

Seven Best Practices for Increasing Efficiency, Availability and Capacity: The Enterprise Data Center Design Guide



Table of Contents

Introduction	3
Seven Best Practices	4
Maximize the return temperature at the cooling units to improve capacity and efficiency	4
Sidebar: The Fan-Free Data Center	8
Match cooling capacity and airflow with IT loads	8
Utilize cooling designs that reduce energy consumption	9
Sidebar: Determining Economizer Benefits Based on Geography	11
Select a power system to optimize your availability and efficiency needs	12
Sidebar: Additional UPS Redundancy Configurations	17
Design for flexibility using scalable architectures that minimizes footprint	18
Increase visibility, control and efficiency with data center infrastructure management	19
Utilize local design and service expertise to extend equipment life, reduce costs and address your data center's unique challenges	21
Conclusion	22
Enterprise Design Checklist	23

Introduction

The data center is one of the most dynamic and critical operations in any business. Complexity and criticality have only increased in recent years as data centers experienced steady growth in capacity and density, straining resources and increasing the consequences of poor performance. The 2011 National Study on Data Center Downtime revealed that the mean cost for any type of data center outage is \$505,502 with the average cost of a partial data center shutdown being \$258,149. A full shutdown costs more than \$680,000.

Because the cost of downtime is so high, availability of IT capacity is generally the most important metric on which data centers are evaluated. However, data centers today must also operate efficiently—in terms of both energy and management resources—and be flexible enough to quickly and cost-effectively adapt to changes in business strategy and computing demand. The challenges data center managers face are reflected in the key issues identified each year by the Data Center Users’ Group (Figure 1).

Efficiency first emerged as an issue for data center management around 2005 as server proliferation caused data center energy

consumption to skyrocket architectures at the same time electricity prices and environmental awareness were rising. The industry responded to the increased costs and environmental impact of data center energy consumption with a new focus on efficiency; however, there was no consensus as to how to address the problem. A number of vendors offered solutions, but none took a holistic approach, and some achieved efficiency gains at the expense of data center and IT equipment availability—a compromise few businesses could afford to make.

In the traditional data center, approximately one-half of the energy consumed goes to support IT equipment with the other half used by support systems (Figure 2). Emerson Network Power conducted a systematic analysis of data center energy use and the various approaches to reducing it to determine which were most effective.

While many organizations initially focused on specific systems within the data center, Emerson took a more strategic approach. The company documented the “cascade effect” that occurs as efficiency improvements at the server component level are amplified through reduced demand on support systems. Using this analysis, a ten-step approach for increasing data center efficiency, called

<i>Spring 2008</i>	<i>Spring 2009</i>	<i>Spring 2010</i>	<i>Spring 2011</i>
Heat Density	Heat Density	Adequate Monitoring	Availability
Power Density	Efficiency	Heat Density	Adequate Monitoring
Availability	Adequate Monitoring	Availability	Heat Density
Adequate Monitoring	Availability	Efficiency	Efficiency
Efficiency	Power Density	Power Density	Power Density

Figure 1. Top data center issues as reported by the Data Center Users’ Group.

Energy Logic, was developed. According to the analysis detailed in the Emerson Network Power white paper, *Energy Logic: Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems*, 1 W of savings at the server component level can create 2.84 W of savings at the facility level.

This paper builds upon several of the steps in Energy Logic to define seven best practices that serve as the foundation for data center design. These best practices provide planners and operators with a roadmap for optimizing the efficiency, availability and capacity of new and existing facilities.

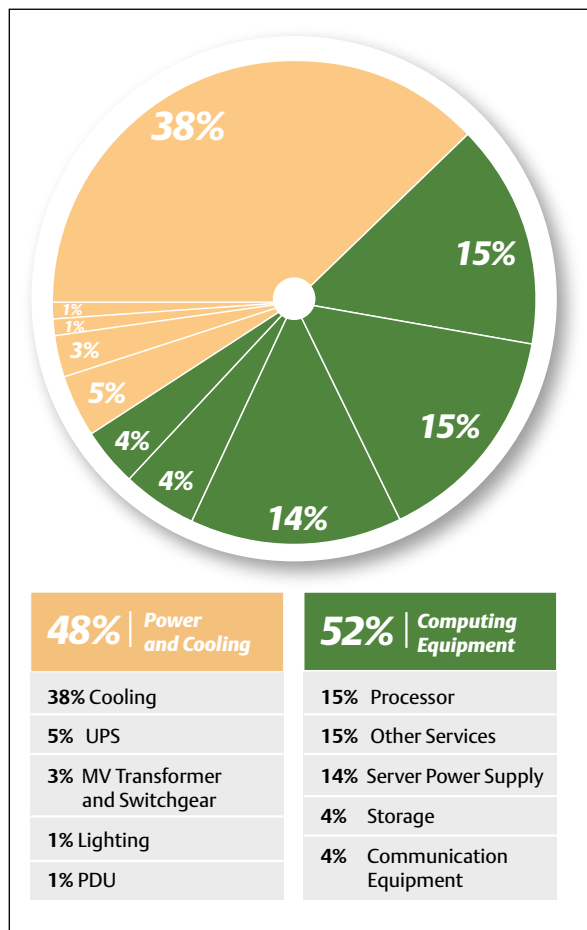


Figure 2. IT equipment accounts for over 50 percent of the energy used in a traditional data center while power and cooling account for an additional 48 percent.

Seven Best Practices for Enterprise Data Center Design

The following best practices represent proven approaches to employing cooling, power and management technologies in the quest to improve overall data center performance.

Best Practice 1: Maximize the return temperature at the cooling units to improve capacity and efficiency

Maintaining appropriate conditions in the data center requires effectively managing the air conditioning loop comprising supply and return air. The laws of thermodynamics create opportunities for computer room air conditioning systems to operate more efficiently by raising the temperature of the return air entering the cooling coils.

This best practice is based on the hot-aisle/cold-aisle rack arrangement (Figure 3), which improves cooling unit performance by reducing mixing of hot and cold air, thus enabling higher return air temperatures. The relationship between return air temperature and sensible cooling capacity is illustrated in Figure 4. It shows that a 10 degree F increase in return air temperature typically results in a 30 to 38 percent increase in cooling unit capacity.

The racks themselves provide something of a barrier between the two aisles when blanking panels are used systematically to close openings. However, even with blanking panels, hot air can leak over the top and around the sides of the row and mix with the air in the cold aisle. This becomes more of an issue as rack density increases.

To mitigate the possibility of air mixing as it returns to the cooling unit, perimeter cooling units can be placed at the end of the hot aisle as shown in Figure 3. If the cooling units cannot be positioned at the end of the hot aisle, a drop ceiling can be used as a plenum to




Figure 3. In the hot-aisle/cold-aisle arrangement, racks are placed in rows face-to-face, with a recommended 48-inch aisle between them. Cold air is distributed in the aisle and used by racks on both sides. Hot air is expelled at the rear of each rack into the “hot aisle.”

