Reliability Analysis of the APC InfraStruXure Power System

By Stephen Fairfax
Neal Dowling
Dan Healey

White Paper #111

MTechnology, Inc.

APC
Legendary Reliability®

Executive Summary

The APC InfraStruXure product line offers an alternative architecture to the central UPS. MTechnology, Inc. used the techniques of Probabilistic Risk Assessment (PRA) to evaluate the reliability of the 40 kW InfraStruXure UPS and PDU with static bypass. The calculations considered the performance of the InfraStruXure in both ideal and real-world conditions. The study also compared the performance of the InfraStruXure architecture to that of a central UPS serving a hypothetical 500 kW critical load in a data center. The results showed that the InfraStruXure architecture was significantly less likely to suffer failure of all loads in the data center, and slightly less likely to experience failure in any one piece of IT equipment. This paper summarizes the key findings of MTechnology's quantitative risk assessment and discusses their implications for facility managers and designers. The findings are presented first, followed by a description of the methods used to analyze the product, and a more detailed discussion of the results.

Findings

1. The calculated reliability of the APC product is comparable to data published by vendors of large, central UPS.

2. Comparing a hypothetical data center served either by a single, 500 kW UPS or by 14 InfraStruXure UPS demonstrated that the APC approach is significantly less likely to suffer a complete system failure. Failures in equipment common to both approaches, such as the ATS, were the most significant cause of system failure.

3. The redundancy provided in the InfraStruXure definitely improves the product's reliability.

4. MTech analyzed the causes and effects of power module failures, and determined that while power module failures will be observed more often, the increase is more than offset by the benefits provided by redundancy.

5. Detailed consideration of common-cause failure mechanisms and potential catastrophic failure modes that could cause a UPS failure did not result in significant reductions of the calculated product reliability.

6. The results are very insensitive to the assumed utility failure rate. This means that the InfraStruXure performs its intended function, and insulates the customer's equipment from the effects of utility transients and outages.

7. While we did not credit APC's products or components with reduced failure rates, APC's modular design and the associated high volume of product allows the utilization of dedicated manufacturing cells that produce products at lower costs and with fewer defects. APC manufactures five power modules when a non-modular design of the same power rating manufactures one. This enables faster reliability growth of the product line.

8. The use of factory-built distribution wiring in the InfraStruXure architecture confers a significant advantage over field-wired distribution systems for centralized UPS products. Distribution wiring introduces multiple opportunities to introduce wiring defects that may eventually cause loss of power to critical loads. Our analysis of the field vs. factory wiring process found that the probability of defects in field-produced systems is 1,500 times higher than the equivalent factory-built system. We did not credit APC with a reduced failure rate, or penalize the central UPS with higher failure rates, in this analysis.

Overview

American Power Conversion Corp. (APC) retained MTechnology, Inc. (MTech) to perform a reliability analysis of the InfraStruXure 40 kW UPS and PDU with static bypass (InfraStruXure). APC wanted to investigate the utility of Probabilistic Risk Assessment (PRA) techniques in understanding the reliability of the product, identifying potential sources of failure, and evaluating the potential for further improvement in the product's reliability and availability. The InfraStruXure uses redundancy in many components to achieve high reliability, and "hot swappable" technology to enable high availability. APC markets the InfraStruXure product line as a scaleable, "pay as you grow" solution sized to serve one or more rows of equipment racks. This strategy is an alternative to the use of one large, central UPS to serve an entire data center.

MTechnology, Inc. (MTech) has been applying formal, quantitative reliability analysis techniques to the 7x24 marketplace since 1997. MTech has adapted PRA techniques to the study of achieving high reliability and high availability in the 7x24 environment by leveraging the decades of experience and millions of dollars invested in reliability growth of the US nuclear power industry. MTech's clients include electric utilities, manufacturers, design firms, and critical facility owners and operators.

MTech performed a detailed analysis of the InfraStruXure 40 kW UPS and PDU with static bypass. Fault tree analysis was the primary technique, supplemented by event tree analysis and Bayesian updating to determine component failure rates from sparse data.

The study included analysis of the product in isolation, analysis in a typical data center environment, and a comparative reliability analysis against a typical central UPS in the same data center. The analysis included a detailed review of the electrical and mechanical design, engineering verification and validation testing, manufacturing techniques, and the performance of the units in actual service. MTech interviewed APC's design engineering team, the product support team, sales and service databases, and senior management. Several of the firm's engineers traveled to APC's design center in Kolding, Denmark, and worked closely with the product's designers and support engineers to verify and extend the mathematical model constructed to evaluate the reliability and availability of the product.

Introduction

The growing reliance on information systems that operate 24 hours per day, 7 days per week has spawned a rapidly growing and developing industry that supplies products and services to this relatively new marketplace. Once dominated by large financial institutions and large mainframe computer-based corporate databases such as airline reservation systems, the need for and utility of on-demand information services now reaches into essentially every business and every office in the world.

Reliability of the electric power supply is an essential foundation for these on-demand services. Electric utility networks are incapable of supplying power of the requisite quality and reliability. The protection

systems in all electric networks are designed to interrupt power in order to protect people and equipment from the effects of accidental contact with energized conductors or equipment failures.

The product of choice for improving the reliability of electric power is the Uninterruptible Power Supply, or UPS. The UPS conditions utility power so that essentially perfect voltage and current is supplied to the protected equipment, called the critical load. The UPS also includes batteries (or other energy storage devices) that keep power flowing to the critical load when the utility fails. UPS have been manufactured for decades, and APC has been producing them since 1984.

While there are many past and present designs for UPS, the InfraStruXure and most of the products aimed at the data center marketplace utilize the double-conversion architecture, as shown in Figure 1. Utility AC power is rectified to DC. The DC bus connects the rectifier to a battery (typically composed of multiple series and parallel strings, not shown here) and to the inverter. The inverter synthesizes an AC voltage free from the effects of spikes, sags, harmonics, and brief utility outages.

The inverter output is connected in parallel with the output of a static bypass switch. The bypass is closed if the rectifier or inverter fails, or if an electrical fault in the critical load requires more current than the UPS can supply.

The collector bus connects the bypass and one or multiple inverter outputs. Some UPS employ multiple inverters to achieve higher power ratings or to provide redundancy. Systems with multiple inverters typically have multiple rectifiers; the assembly of a rectifier and inverter is called a power module.

Items in Figure 1 drawn with dotted lines are optional; for example, some but not all UPS installations utilize parallel battery strings or multiple power modules.

