Update

Technologies and Applications of High-Performance Computing in the Cloud: HPC User Forum, September 15–17, 2014, Seattle, Washington

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

IN THIS UPDATE

This IDC update captures part of the proceedings at the 54th HPC User Forum held in Seattle, Washington, that looked at HPCs in the cloud. Cloud computing – either through public, private, or hybrid platforms – is seen by many as a viable and cost-efficient alternative to traditional HPC ownership and operation. Jim Kasdorf, the outgoing chairman of the HPC User Forum moderated this session, which had commercial presenters summarizing their perspectives on HPC technologies and applications on cloud computing platforms. The speakers at this HPC cloud session included Tim Carroll, Cycle Computing, on clouds as an HPC alternative; David Pellerin, AWS, on HPC clouds and their use cases; Charlie Gonzales, IBM, on cloud customer requirements; Steve Legensky, Intelligent Light, on HPC cloud software requirements; and Alex Sutton, Microsoft, on growing HPC workloads in the cloud.

Cloud computing, which is certainly a growing trend in the commercial space, offers strong potential for a wide range of HPC application as well. However, it is unclear yet what the true technical and financial costs may be in shifting a critical HPC application to a cloud environment. IDC assesses that although there are many successful examples of HPC-based cloud applications, the technology is still relatively new and requires users to consider carefully the costs and benefits of moving their HPC workload into the cloud.

New Developments in HPC Cloud Computing — Tim Carroll, Cycle Computing

"Cloud – really?" should have been the actual title here. I'm an HPC guy because I love this industry. I got into HPC after being in a Web hosting start-up and learning about scale-out servers. What can we do to get more compute to more people? Tom Kendall of the Army Research Lab was looking early on at Linux clusters. Later I went to Platform and then Dell, where I worked with customers like Mark Seager to build the largest-ever test bed cluster. Some customers said, I have this much space, this much power, and this much money; what can I get? Then came the TACC Stampede machine, which was about being available to more researchers at scale than any other machine.

Building HPC systems today is nerve-wracking. Do you build a 100,000-person stadium for five events per year?

Aerospace Corp. uses autoscaling to provide on-demand access on GovCloud for cloud-friendly jobs.

Novartis is testing three promising compounds using computational chemistry codes. The company was using 7,600 cores and needed 10,000 cores to do the work. The company did it on the cloud.

What's next – think 2016. Setting up an account in 2014 is easier than scrounging servers and NICs to create clusters in 2003. Cloud is about collaboration, not just compute. In future years, most people will have a mix of on-premise systems and codes that run on the cloud.

AWS Cloud for HPC and Big Data — David Pellerin, AWS

We're seeing use cases in HPC that make them efficient for cloud. I also want to talk about why HPC matters so much to AWS.

Our GovCloud is unique because it supports ITAR. We have lots of capacity available for HPC customers. Last year, Cycle ran a 156,000-core job on AWS.

We have capabilities such as autoscaling and managed Hadoop and more services that apply to HPC. All of this is built on the Elastic Compute Cloud, and we're adding new hardware and rack types all the time.

Users can use credit card for hours on demand. For longer-term use, such as an HPC cluster running for a bank longer term, users can negotiate contract (reserved) pricing. Spot pricing allows users to bid for unused capacity and be charged at spot price that varies with supply and demand.

Customers are looking for scalability, global collaboration (including letting the data live in the cloud and let people compute on it there), and big data.

HGST is a manufacturer using AWS for HPC, including MD for simulating their disk drives.

There are myriad HPC deployment methods:

- Traditional HPC schedulers
- Born in the cloud tools
- AWS tools and APIs

Use Cases

Cypress Semiconductor uses touch-sensor modeling for its ASICs that sit behind touchscreens. The company needed to run a very large number of simulations. The IT department wouldn't give them priority, but with AWS, IT ran the simulations in parallel.

Reservoir simulation on AWS.

