Best Practices for Sustainable Datacenters

V. Dinesh Reddy
Cognizant Technology Solutions

Brain Setz
University of Stuttgart

G. Subrahmanya V.R.K. Rao
Cognizant Technology Solutions

G.R. Gangadharan
National Institute of Technology Tiruchirappalli

Marco Aiello
University of Stuttgart

Datacenters are an essential utility of most modern organizations. These huge computing infrastructures consume large amounts of power, and their total energy consumption is estimated to be 2 percent of global consumption. Thus, it is important to minimize power usage in the design and operation of datacenters. The authors analyzed seven datacenters in India and the Netherlands and, based on their findings and industry standards, present a set of best practices to improve the datacenters' energy efficiency. Following some of these best practices, these datacenters have achieved 10 to 20 percent improvements in their energy consumption.

Datacenters constitute the core of an organization's information system by centralizing IT operations and equipment. The growth in hyper-scale cloud datacenters is one of the major contributors to the increase in electricity consumption across the globe. Investments in power solutions and datacenter services are expected to grow from $16.60 billion in 2015 to $26.33 billion by the end of 2020 at a compound annual growth rate of 9.67 percent.[1] Energy consumption of datacenters is rising by approximately 7 percent per year, which translates into large carbon emissions.[2] This consumption can be reduced by making datacenter operators aware of the wasted energy, developing utility incentive programs to make datacenters more efficient, and implementing best practices for sustainability.

This article presents the outcome of a study on the best practices toward a high level of scalability and sustainability in datacenters. We analyzed the current practices of seven datacenters in India and the Netherlands using multiple case research methodology. We compared the practices followed in these datacenters against the relevant standards on sustainable datacenters, identified design issues and operation inefficiencies, and provided recommendations for improvements in various operation levels of the datacenters.

Published by the IEEE Computer Society
1520-9202/18/$33.00 ©2018 IEEE

RESEARCH METHODOLOGY

This study is based on the multiple case research approach.[3] The features and infrastructure details of the datacenters used in the study are presented in Table 1. Among the seven datacenters, two are co-location datacenters and provide various services including public cloud services. Three datacenters are privately owned by companies and are used in the financial services sector running banking, financial, and insurance applications. The remaining two datacenters are those of academic institutions.

A uniform and standard data-collection methodology was adopted in each case, which included a standard questionnaire, a review of procedures, benchmarks (to confirm gaps), and interviews with key staff members (with five to eight key personnel having different designations). The standard questionnaire was based on the assessment of the following dimensions: energy efficiency, cooling, thermal and air management, greenness, and network and storage. For each dimension, issues and practices were monitored directly in the datacenter for the purpose of the study. Based on interview transcripts, we developed an ad hoc case study report, which was then distributed and discussed with the interviewees and other staff to gain insights and tailored feedback for the correct understanding of the status of the datacenter.

Table 1. Datacenter configurations.

	Rack space (number of racks)	**Datacenter space (in square meters)**	**Power capacity**	**Security level (zones)**	**Tier**	**Power usage effectiveness (PUE)**
Datacenter 1	5,000	21,300	30 MW	8	4	1.6
Datacenter 2	1,400	3,700	10 MW	6	4	1.7
Datacenter 3	800	1,850	6 MW	6	3	1.6
Datacenter 4	3,000	11,600	20 MW	6	3	1.6
Datacenter 5	1,000	2,800	10 MW	6	3	1.7
Datacenter 6	160	325	450 kW	Redacted	2	1.25
Datacenter 7	100	300	300 kW	Redacted	2	1.25

DATACENTER MANAGEMENT BEST PRACTICES

We considered the entire chain of operations of each datacenter and studied the best practices for sustainability followed in each datacenter across the following dimensions: energy efficiency, cooling, air and thermal management, greenness, and storage and network. Datacenter users are interested in the performance of their applications at different scales of utilization, whereas datacenter operators are interested in the efficiency of the resources. Table 2 summarizes the practices for sustainability followed by the seven datacenters (DC1, DC2, etc.) in our study.

Table 2. Best practices by datacenters for sustainability.

