



**BUILDING A MORE
PROFITABLE
COLOCATION
ENVIRONMENT**

Introduction

Data center management is currently undergoing a period of great change, much of which is being driven by two factors: new technologies and increasing demands. Data center managers are struggling to keep pace with growing capacity needs while working under the constraints of tightened budgets and energy efficiency initiatives. New technologies such as virtualization, cloud computing and data center infrastructure management (DCIM) are transforming data centers into dynamic environments, optimizing the space in ways no one could have predicted only a few short years ago.

Under this new “norm,” data center managers are finding that they now must lower operational expenses, justify equipment upgrades, delay capital expansions and explore ways to reduce energy consumption and meet growing demand without risking downtime, all while optimizing to accommodate future needs. To meet these sometimes conflicting objectives, one option data center managers are increasingly turning toward is third-party colocation.

This white paper is designed to provide colocation facilities with best practices for building a scalable data center infrastructure that will enable them to quickly add customers, while optimizing efficiency at every stage of infrastructure utilization.

Challenges Faced by Colocation Providers

According to results of the Uptime Institute’s 2012 Data Center Survey, “around half of the respondents will turn to third party data center providers to deal with data center growth in the next 12-18 months.” Nearly 62 percent of respondents to a 2012 Data Center Users Group (DCUG) study sponsored by Vertiv™ have considered collocating their data center. Slightly more than 38 percent of that group reported that they did move forward with collocating at least a portion of their data center (Figure 1).

As colocation providers work to meet this new demand for their services, they are driving major changes in the industry. Customer demands, such as faster speed of deployment, energy efficiency, higher availability and capacity-on-demand, are driving colocation providers to incorporate new and innovative technologies; and according to a 2012 Vertiv customer survey, they are installing the technology at a faster pace than enterprise data centers (Figure 2).

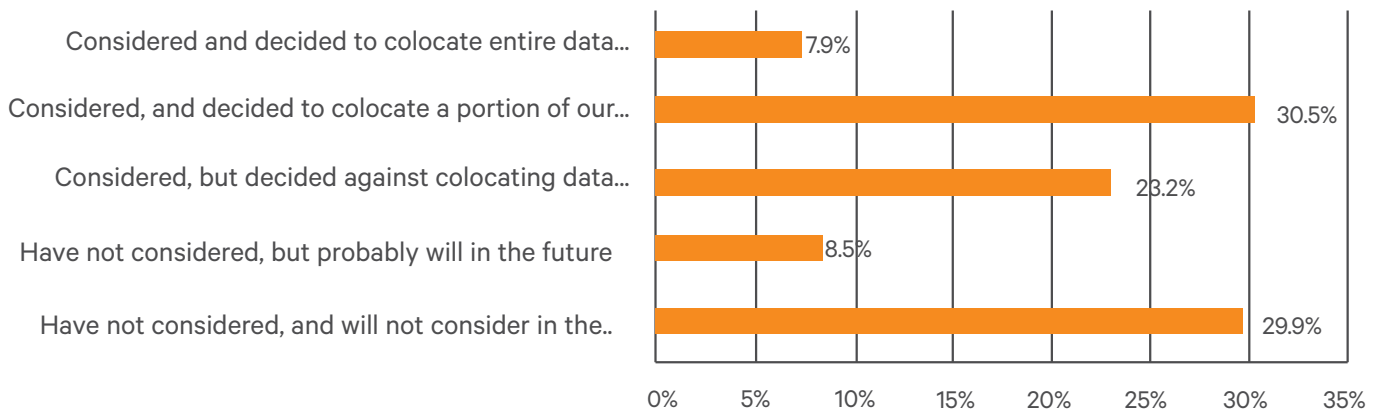


Figure 1: The Data Center Users Group members consider colocation (Fall 2012).

	COLOCATION	ENTERPRISE
DC Power for servers	35%	8%
Transformer-free UPS	27%	6%
High Redundancy in Critical Path	54%	38%
Higher Voltages	27%	15%
Power Containers / Skids	16%	6%
Economization	11%	7%
Modular Chilled Water Plants	30%	17%

Figure 2: Colocation providers installing new technology at higher rate than enterprise data centers (Fall 2012).

