UPDATE

Frontiers in Supercomputing: HPC User Forum, April 2016, Tucson, Arizona

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

IN THIS UPDATE

This IDC update provides near-verbatim notes from talks given on the topic of the frontiers of supercomputing at the 60th meeting of the HPC User Forum, held April 11-13, 2016, in Tucson, Arizona.

"OpenHPC: Community Building Blocks for HPC Systems": Karl Schulz (Intel)

- This is a new initiative.
- Our motivation? Many sites spend a lot of time aggregating large suites of open source projects to provide HPC environments for their users. They need to obtain HPC packages that are missing from – or don't keep pace with – Linux distros. They often leverage a mix of external and in-house tools for provisioning, software installs/upgrades, configuration, and diagnostics. Also, many projects need to engage in continual triage and debugging regarding configuration and installation issues on HPC systems.
- We don't want to tell people how to build their systems. Instead, we want to focus on common packages ("least common denominators"). We also don't want to reinvent the wheel. The goal over time is to foster multiple configuration schemes. OpenHPC 1.0 was released in November 2015, with initial starting components.

Bob Sorensen, Cybersecurity in the Commercial Sector

- Last summer (2015), a government agency asked IDC to do a sanity check: to determine how the agency rates against what's going on in the world as regards the best, typical, and worst practices. The agency also asked us what to avoid. IDC explained and showed a matrix of findings.
- The agency gave us permission to share the full matrix. [Invited audience to ask for a copy and to send comments after reviewing it.]
- We talked with a wide group of companies. [We shared highlights of the findings.]
- The government is now having IDC do a follow-up study to probe more deeply into the practices of best practitioners.

"Keynote: NSCI Update": Dr. Saul Gonzalez Martirena (NSF)

- I was the assistant director of physical sciences in the Office of Science and Technology Policy (OSTP), but am now back at NSF. I was on the team that helped develop the NSCI. I want to update you on where we are and on some of the issues we face.

- Executive order 7/29/15 was focused on "Creating a National Strategic Computing Initiative." This was a whole-of-government, multiagency approach to maximize the benefits of HPC for the United States by increasing capability and capacity and by bringing in new communities.
- How did we get here? There is a steady drumbeat of people raising the issue that we need to do better in this area. NSTC had been recommending the United States have a more cohesive strategy for long-term, basic research, and there were multiple studies from national academies and others.
- One key was to build public-private partnerships with industry and academia.
- There will be a long-term horizon for some issues, such as getting beyond Moore's law.
- HPC is not just about hardware but also about software, workforce development, and more.
- The initiative is also aimed at creating new programs.
- Four principles:
  - Deploy and apply HPC techs broadly.
  - Foster public-private collaboration.
  - Whole of government.
  - Develop comprehensive technical and scientific strategy.
- Objectives:
  - Get to exascale.
  - Coherence between simulations and analytics.
  - Convergence between datacenter computing and scientific computing.
  - Get beyond Moore's law. R&D from QC to neuromorphic and other promising technologies that will take a long time to put onto a road map.
  - Ecosystem – improve the whole HPC ecosystem, including workforce, software, etc.
  - Public-private collaboration.
- Order defines lead agencies: DOE, DoD, and NSF, plus foundational R&D agencies (IARPA, NIST) and deployment agencies (NASA, FBI, NIH, DHS, NOAA) = users that provide feedback and participate in the codesign process.
  - DOE-Science and NNSA are working on a capable exascale computing program that addresses real-world problems that are important to the government.
  - NSF will lead scientific discovery.
  - DoD will focus on data analytics problems.
  - IARPA will focus on future computing technologies, such as neuromorphic and QC.
  - NIST will focus on measurement science.
  - Deployment agencies will participate in the codesign process.
- NSCI timeline:
  - July 30, 2015: NSCI round table at the White House
  - September 15, 2015: RFI on science drivers for capable exascale issued
  - October 20–21, 2015: White House NSCI workshop
  - October 27, 2015: Executive Council delivers implementation plan
  - February 9, 2016: President's FY17 budget request to Congress includes NSCI
- Recently, the second Executive Council met.

