UNDERWRITER CASE STUDY

Uptime Institute Symposium 2012

Financial Benefits of Using Power System Analytics for Mission-Critical Applications Return on Investment (ROI)

By: John Jennings, Director Product Management – Power Analytics
By: John Kintzele, Director Global Operations – Power Analytics
By: Christopher Sticht, Senior Systems Engineer – Power Analytics

The state-of-the-art practices for operation of critical power facilities have some limitations that may result in the loss of thousands of dollars per year. Model-validated simulation software allows determination of true dynamic capacity during both normal and contingency condition. Recovering just 5% of stranded capacity can mean a delay of months or even years until a new facility must be built to meet rising demand translating into millions of dollars of savings! For typical project pricing, the payback period is less than 1.1 years.

~~John Jennings, John Kintzele and Chris Sticht are with Power Analytics Corporation,~~
~~9208 Falls of Neuse Road, Suite 215, Raleigh, NC 27615 USA~~
~~(e-mail: jjennings@poweranalytics.com; jkintzele@poweranalytics.com; and csticht@poweranalytics.com)~~

Financial Benefits of Using Power System Analytics for Mission-Critical Applications

Return on Investment (ROI)

John Jennings, John Kintzele, and Christopher Sticht

***Abstract* – This paper identifies multiple financial benefits of using model-validated analytics software to assess mission-critical power systems. Model-validated analytics allows facility owners and operators to: minimize underutilization of design system capacity (stranded capacity) and realize Future Value (FV) returns on capacity expansion dollars; predict and optimize energy usage and associated costs across multiple system states and source/load profiles; determine the viability, criticality, and appropriate costs of scheduled maintenance, and; increase the situational awareness and coordination between operators and maintainers to reduce downtime costs associated with human error. Taken together, these benefits result in annual operational and capacity expansion savings ranging from thousands to millions of dollars. For typical project pricing, the payback period is less than 1.1 years.**

Keywords: Return on Investment; ROI; Stranded Capacity; Model-Validated Analytics; Maintenance; Savings; Energy Management; Annual Savings; DCIM; Downtime;

I. Introduction

Software providing model-validated analytics of electrical power systems allows a comparison of real-time conditions (generation sources, breaker and switch states, and load distributions) to a computational representation (model) of the as-designed/as-built system in an identical state. This paper explores the financial benefits of using such analytics software to profile and manage power systems in mission-critical facilities. Several considerations are addressed: recovery of unused capacity to extend facility operating life; predictive management of energy usage and costs; maintenance prioritization and costing, and; maintenance planning and coordination to improve operational reliability. These considerations map to the following challenges:

- Physical infrastructure is *oversized* routinely due to uncertainties surrounding true power system capacity.
- Energy usage and cost management is *reactive* because there are no real-time means to simulate and predict the impact of day-to-day system configuration changes.
- Maintenance schedulers often *overestimate* the criticality of maintenance actions based on avoidable uncertainties, driving up costs.
- Operators and maintainers do not possess effective means to coordinate and practice procedures in a unified fashion, opening the door to *human error* and downtime costs.

Model-validated simulations and analyses allow power system owners to greatly reduce costs associated with these challenges.

II. Financial Evaluation – Definitions

A. *Stranded Capacity – "The greenest data center is the one you don't need to build"*

Most data center owners and operators would agree with Matt Warner's statement "The greenest data center is the one you don't need to build." [1] An extension to that thinking would be "or that you don't need to build *yet* due to recovery of stranded capacity".

Over-sizing of data centers is a significant waste of capital and operating costs. According to Gartner, the majority of data centers will never reach their intended capacity. [2] In fact, it is estimated that typical datacenters today could support up to 30% more IT equipment using the same facility and cooling equipment if the infrastructure was properly managed. [3]

While some stranded capacity is inherent to all power system designs, other capacity goes deliberately unused due to uncertainty regarding a system's dynamic capabilities – it is intentionally stranded. The execution of simulations based on model-validated analytics can reduce uncertainty and help identify strategies to recover this second type of stranded capacity. Consequently, the usable life of the facility can be extended and capacity expansion projects may be deferred.