Tweet This:

#Colocation providers adopting new #datacenter technologies at a much faster pace than enterprise counterparts: <http://bit.ly/Yu7kny>

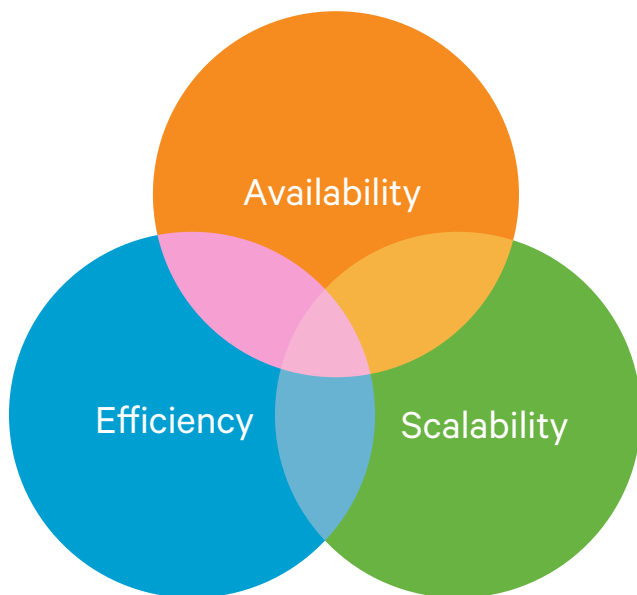
Availability, Efficiency and Scalability

New technologies and approaches to data center design and operation share attributes which drive improvement in three areas: availability, efficiency and scalability. These attributes all contribute to the bottom line by:

1. Reducing costs through a just-in-time approach to infrastructure deployment.
2. Improving uptime by avoiding penalty fees and lowering operating costs through deploying energy efficient technology.

Availability of IT capacity has traditionally been the most important metric on which data centers are evaluated, which is a fact well understood by colocation providers. However, all data centers today must operate efficiently—in terms of both energy and management resources—and be flexible enough to quickly and cost-effectively adapt to both changes in both business strategy and computing demand.

Many in the industry are beginning to realize that using scalable infrastructure technologies for operation and management allows dynamic IT systems and data centers to achieve availability while maximizing opportunities to improve energy and operational efficiencies.



This “trifecta” of availability, efficiency and scalability is especially important for colocation providers, since the data center is their business. Every opportunity to increase availability, efficiency and scalability is an opportunity to increase competitiveness and build the bottom line. For customers, the value proposition of the colocation provider is that they can do it “cheaper” and “faster” while delivering better reliability. To deliver on this value proposition, colocation providers need to do everything they can to optimize the performance of their facilities.

Infrastructure Best Practices for Colocation Facilities

The following strategies expand upon existing data center design approaches and help colocation providers meet their value proposition through a sophisticated, flexible infrastructure that enables them to deploy capacity faster at a lower cost and with greater reliability.

Faster Deployment through Scalability

One of the most important challenges that must be addressed in any data center, especially colocation facilities, is configuring systems to meet current requirements while ensuring the ability to quickly and cost-effectively meet growing and fluctuating performance and capacity needs. In the past, this was partially accomplished by oversizing infrastructure systems and letting the data center grow into its infrastructure over time. As many organizations discovered, this approach can create inefficiencies in terms of both capital and energy costs.

Today, flexibility can be provided by a new generation of infrastructure systems that are designed for simplified configuration and greater scalability, enabling systems to be right-sized during the design phase, with potential for planned growth for dynamic environments. These modular technologies especially help in managing a data

center infrastructure that experiences fluctuations in capacity needs.