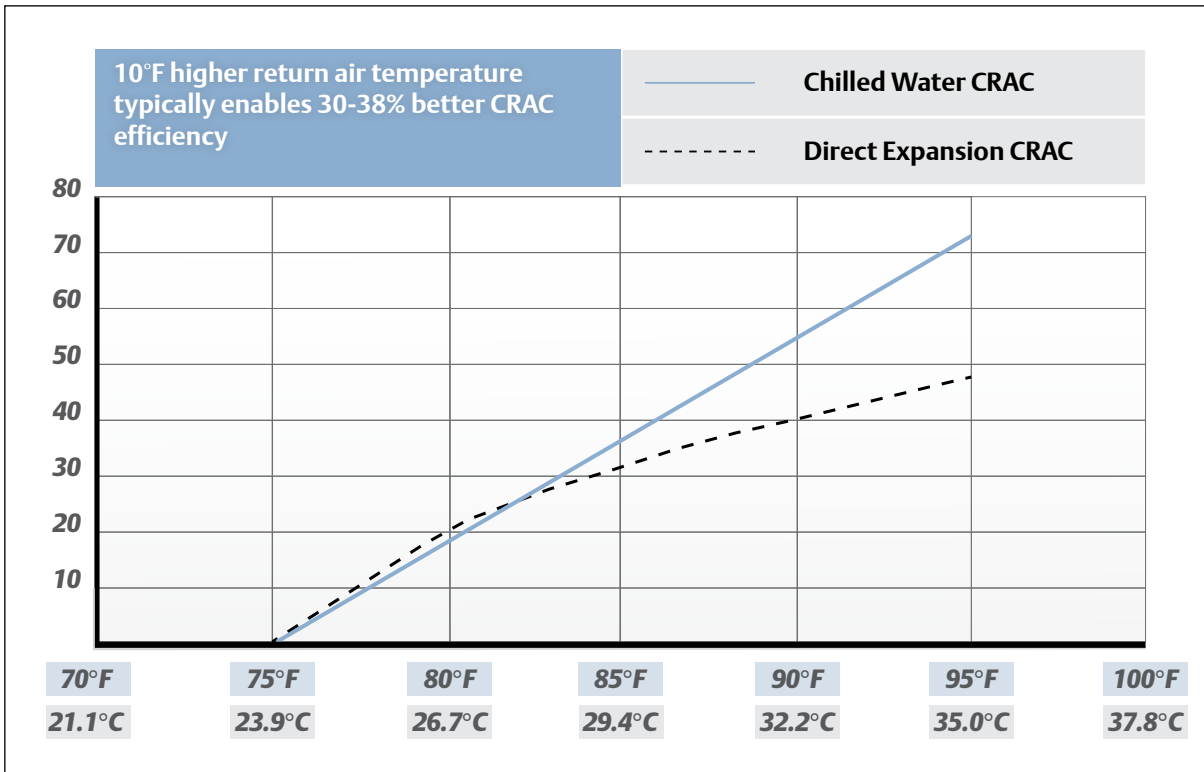


Figure 4. Cooling units operate at higher capacity and efficiency at higher return air temperatures.

prevent hot air from mixing with cold air as it returns to the cooling unit. Cooling units can also be placed in a gallery or mechanical room.

In addition to ducting and plenums, air mixing can also be prevented by applying containment and by moving cooling closer to the source of heat.

Optimizing the Aisle with Containment and Row-Based Cooling

Containment involves capping the ends of the aisle, the top of the aisle, or both to isolate the air in the aisle (Figure 5).

Cold aisle containment is favored over hot aisle containment because it is simpler to deploy and reduces risk during the event of a breach of the containment system. With hot aisle containment, open doors or missing blanking panels allow hot air to enter the cold aisle, jeopardizing the performance of IT equipment (Figure 6). In a similar scenario with the cold aisle contained, cold air leaking into the hot aisle decreases the temperature of the return air, slightly compromising efficiency, but not threatening IT reliability.




Figure 5. The hot-aisle/cold-aisle arrangement creates the opportunity to further increase cooling unit capacity by containing the cold aisle.

Row-based cooling units can operate within the contained environment to supplement or replace perimeter cooling. This brings temperature and humidity control closer to the source of heat, allowing more precise control and reducing the energy required to move air across the room. By placing the return air intakes of the precision cooling units directly in the hot aisle, air is captured at its highest temperature and cooling efficiency

Hot Aisle Containment		Cold Aisle Containment
Efficiency	<ul style="list-style-type: none">Highest “potential” efficiency	<ul style="list-style-type: none">Good improvement in efficiency
Installation	<ul style="list-style-type: none">Requires ducting or additional row cooling unitsAdding new servers puts installers in a very hot space	<ul style="list-style-type: none">Easy to add to existing data centersCan work with existing fire suppression
Reliability	<ul style="list-style-type: none">Hot air can leak into the cold aisle and into the serverNeed to add redundant units	<ul style="list-style-type: none">Leaking cold air into the hot aisle only lowers efficiencyRedundancy from other floor mount units

Figure 6. Cold aisle containment is recommended based on simpler installation and improved reliability in the event of a breach of the containment system.

is maximized. The possible downside of this approach is that more floor space is consumed in the aisle. Row-based cooling can be used in conjunction with traditional perimeter-based cooling in higher density “zones” throughout the data center.

Supplemental Capacity through Sensible Cooling

For optimum efficiency and flexibility, a cooling system architecture that supports delivery of refrigerant cooling to the rack can work in either a contained or uncontained environment. This approach allows cooling modules to be positioned at the top, on the side or at the rear of the rack, providing focused cooling precisely where it is needed while keeping return air temperatures high to optimize efficiency.

The cooling modules remove air directly from the hot aisle, minimizing both the distance the air must travel and its chances to mix with cold air (Figure 7). Rear-door cooling modules

can also be employed to neutralize the heat before it enters the aisle. They achieve even greater efficiency by using the server fans for air movement, eliminating the need for fans on the cooling unit. Rear door heat exchanger solutions are not dependent on the hot-aisle/cold-aisle rack arrangement.

Properly designed supplemental cooling has been shown to reduce cooling energy costs by 35-50 percent compared to perimeter cooling only. In addition, the same refrigerant distribution system used by these solutions can be adapted to support cooling modules mounted directly on the servers, eliminating both cooling unit fans and server fans (see the sidebar, The Fan-Free Data Center).




Figure 7. Refrigerant-based cooling modules mounted above or alongside the rack increase efficiency and allow cooling capacity to be matched to IT load.

The Fan-Free Data Center

Moving cooling units closer to the source of heat increases efficiency by reducing the amount of energy required to move air from where heat is being produced to where it is being removed.

But what if the cooling system could eliminate the need to move air at all?

This is now possible using a new generation of cooling technologies that bring data center cooling inside the rack to remove heat directly from the device producing it.

The first of these systems commercially available, the Liebert XDS, works by equipping servers with cooling plates connected to a centralized refrigerant pumping unit. Heat from the server is transferred through heat risers to the server housing and then through a thermal lining to the cooling plate. The cooling plate uses refrigerant-filled microchannel tubing to absorb the heat, eliminating the need to expel air from the rack and into the data center.

In tests by Lawrence Berkeley National Labs, this approach was found to improve energy efficiency by 14 percent compared to the next best high-efficiency cooling solution. Significant additional savings are realized through reduced server energy consumption resulting from the elimination of server fans. This can actually create a net positive effect on data center energy consumption: *The cooling system decreases data center energy consumption compared with running servers with no cooling.* And, with no fans on the cooling units or the servers, the data center becomes as quiet as a library.

Best Practice 2: Match cooling capacity and airflow with IT loads

The most efficient cooling system is one that matches needs to requirements. This has proven to be a challenge in the data center because cooling units are sized for peak demand, which rarely occurs in most applications. This challenge is addressed through the use of intelligent cooling controls capable of understanding, predicting and adjusting cooling capacity and airflow based on conditions within the data center. In some cases, these controls work with the technologies in Best Practice 3 to adapt cooling unit performance based on current conditions (Figure 8).

