EPRI | ELECTRIC POWER RESEARCH INSTITUTE

Emerging Technologies for Efficient Data Centers

Uninterruptible Power Supply Eco Mode, Liquid Cooling, and Evaporative Cooling

3002001360

Emerging Technologies for Efficient Data Centers

Uninterruptible Power Supply Eco Mode, Liquid Cooling, and Evaporative Cooling

3002001360

Technical Update, December 2013

EPRI Project Manager

B. Fortenbery

ELECTRIC POWER RESEARCH INSTITUTE
3420 Hillview Avenue, Palo Alto, California 94304-1338 ▪ PO Box 10412, Palo Alto, California 94303-0813 ▪ USA
800.313.3774 ▪ 650.855.2121 ▪ askepri@epri.com ▪ www.epri.com

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE FOLLOWING ORGANIZATION PREPARED THIS REPORT:

Electric Power Research Institute (EPRI)

This is an EPRI Technical Update report. A Technical Update report is intended as an informal report of continuing research, a meeting, or a topical study. It is not a final EPRI technical report.

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

Electric Power Research Institute, EPRI, and TOGETHER…SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

Copyright © 2013 Electric Power Research Institute, Inc. All rights reserved.

ACKNOWLEDGMENTS

The following organization prepared this report:

Electric Power Research Institute (EPRI)
3420 Hillview Avenue
Palo Alto, CA 94304

Principal Investigator
M. Sweeney

This report describes research sponsored by EPRI.

This publication is a corporate document that should be cited in the literature in the following manner:

Emerging Technologies for Efficient Data Centers: Uninterruptible Power Supply Eco Mode, Liquid Cooling, and Evaporative Cooling. EPRI, Palo Alto, CA: 2013. 3002001360.

ABSTRACT

This report summarizes research in emerging technologies that improve data center energy efficiency, including evaporative cooling, liquid cooling, and high-efficiency eco mode operation of the uninterruptible power supply. The report describes the efficiency gains of these technologies and their impact on total data center energy use. It also identifies market barriers for each technology and potential next steps to promote adoption of these efficient technologies.

Keywords
Data centers
eco mode
Energy efficiency
Evaporative cooling
Liquid cooling

EXECUTIVE SUMMARY

Data centers are highly energy-intensive facilities that consume a substantial amount of electricity. Industry experts estimate that data centers accounted for 1.7% to 2.2% of the electric power consumption in the United States in 2010 (about 67 to 86 TWh). With the increasing adoption of digital services into everyday business and personal use, the demand for data centers is set to increase for the foreseeable future. This work studies three categories of emerging technologies that can reduce the amount of energy consumed in data centers: evaporative cooling, liquid cooling, and "eco mode" operation of the uninterruptible power supply (UPS).

Liquid cooling and evaporative cooling are two technologies that have the ability to sufficiently cool servers with remarkably lower energy use than conventional cooling systems using vapor compression or chilled water for heat removal. These technologies promise as high as 95% reduction in data center cooling needs in appropriate climates. Cooling signifies the greatest energy consumption in data centers after the servers themselves, and such a dramatic reduction in cooling translates to significant energy savings for the entire data center (up to 40%). With these findings, follow-up efforts could be used to demonstrate the savings that these technologies offer in a variety of climates. A coordinated deployment of these technologies in the field can demonstrate the viability of these systems across the United States and further describe their economics to inform the suitability of energy efficiency incentives in various regions.

After cooling, the UPS is the second-largest overhead energy use in the data center. Many data center operators face business pressure to provide their digital services with extremely high availability, as much as 99.99%. Because of this, data centers often use UPS devices with the best protection from disturbances in input power quality. The best protection is offered by double-conversion UPS architectures, yet this design is among the least efficient. A number of high-efficiency eco mode UPSs have entered the market with the promise to deliver protection from interruption of servers with up to 99% efficiency. This work examines these claims by detailing the performance and efficiency of these designs, along with an analysis of the power quality (PQ) requirements of server power supply units (PSUs) and the availability of high-quality power from the grid.

The findings of this work demonstrate that eco mode operation of UPS devices is able to satisfy the needs of servers without interruption, with up to 18% reduction in total data center energy use. The disparity in performance between different eco mode designs and the reluctance of data center operators to deploy any device that may negatively impact data center availability has highlighted the need for standardized performance metrics that can inform consumers of the trade-off between energy savings and reliability that is offered by the market. A beneficial follow-up to this work is the development of standardized test protocols and appropriate metrics to inform the market of the performance offered by eco mode designs and the sensitivity of server PSUs to such events. This would effectively address the major market barriers and increase the adoption of this high-efficiency technology.

Table ES-1 shows the savings available from the technologies examined in this study as a percent of total data center energy use. These technologies are all commercially available and ready for deployment in new and existing data center installations. Given the ubiquitous nature of data center UPS use and quick payback of the technology, further standards development and technology demonstration offer cost-effective means for a utility to promote the adoption of this

efficient technology. Further, evaporative cooling and liquid cooling each offer tremendous data center energy savings that could be reasonably justified in numerous sites with appropriate utility incentives.

Table ES-1
Total energy savings available from technologies examined

Technology	Energy Savings
Evaporative cooling	< 40%
Liquid cooling	30% to 50%
UPS eco mode	3% to 18%

CONTENTS

***1* INTRODUCTION 1-1**
Overview of Data Center Energy Use 1-1
Cooling Systems 1-2
Conversion Factors 1-3
***2* TECHNICAL ASSESSMENT OF EVAPORATIVE COOLING TECHNOLOGY 2-1**
Thermodynamic Process 2-1
Direct Evaporative Cooling 2-3
Indirect Evaporative Cooling 2-4
Vendor Case Studies and Savings Claims 2-5
Other Evaporative Cooling System Architectures 2-5
Two-Stage Evaporative Cooling 2-5
Hybrid Direct Expansion Systems 2-6
Evaporative Condenser Precooling 2-6
Recommendations for Application of Findings 2-6
Partial List of Vendors 2-7
***3* TECHNICAL ASSESSMENT OF LIQUID COOLING TECHNOLOGIES 3-1**
Submersion Cooling 3-1
Case Studies and Energy Savings 3-2
Vendor Claims 3-3
Phase-Changing Liquid Coolant 3-3
Liquid-Encapsulated Servers 3-4
Device-Mounted Liquid Cooling 3-4
Vendor Claims 3-6
Market Barriers 3-6
Recommendations for Application of Findings 3-6
Partial List of Vendors 3-6
***4* TECHNICAL ASSESSMENT OF UNINTERRUPTIBLE POWER SUPPLY ECO MODE 4-1**
Data Center Critical Power Requirements 4-1
Uninterruptible Power Supply Background 4-2
Eco Mode 4-4
Efficiency Gains 4-4
Eco Mode Performance 4-6
Voltage Sag Immunity of Information Technology Equipment 4-7
Reliability Considerations 4-11
Site Power Quality 4-12
Recommendations for Application of Findings 4-14
***5* CONCLUSIONS 5-1**
***6* REFERENCES 6-1**

_A_ UNINTERRUPTIBLE POWER SUPPLY TESTING PROCEDURE A-1
Purpose A-1
Scope A-1
Definitions A-1
Requirements A-2
Uninterruptible Power Supply and Load Specifications A-2
Test Equipment Specifications A-6
Data Product A-7
Uninterruptible Power Supply Verification and Test Site Survey A-8
Uninterruptible Power Supply Status Verification A-8
Test Site Survey A-8
Test Setup A-9
Functionality Tests A-11
Test Procedure A-11
Test Procedure for Measuring Efficiency in Double-Conversion Mode A-12
Test Procedure for Measuring Efficiency in Eco Mode A-12
Test Procedure for Capturing Switching Transients and Calculating Transfer Time A-12

LIST OF FIGURES

Figure 1-1 Typical data center energy use 1-1
Figure 1-2 Psychrometric chart of ASHRAE guidelines for data centers..... 1-3
Figure 2-1 Psychrometric illustration of evaporative cooling effect 2-2
Figure 2-2 Total data center energy reduction from evaporative cooling 2-3
Figure 2-3 Direct evaporative cooling..... 2-3
Figure 2-4 Indirect evaporative cooling 2-4
Figure 2-5 Two-stage (indirect–direct) evaporative cooling system 2-5
Figure 3-1 Four server racks submersed in liquid coolant 3-2
Figure 3-2 Total data center energy use reduction from submersion cooling..... 3-3
Figure 3-3 Central processing unit–mounted cold plates with in-rack coolant distribution..... 3-5
Figure 4-1 Block diagram of double-conversion (online) uninterruptible power supply architecture 4-2
Figure 4-2 Double-conversion uninterruptible power supply efficiency 4-3
Figure 4-3 Block diagram of standby (offline) uninterruptible power supply 4-3
Figure 4-4 Efficiency of eco mode and double conversion for uninterruptible power supplies.. 4-5
Figure 4-5 Eco mode energy savings..... 4-6
Figure 4-6 Information technology voltage-sag immunity curve (revised 2000) 4-8
Figure 4-7 Measured voltage-sag immunity of server power supply units..... 4-9
Figure 4-8 IEC 62040-3 specification for uninterruptible power supply output performance for general-purpose information technology loads..... 4-10
Figure 4-9 Voltage-sag protection provided by eco mode 4-11
Figure 4-10 Annual rate of voltage sags and interruptions measured at customer-side sites. 4-13
Figure A-1 Block diagram of a double-conversion online uninterruptible power supply..... A-2
Figure A-2 Sample test setup to measure uninterruptible power supply efficiency A-7

LIST OF TABLES

Table 1-1 ASHRAE thermal guidelines for data center environments 1-2
Table 1-2 Conversion factors 1-3
Table 4-1 Data center tier classification 4-1
Table 4-2 Eco mode energy savings 4-6
Table 4-3 Eco mode transfer time specifications 4-7
Table 5-1 Total savings available from technologies examined 5-1
Table A-1 Uninterruptible power supply specifications to be recorded A-4
Table A-2 Reference test load specifications A-6
Table A-3 Test equipment requirements A-6
Table A-4 Uninterruptible power supply status worksheet A-7
Table A-5 Environmental site survey A-8
Table A-6 Uninterruptible power supply parameters to be measured during testing A-10

1
INTRODUCTION

Data centers house the servers that power the digital services that modern society depends on. These facilities can range from server closets to large warehouses whose sole purpose is to house the servers within. Data centers are highly energy intensive, with energy density closer to that of industrial plants than commercial or residential facilities. Energy is required not only for powering the information technology (IT) and communications equipment but also for cooling systems that remove the heat generated by the IT equipment and the power conditioning equipment that provides an uninterrupted supply of power. According to Jonathan Koomey, a Research Fellow at Stanford University, data centers accounted for 1.7% to 2.2% of the U.S. electric power consumption in 2010 (67 to 86 TWh) [1]. Moreover, this industry is continuing to grow as e-mail, social networking, streaming media, and cloud-based services gain popularity and adoption grows.

Overview of Data Center Energy Use

The energy efficiency of a data center is typically reported by its power usage effectiveness (PUE). *PUE* is defined as the ratio of total data center power to the power consumed by the IT equipment, including servers, networking, and data storage. A lower PUE indicates a more efficient data center, approaching the ideal limit of 1.0. Figure 1-1 shows the average breakdown of energy consumption from the results of a Lawrence Berkeley National Laboratory benchmark study, published in the U.S. Environmental Protection Agency (EPA) 2007 Report to Congress on data center energy efficiency [2]. From these results, the IT load represents about half of the electricity used in the data center, with cooling between 30% and 40% and uninterruptible power supply (UPS) losses about 10%.

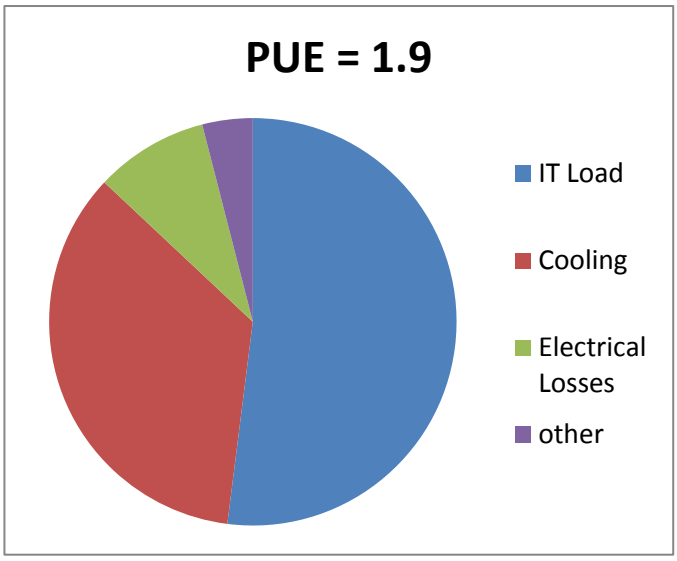


Figure 1-1
Typical data center energy use

As shown in Figure 1-1, the average PUE found in the Lawrence Berkeley National Laboratory survey was 1.9, indicating that, for every watt of IT equipment, 0.9 W was demanded by overhead and support systems [2]. Similarly, the EPA reported an average PUE of 1.9 from a survey of more than 100 data centers for its Energy Star program in 2009 [3].

Recently, some of the best-in-class facilities have reported PUE of less than 1.2. These facilities generally rely on hyper-efficient cooling methods, including economizers that take advantage of "free cooling" by using outdoor air for cooling.

Cooling Systems

Cooling systems typically account for about 30% to 40% of the energy use in data centers, but this can vary greatly, depending on cooling efficiency and climate. Conditioned air is usually used to remove waste heat from servers, and liquid cooling is reserved for supercomputers and similar high-performance or high-density applications. In large data centers, cold air is usually produced from a chilled-water cooling system, whereas smaller data centers rely on direct expansion (DX) air conditioning.

