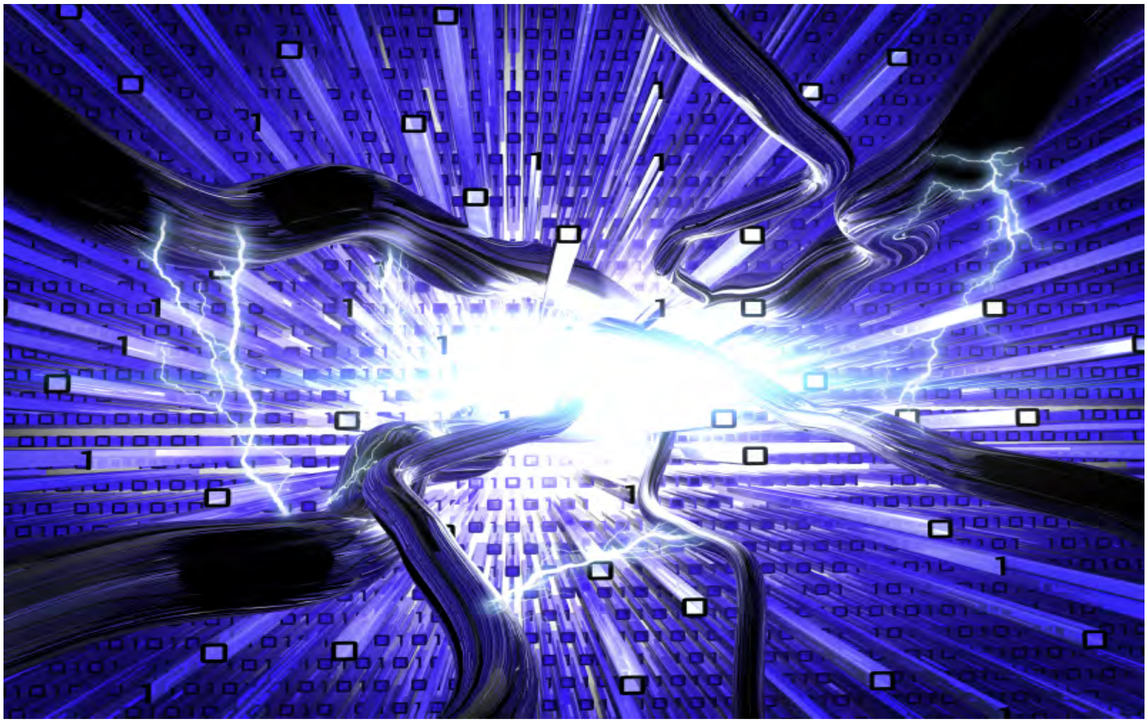

Cost of Data Center Outages

January 2016

Data Center Performance Benchmark Series

Sponsored by Emerson Network Power

Independently conducted by Ponemon Institute LLC

Publication Date: January 2016

Cost of Data Center Outages

Ponemon Institute, January 2016

Table of Contents

Cost of Data Center Outages

Ponemon Institute, January 2016

Part 1. Executive Summary

Ponemon Institute and Emerson Network Power are pleased to present the results of the latest *Cost of Data Center Outages* study*.* Previously published in 2010 and 2013*,* the purpose of this third study is to continue to analyze the cost behavior of unplanned data center outages. According to our new study, the average cost of a data center outage has steadily increased from $505,502 in 2010 to $740,357 today (or a 38 percent net change).

Our benchmark analysis focuses on representative samples of organizations in different industry sectors that experienced at least one complete or partial unplanned data center outage during the past 12 months. Utilizing activity-based costing methods, this year's analysis is derived from 63 data centers located in the United States. Following are the functional leaders within each organization who participated in the benchmarking process:

- Facility manager
- Chief information officer
- Data center management
- Chief information security officer
- IT operations management
- IT compliance & audit
- Operations & engineering

Utilizing activity-based costing, our methods capture information about both direct and indirect costs, including but not limited to the following areas:

- Damage to mission-critical data
- Impact of downtime on organizational productivity
- Damages to equipment and other assets
- Cost to detect and remediate systems and core business processes
- Legal and regulatory impact, including litigation defense cost
- Lost confidence and trust among key stakeholders
- Diminishment of marketplace brand and reputation

Following are some of the key findings of our benchmark research involving the 63 data centers:

- The average cost of a data center outage rose from $690,204 in 2013 to $740,357 in this report, a 7 percent increase. The cost of downtime has increased 38 percent since the first study in 2010.
- Downtime costs for the most data center-dependent businesses are rising faster than average.
- Maximum downtime costs increased 32 percent since 2013 and 81 percent since 2010. Maximum downtime costs for 2016 are $2,409,991.
- UPS system failure continues to be the number one cause of unplanned data center outages, accounting for one-quarter of all such events.
- Cybercrime represents the fastest growing cause of data center outages, rising from 2 percent of outages in 2010 to 18 percent in 2013 to 22 percent in the latest study.

Part 2. Cost Framework

Utilizing activity-based costing, our study addresses eight core process-related activities that drive a range of expenditures associated with a company's response to a data center outage. The activities and cost centers used in our analysis are defined as follows:

- **Detection cost:** Activities associated with the initial discovery and subsequent investigation of the partial or complete outage incident.
- **Containment cost**: Activities and associated costs that enable a company to reasonably prevent an outage from spreading, worsening or causing greater disruption.
- **Recovery cost**: Activities and associated costs that relate to bringing the organization's networks and core systems back to a state of readiness.
- **Ex-post response cost**: All after-the-fact incidental costs associated with business disruption and recovery.
- **Equipment cost**: The cost of new equipment purchases and repairs, including refurbishment.
- **IT productivity loss**: The lost time and related expenses associated with IT personnel downtime.
- **User productivity loss**: The lost time and related expenses associated with end-user downtime.
- **Third-party cost**: The cost of contractors, consultants, auditors and other specialists engaged to help resolve unplanned outages.

In addition to the above process-related activities, most companies experience opportunity costs associated with the data center outage, which results in lost revenue, business disruption and average contribution. Accordingly, our cost framework includes the following categories:

- **Lost revenues**: The total revenue loss from customers and potential customers because of their inability to access core systems during the outage period.
- **Business disruption (consequences)**: The total economic loss of the outage, including reputational damages, customer churn and lost business opportunities.

Figure 1 presents the activity-based costing framework used in this research, which consists of 10 discernible categories. As shown, the four internal activities or cost centers include detection, containment, recovery and ex-post response. Each activity generates direct, indirect and opportunity costs, respectively. The consequence of the unplanned data center outage includes equipment repair or replacement, IT productivity loss, end-user productivity loss, third parties (such as consultants), lost revenues and the overall disruption to core business processes. Taken together, we then infer the cost of an unplanned data center outage.