Intelligent controls enable a shift from cooling control based on return air temperature, to control based on conditions at the servers, which is essential to optimizing efficiency.

This often allows temperatures in the cold aisle to be raised closer to the safe operating threshold now recommended by ASHRAE (max 80.5 degrees F). According to an



Figure 8. Intelligent controls like the Liebert iCOM system can manage air flow and cooling capacity independently.

Emerson Network Power study, a 10 degree increase in cold aisle temperature can generate a 20 percent reduction in cooling system energy usage.

The control system also contributes to efficiency by allowing multiple cooling units to work together as a single system utilizing teamwork. The control system can shift workload to units operating at peak efficiency while preventing units in different locations from working at cross-purposes. Without this type of system, a unit in one area of the data center may add humidity to the room at the same time another unit is extracting humidity from the room. The control system provides visibility into conditions across the room and the intelligence to determine whether humidification, dehumidification or no action is required to maintain conditions in the room at target levels and match airflow to the load.

For supplemental cooling modules that focus cooling on one or two racks, the control system performs a similar function by shedding fans based on the supply and return air temperatures, further improving the efficiency of supplemental cooling modules.

Best Practice 3: Utilize cooling designs that reduce energy consumption

The third step in optimizing the cooling infrastructure is to take advantage of newer technologies that use less energy than previous generation components.

Increasing Fan Efficiency

The fans that move air and pressurize the raised floor are a significant component of cooling system energy use. On chilled water cooling units, fans are the largest consumer of energy.

Fixed speed fans have traditionally been used in precision cooling units. Variable frequency drives represent a significant improvement over fixed-speed as they enable fan speed to be adjusted based on operating conditions. Adding variable frequency drives to the fan motor of a chilled-water precision cooling unit allows the fan's speed and power draw to be reduced as load decreases, resulting in a dramatic impact on fan energy consumption. A 20 percent reduction in fan speed provides almost 50 percent savings in fan power consumption.

Electronically commutated (EC) fans may provide an even better option for increasing cooling unit efficiency. EC plug fans are inherently more efficient than traditional centrifugal fans because they eliminate belt losses, which total approximately five percent. The EC fan typically requires a minimum 24-inch raised floor to obtain maximum operating efficiency and may not be suitable for ducted upflow cooling units where higher static pressures are required. In these cases, variable frequency drive fans are a better choice.

In independent testing of the energy consumption of EC fans compared to variable drive fans, EC fans mounted inside the cooling unit created an 18 percent savings. When EC

fans were mounted outside the unit, below the raised floor, savings increased to 30 percent.

Both options save energy, can be installed on existing cooling units or specified in new units, and work with the intelligent controls described in Best Practice 2 to match cooling capacity to IT load.

Enhancing Heat Transfer

The heat transfer process within the cooling unit also consumes energy. New microchannel coils used in condensers have proven to be more efficient at transferring heat than previous generation coil designs. They can reduce the amount of fan power required for heat transfer, creating efficiency gains of five to eight percent for the entire system. As new cooling units are specified, verify they are taking advantage of the latest advances in coil design.

Incorporating Economizers

Economizer systems use outside air to provide “free-cooling” cycles for data centers. This reduces or eliminates chiller operation or compressor operation in precision cooling units, enabling economizer systems to generate cooling unit energy savings of 30 to 50 percent, depending on the average temperature and humidity conditions of the site.

A fluid-side economizer (often called water-side) works in conjunction with a heat rejection loop comprising an evaporative cooling tower or drycooler to satisfy cooling requirements. It uses outside air to aid heat rejection, but does not introduce outside air into the data center. An air-side economizer uses a system of sensors, ducts and dampers to bring outside air into the controlled environment.




Figure 9: Liebert CW precision cooling units utilize outside air to improve cooling efficiency when conditions are optimal.

The affect of outside air on data center humidity should be carefully considered when evaluating economization options. The recommended relative humidity for a data center environment is 40 to 55 percent. Introducing outside air via an air-side economizer system (Figure 9) during cold winter months can lower humidity to unacceptable levels, causing equipment-damaging electrostatic discharge. A humidifier can be used to maintain appropriate humidity levels, but that offsets some of the energy savings provided by the economizer.

Fluid-side economizer systems eliminate this problem by using the cold outside air to cool the water/glycol loop, which in turn provides fluid cold enough for the cooling coils in the air conditioning system. This keeps the outside air out of the controlled environment and eliminates the need to condition that air. For that reason, fluid-side economizers are preferred for data center environments (see Figure 10).

	<i>Air-Side Economization</i>	<i>Water-Side Economization</i>
PROS	<ul style="list-style-type: none"> • Very efficient in some climates 	<ul style="list-style-type: none"> • Can be used in any climate • Can retrofit to current sites
CONS	<ul style="list-style-type: none"> • Limited to moderate climates • Complexity during change-over • Humidity control can be a challenge; vapor barrier is compromised • Dust, pollen and gaseous contamination sensors are required • Hard to implement in “high density” applications 	<ul style="list-style-type: none"> • Maintenance complexity • Complexity during change-over • Piping and control more complex • Risk of pitting coils if untreated stagnant water sits in econo-coils

Figure 10. Air economizers are efficient but are more limited in their application and may require additional equipment to control humidity and contamination.

Determining Economizer Benefits Based on Geography

Economizers obviously deliver greater savings in areas where temperatures are lower. Yet, designed properly, they can deliver significant savings in warmer climates as well.

Plotting weather data versus outdoor wet-bulb temperature allows the hours of operation for a fluid-side economizer with an open cooling tower to be predicted for a given geography. If the water temperature leaving the chiller is 45 degrees F, full economization can be achieved when the ambient wet-bulb temperature is below 35 degrees F. Partial economizer operation occurs between 35 and 43 degrees F wet-bulb temperature.

In Chicago, conditions that enable full economization occur 27 percent of the year, with an additional 16 percent of the year supporting partial economization. If the water temperature leaving the chiller is increased to 55 degrees F, full economization can occur 43 percent of the year with partial operation occurring an additional 21 percent of the year.

This saves nearly 50 percent in cooling unit energy consumption.

In Atlanta, full economizer operation is possible 11 percent of the year and partial operation 14 percent with a leaving water temperature of 45 degrees F. When the leaving water temperature is increased to 55 degrees F, full economization is available 25 percent of the year and partial economizer operation is available an additional 25 percent of the year. This creates cooling energy savings of up to 43 percent.

Even in a climate as warm as Phoenix, energy savings are possible. With a leaving water temperature of 55 degrees F, full economization is possible 15.3 percent of the time with partial economization available 52.5 percent of the time.

For a more detailed analysis of economization savings by geographic location see the Emerson Network Power white paper *Economizer Fundamentals: Smart Approaches to Energy-Efficient Free Cooling for Data Centers*.

Best Practice 4: Select a power system to optimize your availability and efficiency needs

There are many options to consider in the area of power system design that affect efficiency, availability and scalability. In most cases, availability and scalability are the primary considerations. The data center is directly dependent on the critical power system, and electrical disturbances can have disastrous consequences in the form of increased downtime. In addition, a poorly designed system can limit expansion. Relative to other infrastructure systems, the power system consumes significantly less energy, and efficiency can be enhanced through new control options.