Tweet This:

White paper shows how #colo providers can achieve the “trifecta” of availability, efficiency & scalability <http://bit.ly/Yu7kny>

Modularity and scalability of infrastructure equipment reduces costs and improves deployment speed during construction and at a room- and rack-level when adding capacity. These systems typically feature integrated controls and are

built and pretested in a factory environment. In many cases, they can be deployed in as little as 14 to 16 weeks. The shorter deployment time, combined with the “build it as you need approach” offered by modularity, allows colocation providers to quickly meet fluctuating customer demands while lowering capital expenditure (CapEx) costs.

This sort of scalability can come in many forms. Virtually every component of the physical infrastructure is now available in a modular design. Modularity can be applied to specific infrastructure equipment such as UPS systems that can be expanded by adding power cores, or rack power distribution units with swappable power strip receptacles, power cords and power modules. Modularity can also be achieved holistically by quickly adding containerized enclosures to a data center when capacity is being added on a large scale.

The following list highlights the different forms of scalability, beginning at the component level and ending with room-level.

Software-enabled scalability allows data centers to grow capacity of a base system without additional hardware. Certain UPS units in the market provide capacity-on-demand by using software to grow capacity without adding to the system footprint. UPS units with this sort of software functionality can be sized to current requirements and then easily scaled up to a larger capacity with a simple software key as power needs change. These types of modules are scalable in increments, such as from 40 to 60 to 80 kVA, from 80 to 100 to 120 kVA, or from 400 to 500 to 600 kVA.

Hardware-enabled scalability allows quick power capacity increases with the addition of internal power core hardware assemblies. These core assemblies allow the system to expand for capacity or redundancy in 15 kW increments up to 45 or 90kW within a single cabinet. It is important to look for power cores that incorporate distributed intelligence and scalable power in a common assembly and allow configuration of a completely redundant power and control system, sized to match the capacity of the protected equipment. When power requirements change, data center managers can easily add capacity without increasing the system footprint. This approach allows for right-sizing of the UPS, resulting in improved energy efficiency and reduced power expenditures.

Modular busway is a flexible, robust power distribution system that provides high-density, modular busway power distribution to racks while eliminating distribution cable clutter that restricts airflow. The modular design makes it easy to add capacity simply by adding more receptacles and busway sections. It is an ideal solution for data centers of any size that experience frequent configuration changes and require high power distribution and more effective utilization of space. The modular design of the busway can offer 15 to 30 percent savings in time and installation costs compared to traditional cable and conduit solutions.

 **Tweet This:**

“Modular” doesn’t just mean containers- UPS, chillers & aisle designs can enable incremental #datacenter growth
<http://bit.ly/Yu7kny>

SmartAisle™ is an example of how scalability at the technology level can be incorporated into an aisle-configured package that is simple to deploy and cost effective. These contained or open data center solutions integrate modular power systems, cooling systems with Liebert iCOM™ intelligent controls, and a DCIM system. They offer the potential to deliver 30 percent energy savings and 25 percent more capacity than a traditional design, and at high levels of availability. SmartAisle-based integrated infrastructures are ideal for existing facilities experiencing rapid growth.


Modular chiller systems provide end-to-end, chiller-based cooling at a quality level only achievable in a factory environment. They contain all the mechanical components used in a traditionally-built system, including chillers, chilled water pumps, condenser water pumps, cooling towers, variable frequency drives, controls and isolation valves, water treatment, refrigerant monitoring and room-based IT cooling. The modular approach provides initial and operational cost savings by allowing data center managers to add cooling capacity as needed to match growing IT load.

Modular data centers no longer refer to data centers in shipping containers. They have evolved significantly in recent years – integrating power and cooling infrastructures to meet the changing needs of data center managers. These integrated packages are built and pretested in factory environments, and offer a number of variations on the standard “shipping container,” including AC power enclosures and AC power skids. These solutions are optimized for critical IT and infrastructure equipment, and are designed for maximum availability, security, flexibility and efficiency. Many of these “integrated” offerings are particularly cost-effective due to their pre-configured cooling and power infrastructures, tailored to meet the exact specifications of the deployment.

Lower Operating Costs

While modular systems do provide opportunities to lower CapEx costs, there are a number of technologies and approaches colocation providers can implement that will significantly reduce operational expenditure (OpEx) costs. Most of these focus on the operational efficiency of the data center environment.

