



THE END-TO-END RELIABILITY FORUM™

THE POWER OF WE

A COLLABORATIVE APPROACH TO POWERING AI AND DATA CENTERS

Page 51




CONFERENCE KEYNOTE



Kevin O'Leary

Investor on ABC's Shark
Tank & Chairman of
O'Leary Financial Group



The Impact of UPS Autonomy on Reliability and Safety

Longer autonomy does not increase data center reliability

by Stephen Fairfax

The Impact of UPS Autonomy on Reliability and Safety

Uninterruptible Power Supply (UPS) autonomy time — the interval that a UPS can support its critical load on battery power alone — has historically been determined by the characteristics and limits of electrochemical lead-acid and lithium batteries. Hybrid supercapacitors (HSCs) (sometimes called superbatteries) offer the potential to match UPS autonomy times to operational requirements while improving reliability.

This paper discusses the reliability implications of UPS autonomy (backup) time. Longer autonomy does not increase data center reliability. HSCs offer superior component reliability, while an optimal autonomy time increases system reliability in operation. Human factors tend to decrease data center reliability when autonomy times are longer than necessary.

Required UPS autonomy is typically 30 seconds or less. Some designers assume that increasing battery autonomy time will increase the reliability of the data center's electrical system.

MTechnology, Inc. (MTech) performed a number of detailed reliability studies on this topic^{1,2,3}. Our calculations showed that there was no reliability benefit to UPS battery autonomy times longer than the time required to start and synchronize standby generators. If the generator does not start and run normally, practical battery autonomy times are insufficient for operators to reliably and safely determine the cause(s) of the failure and repair the system.

Reliability is not the only consideration; personnel safety is paramount. Providing longer than required battery autonomy times can create an expectation that operators will attempt to correct a generator fail-to-start (FTS) event. A 5-minute battery places operators under extreme time pressure after a utility failure and a generator FTS. The expectation of attempting a repair under these conditions creates significant hazards to operating personnel and ultimately lowers the reliability of the data center.

Human factors studies demonstrate that even the best-trained and frequently drilled operators are prone to mistakes when placed under extremely stressful


DATA CENTER SOLUTIONS TO ENSURE POWER TO RUN

**We give life to what the OEM
wants to throw away**

- Power Distribution Equipment Repair & Remanufacturing
- Expansive Inventory
- Electrical Safety Solutions
- NETA-Certified Field Services
- Testing & Commissioning
- Disaster/Emergency Repair & Restoration



GroupCBS.com



NEW

grid | Xtreme VR

HOPPECKE
POWER FROM INNOVATION

Ultra Pure Lead Acid Battery for
your UPS Data Center





conditions. The likelihood of a corrective action being the proper or best one is low.

Data center operators are neither trained nor subjected to the sort of realistic simulator training required for airline pilots and nuclear power plant operators. Without this training, stress during outages is increased, and the odds of mistakes are even higher.

Mistakes with multi-megawatt equipment can be costly and dangerous. Errors driven by the expectation of operator intervention could result in injury to operators and/or lasting damage to the data center's electrical systems. These are not theoretical considerations; I have personally investigated many data center failures where operator haste led to inappropriate and harmful actions.

Required battery autonomy

The minimum battery autonomy in modern data centers is relatively short. Failures within the data center's electrical distribution system are generally managed by automated control systems, with AC power re-routed to the affected UPS from redundant pathways. These automated responses take a few seconds at most.

Standby generators and redundant utility sources can respond to utility power failures in 15 seconds or less. For example, a 25 MW array of 40 standby generators in Yuma, Arizona, starts, synchronizes, and provides power to the utility in less than 10 seconds in response to grid frequency fluctuations. Some UPSs with kinetic energy (flywheel) storage use specialized generators that start and begin providing power in less than 4 seconds. Other flywheel-based UPSs can supply power for 10-15 seconds while conventional generators start.

Certain large low-speed engines require some minutes after starting before full load can be applied, but these are not common in data centers.

Very short utility interruptions and power fluctuations are common⁴, so most centers delay starting generators for a few seconds after utility power failure. Including this delay, generator starting and providing power to the UPS and other loads is typically complete in less than

30 seconds. Installations with large arrays of parallel generators and less-than-optimal generator paralleling controls may require up to one minute.

