



UNDERSTANDING THE COST OF DATA CENTER DOWNTIME:

An Analysis of the Financial Impact
on Infrastructure Vulnerability

Executive Summary

Over the course of the past decade, enterprise business has fundamentally changed. Among the many changes experienced, none has been more profound than the increase in reliance on information technology (IT) systems to support business-critical applications. For many of today's enterprises – including banks, telecommunications companies, internet service providers and cloud/co-location facilities – data center throughput has evolved into monetized commodity. No longer simply supporting the internal needs of the organization, data center availability has become essential to many companies whose customers pay a premium for access to a variety of IT applications.

This unprecedented reliance on IT systems has forged an even stronger connection between data center availability and total cost of ownership (TCO). A single downtime event now has the potential to significantly impact the profitability (and, in extreme cases, the viability) of an enterprise. Unfortunately, a severe disconnect exists between IT personnel and their C-suite counterparts with regard to understanding the frequency and the cost of data center downtime.

Recognizing the need to address these misconceptions, Vertiv™ partnered with the Ponemon Institute to conduct two in-depth studies on the perceptions, causes and true monetary costs of data center downtime – totaling thousands of dollars per minute on average – as well as which infrastructure vulnerabilities have the most significant and costly impact on the availability of critical IT systems (see “National Survey on Unplanned Data Center Outages” and “The Cost of Data Center Outages”).

In addition to examining the differing perceptions between the C-Suite and IT staff, this white paper takes a detailed look at the potential “bottom line” costs of data center downtime and examines how power, cooling, monitoring and service inadequacies can contribute to a facility's risk of downtime. It explores specific data center infrastructure vulnerabilities and associated downtime costs, as well as recommendations for fortifying these infrastructures to minimize downtime and achieve the highest possible return on investment (ROI). Finally, it offers a long-term business case for addressing these critical vulnerabilities as well as factors CIOs and IT personnel should consider when prioritizing their actions and investments.

Introduction:

Downtime Perceptions vs. Realities

Since the “dot com” boom (and subsequent bust) of the late 90s and early 2000s, IT networks and data center systems have experienced a resurgence in the central role they play in revenue generation and business growth. From streamlining customer service and networking to facilitating a variety of e-commerce and enterprise IT services, data centers have evolved into business foundations for companies in a wide range of industries. Furthermore, as IT services become increasingly commoditized (via co-location, disaster recovery and cloud computing services), the economic impact of data center operations will continue to grow at an unprecedented rate.

However, even though more enterprises depend on their data centers to support business-critical applications than ever before, significant infrastructure vulnerabilities and misperceptions about the frequency and cost of IT failures have put many companies at increased risk for costly downtime events.

According to a September 2010 Ponemon Institute study commissioned by Vertiv™, misconceptions about the frequency and impact of data center downtime have become commonplace in businesses across the United States. The survey of more than 400 data center and IT operations professionals revealed a widening disconnect in perceptions being perpetuated between the C-suite and “rank-and-file” IT staff:

- Seventy-one percent of senior-level respondents believe their company’s business model is dependent on its data center to generate revenue and/or conduct e-commerce. Only 58 percent of rank-and-file respondents shared this belief.
- Though respondents experienced an average of two downtime events over the two-year period studied (lasting up to 120 minutes apiece, on average), **62 percent of senior-level respondents agreed that unplanned outages did not happen frequently.** Forty-one percent of rank-and-file respondents also agreed with this statement.
- Seventy-five percent of senior-level respondents feel their companies’ senior management fully supports efforts to prevent and manage unplanned outages, while just 31 percent of supervisor-level employees and below agreed with this statement.

- Less than 32 percent of all respondents agreed their company utilizes all best practices to maximize availability of critical IT equipment (40 percent at the executive level; 29 percent at the rank-and-file level).

Based on these findings, it is clear that executive-level respondents are extremely cognizant of the economic importance of their company’s data center operations. This is not surprising, as the core responsibility for senior management and C-level executives (including Chief Information Officers) is to understand how all facets of the business contribute to a company’s growth and performance.

Survey responses also indicated that most of these executives are not as in-tune to the day-to-day data center operations as rank-and-file employees specifically charged with maintaining the company’s IT infrastructure. As such, many of the executives surveyed are not as aware of the frequency of downtime events and the vulnerabilities in their data center infrastructures that are contributing to these events.

Conversely, rank-and-file IT staff are more aware of the frequency of system failures and specific vulnerabilities in their companies’ data center infrastructures than their executive-level counterparts. However, fewer rank-and-file respondents actively acknowledge the role of their companies’ data center operations in generating revenue and/or facilitating e-commerce activity.

On the surface, these findings may appear to be benign examples of how “siloes” work groups can promote disconnects in how common issues are perceived. However, for companies whose profitability is directly tied to the availability of enterprise IT operations, they can lead to dramatic increases in adverse risk for the profitability, and potentially the viability, of a business.

By bridging the perception gap between C-suite executives and rank-and-file IT staff, companies will be better positioned to maximize the availability of critical IT applications without overly inflating a data center’s total cost of ownership. In addition to ensuring the entire organization has an accurate perception of the state of its data center infrastructure, it is critical employees at all levels of the organization have a thorough understanding of the true financial implications of downtime.

These alarming misperceptions about the frequency and impact of data center downtime events triggered the commission of a second study to determine and benchmark the average cost of data center downtime in the United States.

Methodology:
Benchmarking the Cost of Downtime

Data Center Professionals from 41 independent facilities across the country – spanning a variety of organizational responsibilities – were asked to participate in the study. Participating data centers represented a wide variety of industry segments, including financial services, telecommunications, retail (conventional and e-commerce), health care, government and third-party IT services. To ensure that costs were representative of an average enterprise data center, participating data centers were required to have a minimum square-footage of 2,500 ft².

