UPDATE

Cloud-Based Computing: HPC User Forum, April 2016, Tucson, Arizona

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

IN THIS UPDATE

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

Session chair: Paul Muzio, City University of New York

"Jetstream: Adding Cloud-Based Computing to the National Cyberinfrastructure": Matt Vaughn (TACC)

- I'm director of life science computing at TACC.
- Jetstream is funded by NSF.
- Jetstream is a resource to expand the community of users, which benefit from NSF investment in shared cyberinfrastructure.
- It's a production cloud, not just a research platform, with high uptime and high availability for on-demand and interactive computing and configurable architectures.
- Our focus is on ease of use and flexibility for advanced use cases.
- About 300,000 researchers and educators were supported by NSF in 2012-2013.
- Less than 3% have an XSEDE portal account.
- 70% claim to be resource constrained.
- HPC systems are fivefold to sevenfold over-requested; resources are often not well matched to users' needs, and access is not easy.
- Jetstream is aimed especially at the long tail of science — smaller groups and organizations with legitimate science needs.
- There is a need for between four and a few hundred cores at present, but not forever, so no capital expenditure is justified.
- We want to fully customize the OS and configuration.
- There are also cloud-native workflows, with a goal of moving to an academic environment and a need for interactive mode access to analysis.
- Also for people evaluating/experimenting with software not supported on XD systems, plus STEM educators.
- We expect users from diverse work domains, including those with goals for workforce development: HPCS, MSIs, tribal colleges, and higher education institutions in EPSCoR states.

- Platform overview:
  - The average user comes into a web application, or via an API, or directly into the OpenStack APIs.
  - Each production environment at TACC and Indiana University (regions) features 320 nodes, 7,680 cores, 40TB RAM, 640TB local disk, and 960TB storage – plus a U of AZ development environment. VM host configuration plus Ceph storage is present. Each user gets up to 500GB total storage.
  - Jetstream integrates with XSEDE. The user gets a dashboard showing resources in use, past usage, and more. Jetstream supports full programmatic access.
- What comes next?
  - Early operations mode is under way. By the end of March, we anticipate 38 XSEDE projects and 250+ users and expect to open to full operations by the end of May. We are soliciting research allocations now.
  - For organizations using storage on NCSA, Jetstream is working on making it easier to move data to Jetstream storage.

"Use Cases and Best Practices for Cloud-Based Enterprise HPC": David Pellerin (Amazon Web Services)

- The cloud, to me, is a global-scale, reconfigurable system. Cloud is where big compute meets big data (BD). Increasingly, the commercial sector is looking toward the scientific community to see how big data and big compute come together.
- Discovery, analysis, and visualization is simulation driven and data driven.
- The Large Hadron Collider (LHC) has a BD and a big compute problem. When the LHC runs, it spits out big data, and big compute is needed to analyze the data. A lot of this is now done in the cloud. Fermilab used 60,000 cores in January 2016 for one experiment. The commercial sector has the same problem.
- HGST, a disk drive manufacturer, works with Cycle Computing on CAD, CFD, and EDA. The work with big data for manufacturing yields analysis. These simulations now scale to 85,000 cores. HGST is also looking at what the scientific community is doing with BD simulations.
- Novartis is utilizing the cloud for better medicine. Using 87,000 cores at peak, Novartis is able to crunch 39 years of complex chemistry into seven hours. Those are all EP problems, but we're increasingly seeing legacy codes like aerospace codes being run on AWS. Amazon itself needs to run CFD in the cloud. Trek does this for parametric modeling for bicycle design.
- Clusters are tightly coupled, latency-sensitive apps. For example, EC2Grid HPC: loosely coupled, EP EC2, AZsGrids of clusters. They use grid strategy to run a group of parallel clustered HPC jobs.
- Scale matters.
- The cloud comprises multiple datacenters. Today, we have five global regions, with more coming. We add geographic locations because customers need the data to be in their region, or because we need to control latency.
- In any region, there are multiple availability zones (AZs), or clusters, within an urban region. Each AZ could be multiple datacenters, some of them very large. There are 30+ AZs worldwide today.
- Within an AZ, you can provision a virtual private cloud or otherwise configure for your need. In some instances, we configure to be compute centric; for others, we may be storage centric.

