

Special Analysis

How Cloud Computing is Changing HPC Spending

Earl Joseph, Melissa Riddle, and Alex Norton
December 2020

HYPERION RESEARCH OPINION

The use of public clouds to run HPC workloads grew dramatically in 2019 and is projected to have large growth over the next five years. This Special Analysis looks at on-premise spending for HPC combined with spending on cloud computing for HPC workloads. When on-premise spending for HPC is combined with spending in the cloud, the total HPC spending for 2019 was \$31.6 billion and is projected to reach \$47.0 billion in 2024.

The current on-prem HPC spending tracking includes five major categories: servers, add-on storage, middleware, applications, and maintenance services. In this Special Analysis, Hyperion Research is adding in the spending for running HPC workloads in public clouds. Cloud spending has not been separated out into different sub-categories; rather, it is considered in one single line item (everything required to run an HPC workload in the cloud). This can include the compute cost, storage cost (both temporary and long term), data movement charges, ISV licenses, and any additional services from CSPs or other third parties.

As cloud computing continues its strong growth in the HPC market, both in utilization and in spending from end users, it has become a crucial component to many HPC budgets. Historically, Hyperion Research has displayed the broader HPC market to only include the on-prem categories. This year, the cloud portion is being included, creating a wider view of the HPC market to include more of the different components of HPC budgets. Cloud computing, even as recent as a few years ago, was one of the smaller portions of the HPC market, well below add-on storage and just above maintenance services in magnitude. As CSPs have worked to address HPC workloads, and HPC users have started to understand which portions of the HPC application set can be run in the cloud, overall budgets have grown to accommodate more cloud computing spending. Forecasting forward five years, cloud computing is poised to overtake every on-prem category except servers and become the second largest budget allocation in the HPC market, highlighting the shift in the market of users adopting cloud computing to handle an increasing amount of their HPC workloads.

CURRENT SITUATION

Definitions

On-premise HPC spending includes these five categories:

- Servers: Defined as the entire contract for the system, potentially also including built-in storage, middleware, software required to turn on the system, and any NRE required for a particular system.
- Add-on Storage: Defined as any additional storage not included in the server contract.
- Middleware: Defined as middleware software such as compilers, tools, programming languages, etc.
- Applications: Defined as any ISV licenses or additional application spending.
- Maintenance Services: Defined as the services required to maintain existing HPC servers. Note that it excludes other types of professional services.

Public cloud computing for HPC workloads is not broken up into the same categories as on-prem spending tracking but includes all spending, such as computing resources, data movement and storage. Hyperion Research has found, through multiple studies and interviews with users, that roughly two-thirds of spending is on computation, which may include ISV licenses and anything else needed to spin-up a compute environment, while approximately one-third is for storage, which includes long term and archive storage. Public cloud spending, as tracked, includes the entire bill paid to a CSP in order to run HPC applications on a cloud platform. This includes storage, OS, ISV or other application software, the interconnect, and any additional characteristic needed to run a workload in the cloud.

HPC SPENDING TRENDS

Table 1 shows HPC on-premise spending by the five categories listed above, for 2018 to 2024. Total on-premise spending is projected to grow by a 6.7% CAGR from 2019 to 2024. All categories are expected to have a net growth over the next 5 years, with initial decreases in 2020 outweighed by strong growth over the remaining forecast period. Add-on storage is expected to grow the fastest with an 8.3% five-year CAGR, due to high demands from emergent AI workloads and larger simulations requiring expansions in storage capabilities for HPC datacenters. Meanwhile, maintenance services are expected to have the slowest growth (3.4% CAGR).