	DC1	DC2	DC3	DC4	DC5	DC6	DC7
Energy-efficiency practices							
Automation tools	✓	✓	✓	✓	✓		
Virtualization and consolidation	✓	✓	✓	✓	✓	✓	✓
Dynamic voltage and frequency scaling	✓	✓	✓	✓	✓	✓	✓
Handling comatose/zombie servers	✓	✓	✓	✓	✓	✓	✓
Raised thermostat set point	✓			✓		✓	✓
Controlled lighting with sensors	✓	✓		✓			
On-site power plant							
Datacenter infrastructure management tools	✓	✓		✓		✓	✓
Cooling, thermal, and air-management practices							
Central air handler	✓		✓	✓		✓	✓
Liquid cooling	✓						
Sensors for chiller plant	✓	✓	✓	✓	✓	✓	✓
Hot aisle/cold aisle containment	✓	✓	✓	✓	✓	✓	✓
Loop design for chillers	✓		✓	✓	✓	✓	✓
Adjustable speed drive chillers	✓	✓	✓				
Green practices							
Economizers (free cooling)	✓			✓		✓	✓
Reclaimed water for cooling							
Cooling water recirculation						✓	✓
Using renewable resources						✓	✓
Storage and network practices							
Storage tiering	✓	✓	✓	✓	✓	✓	✓
Automation to waste storage management	✓	✓	✓	✓	✓	✓	✓
Centralized control and storage optimization	✓	✓	✓	✓	✓	✓	✓
Software-defined flash storage	✓	✓		✓		✓	✓
Storage pooling and geo-replication	✓	✓	✓	✓	✓	✓	✓

Energy Efficiency Practices

Efficiency is defined as the ratio of the useful work done by a system to the total energy delivered to it. For datacenters, energy efficiency translates to the useful work performed by different subsystems. Following are some of the key steps to achieve energy efficiency in a datacenter.

Automation Tools

Datacenter automation tools help automate tasks such as provisioning, configuration, patching, release management, and compliance. Most of the datacenters we studied rely on automation tools that enable real-time optimization, reduce error rates, and improve the performance of the applications.

Virtualization and Consolidation

Virtualization enables abstracting physical servers in a datacenter facility along with storage, networking, and other infrastructure devices and equipment. Consolidation combines workloads from different machines into a smaller number of systems when servers are under-utilized and consume more energy.[4] All the datacenters in our study are virtualized and use different virtual machine consolidation and placement techniques to reduce power consumption and improve server utilization.

Dynamic Voltage and Frequency Scaling (DVFS)

DVFS reduces the power consumption of a processor on the fly by adjusting clock frequency according to current load, indirectly showing the reduction in the supply voltage.[5] Power saving can be achieved either by scheduling schemes with the capability of dynamic voltage and frequency or by consolidation techniques.

> We observed that decommissioning unused servers results in an energy savings of 50 percent.

Decommission Servers without Computing

Comatose or zombie servers are those that run applications that are no longer required or are unused, yet remain plugged in and operate continuously. Datacenter operators have to audit and root out comatose servers and duplicate applications, which account for up to 30 percent of deployed servers. Our study observed that decommissioning unused servers results in an energy savings of 50 percent.

Controlled Lighting

Installing a lighting control system in conjunction with more efficient fixtures and occupancy sensors can help reduce energy usage. Only three of the datacenters in our study were using resource-friendly timers that dim or shut off lighting when people are not present.

On-Site Power Generation

The critical need for clean and economical sources of energy is transforming datacenters that are primarily energy consumers to energy producers. On-site renewable power generation is an economical and eco-friendly solution for regions with high electricity and low natural gas prices, and for campus-like facilities that can re-utilize excess heating and cooling.[6,7] In these cases, one can utilize the grid power as a backup in combination with on-site generation systems such as solar panels and fuel cells as the primary source. However, none of the datacenters in our study used on-site power generation.

Integrated Monitoring

Energy monitoring allows greater visibility into overall datacenter energy usage while providing solutions to maximize server and infrastructure equipment operating efficiency. Many datacenters use datacenter infrastructure management (DCIM) tools to monitor energy and cooling efficiency and claim 20 percent savings in operational expenses.[8]

Cooling, Thermal, and Air-Management Practices

The datacenter cooling system can be considered a giant air pump where cooled air flows to reach the server inlets.[9] Therefore, energy efficiency of a datacenter can be improved by using better air-management practices in datacenters with raised-floor designs as well as non-raised-floor air-conditioning designs as follows.