- Virtual private cloud = VPC. You can use these instances to create a VPC.
- New instance type: X1 up to 128 vCPUs, up to 2TB RAM. Great for Spark.
- Our network is unique to Amazon. We built it as a very robust software-defined network, which is important for verticals like FSI.
- Cloud formation cluster (CfnCluster): To build and use HPC clusters on AWS.
- AWS EC@ consumption models:
  - On-demand: Pay predictable price per hour
  - Reserved: Low one-time payment for significant discount
  - Spot: Bid for unused capacity, charged at a spot price based on fluctuating demand for time-insensitive or transient workloads (HPC users love spot. Best practices for spot for HPC: stateless, fault tolerance, multi-AZ, loosely coupled, and instance flexibility.)
  - Bidding strategies important for spot (Look at the history of that instance to see if your job is likely to be terminated while running. Some users don't care about the cause and just restart the job. Hedge funds doing liquidity testing, and scientific grid computing jobs, are like this.)
- Altair, ANSYS, and others are rapidly adopting cloud for anytime/anywhere web-based access.
- We've seen customers use more than 100,000 cores. Cycle used 156,000 cores on a run two years ago. GPUs are in very high demand in the cloud for CUDA programming and math solvers in finance and engineering, as well as graphics jobs.
- Memory on node: We decide based on customer use cases, such as in-memory analytics needing a lot of memory per node. We have a proprietary network right down to the silicon, rather than Ethernet or IB. We don't have IB-class latency today.

"Trends in HPC Cloud Computing": Tejas Karmarkar (Microsoft Azure)

- What do users want?
  - Faster product development
  - On-demand compute
  - Seamless integration and workflow
  - Higher fidelity methods
  - Help managing data
  - Elastic capacity when needed
  - Control over compute resources
  - Pay for use
  - No gatekeepers
  - Mobility
- We just ran a CFD job with 5 million elements.
- Our customer, Hendrick Motorsports, ran Star-CCM+ on 224,000 cores. Companies of this size cannot afford to buy this many clusters.
- Cloud providers such as Azure, AWS, and Google achieve true hyperscale. Prices have come way down for utility computing. Users want to move to higher-level services. The challenges are connectivity, intranode latency (to run MPI jobs at scale), and workflow automation.
- We try to localize our networks for the region they're in.

- Commercial ISVs are an important part of the ecosystem. ISVs have been slower to adopt a flexible usage licensing model and cloud computing. No license technology standardization exists – every license server is set up differently.
- The challenges for ISVs are protecting IP, changes to the business model, and legacy codes.
- We work closely with Schlumberger for reservoir simulation as a service.
- Microsoft Azure has seen 500+ features releases in the past year. It sees more than 90,000 new Azure customer subscriptions per month. There are 1.5 million SQL databases running on Azure and 90 trillion objects stored on Azure.
- If your company uses Office 365, there's a connection to Azure.
- Azure Batch is like HPC as a service. We've done testing of an insurance model on 200,000 cores.
- One of the largest European banks uses Azure for compliance management and reporting.
- We call it "no compromise HPC":
  - Low latency
  - Close-to-bare-metal performance
  - Hyperscale
  - Resilient parallel file system
  - High-frequency processors
  - Large fat memory nodes
  - SSD-based local scratch
  - Linux and Windows
  - Provision thousands of cores in minutes
  - Virtualized RDMA – we work very closely with Intel on this
- Azure:
  - ~2.5-3.1 microsecond latency
  - 90% LINPACK efficiency
  - Fast interconnect: QDR IB for intradeployment traffic and 10GbE for standard Azure traffic and network access (soon going to 40GbE)
  - Scratch storage 2GB per node
  - Available in seven regions
  - Visualization VMs with GPU

"Managing Large Production Cloud Deployments": Jason Stowe (Cycle Computing)

- My goal is to give you insights into the benefits of complementing internal hardware with external hardware. Cycle predates the rise of cloud computing. We believe increased access to compute creates new science. Most of our customers spend $10,000-$250,000 per month for real production usage. Most use 100-6,400 cores, with fivefold core hour growth every two years. Most are Global 2000-class customers.
- Perform financial batch, RIS modeling, and pricing.
- BD, NoSQL, and so forth.

