



HYPERION RESEARCH

End-User Research Findings: HPC in Cloud Environments



Steve Conway
Bob Sorensen
March 2018

Our Cloud Research (examples)



HYPERION RESEARCH

Global Study

Special MCS Study
2017 HPC Multi-Client Study: The Use of Public/External Clouds for HPC Workloads, Trends, and Drivers

Bob Sorensen, Earl Joseph, and Steve Conway
September 2017

**Clouds for science
and public
authorities**
SMART 2011/0055


D4 –Final Report

Government Study




April 8, 2013

The opinions expressed in this study are those of the authors and do not necessarily reflect the views of the **European Commission**



HYPERION RESEARCH

Private Study



Pulse Survey:
HPC Private Clouds

Steve Conway
sconway@hyperionres.com
October 2017

What We Often Hear

“My boss says, why can’t you do everything in the cloud?”

“Clouds are only for embarrassingly parallel jobs.”

“It’s much more expensive to run jobs in the cloud.”

“It’s much cheaper to run jobs in the cloud.”

“I’m worried about data security and data loss in the cloud.”



What We (and Users) See Happening

Historic Model



- On premise data center
- Policy driven
- Inelastic
- Inflexible
- Highest data security

Emerging Model



- Seamless hybrid environment
- Policy driven
- Elastic
- Flexible
- Cloud data security concerns

Main Cloud Types

Self-Hosted



- Private cloud
- May extend past firewall
- May be national/global



Externally Hosted



- Public cloud
- Virtualized private cloud (VPC)
- Outside firewall



Pent-Up Demand

All Respondents		
Can you estimate the amount of "pent-up demand" exists for HPC systems (work that your users would like to run, but can't due to limited resources), as a percent of your total yearly workload volume?	None, we are still in the mode of filling up the system	20.4%
	≤ 24%	30.1%
	25-50%	13.6%
	51-100%	16.5%
	101-200%	10.7%
	201-300%	1.0%
	> 300%	1.0%
	Don't Know or Uncertain	6.8%
Government Only		
Can you estimate the amount of "pent-up demand" exists for HPC systems (work that your users would like to run, but can't due to limited resources), as a percent of your total yearly workload volume?	None, we are still in the mode of filling up the system	16.7%
	≤ 24%	16.7%
	25-50%	8.3%
	51-100%	25.0%
	101-200%	16.7%
	201-300%	0.0%
	> 300%	0.0%
Don't Know or Uncertain	16.7%	



Cloud Types for HPC Workloads

TABLE 15

Types of Clouds Used for HPC Workloads in 2017

	Number of Responses	Percent Cloud-Using Sites
Private or internal clouds	39	51.3%
Public or external clouds	26	34.2%
Hybrid private/public clouds	11	14.5%
N=76		

Trends in Cloud Types Used for HPC

**But note that many private clouds are now “virtualized private clouds” offered by Amazon AWS or Microsoft Azure*

TABLE 17

The Change in the Use of Different Types of Clouds for HPC

	2015	2017	Change
Private or internal clouds	70.2%	51.3%	-18.9%
Public or external clouds	25.6%	34.2%	8.6%
Hybrid private/public clouds	4.1%	14.5%	10.3%
N=76			

Note: This is ONLY for sites that use clouds for HPC workloads.

Reasons for Using Public/External Clouds

All Sites

1. Extra capacity (“surges”)
2. Cost-effectiveness
3. Isolate R&D projects
4. Special HW/SW features
5. Management decision
6. No on-premise data center

Government Only

1. Special HW/SW features
2. Extra capacity (“surges”)
3. No on-premise data center
4. Cost-effectiveness
5. Management decision
6. Isolate R&D projects

Running HPC Workloads in Public/External Clouds?

TABLE 12

Use of External Clouds for HPC Jobs in 2017

	Number of Responses	Percentage of Sites
No	81	73.0%
Yes	30	27.0%
Average % of total HPC workload run in public/external cloud: All respondents		9.7%
Average % of total HPC workload run in public/external cloud: <u>Yes</u> respondents only		35.8%
N=111		



% HPC Workload in Public/External Clouds – Next 6-18 Months

All Respondents		
IN THE NEXT 6 to 18 MONTHS -- Approximately what percent of all your HPC application workload will be run on external cloud computing?	0%	40.2%
	1-9%	13.7%
	10-24%	15.7%
	25-50%	11.8%
	> 50%	6.9%
	Unsure / Not Known	11.8%
Government Only		
IN THE NEXT 6 to 18 MONTHS -- Approximately what percent of all your HPC application workload will be run on external cloud computing?	0%	54.5%
	1-9%	9.1%
	10-24%	18.2%
	25-50%	0.0%
	> 50%	9.1%
	Unsure / Not Known	9.1%



