

# A Framework for Auditing Data Center Energy Usage and Mitigating Environmental Footprint

Justin A. Gould  
Purdue University  
West Lafayette, Indiana  
gould29@purdue.edu

## ABSTRACT

Data Science has been a growing field since the 2010s. Between self-driving cars, movie and music recommendations, routing algorithms, and much more, both private companies and universities are making major investments in the field. As the field continues to mature, and we collect more data, the demand to store and analyze them will continue to increase. This increase in data availability and demand for analytics will put a strain on data centers and compute clusters—with implications for both energy costs and emissions. As the world battles a climate crisis, it is prudent for organizations with data centers to have a framework for combatting increasing energy costs and emissions to meet demand for analytics work.

In this paper, I present a generalized framework for organizations to audit data centers energy efficiency to understand the resources required to operate a given data center and effective steps organizations can take to improve data center efficiency and lower the environmental impact.

## KEYWORDS

Energy Efficiency, Energy Audit, Data Center, High-Performance Computing, Environment

## 1 INTRODUCTION

With the increased demand for access to computational resources, the amount of energy required to operate data centers is expected to continue to grow. In 2020, global data center usage was  $> 200TWh$  [1]. According to the Central Intelligence Agency’s 2020 Fact Book, in 2016, the electrical consumption of Poland, Sweden, and Norway were  $142TWh$ ,  $127TWh$ , and  $126TWh$ , respectively [2]. By 2030, global data center electricity usage is predicted to exceed 2 trillion  $TWh$  [1]. In the midst of a global climate crisis, data centers’ emissions and contributions are non-trivial and should be taken seriously to mitigate. The need for a generalized framework for organizations to audit data centers energy efficiency is great.

Data center emissions has implications in the fields of Artificial Intelligence and Machine Learning, too—as training models, and running inference in a production environment, requires access to compute resources, such as graphic processing units (GPU). Open AI’s Tom B. Brown et al.’s work on GPT-3 made reference to energy usage and implications of training models; however, the paper did not offer solutions to this growing problem [3]. The authors reported estimates of the compute resources required to train 17 recent state-of-the-art NLP models. The numbers provided are not accessible to non-subject matter experts and difficult to visualize. I took these estimates and converted the power required to the number of average U.S. homes that can be powered for one

year, depicted in Table 1. For the purposes of this analysis, I assume  $16.876GFLOP/watt$ , per a conservative cursory search of the Top500 List [4] and a similar analysis only for GPT-3 performed by Matthew Burruss [5], and average U.S. home electrical consumption as  $900kWh$  per month, per the U.S. Energy Information Administration [6].

Given the recent popularity of Data Science, and the increased demand for computational resources, previous work has been completed in attempting to create a set of guidelines and standards for organizations to follow to ensure data center efficiency. For example, the Lawrence Berkeley National Laboratory’s Center of Expertise for Energy Efficiency in Data Centers updated its *Data Center Master List of Energy Efficiency Measures* in September 2020 [7]. This document is a collection of best practices for energy efficient data center operations, split into 7 sections: *Data Center Energy Efficiency Management*, *IT Power Distribution Chain*, *IT Equipment*, *Lighting*, *Air Management*, *Cooling the Data Center Space*, and *Central Cooling Plant*. While the report has a large quantity of best practices and action items, many of the recommendations provided are resource-intensive—which may put smaller organizations at a disadvantage—and may not provide enough detail to be clearly understood and actionable. The Master List’s *EEM 1-2.7: Perform an Energy Audit*, for example, does not provide enough detail to carry out an audit. The Master List does, however, link the *DCEE Toolkit’s Energy Assessment Process Manual* [8]. The process manual falls victim to the similar shortfall of the Master List, whereby, quantity of content overshadows specificity and accessibility.

Another example of recent work is the U.S. Department of Energy’s *Best Practices Guide for Energy-Efficient Data Center Design* [9]. This guide is similar to the Master List, as it walks through best practices for various aspects of the data center, such as air management, cooling systems, environmental conditions, and IT systems. While the design guide offers more specificity than the Master List, it is not only nearly a decade old, but lacks the diversity and quantity of best practices provided by the Master List. Most apparent is the absence of guidelines to execute an energy efficiency audit. Furthermore, neither guide provides insight into organizational challenges associated with embarking on an audit, such as: creating an audit team, setting goals and benchmarks, and communication/goal accountability. I aim to address these important aspects of the audit process, and close these gaps, in section 3.

My primary goal for this paper is to leverage, enhance, and complement previous work in this space to provide a specific framework for an energy audit of a data center—with clear suggestions for organizations of varying size and resource availability. This is

**Table 1: Electrical consumption to train state-of-the-art natural language processing models, shown as powering average U.S. homes for 1 year**

