



Cost of Data Center Outages

Sponsored by Emerson Network Power

Independently conducted by Ponemon Institute^{LLC}

Publication Date: February 1, 2011

Cost of Data Center Outages

Benchmark Study of 41 US Data Centers
Ponemon Institute, February 2011

Part 1. Executive Summary

Ponemon Institute and Emerson Network Power are pleased to present the results of the *Cost of Data Center Outages*. The purpose of this benchmark study is to determine the full economic cost of unplanned data center outages and is the second study in a two-part research series on the topic of data center outages. The first study, *National Study on Data Center Outages*, was released in September 2010 and was conducted to determine the frequency and root causes of unplanned data center outages. We believe both studies are important because of evidence that IT leaders are underestimating the economic impact unplanned outages have on their operations.

The Cost of Data Center Outages study is the first benchmark study that attempts to estimate the full costs associated with an unplanned data center outage. This benchmark analysis focuses on representative samples of organizations in the U.S. that experienced at least one complete or partial unplanned data center outage during the past 12 months. The analysis was based on 41 independent data centers located in the United States. Following are the functional leaders within each organization who participated in the study:

- Facility manager
- Chief information officer
- Data center management
- Chief information security officer
- IT compliance leader

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

Our research indicates data center outages have serious financial consequences for an organization. According to the study, the cost of a data center outage ranges from a minimum cost of \$38,969 to a maximum of \$1,017,746 per organization, with an overall average cost of \$505,502 per incident. Other key findings included:

- Total cost of partial and complete outages can be a significant expense for organizations
- Total cost of outages is systematically related to the duration of the outage
- Total cost of outages is systematically related to the size of the data center
- Certain causes of the outage are more expensive than others. Specifically, IT equipment failure is the most expensive root cause. Accidental/human error is least expensive.

Part 2. Cost Framework

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

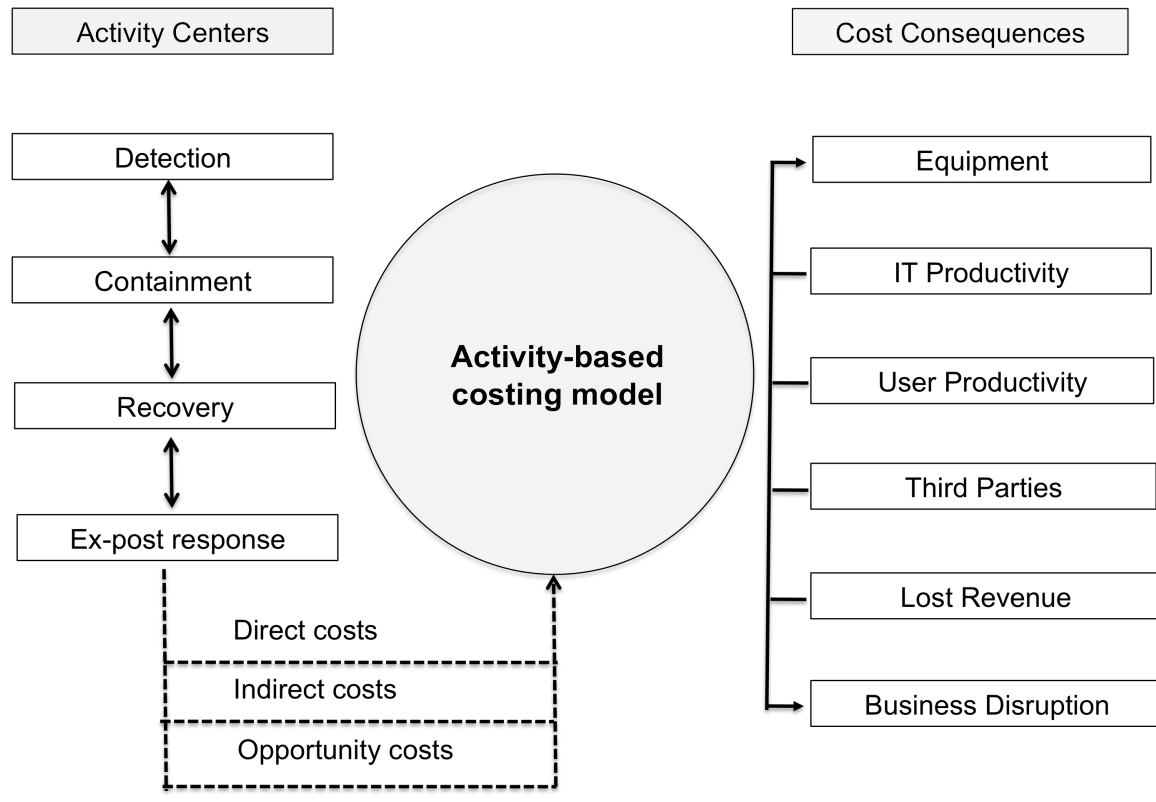
- **Detection cost:** Activities associated with the initial discovery and subsequent investigation of the partial or complete outage incident.
- **Containment cost:** Activities and associated costs that enable a company to reasonably prevent an outage from spreading, worsening or causing greater disruption.
- **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 equipment new 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 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, 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 limited 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 focus 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 July 2010 and fieldwork concluded in October 2010. The recruitment started with a personalized letter and a follow-up phone call to 201 organizations for possible participation in our study. While 53 organizations initially agreed to participate, 41 organizations permitted researchers to complete the benchmark analysis.