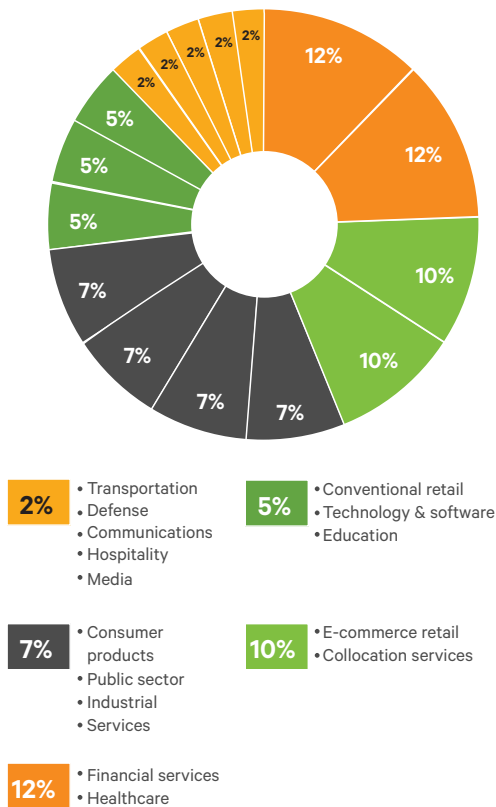


Figure 1: Distribution of participating organizations by industry segment.

Representatives from all levels of the IT staff were asked to participate in the study, including:

- Facility Managers
- Chief Information Officers
- Data Center Management Personnel
- Chief Information Security Officers
- IT Compliance Leaders

To calculate the comprehensive cost of data center downtime, researchers used an activity-based costing model which took into consideration direct, indirect and opportunity costs. As shown in Figure 2, costs were categorized according to internal activity centers and external cost consequences.

Respondents provided direct, indirect and opportunity cost estimates (separately) for a single recent outage based on provided range variables. To ensure reported losses included in the study are as comprehensive as possible, follow up interviews also were conducted to obtain additional information about further revenue losses resulting from data center outages.

Quantifying the Cost of Downtime

The study, completed in 2011, uncovered a number of key findings related to the cost of downtime. Based on cost estimates provided by survey respondents, **the average cost of data center downtime was approximately \$5,600 per minute.**

Based on an average reported incident length of 90 minutes, **the average cost of a single downtime event was approximately \$505,500.** These costs are based on a variety of factors, including but not limited to data loss or corruption, productivity losses, equipment damage, root-cause detection and recovery actions, legal and regulatory repercussions, revenue loss and long-term repercussions on reputation and trust among key stakeholders.

Though direct costs accounted for nearly one third of all costs reported, indirect and opportunity costs – significantly more difficult to perceive for rank-and-file staff – proved to be significantly more costly, accounting for more than 62 percent of all costs resulting from data center downtime.

While business disruption and lost revenue were cited as the most significant cost consequences of downtime, other less obvious costs such as losses in end-user and IT productivity also had a significant impact on the cost of an average downtime event (Figure 3).

Surprisingly, equipment costs were among the lowest costs reported for a downtime event, averaging approximately \$9,000 per incident. This means that **the residual, downstream effects of a data center outage often are far more costly than the costs to detect and remedy the root cause of an outage after it has already occurred.**

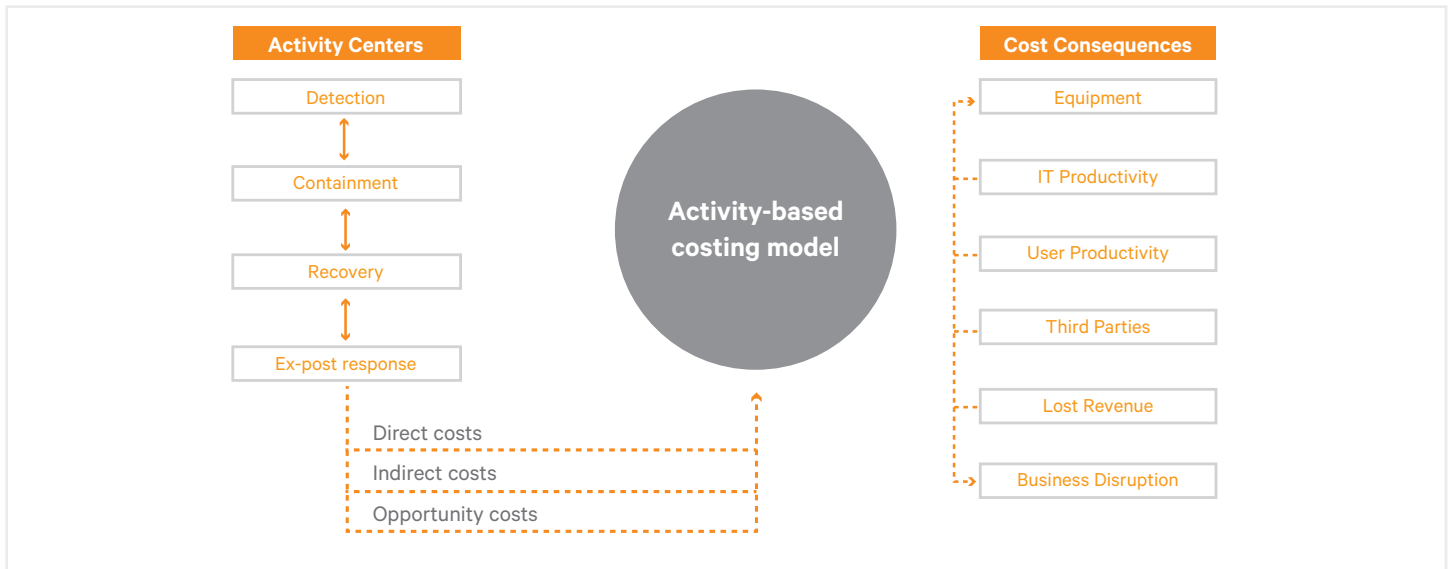


Figure 2: Activity-based cost framework.