Central Air Handler

Efficient airflow can be achieved by eliminating bypass and recirculation air flows, as this is where the airflow is wasted in a datacenter. For efficient airflow management, some of the datacenters in our study made use of custom-designed central air handler systems using variable speed drives. Further, many of the datacenters used loop design, with a median temperature of 10° C to 15° C and load-monitoring sensors for chiller plants.[10]

Hot Aisle/Cold Aisle Containment

In the aisle containment approach (see Figure 1), all the hardware in a row of cabinets faces the same way so that hot air is expelled on one side and cold air blows from the other side. All of the datacenters followed this containment approach that allows the proper flow of cold air to the destination, in turn reducing energy consumption.

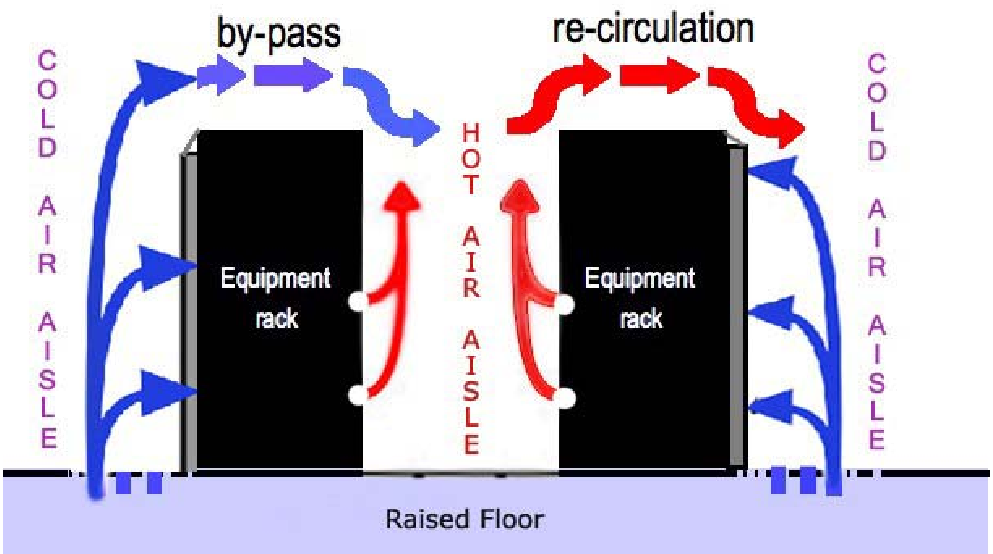

Figure 1. Hot aisle/cold aisle containment in datacenters.

Liquid Cooling Solutions

Liquid cooling solutions provide effective cooling and isolate equipment from the existing cooling system using a liquid in room-level and row-level systems. Only one datacenter in our study was using liquid cooling, as it is more expensive than air-based cooling.

Server Inlet Temperature and Humidity Control

The control of humidity in datacenters is essential to achieve high availability and reduce maintenance costs. The level recommended is around 50 percent or higher. However, datacenters without large high-speed fans can safely operate at 40 percent humidity levels, thus decreasing water and energy consumption.[11] Humidity can be best controlled knowing both inside and outside environmental conditions. Adiabatic humidification technology provides higher efficiency than infrared or isothermal technologies.[12]

Airflow Management

All the datacenters in our study used horizontal, vertical, under-rack panels, and PVC curtains for isolation with the goal of minimizing the recirculation of hot air. All the datacenters also used high-raised floors, overhead cabling, cable grouping, the placing of cable trays below the hot aisle, and cabling within the cabinets and racks to avoid air blockages. The best practice is to have dedicated horizontal airflows rather than a mixture of vertical and horizontal airflows, because dedicated horizontal airflows provide much more uniform distribution. During the inspection of the return air ducts for heating, ventilation, and air conditioning (HVAC), we observed inadequate ceiling height or undersized hot air return plenum in a few of the datacenters. Increasing the size of the return duct to match the air handler avoids this problem. Use of high overhead plenum and several feet of clearance under the raised floor provides better maintenance. Variable frequency fans in the computer room air handler (CRAH) units would allow for self-adjusting, thus resulting in energy savings. Most of the datacenters in our study were operating at baseline temperatures, but raising the baseline temperature would save 4 percent in energy costs with each degree of increase in the set point.[6]

The best practice for airflow management is to have dedicated horizontal airflows rather than a mixture of vertical and horizontal airflows.

