Data Centers
Shared Solutions Marketplace

EVIDENCED-BASED BEST PRACTICES AROUND DATA CENTER MANAGEMENT

LESSONS LEARNED FROM THE PUBLIC AND PRIVATE SECTORS

General Services Administration
Office of Governmentwide Policy
Office of Information, Integrity and Access (ME)

August 11, 2016

Contents

BEST PRACTICES: AT A GLANCE 3
INTRODUCTION 4
Purpose 4
Why *These* Best Practices? How Were They Identified? 5
Using This Document 5
DCOI METRICS 7
Energy Metering 7
Power Usage Effectiveness 11
Virtualization 23
Server Utilization & Automated Monitoring 25
Facility Utilization 28
ADDITIONAL AREAS FOR IMPROVEMENT 30
Management/Operations 30
Sustainability 33
Security 34
Reliability/Uptime 39
Change Management 41
ADDITIONAL TOOLS AND RESOURCES 43

Data Centers
Shared Solutions Marketplace

BEST PRACTICES: AT A GLANCE

Energy Metering	▪ Use a **DCIM** software tool to help manage the data center ▪ **Automate power metering at the device level** to track power demand and usage ▪ Install **wireless sensor technology** to observe real-time data center conditions
Power Usage Effectiveness	▪ Arrange servers in separate **hot and cold aisles** to prevent air commingling ▪ Install **variable speed drives** to match energy usage to workload ▪ Let the supply air reach **higher temperatures** to decrease demand on the chiller ▪ Increase the **allowable humidity range** to lower humidification/dehumidification needs ▪ **Optimize airflow** with blanking panels, grommets, and cable and tile management ▪ Utilize **free cooling** to support energy-intensive mechanical chillers ▪ Use **direct liquid cooling** to cool hot exhaust air more efficiently ▪ **Modernize lighting systems** to cut energy costs
Virtualization	▪ **Right-size** after virtualization to match infrastructure power use to IT workload
Server Utilization	▪ **Turn off zombie servers** to stop wasting money and energy ▪ Deploy **server power management** technology for better power proportionality
Facility Utilization	▪ Use **rack and row based cooling** to avoid stranded capacity
Management/ Operations	▪ Properly **train employees** to increase operational efficiency and effectiveness ▪ **Benchmark** to track performance over time
Sustainability	▪ Procure **Green IT** to conserve energy and reduce a facility's environmental impact
Security	▪ Practice a **defense-in-depth** security approach to protect physical and virtual assets ▪ Develop a **disaster recovery plan** to ensure operations continue in an emergency ▪ Develop a **security awareness curriculum** to reduce vulnerability to man-made threats
Reliability	▪ Increase **automation capabilities** to avoid human error and increase uptime
Change Management	▪ Get **stakeholder support** when implementing change to avert resistance

INTRODUCTION

Purpose

The Data Center Optimization Initiative (DCOI) memorandum, M-16-19, released by the Office of Management and Budget (OMB) lays out a series of targets for data centers that Federal agencies must meet by the end of Fiscal Year 2018. These targets encompass cost savings, consolidation and closure of existing facilities, as well as improvements in five key optimization metrics.[1] The General Services Administration (GSA) Office of Governmentwide Policy (OGP) has created this document and the DCOI Community of Practice to help guide agencies as they work towards reaching those goals.

Help Data Centers Reduce Costs and Improve Performance

These best practices are directed at policy managers, data center program managers, facilities managers, and sustainability officers at pre-existing, brick-and-mortar Federal data centers. It is designed to support them in researching possible next-steps for optimizing their data centers and achieving cost savings, as required by M-16-19.

Help Agencies Identify Data Centers That Are Unable to Improve

This document is further intended to help agency leadership evaluate the appropriateness of dedicating more resources to particular data centers: if a data center cannot satisfactorily improve its performance and ultimately reach the mandatory targets by following some of these best practices in a cost-efficient manner, it may be a prime candidate for helping an agency meet its federal closure and consolidation requirements.

Motivate Discussion That Helps Identify More and Better Practices

Beyond serving as a resource, this document is also meant to help start a conversation about best practices. As the Data Center Shared Services (DCSS) Managing Partner, OGP has created a Community of Practice (CoP) for agencies to discuss the shared services marketplace. This CoP is intended to serve as an outlet for agencies to, among other purposes, leverage and provide feedback on best practices. Data center practitioners are encouraged to use this guide as a starting-point for thinking systematically about their own best practices, and are invited to share their thoughts and perspectives with the community.

[1] The five optimization metrics are: Energy Metering; Power Usage Effectiveness; Virtualization; Server Utilization & Automated Monitoring; and Facility Utilization

Why *These* Best Practices? How Were They Identified?

While there are many good practices surrounding data center operations and management, the practices included here are "high value" options that will enable agencies to make the most progress towards meeting DCOI targets. Best practices in this guide cover the five metrics directly addressed in M-16-19, as well as five additional areas that are critical to being a well-functioning, 21st century data center. A practice is listed only if it has a proven track record of markedly increasing the efficiency and effectiveness of data centers. To illustrate their potential value, each best practice includes a brief description of a data center that benefited from it.

Input on best practices was obtained from public, private, and nonprofit sector sources. Federally-owned and managed data centers were ranked based on data reported to OMB, and top performers in each metric were interviewed about their strengths and weaknesses. Open-source information from technology giants, such as Amazon, Google, and Facebook, as well as leaders in data center technology—including Schneider Electric, Emerson Network Power, and Nlyte Software—were surveyed to determine trends among the world's highest performing data centers, and then assessed to see which practices could be transferrable to the federal context. Standards and guidance issued by advisory organizations like The Green Grid and the Uptime Institute further helped to determine the value potential and technical feasibility of candidate best practices.

Using This Document

What this document is:

- This document is an overview of the best practices that are most commonly credited by high-performing data centers with helping them to achieve their high effectiveness, efficiency, and productivity.

- This document is intended to help Federal data centers begin exploring optimization techniques by using the suggested tools and resources provided for each best practice.

What this document is not:

- This document is not a prescriptive list of practices that all Federal data centers should adopt. Before implementing any of these best practices, a data center should first conduct a thorough assessment of the practice's suitability and feasibility based on their facility's own needs and resources.

- This document is not meant to serve as a technical reference or implementation manual. Rather, it offers a high-level explanation of how and why a best practice leads to improved performance.

- This document is not a complete or permanent list of data center best practices. Individual data centers may have best practices that are not listed here, and as technology changes, best practices should be expected to change.

This document makes use of icons to emphasize best practices with particular characteristics:

 indicates that a best practice can lead to significant cost savings

 indicates that a best practice can help curtail energy usage substantially

 indicates "low hanging fruit," or best practices that are comparatively low-cost and easy to implement

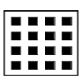 indicates best practices that are typically cost-effective only for data centers able to obtain economies of scale

Finally, a  draws attention to potential roadblocks that might arise when implementing a best practice.

DCOI METRICS

Energy Metering

Definition: Percent of total gross floor area (GFA) in an agency's tiered data center inventory located in tiered data centers that have power metering.
DCOI FY 18 Target: 100%

Use a DCIM Software Tool/Solution

Data Center Infrastructure Management (DCIM) tools represent the convergence of the previously-isolated IT and building facilities functions within an organization. The goal of DCIM is to provide administrators with a holistic view of a data center's performance, so that they can make and monitor fully-informed decisions that will enable energy, equipment, and space to be used as efficiently as possible. Per the Data Center Optimization Initiative, all Federal data centers will be required to have automated monitoring, inventory, and management tools (e.g. DCIM) by the end of fiscal year 2018.

A standard DCIM software package offers capabilities that can be used to improve performance in metering, Power Usage Effectiveness (PUE), server and facility utilization, and virtualization:

DCIM makes benchmarking and analyzing performance easier. DCIM collects and aggregates data reported by energy meters, sensors, and IT equipment in a facility, and then generates easy-to-interpret visualizations and detailed reports. These outputs are useful for monitoring current conditions, becoming aware of and diagnosing problems and inefficiency, and tracking trends over time.

DCIM tools help with capacity and change management, and offer insights that can help improve facility and server utilization. A typical DCIM tool will include inventory capabilities, which are useful for prioritizing equipment refreshes and drawing conclusions about the ability to scale up or down. DCIM can also track and highlight stranded infrastructure capacity, underutilized floor space, and over or under-provisioned servers. This capacity monitoring capability is particularly useful given increased rates of virtualization: a DCIM solution working in conjunction with a virtual machine tool can direct virtualized servers to move amongst physical servers in the most efficient manner possible, accounting for available space and cooling and power capacity.

A DCIM solution can be costly to purchase and implement, with the price dependent on the level of functionality built into the suite. Small-budget data centers that want to obtain the many benefits of DCIM can turn to free or low-cost options available on the internet. These "DCIM-light" products do not include all the capabilities of a fancier package, but typically have the standard tools for inventorying, tracking, and room and systems mapping.