Model Name	GFLOP	Consumption	Number of Homes
T5-Small	$1.8e^{11}$	10,666,034.61 kWh	987.60
T5-Base	$6.6e^{11}$	39,108,793.55 kWh	3,621.18
T5-Large	$2.31e^{12}$	136,880,777.40 kWh	12,674.15
T5-3B	$9e^{12}$	533,301,730.30 kWh	49,379.76
T5-11B	$3.3e^{13}$	1,955,439,678 kWh	181,059.23
BERT-Base	$1.64e^{11}$	9,717,942.64 kWh	899.81
BERT-Large	$5.33e^{11}$	31,583,313.58 kWh	2,924.38
RoBERTa-Base	$1.5e^{12}$	88,883,621.71 kWh	8,229.964
RoBERTa-Large	$4.26e^{12}$	252,429,485.70 kWh	23,373.10
GPT-3 Small	$2.25e^{11}$	13,332,543.26 kWh	1,234.49
GPT-3 Medium	$6.41e^{11}$	37,982,934.34 kWh	3,516.94
GPT-3 Large	$1.37e^{12}$	81,180,374.50 kWh	7,516.70
GPT-3 XL	$2.38e^{12}$	141,028,679.80 kWh	13,058.21
GPT-3 2.7B	$4.77e^{12}$	282,649,917 kWh	26,171.29
GPT-3 6.7B	$1.2e^{13}$	711,068,973.70 kWh	65,839.72
GPT-3 13B	$2.31e^{13}$	1,368,807,774 kWh	126,741.46
GPT-3 175B	$3.14e^{14}$	18,606,304,812 kWh	1,722,806

accomplished via the “full” and “lite” audit frameworks identified in this work.

The “full” audit is geared toward large, mature organizations with the resources required to not only perform deep analyses of many inner-workings of a data center, but the capital needed to address shortcomings identified in the audit.

For smaller organizations with fewer resources or expertise in high-performance computing, I offer a “lite” audit to focus effort on more attainable, less resource-intensive aspects of data center energy efficiency—as well as provide the roadmap to more mature energy assessment guidance to anticipate for future growth of the organization.

## 2 ENERGY AUDIT

This section contains the guidelines which comprise both the “full” and “lite” data center energy audits. The guidelines are binned into six categories, following the naming conventions established in the Master List [7]: *Cooling Air and Air Management*, *Environmental Conditions*, *Global*, *IT Equipment*, *IT Power Distribution Chain*, and *Lighting*.

### 2.1 "Full" vs. "Lite" Audit Frameworks

The appendix offers two tables, depicting the components which make up the “full” and “lite” audits. To reiterate: the “full” audit is intended for large organizations with the resources required to undertake a laborious, and potentially expensive, audit. The “lite” audit framework is meant for smaller organizations with fewer resources, to serve as an accessible option to have the ability to gain an understanding of the state of data center energy efficiency. The action items of the “lite” audit framework generally require fewer resources and expertise than large-scale equipment replacement specified in the “full” audit. To conduct a “lite” audit, only complete items listed in Table 4. To conduct a “full” audit, an organization

must complete items all “lite” (Table 4) and “full” (Table 5) options, depicted in Table 6.

### 2.2 Audit Item Structure

All entries in the proceeding subsections will use the following structure:

- **Audit Level:** At what level (e.g., individual rack, equipment room, entire data center, etc.) shall the analysis cover
- **Item Description:** Description of what the audit item is and how to perform it.
- **Metric:** How to measure the audit item.
- **Desired Output(s) & Goal(s):**
  - (1) Takeaway(s) from the audit item.
- **Action(s) to Take:**
  - (1) What action(s) should follow completion of this audit item (i.e., mitigation strategies to improve efficiency).
- **Audit Type:** “Full” or “Lite” item.

### 2.3 Cooling Air and Air Management

Items related to efficiently managing airflow and cool air within the data center.

#### 2.3.1 Return Temperature Index (RTI)

The Return Temperature Index (RTI) [10] is a key indicator explaining how effectively of the equipment room air-management system is operating. The RTI reflects either net by-pass air ( $RTI < 100\%$ ) or net recirculation air ( $RTI > 100\%$ ).

- **Audit Level:** Perform at the equipment room or entire data center level.
- **Item Description:** Capture the RTI for either an equipment room in a multi-room data center, or roll up to the entire data center level.
- **Metric:** Percentage: (measure of net by-pass air ( $RTI < 100\%$ ) or net recirculation air ( $RTI > 100\%$ ))

- **Desired Output(s) & Goal(s):**
  - (1) Identify which rooms or data centers do not have sufficient airflow management (i.e.,  $RTI \neq 100\%$ ).
- **Action(s) to Take:**
  - (1) Adjust room air management to ensure optimal  $RTI$  (100%).
- **Audit Type:** Full.

### 2.3.2 Alternating Hot and Cold Rack Aisles

. Aisles of server racks should alternate temperature between hot and cold. This can be achieved by ensuring the equipment is installed into the racks to create a front-to-back pattern of airflow, which will bring in cold air from the HVAC-chilled cold aisles and get rid of hot air via the hot aisles located behind the racks [11].

- **Audit Level:** Perform at the individual aisle level.
- **Item Description:** Determine how many aisles of racks are alternating hot and cold temperatures.
- **Metric:** Percentage of the aisles compliant with the alternating hot and cold aisle scheme:  $COMPLIANCE = \frac{\# \text{compliant aisles}}{\# \text{total aisles}}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which rows do not follow the alternating hot and cold aisle scheme.
- **Action(s) to Take:**
  - (1) Re-organize server racks and HVAC system to ensure 100% compliance.
- **Audit Type:** Full.