Figure 1: Activity-Based Cost Account Framework

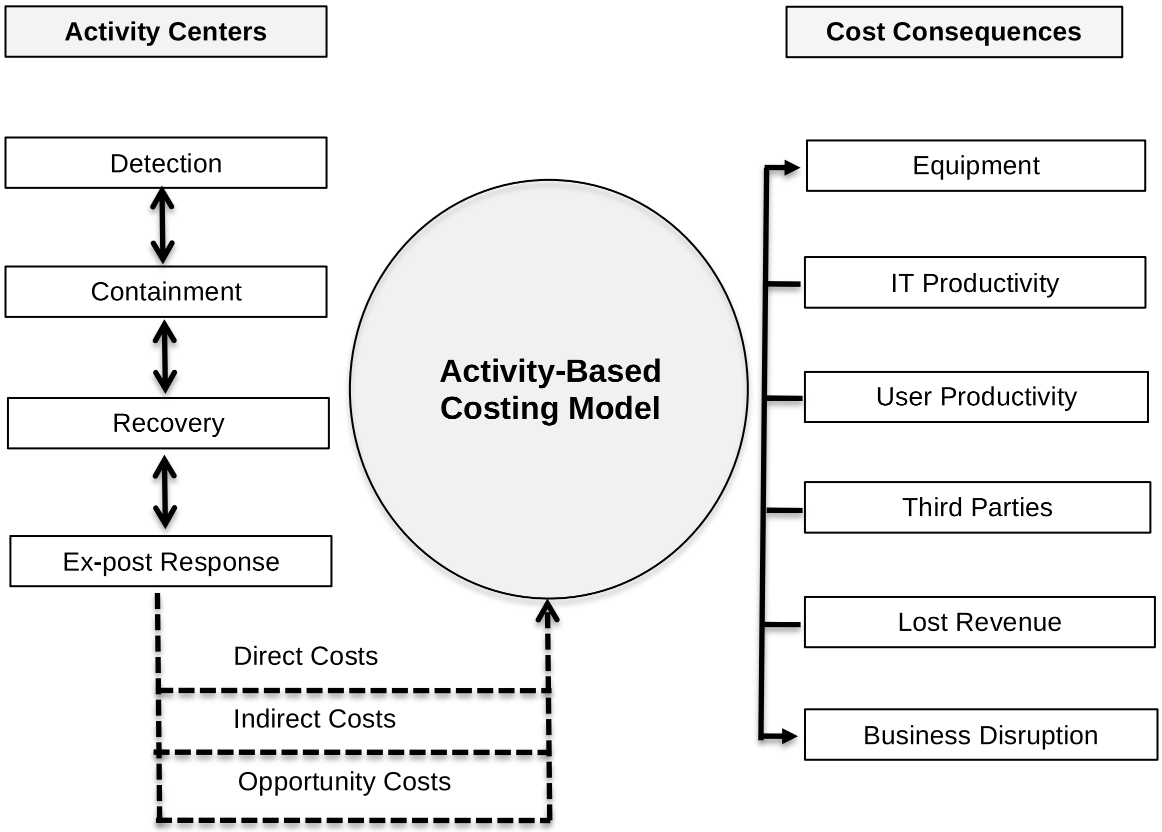

Part 3. Benchmark Methods

Our benchmark instrument was designed to collect descriptive information from IT practitioners and managers of data center facilities about the costs incurred either directly or indirectly as a result of unplanned outages. The survey design relies upon a shadow costing method used in applied economic research. This method does not require subjects to provide actual accounting results, but instead relies on broad estimates based on the experience of individuals within participating organizations.

The benchmark framework in Figure 1 presents the two separate cost streams used to measure the total cost of an unplanned outage for each participating organization. These two cost streams pertain to internal activities and the external consequences experienced by organizations during or after experiencing an incident. Our benchmark methodology contains questions designed to elicit the actual experiences and consequences of each incident. This cost study is unique in addressing the core systems and business process-related activities that drive a range of expenditures associated with a company's incident management response.

Within each category, cost estimation is a two-stage process. First, the survey requires individuals to provide direct cost estimates for each cost category by checking a range variable. A range variable is used rather than a point estimate to preserve confidentiality (in order to ensure a higher response rate). Second, the survey requires participants to provide a second estimate for both indirect cost and opportunity cost, separately. These estimates are calculated based on the relative magnitude of these costs in comparison to a direct cost within a given category. Finally, we conduct a follow-up interview to obtain additional facts, including estimated revenue losses as a result of the outage.

The size and scope of survey items is limited to known cost categories that cut across different industry sectors. In our experience, a survey focusing on process yields a higher response rate and better quality of results. We also use a paper instrument, rather than an electronic survey, to provide greater assurances of confidentiality.

In total, the benchmark instrument contains descriptive costs for each one of the five cost activity centers. Within each cost activity center, the survey requires respondents to estimate the cost range to signify direct cost, indirect cost and opportunity cost. These are defined as follows:

- Direct cost – the direct expense outlay to accomplish a given activity.
- Indirect cost – the amount of time, effort and other organizational resources spent, but not as a direct cash outlay.
- Opportunity cost – the cost resulting from lost business opportunities as a consequence of reputation diminishment after the outage.

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

To keep the benchmark instrument to a manageable size, we carefully limit items to only those cost activities we consider crucial to the measurement of data center outage costs. Based on discussions with learned experts, the final set of items focuses on a finite set of direct or indirect cost activities. After collecting benchmark information, each instrument is examined carefully for consistency and completeness. In this study, four companies were rejected because of incomplete, inconsistent or blank responses.

The study was launched in June 2015 and fieldwork concluded in October 2015. Recruitment started with a personalized letter and a follow-up phone call to 631 US-based organizations for possible participation in our study. All of these organizations are members of Ponemon Institute's benchmark community. This resulted in 78 organizations agreeing to participate. Forty-nine organizations (63 separate data centers) permitted researchers to complete the benchmark analysis.[1]

Six cases were removed for reliability concerns. Utilizing activity-based costing methods, we captured cost estimates using a standardized instrument for direct and indirect cost categories. Specifically, labor (productivity) and overhead costs were allocated to four internal activity centers and these flow through to six cost consequence categories (see Figure 1).

Total costs were then allocated to only one (the most recent) data center outage experienced by each organization. We collected information over approximately the same time frame; hence, this limits our ability to gauge seasonal variation on the total cost of an unplanned data center outage.

[1]The Ponemon Institute's benchmark community is comprised of organizations that have participated in one or more research studies over the past 14 years.

Part 4. Sample of Participating Companies & Data Centers

The following table summarizes the frequency of companies and separate data centers participating in the benchmark study. As reported, a total of 16 industries are represented in the sample.

Our final sample includes a total of 49 separate organizations representing 69 data centers – which is the unit of analysis. A total of six data center organizations were rejected from the final sample for incomplete responses to our survey instrument, thus resulting in a final sample of 63 separate data centers.

Table 1: Tally of cost studies by recruited company & data center

Industries	Companies	Data centers	Rejected	Final sample
Co-location	4	11	1	10
Communications	1	2	0	2
Consumer products	2	2	0	2
eCommerce	7	7	1	6
Education	2	3	0	3
Financial services	6	11	2	9
Healthcare	5	5	0	5
Hospitality	1	1	0	1
Industrial	4	6	0	6
Media	1	2	0	2
Public sector	4	5	1	4
Research	1	1	0	1
Retail	5	5	0	5
Services	5	7	1	6
Transportation	1	1	0	1
Total	49	69	6	63

The following table summarizes participating data center size according to total square footage and the duration of both partial and complete unplanned outages. The average size of the data center in this study is 14,090 square feet and the average outage duration is 95 minutes.