Two cases were removed from our final analysis because those data centers fell below the minimum size requirement of 2,500 square-feet. 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.

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, our final sample includes a total of 36 separate organizations representing 45 data centers – which is our primary unit of analysis. A total of four organizations were rejected from the final sample for incomplete responses to our survey instrument, thus resulting in a final sample of 41 data centers.

Table 1: Description of the final sample of participating data centers					
#	Industry classification	Companies	Data Centers	Rejected*	Total
1	Financial services	5	5	0	5
2	E-commerce retail	5	5	1	4
3	Conventional retail	1	2	0	2
4	Technology & software	1	2	0	2
5	Consumer products	3	3	0	3
6	Transportation	1	1	0	1
7	Defense	1	1	0	1
8	Public sector	2	4	1	3
9	Healthcare	5	5	0	5
10	Industrial	2	4	1	3
11	Communications	1	1	0	1
12	Hospitality	1	1	0	1
13	Services	3	3	0	3
14	Education	1	2	0	2
15	Media	1	1	0	1
16	Collocation services	3	5	1	4
	Totals	36	45	4	41

Pie Chart 1 summarizes the sample of participating companies' data centers according to 16 primary industry classifications.

Pie Chart 1: Distribution of participating organizations by industry segment

Computed from 41 benchmarked data centers




As can be seen above, financial services and healthcare companies are the two largest industry segments representing 12 percent of the sample, respectively. Financial services companies include retail banking, insurance, brokerage and credit card companies. The second and third largest segments are e-commerce retailers and collocation services (both at 10 percent).

Bar 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 7 percent).

Pie Chart 2: Distribution of participating organizations by US geographic region

Computed from 41 benchmarked data centers



The following table summarizes participating data center size according to total square footage and the duration of both partial and complete unplanned outages. In total, 59 percent of participating data centers experienced a complete data center outage, and 41 percent experienced a partial outage.

Table 2: Key statistics on data center size and duration of the outage		
Description	Square footage of the data center	Duration in minutes
Average	10,481	102
Maximum	30,165	250
Minimum	2,557	17

Part 5. Key Findings

Bar Chart 1 reports the cost structure on a percentage basis for all cost activities. As shown, indirect cost accounts for half of the total cost of data center outages. Direct cost represents 38 percent and opportunity loss represents 12 percent of total cost of outages.

Bar Chart 1: Percentage cost structure of unplanned data center outages

Computed from 41 benchmarked data centers





Table 3 summarizes the cost of unplanned outages for all 41 data centers. Bar Chart 2 provides a graphic showing the cost for nine categories in descending order. Please note that cost statistics are derived from the analysis of one unplanned outage incident.

Cost categories	Total	Mean	Median	Minimum	Maximum
Detection	916,245	22,347	16,138	519	48,178
Recovery	856,226	20,884	15,899	-	48,178
Ex-post activities	391,015	9,537	10,261	-	26,332
Equipment	371,586	9,063	6,369	357	52,136
IT productivity	1,743,738	42,530	23,861	-	245,090
End-user productivity	3,945,269	96,226	67,904	1,251	599,000
Third parties	287,331	7,008	6,097	-	21,634
Lost revenue	4,841,270	118,080	-	-	755,077
Business disruption	7,372,922	179,827	98,065	-	912,263
Total cost	\$20,725,602	\$505,502	\$507,052	\$38,969	\$1,017,746

Bar Chart 2 reveals significant variation across nine cost categories. 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 2: Average cost of unplanned data center outages for nine categories


Computed from 41 benchmarked data centers



Line graph 1 provides the total cost of unplanned outages for 16 industry segments included in our benchmark sample. The 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 over \$1 million (communications) to a low of \$84,986 (hospitality).