### 2.3.3 Physical Separation of Hot and Cold Air of Rack Aisles

. To further reduce mixing the cold supply air with the hot exhaust air (e.g., using flexible strip curtains or rigid enclosures). This, in turn, will reduce fan energy and improve chiller efficiency [7]. According to the U.S. Department of Energy [9], this can reduce fan energy requirements by up to 20% – 25%.

- **Audit Level:** Perform at the individual aisle level.
- **Item Description:** Determine how many aisles of racks provide physical barriers to other aisles.
- **Metric:** Percentage of the aisles compliant physical barrier scheme:  $COMPLIANCE = \frac{\# \text{compliant aisles}}{\# \text{total aisles}}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which aisles do not follow the physical barrier scheme.
- **Action(s) to Take:**
  - (1) With non-compliant aisles, install a barrier (e.g., strip curtain, rigid enclosures, etc.) to ensure 100% compliance.
- **Audit Type:** Lite.

### 2.3.4 Structured Cabling for a Rack

. Unstructured cable (see Figure 1) can restrict airflow to and from the rack, allowing for neither cold air to cool the server nor hot air to expel. Use structured cabling (see Figure 1) to avoid this.

- **Audit Level:** Perform at the individual rack level.
- **Item Description:** Determine how many racks within the data center follow acceptable structured cabling best practices.
- **Metric:** Percentage of the racks leveraging structured cabling:  $COMPLIANCE = \frac{\# \text{compliant racks}}{\# \text{total racks}}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which racks have unstructured cabling.




Figure 1: An example of unstructured cabling (left) and structured cabling (right) [12].

- **Action(s) to Take:**
  - (1) With non-compliant racks, re-wire to ensure structured cabling.
- **Audit Type:** Lite.

### 2.3.5 Air Filter MERV Rating Compliance

. The Minimum Efficiency Reporting Values (MERV) rating is the measure of an air filter's ability to capture larger particles between 0.3 and 10 microns ( $\mu\text{m}$ ) [13]. According to the American Society of Heating (ASHRAE), Refrigerating and Air-Conditioning Engineers [14], recommends that any air entering the data center use MERV 11-13 filters, whereas the air continuously present in the data center can use MERV 8 filters.

- **Audit Level:** Perform at the individual filter level.
- **Item Description:** Determine how many air filters within the data center are compliant w.r.t. MERV rating, for an air filter's given purpose (i.e., external air vs. continuous air internal to the data center).
- **Metric:** Percentage of the compliant air filters:  $COMPLIANCE = \frac{\# \text{compliant air filters}}{\# \text{total air filters}}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which air filters are not compliant.
- **Action(s) to Take:**
  - (1) With non-compliant air filters, replace with the appropriate MERV-rated air filter.
- **Audit Type:** Lite.

## 2.4 Environmental Conditions

Items related to environmental factors, such as: data center ambient temperature, rack cooling, etc.

### 2.4.1 Data Center Ambient Temperature

. One of the simplest, yet effective, measures an organization can take to immediately improve a data center's energy efficiency is analyze the data center's ambient temperature. According to a study conducted by the U.S. Environmental Protection Agency and the U.S. Department of Energy [12], many data centers set their temperatures much lower than the suggested range. Some data centers in the study were reported to have ambient temperatures as low as 55°F. The suggested range by ASHRAE [9] is depicted in Table 2. The agencies also report that data centers can save up to 4 – 5% in energy costs for every 1°F increase in server inlet temperature.

**Table 2: Data center ambient temperature ranges suggested by ASHRAE [9]**

Temperature	Class 1 and Class 2 Recommended Range	Class 1 Allowable Range	Class 2 Allowable Range
Low Temperature Limit	64.4°F	59°F	50°F
High Temperature Limit	80.6°F	89.6°F	95°F

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Capture hourly data center ambient temperature ranges.
- **Metric:**
  - (1) Percentage of the time within the suggested ASHRAE range:  $COMPLIANCE = \frac{\# \text{ acceptable ambient temperature readings}}{\# \text{ total ambient temperature readings}}$ .
  - (2) Plot the ambient temperature over time, and aim to be near the high temperature range, depicted in Table 2.
- **Desired Output(s) & Goal(s):**
  - (1) Determine if data center follows ASHRAE guidelines for data center temperature ranges.
- **Action(s) to Take:**
  - (1) Raise the data center's ambient temperature to not only be within the acceptable ASHRAE suggested temperature range, but strive to be near the high temperature limit.
- **Audit Type:** Lite.

#### 2.4.2 Rack Cooling Index(RCI)

. The Rack Cooling Index (RCI) [15] is a key indicator explaining how effectively individual racks are cooled and maintained. The RCI comprises of two metrics:  $RCI_{HI}$  and  $RCI_{LO}$ , where  $RCI_{HI}$  measures the presence or absence of over-temperatures and  $RCI_{LO}$  measures the presence or absence of under-temperatures.

- **Audit Level:** Perform at the individual rack level.
- **Item Description:** Capture the  $RCI_{HI}$  and  $RCI_{LO}$  for a given rack.
- **Metric:** Percentage for both  $RCI_{HI}$  and  $RCI_{LO}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which rooms or data centers do not have sufficient temperature management (i.e.,  $RCI_{type} \neq 100\%$ ).
- **Action(s) to Take:**
  - (1) Adjust rack temperature management to ensure optimal  $RCI_{HI}$  and  $RCI_{LO}$  (100%).
- **Audit Type:** Full.