- An RFI was issued in September:
  - What are the science drivers for exascale?
  - What specific scientific problems and barriers are to be overcome?
  - About 250 responses received: academia 94, industry 8, foreign 2, and others 5, from across many domains.
- An NSCI White House workshop (government, academia, and industry) will inform the NSCI implementation plan.
- Convergence is not simple; there are many potential paths. The key is in embracing diversity in convergence and in exploring different avenues, including recognizing that CMOS is not dead. There is still a ways to go. Cloud computing will likely be an important part of the future. Deeper engagement with the industry sector is needed, not just with HPC vendors but also with industrial end users of HPC. We need to make ROI arguments louder. It will also be important to make HPC easier to use and lower barriers to entry.
- NSCI governance comes from the Executive Council. This is led by the White House and cochaired by the director of the Office of Science and Technology Policy and the director of the Office of Management and Budget. They've met twice.
- NSCI implementation: The first internal implementation plan was delivered October 26, 2015, and informed the President's FY17 budget request. We can't discuss this publicly yet, because budget discussions are still to come. In the budget, DOE NNSA has $95 million, the Office of Science has $190 million, and the NSF has $33 million. In the Office of Science, the ASCR has $154 million, basic energy science has $26 million, and biological and environmental research has $10 million.
- Exascale at DOE is a partnership between the Office of Science and the NNSA. The entire HPC ecosystem supports DOE's missions in national security and science. Key challenges are power consumption, high parallelism, and reliability. Extreme Exascale computing can't be achieved with a business-as-usual approach.
- Exascale Computing Initiative (ECI) and Exascale Computing Project (ECP).
- NSCI at NSF is driven by the frontiers of science:
  - In 2016: community input (a series of workshops already held or scheduled)
  - In 2017: budget of $33 million for pilots and prioritizations, to see what sticks; investment areas of 2017 are low-power computing, post-Moore's law systems, algorithms and arches, workforce development, and big data (BD)
- NSCI sits in a complex of other administrative priorities, such as the Materials Genome Initiative, and the BRAIN Initiative.
- We have achieved a high level of coordination and collaboration across federal agencies. This took two years. We are in the process of doing the same with academia and the private sector.
- DOE does workforce development, as do the other agencies. NSF has a trainee program. There's a workforce development thrust on both the domain science side and the computational science side.
- HPC is key to science and economic prosperity in the United States, and we don't want to fall behind. Agencies work with international partners.

"HPC4Mfg": Peg Folta, LLNL

- Folta was introduced by Suzy Tichenor.
- They are creating an ecosystem to support U.S. manufacturing's adoption of HPC.

- They are working with LLNL, LBL, and ORNL.
- Manufacturing has a poor image ("the 4Ds"): dirty, dark, dangerous, and declining.
- U.S. manufacturing has been slow to adopt technology, but some leading manufacturers have adopted HPC and shown the benefits. The challenge is to move more of the manufacturing industry to do this.
- U.S. national labs have great resources to help. Many manufacturers are energy intensive.
- There is a big divide between national labs and industry. Much of the industry doesn't know how to use HPC or engage with national labs. Advanced Manufacturing Office's HPC manufacturing program partners with national labs and the industry. LLNL leads the program, with LBNL and ORNL partnerships:
  - Phase 1 projects (one year): HPC impact demonstration. AMO funds <$300,000 to labs; industry 20% in kind
  - Phase 2 projects (multiple years): HPC adoption. Mix of industry, consortium, and government funding
- They have streamlined public-private execution. They share what is learned and are working to build an HPC manufacturing community. The goal is to lower the barrier of entry to HPC and make it easier to try out HPC.
- Project examples:
  - **Agenda 2020 — paper manufacturing consortium.** The biggest challenge is drying the paper. Modeling how the press works that presses the water out of the paper saves hundreds of millions of dollars a year by finding ways to prevent water from seeping back into the paper after pressing.
  - **Purdue Calumet — steel manufacturing consortium.** Blast furnace understanding is the same as 50 years ago. The goal is to model different parts of what's happening in the blast furnace, but codes are not efficient. There is a need to model larger problems. LLNL is working to speed up its codes. Phase 1 (six months): 1,000x performance speedup to reduce coke consumption.