Historically, the air conditioning in some data centers was kept quite cold. ASHRAE (formerly the American Society of Heating, Refrigerating and Air Conditioning Engineers) has published guidelines for cooling IT equipment in data centers. Updated in 2008, these guidelines recommend that cooling air in the data center be kept between 64°F and 81°F, with 60% maximum relative humidity (see Table 1-1). Figure 1-2 depicts these guidelines on a psychrometric chart, which is a graphical representation of the thermodynamic properties of moist air at constant pressure. These guidelines make allowances for a wider temperature and humidity range to be used, accommodating free cooling systems that use outdoor air when possible.

Table 1-1
ASHRAE thermal guidelines for data center environments

Guideline		2004	2008
Recommended	Temperature (dry bulb)	68°F to 77°F	64°F to 81°F
	Relative humidity	40% to 55%	42°F to 59° F dew point and 60% relative humidity
Allowable	Temperature (dry bulb)	59°F to 90°F	50° to 95° F
	Relative humidity	20% to 80%	20% to 80%

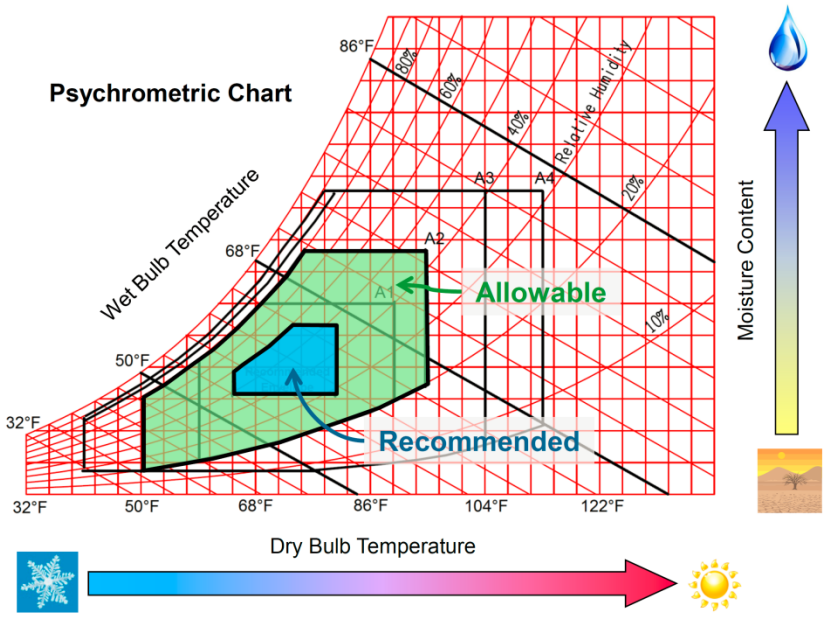


Figure 1-2
Psychrometric chart of ASHRAE guidelines for data centers

By convention, fans are installed to expel warm air from the rear of the server, drawing cool air into the front of the server. The best practice for managing airflow within a data center is to orient the rows of servers so that they are face-to-face and back-to-back. In this way, alternative hot and cold aisles are created in the data center.

Conversion Factors

Table 1-2 lists conversion factors for the units of measure used in this report.

Table 1-2
Conversion factors

Parameter	Conversion
Distance	1 m = 3.3 ft
Temperature	°F = (°C x 9/5) + 32 °C = (°F – 32) x 5/9

2
TECHNICAL ASSESSMENT OF EVAPORATIVE COOLING TECHNOLOGY

In seeking to improve data center energy efficiency, the cooling system is often targeted for improvement because it is typically the largest energy use after the IT equipment itself. Evaporative cooling has emerged as one of the most efficient alternatives to conventional DX (using vapor compression) or chilled-water cooling systems. This section describes the evaporative cooling process, various system architectures, and some of the considerations that must be made in implementing this technology.

Thermodynamic Process

Evaporative cooling is an efficient cooling process in which heat is absorbed by the vaporization (evaporation) of water into the air. This phenomenon is responsible for the cooling effect of perspiration. It is sometimes referred to as *adiabatic cooling* because thermal energy is not removed from the air. Rather, the temperature of the air decreases as thermal energy is absorbed in the change in phase of the water.

Figure 2-1 illustrates the evaporative cooling process overlaid on a psychrometric chart. Beginning from ambient conditions of 80° F and 30% humidity, water vapor is added to the air. This increases relative humidity at constant enthalpy (that is, there is no change in total thermodynamic energy). To achieve 70° air from these conditions requires roughly 1.5 gallons of water be added to the air per hour, assuming 1 ton of total cooling load.

Because thermal energy is not removed, evaporative cooling systems cannot recycle indoor air for cooling, which would cause a buildup of heat and humidity in the data center. Instead, outside air is drawn into the data center and introduced to a large surface area of water. This can be accomplished either by jets that spray a mist of water into the air or by forcing the air through a saturated filter pad. The water absorbs thermal energy from the warm intake air, causing some evaporation, which lowers the temperature of the air. In this way, latent heat of the water vapor is converted to sensible cooling in the data center.

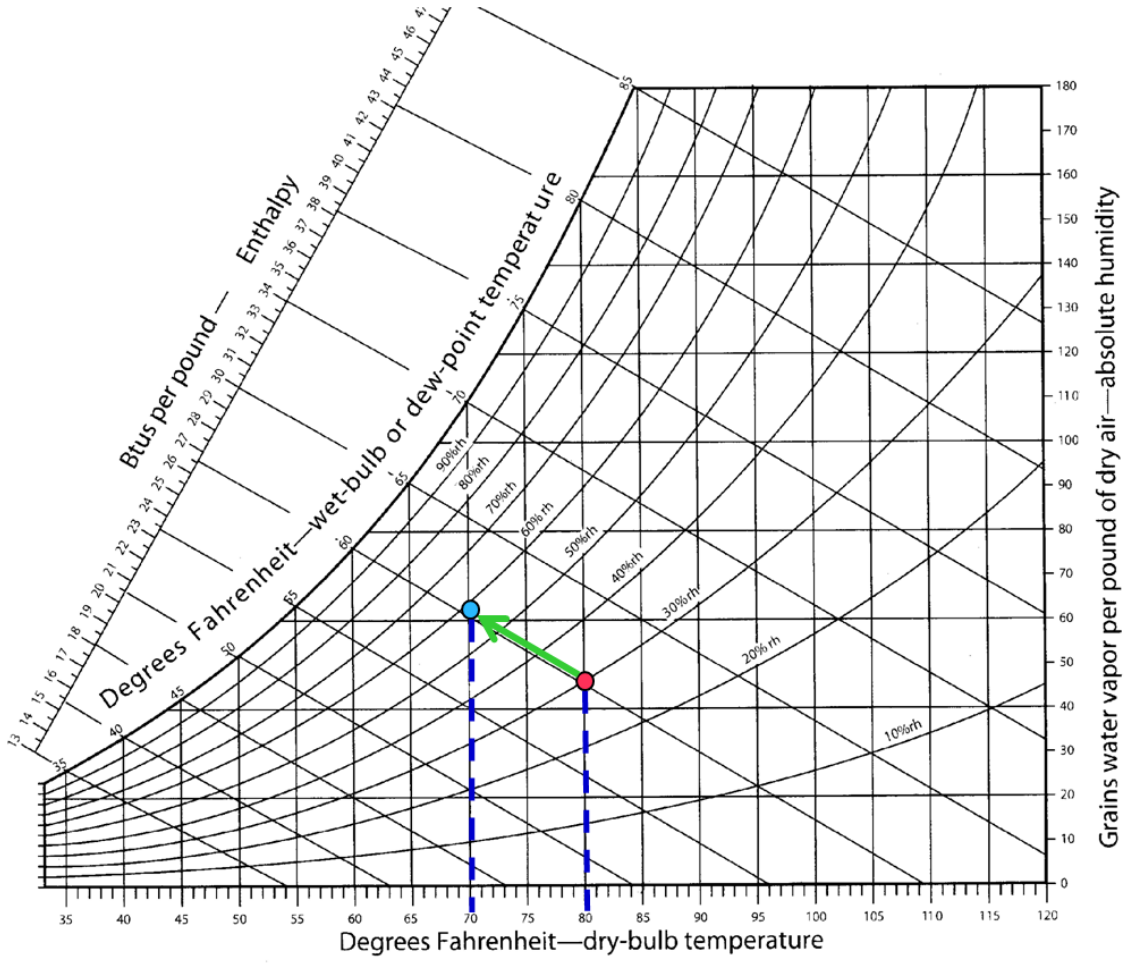


Figure 2-1
Psychrometric illustration of evaporative cooling effect

When relative humidity is high, there is less available capacity for water vapor in the air, limiting the amount of cooling available. For this reason, evaporative cooling is most effective in dry, arid climates, such as the western region of the United States. Evaporative cooling is most effective during summer peak periods because the difference between outdoor dry bulb and wet bulb temperatures is greatest during these periods. Even in humid climates, evaporative cooling aids in reducing summer peak demand.

In some locations, evaporative cooling can be used to provide year-round cooling for the data center. Such an installation can achieve up to 95% savings in cooling energy by eliminating the power-hungry compressors of conventional DX cooling. Considering that cooling is usually 30% to 40% of the data center, evaporative cooling represents nearly 40% savings of total data center energy consumption. Figure 2-2 illustrates the reduction in overall data center energy use achieved by 95% savings in cooling energy from an evaporative cooling system. In this example, PUE decreases from 2.0 to 1.24, for a 38% reduction in total energy consumption.

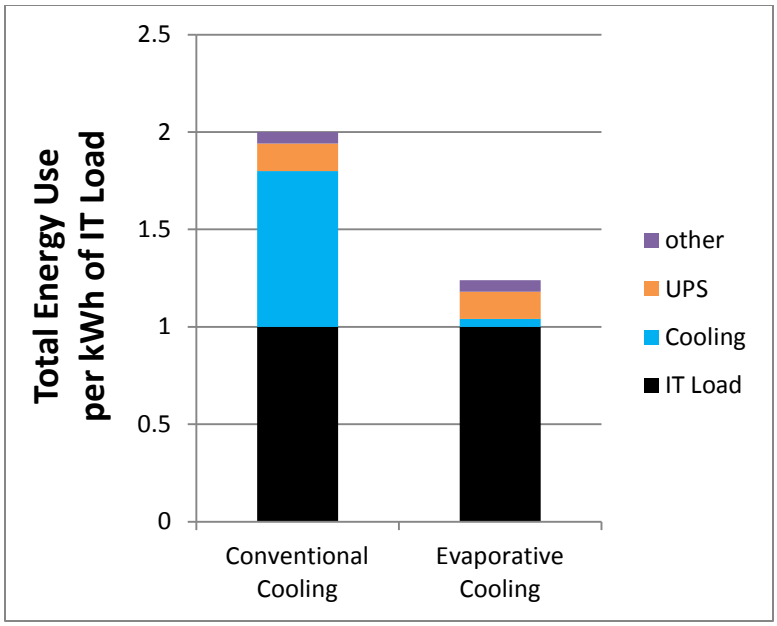


Figure 2-2
Total data center energy reduction from evaporative cooling

Direct Evaporative Cooling

Systems that rely solely on evaporative cooling are called *direct evaporative cooling*, historically known as *swamp coolers*. Direct evaporative cooling systems are commercially available and have been deployed in data centers with some success. However, the effectiveness of these systems is dependent on climate. Figure 2-3 depicts a simple block diagram of a direct evaporative cooling system.

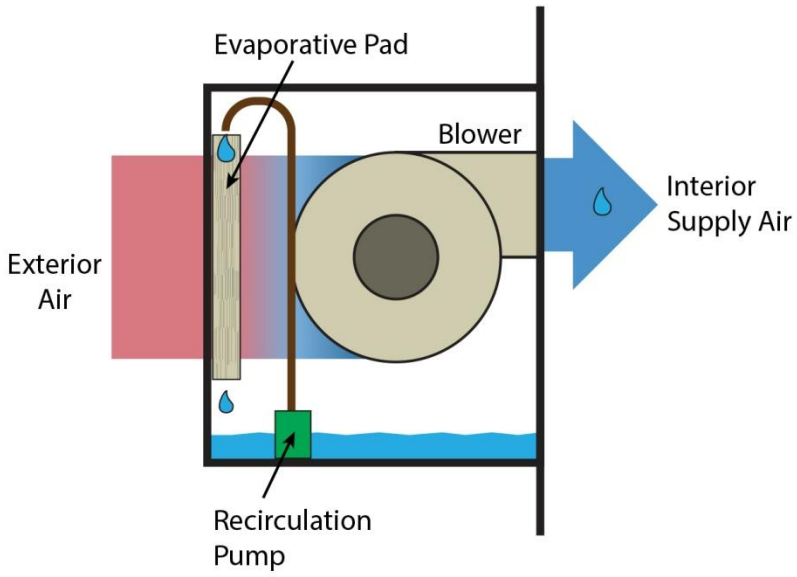


Figure 2-3
Direct evaporative cooling

The major drawback of direct evaporative cooling systems is that they increase the humidity inside the data center. Conventional vapor compression systems, on the other hand, reduce the amount of humidity in the air used for cooling. ASHRAE TC 9.9 guidelines for data centers specify that relative humidity be kept below 80% (see Table 1-1) to prevent condensation from forming on the computing hardware. This limits the amount of cooling available when the humidity of outdoor air is high. In addition, the use of outside air directly in the data center presents the opportunity for introducing contaminants into the data center. Gross particulates are addressed by use of an air filter, which must be cleaned or replaced regularly to ensure proper air flow. Finally, evaporative cooling consumes some amount water, adding to cooling system operating costs.

A vendor in the United Kingdom offers direct evaporative cooling systems for a variety of building types, including data centers. These systems require only 5 W to cool 100 W of IT load and can save up to 90% of the energy required to cool a data center (at locations in the United Kingdom). In total, data centers using these systems consistently exhibit PUE less than 1.3.

One vendor offers direct evaporative cooling systems that inject vapor to cooling air through miniature nozzles. This vendor reports a PUE of 1.06 at a large data center in the Pacific Northwest, which powers a major social networking web site with the help of evaporative cooling.

Indirect Evaporative Cooling

Alternatively, evaporative cooling can be used with an air-to-air heat exchanger, so that humidity is not added to the air entering the data center. Such a design is called *indirect evaporative cooling (IEC)* because evaporative cooling is not used to directly cool the servers. Rather, evaporative cooling removes heat from a working air stream. The product air returned to the data center is cooled through an air-to-air heat exchanger by the working air stream. Figure 2-4 depicts an IEC system.

