DRIVERS AND BARRIERS TO USING HIGH PERFORMANCE COMPUTING IN THE CLOUD

Executive Briefing – September 2016

Prof. Theo Lynn

Anna Gourinovitch

David Kenny

Xiaoning Liang

CLOUD LIGHTNING

INTRODUCTION

Cloud computing is transforming the utilisation and efficiency of IT infrastructures across all sectors. Enterprises regardless of size can access highly scalable hardware and software resources for the overwhelming majority of IT use cases. In the most part, enterprises and consumers can leverage near real-time on-demand deployment and processing to gain a competitive advantage through reduced capital expenditure, faster response times and in some cases, new business models. Historically, cloud computing has not been used for high performance computing (HPC) to the same degree as other use cases for a number of reasons.

This executive briefing is a preliminary report of a larger study on demand-side barriers and drivers of cloud computing adoption for HPC. A more comprehensive report and analysis will be published later in 2016. From June to August 2016, the CloudLightning project surveyed over 170 HPC discrete end users worldwide in the academic, commercial and government sectors on their HPC use, perceived drivers and barriers to using cloud computing, and uses of cloud computing for HPC.

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

SAMPLE

The sample comprised 92 valid responses worldwide. CloudLightning is an EU-funded project and therefore the sample is weighted towards organisations located in the EU (57.6% of responses). Due to the nature of the CloudLightning project, unsurprisingly, academic sites feature heavily (66.3%). These are summarised in

Table 1 and Table 2 below.

Table 1 Respondents by major sector (n=92)

Sector	Number of Responses	% of Responses
Academic	61	66.3%
Commercial	28	30.4%
Government	5	5.4%
Total	***92***	***100%***

** Respondents may work in various sectors*

Table 2 Respondents by region (n=92)

Region	Number of Responses	% of Responses
European Union	53	57.6%
North America	33	35.9%
Asia	5	5.4%
Middle East	4	4.3%
South America	3	3.3%
Africa	1	1.1%
Total	***92***	***100%***

** Respondents may be based in various locations*

To facilitate future comparison with other third party reports, respondents were also profiled by workload/application areas using an adaptation of IDC's worldwide industry/application area category definitions (IDC, 2015). This data is summarised in Table 3. The sample received a large number of respondents making use of HPC for **biological sciences** including genomics (a focus of the CloudLightning project) and other bio-informatics research. A strong response was also received from those using HPC for **university and academic workloads and applications**.

Table 3 Respondents by HPC workloads/application areas (n=92)

Workload/Application Areas	Number of Responses	% of Responses
Biological sciences	52	56.5%
Chemical engineering	9	9.8%
Computer-aided engineering (CAE) and mechanical design and analysis	12	13.0%
Digital content creation and distribution (DCC&D)	7	7.6%
Economic and financial modeling	14	15.2%
Electronic design and analysis/IT (EDA/IT/ISV)	5	5.4%
Geosciences and geoengineering	12	13.0%
Government laboratories and research centers	13	14.1%
National defense	3	3.3%
University and academic	49	53.3%
Weather forecasting and climate modeling	14	15.2%
Other	14	15.2%
Total	***92***	***100%***

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

CURRENT USE OF CLOUD AND FAMILIARITY WITH CLOUD COMPUTING

Over half the respondents reported using high performance computing frequently (defined as more than once per week). However, only **47.8% reported using the cloud for HPC workloads.** This is higher than previously reported studies and may represent a skew in the sample and/or an upwardly-moving trend towards cloud adoption[1].