Table 2: Key statistics on data center size and duration of the outage

Key stats	Data center square footage	Duration in minutes
Average	14,090	95
Maximum	55,000	415
Minimum	1,505	13

Pie Chart 1 summarizes the sample of participating companies' data centers according to 15 primary industry classifications. As can be seen, ecommerce and financial services are the two largest industry segments representing 15 percent and 13 percent of the sample, respectively. Financial service companies include retail banking, payment processors, insurance, brokerage and investment management companies.

Pie Chart 1: Distribution of participating organizations by industry segment
Computed from 63 benchmarked data centers

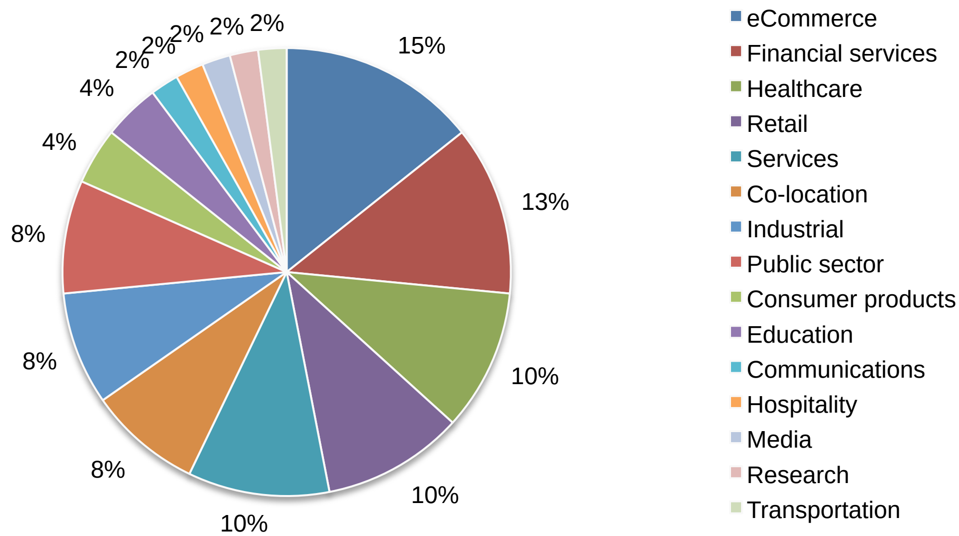

Pie Chart 2 reports the percentage frequency of companies based on their geographic location according to six regions in the United States. The Northeast represents the largest region (at 24 percent) and the smallest region is the Southwest (at 11 percent).

Pie Chart 2: Distribution of participating organizations by US geographic region
Computed from 63 benchmarked data centers

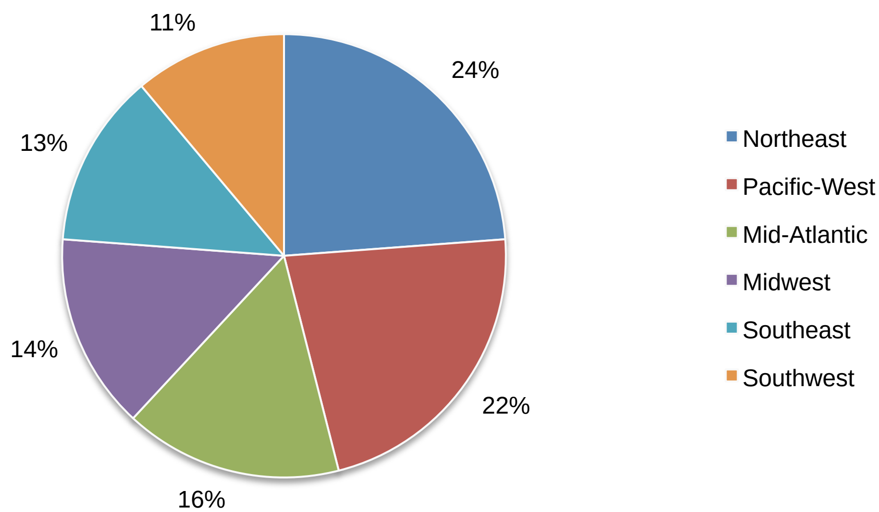

Part 5. Key Findings

Bar Chart 1 reports key statistics on the cost of unplanned outages as reported in 2010, 2013 and 2016. The cost range, median and mean all consistently show substantial increases in cost over time. As shown, the maximum cost has more than doubled over six years from just over $1 million to $2.4 million (a 34 percent increase since the last study). Both mean and median costs increased since 2010 with net changes of 38 and 24 percent respectively. Even though the minimum data center outage cost decreased between 2013 and 2016, this statistic increased significantly over six years, with a net change of 58 percent.

Bar Chart 1: Key statistics on data center outages
Comparison of 2010, 2013 and 2016 results

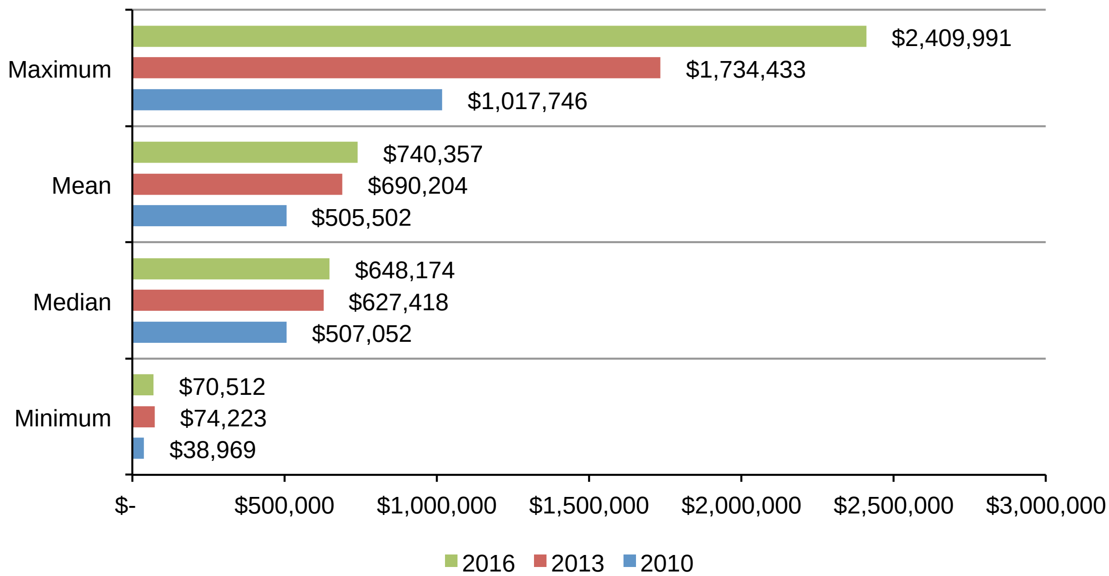

Bar Chart 2 reports the cost structure on a percentage basis for all cost activities for FY 2010, 2013 and 2016. As shown, the cost mix has remained stable over the past five years. Indirect cost represents about half and opportunity loss represents 12 percent of total cost of outages.