Green Practices

A green datacenter incorporates energy-efficient design with high-efficiency power delivery, high-efficiency cooling, and increased utilization of renewable energy sources.[13,14] Some of the approaches and practices for developing green datacenters are as follows.

Economizers (Free Cooling)

Datacenters can achieve significant energy savings either through water-side or air-side economizers. Economizers make an impact when wet bulb temperatures outside the datacenter are less than 13° C for more than 3,000 hours per year. Some of the datacenters in our study claimed up to a 20 percent decrease in energy costs and a 7 percent decrease in maintenance costs since deploying economizers. However, the use of economizers depends on the geography, site conditions, and economizer design.

Reclaimed Water for Datacenter Cooling

The use of reclaimed or "gray" water is neither harmful to the environment nor to human health. Using gray water for cooling is considered eco-friendly because it reduces demands for ground water and does not require energy for the recycling process at waste-water treatment sites.

Cooling Water Recirculation

Using the same water for several cycles of cooling operations reduces water consumption. The water savings improve datacenter energy efficiency and lower the impact on the environment and on potable water supplies while cutting costs.

Renewable Resources

Renewable energy comes from solar panels, wind turbines, or hydroelectric installations. As renewable energy production is intermittent in nature and depends on location, it is often combined with energy storage facilities. Nevertheless, these are still expensive installations. The datacenters in our study were not using green energy.

STORAGE AND NETWORK PRACTICES

Generally, a datacenter is seen as a facility used to house computer systems and associated components, such as telecommunications and storage systems. Some of the best practices followed for storage and network communications of a datacenter are described as follows.

Most of the datacenters we surveyed have centralized control over the servers, storage, and databases for storage optimization. A solid-state drive helps reduce the energy consumption of spinning disks and handles the enormous demand on storage systems. Some of the datacenters in our study were using pooling storage, hybrid storage, and flash cache. Most of the datacenters were using geo-replication for storage backup, e-discovery, and data mapping for archiving, whereas flash storage was used for specific applications in very few of the datacenters in our study.

The datacenters in our study used automation tools to predict network loads and avoid outages. Some were using network virtualization, which allows each customer to have their own network with different controller applications and balances the performance, port utilization, and traffic demands. As software-defined networking (SDN) is the virtualization of networking and storage infrastructure, it offers resource flexibility, optimal resource usage, and scalability.[15] Some of the datacenters had either adopted SDN or are planning to adopt SDN in the near future.

RECOMMENDATIONS FOR DATACENTER OPERATORS AND IT PROFESSIONALS

Datacenter lifecycle management helps owners understand key management tasks, connections between different phases, and the pitfalls that exist in each phase. Generally, the datacenter lifecycle is comprised of five phases: plan and analyze, design, build, operate, and continuous evaluation. For initial phases, it is better to use reference designs to validate the early project choices and develop system concepts. Considering the whole chain of operations, datacenter operations become the base layer with the goal of optimizing not only energy and cost, but also of helping with the long-term planning and provisioning of equipment and resources.

Nowadays, cloud computing is of strategic importance, benefitting both providers and their customers. If a new datacenter is under-utilized, it can act as a cloud provider for other datacenters and customers. To accommodate the growing demands of users and other background processes using the same physical resources, datacenters are required to make optimal use of all the resources by increasing utilization and visibility. Proper selection of virtual machines for migration minimizes the number of power-on nodes. Designing and implementing fast energy-efficient virtual machine allocation and selection algorithms considering multiple resources can result in energy-efficient datacenters.

Datacenter infrastructure management or automation tools can achieve considerable energy savings, from 5 to 20 percent.

Maintaining a separate, direct current feed to power the telecommunications and storage systems directly will reduce energy consumption, real estate costs, and conversion losses. DCIM or automation tools can achieve considerable energy savings, ranging from 5 to 20 percent.[8] Centralized cooling systems in a large-scale datacenter can be optimized by maintaining a median temperature of 10° C to 15° C, using adjustable-speed drive chillers, storing excess thermal energy, and installing energy- and load-monitoring sensors.