Example: An Ohio bank acquired by PNC Financial Services deployed a DCIM suite in its data center to enable better asset management. Based on the real-time data gathered from the DCIM solution, the data center was able to identify and decommission 500 comatose servers, which saved .5 MW of power and added 10 years to the expected life of the facility. The cost savings associated with in-house management of power, cooling, and space assets made it possible for the data center to go from spending 70% of its budget on operations to spending only 40%, freeing up funds for development opportunities.[2]

Recommended Tools and Resources:

- Gartner. "Magic Quadrant for Data Center Infrastructure Management Tools." 2015.
- Taylor, Allen G. "DCIM for Dummies." 2015.
- OpenDCIM's Data Center Infrastructure Management application
- Ralph CMDB Project
- Sunbird. "40 Critical Problems Your DCIM Should Solve." 2015.
- U.S. National Institutes of Health. "Sustainable Data Center Design Guide." 2013.

Automate Power Metering at the Device Level

Executive Order 13693 requires Federal data centers to have advanced energy metering equipment capable of tracking Power Usage Effectiveness (PUE) by the end of fiscal year 2018. The Green Grid, developer of the PUE metric, defines three levels of PUE metering; each successive level calls for increased monitoring and automation. Level Two, intermediate metering, meets the Federal minimum requisite for PUE tracking by requiring daily, automated data collection at the level of the building and some systems. The best practice among high-performing data centers, however, is to measure energy use more frequently and at a deeper level.

Energy-use data helps to track long-term trends, benchmark progress and performance, perform diagnostics, and alert data center operators to impending or ongoing system failures.

[2] Nlyte Software. "Case Study: PNC Financial Services." (PDF)

The more detailed and frequently obtained the data, the more insight into the data center that can be gleaned. Advanced (Level Three) metering encourages automated monitoring of power use at a device level: electricity is tracked from its entry into the facility down to individual components of a data center (PDUs, HVAC equipment, IT equipment, etc.). Such detailed metering enables data center operators to pinpoint trouble spots and find excess capacity, maximizing the possibilities for energy and cost savings (especially when combined with a DCIM tool).

Advanced metering requires the largest upfront capital expenditure. If deploying an advanced network of meters at once is prohibitively expensive for your sized data center, it can still make gradual advances towards this best practice by taking advantage of periodic equipment refreshes to purchase “smart” equipment with built-in metering capabilities. In the interim, Level Two metering practices can be followed.

Example: A software tool with the ability to monitor energy usage at the device level was tested at a data center at Lawrence Berkeley National Laboratory. Because of the precise level of monitoring, researchers were able to observe exactly where power was being demanded and utilized, as well as analyze power usage by device category, location, etc. The study found that by tracking and reducing energy-consumption at the device level, the data center was able to find opportunities for reducing energy use at the systems and building level. For example, by virtualizing just 25 under-utilized servers, the facility could save over $6,000 a year and reduce its CO_2 emissions by 30 tons.[3]

Recommended Tools and Resources:

- U.S. Department of Energy. “Data Center Metering and Resource Guide V2.3.” 2016.
- U.S. Department of Energy. “Metering Best Practices: A Guide to Achieving Utility Resource Efficiency, Release 3.0.” 2015. (PDF)
- ASHRAE. “Data Center Power Equipment Thermal Guidelines and Best Practices.” 2016. (PDF)
- The Green Grid. “PUE: A Comprehensive Examination of the Metric.” 2012. (PDF)
- Lawrence Berkeley National Laboratory. “Data Center Profiler Tools.”
- Lawrence Berkeley National Laboratory. “Data Center Master List of Energy Efficiency Actions.” 2016.

Install Wireless Sensor Technology

The deployment of sensors enables a data center operator to obtain real-time information about power usage, temperature, humidity, and air pressure within the data center. Recent

[3] Lawrence Berkeley National Laboratory and ANCIS Incorporated. “Case Study: JouleX Energy Management (JEM) Solution at Lawrence Berkeley National Laboratory (LBNL).” 2012.

technological advances have made wireless sensor technology the data collection tool of choice for many high performing data centers. Compared to traditional wired sensors, wireless technology is cheaper, easier to install, and easier to move, while still retaining the essential level of reliability.

A wireless mesh network of sensors requires the placement of a metering node wherever power usage or environmental conditions will be measured. The cost of a wireless node can be as little as 1/10th the average cost of a wired node, and the total savings for a full network quickly add up.

Wireless sensors are easier to install than wired sensors, making initial deployment and future expansion quicker and less disruptive to data center operations. Although the installation of power meters typically requires the temporary shutdown of a center's IT equipment, this process can be scheduled around pre-planned downtime.

Wireless mesh networks are valued for their reliability: built-in redundancy controls ensure that the failure of one node does not disable the entire system. An additional benefit of these networks is the absence of wires that face physical routing constraints, and that can impede airflow within the data center.

⚠ Federal data centers considering wireless sensor networks should confirm that such technology is permitted under their Agency's security policies.

Example: A study was conducted to research possible ways of utilizing real-time performance data to improve energy-efficiency in Federal data centers. Wireless sensors that tracked temperature, humidity, power usage, and pressure were installed at the U.S. Department of Agriculture's Tier 3 National Information Technology Center (NITC) Data Center. Based on the data obtained from the sensors, researchers used a custom-designed software tool to identify potential high-payoff energy-efficiency measures. After implementing the suggested measures, the data center achieved significant energy-efficiency improvements, including a 48% reduction in cooling load; a 17% reduction in facility power usage; and a drop in PUE from 1.83 to 1.51.[4]

Recommended Tools and Resources:

- U.S. General Services Administration. "Wireless Sensor Network for Improving the Energy Efficiency of Data Centers." 2012. (PDF)
- Lawrence Berkeley National Laboratory. "Guidelines for Data Center Energy Information Systems. DRAFT." 2014. (PDF)
- U.S. Federal Energy Management Program. "Wireless Sensors Improve Data Center Energy Efficiency." 2010. (PDF)
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

[4] General Services Administration. "Wireless Sensor Network for Improving the Energy Efficiency of Data Centers." 2012. (PDF)

Power Usage Effectiveness

Definition: Proportion of total data center energy used by IT equipment (calculated on a per-data center basis)
DCOI FY 18 Target: ≤ 1.5 (≤ 1.4 for new data centers)

Arrange Servers in Separate Hot and Cold Aisles

Arranging servers in a hot aisle/cold aisle setup is rapidly becoming a top recommended standard in data center management.[5] Hot aisles and cold aisles are formed by taking advantage of the front-to-back airflows common to most servers. To form a cold aisle, the front sides of a row of servers are arranged to face the front sides of another row of servers, with cold air vents positioned between. The back sides of those same servers are then arranged directly opposite the back sides of another row of servers, forming a hot aisle into which warm exhaust air is expelled.

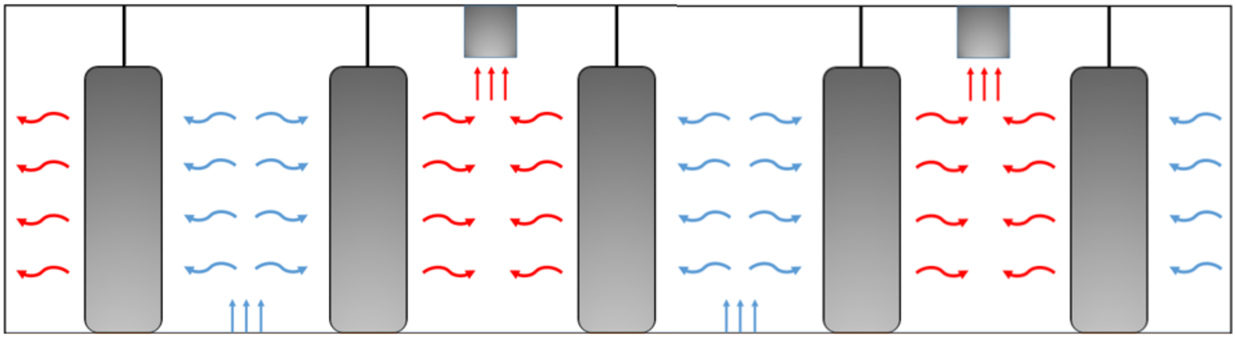

Figure 1: Side View of a Hot Aisle/Cold Aisle Arrangement

Hotspots develop in data centers when the heat in a localized area is not sufficiently cooled, either because the cold air cannot reach it or is not at a low enough temperature. Keeping these airflows separate allows a data center to raise the intake air temperature without worrying about hot spots pushing the ambient room temperature above the allowable level. This setup also keeps the return air hotter, which increases the temperature differential between air as it enters and leaves the cooling system (known as the Delta T), which improves the system’s efficiency. When the cooling system functions more efficiently, a single cooling unit can cool more air than it did before, with possible efficiency gains of 20% or more. This improved efficiency enables data centers to turn off now-unused cooling units for significant savings associated with both operating costs and equipment replacement costs.

---

[5] TIA/ANSI Standard 942

To maximize the benefits from a hot aisle/cold aisle arrangement, a data center should implement either a Hot Aisle Containment System (HACS) or a Cold Aisle Containment System (CACS) to physically separate the cold air being drawn into a server from the hot air being expelled back into the environment. In the absence of a physical barrier, up to 50% of chilled air produced by the cooling systems is wasted when it mixes prematurely with the exhaust air from IT equipment. Also, such systems make it possible to limit cooling to areas within barriers, and thus there is little or no waste cooling in aisles.

The HACS or CACS can be partitioned off with flexible strip curtains or rigid enclosures. HACS and CACS will not work as intended if air can leak between the aisles, so beyond filling gaps in the racks with blanking panels, floor-to-ceiling containment requires barriers both above and below the racks. Data centers should choose only one of these systems—using both HACS *and* CACS is unnecessary and offers no additional benefits.

⚠ Data centers planning to install rigid enclosures must verify that such equipment does not interfere with fire suppression systems or violate building fire codes.

⚠ When deciding between HACS and CACS, consideration should be given to applicable health and safety regulations. In a CACS, the hot exhaust air is permitted to flow freely into the computer room, and room temperatures may rise above what an individual can comfortably and safely work in for significant periods of time.

Example: In response to new state regulations, The University of California needed to identify energy-savings opportunities for its 19,000 square foot computer room (known as the West Room) in the San Diego Supercomputer Center (SDSC), as well as the new 14,5000 square foot East Computer Room then under construction. SDSC managers opted to implement hot aisle containment systems in the rooms: flexible vinyl curtains were installed in the West Room to accommodate the varying sizes of the racks, and lowered the room's PUE from 1.8 to 1.5. The East Room's identical server racks made it a candidate for rigid enclosures which, when deployed, lowered the room's projected PUE from 1.5 to 1.3.[6]

Recommended Tools and Resources:

- ENERGY STAR. "Containment/Enclosures."
- Schneider Electric. "Implementing Hot and Cold Air Containment in Existing Data Centers." 2013. (PDF)
- Schneider Electric. "Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency." 2015. (PDF)
- The Uptime Institute. "Implementing Data Center Cooling Best Practices."
- The Uptime Institute. "A Look at Data Center Cooling Technologies."
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)
- U.S. National Institutes of Health. "Sustainable Data Center Design Guide." 2013.

[6] Emerson Network Power. "San Diego Supercomputer Center at the University of California." 2011. (PDF)

Install Variable Speed Drives

Traditionally in HVAC equipment, electricity is supplied and thus equipment operates at a constant speed, independent of workload requirements. Mechanical throttling devices can be used to conserve energy during low-demand times, but limiting the flow of air or water through the unit has no impact on the substantial amount of power consumed by the motor. In an effort to minimize energy waste from motors running at full-speed during off-peak periods, many high-performing data centers have installed Variable Speed Drives (VSD) on fans and pumps.

A VSD enables variability in the frequency with which electricity is supplied to a motor, which in turn allows the motor to spin at different speeds. With a VSD, a motor can thus react to demand fluctuations and deliver power at a level proportional to what is required. The cooling system is one of the most energy-intensive processes in a data center, and using a VSD to stop chiller fans and pumps from consuming excess energy by default can dramatically improve a facility's PUE. Some data centers will be able to retrofit their Computer Room Air Conditioning (CRAC) or Computer Room Air Handler (CRAH) units with VSD tools provided by the manufacturer, while others should, during an equipment refresh cycle, look into purchasing HVAC units with the technology built-in.

VSDs can also be installed on other motors in the data center that service equipment with fluctuating workloads. Servers, for example, are not constantly called on to perform the same level of computing power. When demand is low, a server does not require the same intensity of cooling to keep its internal temperature within allowable limits: consequently, a VSD will enable the server fan to slow down to the point where it expends less excess energy on unneeded cooling. Many newer servers now come with VSDs built in, so agencies may wish to call for VSDs as they purchase new servers as part of equipment refresh or expansion.

Example: BNY Mellon, a top banking and financial services company, installed variable speed drives on its CRAC units in a 71,000 square foot data center. The VSDs enabled the CRAC fans to match speed to workload demand, which resulted in a 25% drop in fan speed. This speed decrease in turn led to a 58% reduction in fan energy consumption,[7] in line with Energy STAR findings of a proportionate relationship between fan speed and electricity consumption (approximately a 45% drop in power consumption for a 20% reduction in fan speed).[8]

Recommended Tools and Resources:

- ABB. "Technical Guide No. 4: Guide to Variable Speed Drives." (PDF)

[7] ENERGY STAR. "BNY Mellon's Commitment to Reducing Data Center Energy Use." (PDF)
[8] ENERGY STAR. "Variable Speed Fan Drives."

- Carbon Trust. "Variable Speed Drives: Introducing Energy Saving Opportunities for Businesses." 2007. (PDF)
- U.S. Federal Energy Management Program. "Variable-Speed Fan Retrofits for Computer-Room Air Conditioners." 2013. (PDF)
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)
- U.S. National Institutes of Health. "Sustainable Data Center Design Guide." 2013.
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

Permit the Supply Air to Reach Higher Temperatures

It may be advantageous to buy higher priced equipment that functions within a broader operating temperature range to save on cooling costs. Contrary to popular belief, data centers no longer need to run within the 60-70° Fahrenheit range that led to their common "meat locker" moniker. Technological advances have widened the temperature limits within which IT equipment can effectively and reliably operate. The best practice for realizing energy and cost savings is to aim for the high-side of the recommended envelope for data center temperature ranges given by the most current ASHRAE[9] Standards.

When a higher server air intake temperature is permitted, cooling systems are able to run less often and at lower cost, since they no longer must chill air to such low temperatures. Department of Energy-ENERGY STAR estimates that a 1° increase in server intake temperature can lead to 4-5% savings in energy costs. Higher intake temperatures also permit increased water temperatures in chillers, CRACs, and CRAHs, which provides another opportunity to reduce energy consumption.

To prevent unwanted/unneeded cooling, the fan speed of CRAC/CRAH equipment should be based on server air intake temperatures, rather than exhaust air temperatures. Sensors should be placed in front of servers, and not behind them, in order to insure air intake temperatures are accurately measured. In order to continue to obtain efficiency benefits after the initial temperature adjustment, data centers should conduct periodic sensor recalibration.

Before raising intake air temperatures, check manufacturer guidelines for the specific equipment operating in a facility. Older equipment cannot necessarily work reliably in warmer conditions, and may restrict potential temperature increases.

[9] American Society of Heating, Refrigerating, and Air-Conditioning Engineers

If internal server fan controls are not modified prior to making temperature adjustments, higher intake temperatures—and thus higher server temperatures—will force the fans to work harder. Continuously-running fans will undermine any energy savings that could come from higher air intake temperatures.

Example: A data center colocation provider hoping to cut down on its electricity consumption decided to raise the temperature in its facilities in line with new guidance released by ASHRAE. Taking into account the temperatures at which its equipment could reliably operate, over a period of six months the co-lo gradually increased the facility's internal temperature by 9° Fahrenheit. The thermostat adjustment led to a drop in power use by the chillers from .53kW/ton to .32kW/ton, and saved the colocation company over $285,000 annually.[10]

Recommended Tools and Resources:

- ASHRAE. "Data Center Power Equipment Thermal Guidelines and Best Practices." 2016. (PDF)
- ASHRAE. "Data Center Storage Equipment: Thermal Guidelines, Issues, and Best Practices." 2015. (PDF)
- The Uptime Institute. "Implementing Data Center Cooling Best Practices."
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)

Increase the Allowable Humidity and Dew Point Temperature Ranges

As with air intake temperature, the best practice concerning data center humidity ranges is to expand them in line with current ASHRAE standards. Rather than targeting a specific humidity point, data center operators should determine the recommended minimum and maximum absolute humidity level and dew point temperature at which their equipment can reliably function, and allow the level to fluctuate within that envelope.

Most air handling systems have the capability to both humidify—to prevent electrostatic discharge—and dehumidify—to prevent condensation buildup in equipment. Both processes are extremely energy-intensive: dehumidification, for example, typically entails severely cooling the air, and then warming it again to bring it within acceptable operating range. Loosening the controls on humidity and dew point temperature reduces the amount of energy that must be spent on humidification and dehumidification.

[10] Nlyte Software. "Raising Co-location Temperature Safely." 2016. (PDF)

Humidity sensors, like temperature sensors, will naturally fall out of calibration over time, and should be periodically tested and recalibrated. Even if operating correctly, however, fluctuations in humidity levels within a data center may lead different CRAC/CRAH units in the same space to work against one another, with one device humidifying at the same time the other is dehumidifying. The best practice to avoid this energy waste is to operate cooling equipment based off of the same set of humidity sensors.

Data centers that allow lower humidity levels should make certain that their equipment is properly grounded to prevent increased levels of electrostatic discharge that could endanger employees or damage equipment.

Example: The Savannah River Site (SRS), a nuclear reservation in South Carolina, undertook efforts to reduce the environmental impact of its data centers. Sensors installed in a data center to measure operating conditions suggested that three CRAH units could be shut down, and the humidity controls on the other CRAH units could be disabled. Tests conducted after the change found that the humidity levels were still within safe and reliable ranges for the IT equipment; because the regional humidity ranged from 50% to 85% year-round, there was no need to re-enable the humidity controls. By reducing the power draw of humidifiers and chillers, the data center went from expending 402.9 kW on its cooling load to 229.8 kW, and PUE dropped from 4.00 to 2.77.[11]

Recommended Tools and Resources:

- Siemon's Standards Informant. "Data Center Temperature and Humidity Requirements: TIA-942-A."
- ASHRAE. "Data Center Power Equipment Thermal Guidelines and Best Practices." 2016. (PDF)
- ASHRAE. "Data Center Storage Equipment: Thermal Guidelines, Issues, and Best Practices." 2015. (PDF)
- Data Center Journal. "Humidity in the Data Center: Do We Still Need to Sweat It?" 2012.
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)

Optimize Airflow with Blanking Panels, Grommets, and Cable and Tile Management

Suboptimal airflow—whether due to obstructions, leaks, or an unintended mixing of hot and cold air—can force cooling equipment to work harder, cause hotspot buildups, and hurt overall

[11] U.S. Department of Energy. "Retro-Commissioning Increases Data Center Efficiency at Low-Cost: Success at Savannah River Site (SRS) at Low-Cost." 2010. (PDF)

facility performance. Commingling of air is particularly damaging: warmer air at intake may be unable to cool IT equipment to the level needed to guarantee reliability. Proper airflow configuration should separate cold intake air from hot IT exhaust air, and maximize airflow efficiency[12] by removing impediments to air passage through the computer room.

Blanking panels should be installed in every open space within a server rack to stop cold air vented by the cooling equipment from mixing with hot air produced by the IT equipment. Preventing such commingling will lower server inlet air temperatures and keep valuable cold air from leaking into hot aisles, which in turn will lighten the load put on cooling equipment.

Some raised-floor data centers use their sub-floor space as an air plenum. These underfloor plenums should be kept sealed as much as possible, and *brush seal grommets* should be installed where cables enter into the computer room from the plenum. Grommets allow cables to pass through while simultaneously preventing the leaking of cold air into the computer room before it reaches its targeted destination.

Jumbled, disorganized, "spaghetti" cabling and wires can clog air conduits, restricting desired airflow. Best cable management practices entail keeping cables and wires neatly organized and carefully labeled, and implementing a cable mining program to periodically remove old or broken cables. If the configuration of the data center allows it, deploying cables overhead and above the servers—as opposed to beneath a raised floor—can lead to easier cable management. If air is routed through a sub-floor plenum, positioning cables overhead also helps ensure better airflow (and saves on costs, since cables deployed in air plenums must be plenum rated.)

Perforated or vented tiles should be located exclusively in cold aisles, where cold intake air is expected to enter from an underfloor plenum. Optimal placement of these tiles is in front of server racks to allow cold air to vent directly onto the front of the IT equipment. Similarly, any above-floor vents for cold air supply ducts should be oriented in cold aisles.

⚠ Data center operators should be cognizant of possible changes to air pressure that result from rearranging tiles or removing obstacles in sub-floor plenums.

Example: Lawrence Berkeley National Laboratory retrofitted a data center to incorporate airflow best practices. LBL used computational fluid dynamics to create an airflow visualization model, which showed hotspots, leaks, obstructions, and other substandard airflow conditions within the facility. After the retrofit—which entailed such steps as adding blanking panels, patching leaks, and rearranging floor tile openings—substantial benefits were realized, including the elimination of most hot spots and a 21% increase in cooling capacity.[13]

---

[12] Airflow Efficiency is equal to the Total Fan Power (W) over Total Fan Airflow (cfm).
[13] Federal Energy Management Program. "Data Center Airflow Management Retrofit." 2010.

Recommended Tools and Resources:

- Lawrence Berkeley National Laboratory's Data Center Air Management Tool
- ASHRAE. "Data Center Power Equipment Thermal Guidelines and Best Practices." 2016. (PDF)
- Upsite Technologies and Tate. "6 Reasons to Specify Air-Sealing Grommets in the Design of a New Data Center." 2013.
- Schneider Electric. "Planning Effective Power and Data Cable Management in IT Racks." 2015. (PDF)
- Lawrence Berkeley National Laboratory. "Case Study: Opportunities to Improve Energy Efficiency in Three Federal Data Centers." 2014. (PDF)
- ENERGY STAR. "Top 12 Ways to Decrease the Energy Consumption of Your Data Center."
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)

Save Money with Free Cooling Using an Air-side or Water-side Economizer

Free cooling, a technique in which outside air or water is used to cool a building alongside of or in place of mechanical chiller systems, is a growing best practice among data centers looking to reduce costs and save energy. Traditional chillers account for approximately 30% of an average data center's electricity consumption; if servers can be cooled using freely available resources instead, data centers can decommission their energy-intensive chillers at least part of the time for considerable monetary savings. These savings opportunities are not limited to colder climates: free cooling can be used whenever the outside air or water temperature is lower than a data center's return air or water temperature, such as at night or during winter in warmer regions.

Free cooling can be as simple as opening a window and allowing a breeze to flow in—for practical and security reasons, of course, such an approach is not typically appropriate for Federal data centers. Instead, data centers wishing to capitalize on outdoor resources can turn to economizers. An economizer is a tool that is used to decrease energy consumption by mechanical cooling equipment, and can offer data center operators some control over the free cooling process.

An air-side economizer controls the flow of external air brought into a building. The allowable volume is based on real-time conditions reported by exterior sensors, which the economizer then responds to by opening or closing the entrance to an outdoor air duct. There are two types of air-side economizers: a direct air-side economizer brings outside air into the facility, filters and dehumidifies it as needed, and then vents it directly into the computer room. An indirect air-side economizer, in contrast, keeps the colder external air and hotter exhaust air physically separated at all times. This separation is often accomplished by containing the hot air in a pipe as it is chilled.

Data centers interested in adopting air-side economizers must consider local risk factors. Air pollutants like salt or smoke from industrial zones can damage IT equipment. Extreme humidity levels increase the risk of excessive electrostatic discharge or condensation buildup, and the humidity of the outside air is difficult to maintain in the acceptable range in many parts of the U.S. In certain environments, data centers may spend more energy on treating and filtering the air than they can save from decommissioning their chillers.

A water-side economizer works when water stored in an outdoor cooling tower is piped into the data center. The economizer, which is installed on the chillers, cools the water before it reaches the compressor, reducing or even eliminating the need for the compression stage in the cooling process. Like with air-side economizers, water-side economizers rely on sensors to determine whether the temperature of the exterior water is low enough to serve as an effective coolant.

The cost and energy saving opportunities from economization will be endangered if an economizer is not working correctly. The mechanical pieces of an economizer unit should be periodically examined to ensure they are functioning properly. The exterior air and water sensors, like a data center's internal sensors, should also be checked and recalibrated on a regular basis.

Example: During a retrofit of a data center at company headquarters, outdoor gear retailer REI installed an evaporative cooling system to make use of the low-humidity air of Washington State. The evaporative system pumped water into a rooftop cooling tower to be exposed to the lower-temperature outside air, and then returned the now-chilled water into the facility where it was used to cool the hot computer room air. By using evaporative cooling in place of a mechanical chiller over 99% of the time, REI was able to cut energy use by 93% at the facility, and saw a drop in PUE from 2.4 to 1.4.[14]

Recommended Tools and Resources:

- The Green Grid's Free Cooling Estimated Savings Calculator
- The Green Grid. "Updated Air-Side Free Cooling Maps: The Impact of ASHRAE 2011 Allowable Ranges." 2012. (PDF)
- ASHRAE. "2011 Gaseous and Particulate Contamination Guidelines for Data Centers." 2011. (PDF)
- ENERGY STAR. "Air-Side Economizer."
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

[14] CIO.com. "IT Leaders Pursue Data Center Innovation to Beat the Heat." 2014.

- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)
- Emerson Network Power. "Economizer Fundamentals: Smart Approaches to Energy-Efficient Free-Cooling for Data Centers." 2010. (PDF)
- Schneider Electric. "Choosing Between Direct and Indirect Air Economization for Data Centers." 2015. (PDF)
- U.S. National Institutes of Health. "Sustainable Data Center Design Guide." 2013.

Install Direct Liquid Cooling (e.g. Rear Door Cooling)

Traditionally in data centers, hot exhaust air from IT equipment is routed back across the computer room to CRACs or CRAHs, where it is mechanically chilled and re-released. This process requires a significant amount of energy, even in a facility with a hot aisle/cold aisle arrangement. Because liquid is far more efficient at transferring heat than air, more and more data centers are moving away from CRACs/CRAHs in favor of Direct Liquid Cooling (DLC).

A popular method of DLC is Rear Door Cooling (RDC), which takes advantage of the fact that a liquid is more efficient at cooling the closer it is to a heat source. In RDC, flexible hoses transport water or other liquid coolant to coils affixed to doors on the backs of server racks. Due to the front-to-back nature of air movement through the servers, the hot exhaust air is pushed into contact with the door, where it is cooled by the coils. Such a system means that hot air no longer has to be released into the room to return to the cooling equipment.

Fears about leaks stop some data centers from moving to DLC, despite improvements in the reliability of the technology. A common response to this concern is to install underfloor transportation of the coolant, to remove any risk of damage from a faulty hose.

Example: Online marketplace eBay teamed up with Dell and Intel to develop a liquid-cooling solution for its hyper-scale data centers tasked with handling over one billion transactions a day. The project, known as "Triton," resulted in a rack-based, direct contact liquid cooling product. Triton works by circulating building water through the server and out to an evaporative cooling tower, and has advanced mechanisms for identifying and responding to localized water leaks. Thanks to Triton, eBay was able to achieve a 70% increase in search requests per second. Tests found that Triton has a PUE ranging from 1.026 to 1.029, and uses 97% less power than a typical air-to-air heat exchange cooling solution.[15,16]

[15] Enterprise Tech. "eBay's Liquid Cooled Billion-Transactions-a-Day Data Center." June 29, 2016.
[16] Moor Insights and Strategy. "Dell's Liquid Cooling Innovation for Scale-Out Datacenter Environments." 2016. (PDF)

Recommended Tools and Resources:

- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- Pacific Gas and Electric Company. "Data Center Best Practices Guide." 2012. (PDF)

Modernize Lighting Systems

Lighting accounts for approximately 2% of an average data center's energy consumption, and retrofitting a lighting system can improve its efficiency by as much as 90%.[17] Data centers do not typically require full 24/7 internal lighting—as a result of increased automation and remote management capabilities, employees can often limit their physical presence in the data-center to an infrequent, as-needed basis. Automatic light sensors that turn on lights only when an area is occupied helps lower energy consumption and reduces heat generation.

Other best practices for data center lighting include switching to higher-efficiency bulbs, positioning lights above aisles instead of racks (to avoid unnecessary heating of IT equipment), and adopting a multi-level lighting protocol with varying levels of brightness. Currently, LED lights are the best choice for use in a data center: compared to a 60w traditional incandescent bulb, an LED bulb uses approximately 75-80% less energy, offers 35 to 50 times more hours of lighting, has a 75% lower energy cost annually, and emits far less heat.

Example: Evaluations were conducted of three anonymous federal government data centers by Lawrence Berkeley National Laboratory in an effort to uncover opportunities for efficiency improvements. LBL found that best practices for lighting were not being followed in the three federal data centers; it concluded that at an approximate cost of $2 to $6 per square foot of the data center, such measures as reducing the amount of lighting and installing occupancy sensors would pay for themselves in about 2.2 years.[18]

Recommended Tools and Resources:

- U.S. Federal Energy Management Program's Acquisition Guidance for Lighting Products
- ENERGY STAR's Light Bulb Calculator (EXCEL)
- Seimen's Standards Informant. "Data Center Lighting Requirements: TIA-942-A."

[17] Consulting Specifying Engineer. "Intelligent Lighting Improves Data Center Efficiency." 2012.
[18] Lawrence Berkeley National Laboratory. "Case Study: Opportunities to Improve Energy Efficiency in Three Federal Data Centers." 2014. (PDF)

- U.S. Department of Energy. "How Energy-Efficient Light Bulbs Compare With Traditional Incandescents."
- ENERGY STAR. "Lighting Technologies: A Guide to Energy-Efficient Illumination." (PDF)
- Data Center Journal. "Lighting as the Unsung Hero in Data Center Energy Efficiency." 2012.
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

Virtualization

Definition: Ratio of operating systems (OS) to physical servers.

DCOI FY 18 Target: ≥ 4

Right-Size After Virtualization

Because the PUE metric is a measure of the efficiency of a data center's physical infrastructure, facilities that are heavily virtualized may see their PUE increase. This perhaps-surprising outcome stems from the fact that post-virtualization, the amount of electricity consumed by IT equipment drops. By contrast, if power and cooling equipment are not adjusted, they will continue to expend the same amount of energy for this smaller IT load. Subsequently, a best practice after virtualization is to right-size a data center's physical infrastructure (unless there is an expectation for near-term future growth through the addition of more virtualized servers.)

Methods for achieving this right-sizing include installing Variable Frequency Drives on fans and pumps to enable energy use to match lower workloads; using a capacity management system to find specific areas of excess capacity; and switching to more energy efficient equipment. Servers can also be consolidated into a smaller number of racks, reducing the space in a computer room that must be serviced by power and cooling equipment.

If IT power consumption decreases by enough, a data center may be able to decommission one or more of its cooling and/or power units. Such a move leads to maximum cost and energy savings by removing "fixed losses" from unneeded equipment consuming power simply to stay on. Alternatively, a data center could spread the cost of its infrastructure over additional IT equipment by offering to host other agencies' co-located IT equipment.

Example: Researchers at Schneider Electric's Data Center Science Center conducted a study to determine the differences in PUE and cost savings if a hypothetical data center (1 MW, 70% loaded, average electricity bill of $1.4 million) did and did not right-size its physical infrastructure after virtualizing all of its servers. With a starting PUE of 2.00, virtualizing without any adjustments to cooling and power capacity led to a 17% decrease in electricity spending, but increased the PUE to 2.25. Rightsizing the physical infrastructure after the initial virtualization, by contrast, led to an additional 28% decrease in electricity costs, for a total drop of 40% from the original rate. PUE also fell to 1.63 as a result of the diminished amount of excess capacity.[19]

[19] Schneider Electric. "Virtualization and Cloud Computing: Optimized Power, Cooling, and Management Maximizes Benefits." 2012. (PDF)

Recommended Tools and Resources:

- Schneider Electric's Virtualization Energy Cost Calculator.
- Eaton. "Right-Sizing Your Power Infrastructure." 2009. (PDF)
- Schneider Electric. "Impact of Virtualization on Data Center Physical Infrastructure." 2010. (PDF)
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- U.S. National Institutes of Health. "Sustainable Data Center Design Guide." 2013.

Server Utilization & Automated Monitoring

Definition: Percent of time busy (measured as 1 – percent of time spent idle), measured directly by continuous, automated monitoring software, discounted by the fraction of data centers fully equipped with automated monitoring.
DCOI FY 18 Target: ≥ 65%

Turn Off Zombie Servers

Many data centers—intentionally or unintentionally— squander significant amounts of energy and money keeping unused servers powered on. Estimates show that approximately 15-30% of IT equipment in an average data center can be classified as “undead,” or “zombies,” meaning that a server is powered on and consuming energy—up to 50% the amount of a fully-working server—without doing any meaningful work.

Data centers can achieve higher server utilization ratios by turning off undead servers. In order to carry out this best practice, a data center needs a policy in place for tracking every server and application in the facility, and must conduct periodic audits to identify and decommission unused machines.

Data center operators are commonly cautious about retiring servers for fear of inadvertently interfering with a critical application. By keeping detailed records of what is on every server, operators can know exactly what will be impacted if a specific machine is powered down.

Example: Aware of the costs and energy waste of unknown “comatose” servers, banking company Barclays identified and decommissioned over 9,000 servers in a single year. Removal of the unneeded servers freed up rack space (enough to fill nearly 600 server racks) and over 20,000 network ports. In total, the lower power consumption and maintenance needs led to annual savings of $6 million.[20]

Recommended Tools and Resources:

- Uptime Institute’s Comatose Server Decommissioning Starter Kit.
- Uptime Institute’s Comatose Server Savings Calculator.
- EdTech. “How to Kill the Walking Dead Servers.” 2015.

[20] Data Center Dynamics. “Barclays Saves US$6m by Weeding Out Dead Servers.” 2014.

- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.
- ENERGY STAR. "Top 12 Ways to Decrease the Energy Consumption of Your Data Center."

Make Use of Server Power Management Technology

Data centers often run more servers than they need in order to be ready for peak demand periods. On average, approximately 80% of a server's capacity sits unused at any one time. To increase server utilization, a current best practice is to aim for power proportionality, in which the power used by the IT equipment matches the amount of work it is performing. Power proportionality can be accomplished through the use of server power management technology, which changes resources from "always on" to "always available." Executive Order 13693 requires that agencies enable power management features on all non-exempt computers and displays.

Unlike older technology, newer servers are built to withstand thousands of on/off cycles. Provided that the servers in a data center are of sufficiently recent build dates, a data center can use power management technology without worrying about compromising the integrity of its equipment.

Most modern servers come equipped with built-in power management capabilities in the form of a sleep mode, which reduces power to the equipment when it is not being actively used. The Green Grid estimates that using server reduced-power modes can save a data center 20% in operating costs. Data center operators should enable sleep mode on servers which store data that does not need to be always-accessible. Settings can be adjusted as needed to restrict sleep mode to certain times of the day or the week, depending on typical patterns of demand.

Data centers can also take advantage of other power management techniques built-in to their servers. Multicore processor servers, for instance, can be set to shut down one or more of their internal cores when they are not needed.

Other popular forms of power management technology includes Massive Array of Idle Disks (MAID) systems and throttle down drives. MAID reduces power consumption by spinning disk drives only as the workload demands. Throttle down drives reduce energy consumption by reducing the level of power being directed to a server that is running below its designated utilization point.

Example: Social media giant Facebook developed an in-house program for data center-wide power management. The program, called “Dynamo,” was installed on every server and used to report power usage data to a set of controllers, which then capped or uncapped an individual server’s power consumption as needed. Over a period of three years, Dynamo enabled Facebook to accommodate 8% more servers in a single data center without consuming any additional power. On average, Dynamo’s monitoring abilities also prevented around three power outages a month.[21]

Recommended Tools and Resources:

- ENERGY STAR’s PC Power Management Savings Estimator (EXCEL)
- ENERGY STAR. “Not a Techie? Tips for Getting Your CPM Project Off the Ground.”
- Lawrence Berkeley National Laboratory. “Data Center Master List of Energy Efficiency Actions.” 2016.
- Pacific Gas and Electric Company. “Data Center Best Practices Guide.” 2012. (PDF)

[21] Facebook, Inc. “Dynamo: Facebook’s Data Center-Wide Power Management System.” 2016.

Facility Utilization

Definition: Portion of total gross floor area in tiered data centers that is actively utilized for racks that contain IT equipment.
DCOI FY 18 Target: ≥ 80%

Implement Rack and Row Based Cooling Solutions

Data center designers regularly opt to build out excess cooling capacity to be prepared for expected eventual growth. Excess capacity shortens the time it takes to scale up operations, but also hurts a data center's facility utilization ratio as cooling-ready space sits unused by servers. A recent trend among new data centers is to use a modular design that allows for expanding infrastructure capacity on an as-needed basis. Pre-existing data centers without the flexibility of the design stage can improve their performance in this metric by adopting rack and row-based cooling techniques, which reduce stranded capacity by minimizing the area supported by infrastructure.

Options for targeted cooling technologies include row coolers, rack coolers, and rear-door coolers:

- A row cooler operates by capturing the hot air exhausted from a row of servers, cooling it, and then returning the lower-temperature air to the front of the servers, where it can be re-used as intake air. Row-based CRACs are 95% efficient, compared to the typical 80% efficiency of a room-based unit.

- A rack cooler uses a heat exchanger to cool the hot exhaust air in close proximity to the source. A rack cooler removes the need for hot exhaust air to move through the computer room, and so the air is able to recirculate over and over again within the confines of the rack cooler enclosure device.

- A rear-door cooler removes any need to transport air beyond the perimeter of a single rack. In a rear-door cooler system, a door attached to the back of a rack is equipped with coils containing some form of liquid coolant runs. The hot exhaust air is cooled when it exits the rear of the server and makes contact with the coils. As with the rack cooler, this cooled air can be recirculated to the front of the server, where it again becomes intake air.

Example: PC manufacturer Lenovo purchased two preexisting data centers in California and North Carolina. During the process of retrofitting the facilities to meet Lenovo's needs, the company decided to purchase and install a higher-efficiency cooling solution comprised of modular cooling blocks that use approximately 80% less water than a standard chiller system. Because of the modular nature of the product, Lenovo was able to avoid unnecessarily over-provisioning the facilities, saving both money and energy. After the installation, both the data centers were using over 80% less power for their cooling systems than they were before the retrofit.[22]

Recommended Tools and Resources:

- The Green Grid. "Qualitative Analysis of Cooling Architectures for Data Centers." 2011.
- Schneider Electric. "Power and Cooling Capacity Management for Data Centers." 2012.
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

[22] Mission Critical. "Modular Cooling Blocks Revolutionize Data Center Energy Usage." 2016.

ADDITIONAL AREAS FOR IMPROVEMENT

Management/Operations

Properly Train Employees

Due to its complex nature, a data center cannot operate effectively or efficiently if employees are not adequately trained. Training must be a continual process that begins from the day that an employee first comes aboard, and should ensure that every individual is capable of working safely and productively in the data center. All training curriculums should include some form of assessment (for example, written or practical exams) to verify that workers have processed the material and will be able to put it into practice on the job.

All employees, regardless of their job description, should receive training on:

- Safety and emergency procedures
- Facility operating procedures
- Industry standards and best practices
- Compliance with relevant regulatory laws and government codes

Based on their duty requirements and level of expertise, employees should also receive ongoing technical training. For example, the purchase of new equipment might include a vendor-led training session and incorporate time for employees to familiarize themselves with the new technology.

Employees should be encouraged to seek external certifications to supplement in-house training and stay abreast of industry trends. Per the Executive Order 13693 Implementing Instructions, all core Federal data centers are advised to have at least one certified Data Center Energy Practitioner assigned to manage the data center's performance and continued optimization.

> **Example:** Due to the nature of their work, employees in ChARM HER's (an electronic health records service) data center must not only know how to handle the technical and sometimes hazardous responsibilities of working in a data center, but must also be fully versed in patient privacy laws. Employees are given the typical security and practical training upon being hired, and are only permitted to enter the computer room after a probationary period. In addition, however, they receive instruction in the standards of the Health Insurance Portability and Accountability Act, commonly known as HIPAA. Such training means that ChARM's healthcare customers can turn to the cloud without fear of violating patient privacy.[23]

[23] ChARM EHR. "HIPAA Privacy Policy."

Recommended Tools and Resources:

- Data Center Energy Practitioner (DCEP) Training
- ABB. "Ten Ways to Ensure the Safety of Data Center Employees." 2013. (PDF)
- Fortrust. "A Data Center Operations Guide for Maximum Reliability." 2013. (PDF)

Benchmark Performance

Benchmarking a data center allows data center operators, leadership, and other relevant stakeholders to track a facility's performance over time, compare it to other data centers, and find potential areas for improvement. Possible metrics to benchmark beyond those given in M-16-19 include Rack Cooling Index, Return Temperature Index, Energy Reuse Effectiveness, and Airflow Efficiency.

In order to benchmark performance, a data center must have data collection capabilities. Sufficiently sophisticated energy meters and environment sensors can be set to automatically collect data at user-defined intervals. Data centers should determine appropriate intervals based on their needs, bandwidth for reporting and storing data, and manpower for interpreting it. Data centers will want to avoid expending more effort and cost to collect data than they could gain from cost saving opportunities identified by analyzing it.

Data center operators may find it appropriate to collect some types of data—such as temperature and humidity levels—in hourly or even shorter intervals, in order to track real time conditions throughout the day. Weekly or monthly intervals may be more cost effective for areas that do not fluctuate widely over the course of a day, such as power usage.

Benchmarking will be of little value if it is not done against a truly comparable facility. The data centers chosen to benchmark against should be peers in terms of such factors as size, technology, maturity, and units of measurement.

Benchmarking will produce flawed conclusions if meters, sensors, and other monitoring equipment are not taking correct readings. Periodic checks and recalibration of this equipment should be conducted to ensure the highest possible level of accuracy.

Because conditions in a data center are not static, benchmarking should be based on multiple measurements over a period of time, and not a single snapshot.

Example: Salt River Project (SRP), one of the largest public power utilities in the United States, spent several years making infrastructure changes to an Arizona data center in an effort to reduce its carbon footprint. Eventually, SRP decided to benchmark its energy usage to compare itself to other regional data centers, and to track the effectiveness of the efficiency improvements it had already made. An audit determined that the facility's post-improvement PUE was 2.6, and identified further cost and energy savings measures that would lower it to 1.79. The implementation of the suggested changes, such as installing barriers between hot and cold aisles and raising the facility's temperature, enabled the data center to save $53,000 in energy costs annually, and lowered its carbon emissions by 289 tons per year.[24]

Recommended Tools and Resources:

- ENERGY STAR's Portfolio Manager
- The Green Grid's Data Center Maturity Model and Data Center Maturity Model Assessment Tool
- U.S. Department of Energy. "Federal Building Energy Use Benchmarking Guidance." 2014. (PDF)

[24] HP. "Nation's Third-Largest Public Utility Audits and Enhances Data Center Energy Efficiency." 2009.

Sustainability

Purchase Green IT Equipment

The EPA-run ENERGY STAR program helps promote environmental-friendly, energy-efficient products and buildings, in part by giving qualified products an ENERGY STAR label. The Department of Energy's Federal Energy Management Program (FEMP) designation is given to products that rank in the top 25% for energy efficiency among their peers.

Under Federal government regulations, Federal buildings—including data centers—are required to give preference to products that are ENERGY STAR-qualified and/or FEMP-designated. The purchase of such "green" technology helps shrink a data center's carbon footprint, but can also help it achieve cost savings: a green server approved by ENERGY STAR, for instance, uses 30% less energy than an average traditional server.

Federal regulations do not require that facilities immediately replace all of their equipment at once. A data center hoping to achieve both cost and energy savings can make it a policy to purchase green technology over time as periodic equipment refreshes are conducted.

Example: As part of a drive to power 100% of its data centers with renewable energy, Apple began experimenting with various green technologies at its data centers, including biogas fuel cells, solar panels, and wind and water turbines. Apple's North Carolina data center, for example, receives between 60-100% of its energy from a mix of solar energy and biogas fuel, and installed a non-water-intensive cooling system that lowered its water consumption by 20%. In total, Apple data centers avoided consuming nearly 190 thousand metric tons of CO_2e emissions in a single year with the help of green IT.[25]

Recommended Tools and Resources:

- ENERGY STAR "Purchase Energy-Saving Products"
- U.S. General Services Administration. Office of Federal High-Performance Green Buildings.
- U.S. General Services Administration's Sustainable Facilities Tool
- Department of Energy. "Search for Efficient Technologies and Products for Federal Facilities."
- Department of Energy's Energy Cost Savings Calculator for Air-Cooled Electric Chillers
- Pacific Northwest National Laboratory's Rooftop Unit Comparison Calculator
- National Renewable Energy Laboratory's Dynamic Maps, GIS Data, and Analysis Tools
- Lawrence Berkeley National Laboratory. "Data Center Master List of Energy Efficiency Actions." 2016.

[25] Apple. "Environmental Responsibility Report." 2016.

Security

Practice a Defense-in-Depth Approach to Security

Defense-in-depth is a security tactic commonly used by top-performing private sector data centers that house high-profile or sensitive data. It moves beyond the traditional “focus on the perimeter” method of defense by establishing overlapping, multidimensional layers of protection. If designed effectively, each successive level of security will be stronger than the one surrounding it, and able to withstand attacks where earlier defenses failed.

Defense-in-depth is not a one-size-fits-all strategy. Data centers should craft their approach based on a thorough assessment of their needs and capabilities. It would be excessive, say, if a mid-sized, limited-budget data center housing low-sensitivity information procured the same laser beam intrusion detection system that internet giant Google deploys in its computer rooms.

Physical Security

The outermost level of physical security prevents unauthorized individuals from gaining access to the building housing a data center. Typical security measures include alarms, motion-activated security lighting, video surveillance, and perimeter fences equipped with gates, guard shacks, and vehicle entry barriers.

The second zone of defense encompasses the interior of the building, and is responsible for preventing attacks from within. This level may incorporate multiple points of ID verification; visitor management and monitoring procedures; metal detectors; locked doors; and both still cameras as well as video surveillance (especially around all computer room entryways.)

Robust computer room protection modifies many of the tools used to guard the building at large to meet the stricter security requirements of the data center floor. Employees who can freely move around the rest of the facility, for example, should not automatically be granted entry to the computer room. Stronger access requirements, such as biometric identifiers like fingerprints or retina patterns, can exclude all but a select group of individuals from entering the space.

The last line of defense for protecting a data center’s physical assets comes at the cabinet level. An obvious but important best practice is to have a reliable cabinet-locking instrument in place: if an intruder makes it past all other levels of defense, a lock may make the difference in whether or not facility operators have time to respond to a breach.

Virtual Security

A good virtual defense-in-depth strategy will consider the increased interconnectivity data centers have with the outside world as they move assets to the cloud. Similar to protecting physical assets, a virtual security strategy cannot simply defend against external actors. Recent high-profile data breaches highlight the importance of restricting employee access to computer networks in order to reduce the likelihood and scope of intrusions: The IBM Cyber Security Index 2016, for instance, found that 60% of attacks on its clients were carried out by individuals with insider access (and 15.5% of those insider breaches were unintentional). A best practice among top data centers is to define least-privilege access rights, whereby users are initially equipped with a minimum set of privileges, and only granted higher level privileges after a thorough review of their duty needs.

Other essential best practices include keeping all operating systems and applications current with the latest versions and releases and a robust patch management system. Antivirus and intrusion detection software must be up-to-date, and continuously monitoring networks for signs of abnormalities. There should be strong authentication requirements for access to a network (such as using Personal Identity Verification cards or other multi-factor authentication tools.) To minimize vulnerabilities, servers and workstations should be hardened by removing applications unnecessary for the machine and its user to perform their key functions.

Data center operators should stay abreast of alerts and guidance issued by the United States Computer Emergency Readiness Team (a component of the Department of Homeland Security), the FBI, and other government offices that monitor cyber security developments. These entities can serve as partners in the effort to protect a Federal data center's assets, and can provide valuable resources, such as DHS's EINSTEIN firewall. A total awareness program allows the data center to react to what is happening to other Government systems and more broadly in the world, and to have prepared tools to increase data security during critical times.

Example: Despite the protection of being located 220 feet underground, Iron Mountain's Tier III data center in Western Pennsylvania still turns to more traditional methods of security to protect its clients' data. The data center, encompassing over 200 subterranean acres, is protected from the outside by fences, vehicle crash barriers, and a three ton steel gate. Within the facility, armed guards patrol around the clock, closed circuit cameras monitor the interior 24/7, and biometric controls and mantraps limit access to the computer room floors.[26]

Recommended Tools and Resources:

- National Institute of Standards and Technology. "NIST SP 800-82, Guide to Industrial Control Systems." 2011. (PDF)

[26] Iron Mountain. Pittsburgh Data Center.

- U.S. Department of Homeland Security. "Federal Information Security Modernization Act (FISMA)."
- U.S. Department of Homeland Security. "Best Practices for Continuity of Operations (Handling Destructive Malware.)" 2015.
- U.S. Department of Homeland Security's EINSTEIN System.
- U.S. Department of Homeland Security's Continuous Diagnostics and Mitigation program
- U.S. Department of Homeland Security. "Improving Industrial Control Systems Cybersecurity with Defense-in-Depth Strategies." 2009. (PDF)
- U.S. Federal Trade Commission. "Start with Security: A Guide for Business." 2015.
- Anixter. "The Four Layers of Data Center Physical Security for a Comprehensive and Integrated Approach." 2012. (PDF)
- IBM. "2016 Cyber Security Intelligence Index." 2016.

Prepare for Emergencies with Disaster Recovery and Business Continuity Plans

A crucial requisite for being a well-managed data center is having plans in place for responding to various potential disasters. If an emergency strikes, a data center will ideally be prepared to make a rational decision based on predetermined criteria about whether to continue operations as normal, shift to a disaster recovery site, or commence with an orderly shutdown. A good plan should be formed during a non-crisis period with input from individuals with different areas of expertise, and should document the steps that will be taken as soon as an emergency begins or appears imminent. Good planning includes execution of regular system backups and recovery site storage of the backups. The plan(s), along with related emergency equipment like backup generators, should be tested at least once annually as part of an emergency-response drill.

Initial site location is key and establishes the floor level of risk. Awareness is crucial in being prepared for a disaster. A data center should evaluate its level of risk from natural disasters—such as earthquakes, fault lines, tornados, and hurricanes, along with the secondary risks of fires, flooding, and power outages—as well as its risk from man-made threats, such as hacking or terrorism. Data centers should also be cognizant of vulnerabilities created by their immediate surroundings, such as nearby nuclear plants, prisons, or airports. The creation of response plans to these threats can be prioritized based on the likelihood of each.

Data centers should continuously monitor local conditions for situations that could disrupt normal operations. The more warning a data center has about an impending event, the better it will be able to respond. A data center in the path of a blizzard or hurricane, for instance, can take the time to review emergency procedures with employees, make provisions for housing and feeding workers that will stay on site, and shore up any structural vulnerabilities (such as surrounding the facility with sandbags to prevent flooding).

Continuity of Operations Planning (COOP) and frequent practicing can help make data center operations resilient. Communication between management and employees is critical before and during an emergency. Staff can provide valuable input in crafting an emergency response strategy or alerting management to local hazards. Employees should be kept apprised of their roles in disaster recovery plans, and there must be clear lines of communication in place for workers to obtain updates and direction during a crisis. Frequent practice is key to success.

