



Energy Logic 2.0

New Strategies for Cutting Data Center Energy Costs and Boosting Capacity



Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion



Introduction

When Energy Logic was introduced in 2007, data center energy efficiency was just emerging as a serious issue. Increases in data center density and capacity were driving up energy bills while concerns over global warming had spurred a U.S. EPA report on data center energy consumption. The industry responded with a number of tactical approaches, but no cohesive strategy for optimizing efficiency.

Energy Logic filled that gap. It bucked the conventional wisdom of the time, which focused on the energy usage of data center support systems, most notably cooling, while virtually ignoring the efficiency of the IT systems that consume more than half of data center energy and drive the demand for cooling and other support systems. (This oversight is perpetuated by the current reliance on PUE, which is discussed later in this paper in the section PUE Analysis.)

Instead, Energy Logic took an “inside-out” approach that drives improvement in IT efficiency as well as the efficiency of support systems. Through this more strategic approach, Energy Logic was able to leverage the cascade effect that occurs when lower energy consumption at the component and device level is magnified by reducing demand on support systems.

Savings from the original Energy Logic, introduced in 2007, were calculated by constructing a detailed statistical model of a 5,000 square foot (464.5 square meter) data center housing 210 server racks with an average rack density of 2.8 kW. Each Energy Logic strategy was then applied to the model to calculate the impact on energy consumption.

Energy Logic demonstrated that a 1 Watt savings at the processor level produced a 2.84 Watt savings at the facility level through the cascade effect.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion

Through this analysis, Energy Logic demonstrated that a 1 Watt savings at the processor level produced a 2.84 Watt savings at the facility level through the cascade effect. In addition, the ten strategies in Energy Logic, working together, created a 52 percent savings in data center energy consumption while freeing up 65 percent of data center space.

As the original **Energy Logic** paper noted, the analysis focused on efficiency, but the same strategies can be applied to remove constraints to growth and increase data center capacity. As IT budgets tightened following the 2008 economic collapse, the ability to cost-effectively expand capacity was the driving force for many organizations seeking to enhance data center efficiency.

Introduction

- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
- PUE Analysis
- Key Takeaways
- Conclusion



View the original Energy Logic video.





Updating Energy Logic

Data center technologies have advanced since Energy Logic was introduced, creating new opportunities to optimize efficiency and capacity. As a result, Emerson Network Power has developed Energy Logic 2.0 to incorporate the advances in technology and new best practices that have emerged in the past five years.

While Energy Logic has been updated, the core principles remain unchanged:

- The greatest savings are achieved by focusing on the core IT systems that drive data center power consumption, thus leveraging the cascade effect.
- The data center can only operate efficiently when energy consumption changes with demand. Systems that cannot operate efficiently at less than peak demand waste energy.
- It is possible to achieve significant reductions in data center energy consumption without reverting to untried designs or technologies that sacrifice data center performance.

There are actually more similarities than differences between Energy Logic 2.0 and the original Energy Logic. That's a testament to the soundness of the Energy Logic approach and, to a degree, a reflection of the lack of progress that has been made in optimizing data center efficiency: the data center of today looks very much like the data center of 2007 in many respects.

Emerson Network Power has developed Energy Logic 2.0 to incorporate the advances in technology and new best practices that have emerged in the past five years.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis


Key Takeaways

Conclusion

Servers, of course, deliver more processing power with greater efficiency, resulting in higher rack densities. However, few data centers are taking advantage of the highest efficiency components available to them and lack of visibility into real-time data center performance continues to limit the ability of data center managers to leverage all of the optimization opportunities available.

The advances in server processing power are estimated to increase the total power consumption of the 5,000 square foot (464.5 square meters) data center established in the original Energy Logic from 1,127 kW to 1,543 kW. As in the original Energy Logic, slightly more than half of energy is consumed by IT equipment for a PUE of 1.92 (Figure 1).

- Introduction
- Updating Energy Logic**
- Calculating the Impact
- Energy Logic Strategies
- PUE Analysis
- Key Takeaways
- Conclusion



The visibility and control provided by DCIM is so integral to Energy Logic 2.0 that it is impossible to attribute some isolated percent of energy savings to DCIM. DCIM enables multiple Energy Logic strategies.

Figure 1. IT systems account for 52 percent of energy consumption in a typical 5,000 square foot (464.5 square meter) data center.

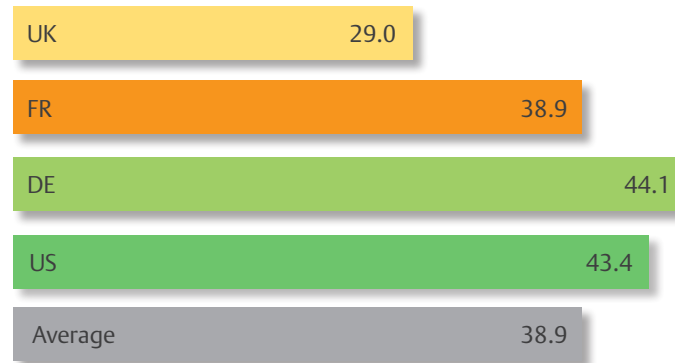




Virtualization is the one area where the market has advanced beyond the original Energy Logic vision. In 2007, the Energy Logic base data center assumed no virtualization while the optimized data center employed virtualization on 20 percent of servers. Today, the average level of server virtualization exceeds 30 percent¹ (Figure 2).

The strategies in Energy Logic 2.0 have been updated to reflect the latest technology and best practices. In addition, two strategies have been revised to reflect new technologies and best practices not available when Energy Logic was first introduced:

- Strategy 4 previously focused on blade servers and their ability to save energy through shared components. In Energy Logic 2.0, blade servers are not called out as a separate strategy. Instead, server consolidation has been integrated into Energy Logic 2.0's recommendations on virtualization to better reflect the fact that consolidation generally happens in concert with virtualization. Strategy 4 in Energy Logic 2.0 now focuses on Information and Communications Technology (ICT) architecture, an emerging best practice that delivers energy savings by optimizing IP connections in the data center.
- The other significant change in Energy Logic 2.0 is that it takes full advantage of data center infrastructure management (DCIM). This may seem like a relatively minor advance — updating the final strategy in Energy Logic from monitoring to DCIM — but the capabilities that DCIM enables are significant compared to monitoring, which is a component of DCIM. In fact, the visibility and control provided by DCIM is so integral to Energy Logic 2.0 that it is impossible to attribute some isolated percent of energy savings to DCIM. DCIM enables multiple Energy Logic strategies.



