

A Quantitative Comparison of UPS Monitoring and Servicing Approaches Across Edge Environments

White Paper 283

Version 1

by Wendy Torell

Executive summary

A fleet of single-phase UPSs distributed geographically across many edge sites presents unique challenges when it comes to monitoring and servicing. In this paper we present key considerations when deciding between managing the fleet of UPSs yourself vs. outsourcing that responsibility to a third-party vendor or partner. A tool is also presented that provides a framework for discussion on quantifying the costs associated with each alternative. We walk through four scenarios and demonstrate how key drivers like age distribution of the fleet and cost of downtime influence which approach makes financial sense.

RATE THIS PAPER



Introduction

UPSs must be managed properly over their lifecycle in order to do the job for which they were intended – to keep critical systems up and running. When a company has UPSs distributed across dozens, hundreds or even more edge computing sites, managing the fleet often becomes challenging. This is primarily because (1) the sites are likely to be geographically dispersed, (2) there is a lack of trained and/or dedicated staff on each site, and (3) the UPSs vary in age across sites.

There are two core functions of effective UPS fleet management:

- **Monitoring** – Most UPSs today can be proactively monitored through either on-premise software or cloud-based apps. UPSs log activity, report problems, and tell you when batteries, which are consumable parts, are near the end of their life. Proactive monitoring is essential for maintaining high availability.
- **Servicing** – There are two general reasons for maintenance activities – (1) proactive or preventive maintenance, where the system is serviced ahead of failures to reduce the likelihood of unplanned downtime and (2) reactive or break/fix maintenance where work is done once a problem is detected in order to make it operational again.

Figure 1 depicts the possibilities for who performs these two core functions. As the diagram illustrates, there are **two approaches for each of the core functions**, or a total of four combinations of management approaches.

1. Perform both monitoring & servicing yourself with internal staff / resources (blue quadrant)
2. Outsource both to a 3rd party vendor or partner (green quadrant)
3. Monitor yourself but outsource the servicing tasks to a 3rd party (light grey quadrant)
4. Outsource the monitoring to a 3rd party but do the service tasks yourself (dark grey quadrant)

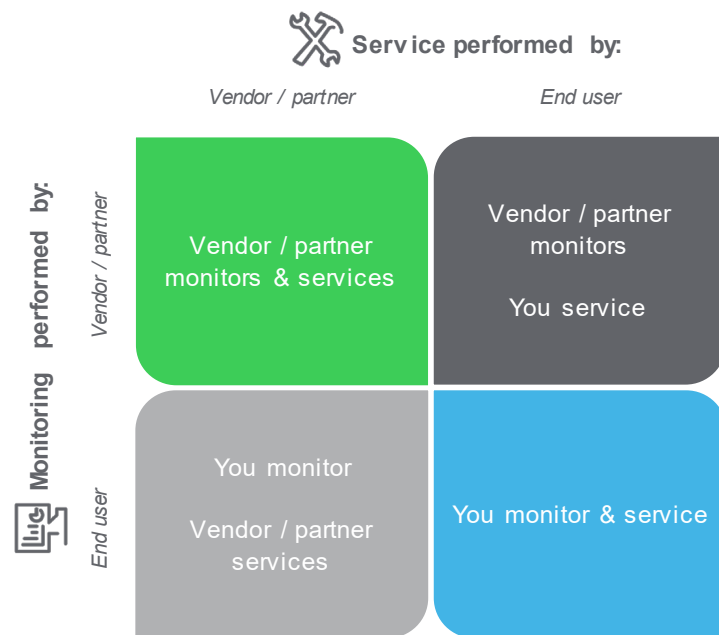


Figure 1
Diagram illustrating the four possible approaches to managing a UPS fleet

In this paper, we first discuss the *qualitative* tradeoffs between the two **monitoring** approaches (self vs. 3rd party) as well as the *qualitative* tradeoffs between the two **servicing** approaches (self vs. 3rd party). Then we present a tool that helps *quantify* the annual operational expense (opex) of managing the fleet entirely yourself (blue quadrant) vs. outsourcing both the monitoring and servicing functions (green quadrant). Through scenarios, we demonstrate the key drivers impacting which approach is financially favorable for a particular UPS fleet. Note, the tool focuses on scenarios where the fleet is either entirely managed internally, or entirely managed with 3rd party resources. The tool does not consider the hybrid scenarios where one function is fulfilled in-house while the other is outsourced, although the framework for analyzing them would be similar.


