, Big Data

Chirag Dekate

HPC QView, technology trends, Big Data, innovation awards program

Lloyd Cohen

HPC market analysis, data analysis and workstations

Mike Thorp

Government account support and special projects

Kurt Gantrish

IDC Government Insights

Charlie Hayes

Government HPC issues, DOE and special studies

Mary Rolph

HPC User Forum conference planning and logistics

About IDC:
IDC HPC Activities

- **Track all HPC servers sold each quarter**
- **4 HPC User Forum meetings each year**
- **Publish 45 plus research reports each year**
- **Visit all major supercomputer sites & write reports**
- **Assist in collaborations between buyers/users and vendors**
- **Assist governments in HPC plans, strategies and direction**
- **Assist buyers/users in planning and procurements**
- **Maintain 5 year forecasts in many areas/topics**
- **Conduct special research studies**

Presentation Topics

- **General economic trends**
- **HPC market results and forecasts**
  - Early results from 30,000 micro HPC surveys
- **New IDC Pulse panel**
- **High Performance Data Analysis (HPDA) update**
- **New ROI report and model**
- **Multiclient study highlights**
- **HPC User Forum update**
- **Innovation award winners**

General Economic Trends

Worldwide IT Spending (% growth, total IT spending)

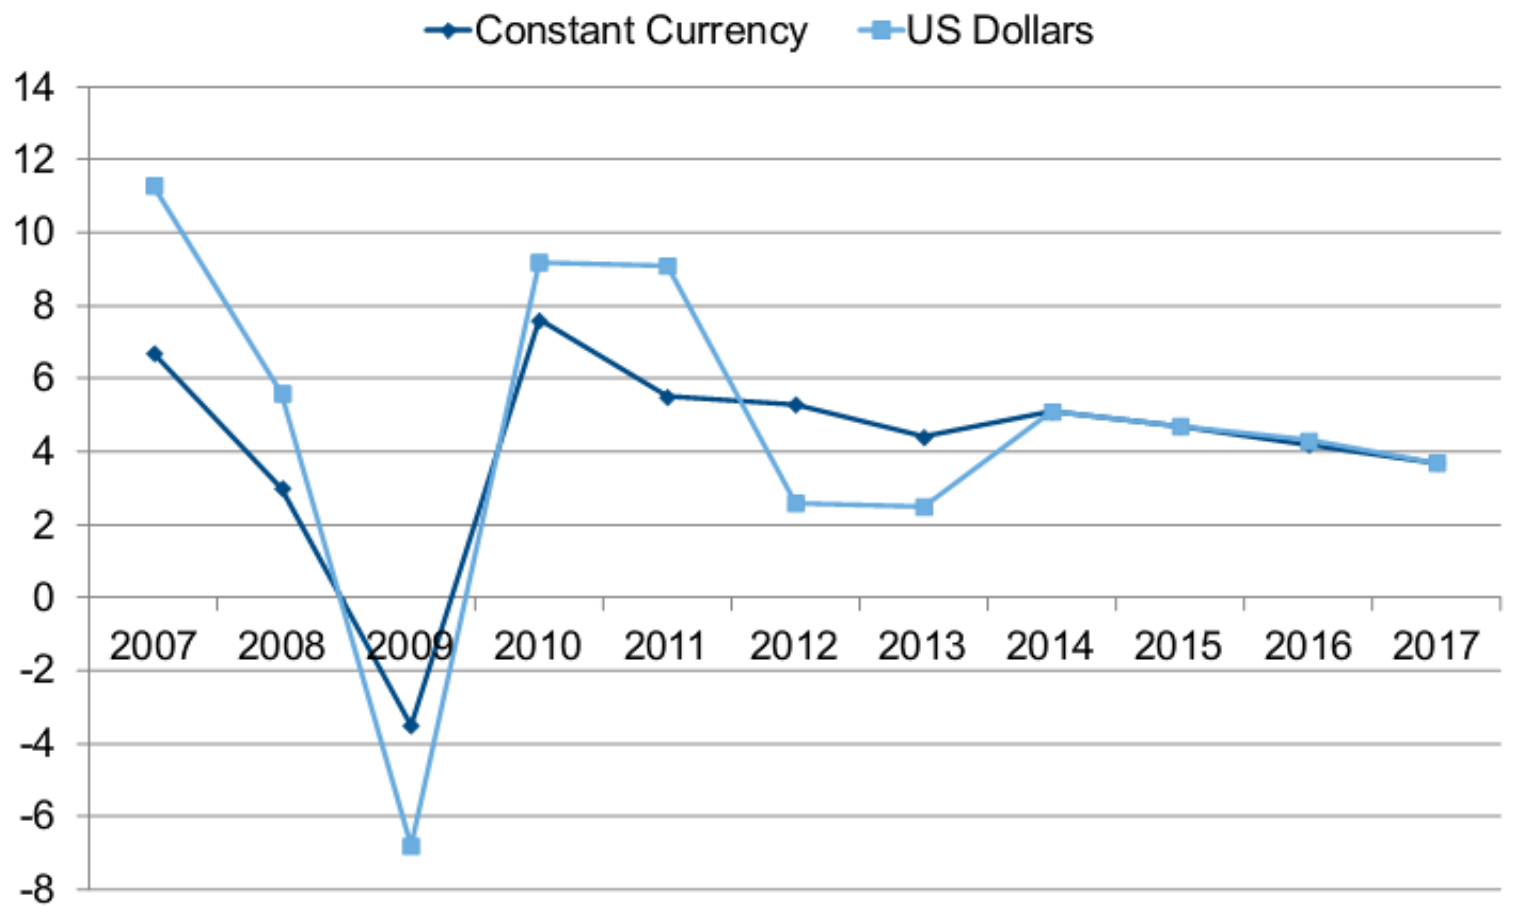

Source: IDC Worldwide Black Book, Q3 2013

General Economic Trends

Worldwide IT Spending Growth

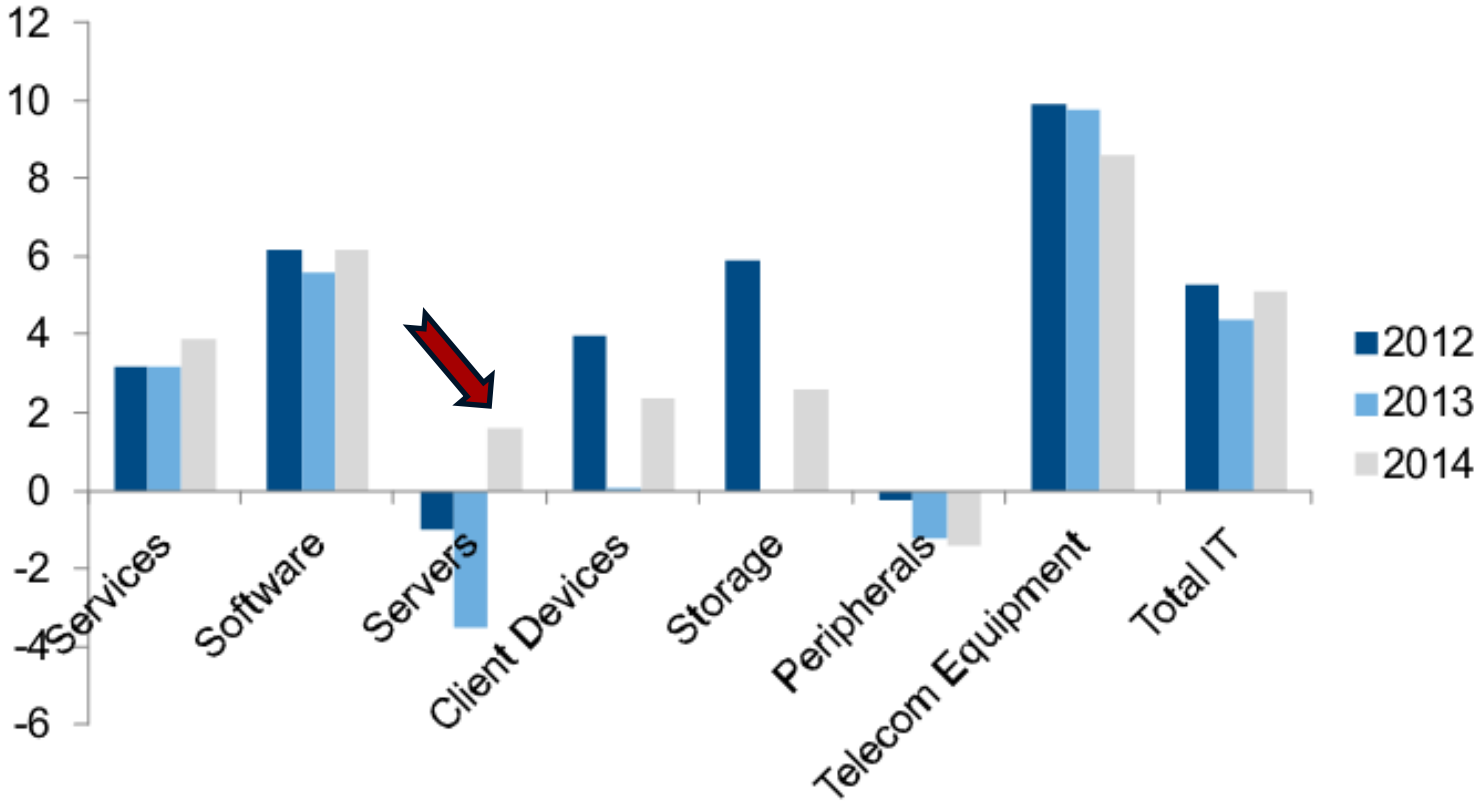

Source : IDC Worldwide Black Book Q3 2013; growth in constant currency

General Economic Trends

Regional IT Growth (%)

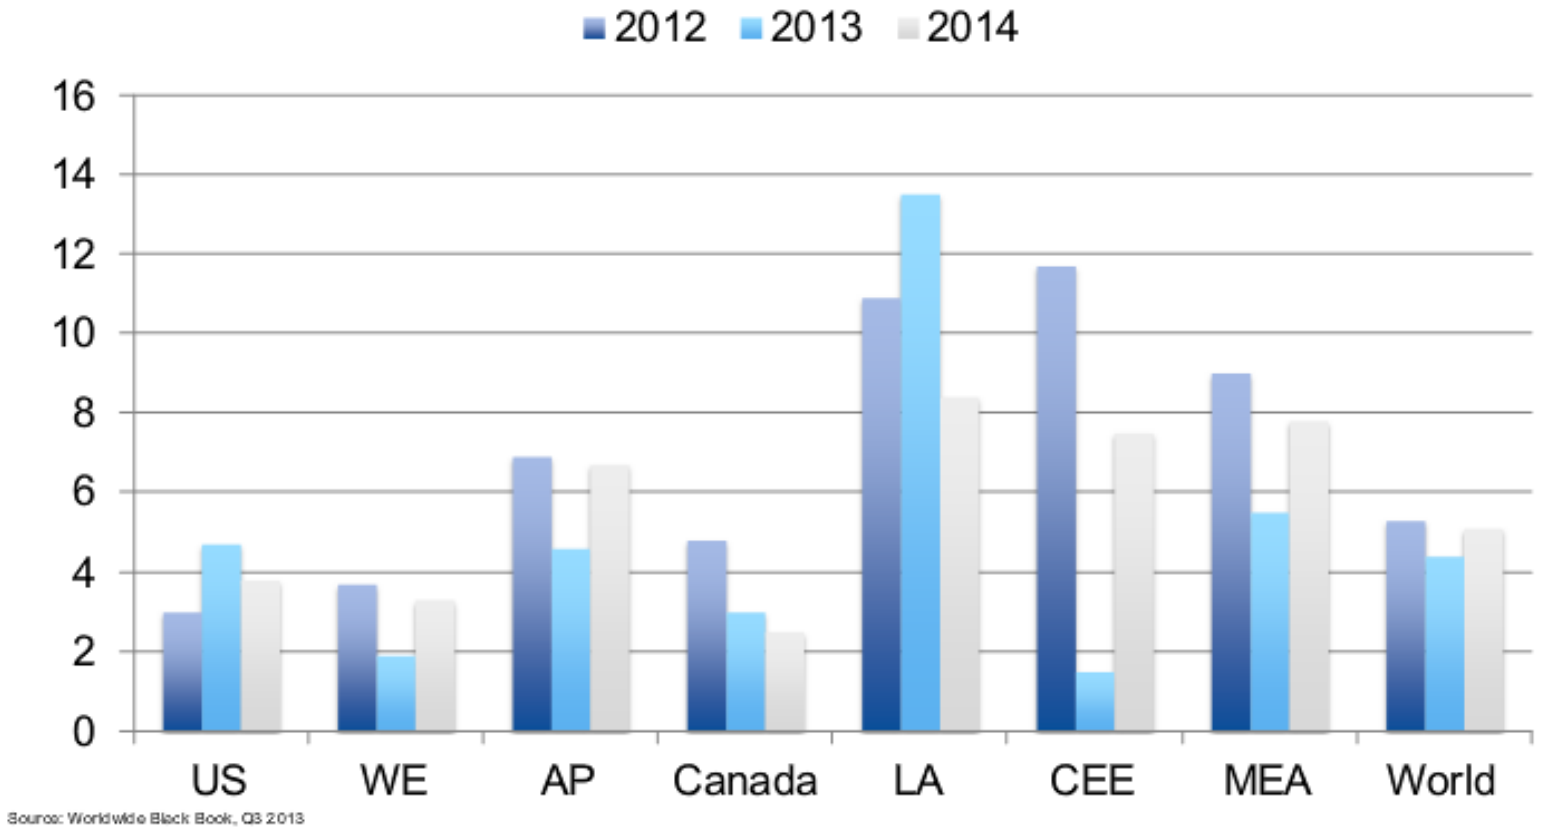

Source: Worldwide Black Book, Q3 2013

General Economic Trends

BRIC GDP Growth (% Real)