Figure 1 *– Double conversion UPS*

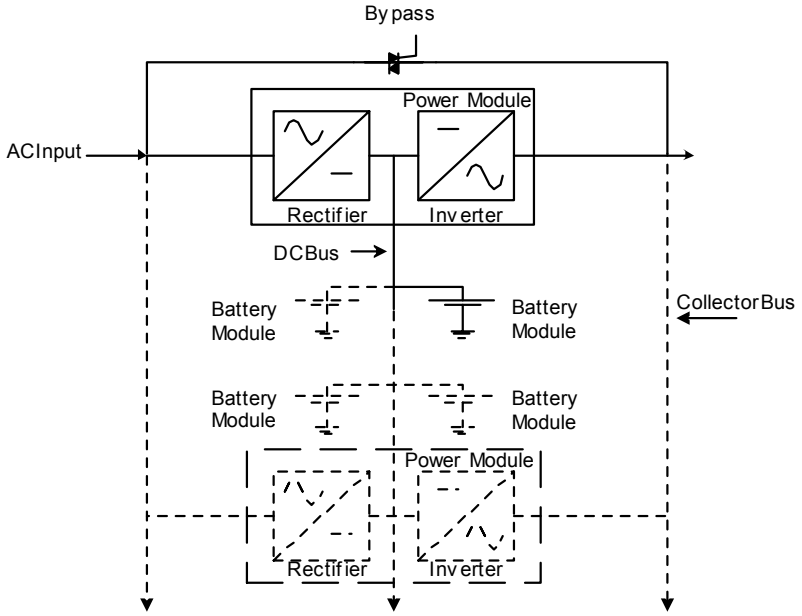

The double conversion UPS architecture pays an efficiency penalty as both conversions result in some losses. The design has gained wide acceptance because it requires no active switching or other positive actions when utility power fails. The battery begins to discharge as soon as the DC bus voltage drops, and the inverter works as it did when utility power was available.

UPS work very well, and a relatively few standard design approaches have come to dominate the field. One problem with introducing a new product to any marketplace is that of demonstrating to the customer that the new product will work as well as, or better than, the older solutions. This problem is acute in the UPS marketplace, as nearly every installation is a custom design, with external equipment, conditions, and operating practices that make it very difficult to compare performance among different installations. The success of the UPS solution presents an additional barrier: failures are relatively rare, and reliable sources of data on failures among various models are rare or non-existent. The reliability claims of most major UPS vendors are equivalent to less than one failure per century of operation, but few data centers or UPS are operated for more than 20 or 30 years.

It is possible to introduce a new product and then observe the number of failures in order to determine its reliability. This approach has many drawbacks. First, the customer becomes the subject of an experiment. Second, since even a poorly designed or manufactured unit might not fail very often, it may take months or, more probably, years of observation before statistically significant differences could be demonstrated. Third, achieving reliability in critical systems (e.g. airplanes, anti-lock brakes, and telephone switches) requires observation of large fleets of essentially identical components over long periods. The present UPS marketplace has evolved to include a substantial number of specially designed data centers. Each data center has a unique design, and the UPS within that data center are exposed to unique operating environments and management practices. UPS vendors have naturally responded with an ever-growing array of custom and customizable solutions that can meet any conceivable design specification for the next custom-built data center.

Surely it would be more efficient and less costly to employ some means to learn about the reliability of a new product before subjecting thousands of customers to potential mistakes that compromise reliability. Further, it would be extremely useful to know which of several competing proposals offers the best reliability for the least cost. The product's designers would very much like to understand which components and sub-systems are most important to the product's overall reliability. The product support engineers, charged with tracking the products' performance in actual use and quickly identifying and implementing changes necessary to correct deficiencies or defects, would benefit from a road map identifying components most likely to fail. Deviations from the predictions of the road map would identify new areas for more intensive investigation and possible remedial action.

Probabilistic Risk Assessment, PRA, was first developed as a response to the exasperation of early rocket engineers, who grew frustrated with the seemingly endless litany of reasons for their cherished vehicles to fail. Mathematical analysis quickly showed that, in a highly interconnected system such as a rocket or a data center, the old adage that a chain is only as strong as its weakest link is no longer true. The chain comes to

resemble a net, one with many weak links and undiscovered threads linking one area to another. Failures in one part of the net place new and different stresses on other parts, which are then more likely to fail. The result is an environment where even minor upsets start a series of cascading failures that end with complete failure of the system.

PRA was applied on a large scale to the US nuclear power industry, first as a means to address public concerns regarding safety. After the events of Three Mile Island (TMI) threatened the viability of the entire multi-billion dollar industry, PRA techniques were embraced and extended to include not only design choices but also operating and maintenance decisions and the effects of management practices. The results have been gratifying. Not only have there been no more incidents resembling TMI, but the fleet of 103 power plants now produces 20% more electricity annually than they did prior to TMI. It is becoming routine for plants to operate for 18 months or 24 months without a single forced outage, shutting down only when it is necessary to refuel. PRA has also informed maintenance practices and demonstrated that many "best practices" in fact unnecessarily increased the risks of component failures and accidents.

PRA is a powerful tool when applied carefully. The process of building the logical model results in a comprehensive review of the decisions, features, and assumptions that shaped the product. The mathematical nature of the calculation limits the appeals to experience and other common logical fallacies that tend to dominate qualitative evaluation of reliability. All too often, claims of "twenty years experience" are roughly equivalent to one year's learning followed by nineteen years of doing the same thing over and over again.

MTech's PRA calculations are routinely challenged, particularly when our client believes that the reliability of the system in question is much higher than our calculated value. A review of the logic in the mathematical model will disclose whether there are any flaws or misunderstandings on the part of either party regarding the functional behavior of the system. Changing the component failure rates to the client's preferred values almost never results in significant changes to the final result. Nearly all UPS contain redundant paths such as the bypass switch. System reliability should not be very sensitive to component failure rates in redundant designs.

The value of PRA is due both to the quantitative results, and its ability to identify the relative contribution of each component to failure. Without quantifiable, reproducible calculations of each component's role in the systems' success or failure, it is simply impossible to allocate resources rationally, much less optimally. The traditional reliance on redundancy to characterize system reliability illustrates this point. Many data center designs are characterized as "N+1" or "N+2" or even "2N" or "2N +1" designs. The implication is that if N components are required for success, there is one, two, twice as many, or even twice plus one as many units available. But clearly not all redundancy makes the same contribution to reliability. Redundant standby generators, with a 1% failure to start per demand, will contribute far more to reliability than redundant dry-type transformers, whose failure rate is so low that the money spent on redundant units can almost invariably be spent to better effect elsewhere. Absent the ability to determine the quantitative contribution of each

component, redundant or not, designers and buyers simply can not make informed decisions regarding the best use of always scarce financial and other resources. PRA is a powerful tool to answer these questions.

There are fundamental questions regarding redundant designs. While redundancy can in principle increase reliability by allowing individual components or subassemblies to fail without causing the system failure, there are significant costs and potentially serious drawbacks. A redundant system has more components, and in general systems with more components will experience more failures. (Twin-engine airplanes experience roughly twice as many engine failures per hour of operation than comparable single-engine airplanes.) There must be very reliable mechanisms in place to identify the failed component and isolate it from the system, or the benefits of redundancy are lost, while the number of component failures has increased.

Some failure modes can affect multiple components simultaneously. Such common cause failures place a significant limit on the benefits of redundancy. Design defects, manufacturing defects, defects introduced during installation, maintenance, or repair, all can result in failure modes where multiple, supposedly independent units fail, often causing the entire system to fail despite the redundant design. Catastrophic failures of some components can damage connected or nearby equipment and cause system failure despite redundant design.