## 2.5 Global

High-level items to monitor overall efficiency of energy use, such as; electrical, water, HVAC, airflow, etc.

#### 2.5.1 Power Usage Effectiveness (PUE)

. The Power Usage Effectiveness (PUE) [9] is a standard metric to measure the overall efficiency of power used by the data center.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Capture the PUE at an organization-defined cadence to track over time.
- **Metric:** The ratio of the total power to run the data center facility to the total power drawn by all IT equipment:  $PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$ .
- **Desired Output(s) & Goal(s):**

- (1) Benchmark PUE per data center; overall current state of data center energy efficiency, per U.S. Department of Energy standards, depicted in Table 3.

- **Action(s) to Take:**

- (1) Once an organization has a benchmark understanding of its PUE, it can use the non-metric best practices to take actions to improve PUE and track it over time. This benchmark can also be used in coöperation with any existing sustainability strategy from the organization, to drive goals for future PUE values (see section 3).

- **Audit Type:** Lite.

#### 2.5.2 Data Center Infrastructure Effectiveness (DCIE)

. The Data Center Infrastructure Effectiveness (DCIE) [9] is the inverse of PUE-defined as the ratio of the total power drawn by all IT equipment to the total power to run the data center facility.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Capture the DCIE at an organization-defined cadence to track over time.
- **Metric:** The ratio of the total power drawn by all IT equipment to the total power to run the data center facility (inverse of PUE):  $DCIE = \frac{1}{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$ .
- **Desired Output(s) & Goal(s):**
  - (1) Benchmark DCIE per data center; overall current state of data center energy efficiency, per U.S. Department of Energy standards, depicted in Table 3.
- **Action(s) to Take:**
  - (1) Once an organization has a benchmark understanding of its DCIE, it can use the non-metric best practices to take actions to improve DCIE and track it over time. This benchmark can also be used in coöperation with any existing sustainability strategy from the organization, to drive goals for future DCIE values (see section 3).
- **Audit Type:** Lite.

#### 2.5.3 Energy Reuse Effectiveness (ERE)

. The Energy Reuse Effectiveness (ERE) [9] is the ratio of the total energy to run the data center facility minus the reuse energy to the total energy drawn by all IT equipment. ERE value range: 0 – ∞. For example, an ERE of 0 means that 100% of the energy brought into the data center is reused elsewhere, outside of the data center control volume.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Capture the ERE at an organization-defined cadence to track over time.
- **Metric:** The ratio of the total energy to run the data center facility minus the reuse energy to the total energy drawn by all IT equipment:  $ERE = \frac{\text{Cooling} + \text{Power} + \text{Lighting} + \text{IT} - \text{Reuse Energy}}{\text{IT Equipment Energy}}$ .
- **Desired Output(s) & Goal(s):**

**Table 3: U.S. Department of Energy recommendations [9] for standard data center energy efficiency metrics**

Metric Name	Standard Score	Good Score	Better Score
Power Usage Effectiveness (PUE)	2	1.4	1.1
Data Center Infrastructure Efficiency (DCIE)	0.5	0.7	0.9
HVAC System Effectiveness (HVACSE)	0.7	1.4	2.5
Airflow Efficiency (AE)	1.25 W/cfm	0.75 W/cfm	0.5 kW/cfm
Cooling System Efficiency (CSE)	1.1 kW/ton	0.8 kW/ton	0.6 kW/ton

- (1) Benchmark *ERE* per data center; overall current state of data center energy efficiency.

- **Action(s) to Take:**

- (1) Once an organization has a benchmark understanding of its *ERE*, it can use the non-metric best practices to take actions to improve *ERE* and track it over time. This benchmark can also be used in cooperation with any existing sustainability strategy from the organization, to drive goals for future *ERE* values (see section 3).

- **Audit Type:** Lite.

#### 2.5.4 HVAC System Effectiveness (HVACSE)

. HVAC System Effectiveness (HVACSE) [9] is the ratio of the annual IT equipment energy to the annual HVAC system energy.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Capture the *HVACSE* annually to track over time.
- **Metric:** The ratio of the annual IT equipment energy to the annual HVAC system energy:  $HVACSE = \frac{kWh/y_{IT}}{kWh/y_{HVAC}}$ .
- **Desired Output(s) & Goal(s):**

(1) Benchmark *HVACSE* per data center; overall current state of data center energy efficiency, per U.S. Department of Energy standards, depicted in Table 3.
- **Action(s) to Take:**

(1) Once an organization has a benchmark understanding of its *HVACSE*, it can use the non-metric best practices to take actions to improve *HVACSE* and track it over time. This benchmark can also be used in cooperation with any existing sustainability strategy from the organization, to drive goals for future *HVACSE* values (see section 3).
- **Audit Type:** Lite.

#### 2.5.5 Airflow Efficiency (AE)

. Airflow Efficiency (AE) [9] is the overall measure of how efficiently air is moved through the data center.