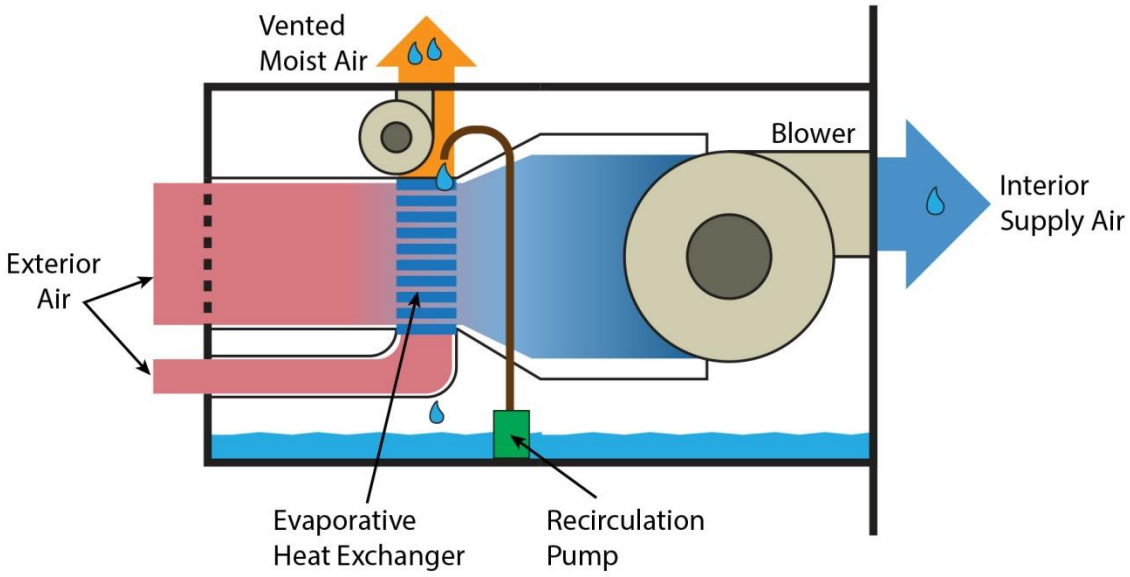


Figure 2-4
Indirect evaporative cooling

IEC systems address some of the largest drawbacks of direct evaporative cooling. Water consumption is generally lower for indirect systems than direct because the working air stream is often a closed loop. With the air-to-air heat exchanger, humid air is not brought into the data center. However, cooling efficiency is slightly less than that of direct evaporative cooling systems. Therefore, IEC systems alone are not able to provide sufficient cooling in warm and humid climates, as is common for most of the eastern half of the United States.

Vendor Case Studies and Savings Claims

One vendor of IEC systems has a number of data center case studies, describing the use of these systems in medium-sized data centers in dry climates in the western United States. These sites demonstrated 85% to 90% savings in cooling energy (about 40% of the total data center energy use), with PUE falling from average values of 1.8–2.0 to 1.1–1.2. Cooling costs, including electricity and water use, are 85% lower, with payback typically in 2–3 years. The vendor cites the following benefits:

- Up to 90% savings in cooling energy
- No humidity added to data center
- Low-maintenance design
- Quiet—less than 70 dBA at a distance of 1 m
- Net zero water impact

Other Evaporative Cooling System Architectures

Two-Stage Evaporative Cooling

For greater cooling capabilities, a two-stage process may be necessary to sufficiently lower air temperature for cooling. Such systems usually use an IEC stage followed by a direct evaporative stage (see Figure 2-5). No humidity is added to the product air by the first-stage IEC, so a direct evaporative cooling stage can be used to reduce air temperature even further.

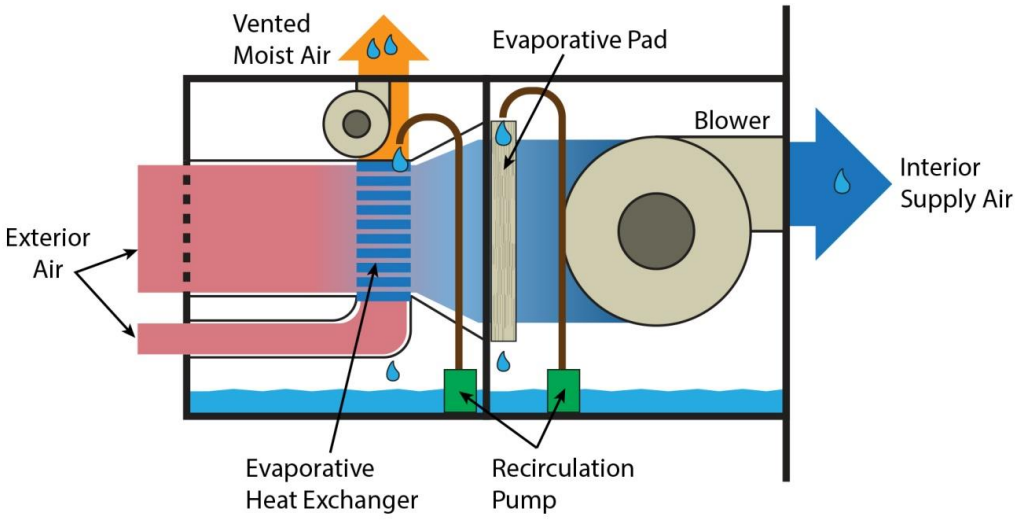


Figure 2-5
Two-stage (indirect–direct) evaporative cooling system

Two-stage systems offer greater cooling capabilities than either direct or IEC alone. However, these systems exhibit similar drawbacks to the single-stage designs, including the addition of humidity to the data center, non-zero water use, and limited cooling capacity available in humid climates.

Hybrid Direct Expansion Systems

For sufficient cooling in humid climates, evaporative cooling can be combined with conventional DX cooling as a hybrid system. Such a design typically uses an IEC as the first cooling stage, followed by a conventional DX or chilled-water cooling stage. Many systems of this type have the ability to recirculate exhaust air when outdoor conditions exceed the temperature and humidity of the return air. This operating mode is essentially the same as traditional cooling system designs with an IEC stage boosting cooling efficiency.

These hybrid systems exhibit all the advantages of IEC over direct evaporative systems—outdoor humidity and contaminants are not introduced into the data center. The conventional cooling stage allows a hybrid system to handle the complete cooling needs of data centers in all climates, including humid climates that are typical for most of the Eastern portion of the United States. Even during hot and humid conditions, the evaporative cooling stage is able to reduce the load on the conventional cooling stage due to the difference between dry bulb and wet bulb temperatures.

One vendor has published a report calculating savings for its hybrid system for large data centers at a number of locations worldwide. The calculations show 59% to 90% savings in cooling costs at locations in all continents and climates, including Abu Dhabi, Beijing, and Moscow. The vendor cites a number of benefits of its systems, including the following:

- Increased efficiency through free cooling using outdoor air
- Evaporative cooling European Seasonal Energy Efficiency Ratio (ESEER) greater than 45
- Fully packaged, complete systems
- Supply and outdoor air completely separated to prevent contamination

Evaporative Condenser Precooling

Another way that evaporative cooling can be used to improve the efficiency of conventional cooling systems is by cooling the condenser coil, where heat is rejected from the compressed refrigerant. These condenser precooling technologies are analogous to hybrid evaporative systems in that evaporative and conventional cooling technologies are used in tandem for enhanced system efficiency. However, these conventional cooling systems do not use outside air, avoiding concerns of adding humidity or contaminants into the data center. Furthermore, evaporative precooling can be a simple and cost-effective retrofit solution for data centers with rooftop DX units.

Recommendations for Application of Findings

This research has shown that evaporative cooling has the ability to significantly reduce the energy used by data center cooling systems. Recent advances have enabled the design of IEC systems that provide most of the savings available from direct evaporative cooling without the negative impacts on air quality and humidity within the data center. In addition, hybrid

evaporative cooling systems allow the use of evaporative cooling to provide all of the cooling required by data centers in all climates.

The results of this work can be used to inform codes and standards for the implementation of evaporative cooling systems—IEC systems in particular—in dry climates where they are most effective. Field demonstrations of hybrid systems can showcase their ability to provide sufficient cooling with significant energy savings in even hot and humid climates. On the whole, evaporative cooling for data centers can be considered for inclusion in utility energy efficiency programs nationwide.

Partial List of Vendors

The following vendors, among others, offer evaporative cooling technologies:

- Applied Air, Aztec
- Coolerado, indirect and hybrid
- EcoCooling (United Kingdom), direct
- Energy Labs, Inc. (ELI)
- Mee Industries, MeeFog (direct, mist jets)
- Munters, Oasis (two-stage, hybrid)
- Trane, Voyager DC

3
TECHNICAL ASSESSMENT OF LIQUID COOLING TECHNOLOGIES

Most servers and desktop computers use air to cool processors, graphics cards, and memory. Yet air is an inefficient heat transfer medium due to its low specific heat, requiring a significant temperature difference between hardware and coolant to remove heat. For example, the mineral oil used in some server cooling systems has 1200 times the heat capacity of air by volume. This has the added benefit of lower power requirements for pumping coolant or moving air. Historically, liquid cooling has been used for high-performance computing applications (supercomputers), over-clocked gaming computers, utility transformers, and traction inverters (electric vehicle drives). With increasing data center loads and more attention focused on reducing cooling system energy use, liquid cooling is gaining attention for cooling data centers of average power density.

To achieve an equal amount of heat removal, air must be cooled much further than liquid coolant. For this reason, liquid cooling systems are able to eliminate powerful chillers and primarily rely on more efficient methods of rejecting heat, including evaporative cooling towers, dry coolers, and facility hot water supply lines. The most efficient systems can realize 95% energy savings over traditional DX (vapor compression) cooling systems. Some of these efficient heat rejection methods rely on outdoor climate—evaporative cooling towers are most effective in dry climates, for instance—and availability of water. Regardless of climate and water availability, liquid coolant enables increasing the temperature of the chilled-water supply, allowing more efficient plant operation in chilled-water systems.

Submersion Cooling

One vendor has recently introduced a novel approach to liquid cooling servers, by immersing entire server racks into coolant (see Figure 3-1). Rather than standing vertically, the server racks are laid down into a bath of nonconductive (insulating) mineral oil. This coolant is compatible with most manufacturers' servers, although before submersion, the server fans should be removed, and the thermal grease should be replaced with a nonsoluble replacement. Heat is rejected from the coolant through a coolant-to-water heat exchanger. When final heat rejection is accomplished with a cooling tower or dry radiator, cooling system energy use can be reduced by as much as 95%. Alternatively, heat can be rejected to a building hot water supply line, recapturing waste energy, or on an existing chilled-water supply.



Figure 3-1
Four server racks submersed in liquid coolant
Source: Green Revolution Cooling, www.grcooling.com

In addition to cooling system energy savings, submersion cooling reduces server power by 10% to 20%. This can be attributed to two causes. First, server fans—which can account for 5% to 10% of server power—are not used and are removed completely from the server before being installed in a submersion cooling system. Second, because this technology is more effective at removing heat, servers have been demonstrated to operate at a lower temperature. In particular, processors have been measured as much as 35°F lower in the coolant bath than when cooled by forced air. This reduces the power consumed by the central processing unit, due to the linear dependence of processor leakage current on operating temperature.

Besides energy savings, submersion cooling offers other advantages over conventional air cooling. First, the liquid coolant removes dust from the surface of electronic components, eliminating potential hot spots. This affects both processor temperatures and reliability. Server reliability is further improved by the elimination of server fans, which can often be a point of hardware failure. Liquid coolant may also be able to provide some thermal ride-through during an interruption in power to the cooling system, due to the added thermal mass of the volume of coolant. In addition, liquid cooling enables increased server power densities. Finally, the vendor claims significant savings in capital expense due to the reduction in backup generation capacity. However, this claim is true only in new data centers for which a generator would be purchased to back up the cooling system.

Case Studies and Energy Savings

The vendor offers two case studies of energy savings from submersion cooling on its website. At one facility, a 100 kW system (four server racks) was installed "for roughly half the price of an air-cooled alternative," including generator costs. System energy use was reduced by 52%, with a reduction in PUE from 1.98 to 1.12 and a 15% reduction in server power. The second facility—a high-performance computing center in Texas—installed a 40 kW, single-rack system with an evaporative cooling tower for final heat rejection. This system yielded about 35% reduction in total energy use over the air-cooled baseline, from the 85% reduction in cooling energy and 10% to 20% reduction in server power (depending on the server model tested). These small-scale case

studies from the vendor illustrate the savings that can be expected from implementing submersion cooling: 80% to 95% cooling and 10% to 20% server energy savings, yielding 30% to 50% overall savings.

Figure 3-2 illustrates the total energy savings available from submersion cooling, including reduction of cooling system energy use and server (IT load) energy use for a typical data center with initial PUE of 2.0.