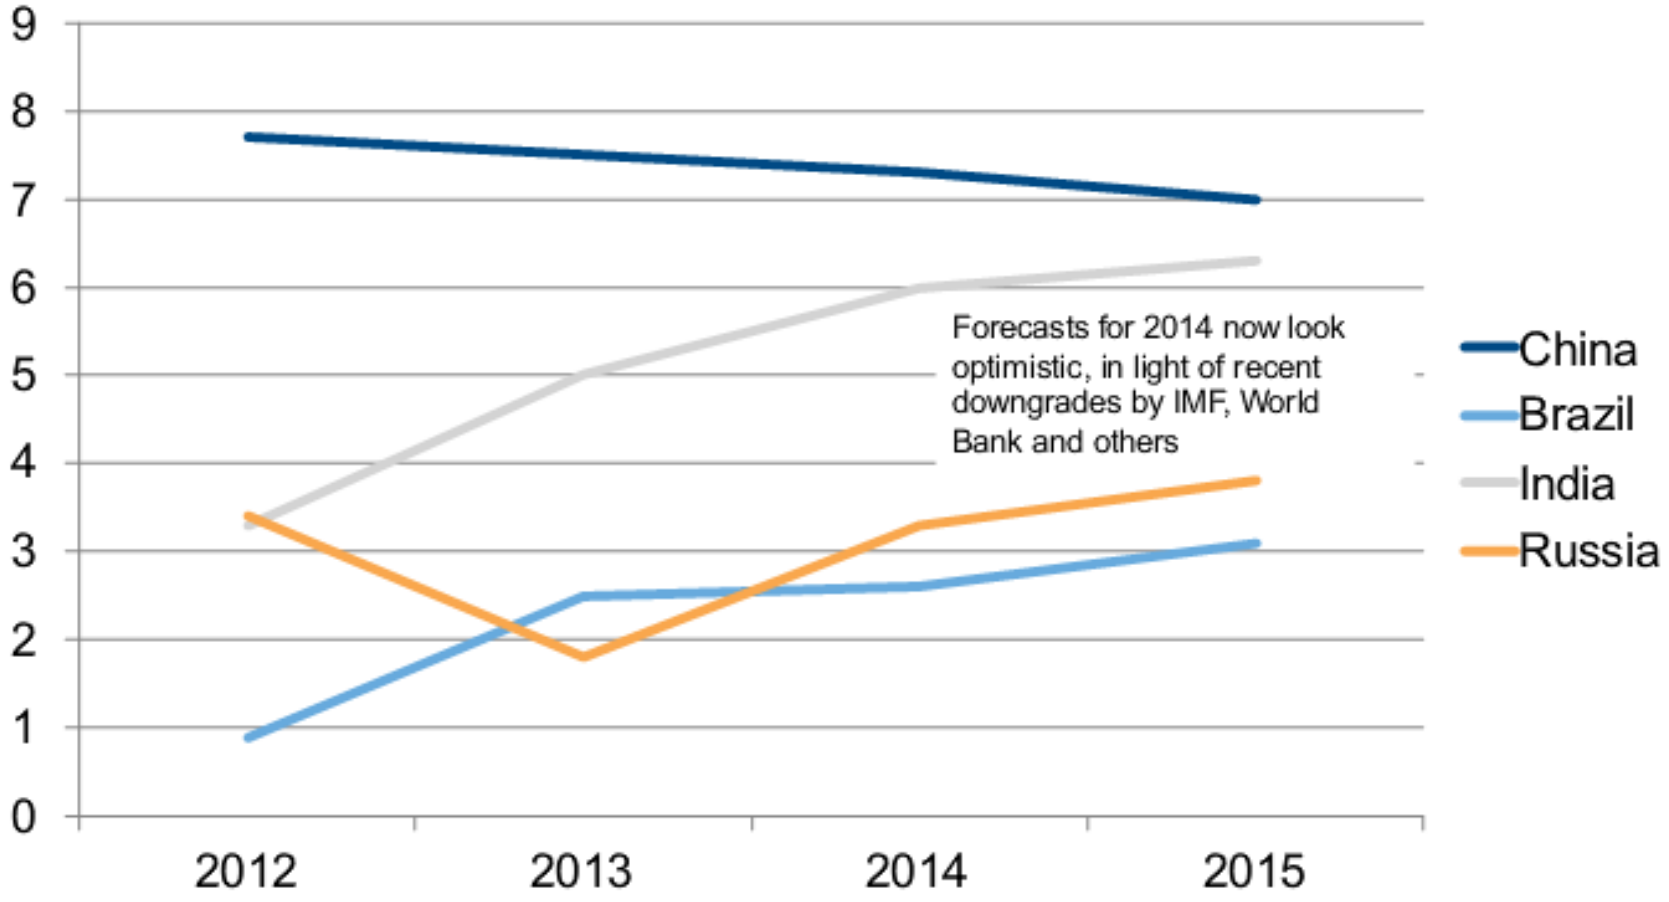

HPC Market Update and Trends

Top Trends in HPC

The global economy in HPC is growing again:

- 2011: grew 8.4% to reach $10.3 billion
- 2012: HPC revenue exceeded $11B
- **From our micro-surveys → 1Q13 -- The lower half of the market is growing well again**

Ongoing challenges for datacenters

- Power, cooling, real estate, system management
- Storage and data management continue to grow in importance

Software hurdles continue to grow

The worldwide Petascale/Exascale Race is at full speed

Big Data and accelerators are hot new technologies

HPC Cloud computing is growing slowly, steadily

2014 Hot Topics in HPC: IDC HPC Research Areas

IDC Analyze the Future

Potentially Disruptive Trends and Technologies

- Growing need to demonstrate/quantify ROI and innovation
- Rise in industrial partnership programs
- Proliferation of High Performance Data Analysis (Big Data using HPC)
- New processors, co-processors and accelerators
- New memory & I/O solutions (flash/SSDs, in-memory processing, etc.)
- New software solutions
- Government programs to help bring to market new capabilities

Special HPC Research Areas & Reports:

- Updating of IDC HPC competitive segments
- End-user based MCS reports: clusters, processors, accelerators, storage, interconnects, system software, and applications
- Emerging markets including China, Russia, Korea, Latin America, etc.
- Pioneering research on ROI from investments in HPC
- SMB and SMS research
- The HPC Innovation Award program
- The evolution of clouds in HPC
- Scaling of software – issues and solutions
- Worldwide exascale initiatives

Research Overview – Parameters Being Collected to Tie to Broader Economic Reports

For each sector we need 4 basic ratios

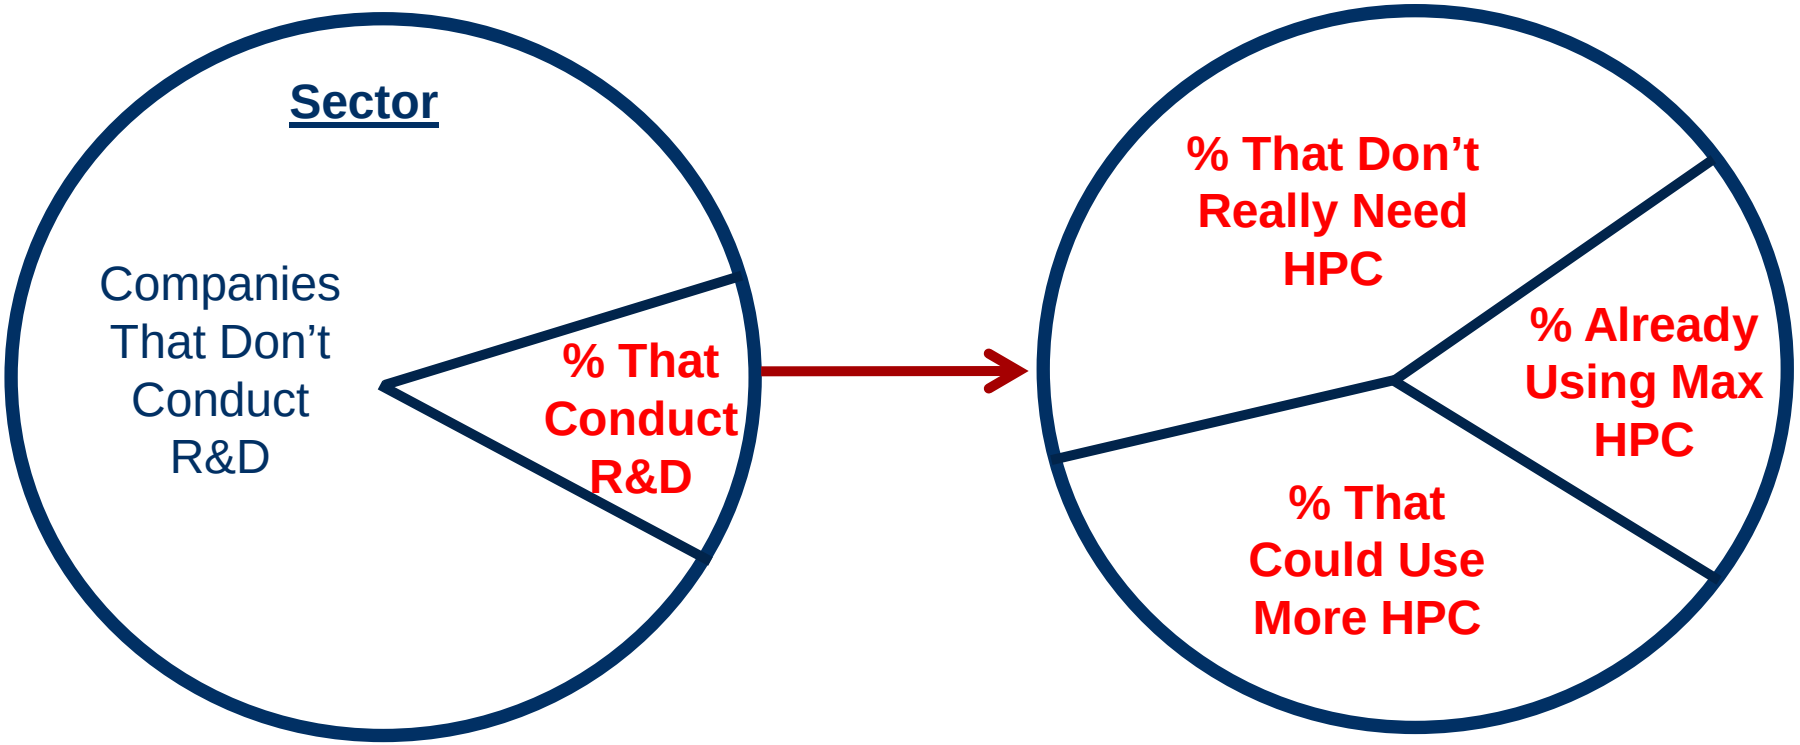

Note: IDC has conducted over 30,000 light phone calls for this data.
We will likely require 5x to 10x more surveys

HPC WW Market Trends ($K): By Competitive Segments

	2010	2011	2012	1H2013	2012/ 2011
Supercomputers	3,475,577	4,370,194	5,654,960	1,725,756	29.4%
Divisional	1,268,735	1,236,684	1,216,187	702,067	-1.7%
Departmental	3,342,747	3,467,271	2,979,230	1,853,790	-14.1%
Workgroup	1,411,264	1,225,910	1,247,366	809,349	1.8%
Total	9,498,323	10,300,058	11,097,743	5,090,962	7.7%

HPC WW Market Trends: By System Units Sold

	2010	2011	2012	1H2013	2012/ 2011
Supercomputers	2,560	2,908	2,400	748	-17.5%
Divisional	3,914	3,724	3,663	2,189	-1.6%
Departmental	20,382	20,625	16,981	11,048	-17.7%
Workgroup	92,988	84,294	81,104	50,963	-3.8%
Total	119,844	111,551	104,148	64,947	-6.6%

HPC WW Market Trends ($K): By OEM

	2011	2012	1H2013
IBM	3,362,098	3,551,723	1,563,771
HP	3,307,427	3,419,554	1,573,560
Dell	1,493,289	1,493,172	737,733
Cray	155,620	353,800	122,220
SGI	225,741	274,693	216,743
Fujitsu	120,351	686,657	65,139
NEC	84,141	64,112	37,106
Appro	135,360	111,648	-
Dawning	102,923	115,359	74,906
Bull	327,536	60,494	40,049
Other	847,140	966,531	659,736
Total	10,300,058	11,097,743	5,090,962