- Why manage a large production cluster on an external cloud? Internal + cloud benefits:
  - **Zero queue wait.** Cloud is just a tool to augment existing servers. Get the simple stuff off of the internal servers that do the difficult stuff. Johnson & Johnson has a job broker system that allocates jobs internally or to the cloud.
  - **Any scale, any time.** Novartis uses 10,600 cloud servers. We do the cloud orchestration for them, plus burst and data workflow management and finding spot cycles. They found three new compounds.
  - Broad Institute has 50,000 cores for machine learning (ML); 8,000 cores were up and running in minutes, up to 51,200 cores in two weeks.
  - Unique requirements are okay.
- Software challenges: OS, libraries, and app versions.
- Hardware challenge example: For a researcher, NASA had to calculate the number of bushes and trees in sub-Saharan Africa. As a satellite came over the region, data was uploaded onto Amazon S3. This is science that could not be done otherwise. Another example: Life Tech needed to do IR data analyses. They applied a Hadoop cluster, then a grid engine.
- Performance goes up over time at the same cost. This is no longer a three-year, same-performance server; optimal app deployment is a moving target. They move to application life cycles that are optimized and rationalized as new capabilities become available. Benchmarking (app performance) is a moving target. For example, HGST uses a lot of AWS all the time for accelerating manufacturing. You want your tools to talk to any cloud.
- Time and cost are the sole metrics when you're bursting externally — you never care about how much power or cooling you have. For example, a FSI firm required a 250,000-hour burst, twice a year, which is too big for an internal cluster. This can finish in one day on the cloud. The firm went from once per month to twice per week, a frequency that would not make economic sense to do internally.
- Faster iterations: HGST, in developing a new semiconductor drive head design, evaluated 1 million possible designs and completed what would have previously been 70 years of simulation in nine hours.
  - Six different materials
  - 70,000 Ivy Bridge cores
  - 90 times faster run time; $9,000
  - HGST would not have done this many iterations without the cloud and had never done this many before
- Every researcher gets his/her own toy. In the past, for someone to win, someone else had to lose access. With cloud, there's no competition for access, so we can achieve faster iteration cycles. Accelerating the answers also accelerates the people.
- Clients include three of the five largest insurers, large hedge funds, big pharma, and SMBs with no internal resources. SMBs greatly outnumber Global 2000 firms in the world, but the Global 2000 organizations use a significant majority of cloud cycles.
- We don't move data back and forth to the cloud, except the workload piece. Internet2 has 10Gbps capability.
- On security, we try to be Switzerland and let our customers dictate things like security levels and protocols. We have many customers with stringent security requirements.

"Status of the 'Missing Middle': Clouds and Classes of Service": Leo Reiter (Nimbix)

- I run technology strategy for Nimbix. We see the market status quo as people using workstations and then large sites; in between is the missing middle. They don't care about the technology – they care only about getting the work done on time. ISVs are focused more on productivity than performance. The HPC space is focused on massive problems, mostly in the public sector. The missing middle is underserved and under pressure, including the need to meet tight deadlines; their problems are also increasing in complexity. It's hard for these missing middle companies to hire HPC talent, which is scarce and expensive. They have shrinking budgets and are pressured to do more with less.
- It's easy to dismiss the needs of the mainstream missing middle. These companies are in just about every vertical. They range from independent entrepreneurs to Fortune 500 firms and include everything in between. Server-based versions of software are different from – and more expensive than – workstation versions. Many missing middle organizations could benefit from HPC, but don't know this. There are new HPC use cases, such as machine learning and other HPDA.
- When these customers come to us, they ask for 500 GPUs and don't realize they're asking for HPC. Digital transformation is reducing the number of physical models with digital models and computing, and we think most of this will happen in the cloud. These firms will say, "Why do we need to invest in all this infrastructure and manage it?"
- Benefits of cloud for the missing middle include:
  - Limitless scale
  - Available right now
  - Etc.
- We've been working with ANSYS on its pay-for-use licensing model for engineering simulation software:
  - Multiple teams globally (users)
  - Needs always changing
  - Difficult to predict scale and spend up front
- Cloud challenges for HPC:
  - Performance not always as good as dedicated HPC
  - Sticker shock possible but can commit to certain number of hours and get discount
  - May still require retaining or acquiring HPC skills; setting up and maintaining HPC workflows is not easy
- Simulation, media, and machine learning already use cloud; ML is ramping up in the cloud age. HPDA, life sciences, and energy will use it soon.
- Missing middle use cases:
  - An enterprise with 100+ engineers running simulation on electronic components; teams in California, India, and China; and need high-speed interconnects and large-memory machines. ANSYS and more.
  - USDA. Intel + Nimbix + Pacific Biosciences and 56 times coverage of long-read PacBio data on the lesser grain borer genome.
  - HUGEdata = data science as a service (start-up). HPDA and machine learning for its customers. Private cloud today, rebuilding it in public cloud. Project-based work.

