HPC Services and Big Data Challenges

Isabel Schwerdtfeger
Systems Services, IBM Global Technology Services
IBM Germany
Hamburg, Germany
schwerdtfeger@de.ibm.com

Abstract— **High performance computing (HPC) systems rely more on a service- based environment than on a strong hardware or technology-related environment, as more challenges arise for large scale HPC systems of the Exascale era. Naturally, the selection of a distinct HPC hardware technology is very important as the future of the HPC applications and Big Data workloads rely on it. However, the deployment concept including the long term operational aspects need to be considered, as they tend to be the next higher cost driver. Due to high digitization needs, large data growth, and new disruptive business models from organizations and businesses, all face a change and challenge to address these new deployment concepts for their future HPC needs. This paper highlights new possible approaches.**

Keywords- HPC; HPCaaS; Data Center; Cloud; Hyperscale data center; Big Data; Exascale; energy efficiency.

I. INTRODUCTION

For a new large-scale supercomputing system ranging from the higher petaflop to pre-exaflop era, the technical design and the deployment solution are a vital purchase criteria, as there is the need for more efficiency while achieving faster time to market. The high amount of power consumption and its costs are limiting factors for a high capital-intensive investment in a supercomputer. After a successful implementation, the operational services concept needs to ensure high availability and cost efficiency of the various technical components for computing, network and storage, that create the overall system. A typical lifetime of a HPC system is five years and consultations about the new systems begin usually at least a year ahead [1]. But, due to shorter timelines of innovative technologies being available, and the pressure for businesses to foster their research and development (R&D) efforts within shorter timeframes to please their stakeholders, more flexible and adaptive HPC services need to be available.

At current, HPC deployment models are differentiated by so-called “on-premise” or “off-premise” implementations, while “on premise” resemble the traditional on-site hardware (HW) installation in a self-owned data center (DC), usually where the rest of IT of an organization or business resides. As more cloud services and Hyperscale DCs become available, there is increased interests to add HPC capacity “off-premise”, as there is the perception that ideal parameters already exist there for this distinctive computing environment.

A standard HPC deployment for an on-premise installation, where the new supercomputing system is placed in an on-site DC center is characterized by best price purchasing of hardware and software (price per performance). The system selection comes together with the technology solution that demonstrates the lowest energy consumption to reduce power costs, i.e. with direct water cooling, and involving an innovative DC concept. In addition, the reliability of standard hardware and software maintenance support by the HPC vendor for the on-site system is critical and needs to be provided over the lifetime of the system.

HPC is custom built that includes writing software to solve cutting-edge problems and is not an IT function but a competitive business advantage for innovation [2]. HPC is now the technology to solve complex mathematical, scientific or engineering problems to foster research and development activities to drive innovation at companies. Therefore, HPC elevates beyond scientific teams to build footprints in enterprises of all industries.

HPC applications, such as for complex product simulations and optimizations, 3D rendering, complex weather prediction, and deep learning, need custom multiprocessor architectures to solve state of the art problems which comes along with a high capital intensive invest.

This paper is structured as follows: Section II describes the development of HPC in the market and its implications for large scale infrastructures. Section III describes a “Private HPC Cloud” deployment model outlining one of the largest DCs in Europe and its containerized solution for the industrialization of HPC as IT as a Service (IaaS). Section IV discusses the possible conjunction of Big Data with HPC and possible outcomes. The acknowledgement and conclusions close the article.

II. HPC MARKET DEVELOPMENT

Today, a high commoditization of hardware components, and a wider acceptance and usage of common software development tools exists. This leads to a high adoption rate from commercial firms and is not limited to existing research organizations: the democratization of HPC happens. However, the key aspects remain unchanged: the need for high energy efficiency for rising petaflop and future exaflop systems; scalable software that is power and failure aware, and data management software that can handle at minimum the “3Vs” of data: volume, velocity, and variety [3]. Thus, overall efficiency at low costs needs to be achieved.