Bar Chart 2: Percentage cost structure of unplanned data center outages
Comparison of 2010, 2013 and 2016 results

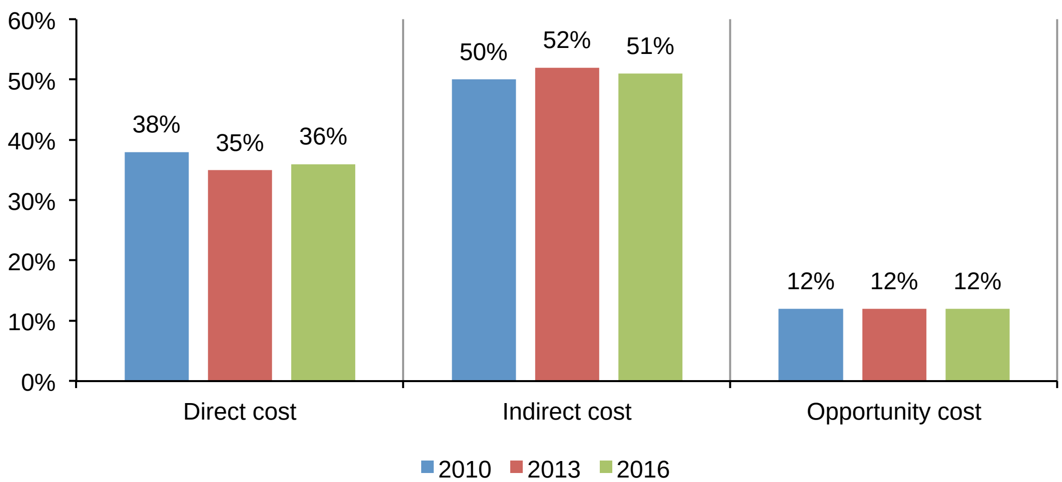

Table 3 summarizes the cost of unplanned outages for all 63 data centers. Please note that cost statistics are derived from the analysis of one unplanned outage incident.

Table 3: Cost summary for unplanned outages

Cost categories	Total	Mean	Median	Minimum	Maximum
Third parties	625,401	9,927	6,970	1,551	27,600
Equipment	597,114	9,478	8,865	1,249	67,783
Ex-post activities	530,964	8,428	11,566	-	36,575
Recovery	1,334,151	21,177	17,570	1,900	58,171
Detection	1,682,856	26,712	22,813	877	69,100
IT productivity	3,898,440	61,880	56,789	6,994	125,600
End-user productivity	8,706,159	138,193	124,551	15,600	456,912
Lost revenue	13,141,737	208,599	197,500	26,591	755,810
Business disruption	16,125,669	255,963	201,550	15,750	812,440
Total cost	46,642,491	740,357	648,174	70,512	2,409,991

Bar Chart 3 reveals a relatively consistent pattern across nine cost categories over five years (three studies). The cost associated with business disruption, which includes reputation damages and customer churn, represents the most expensive cost category. Least expensive involves the engagement of third parties such as consultants to aid in the resolution of the incident.

Bar Chart 3: Comparison of activity cost categories
Comparison of 2010, 2013 and 2016 results
$1,000 omitted

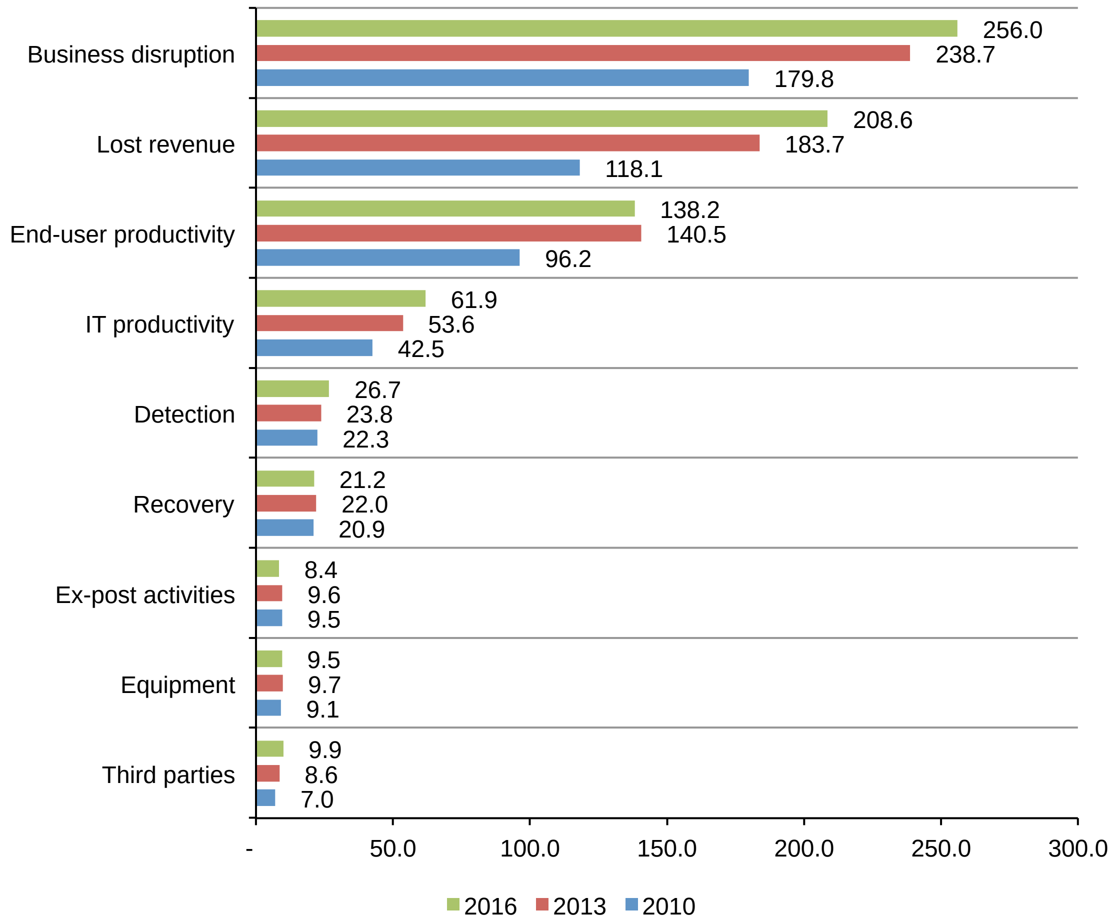

Bar Chart 4 provides the total cost of unplanned outages for the16 industry segments included in our benchmark sample. Analysis by industry is limited because of a small sample size; however, it is interesting to see wide variation across segments ranging from a high of $994,000 (financial services) to a low of $476,000 (public sector).[2]