- **Audit Level:** Perform at the individual fan level.
- **Item Description:** Capture the *AE* at an organization-defined cadence to track over time; roll up to the entire data center level.
- **Metric:** The overall airflow efficiency in terms of the total fan power required per unit of airflow:  $AE = \frac{Total\ Fan\ Power\ (W)}{Total\ Fan\ Airflow\ (cfm)}$ .
- **Desired Output(s) & Goal(s):**

(1) Benchmark *AE* per data center; overall current state of data center energy efficiency, per U.S. Department of Energy standards, depicted in Table 3.
- **Action(s) to Take:**

- (1) Once an organization has a benchmark understanding of its *AE*, it can use the non-metric best practices to take actions to improve *AE* and track it over time. This benchmark can also be used in cooperation with any existing sustainability strategy from the organization, to drive goals for future *AE* values (see section 3).

- **Audit Type:** Lite.

#### 2.5.6 Cooling System Efficiency (CSE)

. Cooling System Efficiency (CSE) [9] is a metric to measure the efficiency of an HVAC system.

- **Audit Level:** Perform at the individual HVAC system level.
- **Item Description:** Capture the *CSE* at an organization-defined cadence to track over time; roll up to the entire data center level.
- **Metric:** The ratio of average cooling system power usage (*kW*) to the average data center cooling load (tons):  $CSE = \frac{\mu\ Cooling\ System\ Power\ (kW)}{\mu\ Cooling\ Load\ (cfm)}$ .
- **Desired Output(s) & Goal(s):**

(1) Benchmark *CSE* per data center; overall current state of data center energy efficiency, per U.S. Department of Energy standards, depicted in Table 3.
- **Action(s) to Take:**

(1) Once an organization has a benchmark understanding of its *CSE*, it can use the non-metric best practices to take actions to improve *CSE* and track it over time. This benchmark can also be used in cooperation with any existing sustainability strategy from the organization, to drive goals for future *CSE* values (see section 3).
- **Audit Type:** Lite.

## 2.6 IT Equipment

Items related to the performance and utilization of servers and other IT equipment.

#### 2.6.1 Identify Unused Operational Servers

. Running unused operational servers-or any IT equipment-can be expensive and wasteful of resources.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Determine which equipment are not used within the data center.
- **Metric:** Percentage of the equipment unused:  $\frac{\#\ unused\ equipment}{\#\ total\ equipment}$ .
- **Desired Output(s) & Goal(s):**

(1) Identify which operational servers are not in use and consider them candidates for retirement.
- **Action(s) to Take:**

- (1) Determine which unused operational servers to retire, to ensure all servers that are drawing energy are in use.

- **Audit Type:** Lite.

### 2.6.2 Monitor Server CPU Utilization

. Understand server performance *w.r.t.* CPU utilization. Ignacio Cano et al. [16] performed a study in which 80% of devices in a sample of VMs have a maximum resource usage that is < 60%, leaving 40% of allocated CPU resources un-utilized. This is a waste of resources, and can be corrected by optimizing server allocation.

- **Audit Level:** Perform at the individual server level.
- **Item Description:** Determine if a server is under- or over-utilized:
  - **Under-utilization** [16]:  $usage \leq 50\%$
  - **Correct utilization** [16]:  $51\% \leq usage \leq 85\%$
  - **Over-utilization** [16]:  $usage \leq 86\%$
- **Metric:** Percentage of the CPU utilized.
- **Desired Output(s) & Goal(s):**
  - (1) Identify which operational servers are under-utilized or over-utilized.
- **Action(s) to Take:**
  - (1) Optimize server performance to avoid under-utilization (poor energy management) and over-utilization (wear and tear).
  - (2) Consolidate under-utilized servers.
- **Audit Type:** Full.

### 2.6.3 Equipment Efficiency

. Understand how performant current IT equipment is *w.r.t.* standard metric [4] of  $GFLOP/W$ .

- **Audit Level:** Perform at the individual server level.
- **Item Description:** Determine how energy efficient a server is.
- **Metric:**  $GFLOP/W$ .
- **Desired Output(s) & Goal(s):**
  - (1) Identify which equipment have a low value of  $GFLOP/W$ .
- **Action(s) to Take:**
  - (1) Procure energy efficient servers to replace inefficient equipment.
- **Audit Type:** Lite.

## 2.7 IT Power Distribution Chain

Items related to the efficiency of systems providing the data center with power.

### 2.7.1 Trace Data Center Power Generation Sources

. Understand from where power fed to the data center are coming. Over time, non-renewable power generation sources should be phased out in favor of replacing with renewable sources.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Determine if an energy supply for the data center is renewable or non-renewable, and how much power is coming from each source.
- **Metric:** Percentage of power from renewable vs. non-renewable sources:  $\frac{\text{Total kWh from renewable}}{\text{Total kWh drawn}}$
- **Desired Output(s) & Goal(s):**

- (1) Identify how electrical power is generated for the data center.

- **Action(s) to Take:**

- (1) Work with the organization's facilities group to create a strategy to phase out non-renewable sources of power in favor of renewable sources. The goal is 100%.

- **Audit Type:** Lite.

## 2.8 Lighting

Items related understanding the impact of lights, lamps, and bulbs throughout the data center.