Clinical trials such as those at Bristol-Myers Squibb.

GE manufacturing uses clouds to improve yield and help design around manufacturing. Users can spin up their own environments. Clouds also support global collaboration and bring the users to the data, so you don't have to send the data to the users.

Calgary Scientific has a tool for the medical imaging world that is an FDA-compliant iPad app for medical imaging supporting global accessibility.

Amazon AppStream uses GPUs in the cloud for app developers, including remote visualization.

AWS has been doing big data plus cloud for a long time, such as for Web log analysis.

FINRA audits financial institutions, and soon this capability will be entirely on AWS. FINRA must analyze 30 billion market events each day.

DNAnexus is working with AWS and Baylor School of Medicine.

The Climate Corporation (a division of Monsanto) uses AWS to help it inform crop insurance pricing, water availability, and more.

Although AWS can't see what's running inside the cloud, the company is confident HPC running on AWS "is a lot."

Predictable Performance and Secured HPC Cloud Services Through IBM Platform Computing — Charlie Gonzales, IBM

Platform's been around 20 years, and the company's flagship products are schedulers. The first was LSF, and then Symphony, which is important in FSI. It supports small clusters to global grids and clouds. It was acquired two years ago by IBM and it is bringing on first customers now.

Customer challenges:

- Spikes in demand
- Twitter – 1 million rows of data per second
- For most sites – it's easy to project capacity and determine whether to buy more hardware or open up an additional or expanded datacenter to send overload work to the cloud
- Limited staff or capex budget
- Need to focus on priority projects
- Deploying HPC apps in the cloud can be harder than it seems
- Performance – can be hard to guarantee service levels
- Fast networks and specialized hardware
- Advanced data management

Benefits include a continual technology refresh, onsite resources, and performance that remains the same for the life of the asset.

Overall, there are privacy and security concerns.

Although a lack of cost transparency has been a problem for some clouds, we address this with LSF.

Revolutionizing Innovation with On-Demand HPC — Steve Legensky, Intelligent Light

We're a software provider that was established in 1984. FieldView is our CFD post-processing suite of tools. We solve problems for customers by collaborating closely with them. Our product started in 1990 for visualization and post-processing on the desktop. But as compute got farther away, we added a true client/server. In 2005, we started working to use the HPC resource to extract the most important part of the data that might be just 3% or so of all the data. Since 2013, we've focused on in situ extraction of this important data to completely avoid writing out the large files.

CFD solver coders are well suited to the cloud, but results files can be large and numerous.

Use Cases

Zipp Speed Weaponry

Advanced bicycle racing wheels cost $1,800 each. Zipp needed reliable simulation results and had to deal with unsteady effects with fast turnaround. The company turned to Intelligent Light and we formed a team with R Systems and CP-adapco and Dell. Our XDB files are 43x smaller than CFD files and have same numerical fidelity. This resulted in a new Zipp wheel that had 100% category revenue growth and added 120 new U.S. jobs.

Daimler (Mercedes-Benz)

Primary focus was on will the car catch on file. We worked with Daimler for vehicle thermal analysis. The company is in Germany, we're in New Jersey, and there are partners elsewhere, all involving multi-terabyte files. Our solution was to partition the big data files into XDBs remotely, and each part of the car is broken into a small file. The modeling went from 30GB/time step to 4GB/time step, and the work was done in 10 days versus the earlier 3 months.

DoD CREATE Program

DoD CREATE program was to certify aircraft handling with various external payloads attached, which required many time steps of unsteady data. Work was computed at Maui HPC Center with post-processing done at Pax River, Maryland.

We don't charge for software licensing on server side (to encourage scaling); we charge only on the client side.

So companies that have not used HPC or CFD can get started on the cloud with the resources they have.