TABLE 1**Projected On-Premise HPC Spending by the Broader HPC Market Areas (\$M)**

	2018	2019	2020	2021	2022	2023	2024	CAGR '19-'24
Server	13,675	13,710	11,846	13,295	15,817	17,942	19,044	6.8%
Add-on Storage (On-Prem)	5,381	5,427	4,772	5,410	6,519	7,577	8,099	8.3%
Middleware	1,590	1,613	1,402	1,576	1,902	2,171	2,317	7.5%
Applications	4,652	4,689	4,062	4,455	5,258	5,862	6,111	5.4%
Maintenance Services	2,248	2,239	1,899	2,040	2,366	2,587	2,643	3.4%
Total On-Prem	27,546	27,678	23,981	26,774	31,862	36,138	38,214	6.7%

Source: Hyperion Research, October 2020

Table 2 shows projected spending to run HPC workloads in public clouds. HPC cloud spending is expected to grow at a 17.6% CAGR from 2019 to 2024, outpacing on-premise growth over the same time period. Cloud spending is also expected to increase each year of the forecast period. Although all on-premise categories are expected to see a dramatic drop in 2020 followed by a recovery, cloud spending is expected to grow 10.0% in 2020. Notably, cloud spending also had exceptionally high year over year growth (58.6%) between 2018 and 2019. This exceptionally high growth year was defined by Hyperion Research as a tipping point, driven by the increase of new HPC cloud users, as well as the increased utilization of current HPC cloud users.

TABLE 2**HPC Public Cloud Spending (\$M)**

	2018	2019	2020	2021	2022	2023	2024	CAGR '19-'24
HPC Cloud Spending	2,466	3,910	4,300	5,300	6,400	7,600	8,800	17.6%
Yearly Growth Rate	--	58.6%	10.0%	23.3%	20.8%	18.8%	15.8%	

Source: Hyperion Research, October 2020

Table 3 adds HPC cloud spending to the on-premise spending to show total HPC spending across the forecast period. The addition of the "cloud line" increases total HPC spending from \$27.6 billion on-

prem in 2018 to \$30.0 billion total (on-prem and cloud) in 2018. The difference between on-prem spending alone and on-prem spending plus cloud spending is most dramatic in 2024, where inclusion of cloud increases total projected HPC spending from \$38.2 billion on-prem to \$47.0 billion total.

TABLE 3

HPC Spending by All HPC Areas (\$M)

	2018	2019	2020	2021	2022	2023	2024	CAGR '19-'24
Server	13,675	13,710	11,846	13,295	15,817	17,942	19,044	6.8%
Add-on Storage (On-Prem)	5,381	5,427	4,772	5,410	6,519	7,577	8,099	8.3%
Middleware	1,590	1,613	1,402	1,576	1,902	2,171	2,317	7.5%
Applications	4,652	4,689	4,062	4,455	5,258	5,862	6,111	5.4%
Maintenance Services	2,248	2,239	1,899	2,040	2,366	2,587	2,643	3.4%
Total On-Prem	27,546	27,678	23,981	26,774	31,862	36,138	38,214	6.7%
HPC Cloud Spending	2,466	3,910	4,300	5,300	6,400	7,600	8,800	17.6%
Total HPC Spending (Cloud & On-Prem)	30,012	31,588	28,281	32,074	38,262	43,738	47,014	8.3%

Source: Hyperion Research, October 2020

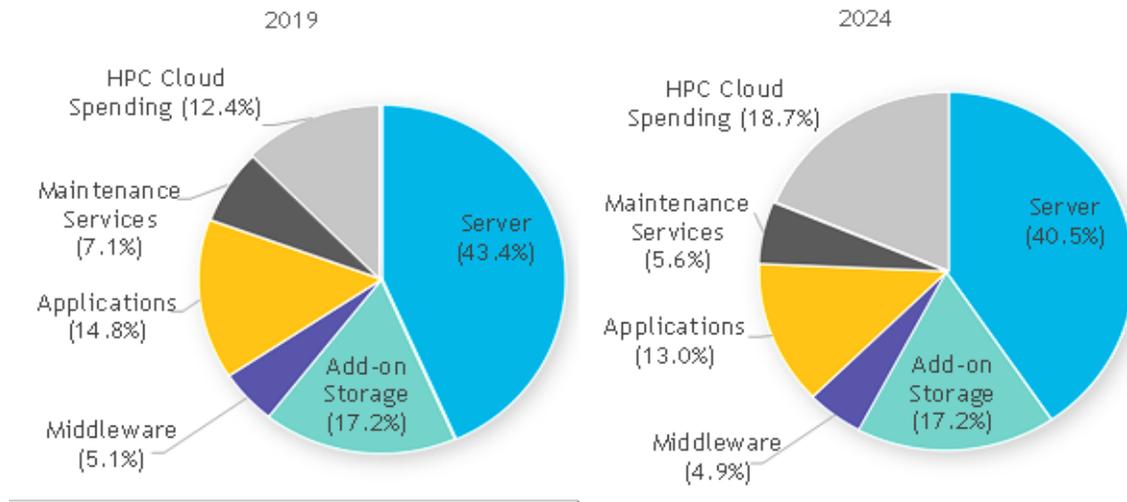
CLOUD CONTINUES TO INCREASE SHARE OF HPC BUDGETS

