Cost to Support Compute Capacity

Data Center Performance Benchmark Series

A Special Analysis for Emerson Network Power

Independently conducted by Ponemon Institute LLC

Publication Date: August 2016

Cost to Support Compute Capacity

Ponemon Institute, August 2016

1. Introduction

Ponemon Institute and Emerson Network Power are pleased to present the results of an original benchmark study to determine average costs to support 1 kW of compute capacity in today's data centers. The purpose of this study is to analyze the major cost components in supporting compute capacity so that organizations can more effectively identify opportunities to reduce costs and make informed decisions about future capacity.

The results of this study are based on data from 41 data centers, representing 31 companies, who reported on their costs in four categories that together comprise total data center costs: Physical Plant, IT Assets, Operating Costs and Energy Costs.

These organizations also reported on data center size, IT load, number of racks and median rack density, enabling the Ponemon Institute to quantify the cost to support 1 kW of capacity for data centers in five size ranges:

- 500–5,000 sq. ft.
- 5,001–10,000 sq. ft.
- 10,001–25,000 sq. ft.
- 25,001–50,000 sq. ft.
- > 50,000 sq. ft.

2. Benchmarking Methods

Our benchmark instrument collected descriptive cost information for 41 data centers. Senior-level IT personnel, managers of data center facilities and finance and accounting specialists provided cost-related information. The total cost of data center operations is organized into four discrete components, as follows.

- **Total cost of plant and fixtures** amortized over 20 years, plus the total cost of power protection and distribution and cooling systems amortized over 10 years.

- **Total cost of IT assets** amortized over 3 years, which includes servers, routers, storage devices, network equipment and telecommunications. In addition to hardware costs, this category includes the cost of software applications.

- **Total cost of data center operations**, which includes both direct and indirect labor costs plus overhead and administration costs. This category also includes software licensing fees, support services and system maintenance.

- **Total energy cost**, which is the cost of power measured in kW usage per year.

Table 1 summarizes costs for the 41 participating data centers.

Table 1: Average costs by category for all participating data centers.

	Mean	Median	Minimum	Maximum
Plant*	255,000	258,000	97,000	409,000
IT assets*	525,000	531,000	168,000	789,000
Operating costs*	2,241,000	2,248,000	609,000	3,500,000
Energy costs*	1,912,000	1,734,000	683,000	3,416,000
Total costs*	4,933,000	4,657,000	1,557,000	8,099,000

*Estimated costs are presented in US dollars

To maintain complete confidentiality, the survey instrument does not capture company-specific information. Research materials do not contain tracking codes or other methods that could link responses to participating companies.

The cost benchmarking process was launched in January 2016 and fieldwork concluded in May 2016. Recruitment started with a personalized letter and a follow-up phone call to 63 data centers that participated in an earlier study on the cost of unplanned outages conducted by Ponemon Institute.[1] All of these organizations are members of Ponemon Institute's benchmark community. This resulted in 31 organizations agreeing to participate.[2]

[1]See: The Cost of Data Center Outages conducted by Ponemon Institute and sponsored by Emerson Network Power published in January 2016.
[2]The Ponemon Institute's benchmark community is comprised of organizations that have participated in one or more research studies over the past 14 years.

3. Key Findings

Table 2 shows annual and monthly average costs per kW and average cost per rack for each of the five data center size ranges. Average annual cost/kW range from $5,467 for data centers in the > 50,000 sq. ft. range to $26,495 for data centers in the 500–5,000 sq. ft. range.

Table 2. Annual and monthly costs per kW and rack by data center size.

Data Center Size (sq. ft.)	No. of Racks	Avg Compute Load (kW)	Avg Rack Density (kW)	Average Cost Per kW		Average Cost Per Rack	
				Annually	Monthly	Annually	Monthly
500 to 5,000	28.5	105	3.5	$26,495	$2,208	$97,614	$8,134
5,001 to 10,000	58	318	5.7	$13,662	$1,135	$74,689	$6,224
10,001 to 25,000	95	620	6.5	$8,464	$705	$55,242	$4,604
25,001 to 50,000	128.5	972	8	$6,734	$561	$50,841	$4,245
> 50,000	183	1,400	7.8	$5,467	$456	$41,825	$3,485

Figure 1 plots the affect of data center size on annual cost/kW of IT load and by rack.

Figure 1: Cost per rack and kW decrease as data center size increases.