### 2.8.1 LED Bulbs

. A quick and easy way an organization can reduce a data center's environmental footprint is by looking at the bulbs used in its lights and lamps. According to the U.S. Department of Energy [17], LED bulbs required 75% less energy than traditional incandescents.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Determine how many light bulbs in the data center are LEDs.
- **Metric:** Percentage of LED bulbs vs. traditional incandescents:  $\frac{\# \text{ LED bulbs}}{\# \text{ Total bulbs}}$
- **Desired Output(s) & Goal(s):**
  - (1) Identify which lamps do not contain LED bulbs.
- **Action(s) to Take:**
  - (1) Replace all non-LED bulbs with LED. The goal is 100%.
- **Audit Type:** Lite.

### 2.8.2 Lighting Control and Dimming

. Outside of replacing all non-LED bulbs with LEDs, one way an organization can reduce energy costs within a data center is to enable lights and lamps to operate at a lower level.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Determine how many lamps in the data center offer dimming.
- **Metric:** Percentage of lamps which support dimming:  $\frac{\# \text{ Lamps offer dimming}}{\# \text{ Total lamps}}$
- **Desired Output(s) & Goal(s):**
  - (1) Identify which lamps do not offer dimming.
- **Action(s) to Take:**
  - (1) Install dimming controls on all non-dimming lamps and lights. The goal is 100%.
- **Audit Type:** Lite.

### 2.8.3 Occupancy Sensors to Control Lights

. This is very closely related to the previous best practice. Sensors to determine if a person is present or absent, and turning lights off when not needed, can help save energy and costs in operating the data center.

- **Audit Level:** Perform at the entire data center level.
- **Item Description:** Determine how many lamps in the data center offer control via occupancy sensors.
- **Metric:** Percentage of lamps which support control via occupancy sensors:  $\frac{\# \text{ Lamps with sensors}}{\# \text{ Total lamps}}$
- **Desired Output(s) & Goal(s):**
  - (1) Identify which lamps do not offer control via occupancy sensors.

Data Center Efficiency Dashboard

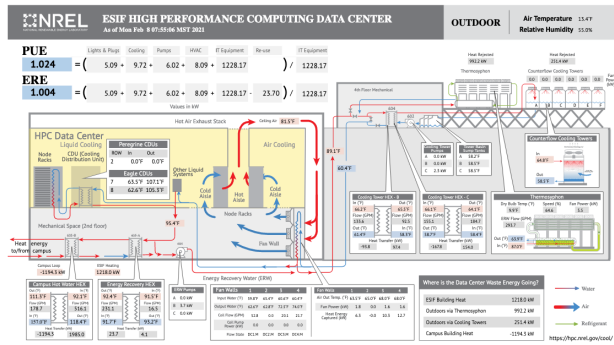


Figure 2: An example of a data center's energy efficiency dashboard, from the National Renewable Energy Laboratory's ESIF High Performance Computing Data Center [21].

- **Action(s) to Take:**

- (1) Install occupancy sensors on all non-equipped lamps and lights. The goal is 100%.

- **Audit Type:** Lite.

### 3 ORGANIZATIONAL CONSIDERATIONS

Having a framework to conduct a data center energy efficiency audit is only the beginning in the journey for an organization to make its data centers more environmentally friendly. An organization must also have a team in place to execute the audit and make changes once benchmark performance data are collected and a strategy to make necessary changes to data centers-typically to contribute to an ongoing sustainability campaign (such as Google's [18], Microsoft's [19], or Facebook's [20] commitments to become more sustainable organizations). The goal of this section is to provide the information necessary to help organizations in this space to ensure a successful energy audit, as well as enable the organization post-audit to make the changes necessary to improve energy efficiency and reduce the environmental impact.

#### 3.1 Creating a Team

The Lawrence Berkeley National Laboratory's Center of Expertise for Energy Efficiency in Data Centers provides some insight on the creation of a team to execute some of the best practices for energy-efficient data centers [7]; however, given the ambiguity of the energy audit criterion described, I wanted to provide suggestions to complement the Lawrence Berkeley National Laboratory's Center of Expertise for Energy Efficiency in Data Centers's suggestions [8].

Roles:

- **Audit Lead:** The Audit Lead will coordinate with Site Lead(s) and Organization Engagement Lead/Liaison to organize audit and communicate results to responsible parties within the organization. The Audit Lead will also distribute work to the Audit Assessors and meet regularly to track audit progress and understand benchmark performance.

- **Site Lead(s):** This is recommended by the Lawrence Berkeley National Laboratory's Center of Expertise for Energy Efficiency in Data Centers [7]. The lab defines this role as: "Data center technical representative that will be the primary person participating in the assessment." This individual will lead the Site Consulting Engineer(s) in any technical work that needs to be completed in the data center for the purposes of the audit, and will report findings back to the Audit Assessor(s) and Audit Lead.
- **Audit Assessor(s):** This is recommended by the Lawrence Berkeley National Laboratory's Center of Expertise for Energy Efficiency in Data Centers [7]. The lab defines this role as: "The energy expert assigned to complete the...assessment; the expert serves as the facilitator for all activities." Depending on the size of the organization, > 1 Audit Assessor may be necessary-especially if multiple data centers are included in the audit.
- **Site Consulting Engineer(s):** At least one Data Engineer per data center to provide information on infrastructure, installing sensors to collect data, handling operational servers, etc.
- **Data Scientist(s):** At least one Data Scientist to collect, parse, visualize, and communicate data required to successfully execute the audit. Suggested action: create dashboards to report, track, and monitor data center performance over time-not only during the audit to benchmark performance. An example of this dashboard can be seen in Figure 2.
- **Organization Engagement Lead/Liaison:** Serves as the point of contact between organization management (e.g., for funding, marketing/communication/organizational publicity), facilities (e.g., for power generation, data center lighting control, etc.), and sustainability (e.g., for contributing to an organization's ongoing sustainability strategy).