Data center professionals have long recognized that while every data center aspires to 100 percent availability, not every business is positioned to make the investments required to achieve that goal. The Uptime Institute

defined four tiers of data center availability (which encompass the entire data center infrastructure of power and cooling) to help guide decisions in this area (Figure 11). Factors to consider related specifically to AC Power include UPS design, module-level redundancy and power distribution design.

UPS Design

There is growing interest in using transformer-free UPS modules in three-phase critical power applications. Large transformer-free UPS systems are typically constructed of smaller, modular building blocks that deliver high power in a lighter weight with a smaller footprint and higher full load efficiency. In addition, some transformer-free UPS modules offer new scalability options that allow UPS modules and UPS systems to be paralleled to enable the power system to grow in a more flexible manner with simple paralleling methods, or internally scalable or “modular” designs. Professionals who value full load

<i>Data Center Infrastructure Tier</i>	<i>Description</i>	<i>Availability Supported</i>
<i>I: Basic Data Center</i>	Single path for power and cooling distribution without redundant components. May or may not have a UPS, raised floor or generator.	<i>99.671%</i>
<i>II: Redundant Components</i>	Single path for power and cooling distribution with redundant components. Will have a raised floor, UPS and generator but the capacity design is N+1 with a single-wired distribution path throughout	<i>99.741%</i>
<i>III: Concurrently Maintainable</i>	Multiple active power and cooling distribution paths, but only path is active. Has redundant components and is concurrently maintainable. Sufficient capacity and distribution must be present to simultaneously carry the load on one path while performing maintenance on the other path.	<i>99.982%</i>
<i>IV: Fault Tolerant</i>	Provides infrastructure capacity and capability to permit any planned activity without disruption to the critical load. Infrastructure design can sustain at least one worst case, unplanned failure or event with no critical load impact.	<i>99.995%</i>

Figure 11. The Uptime Institute defines four tiers of data center infrastructure availability to help organizations determine the level of investment required to achieve desired availability levels.

Characteristic	Transformer-Free	Transformer-Based
Fault Management		+
Low Component Count		+
Robustness		+
Input / DC / Output Isolation		+
Scalability	+	
In the Room / Row	+	
Double Conversion Efficiency	Up to 96%	Up to 94%
VFD (Eco-Mode) Efficiency	Up to 99%	Up to 98%

Figure 12. Comparison of transformer-based and transformer-free UPS systems. For more on VFD mode and other UPS operating modes, see the Emerson white paper, *UPS Operating Modes – A Global Standard*.

efficiency and scalability above all other attributes may consider a power system design based on a transformer-free UPS. However, some transformer-free UPS designs utilize high component counts and extensive use of fuses and contactors, compared to traditional transformer-based UPS, which can result in lower Mean Time Between Failure (MTBF), higher service rates, and lower overall system availability.

For critical applications where maximizing availability is more important than achieving efficiency improvements in the power system, a state-of-the-art transformer-based UPS ensures the highest availability and robustness for mission critical facilities. Transformers within the UPS provide fault and galvanic isolation as well as useful options for power distribution. Transformers serve as an impedance to limit arc flash potential within the UPS itself and in some cases within the downstream AC distribution system. Transformers also help to isolate faults to prevent them from propagating throughout the electrical distribution system.

Selecting the best UPS topology for a data center is dependent on multiple factors such as country location, voltage, power quality, efficiency needs, availability demands, fault management, as well as other factors (Figure 12). A critical power infrastructure supplier who specializes in both designs is ideally suited to propose the optimal choice based on your unique needs.

UPS System Configurations

A variety of UPS system configurations are available to achieve the higher levels of availability defined in the Uptime Institute classification of data center tiers.

Tier IV data centers generally use a minimum of 2 (N + 1) systems that support a dual-bus architecture to eliminate single points of failure across the entire power distribution system (Figure 13). This approach includes two or more independent UPS systems each capable of carrying the entire load with N capacity after any single failure within the electrical infrastructure. Each system provides power to its own independent distribution network, allowing 100 percent concurrent maintenance

and bringing power system redundancy to the IT equipment as close to the input terminals as possible. This approach achieves the highest availability but may compromise UPS efficiency at low loads and is more complex to scale than other configurations.

For other critical facilities, a parallel redundant configuration, such as the $N + 1$ architecture—

in which “ N ” is the number of UPS units required to support the load and “ $+1$ ” is an additional unit for redundancy—is a good choice for balancing availability, cost and scalability (Figure 14). UPS units should be sized to limit the total number of modules in the system to reduce the risk of module failure. In statistical analysis of $N + 1$ systems, a $1 + 1$ design has the highest data center availability,