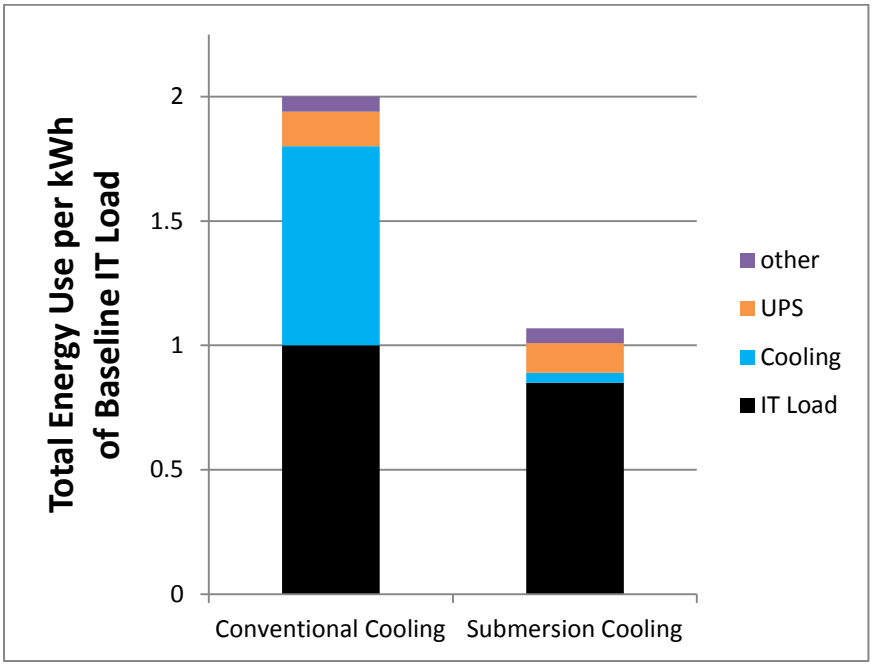


Figure 3-2
Total data center energy use reduction from submersion cooling

Vendor Claims

The vendor cites the following advantages:

- Cooling system power reduced by 95%
- Installation costs reduced by 60% (reduced UPS and backup generator to power cooling)
- Overall data center power reduced 50%
- Reduced server operating temperature
- Increased server density—up to 100 kW per full-size (42-unit) rack

Phase-Changing Liquid Coolant

A major chemicals manufacturer has introduced its own line of nonconductive liquid coolant suitable for cooling servers. One application promoted by this manufacturer is the application of its coolant in submersion cooling. This technology is unique in that the liquid coolant experiences a phase change in the cooling loop. This volatile liquid evaporates after absorbing enough heat from the electronics to reach its boiling point. The saturated vapor rises by natural convection, transferring heat to outside air or facility water, after which the coolant condenses and falls back into the liquid bath.

Unlike the liquid cooling that has been used in utility transformers for decades, these server baths are not sealed, meaning they can be opened for routine maintenance. The two-phase coolant allows passive circulation without the need for pumps or complex plumbing, although few details are given of final heat rejection in a passive configuration. This system inherently loses coolant, but the manufacturer claims that annual coolant loss costs less than the energy that would otherwise be required to power a pump. In addition, the coolant is cited as having reduced greenhouse gas emissions and impact on technician health compared to historical coolants.

Liquid-Encapsulated Servers

This phase-change coolant is used by another company in coolant-encapsulated servers. Offered by a few vendors, encapsulated servers are hermetically sealed and filled with liquid coolant. Coolant is circulated through the server from the bottom, with warm coolant pulled from the top. This system shares many of the same advantages of submersion cooling, including significant reduction in cooling system energy, reduced server operating temperature and power, and elimination of hot spots (improved reliability). The manufacturer cites savings in cooling system consumption of 70% or more, and overall 40% to 60% savings. The primary drawback of this type of liquid cooling is the limited number of server models available, as server encapsulation is not applicable as a retrofit to original equipment manufacturer (OEM) devices. A prototype of this system was demonstrated at the University of Leeds (United Kingdom), and was recently awarded a Green Enterprise IT Award 2013 by the Uptime Institute.

Device-Mounted Liquid Cooling

A more traditional approach to using liquid to cool electronics hardware is the use of cold plates (see Figure 3-3). A cold plate is a liquid-filled heat sink attached directly to the heat-generating component, usually central processing units and graphics processing units but also memory. The liquid coolant is pumped from the cold plate to a heat exchanger for final heat rejection.

Cold plates offer many of the same benefits of other liquid cooling technologies. Again, due to the effectiveness of liquid cooling at removing waste heat, the liquid does not need to be cooled as much as air for sufficient heat removal. This allows use of efficient final heat rejection techniques, such as evaporative cooling towers, in place of vapor compression. In chilled-water cooling systems, liquid cooling allows the chilled-water temperature to be raised for maximum chiller efficiency.

Cold plates are slightly less effective at removing heat than server submersion or encapsulation because the liquid coolant is not in direct contact with the electronics. With cold plates, heat must be transferred from the component (usually the processor) to the cold plate, adding some small thermal resistance to the heat rejection path. Even so, device-mounted liquid cooling is able to save up to 50% of the total energy used by the data center, from the combined improvement in cooling efficiency and reduced server demand (lower processor temperature and removal of server fan).

One vendor of server cold plates offers a few variations to the simple cold plate concept. One product incorporates coolant distribution directly into the server rack. In this system, the coolant pumping and heat exchange modules are built into the server rack, increasing available rack space and facilitating installation and removal of liquid-cooled servers.

A unique system offered by the vendor integrates cold plate and coolant pump into the device-mounted heat sink (see Figure 3-3). The liquid-to-water heat exchanger is built into a fan that replaces the original server fan. In this manner, a closed-loop liquid cooling system is contained within the server. Because conventional air cooling is still necessary for this system, energy savings are lower—up to 10% of the total. However, this device is compatible with a wide range of OEM servers and can be somewhat straightforward to retrofit into existing hardware.



Figure 3-3
Central processing unit–mounted cold plates with in-rack coolant distribution
Source: Asetek, www.asetek.com

A 2013 study by LBNL confirmed the manufacturer's energy savings claims, finding 50% savings in cooling energy plus reduction in server energy. The study used device-mounted liquid cooling to supplement the conventional chilled-water cooling system, yielding a 16% to 24% reduction in total energy use at an already efficient data center.

Vendor Claims

The vendor cites the following advantages:

- Cooling system energy savings exceeding 50%
- Rack-level server density increases of 2.5 times
- Highest server density form factors enabled
- Immediate to 1-year payback cycles typical
- High-temperature data center operation enabled

Market Barriers

Despite the superior heat transfer capabilities of liquid cooling, this group of technologies has not been widely used. Perhaps the most significant market barrier facing liquid cooling is the added complexity of these systems compared to conventional air cooling. This added complexity increases initial expense of the hardware, design, and installation, and it affects system reliability. Reliability concerns are primarily focused on the added complexity of liquid cooling systems. Server reliability is generally improved due to server fan removal (in immersed or encapsulated servers) and reduced processor operating temperature. Independent verification could be used to further prove the reliability of liquid cooling systems through field demonstrations and deployment in real-world data center installations.

Recommendations for Application of Findings

The results of this work can be used in two primary ways. First, the emergence of submersion cooling has created an opportunity for field demonstrations to verify the effect of this technology on server reliability. The energy savings potential of this technology is dramatic, yet this technology faces the perception that liquid cooling affects availability and adds to the cost of operation and maintenance of the data center.

Second, liquid cooling can be considered for utility energy efficiency programs. The inherent efficiency gains of these systems could allow the development of prescriptive incentive programs for the replacement of air-cooled systems with liquid-cooled systems. In data centers that already use chilled water for final heat removal from computer room air conditioners—not uncommon in medium-sized data centers (roughly 100 kW and larger)—replacement of the air conditioner with a liquid cooling system can be somewhat economical. Moreover, liquid cooling can be used to expand the computing capacity of existing data centers in which water-chilled air conditioning is operating near capacity.

Partial List of Vendors

The following vendors, among others, offer liquid cooling technologies:

- Asetek, device-mounted cold plates, rack-integrated coolant distribution
- Clustered Systems, device-mounted
- Green Revolution Cooling, submersion cooling
- Iceotope, encapsulated servers using 3M Novec coolant
- LiquidCool Solutions (formerly Hardcore Computing), encapsulated servers

4
TECHNICAL ASSESSMENT OF UNINTERRUPTIBLE POWER SUPPLY ECO MODE

As modern society becomes increasingly dependent on the communications services that data centers power, the availability of these facilities is required to be kept quite high—sometimes more than 99.99% of the year. The UPS is a critical component of a data center's infrastructure, ensuring that power is delivered to the data center through interruptions in utility power. Given that these devices are energized continuously, often supplying enough power for IT equipment and cooling systems, their efficiency can have a substantial impact on the data center's overall energy use. This section describes high-efficiency eco mode UPS designs that have emerged in the data center UPS market.

Data Center Critical Power Requirements

Depending on the types of services and industries it supports, a data center may be required to have extremely high reliability. One way that data center reliability is measured is by *uptime*—the availability of digital services as a percent of the time of operation of the data center. The Uptime Institute, a data center industry consortium, has defined four tiers of data center availability to provide a standardized methodology for designing and comparing data center performance. Table 4-1 shows the four tiers defined by the Uptime Institute, with tier 4 providing the highest availability of services.

Table 4-1
Data center tier classification

	Tier 1	Tier 2	Tier 3	Tier 4
Capacity redundancy	N	N+1	N+1	N after any failure
Uptime	99.67%	99.75%	99.98%	99.99%

A number of factors impact uptime, but chief among them is availability of power. Many facilities use of a diesel backup generator to protect against utility power outages. However, backup generators are slow to respond to electrical disturbances, requiring seconds to bring up to full power. A UPS is used to respond more quickly than backup generators are able, providing protection from a number of quick power quality (PQ) disturbances, including voltages sags (or dips) and momentary interruptions in power, which can last from a fraction of a line cycle (16.7 ms for 60 Hz supply) to a few seconds. In facilities without backup generators, the data center may rely on the UPS for power through an outage until the UPS energy storage has been depleted.

Voltage sags are the most common PQ disturbance to cause an interruption to electrical equipment. They can last anywhere from one line cycle to tens or hundreds of milliseconds. These events are usually caused by an abrupt increase in load (current) elsewhere on the circuit. Typical culprits may include the inrush current into motors, pumps, or compressors or large resistive loads, but they can also include abrupt increases to source impedance from loose

connections. Voltage sags can be quite common at industrial sites, possibly occurring several times per year at the service entrance to the facility. More often, voltage sags are generated within a building's wiring from large loads such as air conditioning, electric heating, or refrigeration turning on. For this reason, data center air conditioning is often wired on separate circuits from computing equipment so as not to cause additional disruptions.

Uninterruptible Power Supply Background

Historically, data centers have relied on UPS with double-conversion (online) architecture, in which power is continuously supplied to the servers from the battery of the UPS. In this design, all power is rectified from input ac to dc to charge the battery, from which the inverter creates output ac. Thus, the entire output power has been double-converted—from ac to dc and back to ac. Figure 4-1 shows the simple block diagram of double-conversion UPS architecture.

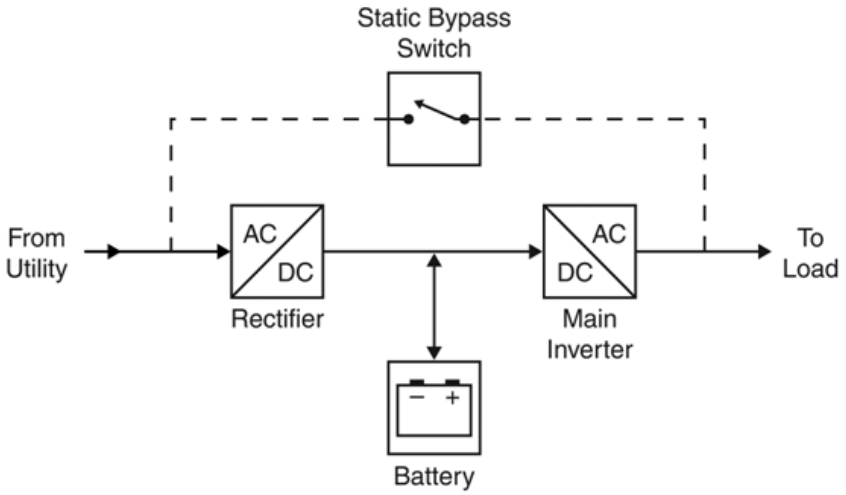


Figure 4-1
Block diagram of double-conversion (online) uninterruptible power supply architecture

The double-conversion architecture is the most reliable, transferring the least amount of disturbance to the output from a disturbance in input PQ. This is because the output is continuously powered from the battery. However, this topology can be somewhat inefficient, due to the losses in the rectifier, battery, and inverter. Figure 4-2 shows the measured efficiency of typical (legacy) double-conversion UPSs.

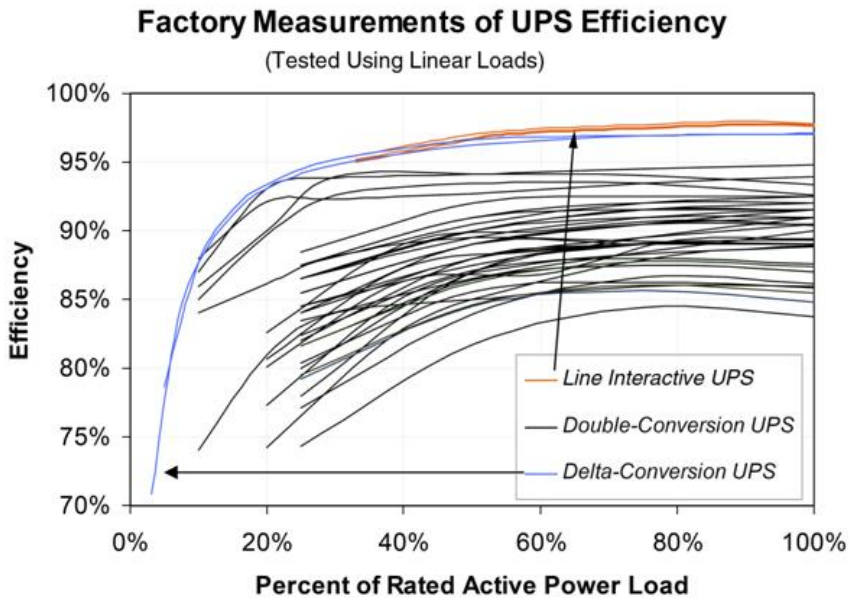


Figure 4-2
Double-conversion uninterruptible power supply efficiency

Other UPS architectures boast higher power conversion efficiency than double-conversion. Standby (offline) UPSs remain in bypass during normal operation and transfer to battery backup when needed (see Figure 4-3). These designs are generally considered too slow in responding to PQ disturbances for use in data centers. Line interactive UPSs are similar to standby devices, but they keep the inverter energized to provide some level of additional filtering and voltage regulation to the output. The line interactive design is generally more efficient than double-conversion architecture and provides greater reliability and higher quality power than standby designs.