Graph 1: Distribution of total cost for 16 industry segments


Computed from 41 benchmarked data centers



Bar Chart 3a compares costs for partial unplanned outages and complete unplanned outages. As can be seen, complete outages are more than twice as expensive as partial outages. Bar Chart 3b compares the average duration (minutes) of the event for partial and complete outages. As shown, complete unplanned outages, on average, last 75 minutes longer than partial outages.


Bar Chart 3a

Average cost for partial & complete outage



Bar Chart 3b


Duration (minutes) of partial & complete outage



Graph 2 shows the relationship between outage cost and duration of the incident. The graph is organized in descending order by duration of the outage in minutes. Accordingly, observation 1 has the shortest duration and observation 41 has the longest duration. The regression line is derived from the analysis of all 41 data centers. Clearly, these results show that the cost of outage is linearly related to the duration of the outage.

Graph 2: Total cost in ascending order by duration of unplanned outages


Computed from 41 benchmarked data centers



Bar Chart 4 reports the minimum, median, mean and maximum cost per minute of unplanned outages computed from 41 data centers. This chart shows that the most expensive cost of an unplanned outage is over \$11,000 per minute. On average, the cost of an unplanned outage per minute is likely to exceed \$5,000 per incident.

Bar Chart 4: Total cost per minute of an unplanned outage


Computed from 41 benchmarked data centers



Graph 3 shows the relationship between data center size as measured by square footage and the total cost of unplanned outages. Observation 1 has the smallest and observation 41 has the largest data centers in square footage, respectively. The regression line is computed from the analysis of all 41 data centers. Similar to the duration analysis above, these results show that the cost of outage is linearly related to the size of the data center.

Graph 3: Total cost in ascending order by the square footage (size) of the data center


Computed from 41 benchmarked data centers



Bar Chart 5 reports the minimum, median, mean and maximum cost per square foot of unplanned outages based on all 41 data centers. This chart shows that the most expensive cost of an unplanned outage is \$246 per square foot. On average, the cost of an unplanned outage per square foot is likely to exceed \$50 per incident.

Bar Chart 5: Total cost per square foot of an unplanned outage


Computed from 41 benchmarked data centers



Pie Chart 3 groups the sample of 41 data centers by the primary root cause of the unplanned outage. The “other” category refers to incidents where the root cause could not be determined. As shown, 29 percent of companies rate UPS system failure as the primary root cause of the incident. Twenty-four percent rate accidental or human error and 15 percent as water, heat or CRAC failure as the primary root cause of the outage. IT equipment failure represents only five percent of all outages studied in this research.

Pie Chart 3: Primary root causes of unplanned outages


Computed from 41 benchmarked data centers



Bar Chart 6 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 UPS system failures. The least expensive root cause appears to be related to accidental/ human errors.

Bar Chart 6: Total cost by root causes of the unplanned outage

Computed from 41 benchmarked data centers



Part 6. Causes of Data Center Outages & Concluding Thoughts

In the *National Survey on Data Center Outages* report we discussed the primary causes of data center downtime and the direct correlation to the current trends and challenges being faced by today's data centers. Below are four key industry drivers that directly impact availability and brief explanations of how they cause downtime.

Increasing IT Demands/Exceeding Data Center Capacity. As demand for IT applications grows and more IT equipment is added to the data center, the supporting IT infrastructure must grow as well. If the IT demand outgrows what the critical infrastructure can support, downtime often is the result. Downtime root cause correlation: UPS Capacity exceeded, IT equipment failure (often tied to thermal events) & PDU/circuit breaker failure.

Rising Rack Densities. With the introduction of blade servers and other high-performance IT equipment, the typical server rack will contain well over 10 kW of IT. High heat densities will require precision cooling closer to the server. However, depending on the cooling design, this also could bring water closer to the server. Downtime root cause correlation: Water incursion, heat-related/CRAC failure and IT equipment failure.

Data Center Efficiency. Data centers consume a lot of power—estimates put it at about 2 percent of global energy consumption—and many managers are evaluating high-efficiency power and cooling technologies that provide cost reductions but may not provide the highest reliability or ideal operating environment. Our analysis of the cost of downtime clearly shows that, while efficiency is important, the risk is too great to chase marginal efficiency gains at the expense of availability, especially in critical data centers. Downtime root cause correlation: UPS failure, heat-related/CRAC failure and IT equipment failure.