    Phase 2 will involve integrating three codes and running on an HPC system. SORAA is a lighting company needing to model GaN crystal growth. If it can do this, it can reduce production costs by 20% and have lighting for next-generation electronics.
- HPC4Mfg involves a two-phase solicitation process. They ask for a two-page concept paper, then a full proposal, and then proceed to award selection. Labs help companies with proposals because, in most cases, companies don't understand HPC. We also help seek potential partners. We link a lab's principal investigator to a submitter to codevelop their proposal. The program provides substantial feedback, even to those not getting proposals funded.
- We heard repeatedly from industry that it did not need to know who to engage with at national labs or anything about HPC.
- The program was launched March-September 2015, with seedling projects. From September 2015 to March 2016 was the inaugural solicitation, where $3 million was awarded to eight companies. They just announced the next solicitation and established summer internships at labs.
- The future will feature a twice-per-year cycle, with solicitations in both spring and fall.
- On an industry day, we'll invite them in and get feedback.
- Much focus is on building the HPC community in manufacturing, such as building internship programs.

- Typical project goals include process optimization, design improvement, or new computational tools.
- One challenge is to target industries that are not yet at the table.
- Goals include:
  - Ensure HPC becomes part of the fabric of manufacturing after the one-year period.
  - Build the ecosystem to support HPC in manufacturing, with access to cycles and talent.
  - Leverage infrastructure to other areas of DOE and other labs.

"Coupled Models for Water Resources Research": Ruth Cheng (ERDC)

- This project is traditional scientific computing related to the ERDC's mission in water resources.
- The Coastal and Hydraulics Laboratory is one of seven ERDC labs.
- ERDC was established after the Great Mississippi Flood of 1927.
- This project is related to storm surges — collecting relevant data (CHL) such as wind, waves, water levels, and currents. We currently have 34 years of data. ERDC collects some of this data.
- The coastal storm modeling system (CSTORM-MS) uses tight, two-way coupling to compute surge waves. The two models, ADCIRC/ADH and STWAVE, can exchange data. The model can be expanded to include more applications.
- The demonstration study was the North Atlantic Coast Comprehensive Study. From Maine to Virginia, 1,150 storm scenarios were created for 3,500 high-fidelity simulations. Each used about 2,400 CPUs for 11 hours, on average, with 100 million CPU hours used in total. Both systems used were from Cray:
  - Total 250TB of disk space
  - Resolutions 10m–100km
  - 3.1 million nodes
  - Time-step size: 1 second
- The CSTORM Production Systems make use of standard Linux/Unix tools.
- We look forward to coupling more applications in parallel in the future. This is big data research.

"Fostering Innovation at AFRL": Lloyd Slonaker (AFRL)

- Our work supports a dual mission: DoD HPC (their scientists and engineers) and AFRL. We collaborate with partners to help address national security challenges.
- We serve DoD users and contractors, and the time on our systems is free to these users. We're fully funded up front. AFRL is one of five DoD HPC centers. Defense Research and Engineering Network (DREN) is a major asset connecting the centers, but it's also available to researchers at various universities. We just went through a 100Gbps upgrade. Also, part of the program in the past few years has been working on software development. We offer consulting, system hosting, access to DoD HPC, hosting and administration, visualization support, applications support, and more.
- Missing middle: Six years ago, we did a survey and found many DoD users didn't even know we existed. Nontraditional users weren't being served, so we started an outreach initiative to entry-level users, to help them access and utilize HPC. We're building relationships with more organizations within DoD. It can take a while to get security clearance, training, etc.

- The Innovation Laboratory is mainly for staff to try new things, bring in vendors, etc. Users weren't allowed in that environment. We decided we had to let users in to try things out too. The lab includes systems, software, and expert support. The lab has three fairly large systems that are somewhat older. We set them up as non-allocated systems, so projects would not need to go through the bureaucracy and red tape. We encouraged users to bring in new ideas. The goal is to grow HPC use by nurturing new projects, new uses, and new capabilities for DoD. New resources include new databases, data analytics, Hadoop, scheduling software, workflows (this is a major development), and virtual machine environments.
- The Advanced Framework for Simulation, Integration and Modeling (AFSIM) enables the use of HPC more for modeling and simulation, rather than workstations. Bringing in the users is relatively new to us. We're especially trying to bring in that missing middle. It also energizes our staff to do something a little different like this.