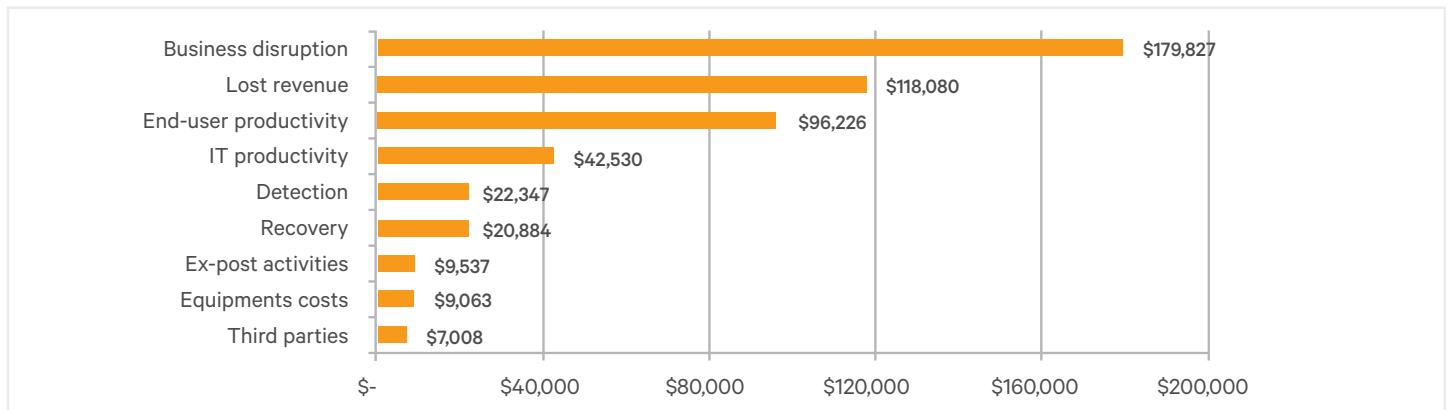


Figure 3: Average cost of unplanned data center outages for nine categories.

When considering that the typical data center in the United States experiences an average of two downtime events¹ over the course of two years, the costs of downtime for an average data center easily can surpass \$1 million in less than two years' time.

For enterprises with revenue models that depend solely on the data centers' ability to deliver IT and networking services to customers – such as telecommunications service providers and e-commerce companies – downtime can be particularly costly, with the highest cost of a single event topping \$1 million (more than \$11,000 per minute).

In total, **the cost of the most recent downtime events for the 41 participating data centers totaled \$20,735,602.**

¹ Downtime events are not limited to total data center outages. Rack- and row-level outages also are factored-in to this aggregate as well as associated downtime costs.

Other key findings from the study included:

- Total cost of **both** partial and total unplanned outages can be a significant expense for organizations (approximately \$258,000 and \$680,000 per event on average, respectively).
- The average recovery time from a total outage was more than twice that of a partial outage (134 and 59 minutes, respectively).
- Total cost of outages is systematically related to the duration of the outage and the size of the data center.
- The leading (and most costly) root causes of downtime reported by respondents were directly related to vulnerabilities in the data center's power and cooling infrastructures.

The Cost of Infrastructure Vulnerability

In addition to revenue costs associated with downtime events, a variety of costs are directly associated with the response activities necessary for restoring service and identifying and addressing the root-cause(s) of the outage. As such, respondents were asked to cite the specific root cause(s) of the most recent outage at their organization as well as all costs associated with identifying and remedying the root cause to restore data center operations.

As evidenced by Figure 4, while a variety of root causes were cited by survey respondents – including UPS system failure (battery), water incursion and IT equipment failures – the majority of root causes can be attributed to vulnerabilities in the data center’s power and cooling infrastructure. These root causes closely mirror those identified by respondents to the initial Ponemon Institute study.

As explored in the Vertiv white paper “Addressing the Leading Root Causes of Downtime,” many of the leading root-causes of downtime can be attributed to a variety of factors – chief among them being the need to “get more from less.” As demands to increase performance and efficiency increased amidst the recent national economic

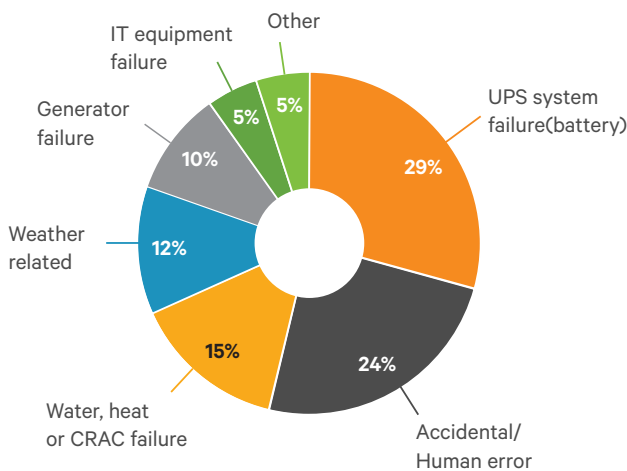


Figure 4: Primary root causes of reported unplanned outages.

recession, data center managers began implementing design strategies that achieved these gains at the cost of exposing critical vulnerabilities in their infrastructures.