In traditional data centers, approximately one-half of the energy consumed goes to support IT equipment with the other half used by support systems (Figure 3). Vertiv™ conducted a systematic analysis of data center energy use and the various approaches to reducing it to determine which were most effective. The analysis documented the “cascade effect,” (Figure 4) which occurs as efficiency improvements at the server component level are amplified through reduced demand on support systems. Using this analysis, Vertiv developed Energy Logic, a vendor-neutral roadmap of 10 strategies that can reduce a data center’s energy use by up to 74 percent.



48% Power and Cooling		52% Computing Equipment	
38%	Cooling	15%	Processor
5%	UPS	15%	Other Services
3%	MV Transformer and Switchgear	14%	Server Power Supply
1%	Lighting	4%	Storage
1%	PDU	4%	Communication Equipment

Figure 3: Energy consumption in a typical data center

According to the analysis, in a data center with a PUE of 1.9, 1 W savings at the server processor create a 2.84 W savings at the facility level as a result of the cascade effect. At higher PUEs, the savings are even greater.

For more information on Energy Logic, see Vertiv’s ebook: “Energy Logic 2.0: New Strategies for Cutting Data Center Energy Costs and Boosting Capacity.”

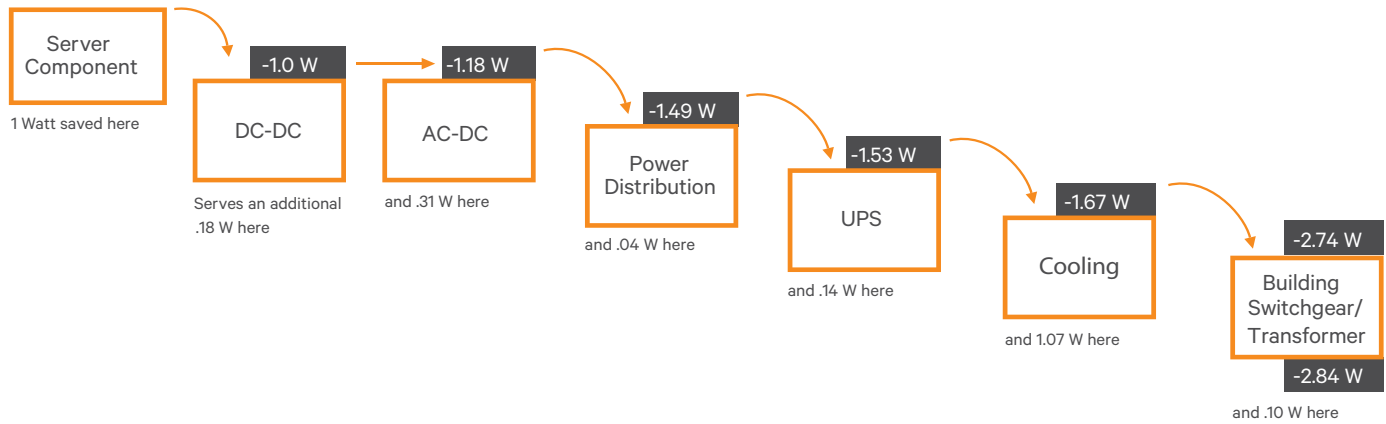


Figure 4: One Watt saved at the server component level results in cumulative savings of about 2.84 Watts in total consumption

Energy Logic - Cooling

A best practice that builds upon the steps in Energy Logic is to optimize the cooling infrastructure. One of the approaches that help accomplish this is the use of “free cooling.” A popular approach in the data center has been economizer systems that use outside air to provide “free-cooling” cycles. This approach has been shown to reduce or eliminate chiller operation or compressor operation in precision cooling units, enabling the economizer systems to generate cooling unit energy savings of 30 to 50 percent, depending on the average temperature and humidity conditions of the site.