The V-index Penetration Rate

September 2011 (%)

Figure 2. According to the V-index, a measure of virtualization penetration by percent of servers from Veeam Software and Vanson Bourne, the virtualization penetration rate as of September 2011 was 38.9 percent.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion



Calculating the Impact

One of the challenges organizations face in optimizing data center performance is balancing what often seem to be competing objectives: manage costs, meet growing demand for compute power and ensure continuous availability. Energy Logic provides the flexibility to use efficiency improvements to cut costs or increase capacity without compromising data center availability, thus allowing data center managers to effectively address what appear to be competing objectives without conflict.

The Cascade Effect

The cascade effect is the linchpin of the Energy Logic strategy, providing a clear focus for data center efficiency initiatives. In a data center with a PUE of 1.9, a 1 W savings at the server processor creates a 2.84 W savings at the facility level (Figure 3). At higher PUEs, the savings will be greater.

The cascade effect is the linchpin of the Energy Logic strategy, providing a clear focus for data center efficiency initiatives.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

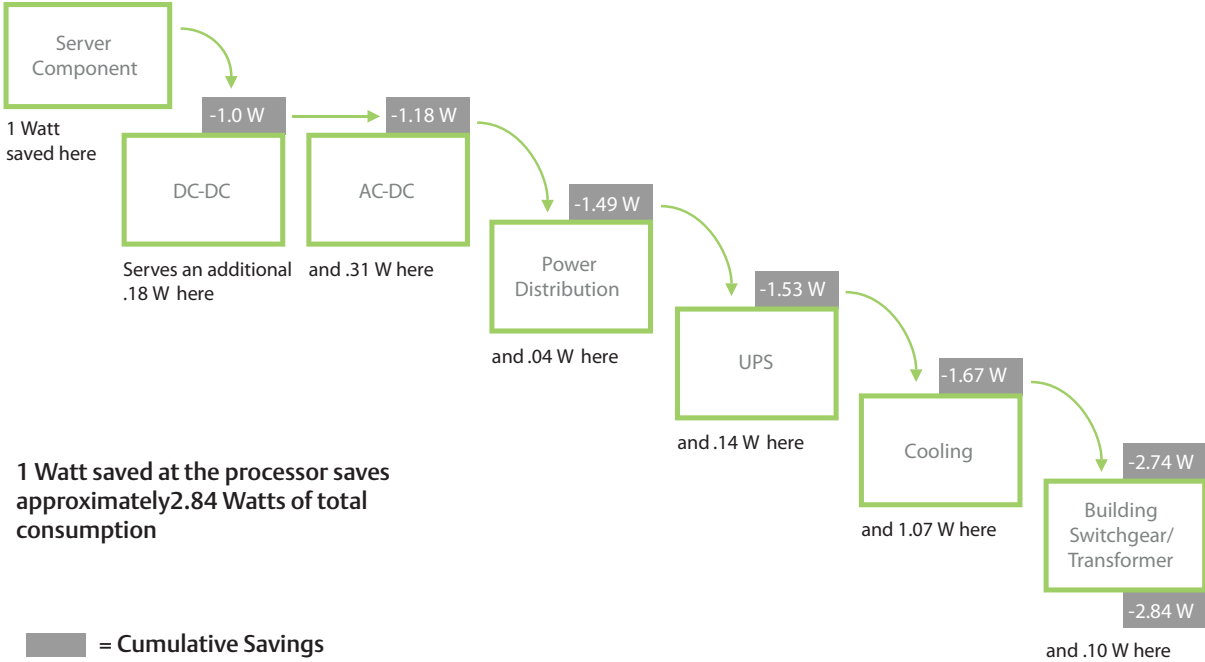
Conclusion



See how Emerson used the original Energy Logic to create a high-efficiency, high-availability data center.



The Cascade Effect



1 Watt saved at the processor saves approximately 2.84 Watts of total consumption

- Introduction
- Updating Energy Logic
- Calculating the Impact**
- Energy Logic Strategies
- PUE Analysis
- Key Takeaways
- Conclusion

Figure 3. Savings at the server component level cascade across support systems to amplify savings. At a PUE of 1.9, 1 Watt saved at the server component level results in cumulative savings of about 2.84 Watts.



Energy Efficiency

The strategies in Energy Logic 2.0 can reduce the energy consumption of a typical 5,000 square foot data center by up to 74 percent. Assuming a facility load of 1,543 kW and energy costs of \$.08 per kilowatt hour, annual energy costs would be cut from \$1,081,334 to \$285,926. At a cost of \$.15 per kilowatt hour the savings are even greater, with a reduction in annual energy costs from \$2,027,502 to \$536,112 (Figure 4).

Capacity

Space, cooling and power are the three most common constraints to data center capacity growth. Energy Logic loosens these constraints by converting efficiency gains into additional capacity. Increases in IT efficiency equate to additional UPS capacity while cooling improvements enable higher rack density, freeing up physical space. By increasing rack density from an average of 5 kW to 12 kW, along with the other increases in efficiency achieved through Energy Logic 2.0, the number of server racks is slashed from 161 to 27, creating the potential for an 83 percent savings in data center space.

Availability

The strategies in Energy Logic have been carefully selected to deliver efficiency gains without compromising data center availability. In some cases data center availability objectives may dictate how certain steps are applied, and facility type may determine where the greatest savings will be realized. Facilities operating at high utilization rates throughout a 24-hour day will want to focus initial efforts on sourcing IT equipment with low-power processors and high-efficiency power supplies. Facilities that experience predictable peaks in activity may achieve the greatest benefit from power management technology and select cooling enhancements such as containment, which, in addition to its efficiency benefits, can extend thermal ride-through time and enable more precise capacity control. Yet, all strategies are designed to be applicable to all data centers and all can be implemented without increasing the risk of downtime.




Figure 4. Annual energy costs before and after Energy Logic. Before costs based on 1,543 kW load. After costs based on 408 kW load.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion

Energy Logic 2.0 Strategies

ENERGY LOGIC			
	LOAD	SAVINGS (KW)	SAVINGS %
Low-Power Processors	1,127	111	10
High-Efficiency Power Supplies	1,016	124	11
Power Management Features	892	86	8
Blade Servers	806	7	1
Server Virtualization	799	86	8
High-Voltage AC Power Distribution	713	20	1
Cooling Best Practices	693	15	1
Variable Capacity Cooling	678	49	4
Supplemental Cooling	629	72	7
Monitoring Optimization	557	15	1
TOTALS	542	585	52

ENERGY LOGIC 2.0			
	LOAD	SAVINGS (KW)	SAVINGS %
Low-Power Components	1,543	172	11.2
High-Efficiency Power Supplies	1,371	110	7.1
Server Power Management	1,261	146	9.4
ICT Architecture	1,115	53	3.5
Server Virtualization and Consolidation	1,062	448	29
Power Architecture	614	63	4.1
Temperature and Airflow Management	551	80	5.2
Variable-Capacity Cooling	471	40	2.6
High-Density Cooling	431	23	1.5
Data Center Infrastructure Management	408	--*	
TOTALS	408	1,135	73.6

*Because DCIM is integral to many Energy Logic 2.0 strategies, it isn't possible in this model to attribute an isolated savings percentage to DCIM.