Fortunately, the risk of many of the leading root causes of downtime can be minimized by observing best practices in infrastructure design and system redundancy, as well as implementing a comprehensive preventive service and maintenance regimen.

In the following sections, this paper will further examine the costs incurred by vulnerabilities in respondents’ power and cooling infrastructures as well as actions and best practices that can be implemented to minimize recovery costs as well as the overall risk of downtime².

Power-Related Outages

According to survey respondents, more than 39 percent of data center outages reported were attributed directly to vulnerabilities in the data center’s power. Among the general root causes of downtime related to power, **UPS related failures (including batteries) proved to be the most costly (\$687,700)** followed by generator failures (\$463,890).

One of the primary reasons power vulnerabilities are so costly for data centers is that a failure in the power infrastructure will likely result in a catastrophic, total unplanned outage. This means that in addition to any direct costs incurred to remedy the cause of the outage, indirect and opportunity costs also will be significant due to the fact that all stakeholders will be affected by the outage.

By definition, Tier I and II data center facilities are not equipped with the technologies needed to isolate a power system failure, such as redundancy, dual power paths and static switches. As a result, the availability of these data centers’ power infrastructures is wholly dependent on the integrity of the facility’s single backup system.

Because Tier I and II data centers can do relatively little to prevent the indirect and opportunity costs incurred by a total data center outage caused by a power failure, making investments that minimize the impact of a power system failure on data center operations is strongly recommended. One of the best ways to do this is to ensure that all power systems are backed by an adequate level of redundancy.

Implementing redundancy allows facility managers to eliminate single points of failure in their power

² NOTE: For detailed recommendations for fortifying data center infrastructures against the most common root-causes of downtime, please refer to the companion white paper “Addressing the Leading Root Causes of Downtime: Technology Investments and Best Practices for Assuring Data Center Availability.”

infrastructures. Because there is always a possibility of equipment failure over time, redundancy ensures that a backup is always in place. While direct costs would still be incurred to repair or replace the failed module, the equipment failure would not have a catastrophic impact on data center availability, and thus the organization would not incur the substantial indirect and opportunity costs associated with a total unplanned outage.

When adding a UPS for redundancy or replacing an existing or failed module, the long-term reliability of the solution should be the highest priority. Some UPS systems, including the Liebert NXL, also are capable of achieving superior performance and availability through redundant components, reduced number of components, fault tolerances for input currents and integrated battery monitoring capabilities.

In addition to establishing redundancy in the power infrastructure, adequate service and maintenance for critical power systems can play a significant role minimizing the risk of power equipment failure. In fact, even a single annual preventive maintenance visit can increase the “mean time between failure” (MTBF) of a UPS unit by more than ten-fold.

Finally, the implementation of comprehensive infrastructure monitoring and management tools such as Liebert Nform, Liebert SiteScan and Alber Battery Monitoring also can minimize the activity costs intrinsic to detecting and recovering from power system failures. Integrating a comprehensive monitoring solution – including battery and branch circuit monitoring – allows IT staff to quickly identify, isolate and address power equipment issues.

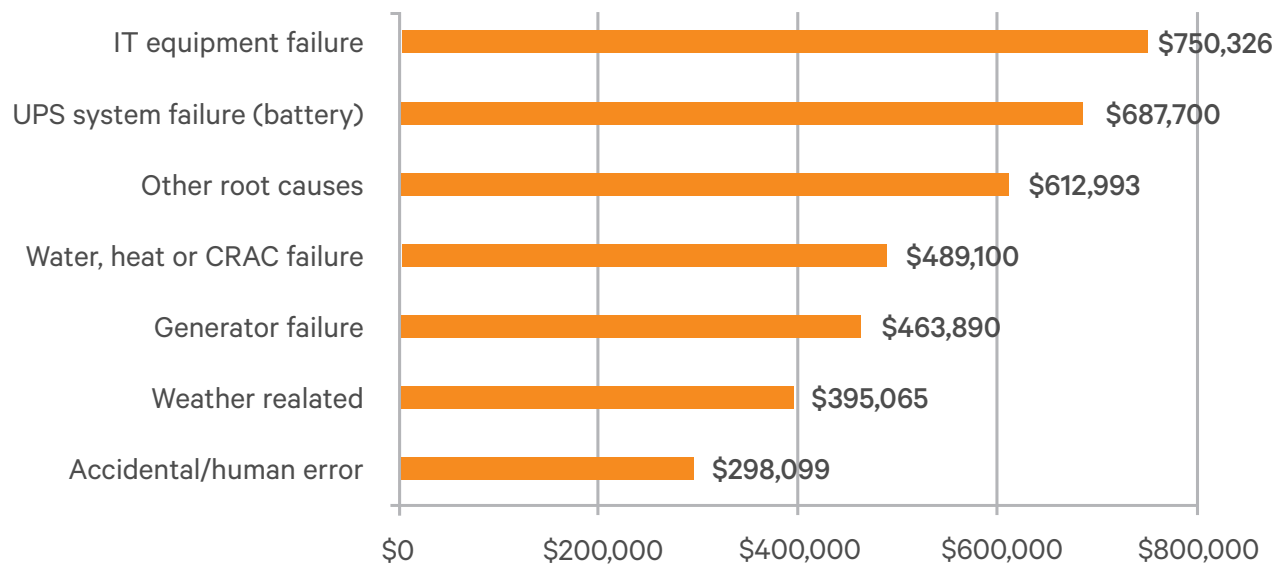


Figure 5: Average total cost by root causes of the unplanned outage.

