FINAL SCIENTIFIC/TECHNICAL REPORT FOR DOE--OFFICE OF SCIENCE

October 27, 2017

PROJECT TITLE: TO CREATE FULL-SCALE PREDICTIVE ECONOMIC MODELS ON ROI AND INNOVATION WITH HIGH PERFORMANCE COMPUTING

COVERING PERIOD: 08/01/2014 to 07/31/2017

APPROVED PROJECT PERIOD: 08/01/2014 to 07/31/2017

SUBMISSION DATE: October 27, 2017

RECIPIENT: IDC RESEARCH, INC.

WEBSITE (IF AVAILABLE): www.hpcuserforum.com/ROI (has the detailed research reports, slides and updated numerical database). Note: The full dataset is available to everyone to conduct their own research and modelling.

AWARD NUMBER: DE-SC0012576

WORKING PARTNERS: None.

COST-SHARING PARTNERS: None.

PI: Dr. Earl Joseph II, IDC (now called Hyperion Research Holdings, LLC)

SUBMITTED BY: Earl C. Joseph II, Ph.D., 612-812-5798, ejoseph@hyperionres.com

DOE PROJECT TEAM:

- **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 Doug Wade 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

PUBLICATIONS/PRESENTATIONS:

- The research has been submitted for publication to Dr. Charla Griffey-Brown, editor, Technology in Society (an Elsevier journal). Abstract requested and submitted.

Presentations & Publications:

- HPC Global Trends, meeting presentation, DOD, October 2016
- Global HPC Trends, DoD CIO Chief of Analytics, October 2016
- HPC Global Trends, meeting presentation, US IC, October 2016
- Recent Trends and Changes in the High End of the Supercomputer Market Space, IDC document for clients, September 2016
- HPC User Forum talk, HPC and industry, University of Oxford (UK), September 2016
- Global HPC Trends, US Army Research Lab in Aberdeen, September 2016.
- Critical HPC Policy Issues, DoD Special Access Program Director, September 2016
- HPC User Forum talk, HPC and industry, Beijing, September 2016
- Global HPC Technology Trends and Market Status, Center or the Advanced Study of Languages, University of Maryland, September 2106.
- IDC HPC Market Update and New ROI with HPC Results, slides from HPC User Forum, Austin, Texas, YouTube, September 2016
- Market Analysis Perspective: Worldwide Technical Computing, 2016. IDC document for clients, August 2016
- Status and Prospects for Big data and Cyber security, DoD Army G6 CTO, August 2016.
- HPC Technology Trends and Policies, DoD Army Strategic Integration Office, August 2106.
- ISC16 breakfast briefing in Germany, June 2016
- ISC European HPC conference, talk on HPC and industry, Frankfurt, Germany, June 2016
- Lenovo in Beijing, China, June 2016
- Alibaba Cloud in Beijing, China, June 2016
- Sugon in Beijing, China, June 2016
- Rescale, LLC in San Francisco, California, Spring 2016
- Blue Waters Symposium in Sun River, Oregon, Summer 2016
- High Performance Computing (HPC): IDC Reports Stunning ROI for Technology, Formtek, June 2016
- Statius and Prospects for HPC Technology and Policy, DoD USAF G6 CTO, May 2016.
- HPC Policy Study, ITIF Panel Discussion, April 2016
- HPC Status Update, High Performance Computing Advisory Committee, Council on Competitiveness, April 2106
- IDC Directions in San Jose and Boston, March 2016
- HPC Technology and Policy Overview, US Government Client, February 2016
- 10 Things CIOs Need to Know about High Performance Computing, IDC document for clients, March 2016
- HPC Accelerates Innovation for Small and Medium-Size Businesses, IDC document for clients, March 2016

- SC annual supercomputing conference, IDC market update presentation, November 2015
- Industrial Applications of High Performance Computing: Best Practices (book), eds. A. Osseyran and M. Giles, CRC Press, 2015
- IDC FutureScape: Worldwide High-Performance Data Analysis 2016 Predictions, IDC document for clients, November 2015
- HPC ROI: Invest a Dollar to Make $500-plus Reports IDC, article by John Russell, HPCwire, November 2015
- Datanami Subscribers Receive the IDC HPC ROI Research Update, article, Datanami, November 2015
- HPC ROI: Invest a Dollar to Make $500-plus Reports IDC, University of Delaware, November 2015
- How To Make HPC Happen in Europe, article by Leonardo Flores Anover and Augusto Burgeuno Arjona, Scientific Computing World, October 2015
- 4 IDC HPC User Forums: Tucson in 2015: Virginia (4/13), Colorado (9/8), Paris (10/12), and Munich (10/15)
- IDC Presents an Update on ROI with HPC, InsideHPC, September 2015
- ISC European HPC conference, talk on HPC and industry, Frankfurt, Germany, June 2015
- Applying HPC to Improve Business ROI, Scientific Computing, December 2014

PROJECT OBJECTIVES:

- The U.S. Department of Energy (DOE), the world's largest buyer and user of supercomputers, awarded IDC Research, Inc. a grant to create two macroeconomic models capable of quantifying, respectively, financial and non-financial (innovation) returns on investments in HPC resources. Following a 2013 pilot study in which we created the models and tested them on about 200 real-world HPC cases, DOE authorized us to conduct a full-out, three-year grant study to collect and measure many more examples, a process that would also subject the methodology to further testing and validation.
- A secondary, "stretch" goal of the full-out study was to advance the methodology from association toward (but not all the way to) causation, by eliminating the effects of some of the other factors that might be contributing, along with HPC investments, to the returns produced in the investigated projects.

  Note: the full data set, along with many analysis tables are provided in the form of an excel workbook at: www.hpcuserforum.com/ROI. We plan to continue posting updated versions of the data set as it grows.

BACKGROUND:

The authors thank DOE for its insights, guidance, 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-SC0012576.

High performance computers, also called supercomputers, have been anecdotally credited with making important contributions to national security, scientific progress, industrial innovation and the quality of human life—from safer automobiles to more predictable severe storms, better cancer-fighting drugs and even more absorbent diapers. A 2017 federally sponsored Hyperion Research study showed that sectors relying on supercomputers now represent about 55% of the U.S. GDP ($9.8 trillion) and 15.2 million American jobs. HPC's strategic importance for scientific and economic competitiveness, increasingly recognized by national and regional political leaders, has brought about larger investments in HPC and triggered a race among the U.S., China, Japan and the EU to build the fastest next-generation supercomputer. Public- and private-sector organizations planning to invest in these costly computers that have existed since the 1960s have had no adequate, direct method for measuring or predicting the returns on HPC investments, however. To justify the investments, which may exceed $150 million for a single computer, they typically have had to rely primarily on second-order effects, such as numbers of supercomputer-related papers published in scientific journals and supercomputer-supported doctoral degrees awarded at universities. To help remedy this situation, the U.S. Department of Energy (DOE), the world's largest buyer and user of supercomputers, awarded IDC Research, Inc. a grant to

create two macroeconomic models capable of quantifying, respectively, financial and non-financial (innovation) returns on investments in HPC resources. Following a 2013 pilot study in which we created the models and tested them on about 200 real-world scientific and industry HPC projects, DOE authorized us to conduct a full-out, three-year grant study to collect and measure many more examples, a process that would also subject the methodology to further testing and validation. This research is focused on the common good and should be useful to DOE, other government agencies, industry, and academia.

WORK PLANNED:

- **Original Research Vision:**

  Phase I (Year 1) – The goal is to create the actual ROI models with a full data set (at least twice as much data as in the pilot study) in Phase I, growing to at least 8x in size by Phase III. This is needed to enable predictive models that support a fuller understanding of the relationships, to provide enough data/analysis to start making predictive results, and to refine the models as needed.

  - The goal is to have enough data to start making statistically sound correlations between industries, between countries and between different sizes of organizations.

  Phase II (Year 2) – Expand the data set by at least 2x more and include more countries and industries. Motivate a larger set of nations to contribute deeper data samples.

  - The goal is to have enough data to enable statistically sound correlations between industries, between countries and between different sizes of organizations – and cross-correlations such as industries by country, and organization size by country.

  Phase III (Year 3) – Expand the data set again by at least 2x more; harden the predictive nature of the models by conducting additional research to "test" correlations; identify non-HPC contributors to investment returns, in order to eliminate their effects and advance the methodology from correlation/association along the path toward isolatingn HPC causation.

  - The goal is to conduct enough research to show both strong statistical correlation – and quasi-causation between investments in HPC and the resulting ROI and innovation.
  - In this phase, the scenario testing should improve to be very robust and directly useful for making national policy decisions.

- **Scope Changes:**

  We found that it was much more difficult to collect thousands of data points than hundreds, and had to refocus our efforts in a number of areas. We were able to collect considerable useful data, but had to

throw out over 50% of the examples collected, due to a lack of trust in the data collected. So we weren't able to conduct the causation analysis phase of the research.

We were able to expand the project by applying the results to full data centers as a test of the robustness of the approach and data. The test cases were NCSA (USA), HLRS (Germany) and Riken (Japan).

- Data collection is more difficult than we had expected:
    - Over 50% of the examples are rejected for various reasons
    - Many researchers agree to participate, but a much smaller number actually participate
    - Some sites still require multiple face-to-face visits to collect reliable results

Solutions added:

- We added a person in China to directly collect ROI examples, each requiring one or more face-to-face meetings
    - Results so far: 126 ROI/ROR examples
- We added a part time person in Japan, where each survey also requires one or more face-to-face meetings
    - Results so far: 79 ROI/ROR examples
- We added a program to motivate HPC center directors to help push their researchers to participate
    - Results so far: Germany has started, Riken & Blue Waters were the test-beds

The term "innovation ROI" isn't well understood by many people in a number of countries

- So we now refer to it as ROR = Return on Research

Three innovations indexes are now being used:

- Based on the importance of the innovation
- How broadly the innovations impact different organizations
- A combined score to create innovation "Class" levels

SIGNIFICANT ACCOMPLISHMENTS AND STATUS:

Hypotheses.

- Our starting hypothesis, based on a limited number of previously quantified historical examples and considerable unquantified and partially quantified anecdotal evidence, was that financial investments in HPC to support scientific and engineering activities often produce attractive returns—in the form of innovation associated with fundamental, pre-competitive research activities (return on research) or in the form of additional revenue and profit associated with post-competitive, commercial/industrial activities involving HPC.
- We further assumed that the units of activity to be evaluated for HPC investment returns would need to be well defined and finite in scope, and that completed projects having specific beginning and end dates would serve this purpose better than ongoing,

less clearly punctuated research activities that in some cases may continue for decades or more.

- Another important assumption was that our macroeconomic models and related methodology would need to be very direct if they were going to be able to capture first-order effects of investments in HPC and improve on existing second-order metrics, such as numbers of peer-reviewed scientific papers published or numbers of doctoral degrees awarded in scientific-technical disciplines. This is why we decided to ask projects leaders themselves to rate the importance and impact of their projects on simple scales we created, and then to have a experts in the projects' domains independently rate the same projects on the same scales, without have seen the ratings made by the project leaders.
- We also assumed that HPC resources would be important contributors to the returns from the investigated projects, but would not be the only contributors. In other words, we assumed that our methodology could closely *correlate* the project returns with their HPC investments but could not isolate the HPC investments as the *exclusive causes* of the returns. To help address this limitation, we directly asked survey respondents how large a part HPC investments played in obtaining their project returns. We also planned to augment the methodology employed in our predecessor 2013 pilot study so that we could advance findings from correlation at least partway toward causation, by identifying and eliminating the effects of other material contributors to the project returns; but for reasons explained elsewhere in this report, we did not achieve this goal and had to postpone it for a future phase of research by us or others.
- Finally, we assumed that collecting large numbers of additional examples of projects would be as easy as gathering the first 200 we collected for the 2013 pilot study. That turned out not to be true, although we raised the total number of evaluated projects to more than 700.

Methodology

- For the post-pilot study research phase that is the main focus of this report, we used the two macroeconomic models we had developed and refined through live testing ("hardened") in the 2013 pilot study. Each of the two models employed easily understood scales which survey respondents (typically PIs of the projects under investigation) used to rate, separately, the importance and impact of their completed projects. We used one model for evaluations of fundamental, pre-competitive research projects (return on research) and the other model for evaluations of post-competitive projects that had financial outcomes.
- Expert reviewers employed the same scales as the respondents to rate the projects, without having seen the ratings of the project participants. Serving as the first-level reviewers were the members of the steering committee of the HPC User Forum (www.hpcuserforum.com), a multinational group of representing leading HPC organizations in government, academia and the private sector. In rare cases where no

one within this group felt qualified enough to rate a completed project, that task was assigned to someone with deeper expertise on the topic in question.

- Self-evaluations by the project participants were then compared with independent evaluations by the expert reviewers, to see how well ratings of the same projects correlated with each other. If ratings of the same projects had often been far different from each other, we would have had to conclude that our models were too flawed to apply to necessarily subjective human judgments.

Main Findings

- The models and methodology continued to work well in the post-pilot study phase of the research. With rare exceptions, correlations remained strong between the self-ratings of project participants and the independent, blinded ratings of domain experts. "Grade inflation" by self-grading project participants proved not to be a notable concern.
- We added more than 500 evaluated projects to the 200 we had collected and measured in the pilot study. We learned, however, that the first 200 represented "low hanging fruit" that was easier to harvest than the next few hundred projects, so even with more effort than we anticipated, we did not add as many projects as we set out to add. To ensure the robustness of the methodology and to enable trustworthy sorting of results by category (nation, scientific/engineering domain), we believe, would require applying it to about 2,000 projects, rather than the 700 we've measured so far.
- Despite the methodological limitation (strong correlation with HPC investments without eliminating less-important other factors contributing to project returns), and despite the shortfall in numbers of evaluated projects compared with our original goal, the research in this phase produced two notable positive results:
    1. When the number of evaluated projects more than tripled from 200 in the pilot study to over 700 in the latest research phase, the average ratings and the average returns for the projects increased in some areas, but overall did not change dramatically. The average additional revenue associated with each $1 investment in HPC resources in the private sector projects has settled so far at $551, and the average added profit associated with each $1 of HPC investment has settled at an equally impressive $52. We expect these numbers would continue to change as more projects are evaluated, but given this history we would not expect major changes to appear.
    2. Other government-supported organizations that had seen the findings of the pilot study engaged Hyperion Research to use the models and methodology we employed in the pilot study to conduct studies of their HPC users (see next section for details). We believe this rapid demand confirms the strong need for a more direct method, of the kind we developed, for documenting the returns on government-supported investments in HPC resources.

Highlights of the findings:

New Findings: Primary Financial ROI Results

Results continue to indicate very substantial returns for investments in HPC:

- **$515 dollars on average in revenue** per dollar of HPC invested.
- **$52 dollars on average of profits** (or cost savings) per dollar of HPC invested.

Organization Type	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $
Academic	1,150.2	41.3
Company	441.8	41.3
Government	1,205.7	141.5
Grand Total	**515**	**52**

IDC Analyze the Future

19

New Findings: Financial ROI Model – By Industry

Results continue to indicate very substantial returns for investments in HPC:

- **2,335 jobs were created across these financial ROI projects**

Q3: Industry	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $	Sum of Total Jobs Added	Count of Accomplishment Type
Academic	1150.2	44.3	14	21
Defense	75.0	5.3	0	2
Financial	725.7	80.5	602	31
Government	1205.7	140.8	42	15
Insurance	71.4		5	1
Life Sciences	160.0	40.9	48	13
Manufacturing	83.0	20.2	678	43
O&G	418.6	46.0	100	10
Retail	30.3	12.3	49	3
Telecomm	10.0	10.0	420	2
Transportation	1804.3	15.6	377	7
Grand Total	**515**	**52**	**2,335**	**148**

IDC Analyze the Future

20

New Findings: Job Creation

Job creation by sector:

Sector	Sum of Total Jobs Added	Average of HPC $K per New Employee
Academic	1,632	$107 K
Government	196	$473 K
Industry	3,920	$266 K
Grand Total	**5,748**	**$270 K**

IDC Analyze the Future

23

New Findings: Innovations – Basic vs. Applied

Sector	Count of Q8: Basic	Count of Q8: Applied	Count of Q8: Basic/ Applied
Academic	265	80	341
Government	37	27	61
Industry	29	92	123
Grand Total	**331**	**199**	**525**

IDC Analyze the Future

28

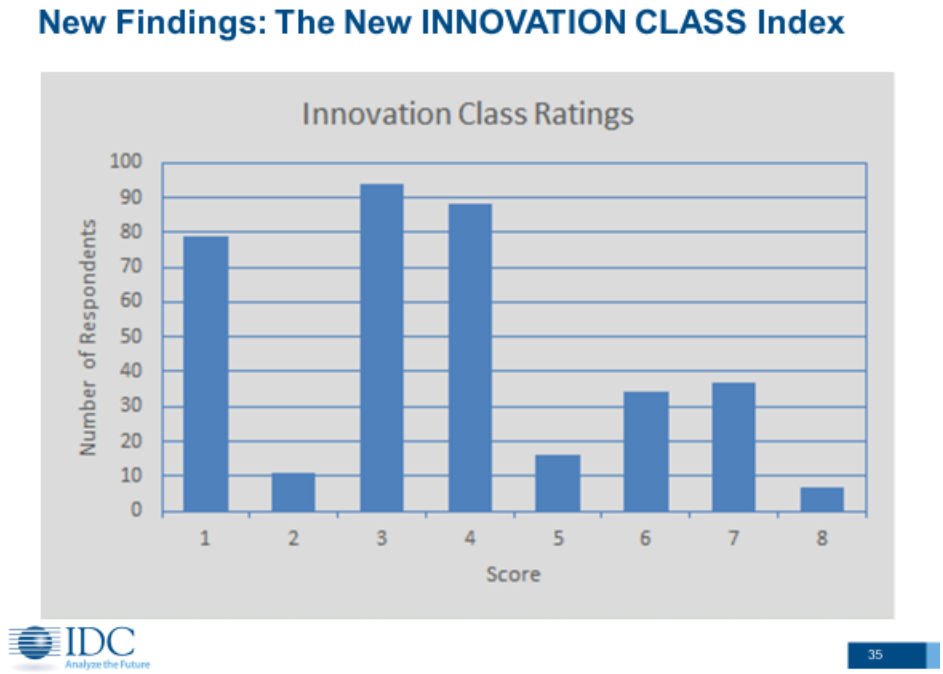

USED BY OTHER ORGANIZATIONS:

- See exhibit 3 for 14 examples of HPC ROI data points.
- See exhibit 4 for the use by NSF/NCSA

As noted above, based primarily on the findings of the pilot study, multiple government-supported organizations in the United States and beyond engaged Hyperion Research to conduct studies using our new methodology on their HPC users. We believe this rapid response to the pilot study findings confirms the strong need for a more direct way to document financial and non-financial returns on government-supported investments in HPC resources. Studies of this kind we have completed to date are the following:

- *An Investigation and Evaluation of the Scientific Results from the Blue Waters Supercomputer System*. Public study for NCSA. 2017
- *The Status and Prospects of European Suppliers of High Performance Computing Products and Services.* Private study for EXDCI (Europe). 2017
- *A Study of the Ripple Effects from the K Computer and the Potential Impacts of the Post-K Computer.* Public study for RIKEN (Japan). 2016
- *The Impact of the High Performance Computing Center Stuttgart on Industrial Partners.* Private study for HLRS (Germany). 2016
- *Funding Models for National Computing Centers in Europe*. Private study for SurfSARA, The Netherlands' national supercomputing center. 2015

PATENTS:

None.

TRAINING AND PROFESSIONAL DEVELOPMENT:

The results should be useful for showing how investments in HPC can help a company or university or government agency improve its capabilities in many ways. We have presented the results to many organizations to help teach how to better look at investments in HPC as well as investments in R&D. Below is a table that helps show these results by the different types of innovation areas:

New Findings: The Mix Of Innovation Types (For The Innovation Projects)

The average cost of an innovation in the INNOVATION ROR projects was $12.7 million:

Q4: Primary Innovation / ROI	Count of Accomplishment Type2	Average of HPC $M per Innovation	Count of Q8: Basic	Count of Q8: Applied
Better Products	95	$3.5 M	22	72
Cost Savings	16	$1.6 M	8	9
Created New Approach	233	$1.2 M	185	50
Discovered Something New	40	$3.0 M	27	13
Helped Society	42	$6.4 M	21	21
Scientific Breakthrough	55	$80.3 M	49	8
Support Research Programs	44	$36.7 M	19	26
Grand Total	**525**	**$12.7 M**	**331**	**199**

IDC Analyze the Future

27

CONCLUSIONS:

This new approach to measuring the economic and scientific value of HPC provides new insights into the value of investing in R&D&E, investing in supercomputing, and how to better optimize these investments. The model provides a tool of looking at national HPC investments, and can show the potential value of these investments. In the future, the team would like to expand the research in these areas:

1. Collect more data points show that comparisons can be made between countries, industries & sectors.

2. Expand the model so that it can be used to PREDICT the future value of a supercomputer investment. For example, to show the potential value of an exascale system, and how it varies by sites.
3. Expand the model to be a tool for measuring the overall value of an HPC datacenter.

Again, the research team wants to express its gratitude and appreciation to DOE for funding this research.

EXHIBITS:

EXHIBIT 1:
THE ROI PILOT STUDY RESULTS (2013)

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

 Oct-17

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

 Oct-17

Project Overview: Why It Is Key To DOE

IDC Analyze the Future

- World scientific leadership and innovation leadership are becoming more dependent on the use of HPC/supercomputers every year
- Economic leadership increasingly directly results from a nation's or an industry's or an enterprise's application of supercomputers in innovative and productive ways
- Many countries are putting in place plans to gain leadership in innovation and economic progress by more broadly applying HPC/supercomputing across many different industries and segments (like China, Russia, Europe, Japan and other Asian countries)

 Oct-17

The Financial ROI Models That Were Developed

IDC Analyze the Future

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

 Oct-17

The Innovation Models That Were Developed

IDC Analyze the Future

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

 Oct-17

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

 Oct-17

Key Findings

- **IDC is able to collect the required data across a broad set of organizations, with enough detail to create the economic models and the innovation index**
- **Early results indicate very substantial returns related to investments in HPC, on average:**
  - $356.5 in revenue per dollar of HPC invested
  - $38.7 of profits/cost savings per dollar of HPC invested
  - The average HPC investment per innovation was $3.1M
- **Note that an additional outcome of this research is an expansive list of HPC success stories**

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

For each sector we need 4 basic ratios

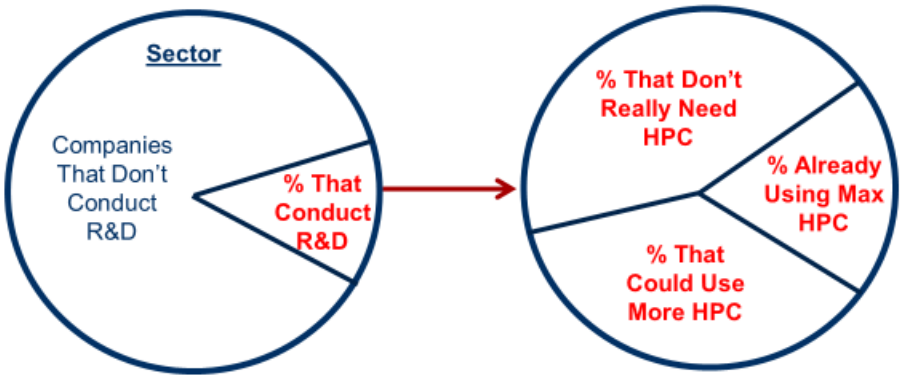

Note: IDC has conducted over 30,000 light phone calls for this data. We will likely require 5x to 10x more surveys

Pilot Study Update

Research Schedule: 9/1/12 to August 2013

1. ✓ **October 2012 – finalize the plan, surveys and overall research approach**
2. ✓ **November 2012 – conduct a limited number of surveys to both see what can be collected and start testing the two models**
3. ✓ **December 2012 – refine the survey guide and models as needed**
   - Refine as needed
4. ✓ **January – March 2013 – start the broad data collection/surveys – and review the initial results**
   - We noticed that we need a larger data set
5. ✓ **April/May 2013 – populate the two models with more data**
   - Refine as needed
6. ✓ **July 2013 – collect a few more data points, and start disseminating the results**
7. ✓ **October 2013 – finalize the report and continue disseminating the results**

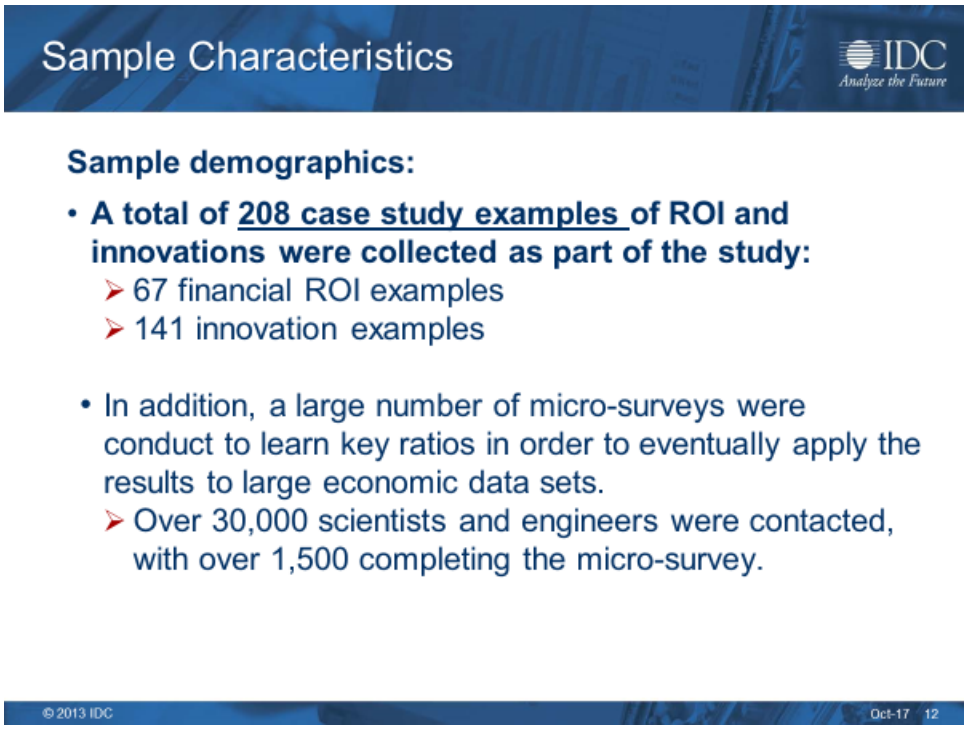
Sample Characteristics
IDC
Analyze the Future
Sample demographics:
• A total of 208 case study examples of ROI and innovations were collected as part of the study:
➢ 67 financial ROI examples
➢ 141 innovation examples
• In addition, a large number of micro-surveys were conduct to learn key ratios in 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.
© 2013 IDC
Oct-17 12