MTech used PRA techniques and software adapted from the nuclear power industry to analyze the InfraStruXure product line and compare its performance to a traditional system. The resulting mathematical models were used to answer some of the key questions. The InfraStruXure utilizes redundancy in nearly all components as a means of achieving high reliability. MTech's analysis showed that there are both costs and benefits to redundancy, and that some sub-systems benefit less from redundancy than others.

Reliability and Availability

This study was primarily concerned with the reliability of the products. Many vendors prefer to discuss availability. The distinction is subtle but important. Reliability is the probability that a system will operate as intended for a given period of time. The time period, also called the mission, must be specified. A 747 is extremely reliable once it takes off; the probability of making a successful landing with no damage to equipment or passengers is much greater than 99.99% for flights of 14 hours or less. For flights of 36 hours, the reliability of the 747 is zero, as it will always exhaust its fuel before the mission is completed.

Availability is the fraction of time that a system will be operational. Availability can be associated with a mission time, or can be expressed as the long-term availability, which is the asymptote of the availability as time goes towards infinity. Availability requires knowledge of the time required to repair the system after a failure. Given equal failure rates, systems that are repaired quickly will spend more time in the operational state than systems that require lengthy repairs, and so will have higher availability.

There are valid reasons to calculate and understand each metric, but MTech believes the reliability, or more specifically Unreliability, the probability of failure during a given mission, is the more valuable metric for data center owners and operators. A system with very high reliability but long time to repair might show the same, or lower, availability as a system that fails frequently but is rapidly returned to service. The financial and other damages associated with a data center power failure are very large no matter how rapidly power is restored; most owners will prefer a more reliable system if they have the information necessary to make an informed choice.

The primary reason to use probability of failure (unreliability) is that our end-user customers find it the most useful metric. Few firms have substantial experience in the mathematical techniques of probabilistic risk assessment, but executives and managers routinely juggle competing proposals that have various degrees of risk. Risk is a function of probability and consequences. Many firms purchase products such as insurance or disaster recovery programs based upon their assessment of risk, the probability of suffering a loss multiplied by the amount of damage they anticipate from such a loss. Most firms that operate data centers will suffer substantial losses in the event of a single outage, and they need to know the likelihood of such an event in order to make informed decisions regarding additional investment or other means of mitigating the risk.

A second reason to use probability of failure is that the metric is constant across the organization. APC has developed a 4-level hierarchy that describes the interaction of various systems in a typical firm. The top level encompasses people, the next processes, the third level information technology, and the bottom layer infrastructure, including electrical power. The experience of a single failure will yield dramatically different results for availability of each layer.

Consider a hypothetical firm that experiences one outage of the UPS system in 10 years of operation:

- The infrastructure layer restores power in 10 minutes. They can then calculate their availability: A = 87599.8 / 87600 = 99.9998%. The infrastructure layer can claim “five nines” availability
- If IT restored applications in 12 hours, then their availability calculation will be A = 87588 / 87600 = 99.99%, and they can claim "four nines.”
- If process or application managers repaired damage to database, restored normal work flow in 2 days, they would calculate A = 87552/87600 = 99.95%, and claim “three nines” availability.

Executives who spend two months soothing clients, filing SEC reports, firing subordinates, hiring and training replacements will probably not calculate their availability, but if they did, they might get A = 86160/87600 = 98.4% and would probably resent being informed that their availability was only “two nines.”

The numbers used above are typical of repair times at various points within a facility, and illustrate that the perceived level of availability depends on the observer's point of view. The probability of failure for the entire organization is once per 10 years, at all levels within the firm. If the system is relatively reliable, the probability of multiple failures will be low, and only the probability of first failure need be considered.

Conduct of the study

The study began with an introduction to the InfraStruXure product line and a detailed examination of the UPS and PDU at APC's offices in Billerica, Massachusetts and East Providence, Rhode Island. APC provided engineering documentation and access to field service personnel in order for MTech to review the fleet field history.

MTech developed a fault tree model of the system. A complete description of the fault tree modeling technique is beyond the scope of this paper, but there are many readily available texts and articles on the subject. The first model considered the APC products in isolation. The utility input was assumed to be perfect, and the customer load was likewise perfect. The model explored how often failures internal to the UPS and PDU would cause loss of the critical load.

UPS vendors or buyers rarely acknowledge the fact that all UPS insert a new component, and hence at least one new failure mode, in the circuit serving a critical load. All actions have both beneficial and negative effects on reliability; the goal is to maximize the former while minimizing the latter. The initial fault tree model served to highlight components and sub-assemblies whose performance had a significant impact on system reliability. MTech then began a more intensive inquiry into those components.

The authors traveled to APC's Kolding, Denmark design facility and spent a week in intensive interviews and discussions with the designers of the product. We examined the product development process, design rules, verification and validation testing, review and QA activities, and field service records. The extensive records for this and earlier, similar products were examined. We presented our initial fault tree analysis for review and comment, and revised it to reflect both mis-understandings on our part and to include more detailed information about the causes of failure, particularly common-cause failures.

Some components have more than one failure mode. The collector bus, batteries, controls, and power modules were modeled with two failure modes: normal and catastrophic failures. Catastrophic failures in these components result in failure of the UPS, while normal failures do not as the components are redundant. One form of catastrophic failure is a component that fails but is not identified as failed. The failed component can cause other components to malfunction, or the failed component may continue to deteriorate until a more significant failure occurs. Physically, there are catastrophic failure modes that result in plasma vented to the UPS interior, which will short multiple power and control circuits and cause a load drop. The fraction of catastrophic failures in a component is a key parameter. We began the study with informed judgment that approximately 1% of all component failures are catastrophic. After completion of the initial modeling effort and review of the model with APC engineers, we adjusted the fraction of catastrophic failures to reflect the actual field data. The 1% ratio of catastrophic to normal failures was reasonably accurate. Figure 2 summarizes the results of this phase of the study.

Figure 2 *– Component contribution to failure: InfraStruXure only, No utility failures*

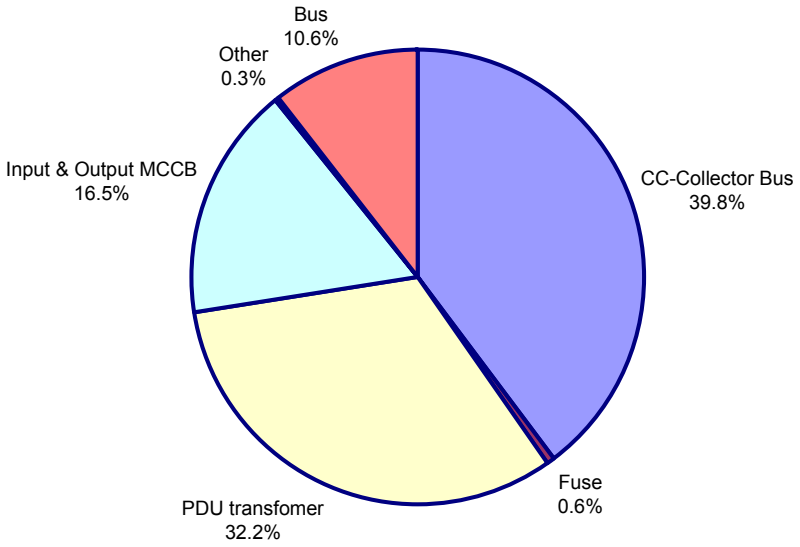

Failures in the PDU transformer and catastrophic failure of the collector bus (the point of parallel connection between power modules and the bypass switch) account for 72% of all expected failures. The input and output molded case circuit breakers account for nearly 17%, despite component failure rates of 1.2 x 10-7 per hour, equivalent to a MTTF of 8.3 million hours.