In addition to other benefits, the recovery of intentionally stranded capacity has a direct financial return in the form of interest on construction dollars for the interval between the original and revised construction dates. The following figure illustrates this relationship.

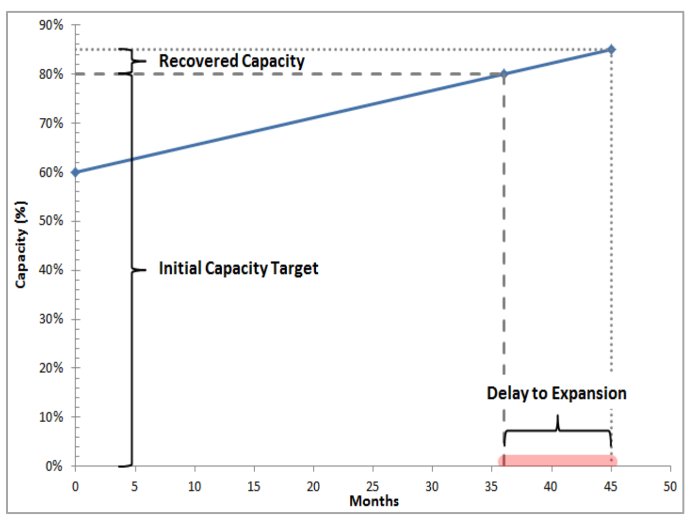

Figure 1. *Capacity Expansion Delay Due to Recovery of Capacity*

John Jennings, John Kintzele and Chris Sticht are with Power Analytics Corporation,
9208 Falls of Neuse Road, Suite 215, Raleigh, NC 27615 USA
(e-mail: jjennings@poweranalytics.com; jkintzele@poweranalytics.com; and csticht@poweranalytics.com)

Capacity management tools, such as model-validated analytics, can be used to properly manage the electrical infrastructure, to better utilize power resources, and to manage electrical consumption. For example, given that most facilities are designed for initial capacity usage of 60%, and planning for construction of new capacity begins at 80% [4], the recovery of even 5% of intentionally stranded capacity can return millions to the bottom line.

B. *Informed Energy Management – Power (Data) to the People*

A successful energy management program is built upon having timely and accurate information regarding the consumption of energy, where and when energy is consumed, and the operational and environmental costs of the electrical infrastructure of the data center. Ideally, the reporting of this information includes a comparison of expected usage to actual usage, and allows simulation of energy strategies to maximize efficiency and minimize costs.

Even in this era of increasingly sophisticated infrastructure management, the energy strategies and policies of mission-critical facilities are often determined based on anecdotal prior experience. The "strategies" are then vetted by committee and tested operationally with no pre-validation or guarantee of success. Alternately, energy strategies can be determined by expensive engineering studies, and the more scenarios (system configurations, external conditions) evaluated, the more the price climbs.

Until recently, there has been no path for those closest to the power system (owners and operators) to assess energy strategies independently and at low cost before testing them in practice. Model-validated analytics with simulation capabilities offer such a path, provided they allow users to compare simulated and actual energy metrics as shown in Figure 2.

Figure 2. *Example of Interface for Energy Simulations [5]*

① Top panel displays actual energy usage and related metrics (power factor, PUE, carbon footprint) from the data historian, with utility rates and penalties schedule applied.

② Bottom panel displays identical information, but using model-validated analytics based on the current simulation.

Software capable of providing a comparison between actual usage and simulated usage allows operators and owners to assess potential savings due to alternate and supplemental generation sources, as well as the interaction of loading profiles with equipment efficiency curves and overall system stability. It can help those responsible for facility operation assess these factors and pre-validate energy strategies under a variety of operating conditions, and with time-of-day resolution.