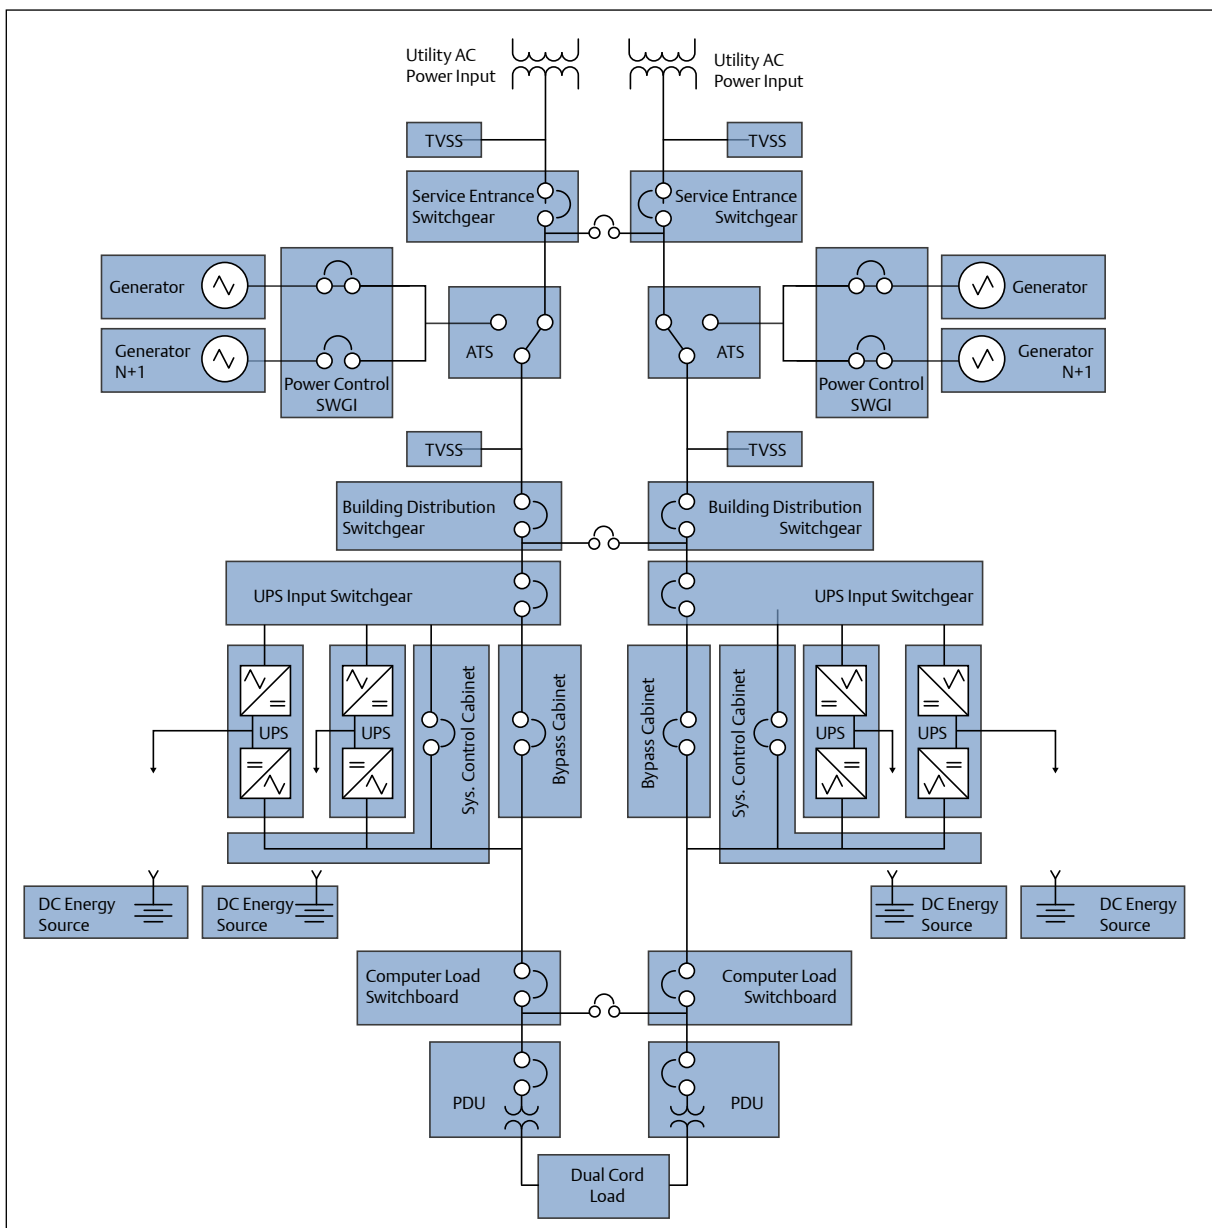


Figure 13. Typical Tier IV high-availability power system configuration with dual power path to the server.

but if there is a need for larger or more scalable data center power systems, a system with up to 4 UPS cores (3+1) has greater availability than a single unit and still provides the benefits of scalability (Figure 15). There are also other parallel configurations available outlined in the sidebar, UPS Redundancy Configurations.

UPS Efficiency Options

Today's high-availability double-conversion UPS systems can achieve efficiency levels similar to less robust designs through the use of advanced efficiency controls.

Approximately 4-6 percent of the energy passing through a double-conversion UPS is used in the conversion process. This has traditionally been accepted as a reasonable price to pay for the protection provided by the UPS system, but with new high-efficiency options the conversion process can be bypassed, and efficiency increased, when data center criticality is not as great or when utility power is known to be of the highest quality.

Here's how it works: the UPS systems incorporate an automatic static-switch bypass that operates at very high speeds to provide a break-free transfer of the load to a utility or




Figure 14. N + 1 UPS system configuration.

backup system to enable maintenance and ensure uninterrupted power in the event of severe overload or instantaneous loss of bus voltage. The transfer is accomplished in under 4ms to prevent any interruption that could shut down IT equipment. Using advanced intelligent controls, the bypass switch can be kept closed, bypassing the normal AC-DC-AC conversion process while the UPS monitors bypass power quality. When the UPS senses power quality falling outside accepted standards, the bypass opens and transfers

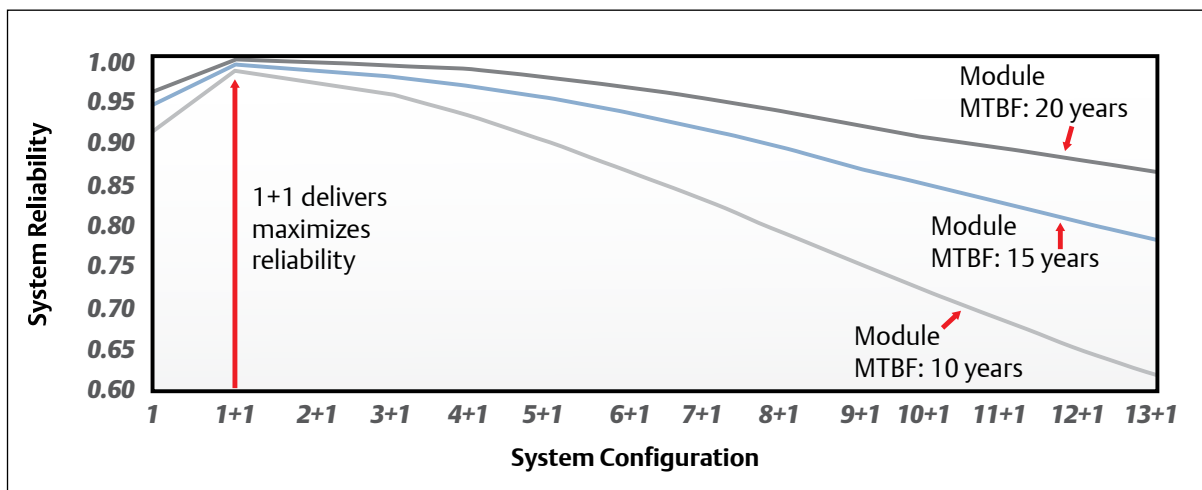


Figure 15. Increasing the number of UPS modules in an N+ 1 system increases the risk of failure.

power back to the inverter so anomalies can be corrected. To work successfully, the inverter must be kept in a constant state of preparedness to accept the load and thus needs control power. The power requirement is below 2 percent of the rated power, creating potential savings of 4-4.5 percent compared with traditional operating modes.

For more on UPS operating modes, please see the Emerson white paper titled: *UPS Operating Modes – A Global Standard*.

Another newer function enabled by UPS controls is intelligent paralleling, which improves the efficiency of redundant UPS systems by deactivating UPS modules that are not required to support the load and taking advantage of the inherent efficiency improvement available at higher loads. For example, a multi-module UPS system configured to support a 500 kVA load using three 250 kVA UPS modules can support loads below 400 kVA with only two modules while

maintaining redundancy and improving the efficiency of the UPS by enabling it to operate at a higher load. This feature is particularly useful for data centers that experience extended periods of low demand, such as a corporate data center operating at low capacity on weekends and holidays (Figure 16).

High-Voltage Distribution

There may also be opportunities to increase efficiency in the distribution system by distributing higher voltage power to IT equipment. A stepdown from 480 V to 208 V in the traditional power distribution architecture introduces minimal losses that can be eliminated using an approach that distributes power in a 4-wire Wye configuration at an elevated voltage typically 415/240V. In this case, IT equipment is powered from phase-to-neutral voltage rather than phase-to-phase. The server power supply receives 240V power, which may improve the operating efficiencies of the server power supplies in addition to the

