IDC Market Impact Study

OCP Global Summit 2025

Chris Drake

Senior Research Director
Compute Infrastructure and Service Provider Trends

November 12, 2025

Overview

1. OCP ecosystem growth & development
2. Server & rack project initiatives
3. IDC's growth expectations for OCP infrastructure
4. On-premises compute bottlenecks
5. Preparing datacenter networking for AI
6. Power consumption growth & challenges
7. Adopting advanced datacenter cooling
8. Ensuring sustainable datacenter growth
9. Takeaways & Predictions

The OCP community and vendor ecosystem continues to expand and mature

OCP community contributions by year

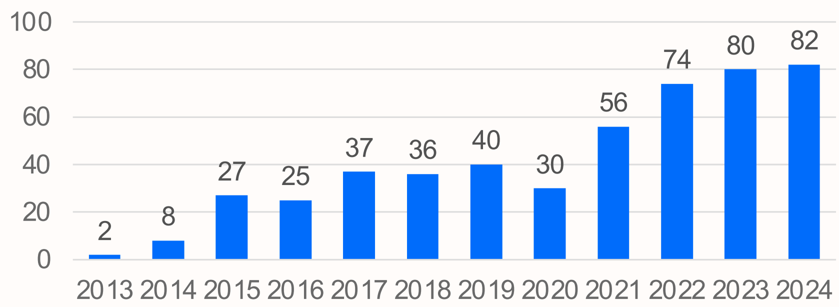

22%

Spending growth on OCP-recognized infrastructure from 2025 to 2029 (CAGR)

OCP corporate members, by region

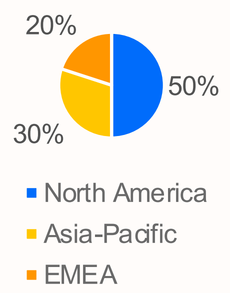

- The Open Compute Project was founded in 2011 by **Facebook**, **Intel**, and **Rackspace**, with **Goldman Sachs** and **Andy Bechtolsheim** joining soon after, to apply open-source principles to datacenter hardware, including servers, storage appliances, scalable rack designs, and other datacenter technologies.
- The OCP is now a large, diverse, and globally distributed community that includes just **over 500 corporate members**, nearly 200 work efforts and thousands of active individual engineers.
- Although initially focused on large hyperscale datacenter operators and their vendors, OCP membership has expanded to include organizations from the telecoms, financial services, manufacturing, and oil and gas industries. Its focus has also expanded beyond "greenfield" new build datacenters to include "brownfield" datacenters (i.e., retrofits and refurbishments).
- In addition to contributing new technology specifications to the OCP Community, many members are **important end users** of OCP-recognized technologies themselves.
- Use of OCP-recognized infrastructure and technology designs goes beyond the organization's formal membership. Worldwide use of OCP-recognized IT infrastructure continues to grow thanks partly to the way hardware designs are **openly shared and available** to any organization that wants to access and implement them.
- The **ecosystem of solution providers** offering OCP-recognized infrastructure continues to grow and develop, encouraging adoption.

Source: IDC

Server and rack project initiatives prioritize interoperability, scalability, and efficiency

The OCP's Server Project provides standardized specifications for compute systems with a scope that covers server chassis and sleds with CPU and GPU, server components and peripherals, mezzanine I/O, hardware accelerators, networking, and electrical and mechanical interfaces

- **Driving server platform interoperability, integration, and energy efficiency**
  - Over time, the priorities of the OCP Community's focus on servers and rack systems have evolved from improving interoperability, integration, and energy efficiency, to preparing server platforms for the growing demands of AI. This includes a strong focus on sustainable growth.
- **Advancing Open Rack capabilities and compatibility**
  - The introduction of **Open Rack v1** (**ORv1**) in 2013 brought numerous advantages over traditional 19-inch racks, including increased energy efficiency, improved installation and maintenance, and a standardized power shelf design for enhanced interoperability and flexibility. Subsequent releases of the Open Rack specification have focused on improving rack design and space utilization, enhancing power capacity, management and efficiency, and ensuring compatibility with new technologies (e.g., liquid cooling). A key consideration in the development of **ORv3** has been the need to prepare server platforms and rack architectures for the rapidly evolving demands of AI.
- **Moving towards an open chiplet economy**
  - The OCP's commitment to ensuring that compute platforms can support diverse, specialized and evolving workload requirements also influence its work on chiplet infrastructure and the wider chiplet market. Here the focus has been on improving customization and performance, enhancing scalability and deployment flexibility, and boosting cost efficiency and manufacturing processes.

Compute platform investments are driving spending across all OCP-recognized infrastructure categories

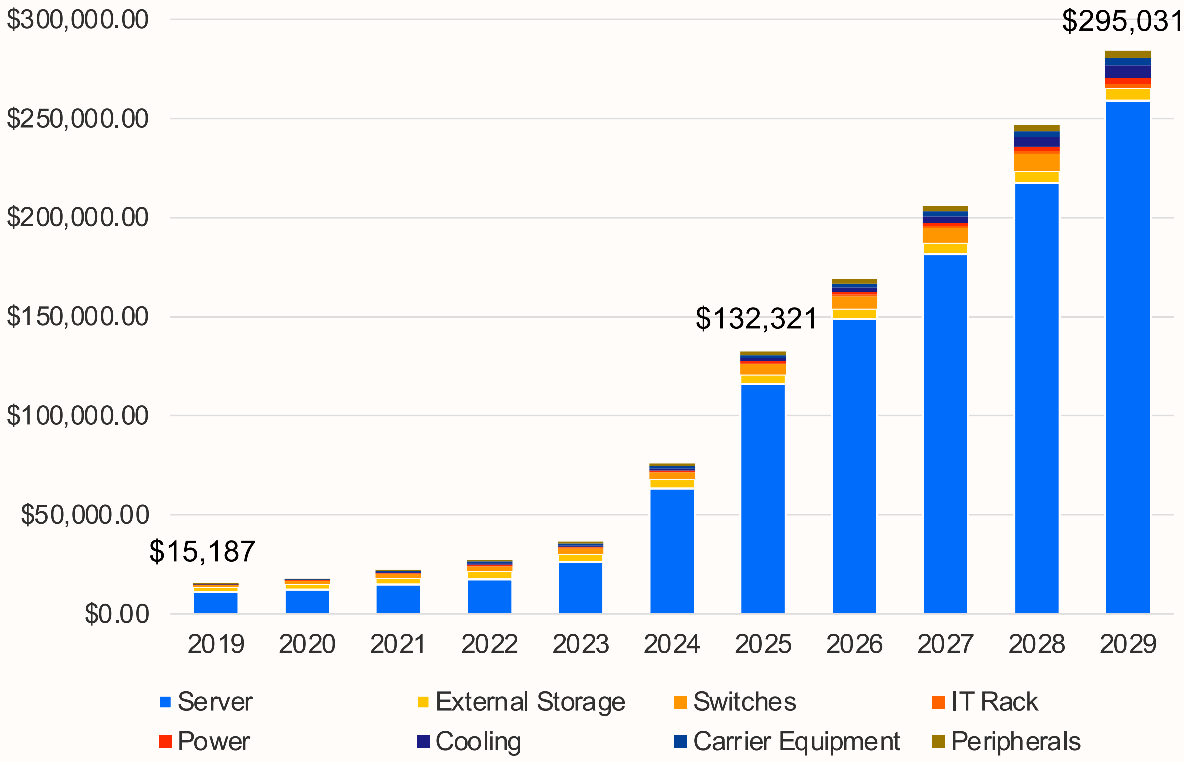

Source: IDC, 2025

- **Total spending** on OCP-recognized IT infrastructure and solutions will grow to US$295 billion in 2029, up from US$132 billion in 2025, reflecting a compound annual growth rate (CAGR) of 22.2%.
- **Server spending** to rise to $$258.9 billion in 2029 (CAGR = 35.7%).
- **External storage** and **switching** equipment represent the next largest segments of the market.
- Categories with the strongest spending growth include **cooling** technologies (CAGR of over 32%), **power** technologies (CAGR of 23.6%) and IT rack infrastructure (CAGR of 23.5%).

