Press Release
For Immediate Release

HITACHI

Hitachi develops a new GPU cold plate design to reduce power consumption at next-generation AI data centers

Efficient GPU cooling and integration with the HMAX Data Center digital service help enable stable data center operations

Tokyo, October 21, 2026 – Hitachi, Ltd. ("Hitachi") has developed a GPU[*1] cold plate[*2] designed to reduce the increasing power consumption of data centers resulting from the widespread adoption of generative AI. At AI data centers, coolant circulating through cold plates captures heat generated by GPUs inside servers. The heat is then rejected to the outside air through pumps, chillers[*3], cooling towers and other equipment (Figure 1). As GPU heat generation increases and the number of installed GPUs grows, the load on the cooling equipment used for heat rejection also rises, leading to greater power consumption.

The prototype cold plate developed by Hitachi adopts a proprietary U-Tunnel channel structure that directs coolant toward the GPU's heat-generating area, unlike conventional structures in which coolant flows along the GPU surface (Figure 2). In this structure, several thousand U-Tunnel channels are arranged in parallel, enabling rapid heat removal in the vicinity of the GPU. After absorbing heat at the heat-generating area, the coolant makes a U-turn and flows toward the outlet. This prevents the cold incoming coolant from mixing with the heated outgoing coolant, thereby maintaining high cooling performance.

This structure enables efficient GPU cooling even with coolant at a higher temperature than in conventional cold plates. This allows the coolant temperature to be kept higher than the outdoor air temperature, enabling more efficient heat rejection to the outside air.
As a result, Hitachi estimates that the technology can reduce the power required for cooling by 30%.[*4]

Hitachi will further develop the cold plate as one of the cooling technologies that help bring out the benefits of its HMAX Data Center[*5] digital service by combining it with integrated cooling system controls. Through greater data center efficiency and stable operations, Hitachi will contribute to building sustainable digital infrastructure.

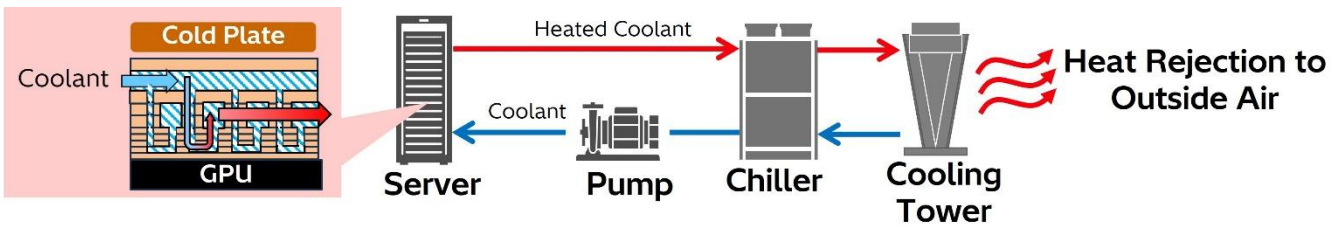


Figure 1: Overview of cooling at an AI data center (illustration)

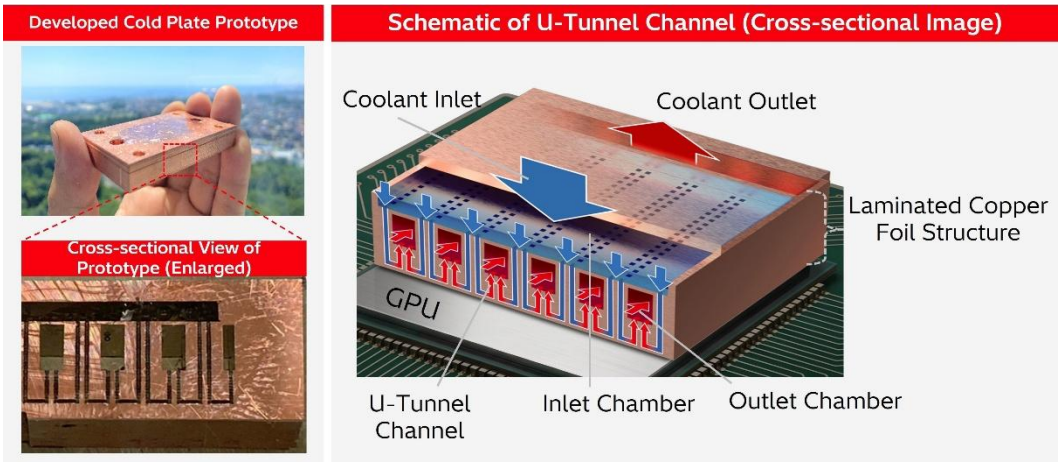


Figure 2: Prototype GPU cold plate and schematic of the proprietary U-Tunnel channels