C. *Quantifying Maintenance Criticality – Pay Only for the Labor You Need*

After energy expenditures, maintenance costs are usually the second largest operational expense for facilities with critical power requirements. Indeed, for some facilities, maintenance costs are the single largest operational expense. Facility managers can have significant impacts on their operating costs by understanding the opportunities that exist to reduce or control maintenance costs. [6]

One important opportunity to reduce costs lies in the decision-making process used to estimate the impact of maintenance procedures on critical loads. Operators of critical power systems often have guidelines defining which maintenance procedures are permitted during periods of peak or near-peak load, and which must occur after-hours. The cost impact of these decisions is obvious; when owners and operators decide that a maintenance action must occur outside normal business hours, costs can increase by as much as a factor of 3x.

When operators cannot quantify the potential impacts of a planned maintenance action, they err on the side of caution, incurring the higher direct labor costs. Lack of hard data also drives extended decision-making processes, with involvement to very high levels and multiple entities in the owner, operator, and even customer organizations. Model-validated analytics can be used to better understand the criticality of each maintenance action, reducing uncertainty and allowing the use of lower-cost labor. As an example of reduced uncertainty, the use of simulations can allow a work order to include verifiable statements of the following types:

- The system configurations planned during maintenance are electrically viable and will not affect critical loads.
- Maintenance redistribution of loads does not result in overloading of any feeder or component.
- Maintenance adjustments to breaker and switch topology continue to support critical electrical/mechanical loads.
- Technicians are aware of all equipment in the work area that will be energized during maintenance, including equipment, feeders, and buses not normally energized.
- All of the above statements remain true for loss of utility and other defined contingency conditions.

The confidence this builds can allow streamlining of the decision-making process, reducing indirect costs in addition to the direct labor savings. Using maintenance labor rates of $4.20 per square foot [7] for an average Tier III / IV data center, and assuming a redistribution of just 10% of maintenance actions from critical to routine, the direct annual savings alone can be in the tens of thousands of dollars.

D. *Improve Maintenance Training and Coordination*

Ponemon Institute studies released in 2010 and 2011 [8] include a number of key findings related to the causes and costs of downtime. These findings make it possible to estimate the average contribution of human error during maintenance to downtime costs. The relevant conclusions for the average data center are:

- Average cost of data center downtime: $5,600 per minute.
- Average reported incident length: 90 minutes
- Average cost of a single downtime event: $505,500.
- Contribution of human error to downtime: 25-50%

Therefore, the average annual downtime costs due to human error can range from $120,000 - $245,000. [9]

Definitive research by the Federal Aviation Administration and others has demonstrated that effective training can reduce incidents due to human error by 51-81%. [10] Effective training requires that:

- Workers understand the big picture of which their activities are a part – they must have "situational awareness". [11]
- Workers have practiced their specific tasks recently and to the level of detail required. [12]
- Workers anticipate and know how to react to contingency conditions. [13]

Model-validated analytics and simulation software helps facility operators address all of these requirements. The use of simulations, particularly those that highlight variances from as-designed behavior, allow operators and maintainers to view both the system-wide and local effects of changes to sources, topology, and loading.

These same simulations allow operators and maintainers to run joint systematic walkthroughs of the intended procedure; as a result, operators can better understand the probable alerts and indications they will receive, and maintainers can understand their local conditions and anticipate refinements to the procedure in advance. In addition, and distinct from current best practices, these joint sessions can also assess the effects of common contingency conditions such as loss of utility, and decide in advance on the appropriate reactions.

The net effect of simulation-driven mutual planning sessions is improved communication and coordination between the parties throughout the procedure, and greatly enhanced ability to react to unexpected conditions.

