


THE DATA CENTER OF THE FUTURE



Sai Ganesh

Specialist 2, Inside Product

Dell Technologies

Sai.ganesh@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

Table of Contents

Introduction	4
Data Center Challenges.....	4
Power Use Effectiveness.....	5
How Is PUE Calculated?	5
Carbon usage effectiveness	6
Efficient data center design considerations.....	6
Continuously measure data center Power Usage Effectiveness	6
Focus on optimizing IT power	6
Optimizing Data Center Space	6
Optimizing Data Center Cooling.....	7
Utilize Data Center Infrastructure Management software tools.....	7
Using AI to optimize energy consumption.....	8
Data center cooling technologies	8
Computer Room Air Conditioner	8
Free air-cooling	10
Liquid cooling	10
Project Natick by Microsoft.....	11
Evaporative cooling.....	12
Ultrasonic humidification.....	12
Electronic Product Environmental Assessment Tool	12
Conclusion.....	13
References	14

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Introduction

As technology advances and innovations take the world to the next level, the importance of data centers is also growing. In recent years, the Internet of Things (IoT), ever-increasing data requirements, and constant cloud adoption due to remote work have led to a shift away from traditional enterprise data center facilities. A data center centralizes an organization's IT activities and equipment into a single infrastructure that can be distributed for redundancy. Data is stored, controlled, and distributed across a variety of devices and locations.

The key components of a data center design are:

- **Compute:** Considered to be the brain of a data center, compute resources usually consist of high-end servers with fast memory and high computing power.
- **Storage:** Critical business data is normally stored in a storage facility, with several copies, on media ranging from tape to SSD's.
- **Networking:** Consists of components that offer Interconnectivity between the various devices inside the data center and the outside world. Networking includes routers, switches, control hubs, etc.

Data centers have evolved from physical infrastructures to virtual infrastructures. Many organizations opt to operate hybrid clouds. A hybrid cloud is a combination of both types of data centers. The data center function and composition have significantly changed over the years. Constructing a data center used to be a 25+ year commitment that had several inefficiencies concerning power/cooling, no cabling flexibility, and limited mobility inside or between data centers. Now it's all about quick deployment, superior performance, and highly efficient architectures.

Data Center Challenges

Ever-growing demand for data centers that can store, and process Exabyte of data has increased the focus on efficiency and management. Energy is a major and fast-growing growing problem for data centers. According to a new study, data centers are responsible for 1 percent of all electricity consumed globally and the power consumption is estimated to amount to 416 terawatts. In 2017, US-based data centers alone used up more than 90 billion kilowatt-hours of electricity. This is roughly equal to 34 massive coal-powered plants generating 500 megawatts each.

These figures are projected to rise exponentially as more data center facilities are constructed each year to meet the growing demand for unstructured data. With 80% of the world's electricity still being generated by burning fossil fuels and other non-renewable resources, the ever-increasing power demands could exhaust these resources and cause a major problem. Fortunately, data center operators are working to meet the needs of consumers while keeping their energy use at reasonable levels.

Streaming video has now changed the game, but the proliferation of artificial intelligence (AI) and Internet-connected technology would transform the digital landscape. AI is the future, and is hungry for power processing. IoT is predicted to reach 20 billion units by 2022.

Organizations are decommissioning their old power-hungry hardware and are embracing modernization by pushing towards more efficient technologies such as flash storage and cloud virtualization. We now

have more efficient data center cooling systems and greener, smarter architectural design practices that have led to more efficient buildings.

Power Use Effectiveness

The Green Grid (TGG), a global alliance committed to advancing energy efficiency in data centers, has developed two widely used metrics; Power Use Effectiveness (PUE) and Carbon Use Effectiveness (CUE).

Energy consumption is a critical problem for data centers. Power consumption ranges from a few kW for a rack of servers in a cabinet to several tens of MW for the largest installations. In the case of high-power data centers, electricity costs are a dominant operating expense, accounting for more than 10% of total cost of ownership (TCO).

Power use effectiveness (PUE) is the most commonly used energy efficiency metric of data center energy efficiency. PUE is calculated as the ratio of total power entering the data center divided by the power used by IT equipment. A PUE ratio describes how efficiently a computer data center uses energy; specifically, how much energy is used by the computing equipment.

How Is PUE Calculated?

