Data Center Power Savings to Achieve Carbon Neutrality

NOBUHIDE HARA[*1] NAOHIKO MATSUDA[*1]
YUSUKE ISOBE[*2] FUMITO KAJITANI[*1]
NAOKI INOUE[*3] HIROKO KITAMOTO[*3]

An increasing number of data centers have been developed in line with the progress of digitalization in recent years. It is an urgent issue to reduce the power consumption of non-IT facilities and significantly improve power usage effectiveness (PUE) in the trend toward carbon neutrality. As such, Mitsubishi Heavy Industries, Ltd. (MHI) participated in the development of a containerized data center with an immersion cooling system that achieves compactness and low PUE for use in multi-access edge computing (MEC) and installed it in Yokohama Hardtech Hub. As a result, PUE of less than 1.07 was achieved. Moving forward, we will propose a system that optimizes the entire system, including a hyperscale data center and progress with the development to meet the needs of our customers.*

* PUE is total power in the data center divided by IT power, and the smaller the value, the better the power savings.

1. Introduction

With the progress of global digitization, the amount of data in 2025 is expected to be approximately 150 times [(1)] greater than that in 2010 and the development of digital infrastructure such as data centers is being promoted.

Figure 1 shows changes of the power consumption of data centers as a percentage of global power consumption. As of 2018, data centers consumed power equivalent to about 1% of total power generation, and the percentage is expected to increase to about 8% by 2030. **Figure 2** shows the breakdown of power consumption in data centers. Electricity provided to servers accounts for the largest share at 57%, and cooling, at 31%, is second to this. As such, in order to achieve carbon neutrality as an equipment manufacturer and realize a sustainable society, efforts to reduce power consumption for cooling and other purposes are necessary.

Regarding chips such as CPUs and GPUs installed on server boards, it is said that products with heat generation exceeding 1,000 W per chip will be distributed in the market in the near future. One of the current products, NVIDIA GeForce RTX3090Ti has a heat flux of about 700 kW/m^2 but cooling performance needs be improved to handle future increases in heat generation. Conventional systems cool the server room and blowing air around chips to cool them. However, when the heat generation per chip exceeds 1,000 W, the cooling capacity of this air-cooling method is not sufficient, and as a result, the allowable temperature of the chip could be exceeded. Therefore, it is expected that local cooling methods that intensively cool the heat-generating parts will be adopted in an increasing number of cases. The following three local cooling methods are being considered.

(1) Rear door cooling

This system has cooling coils and fans in the rear section of the server rack and provides cooling by forcibly exhausting the air in the rack, cooling the exhausted air with the cooling coils, and recirculating the air.

(2) Direct chip cooling

This method provides cooling by supplying refrigerant to the cooling jacket of a cooling

*1 Chef Staff Researcher, Heat Transfer Research Department, Research & Innovation Center, Mitsubishi Heavy Industries, Ltd.
*2 Chef Staff Manager, Business Intelligence & Innovation Department, Technology Strategy Office, Mitsubishi Heavy Industries, Ltd.
*3 Heat Transfer Research Department, Research & Innovation Center, Mitsubishi Heavy Industries, Ltd.

device called a cold plate installed in the heat-generating section of the CPU or GPU. In some cases, cooling is provided by allowing the supplied refrigerant to flow in a single-phase (liquid) state, while in other cases, cooling is provided by phase-changing (boiling) the refrigerant from liquid to gas.

(3) Immersion cooling

This method provides cooling with refrigerant circulated through a tank in which the server is immersed. As in the direct chip cooling method, cooling is provided by allowing the supplied refrigerant to flow in a single-phase (liquid) state in some cases, while cooling is provided by phase-changing the refrigerant from liquid to gas in other cases.