After reviewing our models and preliminary results with APC engineers, we revised the fault tree model, and then extended it to account for the actual working environment of the product. We included utility failures, generator failures to start, and failures in the transfer switch that selects between utility and generator. We examined the effects of electrical faults in the customer's equipment.

This "real world" analysis of the product resulted in some new questions. Should failure of one branch circuit breaker be counted a failure of the product? While molded case circuit breakers are quite reliable, with an MTTF for spurious trips in excess of 8 million hours, there are so many in even a modest data center that circuit breaker failures become a large fraction of the expected failures.

We utilized the work in a previous study of an actual data center to construct a fault tree model of a "typical" data center distribution system. "Typical" is not a good term to apply to data center designs; there is little standardization, and we cannot claim that the example we selected is average, or below average, or better than average. We merely assert that our model was based on an actual, recently constructed data center. The one-line diagram for this model data center is shown in Figure 3 for the case of a central, 500 kW UPS. Figure 4 shows the equivalent one-line diagram for 14 APC InfraStruXure UPS serving the same load. Note that the source and distribution systems are equivalent for both cases.

Figure 3 *– Large UPS one-line diagram for 500 kW data center*

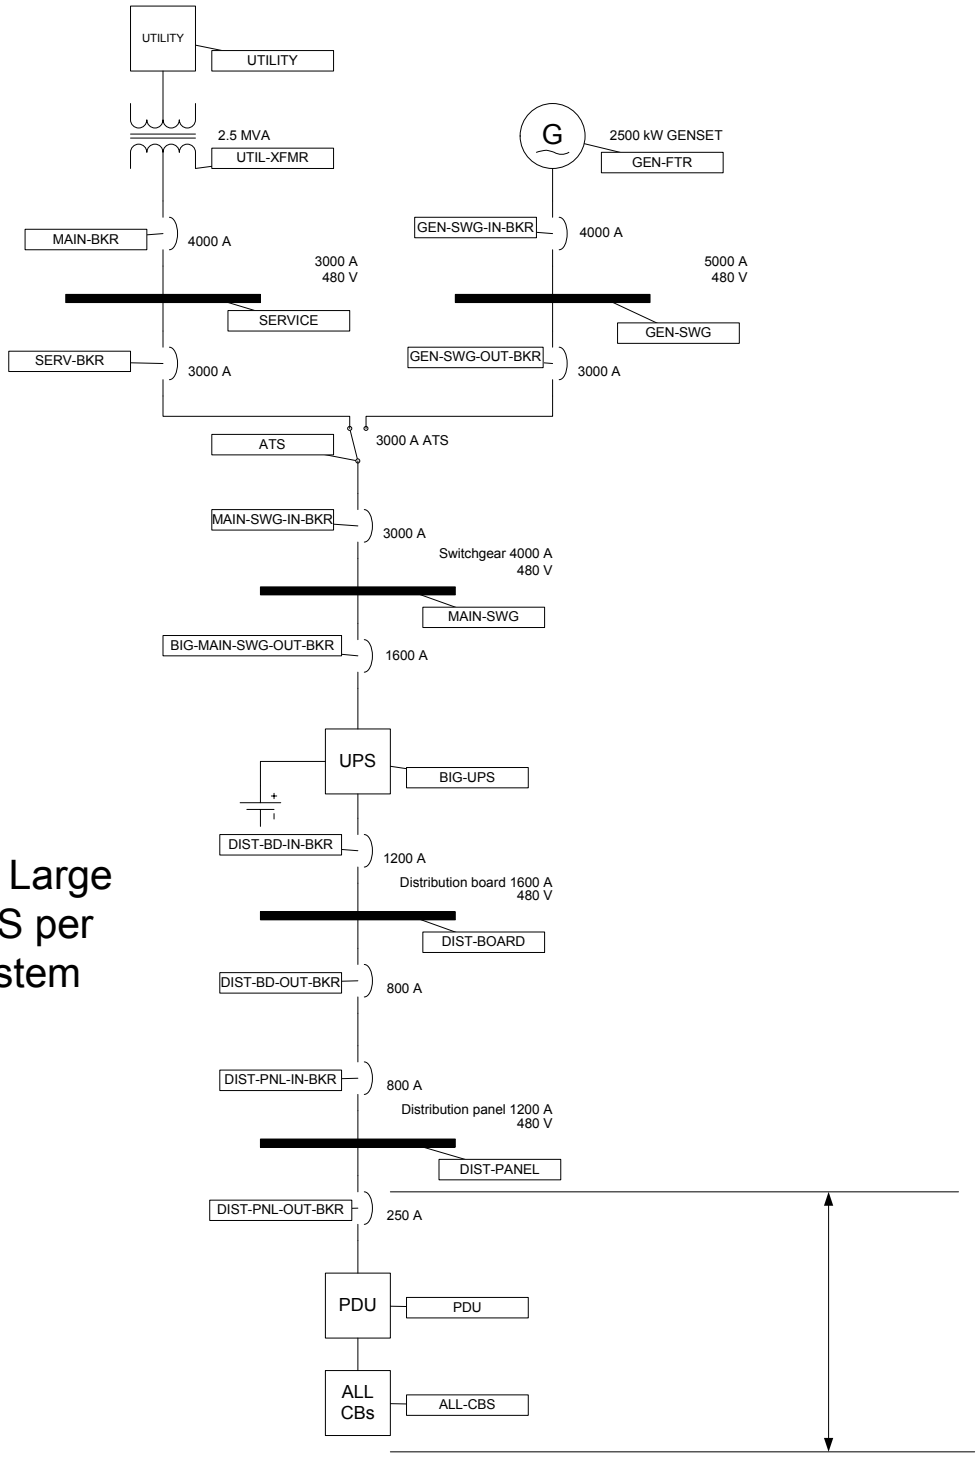

One Large UPS per system

Figure 4 *– APC InfraStruXure UPS in a 500 kW data center*

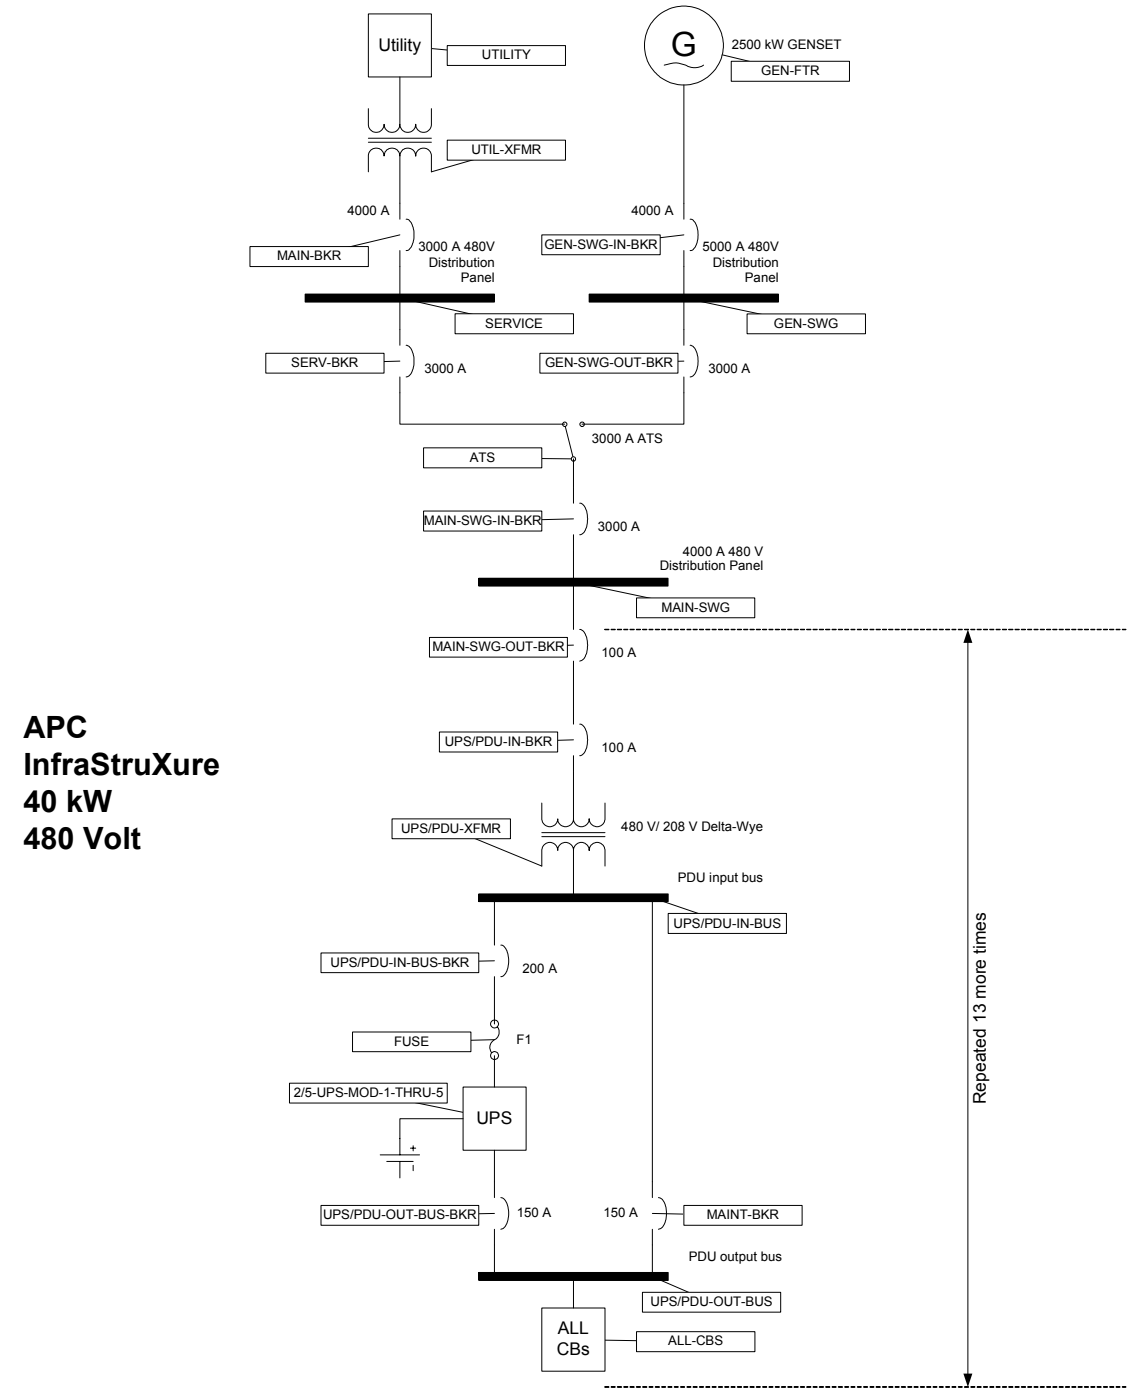

We collected data on failures in large UPS from both vendor's publications and third-party published papers for failure rates in power system components. We utilized fairly common assumptions in determining what constitutes success or failure. Failure of one or more power modules that resulted in a successful transfer to bypass is counted as a success. Failure due to battery exhaustion is not counted as a failure unless the

batteries were exhausted much faster than specified, or failed when the utility failed. We assumed quarterly battery testing and made the optimistic assumption that such testing would identify failed cells or connections with nearly 100% accuracy. Loss of power to the critical load due to operator error was not counted as a failure, although we concluded in a separate analysis that at least some operator errors might be attributed to poor ergonomics or misleading indications.