A data center can lower its chances of needing to use its disaster recovery plan in the first place by strengthening its business continuity ability. Automation and remote management capabilities may allow critical work to be accomplished even if workers cannot make it to the data center's physical site. Continuous data protection and electronically backing up information to an off-site location will enable the data center to move its operations to a remote facility, if need be, or recover faster in case of the total destruction of the data center. Many critical systems should be designed to operate in a mirror image, multi-location environment so that an event or attack on one center or region of the country does not affect the systems that seamlessly maintain operations at other locations.

Example: The National Finance Center's (NFC) response to Hurricane Katrina is applauded as a shining example of successful Federal disaster preparedness. NFC is a Shared Service Provider under the U.S. Department of Agriculture; in 2005, it was responsible for payrolling nearly 600,000 Federal employees. Located in New Orleans, the NFC was directly in the path of what turned out to be the costliest natural disaster in American history. Based on lessons learned from past storms, NFC had detailed disaster recovery and business continuity plans in place, and made it a habit of running frequent emergency response drills. On Friday, August 26th, 2005, the day after Katrina made landfall in Florida, NFC's disaster recovery plan was activated. Over the weekend, staff evacuated the magnetic data tapes to a recovery site in Philadelphia and shut down the data center. By that Monday, NFC was able to begin limited functions from multiple remote locations around the country. The NFC data center eventually moved to a new permanent facility in Denver, but throughout the entire interim period succeeded in never missing a payroll.[27]

Recommended Tools and Resources:

- SearchDisasterRecovery's Data Center Disaster Recovery Plan Template
- Uptime Institute. "Lessons Learned from Superstorm Sandy." 2013.

[27] USDA. "The National Finance Center and Hurricane Katrina." 2006. (PDF)

Create a Security Awareness Curriculum for Employees

A simple way to lessen a data center's vulnerability to man-made threats is to develop a security awareness program. A strong program should educate employees on a variety of security topics, including possible targeted threats (for example, social engineering attempts and phishing scams) and hazards to the facility (such as piggybacking through access points and non-accompanied guests). Employees should know how to protect confidential data, and what sort of information is inappropriate to be disclosed outside of the data center. Employees need to have the capacity to identify threats, and be knowledgeable on the appropriate process for reporting them.

Threat awareness instruction should begin at the employee onboarding process and be periodically reviewed throughout an employee's entire time with the data center.

Example: Cognizant of the massive amounts of private information passing through the hands of its employees, Google works to ensure that all of its workers are trained in security best practices. The on-boarding process of every employee includes a session on security training (including a review of the section of the Code of Conduct highlighting Google's security policy), and workers are expected to take follow-on trainings during their entire time with the company. Employees also receive job-specific security trainings on an as-needed basis, and have the opportunity to attend Google-sponsored lectures and conferences focusing on security.[28]

Recommended Tools and Resources:

- National Cyber Security Alliance. "Stay Safe Online: Train Your Employees."
- Payment Card Industry Security Standards Council. "Best Practices for Implementing a Security Awareness Program." 2014. (PDF)

[28] Google. "Google for Work Security and Compliance Whitepaper."

Reliability/Uptime

Increase Automation Capabilities

A best practice common among data centers that boast the highest percentages of annual uptime hours is to automate processes related to IT workloads, networks, servers, infrastructure, and more. While not all data center management tasks can be automated, doing so where feasible can improve facility performance and allow a data center to shift resources to other needs.

Automation can help a data center avoid unplanned downtime by enabling better awareness of and faster response times to problems. For example, software can continuously scan a data center's systems for abnormalities and, upon detecting an issue, commence with an automatic fix, or—if the problem requires human intervention—alert operators to the situation. Relying on manual reviews of system logs and human-instigated repairs may mean that a data center cannot resolve a problem before it causes an outage.

Industry studies estimate that human error is responsible for anywhere from 20% to 70% of data center outages. By using technology to automate processes that were previously performed manually, data centers can substantially lower their vulnerability to outages caused by avoidable mistakes and focus more of their efforts on other leading causes of downtime—like cyber-attacks and equipment failures—that are harder to control.

While automation technology can require considerable initial costs, the expenses associated with unintended outages are also substantial and growing ever larger; a study conducted by the Ponemon Institute found that from 2010 to 2016, the average cost of a data center outage rose by 38%, to $740,357.

Example: Image publishing company Shutterfly relies on the internet to connect to its customers. After realizing that its growth potential was limited by outdated network switches in its data center, Shutterfly replaced the legacy devices with modern network switches with built-in automation capabilities. The new switches enabled network provisioning tasks to be accomplished automatically, saving the company 3-4 hours of work and $400 per job. The switches allowed the Shutterfly network to handle increased levels of traffic—up to 40% during the holidays—with no associated downtime, largely because the automated switches dramatically lowered the time it took to detect and resolve an error.[29]

[29] Cisco. "Digital Photos Come to Life With Software-Defined Networking." 2016. (PDF)

Recommended Tools and Resources:

- Ponemon Institute. "Cost of Data Center Outages." 2016. (PDF)
- Fortrust. "A Data Center Operations Guide for Maximum Reliability." 2013. (PDF)
- Emerson Network Power. "Addressing the Leading Root Causes of Downtime." 2010. (PDF)

Change Management

Collaborate with Stakeholders to Design and Implement Change

Most data center operators would not consider beginning a major facility change without having a complete inventory of IT systems and equipment. Equally important and frequently overlooked, however, is an inventory of relevant stakeholders. For a Federal data center, stakeholders can include agency leadership, facility employees, the general public, and customers (in the case of inter-agency shared services providers). A stakeholder analysis should identify who is likely to be impacted by a change and how they will be impacted, as well as ascertain which individuals may be needed to help carry it out.

After identification, communication with the impacted stakeholders is crucial not only to ensure that they are not surprised by the change, but also to give them the opportunity to voice their own interests and priorities. Some of these interests may need to be harmonized with the overall change vision, particularly if an individual has the potential to help or hinder its implementation. Good communication also enables the change leader to identify knowledge shortcomings (for instance, a lack of employee familiarity with a proposed new technology acquisition). In the long run, benefits will accrue faster after implementation if all necessary individuals are on board with the plan and hold the essential skills for carrying it out.

Change leaders need to be prepared to face resistance to their plan. Employees may be wary or even hostile to proposed changes—and rightfully so, since some necessary changes, such as increased virtualization, may render positions redundant. Even individuals whose jobs are not at stake may be skeptical of dramatic modifications to facility procedures: some newer best practices (for example, higher server intake temperatures and using liquid to cool IT equipment) defy longstanding beliefs about how to best run a data center. Communication again plays a role in addressing this problem: impacted individuals will want the opportunity to express their concerns, and change leaders need to be able to explain the rationale behind the plan.

Data center managers should be realistic about the results they hope to achieve: stakeholders will not be pleased if promised outcomes never materialize. When making their case to all interested parties, the change owners must outline the milestones they intend to meet, as well as the metrics they will use to define success. During and after the change process, stakeholders should be kept apprised about whether these milestones and target metrics are being met. Frequent conversations and open lines of communications between the individuals implementing the change and the stakeholders who will be impacted will continue to be crucial.

Example: The discovery of mold in a Florida Agency for State Technology (AST) data center housing over 2000 servers for 12 state agencies meant that a planned 18 month move to a new facility had to be accomplished in four months. The move was undertaken successfully at a breakneck pace, which AST largely credits to "excessive communication" with stakeholders: customer agencies, as well as both facilities and IT employees, were kept apprised of the situation through frequent meetings and status updates. While AST drove the operation, the customers were treated as partners, and their input and assistance was sought throughout the entire process. Thanks to cooperation from all stakeholders, the move to the new, mold-free facility was completed with 20 days to spare.[30]

Recommended Tools and Resources:

- U.S. Agency for International Development. "Change Management Best Practices Guide." 2015. (PDF)
- BMC. "ITIL Change Management 101."

[30] GovTech.com. "Florida Buttons Up Major Data Center Move in a Mere 100 Days." June 30, 2016.

ADDITIONAL TOOLS AND RESOURCES

The White House. "M-16-19: Data Center Optimization Initiative (DCOI)." 2016. (PDF)

The White House. "Executive Order 13693: Planning for Federal Sustainability in the Next Decade." 2015.

The Federal Energy Management Program's Federal Laws & Requirements Search

Anxiter. "Data Center Infrastructure Resource Guide." 2012. (PDF)

Schneider Electric. "Guide for Reducing Data Center Physical Infrastructure Energy Consumption in Federal Data Centers." 2016.

Schneider Electric's White Paper Library.

Open Compute Project. "Data Centers."

Lawrence Berkeley National Laboratory. "United States Data Center Energy Usage Report." 2016.