GDP, Supercomputers and HPC

IDC Analyze the Future

TABLE 45

GDP And Supercomputer Spending By Country, Sorted By GDP

	GDP (1)	Average Supercomputer Sales Over Last Five Years (2)	Supercomputers As A Percentage Of GDP	Average 5 year HPC Spending	HPC As A Percentage Of GDP
U.S.	14,270,000	1,276,067	0.0089%	4,464,817	0.0313%
Japan	5,049,000	278,385	0.0055%	651,126	0.0129%
China	4,758,000	67,836	0.0014%	278,480	0.0059%
Germany	3,235,000	203,245	0.0063%	761,309	0.0235%
France	2,635,000	142,209	0.0054%	517,170	0.0196%
U.K.	2,198,000	129,384	0.0059%	478,353	0.0218%
Italy	2,090,000	76,751	0.0037%	338,661	0.0162%
Spain	1,466,000	37,690	0.0026%	138,984	0.0095%
Russia	1,255,000	30,371	0.0024%	75,720	0.0060%
India	1,243,000	19,627	0.0016%	74,780	0.0060%
Australia	920,000	55,411	0.0060%	238,900	0.0260%
Korea	800,300	59,305	0.0074%	284,705	0.0356%
Switzerland	484,100	24,144	0.0050%	94,481	0.0195%
Sweden	397,700	21,314	0.0054%	75,043	0.0189%
Hong Kong	208,800	15,491	0.0074%	67,547	0.0324%

China Supercomputer Growth

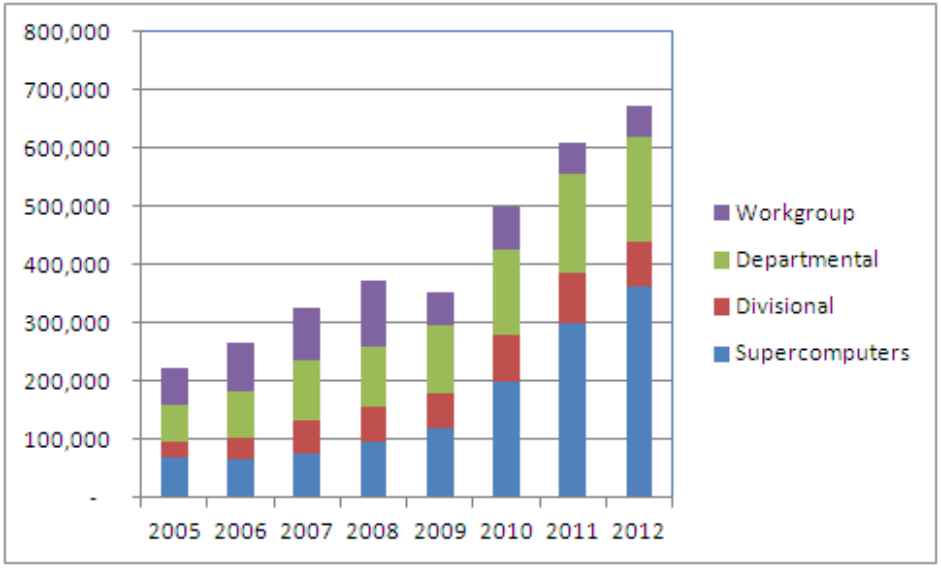

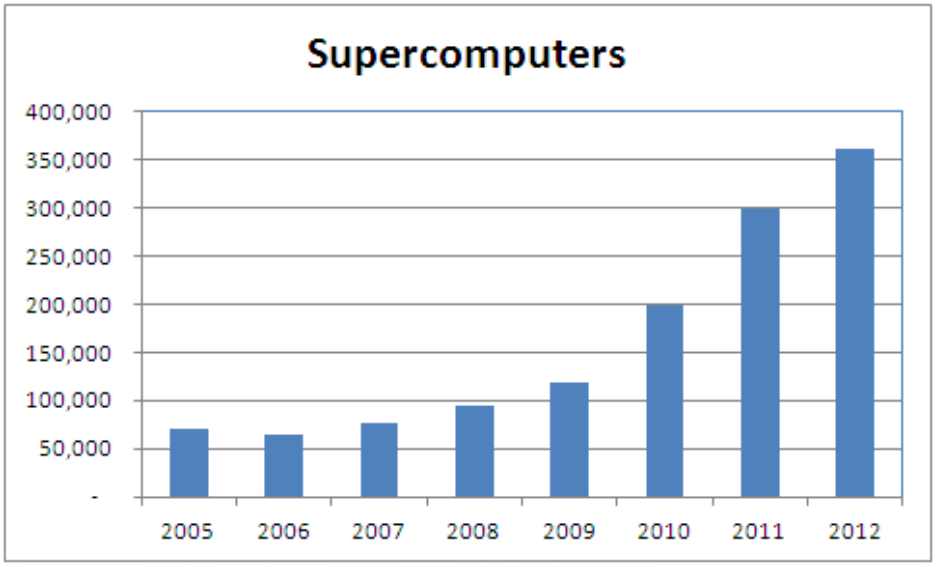

- **The China HPC market wasn't impacted by the recession, and is well under way to reach $1 billion**
- **The China supercomputer segment grew the most heavily since 2007**

HPC Forecasts: By Competitive Segment

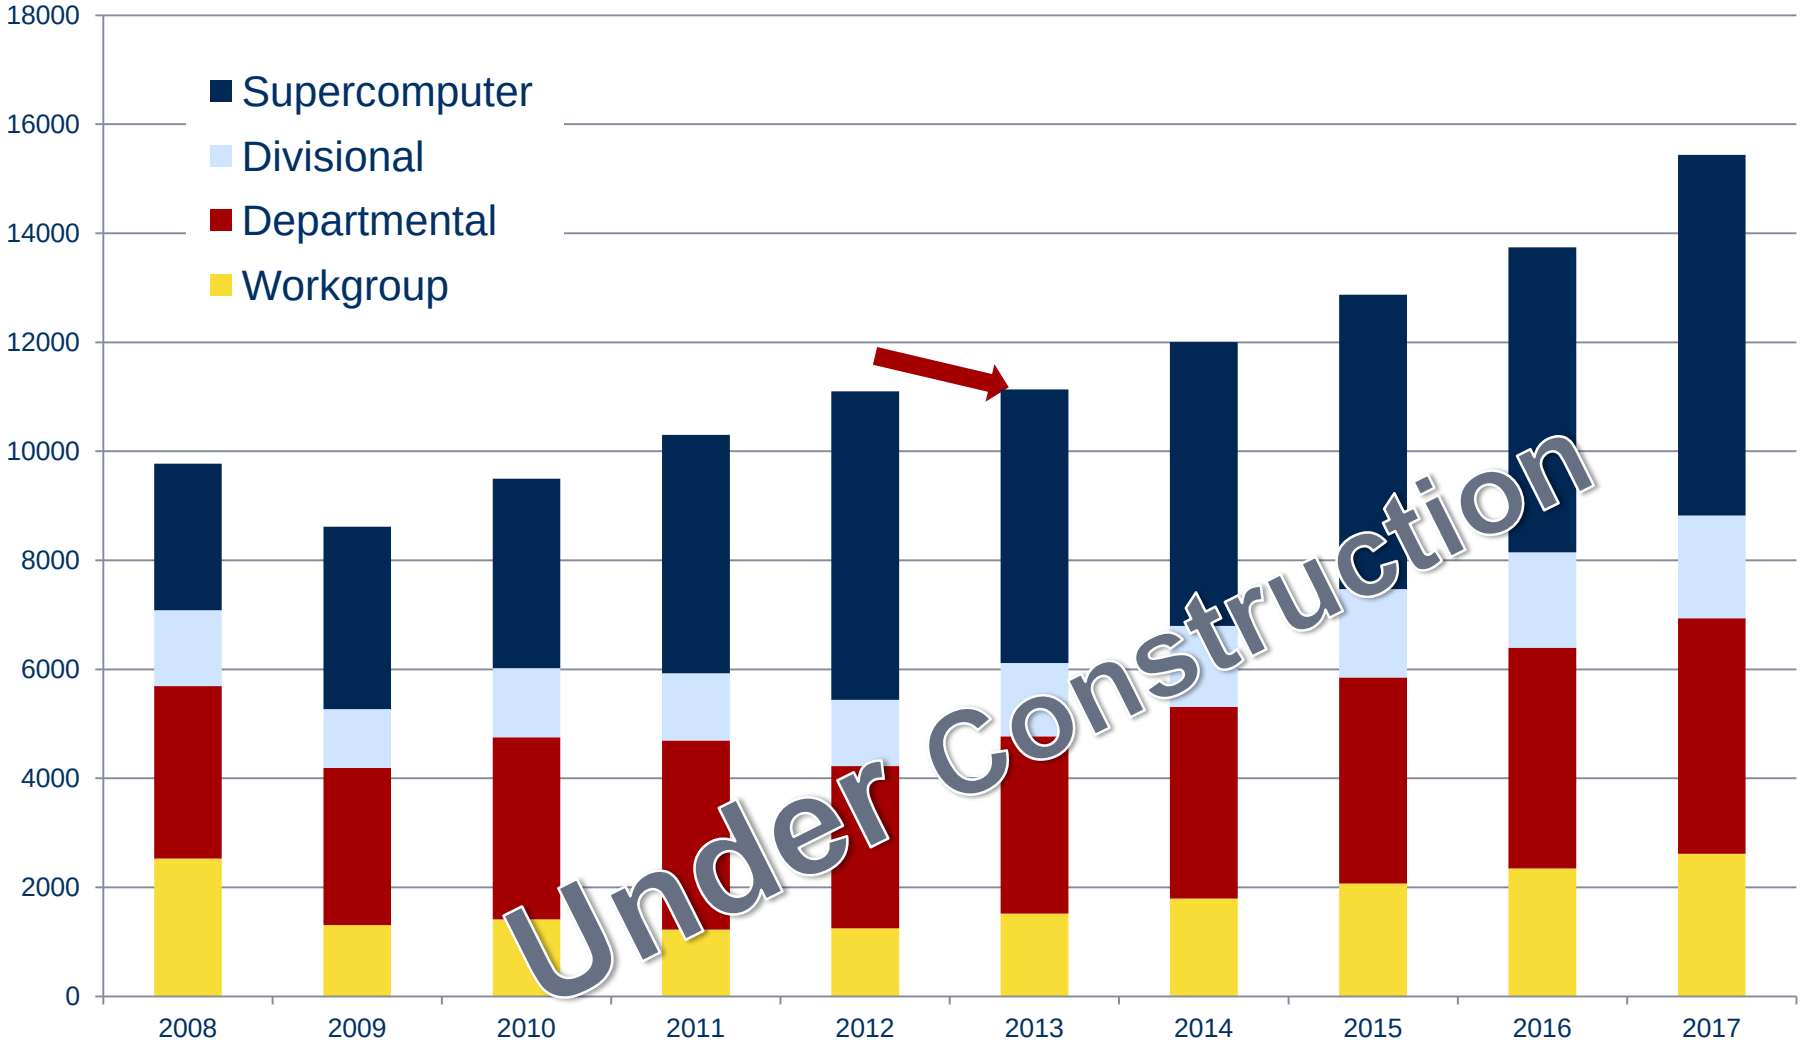

HPC Forecasts ($K): By Verticals/Application Areas

	2012	Est. 2013	2017	CAGR (12-17)
Bio-Sciences	$1,199,980	$1,267,745	$1,839,750	8.9%
CAE	$1,164,471	$1,220,781	$1,832,921	9.5%
Chemical Engineering	$182,006	$189,902	$262,116	7.6%
DCC & Distribution	$585,696	$621,945	$910,431	9.2%
Economics/Financial	$316,397	$339,485	$499,871	9.6%
EDA / IT / ISV	$624,696	$687,245	$1,028,523	10.5%
Geosciences and Geo-engineering	$707,869	$706,911	$958,673	6.3%
Mechanical Design and Drafting	$55,531	$61,482	$86,430	9.3%
Defense	$1,129,225	$1,104,015	$1,511,223	6.0%
Government Lab	$2,396,806	$2,279,607	$2,990,083	4.5%
University/Academic	$2,058,774	$2,029,805	$2,737,149	5.9%
Weather	$486,467	$476,181	$642,012	5.7%
Other	$189,823	$151,927	$141,563	-5.7%
Total Revenue	$11,097,742	$11,137,031	$15,440,744	6.8%

The HPC Market Beyond The Servers: The Broader HPC Market

IDC Analyze the Future

Worldwide HPC Compute, Storage, Middleware, Application and Service Revenues ($M)				
	2011	2012	2017	CAGR (12-17)
Server	10,300	11,098	15,441	6.8%
Storage	3,664	4,059	6,008	8.2%
Middleware	1,147	1,254	1,568	4.6%
Applications	3,370	3,621	4,837	6.0%
Service	1,801	1,877	2,368	4.8%
Total	20,282	21,909	30,223	6.6%

IDC WORLDWIDE HPC END-USER MULTI-CLIENT STUDY (2013)

Six Topical Reports:

- Industries/applications/workloads report
- System software and middleware report
- Storage and interconnect report
- Processors/co-processors/accelerators
- High performance data analysis
- Cloud computing

IDC HPC End-User MCS Study: Study Background

- **Surveys were conducted from January to May 2013**
- **The sample represents 905 installed HPC systems across 139 sites around the world**

Sector/Industry	Percentage
Government	7.9%
Education/Academia	41.7%
a. Energy, petroleum, oil and gas	2.9%
b. Chemical	0.7%
c. Pharmaceutical, life sciences, healthcare	2.2%
d. Financial or economic modeling and BI	20.1%
e. Manufacturing	10.8%
f. IT, electronics and telecommunications	4.3%
g. Transportation and logistics	2.9%
h. Entertainment	2.2%
i. Other	4.3%