III. FINANCIAL EVALUTION – EXAMPLE

In order to quantify the ROI due to implementation of model-validated simulation and analytics software, some assumptions must be made. For the purposes of this financial evaluation, a data center with the following characteristics is employed:

- Square footage: 100,000
- MW of useable UPS output: 5 MW
- Tier Level: Tier III
- Data center cost for build-out: $122 M USD
- Loading at initial operating capability: 60%
- Target loading before new capacity online: 80%
- Achievable loading after capacity recovered: 85%
- Interest rate (APR) after inflation: 1.6%
- Estimated time to reach targeted loading: 36 months

It is further assumed that a new, identical data center will be built at the 36-month point to handle rising demand.

A. *Earnings Due to Recovery of Stranded Capacity*

The use of advanced electrical simulation and analytics software can help to extend usable data center capacity. For the example case, a 5% gain in usable capacity (a change in targeted utilization from 80% to 85%) is modeled. Given the other assumptions above, such a change would result in a construction deferment of 9 months, at the end of which the future value of the original construction budget would be $123,471,832. The ROI in this case would therefore be $1,471,832.

B. *Energy Management Cost Savings*

Electrical simulation and analytics software has tremendous potential for identifying savings in energy costs. Simulations can be used to analyze the electrical system and find ways to reduce demand charges and power factor penalties. For example, an operator can examine the effect of redistributing loads to achieve peak UPS efficiencies. The result of such a simulation for an example case appears below.

Actual Usage	Usage	Cost Factor	Cost	Explanation
On-peak demand (kW)	39,000	0.498	$ 582,660	On-peak demand charge
Off-peak demand (kW)	39,000	0.274	$ 320,463	Off-peak demand charge
On-peak supply usage (kWh)	3,615,117	0.0460	$ 166,295	On-peak supply charge
Off-peak supply usage (kWh)	14,647,848	0.0072	$ 105,465	Off-peak supply charge
Total usage (kWh)	18,262,965	0.0072	$ 131,493	Capacity charge
On-peak power factor (%)	87%	13%	$ 5,070	On-peak PF penalty
Off-peak power factor (%)	81%	19%	$ 7,410	Off-peak PF penalty
Other			$ 1,050	Access charge
		Total 30-Day Cost	***$ 1,319,906***	

Simulation Usage	Usage	Cost Factor	Cost	Explanation
On-peak demand (kW)	38,454	0.498	$ 574,496	On-peak demand charge
Off-peak demand (kW)	38,454	0.274	$ 315,973	Off-peak demand charge
On-peak supply usage (kWh)	3,564,462	0.0460	$ 163,965	On-peak supply charge
Off-peak supply usage (kWh)	14,442,605	0.0072	$ 103,987	Off-peak supply charge
Total usage (kWh)	18,007,068	0.0072	$ 129,651	Capacity charge
On-peak power factor (%)	87%	13%	$ 4,999	On-peak PF penalty
Off-peak power factor (%)	81%	19%	$ 7,306	Off-peak PF penalty
Other			$ 1,050	Access charge
		Total 30-Day Cost	***$ 1,301,427***	

Monthly Savings Due to Energy Simulations	***$ 18,480***
Annual Savings Due to Energy Simulations	***$ 221,755***

Figure 3. *Energy Simulation Cost Savings*

C. Reduction in Overtime Maintenance Service Costs

The reduction in costs due to re-categorization of maintenance criticality and the commensurate reduction in overtime labor costs require further assumptions. For the example data center of 100,000 square feet, and typical Tier III / Tier IV costs of $4.20 per square foot [7], the direct annual maintenance labor costs are estimated as $420,000. In addition, consistent with this same survey of customers, the example case assumes that the breakdown of each level of criticality, based on typical, cautious decision-making practices is:

- 75% of operational maintenance at normal labor rates
- 15% of operational maintenance at 1.5 normal labor rates
- 10% of operational maintenance at 2.0 normal labor rates

These expenditures correspond with the first table of costs, and total $420,000 per the base assumptions. The second table of costs makes the further assumption that increased knowledge of maintenance effects (as detailed in Section II.C.) will allow a 10% redistribution of higher-cost off-hours maintenance to routine maintenance. A final assumption is made concerning indirect costs such as streamlined decision-making, primarily in the form of lower level, expedited approvals for maintenance actions meeting the additional criteria exposed and verified through power system simulations.