Table 4 Respondents using cloud computing for HPC workloads/applications (n=44)

Workload/Application Areas Adopting Cloud Computing	Number of Responses	% of Responses
Biological sciences	24	54.5%
Chemical engineering	1	2.3%
Computer-aided engineering (CAE) and mechanical design and analysis	4	9.1%
Digital content creation and distribution (DCC&D)	5	11.4%
Economic and financial modeling	5	11.4%
Electronic design and analysis/IT (EDA/IT/ISV)	2	4.6%
Geosciences and geoengineering	5	11.4%
Government laboratories and research centers	5	11.4%
National defense	0	0%
University and academic	18	40.9%
Weather forecasting and climate modeling	2	4.6%
Other	9	20.5%
Total	***44***	***100%***

The respondents reported making use of a wide range of cloud deployment models with many making use of more than one. Unsurprisingly **public cloud services** were the most popular **(70.5%)**, followed by **private clouds (65.9%)**; the latter mostly likely reflecting the sensitive nature of the sample's application areas. **15.9%** of cloud adopters reported using **community clouds.**

Moreover, for those who reported using private clouds for HPC, on average, 65.5% of their HPC workloads or application areas were running in private clouds. For public cloud, community cloud, and hybrid cloud users, the average percentages of their HPC in the cloud workloads were 61.1%, 58.9%, and 38.6%, respectively.

Table 5 Share of HPC workload in cloud (n=44)

Cloud Categories	Public Cloud	Private Cloud	Community Cloud	Hybrid Cloud
Number of responses	31	29	7	5
% of responses	70.5%	65.9%	15.9%	11.4%
Average Usage (%)	61.1%	65.5%	58.9%	38.6%

[1] IDC, 2015. Worldwide Broader HPC 2014–2018 Forecast: Servers, Storage, Software, Middleware, and Services.

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

The most commonly used cloud service providers were **Amazon Web Services (76.2%) and Microsoft Azure (28.6%).** A small number of respondents identified IBM (Softlayer) (9.5%) and Google Cloud Platform (9.5%). Over 13 other cloud service providers were identified with only one respondent each. These included academic institutions, government laboratories, national and international shared services for academic institutions such as ICHEC (Ireland), EGI Federated Cloud (EU), Compute Canada (Canada) and a number of smaller players in the cloud service provision market.

Despite the reported adoption rates, the majority of respondents claimed to be familiar with cloud computing with **43.5% claiming to have advanced knowledge of cloud computing or to be cloud computing experts.**

Table 6 Familiarity with cloud computing (n=92)

Statement	Number of Responses	% of Responses
I am a cloud computing expert	18	19.6%
I have advanced knowledge about cloud computing	22	23.9%
I know a lot about cloud computing, but I am still learning	25	27.2%
I have a basic understanding of cloud computing	24	26.1%
I have no understanding of cloud computing	3	3.3%
Total	***92***	***100%***

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

DRIVERS OF CLOUD COMPUTING ADOPTION FOR HIGH PERFORMANCE COMPUTING

The primary driver for use of cloud computing for HPC workloads was reported as access to **extra capacity for overflow or surge workloads.** Clearly the on-demand rapid deployment of IT resources within the cloud computing paradigm is extremely attractive to HPC users. Other popular reported drivers included reduced costs (including capital and operating expenditure) and access to specialised hardware, software and other resources not available in the respondents' on-premise datacentres. Respondents were relatively consistent in what they reported as drivers but also those factors that did not currently influence their decision i.e. government policy and the opportunities to access and work using HPC infrastructure remotely.

Figure 1 Drivers of adoption of cloud computing for HPC workloads (n=92)