Environmental-Related Outages

Along with vulnerabilities in the power infrastructure, environmental vulnerabilities also accounted for a noteworthy portion of the root-causes cited by survey respondents. Fifteen percent of all root causes were directly attributed to thermal issues, including water incursion and IT equipment failures related to heat density and cooling capacity. **The costs associated with detecting and recovering from these failures also was significant, at more than \$489,000 per incident.**

Environmental issues also are a leading cause of IT equipment failures. In fact, though IT equipment failures only accounted for five percent of root causes cited by survey respondents, **these failures incurred the highest overall cost – more than \$750,000.**

In many cases, a single failure can cause a chain reaction of IT equipment failures – requiring extensive detection and recovery efforts to identify the root-cause in addition to the replacement of affected IT equipment. For example, a chilled water leak in the data center’s in-row cooling system can cause the failure of sensitive IT equipment. In addition to identifying and remedying the cooling issue that caused the outage, servers and other damaged IT equipment will need to be replaced.

Also, it is critical to point out that cooling equipment **does not** need to fail to cause an IT equipment failure. Conversely, these failures – typically caused by high heat densities and “hot spots” within the rack – frequently occur as a result of an **inadequate cooling infrastructure rather than a cooling equipment failure.** This further reinforces the importance of an optimized cooling infrastructure.

While some outages relating to the data center’s cooling infrastructure may be more isolated than power-related failures – contributing to both total and partial data center outages – a comprehensive cooling infrastructure remains critical to minimizing downtime events and their associated costs. This is particularly true considering the many connections between a data center’s cooling infrastructure and the viability of critical IT equipment – where cooling systems do not need to fail to cause catastrophic failures and damage sensitive and costly equipment.

Fortunately, there are a number of best practices and investments that can be made to a data center’s cooling infrastructure to minimize the risk of catastrophic equipment failures and associated downtime events. Many of these best practices are explored in the white paper “Addressing the

Leading Root Causes of Downtime,” including:

- Minimizing the risk of water incursion through the use of **refrigerant-based cooling instead of water-based solutions.**
- Eliminating hot spots and high heat densities by bringing precision cooling closer to the load via **row-based precision cooling** solutions.
- Installing **robust monitoring and management solutions with remote monitoring** functionality.
- Fortifying cooling and IT equipment investments with **regular preventive maintenance and service visits.**

While these recommendations embody many of the best practices for maximizing the availability, effectiveness and efficiency of the data center’s cooling infrastructure, some vendors, including Vertiv, now offer facility managers the ability to implement an integrated solution optimized for efficient, high-availability power and cooling performance. These solutions offer all of the aforementioned design best practices, some with the additional benefit of rapid deployment for data center expansion or disaster recovery.

These integrated solutions also offer the added benefit of efficient precision cooling through cold-aisle containment (See Figure 6), maximizing the effectiveness of the integrated cooling solution. These characteristics play a critical role in focusing cooling based on the real-time needs of the equipment housed within the racks, minimizing the risk of hot spots and other faults common in high density computing environments while operating at a high level of efficiency.



Figure 6: Data center solutions to optimize precision cooling, like SmartAisle from Vertiv, address specific needs with rapidly deployable solutions that costeffectively add data center capacity, improve IT control and increase efficiency.

Making the Business Case for Infrastructure Optimization ³

As detailed in the preceding sections, vulnerabilities in a data center's infrastructure can have a dramatic impact on a facility's susceptibility to costly downtime events totaling hundreds of thousands of dollars. However, as this paper has demonstrated, only 29 percent of rank-and-file IT staff members believe that their companies have implemented the technologies and best practices required to minimize the occurrence and impact of data center downtime.

This disconnect begs the obvious question: If executives understand the role of their data centers in generating revenue and sustaining their respective business models, why have many hesitated to make the necessary investments required to fortify their infrastructures against downtime? The likely answer is that, prior to quantifying the cost of data center downtime, most executives could not recognize how downtime prevention speeds the ROI of their infrastructure investments.

As evidenced by the findings of the Ponemon Institute, downtime can result in a variety of long-term recurring costs, which include direct costs associated with identifying and addressing root causes, as well as indirect costs associated with disrupting business-critical operations. While minimizing the risk of downtime events and their overall financial impact may necessitate a significant up-front CAPEX investment, when considering the gains in direct and indirect downtime costs as well as savings gleaned from increases in efficiency that reduce OPEX, select investments can actually speed a business' time-to-ROI while reducing a data center's total cost of ownership over time.

To emphasize this point, one needs only to compare the cost of infrastructure optimization to the average cost and occurrence of downtime over time. It is important to first understand how the cost of downtime impacts the speed to ROI for data center infrastructure investments.

³ NOTE: Though based on real-world scenarios, the costs detailed in this analysis are approximations of market costs for a reference model data center (presented in Appendix A). To obtain a detailed estimate for optimizing your specific data center infrastructure in accordance with the below recommendations, please contact your Vertiv Representative.

Power Infrastructure Optimization

First, consider that a typical unoptimized enterprise data center experiences an average of ten downtime events over a period of ten years, spanning a variety of root causes. At an average per-event cost of just over \$500,000 (including direct costs, indirect costs and opportunity costs), a typical enterprise data center can incur more than \$5 million in downtime costs during this time.

UPS system failure costs accounted for 29 percent of data center outages reported by survey respondents. Extrapolated over ten years, these data centers can expect to incur at least three downtime events related to UPS system failure, at an average total cost in excess of \$2 million in total downtime costs.

Compare this figure to the approximate costs associated with adding UPS redundancy to a 2,500-square-foot data center with 105 high-density racks (1,000 servers) and a facility power draw of approximately 1,200 kW. Adding UPS redundancy to a data center of this size would likely require an initial capital investment of approximately \$250,000 and an annual investment of up to \$15,000 for two annual preventive service visits (increasing the MTBF for UPS systems by up to **23 times**).