2013 MCS Report Highlights:
Industries/Applications/Workloads

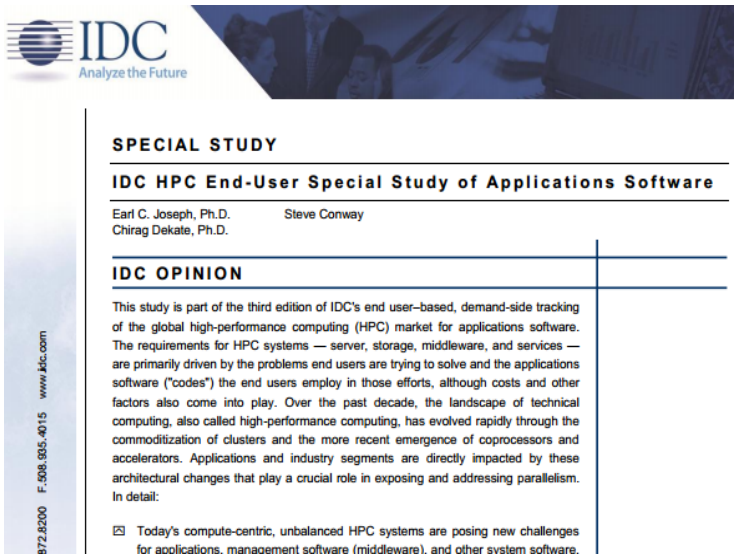

IDC Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of Applications Software

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

This study is part of the third edition of IDC's end user–based, demand-side tracking of the global high-performance computing (HPC) market for applications software. The requirements for HPC systems — server, storage, middleware, and services — are primarily driven by the problems end users are trying to solve and the applications software ("codes") the end users employ in those efforts, although costs and other factors also come into play. Over the past decade, the landscape of technical computing, also called high-performance computing, has evolved rapidly through the commoditization of clusters and the more recent emergence of coprocessors and accelerators. Applications and industry segments are directly impacted by these architectural changes that play a crucial role in exposing and addressing parallelism. In detail:

- Today's compute-centric, unbalanced HPC systems are posing new challenges for applications, management software (middleware), and other system software.

www.idc.com F.508.935.4015 872.8200

Scaling remains a major issue:

- 13.3% of the codes are running on 1 core
- 64.4% of the codes are running on 1 node or less
- Only 5.2% of codes run on >1,000 cores
- Only 0.9% of codes run on >10,000 cores

IDC forecasts that HPC software applications spending will reach $4.8 billion in 2017

2013 MCS Report Highlights:
System Software & Middleware Report

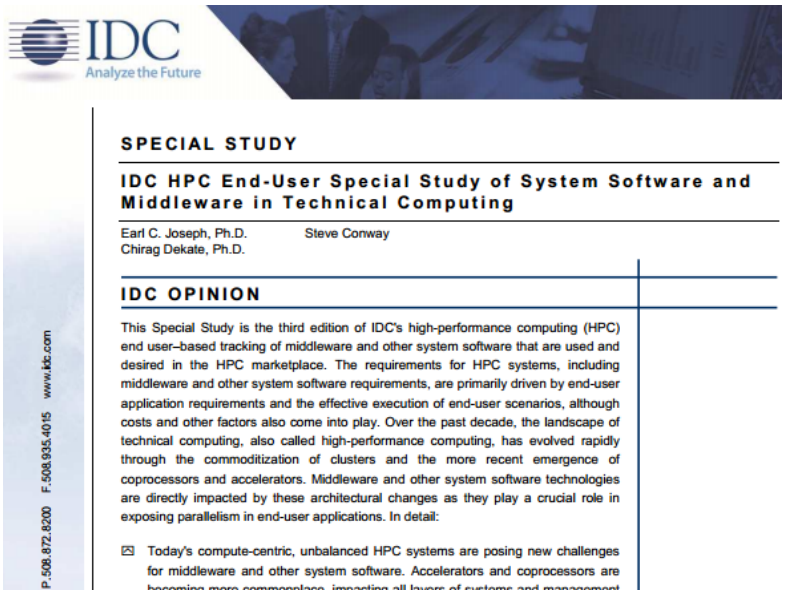

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of System Software and Middleware in Technical Computing

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

This Special Study is the third edition of IDC's high-performance computing (HPC) end user–based tracking of middleware and other system software that are used and desired in the HPC marketplace. The requirements for HPC systems, including middleware and other system software requirements, are primarily driven by end-user application requirements and the effective execution of end-user scenarios, although costs and other factors also come into play. Over the past decade, the landscape of technical computing, also called high-performance computing, has evolved rapidly through the commoditization of clusters and the more recent emergence of coprocessors and accelerators. Middleware and other system software technologies are directly impacted by these architectural changes as they play a crucial role in exposing parallelism in end-user applications. In detail:

- Today's compute-centric, unbalanced HPC systems are posing new challenges for middleware and other system software. Accelerators and coprocessors are

P.508.872.8200 F.508.935.4015 www.idc.com

Key Challenges:

- Increasing system sizes
- Increasing system complexity
- New environments (e.g., cloud)
- Shortage of skilled personnel

IDC forecasts that HPC system software spending will reach $1.5 billion in 2017

2013 MCS Report Highlights:
Storage and Interconnect Report

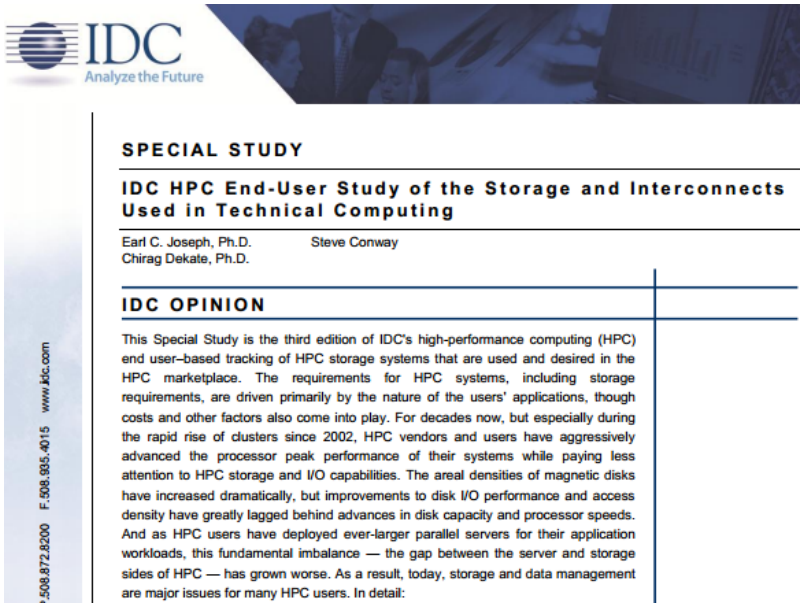

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Study of the Storage and Interconnects Used in Technical Computing

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

This Special Study is the third edition of IDC's high-performance computing (HPC) end user–based tracking of HPC storage systems that are used and desired in the HPC marketplace. The requirements for HPC systems, including storage requirements, are driven primarily by the nature of the users' applications, though costs and other factors also come into play. For decades now, but especially during the rapid rise of clusters since 2002, HPC vendors and users have aggressively advanced the processor peak performance of their systems while paying less attention to HPC storage and I/O capabilities. The areal densities of magnetic disks have increased dramatically, but improvements to disk I/O performance and access density have greatly lagged behind advances in disk capacity and processor speeds. And as HPC users have deployed ever-larger parallel servers for their application workloads, this fundamental imbalance — the gap between the server and storage sides of HPC — has grown worse. As a result, today, storage and data management are major issues for many HPC users. In detail:

T.508.872.8200 F.508.935.4015 www.idc.com

Storage is the fastest-growing HPC segment

- Ethernet variants lead
- InfiniBand is a strong second

IDC forecasts that storage & interconnect spending will reach $6.0 billion in 2017

- This equals the value of the HPC server market in 2000

2013 MCS Report Highlights:
Processors, Coprocessors, Accelerators

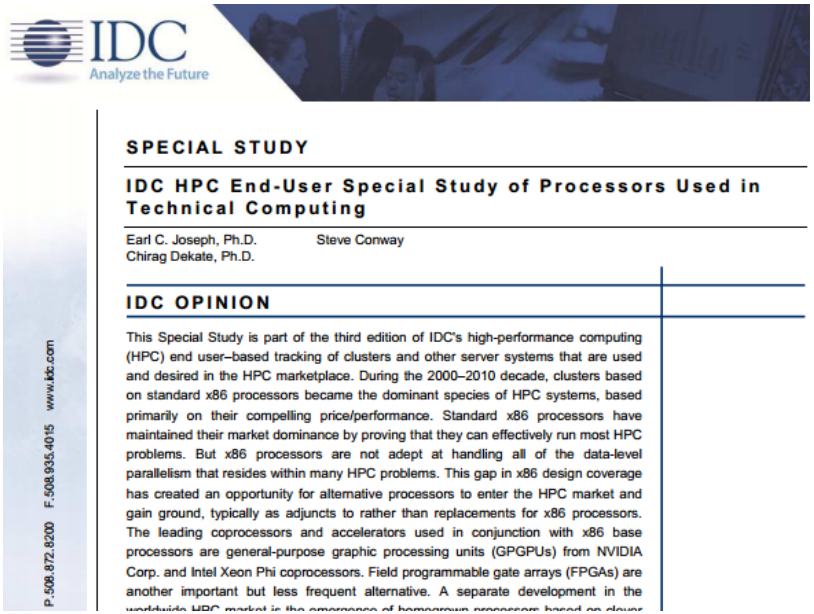

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of Processors Used in Technical Computing

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

This Special Study is part of the third edition of IDC's high-performance computing (HPC) end user–based tracking of clusters and other server systems that are used and desired in the HPC marketplace. During the 2000–2010 decade, clusters based on standard x86 processors became the dominant species of HPC systems, based primarily on their compelling price/performance. Standard x86 processors have maintained their market dominance by proving that they can effectively run most HPC problems. But x86 processors are not adept at handling all of the data-level parallelism that resides within many HPC problems. This gap in x86 design coverage has created an opportunity for alternative processors to enter the HPC market and gain ground, typically as adjuncts to rather than replacements for x86 processors. The leading coprocessors and accelerators used in conjunction with x86 base processors are general-purpose graphic processing units (GPGPUs) from NVIDIA Corp. and Intel Xeon Phi coprocessors. Field programmable gate arrays (FPGAs) are another important but less frequent alternative. A separate development in the

P.508.872.8200 F.508.935.4015 www.idc.com

Coprocessor/accelerator use is growing fast

- 2011: 28.2% of sites and 1.0% of processor parts
- 2013: 76.9% of sites and 3.4% of processor parts
- Still used more often for experimentation than production

CPUs from the embedded world are arriving (Atom, ARM)

Indigenous processor initiatives in Asia

2013 MCS Report Highlights:
High Performance Data Analysis (HPDA)

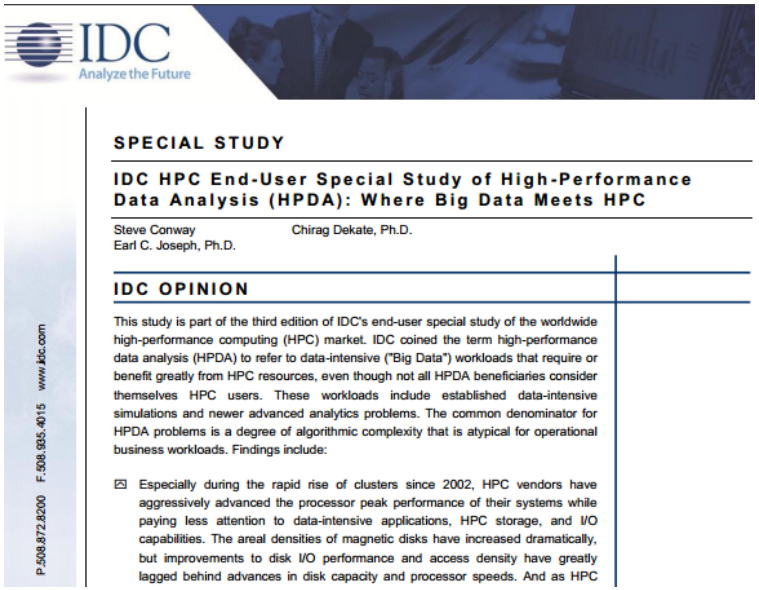

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of High-Performance Data Analysis (HPDA): Where Big Data Meets HPC

Steve Conway Chirag Dekate, Ph.D.
Earl C. Joseph, Ph.D.