"Leveraging HPC for Alzheimer's Research and Beyond at The National Supercomputing Center for Energy and the Environment": Joseph Lombardo (NSCEE)

- This is one of our most important projects. Working with private industry has really worked well for us. We were established in 1991, but NSCEE namedropped and we are the National Supercomputing Institute. The Alzheimer's project involves an $11.5 million NIH grant in partnership with Cleveland Clinic. At the start, it involved clinical trials not needing HPC but moved into genomic sequencing, thereby becoming an HPC problem.
- Altair and Intel are two of our industry sponsors. One of our feet is on the federal side, the other is on the academic side. We serve the University of Nevada system and still do energy research.
- One partner is Switch, a company that builds some of the largest datacenters in the world. We're hosted on the Cherry Creek system, a large Intel system for scientific and economic R&D with 26,000 cores. We maintain a dedicated research network with 100Gbps potential.
- We're trying to evaluate our colocation situation. There's a scientific center at our new science and engineering building at UNLV, but it could not hold our large 60-cabinet Cray supercomputer, so we're at Switch. During the next two years, we'll decide whether colocation is right for us.
- So far, as datacenter director, I really like it. There's been no downtime so far; other than that, it's 5mi from our university location – so we built a 20Gbps connection between the sites. Cisco decided to help out and donated gear, so we now have 2 x 100Gbps connections. I'm trying to move everything I can out to the Switch location. To get from 16,000 to 26,000 cores cost $545,000, but Switch decided to pay for it. There is talk about upgrading to Knights Landing. They really liked our project, collaborating with Cleveland Clinic on neural problems.
- The new Institute for Personalized Medicine houses the Alzheimer's project.
- We struggled with cluster management. Adding Altair's PBS Pro made it all work well and simplified administration. Other projects include:
  - Quantum mechanics of chemical reactions
  - Fracking (in collaboration with LANL)
  - Magnetically dominated jets in gamma ray burst

"Graph Analysis Using GPUs, with Real-World Examples at the British Museum": Brad Bebee (Blazegraph)

- Bebee was introduced by Arno Kolster, PayPal. Bebee is CEO of Blazegraph.
- We were originally an HPC software company focused on graphs. Graphs show up in many types of applications – any time you need to represent complex dependencies in your data – but as they get large, they challenge your computational platform.
- We've been around about 10 years. We are open source, so we get Fortune 500 customers as researchers, and from FSI and life science, three more.
- We've seen a lot of growth in precision medicine (PM). A company in San Francisco, Synapse, is pioneering graphing for precision medicine.
- Healthcare leaders are moving into PM for competitive reasons, etc., but implementation is difficult. Medical records don't really support PM. The challenge is when a physician wants to use genomic testing and often gets a fax, a piece of paper. The question is how you leverage data for precision medicine.
- There are 5,685 hospitals in the United States. The top 10% have 100+ billion EdgeGraph problems (256 K80 GPUs for a 100 billion EdgeGraph integration).
- The second major use case is knowledge graphs: getting the data you want versus just getting back a bunch of text.
- Chemical company use case (material science): combination of public data and proprietary data uses graph to integrate. Integrate 1 billion to 2 billion edges proprietary + 3 billion to 4 billion edges public.
- Large-scale law enforcement application with 40 billion edges. Need to scale to 1 trillion edges.
- Graphs are like other BD challenges, having nonlocality. CPU cores will wait while the graph is being traversed – from CPU cache to main memory results in a fivefold penalty in performance.
- Blazegraph DASL provides a high-level way to author graph analytics and integrate with Apache Spark and Apache Hadoop.
- We've developed the ability for GPU acceleration without code changes, mostly running on Titan, but also working with DARPA. With familiar SPARQL queries, 200–300x acceleration is possible without code changes. [Showed example CPU 54s, GPU 200ms.]
- Did recent work for a defense organization: GPUs were 700–1,800x faster for graphs than baseline Spark CPU configuration.
- Need simpler, smarter algorithms that expose the power of the hardware platforms but that don't require explicit programming.
- DASL, a Blazegraph product: a domain-specific language for graph and machine learning algorithms.
- We've seen a lot of blending between graph analytics and machine learning analytics, for example, people do graph, then add ML, and then do more graphing.
- ResearchSpace is part of the Mellon Foundation's goal of using graphs to analyze cultural heritages, in relation to not only the British Museum but also for the British Museum DBpedia, other British Museum data, and more. Only a small fraction of the British Museum collection can be physically displayed at any one time, and graphing (cataloguing) of the entire collection gives researchers access to the entire collection at all times. You can save graph queries for your research and collaborate with other researchers.