However, sometimes the use of economizer systems can potentially introduce new problems. The effect of outside air on data center humidity should be carefully considered when evaluating economization. While there are ways to maintain appropriate humidity levels, they usually offset some of the energy savings provided by the economizer, or achieve reduced energy consumption at the expense of data center availability. Figure 5 shows the pros and cons associated with the two most common forms of economization found in data centers today.

For more information on economizers, see Vertiv™’s white paper: “Economizer Fundamentals: Smart Approaches to Energy-Efficient Free-Cooling for Data Centers”

AIR-SIDE ECONOMIZATION		WATER-SIDE ECONOMIZATION	
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 5: The pros and cons of the two most common approaches for utilizing outside air to increase energy efficiency in data center thermal management.

One way colocation providers can mitigate the operational limitations or concerns commonly associated with air economizer technology is by installing a cooling system that features technology pumped refrigerant economizer in a traditional direct expansion (DX) unit. This type of air-cooled system uses the same refrigerant circuit, coils and condenser in both economizer and non-economizer modes. It employs a two-phase refrigerant versus a traditional single-phase economizer solution, resulting in a simple and efficient cooling system that maximizes the hours available to use reliable free cooling.

Not only does this DX based system utilize cooler outside air to enable “free cooling”, but these types of systems enable a “just-in-time” approach to cooling deployment. There is no need to install a large chiller plant that is sized for full capacity on day one. By deploying DX units as capacities and server cabinet footprints grow, lower capital costs are realized.

The Liebert DSE™ is an example of this type of cooling system. It incorporates an optional “free cooling” technology feature called Liebert® EconoPhase™. At full capacity, the unit is 114 percent more efficient than an air economizer, and mitigates the operational limitations or concerns commonly associated with air economizer technologies. The Liebert DSE offers an energy-efficient SCOP (seasonal coefficient of performance) rating of 2.9, making it 52 percent more efficient than the ASHRAE 90.1 (2010) minimum requirement of 1.9 for data center cooling units. For more information on the Liebert DSE, visit: VertivCo.com/LiebertDSE

Energy Logic - Power

Another best practice that builds upon the Energy Logic steps is to utilize double-conversion UPS systems that feature an economy mode of operation (eco-mode) to increase energy efficiency and reduce the operating expenses experienced in the data center.

Eco-mode allows UPS systems to achieve the highest possible efficiencies by transferring the UPS to operate on utility power when conditions are ideal. When power problems are detected, the UPS switches back to double conversion mode. Until recently, this transfer was not seamless because when the critical load was being powered through the bypass, the rectifier and inverter were switched off. This resulted in a delay and a notch in the output waveform when the critical load returned to double-conversion mode.

A new approach, active inverter intelligent eco-mode (Figure 6), eliminates this issue by keeping the inverter and rectifier in an active state, ready to accept the load nearly seamlessly when transfer occurs.

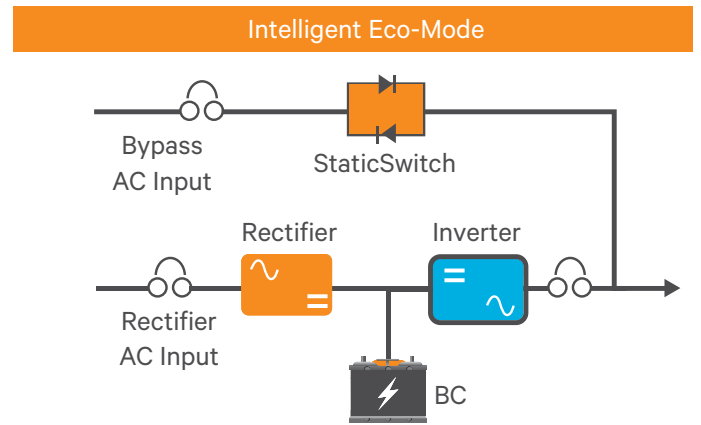


Figure 6: Active inverter intelligent eco-mode offers a more seamless transfer and UPS efficiency level up to 99 percent.