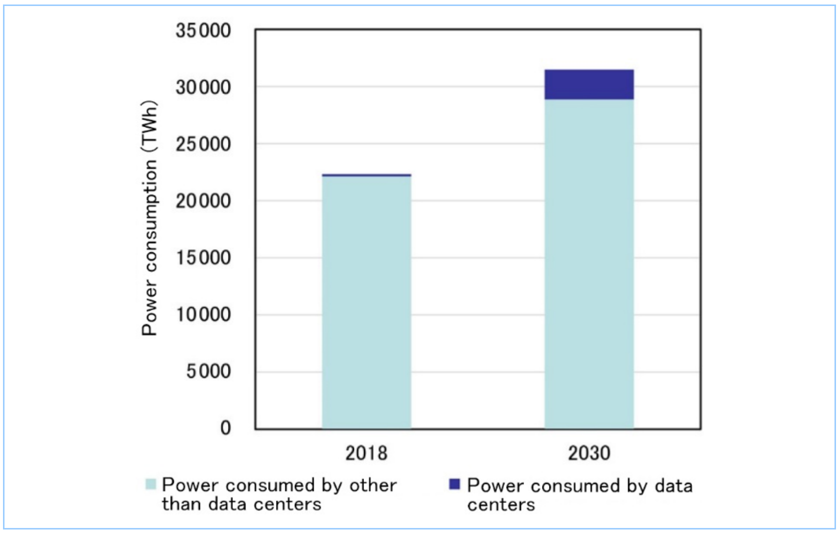

Figure 1 Prediction of changes in power consumption of data centers with respect to global power consumption

Prepared by the authors with data on the forecast of the transition of data center power consumption relative to global power consumption cited from the following.

(1) National Institute of Science and Technology Agency, Low Carbon Society Strategy Center, Impact of the Development of Information Society on Energy Consumption (Vol. 4) - Feasibility Study on Technologies to Reduce Data Center Power Consumption -, 2022, Forecast data in the "As is" columns for 2018 and 2030 in Table 8., (https://www.jst.go.jp/lcs/pdf/fy2021-pp-01.pdf)

(2) Ministry of Economy, Trade and Industry, Agency for Natural Resources and Energy website, [223-1-1] Electricity Consumption in the World, Electricity consumption data for 2018 (cited on August 18, 2022), (https://www.enecho.meti.go.jp/about/whitepaper/2021/html/2-2-3.html)

(3) Japan Nuclear Energy Industries Association, Information and Communication Department, World Energy Outlook 2020 (WEO2020), Summary Introduction (with focus on electricity and nuclear power), 2020, p. 19, Sustainable development scenario case in 2030, (https://www.jaif.or.jp/cms_admin/wp-content/uploads/2020/12/weo_2020)

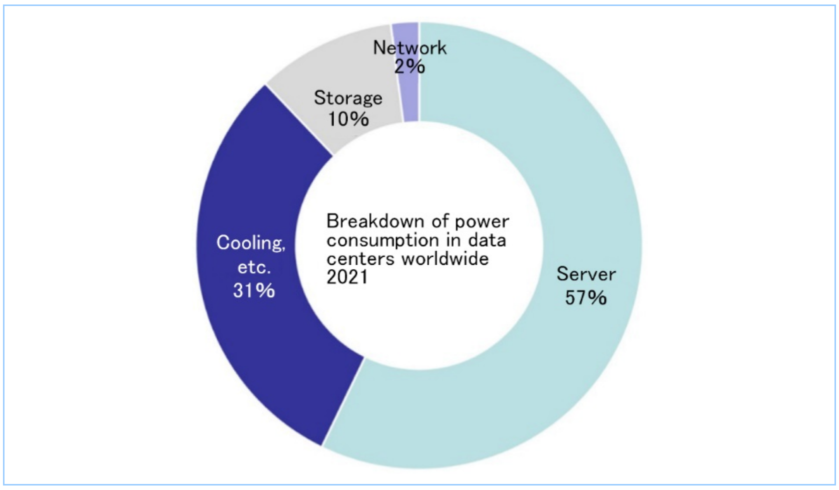

Figure 2 Breakdown of power consumption in data centers

Prepared by the authors with data on the forecast of the breakdown of power consumption in data centers cited from the following IEA analytical data., (https://www.iea.org/commentaries/data-centres-and-energy-from-global-headlines-to-local-headaches)

This report describes the development status of a containerized data center with an immersion cooling system, which is currently under development as part of the MHI Group's efforts to achieve a carbon-neutral society.

2. PoC (Proof of Concept) of containerized data center

2.1 Background and overview of PoC

As mentioned above, in recent years, servers installed in data centers tend to be equipped with highly integrated and high-heat-generating chips such as GPUs, and it has become apparent that conventional cooling methods using air from air conditioners cannot realize sufficient heat removal.