Monitoring

Proactive monitoring of UPSs can reduce downtime, decrease the mean time to recover (MTTR) from failure events, and lower the cost of maintenance tasks. But not all monitoring is the same.

Spectrum of monitoring

Before we get into the considerations of self-monitoring vs 3rd party monitoring, it's important to recognize that there is a spectrum when it comes to the type of monitoring. **Figure 2** illustrates the spectrum that can exist for edge sites. As you move from left to right, the benefits described increase.

Figure 2
Spectrum of monitoring for a distributed fleet of UPSs



- **No proactive monitoring** – Instead of having staff proactively monitor the UPSs, some companies may choose to wait until problems occur and then simply react to those problems. While monitoring costs are clearly reduced in this scenario, downtime and its associated cost will certainly increase since they will depend on the time it takes to find and react to the issues.
- **Physical inspection** – This is the most basic type of monitoring, which relies on staff that are on-premise to perform visual inspection (i.e. looking for indicator lights on or flashing) or listen for audible alarms. When staff are in close proximity to the systems, this is sometimes deemed adequate for basic UPSs, assuming the staff are knowledgeable on the alerts of the systems.
- **Basic remote monitoring** – With a fleet of distributed assets, having a staffed network operations center or NOC (physical or virtual) allows you to have visibility to all distributed UPS assets. With decentralized remote monitoring, the assets are compartmentalized, which makes it more challenging to know the overall health of the fleet. It requires checking of each UPS one-by-one or reacting to intermittent data (i.e. email alerts).
- **Advanced remote monitoring** – This real-time monitoring provides the staff with an aggregated dashboard of all assets, leverages the cloud and a data lake for insights that may help predict failures, and suggests actions to maximize the life of the assets. It provides a complete view of the fleet on one dashboard and predicts component failures for just-in-time parts replacements that maximize service life and minimize maintenance costs.

We recommend advanced monitoring of a distributed fleet. White Paper 237, [Digital Remote Monitoring and Dispatch Services' Impact on Edge Computing and Data Centers](#), provides further discussion on the benefits of advanced remote monitoring over more basic approaches. Note, the discussion of what level of monitoring to choose for your business is different than the decision of *who* performs the work. In this paper, we focus on the implications of doing the monitoring yourself vs. outsourcing the function, and assume advanced monitoring is feasible by you or a 3rd party provider.

Considerations in choosing who performs monitoring function

Later in this paper, we will walk through a tool that quantifies the financial drivers and impacts... but there are also qualitative factors. Below we identify three key considerations that should factor into the decision of monitoring yourself or relying on a 3rd party.

Digital capabilities – Monitoring your fleet yourself requires a physical or virtual NOC that is sufficiently staffed. Many companies already have a NOC in place to monitor the IT assets at the distributed sites. If the staff have the bandwidth to monitor the UPSs (or other physical infrastructure) and are knowledgeable on these systems, then this may be a good option. Effectively doing this, however, requires ensuring all UPSs are connected to the network, for example through network management cards (NMCs), and firmware maintained with the latest updates. While these sound like obvious and straight-forward tasks, they sometimes prove challenging given the distributed nature of the assets, the mixed ages of the assets (you may have sites getting new UPSs each year), and the fact that personnel at these edge sites aren't trained to complete these tasks. If new UPSs added to your fleet aren't network connected, you'll only have a partial view of your fleet, and for the remaining assets, you'll be, in essence, at the left side of the spectrum of **Figure 2**, flying blind. Cyber-security is another related factor. If using a 3rd party to monitor, it is important they can demonstrate their secure development lifecycle (SDL) practices and policies, and that there are minimal points of entry into your network using gateways. There are remote monitoring architectures where the staff doing the monitoring only access data in the cloud vs. connecting locally to your network. In these cases, device data is sent from the gateway to the cloud, in one direction only.