We used standard statistical techniques to combine the disparate failure rates into an estimated rate that we applied to the comparison "central UPS." The result was a failure rate of almost exactly 1 x 10-6 per hour for a large UPS at the collector bus, which is sometimes identified as the critical bus. This corresponds to a Mean Time To Failure (MTTF) of 1 million hours assuming constant failure rates. Our research included, and our results substantially agree with, the data and analysis conducted by Liebert for its 600-series UPS products. [http://www.liebert.com/support/whitepapers/documents/techmtbf.asp, checked on February 19, 2004]

There are relatively few failures reported in large UPS. The Liebert paper discloses 80 failures in a fleet history of 200 million hours, and correctly points out that units serviced by others might not report failures. Given the small number of reported failures, missing only a few would significantly skew the results. We agree with Liebert's methods and conservatism in reporting MTBF "in excess of 1 million hours." The sensitivity of the final probability of load drop will not be strongly affected by improvements in the UPS failure rate, as will be illustrated by the results shown below.

It is important to emphasize that we did not analyze or model the large UPS with the same detail as we used for the APC products. We merely sought reasonable failure rates for the entire UPS for comparison. We found that the module failure rates published by other vendors were roughly comparable to those we had determined for APC power modules. We also found that common-cause failures, including control system failures and catastrophic component failures, are the most frequent source of UPS failure. This was consistent with our detailed analysis of the APC field data and our modeling of catastrophic failure modes.

We constructed fault trees for two hypothetical data centers; one using a single 500 kW UPS, the other using 14 APC InfraStruXure products to serve the same load. We did not model cooling systems, nor did we consider the effects of partial loading of the UPS output.

The results are summarized in Figure 5, for the InfraStruXure architecture, and Figure 6, for the Central UPS architecture. The InfraStruXure system failure rate (when failure is defined as all critical loads in the data center lose power) was approximately 40% lower than that of the central UPS system.