PUE is the relation between the total energy entering a data center and the energy used by IT equipment inside the data center along with the energy consumed by the data center infrastructure to house the IT equipment. (Heating, Ventilation, Air-cooling, Power distribution, lighting).

Hence, the typical PUE formula is as follows:

PUE = the total facility power / the energy used by IT equipment

The formula must be used only to determine the efficiency of a data center over time. It should not be used to compare different data centers.

The PUE ratio can range from 1.0 to infinity. An ideal PUE is 1.0, which means the data center is 100% efficient (i.e. all consumed energy is used only on IT equipment, no power distribution losses). However, it is almost impossible to achieve such a level of efficiency.

PUE	1.2	1.5	2	2.5	3
Efficiency level	Very efficient	Efficient	Average	Inefficient	Very inefficient

According to a study conducted by Uptime Institute research, an average US data center has a PUE of 2.5. However, it is common to find servers with a PUE of 3.3 and higher as well. This means that only 1/3 of the electricity used by data centers is consumed by its IT facilities; 2/3 of that energy is lost.

Industry leaders such as Google and Microsoft are designing data centers for PUE's of 1.2 or better. Google currently maintains an impressive PUE of 1.12 across all its data centers, including all overhead sources, which is very close to the theoretically perfect PUE of 1.0. But analysts now believe that since the PUEs have significantly improved, the low-hanging fruit is gone. Future productivity gains would be minimal and incremental.

Carbon usage effectiveness

Carbon usage effectiveness (CUE) is another metric used to measure energy use and sustainability in data centers. It is calculated by using the following formula:

$$\text{CUE} = \frac{\text{Total CO emissions caused by the Total Data Center Energy}}{\text{IT Equipment Energy}}$$

CUE provides a way to determine opportunities to enhance a data center's sustainability and focus on considering sustainable energy sources.

Efficient data center design considerations

Though optimizing utilization of data center power is a top priority for data center administrators, they continue to face obstacles as power raises the percentage of continuing data center costs.

Gartner forecasts that on-going power costs are increasing at least 10% per year due to cost per kilowatt-hour (kwh) increases, especially for high power density data centers. Approximately 10% of data center operating expenditure (OpEx) is electricity, and power is likely to be about 15% of data center OpEx in the next five years.

Hyperscale data centers, which maintain thousands of servers and can operate servers at higher utilization in infrastructure-efficient spaces, can deliver significant overall energy savings. Companies operating these massive centers, i.e. Google and Microsoft, are also experimenting with AI and renewable energy sources to further improve efficiency and lower costs. Below are a few considerations for efficient data center design:

- **Continuously measure data center Power Usage Effectiveness**

Continue to measure energy use consistently and in regular intervals. Seasonality and changes in outside temperature can have a significant impact on power use and influence a data center's energy savings. A year's worth of monthly data is a good guide to help data center admins determine whether the energy-saving measures put into action have improved PUE value.

- **Focus on optimizing IT power**

Since IT equipment are what ultimately consumes power, data center administrators need to try to lower the amount of power load. Low-power servers are more energy-efficient than conventional servers in data centers. It is estimated that sixty percent of payload power is consumed by servers so taking the following actions to reduce the power they require is crucial:

- Consolidate workloads and eliminate additional overheads
- Virtualize more workloads if possible
- Continue to eliminate Zombie equipment that are powered but doing nothing useful
- Modernize old servers with the latest low power servers.

- **Optimizing Data Center Space**

When architecting a new data center design, it's worth considering a modular design that breaks down the data center into individual modules that can be periodically refreshed to ensure highly efficient systems. Below are some of the best practices that can also be considered while architecting a new data center:

- **Airflow management**
Controlling airflow and limiting the mixture of hot and cold air is critical to lowering of PUE in a data center. Blanking plates can be mounted to fill the gaps where no IT equipment is present.
 - **Align hot and cold aisles**
All server or storage racks should be arranged to face the same direction. These rows can then be positioned into hot and cold aisles with the front of the servers facing one direction to allow for cool air aisles to be created. This will ensure that the cold air is blown to the front of the servers. Hot air from the rear will be directed to hot air aisles.
 - **Introduce aisle containment**
The next natural step to ensure separation between hot and cold air and improve the direction of airflow.
 - **Check flooring**
Raised flooring's main goal is to direct the cold air in one direction, through the racks, and out the other side. Air gaps should be mounted with blanking plates to generate efficiency gains.
- **Optimizing Data Center Cooling**
To achieve a minimum guaranteed level of efficiency, data center operators should ensure the adoption of basic data center cooling best practices such as:
 - **Installing Economizers**
In colder regions, an air economizer can significantly improve the PUE. For example, In most of North America, 40-90% of the cooling can come from outside by using air economizers.
 - **Contain Equipment and Heat**
Isolation structures can be designed to house the data center equipment that produces the most heat and direct the heat out of the data center or to re-use the heat in other parts of the structure.
 - **Optimize Air Conditioning Systems**
There are two main methods of optimizing the air conditioning system – shutting it off periodically, using an external cooling source, such as an air optimizer or constantly varying the temperature, which tends to minimize the overall amount of energy the device uses.
 - **Regulate temperature and humidity level**
Managing temperature and humidity is a significant factor that influences the PUE. With marginally higher humidity bandages and lower temperatures, less energy is consumed for cooling.
 - **Utilize Data Center Infrastructure Management software tools**
Energy efficiency can further be optimized with the use of data center infrastructure management (DCIM) software. These tools monitor, calculate, manage and regulate data center utilization and power consumption of all IT-related equipment (i.e. servers, storage, and network switches) and data center infrastructure components (i.e. power distribution units [PDUs] and computer room air conditioners [CRACs]).

As data center complexity continues to increase, DCIM tools become essential to proactively monitor and automate repetitive tasks to simplify data center control. You will reduce costs, maximize efficiency, leverage the physical room, and boost device performance with the right tool. Some of the most popular DCIM software are Device42, RackTables, Athenta DCIM, ANSYS Fluent, and Sunbird DCIM.

- **Using AI to optimize energy consumption**

The potential for AI to cut data center power consumption is an area of growing interest in the IT industry. Google's **Deepmind** machine learning (ML) capabilities can reduce the amount of energy it uses for cooling by up to 40%.

Earlier this year, Google published details on how by using **Deepmind**, they have created a machine learning system that could help optimize cooling of its data centers and in turn, cut their energy consumption by half. The algorithm is built using a series of neural networks that are trained using data collected from the thousands of data center sensors to monitor changes both within and outside the data center environment and predict how some of these conditions could change in the next hour. The company deployed the algorithm at one of their live data centers and claims the experiment resulted in a 40% reduction in the energy used to cool the facility, leading to the lowest PUE score the site had ever achieved – 1.06

Data center cooling technologies

Given the importance of the data center cooling infrastructure, it is worth taking a moment to analyze some widely used and emerging data center cooling technologies.

- **Computer Room Air Conditioner**

One of the most common features of any data center, computer room air conditioner (CRAC) units are very similar to conventional air conditioners powered by a compressor that draws air across a refrigerant-filled cooling unit. CRAC offers precision cooling and is critical for optimal performance of data center hardware such as servers, storage, networking, and communication systems. Though quite inefficient in terms of energy consumption, the machinery itself is relatively inexpensive. Below are two commonly used CRAC configurations

- **Raised Floor**

Traditionally data centers use raised floor systems to deliver cold air to servers. Cold air from a CRAC pressurized the space below the raised floor. Perforated tiles provided a path for the cool air to exit the CRAC valve and enter the front of the installed racks.




Image source : <https://journal.uptimeinstitute.com/a-look-at-data-center-cooling-technologies/>

- **Cold Aisle/Hot Aisle Design**

Cold Aisle containment environments are used in conjunction with a raised floor. In a Cold Aisle containment system, cool air from air handlers is contained, while hot server exhaust air can return freely to the air handlers. In a Hot Aisle containment system, hot exhaust air is contained and returns to the air handlers, usually via a ceiling return plenum.




Image source : <https://journal.uptimeinstitute.com/a-look-at-data-center-cooling-technologies/>

- **Chilled Water System cooling**

This system uses chilled water to cool air being brought in by computer room air handlers (CRAHs) to cool the data center equipment. Water is supplied by a chiller plant usually located somewhere in the facility.