Bar Chart 4: Distribution of total cost for 15 industry segments
$1,000 omitted

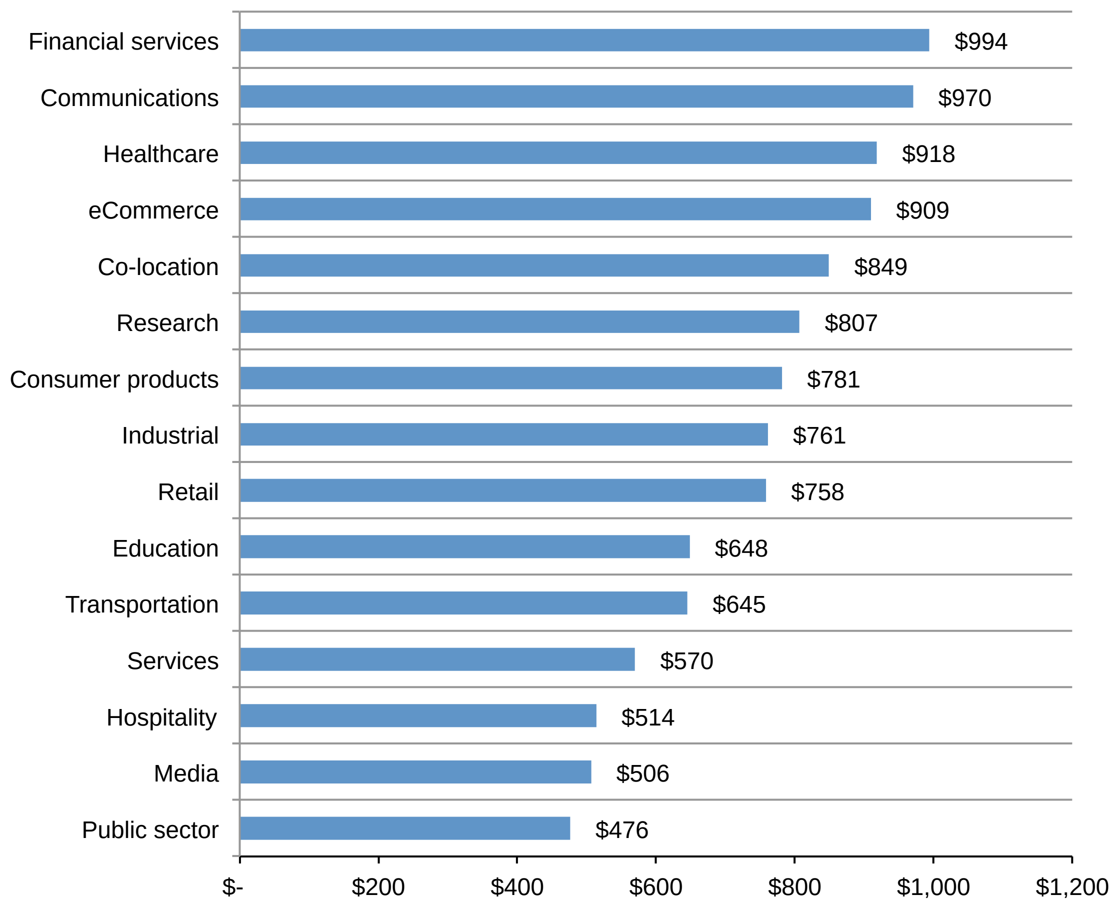

[2] Small sample segments limit our ability to draw inferences about industry differences.

Bar Chart 5 compares costs for partial unplanned outages and complete unplanned outages. All three studies show complete outages are more than twice as expensive as partial outages.

Bar Chart 5: Cost for partial and total shutdown
Comparison of 2010, 2013 and 2016 results

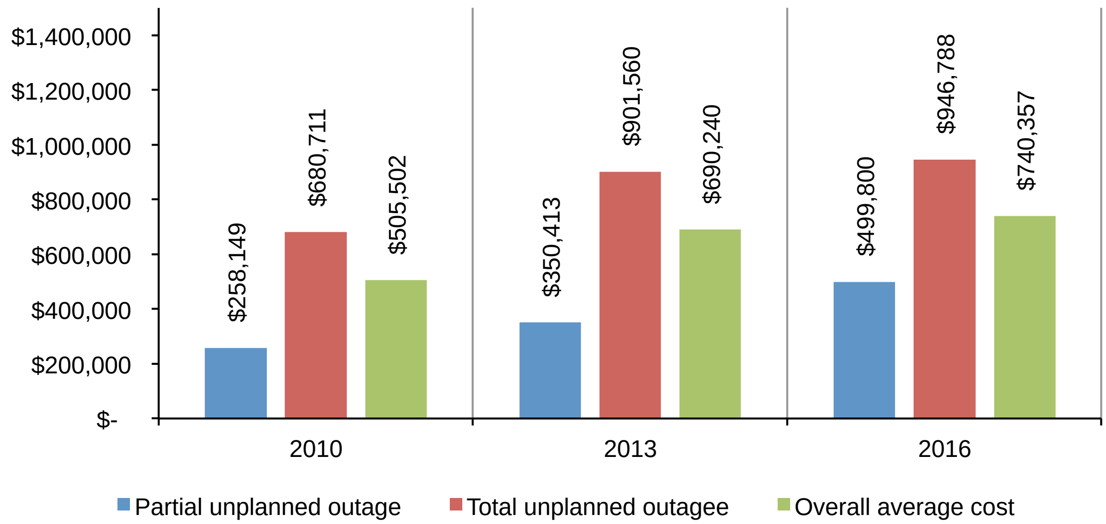

Bar Chart 6 compares the average duration (minutes) of the event for partial and complete outages. In this year's study, complete unplanned outages, on average, last 66 minutes longer than partial outages. It is also interesting to note a U-shape relationship in duration over six years – wherein unplanned outages on average decreased 15 minutes between 2010 and 2013 and increased 11 minutes between 2013 and 2016. One possible reason for the increase in duration time is the rise in cyber attacks, which are difficult to detect and contain.