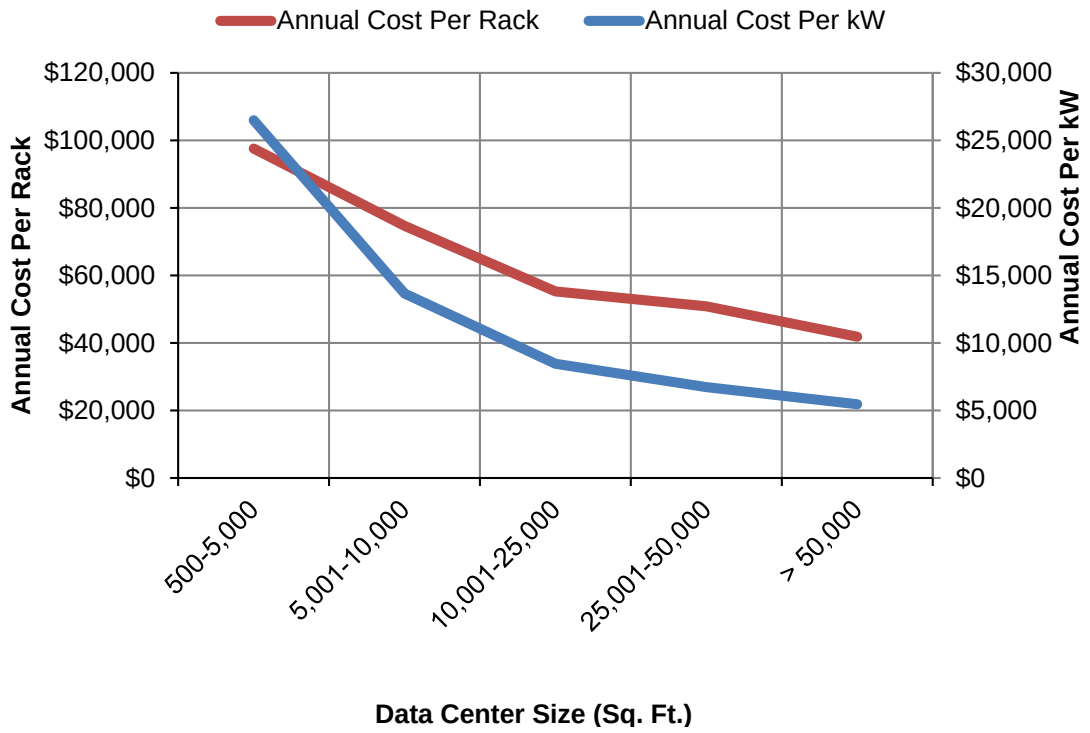

Figure 2 illustrates costs by category for each of the five data center size ranges. Amortized Plant and IT Asset costs account for just 15 to 20 percent of annual costs across all size ranges, while Energy and Operating costs account for 80 to 85 percent of annual costs. In all cases, Operating Costs, which include personnel, administrative, overhead and licensing costs, represent the largest percentage of total costs, accounting for 45 to 49 percent of total costs.

Figure 2: Total costs by category for each of the five size ranges.