Figure 5. Energy Logic 2.0 assumes a 415 kW increase in total data center power consumption compared to the original Energy Logic. Based on that load, the Energy Logic strategies outlined in this paper have the potential to reduce energy consumption to 408 kW, a 73.6 percent improvement in data center efficiency.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies**
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



1. Low-Power Components

The cascade effect rewards savings at the component level, which is why low-power components represent the first step in Energy Logic 2.0. The original Energy Logic highlighted the savings that could be realized by improving processor efficiency, and there are still significant opportunities on that front.

Thermal design power (TDP) is the best metric available for processor efficiency. While server efficiency has improved considerably in the past five years, those improvements have been offset by steady growth in compute capacity. Therefore, server processors consume about the same power as they did when the original Energy Logic was introduced, which was estimated at 91 W.

Processor manufacturers continue to advance the state-of-the-art with high-efficiency processors that consume 40 to 60 W less than standard processors. Independent research studies show these low-power processors deliver the same performance as higher power models (Figure 6).

If processor power is reduced from a 91 W average to a 54 W average, an 11.2 percent (172 kW) reduction in data center energy consumption can be achieved.

In addition, DDR3 and DDR4 RAM represent lower power alternatives to traditional server memory, although the savings these components deliver may be offset by increasing server memory. Replacing mechanical hard disk drives with solid-state drives can also enhance server efficiency.

If processor power is reduced from a 91 W average to a 54 W average, an 11 percent (172 kW) reduction in data center energy consumption can be achieved.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling


Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion

System Performance Results



Source: Anandtech

Figure 6. High-efficiency processors provide performance similar to traditional processors while creating energy savings that cascade across the data center.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
- Low-Power Components**
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



2. High-Efficiency Power Supplies

As with other server components, power supplies in use today remain below the efficiencies currently available. While power supply efficiency has improved since the original Energy Logic assumed 79 percent power supply efficiency, they continue to consume more energy than is necessary. Average power supply efficiency is now estimated at 86.6 percent, well below the 93 percent that is available.

Increasing efficiency from 86.6 percent to 93 percent reduces total data center power consumption by 7.1 percent. In the theoretical 5,000 square foot data center, this amounts to 110 kW—7 percent of the total savings.

As with other data center systems, server power supply efficiency varies depending on load (Figure 7). Some power supplies perform better at partial loads than others, and this is particularly important in dual-corded devices where power supply utilization can average less than 30 percent. Figure 7 shows power supply efficiencies at different loads for two power supply models. At 20 percent load, model A has an efficiency of approximately 88 percent while model B has efficiency closer to 82 percent.

Increasing efficiency from 86.6 percent to 93 percent reduces total data center power consumption by 7 percent.



Watch Emerson experts address the latest trends in power supply efficiency.



Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion




Figure 7. Power supply efficiency at various loads.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies**
 - Server Power Management
- ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
- Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
- Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



3. Server Power Management

Server power management remains an untapped opportunity for reducing data center use because most data centers rarely run at full capacity. This is a problem because, consequently, a facility operating at just 20 percent capacity may use 80 percent of the energy as the same facility operating at 100 percent capacity.

Server power management can significantly reduce the energy consumption of idle servers, but is not utilized in the typical data center because of concerns about response times for “waking” an idle server using power management.

The Green Grid, an industry consortium focused on improving data center resource efficiency, has conducted research² into server power management to identify the chief obstacles to adoption and is developing materials to help educate data center managers on server power management and increase utilization of this technology.

In addition, new research has revealed that the risks of applying power management to older servers may be very low, as these servers contribute little to data center performance. In a presentation at AFCOM Data Center World in 2011, William Carter and John Kuzma of Intel³ presented an analysis of server utilization at one enterprise data center and found that servers installed before 2008 accounted for 60 percent of energy consumption, but only delivered 4 percent of relative performance capability (Figure 8).

Energy Star Servers

The 2007 U.S. EPA report on data center consumption has spawned efforts to establish meaningful data center efficiency metrics and to extend the successful Energy Star program to data center servers.

Specifications for Energy Star servers⁴ have now been published and align with Energy Logic principles, including:

- 1 kW or larger servers must use power supplies that deliver 92 percent efficiency at 50 percent load.
- Servers with two or fewer processors must have an idle power consumption of less than 55 W.
- All three- and four-socket servers must enable power management to reduce power of the processor during times of low utilization.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis


Key Takeaways

Conclusion

Long-term, older servers that contribute little to data center performance need to be identified and consolidated into newer servers; however, power management represents a powerful interim solution to cut server power consumption without additional technology investment. It also represents an ideal long-term strategy for enabling server power consumption to adapt dynamically to changes in data center load.

Data center infrastructure management systems that can collect real-time operating data from rack power distribution systems and consolidate this data with server utilization data provide the visibility to identify stranded capacity and enable safe and effective use of server power management.

Implementing power management can reduce total data center energy consumption by 10 percent, cutting 146 kW from the 1,261 kW load remaining after strategies 1 and 2 have been implemented.



Source Intel Corp.

Figure 8. In a study of an enterprise data center by Intel, older servers consumed 60 percent of the energy but delivered only 4 percent of performance.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management**
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



4. ICT Architecture

Because of the way they have grown, today's data centers typically feature a silo-type architecture with virtualization limited to the individual application layer by line-of-business requirements. Unoptimized network architectures often contain duplication and lack asset tracking and coordination of the network switching/routing infrastructure. This disjointed Information and Communication Technology architecture is similar to a fragmented hard disk in that both efficiency and performance are compromised.

Efficiency gains can also be realized through new structural cabling technology that reduces heat load by speeding information transmission⁵ (Figure 9).

Implementing a cohesive ICT architecture involves establishing policies and rules to guide design and deployment of the networking infrastructure, ensuring all data center systems fall under the same rules and management policies. The network is architected as one end-to-end system owned and managed by IT as a service to the business. Servers and storage shift from line-of-business to IT ownership with limited exceptions for internal HPC-type applications and/or legal requirements. In addition, IT assets are deployed in accordance with a master plan that ensures systems scale to workloads, are located to minimize network size and costs, and take advantage of centralized control through DCIM.