As power consumption in data centers is expected to dramatically increase, in addition to the trend toward carbon neutrality, there is an urgent need to reduce the power consumption of the entire data center, especially that of non-IT equipment, and to significantly improve PUE.

Furthermore, as the scale of equipment is rapidly expanding with the increase in opportunities for server utilization and large-scale data centers, known as hyper-scale data centers, are being constructed, there is also concern about the shortage of suitable sites near large-scale demand sites in the near future in urban areas.

Against this market trend, MHI focused on the combination of a single-phase immersion cooling method (using insulating liquid oil) of servers and free cooling which removes the heat of the liquid with ambient air, and participated in the development and PoC of a containerized data center with compact size and low PUE for use in MEC in the future, together with KDDI Corporation and NEC Networks & System Integration Corporation.

Figure 3 shows the cooling method and equipment configuration presented in this report. By using liquid (oil), which has a higher heat transfer coefficient than gas (air), the ambient air temperature of the server may be higher than that of the conventional general air-cooling method, making it possible to cool the server even in ambient air above 40°C, for example. As a result, a heat pump chiller that produces low-temperature chilled water is not required in principle and PUE can be reduced by the amount of power for compressing the refrigerant inside the chiller, which is the most significant feature of the system.

We were mainly responsible for the development and prototyping of the ambient air-cooling system and the design and construction of the immersion cooling system and the container that contains the system, and worked on a technical solution from the energy demand side perspective, utilizing our knowledge and expertise in layout and space design, which we have cultivated in the cooling design of refrigeration equipment and power generation plants.

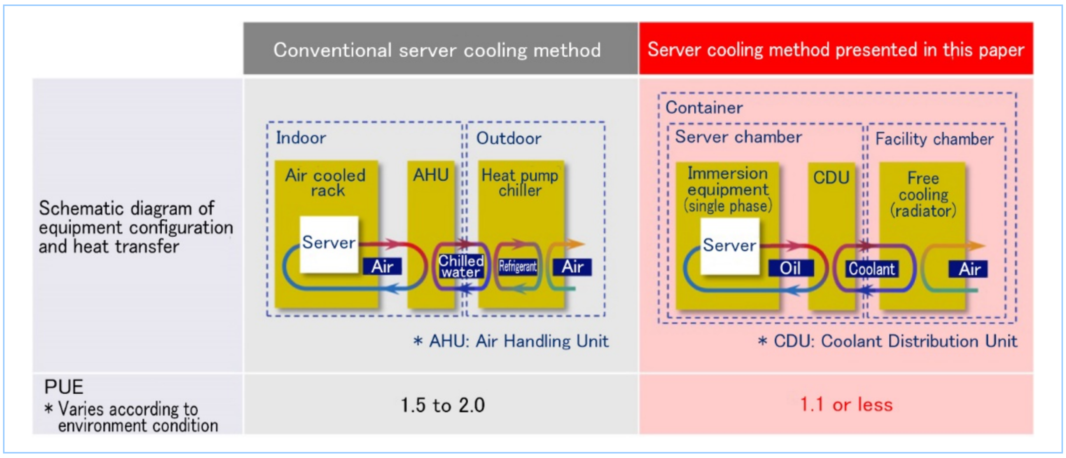

Figure 3　Server cooling method

2.2 MHI's role in PoC

In conducting the PoC, participating companies brought their own technology and products to create a PoC model of the containerized data center. We played a role mainly in the following two items, contributing to the verification of the actual prototype.

(1) Free cooling equipment (radiator)

MHI designed and manufactured a free cooling (radiator) device and its peripherals, which have the function of removing the heat of the coolant sourced from heat exchange with the oil for cooling the server by using ambient air applying our refrigeration design technology. The device was used a large area fin-and-tube heat exchanger to aiming to reduce PUE across the immersion equipment system and to cope with the Japanese climate, especially in summer, so that the server can be cooled even at the maximum ambient air temperature of 42°C. This contributes to keeping the PUE value as low as possible even in high-temperature environments. Note, the redundancy required for data centers is assumed to be ensured by the MEC network (as data processing is secured by multiple containers) and the redundancy of individual containers was not verified in this PoC experiment.