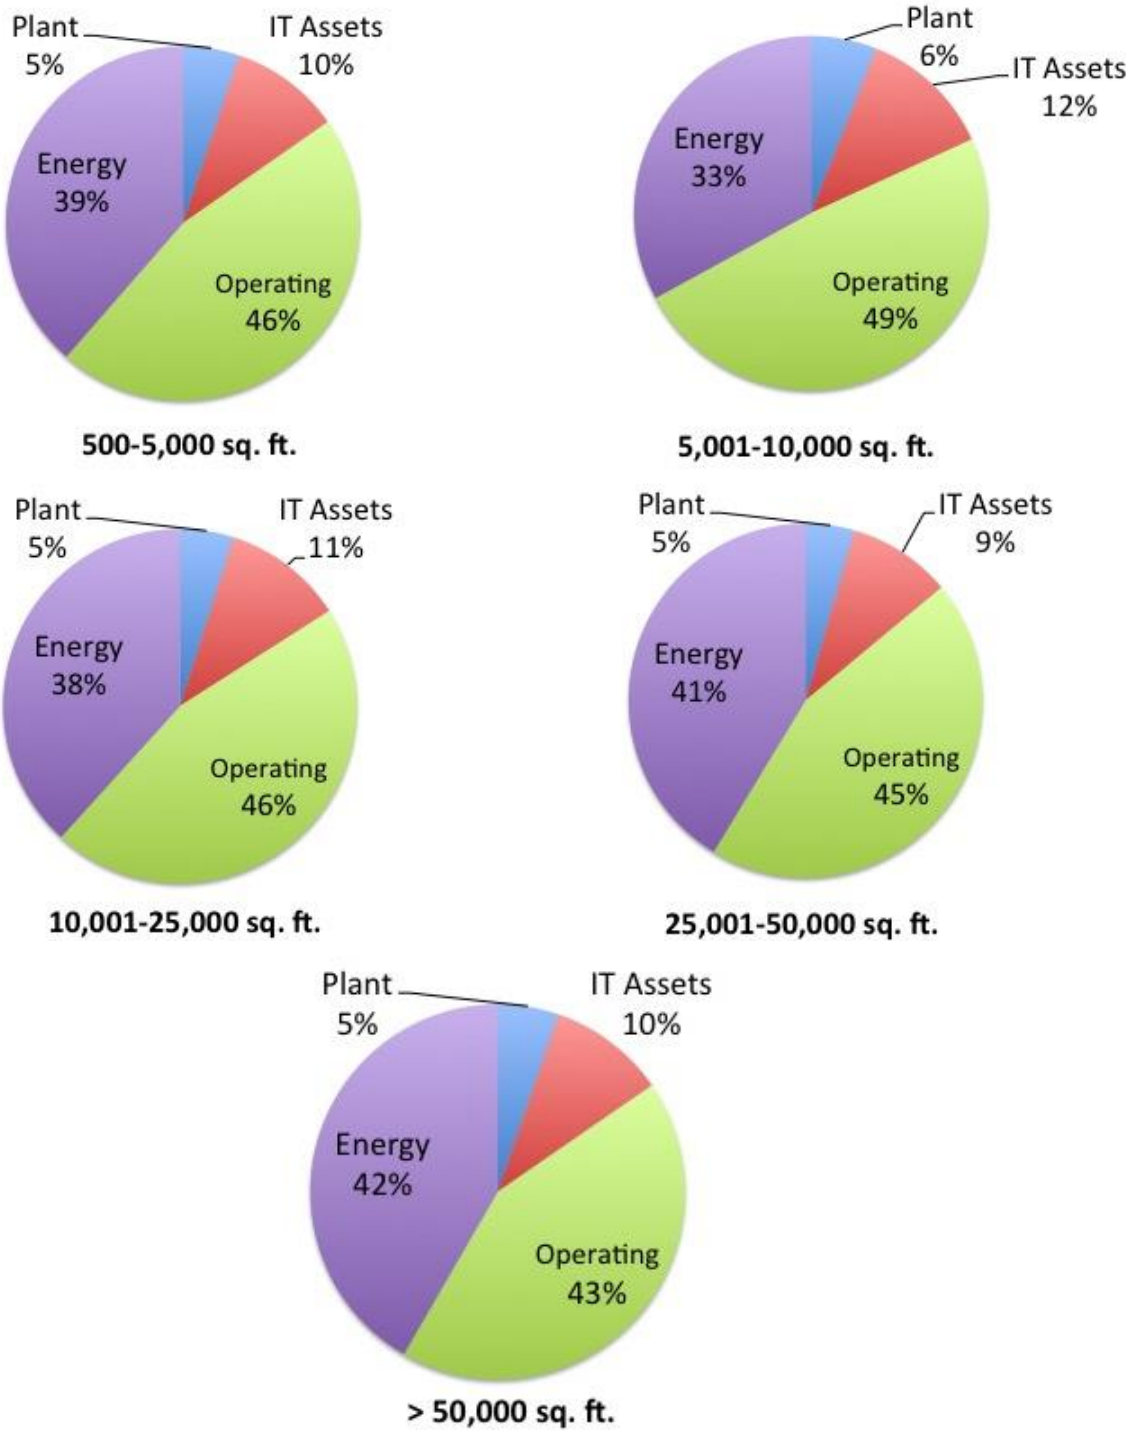

Figure 3: Cost for all categories goes down as data center size increases.