This approach has the potential to contribute an additional 53 kW of energy savings or 3.5 percent of total savings.

Implementing a cohesive ICT architecture involves establishing policies and rules to guide design and deployment of the networking infrastructure, ensuring all data center systems fall under the same rules and management policies.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

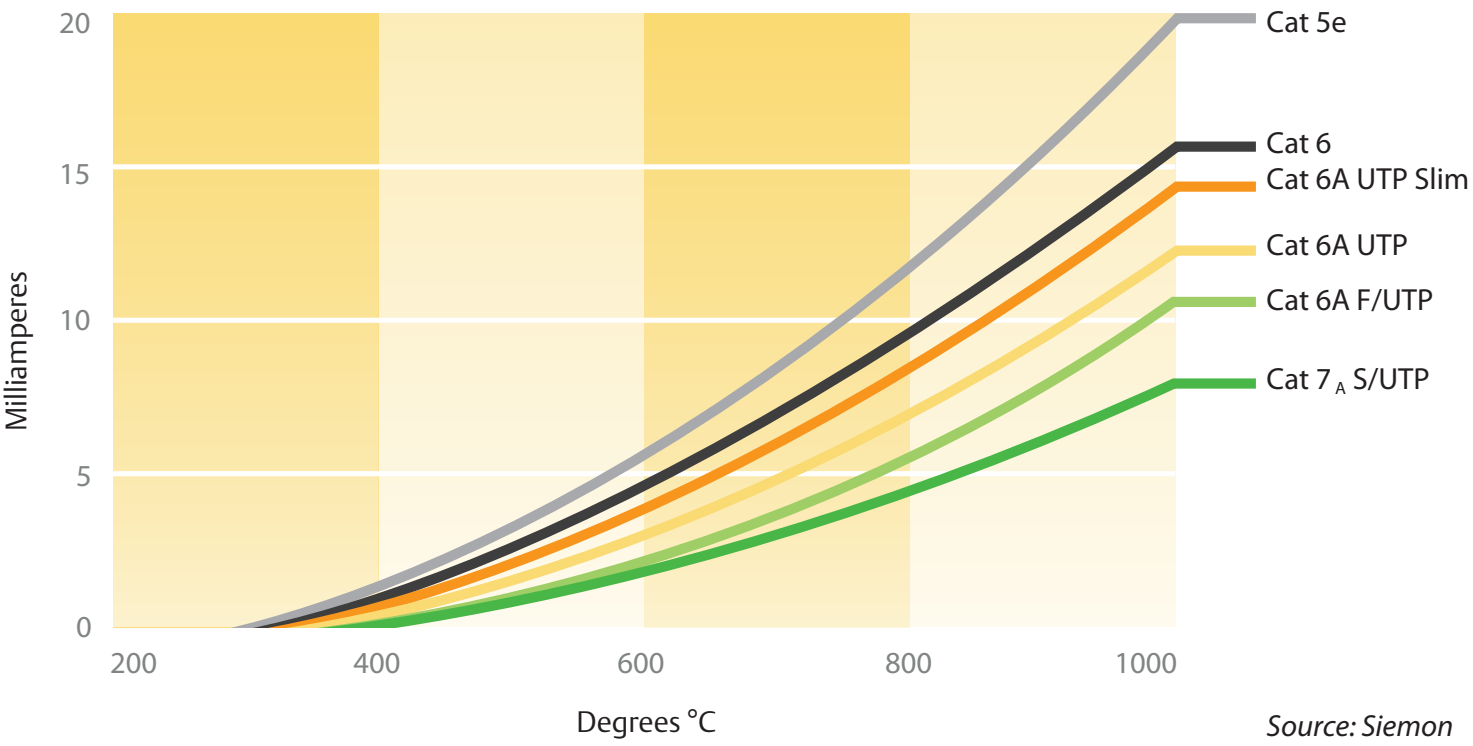
High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



Source: Siemon

Figure 9. Older communication cables generate more heat, reducing data center efficiency.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture**
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



5. Server Virtualization and Consolidation

Virtualization allows older, power-wasting servers to be safely consolidated onto much less hardware. It also increases the ability of IT staff to respond to changing business needs and computing requirements. The static nature of a physical server makes it difficult to respond to varying IT loads; hardware must be over-provisioned to account for peaks in demand, which occur infrequently. By allowing multiple applications on the same server, IT capacity can be more accurately sized to actual demand, significantly reducing the number of servers required to support demand. According to VMware, every server that is virtualized saves 7,000 kWh of electricity and four tons of carbon dioxide emissions per year.⁶

Most data centers have already discovered the benefits of virtualization, but there is often opportunity to go further. For example, increasing from 30 percent server virtualization to 60 percent can provide a 29 percent reduction in data center energy consumption — the largest contribution of any Energy Logic strategy. This represents savings of 448 kW in the 1,543 kW data center.

DCIM can play an important role in helping organizations increase the level of virtualization and manage the virtual environment. It provides visibility into how virtual servers are deployed and the infrastructure capacity supporting those virtual servers. With DCIM, infrastructure capacity can be fully utilized without risking over provisioning.



Ben Matheson of VMware and Jake Smith of Intel discuss server virtualization.



Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



6. Power Architecture


Historically, data center designers and managers have had to choose between availability and efficiency in the data center power system. Double-conversion UPS systems provided the highest availability but could not deliver the same efficiency as line-interactive systems. The conversion process that ensures comprehensive power conditioning and better isolation of sensitive electronics from the power source also introduces some losses into the process.

Now, however, advances in double-conversion UPS technology have closed the gap in efficiency, and new features have been introduced that enable double-conversion UPS systems to achieve efficiencies on par with line-interactive systems.

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. This is accomplished by incorporating an automatic static-switch bypass into the UPS. The bypass operates at very high speeds to provide a break-free transfer of the load to a utility or 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 4 milliseconds to prevent any interruption that could shut down IT equipment. Using intelligent controls, the bypass can be kept closed, bypassing the normal AC-DC-AC conversion process while the UPS monitors bypass power quality.



**Emerson Network Power's
Peter Panfil reviews
different approaches to
UPS "eco-mode."**



Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and
Consolidation

Power Architecture

Temperature and Airflow
Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure
Management

PUE Analysis

Key Takeaways

Conclusion

When the UPS senses power quality falling outside accepted standards, the bypass opens and transfers power back to the inverter so anomalies can be corrected. The inverter must be kept in a constant state of preparedness to accept the load and thus needs control power, and the transfer must occur with no break to prevent an interruption in the critical bus. The power requirement is below 2 percent of the rated power, creating potential savings of 4.1-4.5 percent compared with traditional operating modes.