Decades ago, when a 10 MW data center was considered extremely large, some designs incorporated UPS battery banks that were rated for 15 minutes or longer. The justification was that in the case of a combined utility outage and generator failure, 15 minutes provided time for a graceful shutdown of critical processes and loads. As e-commerce and data centers grew in size and power, eventually evolving to serve global markets seven days a week, 24 hours a day, graceful shutdowns became impractical or impossible in nearly all applications.

The choice of 5-minute or 15-minute battery autonomy was driven more by limitations of battery technology than autonomy time requirements. Batteries have a maximum discharge rating, and the limits of chemical reactions make it difficult to efficiently discharge the battery's stored energy in less than 5 minutes. While some UPSs have utilized heavy truck batteries with a 30-second cranking rating, those batteries were far more efficient when discharged for 5 minutes and more efficient still when discharged over 15 minutes or an hour.

In contrast, flywheel batteries can provide high efficiency and power density for 10- to 15-second autonomy but tend to become heavy and prohibitively expensive for much longer discharge times.

Data centers using battery-backed UPSs generally do not require more than 30 seconds of UPS autonomy. HSCs offer a solution that matches this requirement.



Human factors

In most cases, immediate operator action is neither required nor desired after an electrical failure. The redundant systems and automated controls respond to UPS, distribution, and utility failures. When automated responses fail, operators are confronted with multiple factors that degrade their performance and significantly increase the odds that they will make a mistake.

HyAxiom, Inc. provides high-efficiency, resilient power solutions.

Fueled by natural gas or hydrogen, our stationary fuel cells deliver clean, reliable, and scalable energy, reducing reliance on the grid while lowering overall emissions.

With a focus on uptime, efficiency, and sustainability, HyAxiom helps data centers meet growing energy demands with cost-effective, clean power.



integra

Leading the Way

Building Trust.
Maintaining Excellence



integra

Your Premier National Testing, Adjusting, Balancing, and Commissioning Service Provider for Data Centers!

Integra Testing Services is the nations largest premier testing, adjusting, balancing, controlled environment, and commissioning service company. Integra specializes in data center services such as air and water balancing, building enclosure testing, vibration analysis, mechanical commissioning, electrical commissioning, and many more!

These services ensure proper temperature, humidity, and pressurization within the data center. This is critical in ensuring the operation, longevity, and mitigation of downtime in a data center. Check out 60+ service offerings at integratesting.com to learn more. For assistance reach out to us at info@integratesting.com or (440) 266-5200.

An operator confronted with failure to provide AC power to UPS after the 10-30 seconds required to start the generators is facing one or more of three potential failures: the utility has failed, the generators have failed to start, and/or the control system has malfunctioned.

The performance of operators in military aviation, civil aviation, nuclear power plants, and other critical activities has been studied extensively. Swain and Gutman⁵ authored a classic text that remains relevant over 40 years later. Technology changes, but humans are primarily the same.

Consider the extreme case where utility power fails but no generators start and all UPSs are on battery. While this is an unlikely event, it can and has occurred. It is similar in severity but with even more time stress than a nuclear power plant operator confronted with a large loss of coolant accident (LOCA). Swain and Gutman discussed this scenario.

Following a LOCA, human reliability would be low, not only because of the stress involved, but also because of a probable incredulity response. Among the operating personnel the probability of occurrence of a large LOCA is believed to be so low that, for some moments, a potential response would likely be to disbelieve panel indications. Under such conditions, it is estimated that no action at all might be taken for at least one minute and that if any action is taken, it would likely be inappropriate.

High time-stress only adds to the Human Error Probability (HEP)

When the time available to take corrective action is severely limited ... given that an error has been made and recognized as such, or that corrective action has failed to have its intended effect, the error probability for the following attempted corrective action doubles. Thus, if one working under severe time stress attempts a task that has a HEP of 0.1 and fails on the first attempt, it takes only three more unsuccessful attempts to reach the limiting case of an error probability of 100% on the following attempt. This limiting condition corresponds to the complete disorganization of the individual...

Longer than minimum required battery autonomy encourages operators to respond even after the system has acted properly following a component failure. I have investigated numerous data center power failures where UPS power was supplied to all critical loads after a utility power failure, but the response was not entirely as expected.

A typical case is when one UPS (out of many in the data center) failed to receive AC power but had its load transferred to reserve or redundant units. The data center is in a stable, if not normal state. All critical loads are powered, and cooling systems are operational. There is no real-time pressure. The UPS that fails to receive AC



input power and has no critical load can be investigated and repaired at leisure. Operators rushing to restore AC power to that single UPS caused a cascading failure that affected part or all of the critical loads.