Battery failure is a significant contribution for the central UPS but negligible for the InfraStruXure. We utilized identical battery failure rates for both systems. The disparity is a consequence of our modeling the central UPS with a single string of VRLA batteries. The InfraStruXure utilizes 8 series-parallel strings (4 positive and 4 negative strings) and will operate with multiple failed strings. The InfraStruXure strings are 196 VDC versus typical central UPS strings rated 400 VDC or higher. The higher voltage strings, with more

cells in series, are slightly less reliable than the lower-voltage strings. If the central UPS architecture were implemented with two or more parallel strings of batteries, battery failures leading to critical load failure would be greatly reduced. Even discounting battery failures, the InfraStruXure failure rate (failure defined as loss of all critical loads in the data center) is approximately 18% less than that of the comparable central UPS architecture. The remaining difference is due to the architecture of each system, not differences in component failure rates.

Both the InfraStruXure and Central UPS designs share common vulnerabilities in the equipment from the utility and generator to the transfer switch, as shown in the one-line diagrams (Figures 3 and 4). The InfraStruXure approach loses power to all loads only when the common power infrastructure fails, for example the main entrance bus fails, or the transfer switch fails open. The probability of all 14 InfraStruXure units failing simultaneously due to internal failures is extremely low. In contrast, the central UPS and bypass can fail, and then all loads will fail. This represents an additional source of failure not present in the InfraStruXure architecture. Note that if the definition of failure is changed so that loss of a PDU constitutes failure, the difference in reliability between the two approaches will be further reduced. Failure of InfraStruXure input or output circuit breakers will cause partial load failure, just as failure of the PDU input circuit breaker in the Central UPS system will cause partial load failure.

A second architectural difference lies in the reduction in the number of circuit breakers that can trip and cause the UPS power to fail to reach all loads. After the transfer switch, there are 5 circuit breakers in the central UPS one-line diagram; 2 on the UPS input and 3 on the output. Output circuit breaker failures cause immediate loss of the critical load. Input circuit breaker failures cause loss of load after the UPS batteries are exhausted. While it is theoretically possible to repair a circuit breaker that trips while carrying less than rated current in the 10-45 minutes of UPS autonomy typically available from battery banks, the probability of doing so without causing other failures is very small. The InfraStruXure architecture has only 1 circuit breaker after the transfer switch that will cause all critical loads to fail.

Figure 5 *– Component contribution to failure, InfraStruXure architecture*

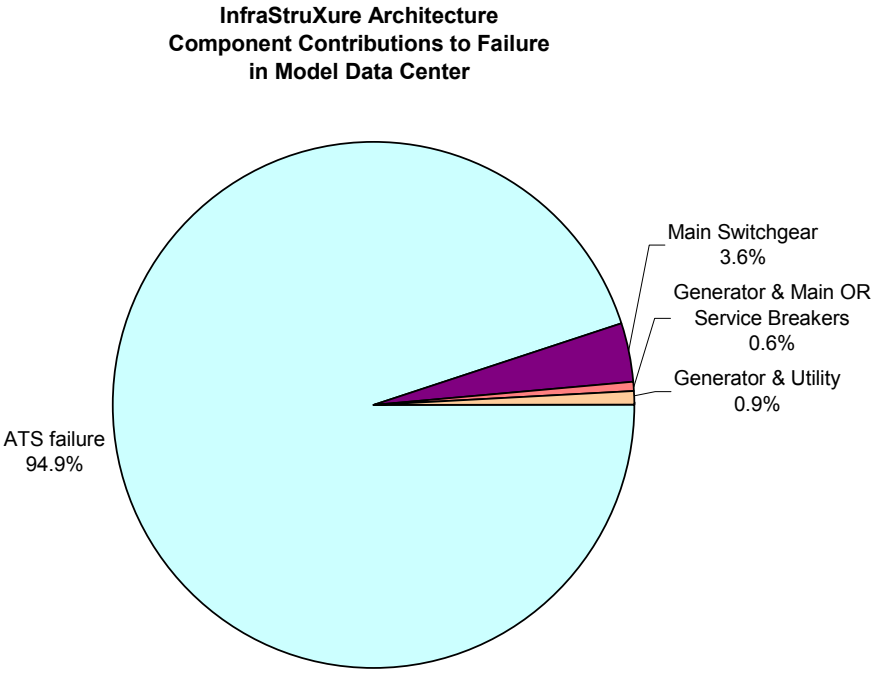

Figure 6 *– Component contributions to failure: Central UPS architecture*

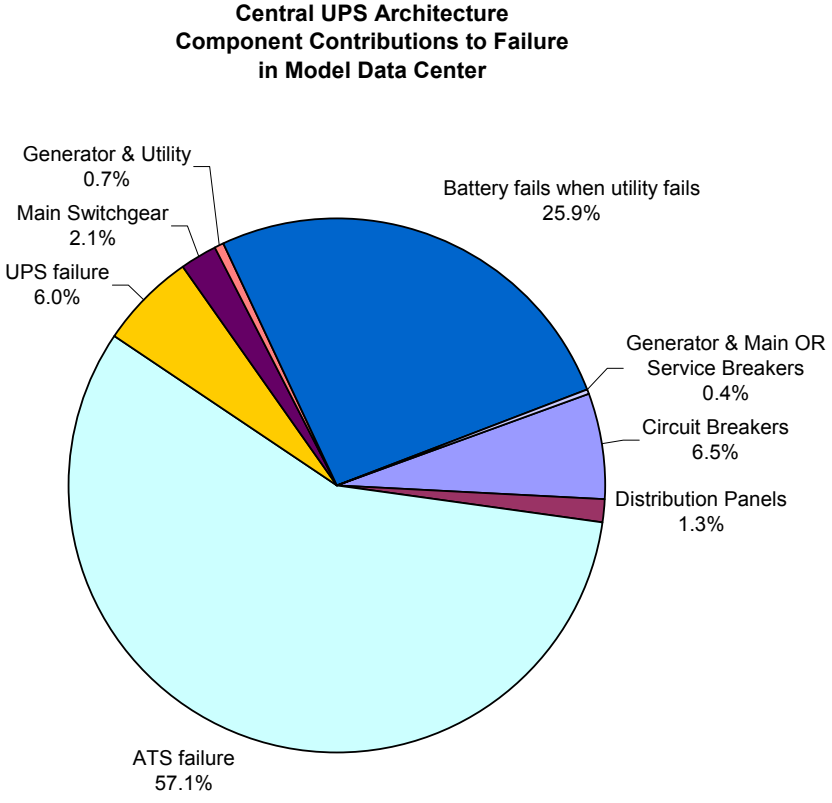

We examined the effects of operator error on UPS failures and concluded that there was not a significant difference in the APC versus large UPS. We note that operators in this hypothetical data center would perform any manipulations of the InfraStruXure UPS 14 times more often than with the large UPS, which standard performance shaping factor analysis would score as a significant decrease in the probability of error in any given operation. The effects of errors, to the extent that they are immediately apparent, would also be reduced in the APC approach.