IDC OPINION

This study is part of the third edition of IDC's end-user special study of the worldwide high-performance computing (HPC) market. IDC coined the term high-performance data analysis (HPDA) to refer to data-intensive ("Big Data") workloads that require or benefit greatly from HPC resources, even though not all HPDA beneficiaries consider themselves HPC users. These workloads include established data-intensive simulations and newer advanced analytics problems. The common denominator for HPDA problems is a degree of algorithmic complexity that is atypical for operational business workloads. Findings include:

- Especially during the rapid rise of clusters since 2002, HPC vendors have aggressively advanced the processor peak performance of their systems while paying less attention to data-intensive applications, HPC storage, and I/O capabilities. The areal densities of magnetic disks have increased dramatically, but improvements to disk I/O performance and access density have greatly lagged behind advances in disk capacity and processor speeds. And as HPC

P.508.872.8200 F.508.935.4015 www.idc.com

HPDA is becoming pervasive

- 67% of surveyed sites run HPDA workloads
- On average, HPDA workloads consume 30% of compute cycles

IDC forecasts robust HPDA growth

- HPDA server revenue will grow from $748.8M in 2012 to $1.4B in 2017
- HPDA storage revenue will approach $1B in 2017

2013 MCS Report Highlights:
Cloud Computing

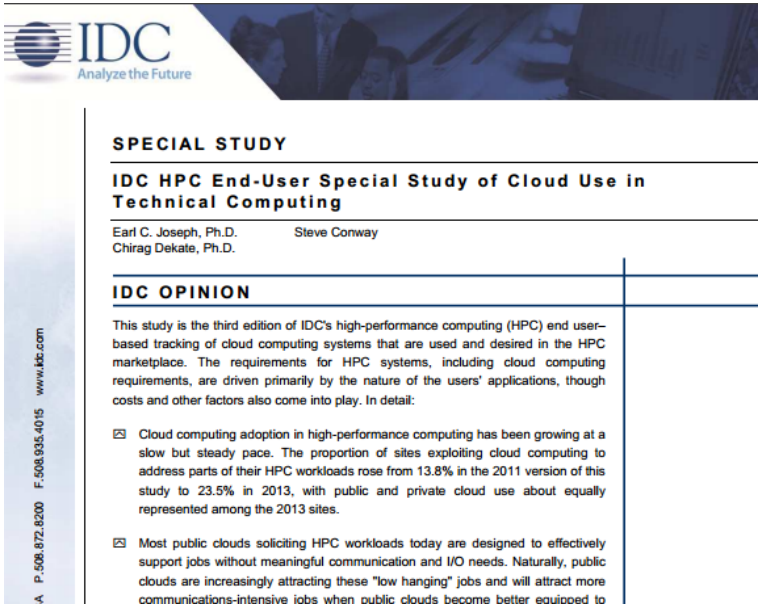

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of Cloud Use in Technical Computing

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

This study is the third edition of IDC's high-performance computing (HPC) end user–based tracking of cloud computing systems that are used and desired in the HPC marketplace. The requirements for HPC systems, including cloud computing requirements, are driven primarily by the nature of the users' applications, though costs and other factors also come into play. In detail:

- Cloud computing adoption in high-performance computing has been growing at a slow but steady pace. The proportion of sites exploiting cloud computing to address parts of their HPC workloads rose from 13.8% in the 2011 version of this study to 23.5% in 2013, with public and private cloud use about equally represented among the 2013 sites.
- Most public clouds soliciting HPC workloads today are designed to effectively support jobs without meaningful communication and I/O needs. Naturally, public clouds are increasingly attracting these "low hanging" jobs and will attract more communications-intensive jobs when public clouds become better equipped to

P.508.872.8200 F.508.935.4015 www.idc.com

HPC Cloud Computing Is Growing Steadily

- From 13.8% of sites in 2011 to 23.5% in 2013
- Private and public cloud computing about equally represented

New HPC & HPDA Technical Computing Pulse Panel

New IDC HPC Pulse Panel Capability

A new IDC capability (available now)

- Rapid answers to urgent questions about the HPC, HPDA markets: new trends, new technologies, emerging markets
- Rifle-shot probes: 3-6 weeks, vs. several months for full-out studies
- IDC developed this capability for key clients
- Many successful Pulses carried out over past 6 months

Methodology: the Pulse panel

- Highly qualified, representative panel of 500+ HPC/HPDA technical experts, funders, buyers, and end-users already in the program
- To develop and qualify the panel, IDC has contacted 30,000 scientists and engineers around the world. Will grow to 100,000.
- Primarily web-based

Example: HPDA Pulse (Highlights)

Pulse conducted Sept/Oct 2013

- Only panelist organizations that use HPDA (Big Data on HPC)
- 62.2% have access to supercomputer-class systems (>$500,000)
  - Respondents also included many SMEs
- 45.9% consider HPDA jobs “very critical” to their missions
- 46.6% have near-real time HPDA requirements (few seconds to few minutes)
  - 19.2% real time to a few seconds
  - 27.4% within a few minutes

HPDA Update

High Performance Data Analysis

- Needs HPC resources
- Intelligent questions / smart algorithms)
- Often near-real time

- Search & pattern discovery
- Simulation & analytics
- Also iterative methods
- Established HPC users + new commercial users

- The 4 V's: volume, variety, velocity, value
- Partitionable & non-partitionable problems
- Regular and irregular data patterns

HPC Adoption Timeline (Examples)

Ford

Morgan Stanley

P&G

SCHRÖDINGER, INC.

GEICO

TOYOTA

DREAMWORKS ANIMATION SKG

Walmart

CMS CENTERS for MEDICARE & MEDICAID SERVICES

PayPal

VISA MasterCard DISCOVER AMERICAN EXPRESS

BOEING

bp

SAMSUNG

MERCK

PING

UNITED STATES POSTAL SERVICE

AIRBUS

Panasonic

MAYO CLINIC

emcien

1960 1970 1980 1990 2000 2012

HPDA User Talks: HPC User Forums, UK, Germany, France, China, U.S.

IDC Analyze the Future

- HPC in Evolutionary Biology, Andrew Meade, University of Reading
- HPC in Pharmaceutical Research: From Virtual Screening to All-Atom Simulations of Biomolecules, Jan Kriegl, Boehringer-Ingelheim
- European Exascale Software Initiative, Jean-Yves Berthou, Electricite de France
- Real-time Rendering in the Automotive Industry, Cornelia Denk, RTT-Munich
- Data Analysis and Visualization for the DoD HPCMP, Paul Adams, ERDC
- Why HPCs Hate Biologists, and What We're Doing About It, Titus Brown, Michigan State University
- Scalable Data Mining and Archiving in the Era of the Square Kilometre Array, the Square Kilometre Array Telescope Project, Chris Mattmann, NASA/JPL
- Big Data and Analytics in HPC: Leveraging HPC and Enterprise Architectures for Large Scale Inline Transactional Analytics in Fraud Detection at PayPal, Arno Kolster, PayPal, an eBay Company
- Big Data and Analytics Vendor Panel: How Vendors See Big Data Impacting the Markets and Their Products/Services, Panel Moderator: Chirag Dekate, IDC
- Data Analysis and Visualization of Very Large Data, David Pugmire, ORNL
- The Impact of HPC and Data-Centric Computing in Cancer Research, Jack Collins, National Cancer Institute
- Urban Analytics: Big Cities and Big Data, Paul Muzio, City University of New York
- Stampede: Intel MIC And Data-Intensive Computing, Jay Boisseau, Texas Advanced Computing Center
- Big Data Approaches at Convey, John Leidel
- Cray Technical Perspective On Data-Intensive Computing, Amar Shan
- Data-intensive Computing Research At PNNL, John Feo, Pacific Northwest National Laboratory
- Trends in High Performance Analytics, David Pope, SAS
- Processing Large Volumes of Experimental Data, Shane Canon, LBNL
- SGI Technical Perspective On Data-Intensive Computing, Eng Lim Goh, SGI
- Big Data and PLFS: A Checkpoint File System For Parallel Applications, John Bent, EMC
- HPC Data-intensive Computing Technologies, Scott Campbell, Platform/IBM
- The CEA-GENCI-Intel-UVSQ Exascale Computing Research Centre, Marie-Christine Sawley, Intel

IDC HPDA Server Forecast

- Fast growth from a small starting point: $1.2B by 2016)
- HPDA ecosystem >$2B in 2016

TABLE 2

IDC Worldwide High Performance Data Analysis (HPDA) Server Revenues

($ Millions)

	2009	2010	2011	2012	2013	2014	2015	2016	CAGR '11-'16
WW HPC Server Sales	8,637	9,498	10,300	11,098	11,397	12,371	13,485	14,621	7.3%
WW HPDA Server Sales	535	603	673	744	786	881	1,109	1,253	13.3%
HPDA Portion	6.2%	6.3%	6.5%	6.7%	6.9%	7.1%	8.2%	8.6%	5.6%

Source: IDC 2013

IDC HPDA Storage Forecast

- Storage is the fastest-growing HPC market (8.4% CAGR, 2011-16) and HPDA storage will grow even faster (18.1% CAGR).

TABLE 2

Worldwide High-Performance Data Analysis Storage Revenue, 2009–2016 ($M)

	2009	2010	2011	2012	2013	2014	2015	2016	2011–2016 CAGR (%)
HPC storage	3,023.0	3,325.9	3,761.5	4,194.0	4,349.8	4,739.1	5,163.2	5,625.3	8.4
Share as total HPC server revenue (%)	35.0	35.0	36.5	37.8	38.2	38.3	38.3	38.5	1.0
HPDA storage	262.2	301.5	343.0	387.0	432.2	519.9	676.5	789.5	18.1
Big Data attach rate (%)	49.0	50.0	51.0	52.0	55.0	59.0	61.0	63.0	4.3

Source: IDC, 2013

IDC 2013 Worldwide HPC End-User Study: HPDA Top Findings

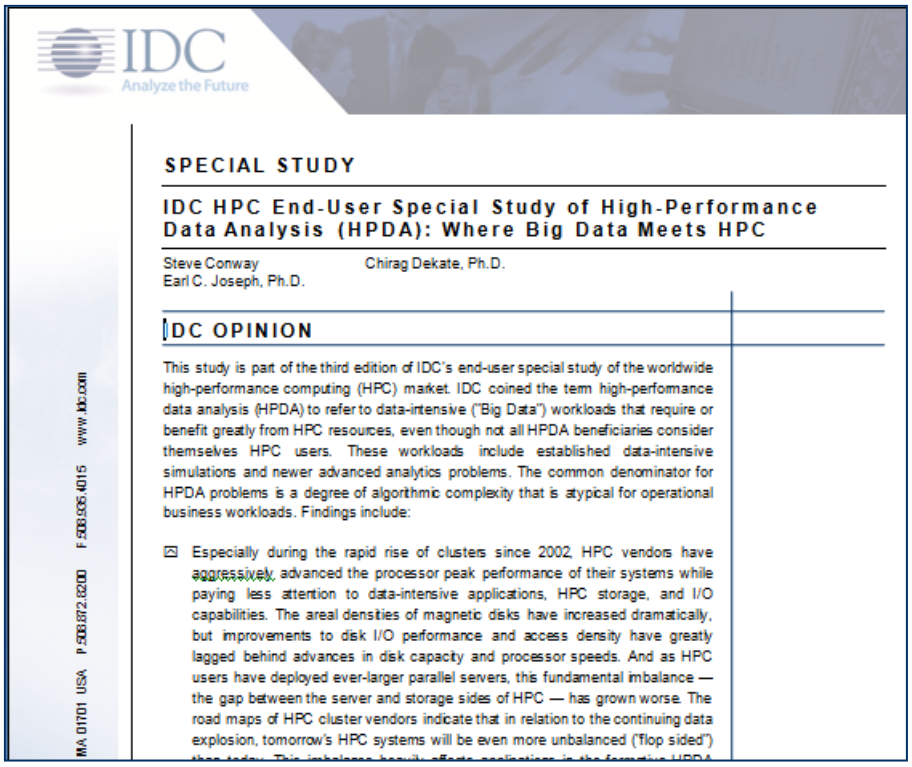

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of High-Performance Data Analysis (HPDA): Where Big Data Meets HPC

Steve Conway
Earl C. Joseph, Ph.D.
Chirag Dekate, Ph.D.