Operators will require about 1 minute to recognize that a UPS is not receiving power after a utility failure and generator start. A 5-minute battery creates an expectation that when confronted with a malfunctioning high-power electrical system, operators will diagnose the problem and manipulate the controls in an attempt to correct the problem in less than 4 minutes. The most probable outcome is operator error. This means that longer than the minimum required autonomy times perversely reduce the reliability of the data center power system in actual operations.

Safety

All data centers develop safety programs that include training in electrical safety, the proper use of Personal Protective Equipment (PPE), and detailed written procedures for safe operation of the multi-megawatt electrical plant.

After a utility power failure and generator failure, certain recovery procedures can be conducted from the control room. Data center control systems are designed with automated responses to generator or distribution system failures. Control room operators cannot react as quickly as these automated systems. Only after the automated responses fail can operators intervene.

Operators in this scenario confront a very challenging problem. Not only have there been two or more failures, but the corrective actions from automated systems have failed. Human factor studies suggest that operators are likely to make mistakes. Misoperation of power systems controlling megawatts of power can be dangerous, especially if controls are not functioning normally. There is a direct threat to safety from operator errors.

If operators enter the generator room or enclosures the situation is more grave. Modern data centers are large and control rooms and operator stations are not usually located close to the noisy, hot generators. Running while racing a 4-minute clock violates multiple data center safety rules. Even if the operators get to


Experience Data Center Design *Excellence*

Kimley-Horn's
Deep Bench of
Expertise Drives
Innovation and
Performance.

Kimley»Horn

WE HAVE THE POWER TO PROTECT OUR FUTURE

When lighting is value engineered, it often comes at the expense of efficiency.
A 0.2 kWh inefficiency per fixture can compound into massive
megawatt-hour energy consumption across an operation.

LANTANA  LED

Reach out to sales@lantanaled.com for your complimentary lighting analysis today and see how much you could save.
After all, the cost of wasted energy is more than just dollars!

the room in time, out of breath and full of adrenaline, they will be tempted to enter without taking the time to don appropriate PPE or retrieve and execute relevant checklists.

Once inside the generator room or far more crowded generator enclosure, operators are facing a malfunctioning multi-megawatt component and electrical system. The one thing they know for certain is that the system is not operating as intended. This means a fundamental mismatch exists between the operator's mental model of how the system operates and the actual system state. This creates a potentially hazardous situation and greatly increases the potential for misoperation.

Reliability

MTech's studies conducted over 26 years consistently found no reliability benefit to battery autonomy times longer than required to start the generators. These findings were consistent with failure investigations where major failures occurred in systems with battery autonomy times as long as 45 minutes. As discussed in the section on Human Factors, this behavior has been observed in multiple industries.

Operators acting under severe time stress and confronting a system that has not responded normally are prone to errors of both omission (not noticing important indications) and commission (attempting a corrective action that does not work and often makes the problem worse).

Legacy lead-acid UPS batteries are typically rated for 1200 discharge/charge cycles and are generally replaced every five years. MTech's calculations 1,2,3 utilized a battery model with continuous cell monitoring and monthly testing. The base case model used a 99% probability of detecting failed cells before they caused a failure during a utility outage. These are optimistic assumptions, yet our calculations showed that non-detectable and detectable battery failures account for more than 83 percent of all double-conversion UPS failures during short utility outages.

Lithium batteries have largely replaced lead-acid batteries in new UPSs and in many retrofits. Like lead-acid batteries, discharge times less than 5 minutes reduce lithium battery efficiency and possibly lifetime.

Reliability data for lithium batteries (mean time to failure, failure modes, detectability of each failure mode, cycle life, etc.) are relatively scarce. This is due in part to the rapidly changing technology and chemistries available. MTech searched in vain for authoritative lithium battery reliability data for over seven years and consulted with National Laboratory personnel who commiserated with us, as they had been equally unsuccessful.

The rapid proliferation of utility-scale Battery Energy Storage Systems (BESS) provides a substantial population of cells and battery strings suitable for developing reliability data, but to my knowledge, there are few if any, published studies. There are numerous examples of failures in these systems.

The newest technology is hybrid supercapacitors, which incorporate elements of electrostatic energy storage (capacitors) and electrochemical storage (batteries.) These products are being developed for applications where very large numbers of cycles are anticipated. The Musashi Energy Solutions' (MES) prismatic hybrid supercapacitors (HSCs) are rated 3.8 volts, over 1 million 100% charge/discharge cycles, and a 15-year service life.⁷ The operating temperature range is -30°C to 70°C. These units can be charged and discharged in 10-30 seconds, far faster than conventional battery technologies.