    - Real-time photorealistic rendering (migenius). On demand. Need lots of GPUs and CUDA cores. eCommerce and design.
- We don't have a lot of DP use cases and don't see a lot of standards yet. A lot of the data used for training (such as images, sounds, and texts) is in the cloud. Challenges: many GPUs, part time.
- In the future, economics will drive many traditional HPC use cases to the cloud.
- Increasing network speeds will reduce barriers for data transfer to the cloud.
- Market demand continues to convince ISVs to offer elastic cloud licensing.
- More workflows are needing instant domain.

"Performing Simulation-Based, Live Decision Making with Cloud HPC": Zack Smocha (Rescale)

- Rescale was founded by two Boeing engineers. They worked on the 787 redesign, where they did only 11 physical tests, compared with 77 such tests on the Boeing 767.
- We have partnerships to offer more than 120 software packages. We have multiple clouds: 30 locations, with varied hardware configurations.
- Our customers are pushing the limits in aerospace, automotive, semiconductors, life sciences, oil and gas, and other industries.
- Automotive today involves multiphysics and ensembles, but all in the design part of the product cycle. In the manufacturing phase, much less simulation is done (e.g., simulation while the product is in service). One customer runs a fleet of helicopters and captures data from sensors on the helicopters to decide how to do maintenance and detect abnormal behavior. Simulation takes days to hours.
- Using real-time trackside data for race strategy takes minutes.
- Real-time, next-move gaming analysis takes seconds.
- Making sure your car doesn't crash has to happen in microseconds.
- In F1 racing, simulation can help you decide at which points in the race to switch tire types to hard, medium, or soft. It can also tell you when to make pit stops. Data is available trackside. You upload the data. Every lap takes 1.5–2 minutes and decisions take perhaps 30 seconds, so a decision can be made in every lap. So a cluster is running the whole race: 3,000 tasks on 1,500 cores. The result is optimum race time. Results, which take 5MB, are downloaded to a PC. Regulations limit the number of simulations you can do during a race.

"Financial Risk Management Analytics in the Cloud": Jeffrey Smart (SunAmerica)

- For variable annuities, much of our market risk spending is moving to cloud providers. As a result, the company will need a much smaller IT staff than in years past.
- Our end users (customers) are people seeking retirement income for life through variable annuities that have investment and life insurance components. The rules are set in the prospectus.
- Business goals are:
    - Risk management/hedging – Monte Carlo simulation
    - Valuation and support GAAP financial reporting
    - Capital management – this led to our first work on the cloud

    - Support product development – evaluate new products and feature changes
- We calculate values and risk profile, understand actual versus expected profit and loss, and identify hedge instruments to offset market risks.
- We experience a constant demand for the latest information.
- We have a scale of $40 billion book runs overnight on 8,000 cores; this higher-compute capacity translates into more flexibility.
- Increasing demands from the business side drive a need for more computing. Reporting deadlines shrink, assets grow, and higher model resolution mean fewer approximations.
- The key technical constraints were that the internal grid was too small when running and too big when not running; as a consequence, we were still paying for the grid even when we weren't using it.
- A multiyear evolution of our computational platform – with all new hardware and new capabilities – is seeing us migrate our internal datacenter to cloud, where we'll run daily production. We'll then decommission our internal platform. We've expanded our capacity such that we have 10 times as many cores, in two time zones, and for more users, and we've cut month-end batch processing from four weeks to four days. Multidepartment use grew sharply, which has driven down costs. We now pool hardware instead of having departments own their own machines.
- Our next step will be to run a high-visibility, semiannual project on the cloud, for which we'll be recognized for success. The CCAR stress test, mandated by the Federal Reserve, will demand a fourfold increase in capacity. The alternative would have been to buy more hardware, which would have been an order of magnitude more expensive and would not have produced equally good results. Technical issues include that it's hard to modify for a larger, external grid; we need to incorporate information security requirements; and the demands add new stress on the network and schedulers. Organizationally, more people and organizations are involved, in new roles, and with a new cost model.
    - Internal cost scores with number cores.
    - External cost scales with core hours and cuts runtime at fixed cost.
- Step three is to run month-end batches on the cloud. This has evolved from a project to a process by increasing reliability, stability, and automation.
- Step four will be to run daily batches on cloud to perform valuation and hedging runs. We will base trading decisions on this, so results must be on time. We must hedge on market opening. Our focus will be on reliability and automation, and we will continue to refine roles and responsibilities.
- The final step will be to decommission the internal grid. This will be the most difficult step in many ways, in spite of the cost reduction and the simpler process (one environment versus two). The binding constraint is that it will still take eight hours for overnight runs.
- Lessons learned include:
    - Make a long-term technical plan, because schedules don't fail gracefully.
    - Revisit roles and responsibilities as the scope grows.

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