We found that there was a substantial sensitivity to the failure rate of the distribution system, which includes the wiring and protective devices between the output of the UPS and the critical load. This problem is common to both the APC and central UPS approach. This led us to a detailed examination of the APC distribution system factory production techniques. We reviewed the process controls and quality assurance practices used by APC in producing the factory-wired distribution systems, and compared them to typical practices for field wiring of data centers with central UPS. (The use of the term "standard" with respect to anything in a modern data center is problematic.) While our analysis showed a very substantial reduction in the rate of wiring defects for factory-wired distribution systems, we did not account for those effects in the results presented here. This paper is intended to provide an "apples to apples" comparison of architectures as opposed to a competitive analysis of particular products.

Amplified Findings and Discussion

MTech found that a data center employing the InfraStruXure architecture would be significantly more reliable than a comparable data center utilizing a single-module UPS with a single battery string. While the redundant subsystems within InfraStruXure successfully reduced the probability of UPS failure, the effects of external systems common to approaches tended to obscure any difference. The PRA demonstrated that utilizing parallel redundant battery strings in the central UPS would reduce but not eliminate the difference in reliability. Most UPS can support the critical load for only a few minutes on battery power alone. Longer-term protection requires a standby generator or other source of power, and an automatic transfer switch (ATS) to select between the standby and utility power sources. MTech's analysis showed that the performance of the ATS was often the limiting item in achieving higher reliability.

The numerical results showed that the model data center using InfraStruXure architecture was roughly 40% less likely to fail all critical loads than a comparable central UPS in the equivalent data center. Adding a redundant battery string for the central UPS improves reliability significantly, but the InfraStruXure architecture was still 18% less likely to fail in one year of operation. Changing the definition of failure will change these results. If failure is defined to include dropping of any single load, due to a branch circuit failure but not UPS failure, the InfraStruXure architecture was 6% less likely to fail. The reduction from 18% to 6% arises solely from the contribution of spurious trips in molded case circuit breakers, which increases the unreliability of both approaches. The uncertainties in input data, field installation quality, and variations among competitive products are large enough to obscure such a modest advantage.

In the interest of fairness, the models comparing the InfraStruXure to a traditional UPS used identical failure rates for all components. In the interest of gaining an advantage over competitors, APC has embarked on a program to improve the reliability of the key components identified by MTech's analyses. APC has changed the construction of the PDU transformer to eliminate several of the most common failure modes. They have strengthened the collector bus and improved the connections to modules. APC already tests 100% of all circuit breakers before installing them in InfraStruXure systems, while some, but by no means all data centers test branch circuit breakers prior to installation. MTech nevertheless used the same circuit breaker failure rate for both systems in the comparison. Based on MTech's analysis of circuit breaker failure modes, APC is investigating the causes of circuit breaker failure and considering new tests designed to better identify those units most likely to fail.

MTech's analysis demonstrates that differences in architecture distinguish the InfraStruXure from the central UPS, not differences in number of components or component reliability. Customers experience the reliability of UPS products when used in the physical environment of a data center. The reliability of the InfraStruXure in this environment was consistently superior to that of the central UPS architecture, although the differences became statistically insignificant if the definition of failure is changed from loss of the entire data center to loss of a single branch circuit. The results identified areas in both architectures where relatively modest changes in component selection or use could result in major improvements in reliability. To our knowledge, this is the first public disclosure regarding the use of formal, quantitative PRA techniques to guide the development and manufacturing practices of a UPS product.

APC literature suggests that the InfraStruXure architecture offers cost and flexibility advantages so compelling that informed customers would choose the product even if the reliability were no better than that of competing products. The analysis of cost of ownership and related issues is available in APC white papers and will not be discussed further in this document. For more information see APC White Paper #37, “Avoiding Costs From Oversizing Data Center and Network Room Infrastructure” and #6, “Determining Total Cost of Ownership for Data Center and Network Room Infrastructure”.

MTech analyzed the manufacturing techniques used for the InfraStruXure and compared them to those of traditional products. A key distinction between the InfraStruXure and central systems is the factory wiring of the distribution system for all InfraStruXure products. The UPS is only a part of the "whole product" that keeps power flowing to critical loads when the utility fails. In traditional data centers, the UPS sits at the edge of the raised floor, or in another room entirely, and custom-built arrays of conduit and wires bring the power to the racks that house the computers and other critical loads. Electricians and other tradespeople must build these custom systems on-site.

The InfraStruXure distribution system wiring is performed entirely at the factory. MTech analyzed the production process in both factory and field settings. Because APC factory procedures make use of calibration, quality control, specialized fixtures and tooling, and automated inspection of all products, MTech

found a remarkable difference between the numbers of expected defects in factory wired versus field wired systems.

Installing a single branch circuit requires multiple steps, from selecting the proper wire or cable, installing it in the conduit (in field-wired assemblies), stripping the wires, connecting to terminal equipment such as power strips or circuit breakers, marking the finished connections, and so forth. MTech analyzed the probability of making an error at each step of the process for both factory and field-produced systems. The analysis used data and methods from military and nuclear reliability sources. The probability of producing field connections with defects was approximately 1,500 times higher than for factory-wired systems. This difference was not included in the comparative reliability analysis.

Not all defects will result in load drops, and some defects, such as mislabeled switches or wires, may go undetected for the life of a system. Mislabeled components are often discovered during IT equipment changes, when opening a branch circuit breaker causes an unexpected load drop for a different piece of equipment. This is often charged as "operator error" but in reality is a consequence of the latent defect introduced during field wiring. The very large difference between factory and field wiring error rates demonstrate an important lesson useful to all customers: standardized products, and standardized factory manufacturing techniques, result in the highest reliability systems. Custom products, custom wiring, and custom operating procedures will surely increase the probability of both common errors and custom problems.

MTech examined APC's manufacturing process and evaluated the effects of their reliability growth management techniques. There is a valid question regarding the design of the InfraStruXure, with five power modules, compared to a competitive system with one power module. Even though the InfraStruXure is designed to operate with any one of the five modules out of service, one can argue that five modules will experience failures five times more often than a single-module system. Is it possible that the redundant design will result in more frequent loss of power to the critical load?

MTech analyzed the causes and effects of power module failures, and determined that while power module failures will be observed more often, the increase is more than offset by the benefits provided by redundancy.

There is a benefit to using multiple modules that was not quantified in this study. Once the volume of sales for a product can justify the expense of a dedicated manufacturing facility, the rate of manufacturing defects declines very significantly. Dedicated personnel, test fixtures, and experience combine to drive out the potential to introduce defects. A dedicated manufacturing cell will quickly evolve to the point that the most common errors are literally impossible to make, and the less common errors are quickly and uniformly detected and corrected before the assembly leaves the plant.

The InfraStruXure modular design means that multiple power modules are shipped with every unit. APC can therefore switch to dedicated power module manufacturing cells more rapidly than a manufacturer of a

competing single-module product. APC can discover defective components and similar deficiencies in shipped products faster because there are more modules in service than in an equivalent single-module system. Finally, APC can more accurately determine the cause of any module failures because customers can quickly and easily replace failed modules, which are then returned to APC for diagnosis and repair. Most single-module designs are repaired in the field, making the process of identifying the root causes of problems vastly more difficult.