Bar Chart 6: Duration for partial and total shutdown (measured in minutes)
Comparison of 2010, 2013 and 2016 results

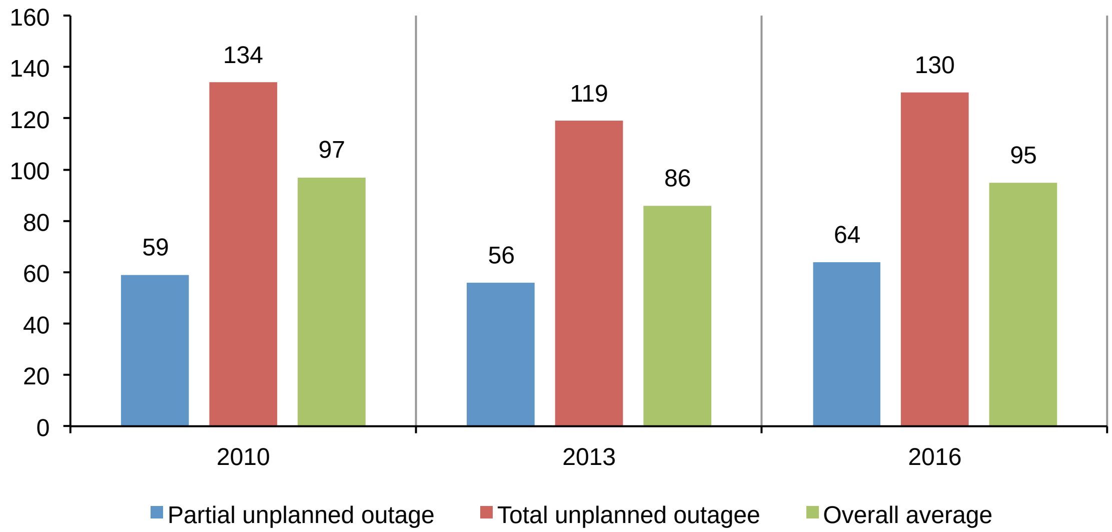

Graph 1 shows the relationship between outage cost and duration of the incident. The graph is organized in ascending order by duration of the outage in minutes. Accordingly, observation 1 has the shortest duration and observation 63 has the longest duration. The regression line is derived from the analysis of all 63 data centers. Clearly, these results show that the cost of outage is linearly related to the duration of the outage**.**

Graph 1: Relationship between cost and duration of unplanned outages
Minutes of down time

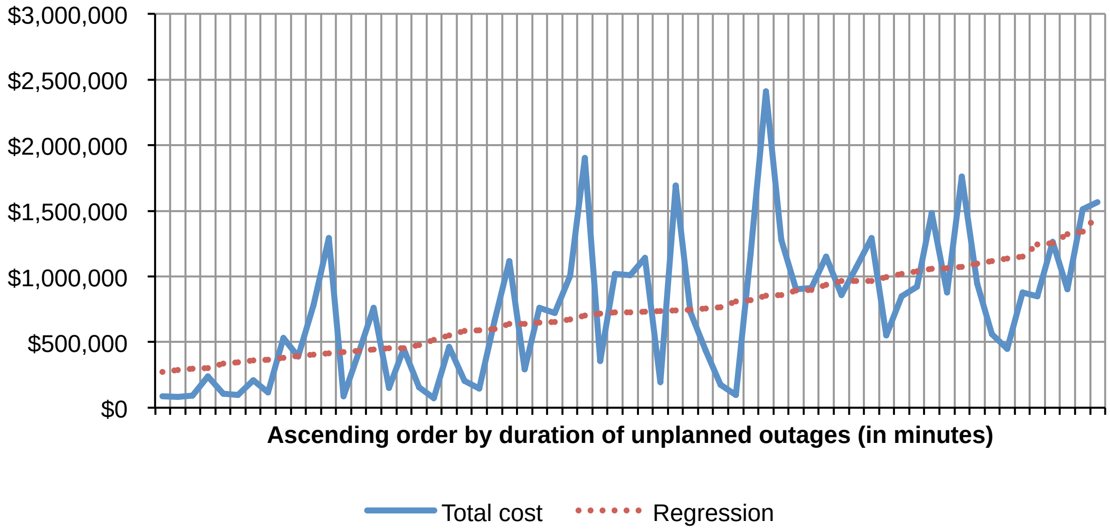

Bar Chart 7 reports the minimum, median, mean and maximum cost per minute of unplanned outages computed from 63 data centers. This chart shows that the most expensive cost of an unplanned outage is over $17,000 per minute. On average, the cost of an unplanned outage per minute is nearly $9,000 per incident.

Bar Chart 7: Total cost per minute of an unplanned outage
Comparison of 2010, 2013 and 2016

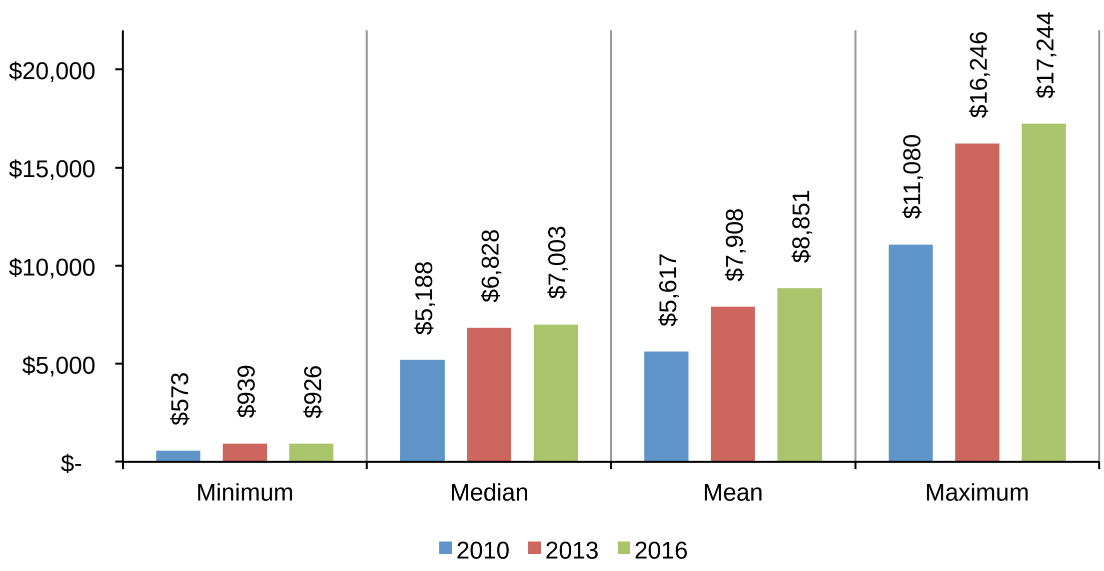

Graph 2 shows the relationship between data center size as measured by square footage and the total cost of unplanned outages. The regression line is computed from the analysis of all 63 data centers. Similar to the duration analysis in Graph 1, these results show that the cost of outage is linearly related to the size of the data center.