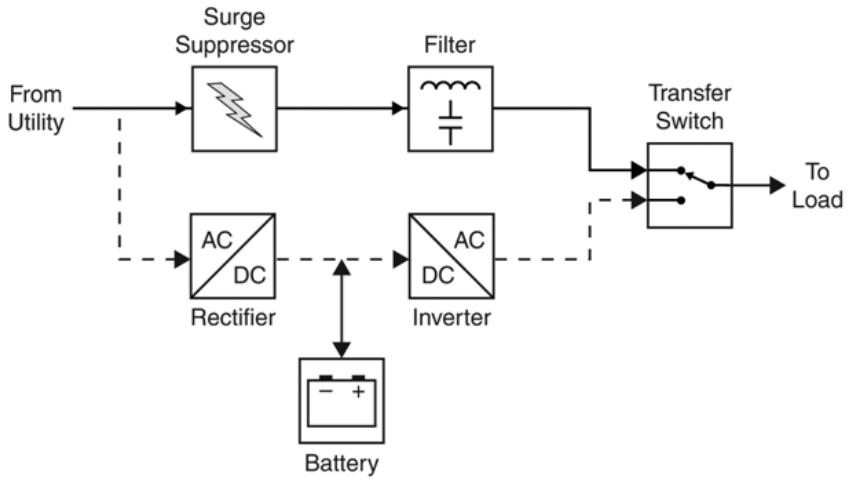


Figure 4-3
Block diagram of standby (offline) uninterruptible power supply

Different performance classifications of UPSs are defined by IEC 62040-3 based on the dependence of their output power on input power in normal mode operation. These classifications are defined by whether voltage and/or frequency of output power is dependent on that of the input power, as follows:

- **Voltage and frequency dependent.** Output is dependent on any changes in input voltage and frequency without conversion or compensation. A standby UPS is an example.
- **Voltage independent (VI).** Voltage is converted or compensated to remain within a narrower tolerance window, but frequency is dependent on input. A line interactive UPS can be classified as voltage independent.
- **Voltage and frequency independent.** Output voltage and frequency is independent of input power, providing protection from variations without depleting energy storage. A double-conversion UPS can be classified as voltage and frequency dependent.

Eco Mode

Recently, a number of UPS manufacturers have advertised a high-efficiency eco mode that is capable of protecting against PQ disturbances with no degradation in performance. Eco mode is sometimes referred to as *multi-mode* by the EPA Energy Star program to distinguish eco mode as a second, normal mode of operation (in addition to double conversion) that protects output equipment. The high efficiency of these designs—typically 98% to 99%—is achieved by operating the UPS in bypass until a disturbance is sensed or input power falls outside a tolerance window. At this point, the UPS transfers to double-conversion mode, in which output power is supplied through the inverter, ready to quickly respond to an interruption in input power.

Eco mode is distinct from the standby mode of a double-conversion UPS in that eco mode is a normal mode of operation that provides protection for PQ events. Eco mode operation is similar to that of standby UPSs, but a couple of differences make this design better suited to the high reliability required of data centers. An eco mode UPS is able to switch to inverter output much more quickly than a standby UPS. In most cases, eco mode is set with a narrower tolerance window for input power, meaning that transfer to double conversion is usually initiated at voltage or frequency disturbances that a standby UPS would tolerate. These under- and over-voltage and frequency limits are often user programmable.

Efficiency Gains

Eco mode is highly efficient, usually 98% to 99%, and it offers significant savings over more common UPS modes of operation. This high level of efficiency is the result of minimal power conversion steps between the grid and server, similar to bypass mode or standby UPS operation. While operating in eco mode, typically the only components energized are the bypass switch (solid-state or electromechanical relay), some sensing and control circuitry, surge protection devices, and perhaps the output transformer. The only components that carry the full output current are the bypass switch and, in some cases, the output transformer. In contrast, double-conversion mode requires the full output current to be carried by the ac–dc rectifier (battery charger) and the dc–ac output inverter. Figure 4-4 shows the efficiency of both eco mode and double-conversion mode of a modern UPS that offers eco mode.

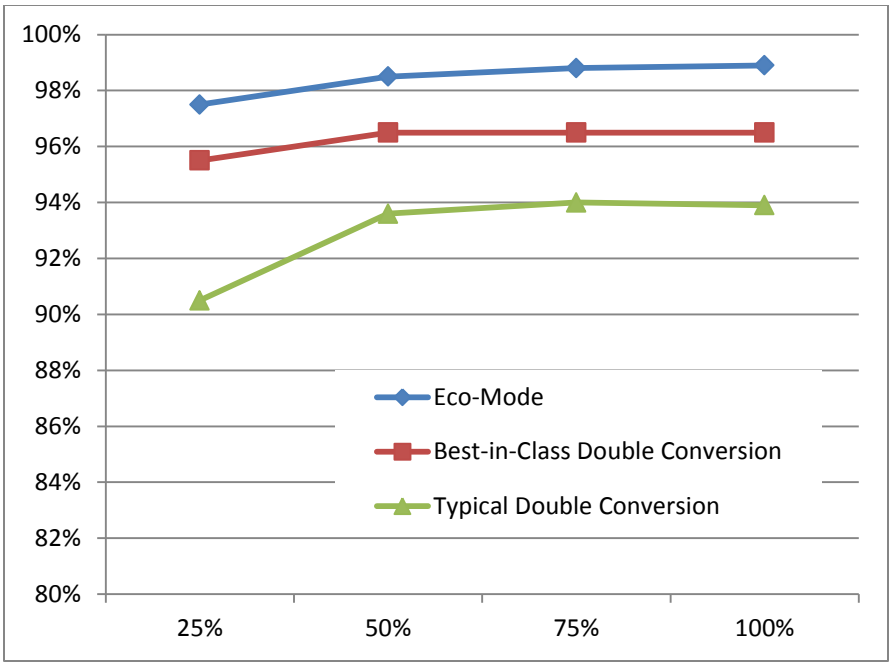


Figure 4-4
Efficiency of eco mode and double conversion for uninterruptible power supplies

As shown in Figure 4-4, an eco mode UPS will exhibit about 4% to 5% higher efficiency when operated in eco mode than in full double-conversion mode. This signifies up to 80% reduction in losses in the UPS power conversion process. According to a survey published in 2007 by Lawrence Berkeley National Lab, UPS losses represent roughly 5% to 10% of data center energy use overall. In addition, losses dissipated in the UPS add heat to the data center, increasing the load on the cooling system. Therefore, reduction in UPS losses is compounded by reduced cooling use.

Equipment as large as the UPS in a data center is often purchased in hardware lifecycles of 10 years or longer, so savings from installation of eco mode in legacy facilities can be even more dramatic. As shown in Figure 4-4, the efficiency of legacy double-conversion UPS equipment may be 80% to 90%, depending on the UPS size and loading level. Table 4-2 shows the impact that using an eco mode UPS would have on PUE and overall energy use in a variety of data center facilities. In total, use of an eco mode UPS will typically save 4% to 13% of the total energy consumed by a data center.

Table 4-2
Eco mode energy savings

Data Center	Poor	Legacy	Typical	Modern	Best in Class
Baseline UPS	80%	85%	92%	94%	96%
PUE	2.5	2.0	1.8	1.6	1.3
Eco mode	98%	98%	98%	98%	99%
Reduction of UPS losses	90%	87%	75%	67%	75%
New PUE	2.04	1.73	1.69	1.53	1.26
PUE improvement	0.46	0.27	0.11	0.07	0.04
Total energy savings	18.4%	13.3%	6.1%	4.1%	3.0%

Figure 4-5 shows the savings available from eco mode operation of the UPS in the five data center classes identified in Table 4 2.

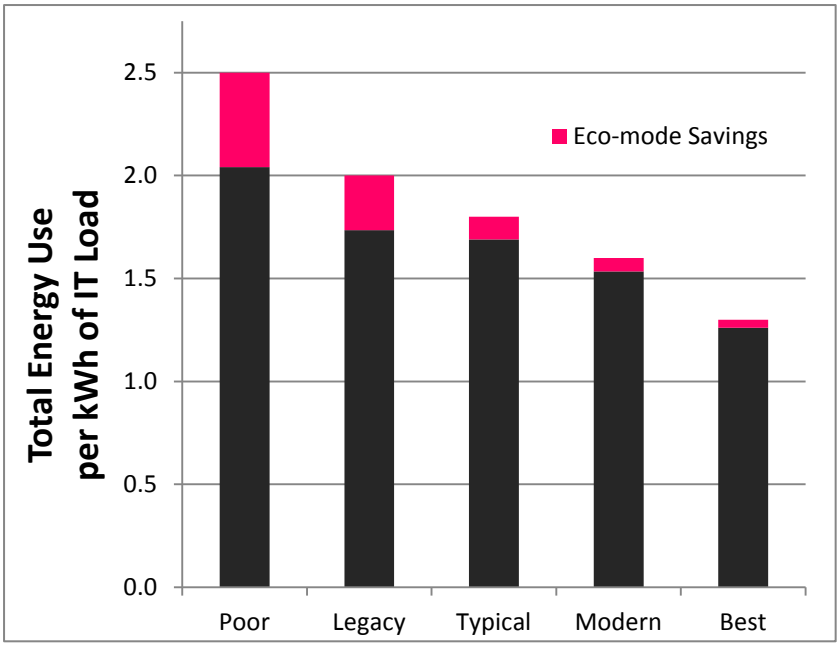


Figure 4-5
Eco mode energy savings

Eco Mode Performance

To transfer from bypass to double conversion, the UPS must complete three steps. First, any disturbance in input voltage must be recognized, requiring precise voltage sensing, analog-to-digital conversion, root-mean-square (RMS) calculation, and quick identification of over- or under-voltage condition. Second, the battery inverter must be energized to prepare for delivery of output power, which involves powering the inverter control circuitry, synchronizing the inverter and the output ac, and energizing the transformer at the output of the inverter. Finally, the transfer switch must disconnect the input (mains) from the load so that the inverter can deliver power to the load in an isolated connection.

There is some inherent delay in this process, creating a transfer time between loss of input power and resumption of high-quality output power. Table 4-3 shows the eco mode transfer time specified for data center class UPSs from three of the largest manufacturers. Vendors A and B specify that transfer to double conversion is accomplished in less than 2 ms, with complete restoration of high-quality output power achieved in less than 4 ms. Other designs may require up to one 60-Hz line cycle (approximately 17 ms) to fully energize the output transformer. One manufacturer has patented its method of energizing the secondary of the output transformer through the bypass circuit, which is connected in zigzag or interconnected-star configuration. This design enables transfer to the battery inverter in less than 4 ms.

Table 4-3
Eco mode transfer time specifications

Specification	Vendor A	Vendor B	Vendor C
Rating	225–750 kVA	225–1100 kVA	250–1100 kVA
Transfer time	< 2 ms	< 2 ms	15 ms

One of the greatest differences in eco mode between different UPS manufacturers is the choice of transfer switch. The slowest component used to switch from transfer to double conversion is a single-pole switch, preventing the battery inverter and ac bypass from being simultaneously connected to the output. More sophisticated switching can be accomplished with electromechanical or solid-state relays. These devices are used to connect the bypass or the battery inverter to the output, and often a separate relay is used for each path. Solid-state (thyristor-based) devices are generally much faster than electromagnetic relays and will provide superior transfer performance from bypass to double conversion.

Currently, there is no standard metric for evaluation of eco mode transfer time performance. The transfer time performance of a UPS is listed on the data sheet and cited by the manufacturer in marketing material, but there is no standard test procedure or independent verification of UPS performance. Without confirmation of adequate protection from PQ disturbances, data center operators may have some concern that eco mode operation will negatively impact server uptime. Validation of performance may be the greatest market barrier faced by eco mode. Development of applicable codes and standards for testing and certification of eco mode transfer time are required to alleviate users' concerns and accelerate adoption of this efficient technology in the market.

Voltage-Sag Immunity of Information Technology Equipment

Eco mode's quick transfer performance (compared to standby or offline UPS designs) enables high-quality output power suitable for servers. To prevent interruption of server operation, the transfer characteristics of the UPS should be considered in relation with the ride-through capabilities of the server PSU. Because a server PSU contains some energy storage, as both capacitors and inductors in the ac–dc conversion stage, it is able to "ride through" some PQ disturbances, including short-duration voltage sags and momentary interruptions in power.

The Information Technology Industry Council (ITIC) (formerly the Computer & Business Equipment Manufacturers Association) developed and published a PQ event immunity curve for IT equipment, as a specification to be used by equipment manufacturers and power providers.

Figure 4-6 shows the voltage-sag immunity curve updated in 2000 by ITIC to reflect modern technology. The two blue lines represent the limit of voltage deviation and event duration that can be tolerated without interruption. Equipment should continue to operate without interruption for all input power disturbances that fall between these limits. IT equipment that is designed to meet this specification will ride through momentary interruptions in power (0 V) for up to 20 ms.

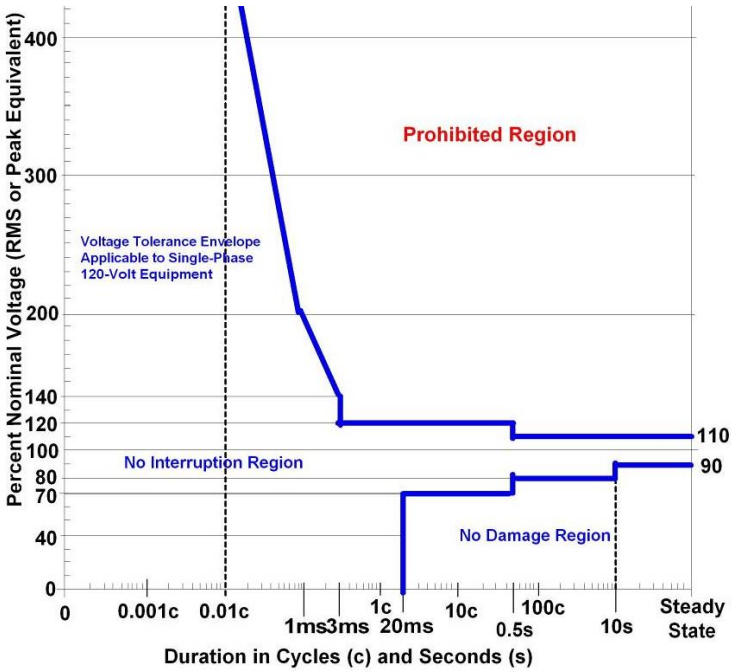


Figure 4-6
Information technology voltage-sag immunity curve (revised 2000)

However, server power supply ride-through is dependent on loading level. With redundant PSUs or over-provisioned capacity, PSU ride-through increases. Legacy equipment may have been oversized to the point that servers could ride through multiple cycles of power interruption. Knowing this, some PSU manufacturers no longer design for 20 ms of ride-through. The ride-through capability of some low-cost server PSUs has been reported to be as low as 10 ms.