IDC OPINION

This study is part of the third edition of IDC's end-user special study of the worldwide high-performance computing (HPC) market. IDC coined the term high-performance data analysis (HPDA) to refer to data-intensive ("Big Data") workloads that require or benefit greatly from HPC resources, even though not all HPDA beneficiaries consider themselves HPC users. These workloads include established data-intensive simulations and newer advanced analytics problems. The common denominator for HPDA problems is a degree of algorithmic complexity that is atypical for operational business workloads. Findings include:

- Especially during the rapid rise of clusters since 2002, HPC vendors have aggressively advanced the processor peak performance of their systems while paying less attention to data-intensive applications, HPC storage, and I/O capabilities. The areal densities of magnetic disks have increased dramatically, but improvements to disk I/O performance and access density have greatly lagged behind advances in disk capacity and processor speeds. And as HPC users have deployed ever-larger parallel servers, this fundamental imbalance — the gap between the server and storage sides of HPC — has grown worse. The road maps of HPC cluster vendors indicate that in relation to the continuing data explosion, tomorrow's HPC systems will be even more unbalanced ("flop sided")

MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

- 67% of the sites perform HPDA work (data-intensive simulation and/or advanced analytics).
- On average, HPDA consumes 30% of compute cycles.
- Major pain points worth 10-15% premium pricing:
  - Higher-performance interconnects between nodes
  - Higher-performance external I/O and storage

Big Science: Big Data Challenges Are Growing…

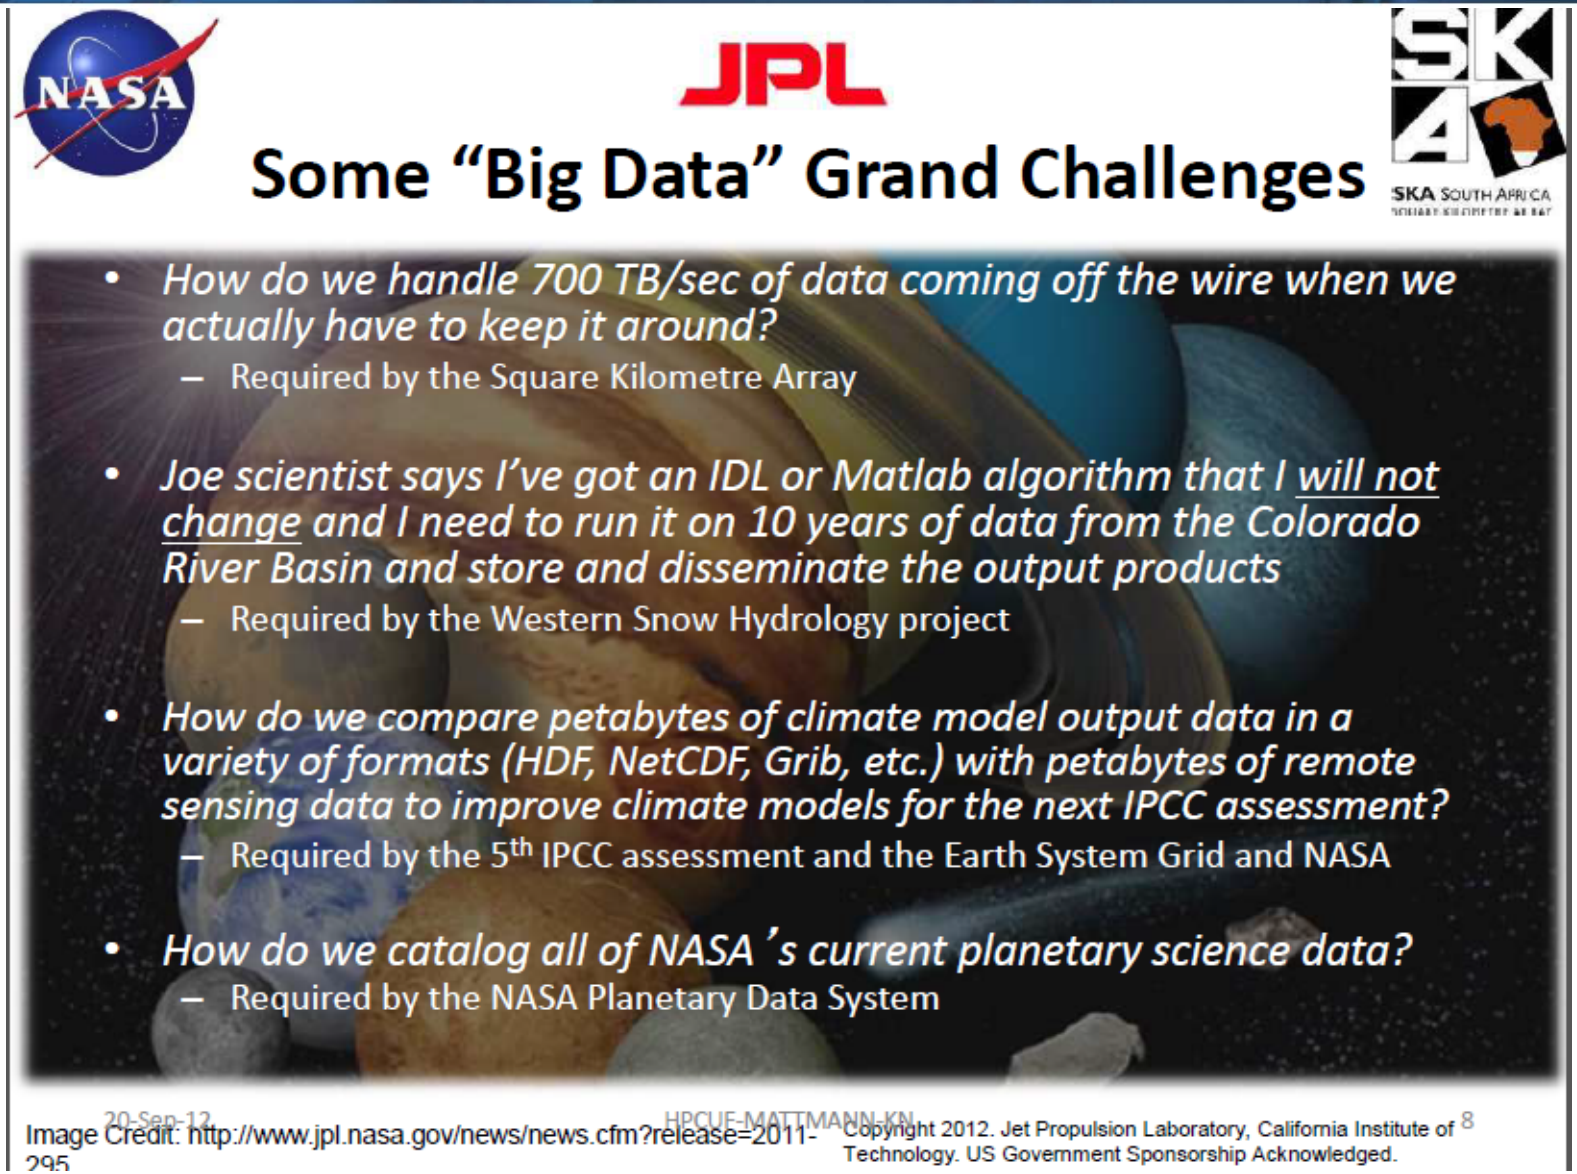

Anomaly Detection / Remediation via Pattern Discovery

- Fraud/errors
- Identity resolution
- Anti-terrorism, anti-crime
- Cyber security

U.S. Postal Service

Total Revenue Protection Program

- Processing Requirements
  - Rate
    - 4 billion mail scans per day peak (74,000 per second)
  - Geographic Scope
    - Incoming mail from 275 Processing and Distribution Centers
    - Outgoing mail to 33,000 postal operated facilities
  - Objective
    - To find, track and reject mail pieces due to:
      - Duplicate postage
      - Short Pay
      - Ineligible Discounts

U.S. Postal Service

Why Real Time Fraud Detection?

Save time… print your postage online.

Print exact postage for letters and packages using just your PC and printer.

Print Postage Stamps

- Print any denomination
- Use for letters or packages
- Never run out of stamps again

DETAILS

Stamps.com… Your own personal Post Office open 24 hours a day.

Developed in conjunction with the United States Postal Service,℠ Stamps.com is a revolutionary software-based service that allows you to calculate and print official USPS postage right from your PC.

NO ADDITIONAL HARDWARE REQUIRED. Stamps.com even keeps track of all your postal spending using your client codes, and can even recommend optimal delivery methods, formats and more. Plus, Stamps.com gives you postage discounts you can't even get at the Post Office or with a postage meter.

U.S. Postal Service

TRP Results using MCDB &TimesTen

Pre-MCDB

1. 509 row inserts per second (RIPS)
2. Direct path load option a partial solution (2000 RIPS)
3. 275 Million Transactions per 15 hour processing window created backlog during peak processing windows
4. Revenue Protection performed as a batch data warehouse process, run 3 – 12 hours after Mailpiece scan

With MCDB Deployed

1. 190,222 RIPS (3 Threads)
2. 1,091,018 RIPS (18 Threads)
3. Processed 4 B Transactions in less than 6 hours
4. Revenue Protection is performed in real-time upon first scan

MCDB = memory-centric database

10

Life Sciences / Materials Science

- Drug discovery
- Genome mapping/comparison
- Health care management
- Personal / outcomes-based medicine
- Materials science

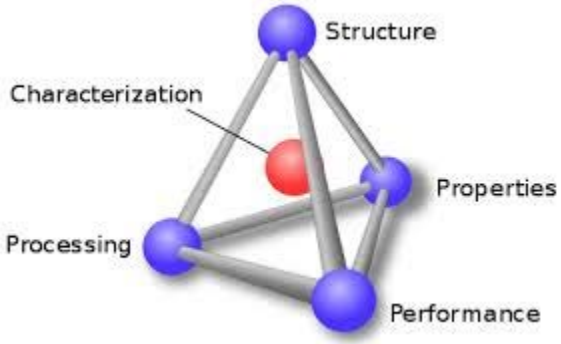

Schrödinger: Cloud-based Lead Discovery for Drug Design

IDC
Analyze the Future

Metric from March, 2012	Count
Compute Hours of Work	109,927 hours
Compute Years of Work	12.55 years
Total # Cores/Servers	51132 cores, 6742 servers
Infrastructure Value	~ $20,000,000 (£12,500,000)
AWS Regions	All (7: us-east, us-west1, us-west2, eu-west, sa-east, ap-northeast, ap-southeast)

Using CycleCloud & AWS:
Impossible run in 3 hours for $4,828/hr
Today's pricing < $1,000/hr (£625/hr)

CYCLECOMPUTING

Genome Sequencing/Analysis

Shotgun genomics

- Collect samples
- Extract DNA
- Feed into sequencer
- Computationally analyze
- A lab can generate ~100 Gbp in ~1 week for $10k (£6.3K)

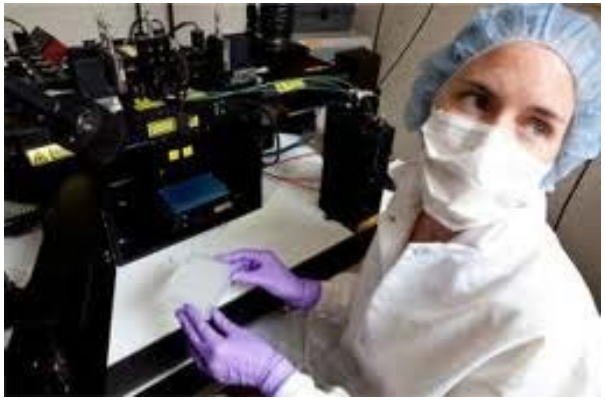

Real-World Use Cases

- Real-time pathogen analysis
- Cancer genome analysis => diagnosis & treatment
- Drug resistance in HIV
- Gene expression analysis in agricultural animals
- Microbial community change in response to agriculture or global climate change
- Gene discovery & genome sequencing in non-model organisms

Courtesy Titus Brown, Michigan State University

Outcomes-Based Medical Diagnosis and Treatment Planning

- Enter the patient's history and symptomology.
- While patient is still in the office, sift through millions of archived patient records for relevant outcomes.
- Provider considers the efficacies of various treatments for "similar" patients (but is not bound by the findings).
- Ergo, this functions as a powerful decision-support tool.
- Benefits: better outcomes + rein in costly outlier practices

Iterative Methods (Cumulative Data)

- Parametric modeling (product design)
- Stochastic modeling (financial)
- Ensemble modeling (weather/climate)

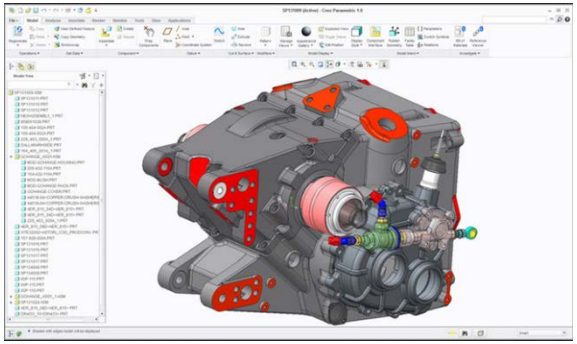

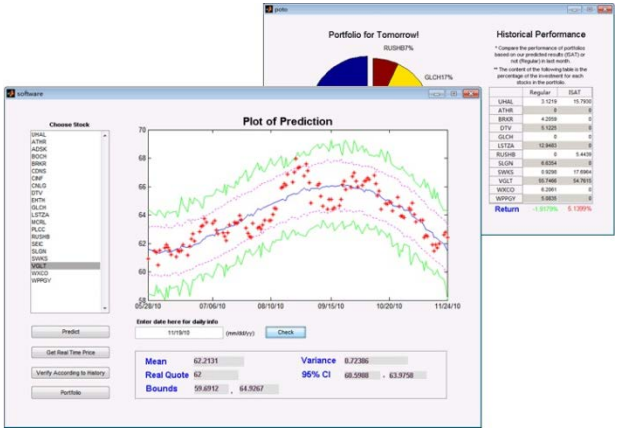

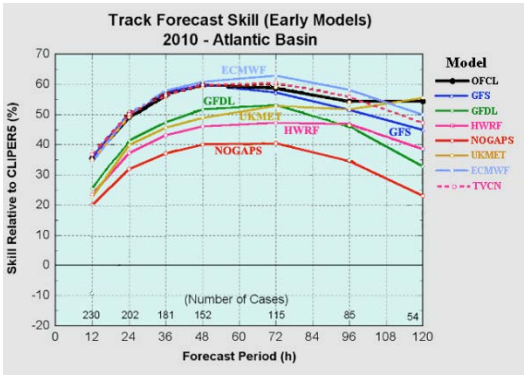

Summary: HPDA Market Opportunity

- **HPDA: simulation + newer high-performance analytics**
  - IDC predicts fast growth from a small starting point
- **HPC and high-end commercial analytics are converging.**
  - Algorithmic complexity is the common denominator
- **Economically important use cases are emerging**
  - Which ones will become attractive markets?
- **No single HPC solution is best for all problems**.
  - Clusters with MR/Hadoop will handle most but not all work (e.g., graph analysis)
  - **IDC believes our growth estimates could be conservative**

A New IDC Study:

Creating Economic Models For HPC and ROI And for HPC And Innovation

Grant References

The authors thank DOE for its insights and guidance on and funding of this grant-based research project.

This study is based upon work funded by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, and the National Nuclear Security Administration, under award number DE-SC0008540.

DOE Program Managers:

Christine Chalk, 301-903-5152, christine.chalk@science.doe.gov, and Barbara J. Helland, 301-903-3127, barbara.helland@science.doe.gov, U.S. Department of Energy Office of Science, Germantown Building, 1000 Independence Avenue, S.W. Washington, D.C., 20585-1290.
and Bob Meisner in National Nuclear security Administration (NNSA).

Administrator/Contracting Officer:

Warren Riley, 630-252-2485, warren.riley@ch.doe.gov, U.S. Department of Energy Office of Acquisition and Assistance, 9800 South Cass Avenue, Argonne, Illinois, 60439

IDC Reporting:

- Principal investigator: Earl C. Joseph, Ph.D., 612-812-5798, ejoseph@idc.com
- Senior technical project manager: John Daly, 508-935-4643, jdaly@idc.com

Background: Project Overview

A pilot study that describes how HPC investments are related to improved economic success and increased scientific innovation

The study included creating two unique models:

1. A macroeconomic model which depicts how HPC investments result in economic advancements in the form of ROI, growth and jobs

2. An "Innovation Index" that measures and compares innovation levels, based on the level of applying HPC computing resources towards scientific and technical advancement

The Financial ROI Models That Were Developed

The Financial ROI models include:

1. **ROI based on revenues/GDP generated, divided by HPC investment**
2. **ROI based on profits generated, divided by HPC investment**
3. **ROI based on jobs created (and the HPC investment required per job created)**

The ROI models were tested for variances by:

- Industry sector
- Country
- Organization size

The Innovation Models That Were Developed

The Innovation models are of two main types:

1. **Basic Research / Major Innovations**
2. **Applied Research / Incremental Innovations**

These are captured as:

- Innovations in government & academia
- Innovations in industry

The Innovation models can be sorted for variances by:

- Industry sector
- Country
- Organization size
- Government, Industry and Academia

The Innovation Index Scale

IDC
Analyze the Future

10 = One of the top 2 to 3 innovations in the last decade

9 = One of the top 5 innovations in the last decade

8 = One of the top 10 innovations in the last decade

7 = One of the top 25 innovations in the last decade

6 = One of the top 50 innovations in the last decade

5 = It had a major impact and is useful to many organizations

4 = A minor innovation that is useful to many organizations

3 = A minor innovation or only useful to 2 -3 organizations

2 = A minor innovation or only useful to 1 organization

1 = An innovation that is recognized ONLY by experts in the field

Sample Characteristics

Sample demographics:

- **A total of 208 case study examples of ROI and innovations were collected as part of the study:**
  - 67 financial ROI examples
  - 141 innovation examples

- In addition, a large number of micro-surveys were conduct to learn key ratios in order to eventually apply the results to large economic data sets.
  - Over 30,000 scientists and engineers were contacted, with over 1,500 completing the micro-survey.

Key Findings:
Primary Financial ROI Results

1. **IDC is able to collect the required data across a broad set of organizations with enough detail to create the two economic models and the innovation index**

2. **Early results indicate very substantial returns for investments in HPC:**
   - **$356 dollars on average in revenue** per dollar of HPC invested.
   - **$38 dollars on average of profits** (or cost savings) per dollar of HPC invested.

Key Findings:
The Financial ROI Model – By Sector

3	Organization Size: People	(All)
4	Organization Size in $M	(All)
5	Organization Size (S,M,L)	(All)
6	Industry	(All)
7	Innovation Level	(All)
8	Country	(All)
9	Years Before 1st Return	(All)
10	Applied	(All)
11	Basic	(All)
12	Accomplishment Type	Financial R
13	Total R&D	(All)
14	Employee Growth	(All)
15	Organization Type	(All)
16	Total HPC Investment	(All)

19	Sector	Count	Sum of Employee Growth	Years Before 1st Return	Average of Revenue $ per HPC $	Average of Profit $ per HPC $
20	Academic	12	2	1.8	37.4	70.8
21	Government	4	10	1.4	9.2	3.9
22	Industry	51	1,157	1.9	462.4	36.4
23	**Grand Total**	**67**	**1,169**	**1.9**	**356.5**	**38.7**

Key Findings:
The Financial ROI Model – By Country

3	Organization Size: People	(All)	
4	Organization Size in $M	(All)	
5	Organization Size (S,M,L)	(All)	
6	Industry	(All)	
7	Innovation Level	(All)	
8	Years Before 1st Return	(All)	
9	Applied	(All)	
10	Basic	(All)	
11	Sector	(All)	
12	Accomplishment Type	Finan	ROI
13	Total R&D	(All)	
14	Employee Growth	(All)	
15	Organization Type	(All)	
16	Total HPC Investment	(All)	

19	Country	Count	Sum of Employee Growth	Average Years Before 1st Return	Average of Revenue $ per HPC $	Average of Profit $ per HPC $
20	China	3	30	1.3	8.7	5.4
21	France	4		5.1	621.7	125.0
22	UK	31	896	1.6	366.5	26.7
23	US	27	243	1.8	373.3	49.8
24	Italy	2		1.0	10.0	7.5
25	**Grand Total**	**67**	**1,169**	**1.9**	**356.5**	**38.7**

Pilot Study Results: Innovation

Key Findings: The Innovation Areas For The 141 Innovation Data Examples

Industry	(All)			
Years Before 1st Return	(All)			
Applied	(All)			
Basic	(All)			
Accomplishment Type	Innovation			
Country	(All)			
Total R&D	(All)			
Organization Type	(All)			

Primary Innovation / ROI Area	Count	Sum of Total HPC Investment	Average Years Before 1st Return	Average of HPC $M per Innovation
Better Products	54	$114 M	1.9	$4.2 M
Created New Approach	40	$15 M	1.2	$0.4 M
Discovered Something New	20	$46 M	1.8	$2.7 M
Helped Society	11	$66 M	1.0	$6.0 M
Cost Saving	6	$180 M	1.3	$2.1 M
Major Breakthrough	5	$3 M	3.2	$1.1 M
Helped Research Program	5	$71 M	1.5	$14.3 M
Grand Total	**141**	**$497 M**	**1.6**	**$3.1 M**

Key Findings: Investments Per Innovation

5. **The average HPC investment per innovation was $3.1 million.**
   - Overall $497 million in HPC investments were made to generate the 141 innovations in the pilot study.
   - With many at under $1 million per innovation.

Key Findings:
The New Innovation Index Scores

The average innovation rating = 5.0

- 4.4 for the 67 basic research/major innovations
- 5.5 for the 74 applied research/incremental innovations

10 = One of the top 2 to 3 innovations in the last decade
9 = One of the top 5 innovations in the last decade
8 = One of the top 10 innovations in the last decade
7 = One of the top 25 innovations in the last decade
6 = One of the top 50 innovations in the last decade
5 = It had a major impact and is useful to many organizations
4 = A minor innovation that is useful to many organizations
3 = A minor innovation or only useful to 2 -3 organizations
2 = A minor innovation or only useful to 1 organization
1 = An innovation that is recognized ONLY by experts in the field

Key Findings: The New Innovation Index Scores – For All 141 Innovations

FIGURE 12

HPC Innovation Index Scale Results: All Respondents

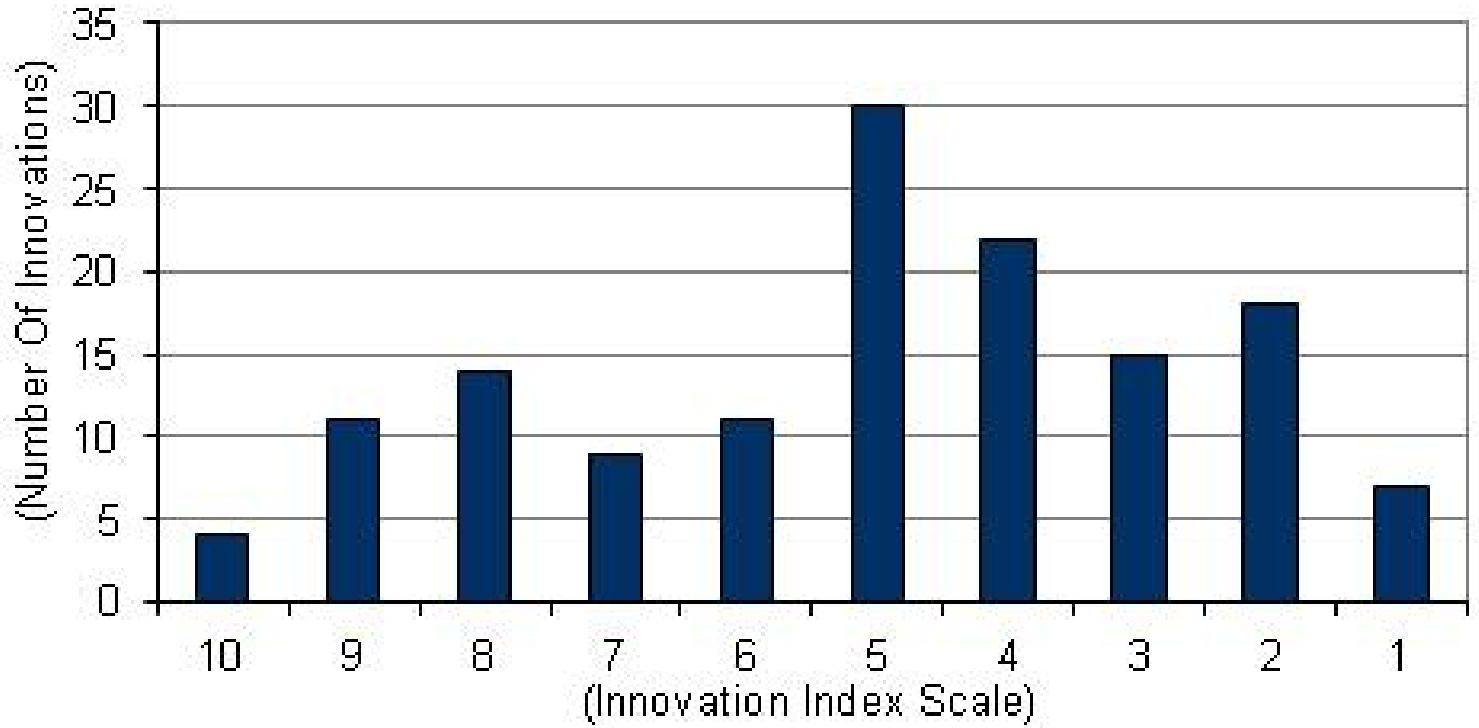

N = 141

Source: IDC 2013

Key Findings:
The Innovation Index By Country

Sector	(All)			
Accomplishment T	Innova n			
Years Before 1st R	(All)			
Jobs Added	(All)			
Total R&D	(All)			
Employee Growth	(All)			
Organization Type	(All)			
Basic/ Applied	(All)			
Innovation Level	(All)			

Country	Count of Basic	Count of Applied	Average Innovation Level	Average of HPC $M per Innovation
China	3.0	10.0	6.8	12.0
France		4.0	8.5	17.1
India	1.0		8.0	
UK	43.0	7.0	3.5	1.1
US	20.0	48.0	5.6	2.4
Italy		5.0	4.0	0.1
Grand Total	**67.0**	**74.0**	**5.0**	**3.1**

Key Findings: Success Stories

Note that an additional outcome of this research is an expansive list of HPC success stories

- These can be used to help explain the importance of HPC to funding bodies, key decision makers and the broader public
- IDC is writing up a number of them for broader dissemination

Download the report and models at:

www.hpcuserforum.com/ROI

HPC User Forum Update

HPC User Forum Steering Committee

IDC Analyze the Future

Newest Members:

- **Swamy Akapasu, General Motors**
- **Sharan Kalwani, Fermi Lab**
- **Keith Gray, BP**
- **Jysoo Lee, NISN (Korea)**
- **Suzy Tichenor, ORNL**

STEERING COMMITTEE

James Kasdorf
Chairman
Pittsburgh Supercomputing Center

Rupak Biswas
NASA Ames
Vice Chairman

Earl Joseph
IDC, Executive *Director*

Vijay Agarwala
Penn State University

Swamy Akasapu
General Motors

Alex Akkerman
Ford Motor Company

Doug Ball
The Boeing Company

Jeff Broughton
NERSC/Lawrence Berkeley National Lab

Paul Buerger
Avetec

Sharan Kalwani
Fermilab

Chris Catherasoo
California Institute of Technology

Jack Collins
National Cancer Institute

Steve Conway
IDC Research
Vice President

Steve Finn
Cherokee Information Services

Merle Giles
NCSA/University of Illinois

Keith Gray
BP

Doug Kothe
Oak Ridge National Laboratory

Jysoo Lee
NISN, Korea

Paul Muzio
City University of New York

Michael Resch
HLRS, University of Stuttgart

Vince Scarafino
Industry Expert

Suzy Tichenor
Oak Ridge National Laboratory

Recent Meetings

50th HPC User Forum: Boston (Sept. 2013)

51st HPC User Forum: Seoul (October 2013)

2014 HPC User Forum Meetings

April 5-7, 2014

Santa Fe, NM

- Industrial outreach programs
- ROI from HPC investments
- Storage innovation
- x86 processor alternatives
- High performance data analysis
- Computing for national security
- Global exascale initiatives
- IDC market update and forecast

September 15-15,2014: Seattle, WA

International Meetings Will Be Announced Soon

SC13
HPC Innovation Award Winners

We Are Collecting A Large Set Of HPC ROI Examples

We invite users to submit their examples at: www.hpcuserforum.com/innovationaward/

New HPC Innovation Award Program: www.hpcuserforum.com/innovationaward/

Sponsors – Thanks!

IDC
Analyze the Future

Gold Sponsor

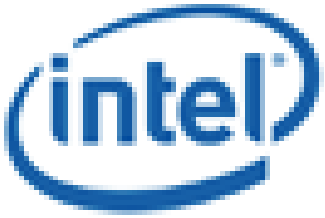

Silver Sponsors

Bronze Sponsors

HPC Award Program Goals

#1 Help to expand the use of HPC by showing real ROI examples:

1. Expand the "Missing Middle" – SMBs, SMSs, etc. by providing examples of what can be done with HPC
2. Show mainstream and leading edge HPC success stories

#2 Create a large database of success stories across many industries/verticals/disciplines

- To help justify investments and show non-users ideas on how to adopt HPC in their environment
- Creating many examples for funding bodies and politicians to use and better understand the value of HPC → to help grow public interest in expanding HPC investments
- For OEMs to demonstrate success stories using their products

Users Have to Submit the Value of the Accomplishment

Users are required to submit the value achieved with their HPC system, using 3 broad categories, following a very specific set of guidelines:

a) Dollar value of the HPC usage

- e.g., made $$$ in new revenues, saved $$$ in costs, made $$$ in profits, etc.

b) Scientific or engineering accomplishment

- e.g. discovered how xyz really works, develop a new drug that does xyz, etc.

c) Value to society as a whole

- e.g. ended nuclear testing, made something safer, provided protection against xyz, etc.

… and the investment in HPC that was required

The Judgment Process -- Clear, Fair And Transparent

The ranking of the accomplishments are done by only HPC USERS, following very specific rules.

A three step process is proposed:

1. First the submission has to be complete with a clear “value” shown
   - A number of the submissions were good, but needed a little more information – we have invited them to apply for the fall award
2. Secondly, an assessment is made to see that it is a realistic assessment of the value/returns
   - By the HPC User Forum Steering Committee
3. Then in cases where the value isn’t clear, or a deeper technical depth is required -- the final evaluation is by experts in the specific area/discipline

The Winners At SC13

Site	Lead	Area	Country
Imperial College London & NAG	NAG Hector CSE Team	Innovation	U.K
Spectraseis Inc, Denver, USA, & CADMOS, Univ. of Lausanne, Switzerland	Igor Podladtchikov & Yury Podladchikov	Both	U.S./ Switzerland
HydrOcean / ECN	David Le Touzé	Innovation	France
The Procter and Gamble Company	Kelly L. Anderson	Innovation	U.S.
Southern California Earthquake Center	SCEC Community Modeling Environment Collaboration	Innovation	U.S.
GE Global Research	Aero Acoustics team	Innovation	U.S.

The Winners At SC13

Site	Lead	Area	Country
Queen Mary University of London and, NAG	NAG Hector CSE Team	Innovation	U.K
EDISON Project - KISTI/NISN	Dr. Kumwon Cho	Both	S.Korea
Facebook	Avery Ching	Innovation	U.S.
Ford Werke GMBH	Dr. Burkhard Hupertz, Alex Akkerman	Innovation	Germany
Intelligent Light	Dr. Earl P.N. Duque	Innovation	U.S.
Oak Ridge National Lab		Innovation	U.S.
Princeton University	Dr. William Tang	Innovation	U.S.
GE Global Research	Dr. Masako Yamada	Innovation	U.S.

The Trophy For Winners

HPC Users – The Next Submission Deadline is April 15th

We invite users to submit

their ROI / achievement examples

at:

www.hpcuserforum.com/innovationaward/

In Summary

CUSTOMER NEEDS AND STRATEGIES

Buyer Insights: ROI Is Becoming More Important for Obtaining Funding

Advancing Science and National Security May No Longer Be Enough

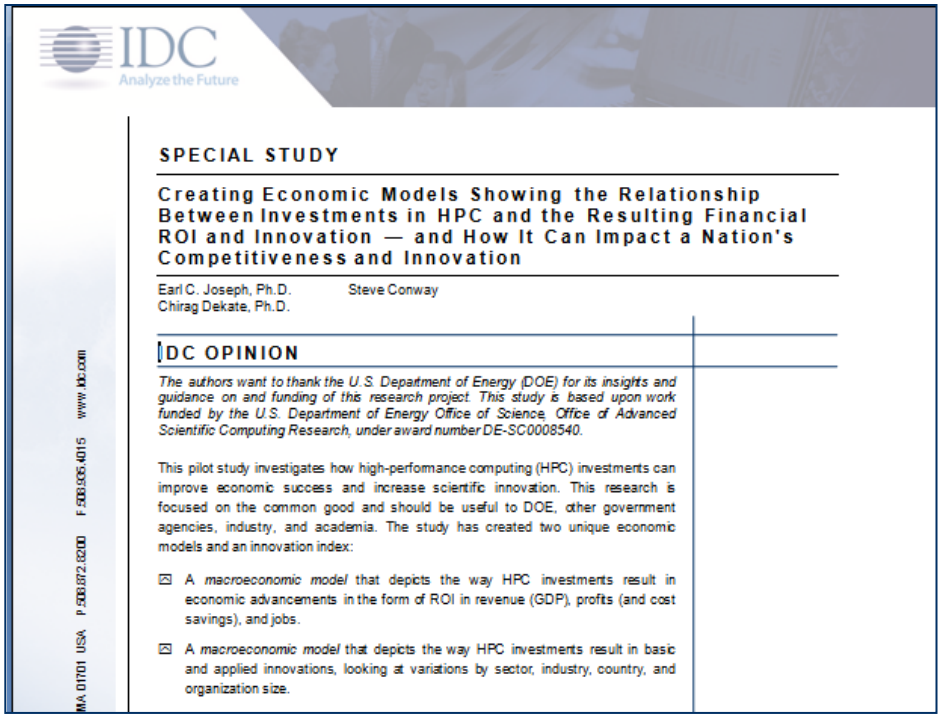

IDC
Analyze the Future

SPECIAL STUDY

Creating Economic Models Showing the Relationship Between Investments in HPC and the Resulting Financial ROI and Innovation — and How It Can Impact a Nation's Competitiveness and Innovation

Earl C. Joseph, Ph.D. Steve Conway
Chirag Dekate, Ph.D.

IDC OPINION

The authors want to thank the U.S. Department of Energy (DOE) for its insights and guidance on and funding of this research project. This study is based upon work funded by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, under award number DE-SC0008540.

This pilot study investigates how high-performance computing (HPC) investments can improve economic success and increase scientific innovation. This research is focused on the common good and should be useful to DOE, other government agencies, industry, and academia. The study has created two unique economic models and an innovation index:

- A *macroeconomic model* that depicts the way HPC investments result in economic advancements in the form of ROI in revenue (GDP), profits (and cost savings), and jobs.
- A *macroeconomic model* that depicts the way HPC investments result in basic and applied innovations, looking at variations by sector, industry, country, and organization size.

MA 01701 USA P 508.872.8200 F 508.935.4015 www.idc.com

- The Cold War arms race is becoming an economic race
- HPC is a proven accelerator of economic competitiveness
- High-end supercomputers can cost $200-500 million.
- ROI can be a scientific advance or corporate profit, new jobs
- More large HPC centers have industry outreach programs
- IDC's HPC ROI pilot study for DOE quantified 208 examples

Buyer Insights: The Exascale Race Is Off and Running

The Outcome Could Shift Global HPC Leadership Positions

- Some original goals (e.g., 20MW) will be pushed back beyond 2020.
- The U.S. effort has been slow to ramp up
- Software will be at least as important as hardware
- Will the extreme high end split off from mainstream HPC?

Buyer Insights: High Performance Data Analysis (HPDA) Will Transform HPC

Architectures Will Need to Shift Away from Extreme

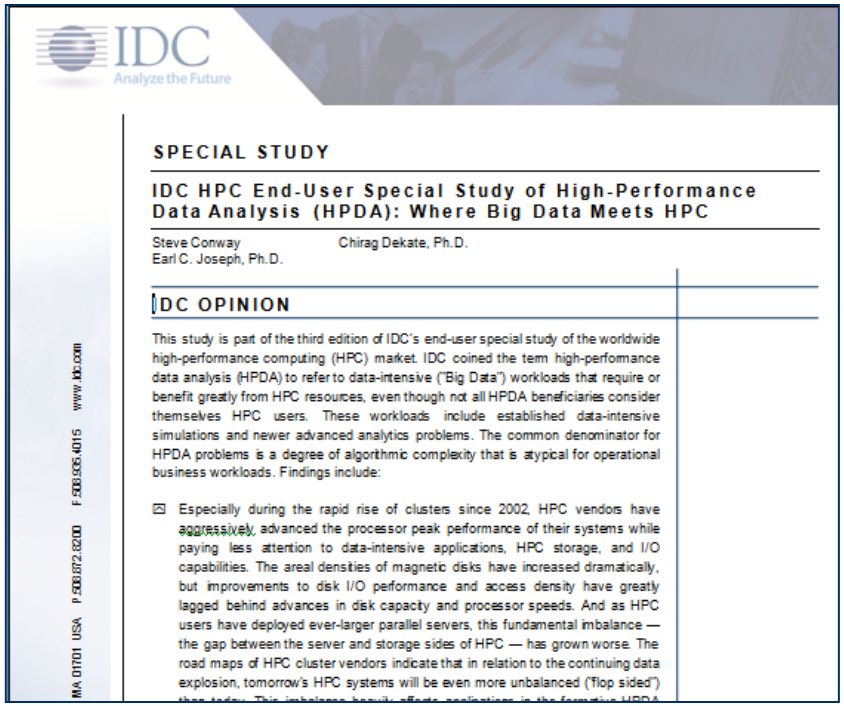

IDC
Analyze the Future

SPECIAL STUDY

IDC HPC End-User Special Study of High-Performance Data Analysis (HPDA): Where Big Data Meets HPC

Steve Conway Chirag Dekate, Ph.D.
Earl C. Joseph, Ph.D.

IDC OPINION

This study is part of the third edition of IDC's end-user special study of the worldwide high-performance computing (HPC) market. IDC coined the term high-performance data analysis (HPDA) to refer to data-intensive ("Big Data") workloads that require or benefit greatly from HPC resources, even though not all HPDA beneficiaries consider themselves HPC users. These workloads include established data-intensive simulations and newer advanced analytics problems. The common denominator for HPDA problems is a degree of algorithmic complexity that is atypical for operational business workloads. Findings include:

- Especially during the rapid rise of clusters since 2002, HPC vendors have aggressively advanced the processor peak performance of their systems while paying less attention to data-intensive applications, HPC storage, and I/O capabilities. The areal densities of magnetic disks have increased dramatically, but improvements to disk I/O performance and access density have greatly lagged behind advances in disk capacity and processor speeds. And as HPC users have deployed ever-larger parallel servers, this fundamental imbalance — the gap between the server and storage sides of HPC — has grown worse. The road maps of HPC cluster vendors indicate that in relation to the continuing data explosion, tomorrow's HPC systems will be even more unbalanced ("flop sided") than today. This imbalance heavily affects applications in the formative HPDA

MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

- Static searches will give way to dynamic pattern discovery
- 67% of surveyed HPC sites are running big data workloads
- HPDA includes data-intensive simulation & advanced analytics
- Most HPDA work will happen on clusters, but not everything
- Data movement is a big challenge
- Storage vendors will benefit

Essential Guidance: General

- Existing major challenges will remain inadequately addressed:
  - Weak application performance improvements
  - Highly parallel programming
  - System imbalance (the "memory wall")
  - Power and space usage
  - Software licensing costs
  - Ease-of-use – dealing with the growing system complexity
- HPC server growth will continue to outpace enterprise server growth through 2016
  - The HPC community will need to present strong ROI cases to compete for pressured government budgets

Essential Guidance

Major Customer Pain Points = Opportunity Areas For Vendors

New challenges affecting HPC data centers -- the increase in CPUs, heterogeneous processing, server units and data is creating significant IT challenges in:

- Managing HPC system complexity
- Balancing user needs & peak/Linpack performance expectations
- Managing data volumes and types
- Specifying and managing storage
- Planning for power/cooling and space
- **Application scaling and hardware utilization**
- Making optimal use of new processor and system designs
- Finding the appropriate role for private/public cloud computing

In Conclusion: Why HPC Is Projected To Keep Growing

IDC Analyze the Future

1. **The low half of the market is finally back to a recovery mode**
2. **It has become a competitive weapon**
3. **Governments view HPC leadership as critical**
   - For national pride, but more importantly for economic prosperity
   - It use to be 1 large supercomputer – now its multiple ones
4. **There are very critical HPC issues that need to be solved**
   - Global warming, alternative energy, safe NE, financial disaster modeling, healthcare, homeland security, …
   - And 3D movies and large scale games are fun
5. **The combination of big data and HPC is creating many new opportunities**
6. **At the same time, "live" science and "live" engineering costs have escalated**
   - And time-to-solution is months faster with simulations

Questions?

Please email:
hpc@idc.com

Or check out:
www.hpcuserforum.com