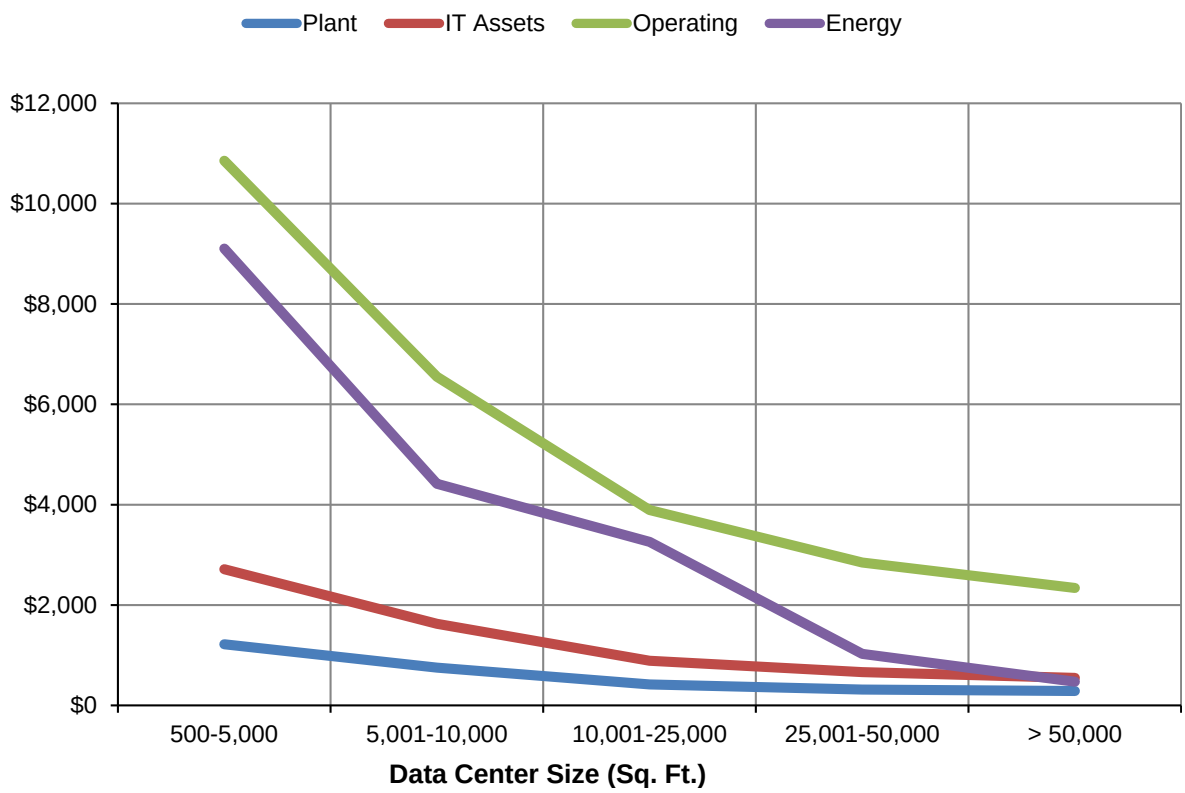

Costs were also analyzed by rack density. While larger data centers had higher average rack densities, an analysis of data centers in the same size range with different rack densities supports a correlation between higher rack density and lower costs. This analysis is presented in the following section.

Figure 4: Cost per kW decreases as rack density increases.

Eight of the data centers included in the study were colocation facilities. Three of the eight were in the 10,001-25,000 sq. ft. range, three in the 25,001-50,000 sq. ft. range, and two in the over 50,000 sq. ft range. While sample sizes are limited, Figure 5 shows average costs for these data centers compared to other comparably sized data centers that participated in the study.

Figure 5: Cost per kW for colocation facilities versus all others.