(2) Design and construction of container

The notable point of the container design and constructing for data center in this test was the success of storing all data center functions in a 12-foot container, including a system with a total server capacity equivalent to 50 kVA. While there have been cases where conventional server cooling systems have been used and stored in a 20-foot container or the like, however, we assumed MEC and aimed for the smallest possible footprint in this test.PoC In addition, the equipment layout was designed to allow easy server maintenance including loading and unloading of the servers, taking into account the flow lines of workers and ventilation in each chamber. The following lists the points for which we made special efforts in the layout.

- Minimization of facility chamber space by sharing its corner dead space with that of server chamber
- Downsizing of fan for free cooling equipment (radiator)
- Integration of heat exchanger for free cooling equipment (radiator) and panel door
- Consolidation of air exhaust ports for the server and facility chambers
- Considering space for maintenance operation of immersion equipment such as loading/unloading of servers and worker flow lines

Figure 4 shows a schematic structure of the containerized data center.

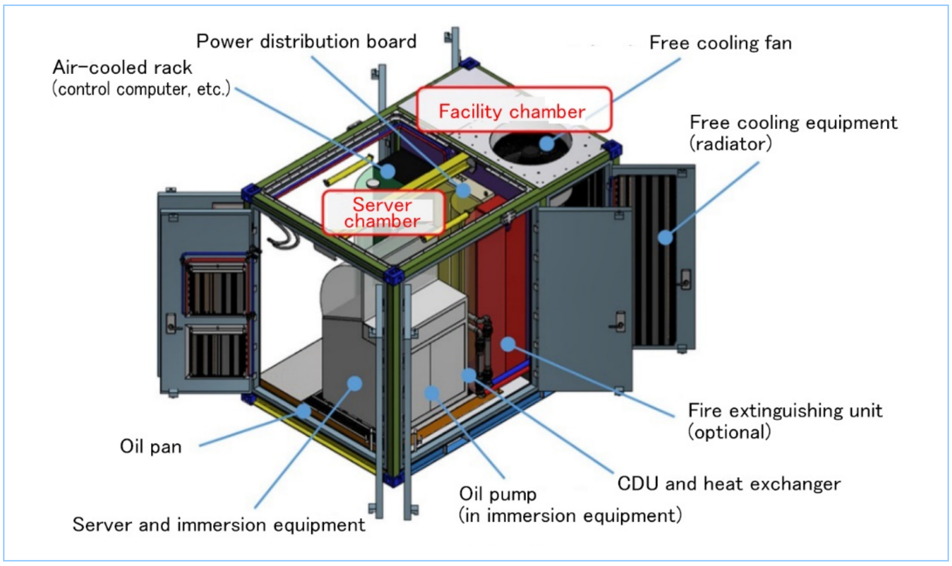

Figure 4 Structure of containerized data center

2.3 Content and result of PoC test

The containerized data center test equipment described in the previous section was installed in the co-creation facility: Yokohama Hardtech Hub, the co-creation facility located in the Honmoku Plant of our Yokohama Works and test was carried out to operate IT equipment such as servers and immersion cooling system equipment equivalent to a total of 24 U/50 kVA. **Figure 5** shows the installed PoC equipment. Due to containerization, the on-site installation and power connection work were carried out in a single day.

After installation, various test settings were made to verify the cooling performance of the server and it was and measured the PUE for a certain period of time to perform verification over about eight months. The main verification contents are as follows.

Figure 5 PoC equipment installed at Yokohama Hardtech Hub

(1) Verification of heat removal performance

The electric heater with a power total of 50 kVA was installed inside the immersion equipment instead of real servers to apply a heat load of that power output to the immersion cooling system to perform verification of the heat transfer from the server section, through the coolant distribution unit (CDU) and the free cooling section (radiator), to the ambient air, as shown in **Figure 6**. As a result, it was confirmed that all the server heat of 50 kVA in total was released into the ambient air, the oil temperature was below the upper-temperature limit of the server and verified that the entire immersion cooling system operated soundly.