The use of eco-mode has been gaining support in the industry. The Green Grid, a global consortium of companies, government agencies and educational institutions dedicated to advancing resource efficiency in data centers and business computing ecosystems, included eco-mode in its Data Center Maturity Model. The Environmental Protection Agency (EPA) also mentioned eco-mode in its recently introduced ENERGY STAR® for Uninterruptible Power Supplies specifications as one of the operating modes for efficiency improvements.

Greater Reliability

Reliability and availability concerns are very familiar to colocation providers. They basically come down to economics. The increasing reliance on IT systems to support business-critical applications has forged an even stronger connection between data center availability and total cost of ownership (TCO). One significant outage can be so costly that it wipes out years of savings achieved through incremental efficiency improvements, and can severely impact a colocation provider's reputation and bottom line.

The actual costs associated with an unplanned outage are stunning. A 2011 Ponemon Institute study, sponsored by Vertiv™, revealed that the mean cost of a data center outage downtime event is \$505,502, with the average cost of a partial data center shutdown running \$258,149 and a full shutdown costing more than \$680,000.

For more information on the causes and cost of downtime, see Vertiv's white papers:

- "Understanding the Cost of Data Center Downtime"
- "Addressing the Leading Root Cause of Downtime."

There are a number of ways colocation providers can achieve greater reliability and meet customer expectations for 24/7 availability and accessibility. One very practical step is to utilize service expertise and regularly scheduled maintenance to extend equipment life.

For established facilities, **preventive maintenance** has proven to increase system reliability. Vertiv 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 Failure (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.

Tweet This:

Avg. cost of a
#datacenter
outage = \$505,502;
partial shutdown =
\$258,149;
full shutdown =
\$680,000
<http://bit.ly/Yu7kny>

At the UPS level, dedicated **battery monitoring** is critical to preventing outages. According to the Ponemon study, UPS battery failure is the leading cause of downtime, with 65 percent of respondents citing it as the cause of their unplanned data center outages (Figure 7). 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.

Deploying a **data center infrastructure management** (DCIM) platform further optimizes availability. DCIM technologies enable the unified management of the data center’s IT and facilities infrastructure, giving colocation providers unprecedented insight and action to achieve optimal data center performance and easily add capacity to meet fluctuating customer capacity demands.

From a single, unified, dynamic platform, colocation providers are able to monitor and control the critical infrastructure (IT and facility equipment), manage capacity, track inventory, plan changes, visualize configurations, analyze and calculate energy usage, optimize cooling and power equipment and account for the impact of virtualization on the equipment supporting it. Users are able to manage the system as a whole, instead of in parts, resulting in higher utilization of existing infrastructure, increased personnel efficiency and lower operating costs.




Figure 7: Respondents to the Ponemon Institute survey most frequently cited UPS battery failure as the cause of unplanned data center outages.

Conclusion

As data center managers struggle to keep pace with growing capacity needs while working under the constraints of tightened budgets and energy efficiency initiatives, many are considering third-party colocation facilities as an alternative to building out their own environments.

In order for colocation facilities to capitalize on this increased interest and customer demand, it is important that they deliver on their value proposition of doing it “cheaper” and “faster” while delivering better reliability than an organization can deliver on its own. To accomplish this, they need a sophisticated, scalable data center infrastructure that will enable them to quickly add customers, while optimizing efficiency at every stage of infrastructure utilization.

There are a number of new and proven data center infrastructure technologies available that advance the flexibility of data center environments, increase efficiency and safeguard availability. These technologies will help colocation providers meet their value proposition by enabling them to deploy faster at a lower cost and with greater reliability.



VertivCo.com | Vertiv Headquarters, 1050 Dearborn Drive, Columbus, OH, 43085, USA

© 2016 Vertiv Co. All rights reserved. Vertiv and the Vertiv logo are trademarks or registered trademarks of Vertiv Co. All other names and logos referred to are trade names, trademarks or registered trademarks of their respective owners. While every precaution has been taken to ensure accuracy and completeness herein, Vertiv Co. assumes no responsibility, and disclaims all liability, for damages resulting from use of this information or for any errors or omissions. Specifications are subject to change without notice.