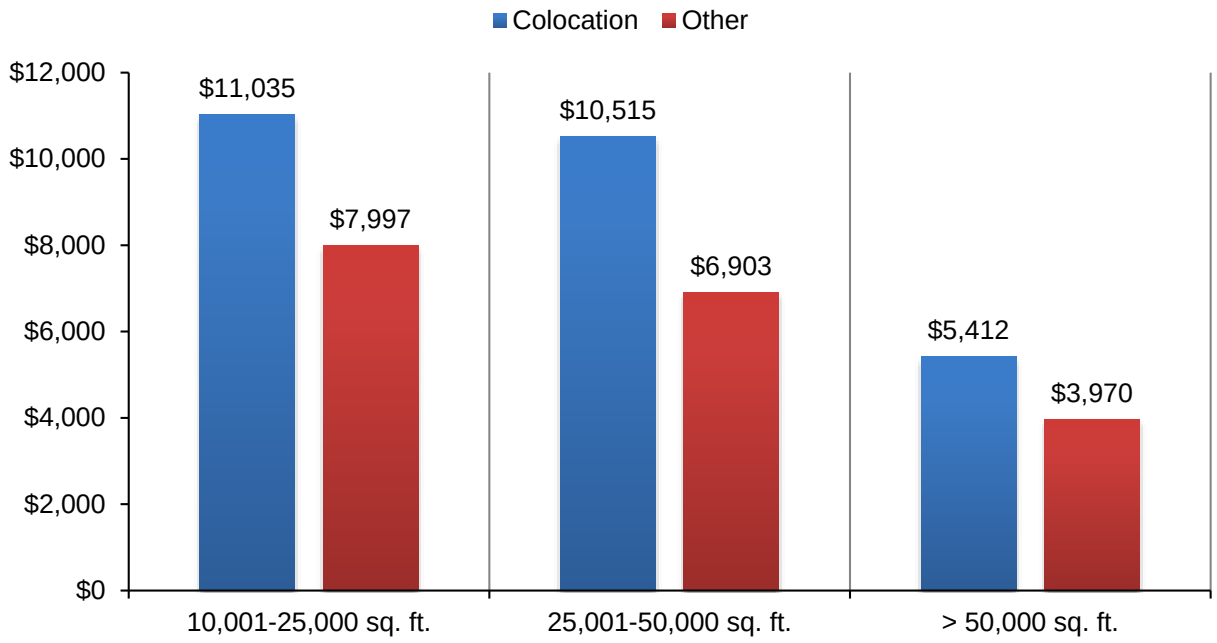

4. Analysis of Findings

Following are some of the key findings from this data:

- When IT Asset and Plant costs are amortized, Energy and Operating costs together account for 80 percent or more of annual data center costs and represent a significant investment for the organizations they support. These costs appear even higher when typical utilization rates, which have been estimated in other studies at 6–12 percent, with 30 percent of servers defined as comatose, are factored in. Efforts to increase utilization, decommission comatose servers and increase operating and energy efficiency can generate significant annual savings.

    - While energy costs per kW do decrease as data center size increases, there does seem to be significant opportunities to reduce energy costs across all size ranges. Achieving the 50 percent reduction in energy costs demonstrated in Energy Logic in a 10,000–25,000 sq. ft. data center could reduce cost to support 1kW of compute from $8,464 to $6,834 per kW. At an average IT load of 620 kW, that translates into an annual savings of $1,010,600.

    - Energy efficiency has received significant attention within the industry, yet the data suggests personnel productivity also presents an opportunity. Improving productivity of personnel to achieve a 20 percent reduction in operating costs reduces average cost to support 1 kW of compute for a 10,000–25,000 sq. ft. data center from $8,464 to $7,685. At an average load of 620 kW, this translates into annual savings of $482,980.

- Costs per kW and rack decrease significantly with data center size. This suggests that smaller data centers have the strongest economic case for moving capacity to the cloud or colocation to leverage economies of scale. Midsize and large data centers may have valid reasons for outsourcing compute; however, they should evaluate the impact of the cost to support existing capacity in their decision. Unless increased dependence on cloud services will result in a reduction in labor and overhead costs, it may result in higher cost/kW to support existing capacity. In many cases, there appears be an opportunity to add capacity by increasing rack density without significantly impacting total annual costs.

- The variance in rack density within each size range did not enable a statistical analysis of the impact of rack density within each data center size range. However, an analysis of similar cases within each category does illustrate how higher rack densities can reduce the cost to support the IT load on a kW basis.

    - In the 5,001–10,000 sq. ft. size range, one data center supported a 351 kW load across 78 racks with a median rack density of 4.5 kW. Another supported a 425 kW load across 50 racks with a median rack density of 8.5 kW. In the first case, the cost/kW was $14,376—5 percent higher than the average for the size range. The second case had an average cost of $11,870 per kW, 11 percent lower than the average for the size range.

- In the 10,001–25,000 sq. ft. size range, one data center supported a 401 kW IT load across 89 racks with a median rack density of 4.5 kW, while another supported a 514 kW IT load across 79 racks with a median rack density of 6.5 kW. In the first case, the cost to support 1 kW was $10,960, 25 percent higher than the $8,464 average for the size range. The second case had a cost of $8,486 per kW, which is essentially the average for the category.