3 Units @ 25% Load Each = 91.5% Efficiency



2 Units @ 38% Load Each = 93.5% Efficiency

- Stand-by unit idled
- Schedule based on load operations
- Distribute operations to all units

Figure 10. Firmware intelligence for multi-module UPS systems increases efficiency by idling unneeded inverters or whole modules.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture**
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion



Another newer function enabled by UPS controls, intelligent paralleling, may contribute additional savings. Intelligent paralleling 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 250 kVA with only two modules. Deactivating one module maintains redundancy and improves the efficiency of the remaining modules by enabling them to operate at a higher load (Figure 10).

In the distribution systems, high-efficiency transformers can minimize losses between the UPS and the rack. DC power may also be a viable option for optimizing the power architecture. DC power provides efficiencies comparable to the highest efficiency AC power systems while delivering modular, redundant paralleling and full isolation from the grid in all normal operating modes.

DCIM can support intelligent control in the power system by providing a holistic view of data center energy consumption and server utilization to fine-tune eco mode operation or determine whether intelligent paralleling can be employed.

Implementing eco-mode on the UPS and optimizing the power path from the UPS to servers and other connected equipment can reduce data center power consumption by 4 percent of the total, or 10 percent of the 614 kW load remaining after other strategies have been implemented.

DCIM can support intelligent control in the power system by providing a holistic view of data center energy consumption and server utilization.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



7. Temperature and Airflow Management

The original Energy Logic advocated use of hot-aisle/cold-aisle rack arrangement and stressed the importance of sealing gaps in the floor. In addition, it suggested exploring the opportunity to raise chilled water temperatures to 50° F (10° C). This reduced overall facility energy costs by 1 percent with virtually no investment in new technology. Energy Logic 2.0 builds on those best practices, which are now firmly established in most enterprise data centers, to take temperature, humidity and airflow management to the next level through containment, intelligent controls and economization.

From an efficiency standpoint, one of the primary goals of preventing hot and cold air from mixing is to maximize the temperature of the return air to the cooling unit. The relationship between return air temperature and sensible cooling capacity is illustrated in Figure 11. It shows that a 10° F (5.6° C) increase in return air temperature typically results in a 30 to 38 percent increase in cooling unit capacity, depending on the type of system.⁷ This additional capacity translates into more efficient operation in a DX system by running the compressors “unloaded” and in a chilled water system by running the chillers unloaded and the pumps at reduced flow. It can also increase the sensible heat ratio (SHR) of the precision cooling units to closer to 1, avoiding unnecessary dehumidification.

The racks themselves arranged in a hot-aisle/cold-aisle configuration 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.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion

10° F higher return air temperature typically enables 30-30% better CRAC efficiency

Chilled Water CRAC
Direct Expansion CRAC




Figure 11. The efficiency of cooling units improves as the temperature of the return air increases.

Containment involves capping the ends of the aisle, the top of the aisle, or both to isolate the air in the aisle (Figure 12). 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. Even worse, in the absence of a vapor barrier, external humidity can condense on racks and server inlets. 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. In fact, cold-aisle containment can improve system availability by extending ride-through times in the event of a cooling problem. Containment also enables more precise fan speed control, higher chilled water feed temperature and maximum economizer utilization.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management**
 - Variable-Capacity Cooling
 - High-Density Cooling
 - Data Center Infrastructure Management
- PUE Analysis
- Key Takeaways
- Conclusion




Figure 12. Containment improves cooling system efficiency by allowing return air temperatures to be increased and creates an ideal environment for precise supply-air temperature control.

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. 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 is maximized. The possible downside of this approach is that more floor space is consumed in the aisle, but the space savings realized by Energy Logic 2.0 should minimize this concern.

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. The controls ensure the optimum combination of compressor/chiller capacity and air flow and often allow temperatures in the cold aisle to be raised closer to the safe operating threshold recommended by ASHRAE (max 80.5° F; 27° C) for class A1-A4 data centers. According to an Emerson Network Power study, a 10° F (5.6° C) increase in cold-aisle

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and
Consolidation

Power Architecture

**Temperature and Airflow
Management**

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure
Management

PUE Analysis

Key Takeaways

Conclusion



temperature can generate a 20 percent reduction in cooling system energy usage. However, raising temperatures too high can cause server fans to work harder, improving PUE but actually increasing overall energy consumption. The DCIM system can help identify the optimum cold-aisle temperature by comparing temperatures with cooling system and server energy consumption patterns.

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 ensure all cooling fans are running at the same speed for optimum efficiency and 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 at target levels and match airflow to the load.

Various types of economizer systems should also be considered for any data center seeking to optimize efficiency. Operating at the upper temperature limits of the ASHRAE guidelines makes an economizer system even more efficient. Economizers use outside air to provide “free-cooling” cycles for data centers. This reduces or eliminates chiller or compressor operation in precision cooling units and produces significant energy savings in a wide variety of climates. Emerson Network Power has evaluated economizer operation in various climates and determined that even in a warm weather climate like Atlanta, Georgia, full economization is available 25 percent of the year, with partial economizer operation achieved another quarter of the year.⁸

A fluid-side economizer (often called water-side) can be integrated with the chiller or work in conjunction with a heat rejection loop comprising an evaporative cooling tower or drycooler. 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, an exhaust fan and dampers to bring outside air into the controlled environment.



See cold-aisle containment in action at Bay Area Internet Solutions.



Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



The affect of outside air on data center humidity should be carefully considered when evaluating economization options. Introducing outside air via an airside economizer system 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 to handle the room load without compressors. 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.

A new version of fluid-side economizer uses the refrigerant in a DX air-cooled system as the “free cooling” fluid during cold weather. This eliminates a number of inefficiencies associated with other economizer systems, such as extra fans, pumps or coils.

Optimizing the cooling system, which has already had its load reduced significantly by other Energy Logic 2.0 strategies, results in an additional 5.2 percent reduction in energy consumption.

Even in a warm weather climate like Atlanta, Georgia, full economization is available 25 percent of the year, with partial economizer operation achieved another quarter of the year.⁸

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



8. Variable-Capacity Cooling

Like the IT systems themselves, cooling must be sized to handle peak load conditions, which occur rarely in the typical data center. Therefore, cooling systems must be able to operate efficiently at partial loads. This was a challenge for legacy cooling units that were either on or off. Today's cooling systems implement a number of technologies that improve efficiency at part loads, including some of the technologies used for temperature and airflow management covered in the previous section.