 ISBN: 978-1-61208-567-8

Industrial HPC is the driving engine of the market growing at 6.8 percent over the forecast period, while education/government remain flat at ca. 3 percent until 2020 [4]. The HPC service market is ca. $4 billion in 2015 and is expected to grow to $5.5 billion by 2020, growing at a Compound Annual Growth Rate (CAGR) of 6.25 percent during the forecast period [5]. The HPC market in total was approximately $28 billion in 2016 and is to grow to over $35 billion globally until 2020 [6]. New era of Exascale computing and cloud-based HPC will offer lucrative opportunities for the market players in coming five years [7]. Emerging economies in regions such as Asia Pacific, and Latin America offer several untapped and unexplored opportunities. North America accounts for more than 45 percent of the total HPC market, followed by Europe with 26 percent market share [5].

The current landscape of large HPC deployments shows mostly “on- premise” installations, custom built in a DC with a range of 1 to 5 Megawatt (MW) power capabilities, in a CAPEX model, with limited adaption capability to a constantly changing IT market. A solution based on a price per core per hour or per Teraflop per Month in an “Infrastructure-as-a-Service” (IaaS) for HPC would remove the high CAPEX investment hurdle to enable HPC for small and medium enterprises (SMEs) and for industrial HPC applications. This accompanied with scalable services with a low consulting cliff, and either fully managed or self-serviced, would spur research and development efforts off.

A new discussion evolves in the HPC market among the existing Hyperscale DCs that are used by large internet companies such as Google, Facebook, and major cloud vendors. One might think that those facilities should be ideal places to host large scale HPC applications. However, when looked at it in detail, the following issues emerge: Servers and direct attached storage are the basic unit and data is widely spread. The hardware is not built with redundant components: if a failure occurs, the workload moves simply to another server [6]. Therefore, usually no or very little custom architecture design is available for applicable HPC requirements. In general, a Hyperscale facilities offers usually no given performance guarantee, nor provide benchmark capabilities. The HPC Cloud is today 3 percent of the total HPC market [4]. Thus, Hyperscale DCs fall short in addressing this growing market.

III. Private HPC Cloud Deployment Model

As the democratization of HPC is already underway, the same happens for the data centers. This is called “The industrialization of data centers”. The goal of this initiative is to become Europe's number one in terms of cost-efficiency, security, flexibility and sustainability. To reach it, the Lefdal Mine Data Center (LMD) is using standardized DC infrastructure based on Rittal's modular and standardized “RiMatrix S” data center portfolio [7]. “The Norwegian Solution” has developed out of the Lefdal Mine in a unique data center concept. Low cost and modularity in a scalable, green and secure facility [8]. In fact, this DC is the largest of the world. It spans 120.000 square meter, is fueled by 100% renewable energy (only wind and water power), direct water-cooled from the local fjord, and a five-level installation in an old mineral mine near Måløy (550 kilometers north-east of Oslo) in Norway. Fig. 1 shows the picture and the dimensions of LMD in comparison to the New York Statue of Liberty, a commercial airplane, a truck and a car vehicle.

Figure 1. Lefdal Mine Datacenter [7]

The solution offering of IBM’s Private Cloud Service for Petascale to Exascale infrastructures for HPC includes the use of LMD, maximizing lowest energy costs and highest expandability at the same time. LMD in Norway aims to exploit Norway’s and the Lefdal Mine’s nature given advantages in terms of location, green power, the low cost of power and cooling, and a stable political environment.

An optimal floor space for "HPC as a Service" (HPCaaS) and a co-location space to cover the extensive capacity requirements of HPC resources in a highly energy- and cost-efficient way. The mine is already there. Cost of land is limited and the investment to be made to secure redundant power and fiber infrastructure and to build out the cooling solution is extremely low. The mine gives natural EMP security and there is less need of perimeter fencing and other investments in physical security. The construction cost per MW is leading in Europe: 30-60% lower than standard DC build out, yet bringing a Tier 3 product to the market [9]. The mine as a IT facility started its operation in the third quarter of 2016.

The high standardization and leveraging economies of scale arise using the standardized modular container solution from Rittal. The modules are suitable to be transported directly into the mine and can be fitted for instance in either containers or secure rooms depending on the protection level needed by the customer [7].

As standard and local client DCs have limitations for the density of racks, many businesses and organizations are not able to fully stack the hardware as tight as possible per rack. Up to 50 KW/rack with leading low cost of power, cooling and space [8]. Hence, LMD is an ideal location for HPC environments using cold- water for cooling. In addition, standard commodity air-cooled server, storage and switches can be used in any custom designed architecture, as the overall energy costs for the operations is at least 30 percent lower than in any standard German IT facility. The key measurement parameter to compare this is the Power Usage Effectiveness (PUE) value: for LMD the PUE is under 1.15. In contrast, current IT operation centers, according to general market information, range from a PUE of 1.35 up to 1.80.

 ISBN: 978-1-61208-567-8

Most of the companies today are interested in using HPC technologies but not necessarily managing it [10]. The combination of IBM Global Services and the offering of the DC capabilities fulfils this desire. Fig. 3 illustrates the solution by IBM in an integrated way, making use of the overall efficiency at low cost which is essential to achieve benefits in terms of high energy efficiency for the need of future Exaflops systems in Europe [11].

IV. BIG DATA AND HPC AS A SERVICE

The growth of substantial amounts of digital data, files, objects or any other digitized information, relate to the term "Big Data". All those types can be characterized by the "5 Vs of Data": Volume, Velocity, Variety, Value and Veracity [12]. Driving needs evolve among businesses and organization to seek more evidence and knowledge out of existing and newly generated data. Available analytics software enables applications to discover new insights of structures, behaviors, trends and relationships. Big Data is extremely complex to deal with via traditional approaches, and requires real-time or almost real-time analysis. [13].

For HPC, open-source software developments are vital to expand the usage and further democratization of its use. OpenStack is a free and open-source software platform for cloud computing, mostly deployed as IaaS [14]. Big Data has changed the way people understand and harness the power of data, both in the business and research domains [15]. When HPC and OpenStack marry, Big Data comes alive within Supercomputing.

Businesses and organizations that are not familiar with HPC, may ask if a big data job could be run on existing HPC infrastructures. The answer is yes, as it is just another type of job vs. a traditional MPI job. [15] The challenges lie within the design of big data systems for I/O libraries and communication. [15]. Many Big Data workloads are running nowadays their new applications on Apache Spark, achieving twice the performance of traditional systems [17], thus supporting the rising demands of real-time or near-real-time analytics.

HPC workloads have varying data, compute and latency requirements. They can have either light data, running smoothly and fast via conventional networking and communication networks. Or they can be heavy workloads requiring more detailed designed communication and high-bandwidth networks. The HPC architectures are either designed in a cluster or in a lightweight grid of a distributed computing system. Fig. 2 outlines the four dimensions of HPC and High Performance Data Analytics (HPDA), also interchangeably used as term for Big Data, of different workloads.

It is critical that data-intensive computing middleware (such as Hadoop and Spark) process such data and are diligently designed, with high performance and scalability to meet the growing demands of Big Data applications. Big Data applications are found in many industries, i.e., in healthcare, bio-medical research, Internet search, finance, and scientific computing. Therefore, mainly the same organizations and businesses that use HPC today have also Big Data requirements. These technologies and applications play a vital role to gain meaningful insights for society and economies.

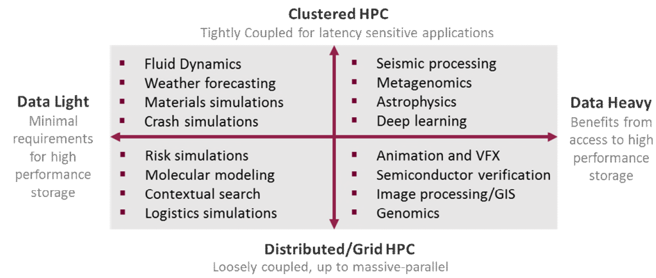

Figure 2. Dimensions of HPC Workloads [17]

Using Apache Spark as a key technology in Big Data leads to certain benefits:

- In-memory large-scale distributed processing;
- Uses distributed file system such as Hadoop Distributed File System (HDFS), which supports automatic data distribution across computing resources;
- Language supports the operations needed to implement the algorithm;
- Good for similar repeated analysis performed on the same large data sets [18].

It can be argued that the combination of Big Data and HPC brings together best of both worlds or it becomes a threat of a takeover to each other, as the demand for Big Data might overwhelm the distinctive HPC market. As a turnkey solution, HPCaaS should demonstrate the high value while delivering this at low costs. Natural benefits come along with the newest available technologies as in the computing processor units (CPUs), either from Intel's x86- or IBM Power-based processors, from many-core or accelerators (GPGPUs), such as from Nvidia P100 or Intel Xeon Phi x200, high performance Infiniband networks, and file-systems such as Spectrum Scale, Lustre, or CEPH based storage.

V. CONCLUSION

This short paper summarized the current state of the HPC market and trend for the next five years. In parallel, as the democratization of HPC occur, the industrialization of DCs evolve, utilizing from the Earth the nature capabilities in a better way for IT demands; thus, being at the same time a novum to standard deployments of traditional DCs sites, matching the new requirements for the era of exaflop systems. As many Big Data applications need more computing capabilities (and more power consumption) to achieve real-time accessibility and more ease-of-use, HPC technologies can be found as the underlying infrastructure. However, the overall operational model must be considered so that further benefits are achieved to reduce increasing costs of those very large systems.

The outlook remains stable for HPC:

 ISBN: 978-1-61208-567-8

- The vendors in this market are continuously focusing on developing innovative technologies and solutions to boost performance and reduce the overall utility cost.
- The major focusing areas are Exascale computing, hot water cooling, networking technologies and embedded GPU and processors.
- The HPC market is no longer limited to the on-premise.
- The battle for HPC is between Intel and NVIDIA for the massive number crunching and data moving work that is the hallmark of HPC.

ACKNOWLEDGMENT

To Martin Kipping, Director International IT-Projects from Rittal, Friedhelm Loh Group and Dr. Sebastian Ritz, CEO of iNNOVO Cloud GmbH, for their inspiration and solution competence.

REFERENCES

[1] R. Blake, F. Robin and M. Sbrighi, PRACE 1IP Work Package 8: HPC Systems Procurement Best Practice, [Online]. Available from: http://www.prace-ri.eu.

[2] S. Conway, “10 Things CIOs Need to Know About High-Performance Computing”, IDC Opinion, Mar. 2016.

[3] D. Soubra, “The 3V´s that define Big Data”, Data Science Central, Jul., 2012, URL: http://www.datasciencecentral.com/forum/topics/the-3vs-that-define-big-data [accessed: 2017-04-15].

[4] M. Feldman, 2016, ”Intersect360 Publishes New Five-Year HPC Market Forecast”, Intersect360 Inc., URL: https://www.top500.org/news/intersect360-publishes-new-five-year-hpc-market-forecast/ [accessed: 2016-11-25].

[5] MarketsandMarkets, “High Performance Computing Market, Global Forecast to 2020”, 2016.

[6] F. Moore, 2016, “Storage Outlook”, Horison Information Strategies, URL: https://indico.bnl.gov/conferenceOtherViews.py?view=standard&confId=1955#20160828 [accessed: 2016-10-15].

[7] Rittal GmbH & Co. KG Press release, 2015, The industrialisation of data centre starts with Rittal, URL: https://www.rittal.com/com-en/content/en/unternehmen/presse/pressemeldungen/pressemeldung_detail_49856.jsp, [accessed: 2015-08-25].

[8] Lefdal Mine Data Center Company Website, 2016, URL: http://www.lefdalmine.com, [accessed: 2017-01-02].

[9] M. Andersson, “Lefdal Mine Data Center, Concept Description”, Aug. 2015, unpublished.

[10] K. Elamrawi, 2016, “Brightskies Technologies: Shaping the Computing Industry of Tomorrow”, URL: http://high-performancecomputing.cioreview.com/vendor/2016/brightskies_technologies, [accessed: 2016-11-25].