- In the 25,001–50,000 sq. ft. size range, a data center supporting a 675 kW IT load across 150 racks with a median density of 4.5 kW/rack had a cost/kW of $7,890, 15 percent higher than the average for the size range. Another data center in that range supported 1.13 MW across 133 racks with a median density of 8.5 kW/rack. In that case, cost/kW dropped to $4,798, 33 percent below the average.

5. Participant Profile

The following table summarizes the frequency of companies and separate data centers participating in the benchmark study. As reported, a total of 13 industries are represented in the sample.

Our final sample includes a total of 34 separate organizations representing 41 data centers. A total of three data centers were removed from the final sample for incomplete responses to our survey instrument.

Table 3: Breakdown of participants by industry

Industries represented	Companies	Data Centers
Colocation	2	8
Communications	1	1
Consumer products	1	1
eCommerce	5	6
Education	1	1
Financial services	5	8
Healthcare	3	3
Hospitality	1	1
Industrial	3	3
Public sector	2	2
Research	1	1
Retail	3	3
Services	3	3
Total	31	41

The following pie chart summarizes participating data center size according to total square footage. The average size of data centers in this study is 18,201 square feet.

Figure 6: Distribution of participating data centers by square footage
Computed from 41 benchmarked data centers

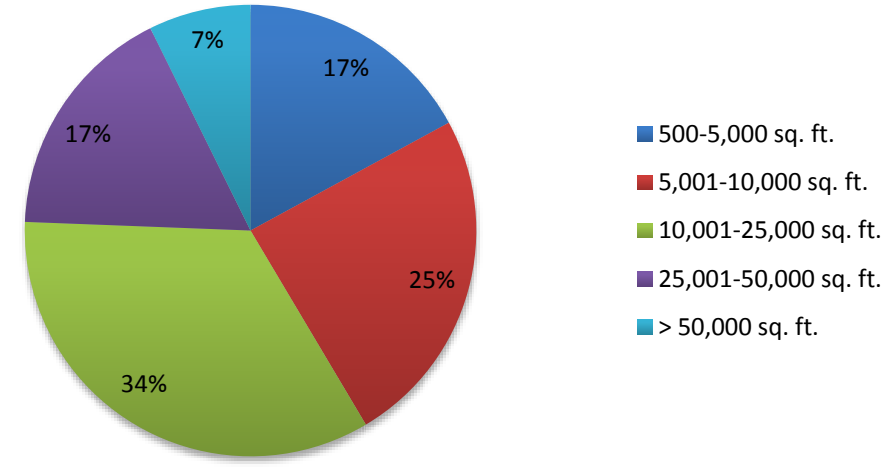

Figure 7 summarizes the sample of participating companies' data centers according to 13 primary industry classifications. As can be seen, co-location and financial services are the two largest industry segments, each representing 19 percent of the sample. Financial service companies include retail banking, payment processors, insurance, brokerage and investment management companies.

Figure 7: Distribution of data centers by industry segment
Computed from 41 benchmarked data centers

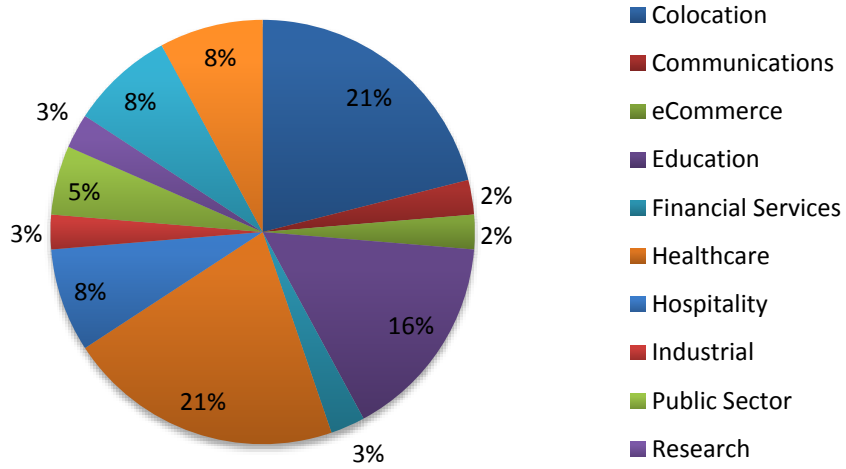

Figure 8 reports the percentage frequency of companies based on their geographic location according to six regions in the United States. The Northeast and Pacific-West represent the largest U.S. regions (both at 23 percent). The smallest regions are the Southwest and Southeast (both at 12 percent).