Staffing – When it is someone's part time job or "side job", things are likely to be missed, resulting in downtime that could have been avoided. Ideally the systems are being monitored by a central full-time resource that is proactive in ensuring resolution to notifications and alerts before service disruptions occur. For some companies, having their own NOC staff is worth the investment because of their business philosophies and desire to remain in control, and to have complete visibility and flexibility with how to manage the assets. Some companies already have a NOC to monitor their IT/telecom equipment, so adding UPSs (and possibly other physical infrastructure) to the list of equipment to monitor does not become a major investment or burden. For others, offloading it to a 3rd party provides peace of mind, as we discuss next.

Expertise & peace of mind – Subject matter expertise is an important consideration. Having the staff and infrastructure to monitor the fleet isn't enough. The delegated staff must know the ins and outs of the systems well enough to take the appropriate actions when a problem occurs. With visual inspection monitoring, consider retail store fronts where it is often a store manager or salesperson that goes into the break room to deal with the beeping UPS. If they are unfamiliar with the system and what the alert means, they will likely either ignore it (alarm fatigue commonly sets in when the same alarm goes off repetitively) or take an action that has a negative impact. If you chose to monitor yourself, make sure the vendor can offer

adequate technical training to staff, and that they have technical support resources to help when problems arise. With 3rd party monitoring, ensure the vendor or partner has subject matter experts attending to your fleet around the clock, addressing alarms before downtime problems occur. Although hard to quantify, the peace of mind that can come from knowing someone is looking after your fleet, keeping track of assets of varying ages, and mitigating downtime risks 24x7 is worth the investment.


Servicing

Performing service on a distributed fleet presents another set of challenges. While there is some overlap in the challenges faced with the functions of monitoring and servicing, it is very possible that a business chooses to do the monitoring themselves but outsources the servicing to a 3rd party.

Types of servicing

Before we discuss the considerations in choosing whether to service yourself or outsource, it's important to clarify the types of preventive service that may be implemented at edge sites. **Figure 3** illustrates the spectrum of servicing. Again, as you move from left to right, the benefits increase.

Figure 3
Spectrum of servicing for
a distributed fleet of UPSs



- **Run to fail** – Some companies choose this low-cost maintenance method since there is no recurring service contract. The UPSs are allowed to operate until they fail. When a system fails, a service ticket is placed, and the system is restored to operation or replaced. Costs are only incurred when a failure event occurs. This reactive servicing strategy, however, means the likelihood of downtime and therefore downtime cost is increased.
- **Calendar-based** – Regularly occurring maintenance visits (often annually) are used to inspect systems and proactively perform any necessary parts replacement to prevent future failures from occurring. Since they are scheduled in advance, the sites can plan for the service visit to minimize impact on daily operations.
- **Condition-based** – Sensor data from the UPSs are used to predict when a maintenance activity is necessary, for just-in-time servicing. This type of maintenance minimizes unscheduled downtime, maximizes the service life of the UPS battery, and ensures reliable operation. An example is waiting for a battery to fail its self-test before replacing the battery. Once that test is failed, you have weeks to months left to proactively replace it.

For more information on the differences of these service types and the benefits of condition-based maintenance, see white paper, [Strategies for Maintaining Electrical Distribution Equipment](#). We recommend condition-based whenever possible, whether it is serviced by your own staff, or by a vendor, or partner's staff.

Considerations in choosing who performs servicing function

In this section we discuss factors to consider when choosing your servicing approach (self-serviced or 3rd party serviced), that impact the effectiveness of proactive as well as remediation maintenance of your UPS fleet.

Logistics of getting technicians where they are needed – With geographic dispersity, scheduling technicians to arrive when they are needed at each site to perform maintenance can be a cost-prohibitive logistics nightmare, depending on the number of sites, technicians available, and the distance between sites. Staffing can also be a challenge when doing it yourself, as the number of field technicians needed may change as the assets age. If you plan to service yourself, make sure you understand the service expectations of the systems and staff accordingly. With assets that are 3rd party managed, you generally have a contract guaranteeing¹ a response time so expectations are clear, and problems are resolved quickly (often as early as next business day parts delivery). Vendors focused on this type of dispatch service have the benefit of scale and can streamline their operations.