Field repairs represent another potential source of failure that was not examined in this study. Field repairs are much more likely to introduce new defects than factory repairs. All UPS manufacturers subject every production unit to a battery of tests, including high potential testing, functional testing at the extremes of the specified environment, and so forth. Field repairs can rarely be tested with the same degree of rigor or completeness.

Conclusions

The PRA analyses conducted by MTech on the InfraStruXure architecture demonstrated modest reliability advantages when the product is compared to a hypothetical central UPS in an equivalent data center. The difference in probability of failure per year of operation is a strong function of the definition of failure and the design of the central UPS battery bank, but in each case that we examined, the InfraStruXure architecture provided superior reliability. Failures arising from equipment in the data center were always more significant than failure of either UPS.

APC found the study worthwhile even though it did not demonstrate a compelling reliability advantage for the InfraStruXure products in all cases. The original motivations were to better understand the causes of failure and to identify components or processes where additional investments could be expected to produce the greatest improvements in reliability. The study was an unqualified success when measured by these standards.

The PRA study showed APC how best to invest in better components and testing to raise the reliability of the InfraStruXure significantly. Improvements in just 3 kinds of components may reduce the frequency of random failures by a factor of 10 or more. An InfraStruXure built with generic components is not significantly more reliable than a competitor's product built from equivalent components. As a result of this study, APC has devoted additional attention and resources to key components in the areas that PRA identified as causing the vast majority of product failures.

The study also documented the very substantial benefits that accrue from factory manufacturing and testing of distribution system wiring. APC-built distribution wiring is more than 1,500 times less likely to have latent defects than field-built custom distribution systems. The benefits of modular design are also significant and allow APC to utilize the most modern, highest-quality manufacturing techniques, reducing cost while simultaneously increasing product reliability.

APC and MTech learned a great deal during the conduct of this study. Assumptions about the roles and effects of different design choices and manufacturing techniques were questioned and revised when confronted with mathematical analyses and relatively straightforward logical arguments. PRA is a powerful addition to the techniques of design engineers, field engineers, and manufacturing operations. Applied consistently and thoughtfully, investments in PRA will result in more reliable products. The techniques that calculate the reliability of nuclear power plants are an important resource for the rapidly evolving critical power industry, where the economic and human costs of failure are growing rapidly.

PRA provides an unbiased method to compare disparate products and system architectures. The results of careful PRA studies enable customers to make informed decisions regarding which product or architecture is best suited to their particular application. Well-informed customer decisions will enable the industry to most rapidly improve the reliability, cost, and performance of products, and to tailor the reliability and features to customer needs. PRA offers a tool to assist in making rational, informed choices for manufacturers, data center designers, and customers.

References

1. IEEE, Inc. IEEE Guide to the Collection and Presentation of Electrical, Electronic, and Sensing Component Reliability Date for Nuclear-Power Generating Stations. New York: IEEE Press, c1977. [IEEE Nuclear Reliability Data Manual.]

2. IEEE, Inc. IEEE Recommended Practice for Design of Reliable Industrial and Commercial Power Systems. New York: IEEE Press, 1988. [IEEE Gold Book: Power Systems Reliability.]

3. Hale, Peyton and Arno, Robert "Survey of Reliability Information for Power Distribution, Power Generation, & HVAC Components for Commercial, Industrial, & Utility Installations", IEEE Industrial and Commercial Power Systems Technical Conference, 2000.

4. Kumamoto, Hiromitsu, and Henley, Ernest J. Probabilistic Risk Assessment and Management for Engineers and Scientists. 2nd Ed. New York: IEEE Press, 1996.

5. Kusko, Alexander. Emergency/Standby Power Systems. New York: McGraw-Hill, 1989.

6. Military Handbook: Reliability Prediction of Electronic Equipment. MIL-HDBK-217F, Wash., DC: U.S. Dept. of Defense, January 1990.

7. Ramakumar, Ramachandra. Engineering Reliability: Fundamentals and Applications. Upper Saddle River: Prentice Hall, 1993.

8. Sanders, Mark S., and McCormick, Ernest J. Human Factors in Engineering and Design. 6th Ed. New York: McGraw-Hill, 1987.

9. Snevely, Rob. Enterprise Data Center Design and Methodology. Palo Alto: Sun Microsystems Press, A Prentice Hall Title, 2002.

10. Swain, A.D., and Guttman, H.E. Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications (THERP) Final Report. NUREG/CR-1278-F, Wash., DC: U.S. Nuclear Regulatory Commission, August 1983.

About the Authors:

Steve Fairfax is President of MTechnology, Inc. Steve joined MTech in 1997, but he has been working with multi-megawatt power systems since his undergraduate days at MIT, where he helped build and operate a 200 MW power system for a tokamak fusion reactor. He began full-time study of power system reliability while working as Managing Engineer for Failure Analysis Associates. He served as head of engineering and operations for the Alcator C-MOD nuclear fusion reactor during its design and initial operation at the MIT Plasma Fusion Center, and as principal engineer in Boston-area firms. Mr. Fairfax holds Master's Degrees in both Physics and Electrical Engineering from MIT.

Neal Dowling is a senior engineer at MTechnology, Inc. He performs fault tree analysis and related modeling and simulation, develops and tests new power supply and switch technology, and supervises the operation and maintenance of MTech's 400 kW fuel cell power plant facility. Neal worked at several Boston-area medical device manufacturers prior to joining MTech. His expertise includes development and maintenance of firmware and software for critical functions, FDA compliance, and analog and digital design. Neal holds Bachelors and Masters degrees in Electrical Engineering from MIT.

Dan Healey is a senior engineer at MTechnology, Inc. He specializes in human factors analysis and the applications of PRA techniques to operations and maintenance activities. Dan served as Director of Engineering at several Boston-area firms, overseeing product development for semiconductor processing, medical equipment, robotics, and electro-optical systems. Dan holds a Bachelors degree in Electrical Engineering from the University of Rochester with additional graduate work in optics and programming. He is presently a special student at Harvard studying management of technology and software development.

MTechnology, Inc. provides power systems engineering for the 21st century. The firm offers consulting, testing, product development, and prototype fabrication services.
MTech performs probabilistic risk analysis of electric power systems, design reviews, root cause failure analysis, and provides expert testimony in both regulatory and litigation settings. MTech offers consultation on risk-informed system design, operations, maintenance, upgrades, and reliability growth management. Clients frequently realize substantial savings on capital and operating expenses while simultaneously increasing reliability. MTech's facilities include a 5,000 square foot test and laboratory facility with ability to operate 500 kW continuous loads and multi-megawatt pulsed loads. MTech has worked on high-reliability distributed generation projects spanning technologies from reciprocating engines to fuel cells. The firm's clients include electric utilities, designers and engineers, critical facility owners and operators, and manufacturers serving the 7x24 mission critical industry.