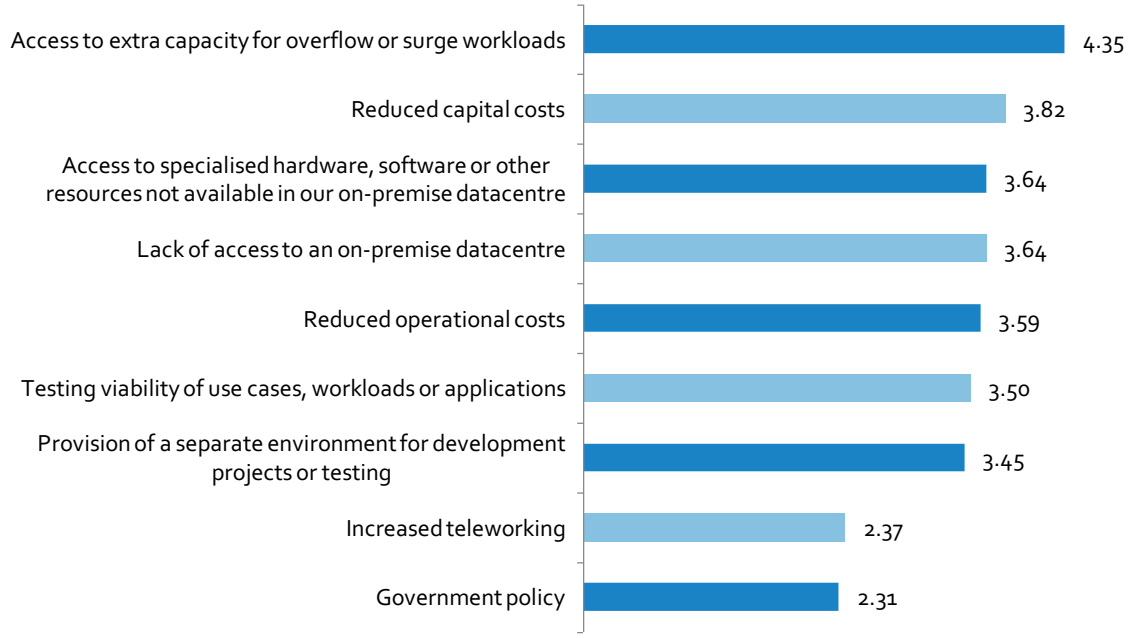

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

BARRIERS TO CLOUD COMPUTING ADOPTION FOR HIGH PERFORMANCE COMPUTING

Trust in cloud computing would appear to be a significant barrier to adopting cloud computing for HPC workloads. Data management concerns dominate the responses. This is not surprising given the large number of bio-science and university and academic respondents within the sample. The main technical barriers relate to **communication speeds.** This reflects a perceived lack of cloud infrastructure capable of meeting the communications and I/O requirements of high-end technical computing. Government policy is again ranked low it would seem it is neither a driver nor a barrier. Unsurprisingly availability and capital expenditure are not barriers reflecting their positive impact on adoption.

Figure 2 Barriers to cloud computing adoption for HPC workloads (n=92)

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

SUMMARY

- There is unlikely to be a full shift of high performance computing workloads to the cloud in the short term however there is evidence of demand to meet the capacity limitations of internal infrastructures including use cases for testing the viability of the cloud or specific software for various use cases. This is consistent with previous research.

- There is demand for cloud computing for high performance computing but trust concerns remain. This includes data protection, control, privacy and general security. Cloud service providers need to deploy and communicate compliance with relevant national and international regulations for data protection and control by application/workload area.

- The economic arguments for adoption would seem to be accepted however there are still technical concerns regarding the limitations of communication speeds in the cloud for high performance computing. Cloud service providers need to make more capable interconnects, larger memory and a wider range of heterogeneous resources (including GPUs, many integrated cores (MICs), and data flow engines (DFEs)) available in the cloud to address these concerns.

- Government policy may be viewed as ineffectual within the context of the use of high performance computing in the cloud. Given in many jurisdictions, the government is a major funding source for high performance computing infrastructure, it is arguable that cloud computing should be promoted by government agencies to support the principle of economy, efficiency and effectiveness.

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NUMBER 643946

About CloudLightning

CloudLightning is a three-year €4m EU-funded research project to address energy efficiency and high performance in cloud computing. It proposes a novel cloud management and delivery architecture based on the principles of self-organisation and self-management that shifts the deployment and optimisation effort from the consumer to the software stack running on the cloud infrastructure. The goal of the project is to address this inefficient use of resources and consequently to deliver savings to the cloud provider and the cloud consumer in terms of reduced power consumption and improved service delivery, with hyperscale systems particularly in mind. The CloudLightning solution will be demonstrated in the three application domains - genome processing, oil and gas exploration, and ray tracing.

For more information, please visit http://cloudlightning.eu.

Acknowledgement

This work is partially funded by the European Union's Horizon 2020 Research and Innovation Programme through the CloudLightning project (http://cloudlightning.eu/) under Grant Agreement Number 643946.