Microsoft Azure for HPC — Alex Sutton, Microsoft

My group is responsible for driving cloud-based solutions across Microsoft. We've seen large-scale batch processing and HPC as the next major cloud workloads. Our mission is to radically simplify HPC in the cloud; so users can run jobs, not manage clusters. We use the same tools for public and private clouds.

We're also investing in RDMA for HPC workloads. Many people don't know Linux, not just Windows, runs well on Azure.

LEARN MORE

Related Research

Additional research from IDC in the technical computing hardware program includes the following documents:

- *Worldwide Broader HPC 2014-2018 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #248835, June 2014)
- *When Massive Data Never Becomes Big Data* (IDC #lcUS24922014, June 2014)
- *Worldwide Technical Computing Server 2014-2018 Forecast* (IDC #248779, May 2014)
- *Perspectives on High-Performance Data Analysis: The Life Sciences* (IDC #248348, May 2014)
- *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)
- *Disruptive Technologies in High-Performance Computing: HPC User Forum, April 7-9, 2014, Santa Fe, New Mexico* (IDC #248112, April 2014)
- *Advances in Processors, Coprocessors, and Accelerators in High-Performance Computing: HPC User Forum, April 7-9, 2014, Santa Fe, New Mexico* (IDC #248111, 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)
- *Worldwide HPC Public Cloud Computing 2014-2017 Forecast* (IDC #247846, April 2014)
- *Summary of IDC's 2014 Research in the Use of HPC by Oil and Gas Organizations* (IDC #247704, March 2014)
- *IBM Sale to Lenovo Opens Opportunity for Other HPC Vendors* (IDC #lcUS24694314, February 2014)
- *IDC's Worldwide High-Performance Computing Predictions 2014* (IDC #WC20140211, February 2014)

- *Seagate Looking for the X Factor in Its Acquisition of Xyratex* (IDC #IcUS24555413, December 2013)
- *Micron Demonstrates Technologies to Address Emerging Challenges in Big Data Applications* (IDC #244843, December 2013)
- *Market Analysis Perspective: Worldwide HPC, 2013 – Directions, Trends, and Customer Requirements* (IDC #244742, December 2013)
- *HPDA Pulse: 2013 Software and Consulting Market Analysis* (IDC #244513, November 2013)
- *HPDA Pulse Results: 2013 Hardware and Storage Market Analysis* (IDC #244493, November 2013)
- *HP FY13: Revenue Declines Abate on Stronger Core Business* (IDC #IcUS24466413, November 2013)
- *Catalyst Supercomputer Heralds Shift to More Balanced Architectures* (IDC #IcUS24437513, November 2013)
- *China Eyes 10,000-Fold Data Reduction for Internet of Things* (IDC #IcUS24392513, October 2013)
- *HPC User Forum, October 2013, Seoul, Korea* (IDC #243786, October 2013)
- *Tools and Techniques for Technical Computing in Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243778, October 2013)
- *Perspectives on Quantum Computing: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243777, October 2013)
- *National and International Initiatives: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243776, October 2013)
- *Issues in High-Performance Computing: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243775, October 2013)
- *High-Performance Data Analysis in the Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243774, October 2013)
- *Chinese Research in Processor Designs for High-Performance Computing and Other Uses* (IDC #243502, October 2013)
- *World's Fastest Supercomputer Set to Reach Customer in October 2013* (IDC #IcUS24300913, September 2013)
- *The Broader HPC Market 2012-2017 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #242742, August 2013)
- *IDC's Worldwide Technical Server Taxonomy, 2013* (IDC #242725, August 2013)
- *China Regains Top Supercomputer Title* (IDC #IcUS24190613, June 2013)
- *10 Things CIOs Should Know About High-Performance Computing* (IDC #241565, June 2013)
- *Worldwide High-Performance Data Analysis 2013-2017 Forecast* (IDC #241315, June 2013)
- *Top Issues for HPC Sites: HPC User Forum, April 29-May 1, 2013, Tucson, Arizona* (IDC #241463, June 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-insights-community.com
www.idc.com