Unpredictable expenses when systems fail – Choosing to do servicing of the UPS fleet yourself means that when a system fails, there is a cost incurred to service the system. This includes the cost to get the field technician to the site to do the repair (sometimes referred to as truck roll or dispatch expense which includes the technicians time and travel), plus, if the UPS is out of warranty, the cost of the necessary parts. A common dispatch cost is \$1000², so it can add up fast with a fleet. This irregular cost that occurs when emergencies and failures happen is OK for some companies (as you only pay for services you need), but for others, there is a preference towards predictable recurring operating expenses, as is typically the case with a 3rd party service contract. With a vendor or partner service contract, you are able to amortize the parts replacement expenses by paying a portion of it each year.

Expertise & confidence of onsite staff performing work – While there are typically plenty of onsite personnel at these distributed edge sites (i.e. a retail storefront), they generally have different core responsibilities, and are not trained or skilled in the maintenance of UPSs and other physical infrastructure systems. Taking someone away from their primary role at the sites (i.e. selling, managing inventory) to deal with these maintenance tasks not only has a business opportunity cost associated with it, but increases the risk of downtime from lack of familiarity with the system. Employees that are assigned the role of diagnosing and performing the service tasks must be trained. Training programs are generally available from vendors or partners.

Mixed age fleet complexity – When you have a fleet of thousands of UPSs, there is a high likelihood that that fleet is of mixed age, ranging from new UPSs to ones approaching end of life. Rarely are they all uniform age. It is also common that a fleet is made up of different models and sizes across your varying edge computing sites. Maintenance requirements and risks of downtime change as UPSs age, and different models may have different life expectancies (such as li-ion batteries vs VRLA batteries) and therefore service needs. As a result, making sure all systems are maintained adequately to prevent unexpected downtime can be challenging and time consuming. Effective UPS monitoring software (DCIM) aids in this, by providing aggregated dashboard views of the fleet, and alerting you to what systems need prioritized maintenance activity.

¹ Within the contract, clear financial consequences are documented, should the 3rd party not meet the service level agreement (SLA)

² <https://techsee.me/blog/reduce-truck-rolls/>

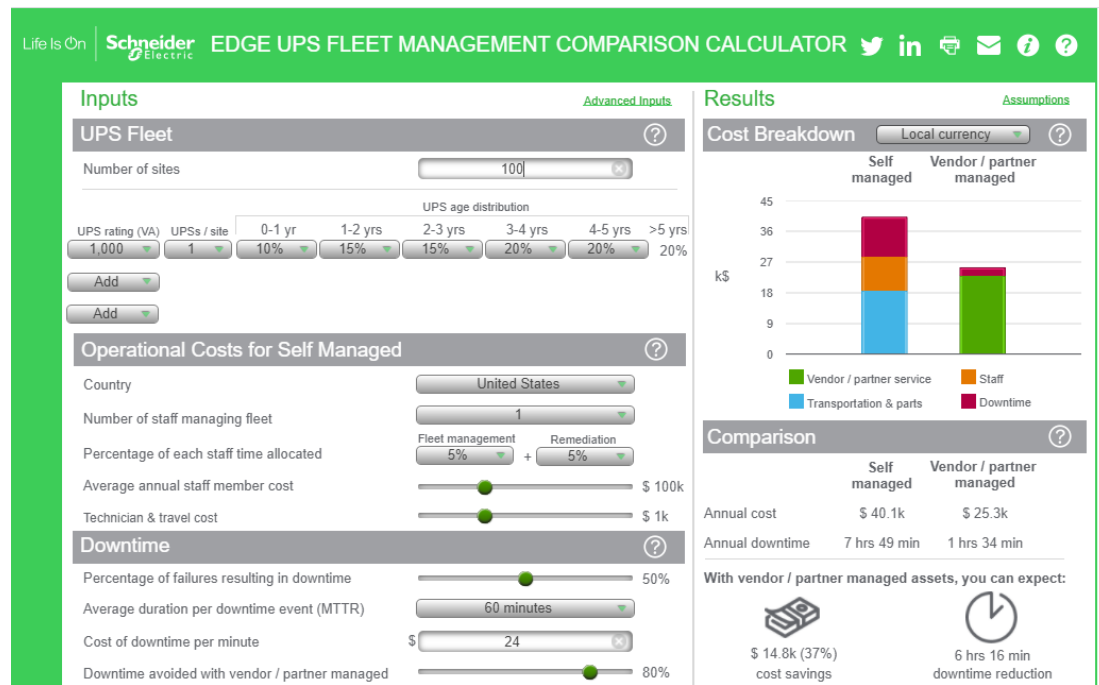
Access to parts – With break/fix problems, a customer may be able to scramble and get a technician to the site quickly, but if they can't obtain the parts in a timely fashion, the service technician is unable to perform the needed work. It is helpful to have regional distribution centers with spare parts to enable faster procurement in these cases. Generally, 3rd party service/dispatch companies have scale to do this, so you don't have to have the burden of parts procurement. Furthermore, they will put it in their contract, so you are guaranteed resolution in a specified time frame. Batteries are the most common item in need of replacement, so if maintaining the fleet with your own service personnel, ensure you have inventory that is accessible.