On chilled water cooling units, the fans that move air and pressurize the raised floor are the largest consumer of energy. While variable frequency drives represent a significant improvement over traditional fixed-speed fans, electronically commutated (EC) plug fans may provide an even better option for increasing cooling unit efficiency. EC plug fans are inherently more efficient than traditional centrifugal fans. Both variable frequency drive and EC plug fans can be installed on existing cooling units or specified in new units, and work with intelligent controls discussed previously.

The use of variable capacity compressors on both DX and chilled water systems can increase the efficiency of any system by allowing it to operate at higher efficiencies when not operating at full load. There are several approaches to providing variable capacity in a direct expansion CRAC unit. The two most common are four-step compressor unloading, and Digital Scroll™ compressor technology.

The concept of four-step compressor unloading works by shutting off the flow of refrigerant to some of the cylinders within the system; thereby, minimizing the need to cycle compressors on and off to control capacity. Because unloading essentially changes the compressor operating point, it enables the cooling system to operate more efficiently at lower capacities. In this case the system is able to provide four stages of operation to accommodate changes in capacity.

Digital Scroll compressors are not limited to four stages of operation and instead can precisely adjust their output to the load. Digital Scroll technology allows the compressor to never be cycled off. It reduces power consumption linearly as it modulates capacity, resulting in optimum system performance and control.

Optimizing the cooling system to work more efficiently at partial loads can reduce data center energy consumption by an additional 2.6 percent.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



9. High-Density Cooling

Traditional room-cooling systems have proven very effective at maintaining a safe, controlled environment for IT equipment. However, optimizing data center energy efficiency requires moving from traditional data center densities (2 to 3 kW per rack) to an environment that can support much higher densities.

This requires implementing an approach to cooling that shifts some of the cooling load from traditional CRAC units to supplemental cooling units mounted above or alongside equipment racks that pull hot air directly from the hot aisle and deliver cold air to the cold aisle. They reduce cooling costs by bringing cooling closer to the source of heat, reducing the fan power required to move air. They also use more efficient heat exchangers and deliver only sensible cooling, which is ideal for the dry heat generated by electronic equipment. They can even be applied directly to the rear of the rack, effectively removing heat before it enters the data center.

Refrigerant is delivered to the supplemental cooling modules through an overhead piping system, which once installed allows cooling modules to be easily added or relocated as the environment changes. This enables supplemental cooling to be deployed in select zones in the data center.

Optimizing data center energy efficiency requires moving from traditional data center densities (2 to 3 kW per rack) to an environment that can support much higher densities.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



Supplemental cooling can also be performed by water-based systems within the space supporting cooling modules at the rear of the rack. Water-based systems do introduce some additional risk compared to pumped refrigerants, which convert to gas at room pressures. Therefore, a leak detection system should be used in conjunction with water-based cooling on or near equipment racks.

Additionally, the refrigerant delivery system installed to support supplemental cooling systems also can support the next generation of cooling, which promises to eliminate the need for server fans by removing heat directly from the server. Heat is transferred from the processors through heat risers to the server housing and then through a thermal lining to the cooling plate, which eliminates the need to expel air from the rack and into the data center. This can actually create a net positive effect on data center energy consumption in that the cooling system uses less energy than the server fans it makes unnecessary.

High-density cooling creates an additional 1.5 percent reduction in the energy consumption of the base data center, driving consumption from 423 kW down to 408 kW.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



10. Data Center Infrastructure Management

Data center infrastructure management technology is capable of collecting, consolidating and integrating data across IT and facilities systems to provide a centralized real-time view of operations that can help optimize data center efficiency, capacity and availability. From the ability to generate alarms when conditions move outside preset limits to providing the visibility to identify and remove stranded server capacity while safely tapping unused infrastructure capacity, DCIM brings together data from disparate systems to create a unified view of the data center. DCIM also delivers significant operational efficiencies by providing auto-discovery of data center systems and simplifying the process of planning for and implementing new systems.

Without the real-time visibility provided by a DCIM system, data center personnel have been forced to manage ultra-conservatively to avoid any situations that could increase the potential for downtime. This has led to systems that are oversized and underutilized, increasing both capital and operational expenses. With the advent of real-time management and control, organizations can take full advantage of the technologies available to them, such as server power management and virtualization, while ensuring the systems in the data center are using energy wisely.

Because DCIM is integral to many Energy Logic 2.0 strategies, it isn't possible in this model to attribute an isolated savings percentage to DCIM.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

Low-Power Components

High-Efficiency Power Supplies

Server Power Management

ICT Architecture

Server Virtualization and Consolidation

Power Architecture

Temperature and Airflow Management

Variable-Capacity Cooling

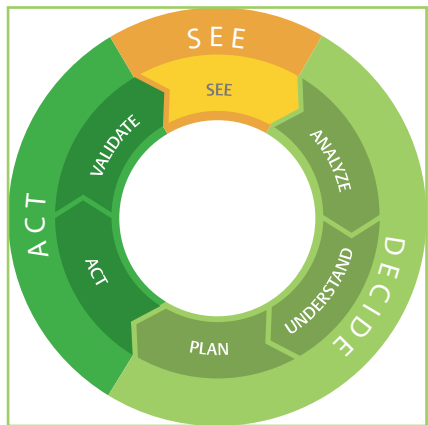
High-Density Cooling

Data Center Infrastructure Management

PUE Analysis

Key Takeaways

Conclusion



Without the real-time visibility provided by a DCIM system, data center personnel have been forced to manage ultra-conservatively to avoid any situations that could increase the potential for downtime.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
 - Low-Power Components
 - High-Efficiency Power Supplies
 - Server Power Management
 - ICT Architecture
 - Server Virtualization and Consolidation
 - Power Architecture
 - Temperature and Airflow Management
 - Variable-Capacity Cooling
 - High-Density Cooling
- Data Center Infrastructure Management**
- PUE Analysis
- Key Takeaways
- Conclusion

Figure 13. DCIM promises closed-loop control based on real-time data from facilities and IT systems.



PUE Analysis

While PUE has become the most widely used metric for measuring data center efficiency, Energy Logic highlights the limitations of PUE and reinforces the importance of taking a broader view of data center performance.