- Cyber is the biggest growth space we're seeing outside of bioscience. This creates very large graph challenges, for which we use the DASL tool.
- It costs about $140,000 for 1GTEP of performance on Urika, as opposed to $14,000 on DASL, according to Bebee. We think the right way forward is to combine a large shared memory system with GPUs. We're talking with SGI about doing this with UVs on a very large implementation.

"NSCI Discussion": Bob Sorensen (IDC)

- Sorensen briefly recapped Martirena's morning talk.
- From my perspective, the R&D aspects of NSCI are the ones we're most used to and that government is least used to. The first big challenge is funding. The probability of the FY17 budget being funded this year is low, especially with a departing president.
- The initiative needs a quick success. There are lots of compelling national priorities ahead of NSCI. NSCI also needs to address a fundamental dichotomy: Ensure U.S. government agencies have access to the HPC system they need while doing the same for the private sector.
- Another challenge is that everyone talks about workforce development, but no one does much about it. HPC needs to be made an interesting, compelling field to be in.
- The U.S. government is not very experienced in fostering economic competitiveness. In my opinion, the government needs to learn how to foster an entrepreneurial spirit.
- Audience comments:
  - It's all B.S. Government is not going to influence which chips get built or what the capabilities of these chips are. Vendors don't want the government telling them how to design their products and attaching strings to contracts. The government has tried for years to pick winners and losers and this hasn't succeeded. The government should leave the market alone and buy systems, not try to force technology directions and development.
  - I don't think many companies will advance in using HPC until we solve the ISV problem. Unless ISVs see larger-scale HPC as a lucrative market, they will not have the incentive to scale their software. Companies that depend on ISV software will continue to be stymied in using HPC.
  - Education and workforce development are really key, but it's not limited to computational scientists. Executives also need to know enough about HPC to make investment decisions. The government has been poor about handing out money and incentivizing it with metrics. And there has to be some sort of software investment. And it has to be focused on sustained performance, not peak performance such as exaflop. Exaflop for what should be the question.
  - Most HPC centers offer training and would welcome government funding to support and expand this.
  - DOE has no marketing budget to alert companies to the value of HPC and how to access it. That's a barrier.
  - ISVs won't invest in porting their applications to different HPC platforms unless they see a possibility of a good return. The other item is that if you're a small manufacturing or engineering company, will you take one of your engineers and send them off to a training program?
  - I'm coming from industry. I suspect right now there are 40 different visions of what NSCI success looks like. NSCI needs to paint a vision of what success looks like in light of the initiative's objectives. Until you do this, everyone will continue going down a different road to a different destination.

- The government seems to be guessing what direction the market will take, but vendors have this mapped out for 5–7 years ahead already and will resist changing course. ISVs might modernize their code once, but they won't do it every 2–3 years. A modernization has to last 20–30 years. I don't think sending someone for a two- to three-day course will help.

LEARN MORE

Related Research

- *Worldwide HPC Server Forecast, 2016–2020* (IDC #US41318216, June 2016)
- *NSCI Update: Rolling Up Its Sleeves and Getting to Work* (IDC #lcUS41072416, March 2016)
- *EU Consortium Gathers Core Capabilities to Build Exascale HPC Prototype* (IDC #lcUS41048816, February 2016)
- *Next Steps for the NSCI: Looking to Ensure a Long and Lively Life Span* (IDC #lcUS40980816, January 2016)
- *IDC Study: U.S. Private Sector Cybersecurity Best Practices* (IDC #US40688815, January 2016)
- *Baidu's Deep Learning Efforts: Notable Progress on Many Fronts* (IDC #lcUS40704215, December 2015)
- *Worldwide HPC Server Forecast Update, 2015–2019* (IDC #259950, November 2015)
- *The U.S. National Strategic Computing Initiative as a "Moonshot": Taking Its First Small Steps* (IDC #259288, September 2015)
- *White House Announces Strategic HPC Plan: A Good Start on a Long Road* (IDC #258194, August 2015)
- *Global HPC Market Dynamics in 2013* (IDC #248137, April 2014)
- *Industrial Partnership Programs and High-Performance Computing: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico* (IDC #248113, April 2014)
- *International Perspectives on Industrial High-Performance Computing Partnerships: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico* (IDC #248122, April 2014)
- *Catalyst Supercomputer Heralds Shift to More Balanced Architectures* (IDC #lcUS24437513, November 2013)

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.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com