TODAY: % Big Data Applications Run in External Cloud

All Respondents		
Approximately what percent of all your Big Data application workload is run on external cloud computing TODAY?	0%	54.9%
	1-9%	10.8%
	10-24%	13.7%
	25-50%	5.9%
	> 50%	7.8%
	Unsure / Not Known	6.9%
Government Only		
Approximately what percent of all your Big Data application workload is run on external cloud computing TODAY?	0%	45.5%
	1-9%	0.0%
	10-24%	36.4%
	25-50%	0.0%
	> 50%	9.1%
	Unsure / Not Known	9.1%



Top Limitations of External Clouds

All Sites

1. Data upload/access times
2. Data security

Government Only

1. Data upload/access times
2. Data security



**EXPECT
DELAYS**



DATA BREACH

Ease of decision for workloads on cloud

Commentary

- Most sites (55.8%) say it is relatively easy to decide which jobs to run in a cloud.
- But 1 in 5 (19.2%) disagree.

In your organization (or those you sell to), how easy is it to decide which technical computing workloads to run in a cloud

	Percent	Respondents
Very easy	17.3%	9
Somewhat easy	38.5%	20
Somewhat difficult	19.2%	10
Very difficult	11.5%	6
Not sure/don't know	13.5%	7
		54

Ease of decision—private/public cloud

Commentary

- 52% say it's relatively easy to decide which jobs to run in public vs. private clouds.
- But 1 in 3 sites (32.7%) say the decision is difficult.

In your organization (or those you sell to), how easy is it to decide which technical computing workloads to run in PUBLIC VS. PRIVATE cloud environment?

	Percent	Respondents
Very easy	21.2%	11
Somewhat easy	30.8%	16
Somewhat difficult	17.3%	9
Very difficult	15.4%	8
Not sure/don't know	15.4%	8
		52

Regulated industry?

Commentary

- 3 in 10 sites say they're in regulated industries.

Is your organization (or those you sell to) in a highly regulated industry?

	Percent	Respondents
Yes	30.8%	16
No	50.0%	26
Not sure/don't know	19.2%	10
		52

Most frequent technical workloads for cloud

Commentary

- EP workloads dominate a bit more as the MOST IMPORTANT cloud workloads of the sites.

Which of the same technical computing workloads does your organization (or those you sell to) run MOST OFTEN in a cloud environment? (Choose ONLY ONE)

	Percent	Respondents
"Embarrassingly parallel" workloads	33.3%	17
Communications- and I/O-intensive workloads	27.5%	14
Other	15.7%	8
Not sure/don't know	23.5%	12
		51

Final decisions on workloads for cloud

Commentary

- PIs and CTOs/R&D leaders stand out as the most important final decision-makers.

In your organization (or those you sell to), who makes the final decision about which technical computing workloads get sent to cloud environments (Choose ALL that apply)

Percent

Respondents

	Percent	Respondents
Principal investigator (PI)	27.5%	14
CTO/R&D leader	11.8%	6
R&D member	7.8%	4
CIO	7.8%	4
Data center manager/other data center official	5.9%	3
IT manager	5.9%	3
Data scientist	3.9%	3
CEO	3.9%	2
Chief data officer	2.0%	1
Line of business (LOB) manager	2.0%	1
Product designer	0.0%	0
CFO	0.0%	0
Vendor/partner	0.0%	0
Other	2.0%	1
Not sure/don't know	19.6%	10
		51

Factors to cause increase in cloud use

Commentary

- A whole constellation of related factors is expected to drive cloud growth in the next 2 years, with reduced costs leading the way.

What factors will cause your organization to expand cloud use in the next 2 years? (Choose ALL that apply)

	Percent	Responses
Cheaper to run some workloads in a cloud	39.2%	20
Need for more capacity to augment on-premise	31.4%	16
Need to avoid more capital expense (CAPEX)	23.5%	12
Need to limit/avoid buying more on-premise HPC	21.6%	11
Management directive to do more on cloud	19.6%	10
Need to isolate R&D projects	17.6%	9
Need to limit/avoid operating costs (OPEX)	13.7%	7
Need for special hardware/software we don't have	11.8%	6
We have no on-premise HPC resources	5.9%	3
Not sure/don't know	17.6%	9
		51



HYPERION RESEARCH

Questions?

Steve Conway

sconway@hyperionres.com

Bob Sorensen

bsorensen@hyperionres.com