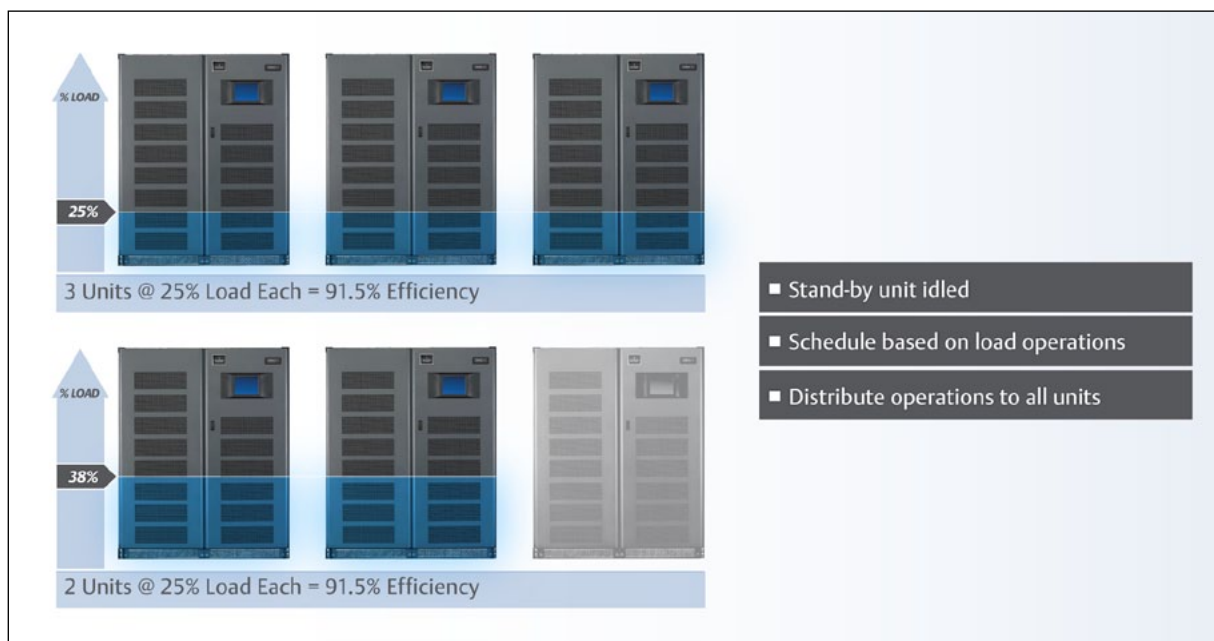


Figure 16. Firmware intelligence for intelligent paralleling increases the efficiency of a multi-module system by idling unneeded inverters or whole modules and increases capacity on demand.

lower losses in the AC distribution system. However, this higher efficiency comes at the expense of higher available fault currents and the added cost of running the neutral wire throughout the electrical system. For more information on high-voltage power

distribution, refer to the Emerson Network Power white paper *Balancing the Highest Efficiency with the Best Performance for Protecting the Critical Load*.

Additional UPS Redundancy Configurations

The N+1 architecture is the most common approach for achieving module-level redundancy. However, there are other variations that can prove effective.

A Catcher Dual Bus UPS configuration (a variation of Isolated Redundant) is a method to incrementally add dual-bus performance to existing, multiple, single-bus systems. In this configuration, the static transfer switch (STS) switches to the redundant “catcher” UPS in the event of a single UPS system failure to a single bus. This is a good option for high density environments with widely varying load requirements (Figure 17.)

The 1+N architecture, common in Europe, is also becoming more popular globally. This is sometimes referred to as a “Distributed

Static Switch” system, where each UPS has its own static switch built into the module to enable UPS redundancy (Figure 18). When comparing N+1 or 1+N designs it is important to recognize that an N+1 design uses a large static switch which must be initially sized for end-state growth requirements while 1+N systems require the switchgear be sized for end-state growth requirements. However in 1+N systems the use of distributed static switches, rather than the single large switch, reduces initial capital costs.

An experienced data center specialist should be consulted to ensure the selected configuration meets requirements for availability and scalability. There are tradeoffs between the two designs such as cost, maintenance and robustness which need to be understood and applied correctly.

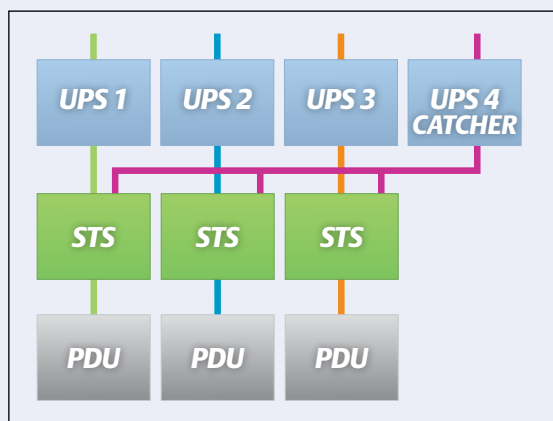


Figure 17. Catcher dual-bus UPS system configuration

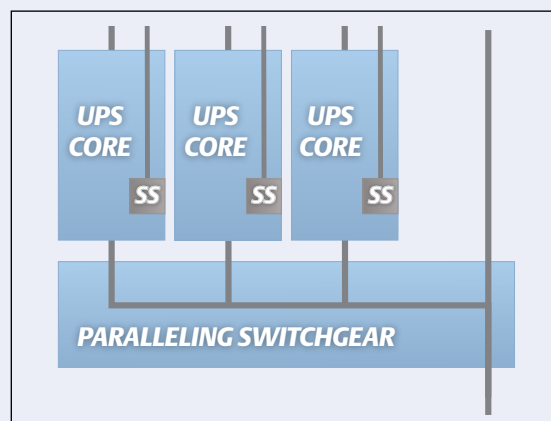


Figure 18. 1+N UPS redundancy configuration.

Best Practice 5. Design for flexibility using scalable architectures that minimizes footprint

One of the most important challenges that must be addressed in any data center design project is configuring systems to meet current requirements, while ensuring the ability to adapt to future demands. In the past, this was accomplished by oversizing infrastructure systems and letting the data center grow into its infrastructure over time. That no longer works because it is inefficient in terms of both capital and energy costs. The new generation of infrastructure systems is designed for greater scalability, enabling systems to be right-sized during the design phase without risk.

Some UPS systems now enable modularity within the UPS module itself (vertical) across modules (horizontal) and across systems (orthogonal). Building on these highly scalable designs allows a system to scale from individual 200-1200 kW modules to a multi-module system capable of supporting up to 5 MW.

The power distribution system also plays a significant role in scalability. Legacy power distribution used an approach in which the UPS fed a required number of power distribution units (PDUs), which then distributed power directly to equipment in the rack. This was adequate when the number of racks and servers was relatively low, but today, with the number of devices that must be supported, breaker space would be expended long before system capacity is reached.

Two-stage power distribution creates the scalability and flexibility required. In this approach, distribution is compartmentalized between the UPS and the server to enable greater flexibility and scalability (Figure 19). The first stage of the two-stage system provides mid-level distribution.