#### 3.2 Setting Energy Efficiency Goals and Benchmarks Post-Audit

Once the initial data center energy efficiency audit is complete, the organization will understand the environmental impact its data center(s) are having. At this time, it is prudent for the organization to make changes necessary to improve energy efficiency and reduce the environmental footprint of the data centers. Outside of the taking the recommended actions based on the result of the items in the audit (described in length in the previous section), the organization should continue to track and monitor its energy efficiency at a regular cadence (quarterly reporting is recommended, per [18]; however, a far more granular cadence of collecting these data can enable the organization to leverage techniques like machine learning to optimize the data center's parameters to maximize energy efficiency [22]).

An important consideration to make is setting annual goals for energy efficiency improvement, as well as reporting mechanisms to communicate these improvements. As seen in the examples from Google [18], Microsoft [19], and Facebook [20], these organizations have set ambitious environmental goals for the near future to reach carbon-free energy generation, carbon negativity, and net zero GHG emissions, respectively. However, smaller organizations-and



higher education institutions-can set similar goals to improve its environmental footprint. For example, Purdue University is working toward its Sustainability Master Plan [23], consisting of goals related to energy, water, materials/waste, buildings, and the grounds of the University to reduce the impact operations have on the environment. To demonstrate that a data center energy efficiency initiative contributes to a strategy such as this:

Sustainability Master Plan Initiatives:

- **Energy [24]:**

- *E-1 – Cut Carbon Emissions in Half*: "Reduce...carbon emissions by 50% by FY25, with FY11 as the baseline year."
  - (1) Improving data center energy efficiency (decreasing PUE, for example) will reduce the required power drawn from power sources potentially contributing to an organization's carbon emissions. Actions as trivial as dimming data center lights when not in use contributes to this goal.
  - (2) Over time, matriculating a data center's power generation sources from non-renewable sources to renewable sources can reduce an organization's overall carbon emission footprint.
- *E-2 – No Net Gain*: "Cap total energy consumption at FY11 levels."
  - (1) Improving data center energy efficiency (decreasing PUE, for example) will reduce the required power drawn from power sources potentially contributing to an organization's carbon emissions. Actions as trivial as dimming data center lights when not in use contributes to this goal.
  - (2) Over time, matriculating a data center's power generation sources from non-renewable sources to renewable sources can reduce an organization's overall carbon emission footprint.

- **Water [25]:**

- *W-1 - Reduce Water Consumption by 30%*: "Reduce potable water consumption inside buildings and for irrigation by 30% by FY25 on a gallon per square foot basis, with FY11 as the baseline year."
  - (1) Employing best practices to improve rack cooling efficiency and optimizing operational servers in use can reduce the amount of chilling required for the data center, thus decreasing water consumption.