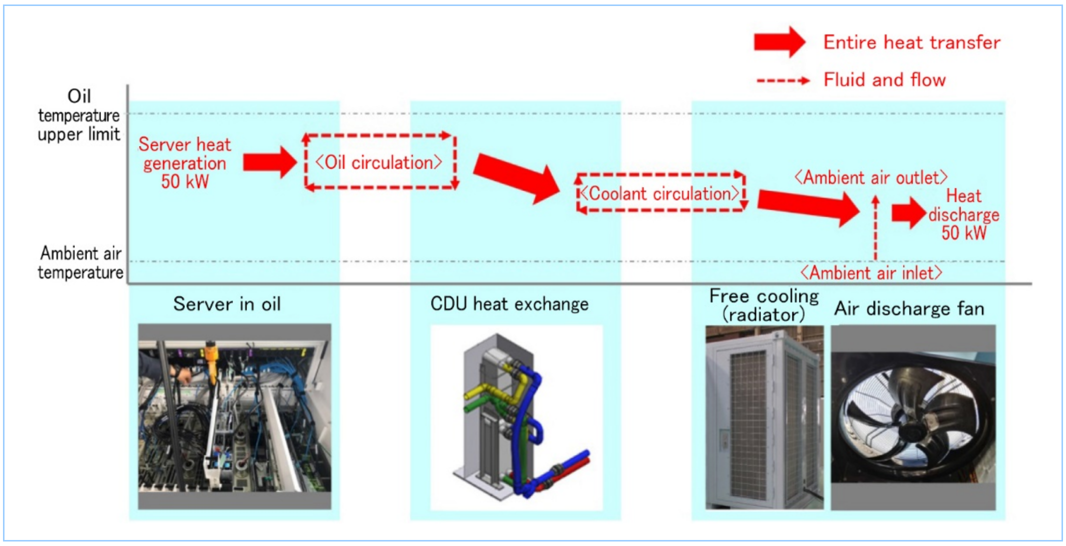

Figure 6 Temperature map of heat removal test

(2) Verification of power-saving operation (low-PUE operation)

A total of 50 kVA of heating elements were installed in the immersion equipment by a real server and an electric heater in the outdoor environment of the Yokohama Hardtech Hub, and measured various power consumptions when performing simulated arithmetic processing. As a result, the PUE of 1.07 or less was achieved compared to the planned target PUE value of 1.1 or less. By conducting multiple tests and measuring PUE values, it was confirmed the reproducibility of this result.

In this test, the oil flow rate was controlled to keep the oil temperature within the appropriate range and the heat transfer area of the radiator was increased to ensure robustness. Thus, the difference between the oil temperature and the ambient air temperature could be kept almost constant and therefore, no significant changes in airflow and coolant flow rates were required and the PUE value is expected to remain constant throughout the year. As a further improvement, the PUE value can be further reduced by actively controlling and adjusting the power output of cooling equipment when the outdoor air temperature is low. In anticipation of the commercialization of containerized data centers, we will aim for a more usable immersion cooling system by providing multiple modes according to environmental conditions and the user's operation policy, as shown in **Figure 7**. The oil temperature priority mode changes the fluid flow rate under constant conditions, resulting in an almost constant PUE throughout the year. The PUE (power saving) priority mode sets the oil temperature to the maximum allowable temperature and controls the fluid flow rate according to the ambient air temperature (e.g., air flow rate can be reduced when the ambient air temperature is low), thus achieving a lower PUE.