Graph 2: Relationship between cost and data center square footage

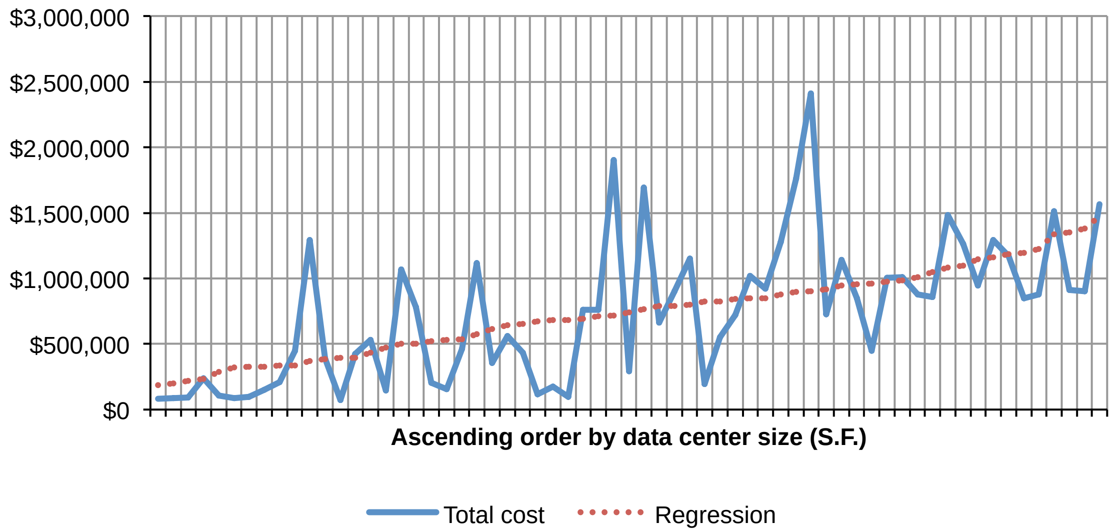

Bar Chart 8 reports the mean cost per square foot of unplanned outages based on all 63 data centers according to quartile. This chart shows that the most expensive cost of an unplanned outage is $99 per square foot for the smallest quartile of companies. The lowest average is $50 for larger organizations.

Bar Chart 8: Unplanned outage cost per square foot of data center by quartile
Average data center SF is reported for each quartile

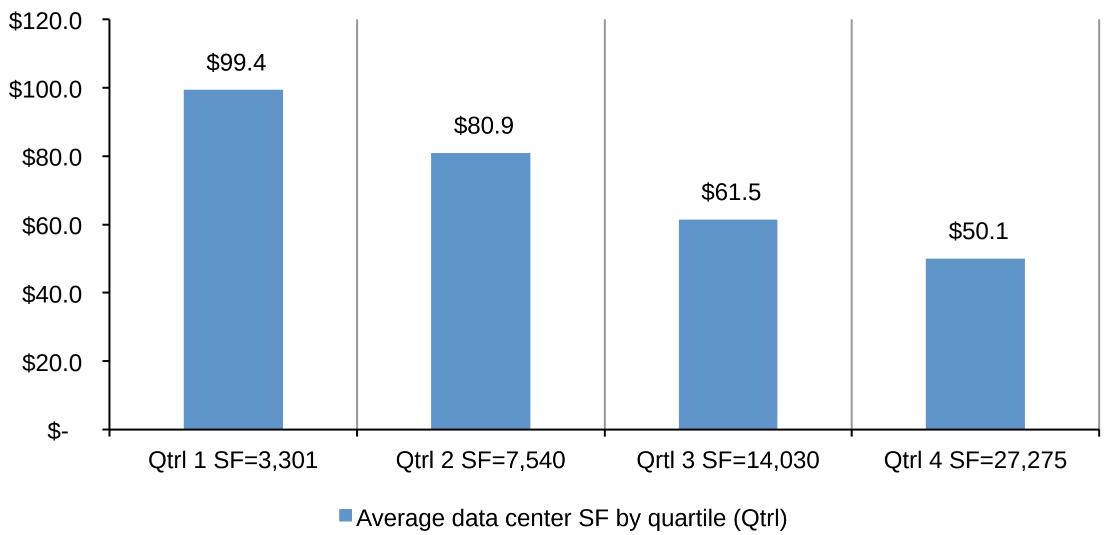

Bar Chart 9 analyzes the sample of 63 data centers by the primary root cause of unplanned outage. The “other” category refers to incidents where the root cause could not be determined. As shown, 25 percent of companies cite UPS system failure as the primary root cause of the incident. Twenty-two percent cite accidental or human error and cyber attack as the primary root causes of the outage.

For human error, this is the same percentage as found in 2013, indicating no progress in reducing what should be an avoidable cause of downtime. The 22 percent figure for cyber crime represents a 20 percent increase from 2013 and a 167 percent increase from 2010. IT equipment failure represents only four percent of all outages studied in this research.

Bar Chart 9: Root causes of unplanned outages
Comparison of 2010, 2013 and 2016 results

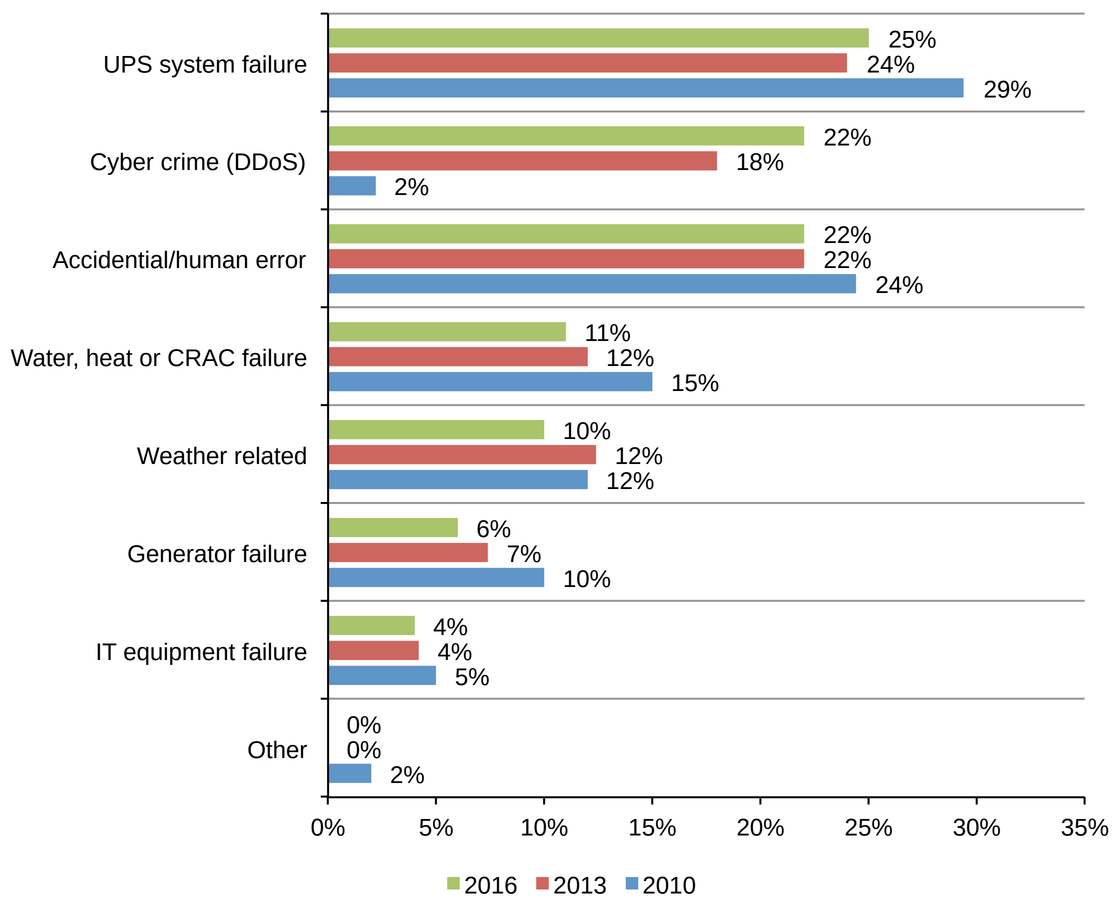

Bar Chart 10 reports the average cost of outage by primary root cause of the incident. As shown below, IT equipment failures result in the highest outage cost, followed by cyber crime. The least expensive root cause appears to be related to weather followed by accidental/ human errors.