[11] I. Schwerdtfeger, “HPC Innovation Services“, IBM Deutschland GmbH, Mar. 2017, unpublished.

[12] B. Marr, 2014, “Big Data: The 5 Vs Everyone Must Know”, URL: https://www.linkedin.com/pulse/20140306073407-64875646-big-data-the-5-vs-everyone-must-know [accessed: 2016-05-12].

[13] S. Yin, O. Kaynak, 2015, “Big Data for Modern Industry: Challenges and Trends”, in Proceedings of the IEEE Volume: 103 Issue: 2, Print ISSN: 0018-9219, Mar. 2015, URL: http://ieeexplore.ieee.org/document/7067026/?reload=true&arnumber=7067026 [accessed 2017-04-30].

[14] Wikipedia, 2016, URL: https://en.wikipedia.org/wiki/OpenStack [accessed: 2016-12-15].

[15] D. K. (DK) Panda, 2017, “Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Big Data Processing”, http://www.cse.ohio-state.edu/~panda, Talk at HPC Advisory Council -Switzerland, Apr. 2017.

[16] M. Christensen, 2017, “OpenPower Foundation Overview”, Talk at HPC Advisory Council -Switzerland, Apr. 2017.

[17] S. Ritz, 2017, “HPC as a Service”, iNNOVO Cloud GmbH Presentation, May 2017, unpublished.

[18] S. Saba, J. Kowalkowski, and M. Paterno, 2016, “Exploring the Performance of Spark for a Scientific Use Case, Saba Sehrish (Fermi National Accelerator Laboratory), HPBDC 2016, The 2nd IEEE International Workshop on High-Performance Big Data Computing, May 2016, URL: http://web.cse.ohio-state.edu/~lu.932/hpbdc2016/slides/hpbdc16-saba.pdf [accessed: 2017-04-12].

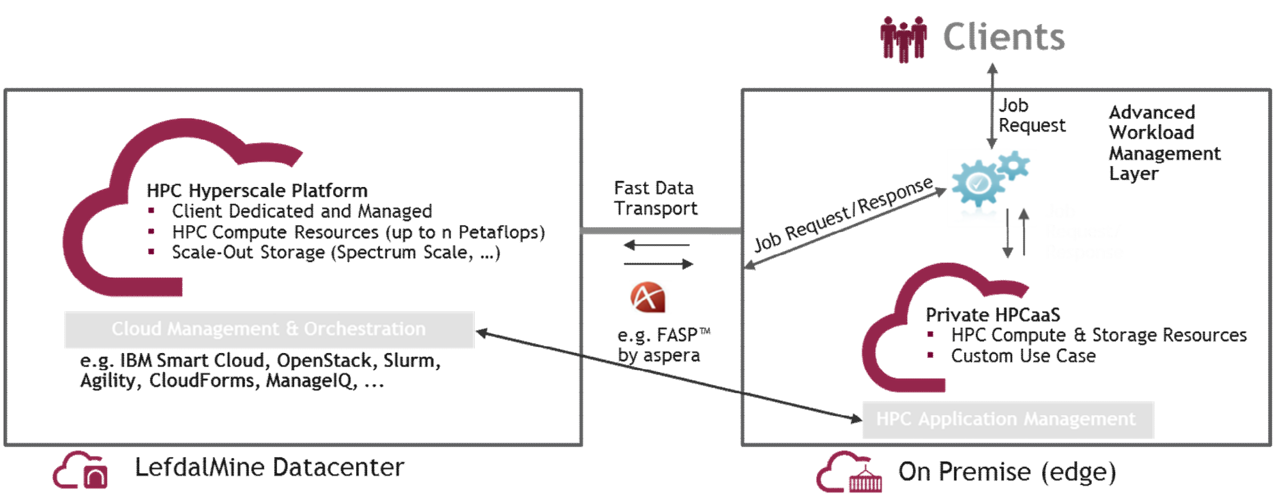

Figure 3. Private Cloud solution for HPCaaS, IBM Deutschland GmbH [11].