Consider that the Energy Logic 2.0 base data center has a PUE of approximately 1.9 before any of the corrective actions are taken. If an organization was only to adopt the first five Energy Logic 2.0 actions, choosing to transform their data center by replacing older servers with new high-efficiency, virtualized servers with power management enabled, overall data center power consumption would be reduced by about 650 kW, while maintaining the current level of data center performance. Yet, the PUE for this facility would actually get slightly worse, moving from 1.92 to 1.94. Conversely, if the organization chose to begin with the power and cooling systems, implementing Strategies 6, 7, 8 and 9, it would achieve energy savings of only about 200 kW but the PUE would improve.

Implementing all 10 Energy Logic strategies increases the efficiency of IT and support systems, while reducing the load on support systems, reducing PUE from 1.92 to 1.28.

Implementing all 10 Energy Logic strategies increases the efficiency of IT and support systems, while reducing the load on support systems, reducing PUE from 1.92 to 1.28.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion

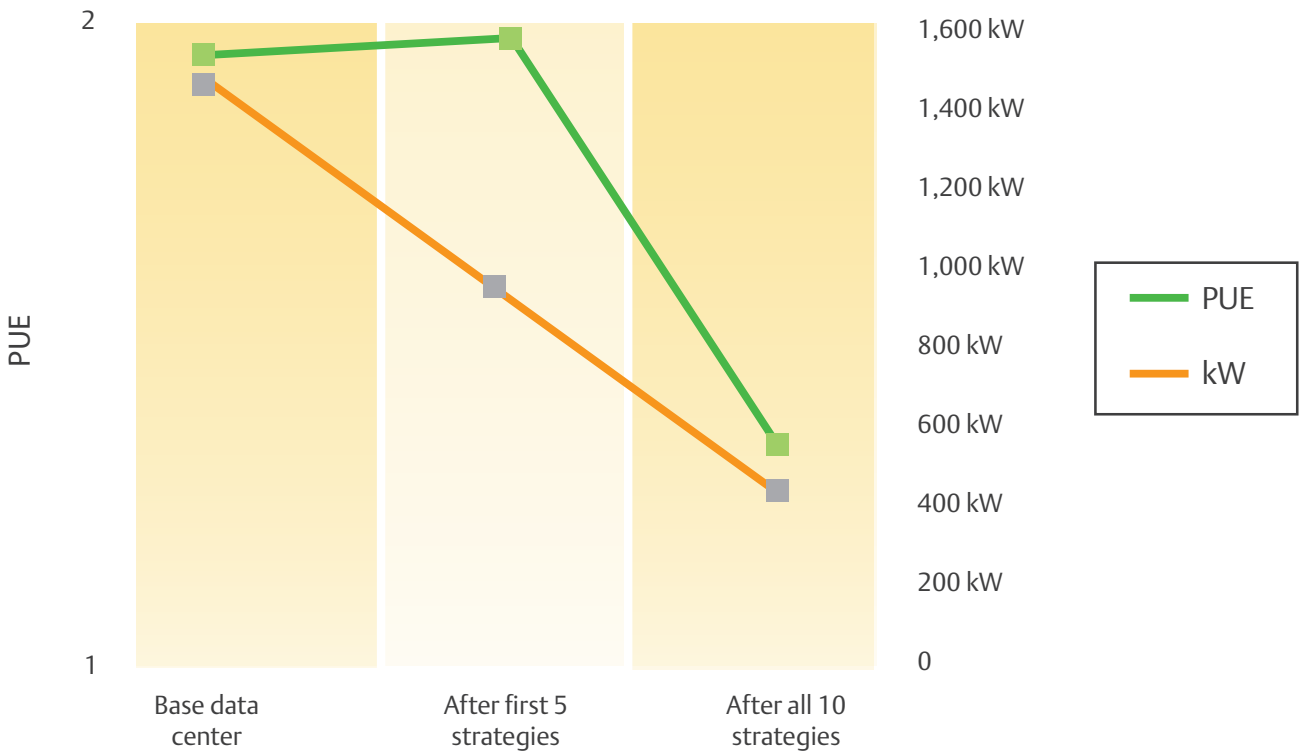


Figure 14. The first five strategies in Energy Logic 2.0 produce a dramatic reduction in energy consumption but a small increase in PUE.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
- PUE Analysis**
- Key Takeaways
- Conclusion



Key Takeaways from Energy Logic

Energy Logic creates a clear roadmap for driving dramatic reductions in energy consumption without jeopardizing data center performance. As this analysis demonstrates, an organization that systematically adopted the Energy Logic 2.0 roadmap could achieve a better than 70 percent reduction in energy consumption while removing constraints to growth.

In support of Energy Logic 2.0, Emerson Network Power has created the Energy Logic 2.0 Cascading Savings Calculator at www.EfficientDataCenters.com. Entering compute load and PUE for a particular facility allows users of the calculator to see the impact of individual Energy Logic 2.0 strategies on compute load, PUE and energy costs.

Not every organization is in a position to adopt every Energy Logic 2.0 strategy. These organizations can still benefit from the Energy Logic 2.0 analysis using these four lessons as a guide when considering data center changes.

1. Leverage the cascade effect

Because support systems account for a relatively high percentage of data center energy consumption, but contribute nothing directly to data center output, it is tempting to attack these systems first. But the load on support systems is determined by the IT load; efficiency improvements in IT systems are amplified in the support systems. In addition, a relatively high percentage of IT systems contribute almost nothing to data center output. This turns the cascade effect on its head — these non-productive IT systems not only waste the energy they consume, they waste the energy of the power and cooling systems that support them. In fact, in the base Energy Logic data center, an unutilized 200 W server actually consumes 381 W of power. That means every Watt of stranded server capacity is actually wasting 1.91 W of energy (Figure 15).

These non-productive IT systems not only waste the energy they consume, they waste the energy of the power and cooling systems that support them.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