Based on these numbers, when extrapolating these investments over ten years, the total investment in strengthening this data center's UPS systems infrastructure would be approximately \$400,000. Compared to the average total cost of downtime events caused by a UPS systems failure as reported by respondents (\$687,000), **ROI is easily achieved through the prevention of a single UPS-related downtime event**. Furthermore, over a period of ten years, ROI can be achieved three-fold in potential downtime costs alone, not considering gains in efficiency and OPEX associated with reactive service visits.

Cooling Infrastructure Optimization

A similar analysis can be conducted with regard to the optimization of a data center's cooling infrastructure. Data center outages related to failures or inadequacies of critical cooling systems accounted for approximately 20 percent of reported outages, including IT equipment failures. Collectively, the average cost of these root causes was approximately \$554,000. This means that if an average data center experiences ten downtime events over a period of ten years, an average of two events (with an average total cost of more than \$1.1 million in downtime costs) will be related to vulnerabilities in the data center's cooling infrastructure.

To contrast these costs with the cost of infrastructure optimization, one can revisit the aforementioned “model” data center. In this case, the model data center is assumed to rely on eight chilled-water based cooling solutions servicing load from the data center’s IT equipment, UPS and PDU systems, as well as building egress and human load.

Based on these parameters, it is strongly recommended that data center managers invest in an assessment of their data center space. These services can range from a data center audit performed by trained service representative (often free as part of an existing service agreement) or a more comprehensive thermal assessment complete with CFD modeling (approximately \$12,000 for the baseline data center in Appendix A) which unveils a clear picture of vulnerabilities in a data center’s cooling infrastructure and areas where significant efficiency gains can be achieved through cooling optimization. Often, such assessments conclude that additional equipment investments can be postponed by optimizing the configuration of cooling systems, racks and IT equipment.

By optimizing a data center’s existing cooling infrastructure via a cold-aisle containment strategy (costing as little as approximately \$15,000 for a partitioned containment

solution), data center managers and dramatically enhance the effectiveness of their cooling equipment with the added benefit of significant gains in energy savings. The addition of intelligent controls (Liebert® iCOM™) and remote monitoring to a contained infrastructure (approximately \$80,000 for the baseline data center presented in Appendix A) can further enhance cooling efficiency by at least 12 percent and ensure that all IT equipment is being adequately and precisely cooled based on real-time heat densities (see Figure 7). Finally, investing in ongoing preventive maintenance and service for the equipment (an approximate annual investment of \$2,000) and installation of a comprehensive leak detection solution for all cooling units (approximately \$5,000) is recommended.

Over ten years, the total investment in strengthening this data center’s cooling infrastructure would be approximately \$135,000 (\$115,000 in year one). Compared to the average total cost of a single downtime event caused by IT systems failure or thermal-related outages as reported by respondents (\$554,000), these investments can easily be justified if they prevent even a single thermal-related downtime event.

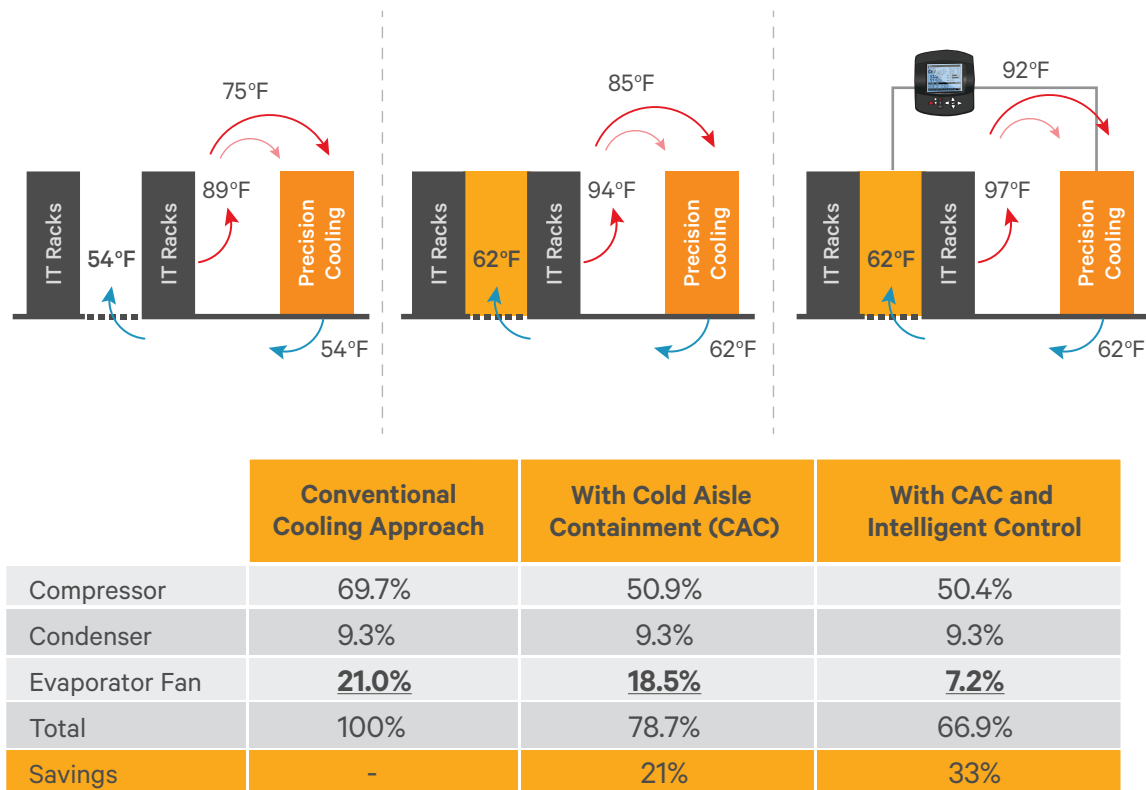


Figure 7: Dynamic control provides an additional 15 percent increase in total system efficiency over cold aisle containment alone.