Facility square footage
100,000
Mean maintenance labor costs per square foot
$4.20
Typical annual maintenance labor costs
$420,000

Current maintenance costs with conservative allowances for uncertainties			
Type of maintenance	**Hourly rate factor**	**% of annual budget**	**Expenditures**
Routine (normal labor rates)	1.0	75.0%	$ 268,085
After-hours (over-time rates)	1.5	15.0%	$ 80,426
Weekend (double-time rates)	2.0	10.0%	$ 71,489
			$ 420,000

Reduced maintenance costs due to informed analysis			
Type of maintenance	**Hourly rate factor**	**% of annual budget**	**Expenditures**
Routine	1.0	85.0%	$ 303,830
After-hours	1.5	12.5%	$ 67,021
Weekend	2.0	2.5%	$ 17,872
			$ 388,723

Direct annual savings due to reduced labor expenditures	*$ 31,277*
Indirect annual savings due to streamlined decision processes	*$ 20,000*
Total annual savings due to informed maintenance criticality assessement	***$ 51,277***

Figure 4. *Reduction in Costs Due to Reassignment of Critical Maintenance*

The savings for the data center in the example amount to about $51,000 dollars on a yearly basis. Amounts for other data centers will scale directly with the overall labor budget, the percentages of procedures conducted at each level of criticality, and the findings of the simulations performed using model-validated analytics.

D. Reduction in Downtime Costs Due to Human Error

As discussed previously, the annual downtime costs due to human error can range from $120,000 - $245,000 for a typical data center. However, effective training can reduce human error between 51-81%. Taken together, these metrics allow an estimate of yearly savings due to reduction of human error of $61,000 to $198,000, or a mean savings of $130,000.

E. Example ROI Summary

The earnings and savings from model-validated analytics software fall into two categories: per-project savings, and recurring annual savings. A summary of the financial benefits for the example case is shown in Table I.

TABLE I. ROI SUMMARY

Item	**ROI Financial Evaluation [a]**		
	Category	***Per Project Savings $ (000s)***	***Annual Savings $ (000s)***
1	Recovery of stranded capacity and delay in capacity expansion	$1,475	N/A
2	Energy use prediction and optimization	N/A	$222
3	Reduction in maintenance costs due to viability/criticality assessments	N/A	$51
4	Median reduction in downtime due to human error	N/A	$130
5	Total of annual savings	N/A	$403
6	Total savings over 4 years	$3,087	

a. Applies to this specific example – an ROI calculator is available on the Power Analytics website

***Summary of financial benefits**: The median cost for the purchase and installation of model-validated analytics software is $450,000. Based on the annual savings alone, the average payback period for this example is 1.1 years. When amortized earnings resulting from the recovery of stranded capacity are included, the payback can drop to just months.*

IV. Implementing Model-Validated Analytics

A. Overview

Effective implementation and usage of model-validated analytics software is important to ensure accurate information is used to make decisions that affect the data center facility.

B. Steps Required for Implementation

In order to implement model-validated analytics software for an electrical power network, three steps must occur:

1. Create a model of the electrical network with all electrical characteristics of each component as shown in Figure 5.

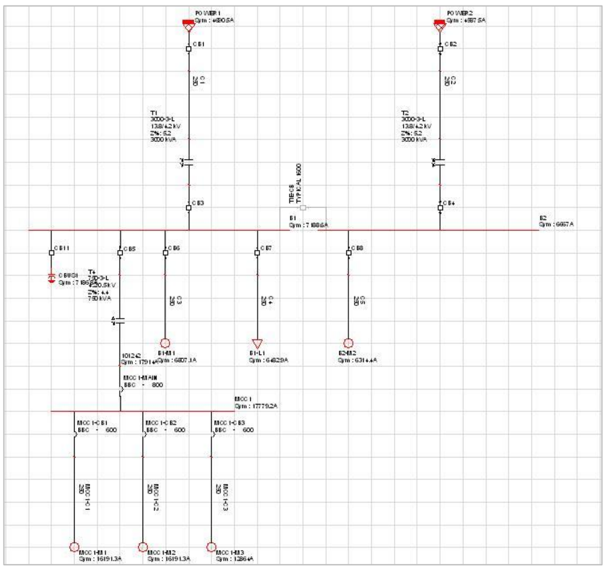

Figure 5. *Step 1 – Create a model of the electrical network*

2. Collect the data in real-time for the entire system as shown in Figure 6.

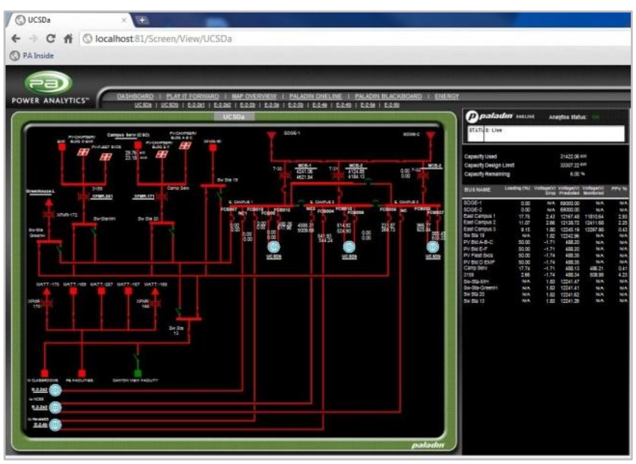

Figure 6. *Collect data and display in a user-friendly format*

3. Perform analytics and simulations based on a comparison of the predictive model, and actual real-time data from the system as shown in Figure 7. The resulting power and performance metrics can be used to optimize the power usage and loading and identify issues prior to problems occurring.

Figure 7. *Step 3 – Perform model-validated analytics*

V. Conclusion

The techniques discussed in this paper allow data center owners and operators to gain greater insight into the current and future state of their facilities and take preemptive steps to optimize capacity expansion plans, energy usage, and maintenance activities. Significant earnings and savings can be realized in the process. For example, model-validated analytics can:

- Help identify and reduce stranded capacity, resulting in extended facility life and future value earnings on construction dollars.
- Help operators participate in the creation of energy management strategies, resulting in enhanced identification of possible savings and more nimble responses to actual conditions.
- Help reduce maintenance labor costs by improving operator insight into actual maintenance effects and criticality.
- Help minimize downtime costs due to human error through "what-if" scenarios and training.

Model-validated electrical simulation and analytics software is a powerful tool for studying operating scenarios that have significant financial benefit to the data center owner/operator. While the benefits of such software are only beginning to be realized, the importance of model-validated analytics and simulations will continue to grow as data center electrical power systems become more complex and customer expectations on pricing and service levels continue to rise.

Acknowledgment

The authors gratefully acknowledge the contributions of all customers that provided input to identify the significant challenges and opportunities in cost savings for mission-critical applications. Also, we offer special thanks to Power Analytics CTO Kevin Meagher; and VP of Sales Doug Whitmer, for coordinating the effort to present the data in a clear and concise format.

REFERENCES

[1] Infrarati, May, 2011, Green Press, http://infrarati.wordpress.com/2011/05/15/stranded-capacity-a-hurdle-to-get-your-data-center-green/.

[2] Cappuccio, D.J. "Data Center Efficiency and Capacity: A Metric to Calculate Both", Gartner, September 2009.

[3] Power and Cooling Capacity Management for Data Centers, Neil Rasmussen. http://www.slideshare.net/datacenters/power-and-cooling-capacity-management-for-data-centers.

[4] Power Analytics Customer Surveys and Questionnaires, Q1, 2012.

[5] Power Analytics Paladin Live Energy Management Interface graphics. Raleigh, NC, 2011.

