



AI Computing Loads Reshape Power Systems: Shenzhen Launches the Greater Bay Area Computing–Electricity Coordination Innovation Consortium

Yimin Deng,¹ Jun Shen,^{1,2,*} Ran Tian,¹ Tao Yang,¹ Shaohua Yang,³ Borong Lin,⁴ Yi Ding,⁵ Yiqing Cai,⁶ Huiyu Ouyang,⁷ and Xingying Chen⁸

¹School of Mechanical Engineering, Beijing Institute of Technology, Beijing 100081, China

²Jiangsu Intelligent Energy Use and Low Carbon Technology Research Institute, Nanjing 211100, China

³Beijing Institute of Technology, Zhuhai 519088, China

⁴Key Laboratory of Eco Planning & Green Building, Ministry of Education, School of Architecture, Tsinghua University, Beijing 100084, China

⁵College of Electrical Engineering, Zhejiang University, Zhejiang 316021, China

⁶China electricity council, Beijing 100761, China

⁷Shenzhen Energy Innovation Technology Co. Ltd., Shenzhen 518045, China

⁸School of Electrical and Power Engineering, Hohai University, Nanjing 210024, China

*Correspondence: jshen@bit.edu.cn

Received: September 20, 2026; Accepted: September 22, 2026; Published Online: September 23, 2026; <https://doi.org/10.59717/ipj.energy-use.2026.100075>

© 2026 The Author(s). This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by/4.0/>).

Citation: Deng Y., Shen J., Tian R., et al. (2026). AI Computing Loads Reshape Power Systems: Shenzhen Launches the Greater Bay Area Computing–Electricity Coordination Innovation Consortium. *Energy Use* 2:100075.

On September 15, 2026, the 2026 Energy Green Development Conference and the 2026 International Digital Energy Expo opened jointly in Shenzhen, Guangdong Province. This was the first time the two events were held as a single conference-and-exhibition. At the conference, North China Electric Power University, together with the Shenzhen Power Supply Bureau, Shenzhen Energy Group, and other partners, launched the Greater Bay Area Computing–Electricity Coordination Innovation Consortium. According to public materials, the Consortium will work across six areas, namely policy research, technology development, standard-setting, demonstration projects, technology transfer, and talent training. The goal is to link planning, data, standards, and demonstration work that has so far been scattered across grid operators, computing and telecom firms, energy companies, and universities.

Jianbang Zheng, Vice Chairman of the Standing Committee of the National People's Congress and Chairman of the Central Committee of the Revolutionary Committee of the Chinese Kuomintang (RCCK), opened the conference. Addresses were subsequently given by Yun Li (Deputy Secretary of the CPC Shenzhen Municipal Committee and Mayor of Shenzhen), Kun Yang (Secretary of the Party Committee and Executive Vice Chairman of the China Electricity Council), Mingbin Ji (Director, General Manager, and Deputy Secretary of the Party Committee of China Southern Power Grid Co., Ltd.), Norela Constantinescu (Deputy Director of the Innovation and Technology Centre of IRENA), and Deshun Liu (Chief Engineer of the National Energy Administration). Also in attendance were Jizhen Liu, Academician of the Chinese Academy of Engineering and Chief Expert at the National Energy Development Strategy Research Center jointly established by the RCCK Central Committee and North China Electric Power University, together with Jiaqiang Liu and Zehua Ouyang, Vice Chairmen of the RCCK Central Committee. Xingying Chen, Vice Chairwoman of the RCCK Central Committee, presided over the opening ceremony. This year's conference gave "energy plus computing power" a central place on the agenda. As artificial intelligence expands rapidly, computational demand is turning into a new kind of high-density electrical load. The key question is no longer just how much electricity is needed, but whether the power system can deliver it at the right place and time, with enough capacity, stability, and low-carbon credentials.

From Growth in Electricity Consumption to Localized Grid Constraints

The International Energy Agency (IEA) estimates that data centers worldwide consumed approximately 415 TWh of electricity in 2024, accounting for about 1.5% of the global electricity consumption. Under its base-case scenario, this figure is projected to rise to approximately 945 TWh by 2030, with artificial intelligence among the most significant drivers of this growth.¹ At the roundtable dialogue on "Artificial Intelligence and Green Energy," chaired by Haisheng Chen, Director of the Institute of Engineering Thermophysics, Chinese Academy of Sciences, and featuring six academicians and experts, Xiang Gao, Jinyue Yan, Yuxin Peng, Guanyi Chen, Bofeng Bai, and Vladimir Terzija noted that considerable uncertainty remains in these figures, and that a global share does not imply that all regions face equal pressure. The more immediate challenge for power grids lies in the spatial concentration of load. Large data centers tend to cluster in a limited number of regions, and the construction cycle for new computing facilities can be shorter than the expansion cycle for transmission, substation, and generation-side infrastructure. As a result, the energy challenges arising from computing-power growth manifest first as constraints on local interconnection capacity, peak load, and supply reliability, rather than merely as growth in annual electricity consumption.

In a presentation on computing–electricity coordination and national security, Wang Peng, Executive Dean of the National Energy Development Strategy Research Institute at North China Electric Power University, several issues were identified the mismatch between the construction timeline of computing centers and the commissioning schedule of power supply, the insufficient willingness to participate in demand-side regulation, and the still-unresolved market mechanisms as key obstacles that current computing–electricity coordination efforts must address. Several participants argued that this reflects growing strain on the conventional, one-directional planning model of "load proposes demand, grid expands to meet it". If the scale, location, and growth rate of computing loads cannot be incorporated into generation and grid planning in advance, relying solely on after-the-fact capacity expansion may result in longer interconnection cycles and higher system costs.

At the same time, experts at the conference cautioned that treating data




Figure 1. Conceptual computing–power–electricity coordination architecture for high-density AI loads.

achievements address bottlenecks at different levels: higher-capacity transmission and distribution infrastructure can alleviate local inter-connection constraints. The emergency response system enhances resilience under extreme conditions, and grid-forming converters address voltage and frequency support in the low-inertia systems that result from high shares of renewable energy penetration.⁵ None of these can substitute for the others. Grid-forming control can improve stability-support capability but does not, by itself, increase the long-term availability of electricity. Enhancing transmission capacity can expand inter-connection space, but this is not equivalent to reducing the carbon intensity of the electricity consumed by data centers.

Alongside infrastructure, load-side flexibility is equally essential. Several speakers emphasized that, for AI data centers, the object that can genuinely be dispatched is not the abstract notion of “computing power” but rather computing tasks with differing latency tolerances and power characteristics, together with the cooling, energy-storage, and backup power equipment coupled to them. Related research indicates that spatiotemporal load shifting can help data centers respond to electricity prices, renewable energy output, and carbon-intensity fluctuations, but its system-level value depends heavily on the data center’s location, season, workload structure, and the proportion of flexible load.^{3,4} This implies that computing–electricity coordination should not pursue a single efficiency metric, such as a data center’s own power usage effectiveness (PUE), in isolation, but must also consider whether its load profile aligns with grid supply–demand conditions and the availability of low-carbon power sources.

Discussions during the conference on market mechanisms for computing–electricity

centers uniformly as inflexible, rigid loads is likewise inaccurate. Existing research indicates that certain computing tasks with temporal elasticity, such as batch processing and model training. It can be deferred, and that some cross-regional computing tasks also possess potential for spatial migration. Energy storage, backup power, and cooling systems provide additional dimensions for load regulation.^{2–4} However, this flexibility is not unconditional, since tasks involving low-latency inference, mission-critical online services, and those bound by service-level agreements are generally difficult to shift substantially in time, while cross-regional migration remains constrained by network bandwidth, data-compliance requirements, and communication latency. Accordingly, the key issue for computing–electricity coordination is not to presume that computing load is inherently adjustable, but rather to identify which loads can be reliably regulated, over what time scales, and at what cost.

From Grid Infrastructure to Market Mechanisms

The conference featured the release of a cluster of achievements, including key technologies for superconducting-cable transmission, a national-level power emergency response system for the Guangdong–Hong Kong–Macao Greater Bay Area, and grid-forming energy storage. These

coordination proposed that computing resources should be made responsive to electricity-price and power supply–demand signals and should participate in demand response and ancillary services under unified rules. This layer determines whether the technologies described above can form a closed loop: the grid needs to be able to communicate local capacity, the degree of temporal tightness, and security constraints. the computing side needs to provide a verifiable load baseline, response speed, and duration. Settlement mechanisms need however to distinguish genuine regulation capability from nominal commitments. Without such information and rules, equipment capability alone does not necessarily translate into a dispatchable resource. As a load-side entity, the AI data center engages in bidirectional interactions of power, information, and value with the grid-infrastructure layer, the coordination-mechanism layer, and the collaborative-actor layer. The technical modules shown in figure 1 illustrate potential coordination relationships and do not imply that all modules have already been integrated into a unified engineering system.

The conference was held concurrently with the launch ceremony of the international journal *Energy Use*, which focuses on patterns of energy use and their coupling with energy systems, echoing the thematic focus of this year’s conference. The park-level Energy-Use Net practice reported by the

journal in 2025 shows clear continuity with the computing–electricity coordination topics addressed here: the park-level scenario is concerned with how distributed generation, energy storage, and end-use loads move from physical interconnection toward coordinated dispatch and value interaction, whereas computing–electricity coordination extends this object to artificial intelligence loads that are larger in scale, subject to more stringent reliability requirements, and more spatially concentrated. What the two share in common is not a specific set of equipment, but the broader shift of energy systems from supplying energy on demand toward jointly optimizing supply capability, network constraints, and energy-use behavior.

Effectiveness Must Be Validated by Operational Outcomes, Not by the Concept Itself

From the perspective of energy systems, whether computing–electricity coordination is effective must ultimately be validated by measurable operational outcomes. At least four categories of indicators warrant tracking: whether the interconnection cycle and connectable capacity for new computing loads improve; how much load can be reliably curtailed or shifted during peak periods; whether load shifting increases the utilization rate of renewable energy and reduces system operating costs or carbon emissions; and, under frequency, voltage, and emergency conditions, how quickly, for how long, and how repeatably data centers can provide response. Trade-offs may also exist among these objectives. For example, more aggressive peak-load shifting may increase task latency, while pursuing round-the-clock low-carbon power matching may raise the cost of energy supply and storage.^{4,5}

This also indicates that grid capacity expansion and load flexibility should not be regarded as mutually exclusive alternatives. Since mission-critical computing tasks that lack adjustability, sufficient generation and network capacity remain a basic prerequisite. For tasks with temporal or spatial elasticity, demand response can defer part of the capacity-expansion need, reduce peak pressure, and improve matching with renewable energy. Genuine computing–electricity coordination integrates “adding supply” and “managing demand” within a single planning and dispatch framework, rather than substituting the latter for the former.

The conference also included a parallel forum on “Space Manufacturing and Space Energy”. The lunar-surface energy systems presented by Zuankai Wang (associate Vice President of The Hong Kong Polytechnic University), Jianzhong Liu (Institute of Geochemistry, Chinese Academy of Sciences), and Guihua Tang (Xi’an Jiaotong University), among others, provide an extreme-case test of the computing–electricity coordination logic discussed above. On terrestrial grids, computing load still accounts for only a small share of total electricity consumption, leaving room to maneuver between capacity expansion and flexible dispatch. On a lunar base, this margin is compressed to zero. Because Earth–Moon transport costs are extremely high, in-situ resource utilization (ISRU) on the lunar surface becomes a prerequisite for long-duration habitation, and the energy system itself constitutes the scarcest infrastructure. The rate at which oxygen, silicon, and metals can be produced from lunar regolith oxides via molten electrolysis or high-temperature hydrogen reduction is directly constrained by the sustained heat-flux density and electric-power quality that the energy system can provide, with no backup grid available as a fallback. Recent studies identify in-situ resource utilization, energy storage across the extended lunar night, and

multi-energy system integration as core issues in the energy design of lunar bases. This pushes all three levels of the terrestrial computing–electricity coordination discussion to their limits simultaneously: at the infrastructure level, there is no room for capacity expansion, as the power envelope of photovoltaic arrays or small-scale concentrated solar-thermal reactors constitutes a fixed ceiling. Load-side flexibility is no longer optional but a condition for survival. Tasks such as ISRU, manufacturing, computation, and life support must share a single energy budget according to priority, and waste heat must likewise be used, through cascaded heat exchange, to preheat regolith or maintain habitat conditions. Market mechanisms are replaced by engineered dispatch rules, and although no price signal exists, load prioritization and real-time power reallocation perform a resource-allocation function analogous to that of terrestrial electricity-pricing mechanisms. This comparison suggests that what computing–electricity coordination must resolve, is not a problem unique to terrestrial grids, but that the more general energy-system proposition that “supply capability is limited, and load must participate in dispatch.”

Artificial intelligence is transforming data centers from conventional large consumers within the energy system into a new category of load that must be jointly addressed through grid planning, market mechanisms, and operational control. The Greater Bay Area Computing–Electricity Coordination Innovation Consortium has established a collaborative platform linking electric power and computing power systems, which have long remained relatively independent of one another. However, the establishment of the Consortium alone is not sufficient to demonstrate that computing–electricity coordination has become effective: its ultimate value will depend on whether adjustable computing capacity, together with cooling and energy-storage capabilities, can be converted into measurable, verifiable, and settleable grid resources, and on whether interconnection capacity, system security, renewable energy utilization, and energy costs can be improved while safeguarding the quality of computing services.

REFERENCES

1. International Energy Agency. (2025). Energy and AI. IEA.
2. Guo C., Luo F., Cai Z., Dong Z.Y. (2021). Integrated energy systems of data centers and smart grids: State-of-the-art and future opportunities. *Appl. Energy* **301**:117474. DOI:10.1016/j.apenergy.2021.117474
3. Radovanović A., Koningstein R., Schneider I., et al. (2023). Carbon-Aware Computing for Datacenters. *IEEE TPS* **38**:1270–1280. DOI:10.1109/TPWRS.2022.3173250
4. Riepin I., Brown T. and Zavala V.M. (2025). Spatio-temporal load shifting for truly clean computing. *Adv. Appl. Energy* **17**:100202. DOI:10.1016/j.adapen.2024.100202
5. Hasan M.M., Razmi D., Babayomi O., et al. (2025). Advanced control and protection strategies for grid-forming inverters in microgrids—A review. *Int. J. Electr. Power Energy Syst.* **172**:111297. DOI:10.1016/j.ijepes.2025.111297

FUNDING AND ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. U22B20112) and the 3xd Explorer prize.

DECLARATION OF INTERESTS

Xingying Chen is the Editor-in-Chief of Energy Use and was recused from reviewing or making final decisions on this manuscript. Peer review was handled independently of this member and their research group.