

Special Report

A Growing and Changing HPC Applications Landscape

Melissa Riddle and Mark Nossokoff

April 2022

HYPERION RESEARCH OPINION

The application software landscape is quickly evolving along with HPC workloads. Independent software vendor (ISV) applications, as opposed to open-source or home-grown, have traditionally been considered the gold standard as the source for many HPC applications and were frequently cited as HPC users' top applications. While ISV revenues continue to rise, open-source application use is growing as well. Although HPC users can be reluctant to change the applications used at their site, developments such as the onset of AI, new types of hardware, and the cloud have spurred many users to explore new applications they may not have considered otherwise. Overall, the HPC application software landscape is rapidly developing in tandem with new HPC infrastructure and use cases.

The overall HPC market is forecasted to grow across all broad market areas, including the application software area, which is projected to grow with a CAGR of 7.4% in 2021-2025.

The application software landscape is also undergoing some fundamental changes:

- ISV revenues are increasing proportionally to growing HPC workloads
- Open-source application usage is increasing as percent of HPC workloads
- Growing use of AI and HPDA is encouraging HPC sites to adopt new applications
- New types of hardware are showing up across all parts of the HPC ecosystem, presenting many new opportunities
- Cloud growth has spurred fundamental changes in application use when users' preferred applications are unavailable in the cloud
- A growing number of applications are being utilized widely across multiple verticals
- Programming languages are becoming an increasingly important part of HPC workflows

Note: This page is intentionally blank.

CURRENT SITUATION

The on-premises HPC broader market is composed of servers, storage, middleware, applications, and services. This market is expected to grow steadily through 2024, followed by a slight decline in 2025 due to the expected slowdown of exascale system acceptances. All of the HPC broader market categories are expected to exhibit this growth pattern, including applications. HPC applications are expected to represent more than \$6B in 2025, a 7.4% five-year CAGR.

TABLE 1

Worldwide HPC Broader Market Forecast (\$M)

	2019	2020	2021	2022	2023	2024	2025	CAGR '20 - '25
Server	\$13,368	\$13,523	\$15,097	\$17,149	\$18,736	\$20,071	\$19,651	7.8%
Storage	\$5,288	\$5,053	\$6,145	\$7,097	\$7,943	\$8,578	\$8,433	10.8%
Middleware	\$1,572	\$1,484	\$1,776	\$1,990	\$2,143	\$2,266	\$2,187	8.1%
Applications	\$4,569	\$4,295	\$5,090	\$5,674	\$6,076	\$6,373	\$6,144	7.4%
Service	\$2,181	\$2,006	\$2,334	\$2,477	\$2,526	\$2,530	\$2,352	3.2%
Total HPC Broader Market Forecast	\$26,979	\$26,361	\$30,442	\$34,386	\$37,425	\$39,819	\$38,766	8.0%

Source: Hyperion Research, 2022

ISV Application Growth

Since ISVs are the only applications that are revenue-generating, ISV revenues are synonymous with all application revenues (see Table 1). ISV fees per site can fluctuate substantially over time because ISV software is often licensed (e.g., per use or per core) rather than purchased. The anticipated ISV license growth is primarily being fueled by expanding HPC workloads. The average number of systems per site now exceeds 14 servers and is continuing to rise to meet demands for more differentiated HPC resources. This translates to increased licensing fees to use the same applications.

Many users are willing to support these rising costs because changing applications can be difficult and costly. In addition, less experienced HPC users, such as new entrants from the traditional IT datacenter space, are likely to select ISV software due to ease of use and more robust support offerings. In general, ISV software is still perceived as superior to open-source in terms of reliability, support, and time to solution.

Open-Source Application Growth

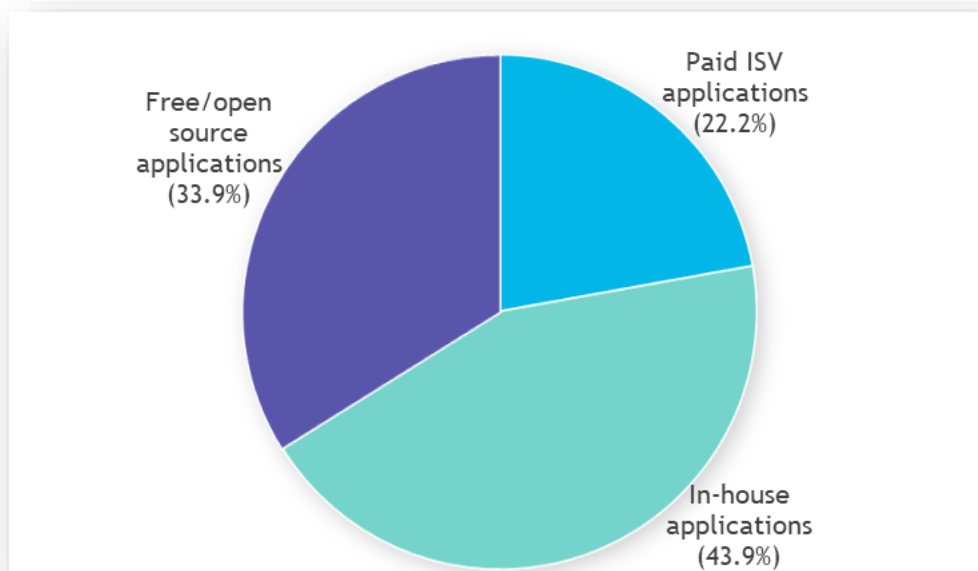
Since open-source applications do not generate revenues, the most meaningful metrics for this category relate to usage patterns. According to recent survey data, more than 90% of HPC users use

open-source software at least some of the time. Over the last three years, the average HPC user reported using open-source software more often than ISVs, with about 30-40% of an HPC workload consisting of open-source applications (see Figure 1).