Bar Chart 10: Total cost by primary root causes of unplanned outages
Comparison of 2010, 2013 and 2016 results
$1,000 omitted

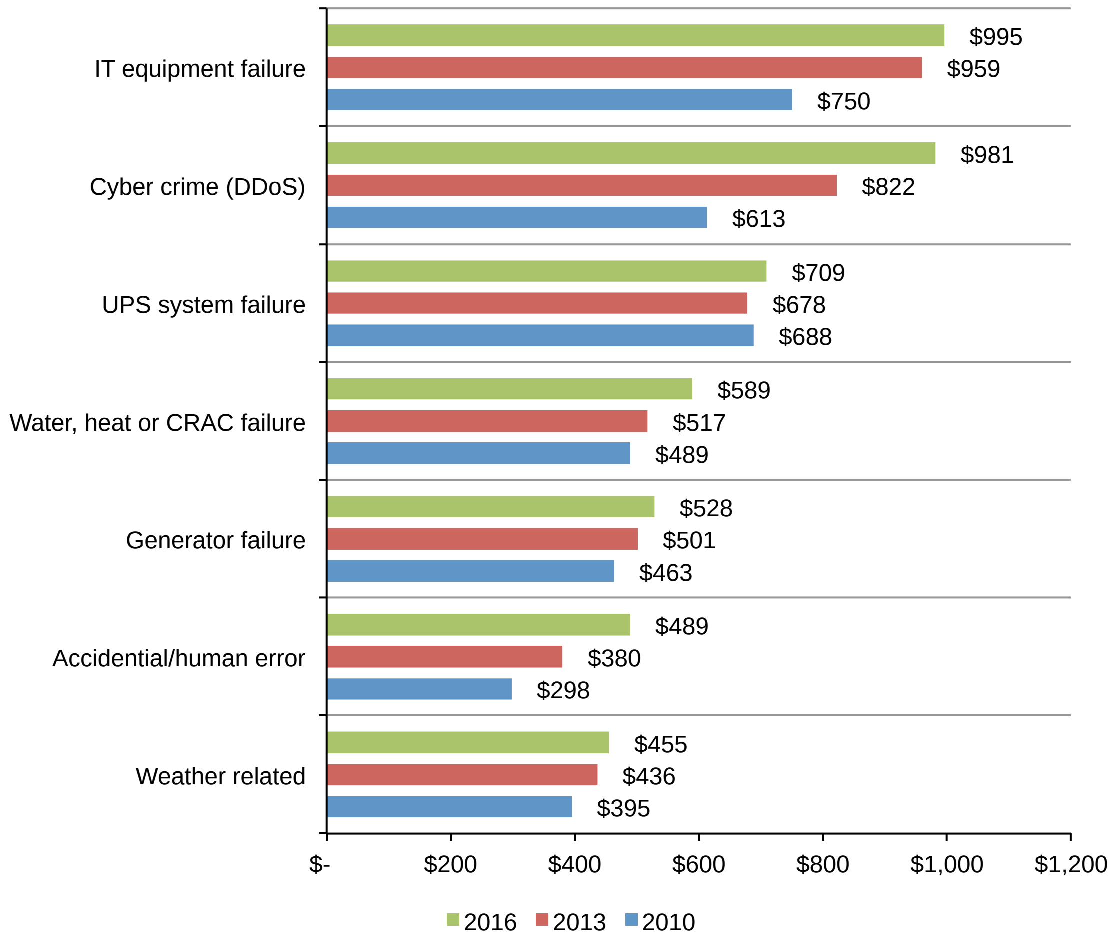

Part 6. Conclusion

The *2016 Cost of Data Center Outages* study builds upon the previous two studies to provide historical perspective on the impact of data center downtime on businesses in various industries. The findings reveal incremental, and somewhat predicable, increases in the cost of downtime. Consequently, it can be easy to underestimate the change the industry has endured since 2010.

When the first *Cost of Data Center Outages* study was published in 2010, Facebook reached 500 million active users. Today the number has grown to 1.5 billion, necessitating a massive data center development program.

Similar rates of growth have occurred in other data center dependent businesses. According to research conducted by BI Intelligence, less than 500,000 smartphones were shipped globally in 2011. That number was expected to be over 1.5 billion for 2015, with 64 percent of US adults now owning a smartphone.

Cloud computing is in the midst of a similar growth spurt today. Goldman Sachs projects a 30 percent CAGR between 2013 and 2018. The Internet of Things will likely drive the next wave of growth. Specifically, IDC predicts the global IoT market will grow to $1.7 trillion in 2020 from $655.8 billion in 2014.

All of these developments mean more data flowing across the Internet and through data centers—and more opportunities for businesses to use technology to grow revenue and improve business performance. The data center will be central to leveraging those opportunities.

As organizations like Facebook invest millions in data center development, they are exploring new approaches to data center design and management to both increase agility and reduce the cost of downtime. However, as this study shows, costs continue to rise and the causes of data center downtime are very similar today to what they were six years ago—with the notable exception of the dramatic rise in cyber attacks, which represents a major challenge for data center operators in the coming years.

This predictable historical pattern highlights the reality that, while there has been a need to move quickly to support the dramatic changes that occurred in social media, mobile devices and cloud services, the lifespan of a typical data center is relatively long and will extend across multiple significant cultural and technology changes. That means data centers designed before anyone was talking about the Internet of Things will be expected to adapt and play a critical role in supporting the Internet of Things.

For this most recent study, we are seeing trends in cost and causes of downtime that reflect a relatively stable industry dealing with many of the same issues as in 2010. But we also understand that there is a strong undertow of change occurring in the industry and we expect this change to begin to surface in future *Cost of Data Center Outages* research.

Part 7. 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 purpose of this study is descriptive rather than normative inference. The current study draws upon a representative, non-statistical sample of data centers, all US-based entities experiencing at least one unplanned outage over the past 12 months. 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 benchmark group of over 600 organizations, all believed to have experienced one or more outages over the past 12 months. Sixty-three data centers provided usable benchmark surveys. Non-response bias was not tested so it is always possible companies that did not participate are substantially different in terms of the methods used to manage the detection, containment and recovery process, as well as the underlying costs involved.

- **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. Industry classification relies on self-reported results.

- **Unmeasured factors**: To keep the 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 at this time.

- **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.

If you have questions or comments about this report, please contact us by letter, phone call or email:

Ponemon Institute LLC
Attn: Research Department
2308 US 31 North
Traverse City, Michigan 49686 USA
1.800.877.3118
research@ponemon.org

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.