To examine this claim, the voltage-sag immunity of four typical server PSUs was measured in the laboratory. Figure 4-7 displays the results of this testing, showing the minimum voltage required for uninterrupted operation versus the duration of the PQ event. As shown in the figure, these devices were able to ride through longer than 10 ms at 0 V. These results indicate that UPSs with quick eco mode transfer time would completely satisfy the needs of these power supplies. However, UPSs with transfer time of 15 ms could cause interruption of PSU D at full load and PSU B at 100% and 75% loading. For this reason, the transfer time should be considered when selecting an eco mode UPS for operation of servers without interruption.

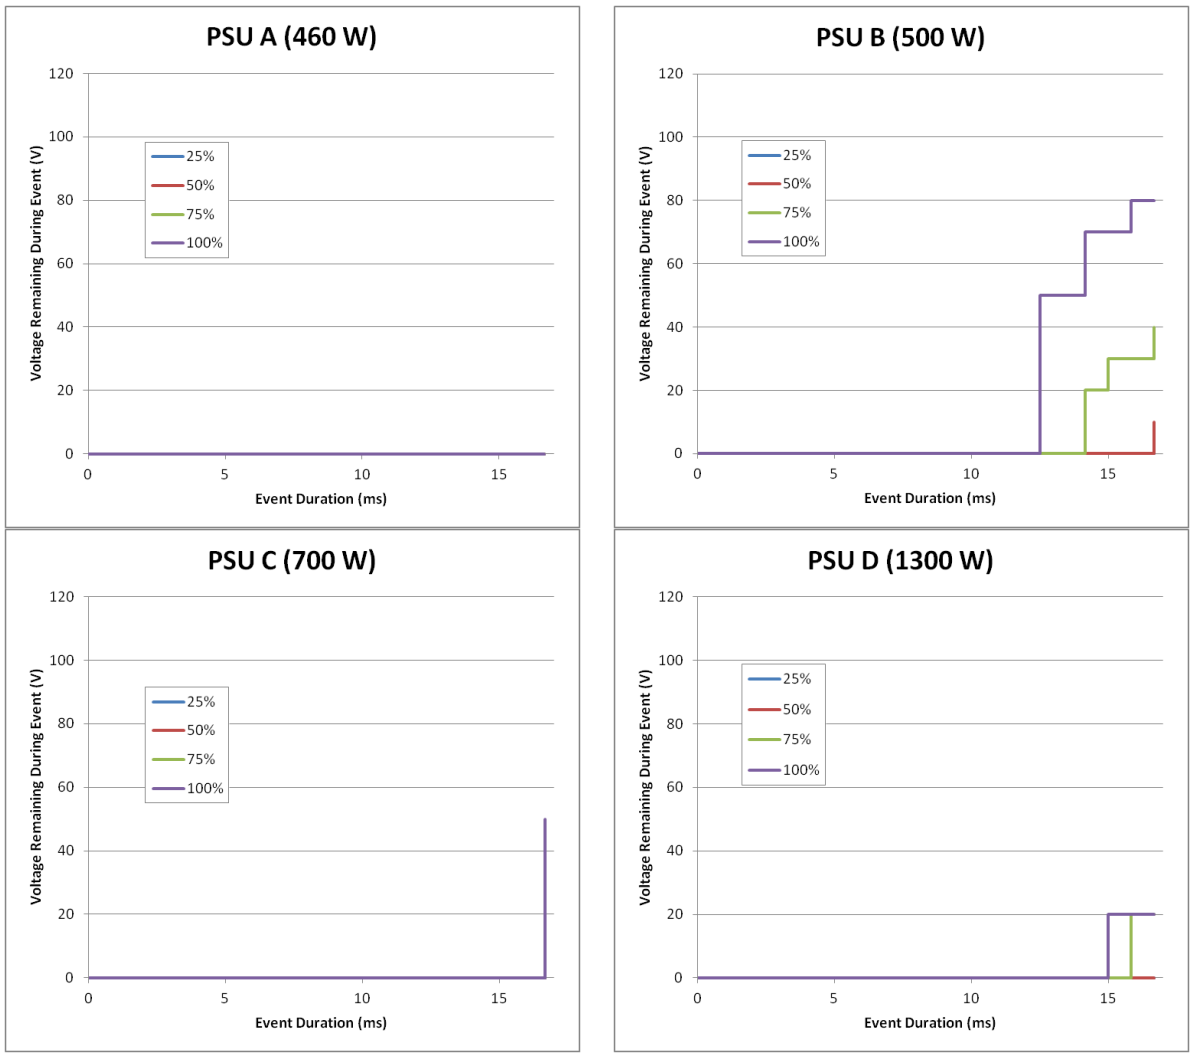


Figure 4-7
Measured voltage-sag immunity of server power supply units

This change in the ride-through capabilities of PSUs is reflected in the IEC 62040-3 (2011 revision) specifications for output performance for UPSs that serve general-purpose IT loads. These specifications call for the UPS to allow no more than 10 ms of interruption of power to IT loads, as shown in Figure 4-8.

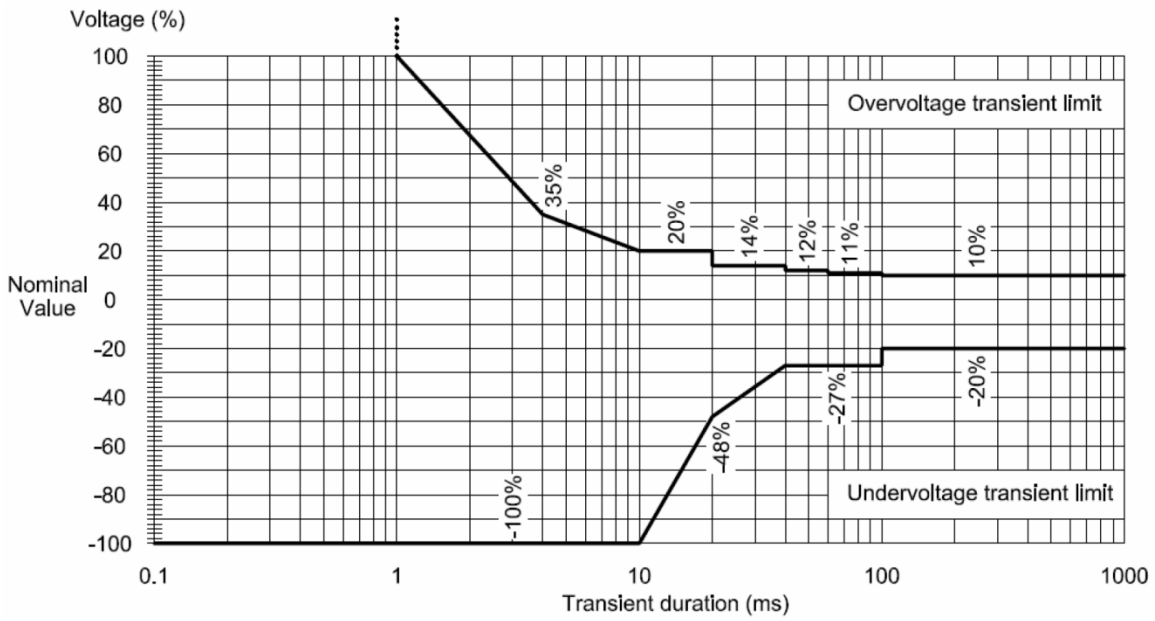


Figure 4-8
IEC 62040-3 specification for uninterruptible power supply output performance for general-purpose information technology loads

To ensure uninterrupted operation of servers, the UPS must be able to restore high-quality output power in less time than the ride-through capabilities of the server PSU. As shown in Table 4-3, some eco mode products are capable of meeting the needs of all server power supplies. However, some UPSs may not be able to transfer within the 10 ms ride-through of low-cost server PSUs. When selecting an eco mode UPS, the transfer time capability of the UPS and ride-through of the PSUs should be considered to prevent interruption of the servers. Figure 4-9 shows the range of fast voltage-sag events that an eco mode UPS is able to protect against, given a transfer time of 4 ms.

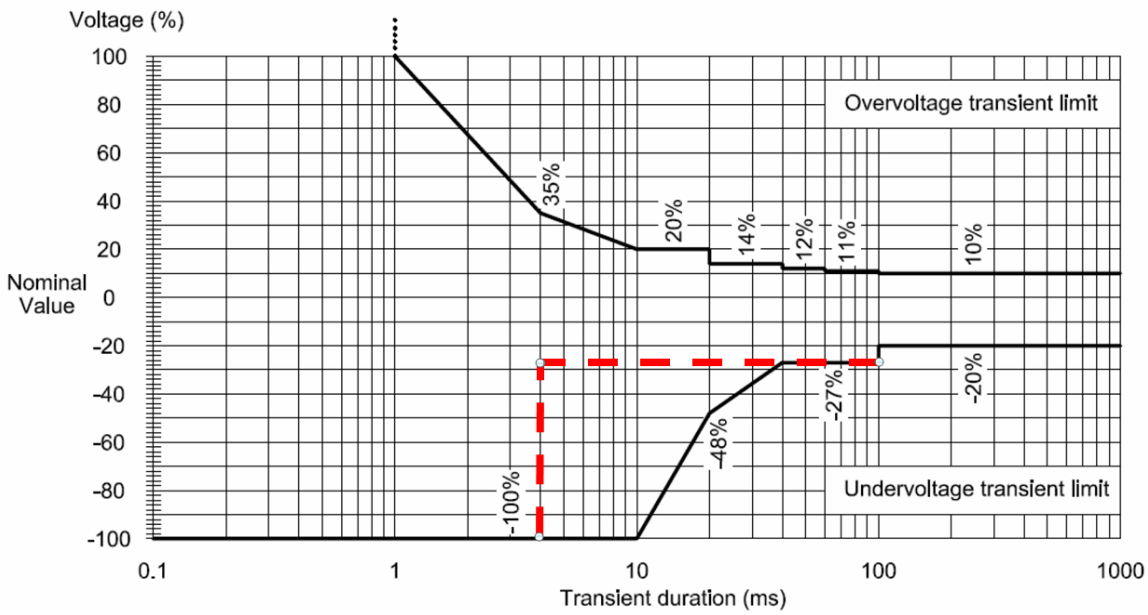


Figure 4-9
Voltage-sag protection provided by eco mode

Currently, there is no standard metric for characterizing the transfer time capabilities of eco mode UPSs. For this reason, there is some divergence in opinion whether eco mode is able to satisfy the PQ immunity requirements of server PSUs. This has created some doubt among data center operators regarding the reliability of eco mode, creating a market barrier to the adoption of this efficient technology. This barrier could be addressed through development of a standard metric and test protocol to determine the transfer time performance from eco mode to double conversion. Such a metric could be incorporated onto product labeling, similar to the 80 PLUS computer PSU efficiency program.

The response of any static transfer switch (STS) in the data center must also be considered when using eco mode operation in a data center. These STSs are used to transfer a server PSU between redundant power sources if the PSU has only one input power cord. The STS is designed to transfer from one primary source to another when a PQ disturbance has been sensed. If the STS switching threshold is set to be too sensitive, the STS may sense a disturbance in power and transfer the PSU to the alternative power source before the eco mode UPS has switched to double conversion. If eco mode is to be used in conjunction with STS, the transfer conditions of both devices should be set to avoid this misbehavior.

Reliability Considerations

Although eco mode exhibits extremely high efficiency, it provides somewhat less protection to IT equipment than full double-conversion mode. Most eco mode UPS designs provide no isolation between input and output. This allows high-frequency harmonic currents from devices in the data center to be transmitted to other loads on the circuit and vice versa. In addition, servers may be exposed to somewhat greater transient voltage disturbances (surges) from lighting discharge and utility capacitor banks. (Used to correct for power factor, utility capacitors can cause a ringing voltage disturbance when connected in and out of the circuit.)

Generally, multiple methods of voltage surge protection are used in concert to suppress these transient disturbances. Surge protection can be located at the service entrance to the customer facility, the front end of the UPS, and/or the input to the server PSU. Even in eco mode, the UPS usually uses some input filtering and surge protection to protect IT equipment from noise, but the isolation provided in double-conversion mode helps limit the instantaneous disturbance transmitted to servers. Therefore, eco mode could impact the lifetime of capacitors and metal oxide varistors, which are used for surge protection in the server PSU.

Similarly, the lifetime of the battery and inverter may be impacted from intermittent use in eco mode. When not used constantly, these devices experience greater thermal cycling. If these components are designed for continuous use in double-conversion mode, the repeated thermal expansion could lead to earlier failure. Eco mode operation may require more frequent maintenance of the backup battery, perhaps calling for quarterly verification of battery health and eco mode transfer capabilities.

Finally, fault protection capabilities may suffer due to the difficulty of differentiating between output fault and loss of input power while in eco mode. However, a well-designed UPS should be able to differentiate between the two, transferring to battery backup on loss of input power and remaining in bypass mode during an output fault to trip upstream circuit breakers.

Site Power Quality

The energy savings available from eco mode operation is related to the PQ experienced at the data center. Sites with numerous PQ disturbances will require the UPS to transfer to full double-conversion mode for a greater portion of its operation. For example, one UPS vendor has presented field data from a site where an eco mode UPS operates in double-conversion mode for as much as 20% of the year, which was attributed to capacitor switching events on the distribution feeder.