A tool to quantify the trade-offs

The amount of time that employees spend managing and responding to UPS fleet alarms and repair requirements, as well as costs associated with potential downtime are important quantifiable variables that influence the decision to delegate the tasks to a 3rd party or not. We developed a tool, [Edge UPS Fleet Management Comparison Calculator](#), to serve as a framework for discussing the variables that impact cost, and to help quantify the costs associated with managing a fleet of single-phase UPSs geographically distributed across edge computing sites. The tool lets you compare the financial impact of managing it yourself vs. using a 3rd party vendor or partner to manage it for you.

Figure 4 illustrates the tool. Note, we only analyzed two of the four approaches of **Figure 1** (green and blue quadrants); The hybrid approaches (light and dark gray quadrants) where one function is outsourced and the other is kept in-house requires further data analysis, although the framework and drivers would be similar.

Figure 4
Tool to help quantify the tradeoffs between managing the fleet yourself vs. using a 3rd party, [Edge UPS Fleet management Comparison Calculator](#)



Analysis methodology and assumptions

The dynamic model was developed to evaluate annual cost differences in managing a fleet yourself vs. outsourcing it, considering key attributes. This is a model for typical distributed UPS fleets found in different environments such as retail stores, healthcare facilities, or school campuses. The model calculates four types of costs:

- **Vendor / partner service cost** – When you choose a 3rd party vendor or partner to service your fleet of UPSs, the cost is generally structured as an annual contract. This price is scaled based on the power capacity (rating) of the system and the age of the system. Typical prices based on Schneider Electric's service offer were used in the model.
- **Transportation & parts cost** – When a fleet is serviced with internal staff, the cost of getting the technician to the site and the parts needed for the repair factor into the cost. The tool lets you set the transportation or truck roll cost. Part costs are assumed based on the cost of battery replacement (which varies based on the size of the UPS selected). The frequency of these recurring costs is based on the typical percentage of the UPS population that fails each year. In the advanced inputs of the tool, you can adjust these percentages based on actual data from your fleet.
- **Staff cost** – When you manage assets yourself, there are staff costs associated with the time spent monitoring and maintaining the fleet. You may have one or more people responsible for managing the fleet, but often, they are not dedicated only to those tasks. The tool lets you assign the percent of staff time allocated to the fleet, as well as define the fully loaded staff cost (salary + benefits).
- **Downtime cost** – Although less tangible than the other direct costs described above, downtime cost has the potential for significant impact on the business. The model lets you define what percentage of the expected UPS failures will result in downtime, and how much it costs your business every minute your systems are down. Downtime costs exist for both alternatives to managing your fleet, but the tool allows you to set a percent of downtime avoided when you have a 3rd party vendor watching your assets 24x7.