While all categories of HPC spending are growing, some are growing faster than others and causing the distributions of HPC budgets to evolve. Cloud spending accounted for only 8.2% of all HPC spending in 2018 and grew to 12.4% of all HPC spending the following year. By the end of 2024, cloud spending is expected to represent nearly a fifth of HPC budgets (18.7%). In turn, most on-prem categories will represent slightly smaller proportions of overall budget. Servers are expected to be the most negatively impacted, dropping from 43.4% of overall budget in 2019 to 40.5% in 2024. Add-on storage is the only on-prem category expected to remain proportional to overall HPC budget over the forecast period. Figure 1, below, highlights the shifting in budget allocation from 2019 to 2024.

Since 2019 was a tipping point for HPC cloud spending, this category grew dramatically between 2018 and 2019 and caused all on-prem spending categories to have a lower relative budget share during the same period.

FIGURE 1

Overall HPC Spending in 2019 and 2024



Source: Hyperion Research, October 2020

Table 4 shows the historical spending mix changes from 2018 to 2019 and then the projected spending mix for 2024.

TABLE 4

Changes in HPC Spending

	2018 % Share of Total HPC Spend	2019 % Share of Total HPC Spend	2024 % Share of Total HPC Spend
Server	45.6%	43.4%	40.5%
Add-on Storage (On-Prem)	17.9%	17.2%	17.2%
Middleware	5.3%	5.1%	4.9%
Applications	15.5%	14.8%	13.0%
Maintenance Services	7.5%	7.1%	5.6%
Total On-Prem	91.8%	87.6%	81.3%
HPC Cloud Spending	8.2%	12.4%	18.7%
Total HPC Spending (Cloud & On-Prem)	100.0%	100.0%	100.0%

Source: Hyperion Research, October 2020

FUTURE OUTLOOK

Cloud computing for HPC applications has grown substantially over the past few years, highlighted by a substantial shift in the market that occurred between 2018 and 2019 and a resilience to the dramatic economic changes of 2020. As Hyperion Research looks forward to the next five years, there is an anticipated major growth year, similar to 2019, although the exact timing and magnitude of the next major market growth year is difficult to predict. *Please see Hyperion Research's recent SC20 Market Update Briefing for an example of that major growth year (<https://hyperionresearch.com/virtual-briefing/>).*

As CSPs continue to improve their platforms to address a growing set of HPC applications, HPC users are looking to the cloud to handle more of their HPC applications, especially AI workloads. The hybrid model, where users take advantage of both on-prem capabilities and cloud platforms, is poised to be a popular resource model in the next few years. Further, as mentioned above, HPC in the cloud has remained strong even during the challenges experienced during the covid-19 pandemic and the related economic issues of 2020. In fact, covid-19 has accelerated the growth of cloud computing for HPC as the value proposition of elastic resources has grown in wake of unpredictable budgets and government restrictions on in-person workspaces.

The HPC market remains a growth market, with all major tracked sectors exhibiting growth over the forecast period. By enabling users to expand their resource pools without on-prem additions, the cloud has altered the course of the HPC market.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue

St. Paul, MN 55102

USA

612.812.5798

www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2020 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.