Need for Infrastructure Management and Control. The data center manager's requirements of improving availability, increasing efficiency and planning for capacity all can be addressed through infrastructure management. Monitoring the float charge of a battery, knowing optimal placement of a new server to even having a people-free facility with remote resolution all are aspects of successful infrastructure management. Downtime root cause correlation: UPS failure and battery failure, heat-related/CRAC failure, IT equipment failure and accidental EPO/human error.

The findings of our research suggest unplanned data center outages present a difficult and costly challenge for organizations. In general, failed equipment, data center mishaps and insufficient resources exacerbate the frequency and duration of unplanned outages.

The challenge for data center management is to communicate effectively to senior leadership the urgent need to implement power, cooling and monitoring systems that increase availability and ensure the performance of mission critical applications.

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 U.S.-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 reference group of over 400 separate organizations, all believed to have experienced one or more outages over the past 12 months. Forty-one 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.
- **Estimated cost results.** The quality of survey research is based on 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.

Appendix 1: Summarized cost data for 41 benchmarked organizations

The following table summarizes the total cost of unplanned outages for 41 data centers. The activity cost column summarizes detection, containment, recovery, and ex-post response costs. The cost consequences column combines third party, IT productivity, end-user productivity, and business disruption costs.

Data centers	Activity costs	Cost consequences	Total cost	S.F. Size	Duration in minutes
1	1,107	37,862	38,969	4050	68
2	25,992	58,994	84,986	5124	44
3	49,132	85,209	134,341	8876	55
4	22,988	129,634	152,622	4088	150
5	105,010	63,752	168,762	6841	76
6	1,996	174,395	176,391	5881	34
7	43,688	133,210	176,898	3980	59
8	81,015	100,851	181,866	6057	25
9	46,455	142,003	188,458	4300	17
10	6,107	187,955	194,062	8491	18
11	6,219	195,885	202,104	5600	37
12	40,585	129,178	169,763	8343	19
13	42,060	237,621	279,681	16528	36
14	32,209	183,287	215,496	6500	25
15	24,742	211,379	236,121	6400	24
16	75,012	306,049	381,061	10000	93
17	54,890	285,770	340,660	3100	60
18	105,558	894,638	1,000,196	6372	212
19	88,925	269,825	358,750	7425	65
20	40,812	322,424	363,236	5664	72
21	4,008	486,406	490,414	24300	46
22	90,618	426,515	517,133	9876	137
23	100,128	406,924	507,052	2965	51
24	107,879	492,540	600,419	8000	126
25	29,320	593,981	623,301	5670	170
26	90,157	539,638	629,795	2557	200
27	92,036	553,720	645,756	11776	82
28	36,215	616,196	652,411	16500	185
29	76,098	503,892	579,990	13568	66
30	76,119	897,928	974,047	5789	250
31	47,239	710,199	757,438	8300	113
32	16,122	802,880	819,002	13900	173
33	19,601	815,234	834,835	10976	128
34	54,428	918,113	972,541	7500	141
35	28,411	755,706	784,117	18000	154
36	53,975	816,522	870,497	16190	159
37	11,674	954,461	966,135	25000	99
38	49,124	965,896	1,015,020	19560	150
39	79,496	938,250	1,017,746	22110	249
40	98,071	695,872	793,943	30165	178
41	108,265	521,322	629,587	23400	174

Following are the percentage direct, indirect, and opportunity costs for all cost activities using in our activity-based costing framework.

Activity-based cost loadings by activity center and related cost consequences				
Cost activities and cost consequences	Direct cost	Indirect cost	Opportunity cost	Total
Detection cost	52%	48%	0%	100%
Equipment cost	60%	40%	0%	100%
IT productivity loss	23%	77%	0%	100%
User productivity loss	22%	78%	0%	100%
Third-party loss	35%	41%	24%	100%
Recovery cost	22%	78%	0%	100%
Ex-post response costs	53%	47%	0%	100%
Customer turnover	55%	31%	14%	100%
Lost revenue	33%	26%	41%	100%
Reputation and brand loss	24%	30%	45%	100%
Overall contribution	38%	50%	12%	100%

If you have questions or comments about this executive summary or you would like to obtain a full 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.