In order to simplify the tool, we make several assumptions about the fleet being analyzed. These assumptions include the following:

- UPSs have VRLA batteries with a life expectancy of 3-5 years.
- UPSs are remotely monitored – i.e. connected to monitoring apps through the network.
- UPSs are covered under factory warranty for 3 years (typical), so the price of 3rd party management increases after year 3 to factor in the cost of parts.
- For self-servicing, when a repair is needed for 3+ year old assets, the model assumes the parts cost is equivalent to the replacement battery cost, since that is the most common repair type. Note, this cost is generally included with 3rd party service contracts.

Scenario results

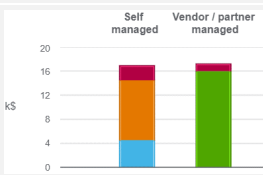



We analyzed a distributed fleet of 100 UPSs, each rated at 1000VA. As we move through each scenario, we vary one attribute (age, downtime cost, staff) to demonstrate the impact of these key drivers on the annual costs. **Table 1** defines the attributes for the 4 scenarios analyzed. Changed values are bolded.

Table 1
4 scenarios ana-

Attribute	Scenario 1 0-4 yr assets	Scenario 2 2-6 yr assets	Scenario 3 No downtime cost	Scenario 4 More staff
Number of sites	100			
UPS rating (VA)	1000 VA			
Age of UPSs	0-1 year (25%) 1-2 years (25%) 2-3 years (25%) 3-4 years (25%)	2-3 years (25%) 3-4 years (25%) 4-5 years (25%) >5 years (25%)	2-3 years (25%) 3-4 years (25%) 4-5 years (25%) >5 years (25%)	2-3 years (25%) 3-4 years (25%) 4-5 years (25%) >5 years (25%)
Staff for self-managed	1 staff, 10% of time allocated	1 staff, 10% of time allocated	1 staff, 10% of time allocated	2 staff, 10% of time allocated
Annual staff cost	\$100k/person			
Technician & travel cost	\$1000/failure			
Failures resulting in downtime	50%			
Downtime per event	60 minutes			
Downtime avoided with 3 rd party	0%			
Downtime cost	\$20/minute	\$20/minute	\$0/minute	\$0/minute

Table 2 summarizes, for each scenario, the cost differences of outsourcing to a 3rd party vs. managing the fleet yourself. As the table shows, the cost difference between the two can vary significantly, and the 3rd party outsourcing is not always the most cost-effective approach. It is therefore very important to have a good understanding of your internal management costs and characteristics of your fleet before performing such an analysis to ensure meaningful results.

Table 2
Analysis results

Attribute	Scenario 1 0-4 yr assets	Scenario 2 2-6 yr assets	Scenario 3 No downtime cost	Scenario 4 More staff
Self-managed cost	\$16.9k	\$45.2k	\$33.2k	\$63.2k
3 rd party managed cost	\$17.2k	\$32.0k	\$26.0k	\$26.0k
Cost breakdown				
Legend	■ Vendor / partner service ■ Staff ■ Transportation & parts ■ Downtime			
% difference	2% premium	29% savings	22% savings	59% savings

Key drivers

There are three key drivers that determine which approach to monitoring and servicing makes financial sense for your business: (1) Age distribution, (2) cost of downtime, and (3) operational costs of managing the fleet yourself.

Age distribution – In scenario 1, the age of the UPSs range from 0-4 years old. In scenario 2, the only attribute we changed was the age distribution – to 2-6 years old. This distribution is important because the cost of vendor / partner service contracts vary based on the warranty status of the UPSs. In-warranty contracts often already include parts coverage due to factory warranty policy. Older, out of warranty UPSs no longer carry parts coverage, and that cost will be represented in out-of-warranty pricing. For UPSs with VRLA batteries, typically the factory warranty is for the first three years of installation.

This age distribution also impacts the number of UPSs that fail in your fleet. **Table 3** represents the default assumed percentage of fleet population that fails in the tool. For the purposes of this analysis, we define a “failure” as an event requiring physical human intervention to resolve the issue. These failure events, although inclusive of critical component failures, are driven by the battery service life, since VRLA batteries are consumables that have a typical life expectancy of 3-5 years. This is why the failures increase so drastically above 3 years.