Investments in GPU-accelerated AI servers are a key contributor to AI infrastructure spending growth

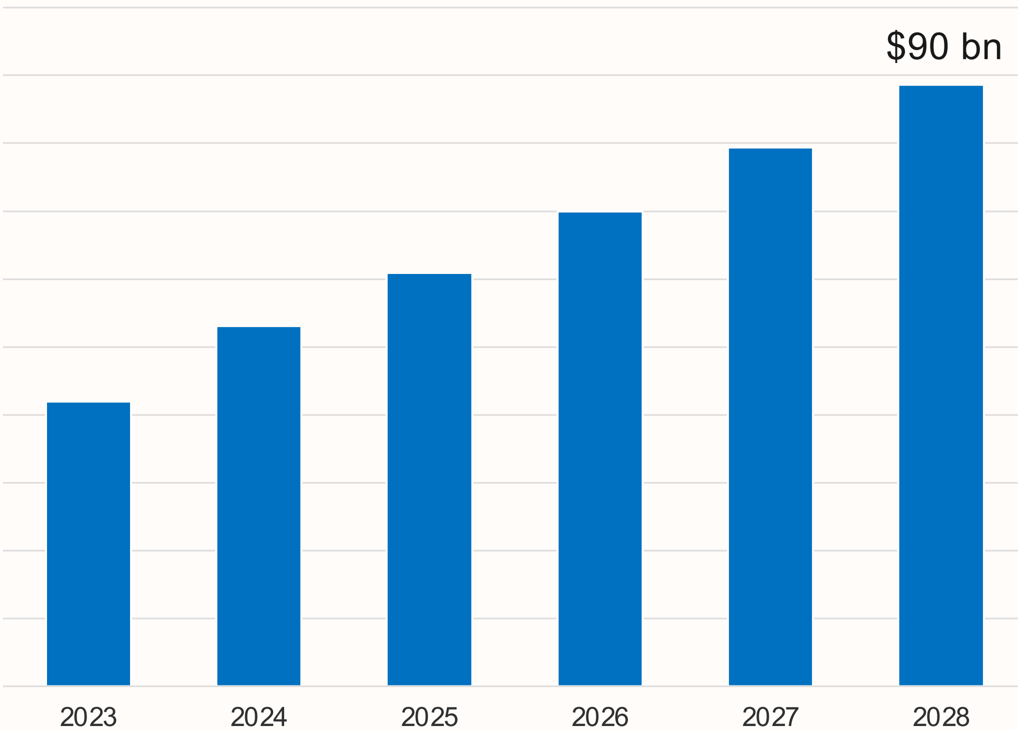

Source: IDC, 2024, Worldwide Artificial Intelligence Infrastructure Forecast, 2023–2028

- Strong growth is envisaged for the deployment of AI infrastructure by **enterprises and service providers**, with total spending expected to reach approximately $90 billion by 2028.
- One of the biggest contributors to this growth will be the deployment of **GPU-accelerated AI servers** by the **hyperscalers** and other organizations.

A major focus for the OCP includes driving standards for AI accelerators and AI clusters

OCP projects and initiatives such as Open Systems for AI and Open Cluster Designs for AI aim to drive standardizations that will help accelerate the market for future generations of AI clusters and the datacenters that host them

Developing a standardized, modular system for AI accelerators

- The Open Accelerator Infrastructure is a sub-project within the OCP's Server Project that aims to create a standardized, modular system for AI accelerators. It was announced in March 2019, with **Facebook**, **Microsoft**, and **Baidu** contributing the OCP Accelerator Module (OAM) design specification to define a common form factor for compute accelerator modules. Key objectives include the simplification of high-speed communication links between accelerators.

Investing in a multi-vendor supply chain for AI cluster infrastructure

- The Open Systems for AI Strategic Initiative aims to develop standardizations for AI clusters and the datacenter facilities that host them, delivering a complete open data center ecosystem. Goals include enhancing the efficiency and sustainability of AI clusters and datacenters and enabling the development of a multi-vendor supply chain. Contributors include **NVIDIA** with its GB200-NVL72 rack and **Meta** with its Catalina and Clemente high-density, liquid-cooled rack system.

Addressing the challenges of running AI clusters at scale

- With its emphasis on enhancing efficiency and sustainability, the Open Systems for AI Initiative addresses the various challenges associated with running AI clusters at scale. These include the need for specialized silicon, higher bandwidth and lower latency interconnects.

On-premises compute bottlenecks are widespread and include CPU speeds, memory, and network limitations

Q. What is the greatest resource bottleneck or limitation for your on-premises compute/server infrastructure?

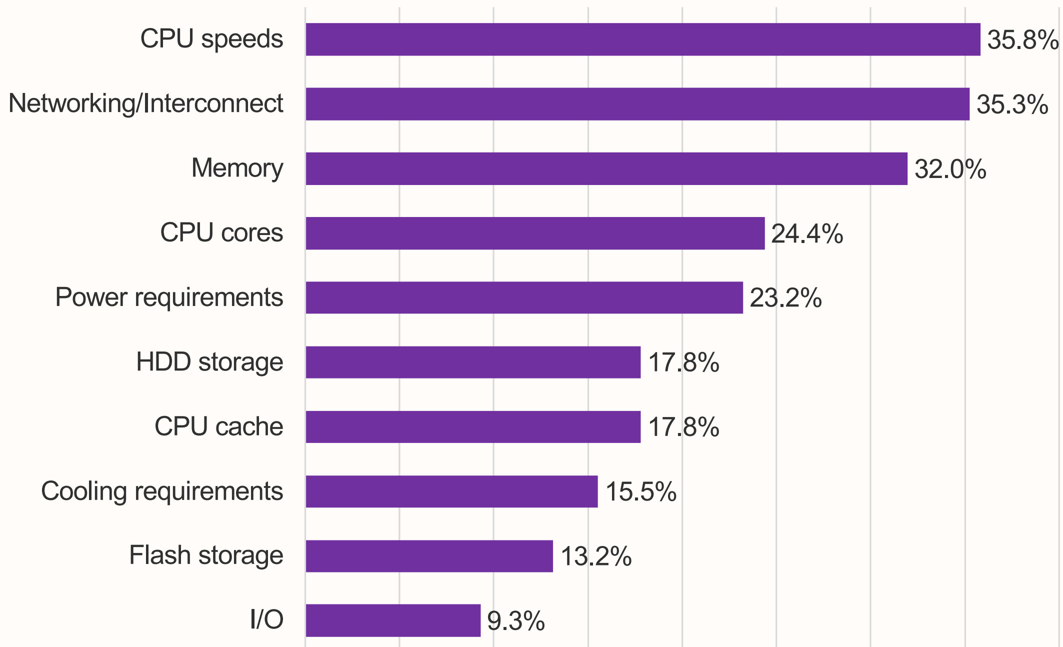

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 882

- More than one third of organizations identify **CPU speeds** as the biggest resource bottleneck or limitation for their on-premises compute infrastructure. CPU speeds can be affected by different factors, including transmission delays and heat accumulation.
- Other cited challenges include **memory limitations**, **CPU cores**, and challenges related to **networking/interconnect** and **power** and **cooling** requirements.
- **Server performance** can be undermined by network latency while memory limitations can negatively affect server performance by contributing to slower response times, reduced scalability, and higher levels of power consumption.

Networking Technologies

Ineffective server network interface technology continues to drive multiple challenges and responses

→ **Q. In the past 2 years, what challenges has your organization faced in relation to its server network interface technology?**

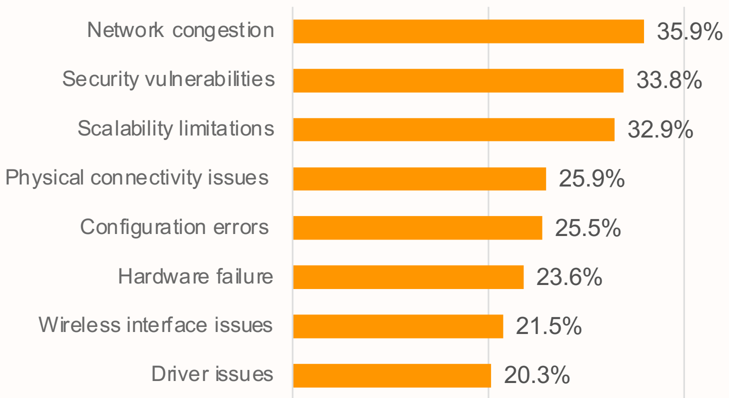

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 882

→ **Q. With which technologies and strategies has your organization addressed its server network interface challenges in the past 2 years?**

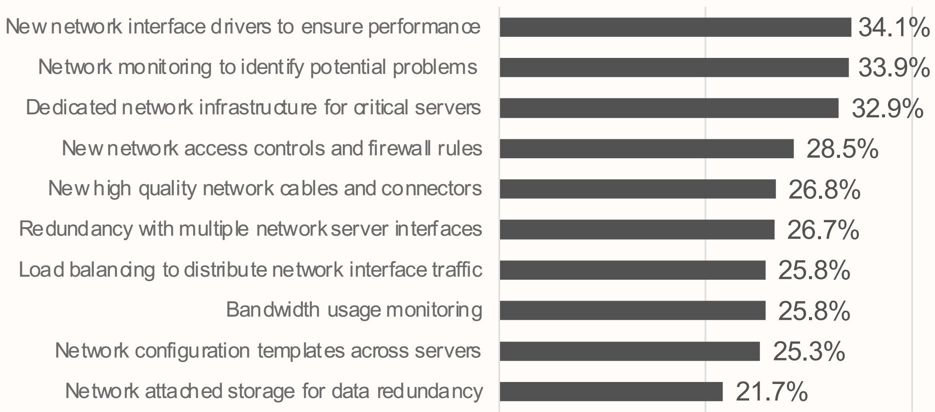

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 882

- One third of IT organizations cite network congestion (e.g., too many devices accessing the network simultaneously), security vulnerabilities (e.g., inadequate access controls), and scalability limitations (e.g., inability to handle growing bandwidth requirements) as the biggest server network interface challenges. Other challenges include physical connectivity issues (e.g., damaged cables), configuration errors, and hardware failure.
- Ineffective server network interface technology can contribute to a host of problems for organizations, including packet loss, slow network speeds and disrupted connectivity, and potential server failure. Costs can be high, and it is essential to have a proper strategy that includes optimizing and upgrading hardware (e.g., optimizing and adopting higher-speed NICs), proactively managing and monitoring network capacity and performance, and adopting modern network technologies, including automation and software-defined networking (SDN).

OCP project initiatives are preparing datacenter networking for the AI era

The OCP's Networking Project aims to establish a fully disaggregated and open networking stack that encompasses hardware and software. Its scope covers switch motherboard hardware, hardware abstraction interfaces allowing choice of switch OS, Linux-based operating systems in collaboration with the Linux Foundation, and developer tools, configuration management and bare metal provisioning, open integration and connectivity, and SDNs

→ **Enabling network software interoperability**

- Early project milestones include the **Open Network Install Environment** (ONIE) Project, which provides an environment for provisioning different operating systems onto bare metal network switches, and the **Switch Abstraction Interface** (SAI) Sub-Project, which is committed to enabling network software to run on different hardware platforms.

→ **Addressing the need for next-generation network fabrics and interconnects**

- Recent areas of focus for the Networking Project include the advancement of open hardware designs, higher-speed optics, and switch-to-switch connectivity. Work by **Meta**, **AMD**, **NVIDIA** and others to support the development of next-generation network fabrics, and efforts by **Broadcom**, **Marvell**, **Celestica**, **Credo Semiconductor** and others to advance the introduction of 800G optical interconnects, are driven by the need to prepare network infrastructure for the future demands of AI.

→ **Readying Ethernet technology for Scale-Up networking**

- The OCP's 2025 Global Summit saw the unveiling of a **new networking-focused working group** to promote the development of Ethernet for Scale-Up networking. The new OCP Ethernet for Scale-Up Networks (ESUN) working group brings together **Broadcom**, **AMD**, **NVIDIA**, **Cisco**, **HPE** and others, and will initially focus on L2/L3 Ethernet framing and switching, engaging with industry bodies such as the Ultra Ethernet Consortium to agree on open standards and best practices.

Power Technologies

Between 2025 and 2029, datacenter IT energy use will grow by 15.8% (CAGR)

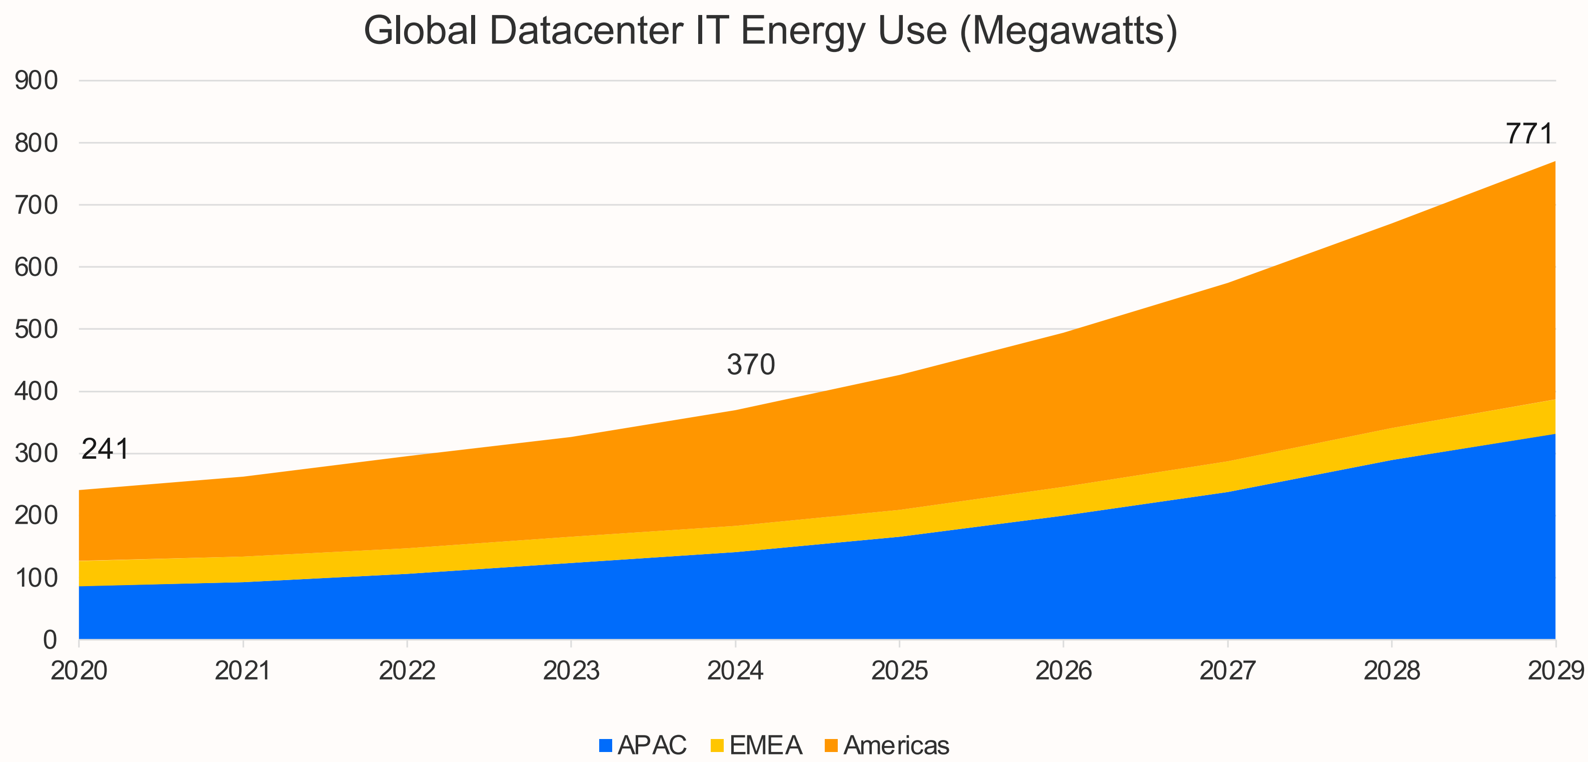

Source: Datacenter Facilities Index, 2H25, IDC, November 2025

Average power consumption and power density per rack is also rising rapidly

- In the last ten years, **average power consumption** of **processors** and **accelerators** has nearly tripled

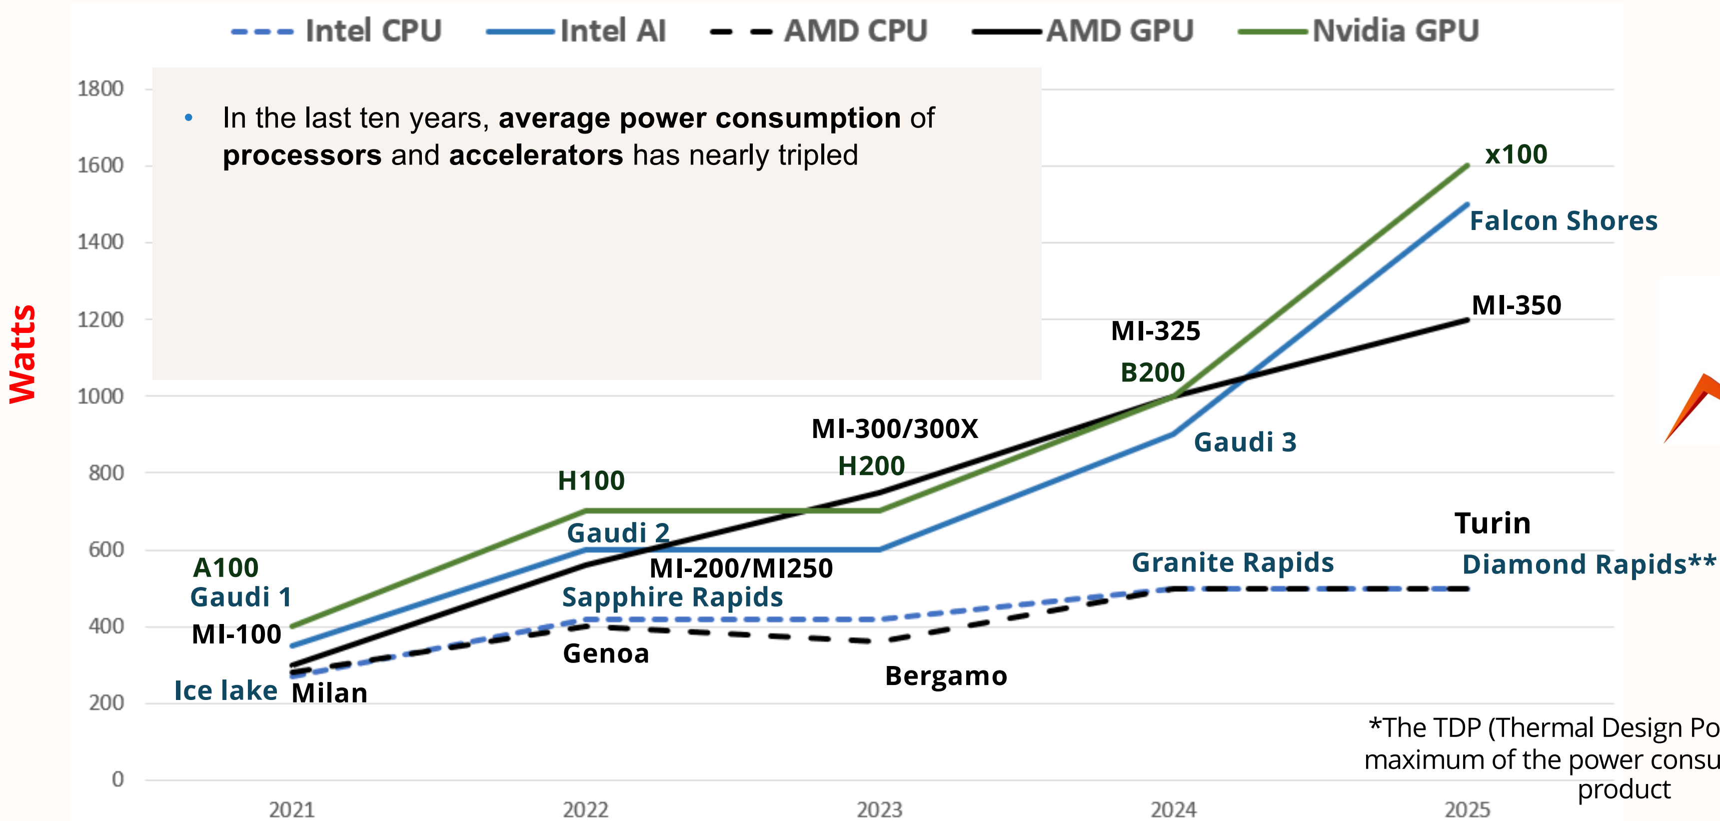

*The TDP (Thermal Design Power) refers to maximum of the power consumption of the product

Source: IDC, 2024, ** Diamond Rapids is based on estimation

Average power consumption and power density per rack is also rising rapidly

- In the last ten years, **average power consumption** of **processors** and **accelerators** has nearly tripled
- At the same time, the **average power density per rack** continues to grow. As more examples of **100kW+ rack densities** emerge, this has huge implications for the **design of datacenters** and their **supporting infrastructure**, including power, cooling, and networking

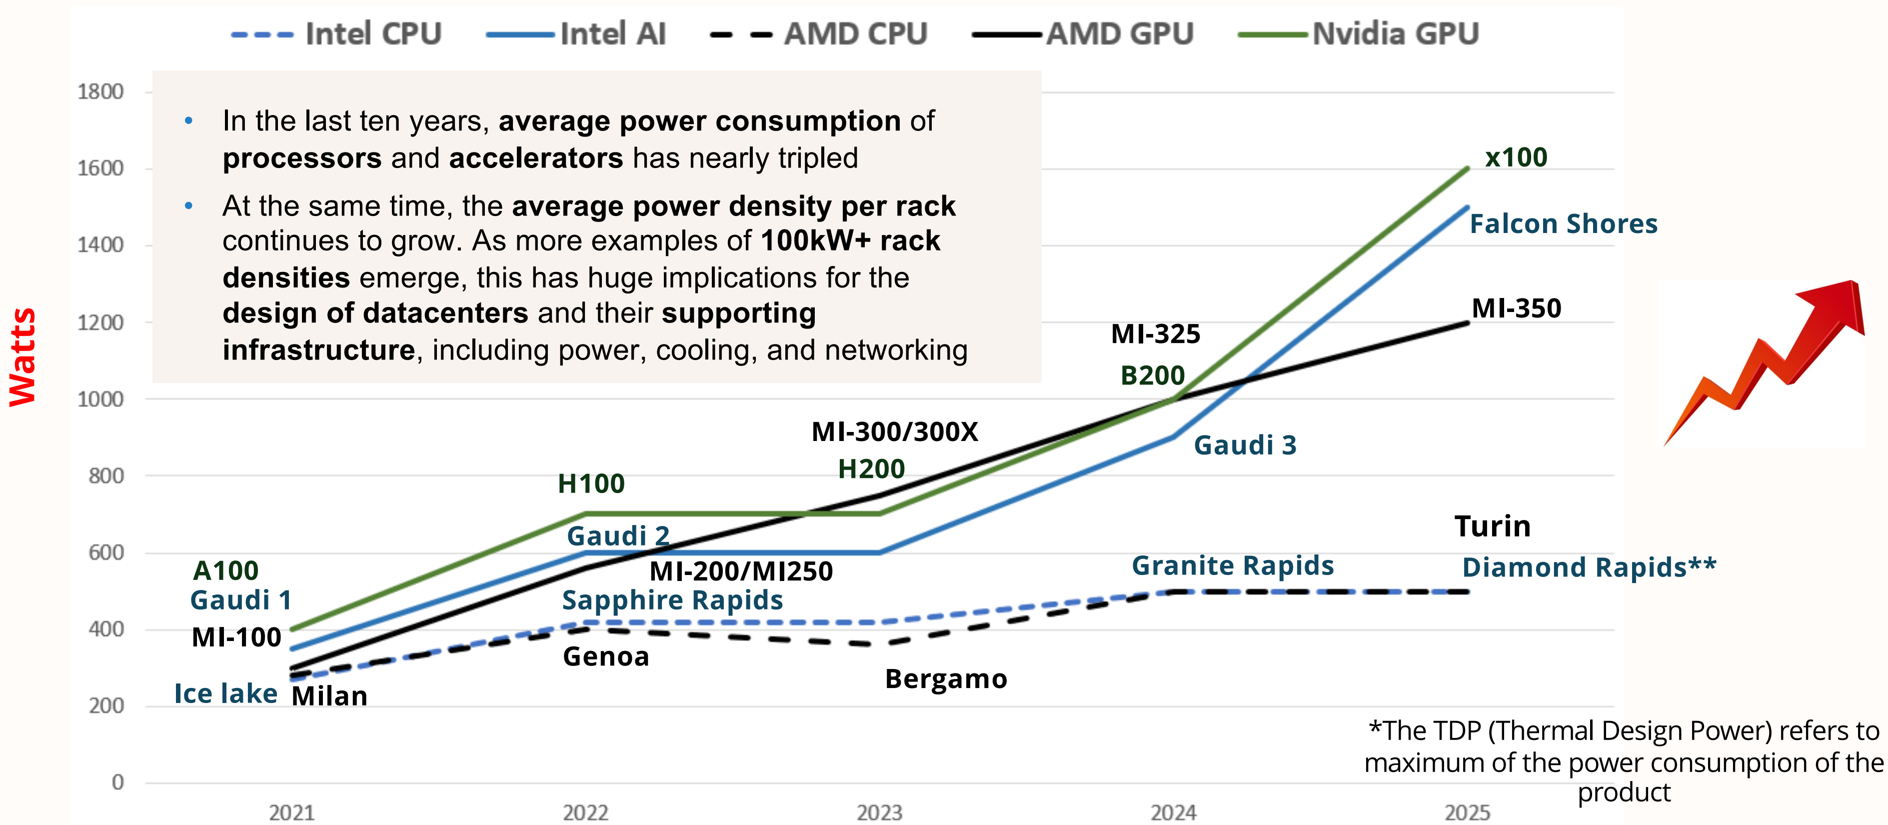

Source: IDC, 2024, ** Diamond Rapids is based on estimation

Rising costs, quality issues, and capacity constraints are among the biggest power usage and supply concerns

Q. What are your organization's main concerns related to datacenter power consumption and/or power supply?

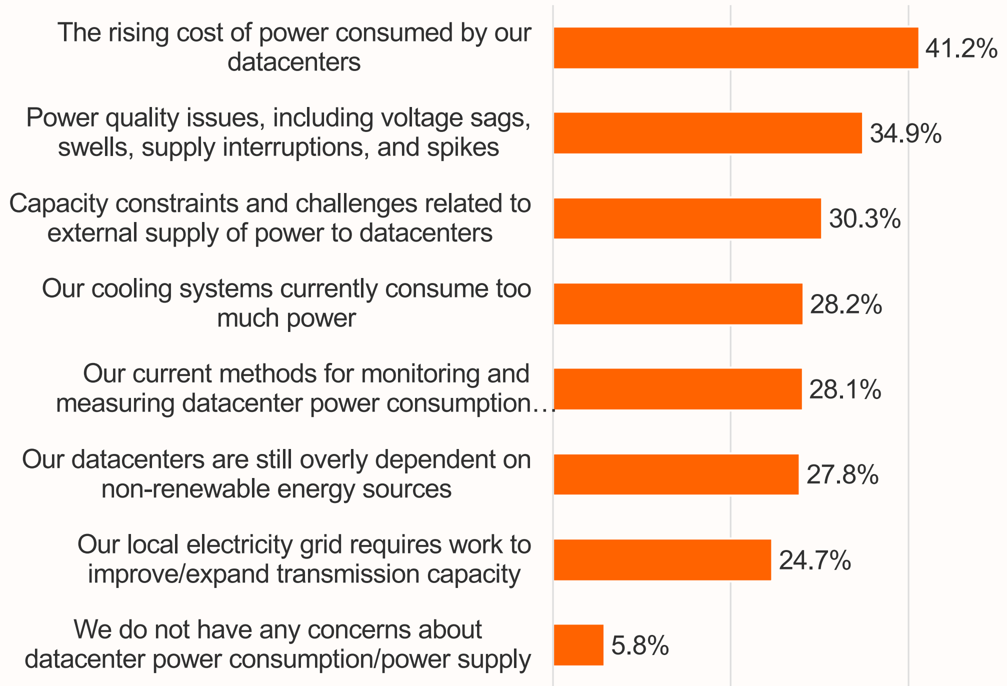

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 882

- Over 41% of organizations identify the rising **cost of power** required by their datacenters as the biggest power-related concern. **Power quality** issues and **capacity constraints** also feature prominently, identified by 34.9% and 30.3% of organizations, respectively.
- Other concerns include an **overdependence on non-renewable energy** sources and the use of unsustainable methods for **monitoring and measuring** datacenter power consumption. Sometimes monitoring strategies focus narrowly on operational energy and do not account for total datacenter energy use.
- Just 5.8% of organizations say they have no concerns about datacenter power consumption and/or power supply.

OCP initiatives include increasing power densities and improving energy efficiency

The OCP's commitment to datacenter power infrastructure spans several OCP projects, including the Data Center Facility Project, the Rack & Power Project, and the DCF Sustainability Project. Key areas of focus include increasing power densities, enhancing energy efficiency, and improving power distribution within datacenters by developing and encouraging the adoption of open-source, standardized hardware designs

Enhancing rack energy usage and efficiency with ORv3

- Key milestones in the OCP's approach to datacenter power include its adoption of the ORv3 specification, which extends to power shelves, battery backup units, and power monitoring interfaces. The specification incorporates **Harting**'s ORv3 plug-and-play connector for power shelves, which is designed to simplify installation and maintenance and boost energy efficiency.

Moving towards higher-voltage power distribution

- Another major development is the promotion of higher-voltage power distribution systems. At the OCP's EMEA Summit in April 2025, **Google** announced the introduction of a +/-400 volts direct current (VDC) power delivery system, which can support up to 1 MW of power per rack.

Improving power consumption management

- The 2025 Global Summit saw growing support for 800 VDC datacenter power architectures. The initiative aims to drive innovations that result in high-efficiency, scalable power delivery for next-generation AI workloads. **NVIDIA** is leading the transition to 800-volt power distribution systems, which will be capable of supporting 1 MW IT racks and beyond. To accelerate adoption, NVIDIA is working with partners that include silicon providers such as **Analog Devices**, **Infineon**, and **STMicroelectronics**, power system component providers such as **Delta**, **Flex Power**, and **LiteOn**, and datacenter power system providers, **Eaton**, **Schneider Electric**, and **Vertiv**.

Cooling Technologies

Most organizations plan to increase their use of liquid cooling over the next 12 months

Q. In the next 12 months, does your organization plan to increase, decrease, or maintain its usage of liquid cooling?

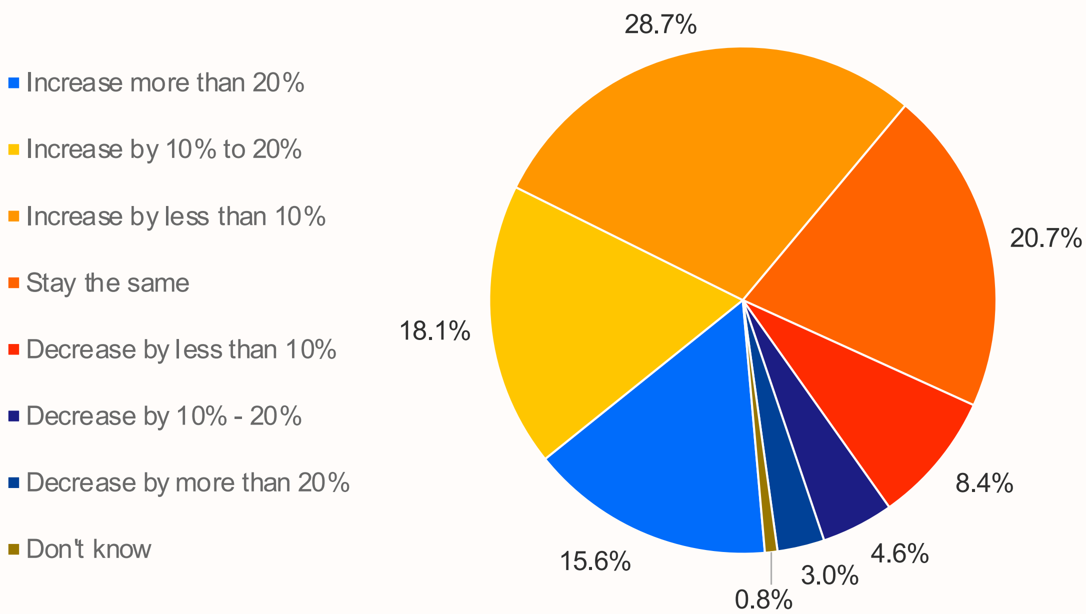

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 235

- Most organizations already using liquid cooling plan to increase their use of liquid cooling by 10% or more over the next 12 months. 15.6% of organizations anticipate their use of liquid cooling increasing by more than 20%. This figure is even **higher in the Asia Pacific region** where over 30% of organizations expect their use of liquid cooling to increase by over 20%.
- 20.7% of organizations expect their use of liquid cooling to stay the same over the next 12 months, while just 16% of organizations anticipate a decrease in their use of liquid cooling over the next year. Latin American organizations account for a significant share of those planning to decrease their use of liquid cooling, as well as those in the manufacturing and education sectors.

Investments in liquid cooling are being driven by multiple requirements and concerns

Q. For what reasons will your organization increase its usage of liquid cooling over the next 12 months?]

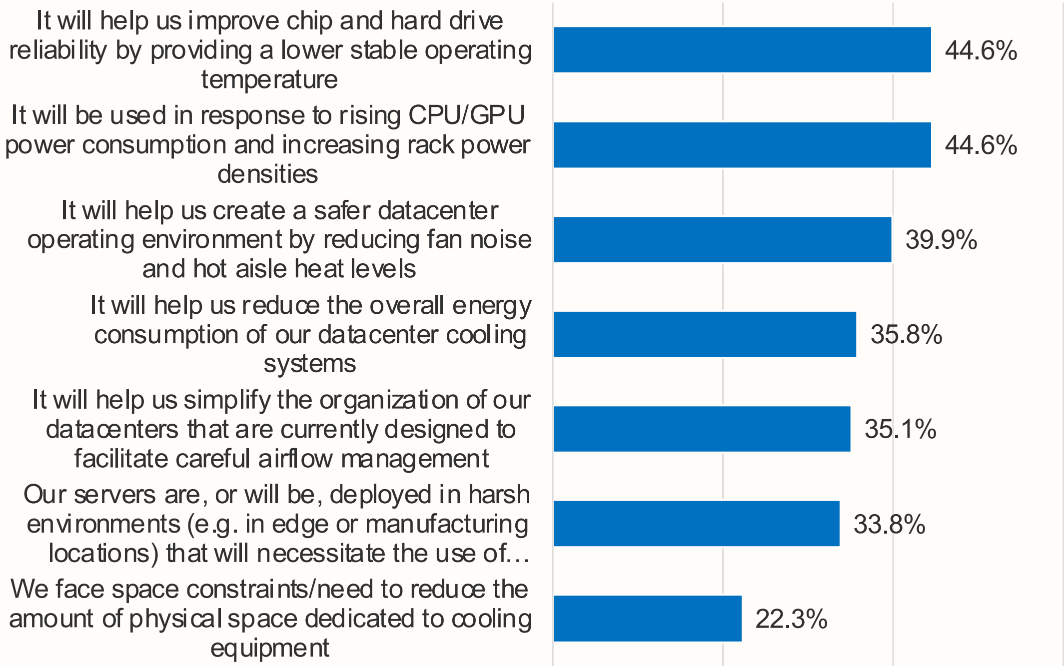

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 148

- Two of the biggest drivers that are influencing organizations' decision to increase their use of liquid cooling are the **need to improve chip and hard drive reliability** by providing a lower stable operating temperature (cited by 44.6% of organizations), and the need to respond to **rising CPU/GPU power consumption and increasing rack densities** (44.6% of organizations).
- Other factors driving the increased use of liquid cooling include a commitment to **creating safer datacenter operating environments** by reducing fan noise and hot aisle heat levels (39.9% of organizations) and efforts to **reduce the overall energy consumption** of datacenter cooling systems (35.8%).

The OCP is advancing standards for liquid cooling and preparing datacenter cooling for AI

The OCP's Cooling Environments Project drives the adoption of open standards for datacenter cooling technologies, including liquid cooling. The Project focuses on several functional areas including Cooling Distribution Units (CDUs), read door heat exchangers, cold plate technologies, immersion cooling, heat reuse, and hardware management for liquid cooling workstreams

- **Driving standards for liquid cooling designs and implementation**
  - Specific project milestones include the development and publication of specifications for different liquid cooling technologies (e.g. cold plates and immersion cooling) and the formation of guidelines for integrating and deploying those technologies.
- **Growing ecosystem of vendors offering OCP-recognized cooling solutions**
  - Recent months have seen server manufacturers such as **Dell** and **Supermicro**, and cooling solution vendors such as **Gigabyte**, **Wiwynn**, **Submer**, **Accelsius** and **Compal** all announce new liquid cooling solutions based on OCP specifications. **LiquidStack**, offers solutions for a range of workloads, including immersive liquid cooling solutions for edge environments..
- **Advancing cooling solutions for AI and high-density environments**
  - At the 2025 Global Summit, the OCP announced a new collaboration with **ASHRAE**, which is committed to establishing standards and best practices for liquid cooling technologies that support AI and high-density environments. It was also at the 2025 Global Summit that **Google** officially announced the contribution of its fifth-generation CDU design, Project Deschutes, to the OCP community. This move aims to standardize and accelerate the industry-wide adoption of liquid cooling solutions for high-density AI and machine learning workloads.

Sustainability

Sustainable growth for IT and datacenter infrastructure has become a major priority for most organizations

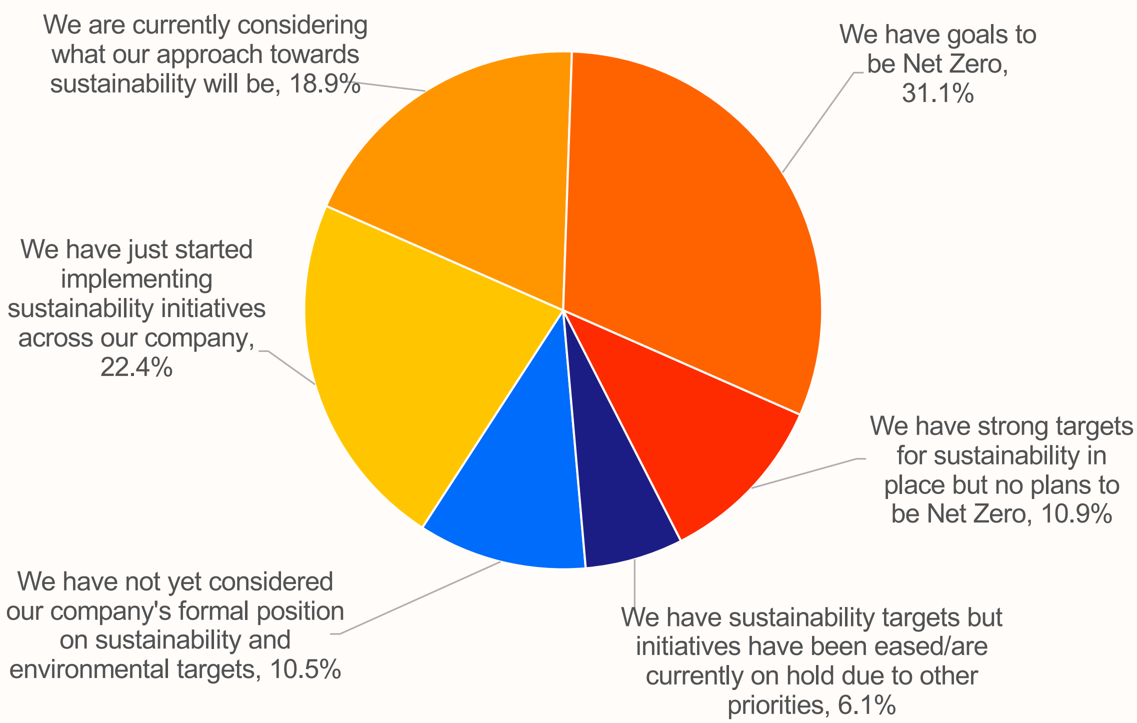

Source: IDC, 2025

- The International Energy Agency predicts that the total **datacenter power consumption** will reach 3% of the global electricity supply by 2030. Rising energy use is fueling a commitment to the use of **renewable energy** sources, improved **energy management** techniques, and more efficient server and datacenter designs.
- Interest in sustainable datacenter growth extends to concerns about datacenter **carbon emissions**, the proliferation of **"e-waste"** (obsolete servers and other hardware), and the impact of datacenter growth on **water** resources, **land** and **construction materials**.
- Over 31% of organizations say they have **goals to be Net Zero**, while a further 10.9% of organizations have strong targets for sustainability but no plans to be Net Zero.

Sustainability policy drivers include regulation and the expectations of customers, investors, and partners

- One of the biggest drivers of the adoption of sustainability initiatives, cited by over 24% of organizations, is the **need to comply with regional and national regulation**, including targets for carbon emissions.
- The second biggest driver of sustainability initiatives is **the impact of industry-specific regulation**.
- Other important sustainability initiative drivers include the **push for more efficient operations**, **cost efficiencies** and **returns on investment**(RoI), and responding to demand from customers, investors, and ecosystem partners.

Q. What are the main reasons behind your company's adoption of sustainability initiatives?

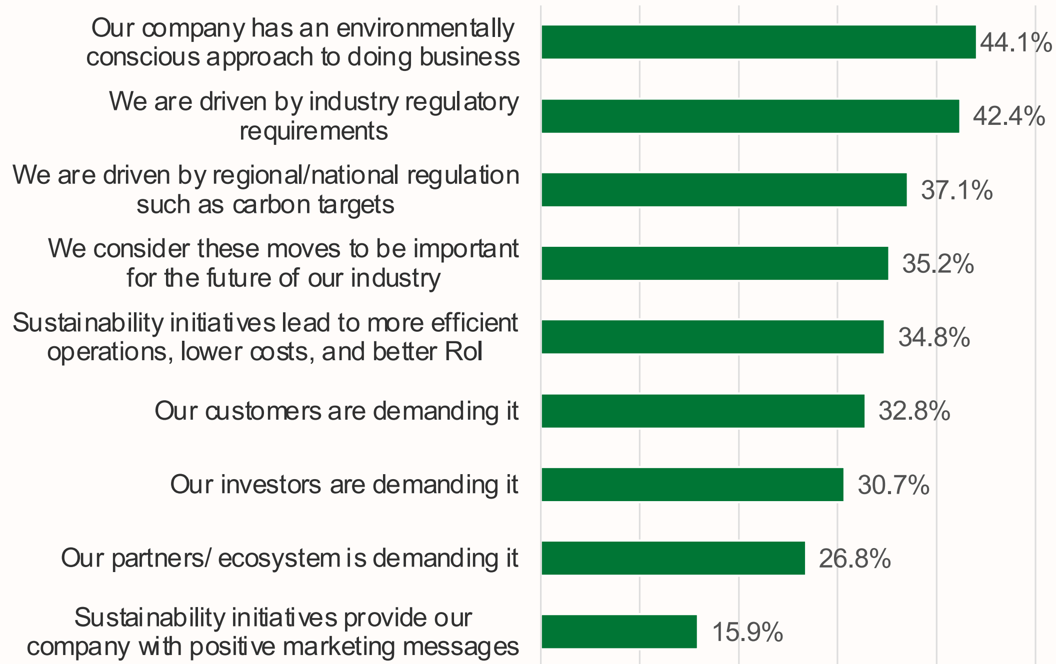

Source: Enterprise Infrastructure Pulse Survey 2025, May 2025, n = 735

Various OCP initiatives focus on encouraging sustainable growth for datacenters

The OCP's Sustainability Project aims to advance best practices that promote the reusability and circularity of datacenter resources, with OCP member initiatives focusing on monitoring, measuring, and reducing carbon emissions, and encouraging greater efficiency in the use of datacenter resources such as energy, water, land, and datacenter construction materials

Advancing sustainable and scalable infrastructure for AI and HPC workloads

- The OCP's collaboration with **Infrastructure Masons** (iMasons) aims to promote sustainable and scalable digital infrastructure for AI and HPC workloads. In October 2025, the OCP and iMasons announced a common, open framework for reporting the carbon impact data associated with datacenter equipment. The Embodied Carbon Disclosure – OCP Base Specification includes contributions from **Google**, **Meta**, **Microsoft**, and **Schneider Electric**, and aims to reduce duplication, align reporting practices, and harmonize carbon disclosure practices.

Promoting the development "green" construction materials for datacenters

- Another OCP sustainability initiative is the "green", low carbon concrete for datacenters initiative. Announced in 2024, the collaboration involves **AWS**, **Google**, **Meta** and **Microsoft** and is focused producing a range of different concrete solutions, the most advanced of which would achieve a 50%+ reduction in greenhouse gas emissions compared with existing concrete production techniques.

Open sourcing designs to share sustainability gains

- Individual OCP members are also driving sustainability moves by open sourcing their technologies. France's OVH Cloud open sourced its liquid cooling solutions to help other datacenter operators become more sustainable. OVH Cloud's datacenter sustainability strategy includes extending the lifespan of components, prioritizing "water frugality" and using brownfield locations for new builds.

Takeaways & Predictions

Key takeaways and expectations for 2026

→ **Preparing datacenters for large-scale AI and HPC workloads**

- The OCP will continue to drive efforts to prepare IT and datacenter infrastructure for AI and HPC workloads, especially those deployed at scale. Key focus areas include the development of open-source hardware specifications for AI clusters and the continued evolution of power and cooling systems to support growing IT power densities. The recently announced Open Data Center for AI strategic initiative will drive deeper collaboration to address datacenter infrastructure challenges, including those related to power, cooling, and management telemetry.

→ **Extending open standards to emerging and adjacent technologies**

- The OCP's commitment to modularity and open standards will continue to be applied to new and adjacent areas, including firmware, chiplet production, and the development of composable memory technologies. Key priorities here include enhancing integration potential and reducing the challenges associated with vendor lock-in. They also include cutting costs for tech buyers, improving platform scalability, and supporting organizations' ability to innovate.

→ **Increasing focus on autonomous operations and automated processes**

- The scale and complexity of datacenters will increase as their infrastructure evolves to support the growing demand of AI and HPC workloads. This will require a stronger focus on autonomous operations and automated processes. This is expected to become an increasingly important focus area for the OCP community, whose collaborative efforts will encourage the standardization of communication protocols for orchestration, monitoring and control (telemetry).

Key takeaways and expectations for 2026

- **Advancing sustainability initiatives and "green" datacenter development**
  - The strong commitment to furthering sustainability initiatives and reducing the environmental impact of datacenters will continue in 2026, with specific focus areas including improved thermal management, datacenter heat waste reuse, the development of "green" datacenter construction materials, and more effective ways to monitor and measure the carbon impacts of specific workloads.

- **Continued expansion in edge and non-traditional datacenter environments**
  - The OCP will continue to extend its focus and the application of principles such as open design and modularity from hyperscale datacenters to edge computing environments. This includes a commitment to the development of compact, energy efficient servers that can be deployed in telecommunications networks to support localized 5G, IoT and AI-based services, or in remote or environmentally challenging locations, closer to points of data generation and collection.

- **OCP community growth and increased project collaboration**
  - The 2026 the OCP community will see further growth and expansion, continuing the year-on-year trend that we have seen for several years. In addition to gaining new members, the OCP is expected to see increased collaboration from its members both in terms of participating in and contributing to its various projects and subprojects. The trend whereby multiple vendors collaborate on specific project contributions is expected to continue.

Primary analyst

Chris Drake

Senior Research Director
Compute Infrastructure and
Service Provider Trends
cdrake@idc.com

Contributing analysts

Leon Kao

Senior Research Manager
Compute Infrastructure and Service Provider Trends
lkao@idc.com

Kuba Stolarski

Research Vice President
Compute Infrastructure and Service Provider Trends
kstolarski@idc.com