In the early 1990s, EPRI began its Distribution Power Quality (DPQ) project to evaluate the frequency and severity of PQ events. Between 2000 and 2002, a follow-on project, referred to as *DPQ II*, was conducted to update the results of the first DPQ study. The project characterized PQ in terms of short-duration variations such as voltage sags, voltage swells, and voltage interruptions. These were based on analysis of data from 480 PQ monitors at different locations on the power system at 13 utilities across the United States.

Figure 4-10 shows the annual rate of voltage sags and interruptions, shown as the number of events recorded annually by voltage and duration of event. These results show the average of all the data measured at the distribution feeder and low-voltage customer, as aggregated in 60-second periods. Figure 4-10 shows that the average site experienced 0.7 incidents per year in which the minimum voltage during a 60-second window was less than 70% of the site's base voltage for one cycle or less. This represents the number of times that a sensitive server PSU could be interrupted by an eco mode UPS with slow transfer time to double-conversion mode.

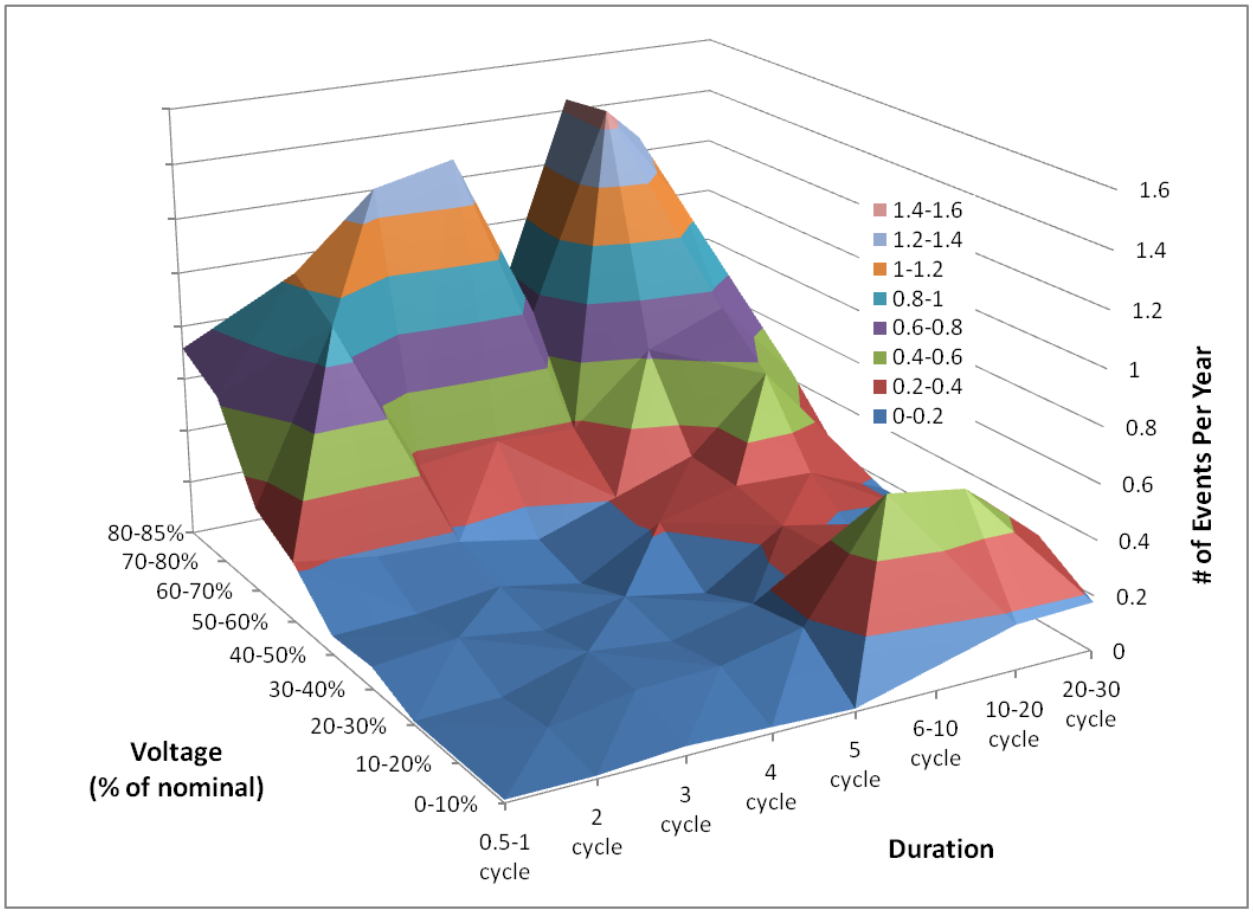


Figure 4-10
Annual rate of voltage sags and interruptions measured at customer-side sites

A metric that is useful in measuring the number of short-duration voltage variations is the system average RMS (variation) frequency index (SARFI). SARFIITIC represents the number of events that fall outside of the ITIC voltage-sag tolerance curve (shown in Figure 4-6). The DPQ II project found that distribution feeder and customer sites experienced 24 interruptions annually that fell outside the ITIC curve. To estimate the number of hours of eco mode operation, consider an eco mode UPS that experiences 24 interruptions in one year. If the UPS were configured to remain on battery backup for 5 minutes before switching back to eco mode after an interruption, the UPS would operate on battery backup for only 120 minutes, or 2 hours, of the entire year. This signifies that eco mode is available for use 99.95% of the year.

To achieve maximum energy savings, a data center would choose a site with few PQ disturbances so that eco mode could be utilized for > 99% of the year. The Green Grid, a data center industry consortium, compared these site-dependent savings to *free cooling*, which is the use of outdoor air in place of traditional vapor compression air conditioning to cool a data center. This concept of free critical power could be used to attract data centers to sites where UPS losses can be minimized through year-round reliance on eco mode.

Recommendations for Application of Findings

The findings presented in this project can be used to promote adoption of this efficient technology in two ways. First, there is a need for codes and standards to verify the compatibility of eco mode UPSs with server PSUs. A UPS standard could be developed that specifies the maximum transfer time allowed for transfer from bypass to battery backup. A standardized test protocol is needed, along with a defined metric that communicates more details about the maximum power disturbance presented by the UPS in transfer to battery backup. In addition, power supply standards could specify the minimum PQ disturbance that a PSU can ride through without disruption. With standard test protocols, performance metrics, and design codes, the operation of the PSU without disruption from an eco mode UPS can be verified.

Second, a field demonstration of eco mode devices would illustrate the impact of real-world PQ disturbances on the reliability and energy savings offered by eco mode. Such a study could verify the number of hours of eco mode operation given the PQ experienced at a site. The results of a field demonstration would assess the relationship between eco mode energy savings and site PQ, informing the potential for development of utility energy efficiency programs.

5
CONCLUSIONS

As demand for digital entertainment and cloud services continues to increase, data center operators are striving to improve the efficiency of their facilities and server hardware. This report describes three technologies that have the potential for significant energy savings. Cooling systems often comprise 30% to 40% of the total energy use in a data center and are frequently targeted for energy efficiency upgrades. Evaporative cooling boasts up to 90% savings over conventional cooling systems (nearly 40% of data center total), and recent advances in indirect evaporative, two-stage, and hybrid DX systems have enabled the use of this efficient technology in all climates. Because of its superior heat removal effectiveness, liquid cooling offers up to 50% savings of total data center energy use from nearly 95% reduction in cooling and 10% to 20% lower server demand.

After cooling, UPS losses are usually the next greatest consumer of overhead energy in the data center, generally 5% to 10% of total energy use. To provide uninterrupted services, data center UPSs typically use double-conversion architecture, which can exhibit efficiency as low as 80% in legacy facilities. Eco mode is a high-efficiency operating mode in which the UPS is able to transfer to battery backup quickly enough that server equipment is not interrupted. Eco mode is usually 98% to 99% efficient, yielding roughly 15% savings in total energy in legacy facilities. Some products can achieve transfer to battery backup in less than 4 ms, well within the ride-through capabilities of most modern server power supplies. The greatest market barrier that this efficient technology faces is data center operators' reliability concerns. This project shows that eco mode is able to provide sufficient protection to prevent server interruptions, but not all eco mode products perform equally. Before deploying an eco mode UPS, a data center operator should consider the transfer time exhibited by the UPS and the ride-through capabilities of the server PSUs.

Table 5-1 shows the savings available from the technologies examined in this research as a percent of total data center energy use. These technologies are all commercially available and ready for deployment in new and existing data center installations. Given the ubiquitous nature of data center UPS use and quick payback of the technology, further standards development and technology demonstrations offer a cost-effective means for a utility to promote the adoption of this efficient technology. In addition, evaporative cooling and liquid cooling each offer tremendous data center energy savings that could be reasonably justified in numerous sites with appropriate utility incentive.

Table 5-1
Total savings available from technologies examined

Technology	Energy Savings
Evaporative cooling	< 40%
Liquid cooling	30% to 50%
UPS eco mode	3% to 18%

6
REFERENCES

1. Jonathan Koomey. *Growth in Data Center Electricity Use 2005 to 2010*. Analytics Press, Oakland, CA (2011). http://www.analyticspress.com/datacenters.html.
2. U.S. Environmental Protection Agency, Energy Star Program. *Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431*. Prepared by Lawrence Berkeley National Laboratory. LBNL-363E. August 2007.
3. Alexandra Sullivan. "Energy Star for Data Centers." *The Green Grid Technical Forum*. 4 February 2010. San Jose, CA.
4. Voltech PM6000 Documentation. http://www.voltech.com/products/poweranalyzers/PM6000.aspx.
5. LeCroy Waverunner Documentation. http://www.lecroy.com/files/pdf/LeCroy_WaveRunner_6000A_Datasheet.pdf.
6. IEC 62040-3 Edition 2.0, *Uninterruptible Power Systems (UPS), Part 3: Method of Specifying the Performance and Test Requirements.*
7. *Latest Opportunities for Efficiencies in Electronics Power Conversion Devices.* EPRI, Palo Alto, CA: 2011. 1021975.

A
UNINTERRUPTIBLE POWER SUPPLY TESTING PROCEDURE

This appendix is adapted from the "Double-Conversion On-Line and Delta-Conversion On-Line Uninterruptible Power Supply Efficiency Testing Plan," appendix of the Electric Power Research Institute (EPRI) report *Latest Opportunities for Efficiencies in Electronics Power Conversion Devices* (1021975) [7].

Purpose

Uninterruptible power supplies (UPSs) are used worldwide in residential, commercial, and industrial applications to protect critical equipment from voltage sags and interruptions. This test plan is intended to measure the efficiency of a double-conversion online UPS. It prescribes a uniform set of testing methods to determine the following:

- The efficiency of a UPS while it operates in its normal and eco modes
- The time needed for a UPS in eco mode to transfer its load to energy storage after a voltage interruption or voltage sag
- The transient generated by a UPS's static bypass switch (see Figure A-1) in eco mode as it transfers the load to battery power

This testing procedure provides guidelines to conduct a witness test in a manufacturer's facility.

Scope

The test procedure includes the following sections:

- Purpose
- Definitions
- Requirements
- Data product
- UPS verifications and test site survey
- Test setup
- Test procedure
- References

Definitions

Double-conversion online UPS. The double-conversion UPS powers the load through a rectifier and an inverter (See Figure A-1). In this normal mode of operation, the load is continuously supplied by the rectifier–inverter combination. When the main ac supply voltage is interrupted, the UPS powers the load through the battery–inverter combination.

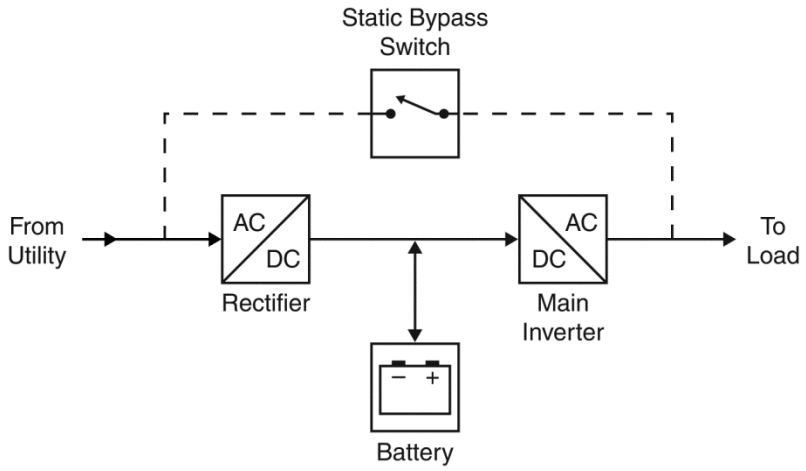


Figure A-1
Block diagram of a double-conversion online uninterruptible power supply

Eco mode. *Eco mode* refers to a high-efficiency mode in which the UPS load is powered through the mains through the static bypass switch as opposed to the double-conversion process.

Static bypass switch. This switch is used to transfer the load between ac supply voltage and battery–inverter voltage. When this switch is closed, the UPS will operate in economy mode and power the load directly with the ac supply voltage.

Switching or transfer time. The time necessary for the UPS to switch from the line to battery power. For the purpose of this test plan, it is the time (in cycles) to take the UPS out of eco mode and back into delta/double-conversion after a supply voltage interruption.

Switching Transient: Any momentary rise in the UPS output voltage or current as a result of switching from eco mode to delta/double-conversion after a supply voltage interruption.

Uninterruptible power supply (UPS). A device used to provide ac power to a load (servers, computers, and so on) in the event that the supply voltage is interrupted (voltage sag, momentary interruption, complete power loss). A UPS is composed of a rectifier and inverter; when power to the rectifier is interrupted, batteries supply power to the inverter, which provides ac power to its loads.

Requirements

Uninterruptible Power Supply and Load Specifications

The following are required for witness testing:

- Extensive documentation of the UPS to be tested, which includes but is not limited to the input and output voltages, operating mode, and UPS type
- Extensive documentation of the environmental conditions in which the UPS will be tested
- Documentation of all available specifications of the load that the UPS is powering

Table A-1 includes UPS parameters that should be recorded to provide a complete profile of the device.