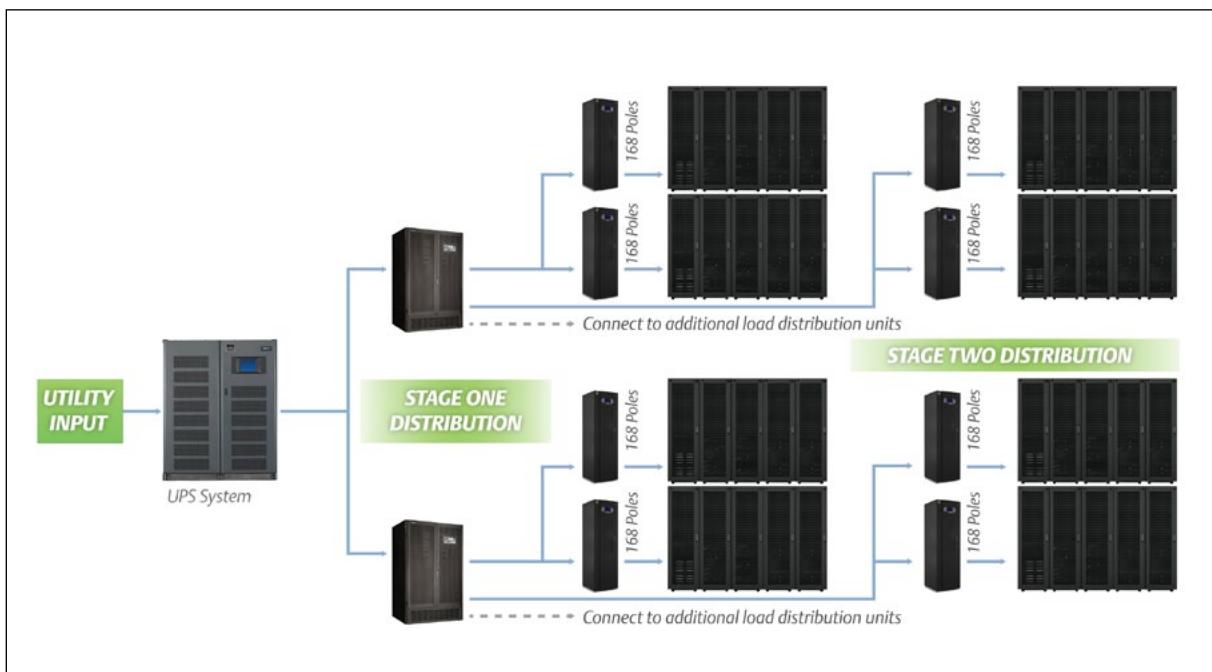


Figure 19. A two-tier power distribution system provides scalability, flexibility when adding power to racks.

The mid-level distribution unit includes most of the components that exist in a traditional PDU, but with an optimized mix of circuit and branch level distribution breakers. It typically receives 480V or 600V power from the UPS, but instead of doing direct load-level distribution, it feeds floor-mounted load-level distribution units. The floor mounted remote panels provide the flexibility to add plug-in output breakers of different ratings as needed.

Rack-level flexibility can also be considered. Racks should be able to quickly adapt to changing equipment requirements and increasing densities. Rack PDUs increase power distribution flexibility within the rack and can also enable improved control by providing continuous measurement of volts, amps and watts being delivered through each receptacle. This provides greater visibility into increased power utilization driven by virtualization and consolidation. It can also be used for charge-backs, to identify unused rack equipment drawing power, and to help quantify data center efficiency.

Alternately, a busway can be used to support distribution to the rack. The busway runs across the top of the row or below the raised floor. When run above the rack, the busway gets power distribution cabling out from under the raised floor, eliminating obstacles to cold air distribution. The busway provides the flexibility to add or modify rack layouts and change receptacle requirements without risking power system down time. While still relatively new to the data center, busway distribution has proven to be an effective option that makes it easy to reconfigure and add power for new equipment.

Best Practice 6: Enable data center infrastructure management and monitoring to improve capacity, efficiency and availability

Data center managers have sometimes been flying blind, lacking visibility into the system performance required to optimize efficiency, capacity and availability. Availability monitoring and control has historically been used by leading organizations, but managing the holistic operations of IT and facilities has lagged. This is changing as new data center management platforms emerge that bring together operating data from IT, power and cooling systems to provide unparalleled real-time visibility into operations (Figure 20).

The foundation for data center infrastructure management requires establishing an instrumentation platform to enable monitoring and control of physical assets (Figure 21). Power and cooling systems should have instrumentation integrated into them and these systems can be supplemented with additional sensors and controls to enable a centralized and comprehensive view of infrastructure systems.

At the UPS level, monitoring provides continuous visibility into system status, capacity, voltages, battery status and service events. Power monitoring should also be deployed at the branch circuit, power distribution unit and within the rack. Dedicated battery monitoring is particularly critical to preventing outages. According to Emerson Network Power's Liebert Services business, battery failure is the number one cause of UPS system dropped loads. A dedicated battery monitoring system that continuously tracks internal resistance within each battery provides the ability to predict and report batteries approaching end-of-life to enable proactive replacement prior to failure.

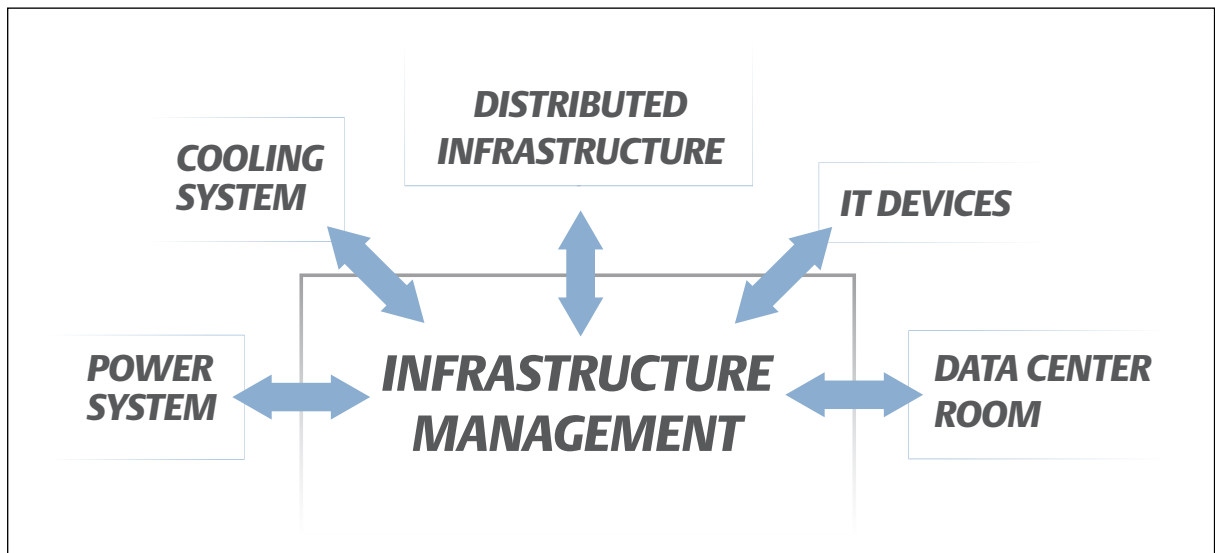


Figure 20. True Data Center Infrastructure Management (DCIM) optimizes all subsystems of data center operations holistically.

Installing a network of temperature sensors across the data center can be a valuable supplement to the supply and return air temperature data supplied by cooling units. By sensing temperatures at multiple locations, the airflow and cooling capacity can be more precisely controlled, resulting in more efficient operation.

Leak detection should also be considered as part of a comprehensive data center

monitoring program. Using strategically located sensors, these systems provide early warning of potentially disastrous leaks across the data center from glycol pipes, humidification pipes, condensate pumps and drains, overhead piping and unit and ceiling drip pans.