There are multiple reasons for this trend, such as growing awareness of, and communities around, open-source applications that make them easier for new users to adopt. Another important factor is the rise of emerging technologies such as GPUs and cloud use (see *Applications in the Cloud*, below). Many ISVs have not yet ported their code to run on accelerators, causing some users to turn to new applications to take advantage of new hardware technologies. For these users, open-source is a popular choice because developing in-house code requires an investment that many users may not be ready to commit to.

FIGURE 1

Percent Usage of ISV, In-House, and Open-Source Codes in Average HPC Workload



Note: n=194

Source: Hyperion Research, 2022

Motivators for Application Growth and Changing Usage

Four important trends are affecting HPC application usage:

- Widespread AI and HPDA growth
- Cloud usage
- Sector and vertical usage patterns
- Programming languages

Applications and AI-HPDA

AI-HPDA use is widespread among HPC sites. 80% of HPC sites today report running AI, big data, or data analytics applications and this proportion is expected to grow to 94% over the next two years. Many HPC users who are expanding from traditional modeling/simulation to AI-HPDA adopt new applications for these purposes.

In the most recent Hyperion Research Multi-Client Study, HPC users were asked to list their top two successful AI/big data applications, if any. Several answers were cited across multiple verticals: Hadoop, MATLAB, R, Python, PyTorch, and TensorFlow. Notably, some of these are actually programming languages, not applications (see *Programming Languages*, below). Overall, these responses show the preference for open-source applications (and free tools, middleware, etc.) for AI-HPDA use.

Applications in the Cloud

Increasing cloud use is affecting application usage patterns for both ISV and open-source applications. When preferred ISV applications are available in the cloud, many HPC users opt to maintain continuity in their on-premises and in-cloud workloads. This generates cloud ISV revenues additional to the on-premises ISV revenues discussed previously.

However, ISV licenses are often unavailable or cost prohibitive in the cloud. This drives some users to adopt new applications for cloud use, whether or not they continue to use their on-premises ISV. Typically, these users select open-source applications because they are perceived as low-risk for experimentation in the cloud. Alternatively, some users plan to wait for their preferred applications to be offered in the cloud and keep HPC workloads exclusively on-premises in the meantime.

Applications are expected to play an important role over the next several years in determining which HPC workloads, and at what scale, will be run in the cloud.

Applications by Sector and Vertical

ISV and open-source software use differs greatly by sector (Industry, Government, Academia) due to factors such as varying software budgets and security requirements:

- Academic sites tend to have more limited budgets, so they depend on open-source applications more frequently (50% of their workloads) while using ISVs more sparingly (14%).
- Government sites also use less ISV software (14%) due to their reliance on in-house codes (52%).
- Industry sites use less open-source (27%) and instead invest in ISV software (28%).

Within the Industry sector, some verticals use more open-source applications than others. In the most recent Multi-Client Study, users were asked to list their top three applications. All verticals surveyed had users who responded with open-source applications among their top three, but this was particularly common among Academic, Bio-Sciences, Economics/Financial, EDA/IT/ISV, and Weather sites. Industrial users tend to use more ISV applications than other segments.

Recent survey data of the top three applications per site has also shown a growing number of applications that are used across multiple verticals. This trend is seen across both ISV and open-source applications. Some of these programs are not traditionally associated with a single vertical, such as programming languages (see *Programming Languages*, below) or software designed for AI and High-Performance Data Analysis.

For example, big data ISV program Spark was listed in the top three at sites in DCC & Distribution, Economics/Financial, EDA/IT/ISV, and Government Lab. Other broadly cited applications have traditionally been associated with one vertical but are being adapted for use in other verticals. For example, open-source molecular dynamics program GROMACS has traditionally been associated with biochemical applications but is now listed among the top three applications at sites in Academic, Bio-Sciences, EDA/IT/ISV, Government Lab, and Weather verticals. These interdisciplinary applications (and more) show how HPC workloads are becoming more varied and novel ways of using existing applications are being discovered.

Programming Languages

Programming languages and applications have traditionally been thought of as distinct categories but are starting to blur for some users. In the most recent Multi-Client Study, 7% of HPC users included a programming language in their list of top three applications. Among these responses, MATLAB was the most common, followed by R and Python. Although this is a minority of users, it still reflects the growing importance of programming languages at HPC sites.

According to recent survey data, each HPC site surveyed uses about 6 different programming languages on average, and this number has been rising. At least half of sites reported using Python (78%), C/C++ (74%), MPI (61%), and OpenMP (50%). Fortran and its variants are also still used at many sites. Overall, this data shows both the breadth of programming languages being leveraged for HPC and the growing trend of considering a programming language as an application in its own right.

FUTURE OUTLOOK

ISV applications were traditionally the cornerstone for many HPC sites, but open-source applications are becoming a more feasible option for many HPC users. Both HPC users and ISV vendors would be wise to keep an eye on the growing communities around relevant open-source software as these are important indicators of application adoption and support. As ISV licensing fees increase along with HPC workloads, some users may experiment with open-source alternatives to save costs.

Emerging technologies will also play an important role in the future of HPC applications. The rise of AI and HPDA is encouraging many traditional modeling/simulation users to adopt new applications dedicated to these purposes. Continued cloud and GPU growth are also expected to continue to change the application makeup at HPC sites. However, if CSPs and OEMs cannot make desired applications readily available on their platforms, applications may in turn become a formidable barrier to increasing cloud and GPU use.

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 and vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides 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.hpcuserforum.com and www.HyperionResearch.com

Copyright Notice

Copyright 2022 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.hpcuserforum.com or 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.