- **Free air-cooling**

A cost-effective way of using low external air temperatures to cool data centers, free air-cooling systems leverage outdoor free air instead of traditional CRAC systems. While outdoor air still needs to be filtered and moisturized, it requires much less energy. The data center's location plays a critical role here since low outdoor air temperature is a critical source of cooling.




Image source : <https://www.akcp.com/blog/data-centers-free-air-cooling/>

- **Liquid cooling**

Energy reuse is key for data centers to become carbon neutral. To do this, it is necessary to create a closed-loop that transforms waste energy into usable thermal energy. Liquid as a medium is more flexible and sustainable as it can absorb, transfer and reject thermal energy in data centers. The decisive factor to ensure energy reuse possible is a sufficient **temperature difference (ΔT)**.

Since air-based systems are configured to allow relatively limited temperature windows for regular operation, temperatures for reuse are often too low to enable direct reuse scenarios. Adoption of liquid-cooled IT in data centers allows for more efficient utilization or reduction of the data center footprint. This can save 6-45% of the IT energy footprint in the facility and can also reduce emergency power requirements.

Liquid technologies available today can be roughly divided into three categories: cooling at the room (CRAC), chip level, and immersion.

- **Room cooling**
CRAC can be water cooled. This air-cooling set-up is the least economical, but in some cases unavoidable due to legacy systems which require cooling or due to partially liquid cooled IT systems.
- **Direct-to-chip liquid cooling**
This method uses purpose-built coolers which combine cold plates and pumps that are mounted directly onto the motherboards and coolant liquid is delivered directly to

disperse heat. Since this system cools processors directly, it's one of the most effective forms of server cooling system.

- **Immersion System**

An innovative new data center liquid cooling solution, an immersion system submerges hardware into a bath of non-conductive, non-flammable dielectric fluid. There is hardly any energy loss, making IT equipment very energy efficient.

Project Natick by Microsoft

A few years ago, Microsoft sank a data center off the coast of Orkney in a wildly ambitious experiment. Microsoft used a series of sensors to remotely monitor conditions inside and outside the vessel, in case any maintenance issues cropped up, as it is designed to be unmanned.

Deploying a data center on the seabed enabled Microsoft to get around some issues with building a facility on land and at scale, such as having to build data centers in major metro areas where there is a little fiber connectivity. Microsoft states, "Half of the world's population living within 200km of an ocean, and undersea fiber cables are already connecting our world" and thus it is a better cooling option. An underwater data center avoids the need for costly mechanical cooling equipment, and it can also be more environmentally sustainable than traditional on land data center infrastructure. In the future, subsea builds can significantly minimize data transfer times for consumers and decrease energy consumption by a huge margin.




Image source : <https://news.microsoft.com/innovation-stories/project-natick-underwater-datacenter/>

- **Evaporative cooling**

Evaporative cooling reduces heat through evaporation of water. The system exposes hot air to water which causes the water to draw heat out of the air and evaporate. Compared to traditional mechanical cooling systems, evaporative cooling generally uses significantly less electricity.

- **Ultrasonic humidification**

Some humidity is required for data centers to function effectively and avoid harm to computers and servers. Ultrasonic humidification uses ultrasound to create moisture, using 90 percent less energy than conventional methods such as resistance steam humidifiers.

Electronic Product Environmental Assessment Tool

The Electronic Product Environmental Assessment Tool (EPEAT) program provides independent verification of manufacturers' claims and the **EPEAT online Registry** lists these sustainable products which can be used by organizations to ensure manufacturer's claims. National governments, including the United States, and thousands of private and public institutional purchasers around the world use EPEAT as part of their sustainable procurement decisions.

The **Green Electronics Council (GEC)** oversees this flagship program, including maintaining the integrity of the EPEAT framework. EPEAT is one example of how GEC helps IT consumers around the world, cultivating a demand for sustainable IT goods in order to accomplish our vision of a world with sustainable IT.

EPEAT-registered products must follow environmental sustainability requirements that address material selection, supply chain greenhouse gas emission reduction, circularity and product durability design, energy efficiency, end-of-life management and IT benchmark performance.

EPEAT currently covers the following product categories:

- Computers & Displays
- Imaging Equipment
- Mobile Phones
- Photovoltaic Modules and Inverters (PVMI)
- Televisions
- Servers

The GEC Digital Environmental Benefits Calculator helps purchasers calculate and monitor environmental benefits obtained by purchasing electronic goods protected by the EPEAT Ecolabel. The benefits calculator also enables purchasers to estimate how they can achieve further environmental reductions by using products longer and responsibly recycling them when taken out of operation.