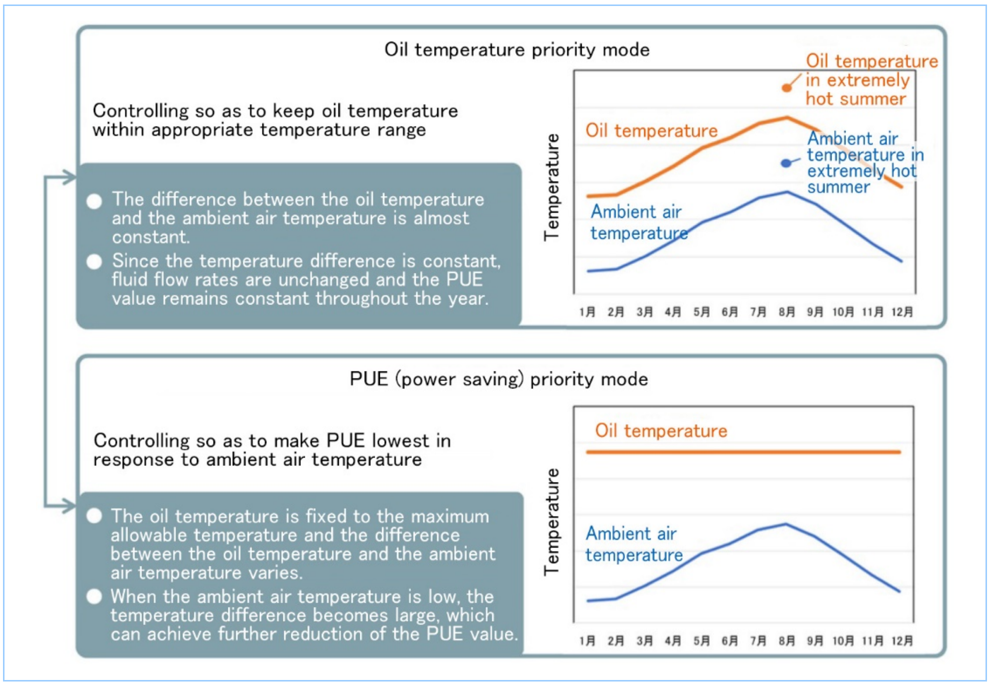

Figure 7 Control mode suitable for immersion cooling system

2.4 Development of immersion cooling system

In this PoC, we focused on the combination of a single-phase immersion cooling method and free cooling to reduce PUE as much as possible in a data center. In addition, with a view to future use in MEC, we accommodated a complete set of data center functions in a small 12-foot container and verified its usability using actual equipment.

As mentioned above, the technology to control the flow rate of each fluid used in the immersion cooling system according to the ambient air temperature to reduce PUE is applicable to all data centers, including hyperscale data centers, not limited to containerized data centers. Moving forward, we are planning to realize the commercial introduction of this technology after verifying the stability, redundancy and high availability of the immersion cooling system for the horizontal deployment of the immersion cooling system for hyperscale data centers.

To make it possible to utilize the cooling system for containerized data centers for society implementation as soon as the market of their use in MEC is established, we intend to further upgrade the functionality and usability in addition to the redundancy required by data centers.

3. Future prospect

Figure 8 shows an overview of hyperscale data center components. Various technologies such as generators, UPS, server racks, cooling equipment, and software to control them, as well as their integration, are required, and many of the components of the data center are related to MHI Group. In addition, the hyperscale data center can use renewable energy sources as its power source. Since the power demand of data centers is expected to further increase in the future, we intend to develop a data center that contributes to the achievement of carbon neutrality by optimizing the entire system and realizing low PUE.

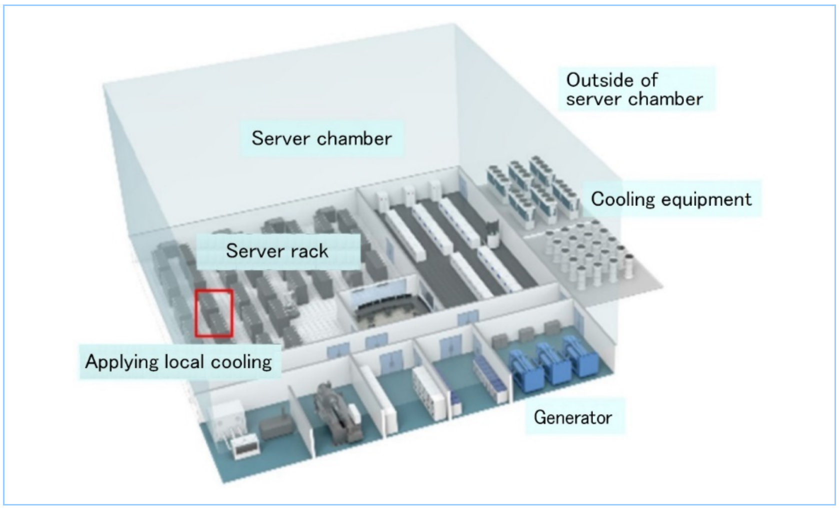

Figure 8　Overview of hyperscale data center components
This figure shows an overview of hyperscale data center components. Various technologies such as generators, UPS, server racks, air conditioners, cooler and software to control them are required, and many of the components are related to MHI Group.

References

1. United States International Trade Commission, Data Centers Around the World: A Quick Look, (2022), (https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf)
2. Techpowerup Website, (https://www.techpowerup.com/gpu-specs/geforce-rtx-3090-ti.c3829)