IDC
Analyze the Future
Pilot Study Results: Financial ROI
© 2013 IDC
Oct-17 13

Key Findings: Primary Financial ROI Results

IDC Analyze the Future

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

IDC Analyze the Future

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)

	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

IDC Analyze the Future

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**

Key Findings: Additional Financial ROI Results

IDC Analyze the Future

3. **The average number of years before returns started was 1.9 years.**

4. **The overall average HPC investment cost for jobs created was $93,000.**
   - 42 sites reported job creation. On average 29.8 jobs were created from HPC investments at these sites.
   - A total of 1,251 jobs were created from the 42 financial ROI examples

Pilot Study Results: Innovation

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

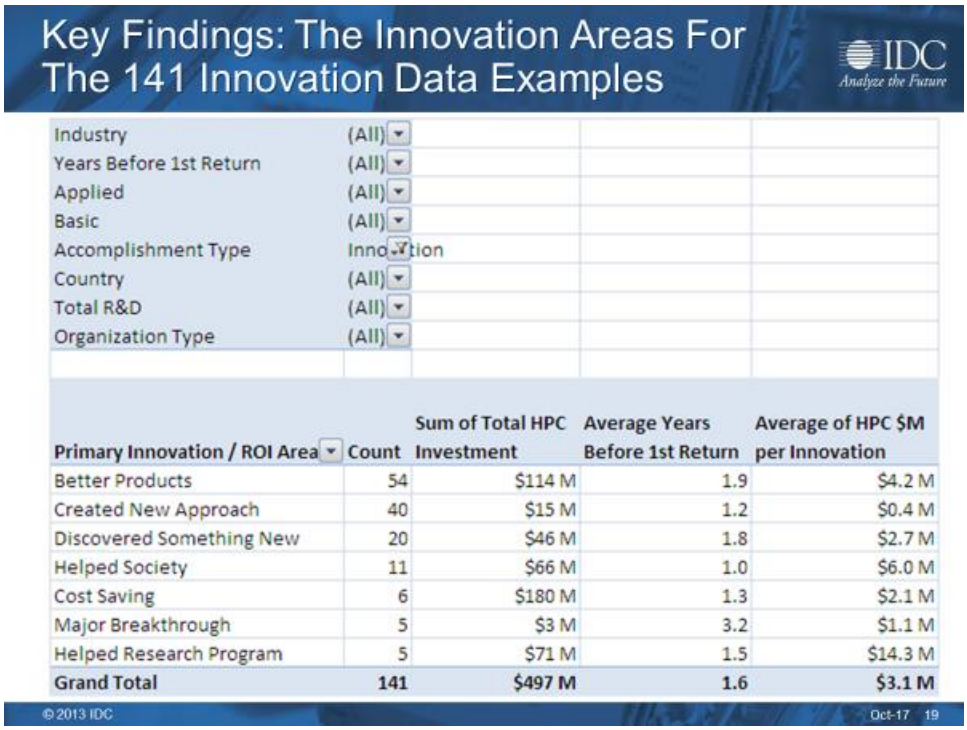

Industry	(All)
Years Before 1st Return	(All)
Applied	(All)
Basic	(All)
Accomplishment Type	Innovation
Country	(All)
Total R&D	(All)
Organization Type	(All)

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

IDC Analyze the Future

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

IDC Analyze the Future

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

IDC Analyze the Future

FIGURE 12

HPC Innovation Index Scale Results: All Respondents

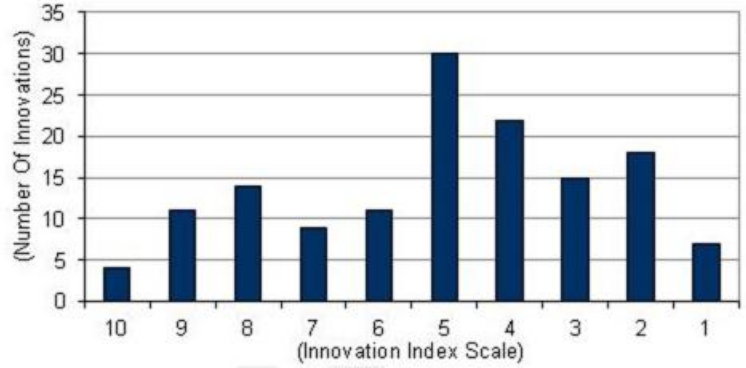

N = 141

Source: IDC 2013

Key Findings: The New Innovation Index Scores – The 67 Basic Research Innovations

IDC Analyze the Future

FIGURE 13

HPC Innovation Index Scale Results: Basic Innovations

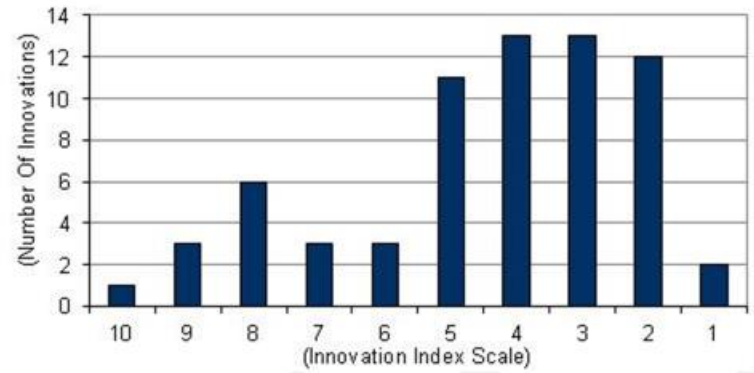

N = 67

Source: IDC 2013

Key Findings: The New Innovation Index Scores – The 74 Applied Research Innovations

IDC Analyze the Future

FIGURE 14

HPC Innovation Index Scale Results: Applied Innovations

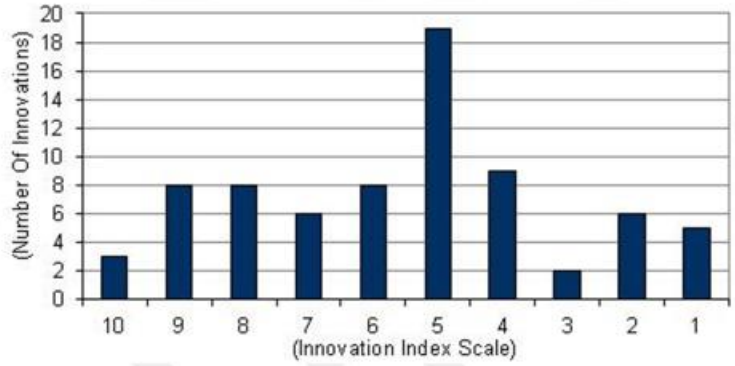

N = 74

Source: IDC 2013

Key Findings: Locations of the Innovations

IDC Analyze the Future

6. **Academic and industrial sites represented the bulk of the innovations in the sample.**
   - Most of the basic research innovations were in academia, while most of the applied research innovations were in industry.

7. **Government innovations were higher on the innovation index scale (averaging 7.0).**
   - Innovations in industry ranked lower at 5.7.
   - Academic innovations averaged 3.9.

Key Findings: The Innovation Index By Country

IDC Analyze the Future

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: Size Of the HPC Investments

IDC Analyze the Future

8. **The total amount of HPC investments made in the sample was $496.7 million.**
   - Government sites tend to invest more per innovation – at $4.4 million in HPC on average.
   - Academic sites invested less – averaging $2.5 million in HPC per innovation.
   - Industrial sites invested on average $3.8 million in HPC per innovation.

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.

Some Notes About The Data: Outliers

Note: there were two sites with considerably higher returns that were removed from the dataset as outliers:

1. **$10,057 revenue dollars per dollar of HPC invested at a finance company**
   - If included, this raises the overall revenue ROI from $365.5 to $536.2
2. **$1,350 in profit dollars per HPC dollar invested at an oil & gas company**
   - If included, this raises the overall profit ROI from $38.7 to $65.5

Dissemination Program

- **The report and excel models are posted at: www.hpcuserforum.com/ROI**
- **Once approved by DOE, IDC will send the report and models to ~4,500 people in the broader HPC community**
- **IDC will brief the community at SC13**
  - On Tuesday during the IDC breakfast briefing
  - On Wednesday, at 5:30 – 7:00pm, during the Cost-Benefit Quantification for HPC: An Inevitable Challenge BOF
- **Press release**
- **Many briefings have already been given – does DOE have additional briefings that IDC should conduct?**

Future ROI Research Plans (Proposed)

Phase I (Year 1) – The goal is to create the actual ROI models with a full data set at least 2x in Phase I, growing to at least 8x in size by Phase III. This is needed to create predictive models a fuller understanding of the relationships, to provide enough data/analysis to start making predictive results, and to refine the models as needed.

- The goal is to have enough data to start making statistically sound correlations between industries, between countries and between different sizes of organizations.

Phase II (Year 2) – Expand the data set by at least 2x more, and including more countries and industries. Motivate a larger set of nations to contribute deeper data samples.

- The goal is to have enough data to make strong statistically sound correlations between industries, between countries and between different sizes of organizations – and cross-correlations like industries by country, and organization size by country.

Phase III (Year 3) – Expand the data set again by at least 2x more, and focus on hardening the predictive nature of the models by conducting additional research to "test" correlations. At this phase it the tie to large general economic data sets should be robust enough to be able to hand-off to other government organizations.

- The goal is to conduct enough research to show both strong statistical correlation – and causation between investments in HPC and the resulting ROI and innovation.
- In this phase the scenario testing should improve to be very robust and directly useful for making national policy decisions.

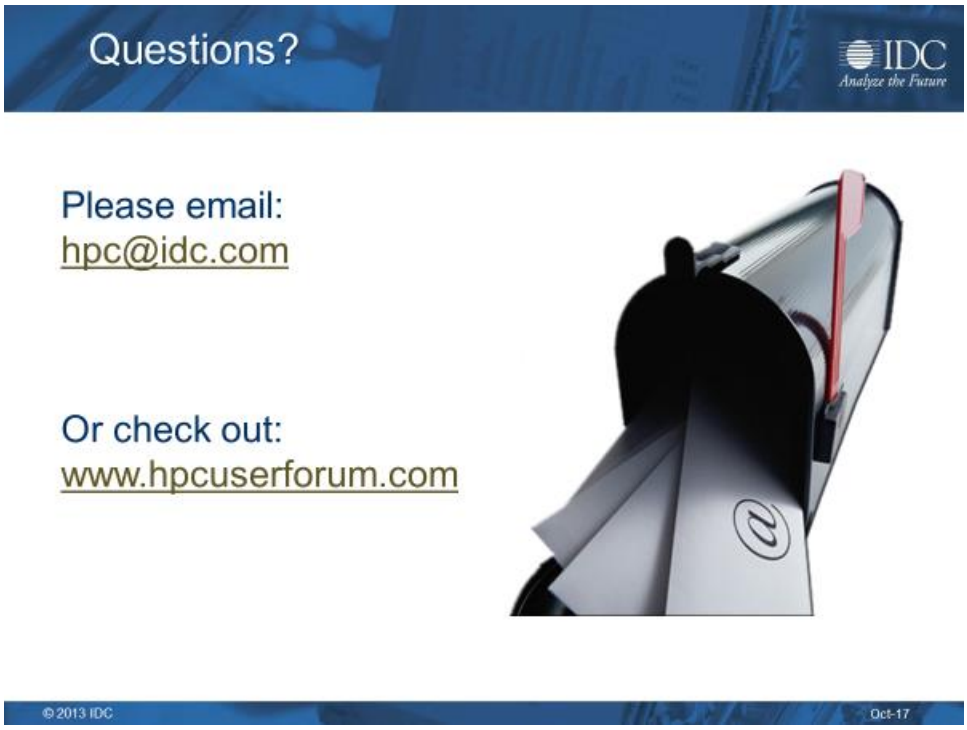
Questions?
IDC
Analyze the Future
Please email:
hpc@idc.com
Or check out:
www.hpcuserforum.com
© 2013 IDC
Oct-17

EXHIBITS:

EXHIBIT 2:
THE MAIN ROI STUDY RESULTS
AS PRESENTED TO DOE IN DECEMBER 2016

IDC
Analyze the Future

IDC HPC ROI Research Update:

Economic Models For Financial ROI And Innovation From HPC Investments

Earl Joseph, ejoseph@idc.com
Steve Conway, sconway@idc.com
Robert Sorensen, bsorensen@idc.com

December, 2016

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-SC0012576.

- **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 Doug Wade 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

IDC Analyze the Future

Project Overview: Why It's Key To DOE

- World scientific leadership and innovation leadership are becoming more dependent on the use of HPC/supercomputers every year
- Economic leadership increasingly results from a nation's or an industry's or an enterprise's application of supercomputers in innovative and productive ways
- Many countries are putting in place plans to gain leadership in innovation and economic progress by more broadly applying HPC/supercomputing across many different industries and segments (like China, Russia, Europe, Japan and other Asian countries)
- And it supports the new White House NSCI initiative

IDC Analyze the Future

Research Background

- **2013 pilot study:**
  - Tested 3 approaches and set the models
  - Populated the models with 208 cases of scientific innovation and industrial ROI (it is now at 673 examples)
  - Created a new innovation index
- **Three-year study: sponsored by DOE Science/NNSA**
  - Refine the models as needed
  - Collect many more RPOI examples/data points:
    - Dense collection in the U.S., plus around the world
    - Focus on Japan, China and Germany
    - Drive to get more job creation data
  - Drive towards publication

IDC Analyze the Future

Background: Project Overview

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

The study includes large scale data collection to populate two unique models:

1. A macroeconomic model which depicts how HPC investments result in economic advancements in the form of ROI, GDP growth and jobs
2. **Now Three** Innovation Indexes that measure and compares innovation levels, based on the level of applying HPC computing resources towards scientific and technical advancement

IDC Analyze the Future

The Financial ROI Models Used

The Financial ROI models:

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

The ROI models show variances by:

- Industry & sector
- Country
- Organization size

IDC
Analyze the Future

Changes Made in 2016

Data collection is more difficult than we had expected:

- Over 50% of the examples are rejected for various reasons
- Many researchers have agreed to participate, but a much smaller number actually participate
- Still requires multiple face-to-face visits to collect results

Solutions added:

- Added a person in China to directly collect ROI examples (Mengland Shi), each requires multiple face-to-face meetings
  - Results so far: 126 ROI/ROR examples
- Added a part time person in Japan (Nishi Katsuya), each survey requires multiple face-to-face meetings
  - Results so far: 79 ROI/ROR examples
- Added a program to motivate HPC center directors to help push their researchers to participate
  - Results so far: Germany has started, Riken & Blue Waters were the test-beds

IDC Analyze the Future

7

Changes Made in 2016

- The term "innovation ROI" isn't well understood by many people, in a number of countries
  - So we now refer to it as ROR = Return on Research

IDC Analyze the Future

8

The Updated Innovation Models Used

Three innovations indexes are now being used:

1. **Based on the importance of the innovation**
2. **How broadly the innovations impact different organizations**
3. **A combined score to create innovation "Class" levels**

The innovations are also sorted by the primary area:

1. Better Products
2. Major Scientific Breakthrough
3. Cost Saving
4. Created New Approach
5. Discovered Something New
6. Helped Society
7. Helped Research Program

IDC Analyze the Future

Changes Made in 2016: New Innovation "Class" Scale

- Expanded the innovation indexes by adding a new innovation "Class" scale:

Class 1 innovations – One of the top 2-3 innovations in a field over the last ten years PLUS useful to over 10 organizations

Class 2 innovations -- One of the top 5 innovations in a field over the last ten years PLUS useful to over 10 organizations

Class 3 innovations – One of the top 5 innovations in a field over the last ten years PLUS useful to at least 5 organizations

Class 4 innovations – One of the top 10 innovations in a field over the last ten years PLUS useful to at least 5 organizations

Class 5 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at over 10 organizations

Class 6 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at least 2 organizations

Class 7 innovations – One of the top 50 innovations in a field over the last ten years PLUS useful to at least 2 organizations

Class 8 innovations – The rest of the innovations in the study

IDC Analyze the Future

10

Changes Made in 2016: New Innovation "Class" Scale – Blue Waters Example

FIGURE 8

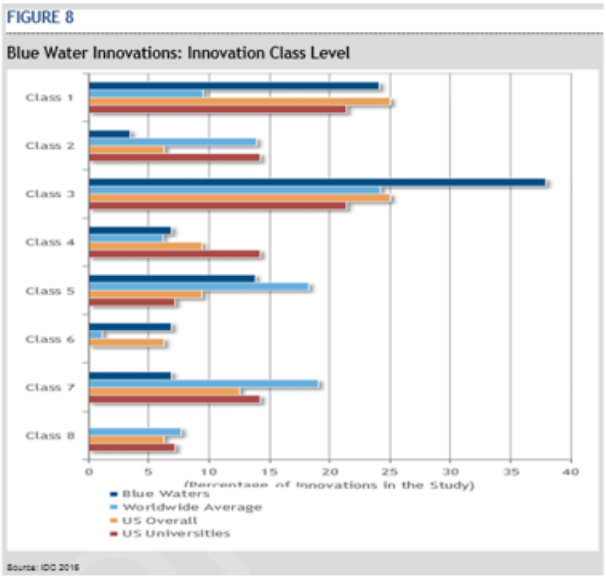

IDC Analyze the Future

11

The Innovation Models Used: Additional Data Collected

Plus the innovations are sorted by:

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

The Innovation models can be sorted for variances by:

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

IDC Analyze the Future

December 2016
Research Update

IDC Analyze the Future

13

Current Database Demographics

- Total ROI & ROR examples (673):
  - Financial ROI examples = 148
  - Innovation ROR examples = 525

Accomplishment Type	Count of Accomplishment Type
Financial ROI	148
Innovation	525
Grand Total	**673**

Note: This is up from July 30th, 2015 when we were at 329

Note: We are finally at a stage in the research that we feel comfortable starting the publication process!

IDC Analyze the Future

14

Current Database Demographics

- By country:

Country	Count of Accomplishment Type
China	126
France	17
India	1
UK	120
US	304
Italy	9
Germany	10
Canada	1
Slovenia	2
Netherlands	1
Australia	2
Japan	79
South Africa	1
Grand Total	**673**

IDC Analyze the Future

15

Current Database Demographics

- By sector:

Sector	Count of Accomplishment Type
Academic	362
Government	76
Industry	235
Grand Total	**673**

IDC Analyze the Future

16

Current Database Demographics

- By industries:

Q3: Industry	Count of Accomplishment Type
Academic	363
Agriculture	1
Defense	3
Entertainment	6
Environmental Safety	1
Financial	41
Government	76
Insurance	1
Life Sciences	31
Manufacturing	114
O&G	21
Retail	3
Telecomm	4
Transportation	8
Grand Total	**673**

IDC Analyze the Future

17

December 2016
Study Results:
The Financial ROI From HPC

IDC
Analyze the Future

18

New Findings: Primary Financial ROI Results

Results continue to indicate very substantial returns for investments in HPC:

- **$515 dollars on average in revenue** per dollar of HPC invested.
- **$52 dollars on average of profits** (or cost savings) per dollar of HPC invested.

Organization Type	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $
Academic	1,150.2	41.3
Company	441.8	41.3
Government	1,205.7	141.5
Grand Total	**515**	**52**

IDC Analyze the Future

19

New Findings: Financial ROI Model – By Industry

Results continue to indicate very substantial returns for investments in HPC:

- **2,335 jobs were created across these financial ROI projects**

Q3: Industry	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $	Sum of Total Jobs Added	Count of Accomplishment Type
Academic	1150.2	44.3	14	21
Defense	75.0	5.3	0	2
Financial	725.7	80.5	602	31
Government	1205.7	140.8	42	15
Insurance	71.4		5	1
Life Sciences	160.0	40.9	48	13
Manufacturing	83.0	20.2	678	43
O&G	418.6	46.0	100	10
Retail	30.3	12.3	49	3
Telecomm	10.0	10.0	420	2
Transportation	1804.3	15.6	377	7
Grand Total	**515**	**52**	**2,335**	**148**

IDC Analyze the Future

20

New Findings: Financial ROI Model – By Country

Country	Average of Revenue $ per HPC $	Average of Profit or Cost Saving $ per HPC $	Sum of Total Jobs Added	Count of Accomplishment Type
China	9	3	196	21
France	593	81	30	6
UK	635	48	1376	47
US	373	39	680	50
Italy	10	8	0	3
Germany	15	16	25	7
Slovenia		65	0	1
Japan	2709	278	28	13
Grand Total	**515**	**52**	**2,335**	**148**

Note: The data set today isn't complete enough to make full country-to-country comparisons. There are also differences by country on how sites calculate returns and costs, this causes variations in the data.

IDC Analyze the Future

21

New Findings: Job Creation

The overall average HPC investment cost per job created was $270K

- 225 sites reported job creation
  - A total of 5,748 jobs were created from these examples (at the 225 sites):
    - 2,335 jobs were created from the financial ROI examples
    - 3,413 jobs were created from the innovation ROR examples
  - On average 25.6 jobs were created from the HPC projects at these sites

IDC Analyze the Future

22

New Findings: Job Creation

Job creation by sector:

Sector	Sum of Total Jobs Added	Average of HPC $K per New Employee
Academic	1,632	$107 K
Government	196	$473 K
Industry	3,920	$266 K
Grand Total	**5,748**	**$270 K**

IDC Analyze the Future

23

New Findings: Job Creation

Job creation by industry:

Q3: Industry	Sum of Total Jobs Added	Average of HPC $K per New Employee
Academic	1,632	$107 K
Financial	857	$255 K
Government	196	$473 K
Insurance	5	$175 K
Life Sciences	134	$88 K
Manufacturing	1,920	$365 K
O&G	158	$291 K
Retail	49	$152 K
Telecomm	420	$179 K
Transportation	377	$238 K
Grand Total	**5,748**	**$270 K**

IDC Analyze the Future

24

New Findings: Job Creation

Job creation by country:

Country	Sum of Total Jobs Added	Average of HPC $K per New Employee	Count of Accomplishment Type
China	326	$206 K	126
France	90	$59 K	17
India	-		1
UK	1,781	$344 K	120
US	3,294	$97 K	304
Italy	-		9
Germany	26	$80 K	10
Canada	-		1
Slovenia	40	$3 K	2
Netherlands	16	$750 K	1
Australia	-		2
Japan	165	$608 K	79
South Africa	10		1
Grand Total	**5,748**	**$270 K**	**673**

IDC Analyze the Future

25

December 2016 Study Results: Innovation ROR (Return On Research)

IDC Analyze the Future

26

New Findings: The Mix Of Innovation Types (For The Innovation Projects)

The average cost of an innovation in the INNOVATION ROR projects was $12.7 million:

Q4: Primary Innovation / RO	Count of Accomplishment Type2	Average of HPC $M per Innovation	Count of Q8: Basic	Count of Q8: Applied
Better Products	95	$3.5 M	22	72
Cost Savings	16	$1.6 M	8	9
Created New Approach	233	$1.2 M	185	50
Discovered Something New	40	$3.0 M	27	13
Helped Society	42	$6.4 M	21	21
Scientific Breakthrough	55	$80.3 M	49	8
Support Research Programs	44	$36.7 M	19	26
Grand Total	**525**	**$12.7 M**	**331**	**199**

IDC Analyze the Future

27

New Findings: Innovations – Basic vs. Applied

Sector	Count of Q8: Basic	Count of Q8: Applied	Count of Q8: Basic/ Applied
Academic	265	80	341
Government	37	27	61
Industry	29	92	123
Grand Total	**331**	**199**	**525**

IDC Analyze the Future

28

New Findings: Innovations – By Country

Country	Count of Q8: Basic	Count of Q8: Applied	Count of Q8: Basic/ Applied
China	72	33	105
France	1	10	11
India	1		1
UK	53	20	73
US	161	96	254
Italy		6	6
Germany	1	2	3
Canada	1		1
Slovenia		1	1
Netherlands		1	1
Australia	1	1	2
Japan	40	28	66
South Africa		1	1
Grand Total	**331**	**199**	**525**

Note: The data set today isn't complete enough to make country-to-country comparisons.

