*Asian Journal of Innovation and Policy (2025) 14.1:103-114*
*DOI: https://doi.org/10.7545/ajip.2025.14.1.103*

# Public-Sector HPC Support in the AI Era: Policy Lessons from Korea's Innovation Support Program

**Hyungwook Shim *, Minho Suh****

**Abstract** This study examines Korea's HPC Innovation Support Program, a public initiative that provides computational resources through Grand Research and Creative Research tracks. The program has broadened national research capacity but also reveals disparities, with strong utilization in fields such as earth sciences and chemistry and relatively low use in electrical and electronic engineering. These imbalances suggest the need for complementary measures including training, software support, and interdisciplinary collaboration. Beyond resource provision, the program represents an institutional mechanism for balancing equity and competitiveness in Korea's research ecosystem. Although smaller in scale than HPC policies in the United States, Japan, or China, its transparent and structured procedures offer a meaningful benchmark model. Ultimately, the program should be seen not only as a technical infrastructure but as a strategic policy tool that strengthens equity, efficiency, and international competitiveness in the global AI–HPC era.

**Keywords** HPC, AI, HPC innovation support program, Republic of Korea

## I. Introduction

With the advent of the AI revolution, the demand for High-Performance Computing (HPC) has increased dramatically. The training and inference of deep learning and generative AI models require vast computational resources, making HPC indispensable for modern scientific and industrial research. Large language models (LLMs) such as GPT-4 and PaLM-2 demand computing performance on the scale of hundreds of TFLOPS, which is only feasible in HPC environments.

---

Submitted, May 17, 2025; 1st Revised, August 27, 2025; Accepted, September 21, 2025
* Ph.D, Division of National Supercomputing Center, Korea Institute of Science and Technology Information, Daejeon, Korea; shw@kisti.re.kr
** Corresponding, Center Director, Division of National Supercomputing Center, Korea Institute of Science and Technology Information, Daejeon, Korea; mhsuh@kisti.re.kr

HPC systems, based on parallel and distributed computing as well as supercomputing technologies, enable large-scale computation by leveraging tens to thousands of GPUs and TPUs. These systems maximize processing efficiency through data and model parallelism, making it possible to accelerate research across diverse domains. Beyond AI model training, HPC also provides the computational backbone for simulations and real-time analysis in science and engineering.

The applications of HPC are rapidly expanding across industries. In fields such as autonomous driving, financial risk prediction, and climate modeling, real-time computation is essential. HPC also plays a transformative role in the life sciences, finance, and environmental sciences by supporting genomics, drug discovery, risk management, and climate prediction. As AI technologies converge with multiple industries, the importance of HPC infrastructure in optimizing and testing increasingly complex models will only grow. The availability of HPC resources, their scale, and the presence of skilled personnel have therefore become decisive factors in shaping the AI technology gap between nations and firms.

In response to these challenges, the Republic of Korea has established the HPC Innovation Support Program, which provides free access to HPC resources for startups, small and medium-sized enterprises (SMEs), and academic institutions. While large corporations often maintain private HPC infrastructures, most SMEs and research organizations lack such capacity. This study, therefore, goes beyond a descriptive introduction and critically analyzes the program’s institutional design, performance, and limitations. By examining issues of equity, efficiency, and international comparability, the study aims to assess the program’s role in Korea’s national innovation system and to derive policy implications that may inform both domestic improvements and global benchmarking efforts.

## II. Literature Review

This study analyzes relevant prior research to identify the trends in existing studies and derive the limitations of previous work, thus reviewing the direction and academic contributions of this paper. The reviewed studies primarily focused on the construction and operation of HPC infrastructure, HPC education and training, and the impact of HPC on industrial and economic development. However, research on HPC support programs in the public sector remains relatively scarce. Consequently, this study aims to examine the HPC support programs of the Korean government and derive policy implications through comparisons with leading countries. Below are some key prior studies related to this research.

MacLachlan (2020) analyzed the challenges and obstacles encountered in the process of establishing an HPC center at the university level for the first time and proposed strategies and solutions to address these issues. Specifically, the study discusses elements such as securing legitimacy, cost models, expected use cases, resource planning, personnel allocation, user engagement, and success metrics, providing practical guidance for universities aiming to operate new HPC centers (MacLachlan, 2020). Škrinárová (2016) investigated HPC education strategies using systems such as grids, clouds, and HPC clusters. By analyzing HPC research strategies in Europe and the United States, the study assessed the need for education and training and proposed a strategic model for HPC education and training. Based on this, it designed HPC curricula for undergraduate, master's, and research programs and suggested e-learning-based HPC cluster practice methods. Furthermore, the educational program was operated in Slovakia's HPC infrastructure in a cluster environment (Škrinárová, 2016). Richards (2011) presented an educational roadmap for researchers and domain experts new to HPC resources. To address the issue where researchers lacked knowledge of key HPC concepts, principles, and resource utilization methods, this study developed an educational roadmap that outlined essential HPC concepts, technologies, and tools. The roadmap includes various learning materials and reference resources, enabling new HPC users to learn based on their needs (Richards, 2011). Joseph (2013) analyzed the impact of HPC investment on economic success and scientific innovation from the perspective of public benefit. The study developed two macroeconomic models and an innovation index. The first model explains how HPC investment impacts economic development, including GDP growth, cost reduction, and job creation, while the second model compares how HPC promotes basic and applied innovation across industries, nations, and organizational sizes. Additionally, it developed a new innovation index to quantitatively assess HPC's impact on economic and scientific progress (Joseph, 2013). Existing studies have primarily focused on analyzing HPC infrastructure construction, education and training, and economic and scientific contributions. However, research on HPC support programs in the public sector is limited, especially in the context of Korea. Most case studies published to date focus on leading HPC nations such as the United States and Europe, resulting in a lack of research that can guide countries like Korea, which are in the early stages of HPC adoption and operation, in setting policy directions.

For comparative analysis of HPC support programs across countries, this study selected representative institutions that provide HPC resources to external researchers. The three key institutions are the NERSC (National Energy Research Scientific Computing Center) at ORNL (Oak Ridge National Laboratory) in the United States, AIRC (Artificial Intelligence Research Center) at AIST in Japan, and NSCC (National Supercomputing Center) in China. First, NERSC supports research missions of the DOE (Department of Energy) and requires researchers to submit proposals during a designated period each year. When applying for resources, researchers must submit an ERCAP proposal form, including project goals, required resources, and expected timeframes. Resources allocated by NERSC include CPU-only nodes and GPU-accelerated nodes of Perlmutter. NERSC provides various HPC utilities, programming libraries, development tools, debuggers, profilers, and data and visualization tools. Furthermore, NERSC supports researchers by offering training events and documentation of best practices. AIRC promotes AI and big data research by sharing HPC resources with external researchers and institutions. In 2018, AIRC launched the AI Bridging Cloud Infrastructure (ABCI), a cloud platform with world-class AI processing capabilities, enabling efficient deep learning, machine learning, and big data analysis. NSCC operates a network of supercomputing centers located in various cities across China, providing HPC resources to a wide range of external users. External users can check information and proceed with the application process through NSCC's official website. Selected projects receive support not only for HPC resources but also for technical assistance, software optimization, and data management (Antypas, 2012; Rikters, 2024; Yang, 2024).

Although these programs are similar in terms of supported resources, application methods, and goals, they differ in that they do not disclose selection criteria and procedures for applying to projects. Countries adopting or operating HPC systems should plan policies and support programs for balanced resource distribution across industries. Benchmarking the approaches and procedures of advanced countries in HPC development could provide an economically and efficiently viable alternative. Therefore, this paper analyzes Korea's representative HPC support programs and aims to derive policy implications by comparing them with the cases of other countries.

## III. Introduction to the National Supercomputing Center

The Korea Supercomputing Center (KSC) was designated by presidential decree under the "Act on Utilization and Fostering of National Supercomputers" to promote the development and utilization of national HPC. As the designated national supercomputing institution, KSC is responsible for establishing fundamental plans and key policies, forecasting resource demand and acquisition, conducting research and development, and fostering international cooperation. Based on the enforcement decree of this act, the Korea Institute of Science and Technology Information (KISTI) has been designated as the managing institution for KSC. Since its establishment in 1988, KSC has been building and operating supercomputing resources for over four decades. Currently, Supercomputer No. 5, Nurion, is in operation. At the time of its deployment in 2018, Nurion achieved a theoretical peak performance of 25.7 PF and was ranked 11th on the TOP500 list of the world's most powerful supercomputers. Nurion consists of two types of compute nodes: Knights Landing and Skylake, and was manufactured by Cray Inc. Key characteristics of Nurion include its high-performance multi-core architecture, with approximately 8,300 many-core CPUs, each capable of delivering 3 teraflops per socket. Additionally, it features a high-speed interconnect, large-scale storage, and a Burst Buffer to efficiently manage large-scale I/O requests generated by applications. To facilitate efficient utilization by researchers, KSC provides a variety of commercial software packages, including MSC ONE for structural mechanics, ANSYS CFX and FLUENT for computational fluid dynamics (CFD), and GAUSSIAN for chemistry and biotechnology applications. Furthermore, KSC offers technical support services, covering batch job execution, system access, application software usage, compiler support, environment configuration, and data migration (Shim, 2023; Lee, 2011; Ko, 2022).

## IV. Introduction of HPC Innovation Support Program

### 1. Objectives

The HPC Innovation Support Program is designed to address grand scientific challenges and national issues that were previously infeasible due to computational resource limitations. Additionally, it aims to support data-driven research and collaborative studies integrating AI and HPC.

### 2. Research Areas and Scope

The program consists of two primary support tracks: Grand Research Program that supports large-scale and medium-scale computational projects. Large-scale projects require simultaneous use of at least 100,000 CPU cores to solve a single computational problem. Medium-scale projects also require extensive computational resources, but on a relatively smaller scale. Creative Research Program is for supporting research in computational science, manufacturing innovation, and AI.

## 3. Eligibility Criteria

The eligibility requirements for each program are as follows. Grand Research Program is open to Ph.D.-level researchers affiliated with domestic industry, academia, or research institutions. Applicants must require dedicated computational resources of either 1,500 nodes (large-scale projects) or 500 nodes (medium-scale projects) to solve a single computational problem. Creative Research Program is open to Ph.D.-level researchers in computational science, manufacturing innovation, and AI from domestic industry, academia, or research institutions.

## 4. HPC Resource Allocation

The HPC resources provided to users are allocated in units of quotas, as summarized in Table 1. The node hours for each system are converted as follows: KNL – 6,400 node hours per quota, SKL – 3,200 node hours per quota, and Neuron – 534 node hours per quota. For large-scale research projects, 524 quotas are allocated, while medium-scale projects receive 175 quotas. In the case of creative research, each project is allocated 50 quotas. The support period for projects in each category is set to 1 month, 3 months, or 1 year.

**Table 1. Summary of Support Details (Unit: Quotas)**

| | Grand Challenge Research | | Creative Research Projects |
|---|---|---|---|
| | Large-Scale Projects | Medium-Scale Projects | |
| Eligible Institutions | Universities, Government Research Institutes, Industries | | |
| Fields | No Restrictions | | |
| Total Quotas | 1,572 | 525 | 5,898 |
| Quotas per Project | 524 | 175 | 50 |
| Support Duration | 1month | 3month | 12month |

Source: The authors

## 5. Application and Selection Process

Applications for the program are submitted through the KSC website, where the principal investigator completes and uploads the application form. Each researcher can apply for up to two projects per year. Additionally, researchers who achieve outstanding research outcomes (within the top 5% of JCR rankings) are eligible for one additional project for the following year. The selection process for projects is illustrated in Figure 1.

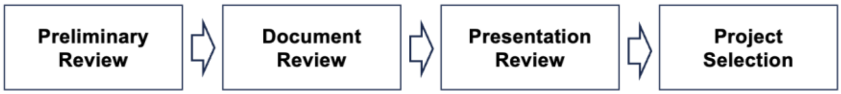


**Figure 1. Selection process**

In the preliminary review stage, the application is assessed for eligibility, completeness, and relevance to the designated research fields. If any deficiencies are identified, applicants are given a specified period to make necessary revisions.

During the document review stage, an evaluation committee composed of experts in relevant fields is formed, and an online review is conducted. Evaluators are selected from a pool of approximately 250 experts based on their expertise and suitability. The evaluation criteria, as outlined in Table 2, include research excellence, applicability, expected outcomes, and potential impact. The final score is determined by calculating the arithmetic mean of the scores assigned by three committee members.

An additional face-to-face evaluation is not conducted for Creative Research Program. However, for Grand Research Program, if a proposal receives a score of 85 or higher in the document review, it proceeds to the face-to-face evaluation stage. The evaluation criteria and scoring system remain the same as in the document review stage.

**Table 2. Evaluation Criteria and Weighting**

| Evaluation Category | Sub-Criteria |
|---|---|
| Scientific Excellence | Necessity of the research |
| | Originality and leadership |
| | Clarity and ambition of objectives |
| Suitability of Utilization | Justification for HPC usage |
| | Adequacy of requested resources |
| Expected Impact | Contribution to science and economy |

Source: The authors

## 6. Program Performance and Results

The operational status and outcomes of the program are summarized in Tables 3 and 4. Table 3 presents an analysis of the resource allocation across various research fields in 2024. The supported research areas are broadly categorized into Grand Research and Creative Research, and the number of selected and non-selected projects, allocated resources, and selection rates for each category are compared. In the Grand Research category, a total of 11 projects were submitted, of which 10 were selected, resulting in a high selection rate of 91%. The allocated resource volume amounted to 2,411 quotas, while the requested resources for non-selected projects totaled 524 quotas, indicating a selection rate of 91% based on the number of projects and 82% based on resource volume. In contrast, the Creative Research category received a total of 422 project applications, with 248 projects selected, resulting in a selection rate of 59%. The allocated resource volume in this category was 7,476 quotas, whereas the requested resources for non-selected projects totaled 6,312 quotas, yielding a selection rate of 54% based on resource volume. Overall, out of 433 project applications, 258 (60%) were selected, with a total allocated resource volume of 9,887 quotas, corresponding to a selection rate of 59% based on resource volume. These results suggest that Grand Research programs are prioritized, while the competition for resources is more intense in the Creative Research category.

**Table 3. HPC Resource Allocation (2024)**

| Research Category | Selected Projects | Rejected Projects | Allocated Quotas | Unselected Quotas | Selection Rate (Projects) | Selection Rate (Resources) |
|---|---|---|---|---|---|---|
| Grand Research | 10 | 1 | 2,411 | 524 | 91% | 82% |
| Creative Research | 248 | 174 | 7,476 | 6,312 | 59% | 54% |
| Total | 258 | 175 | 9,887 | 6,836 | 60% | 59% |

An analysis of SCI(E)-indexed journal publications from 2020 to 2024 (Table 4) reveals a gradual decline in research output. The number of published papers increased from 205 in 2020 to 238 in 2021 but then declined steadily in the following years: 212 papers in 2022, 192 in 2023, and 157 in 2024. Several factors may have contributed to this decline, including changes in the research environment, adjustments in resource allocation policies, and fluctuations in project selection rates. Notably, the 157 publications in 2024 represent a 23% decrease compared to 2020, highlighting the need for a thorough investigation into the underlying causes and the development of strategies to address this issue.

**Table 4. Publication Trends (2020–2024)**

| Year | 2020 | 2021 | 2022 | 2023 | 2024 |
|---|---|---|---|---|---|
| Publications | 205 | 238 | 212 | 192 | 157 |

Source: The authors

An analysis of the average monthly system utilization rate by research field in 2024 (Table 5) shows that earth science had the highest utilization rate at 24.7%, followed by chemistry (19.3%), physics (15.7%), and materials (14.6%).

In contrast, mechanical engineering (8.0%), energy/resources (7.7%), and chemical engineering (4.9%) recorded relatively lower utilization rates, while electrical/electronic engineering (2.0%) and other fields (3.1%) had significantly lower usage levels. These results are likely directly related to the computational demands of each research field. For example, earth sciences involve large-scale simulations such as climate modeling and earthquake prediction, which require substantial computational resources, leading to higher utilization rates. In contrast, fields like electrical and electronic engineering may have fewer computation-intensive research projects. To optimize the research environment moving forward, a resource allocation strategy that considers the specific needs of each field is necessary. Additionally, support measures should be developed to enhance utilization rates in underutilized fields.

**Table 5. HPC Utilization by Research Field (2024)**

| Field | Utilization Rate (%) |
|---|---|
| Earth Science | 24.7 |
| Chemistry | 19.3 |
| Physics | 15.7 |
| Materials | 14.6 |
| Mechanical Engineering | 8 |
| Energy/Resources | 7.7 |
| Chemical Engineering | 4.9 |
| Electrical Engineering | 2 |
| Other | 3.1 |

Source: The authors

## V. Discussion

The findings of this study demonstrate that the HPC Innovation Support Program functions as an important policy instrument within Korea's national innovation system. While previous research has largely emphasized HPC infrastructure development and economic impacts, this study highlights the institutional mechanisms through which public-sector support contributes to technological competitiveness. In this regard, HPC policy is not only a technical foundation enabling large-scale computation but also an institutional tool tha determines the distribution of resources and the reconfiguration of research capacity.

A key issue identified in the analysis concerns the rationale and institutional design of the program. Unlike some advanced countries where resource allocation is decentralized or market-driven, Korea's program is characterized by state-led eligibility requirements, evaluation criteria, and selection procedures. This approach ensures equity and procedural transparency, but it also presents challenges in terms of flexibility and scalability in responding to rapidly growing computational demands. Future comparative studies with cases from the United States, Japan, and China are needed to examine how such institutional features shape program outcomes.

Another important finding relates to disparities in utilization across scientific fields. High demand in earth sciences and chemistry contrasts with limited use in electrical and electronic engineering. While this study suggested possible structural factors such as shortages of specialized personnel, software adoption issues, and disciplinary cultures, these remain hypotheses without strong empirical backing. Future research should therefore incorporate survey-based evidence or case-level data to test these explanations more rigorously. By linking resource use patterns to measurable disciplinary constraints, policymakers could identify targeted interventions to narrow utilization gaps.

From a methodological perspective, program success should not be judged solely by descriptive indicators such as project counts or publication numbers. Establishing causal relationships between HPC resource allocation and research outcomes is essential for stronger academic and policy contributions. Comparative analyses of participating versus non-participating groups, as well as sectoral and disciplinary differences, could help clarify whether access to HPC resources leads to measurable improvements in research productivity, innovation, or competitiveness. Such evidence would allow policymakers to design governance strategies that balance efficiency with equity.

Finally, international comparison must move beyond simple reference to peer institutions and instead focus on their governance models and complementary support mechanisms. For example, NERSC in the United States combines competitive proposals with extensive user support and training; Japan's AIRC emphasizes AI-focused infrastructure and industry collaboration; and China's NSCC integrates large-scale resource provision with technical assistance and software optimization. These non-resource policies—such as education, workforce development, and technical services—provide a structural benchmark for Korea as it develops complementary strategies to resource allocation. When placed in the broader context of intensifying global AI–HPC competition, the Korean program can thus be understood not only as domestic infrastructure but as a strategic instrument. Strengthening international collaboration, expanding workforce training, and incorporating economic impact assessments will position Korea as a more competitive and inclusive participant in the global HPC ecosystem.

## VI. Conclusion

This study analyzed Korea's HPC Innovation Support Program, showing that its dual-track design of Grand Research and Creative Research has enhanced national research capacity and supported diverse scientific achievements. Beyond infrastructure provision, the program serves as an institutional tool to balance equity and competitiveness in the research ecosystem.

However, persistent disparities in utilization across fields and the decline in publication outputs raise concerns about sustainability. The concentration of resources in certain domains and underutilization in others suggests that resource expansion alone is insufficient. Future policies should incorporate complementary measures such as domain-specific training, software optimization, and incentives for interdisciplinary collaboration.

Korea's HPC policy, though modest compared with the United States, Japan, and China, is distinguished by its transparent eligibility and institutional clarity. These features, reinforced through international cooperation and economic impact assessment, can serve as a benchmark for other nations, including developing economies.

Ultimately, the program should be positioned as a strategic instrument in the global AI–HPC era. Future policy development should focus on four pillars: balanced resource distribution, improved evaluation frameworks, expanded international collaboration, and workforce development. By advancing these directions, Korea can establish a more competitive and inclusive HPC policy model.

## Acknowledgment

This research was supported by the Korea Institute of Science and Technology Information (KISTI) (No. K25L2M1C1).

## Reference

MacLachlan, G., Hurlburt, J., Suarez, M., Wong, K.L., Burke, W., Lewis, T., and Ensor, B. (2020). Building a shared resource HPC center across university schools and institutes: a case study, arXiv preprint arXiv:2003.13629.

Škrinárová, J., and Vesel, E. (2016, November). Model of education and training strategy for the high performance computing. In 2016 International Conference on Emerging eLearning Technologies and Applications (ICETA), 315-320. IEEE.

Richards, M., and Lathrop, S. (2011, July). A training roadmap for new HPC users. In Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery, 1-7.

Joseph, E.C., Conway, S., and Dekate, C. (2013). Economic model for a return on investment analysis of United States government high performance computing (HPC) research and development (R & D) investment (No. 243-296), IDC Research Inc., Framingham, MA (United States).

Antypas, K. (2012). NERSC 2011: High performance computing facility operational assessment for the National Energy Research Scientific Computing Center.

Rikters, M., and Miwa, M. (2024, November). AIST AIRC systems for the WMT 2024 shared tasks. In Proceedings of the Ninth Conference on Machine Translation, 286-291.

Yang, H., Liu, L., and Wang, G. (2024). Does large-scale research infrastructure affect regional knowledge innovation, and how? A case study of the National Supercomputing Center in China, Humanities and Social Sciences Communications, 11(1), 1-20.

Shim, H., Ko, M., Choe, Y., and Hahm, J. (2023). A study on institution improvement plans for the national supercomputing joint utilization system in South Korea, International Journal of Advanced Computer Science and Applications, 14(3).

Lee, S.M., Kim, J., Kim, M., Kim, H., and Choi, N.J. (2011, July). Industrial HPC activities in Korea. In 2011 International Conference on High Performance Computing & Simulation, 814-818. IEEE.

Ko, M., Kim, M., and Park, S.U. (2022). An economic ripple effect analysis of domestic supercomputing simulation in the industrial sector, Journal of Information Science Theory and Practice, 10(spc), 66-75.