Conclusion

No industry, organization, or individual company is perfect. The fact that industry giants such as Microsoft, Google, and others are now making giant leaps towards data center sustainability to benefit our environment with green and sustainable data center design.

While designing one must make sustainability a priority and not only worry about energy bills but also think holistically about one's IT energy consumption. They must proactively replace older generation hardware with **EPEAT Gold devices** and empower their employees to continually challenge the use of energy resources and move towards efficient and sustainable practices.

Also, while designing one must adopt proven best practices that offer guaranteed efficiency and sustainability. Set up a **DCIM** system that can proactively monitor your data center and offer recommendations to improve the data center **PUE**. DCIM's can also be configured to automate some of the temperature control settings thus avoiding manual human errors which in turn will benefit the efficiency score. Take part in industry environmental groups to continue to learn the best practices of others and to share the best practices you discover.

In this article, we discussed a variety of issues faced by data centers, the current status, and the future technologies that can make data centers sustainable. We also highlighted some data center architecture design best practices that can be considered. We covered data center cooling technologies and how one can benefit by choosing the right fit for their design based on the data center location and resources available. We hope that this article will provide readers with interesting possibilities to explore in the realization of a view of the future of highly scalable, well-controlled, and energy-efficient, distributed virtualized data centers.

References

- Data center design history
 - https://en.wikipedia.org/wiki/Data_center#History
- Data center evolution
 - <https://www.cse.iitb.ac.in/~puru/courses/spring19/cs695/downloads/dcevolve.pdf>
(Cse.iitb.ac.in)
 - The Evolution of Data Center Architecture and Where We are Now
(<https://www.analyticsinsight.net/the-evolution-of-data-center-architecture-and-where-we-are-now/>)
 - The Evolution of Data Centers | Lifeline Data Centers
(<https://lifelinedatacenters.com/data-center/the-evolution-of-data-centers/>)
- Data center energy consumption report
 - Yevgeniy Sverdlik - Study: Data Centers Responsible for 1 Percent of All Electricity Consumed Worldwide (<https://www.datacenterknowledge.com/energy/study-data-centers-responsible-1-percent-all-electricity-consumed-worldwide>)
- Data center design best practices
 - Ali Marashi - Improving Data Center Power Consumption & Energy Efficiency
(<https://www.vxchnge.com/blog/growing-energy-demands-of-data-centers>)
 - Christy Pettey - 5 Steps to Maximize Data Center Efficiency
(<https://www.gartner.com/smarterwithgartner/5-steps-to-maximize-data-center-efficiency/>)
- Power Usage Effectiveness
 - PUE: A Comprehensive Examination of the Metric | The Green Grid
(<https://www.thegreengrid.org/en/resources/library-and-tools/20-PUE:-A-Comprehensive-Examination-of-the-Metric>)
 - What Is Power Usage Effectiveness? (<https://www.colocationamerica.com/blog/what-is-pue>)
- Carbon Usage Effectiveness
 - <https://airatwork.com/wp-content/uploads/The-Green-Grid-White-Paper-32-CUE-Usage-Guidelines.pdf>
- Green Data center
 - https://docs.media.bitpipe.com/io_10x/io_102267/item_1306461/Focus_Green_datacentre.pdf
- Data center cooling technologies
 - Kevin Heslin - A Look at Data Center Cooling Technologies
(<https://journal.uptimeinstitute.com/a-look-at-data-center-cooling-technologies/>)
- Free air-cooling
 - <https://www.akcp.com/blog/data-centers-free-air-cooling/>
- Liquid cooling technology
 - <https://www.asperitas.com/whitepapers/the-datacentre-of-the-future/>
- DeepMind AI Reduces Google Data Centre Cooling Bill by 40%
 - <https://deepmind.com/blog/article/deepmind-ai-reduces-google-data-centre-cooling-bill-40>

- Project Natick – subsea data center
 - <https://news.microsoft.com/innovation-stories/project-natick-underwater-datacenter/>
- EPEAT
 - https://en.wikipedia.org/wiki/Electronic_Product_Environmental_Assessment_Tool
 - Registry - <https://www.epeat.net/>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.