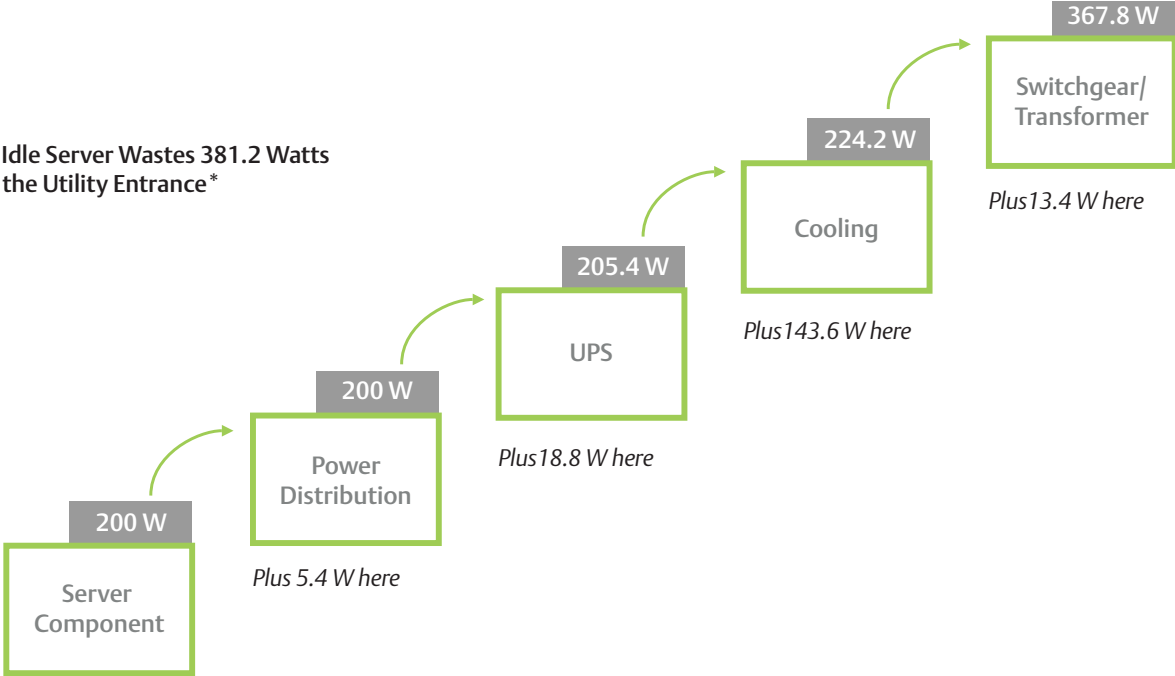
PUE Analysis

Key Takeaways

Conclusion

“Reverse” Cascade Effect

A 200 Watt Idle Server Wastes 381.2 Watts of Power at the Utility Entrance*



200 W Idle Server wastes 200 W

* Assumes a PUE of 1.91
Excludes site and building generator, natural gas/diesel, water, fire, security, etc.

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
- PUE Analysis
- Key Takeaways**
- Conclusion

Figure 15. An unused 200 W server actually consumes 381 W of facility power when support system energy consumption is considered.



2. Don't compromise availability and flexibility for efficiency

Data center energy consumption has created a problem for data center-dependent organizations — and an opportunity for companies seeking to market solutions to that problem. Unfortunately, many of these “solutions” put efficiency above availability and flexibility, which is both dangerous and unnecessary. Energy Logic 2.0 demonstrates that huge reductions in data center energy consumption are possible using proven technologies that do not impact the data center's ability to deliver services.

3. Higher density equals better efficiency

While many of today's data center managers have spent the majority of their careers managing facilities with densities well below 5 kW per rack, the servers and support systems of today are not only capable of being deployed at much higher density, they are designed for it. While a density of 20 kW per rack may make some data center managers nervous, today's data center can be designed to safely and effectively support that density — with room to grow.

4. Capacity is the flip side of efficiency

Despite rising costs, electricity is still relatively cheap in some areas. That has prevented some organizations from aggressively moving to optimize efficiency. However, the Energy Logic roadmap is more than a solution to rising energy costs; it is a solution to increasing demand for compute capacity. Energy Logic can eliminate the need for expensive build-outs or new facilities by removing constraints to growth as the demand for compute and storage capacity continue to grow.

Energy Logic 2.0 demonstrates that huge reductions in data center energy consumption are possible using proven technologies that do not impact the data center's ability to deliver services.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion



Conclusion

The original Energy Logic brought clarity to data center energy efficiency discussions by taking a statistical approach that highlighted the overriding importance of IT equipment in determining data center efficiency. This was an important development in the evolution of data center efficiency; however, economic conditions and a lack of real-time visibility into IT and facilities systems has limited the progress that has been made in optimizing data center performance. What gains have been made have largely been consumed by increases in capacity.

Energy Logic 2.0 illustrates the potential that still exists to optimize the data center, showing how the energy consumption of a “typical” 5,000 square foot data center could be cut by more than 70 percent using available technologies. With efficiency remaining a priority and the introduction of a new generation of management systems that provide greater visibility and control of data center systems, the time is now for the industry to begin making large strides in reducing the overall energy consumption of data centers.

While a density of 20 kW per rack may make some data center managers nervous, today’s data center can be designed to safely and effectively support that density—with room to grow.

Introduction

Updating Energy Logic

Calculating the Impact

Energy Logic Strategies

PUE Analysis

Key Takeaways

Conclusion



References

1. [V-index.com](#), Veeam Software and Vanson Bourne
2. [A Roadmap for the Adoption of Power-Related Features in Servers](#), The Green Grid
3. [Assess IT Efficiency Using Server Utilization Effectiveness](#), Intel
4. [Energy Star Program Requirements for Computer Servers](#), EPA
5. [IEEE802.3at PoE Plus Operating Efficiency](#), Siemon
6. [How VMware Virtualization Right-Sizes IT Infrastructure to Reduce Power Consumption](#), VMware
7. [Seven Best Practices for Increasing Efficiency, Availability and Capacity: The Enterprise Data Center Design Guide](#), Emerson Network Power
8. [Economizer Fundamentals: Smart Approaches to Energy-Efficient Free Cooling for Data Centers](#), Emerson Network Power

- Introduction
- Updating Energy Logic
- Calculating the Impact
- Energy Logic Strategies
- PUE Analysis
- Key Takeaways
- Conclusion**

Emerson Network Power.

The global leader in enabling *Business-Critical Continuity*TM

- AC Power
- Embedded Computing
- Infrastructure Management & Monitoring
- Connectivity
- Embedded Power
- Outside Plant
- DC Power
- Industrial Power
- Power Switching & Controls

EmersonNetworkPower.com

- Precision Cooling
- Racks & Integrated Cabinets
- Services

Cost savings, efficiency and analysis provided for illustration and estimation purposes only. Actual experience may vary based on various factors, including but not limited to site design, site conditions, environmental parameters, annual operating cost estimates. Comparisons based on publically available information. Emerson Electric Co. and/or its affiliates make no representations or warranty about the accuracy, reliability, completeness, or timeliness of the materials and disclaim any and all liability for damages resulting from use of this information or for any errors or omissions. Specifications subject to change without notice. © 2012 Emerson Electric Co. All rights reserved. Business Critical Continuity, Emerson Network Power, and the Emerson Network Power logo are trademarks and/or service marks of Emerson Electric Co. All other product names, brand names, and/or company names are the property and/or trademarks of their respective owners.