## REFERENCES

- [1] How to take responsibility for your data centre's environmental impact. Also available as <https://www.capgemini.com/2020/01/the-more-sustainable-data-center/>.
- [2] Electricity - consumption (kwh) 2020 country comparisons, ranks, by rank. Also available as [https://photius.com/rankings/2020/energy/electricity\\_consumption\\_2020\\_0.html](https://photius.com/rankings/2020/energy/electricity_consumption_2020_0.html).
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. Also available as <https://arxiv.org/abs/2005.14165>.
- [4] November 2020 | top500. Also available as <https://www.top500.org/lists/green500/2020/11/>.
- [5] Matthew Burruss. The (un)ethical story of gpt-3: Openai's million dollar model | matthew burruss. Also available as <https://matthewpburruss.com/post/the-unethical-story-of-gpt-3-openai-million-dollar-model/>.
- [6] Frequently asked questions (faqs) - u.s. energy information administration (eia). Also available as <https://www.eia.gov/tools/faqs/faq.php?id=97&t=3#:~:text=How%20much%20electricity%20does%20an,about%20914%20kWh%20per%20month>.
- [7] Data center master list of energy efficiency actions | center of expertise for energy efficiency in data centers. Also available as <https://datacenters.lbl.gov/resources/data-center-master-list-energy>.
- [8] Department of energy(doe) energy assessment process manual | center of expertise for energy efficiency in data centers. Also available as <https://datacenters.lbl.gov/resources/assessment-resources-energy-training-assessment-process-manual>.
- [9] Best practices guide for energy-efficient data center design | department of energy. Also available as <https://www.energy.gov/eere/femp/downloads/best-practices-guide-energy-efficient-data-center-design>.
- [10] Magnus Herrlin. Airflow and cooling performance of data centers: Performance metrics. *ASHRAE Trans*, 114, 01 2008. Also available as [http://www.ancis.us/images/SL-08-018\\_Final.pdf](http://www.ancis.us/images/SL-08-018_Final.pdf).
- [11] Effects of airflow on the thermal environment and energy efficiency in raised-floor data centers: A review - sciencedirect. Also available as <https://www.sciencedirect.com/science/article/abs/pii/S0048969719337428>.
- [12] Top 12 ways to decrease the energy consumption of your data center | energy star. Also available as <https://www.energystar.gov/buildings/tools-and-resources/top-12-ways-decrease-energy-consumption-your-data-center>.
- [13] What is a merv rating? | indoor air quality (iaq) | us epa. Also available as <https://www.epa.gov/indoor-air-quality-iaq/what-merv-rating-1>.
- [14] Gaseous and particulate contamination guidelines for data centers. Also available as [https://tc0909.ashraetcs.org/documents/ASHRAE\\_Contamination\\_Whitepaper\\_30\\_July\\_2009.pdf](https://tc0909.ashraetcs.org/documents/ASHRAE_Contamination_Whitepaper_30_July_2009.pdf).
- [15] M.K. Herrlin. Rack cooling effectiveness in data centers and telecom central offices: The rack cooling index (rci). *ASHRAE Transactions*, 111:725–731, 01 2005. Also available as [http://www.ancis.us/images/RCI--for\\_Website.pdf](http://www.ancis.us/images/RCI--for_Website.pdf).
- [16] Ignacio Cano, Lequn Chen, Pedro Fonseca, Tianqi Chen, Chern Cheah, Karan Gupta, Ramesh Chandra, and Arvind Krishnamurthy. Adares: Adaptive resource management for virtual machines, 2018. Also available as <https://arxiv.org/abs/1812.01837>.
- [17] Led lighting | department of energy. Also available as <https://www.energy.gov/energysaver/save-electricity-and-fuel/lighting-choices-save-you-money/led-lighting>.
- [18] Progress | google sustainability. Also available as <https://sustainability.google/progress/#data-centers>.
- [19] Environmental sustainability | microsoft csr. Also available as [https://www.microsoft.com/en-us/corporate-responsibility/sustainability?activetab=pivot\\_1:primaryr3](https://www.microsoft.com/en-us/corporate-responsibility/sustainability?activetab=pivot_1:primaryr3).
- [20] Sustainable data centers - facebook sustainability. Also available as <https://sustainability.fb.com/innovation-for-our-world/sustainable-data-centers/#section-Energy>.
- [21] Data center efficiency dashboard | high-performance computing | high-performance computing | nrel. Also available as <https://www.nrel.gov/hpc/cool.html>.
- [22] Data centers | google sustainability. Also available as <https://sustainability.google/progress/projects/machine-learning/>.
- [23] Administrative Operations Communications. 2020 physical facilities sustainability master plan - physical facilities - purdue university. Also available as <https://www.purdue.edu/physicalfacilities/sustainability/sustainability-master-plan/index.html>.
- [24] Administrative Operations Communications. Energy - physical facilities - purdue university. Also available as <https://www.purdue.edu/physicalfacilities/sustainability/sustainability-master-plan/energy.html>.
- [25] Administrative Operations Communications. Water - physical facilities - purdue university. Also available as <https://www.purdue.edu/physicalfacilities/sustainability/sustainability-master-plan/water.html>.



**Table 4: Items included only in the "lite" data center energy efficiency audit.**

Item Title	Audit Type
Physical Separation of Hot and Cold Air of Rack Aisles	Lite
Structured Cabling for a Rack	Lite
Air Filter MERV Rating Compliance	Lite
Data Center Ambient Temperature	Lite
Power Usage Effectiveness	Lite
Data Center Infrastructure Efficiency	Lite
Energy Reuse Effectiveness	Lite
HVAC System Effectiveness	Lite
Airflow Efficiency	Lite
Cooling System Efficiency	Lite
Identify Unused Operational Servers	Lite
Equipment Efficiency	Lite
Identify Power Sources	Lite
LED Bulbs	Lite
Lighting Control and Dimming	Lite
Occupancy Sensors to Control Lights	Lite

**Table 5: Items included only in the "full" data center energy efficiency audit. NOTE: To complete a "full" audit, the items from "lite" an "full" must both be completed.**

Item Title	Audit Type
Return Temperature Index	Full
Alternating Hot and Cold Rack Aisles	Full
Rack Cooling Index	Full
Monitor Server CPU Utilization	Full

**Table 6: All audit items, with their audit type classification.**

Item Title	Audit Type
Return Temperature Index	Full
Alternating Hot and Cold Rack Aisles	Full
Physical Separation of Hot and Cold Air of Rack Aisles	Lite
Structured Cabling for a Rack	Lite
Air Filter MERV Rating Compliance	Lite
Data Center Ambient Temperature	Lite
Rack Cooling Index	Full
Power Usage Effectiveness	Lite
Data Center Infrastructure Efficiency	Lite
Energy Reuse Effectiveness	Lite
HVAC System Effectiveness	Lite
Airflow Efficiency	Lite
Cooling System Efficiency	Lite
Identify Unused Operational Servers	Lite
Monitor Server CPU Utilization	Full
Equipment Efficiency	Lite
Identify Power Sources	Lite
LED Bulbs	Lite
Lighting Control and Dimming	Lite
Occupancy Sensors to Control Lights	Lite