Age	Scenario 1 0-4 yr assets
0-1 year	2%
1-2 years	1%
2-3 years	2%
3-4 years	10%
4-5 years	30%
>5 years	35%

Table 3
Typical percentage of fleet “failures” by age of asset

As the typical fleet matures, and/or becomes more mixed in terms of age, the burden of in-house staff tends to increase. This is why, in general, the later in the lifecycle the UPS fleet is, the greater the value of outsourcing to a 3rd party vendor or partner.

Cost of downtime – Not all failures result in downtime, such as an alert that your UPS battery is approaching end of life – this is a proactive alarm that, if acted upon in a timely fashion, can be resolved with no downtime. Some events, however, are unexpected and lead to downtime. When managing a large fleet of UPSs yourself, the remote nature of a distributed fleet can lead to missed alerts and/or longer time frames to resolve issues. This can increase the percentage of failures that lead to downtime. We assume 50% of failures for a self-managed fleet lead to downtime. When your UPS loads are down, what does it cost your business? According to independent data protection and security research firm, Ponemon, the largest share of downtime cost is business disruption – a category that includes reputational damage and customer churn. Revenue loss took second place in the firm’s research. And the third largest financial pain associated with incidents was end-user productivity. Since downtime cost is an indirect cost, some don’t want to include it in their financial analysis. From scenario 2 to scenario 3, we simply changed the cost of

downtime per minute to \$0, to remove it from the equation. As **Table 2** illustrates, the savings went down slightly, from 29% to 22%, when we didn't factor in downtime.

Operational costs of managing the fleet yourself – When managing a fleet yourself, costs include the staff performing the monitoring functions, as well as technician / travel costs for repairs or parts replacements like batteries. The higher the volume of UPSs in the fleet, the more likely it is that more than one person is responsible for managing the fleet. Often times, people aren't dedicated to this function, so its important to understand the cost of these staff members, and how much of their time is spent on these tasks. From scenario 3 to 4, we adjusted the number of staff from one to two, to demonstrate the impact that adding one additional part-time resource can have. As **Table 2** illustrates, the savings grew to 59%. The technician/travel "truck roll" cost also impacts the result. We fixed our assumption at \$1000, but as your internal costs for repairs grows, again, the savings would grow.

Conclusion

Companies with fleets of UPSs distributed across many edge computing sites often face unique challenges in cost-effectively managing them, while ensuring high availability. Proactive monitoring *and* maintenance of the assets are necessary to achieve this.

Making the decision to manage the fleet yourself or outsource the functions to a 3rd party vendor or partner should factor in both qualitative and quantitative differences. Qualitative factors, such as having resources and expertise available, comfort-level in working with the system, and peace of mind are important considerations. Quantitative (financial) differences exist in terms of staffing cost, downtime cost, and service cost. Schneider Electric's TradeOff Tool, [Edge UPS Fleet Management Comparison Calculator](#), provides a framework for discussion and allows for simple comparison for your unique fleet of UPSs. Through examples, we demonstrated that in many cases, 3rd party management is a more cost-effective alternative.


About the author

Wendy Torell is a Senior Research Analyst at Schneider Electric's Data Center Science Center. In this role, she researches best practices in data center design and operation, publishes white papers & articles, and develops TradeOff Tools to help clients optimize the availability, efficiency, and cost of their data center environments. She also consults with clients on availability science approaches and design practices to help them meet their data center performance objectives. She received her bachelor's of Mechanical Engineering degree from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

RATE THIS PAPER







-  [Digital Remote Monitoring and Dispatch Services' Impact on Edge Computing and Data Centers](#)
White Paper 237

-  [Solving Edge Computing Infrastructure Challenges](#)
White Paper 277

-  [Three Types of Edge Computing Environments and their Impact on Physical Infrastructure Selection](#)
White Paper 278

-  [Essential Guidance on DCIM for Edge Computing Infrastructure](#)
White Paper 281

-  [Browse all white papers](#)
whitepapers.apc.com

-  [Edge UPS Fleet Management Comparison Calculator](#)
TradeOff Tool 27

-  [Browse all TradeOff Tools™](#)
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm