



Technical Research Study



Smarter Data Centers: The Future of Power and Cooling

Harnessing AI, intelligent power management, and advanced cooling technologies for more efficient and sustainable data centers.

Executive Summary

Data centers are essential for the digital economy. They have returned to popularity as an alternative to cloud-based services for organizations that prefer to keep their data on-premises. Additionally, organizations are learning that having their infrastructure only in the public cloud might not be the most cost-effective solution. This renewed focus on data centers brings new challenges: reducing energy consumption while still satisfying demands for more powerful compute and infrastructure, selecting the right cooling systems for now and the future, and rising operational costs that can directly impact a data center operation's bottom line.

This study explores the latest trends and best practices in data center power and cooling efficiency, emphasizing the need to balance performance and costs. According to the International Energy Agency (IEA), global data center electricity use is expected to double by 2026,¹ highlighting the urgency for more efficient energy management. Cooling technology also presents a difficult choice, especially for older data center systems. When deciding between air cooling and liquid cooling, each method has its own distinct advantages and disadvantages.

This study also discusses the role that AI, generative AI (GenAI), and smart-connected devices (that is, network-based or cloud-connected devices) have in optimizing energy usage, managing cooling technologies, and enhancing data center operations management. These technologies enable real-time monitoring for data center operations that can significantly improve efficiency and increase cost savings.

Prowess Consulting was commissioned by Dell Technologies to review and discuss the issues of power and cooling. Among server OEMs, Dell Technologies stands out for its innovative approaches to smart power and cooling management. Dell™ OpenManage™ Enterprise Power Manager and a diverse array of Dell™ Smart Power and Cooling solutions, including air and liquid cooling, demonstrate Dell Technologies' significant advancements in monitoring and, as an option in Dell™ PowerEdge™ servers, managing energy efficiency and operational reliability. By making use of AI-based fuzzy logic controllers and real-time data integration, solutions from Dell Technologies provide data center operations with the tools needed to achieve optimal performance and find cost savings. This includes forward thinking about how technology platforms in existing and new data centers can address ever-increasing demands for system density and compute power. Adopting such strategies is crucial for meeting the challenges of today's AI-driven landscape and ensuring the long-term viability of data center operations.

Introduction

Data centers deliver critical services for businesses and consumers. However, data centers that were originally designed for yesterday's technologies must evolve. As demand has only grown for compute, storage capacity, and connectivity, data centers face significant challenges in energy consumption and operational efficiency. IEA reports that global data center electricity use is expected to double by 2026, reaching 2,000 terawatt-hours per year.¹ This increase in energy consumption poses serious environmental threats, including increased carbon dioxide emissions (greenhouse gases) and water consumption for cooling data center systems.⁴ Aside from environmental costs, these factors can also result in greater operational costs.

This analysis examines current trends and best practices in power and cooling efficiency while providing recommendations to enhance energy management and reduce environmental impact.

Highlights



Impact of Energy Consumption from Traditional Data Centers

Global data center electricity use is set to double by 2026.¹

Traditional data center energy consumption accounts for around 1–1.5% of global electricity use.²

Some of the largest data centers use 40 MW of electricity in one facility, enough energy to power 32,000 homes.³

An increase in energy consumption poses serious environmental impacts, including increased CO₂ emissions and water consumption for data centers.⁴

Energy and Water: Pinpointing Efficiency Trends for Data Centers

The needs associated with modern data centers are growing beyond those of older, traditional data centers: mainly an increasing demand for energy and water. Planning for and implementing efficiency measures for energy and water can represent a significant savings for data center operations.

Impact of Energy Consumption from Traditional Data Centers

Today's IT initiatives require increasing amounts of power and cooling. Traditional data center energy consumption accounts for around 1–1.5% of global electricity use, according to some estimates.² Some of the largest data centers use 40 MW of electricity in one facility, enough energy to power 32,000 homes.³ Power usage can be influenced by several factors including the number and type of servers, their workloads and utilization rates, and the efficiency of a data center's power distribution and conversion systems. The overarching result is that this demand for energy will only grow, straining electrical grids and resources. For this reason, energy efficiency among data centers is an important goal.

Cooling Methods and Their Effects on Energy Consumption

While cooling requirements can contribute significantly to power consumption, they are essential to prevent data center servers and equipment from overheating. Depending on the design and location of the data center, cooling systems can consume up to 40% of the total energy usage.⁵

The choice of cooling methods can greatly affect energy usage. Traditional air-based cooling systems use air tubes to cool the data center. Racks are separated by hot and cold corridors, with air pushed by fans entering the cold corridor and passing through the front of the equipment. The same air is sent to the back of the cabinet, where it comes out warm, and it is then sent back to the tubes. Air cooling alone has its limitations, as some experts have noted there is only so much heat that can be removed using air.⁶

Hybrid cooling and liquid cooling offer alternatives to air cooling alone. Hybrid cooling uses a combination of air and liquid cooling through heat exchangers, air chillers, and cooling towers. This approach has certain advantages over using air cooling alone, although increased system complexity can also be seen as a disadvantage.

Liquid cooling, whether implemented in a hybrid approach or exclusively, has grown as a desirable choice because of its effectiveness, but it also stands out because it requires consumption of another critical resource for data centers. Water consumption is primarily driven by the type of cooling system implemented. For example, hybrid cooling systems such as chillers and cooling towers require large quantities of water, which can lead to significant costs and environmental impacts. Direct liquid cooling technology, with coolant liquids circulating through cold plates or through the immersion of components, reduces the need for air conditioning and can provide a savings in cost. However, the quality and availability of water sources, which can vary based on location and climate, can play a crucial role in water usage and can impact the cost of data center operations.

Table 1 compares the different choices for cooling technologies, in addition to their advantages and disadvantages.

Table 1 | Comparison of cooling technologies for data centers

Type	Mechanics	Components	Processes	Advantages	Disadvantages
Air cooling	Uses fans to circulate air through the data center, passing over heatsinks attached to servers to dissipate heat	Includes computer room air conditioners (CRAC), computer room air handlers (CRAH), and direct expansion (DX) systems	The hot air expelled by servers is cooled by CRAC/CRAH units and recirculated back into the room	<ul style="list-style-type: none"> Typically requires a lower initial investment compared to liquid cooling systems Easy to implement and maintain, with few specialized components 	<ul style="list-style-type: none"> Becomes less efficient at higher densities due to limited heat-transfer capacity Requires high energy consumption due to the need for continuous operation of fans and CRAC units

Type	Mechanics	Components	Processes	Advantages	Disadvantages
Hybrid cooling	Utilizes both air and liquid cooling methods to enhance cooling efficiency	Incorporates rear-door heat exchangers, liquid-to-air heat exchangers, and in-row or in-rack cooling units	Servers are cooled by both direct air flow and liquid cooling systems that absorb excess heat	<ul style="list-style-type: none"> Improved efficiency compared to traditional air cooling Capable of handling high densities Can be implemented in stages, allowing for a gradual transition from air to liquid cooling 	<ul style="list-style-type: none"> More complex to manage and maintain due to the integration of both cooling types Higher upfront costs compared to a purely air-cooling system
Liquid cooling	Uses liquid coolant directly on heat-generating components like CPUs and GPUs	Includes cold plates, rear door heat exchangers, coolant distribution units (CDUs) in row or in rack, and pumps	Liquid circulates through cold plates attached to server components, absorbing heat and transferring it away from the equipment	<ul style="list-style-type: none"> Highly efficient at removing heat Capable of supporting very high server densities Reduces the need for extensive air conditioning, which can lead to significant energy savings 	<ul style="list-style-type: none"> High initial investment for setup and infrastructure Requires specialized maintenance and handling to prevent leaks and ensure proper functioning

Adopting Advanced Cooling Systems for Energy Efficiency

Among the most promising approaches to reducing resource consumption are advanced cooling systems. Because cooling systems account for a significant portion of power and water usage in data centers, they are now facing increasing loads due to the extensive demands of AI applications.

One cost-effective cooling solution is free cooling, which utilizes outside air when temperatures are sufficiently low to cool data centers, thereby reducing the reliance on mechanical cooling. Economizer systems use ambient air for cooling, enabling systems to switch over to traditional methods when necessary, which can help significantly reduce energy consumption and carbon emissions, particularly in regions with cold or moderate climates.

An alternative solution, the direct-to-chip method, applies coolant directly on server components via a cold plate that covers the processor and conducts heat away using liquid tubes. Immersion cooling, on the other hand, submerges servers in non-conductive liquids. These methods provide superior heat dissipation and energy savings compared to traditional air cooling, and they are particularly suitable for high-density and high-performance computing, such as AI applications, which can generate more heat than conventional applications. These alternatives to traditional liquid cooling techniques reduce water consumption by eliminating chillers and cooling towers, which typically require large quantities of water. However, immersion can be particularly costly because it requires a custom solution.

Integrating AI and smart-connected device technologies can greatly add to the performance of advanced cooling solutions. AI programming, tied to data center operations, can help reduce energy waste. Smart sensors, connected via network or cloud-based systems, can be placed among server racks, inside white space regions, and inside liquid coolant conduits. Smart sensors monitor temperature and humidity in real time, allowing for dynamic adjustments to cooling systems. These technologies enable more granular and intelligent control over cooling operations, enhancing both reliability and sustainability.

Costs and Concerns When Choosing Efficient Cooling Improvements

Data centers must manage the heat generated by high-performance equipment to prevent overheating and help ensure optimal performance. The choice of cooling systems is critical not only for maintaining operational efficiency but also for reducing energy consumption and meeting sustainability objectives. Not all cooling technologies are applicable for every data center use.

For example, while free cooling can enable a data center to use significantly less power by utilizing sufficiently cool outside air, this method has its limitations. Outside air requires adequate air filtration, humidity control, and security measures to prevent contamination and intrusion, all of which can contribute to increased costs and maintenance for a cooling source that appears “free” on the surface.

Liquid cooling methods also pose certain challenges, including planning, higher initial costs, potential risks of leaks, and compatibility issues with existing equipment and infrastructure. Older data center operations are often concerned with the overall effort and cost required to convert from air cooling techniques to liquid cooling, as more than 60% of today’s data centers do not use liquid cooling.⁷

Table 2 | Comparison of the efficiency, cost, and feasibility of cooling technologies in data centers

Type	Efficiency	Cost	Feasibility
Air cooling	Less efficient and struggles with high-density server configurations	Low initial costs and high ongoing energy and operational costs	Highly feasible for low- to medium-density data centers, with straightforward implementation and maintenance
Hybrid cooling	More efficient than air cooling and capable of handling medium to high server densities	Moderate initial cost with improved operational cost efficiency over time	Feasible for medium- to high-density data centers and suitable for a gradual transition
Liquid cooling	Most efficient approach, capable of supporting extremely high server densities	High initial costs with significant long-term savings in energy and operational expenses	Best suited for high-density, high-performance computing environments and requires investment in specialized infrastructure and maintenance

Lastly, while AI and smart sensor technologies can be deployed to optimize cooling performance and efficiency, these technologies require additional investment, training, and security measures to help ensure proper functioning and protection from cyberattacks. The optimal choice depends on several factors such as the size, location, and purpose of the data center, in addition to the budget, goals, and preferences of data center operations.

The Risks of Doing Nothing

As data centers spring up across the globe to handle the growing demand for information in an increasingly interconnected world, their need for vast amounts of power and water necessitates planning for energy use and watershed impact. Failure to do so could be catastrophic, as by doing nothing, environmental changes could directly impact data center operations. About 20% of data centers in the United States already rely on watersheds that are under moderate to high stress from drought and other factors.⁸ The Colorado River, for example, has become a critical and contentious resource for power, irrigation, and water supplies. A megadrought has impacted power generation, forcing states and power companies to reconsider their reliance on hydroelectricity.⁹

Many large tech companies are looking at ways to replenish or reuse water supplies by building less in drought-stricken areas and considering locations for data centers with better access to power.