*1 GPU: Abbreviation for graphics processing unit, a parallel computing processor particularly suited to repeatedly performing the large-scale matrix calculations required for AI training.
*2 Cold plate: A cooling component placed in contact with a heat-generating device, such as a semiconductor, to capture heat using coolant circulating inside the component.
*3 Chiller: Equipment that uses a refrigeration cycle to capture heat from coolant and release it externally.
*4 The estimate was calculated based on assumed conditions for a data center in a specified region, including outdoor temperature, and changes in the operating rate of equipment used to facilitate heat rejection to the atmosphere. Results vary depending on conditions.
*5 Hitachi launches HMAX Data Center, a suite of solutions for integrated data center infrastructure operations : Hitachi Global

Background and issues

The rapid expansion of generative AI has increased the number of data centers equipped with large numbers of GPUs. As a result, GPU cooling has become a key factor affecting both operational quality and power efficiency.

At data centers, coolant circulating through cold plates captures heat generated by GPUs. Chillers, cooling towers and other equipment then reject the heat to the outside air. Consequently, substantial amounts of power are consumed not only by GPU computations but also by cooling equipment.

In recent years, expectations have grown for warm water cooling using coolant at approximately 45 degrees Celsius. Higher-temperature coolant can reduce cooling power consumption by enabling efficient heat rejection using the temperature difference between the coolant and the outside air. However, warm water cooling requires the coolant to reach the heat-generating area before becoming excessively warm inside the cold plate.

Conventional cold plates generally use a structure in which coolant flows along the surface of the GPU. Because the coolant temperature increases as it moves through the latter sections of the channel, it is difficult to supply lower-temperature coolant uniformly to the heat-generating area.

Capturing more heat from a GPU also requires channels to be divided into smaller passages to increase the heat-transfer surface area. However, smaller channels impede coolant flow and increase pressure drop,[*6] resulting in greater pump power requirements for coolant circulation.

*6 Pressure drop: A decrease in pressure that occurs as a liquid flows through a channel. Greater pressure drop requires more power to circulate the liquid.

Features of the technology developed to solve these issues

Hitachi developed the GPU cold plate to reduce data center power consumption by suppressing pressure drop while maintaining the cooling performance required for warm water cooling. The principal features of the cold plate are as follows.

1. U-Tunnel channels that deliver coolant directly to the vicinity of the GPU's heat-generating area

The developed U-Tunnel channels direct coolant supplied at the inlet toward the vicinity of the GPU's heat-generating area. After reaching this area and absorbing heat, the coolant makes a U-turn and exits through the channel outlet.

Separating the incoming coolant from the heated outgoing coolant suppresses the decline in cooling performance caused by coolant mixing, enabling efficient GPU cooling.

2. Several thousand U-Tunnel channels arranged in parallel to achieve both high cooling performance and low pressure drop

Instead of directing coolant along the surface of the GPU, as in conventional structures, the new design directs coolant through the thickness of the cold plate toward the GPU's heat-generating area (Figure 3).

Shortening each channel reduces the friction experienced by the coolant as it flows through the channel, thereby reducing pressure drop. In addition, arranging several thousand U-Tunnel channels in parallel provides a greater total channel cross-sectional area within the limited dimensions of the cold plate than conventional structures.

This design makes it easier for coolant to flow while achieving both high cooling performance and low pressure drop.