Figure 8: Distribution of data centers by U.S. geographic region.
Computed from 41 benchmarked data centers

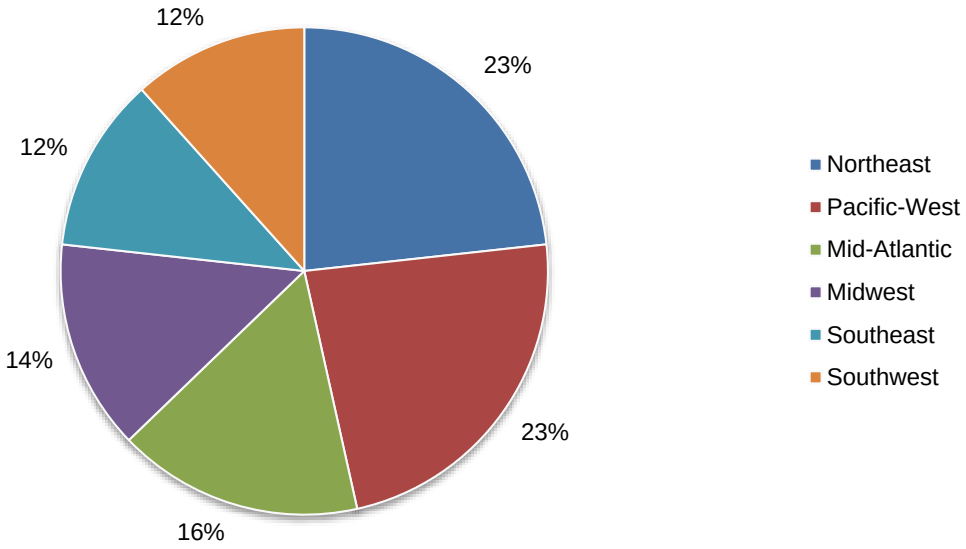

6. Caveats

This study utilizes a confidential and proprietary benchmark method that has been successfully deployed in earlier Ponemon Institute research. However, there are inherent limitations to benchmark research that need to be carefully considered before drawing conclusions from findings.

- **Non-statistical results**: The current study draws upon a representative, non-statistical sample of data centers, all US-based entities. Statistical inferences, margins of error and confidence intervals cannot be applied to these data given the nature of our sampling plan.
- **Non-response**: The current findings are based on a small representative sample of completed case studies. An initial mailing of benchmark surveys was sent to a proprietary group of organizations, resulting in 41 usable data center surveys. Non-response bias was not tested so it is always possible companies that did not participate are substantially different in terms of data center costs and other performance indicators.
- **Sampling-frame bias**: Because our sampling frame is judgmental, the quality of results is influenced by the degree to which the frame is representative of the population of companies and data centers being studied. It is our belief that the current sampling frame is biased toward companies with more mature data center operations.
- **Company-specific information**: The benchmark information is sensitive and confidential. Thus, the current instrument does not capture company-identifying information. It also allows individuals to use categorical response variables to disclose demographic information about the company and industry category.
- **Unmeasured factors**: To keep the benchmark survey concise and focused, we decided to omit other important variables from our analyses such as leading trends and organizational characteristics. The extent to which omitted variables might explain benchmark results cannot be estimated.
- **Extrapolated cost results**: The quality of survey research is based upon the integrity of confidential responses received from benchmarked organizations. While certain checks and balances can be incorporated into the survey process, there is always the possibility that respondents did not provide truthful responses. In addition, the use of a cost estimation technique (termed shadow costing methods) rather than actual cost data could create significant bias in presented results.

Please contact research@ponemon.org **or call us at 800.877.3118 if you have any questions.**

Ponemon Institute

Advancing Responsible Information Management

Ponemon Institute is dedicated to independent research and education that advances responsible information and privacy management practices within business and government. Our mission is to conduct high quality, empirical studies on critical issues affecting the management and security of sensitive information about people and organizations.

As a member of the **Council of American Survey Research Organizations (CASRO),** we uphold strict data confidentiality, privacy and ethical research standards. We do not collect any personally identifiable information from individuals (or company identifiable information in our business research). Furthermore, we have strict quality standards to ensure that subjects are not asked extraneous, irrelevant or improper questions.