Following these recommendations for chiller plants will have quick return on investment in the order of two to three years in terms of energy savings and diagnosis. Hot and cold aisle containment, increasing the datacenter supply air temperatures, using waterside economizers, and increasing the room temperature reduce the need for cooling provided by CRAH units. Further, using a higher temperature-chilled water supply can provide sufficient cooling for a datacenter and reduces the number of hours of compressor-based cooling.

To improve power usage effectiveness (PUE), [16] it is important to understand normal power consumption and find unusual conditions across the datacenters. It is also important to develop compute efficiencies by server type, adjust operations according to peak power utilization, and shift resource usage based on known use profiles. Operators need to correlate infrastructure investments more closely to actual application requirements.

In short, optimizing cooling plant, operational parameters of the datacenter, interruptible power supply load, and zombie servers along with controlled lighting, continuous monitoring, and proper airflow management can improve the energy efficiency of a datacenter. Placing archived data on slower, larger drives that use less power can also save associated energy costs.

In our study, datacenter 3 (DC3) claimed a reduction of 20 percent in their PUE score by following best practices in cooling, thermal and air management, and storage and network practices. DC1 claimed a PUE reduction of 20 percent by strictly adhering to energy-efficiency practices such as virtualization, consolidation, and automation tools. DC2 and DC4 claimed a PUE reduction of 10 percent by effectively following the best practices of energy efficiency and cooling, thermal, and air management.

Table 3 summarizes various implementation issues in the dimensions we studied that lead to inefficiencies and the solutions/practices to avoid these problems and improve datacenter efficiency.

Table 3. Issues and solutions in various dimensions of datacenters.

Dimension	Implementation issues	Solutions/practices
Energy efficiency	Some virtual machines are allocated more resources than requested to achieve high performance, leading to inefficient use of resources.	Virtualization and consolidation; decommission servers without any computing
	There are several power conversion steps while delivering power to IT equipment, leading to losses in power distribution.	On-site power plant or direct DC feed to power ICT equipment
	Most datacenters do not monitor detailed energy use, which can help operators find the actual point of losses and inefficiencies.	Installing datacenter infrastructure management (DCIM) tools; detailed submetering at the component level
	Power savings and performance requirements can lead to service level agreement (SLA) violations.	Placement policies supporting performance and SLAs

Cooling, thermal, and air management	High-density racks might result in one or more areas of excess temperature known as "hot spots," which result in equipment damage.	A global, ITIL-defined capacity plan
	Poor air management can lead to bypass and recirculation air flow, which reduces the cooling efficiency.	Placing perforated tiles near hot spots; utilizing blanking panels
	Air blockages lead to poor flow of cool air to the server inlets.	Using hot aisle/cold aisle layout, PVC curtains for isolation, seal cable cutouts, or other openings in under-floor distribution systems; using overhead cabling technology
	High heat loads from racks restrict space utilization.	Central air handler, loop design for chillers, adjustable-speed drive chillers and economizers
	Improper configurations lead to inefficient use of chillers.	Installing energy- and load-monitoring sensors
	Oversized cooling infrastructure limits operating capacity.	Increasing the datacenter supply air temperatures and room temperature
	Inadequate duct size results in poor airflow.	Using waterside economizer and higher temperature-chilled water
Greenness	Datacenters consume a lot of water but fail to recycle or use water in an efficient manner.	Cooling water recirculation and using reclaimed water
	Recycling of electronic equipment and other materials are not in place.	Recycling ICT equipment through certified authorities
Network and storage	Storage over-provisioning and massive volumes of redundant data lead to inefficient use of storage and consume more energy.	Automating waste storage management

	Increasing virtual machines consumes more storage.	Storage tiering centralized control and optimization
	High throughput and performance requirements limit the efficient use of storage and network.	Joint host-network traffic management and power optimization in software-defined networks (SDNs); software-defined flash storage and storage pooling
	Connectivity issues might lead to outages.	Using traffic engineering techniques to alleviate loss and reduce congestion

CONCLUSION

This study describes best practices to improve operational efficiency and ways to create a sustainable datacenter. We studied issues in seven datacenters in India and the Netherlands and described the best practices to deal with these issues. Implementing and fine-tuning these practices in new and existing facilities will improve datacenter efficiency.