[6] Facilities Issues, "Benchmarking Your Maintenance Costs", FMLINK Group, LLC, 2009.

[7] A Simple Model for Determining True Total Cost of Ownership for Data Centers, Uptime Institute, Version 2.1, March 31, 2008. By Jonathan Koomey, Ph.D., with Kenneth Brill, Pitt Turner, John Stanley, and Bruce Taylor.

[8] Emerson Network Power, "Understanding the Cost of Data Center Downtime: An Analysis of the Financial Impact on Infrastructure Vulnerability". Liebert Corporation LLC, 2011. SL-24661 R05-11.

[9] Ponemon Institute, Calculating the Cost of Data Center Outages, February, 2010 and February, 2011.

[10] FAA Human Factors, Team Performance / CRM Effectiveness Evaluation, www.hf.faa.gov.

[11] D. Gaba, Simulation Center for Crisis Management Training in Health Care, VA Palo Alto and Stanford University.

[12] Training Requirements in OSHA Construction Industry Standards and Training Guidelines, www.osha.gov.

[13] Robert B. Kelly, Industrial Emergency Preparedness, Wiley, 1989.

BIOGRAPHIES

John Jennings has been an IEEE member since 1989 and is an active member of the IEEE 1584, IEEE 1683, and IEEE 2030 working groups. John currently works at Power Analytics as the Director of Product Management. He holds a BSEET from Southern Polytechnic State University in Marietta, GA. John served 6 years in the US Navy as a naval nuclear operator. He has a background in the electrical distribution and control industry with ABB, Eaton Electrical, and Schneider Electric. John also has experience in the pulp and paper industry, pumping applications, and Nuclear Power Industry with Bechtel. John is lead auditor NQA-1 qualified and has published several papers on arc flash hazards and product development.

John Kintzele is Power Analytics' Director of Services, responsible for the project teams implementing Paladin Live model-validated analytics at customer sites worldwide. John has worked as an engineer, project/program manager, consultant, and executive in the public and private sectors. While with NASA/JPL, he authored the formal reliability/availability requirements for the International Space Station, named the rescue craft (the "Assured Crew Return Vehicle"), and evaluated the effect of mechanical and thermal perturbations on experimental microgravity requirements. Later, with Schneider Electric, John was responsible for the product safety evaluations of cross-channel and brand-label power and automation hardware and software products being introduced into the North American market. He has a BSE in Mechanical and Aerospace Engineering from Princeton University.

Christopher Sticht is a specialist in utility system planning, planning software and Smart Grid. Chris has extensive background in utility system planning, system design, operations, and protection. His experience includes work on transmission systems, distribution systems, substations, and data center electrical systems. He has managed large engineering and design teams, and served as a consultant and subject matter expert at two power-flow software companies. He holds a MSEE from the University of Washington and a BSEE from Georgia Institute of Technology.

NOTE

John Jennings, Christopher Sticht, and John Kintzele were employed at Power Analytics Corporation at the time this paper was published. Power Analytics is a privately held developer of software solutions for the design, simulation, deployment, and preventative maintenance of complex electrical power systems. Founded in 1983, its software products are used by thousands of commercial, industrial, governmental, and military customers worldwide to protect more than $100 billion in customer assets. Primary offices are located in San Diego, CA and Raleigh, NC with over 30 sales, distribution, and support offices located in North America, South America, Europe, Asia, and Africa.

Power Analytics' Paladin® software suite helps organizations ensure that their electrical power infrastructure is optimally designed (Paladin DesignBase™), performs precisely as intended in terms of reliability and energy efficiency (Paladin Live™), and operates flawlessly as organizations make real-time transitions between public and on-premise power sources (Paladin SmartGrid Power Management System™). By continually calibrating operating conditions with the original design CAD model, the Paladin software suite offers the only real-time analytics solution for diagnosing electrical power problems or energy inefficiencies at their earliest stages.