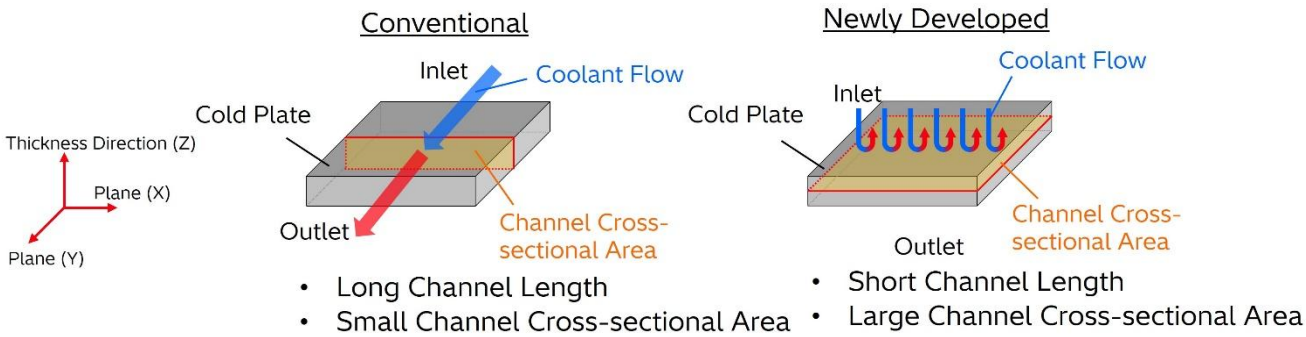


Figure 3: Comparison of the conventional structure and the newly developed structure

Verified results

Measurements using the prototype cold plate demonstrated a thermal resistance of 0.011 K/W[*7], [*8] and a pressure drop of 27.6 kPa at a coolant flow rate of 4 L/min. The results confirmed that the prototype provides higher cooling performance than a conventional structure while using the same cooling power.[*9]

This technology enables efficient heat removal from the GPU even when using coolant at a higher temperature than in conventional systems. This allows the coolant temperature to be kept higher than the outdoor air temperature, thereby extending the period during which outdoor air can be actively used for heat rejection throughout the year.

Using outdoor temperature conditions in Tokyo and other representative cities, Hitachi estimated changes in the annual chiller operating rate. The results indicate that the technology could reduce the power required to reject heat captured from GPUs to the outside air by 30%.

*7 Thermal resistance: A measure of resistance to heat transfer. A lower value indicates better cooling performance. In this news release, thermal resistance is defined as the temperature difference between the surface of the heat-generating component and the coolant outlet, divided by the amount of heat generated.
*8 The cooling performance target assumes GPU heat generation of several kilowatts.
*9 Cooling power: The product of coolant flow rate and pressure drop.

Looking ahead

Starting with the GPU cold plate, Hitachi will work to improve the performance of the entire cooling infrastructure, including pumps, chillers and piping. Hitachi will also integrate this cooling infrastructure with its HMAX Data Center digital service and further develop it as one of the technologies for improving operational efficiency.

Under the One Hitachi approach, the company will combine its strengths in IT, OT, or operational technology, and products with data and digital technologies that optimize equipment operations. Through these efforts, Hitachi will pursue both higher performance and greater efficiency at data centers and contribute to building sustainable digital infrastructure.

A portion of these results is scheduled to be presented at InterPACK 2026, organized by the American Society of Mechanical Engineers, or ASME, in San Diego from October 26 to 29, 2026.

Trademark Notice: All trademarks and product names are the property of their respective owners.

About Hitachi, Ltd.
Through its Social Innovation Business (SIB) that brings together IT, OT(Operational Technology) and products, Hitachi aims to be a global leader in continuously transforming social infrastructure through digital, contributing to a harmonized society where the environment, wellbeing, and economic growth are in balance.
Hitachi operates worldwide across four sectors – Digital Systems & Services, Energy, Mobility, and Connective Industries – as well as a Strategic SIB Business Unit focused on new growth areas. With Lumada at its core, Hitachi creates value by combining data, technology and domain knowledge to solve customer and social challenges. Revenues for FY2025 (ended March 31, 2026) totaled 10,586.7 billion yen, with 606 consolidated subsidiaries and approximately 290,000 employees worldwide. Visit us at www.hitachi.com.

Contact
Hitachi, Ltd.
Research & Development Group
https://www8.hitachi.co.jp/inquiry/hitachi-ltd/hqrd/news/en/form.jsp