Communication with a management system or with other devices is provided through interfaces that deliver Ethernet connectivity




Figure 21. Examples of collection points of an instrumented infrastructure, the first step in DCIM enablement.

and SNMP and telnet communications, as well as integration with building management systems through Modbus and BACnet. When infrastructure data is consolidated into a central management platform, real-time operating data for systems across the data center can drive improvements in data center performance:

- **Improve availability:** The ability to receive immediate notification of a failure, or an event that could ultimately lead to a failure, allows faster, more effective response to system problems. Taken a step further, data from the monitoring system can be used to analyze equipment operating trends and develop more effective preventive maintenance programs. Finally, the visibility and dynamic control of data center infrastructure provided by the monitoring can help prevent failures created by changing operating conditions. For example, the ability to turn off receptacles in a rack that is maxed out on power, but may still have physical space, can prevent a circuit overload. Alternately, viewing a steady rise in server inlet temperatures may dictate the need for an additional row cooling unit before overheating brings down the servers.
- **Increase efficiency.** Monitoring power at the facility, row, rack and device level provides the ability to more efficiently load power supplies and dynamically manage cooling. Greater visibility into infrastructure efficiency can drive informed decisions around the balance between efficiency and availability. In addition, the ability to automate data collection, consolidation and analysis allows data center staff to focus on more strategic IT issues.
- **Manage capacity.** Effective demand forecasting and capacity planning has become critical to effective data center

management. Data center infrastructure monitoring can help identify and quantify patterns impacting data center capacity. With continuous visibility into system capacity and performance, organizations are better equipped to recalibrate and optimize the utilization of infrastructure systems (without stretching them to the point where reliability suffers) as well as release stranded capacity.

DCIM technologies are evolving rapidly. Next-generation systems will begin to provide a true unified view of data center operations that integrates data from IT and infrastructure systems. As this is accomplished, a true holistic data center can be achieved.

Best Practice 7: Utilize local design and service expertise to extend equipment life, reduce costs and address your data center's unique challenges

While best practices in optimizing availability, efficiency and capacity have emerged, there are significant differences in how these practices should be applied based on specific site conditions, budgets and business requirements. A data center specialist can be instrumental in helping apply best practices and technologies in the way that makes the most sense for your business and should be consulted on all new builds and major expansions/upgrades.

For established facilities, preventive maintenance has proven to increase system reliability while data center assessments can help identify vulnerabilities and inefficiencies resulting from constant change within the data center.

Emerson Network Power analyzed data from 185 million operating hours for more than 5,000 three-phase UPS units operating in the data center. The study found that the UPS Mean Time Between Failures (MTBF) for units

that received two preventive service events a year is 23 times higher than a machine with no preventive service events per year.

Preventive maintenance programs should be supplemented by periodic data center assessments. An assessment will help identify, evaluate and resolve power and cooling vulnerabilities that could adversely affect operations. A comprehensive assessment includes both thermal and electrical assessments, although each can be provided independently to address specific concerns.

Taking temperature readings at critical points is the first step in the thermal assessment and can identify hot spots and resolve problems that could result in equipment degradation. Readings will determine whether heat is successfully being removed from heat-generating equipment, including blade servers. These readings are supplemented by infrared inspections and airflow measurements. Cooling unit performance is also evaluated to ensure units are performing properly. Computational Fluid Dynamics (CFD) is also used to analyze air flow within the data center.

The electrical assessment includes a single-point-of-failure analysis to identify critical failure points in the electrical system. It also documents the switchgear capacity and the current draw from all UPS equipment and breakers, as well as the load per rack.

Conclusion

The last ten years have been tumultuous within the data center industry. Facilities are expected to deliver more computing capacity while increasing efficiency, eliminating downtime and adapting to constant change. Infrastructure technologies evolved throughout this period as they adapted to higher density equipment and the need for greater efficiency and control.

The rapid pace of change caused many data center managers to take a wait-and-see attitude to new technologies and practices. That was a wise strategy several years ago but today those technologies have matured and the need for improved data center performance can no longer be ignored. Proper deployment of the practices discussed can have immediate TCO improvements – from capital benefits, to amazing energy efficiency gains to ease of computing adaptations.

In the cooling system, traditional technologies now work with newer technologies to support higher efficiencies and capacities. Raising the return air temperature improves capacity and efficiency while intelligent controls and high-efficiency components allow airflow and cooling capacity to be matched to dynamic IT loads.

In the power system, high efficiency options work within proven system configurations to enhance efficiency while maintaining availability. Power distribution technologies provide increased flexibility to accommodate new equipment, while delivering the visibility into power consumption required to measure efficiency.

Most importantly, a new generation of infrastructure management technologies is emerging that bridges the gap between facilities and IT systems, and provides centralized control of the data center.

Working with data center design and service professionals to implement these best practices, and modify them based on changing conditions in the data center, creates the foundation for a data center in which availability, efficiency and capacity can all be optimized in ways that simply weren't possible five years ago.

Data Center Design Checklist

Use this checklist to assess your own data center based on the seven best practices outlined in this paper.

- ☐ Maximize the return temperature at the cooling units to improve capacity and efficiency
Increase the temperature of the air being returned to the cooling system using the hot-aisle/cold aisle-rack arrangement and containing the cold aisle to prevent mixing of air. Perimeter cooling systems can be supported by row and rack cooling to support higher densities and achieve greater efficiency.
- ☐ Match cooling capacity and airflow with IT loads
Use intelligent controls to enable individual cooling units to work together as a team and support more precise control of airflow based on server inlet and return air temperatures.
- ☐ Utilize cooling designs that reduce energy consumption
Take advantage of energy efficient components to reduce cooling system energy use, including variable speed and EC plug fans, microchannel condenser coils and proper economizers.
- ☐ Select a power system to optimize your availability and efficiency needs
Achieve required levels of power system availability and scalability by using the right UPS design in a redundant configuration that meets availability requirements. Use energy optimization features when appropriate and intelligent paralleling in redundant configurations.
- ☐ Design for flexibility using scalable architectures that minimizes footprint
Create a growth plan for power and cooling systems during the design phase. Consider vertical, horizontal and orthogonal scalability for the UPS system. Employ two-stage power distribution and a modular approach to cooling.
- ☐ Enable data center infrastructure management and monitoring to improve capacity, efficiency and availability
Enable remote management and monitoring of all physical systems and bring data from these systems together through a centralized data center infrastructure management platform.
- ☐ Utilize local design and service expertise to extend equipment life, reduce costs and address your data center's unique challenges
Consult with experienced data center support specialists before designing or expanding and conduct timely preventive maintenance supplemented by periodic thermal and electrical assessments.

Emerson Network Power

1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
800.877.9222 (U.S. & Canada Only)
614.888.0246 (Outside U.S.)
Fax: 614.841.6022
EmersonNetworkPower.com
Liebert.com

While every precaution has been taken to ensure accuracy and completeness in this literature, Liebert Corporation assumes no responsibility, and disclaims all liability for damages resulting from use of this information or for any errors or omissions.

© 2011 Liebert Corporation. All rights reserved throughout the world. Specifications subject to change without notice.

All names referred to are trademarks or registered trademarks of their respective owners.

®Liebert and the Liebert logo are registered trademarks of the Liebert Corporation. Business-Critical Continuity, Emerson Network Power and the Emerson Network Power logo are trademarks and service marks of Emerson Electric Co. ©2011 Emerson Electric Co.

SL- 24664 R08-11 Printed in USA

Emerson Network Power.

The global leader in enabling Business-Critical Continuity™.

- | | | | |
|----------------|--|------------------------------|-------------------------------|
| ■ AC Power | ■ Embedded Computing | ■ Outside Plant | ■ Racks & Integrated Cabinets |
| ■ Connectivity | ■ Embedded Power | ■ Power Switching & Controls | ■ Services |
| ■ DC Power | ■ Infrastructure Management & Monitoring | ■ Precision Cooling | ■ Surge Protection |

EmersonNetworkPower.com