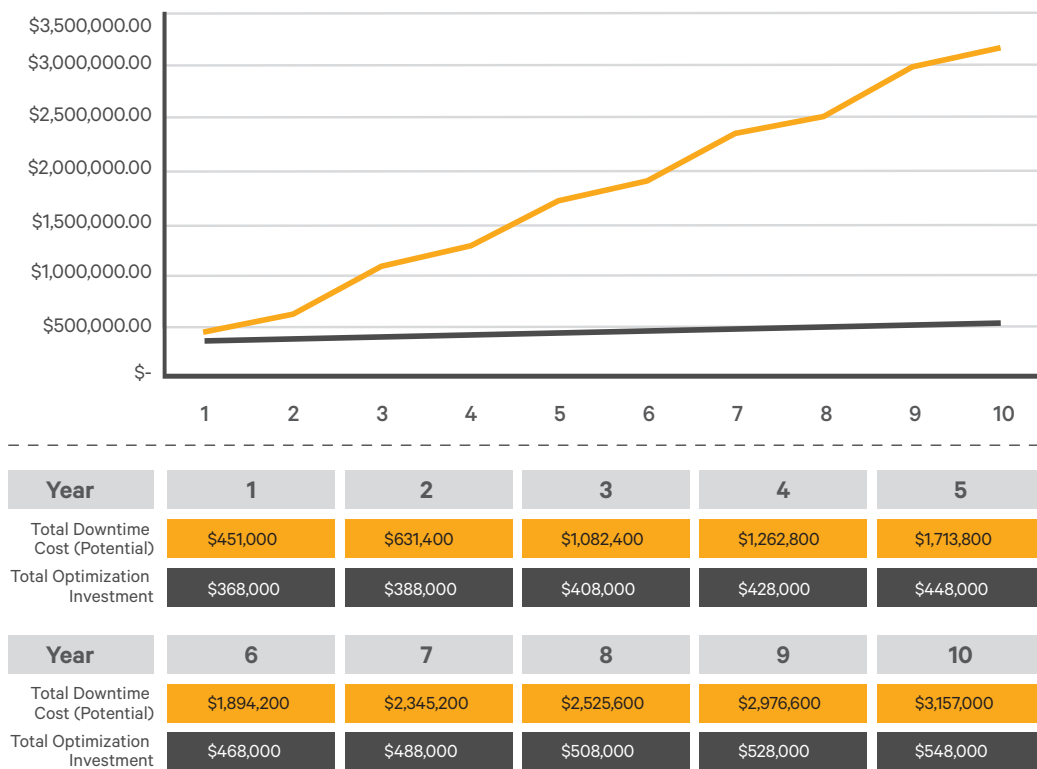


Figure 8: Potential downtime costs (orange) compared to CAPEX and ongoing service investment for power and cooling infrastructure optimization (dark gray).

Furthermore, as in the case of power infrastructure optimization, over a period of ten years, ROI can be achieved several times over when considering potential downtime costs as well as significant gains in energy efficiency – cutting cooling-related energy usage by as much as 33 percent.

Other Opportunities for Optimization

In addition to vulnerabilities in the data center’s power and cooling infrastructure, accidents and human errors also can cause costly downtime events.

Twenty-four percent of study respondents cited human error as the primary cause of their most recent downtime event, with downtime caused by human error accounting for nearly \$300,000 in downtime costs per incident. Over a period of ten years, downtime events related to human errors and/or accidents can easily cost an organization in excess of \$600,000.

Fortunately, best practices to minimize the risk of downtime events caused by human error are among the least expensive to implement. As explained in the white paper

“Addressing the Leading Root Causes of Downtime,” recommended actions for minimizing the occurrence of human errors and Accidental Emergency Power Off (EPO) events include:

- Shielding Emergency OFF buttons
- Strictly enforcing food and drink policies
- Avoiding contaminants
- Establishing secure access policies
- Performing ongoing personnel training
- Promoting consistent standards for operation
- Labeling all components accurately
- Documenting maintenance procedures

According to experts from Vertiv's Liebert Services business, **implementing these recommended actions would cost approximately \$3,500.** When considering the high overall cost of downtime, such investments represent a nominal cost that can easily achieve an ROI of more than a hundred-fold by preventing a single error or accident.

A Comprehensive Comparison

To put all of these calculations into greater perspective, vulnerabilities in a data center's UPS and cooling infrastructure, as well as human error and accidental EPO events, collectively account for nearly three quarters of the root causes of downtime reported by survey respondents with **an average cost of more than \$450,000 per incident.** As such, for data centers experiencing an average of ten major or minor downtime events over a period of ten years, UPS, cooling and human error-related outages can be expected to account for at least seven major or minor downtime events, with **an average total cost in excess of \$3.15 million.**

As illustrated in Figure 8, the ROI of infrastructure optimization can be immediately realized when comparing the potential cost of downtime to the approximate cost of recommended investments capable of minimizing the risk for these root causes: **\$548,000 including ten years of preventive maintenance of power and cooling equipment; \$368,000 in Year One.**

Furthermore, when considering the additional efficiency gains achieved as a result of these changes, the return on investment in power and cooling infrastructure optimization is particularly evident, especially when considering long-term savings in indirect and opportunity costs unique to reoccurring downtime events.