The higher cell voltage means roughly half the number of series elements compared to a chemical battery. The cycle rating is roughly 1,000 times higher and far greater than anticipated in nearly any data center. A rating of 100% depth-of-discharge decreases the size and cost of the battery bank. The 15-year service life matches the design life of modern data centers. In addition, the wide operating temperature range could translate to improved reliability in a controlled environment or



reduced cost by eliminating climate control systems. The high discharge current rating of the new products makes them well-suited for a 10-second or 30-second UPS autonomy time.

These characteristics suggest that HSCs should offer higher reliability than chemical batteries while providing optimal battery autonomy time. Field reliability data from deploying these devices to support extremely peaked AI computing loads should be available in the near future.

Summary

The practice of specifying UPS battery backup times that are much longer than the time required to start generators is the product of at least two factors that are no longer applicable in modern data centers:

LEGENCE

DEMANDING CONDITIONS DEMAND SMARTER SOLUTIONS

Speed, scale, and precision are critical in today's data center market. Legence offers custom prefabrication solutions to accelerate your deployments anywhere in the U.S. With over 450,000 square feet of fabrication capacity, we're equipped to handle projects of any size and complexity.

PARTNERSHIP YOU CAN TRUST

Legence is your trusted partner for optimizing mission-critical building systems. We specialize in complex HVAC, process piping, and other MEP systems, ensuring the seamless operation of data center facilities. Our unmatched expertise, combined with a collaborative approach, ensures your mission-critical systems perform flawlessly.

WeAre**Legence**.com

100+

Data Center clients partner with Legence for their high-performance building needs.

Contact us for
a consultation.

SCAN ME



JACKSONVILLE TEXAS



Dry cooler



Minimized
water consumption



Minimized
power consumption



Adiabatic technology

We will be at **24 7XChange** June 8 -11 2025, JW Marriott Grande Lake, Orlando, FL



- The controlled shutdown of critical IT processes and equipment is no longer practical.

- Chemical batteries are generally not rated for discharge times of less than 5 minutes.

The argument that 5-minute or longer autonomy allows operators to take corrective actions does not withstand a careful examination of human factors. The outcomes are often negative when decisions are made under high-time pressure and amid multiple failures. This analysis shows that longer autonomy times can reduce data center reliability. Longer than necessary autonomy times also increase the potential for equipment damage and injuries to personnel. Hybrid supercapacitors offer the potential to match autonomy time to actual requirements driven by time to start generators or switch utility feeds. Field demonstration of cycle lifetime and reliability of these new products will offer data center designers and owner/operators the possibility of utilizing UPS energy storage that will last the facility's lifetime with high reliability and minimal maintenance.

¹ <https://powertechniquesinc.com/wp-content/uploads/2015/08/Active-Power-WP-103-Reliability-Assessment.pdf>

² https://www.7x24exchange.org/prescall/uploads/contentshare/OQDY-WP115-Mitigating-Risk-of-UPS-System-Failure_web.pdf

³ MTech conducted more than 200 studies of data center double-conversion UPS from 1997-2023; most were performed under NDA and not published.

⁴ https://powersecure.com/wp-content/uploads/2020/05/CaseStudy_MCAS_v07v2_8.5X11_1-30-2020_compressed.pdf

⁵ *Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications*, A. D. Swain, H. E. Guttman, NUREG/CR- 1278, August 1983.

⁶ <https://musashienergysolutions.com/products/#cells>

⁷ <https://musashienergysolutions.com/flex-and-musashi-partner/>

Stephen Fairfax is Data Center Consultant at MTechnology. He can be reached at fairfax@oresme.net.

RELIABLE, STANDARD & CUSTOM HVAC PRODUCTS



DATA CENTER SOLUTIONS

The Metraflex Company makes a variety of engineered piping products including energy-efficient **strainers** and **pump connectors**, **flexible connectors**, **expansion joints**, **valves** and more.

Metraflex
for pipes in motion

FEATURED AT 7x24 EXCHANGE

www.Metraflex.com • 312-738-3800

Thank You!

2025 SPRING CONFERENCE CORPORATE LEADERSHIP PROGRAM PARTNERS (at press time)

MARQUIS



GOLD



SILVER



BRONZE



MEDIA