AI and Smart-Connected Devices in Data Center Management

AI and smart-connected devices are transforming data center management by enhancing performance, reducing energy consumption, and improving reliability. By making use of real-time data, predictive analytics, and automation, AI and smart-connected devices provide data center operations with powerful tools to achieve greater efficiency and reduce costs.

The Role of AI and Smart-Connected Devices in Optimizing Energy Usage

AI and smart-connected technologies assist data center operations through data-driven insights and automation. AI applications can analyze historical and real-time data from smart sensors, logs, and user feedback to identify patterns and anomalies, and they can provide tools that can anticipate and help prevent issues like equipment failures, security breaches, and service disruptions. Additionally, these technologies can forecast energy consumption based on workloads, weather, and electricity prices, allowing for optimized energy usage and procurement and helping reduce carbon footprints and operational costs.

Proactive maintenance and predictive power management are critical applications in data centers. By continuously monitoring the health and performance of equipment, AI-based systems can alert operations to potential issues, enabling preventive maintenance and helping avoid costly breakdowns. Predictive power management optimizes power distribution by adjusting voltage, frequency, and current, and by switching between power sources such as grid, battery, or renewable energy, enhancing power efficiency and resilience while helping reduce costs.

Adaptive cooling can help illustrate the transformative impact of AI and smart-connected devices on data centers. AI-enabled applications can optimize temperature, humidity, and airflow in coordination with workload and power consumption data, reducing cooling energy and water usage while maintaining optimal conditions for equipment. This advancement can enable more intelligent and efficient data center operations.

Smart-Connected Devices and Real-Time Monitoring

Smart-connected devices play a pivotal role in enhancing data center management through real-time monitoring and continuous data collection. These devices and physical sensors connect to the network, enabling the tracking of various parameters such as environmental conditions, power usage, and equipment performance. This constant flow of data is critical for optimizing data center operations.

Continuous Data Collection

Real-time monitoring in data centers relies heavily on continuous data collection from physical sensors. Environmental monitoring is a primary focus, with sensors tracking temperature, humidity, and airflow to help maintain optimal conditions and prevent overheating. Temperature sensors are strategically placed within server racks, walls, ceilings, and cooling conduits to provide a comprehensive thermal profile of the data center. These sensors can be either contact sensors, which measure temperature by direct touch, or non-contact sensors, which detect infrared radiation. Each type has its advantages, with contact sensors offering greater accuracy and non-contact sensors providing more flexibility.



Beyond heat detection and temperature profiling, water-detection sensors and leak monitoring are another tool for continuous data collection. Besides alarms that indicate the presence of moisture, sensor technology can predict moisture changes, pipe or valve failures, and other environmental changes to quickly alert data center operations.

Power usage monitoring is another crucial tool, with sensors measuring the power consumption of servers and other equipment. This data can help organizations efficiently manage their energy use, identifying areas where energy savings can be achieved and helping ensure power distribution remains balanced across the data center. Additionally, equipment performance monitoring tracks the health and functionality of data center hardware, enabling proactive maintenance and reducing the risk of unexpected failures.

Alerts and Notifications

Data centers using smart-connected devices benefit from consolidated alert and notification systems. Anomaly-detection algorithms analyze sensor data in real time to identify deviations from normal operating conditions. When anomalies are detected, automated reporting systems generate alerts, notifying operators of potential issues before they escalate. This real-time incident-response capability enhances the reliability and resilience of data center operations.

Automated reporting not only provides immediate alerts but also offers detailed insights into the nature and cause of anomalies. Notifications help operators quickly diagnose and address issues, which can help reduce downtime and help maintain optimal performance. Furthermore, integrating AI and smart-connected devices can enhance these systems, allowing for more precise predictions and more effective responses. Refer to the next section for a more detailed discussion on sensors and data integration.

In summary, smart-connected devices and real-time monitoring are integral to modern data center management. Continuous data collection, coupled with sophisticated alert and notification systems, can help ensure that data centers operate at peak efficiency.

Sensors and Data Integration

The increasing complexity and scale of data centers necessitate sophisticated methods for monitoring and managing various operational aspects. Smart-connected sensors play a crucial role in providing comprehensive data tracking. These sensors, which include temperature, humidity, power, moisture or leak detection, and vibration sensors, can offer a granular view of a data center environment, allowing operators to optimize performance, reduce costs, and enhance security.

Data integration makes use of the vast amounts of data generated by sensors. Combining data from various sources—such as historical data, real-time data, and user feedback—onto unified platforms makes use of comprehensive analysis and actionable insights. Historical data analysis enables data centers to compare current performance with past metrics, identifying trends and potential areas for improvement. Real-time data integration facilitates remediation of emerging issues, which helps ensure that data centers can maintain optimal performance and reliability.

One of the significant advantages of data integration is implementation with automated control systems. These systems can dynamically adjust various parameters based on real-time data inputs. For instance, dynamic cooling adjustments regulate temperature and airflow based on workload, ambient conditions, and user preferences, which can result in significant energy savings and improved performance. Such systems not only enhance operational efficiency but also contribute to sustainability efforts by helping reduce a data center's carbon footprint.

Smart Adjustments for Optimal Performance

Optimizing data center performance increasingly relies on automated control systems that dynamically adjust various operational parameters based on real-time data. These systems can significantly enhance operational efficiency and sustainability, aligning with the latest advancements in smart technology and data analytics.

Dynamic Cooling Adjustments

Dynamic cooling systems are essential for maintaining optimal thermal conditions within data centers. These systems utilize data from temperature and humidity sensors to make real-time adjustments to fan speeds, cooling unit settings, and air distribution. By responding to workload demands and ambient conditions, dynamic cooling can help ensure that equipment operates within safe temperature ranges, preventing overheating and reducing energy consumption. This approach can not only enhance the reliability of data center operations, it can also contribute to significant cost savings and a reduced carbon footprint.



Power Management and Intelligent Power Scaling

Advanced power management systems are designed to monitor, manage, and optimize power consumption across data center infrastructure. These systems can set power policies, cap usage, and dynamically allocate power based on real-time needs. By integrating data from power sensors and historical usage patterns, intelligent power scaling can adjust power distribution to efficiently match current demand. This can help ensure that energy resources are used optimally, which can lead to lower operational costs and improved performance metrics. Efficient power management is crucial for sustaining the scalability and resilience of modern data centers.

Continuous Improvement and Real-Time Response

Automated control systems benefit from feedback loops that enable continuous improvement and improved real-time response. By integrating data from multiple sources, including historical data and real-time sensor monitoring, these systems can identify trends, patterns, and anomalies. Continuous data analysis allows for informed adjustments and optimizations, enhancing a system's ability to learn and adapt over time. Real-time incident response capabilities help ensure that any emerging issues are swiftly addressed, maintaining high levels of performance and reliability. AI-driven systems that can assess and analyze feedback from these data points are pivotal for achieving ongoing operational excellence and sustainability in data center management.

The implementation of smart adjustments through automated control systems in data centers leads to significant improvements in efficiency, reliability, and sustainability. These systems' ability to dynamically manage cooling, power, and overall operational parameters positions them as essential tools for modern data center operations. This foundation sets the stage for exploring specific solutions and advancements provided by leading technology providers.

Dell Technologies: A Standout for Smart Power and Cooling

As we considered the role of power and cooling in the IT sector, Prowess Consulting observed how Dell Technologies sets itself apart in power and cooling management for servers in the data center. By implementing advanced technologies and intelligent systems, Dell Technologies offers significant improvements in performance, energy efficiency, and sustainability for data center operations. Among these technologies are management software for data center operations, power and cooling advancements for server racks, and comprehensive solutions for data center and IT operations. In the next section, we examine Dell Smart Power and Cooling solutions, in addition to PowerEdge servers, and we take a look at how these solutions contribute to enhanced data center management.

Smart Solutions

Dell Technologies focuses on a holistic approach to power and cooling optimization for data center operations through integrated, intelligent tools. The Integrated Dell™ Remote Access Controller (iDRAC), embedded in every PowerEdge server, enables real-time monitoring and control of individual servers. Dell OpenManage Enterprise Power Manager extends iDRAC and centralizes management of all PowerEdge servers within the data center, providing unparalleled insight into power usage, carbon emissions, and thermal conditions. With the OpenManage Enterprise Power Manager plugin, iDRAC users gain oversight into granular power and thermal monitoring, with the ability to report, cap, and optimize power consumption at the server, rack, row, or full data center level—driving energy efficiency and cost savings. Dell APEX™ AIOps provides an AI-driven observability and management software-as-a-service (SaaS) suite, building on the technologies provided by iDRAC and OpenManage Enterprise Power Manager. The ability to monitor and control up to 8,000 devices with OpenManage Enterprise Power Manager, regardless of form factor, helps ensure comprehensive power management across diverse and distributed data center environments. Through sensor and data point collection, this extensive monitoring capability also enables data center operations to optimize power distribution and improve energy efficiency, contributing to reduced operational costs and lowering environmental impact.

Dell Technologies offers a wide range of cooling options to meet the demands of modern data centers while updating the efficiencies of older centers. Through air cooling and liquid cooling, rear door heat exchangers (RDHx) on the rack or in-row, and rack enclosures for data centers with traditional and non-traditional spaces, Dell Technologies can provide expert guidance to select the right infrastructure. OpenManage Enterprise Power Manager not only optimizes power usage, reducing the heat load on cooling systems, it also provides insights into power and temperature metrics for proactive adjustments, and it offers improved cooling efficiency by reducing heat generation. iDRAC accesses information about the health of each individual PowerEdge server, promoting more efficient cooling processes. Smart devices and sensors also enhance cooling management. The proprietary Dell Technologies AI-based fuzzy logic controller, embedded in iDRAC with extensive smart sensor integration, further enhances cooling efficiency by fine-tuning fan speeds to match the exact cooling needs of the data center. This innovative approach can not only maintain the ideal operating temperature for servers, but also reduces energy consumption associated with air cooling. By integrating advanced cooling technologies, Dell Technologies helps ensure data centers operate efficiently under varying workloads and environmental conditions.

Dell Technologies: More Than Servers

Dell Technologies exemplifies effective smart power and cooling solutions. These technologies provide data center operations with the tools to achieve optimal performance, energy efficiency, and sustainability. The integration of advanced monitoring, control capabilities, and AI-driven optimizations sets Dell Technologies apart as a leader in the field, paving the way for more intelligent and cost-efficient power and cooling management for data centers. With combined technology plus advisory services, Dell Technologies offers more than servers, but also solutions for organizations considering optimizing their current data centers or those planning for the future of data center operations.

Conclusion

By adopting advanced power and cooling technologies and integrating AI and smart-connected devices, administrators can optimize energy usage and subsequently reduce costs. Learn more about how [Dell Smart Power and Cooling solutions](#) can improve data center operations and costs for your organization. Additionally, read Prowess Consulting's analysis concerning energy efficiency, sustainability, IT infrastructure, and data centers with our featured resource, "[Grow and Innovate on an Energy-Efficient, Sustainable IT Infrastructure](#)."

¹ International Energy Agency. "[Electricity 2024](#)." January 2024.

² International Energy Agency. "[Data Centres and Data Transmission Networks](#)." July 2023.

³ EnergyStar. "[Is Energy Efficiency in Data Centers Still Important?](#)" Accessed August 2024.

⁴ TechTarget. "[Assess the environmental impact of data centers](#)." July 2023.

⁵ US Department of Energy. "[DOE Announces \\$40 Million for More Efficient Cooling for Data Centers](#)." May 2023.

⁶ Data Center Dynamics. "[Air cooling will never go away](#)." November 2023.

⁷ Digital Realty. "[The Future of Data Center Cooling: Innovations for Sustainability](#)." Accessed August 2024.

⁸ Md Abu Bakar Siddik, Arman Shehabi, and Landon Marston. "[The environmental footprint of data centers in the United States](#)." *Environmental Research Letters*. May 2021.

⁹ Slate. "[The Colorado River's Urgent Lesson for Energy Policy](#)." February 2023.



The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.

Results have been simulated and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.

Copyright © 2024 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.