Table A-1
Uninterruptible power supply specifications to be recorded

UPS Specifications	
Family name	
Model number	
Power rating	
Rating, kVA	
Rating, kW	
Input ac specifications	
Single or three phase	
Power factor	
Input frequency, Hz	
Frequency range, Hz	
Input voltage	
General specifications	
UPS technology	
Battery specifications	
Battery test type	
Quantity (batteries or cells or jars)	
Recharge time	
Battery technology	
Output ac specifications	
Output voltage	
Output waveform	
Nominal current, amps	
User interface	
Display type	
Communications	
Communications options	
Software and cable	

Table A-1 (Continued)
Uninterruptible power supply specifications to be recorded

UPS Specifications	
Physical data	
Form factor	
Unit height, in. (mm)	
Unit width, in. (mm)	
Unit depth, in. (mm)	
Unit weight, lb (kg)	
Shipping height, in. (mm)	
Shipping width, in. (mm)	
Shipping depth, in. (mm)	
Shipping weight, lb (kg)	
Environmental	
Operating temperature, minimum °F (°C)	
Operating temperature, maximum °F (°C)	
Storage temperature, minimum °F (°C)	
Storage temperature, maximum °F (°C)	
Relative humidity	
Operating elevation, ft (m)	
Storage elevation, ft (m)	
Sound emission/audible noise	
Heat dissipation	
Cooling	
Options	
Optional items	
Agency/Certification/Conformance	
Agency approval	

Use Table A-2 to record all available details of the load that will be used during the UPS efficiency testing.

Table A-2
Reference test load specifications

Value	Parameter
	Manufacturer
	Model
	Power rating
	Load type (resistive, capacitive, inductive, electronic)
	Voltage range
	Current range
	Peak current
	Frequency range
	Crest factor range

Test Equipment Specifications

If the manufacturing facility has a test setup already on site, it must be documented to ensure that is fulfills certain requirements. Table A-3 lists requirements pertaining to test equipment.

Table A-3
Test equipment requirements

Equipment	Requirements
Power meter	• Simultaneously measures three phases of the following: – UPS RMS input voltage, current, and power – UPS RMS output voltage, current, and power • < 0.5% uncertainty • Current calibration • Must have the ability to store data or dump data to a storage medium (such as a flash drive or universal serial bus hard disk drive)
Current transformers	• Used with the power meter to measure RMS current at the input and output of the UPS • < 0.5% uncertainty • Current calibration
Oscilloscope	• Uses two channels to monitor RMS voltage and RMS current simultaneously • Current calibration • Must have the ability to save waveforms captured, preferably in a .csv file but a .jpeg file is acceptable
Thermal camera	• Used to measure temperatures of UPS components to verify steady-state operation • Current calibration • < 2% degree uncertainty • Must have a software package that allows for analysis of all photos taken

After ensuring that the equipment fulfills the requirements listed in Table A-3, use Table A-4 to document all test equipment to be used to monitor and record the RMS input and output parameters of the UPS to be tested.

Table A-4
Uninterruptible power supply status worksheet

Value	Parameter
	Input voltage, phase-to-phase for all three phases
	Input current for all three phases
	Inverter voltage, phase-to-phase for all three phases
	Output current for all three phases
	Inverter frequency
	Battery voltage
	Battery current
	Load (kVA)
	Total operating hours
	Alarms or faults (document all current alarms and/or faults)

The block diagram in Figure A-2 illustrates a sample test setup using a Voltech PM6000 meter [4] to simultaneously measure input and output parameters and calculate UPS efficiency.

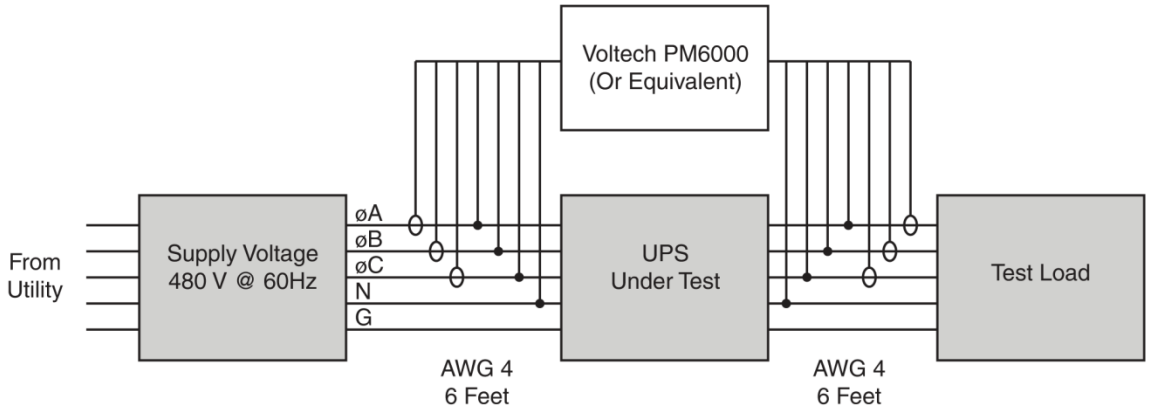


Figure A-2
Sample test setup to measure uninterruptible power supply efficiency

Data Product

The data product of this test plan includes measured efficiencies of UPSs in the range of 500 VA to 1 MW while operating in double-conversion mode. When applicable, two additional measurements with respect to eco mode should be made: 1) the time used to switch from eco mode to battery power (should be less than one cycle) and 2) the switching transient generated by the UPS operating in eco mode as it switches to battery power

Uninterruptible Power Supply Verification and Test Site Survey

The following procedure will be used to ensure that all UPSs to be tested are operating as specified by the original equipment manufacturers (OEMs). In accordance with IEC 62040-3 [6], one should observe and record the environmental conditions listed in Table A-1, when applicable.

Uninterruptible Power Supply Status Verification

The status of the UPS will be verified to ensure proper operations. To help ensure that the UPS is operating properly, EPRI or facilities personnel will determine the UPS status by documenting the parameters identified in Table A-4.

Test Site Survey

In an effort to ensure the greatest degree of repeatability, accuracy, and comparability of captured data between sites, many variables must be considered. Not all these variables will apply to every site, but the investigator should take care to identify and record all conditions that apply. Complete Table A-5 when conducting a site survey.

Table A-5
Environmental site survey

Value	Parameter
	Relative humidity
	Altitude (m)
	Ambient temperature (˚C)
	Airflow across UPS under test
	Pollution degree (see note 1)
	Exposure to vibration
	Damaging fumes
	Moisture
	Dust
	Abrasive dust
	Steam

Table A-5 (Continued)
Environmental site survey

Value	Parameter
	Explosive mixtures of dust or gases
	Salt air
	Weather or dripping water
	Extreme changes in temperature
	Cooling water containing acid or impurities, which may cause scale, sludge, electrolysis, or corrosion of the converter parts that are exposed to the water
	Strong electromagnetic fields
	Radioactive levels above those of the natural background
	Fungus, insects, vermin, and so on
	Restriction of ventilation
	Radiated or conducted heat from other sources
	Battery service conditions

Note 1: Pollution degree is a parameter of IEC 60664-1, from which the following is derived:

- Pollution degree 1 applies where there is no pollution or only dry, nonconductive pollution.
- Pollution degree 2 applies where there is only nonconductive pollution that might temporarily become conductive due to occasional condensation.
- Pollution degree 3 applies where a local environment within the equipment is subject to conductive pollution or to dry, nonconductive pollution that could become conductive due to expected condensation.

Test Setup

The goals of the test are 1) to gather reliable efficiency data of UPSs of various capacities and technologies operating at 100%, 75%, 50%, 25%, or 0% of the rated maximum load in double-conversion and eco modes and 2) to create a data set that is repeatable and comparable between devices.

For the purposes of measuring efficiency, this procedure recommends the Voltech PM6000 meter or a similar device to monitor and record input and output values simultaneously. As shown in Figure A-1, the instrument should be set up so that it will be able to simultaneously monitor the values listed in Table A-6.

Table A-6
Uninterruptible power supply parameters to be measured during testing

Value	Parameter
	Input voltage: phases a, b, c (vrms)
	Input current: phases a, b, c (arms)
	Input power (watts)
	Input apparent power (volt-amps)
	Input reactive power (var)
	Input power factor
	Input crest factor
	Input Vthd
	Input Ithd
	Output voltage: phases a, b, c (vrms)
	Output current: phases a, b, c (arms)
	Output power (watts)
	Output apparent power (volt-amps)
	Output reactive power (var)
	Output power factor
	Output crest factor
	Output Vthd
	Output Ithd

In accordance with IEC 62040-3, the combination of instruments and probes used for the measurement of UPS efficiency shall provide true RMS measurements of the active input and output power, with an uncertainty at full rated load of less than or equal to 0.5% at the 95% confidence level, notwithstanding that voltage and current waveforms can include harmonic components.

To measure the switching transient generated by the UPS during transfer to battery power in eco mode, this procedure recommends an oscilloscope capable capturing a waveform in response to a trigger (such as the LeCroy Waverunner [5]). The oscilloscope must also be capable of measuring output voltage and current simultaneously, which requires at least two channels. This same setup can be used to then measure the time required to switch to battery power.

Functionality Tests

Before officially starting the tests, measurements will be taken while the UPS is placed in the double-conversion mode. This step accomplishes the following:

- Ensures that the target input voltage and intended load percentages are achievable
- Establishes baselines for each test
- Allows for documentation of the setup parameters

Requirements for the baseline, steady-state measurements and setup parameters are the following:

- **Environmental conditions.** The ambient temperature shall be between 20°C and 30°C.
- **Operational and electrical conditions.** The following requirements apply for each measurement:
  - The UPS must operate in double-conversion mode.
  - The UPS and the load shall have been operated for a sufficient length of time to reach a steady-state condition, which is the length of time determined during temperature rise tests plus 25%. Temperature can be measured using an infrared camera while looking at the inverter.
  - Each load condition (100%, 75%, 50%, 25%, and 0%) shall be within the range of 95% to 105% of the prescribed load.
  - All UPS subsystems intended to be operational in double-conversion mode shall be activated.
  - The ac voltage input to the UPS shall be at 97% to 103% of the rated voltage and 99% to 101% of the rated frequency.
- **Initial measurements.** The reference test load (100% of the UPS rated load) shall be applied to the output of the UPS. Confirm that the input/output voltages are at the values specified by the manufacturer. Repeat this step for 75%, 50%, 25%, and 0% load.

Test Procedure

These tests will accomplish the following:

1. Measure the efficiency of a UPS device under 100%, 75%, 50%, 25%, and 0% of its rated load.
2. Record the transient generated by the UPS operating in eco mode as it switches to battery power after the main voltage is interrupted.
3. Using the waveform captured in step 2, determine the switching time (the time, in cycles, required to take the UPS out of eco mode and back into delta/double-conversion after a supply voltage interruption).

Test Procedure for Measuring Efficiency in Double-Conversion Mode

The test procedure consists of the following steps:

1. Set up the meter so that it can accurately monitor all parameters in Table A-6.
2. Set the reference test load to supply a 100% load to the UPS.
3. Apply power to the UPS device in double-conversion mode and allow for sufficient stabilization (as described in IEC 62040-3).
4. Measure and record the values of the parameters listed in Table A-6.
5. Repeat step 4 no less than three times and no more than 15 minutes apart.
6. To calculate efficiency, divide output power over input power for each reading and record the average of the multiple measurements as the efficiency value.
7. Repeat steps 3 through 6 for 75%, 50%, 25%, and 0% load.

Test Procedure for Measuring Efficiency in Eco Mode

Place the UPS in eco mode and perform the following steps:

1. Set up the meter so that it can accurately monitor all parameters in Table A-6.
2. Set the reference test load to supply a 100% load to the UPS.
3. Apply power to the UPS device in double-conversion mode and allow for sufficient stabilization (as described in IEC 62040-3).
4. Measure and record the values of the parameters listed in Table A-6.
5. Repeat step 4 no less than three times no more than 15 minutes apart.
6. To calculate efficiency, divide output power over input power for each reading and record the average of the multiple measurements as the efficiency value.
7. Repeat steps 3 through 6 for 75%, 50%, 25%, and 0% load.

Test Procedure for Capturing Switching Transients and Calculating Transfer Time

Perform the following steps:

1. Use the oscilloscope and the appropriate probes to monitor the UPS output voltage and current.
2. Set the reference test load to supply a 100% load to the UPS.
3. Apply nominal power to the UPS in eco mode and allow for sufficient stabilization.
4. Set the scope so that it will capture the transient generated by the switch and the time needed to provide battery power to the load.
5. Repeat steps 3 and 4 for 75%, 50%, 25%, and 0% load.

Export Control Restrictions

Access to and use of EPRI Intellectual Property is granted with the specific understanding and requirement that responsibility for ensuring full compliance with all applicable U.S. and foreign export laws and regulations is being undertaken by you and your company. This includes an obligation to ensure that any individual receiving access hereunder who is not a U.S. citizen or permanent U.S. resident is permitted access under applicable U.S. and foreign export laws and regulations. In the event you are uncertain whether you or your company may lawfully obtain access to this EPRI Intellectual Property, you acknowledge that it is your obligation to consult with your company's legal counsel to determine whether this access is lawful. Although EPRI may make available on a case-by-case basis an informal assessment of the applicable U.S. export classification for specific EPRI Intellectual Property, you and your company acknowledge that this assessment is solely for informational purposes and not for reliance purposes. You and your company acknowledge that it is still the obligation of you and your company to make your own assessment of the applicable U.S. export classification and ensure compliance accordingly. You and your company understand and acknowledge your obligations to make a prompt report to EPRI and the appropriate authorities regarding any access to or use of EPRI Intellectual Property hereunder that may be in violation of applicable U.S. or foreign export laws or regulations.

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI's members represent approximately 90 percent of the electricity generated and delivered in the United States, and international participation extends to more than 30 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; and Lenox, Mass.

Together…Shaping the Future of Electricity

© 2013 Electric Power Research Institute (EPRI), Inc. All rights reserved.
Electric Power Research Institute, EPRI, and TOGETHER…SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

3002001360

Electric Power Research Institute
3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 • USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com