IDC Analyze the Future

29

New Findings: The Innovation IMPORTANCE Index

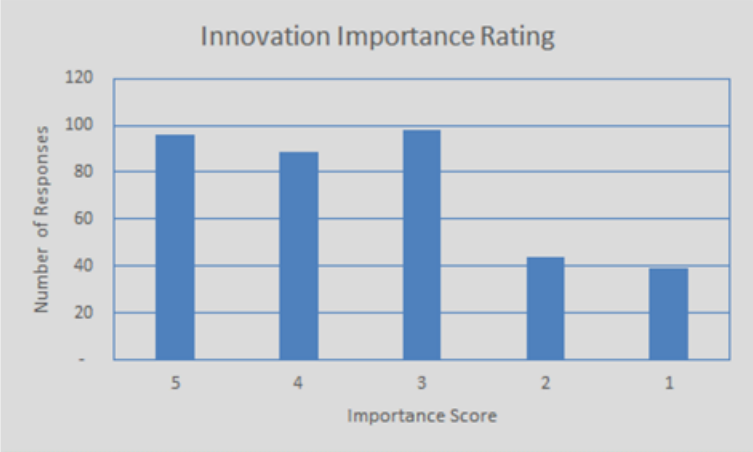

5. One of the top 2 to 3 innovations in the last decade
4. One of the top 5 innovations in the last decade
3. One of the top 10 innovations in the last decade
2. One of the top 25 innovations in the last decade
1. One of the top 50 innovations in the last decade

IDC Analyze the Future

30

New Findings: The Innovation IMPORTANCE Index

Average IMPORTANCE innovation rating = 3.43

Q9: New Innovation Index #1 Importance	Count of Accomplishment	Average of HPC $M per Innovation	Sum of Total Jobs Added
5	96	39.0	1,487
4	89	31.5	256
3	98	1.1	439
2	44	23.0	1,029
1	39	0.9	58
Grand Total	**366**	**19.1**	**3,268**

5. One of the top 2 to 3 innovations in the last decade
4. One of the top 5 innovations in the last decade
3. One of the top 10 innovations in the last decade
2. One of the top 25 innovations in the last decade
1. One of the top 50 innovations in the last decade

IDC Analyze the Future

31

New Findings: The Innovation IMPACT Index

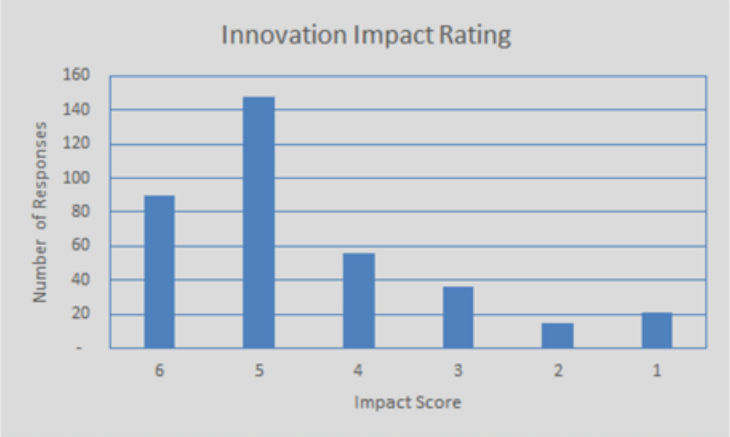

6. It is useful to over 50 organizations
5. An innovation that is useful to 10 to 49 organizations
4. An innovation that is useful to 6 to 10 organizations
3. An innovation useful to 2 to 5 organizations
2. An innovation only useful to 1 organization
1. An innovation that is recognized ONLY by experts in the field

IDC Analyze the Future

32

New Findings: The Innovation IMPACT Index

Average IMPACT innovation rating = 4.54

Q10: New Innovation Index #2 No. Org's Impacted	Count of Accomplishment Type	Average of HPC $M per Innovation	Sum of Total Jobs Added
6	90	9.8	1,260
5	148	18.2	893
4	56	17.9	38
3	36	1.3	49
2	15	0.8	9
1	21	81.3	1,020
Grand Total	**366**	**19.1**	**3,268**

6. It is useful to over 50 organizations
5. An innovation that is useful to 10 to 49 organizations
4. An innovation that is useful to 6 to 10 organizations
3. An innovation useful to 2 to 5 organizations
2. An innovation only useful to 1 organization
1. An innovation that is recognized ONLY by experts in the field

IDC Analyze the Future

33

Changes Made in 2016: New Innovation "Class" Scale

- Expanded the innovation indexes by adding a new innovation "Class" scale:

Class 1 innovations – One of the top 2-3 innovations in a field over the last ten years PLUS useful to over 10 organizations

Class 2 innovations -- One of the top 5 innovations in a field over the last ten years PLUS useful to over 10 organizations

Class 3 innovations – One of the top 5 innovations in a field over the last ten years PLUS useful to at least 5 organizations

Class 4 innovations – One of the top 10 innovations in a field over the last ten years PLUS useful to at least 5 organizations

Class 5 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at over 10 organizations

Class 6 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at least 2 organizations

Class 7 innovations – One of the top 50 innovations in a field over the last ten years PLUS useful to at least 2 organizations

Class 8 innovations – The rest of the innovations in the study

IDC Analyze the Future

34

New Findings: The New INNOVATION CLASS Index

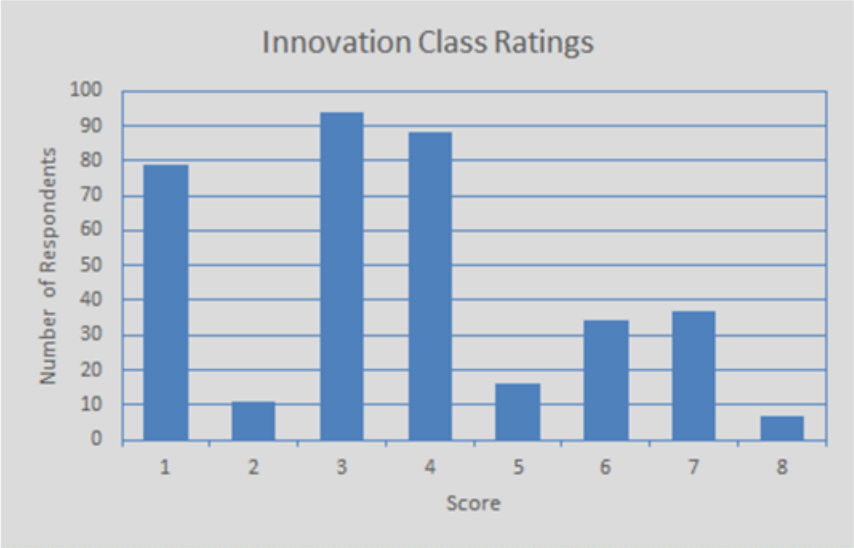

IDC
Analyze the Future

35

New Findings: The INNOVATION CLASS Index

Average innovation Class rating = 3.64

Innovation Class	Count of Accomplishment Type	Average of HPC $M per Innovation	Sum of Total Jobs Added
1	79	33.2	1,449
2	11	125.9	28
3	94	30.6	266
4	88	1.1	434
5	16	1.7	14
6	34	4.0	15
7	37	21.7	1,063
8	7	0.8	-
Grand Total	**366**	**19.1**	**3,268**

IDC Analyze the Future

36

New 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 will be writing up a number of them for broader dissemination

IDC Analyze the Future

37

Example:
Use By Others

IDC
Analyze the Future

38

Example of Results Used In Reports/Studies: Blue Waters Assessment Report

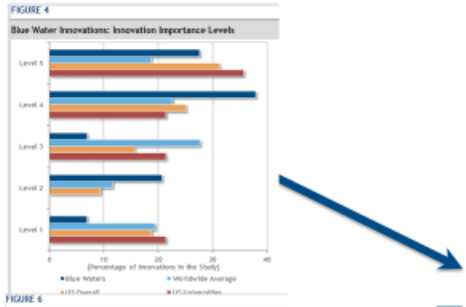

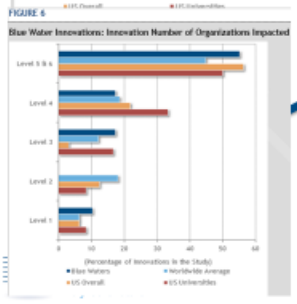

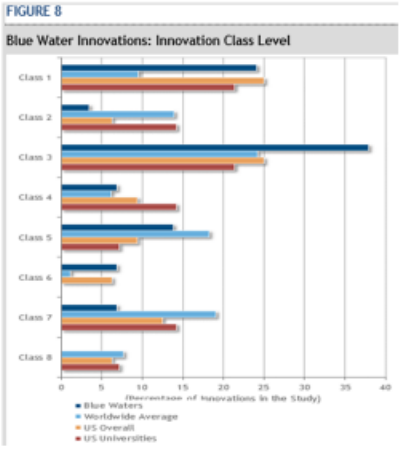

39

Example of Results Used In Reports/Studies: Riken K-Computer Assessment Report

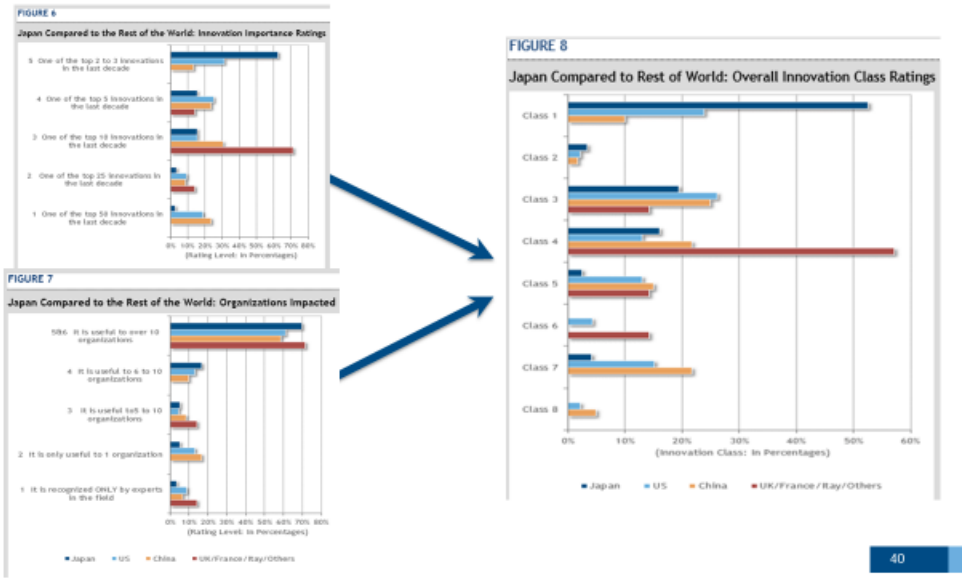

Ideas For Future Research

IDC Analyze the Future

42

Future Research Ideas

1. **Expand the models to be used as a predictive tool**
   - Forecast the potential value of a new supercomputer (e.g. exascale systems)
   - In both scientific ROR terms and financial ROI terms
   - We tested this on the Post-K computer at Riken

IDC Analyze the Future

43

Future Research Ideas

2. **Expand the models to measure the value of an existing HPC system**
   - To show the value of a large supercomputer
   - And perhaps to an existing HPC center
   - We tested this in the Riken study and have a basic approach

IDC Analyze the Future

44

Dissemination Program

- **New results, along with the data file are published at: www.hpcuserforum.com/ROI**
- **The pilot report and excel models are also posted at: www.hpcuserforum.com/ROI**
- **IDC sends the new results to ~6,500 people in the broader HPC community**
- **IDC will brief the HPC community at SC16**
  - On Tuesday during the IDC breakfast briefing
  - And at other meetings
- **Does DOE have additional briefings that IDC should conduct?**

IDC Analyze the Future

45

Web Page: www.hpcuserforum.com/ROI

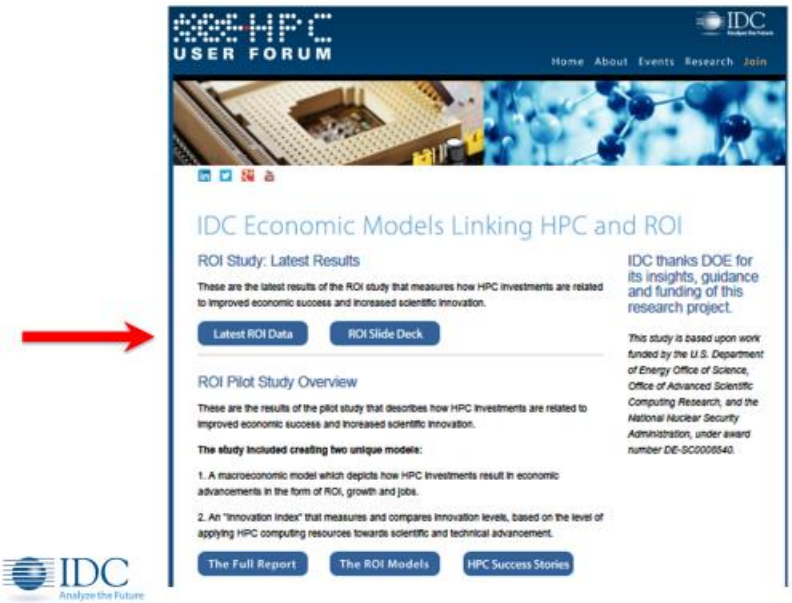

Next Steps: Short Term

Keep pushing to obtain more data points/ROI cases

- With more from Germany, China, France & Japan
- Drive to get better and more complete job creation data

Distribute the results more broadly

- Publish the new results
  - We have started the conversations with key publishers – the goal is publication of the findings in a juried business journal or similar (such as Harvard Business Review)
- More press releases and reach out to business focused groups
  - Now that we have a much stronger data set
- 4 HPC User Forums planned in 2017, plus ISC17 & SC17
- Looking at running a series of ROI success stories with HPCWire

IDC Analyze the Future

47

Questions?

Please email:
hpc@idc.com

Or check out:
www.hpcuserforum.com

IDC Analyze the Future

EXHIBITS:

EXHIBIT 3:
HPC ROI SUCCESS STORIES (2016)

Special Study

Real-World Examples of Supercomputers Used For Economic and Societal Benefits: A Prelude to What the Exascale Era Can Provide

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

IDC OPINION

Since its introduction in the 1960s, high performance computing (also called supercomputing) has made enormous contribution to America's scientific, engineering and industrial competitiveness, as well as to homeland security and other government missions.

Supercomputers have played crucial roles in U.S. government agencies and departments. But that's just part of the story. Supercomputers have already made cars and planes much safer, more fuel efficient and environmentally friendly. They are crucial aids in discovering and extracting new sources of oil and gas, and for developing alternative energy sources. They have enabled the weather community to create more accurate predictions of severe storms that can devastate lives and property. They are heavily relied by industries ranging from financial services to medicine and health care, entertainment, consumer products, and more recently by Internet companies.

In short, high performance computing has become indispensable for maintaining America's national security and economic competitiveness. An IDC study showed that 97% of companies that had adopted supercomputing say that could no longer compete or survive without it.

That's why other nations and global regions including China, Europe, Japan and Russia, to name a few, are racing ahead and have created national programs that are investing large sums of money to develop exascale supercomputers for use later in this decade or early in the next decade.

The new return-on-investment (ROI) examples in this report underscore the benefits of providing American scientific and industrial researchers with access to the latest generations of powerful supercomputers housed at national centers and laboratories.

IDC believes that the United States needs to maintain and advance its global standing in high performance computing, by providing adequate funding and commitment to achieve exascale capability. Without these investments, the competitiveness of key U.S. industries will be substantially reduced, and it will compromise America's national security capabilities.

REAL-WORLD EXAMPLES OF SUPERCOMPUTERS USED FOR ECONOMIC AND SOCIETAL BENEFITS

EXECUTIVE SUMMARY

This report provides examples of how U.S. private- and public-sector organizations have been using supercomputers to achieve breakthroughs of major scientific or economic importance. These achievements, many of which were accomplished through access to very powerful supercomputers and HPC experts at U.S. national laboratories, already have saved American companies, many millions of dollars and have the potential to save many billions of dollars.

Over the last two decades, DOE has paved the way with breakthrough HPC architectures and software programming models. Other nation states are now trying to take the lead in advanced HPC technologies. If that happens, the next wave of scientific breakthroughs and industrial innovations in many industries could come from outside the US and US will miss out on the next wave of economic expansion.

Without these achievements and the supercomputer access they depend on, these American companies and research organizations would be dramatically less competitive than they are today. They would almost surely be able to provide fewer jobs and make smaller contributions to the nation's economy, now and in the future. Examples of the impact of supercomputers in US industry include:

- GE used a supercomputer to reveal a new aspect of turbine behavior that is already providing GE with a competitive advantage in fuel efficiency. Every 1% reduction in fuel consumption saves users of these products $2 billion/year.
- BMI utilized supercomputers and computational models to design components that could save 1.5 billion gallons of fuel and $5 billion in fuel costs per year.
- Supercomputers were used to generate more oil and gas from US reserves. This promises billions of dollars per year in savings and reduction of US dependence on foreign energy.
- Automotive and engine manufacturers are using high performance computing to develop next-generation engines that use less fuel. These fuel savings are estimated at more than $1 billion per year.

An Example of Reducing Fuel Consumption of Our Trucking Fleet

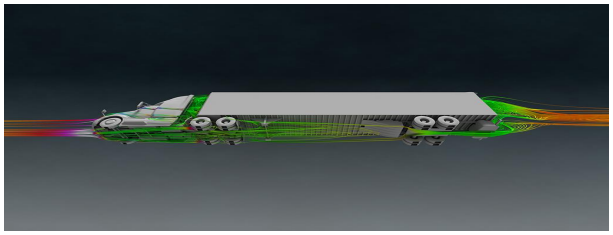

Source: BMI, Oak Ridge National Laboratory, 2014

Supercomputers have played crucial roles in U.S. government agencies and departments. In addition to their use in classified government, including a recognized role in helping to end the Cold War, supercomputers enable NOAA's national weather and severe storm predictions that have help save lives and billions of dollars in property. Supercomputer make possible the development of advanced military aircraft and spacecraft, and are essential for planning, executing and improving NASA missions. Supercomputers were used to identify the causes of the NASA "Challenger" tragedy and the collapse of the World Trade Center towers on 9/11. Supercomputers are being evaluated today to transform the ability of CMS and other agencies to catch fraud. USPS is now using supercomputers to alleviate costly errors and fraud. Examples include:

- Seismic simulations are used to generate advanced hazard maps and this research team developed software that significantly advances this map-making ability. The potential saving is in billions of dollars in preventing property loss/damage, along with protecting lives.
- Researchers from the Centers for Disease Control (CDC) created a far more detailed model of the hepatitis C virus, a major cause of liver disease. Annual health care costs associated with this virus are $9 billion in the U.S. alone.

- For the first time researchers developed a computer model that comprehensively simulates the human heart down to the cell level. This innovation has significant potential for saving health care costs by reducing heart disease and improving heart health. This research has strong potential for helping to reduce coronary heart disease, which costs the United States over $100 billion each year.

Using HPC to Improve Cancer Treatments

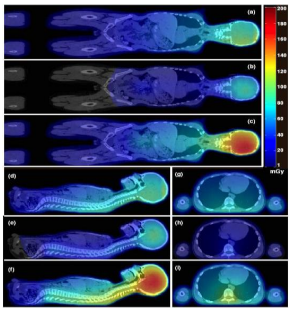

Source: Mary Bird Perkins Cancer Research Center , Louisiana State University, 2014

But that's just part of the story. Without supercomputers, detecting today's sophisticated cyber security breaches, insider threats and electronic fraud would be impractical. In short, high performance computing has become indispensable for both maintaining America's national security and economic competitiveness.

That's why other nations and global regions including China, Europe, Japan and Russia, are racing ahead and have created national programs that are investing large sums of money to develop exascale supercomputers. What this global race is really about is supremacy in supercomputing and in all the disciplines and markets that depend heavily on this game-changing technology.

HPC CASE STUDY SUMMARIES

These shining examples illustrate the importance of continuing to provide this diverse community of HPC users with access to first class supercomputers whose capabilities are on a par with those of other leading nations and global regions, including exascale supercomputers as soon as those begin to appear around the world. Anyone who doubts the importance of supercomputers for scientific and economic competitiveness should consider the examples in this report:

- The Centers for Disease Control (CDC) used a supercomputer to advance understanding of the hepatitis C virus, a major cause of liver disease. This paves the way for researchers to discover new therapies for combating the virus. Annual health care costs associated with this virus are estimated to be $9 billion in the U.S. alone.
- The Mary Bird Perkins Cancer Center (Baton Rouge, LA) made important advances that could lower the incidence of second cancers caused by radiation in children receiving radiation therapy. The collaborators saved more than $12 million by using high performance computing. The researchers estimate this achievement has accelerated radiation toxicity research by more than a decade.
- Researchers at the Salk Institute (San Diego) are using supercomputers at the nearby NSF-funded San Diego Supercomputer Center to investigate how the synapses of the brain work. Their research has the potential to help people suffering from mental disorders such as Alzheimer's, schizophrenia and manic depressive disorders.
- About 600,000 people die of heart disease in the United States every year–that's 1 in every 4 deaths. Scientists from DOE's Lawrence Livermore National Laboratory (LLNL) modeled the human heart in much greater detail than before, using one of the world's most powerful supercomputers. This advance lays the foundation for progress in preventing and treating heart disease.
- Turbines are literally responsible for keeping the lights on, since most of the world's electricity is generated by turbines. And jets couldn't stay aloft without them. Recently, GE used a supercomputer to reveal a new aspect of turbine behavior that is already providing GE with a competitive advantage.
- Thanks to new research on a supercomputer Ramgen will begin testing a 13,000-horsepower CO_2 compressor this year. This compressor is projected to reduce the capital costs of CO_2 compression by 50 percent and produce a minimum of 25 percent savings in operating costs. Applying these cost savings to a new 400-megawatt clean coal plant would result in capital cost savings of approximately $22 million and an annual operating cost savings of approximately $5 million.
- A BMI Corp. SmartTruck technology developed on a supercomputer could save 1.5 billion gallons of diesel fuel and $5 billion in fuel costs per year. This technology is now in use.
- Boeing Corporation saved many millions of dollars by using supercomputers. Boeing physically tested 77 prototype wing designs for the 767 aircraft, but for the new Boeing 787 Dreamliner only 11 wing designs had to be physically tested (a 7 fold reduction in the needed amount of prototyping), mainly because over 800,000 hours of computer simulations on supercomputers had drastically reduced the amount of needed physical prototyping.
- Seismic simulations were used to generate a hazard map, with the potential saving of many lives and properties. Upon completion of California state-wide seismic hazard map savings will be many billions dollars. These National Seismic Hazard Maps will help set building codes and insurance rates, as well as provide short-term forecast of the frequency of damaging earthquakes in California over a specified time span.
- Supercomputing has led to significant improvements in the Navy's effort to improve tropical cyclone intensity prediction.

- Manufacturers will have valuable new information that will ultimately help them design better engines more quickly and at a lower cost. These new models will allow researchers to stretch uses of the models beyond what can currently be done experimentally, testing out theoretical innovations such as low temperature combustion
- RENCI used supercomputing to compute new data in order to update coastal floodplain maps for North Carolina. These maps are required by FEMA for local communities and municipalities to be eligible for flood insurance coverage under the National Flood Insurance Program (NFIP). The modeling program was so successful in North Carolina that it was used by FEMA and the US Army Corps of Engineers to provide data to update coastal floodplain maps from the Virginia-North Carolina state line to the Delaware Bay.
- The low-cost, post-processing system powered by supercomputing, reduces the need for observation wells and has demonstrated commercial success in oil and gas recovery, carbon capture and sequestration and geothermal energy. The system is already in use to track injected carbon dioxide in several energy exploration projects
- Researchers are using simulations to decrease the materials scrapped during the continuous casting process. Decreasing the material scrapped due to defects such as cracks, even by a small percentage, results in a large net savings to steel manufacturers and customers. Based on the roughly 100 million tons of steel produced each year in the U.S. and approximately $400 per ton net cost of scrapping, a one percent reduction in yield loss would save about $400 million per year.

TABLE OF CONTENTS

LIST OF FIGURES

SUPERCOMPUTING SUCCESS STORIES

A) IMPROVING HEALTH AND QUALITY OF LIFE

Finding Cures for Hepatitis C (Cornell University and CDC)

What Was Accomplished

Researchers from the Centers for Disease Control (CDC) created a far more detailed model of the hepatitis C virus, a major cause of liver disease. Annual health care costs associated with this virus are estimated to be $9 billion in the U.S. alone. By scaling the virus model from a desktop computer in Atlanta to an HPC cluster at Cornell, researchers improved our understanding of the virus's networks of amino acids, paving the way for new therapies for combating this devastating virus.

Savings/Potential Savings: $9 billion per year; $360 billion over 40-year patient lifetimes.

FIGURE 1

Visualization of a Hepatitis C Virus Network (Campo et al).

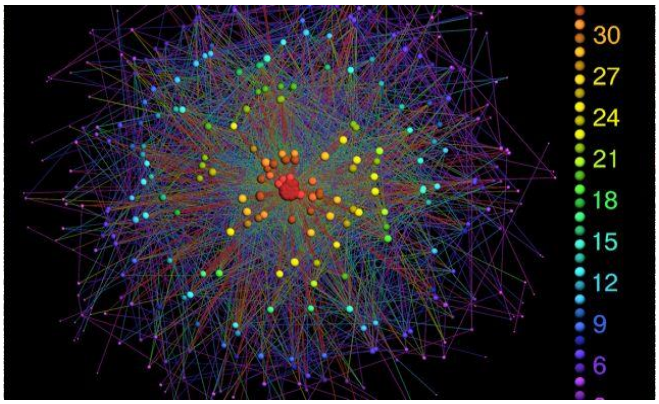

Source: Cornell University, 2014/CDC Campo et al.

Organization

The Centers for Disease Control and the Cornell University Center for Advanced Computing.

Supporting Organizations

The HPC cluster was funded by the National Science Foundation and was executed through industrial partnerships with MathWorks (Natick, MA) and Dell (Round Rock, TX).

Industry

Education, Scientific Research, Engineering

ROI or Innovation Description

With the cost per liver transplant in the range of $280,000 for the first year alone, liver transplantations due to hepatitis C cost nearly $300 million per year in the U.S. The average lifetime cost for hepatitis C without a liver transplant has been estimated at $100,000 per person. Assuming that 80% of the 4.5 million Americans believed to be infected by this virus develop chronic liver disease, the annual health care costs for the U.S. population affected with chronic hepatitis C total about $9 billion and the lifetime costs for this group of 3.6 million persons, assuming average survival of 40 years, amount to a staggering $360 billion in today's dollars

Type of Innovation

This basic Innovation was rated among the top 50 innovations in the last decade.

Nature of Research

Through faster computations (more than 175 times speed-up), a better understanding of networks of coordinated amino-acid variation has opened the door for the discovery of new therapeutic targets for the hepatitis C virus (HCV). Over 500,000 jobs ran on the Cornell system over two years, generating new scientific insights and publications in condensed matter physics, gravitational wave detection, biomedical imaging, orthopedics, neuroscience, and optics.

Broader Impact

According to the world health organization (WHO), about 150 million people are chronically infected by HCV, with 3 to 4 million new infections every year. 75% to 85% of newly infected individuals develop chronic disease with long-term complications, including liver cirrhosis and liver cancer. Since to date no vaccine is available and more than 350,000 patients die yearly from hepatitis C-related diseases, the virus poses a major health concern. Supercomputer simulations can significantly accelerate research on viruses like hepatitis C and help develop effective cures.

Lowering the Cost of Advanced Cancer Care (Mary Bird Perkins Cancer Center and Louisiana State University)

What Was Accomplished

In this research, HPC-driven cancer research enabled better outcomes for long-term survivors of cancer. Advances in cancer detection and treatment have led to large improvements in survival rates. But with increased survival comes an increased need to minimize long-term treatment-related negative effects. In particular, children receiving radiation therapy are more susceptible to radiation-caused secondary cancers (carcinogenesis) later on. Researchers have conducted clinical trials with the help of supercomputer simulations, and these trials are helping to improve success rates for long-term, advanced cancer care. The simulation driven research has generated valuable data that is being used to inform clinical and health policy decisions

Savings/Potential Savings: $12 million in research costs and 6 months in constructing a new $125 million therapy center.

FIGURE 2

Doses For CSI And Boost Treatment Superimposed On A Patient's CT Image

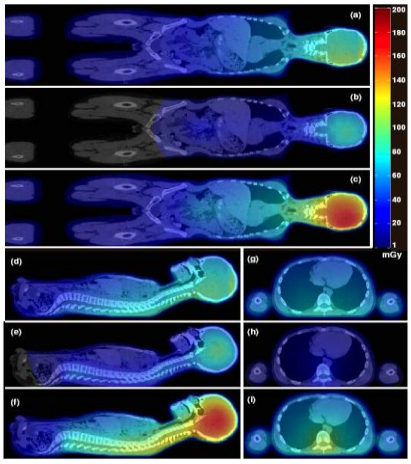

Source: Mary Bird Perkins Cancer Research Center, Louisiana State University, 2014

Organization

The Mary Bird Perkins Cancer Research Center and Louisiana State University. The Mary Bird Perkins Cancer Center, founded in 1971, is the region's largest provider of radiotherapy, community-owned (not-for profit) organization based in Baton Rouge, LA.

Supporting Organizations

The National Institutes of Health (NIH) and the Department of Defense.

Industry

Healthcare, Education, Scientific Research,

ROI or Innovation Description

More than $12 million was saved in research costs by using HPC simulation. Even more important, this simulation-driven research has accelerated progress in the field of advanced cancer care. In addition, the ability to simulate therapy equipment saved more than 6 months on the construction of a new $125 million proton therapy facility and has accelerated radiation toxicity research by more than a decade.

Type of ROI

A conservative estimate of future monetary gains as a result of removing obstacles to founding new proton theory centers is on the order of $300 million. The research has also enabled estimated cost savings of $15 million.

Nature of Research

Predictions of exposure to charged particle radiation are commonly performed for patients receiving radiotherapy. Researchers studied the physical interactions and bioeffect-modeling approaches to predict radiation-produced toxicity in patients. Computer modeling approaches were used in several aspects of the research, to predict radiation dosages and radiation-caused risk for developing second cancers. These approaches applied Monte Carlo methods and supercomputing techniques.

The researchers simulated a pediatric brain tumor treatment to calculate stray radiation. The largest proportions of this risk were assumed by the skin and then the thyroid. The research revealed that using proton beams reduced the total lifetime risk of second cancer from over 30% after photon therapy to about 5% for proton therapy.

The researchers estimated the excess lifetime risk of second-cancer fatality from stray radiation for pediatric male patients who receive this treatment at 4% to 5%, depending on the nature of proton treatment. They note that while this risk is small compared with the benefits of the radiotherapy, it is not negligible. As such, it is important to continue attempts to reduce stray radiation exposure as much as possible. Ultimately, the team aims to generate a high-quality evidence to support clinical decision making as to whether a patient receives proton or photon therapy

Broader Impact

This research was fundamental in demonstrating that computer-based clinical trials can be a cheaper, faster way to generate scientific breakthroughs. In this research, simulation-driven cancer research enabled better outcomes for long-term survivors of cancer. Additionally this research enabled the participation of students and post docs in mentored research training.

Finding Cures for Alzheimer's, Schizophrenia and Manic Depressive Disorders (Salk)

What Was Accomplished

This research has the potential to greatly help people suffering from mental disorders such as Alzheimer's, schizophrenia and manic depressive disorders.

Researchers at the Salk Institute (San Diego, CA) are using supercomputers at the nearby NSF-funded San Diego Supercomputer Center to investigate how the synapses of the brain work. In addition, the use of supercomputers is helping to change the very nature of biology – from a science that has relied primarily on observation to a science that relies on high performance computing to achieve previously impossible in-depth quantitative results.

Savings/Potential Savings: This research could ultimately help reduce the overwhelming cost for treatment and long-term care of brain related disorders.

FIGURE 3

GENISIS of a Purkinje Cell

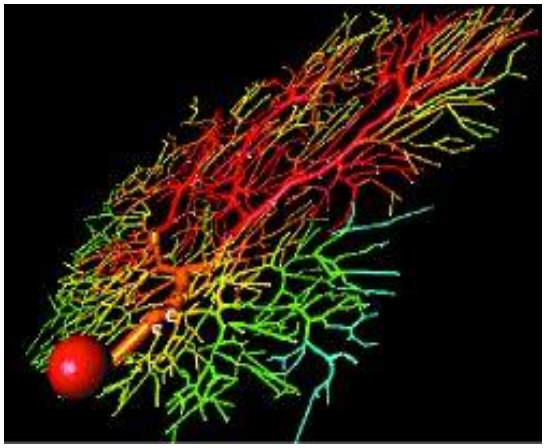

Source: Salk Institute, 2013

Organization

The Salk Institute

Industry

Healthcare, Scientific Research,

ROI or Innovation Description

Significant progress in the modeling of chemical transmission across nerve synapses, which involves an array of complex electrochemical processes

Nature of Research

The research focuses on the events that occur when a neuron, or nerve cell, sends its chemical message across a synapse to influence another neuron. This is a key way in which messages are communicated along neural pathways. This involves modeling the release of neurotransmitter, its diffusion across the synapse, and its binding to receptors to generate currents in the receiving neuron.

The improvements resulted in a 150-fold speedup of simulations, and the speedups enabled an increase in simulation complexity. Researchers can develop 3D

representations of cellular structures, with resolution down to the electron microscope level.

Broader Impact

The supercomputing-simulation driven approach is one of the key ways in which Salk Institute is building computational bridges between brain levels from the biophysical properties of synapses and function of neural systems. This research could ultimately help reduce the overwhelming cost for treatment and long-term care of brain related disorders. Modeling driven precision in circuit information will help researchers understand the scale and scope of problems while enabling them to test and develop targeted therapies that are ultimately more effective.

Developing Lower Cost Heart Disease Treatments (LLNL)

What Was Accomplished

For the first time researchers developed a computer model that comprehensively simulates the human heart down to the cell level. This innovation has significant potential for saving health care costs by digitally screening drugs and drug candidates for cardiotoxicity, reducing mortality from heart disease and improving the effectiveness of heart therapies. The researchers were able to simulate a heartbeat per second, compared with prior modeling that took at much as 45 minutes to simulate a single heartbeat.

Savings/Potential Savings: This research has strong potential for improving the lives of patients with coronary heart disease, which costs the United States over $100 billion each year.

This research has strong potential for helping to reduce coronary heart disease, which costs the United States over $100 billion each year.

FIGURE 4

Computer Visualization Of the CARDOID Heart Model

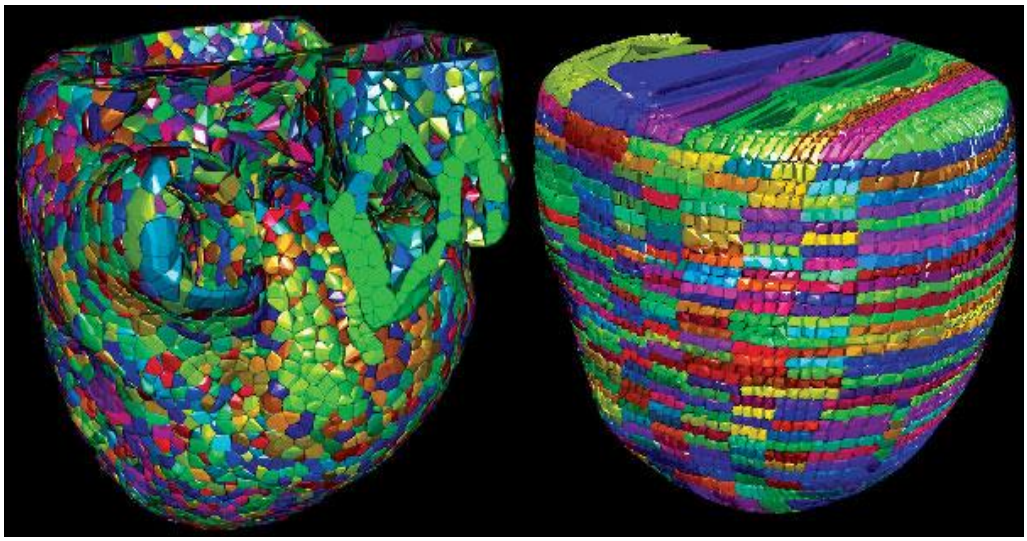

Source: LLNL, 2014

Organizations

Lawrence Livermore National Laboratory (LLNL) and IBM Research

Industry

National security and information technology

ROI or Innovation Description

The research partners discovered new information about the functioning of the human heart that has significant potential for alleviating the health and monetary costs of heart disease.

Type of ROI

Technical innovations with strong potential for health and monetary savings in human societies.

Nature of Research

Computational models of cardiac cells have been in developed since the 1960s. However developing a comprehensive model of the human heart is complex, because the heart is composed of cells that are discrete and individual and yet are connected as part of an excitable system. Because each cell in the heart muscle needs to be simulated at the same time, modeling all this activity in detail would take too long even on the largest supercomputers. Starting with an existing cell model, LLNL and IBM researchers have for the first time demonstrated a multi scale model that comprehensively simulates the human heart in real time at resolution approaching the cellular level..

To use the laboratory's powerful supercomputer more efficiently, the researchers created new ways to divide the heart tissue into small pieces of equal work units that could be evenly distributed across the supercomputer's 1,572,864 computing elements. The researchers were able to model the human heart in unprecedented detail and at unprecedented speed by developing novel software and methods for exploiting one of the world's most powerful supercomputers located the Lawrence Livermore National Laboratory

Broader Impact

This innovation has significant potential for saving health care costs by reducing heart disease and improving heart health. According to the Centers for Disease Control and Prevention:

- About 600,000 people die of heart disease in the United States every year–that's 1 in every 4 deaths.
- Heart disease is the leading cause of death for both men and women. More than half of the deaths due to heart disease in 2009 were in men.
- Coronary heart disease is the most common type of heart disease, killing more than 385,000 people annually.
- Every year about 715,000 Americans have a heart attack. Of these, 525,000 are a first heart attack and 190,000 happen in people who have already had a heart attack.
- Coronary heart disease costs the United States $108.9 billion each year. This total includes the cost of health care services, medications, and lost productivity.

B) REDUCING FUEL COST AND CO_2 EMISSIONS

Improving Fuel Savings in Jet Engines (General Electric)

What Was Accomplished

Each year, about $200 billion worth of fuel is consumed globally on GE's gas turbine products, both aircraft engines and land-based gas turbines used for the production of electricity. Every 1% reduction in fuel consumption therefore saves the users of these products $2 billion/year.

Savings/Potential Savings: Multiple billions of dollars per year.

FIGURE 5

A Visualization of A 4 Stage Low-Pressure Turbine

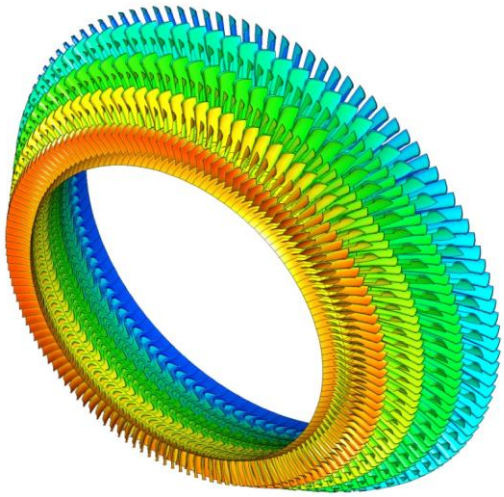

Source: GE Global Research, 2013

Organization

Scientists and Researchers at the GE Global Research Center collaborated with experts at Oak Ridge National Lab to conducted the jet engine simulations

Supporting Organizations

Supercomputers hosted at the Department of Energy's Oak Ridge National Laboratory were used in this research.

Industry

Manufacturing, Aerospace, Automotive, Consumer Products

ROI or Innovation Description

Scientific advance useful for practical engineering of aircraft

Type of ROI

The models employed would not have been practical or in many cases even possible without the supercomputer.

Type of Innovation

GE observed a phenomenon spanning multiple blade rows of the engine that was not observable before. This was rated as one of the top 10 innovations in the last decade.

Nature of Research

For more than a century General Electric (GE) has been building turbomachines and is currently a major producer of turbines for the electric power generation and aircraft engine industries. More recently GE has begun to utilize supercomputers and computational models to accelerate its R&D in turbomachinery.

GE researchers use supercomputers to model unsteady flows from industrial turbines that are deployed in jet engines, power stations and beyond. Computational models enable designers to make advanced and fine-grained adjustments to the turbomachinery while generating greater operating and fuel efficiencies. GE researchers were able to investigate unsteady flows and design solutions that met several key design criteria. The researchers were also able to simulate the turbomachinery in 3D and 4D (3D + a time dimension) and were able to study turbomachinery design impacts over time.

Developing turbomachinery is an extremely complex and globally competitive business with multiple top-tier corporations competing with GE to sell jet engines and gas turbines in different application domains. Consequently any company that can achieve even 1% improvement in efficiency can potentially cause market disruption, as the resultant efficiencies would overtime add up to enormous cost savings to customers, thus providing market advantage.

Broader Impact

The studies performed by GE are targeted towards improving the aerodynamic efficiency of jet engines. Improvements to jet engines can be leveraged to improve the efficiency of land based gas turbines which GE also manufactures. Each year, about $200 billion worth of fuel is consumed globally on GE's gas turbine products, both aircraft engines and land-based gas turbines used for the production of electricity. Every 1% reduction in fuel consumption therefore saves the users of these products $2 billion/year.

Saving Fuel Cost on Long Haul Trucks (BMI SmartTruck)

What Was Accomplished

BMI utilized supercomputers and computational models to design components that could save 1.5 billion gallons of diesel fuel and $5 billion in fuel costs per year.

The big rigs barreling down America's highways day and night are essential to the country's economy – they carry 75% of all U.S. freight and supply 80% of its communities with 100% of their consumables. Thus far, long haul trucks have often

averaged 6 MPG or less resulting in very high fuel costs and considerable impact to the environment with emissions of more than 423 million lbs of CO_2. The BMI engineers combined their aerospace and race care design expertise with the power of supercomputers to create a new, long haul "SmartTruck" to increase fuel efficiency and reduce carbon emissions

Savings/Potential Savings: 1.5 billion gallons of diesel fuel and $5 billion in fuel costs per year.

FIGURE 6

Supercomputing Simulations Of The SmartTruck UnderTray System

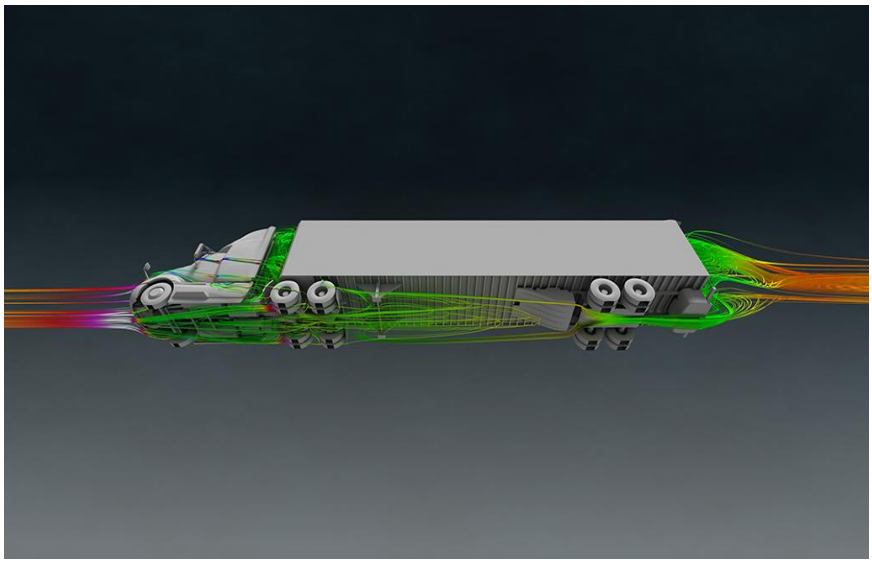

Source: BMI, Oak Ridge National Laboratory, 2013

Organization

BMI (Greensville, SC)

Industry

Manufacturing, Aerospace, Automotive, Consumer Products

ROI or Innovation Description

Since 2011 SmartTruck Systems of Greenville, S.C., has sold more than 25,000 UnderTray Systems to trucking fleets in North America. The award-winning, EPA-certified UnderTray Systems are trailer add-on components that improve highway fuel efficiency by more than 10 percent and save an estimated $5,000 annually in fuel costs per truck. If all of the 1.3 million Class 8 trucks in the country were configured with just the minimum package of new components, the U.S. could annually save almost 1.5 billion gallons of diesel fuel, reduce CO_2 by 16.5M tons and save more than $6B in fuel costs.

Nature of Research

Using BMI's SmartTruck UnderTray System to improve the aerodynamics of 18-wheeler (Class 8) long-haul trucks, the typical big rig can achieve fuel savings of between 7 and 12 %. Using supercomputers BMI simulated UnderTray System, including components such as aerodynamic wheel fairings and special airflow directing mechanisms. Computational models run on supercomputers were also used to design and develop the rear diffuser, the biggest component of the UnderTray System.

Running design simulation models on the supercomputers allowed BMI to jump from concept to a design that could be turned over to a manufacturer in 18 months instead of the more usual 3½ years. BMI's goal was to design add-on parts for existing trucks and trailers to make them more aerodynamic. By reducing drag BMI boosted fuel efficiency and cut the amount of carbon that's being dumped into the environment.

Extensive testing was used to validate BMI's CFD simulations, demonstrating gains of 6.8%. BMI is working with NASA to improve additional techniques and newly developed optimization techniques,

BMI also earned one of the key industry awards and citation as top 20 products of the year for products made possible simulation work performed on supercomputers.

Broader Impact

BMI currently employs more than 25 workers throughout the U.S. and works with materials suppliers nationwide. UnderTray Systems have a quick payback period and a favorable return on investment, so fleets have strong incentive to upgrade trailers with the components, which are manufactured entirely in the USA and are 100 percent recyclable.

BMI also plan to utilize similar modeling and supercomputer driven design methodology to further improve the aerodynamics of both existing tractors and trailers.

Making America Energy Independent By Improving Oil and Gas Exploration (LLNL, Industry Partners)

What Was Accomplished

Supercomputers were used to develop an advanced reservoir-monitoring technology in order to generate more oil and gas from US reserves. This will help reduce the US dependence on foreign sources of energy.

Using the supercomputing power of its HPC machines, researchers and engineers in collaboration with LLNL developed a technique that integrates separate measurement data to predict subsurface fluid distribution, temperature and pressure. The system obviates the need for numerous observation wells, and the low-cost system works well for oil and gas recovery; carbon capture and sequestration; and geothermal energy.

Savings/Potential Savings: Billions of dollars per year in savings and reduction of US dependence on foreign energy.

FIGURE 7

Extracting More Oil from Reserves in the Gulf of Mexico

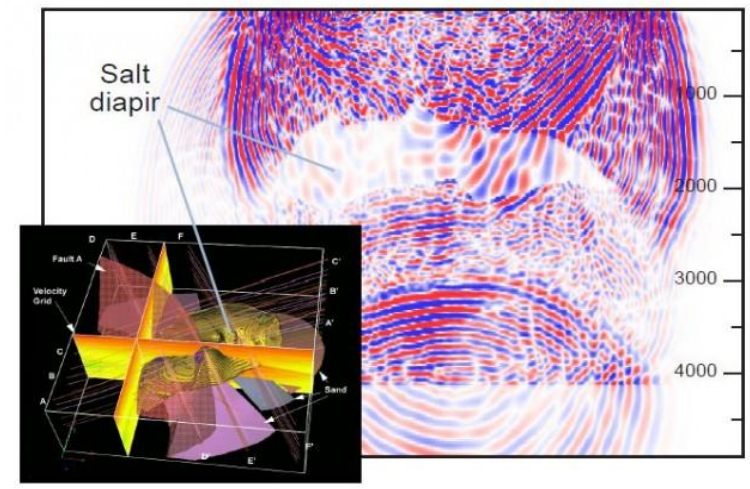

Source: LLNL, 2014

Organization

LLNL and various US industrial partners.

Industry

Oil and Gas industry

ROI or Innovation Description

The system developed by LLNL and its partners is already in use to track injected carbon dioxide in projects in multiple locations. Through modeling and simulation, private sector participants have improved well recovery and reduced failure risk.

Type of ROI

Process improvement resulting in multi fold speed up and increased safety

Nature of Research

Oil and gas companies traditionally have been some of the most aggressive and advanced users of Supercomputing. Most of the O&G companies use supercomputing for geophysical processing of 3D and 4D seismic imaging volumes for exploration and production. O&G companies also harness supercomputing for reservoir management, development of new compounds and chemicals, and planning drilling efforts.

The lab has partnered with oil and gas companies to develop tools to improve recovery with low cost and low environmental impact. Central to these efforts is a tool that directly assesses the probability of key reservoir characterization and sub-surface plumes and fluid distributions. The low-cost, post-processing system powered by supercomputing, reduces the need for observation wells and has demonstrated commercial success in

oil and gas recovery, carbon capture and sequestration and geothermal energy. The system is already in use to track injected carbon dioxide in multiple projects.

Broader Impact

The need for greater energy independence, a sustainable environment, and reduced conventional production has driven innovation and enterprise in the global oil and gas industry. Through advanced modeling and simulation, private sector participants have improved well recovery, reduced operating costs, and reduced failure risk. These early results suggest higher promise for HPC in improving domestic energy production with reduced environmental consequence.

Developing Innovation Performance Engines for Alternative Fuels (ANL)

What Was Accomplished

Automotive and engine manufacturers are using high performance computing to develop of next-generation engines that use less fuel. A key competitive differentiator in designing engines is development of fuel injector technology that can utilize a new generation of bio fuels. Through the use of large scale supercomputers and computational models engineers are developing more robust fuel spray and combustion models for predictive engine simulations.

Savings/Potential Savings: Estimated at more than $1 billion per year.

FIGURE 8

State of the Art Fuel Spray Modeling

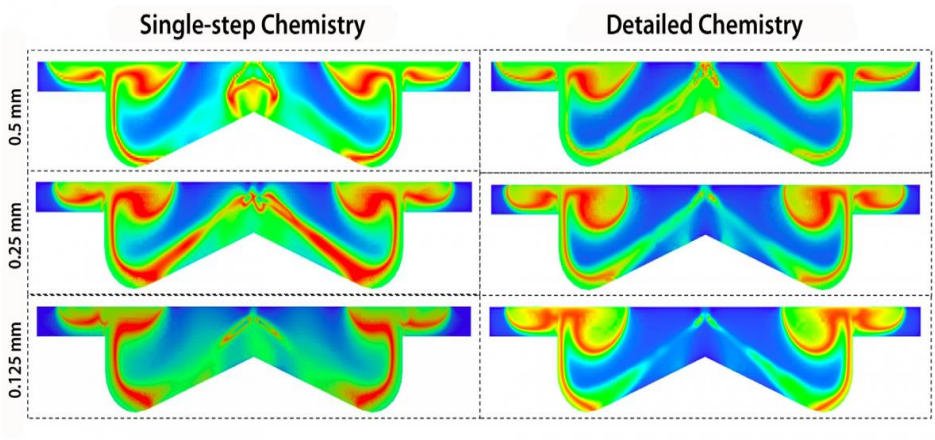

Source: ANL, 2014

Organization

Argonne National Laboratory

Industry

Automotive Engineering, Manufacturing

ROI or Innovation Description

Using larger and detailed simulations at a faster rate, manufacturers will have valuable new information that will ultimately help them design better engines more quickly and at a lower cost. These new models will allow researchers to stretch uses of the models beyond what can currently be done experimentally, testing out theoretical innovations such as low temperature combustion

Type of ROI

Innovative Modeling Process

Nature of Research

High-fidelity engine simulations allow researchers to develop detailed and intricate models of a virtual combustion chamber to better understand how injected fuel, air and combustion products intermingle inside an engine. The fuel injector, for instance, atomizes liquid fuel into fuel vapor, changes that require complicated fuel spray, two-phase flow, and heat transfer modeling. Turbulence produced by the chaotic motion inside engine cylinders due to the fuel spray necessitates the inclusion of these high-fidelity computational fluid dynamics approaches to describe the spray of the fuel into the combustion chamber. This complicated process involving fluid dynamics, chemistry and combustion determines engine performance and emissions characteristics, but it would be impossible to observe and understand without the aid of computer models.

Advanced large scale supercomputing resources allows researchers and engineers to simulate the activity inside an engine at a much finer scale. The space simulated in a computational fluid dynamics models is divided into cells and the smaller the cell, the more accurate the model. Industrial models typically use millimeter resolution, but the scaled-up high resolution models created at Argonne can simulate at a micrometer scale -- a thousand times more precise.

Broader Impact

In addition to improving the scale at which physics and chemistry can be simulated and observed, the high resolution engine models will also be valuable testing ground for the future of alternative fuel. With different types of fuel options including biofuels based on corn, soy, and other new sources, their performance and emissions in different engine designs can be simulated virtually before they are tested in the laboratory. Moreover, with the ability to execute fine grained engine simulations at a faster rate, manufacturers will be able design better engines more quickly and at a lower cost

C) IMPROVING MANUFACTURING AND GLOBAL COMPETITIVENESS

Saving Energy While Increasing Profits for Casting Steel (UIUC, NCSA)

What Was Accomplished

Casting steel with fewer defects has the great benefits of making safer steel products and reducing the cost of steel.

The process of continuous casting steel -- the process that turns molten steel into sheets and that is used to produce 92 percent of the world's steel. With annual global steel production at almost 1.5 billion tons (100 million tons in the United States, 96 percent of that through continuous casting), steel production accounts for an important fraction of the total energy consumed and greenhouse gases produced in the world. Even small improvements to this process, can have a profound benefit to society.

Savings/Potential Savings: More than $400 million per year, and keeping many jobs in the US.

FIGURE 9

Modeling Of Metallurgical and Mechanical Processes In Continuous Casting

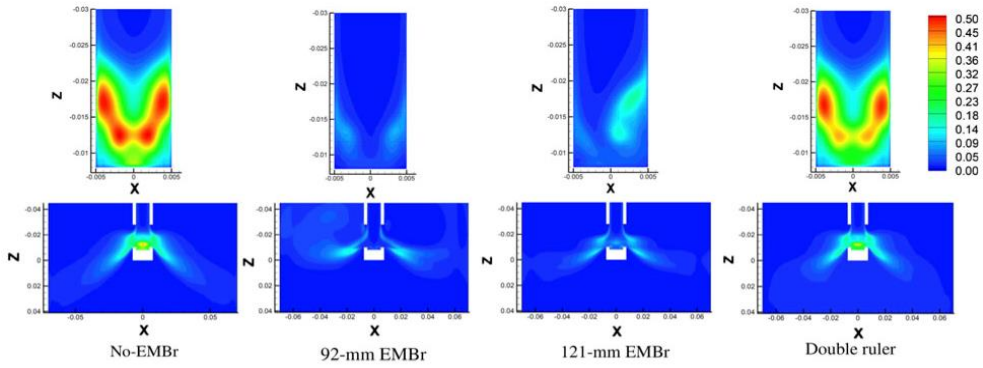

Source: CCC, The University of Illinois, 2014

Organization

This project is supported by the Continuous Casting Consortium at the University of Illinois and the National Center for Supercomputing Applications (NCSA)

Industry

Manufacturing, Scientific Research, Industrial Products

ROI or Innovation Description

Decreasing the material scrapped due to defects such as cracks, even by a small percentage, results in a large net savings to steel manufacturers and customers. Based on the roughly 100 million tons of steel produced each year in the U.S. and at

approximately $400 per ton net cost of scrapping, a one percent reduction in yield loss would save about $400 million per year

Type of ROI

Engineering and manufacturing advances and potentially million of dollars in cost savings

Nature of Research

Continuous casting is an important commercial process to solidify molten metal into semi-finished billets, blooms, or slabs for subsequent rolling. It produces over 95% of steel in the world today, so small improvements have a huge impact. Many defects in final steel products originate in this process. Many defects arise due to problems with the flow pattern, such as level fluctuations that lead to inclusion entrapment. For example, argon gas is often injected into the nozzle to prevent clogging and to help carry away inclusions, in order to improve the steel quality, and also to control the flow. But, this gas also changes the flow pattern, which could be detrimental if not properly taken into account by changing other process variables.

The molten steel flow is also controlled by the nozzle geometry and by Electromagnetic Braking. A magnetic field is generated by applying a direct current to coils, and since molten steel is a conductor, a current field is generated in the fluid, which causes a force that affects flow in the mold cavity. The solidifying shell is very brittle, and if subjected to excessive mechanical or thermal distortion, can form internal cracks, called hot tears. These cracks can lead to defects in the final product, or to expensive and dangerous breakouts, where the solidifying shell tears open and spills molten metal over the casting machine and plant floor. CCC uses supercomputers to improve understanding of this complex process using comprehensive computational models, and to apply those models to find operating conditions to improve the process.

Broader Impact

During the casting process itself, large cracks may cause “breakouts” (where molten steel pours through the crack in the partly-solidified casting, covering the region with superheated molten steel) which are very dangerous (and costly) for the steel plant workers. Undetected defects worsen the mechanical properties of steel products. Increasing casting speed and decreasing spray cooling to conserve just 10% more of the internal energy of the strand would produce energy savings during reheating of $350 million per year, (based on $0.06 per kWH, 100 million tons/year in US) and associated decrease in emissions

Lowering the Fuel Burn and Emissions of Airline Travel (Boeing)

What Was Accomplished

With advanced supercomputers, the Boeing team is able to refine its simulations further and bring safer, more efficient aircraft to market sooner and cheaper.

Boeing researchers used the supercomputers to validate aerodynamics codes for airplane design, saving substantial R&D time that otherwise would be spent calculating solutions. The team validated and improved several aerodynamics codes, saving the

company time and money and likely influencing the process by which next-generation Boeing aircraft are designed and manufactured.

Savings/Potential Savings: HPC use has already saved the aerospace industry multiple tens of billions of dollars over time, and has kept the industry and many jobs in the US. This newest achievement could save additional billions.

FIGURE 10

Aircraft Simulation By Boeing

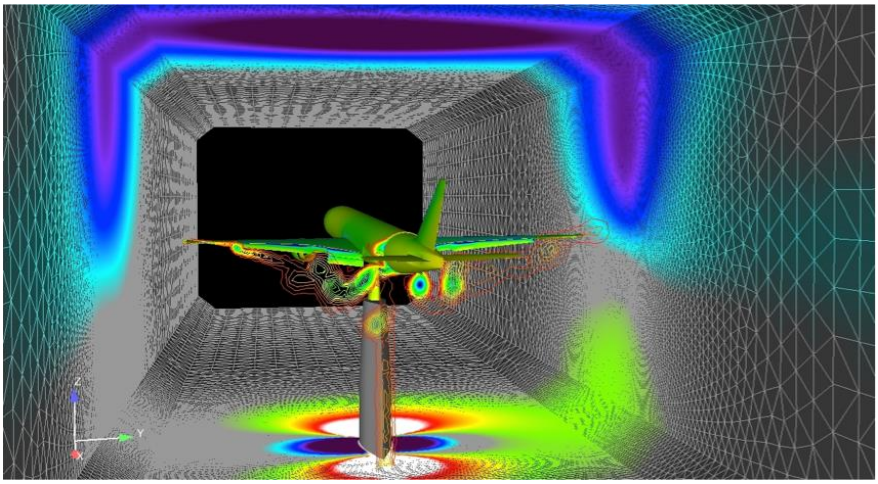

Source: Boeing, 2013

Organization

Boeing

Industry

Aircraft. Aerospace

ROI or Innovation Description

The team validated and improved several aerodynamics codes, saving the company time and money and likely influencing the process by which next-generation Boeing aircraft are designed and manufactured

Type of ROI

Scientific

Type of Innovation

Engineering research innovation

Nature of Research

Researchers at Boeing, utilize simulations to assist in efficient computational modeling and design of passenger and military aircraft. Using HPC resources, enables Boeing to reduce design cost, accelerate R&D while using the best tools and capabilities

towards improving the safety and efficiency of modern aircrafts. HPC simulation allows Boeing researchers to improve diverse design characteristics of aircraft components for example wing design and wind tunnel testing.

Computational modeling of different aircraft components in combination with supercomputers have enabled a 50% reduction in wind tunnel testing for lines development. The computers simulations on the supercomputers very closely matched the wind tunnel results of the tests that Boeing performed in several of its facilities around the world

Supercomputing is now used to design a significant chunk of an aircraft from cab design, cabin noise, interior air quality, high speed wing design, wing tip design, exhaust system design, engine bay thermal analysis, vertical tail and aft body design and much more.

Computational modeling techniques have contributed greatly to the development of new products. Boeing anticipates that the twin factors of need to reduce development cost and cycle time will drive greater dependence on computer simulations. Computational models of different components enabled researchers to accelerate innovation and design cycle times while enabling them to explore more variables, thus enabling development of better products.

Broader Impact

The Boeing team can refine its simulations further and bring safer, more efficient aircraft to market sooner and cheaper, all the while keeping America competitive in a global economy. Boeing Corporation physically tested 77 wing designs for the 767 aircraft, but for the new Boeing 787 Dreamliner only 11 wing designs had to be physically tested mainly because over 800,000 hours of computer simulations on supercomputers had drastically reduced the amount of needed physical prototyping. For the latest 747-8 aircraft only 6 wing designs had to be physically tested.

Developing Future Supersonic Turbomachinery (RamGen)

What Was Accomplished

Computer simulations have helped Ramgen design highly efficient CO_2 compressors that can potentially reduce the capital costs of CO_2 compression by 50 percent and produce a minimum of 25 percent savings in operating costs.

Ramgen Power Systems, a small, Seattle-based energy research and development (R&D) firm, is developing a novel gas compressor system based on shock-wave technology used in supersonic flight applications. These emerging technologies will have significant impact on the broader turbomachinery industry.

Savings/Potential Savings: For each traditional 400 MW clean coal plant, capital cost savings of approximately $22 million and an annual operating cost savings of approximately $5 million..

FIGURE 11

High Resolution Simulation Showing Complex Structures.

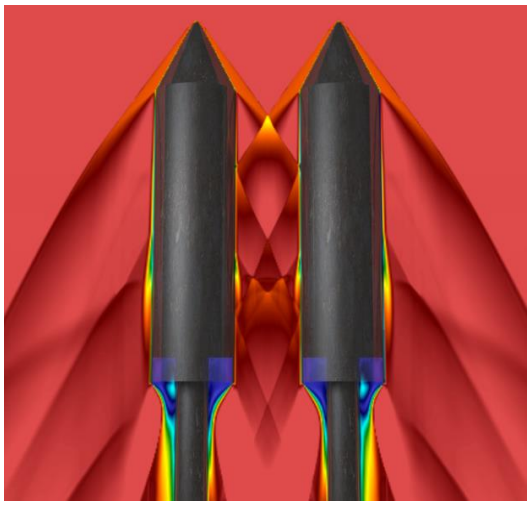

Source: Ramgen, ORNL, 2013

Organization

Collaboration between scientists and engineers from Ramgen under funding from Department of Energy and partner Dresser-Rand

Industry

Energy, Aerospace, Manufacturing

ROI or Innovation Description

Engineering advance

Type of ROI

$4 million savings in gas compressor system development.

Nature of Research

Ramgen demonstrated an approach combining high performance computing based computational models to optimize turbomachinery and extensive testing based validation. Simulation models enabled Ramgen to accelerate R&D resulting in significant reduction in time required to optimize technology performance for commercial applications. Realization of these complex and advanced technologies requires groundbreaking advances in fluid dynamics simulations at extremely high speeds.

Ramgen is utilizing shockwave compression, similar to phenomena encountered in supersonic ramjet inlets, to achieve very high levels of turbomachinery performance. Ramgen executed thousands of design combination simulations, with up to 50 parameters each, to find the optimal designs. Using supercomputers to execute their

design simulations, Ramgen cut the projected time from concept to commercialization by at least 2 years and reduced the cost by over $2M.

Additionally the simulation driven optimization process has been perfected to require just 8 hours, down from previous requirements of months. Combined effect of these advances is enabling Ramgen to advance the design of turbomachinery to a timeframe that simply would not have been possible without supercomputers and simulation models. The supercomputing driven accelerated R&D will enable Ramgen to test a 13K horsepower CO_2 compressor. By applying HPC to design optimization, Ramgen's work has the potential to reduce decades of traditional development process for aerodynamic refinement of shockwave compression based turbomachinery.

Broader Impact

Ramgen's HPC based designs of shockwave compression technology at the scale, represent a paradigm shift in how new turbomachinery is developed. This process can be scaled to different types of gases as well. These technologies can now enter the marketplace sooner. This compressor technology is projected to reduce the capital costs of CO_2 compression by 50 percent and produce a minimum of 25 percent savings in operating costs. Applying these cost savings to a traditional 400 MW clean coal plant would result in capital cost savings of approximately $22 million and an annual operating cost savings of approximately $5 million, per plant.

D) Innovative Disaster Mitigation And Recovery

Saving Lives and Reducing Cost from Tropical Cyclones (DOD MOD)

What Was Accomplished

Without supercomputer resources or computational models, it would be impossible to reasonably predict potentially life-threatening hurricanes and cyclones.

A new version of the software enabled researchers to more accurately predict the track, intensity and size of tropical cyclones with nearly a 120 hour lead time. This new 5-day lead time allows more time to activate emergency evacuations and develop contingency plans.

Savings/Potential Savings: Avoiding hundreds of millions of dollars annually in evacuation costs, as well as the loss of human lives from less-accurate forecasts.

FIGURE 12

Predictions Of Hurricane Irene (Computer Visualizations)

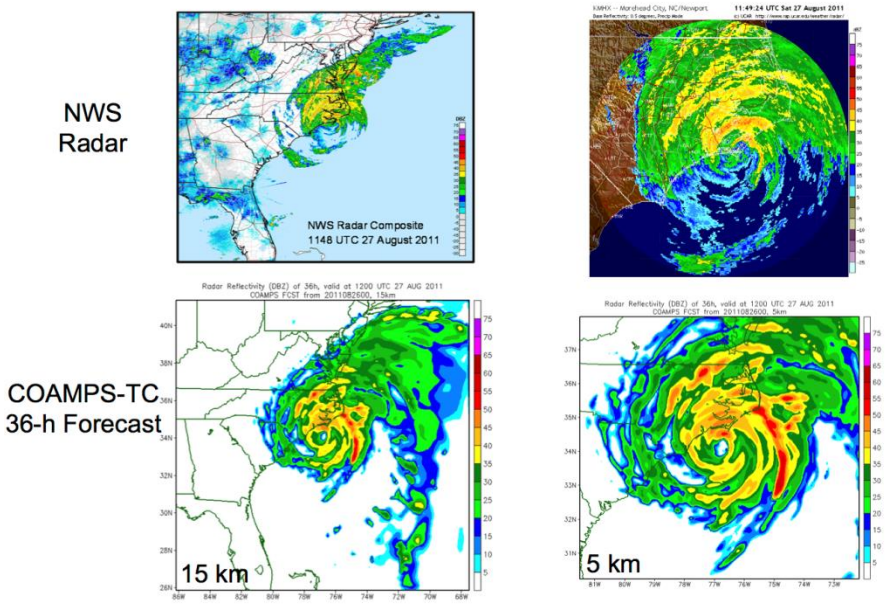

Source: Naval Research Laboratory, 2014

Organization

The Coupled Ocean/Atmosphere Mesoscale Prediction System for Tropical Cyclones (COAMPS-TC™) has been developed by the Naval Research Laboratory (NRL) Marine Meteorology Division in Monterey, California

Supporting Organizations

DOD HPC Modernization program

Industry

Meteorology, weather forecasting, military

ROI or Innovation Description

Significantly improved numerical model for tropical cyclone forecasting. Documented savings of $6 million in one event alone, with far greater potential savings

Type of ROI

Scientific Research and Savings

Nature of Research

The software developed by the Naval Research Laboratory, is designed specifically for forecasting tropical cyclones. The COAMPS-TC model contains a representation of dissipative heating near the ocean surface, which has been found to be important for tropical cyclone intensity forecasts. The system also contains a flexible nesting design

that has proven useful when more than one storm is present in a basin at a given time as well as special options for moving nested grid families that independently follow individual tropical cyclone centers of interest.

The system has options to predict the ocean circulation, temperature, and salinity, and ocean surface waves as well as the interactions between the atmosphere, ocean circulation, and waves (Multiple Paragraphs Detailed Description of what was done… Technology detail oriented)

Broader Impact

This has led to significant improvements in the Navy's effort to improve tropical cyclone intensity prediction (COAMPS-TC was the best numerical model for intensity in the Atlantic during the recent 2010-2012 seasons).

In addition, this research will lead to new capabilities in the form of ensemble forecasts, providing the Navy with probabilistic forecasts of tropical cyclone intensity and structure for the first time. It is also expected this research will help motivate new field campaigns that focus on the key measurements needed to further advance our understanding of the convective structure and dynamics of these systems as well as provide validation.

Providing Advanced Warning of Storm Surge and Waves (RENCI)

What Was Accomplished

Supercomputers are used to provide more detailed predictions of coastal storm impacts for vulnerable coastal areas. The University of North Carolina (UNC) and collaborators at Seahorse Coastal Consulting have developed the ADCIRC Surge Guidance System (ASGS) to enable federal and academic research groups to provide high-resolution predictions of storm surge and waves for vulnerable parts of the US Atlantic and Gulf of Mexico coasts. ASGS automates the running of the ADCIRC tide and storm surge model, providing a robust framework for real-time computations in HPC environments. The predictions of storm surge inundation and wind-wave characteristics are made available to forecasting experts and emergency management groups for impacts assessments, evacuation planning, pre-positioning of supplies and response personnel, for search and rescue, and for other event-based decision support as needed. The products provide an additional source of critical information to decision makers, augmenting official guidance from the National Weather Service

Savings/Potential Savings: The costs of the computational resources needed for detailed, high-resolution coastal storm impact predictions inhibits operational centers from directly running the finite element ADCIRC model. Because of RENCI and UNC's work. operational centers have access to additional, much higher resolution geospatial information, potentially affecting large parts of the coastal population and infrastructure. The advancements have led to discussions with the National Hurricane Center to make this type of information easily accessible to storm surge forecasters.

FIGURE 13

Coastal Emergency Risks Assessment Site: Forecast of Hurricane Irene

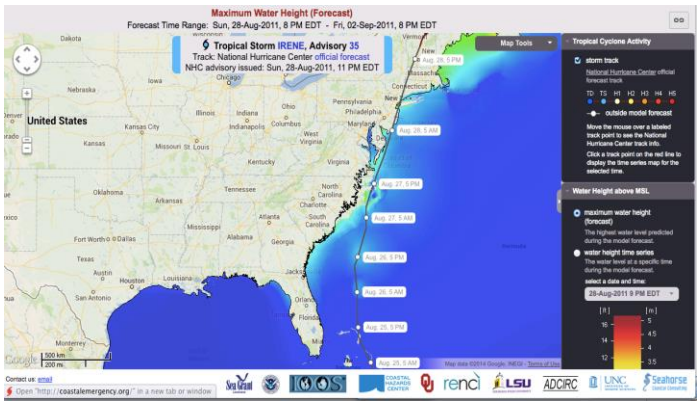

Source: Renaissance Computing Institute, 2014

Organization

The research was conducted at Renaissance Computing Institute and UNC

Industry

Weather Forecasting, Oceanography, Coastal Storms Impacts and Risk Assessment

ROI or Innovation Description

This accomplishment has proven to be an innovation in real-time modeling of natural systems and has great potential to generate financial ROI by substantially improving our ability to rapidly and accurately predict the impacts of severe storms in coastal areas. It has enabled more effective awareness and dissemination of destructive impacts of storm surge and waves during events. Together, these greatly improve the ability of coastal communities to respond and be more resilient to catastrophic tropical cyclones and extratropical storm systems

Type of ROI

Innovation and Scientific Research in coastal hazards prediction

Nature of Research

RENCI's high-performance computing resources have enabled the development of a robust forecast system with dedicated supercomputer cycles, which for the last four years has been an important tool during East Coast tropical and extratropical storms. Typical operational resolutions of other products, such as the National Weather Service's surge model, show coastal detail at 2-5 kilometers. The ASGS system is a major leap forward in the evolution of tools and software used to model natural systems. The research involves not only how to rapidly compute predictions, but also how to most effectively communicate the results to decision makers. During Hurricane Sandy (2012), the U.S. Coast Guard (USCG) extensively used the ASGS output graphics, and as Vice Admiral R. C. Parker said, "significantly closed the gap that exists within our current suite of weather prediction tools." During Hurricane Irene in 2011, the USCG

used information obtained through this modeling effort in deciding to move its fleet Command and Control from Portsmouth, VA, to St. Louis, MO. The model allowed the USCG to understand that even though Hurricane Irene was a relatively weak storm, coastal flooding and inundation was likely to be widespread. Events unfolded such that the Portsmouth USCG base was flooded and lost power shortly after staff evacuated. The modeling work related to Hurricane Irene was recognized with the U.S. DHS's Science and Technology Impact Award in 2012

Broader Impact

The software advances in ASGS and the use of high-performance computing assets for coastal storm predictions and risk awareness can be applied to understanding the full spectrum of coastal risks and to research into the sustainability of coastal infrastructure. The ASGS system provides the computational framework for extending the research into determining optimal, risk-based evacuation routes and shelter locations during hurricanes, as well as to civil infrastructure planning assessments for the New York Harbor area in the aftermath of Hurricane Sandy

Better Understanding and Forecasting of Earthquakes (San Diego)

What Was Accomplished

A research team led by SDSC Computational Scientist Yifeng Cui has accelerated seismic simulations to generate advanced hazard maps and this research team developed software that significantly advances this map-making ability.

Savings/Potential Savings: Billions of dollars in preventing property loss/damage, along with protecting lives

FIGURE 14

Simulating Earthquake ground motions relevant to building engineering design

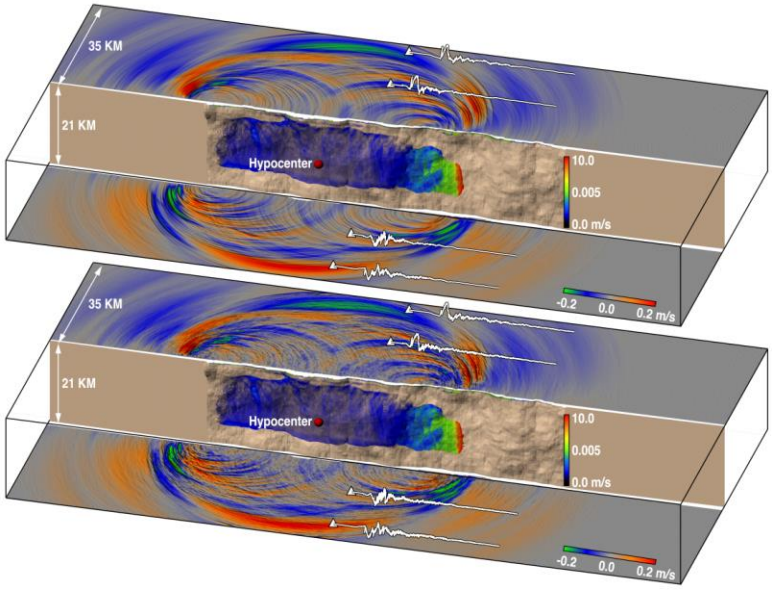

Source: Cui et al. 2013

Organization

This research was carried out by researchers in University of San Diego, San Diego State University, and the broader Southern California Earthquake Center consortium.

Supporting Organizations

Southern California Earthquake Center, Oak Ridge National Laboratory and National Center for Supercomputing Application.

Industry

Scientific Research, Seismic Simulations

ROI or Innovation Description

Researchers simulated realistic 0-10 hertz ground motions on a mesh comprising 443 billion elements in a calculation that includes both small-scale fault geometry and media complexity at a model size far beyond what has been done previously

Type of ROI

Innovative Earthquake Simulation Optional

Nature of Research

In Southern California, it's only a matter of time before a massive earthquake along the San Andreas fault strikes. To help predict the potential destructive power of this quake, supercomputers are being called on to model how waves are generated by and travel

outward from the fault, and then interact with complex 3D geological structures, such as sedimentary basins in the Los Angeles region.

Realistic 0-10 Hz earthquake ground motions were simulated, the largest-ever earthquake simulation performed. The code sustained 2.33Petaflop/s on the supercomputer, resulting in a speedup in key strain tensor calculations critical to probabilistic seismic hazard analysis by a factor of 110.

This achievement makes a California state-wide hazard model a goal reachable with existing supercomputers

Broader Impact

Seismic simulations were accelerated to generate a more advanced hazard map, with the potential saving of thousands of lives and properties upon completion of California state-wide seismic hazard map. Savings could be many billions dollars. Incorporating the simulation results into the U.S. Geological Survey's National Seismic Hazard Maps will set building codes and insurance rates, as well as provide short-term forecast of the frequency of damaging earthquakes in California over a specified time span.

Producing realistic seismograms at high frequencies relevant to building engineering design, up to 10 hertz, will require a major increase in supercomputing power with exascale hardware capabilities. Moreover, to realistically calculate seismic hazards across huge expanses of California and other tectonically active regions requires rapid recompilation of the hazard map to reflect short-term probability variations provided by operational earthquake forecasting.

EXHIBITS:

EXHIBIT 4:
USE OF THE HPC ROI RESULTS BY OTHER ORGANIZATIONS: NCSA/NSF (2016)

Special Study

An Investigation and Evaluation of the Scientific Results from the NCSA Blue Waters Supercomputer System

Earl C. Joseph, Ph.D.
Robert Sorensen
Steve Conway
Kevin Monroe

IDC OPINION

In the U.S. and other developed countries, leadership-class supercomputers have played a major role in advancing science, boosting industrial competitiveness, and improving the quality of daily life for average citizens. NCSA's Blue Waters supercomputer is well known throughout the world and, as this study shows, it has enabled many impressive innovations that would not have been possible with typical, less-capable supercomputers.

During the past five years, national political leaders in the U.S., Europe, and China have recognized the ability of leadership-class supercomputers to help transform their scientific innovativeness, economies, their societies, and their understanding of the natural world.

The intent of this study was to investigate the scientific returns from research projects conducted on the Blue Waters supercomputer system at NCSA. In this study IDC, collected quantitative and qualitative information of the value of scientific work done on the Blue Waters supercomputer. In pursuit of this goal, IDC leveraged its related research work, including but not limited to the methodology we developed and the comparative data we gathered in studies for the Department of Energy on quantifying the value/impact of HPC-enabled scientific innovations. To make the research more specific to Blue Waters, IDC surveyed a set of key Blue Waters researchers on the value of Blue Waters, using a specific survey framework. In total, IDC surveyed 31 Blue Waters-supported research projects in this study. The list of the participants and their projects can be found in the appendix.

To help show the value on the innovations, IDC used a new innovation ranking scale -- a modification of the scale we have been using successfully in the DOE ROI studies -- that is based on two complementary rankings of innovations, combined to form an overall ranking. We believe this new innovation index provides a useful way of looking at and ranking HPC-enabled innovations around the world.

- On both the scales, IDC has used successfully in our DOE studies—ranking the importance and the impact of HPC-enabled scientific innovations—as well as on the third scale IDC developed for the present study, Blue Waters scored very high on both of these scales.
- On the combined, innovation class scale, Blue Waters-enabled innovations as a group substantially outranked IDC's global data base of scientific innovations enabled by supercomputers.

Key Observations:

- Not once during the 31 interviews IDC conducted with Blue Waters users for this study did any of them point to shortcomings of the supercomputer. On the contrary, all of the researchers had praise for the system and enthusiastically reported on the progress it has enabled for their work.
- NCSA did an unusually thorough job of preparing for Blue Waters, first and foremost by working closely with more than 20 users group for two years to identify requirements that would help determine the most appropriate architecture and configuration for the new leadership-class supercomputer. This extreme user focus contrasts sharply with the "build-it-and-they-will-come" approach some other national HPC centers take to future supercomputers. As a result of NCSA's thoroughness, Blue Waters has been exceptionally appropriate and productive for its NSF user community.

    - IDC applauds NCSA's bold decision not to optimize Blue Waters for superior performance on the narrow benchmark test used to determine rankings on the semi-annual Top500 supercomputers list.
    - If more leadership-class HPC sites resisted this political temptation, vendors would be more motivated to design HPC systems that are applicable to a broader range of user needs.
- The innovations produced so far by Blue Waters users are crucially important as a group because they constitute substantial steps forward in major disciplines—even though it's far too soon to estimate their long-term value for scientific and industrial research. It's probably safe to say that they will contribute to a series of other notable discoveries over time.
- A substantially higher percentage of the Blue Waters-enabled innovations qualified for Class 1 & 3 innovation levels, and substantially more Blue Waters-enabled innovations ranked in the top three classes than was true of IDC's global data base of innovations.

In IDC's opinion, these findings confirm that Blue Waters has proven to be an exceptionally—and in some case uniquely—competent platform for accelerating scientific innovation. The Blue Waters-enabled innovations described and ranked in this study will produce strong benefits for the scientific disciplines they belong to. They also have great potential for benefiting U.S. industry and American society as a whole over time.

EXECUTIVE SUMMARY

In order to accelerate scientific advances and economic growth, nations around the world are increasingly investing in large, leadership-class supercomputers. These computers provide a tool that can magnify researchers' capabilities and greatly improve their efficiency in conducting science. When combined with advancements in application design, big data, and top researchers, these systems have become the backbone for new scientific discoveries.

Key Findings

- Where scientific productivity is concerned (as opposed to standard benchmark tests), Blue Waters is certainly one of the top few supercomputers in the world. Blue Water's usefulness and productivity is notable in large part because NCSA worked so closely with its user community to help determine what kind of leadership-class supercomputer would best support the demanding requirements for their scientific work.
- This study confirms that the time and money invested in Blue Waters has been paying off handsomely. Not once during the 31 interviews IDC conducted with Blue Waters users for this study did any of them point to shortcomings of the supercomputer. On the contrary, all of the researchers had praise for the system and enthusiastically reported on the progress it has enabled for their work.
- IDC has seldom encountered such a diverse group of scientific researchers who all agree on something—in this case, the value of the Blue Waters supercomputer.

Surveyed experts cited a broad range of research projects for which the Blue Waters supercomputer has been critical. A few examples will suffice here—many more appear later on in this document.

- "The computational resources at Blue Waters have been instrumental in allowing our group to perform large-scale simulations which have **allowed us to study higher dimensional effects in inertial fusion plasmas**…[Our] simulations performed on Blue Waters have been published in many high impact journals, including *Nature*."
- "Blue Waters is the only place where massively parallel sparse solver technology can be tested. This exciting technology advancement will lead to a massive leap in terms of advances in design and manufacturing, and understanding the properties of the Earth subsurface to **allow a major breakthrough in oil exploration.**"
- "Our results show the potential for **substantial reduction in Ebola spread..**."
- "Newly discovered flow features identified in our simulations indicate mechanisms that may play an important role in discerning **the most devastating, tornado-producing thunderstorms**.
- "We also implemented a **new model for black hole growth**, mergers, and feedback to realistically simulate the properties of the higher mass galaxies."
- "This will be **the first systems-level simulator that targets a specific microbe (E. coli) and will be able to simulate populations of cells with a resolution ranging from individual gene concentrations to population dynamics**.
- "We address the problem of **decision making with uncertainty in the context of climate change** that affects the world in a non-uniform manner (more warming in the polar regions)."

Blue Waters is the only place where massively parallel sparse solver technology can be tested.

Leadership-class supercomputer performance features frequently mentioned in other IDC studies include the ability to:

- Handle large three-dimensional simulations
- Analyze complex physicals systems with a high degree of precision
- Perform critical calculations in real time
- Conduct advanced multiscale and multiphysics simulations
- Reduce the need for empirical testing (including animal experimentation)
- Use data sets large enough to accurately simulate physical environments with a high degree of reliability and fidelity

Benefits of a Leadership Class Supercomputer

- **New Science.** Some stressed that a more powerful system will be needed to enable new science. Areas cited as ripe for new scientific development included personalized medicine, computational astrophysics, molecular dynamics, and social simulation.
- **Improved Simulations.** Others stated that a larger system would drive even better capabilities in existing areas of computational studies, citing the benefits of a larger, more capable system to provide more accurate simulations, enhanced simulation resolution, reduced time to solution, and greater opportunities for advanced multiscale research.
- **Less Expensive Access.** A third group indicated that a larger system would allow more users to benefit from the system, thereby increasing scientific productivity and cost-effectiveness.

Measuring the Scientific Benefits from Blue Waters

World scientific leadership and innovation leadership are becoming more dependent on the use of HPC/supercomputers every year. Economic leadership increasingly directly results from a nation's or an industry's or an enterprise's application of supercomputers in innovative and productive ways. Many countries/regions (such as the U.S., Japan, China, Russia, and Europe) are putting in place plans to gain leadership in innovation and economic progress by more broadly applying HPC/supercomputing across many different research areas.

In order to better quantify the overall impact of leadership class supercomputers for scientific innovations, IDC uses a rating system that measures both the *importance* and the *impact* of each innovation in the existing data set IDC collected for prior DOE studies, plus the new Blue Waters innovations collected in this research study. The new overall innovation ranking was created based on a combination of the two complementary rankings of the innovations, in the form of 8 class levels of innovations.

IDC Observations

- NCSA is doing an excellent job serving the needs of the researchers on Blue Waters.
- NCSA did an unusually thorough job of preparing for Blue Waters, first and foremost by working closely with more than 20 users group for two years to identify requirements that would help determine the most appropriate architecture and configuration for the new leadership-class supercomputer. This extreme user focus contrasts sharply with the "build-it-and-they-will-come" approach some other national HPC centers take to future supercomputers. As a result of NCSA's thoroughness, Blue Waters has been exceptionally appropriate and productive for its NSF user community. IDC applauds NCSA's bold decision not to optimize Blue Waters for superior performance on the narrow benchmark test used to determine rankings on the semi-annual Top500 supercomputers list. If more leadership-class HPC sites resisted this political temptation, vendors would be more motivated to design HPC systems that are applicable to a broader range of user needs.
- The innovations produced so far by Blue Waters users are crucially important as a group because they constitute substantial steps forward in major disciplines—even though it's far too soon to estimate their long-term value for scientific and industrial research. It's probably safe to say that they will contribute to a series of other notable discoveries over time.
- A substantially higher percentage of the Blue Waters-enabled innovations qualified for Class 1 & 3 innovation levels, and substantially more Blue Waters-enabled innovations ranked in the top three classes than was true of IDC's global data base of innovations.

Substantially more Blue Waters-enabled innovations ranked in the top three classes than was true of IDC's global data base of innovations.

TABLE OF CONTENTS

TABLE OF CONTENTS — Continued

LIST OF TABLES

LIST OF FIGURES

IN THIS STUDY

Methodology

The intent of this study was to investigate the scientific returns from research projects conducted on the Blue Waters supercomputer system at NCSA.

IDC assisted NCSA by gathering quantitative and qualitative information of the value of scientific work done on the Blue Waters supercomputer. In pursuit of this goal, IDC leveraged its related research work, including but not limited to the methodology we developed and the comparative data we gathered in our studies for the Department of Energy on quantifying the value/impact of HPC-enabled scientific innovations.

To make the research more specific to Blue Waters, IDC surveyed a set of key Blue Waters researchers on the value of Blue Waters, using a specific survey framework.

IDC used a new innovation ranking scale -- a modification of the scale we have been using successfully in the DOE ROI studies -- that is based on two complementary rankings of innovations, combined to form an overall ranking. We believe this new innovation index provides a useful way of looking at and ranking innovations around the world.

IDC surveyed 31 Blue Waters research projects in this study. The list of the survey participants and their projects can be found in the appendix. Their inputs were reviewed by selected members of the IDC ROI review committee, which includes:

- Paul Muzio, City University of New York
- Rupak Biswas, NASA Ames
- Vijay Agarwala, Virginia Tech.
- Swamy Akasapu, General Motors
- Alex Akkerman, Ford Motor Company
- C. Scot Atkins, Industry Expert – Advanced Analytics/Dense Supercomputing
- Doug Bail, The Boeing Company
- Jeff Broughton, NERSC/Lawrence Berkeley National Lab
- Paul Buerger, Avetec
- Clayton Chandler, Credit Suisse Group AG
- Sharan Kalwani, Fermilab
- Chris Catherasoo, California Institute of Technology
- Simon Burbidge, Imperial College London
- Jack Collins, National Cancer Institute
- Steve Finn, Emagine IT
- Keith Gray, BP
- James Kasdorf, Pittsburgh Supercomputing Center
- Arno Kolster, PayPal
- Doug Kothe, Oak Ridge National Laboratory
- Jysoo Lee, KAUST
- David Martin, Argonne National Laboratory

- Michael Resch, HLRS, University of Stuttgart
- Ryan Quick, PayPal
- Stephane Requena, GENCI
- Vince Scarafino, Industry Expert
- Suzy Tichenor, Oak Ridge National Laboratory

IDC's HPC ROI and Innovation Measuring Approach

For our studies for DOE, IDC created a set of ROI macroeconomic models to quantify the impact of investments in HPC. The models include both the financial ROI (return on investment) returns and the innovation ROR (return on research) returns from projects done on supercomputers. The results of this research for DOE are published at: www.idc.com/ROI

Our 2013 pilot ROI study for DOE investigated how high-performance computing (HPC) investments can improve economic success and increase scientific innovation. This research was focused on the common good and should be useful to all HPC centers around the world. The study 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.
- A new *innovation index* that provides a means of measuring and comparing innovation levels.

For the Blue Waters study, IDC modified our DOE study methodology as described above in the Methodology section of this report.

SITUATION OVERVIEW

The world's most powerful computers, called leadership-class supercomputers, have contributed enormously to scientific advances, national security, economic progress, and the quality of life, Leadership-class supercomputers are indispensable for many applications that are important for the societies they serve, such as the following:

- Predicting severe storms that can devastate lives and property
- Providing accurate daily/weekly weather forecasts needed by the transportation, agricultural, and tourism industries
- Better understanding fundamental scientific principles
- Improving the design and safety of power plants and developing technologies to exploit alternative energy sources
- Detecting sophisticated cyber security breaches, insider threats, and electronic fraud.
- Better understanding natural processes
- Advancing fundamental science as a prerequisite for later advancements in applied research and development, including industrial R&D

International Recognition of Supercomputing's Strategic Value

Political leaders around the world have increasingly recognized the transformational power of leadership-class supercomputers and are supporting initiatives to enable their countries to compete effectively in the worldwide race to develop future supercomputers with unprecedented speed, called "exascale" computers, in the next 5-6 years. U.S. President Obama, European Commission leaders, and Chinese leaders all have endorsed initiatives to reach this goal. Japan is the only other country with the technical ability to accomplish this. In IDC's opinion, the U.S. and Japan have the greatest experience in designing leadership-class supercomputers that are capable of supporting a broad spectrum of challenging scientific and industrial research problems.

SURVEY FINDINGS

Blue Waters Scientific Innovation Success Stories (Return on Research)

IDC surveyed 31 researchers/projects that use the Blue Waters supercomputer. Descriptions of the projects are given in the tables in this section of the report. Fuller information about these researchers and their projects appears in the appendix.

Project Titles/Research Areas

- Simulation-Based Policy Analysis to Reduce Ebola Transmission Risk in Air Travel
- Hydrogen Under Extreme Conditions
- Petascale Quantum Simulations of Nano Systems and Biomolecules
- Custom Genotyping Chip for African Populations
- Dynamic Coarse-Grained Models for Simulations of Large-Scale Biophysical Phenomena
- Instrumenting Human Variant Calling Workflow on Blue Waters
- Sensing the Environment: a Glimpse into the Microscopic Mechanisms of Pain
- Simultaneous VLBA Polarimetric Observations
- VY CMa II: Component-Level Polarization Analysis
- Unlocking the Mysteries of the Most Violent Tornadoes
- Simulating the Earliest Galaxies with Enzo and Blue Waters
- ArcticDEM a White House Initiative to Produce a High-Resolution, Time-Dependent Elevation Model of the Arctic using Blue Waters
- Petascale Particle-in-Cell Simulation of Kinetic Effects in High Energy Density Plasmas
- High Resolution Earth System Modeling Using Blue Waters Capabilities
- Advancing Genome-Scale Phylogenomic Analysis
- Nuclear-Electronic Orbital Calculations on Molecular Systems
- Particle Tracking and Turbulent Dispersion at High Reynolds Number on Blue Waters
- Evolution of the Small Galaxy Population from High Redshift to the Present
- Quantum-classical Path Integral Simulations of Ferrocene-Ferrocenium Charge Transfer in Solution
- Policy Responses to Climate Change in a Dynamic Stochastic Economy
- Advanced Nanoelectronic Device Design with Atomistic Simulations
- Modeling Heliophysics and Astrophysics Phenomena with a Multi-Scale Fluid-Kinetic Simulation Suite
- Big Data on Small Organisms: Petascale Simulations of Data-driven, Whole-cell Microbial Models
- Direct Numerical Simulation of Fully Resolved Vaporizing Droplets in a Turbulent Flow
- Sparse Matrix Factorization In Solid Mechanics And Geophysics on CPUs and GPUs
- Core-Collapse Supernovae through Cosmic Time
- Location- Specific Space Weather Hazard to Electric Power Grids Calculated on a Global Scale
- Ensembles of Molecular Dynamics Engines for Assessing Force Fields, Conformational Change, and Free Energies of Proteins and Nucleic Acids
- Lattice QCD
- Predicting Protein Structures with Physical Petascale Molecular Simulations

Blue Waters Scientific Successes by Innovation Area

IDC sorted the Blue Waters scientific innovations into six high-level categories that our research experience tells us are better suited to communicating value to non-technical audiences. (Note: With fundamental science innovations such as these, it is sometimes difficult to limit their potential impact to a single category, so some of these advances might just as well have been placed in a different category.)

1. Better Products
2. Major Scientific Breakthrough
3. Cost Saving
4. Created New Approach
5. Discovered Something New
6. Helped Society

Blue Waters Successes: Supporting Better Future Products

The five project examples described in Table 1 involved innovations with strong potential to enable companies to produce better future products in a range of industries and market sectors, including:

- Electronics
- Medicine/health care
- Geospatial mapping
- Fusion energy
- Oil and gas

Key Quotes

"**The computational resources at Blue Waters have been instrumental in allowing our group to perform large-scale simulations which have allowed us to study higher dimensional effects in inertial fusion plasmas**, and to make quantitative comparisons between simulations and experiments in plasma-based accelerator. In 2015, simulations performed on Blue Waters have been published in many high impact journals, including *Nature*."

"Blue Waters is the only place where massively parallel sparse solver technology such as WSMP can be tested. This exciting technology advancement will lead to a massive leap in terms of advances in design and manufacturing, and understanding the properties of the Earth subsurface to **allow a major breakthrough in oil exploration.**"

Table 1

Blue Waters Scientific Successes: Supporting Better Future Products

Project Title	Project Description
Petascale Quantum Simulations of Nano Systems and Biomolecules	This project builds on a set of open-source petascale quantum simulation tools being developed via the PetaApps program; these tools enable quantum simulations of unprecedented size in the broad area of nanoscience and nanotechnology. The goals of this project are to radically advance the fundamental understanding of key nanoscale phenomena and processes in the following areas: (i) graphene-based electronics, which will likely form the basis for new generations of ultrafast, ultradense electronic circuitry, (ii) catalytic growth of carbon nanotubes, aiming at discovering the catalysts and conditions that would lead to controlled growth of carbon nanotubes with desired electronic properties, (iii) DNA-based sensor structures for detection of DNA damage, and (iv) metal-induced accelerated onset of Alzheimer's disease.
ArcticDEM a White House initiative to produce a high-resolution, time-dependent elevation model of the Arctic using Blue Waters	Last summer, President Obama announced that a complete DEM of the Arctic will be completed by next year. The ArcticDEM will be a continous, 8m and then 2m DEM mosaic of all land masses above 60 degrees North (above the latitude of SRTM), constructed from DigitalGlobe Inc. satellite imagery. SETSM is the primary software for DEM extraction for ArcticDEM and we have been working feverishly to ramp up data processing on the Blue Waters Supercomputer while simultaneously continuing SETSM development and post-processing, mosaicking and registration methods. Nearly all of Alaska, Greenland, Scandanavia and large parts of Russia and Canada have already been processed to 8m. This enables data filtering and algorithm development. Once complete, the resolution will be increased to 2m for the final product.
Petascale Particle-in-Cell Simulation of Kinetic Effects in High Energy Density Plasmas	The UCLA Simulation Group has been on the frontier of high performance kinetic simulation of plasmas since its early days with its founder John Dawson. Over the subsequent 40+ years, the group has developed a full suite of codes (with various models for fields and particles) to study kinetic effects in plasmas for a variety of applications, including plasma-based accelerators, laser fusion, and basic plasma physics. The computational resources at Blue Waters have been instrumental in allowing our group to perform large-scale simulations which have allowed us to study higher dimensional effects in inertial fusion plasmas, and to make quantitative comparisons between simulations and experiments in plasma-based accelerator. In 2015, simulations performed on Blue Waters have been published in many high impact journals, including Nature. These simulation results and our code development efforts, including the latest timing results on GPU's and the Intel Xeon Phi processors, will be presented.
Advanced Nanoelectronic Device Design with Atomistic Simulations	Advanced Nanoelectronic Device Design with Atomistic Simulations
Sparse Matrix Factorization in Solid Mechanics And Geophysics on CPUs and GPUs	Sparse matrix factorization is a critical algorithm in many science, engineering, and optimization applications. The performance of the massively parallel direct multi-frontal solver Watson Sparse Matrix Package (WSMP) for solving large sparse systems of linear equations arising in an implicit finite element method in solid mechanics and an inversion problem in geophysics was evaluated on Blue Waters and achieved new records in sparse matrix factorizations. We performed fullscale benchmarking tests up to 65,536 cores (100 Tflop/s) on XE6 nodes using assembled global stiffness matrices and load vectors with 5 million to 40 million unknowns. Blue Waters is the only place where massively parallel sparse solver technology such as WSMP can be tested. This exciting technology advancement will lead to a massive leap in terms of advances in design and manufacturing, and understanding the properties of the Earth subsurface to allow a major breakthrough in oil exploration.

Table 1

Blue Waters Scientific Successes: Supporting Better Future Products

Project Title	Project Description

Source: IDC 2016

Blue Waters Successes: Major Scientific Breakthroughs

The project examples described in Table 2 involved scientific innovations with strong potential to benefit society at large and deepen human understanding of the natural world in the following areas and others:

- Epidemic diseases
- Bacteriology
- Severe weather prediction
- Cosmology
- Space exploration
- Air travel (wind shear)

Key Quotes

"Our results show the potential for **substantial reduction in Ebola spread** by changing current [aircraft] boarding and disembarkation procedures. This approach leads to high computational cost, which is handled through massive parallelization on Blue Waters."

Our results show the potential for substantial reduction in Ebola spread

"[The project involves] the application of CG models and software to the study of model systems such as **the nucleation and growth of the HIV-1 capsid protein lattice**…with specific emphasis placed on the use of the Blue Waters platform."

"Newly discovered flow features identified in our simulations indicate mechanisms that may play an important role in discerning **the most devastating, tornado-producing thunderstorms** from those that are much more common and less destructive.

"Further progress was made concerning **tropical cyclone behavior**."

"A substantial algorithmic change in the interpolation of fluid particle velocities on a distributed domain has led to a **dramatic improvement in the scalability of the particle tracking algorithm at 262,144 cores**. We present a number of important results, such as the Lagrangian acceleration autocorrelation, and the scaling of velocity increment statistics conditioned upon the dissipation rate."

"We also implemented a **new model for black hole growth**, mergers, and feedback to realistically simulate the properties of the higher mass galaxies."

"This will be **the first systems-level simulator that targets a specific microbe (E. coli) and will be able to simulate populations of cells with a resolution ranging from individual gene concentrations to population dynamic**s. The Blue Waters Supercomputer with its unique architecture, large-scale simulation capabilities and professional support staff provides the ideal platform to achieve this ambitious goal."

Table 2

Blue Waters Scientific Success: Major Scientific Breakthroughs

Project Title	Project Description
Simulation-Based Policy Analysis to Reduce Ebola Transmission Risk in Air Travel	Air travel is a major cause of the spread of infections. This led to calls for ban on air travel during the recent Ebola outbreak. However, such bans have serious human and economic consequences. Our project is identifying policies and procedures that can provide the same benefit without the negative consequences. This is accomplished using a fine-scale model that tracks individual passenger movement in airplanes. Inherent uncertainties in human behavior make it difficult to accurately predict the consequences of any particular policy choice. Instead, it is more fruitful to determine vulnerabilities under a variety of possible scenarios. We parameterize the sources of uncertainty and perform simulations to cover the range of uncertainties. This approach leads to high computational cost, which is handled through massive parallelization on Blue Waters. Our results show potential for substantial reduction in Ebola spread by changing current boarding and disembarkation procedures.
Dynamic Coarse-Grained Models for Simulations of Large-Scale Biophysical Phenomena	Atomic-resolution computer simulations can provide highly detailed information about biomolecular systems, but are restricted to time- and length-scales that may be insufficient for the study of certain biological phenomena. One alternative technique is the use of “coarse-grained” (CG) models, where degrees of freedom are removed to generate simpler molecular representations while retaining the essential characteristics of the detailed system. Due to their computational efficiency, coarse-grained models provide an appealing tool for the study of relatively large-scale biomolecular phenomena. The application of CG models and software to the study of model systems such as the nucleation and growth of HIV-1 capsid protein lattice will be discussed, with specific emphasis placed on the use of the Blue Waters platform.
Unlocking the Mysteries of the Most Violent Tornadoes	Devastating, long-lived tornadoes are rare, but the death and destruction they cause is significant. We reviewed nearly four years of our work on Blue Waters, focusing on recent supercell thunderstorm simulations in which long-lived, violent tornadoes occur. Newly discovered flow features identified in our simulations indicate mechanisms that may play an important role in discerning the most devastating, tornado-producing thunderstorms from those that are much more common and less destructive. Recent simulations run with 20-meter isotropic grid spacing reveal the internal structure of a long-lived, devastating tornado, which undergoes morphological transitions from a narrow, single-celled tornado to a wide, multiple-vortex tornado.
High Resolution Earth System Modeling Using Blue Waters Capabilities	This project addresses uncertainties associated with numerical modeling of Earth’s climate system and with modeled present and future climate change by conducting high-resolution simulations with the Community Earth System Model (CESM). Pushing model resolution, our group is conducting the first-of-its-kind simulations including the fully-coupled ¼° atm/lnd coupled to the 1° ocn/ice configuration and the fully-coupled ¼° atm/lnd coupled to a 1/10° ocn/ice RCP8.5 scenarios. The former set is used to assess climate sensitivity, 20th Century climate changes, and future climate changes while the latter simulation allows exploration of very high resolution impacts on regional features and climate change processes.
Nuclear-Electronic Orbital Calculations on Molecular Systems	The nuclear-electronic orbital (NEO) method treats electrons and select nuclei quantum mechanically on the same level using an orbital-based formalism with the goal of obtaining a computationally tractable method that includes non-Born-Oppenheimer effects as well as nuclear quantum effects. The NEO method is ideal for studying chemical phenomena such as proton-coupled electron transfer (PCET) because the timescale for proton tunneling is often faster than the timescale for electronic transitions, thereby leading to a breakdown of the Born-Oppenheimer approximation. In applications of the NEO method to PCET, all electrons and one or a few protons are treated quantum mechanically, and a mixed nuclear-electronic time-independent Schrödinger equation is solved using explicitly correlated wavefunctions. Recent advances to the NEO method involving wavefunction and density functional theory will be discussed and benchmarking calculations on small molecules for basis sets will be presented.

Table 2

Blue Waters Scientific Success: Major Scientific Breakthroughs

Project Title	Project Description
Particle tracking and turbulent dispersion at high Reynolds number on Blue Waters	As reported last year, a half-trillion-grid points simulation of turbulence performed on Blue Waters has provided clear evidence of extreme events where local measures of the deformation and rotation of local fluid elements are much stronger than and have topological properties different from previously thought. In this talk we focus on turbulent dispersion at high Reynolds number in the Lagrangian reference frame of an observer moving with the flow. A substantial algorithmic change in the interpolation of fluid particle velocities on a distributed domain has led to a dramatic improvement in the scalability of the particle tracking algorithm at 262,144 cores. We present a number of important results, such as the Lagrangian acceleration autocorrelation, and the scaling of velocity increment statistics conditioned upon the dissipation rate. Particle pair trajectories are also followed both forward and backward in time, which provide unique insights for the physics of enhanced mixing in turbulence.
Evolution of the small galaxy population from high redshift to the present	Creating robust models of the formation and evolution of galaxies requires the simulation of a cosmologically significant volume with sufficient resolution and subgrid physics to model individual star forming regions within galaxies. This project is undertaking this modelling with the specific goal of interpreting Hubble Space Telescope observations of high redshift galaxies. To do this modelling, we are using the highly scalable N-body/Smooth Particle Hydrodynamics code, ChaNGa, based on the Charm++ runtime system on Blue Waters to perform a simulation of a 25 Mpc cubed volume of the Universe with a physically motivated star formation/supernovae feedback model. We also implemented a new model for black hole growth, mergers, and feedback to realistically simulate the properties of the higher mass galaxies.
Modeling Heliophysics Phenomena With A Multi-Scale Fluid-Kinetic Simulations Suite	The investigated physical phenomena occurring when the solar winds (SW) interacts with the local interstellar medium (LISM). These problems include (1) issues related to the mixing of the SW and LISM plasma at the heliopause (the boundary of the heliosphere, the spherical region around the Sun that is filled with solar magnetic fields and the outward-moving solar wind consisting of protons and electrons) in particular due to the heliosphere instability and magnetic reconnection at its surface; (2) the influence of the heliosphere on the observed anisotropy of teraelectronvolt (TeV) galactic cosmic rays; (3) the dynamic effect of non-thermal ions, and (4) the influence of time-dependent phenomena on the energetic neutral atom flux observed by the Interstellar Boundary Explorer (IBEX) space mission.
Big Data on Small Organisms: Petascale Simulations of Data-driven, Whole-cell Microbial Models	This project aims to develop the next-generation of genome-scale, data-driven models for microbial organisms. The project first focuses on the most-studied microbe, the gram-negative bacterium Escherichia coli. Accurate prediction of microbial fitness and cellular state can have profound implications to the way we test hypotheses that are directly related to health, social or economic benefits. This will be a boon for the development and training of the next generation of data-driven predictive methods in molecular and cellular biology. This will be the first systems-level simulator that targets a specific microbe (E. coli) and will be able to simulate populations of cells with a resolution ranging from individual gene concentrations to population dynamics. The Blue Waters Supercomputer with its unique architecture, large-scale simulation capabilities and professional support staff provides the ideal platform to achieve this ambitious goal.
Direct Numerical Simulation of Fully Resolved Vaporizing Droplets in a Turbulent Flow	This numerical study employs direct numerical simulations (DNS) to examine the two-way interactions between freely moving vaporizing droplets and isotropic turbulence. The droplets are fully resolved in 3-D space and time, i.e., not treated as point particles, and all the scales of the turbulent motion are resolved down to the smallest relevant length- and time-scales (the Kolmogorov scales). The emphasis is on the two-way exchange of mass, momentum, and energy between the vaporizing droplets and surrounding turbulent gas. The turbulence is assumed to be isotropic as a first step before considering turbulent shear flows in future studies.

Table 2

Blue Waters Scientific Success: Major Scientific Breakthroughs

Project Title	Project Description
Core-Collapse Supernovae Through Cosmic Time	We study the explosive deaths of massive stars, supernovae, and their contribution to the evolution of the elemental content of the universe. Core-collapse supernovae (CCSNe) are tightly coupled multi-physics events without natural symmetry and require physically and spatially detailed 3D simulations to resolve. The progenitor stars of CCSNe vary in mass, heavy element composition, rotation, and other parameters that affect how the explosion develops, or if it develops at all, and the elemental abundances of the ejecta. Over the course of our project we will compute CCSN models that broadly cover the range of masses and compositions representative of massive stars throughout the history of the universe. In prelude to that survey, we are conducting studies with Blue Waters that will assess the impact of spatial resolution on our simulations.

Source: IDC 2016

Blue Waters Successes: Cost Savings

The project examples described in Table 3 involved scientific innovations with strong potential to benefit broad areas of science and later on to reduce R&D costs for industry in areas including (but not limited to) the following:

- Pharmaceuticals
- Applied biology
- Electronics

Key Quotes

"Harnessing the computation capabilities of Blue Waters, **we explored several pathways of activation [for novel analgesics]** and characterized ion channel conductance and selectivity. Our calculations reveal a novel mechanism for sensing temperature and osmolarity."

"[Our] **new method (HIP-HOP) for classifying sequences into gene families**…substantially improves the accuracy compared to all current alternative methods (including BLAST). [Our version of BAli-Phy] shows scalability to 10,000 sequences, whereas the original implementation could only analyze about 100 sequences.

Table 3

Blue Waters Scientific Successes: Cost Savings

Cost Savings	
Project Title	Project Description
Sensing the environment: a	TRP channels are central to environmental sensation in animals, fungi, and unicellular eukaryotes. Clarifying how TRP channels convert physical and chemical stimuli from the environment into the

Table 3

Blue Waters Scientific Successes: Cost Savings

Cost Savings	
Project Title	Project Description
glimpse into the microscopic mechanisms of pain	allosteric signals underlying channel activation is key to understanding how they control cell excitability in both physiological and pathological conditions. Their relevance in the molecular pathways mediating pain makes them promising targets of novel classes of analgesics. Building on the structural information made recently available for TRPV1 thanks to a series of cryo-microsopy experiments, we performed free energy (metadynamics) simulations on models of TRPV1 embedded in a lipid bilayer. Harnessing the computation capabilities of Blue Waters, we explored several pathways of activation and characterized ion channel conductance and selectivity. Our calculations reveal a novel mechanism for sensing temperature and osmolarity.
Advancing Genome-Scale Phylogenomic Analysis	Research Highlights: 1. Scalable versions of BAli-Phy (3), a Bayesian method for statistical co-estimation of multiple sequence alignments and trees so that they can analyze large datasets (our current implementations show scalability to 10,000 sequences, whereas the original implementation could only analyze about 100 sequences). 2. A new method (HIP-HOP) for classifying sequences into gene families, that substantially improves the accuracy compared to all current alternative methods (including BLAST (1)). 3. New supertree methods with improved accuracy and scalability that will enable the development of divide- and-conquer methods for estimating large phylogenetic trees.
Quantum-classical path integral simulations of ferrocene-ferrocenium charge transfer in solution	Condensed phase electron transfer reactions play a vital role in most biological and synthetic energy transfer pathways. The costs of quantum mechanical calculations scale exponentially with system size; thus, performing highly accurate simulations of realistic condensed phase reactions is notoriously demanding. Quantum-classical path integral (QCPI) is a recently developed highly parallelizable methodology designed to efficiently and accurately simulate the dynamics of a quantum system immersed in a condensed phase environment. It combines a classical treatment of the environment with a path integral representation of the system. The local nature of the quantum paths allows for the system-solvent interaction to be treated exactly, free of approximation. Using Blue Waters' resources, we utilized QCPI to simulate the charge transfer process of the ferrocene-ferrocenium pair in solution with unprecedented accuracy.

Source: IDC 2016

Blue Waters Successes: Creating New Approaches

The project examples described in Table 4 involved innovations that created new and better approaches to scientific investigation in the following domains:

- Astronomy/astrophysics
- Cosmology
- Economic impacts of climate change policy

Key Quotes

"We developed new Quantum Monte Carlo simulation methods to [support] understanding [of] the giant planets, Jupiter and Saturn. However, the observations [from our two experiments] do not agree with each other, **This motivated us to repeat our earlier calculations on Blue Waters**."

"Accurately modeling the **formation and evolution of galaxies over the lifetime of the universe** presents tremendous technical challenges. [We] present simulations (and accompanying published results) that the Enzo collaboration has recently done on the Blue Waters supercomputer."

"We address the problem of **decision making with uncertainty in the context of climate change** that affects the world in a non-uniform manner (more warming in the polar regions)."

Table 4

Blue Waters Scientific Successes: Creating New Approaches

Project Title	Project Description
Hydrogen Under Extreme Conditions	The properties of hydrogen and helium are important for understanding the giant planets, Jupiter and Saturn, but experiments under the relevant conditions are challenging. We have developed new Quantum Monte Carlo simulation methods to treat such systems and using them, have studied molecular dissociation in liquid hydrogen and have observed clear evidence of an extra liquid-liquid phase transition. During the past year, two experiments have reported observations of the transition we predicted, however, the observations do not agree with each other, differing in pressure by a factor of two. This motivated us to repeat our earlier calculations on Blue Waters. It is essential for progress in the high pressure community to resolve the difference between the experiments and computation. After validation, the method can be used with more confidence in modeling the wide variety of astrophysical objects being observed, composed largely of hydrogen and helium under extreme conditions.
Simulating the Earliest Galaxies with Enzo and Blue Waters	Galaxies are complex - many physical processes operate simultaneously, and over a huge range of scales in space and time. As a result, accurately modeling the formation and evolution of galaxies over the lifetime of the universe presents tremendous technical challenges. In this talk I will describe some of the important unanswered questions regarding galaxy formation, discuss in general terms how we simulate the formation of galaxies on a computer, and present simulations (and accompanying published results) that the Enzo collaboration has recently done on the Blue Waters supercomputer. In particular, I will focus on the transition from metal-free to metal-enriched star formation in the universe, as well as the luminosity function of the earliest generations of galaxies and how we might observe it with the upcoming James Webb Space Telescope.
Policy Responses to Climate Change in a Dynamic Stochastic Economy	We extend our Integrated Assessment Model framework, called DSICE (Dynamic Stochastic Integration of Climate and the Economy), for evaluating alternative policy responses to future climate change under both economic and climate uncertainty. We incorporate five interacting climate tipping points into DSICE, and find that the present social cost of carbon (SCC) increases nearly eightfold. Moreover, passing some tipping points increases the likelihood of other tipping points occurring, so that the SCC increases abruptly. The optimal mitigation policy requires zero industrial emission after this midcentury and leads to a path of global average temperature less than 1.5 degree Celsius. We also examine the impact of Bayesian learning of uncertain critical parameters (e.g., climate sensitivity) to decision rules. Furthermore, we address the problem of decision making with uncertainty in the context of climate change that affects the world in a non-uniform manner (more warming in the polar regions).

Table 4

Blue Waters Scientific Successes: Creating New Approaches

Project Title	Project Description

Source: IDC 2016

Blue Waters Successes: Discovered Something New

The project examples described in Table 5 involved innovations centered on new discoveries in the following domains:

- Biology

Key Quotes

"No MD [molecular dynamics] simulation can yet fold a protein larger than around 120 amino acids. But most biologically important proteins are bigger. With this Blue Water allocation, we hope to **achieve the computational folding of larger proteins…that can tell some of biology's most interesting stories**."

Table 5

Blue Waters Scientific Successes: Discovered Something New

Project Title	Project Description
Simulating the Earliest Galaxies with Enzo and Blue Waters	Galaxies are complex - many physical processes operate simultaneously, and over a huge range of scales in space and time. As a result, accurately modeling the formation and evolution of galaxies over the lifetime of the universe presents tremendous technical challenges. In this talk I will describe some of the important unanswered questions regarding galaxy formation, discuss in general terms how we simulate the formation of galaxies on a computer, and present simulations (and accompanying published results) that the Enzo collaboration has recently done on the Blue Waters supercomputer. In particular, I will focus on the transition from metal-free to metal-enriched star formation in the universe, as well as the luminosity function of the earliest generations of galaxies and how we might observe it with the upcoming James Webb Space Telescope.
Policy Responses to Climate Change in a Dynamic Stochastic Economy	We extend our Integrated Assessment Model framework, called DSICE (Dynamic Stochastic Integration of Climate and the Economy), for evaluating alternative policy responses to future climate change under both economic and climate uncertainty. We incorporate five interacting climate tipping points into DSICE, and find that the present social cost of carbon (SCC) increases nearly eightfold. Moreover, passing some tipping points increases the likelihood of other tipping points occurring, so that the SCC increases abruptly. The optimal mitigation policy requires zero industrial emission after this midcentury and leads to a path of global average temperature less than 1.5 degree Celsius. We also examine the impact of Bayesian learning of uncertain critical parameters (e.g., climate sensitivity) to decision rules. Furthermore, we address the problem of decision making with uncertainty in the context of climate change that affects the world in a non-uniform manner (more warming in the polar regions).
Direct Simulation Of Dispersed Liquid Droplets	The objective of our research is to enhance understanding of the two-way interactions between liquid droplets and a turbulent flow by performing direct numerical simulations (DNS). The freely moving deformable liquid droplets are fully resolved in three spatial dimensions and time and all the scales of

Table 5

Blue Waters Scientific Successes: Discovered Something New

Project Title	Project Description
In Isotropic Turbulence	the turbulent motion are simultaneously resolved down to the smallest relevant length and time scales. Our DNS solve the unsteady 3D Navier–Stokes and continuity equations throughout the whole computational domain, including the interior of the liquid droplet. A DNS of single phase isotropic turbulence at Reynolds number 300 requires a grid of 2,0483 mesh points, and about 12 hours on 65,536 processors to cover seven large eddy turnover times. Given the requirements outlined above, Blue Waters is a necessary tool for our research.
Ensembles of Molecular Dynamics Engines for Assessing Force Fields, Conformational Change, and Free Energies of Proteins and Nucleic Acids	This project uses computer simulation methods within the AMBER to simulate the atomic motions of proteins and nucleic acids in their native environment to better understand the structure, dynamics and interactions among various bio-molecules. Preliminary work on the Blue Waters Supercomputer has demonstrated reproducible and complete sampling of the conformational distributions of RNA tetranucleotides, tetraloop RNA structures, and also the internal portion of a B-DNA helix. With complete sampling, it is now possible to reliably assess, validate, and improve the applied force fields. Preliminary simulations on Blue Waters, reducing times to solution from months-years to days-weeks, show that although it is possible to converge the conformational distributions of nucleic acids like RNA tetranucleotides and tetraloops, one does not find conformational distributions that completely agree with experiments.
Lattice QCD	Calculations of QCD must support large experimental programs in high energy and nuclear physics. QCD is a strongly coupled, nonlinear quantum field theory; lattice QCD is a first principles calculation tool that requires large scale computer power. A central goal of nuclear physics is to predict new bound states of quarks, properties of glue balls and exotic states that are not predicted by quark models.
Predicting Protein Structures with Physical Petascale Molecular Simulations	We have a new computational method (MELD) for predicting protein structures. Unlike others, it is based on physics of atomic interactions. The power of physical methods is that they help us understand how proteins perform their biological actions as catalysts, motors, pumps, transporters and transducers of energy and light. In the past, physics-based structure prediction was much too computationally costly. Our work with MELD so far has shown that we can use atomistic molecular dynamics to fold 15 out of 20 small proteins to their correct native structures, and to fold 3 small proteins with very high accuracy. No MD simulation can yet fold a protein larger than around 120 amino acids. But, most biologically important proteins are bigger. With this Blue Water allocation, we hope to achieve the computational folding of larger proteins. With Blue Waters' GPU resources, we hope to bring the power of atomistic MD to folding and mechanisms that can tell some of biology's most interesting stories.

Source: IDC 2016

Blue Waters Successes: Helped Society

The project examples described in Table 6 involved innovations centered on new discoveries in the following domains:

- Ethnology
- Biology/Medicine
- Weather

Key Quotes

"This computational project aimed to produce genomic variant calls for the design of a cost-effective genotyping chip that would **capture the genetic diversity in populations of African origin**, including African-Americans… We also demonstrated, in a production-grade project, the capability of Blue Waters to **conduct high-throughput analysis of human genomes**."

"If **whole genome sequencing and analysis** become part of the standard of care in many hospitals within the next few years, then human genetic variant calling will need to be performed on hundreds of incoming patients on any given day… We [are investigating] the kinds of computational bottlenecks that can be expected, as well as the tools and methods to overcome them."

"The historical record suggest that extreme space weather is likely to impact the Earth again in the future…We are…[calculating] location specific **space weather hazards to electric power grids**. Blue Waters permits us to account for the Earth's topography, oceans, variable composition of the lithosphere as well as variable ionospheric composition and source conditions according to time, altitude, and position around the globe."

Table 6

Blue Waters Scientific Success

Project Title	Project Description
Custom Genotyping Chip for African Populations	This computational project aimed to produce genomic variant calls for the design of a cost-effective genotyping chip that would capture the genetic diversity in populations of African origin, including African-Americans. This work will enable the identification of genetic variation specific to African populations, which will help better understand the links between genotype and disease in people of African origin, and thus extend the principles of personalized medicine to these underserved populations. It will also permit deeper study of African genetic diversity, which will bring important insights into the history and evolution of humans in general. We also demonstrated, in a production-grade project, the capability of Blue Waters to conduct high-throughput analysis of human genomes. Lots of benchmarking data were collected, and the computational workflow was hardened with many quality control steps to ensure delivery of correct results. The code is posted on GitHub, to be shared with the community.
Instrumenting Human Variant Calling Workflow on Blue Waters	High throughput Human Variant Calling Workflow on BlueWaters. If whole genome sequencing and analysis become part of the standard of care in many hospitals within the next few years, then human genetic variant calling will need to be performed on hundreds of incoming patients on any given day. At this scale, the standard workflow widely accepted in the research and medical community, will use thousands of nodes at a time and have I/O bottlenecks that could affect performance even on a major cluster like BlueWaters.
Location- Specific Space Weather Hazard to Electric Power Grids Calculated on a Global Scale	The historical record suggest that extreme space weather is likely to impact the Earth again in the future. However, modern electrotechnologies will be affected by space weather to a much larger degree than in the past. We are using a global Maxwell's equation model of the Earthionoshere waveguide to calculate location specific space weather hazards to electric power grids. Specifically, we are calculating and analyzing electromagnetic field behavior during recent geomagnetic storm in March 2015. Blue Waters permits us to account for the Earth's topography, oceans, variable composition of the lithosphere as well as variable ionospheric composition and source conditions according to time, altitude, and position around the globe. Blue Waters also allows us to calculate

Table 6

Blue Waters Scientific Success

Project Title	Project Description
	and analyze ground-level electromagnetic fields spans over time-spans of hours over microsecond time resolution (as required by algorithm).

Source: IDC 2016

BLUE WATERS SCIENTIFIC INNOVATION RANKINGS

The Innovation Metrics Used In the Study

In order to better quantify the overall impact of leadership class supercomputers for scientific innovations, IDC uses a rating system that measured both the *importance* and the *impact* of each innovation in the existing data set, plus the new Blue Waters innovations collected in this research study. The new overall innovation ranking was created based on a combination of the two complementary rankings of the innovations:

1) The IMPORTANCE this innovation compared to all other innovations in this field over the last ten years:

5. One of the top 2 to 3 innovations in the last decade
4. One of the top 5 innovations in the last decade
3. One of the top 10 innovations in the last decade
2. One of the top 25 innovations in the last decade
1. One of the top 50 innovations in the last decade

2) The IMPACT of this innovation to multiple organizations:

5. An innovation that is useful to over 10 organizations
4. An innovation that is useful to 6 to 10 organizations
3. An innovation useful to 2 to 5 organizations
2. An innovation only useful to 1 organization
1. An innovation that is recognized ONLY by experts in the field

3) Combing these measures, IDC's used these overall innovation ratings for this project:

- Class 1 innovations – One of the top 2-3 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 2 innovations -- One of the top 5 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 3 innovations – One of the top 5 innovations in a field over the last ten years PLUS useful to at least 5 organizations
- Class 4 innovations – One of the top 10 innovations in a field over the last ten years PLUS useful to at least 5 organizations
- Class 5 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at over 10 organizations
- Class 6 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 7 innovations – One of the top 50 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 8 innovations – innovations that are typically only useful to one organization and/or are minor innovations in a given research field

Blue Waters Innovation Rankings

Innovation Importance Levels

Figure 1 shows the scale of innovation importance (explained above), where level 5 is most important and level 1 least important. As the figure indicates, many of the innovations enabled by Blue Waters were in the high-importance levels 4 and 5 categories.

FIGURE 1

Blue Water Innovations: Innovation Importance Levels

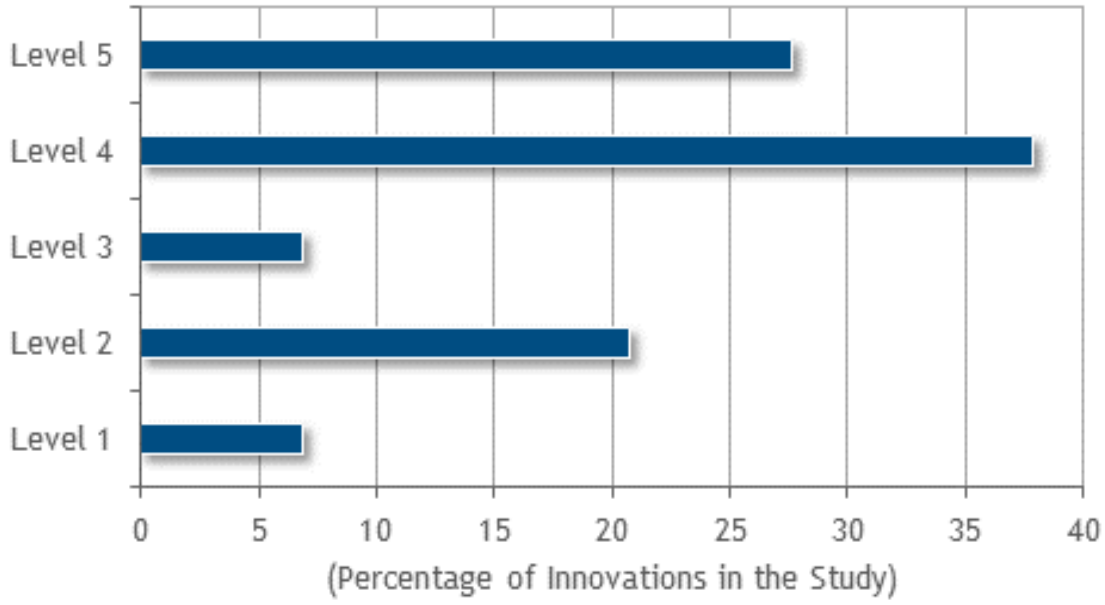

Source: IDC 2016

Innovation Importance Levels: Compared to Worldwide Averages

The profile of the worldwide sites in this study shown in Figures 2 and 3, and in Table 7.

FIGURE 2

Profile of Worldwide Participants in the Study: By Sector

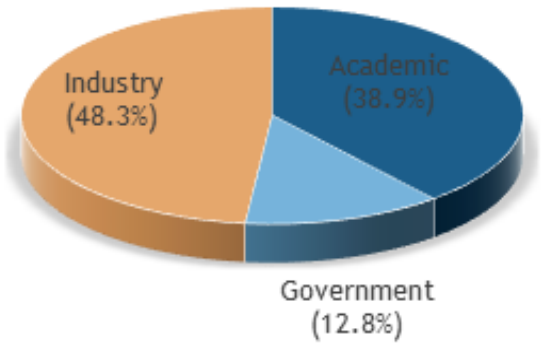

Source: IDC 2016

Figure 3 shows the survey sample by the size of the organization. For this research, IDC uses these three broad categories:

- A small organization has less than 2,000 employees or less than $250 million a year in revenues or budget size.
- A large organization has greater than 10,000 employees or over $500 million in yearly sales or budget size.
- A medium organization are those that are in between the first two groups (at least 2,000 employees or at least $250 million a year, but less than 10,000 employees or less than $500 million a year).

FIGURE 3

Profile of Worldwide Participants in the Study: By Size

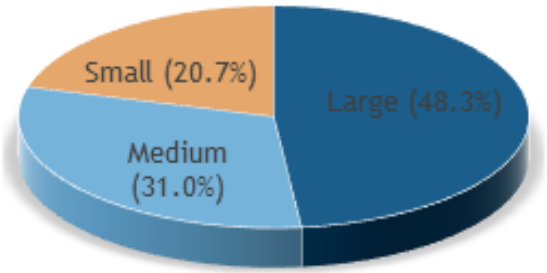

Source: IDC 2016

Table 7

Mix Profiles of Worldwide participants in the Study

Sector	Organization Size	Count of Applied Innovations	Count of Basic Innovations	Total Number of Innovations
Academic	Large	20	10	30
	Medium	4	31	35
	Small	8	6	14
Academic Total		32	47	79
Government	Large	5	5	10
	Medium	2	3	5
	Small	5	6	11
Government Total		12	14	26
Industry	Large	48	10	58
	Medium	17	6	23

Table 7

Mix Profiles of Worldwide participants in the Study

Sector	Organization Size	Count of Applied Innovations	Count of Basic Innovations	Total Number of Innovations
	Small	15	2	17
Industry Total		80	18	98
Grand Total		124	79	203

Source: IDC 2016

Figure 4 compares the importance of the scientific innovations enabled by Blue Waters with a global collection of examples IDC assembled for our DOE research studies. Compared to all worldwide sites, a significantly higher percentage of the Blue Waters-enabled innovations fell into the top two levels of importance, and far fewer Blue Waters innovations fell into the level 1 category of least importance.

Comparisons are also shown for all US innovations, and for US academic only innovations. Looking at the top two rating categories, Blue Waters is similar to other US academic sites. And way fewer Blue waters innovations are in the level 1 category.

FIGURE 4

Blue Water Innovations: Innovation Importance Levels

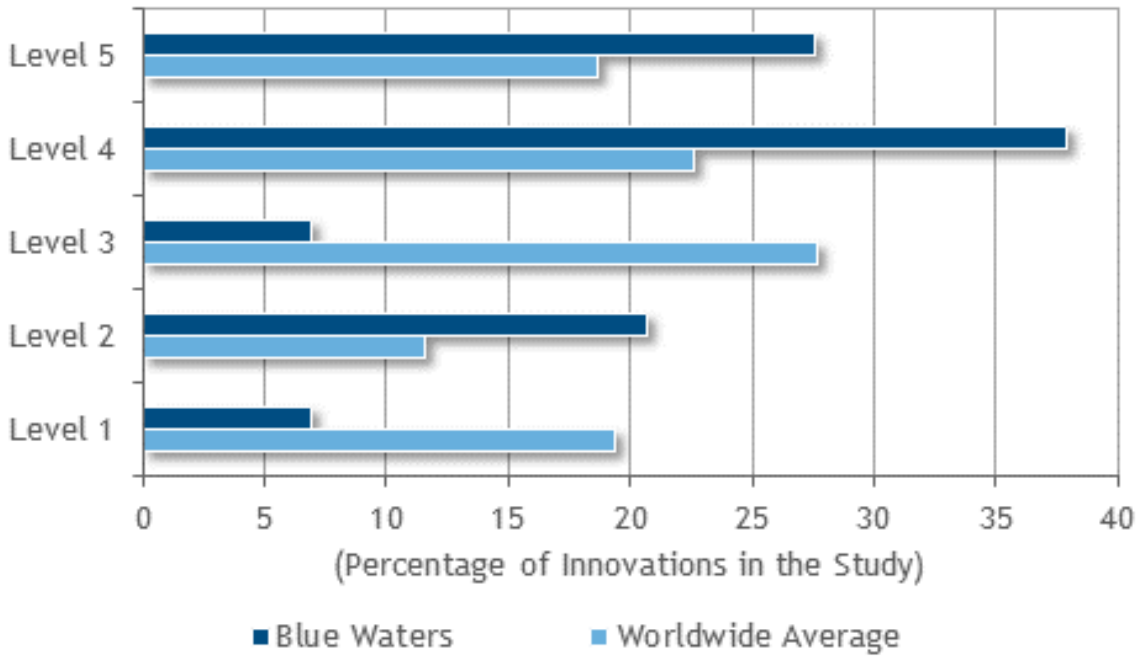

Source: IDC 2016

Innovation Impact Levels

Figure 5 displays the impact levels of the Blue Waters-enabled scientific innovations. The figure shows that more than half of the Blue Waters innovations qualified for levels 5 or 6, representing innovations having the largest impacts to multiple organizations.

FIGURE 5

Blue Water Innovations: Innovation Number of Organizations Impacted

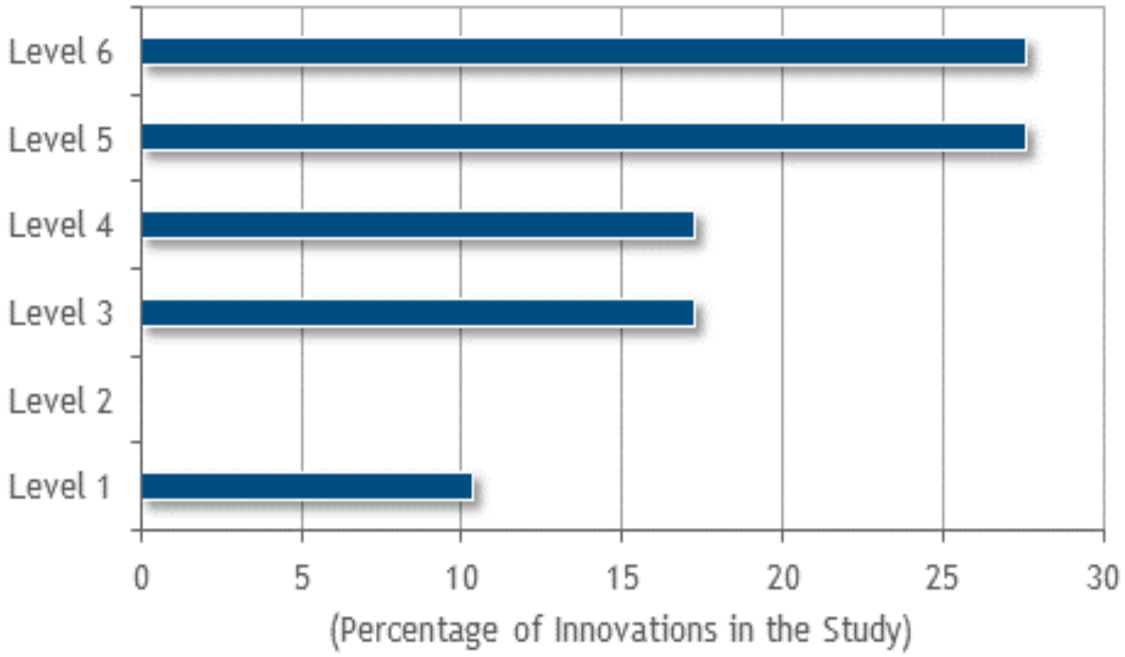

Source: IDC 2016

Innovation Impact Levels: Compared to Worldwide Averages

A notably higher percentage of scientific innovations conducted using Blue Waters qualified for impact levels 5 and 6 than was true for IDC's global data base of supercomputer-enabled scientific innovations (see Figure 6).

Comparisons are also shown for all US innovations, and for US academic only innovations. Looking at the top rating category, Blue Waters is similar to other US academic sites.

FIGURE 6

Blue Water Innovations: Innovation Number of Organizations Impacted

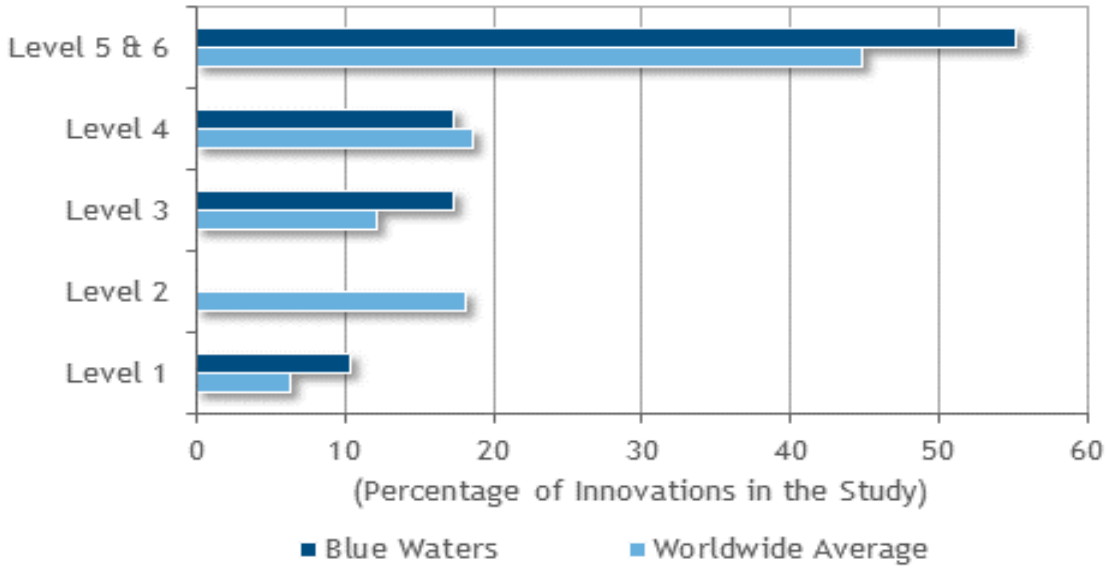

Source: IDC 2016

Innovation Overall "Class" Levels:

Figure 7 shows the rankings of Blue Waters-enabled innovations when both importance and impact are considered, using the following definitions:

- Class 1 innovations – One of the top 2-3 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 2 innovations -- One of the top 5 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 3 innovations – One of the top 5 innovations in a field over the last ten years PLUS useful to at least 5 organizations
- Class 4 innovations – One of the top 10 innovations in a field over the last ten years PLUS useful to at least 5 organizations
- Class 5 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at over 10 organizations
- Class 6 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 7 innovations – One of the top 50 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 8 innovations – innovations that are typically only useful to one organization and/or are minor innovations in a given research field

In this hybrid classification method IDC created for the present study, 62% of the Blue Waters-enabled innovations ranked in the top three classes.

FIGURE 7

Blue Water Innovations: Innovation Class Level

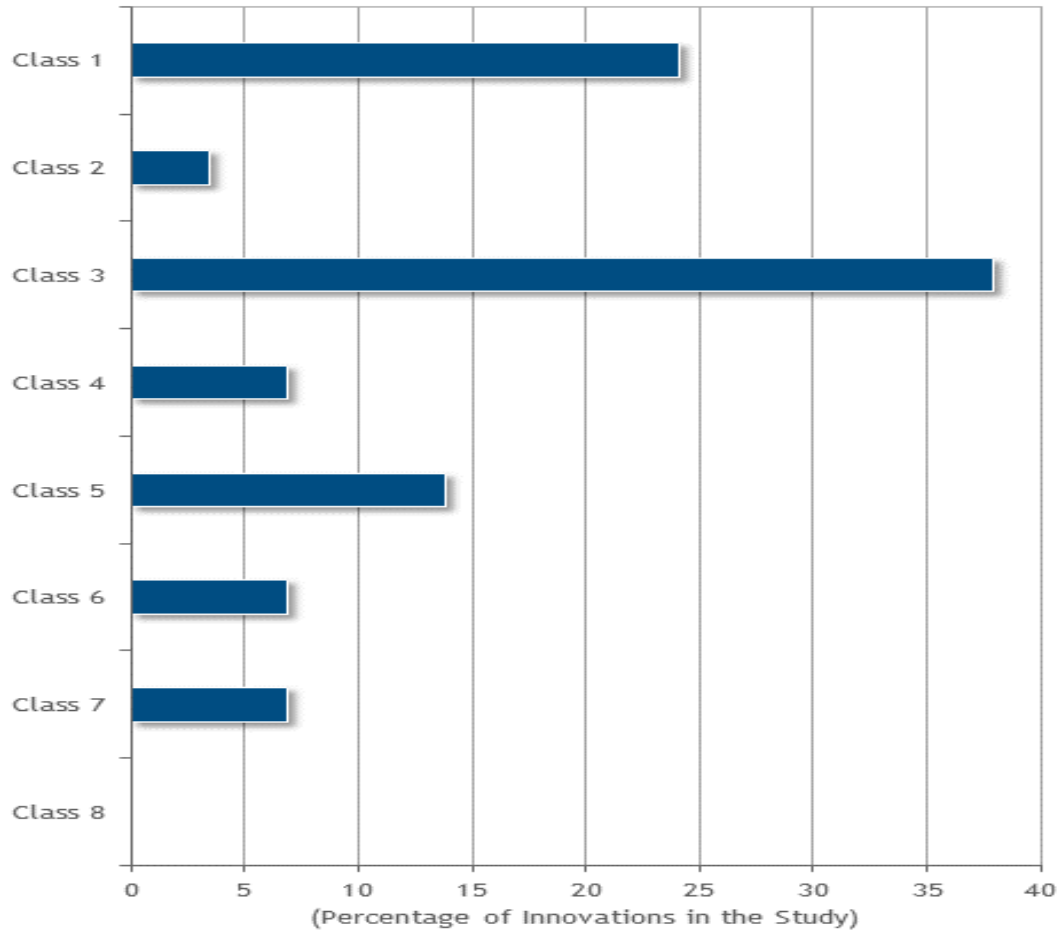

Source: IDC 2016

Innovation Overall "Class" Levels: Compared to Worldwide Averages

Figure 8 shows that a substantially higher percentage of the Blue Waters-enabled innovations qualified for Class 1 & 3 innovation levels, and substantially more Blue Waters-enabled innovations ranked in the top three classes than was true of IDC's global worldwide data base of innovations.

Comparisons are also shown for all US innovations, and for US academic only innovations. Looking at the top three innovation class rating categories, Blue Waters well ahead of other US organizations and other US academic sites. And way fewer Blue waters innovations are in the class levels 1 and 2 categories.

FIGURE 8

Blue Water Innovations: Innovation Class Level

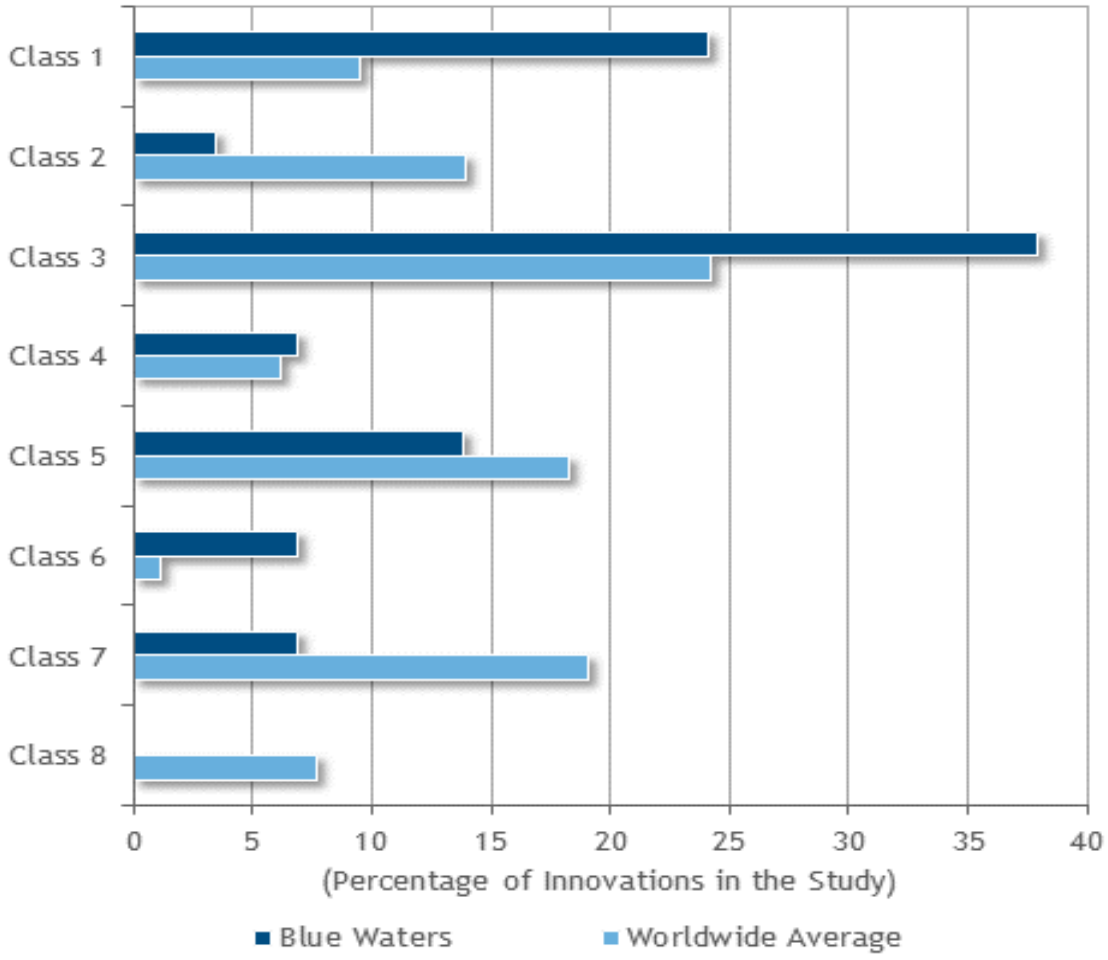

Source: IDC 2016

IDC OBSERVATIONS

The Blue Waters supercomputer at NCSA is a pioneering achievement in HPC design. And where scientific productivity is concerned (as opposed to standard benchmark tests), Blue Waters is certainly one of the top few supercomputers in the world. As noted earlier, Blue Water's usefulness and productivity is notable in large part because NCSA worked so closely with its user community to help determine what kind of leadership-class supercomputer would best support the demanding requirements for their scientific work.

This study confirms that the time and money invested in Blue Waters has been paying off handsomely. Not once during the 31 interviews IDC conducted with Blue Waters users for this study did any of them point to shortcomings of the supercomputer. On the contrary, all of the researchers had praise for the system and enthusiastically reported on the progress it has enabled for their work. IDC has seldom encountered such a diverse group of scientific researchers who all agree on something—in this case, the value of the Blue Waters supercomputer.

LEARN MORE

Related IDC Research

- Worldwide HPC Server Forecast, 2016–2020 Jun 2016 Doc #US41318216 Earl C. Joseph, Ph.D., Steve Conway, Robert Sorensen, Kevin Monroe
- NSCI Update: Rolling Up Its Sleeves and Getting to Work Mar 2016 Doc #lcUS41072416 Robert Sorensen
- EU Consortium Gathers Core Capabilities to Build Exascale HPC Prototype Feb 2016 Doc #lcUS41048816 Robert Sorensen
- Next Steps for the NSCI: Looking to Ensure a Long and Lively Life Span Jan 2016 Doc#lcUS40980816 Robert Sorensen
- IDC Study: U.S. Private Sector Cybersecurity Best Practices Jan 2016 Doc#US40688815 Steve Conway, Sean Pike
- Baidu's Deep Learning Efforts: Notable Progress on Many Fronts Dec 2015 Doc#lcUS40704215 Robert Sorensen
- Worldwide HPC Server Forecast Update, 2015-2019 Nov 2015 Doc #259950 Earl C. Joseph, Ph.D., Steve Conway, Robert Sorensen
- The U.S. National Strategic Computing Initiative as a "Moonshot": Taking Its First Small Steps Sep 2015 Doc #259288 Flash Earl C. Joseph, Ph.D., Steve Conway, Robert Sorensen
- White House Announces Strategic HPC Plan: A Good Start on a Long Road Aug 2015 Doc #258194 Flash Robert Sorensen, Earl C. Joseph, Ph.D., Steve Conway
- Global HPC Market Dynamics in 2013 Apr 2014 Doc #248137 Chirag Dekate, Ph.D., Earl C. Joseph, Ph.D., Steve Conway
- Industrial Partnership Programs and High-Performance Computing: HPC User Forum, April 7-9, 2014, Santa Fe, New Mexico Apr 2014 Doc #248113 Earl C. Joseph, Ph.D., Steve Conway, Chirag Dekate, Ph.D.
- International Perspectives on Industrial High-Performance Computing Partnerships: HPC User Forum, April 7-9, 2014, Santa Fe, New Mexico Apr 2014 Doc #248122 Steve Conway, Earl C. Joseph, Ph.D., Chirag Dekate, Ph.D.
- Catalyst Supercomputer Heralds Shift to More Balanced Architectures Nov 2013 Doc #lcUS24437513 Steve Conway, Earl C. Joseph, Ph.D.
- National and International Initiatives: HPC User Forum, September 2013, Boston, Massachusetts Oct 2013 Doc # 243776 Steve Conway, Chirag Dekate, Ph.D., Earl C. Joseph, Ph.D.
- High-Performance Data Analysis in the Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts Oct 2013 Doc # 243774 Steve Conway, Chirag Dekate, Ph.D., Earl C. Joseph, Ph.D.
- World's Fastest Supercomputer Set to Reach Customer in Oct 2013 Sep 2013 Doc #lcUS24300913 Steve Conway, Earl C. Joseph, Ph.D.
- Worldwide High-Performance Data Analysis 2013-2017 Forecast Jun 2013 Doc #241315 Steve Conway, Chirag Dekate, Ph.D., Earl C. Joseph, Ph.D.
- Top Issues for HPC Sites: HPC User Forum, April 29-May 1, 2013, Tucson, Arizona Jun 2013 Doc #241463 Chirag Dekate, Ph.D., Steve Conway, Earl C. Joseph, Ph.D.
- Livermore Lab Expands Industry Partnerships: Economic Security Is Vital for National Security Mar 2013 Doc #240232 Steve Conway, Earl C. Joseph, Ph.D., Chirag Dekate, Ph.D.
- Advanced Research at the South Ural State University Supercomputing Center Dec 2012 Doc #238225 Chirag Dekate, Ph.D., Steve Conway, Earl C. Joseph, Ph.D.

- IDC Special Study for NCSA/NSF on Industrial Use of HPC, Presented at Supercomputing 2012 Nov 2012 Doc #237984 Steve Conway
- HPC End User Site Update: RIKEN Advanced Institute for Computational Science Earl C. Joseph, Ph.D., Steve Conway, Chirag Dekate, Ph.D. Mar 2012 Doc # 233690

APPENDIX

Participants In The Survey

Table 8

Survey Participants

Investigator	Project Name
Ashok Srinvassan	Simulation-Based Policy Analysis to Reduce Ebola Transmission Risk in Air Travel
David Ceperley	Hydrogen Under Extreme Conditions
Jerry Bernholc	Petascale Quantum Simulations of Nano Systems and Biomolecules
Gerrit Botha	Custom Genotyping Chip for African Populations
John Grime	Dynamic Coarse-Grained Models for Simulations of Large-Scale Biophysical Phenomena
Liudmila Mainzer	Instrumenting Human Variant Calling Workflow on Blue Waters
Vincenzo Carnevale	Sensing the environment: a glimpse into the microscopic mechanisms of pain
Athol Kemball	Simultaneous VLBA polarimetric observations of the v={1,2} J=1-0 and v=1, J=2-1 SiO maser emission toward VY CMa II: component-level polarization analysis
Leigh Orf	Unlocking the Mysteries of the Most Violent Tornadoes
Brian O'Shea	Simulating the Earliest Galaxies with Enzo and Blue Waters
Paul Morin	ArcticDEM a White House initiative to produce a high-resolution, time-dependent elevation model of the Arctic using Blue Waters
Frank Tsung	Petascale Particle-in-Cell Simulation of Kinetic Effects in High Energy Density Plasmas
Susan Bates	High Resolution Earth System Modeling Using Blue Waters Capabilities
Tandy Warnow	Advancing Genome-Scale Phylogenomic Analysis
Kurt Brorsen	Nuclear-Electronic Orbital Calculations on Molecular Systems
PK Yeung	Particle tracking and turbulent dispersion at high Reynolds number on Blue Waters
Tom Quinn	Evolution of the small galaxy population from high redshift to the present
Nancy Makri	Quantum-classical path integral simulations of ferrocene-ferrocenium charge transfer in solution
Yongyang Cai	Policy Responses to Climate Change in a Dynamic Stochastic Economy
Gerhard Klimeck & Jim Fonseca	Advanced Nanoelectronic Device Design with Atomistic Simulations

Table 8

Survey Participants

Investigator	Project Name
Nikolai Pogorelov	Modeling Heliophysics and Astrophysics Phenomena with a Multi-Scale Fluid-Kinetic Simulation Suite
Ilias Tagkopoulos	Big Data on Small Organisms: Petascale Simulations of Data-driven, Whole-cell Microbial Models
Said Elghobashi	DIRECT SIMULATION OF DISPERSED LIQUID DROPLETS IN ISOTROPIC TURBULENCE
Seid Koric	SPARSE MATRIX FACTORIZATION IN SOLID MECHANICS AND GEOPHYSICS ON CPUS AND GPUS
Eric Lentz	Core-Collapse Supernovae Through Cosmic Time
Jamesina Simpson	Location- Specific Space Weather Hazard to Electric Power Grids Calculated on a Global Scale
Thomas Cheatham	Ensembles of Molecular Dynamics Engines for Assessing Force Fields, Conformational Change, and Free Energies of Proteins and Nucleic Acids
Robert Sugar	Lattice QCD
Ken Dill	Predicting Protein Structures with Physical Petascale Molecular Simulations

Source: IDC 2016

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.