ACKNOWLEDGMENTS

This research received funding from the Netherlands Organization for Scientific Research (NWO) in the framework of the Indo Dutch Science Industry Collaboration program in relation to project NextGenSmart DC (629.002.102).

REFERENCES

1. *Global Datacenter Power Market—Growth, Trends, and Forecast (2015-2020)*, report, Mordor Intelligence, August 2016.
2. P. Jones, *DCDi Industry Census 2014*, Data Center Dynamics, 18 July 2014; http://www.datacenterdynamics.com.
3. R.K. Yin, *Case Study Research: Design and Methods*, Sage Publications, 2009.
4. T.C. Ferreto et al., "Server Consolidation with Migration Control for Virtualized Data Centers," *Future Generation Computer Systems*, vol. 27, no. 8, 2011, pp. 1027–1034.
5. G. Magklis et al., "Dynamic Frequency and Voltage Scaling for a Multiple-Clock-Domain Microprocessor," *IEEE Micro*, vol. 23, no. 6, 2003, pp. 62–68.
6. N. Ahuja, "Datacenter Power Savings through High Ambient Datacenter Operation: CFD Modeling Study," *Proc. 28th Ann. IEEE Semiconductor Thermal Measurement and Management Symposium* (SEMI-THERM), 2012, pp. 104–107.
7. J. Tu et al., "Dynamic Provisioning in Next-Generation Data Centers with On-site Power Production," *Proc. 4th Int'l Conf. Future Energy Systems* (e-Energy), 2013, pp. 137–148.
8. P. Gilbert, K. Ramakrishnan, and R. Diersen, *From Data Center Metrics to Data Center Analytics: How to Unlock the Full Business Value of DCIM*, white paper, CA Technologies, 2013.
9. M. Beitelmal and C. Patel, *Model-Based Approach for Optimizing a Data Center Centralized Cooling System*, technical report, HP Labs, 2006.
10. T. Hartmann, "'LOOP' Chiller Plant Design Dramatically Lowers Chilled Water Costs," *Renewable and Advanced Energy Systems for the 21st Century*, Renewable and Advanced Energy Systems for the 21st Century, 1999.

11. *Data Center Operational Efficiency Best Practices—An IBM Study*, report, IBM Global Technology Services, 2012.
12. J. Colby and J. Herzing, "The Why, What, and How of Datacenter Humidification," *Engineered Systems*, 1 February 2016; https://www.esmagazine.com/articles/97562-the-why-what-and-how-of-data-center-humidification?v=preview.
13. Q. Gu et al., "A Categorization of Green Practices Used by Dutch Data Centers," *Procedia Computer Science*, vol. 19, 2013, pp. 770–776.
14. S. Murugesan and G.R. Gangadharan, *Harnessing Green IT: Principles and Practices*, Wiley-IEEE Computer Society Press, 2012.
15. D. Drutskoy, E. Keller, and J. Rexford, "Scalable Network Virtualization in Software-Defined Networks," *IEEE Internet Computing*, vol. 17, no. 2, 2013, pp. 20–27.
16. V. Dinesh Reddy et al., "Metrics for Sustainable Data Centers," *IEEE Trans. Sustainable Computing*, vol. 2, no. 3, IEEE Transactions on Sustainable Computing, 2017, pp. 290–303.

ABOUT THE AUTHORS

V. Dinesh Reddy is currently associated with Cognizant Technology Solutions. His research interests include cloud computing and energy efficient technologies. Contact him at dineshvemula@gmail.com.

Brian Setz is a PhD student at University of Stuttgart. His research interests include cloud computing, Internet of Things, and datacenter efficiency. Contact him at briansetz@gmail.com.

G. Subrahmanya V.R.K. Rao is currently associated with Cognizant Technology Solutions as Associate Vice President of Technology. He is a Senior Member of IEEE and ACM. Contact him at subrahmanyavrk.rao@cognizant.com.

G.R. Gangadharan is an associate professor at National Institute of Technology, Tiruchirappalli, India. His research interests focus on the interface between technological and business perspectives. He is a Senior Member of IEEE and ACM. Contact him at geeyaar@gmail.com.

Marco Aiello is a professor of service computing, University of Stuttgart. His research interests include distributed systems, smart energy systems, pervasive computing, service-oriented computing, and spatial reasoning. Contact him at aiellom@ieee.org.