Investment Prioritization: Evaluating Existing Infrastructure

While the recommended actions outlined in this paper are critical to minimizing the risk of the leading root causes of downtime (as well as their associated costs), many enterprises may wish to prioritize these investments over time. These decisions are often based on a variety of factors, including CAPEX and OPEX required for comprehensive optimization, the criticality of data center operation and the impact of planned downtime on data center operations.

If a comprehensive infrastructure overhaul is not feasible, spreading out investments over time can be an effective way to balance short-term CAPEX/OPEX with the long-term cost

and risk of the leading root causes of downtime, center operations. For example, many of the recommended actions for safeguarding against human error and accidental EPO represent "low hanging fruit" and are relatively inexpensive to execute. As a result, some data centers may choose to complete these and other minimally invasive optimizations (such as row partitioning) first, and plan for more intensive optimizations based on available resources and a required time-to-ROI.

However, regardless of whether an enterprise decides to complete an infrastructure overhaul or space out these updates over time, many overlook the need to complete comprehensive assessments of their existing infrastructures, a critical step that can help to avoid unnecessary investments that yield little additional value in terms of availability or efficiency.

As highlighted in "Addressing the Leading Root Causes of Downtime: Technology Investments and Best Practices for Assuring Data Center Availability" White Paper from Vertiv, a comprehensive assessment of the facility as well as all thermal and electrical systems can offer detailed insight into how an existing data center can be optimized for efficiency without compromising the availability of critical systems.

In addition to the performance of a data center's power and cooling systems, data center assessments also take into consideration a variety of additional factors not tied directly to equipment performance that can impact the availability and performance of critical systems, including heat densities in racks and rows, raised floor obstructions and arc flash vulnerabilities in the data center's electrical infrastructure.

Based on the assessment performed by specially trained service personnel, the data center manager can clearly assess where capital investments are required (including redundant power systems and precision cooling equipment designed for high-density environments) and where existing infrastructure can be adjusted or optimized in accordance with best practices to minimize the risk of data center downtime.

Conclusion

As evidenced by the findings of the Ponemon Institute, a single downtime event now has the potential to significantly impact the profitability (and, in extreme cases, the viability) of an enterprise. This trend can be attributed to a variety of economic trends, evolving business practices and the emergence of revenue streams that are wholly dependent on the availability of critical IT systems.

With an average downtime cost for an enterprise data center totaling thousands of dollars per minute, it is vital to close the widening disconnect between IT personnel and their C-suite counterparts. An effective way to achieve this goal is to promote a thorough understanding of the frequency, cost and causes of data center downtime.

Left unattended, an inadequate data center infrastructure will contribute to recurring downtime events and result in significant financial losses as well as permanent damage to a company's reputation and customer goodwill. While identifying these vulnerabilities and addressing them based on some of the aforementioned best practices may require a significant up-front cost, when contrasting these investments with the potential "bottom line" costs of data center downtime, data center professionals can gain a clear understanding of how direct and indirect costs can impact revenue over time.

Appendix A: Infrastructure Assumptions for Model Data Center (Pre-Optimization)

The 2,500-square-foot hypothetical data center has 105 racks with average density of 5.6 kW each. The racks are arranged in a hot-aisle/cold-aisle configuration. Cold aisles are four feet wide, and hot aisles are three feet wide. Based on this configuration and operating parameters, average facility power draw was calculated to be 1,127 kW.

Following are additional details used in the analysis:

Servers

- Age is based on average server replacement cycle of 4-5 years.
- Processor Thermal Design Power averages 91W/ processor.
- All servers have dual redundant power supplies. The average DC-DC conversion efficiency is assumed at 85% and average AC-DC conversion efficiency is assumed at 79 percent for the mix of servers from four-years old to new.
- Daytime power draw is assumed to exist for 14 hours on weekdays and 4 hours on weekends. Night time power draw is 80 percent of daytime power draw.

Storage

- Storage Type: Network attached storage.
- Capacity is 120 Terabytes.
- Average Power Draw is 49 kW.

Communication Equipment

- Routers, switches and hubs required to interconnect the servers, storage and access points through Local Area Network and provide secure access to public networks.
- Average Power Draw is 49 kW.

Power Distribution Units (PDU):

- Provides output of 208V, 3 Phase through whips and rack power strips to power servers, storage, communication equipment and lighting. (Average load is 539kW).

- Input from UPS is 480V 3-phase.
- Efficiency of power distribution is 97.5 percent.

UPS System

- One double conversion 750 kVA UPS with input filters for power factor correction (power factor = 91 percent).
- The UPS receives 480V input power for the distribution board and provides a 480V, 3 Phase power to the power distribution units on the data center floor.
- UPS efficiency at part load: 92.5 percent.

Cooling system

- Cooling System is chilled water based.
- Total sensible heat load on the precision cooling system includes heat generated by the IT equipment, UPS and PDUs, building egress and human load.
- Cooling System Components:
 - Eight 146 kW chilled water based precision cooling system placed at the end of each hot aisle. Includes one redundant unit.
 - The chilled water source is a chiller plant consisting of three 200 ton chillers (n+1) with matching condensers for heat rejection and four chilled water pumps (n+2).
 - The chiller, pumps and air conditioners are powered from the building distribution board (480V 3 phase).
 - Total cooling system power draw is 429 kW.

Building substation:

- The building substation provides 480V 3-phase power to UPS's and cooling system.
- Average load on building substation is 1,099 kW.
- Utility input is 13.5 kVA, 3-phase connection.
- System consists of transformer with isolation switchgear on the incoming line, switchgear, circuit breakers and distribution panel on the low voltage line.
- Substation, transformer and building entrance switchgear composite efficiency is 97.5 percent.

