CompuDynamics™
UPTIME IN NO TIME

White Paper: Essential Elements for Data Center Optimization

What Every Government and Enterprise Facility Operator Should Know

By Dan Ephraim, Director of Sales and Marketing, Compu Dynamics

Prologue

Data center managers, regardless if they are in the private or public sector, must become more aware of the applications running in their facilities, as well as the hardware supporting those applications. Critical utility input demands can provide data center managers with much of the information they need to successfully forecast and right-size their power and cooling infrastructure, and thereby operate a truly efficient and optimized facility. No data center optimization strategy can succeed without a comprehensive and precise understanding of the system performance requirements.

Overview

Given the widespread adoption of cloud computing, the explosive growth of Over-the-Top (OTT) content and the increased use of virtualization in both federal and large enterprise networks' compute and storage infrastructure, the days of the traditional data center — with its unfettered waste of electricity, and tape measure and spreadsheet-driven thermal management and capacity planning — are long gone. This white paper will address some of the essential elements for the optimization of the data center environment, examining issues concerning airflow management, Data Center Infrastructure Management (DCIM) tools, critical power management, Power Usage Effectiveness (PUE) and critical power capacity matching, in addition to operational best practices.

New federal initiatives and regulations, including the Data Center Optimization Initiative (DCOI), which mandates the planning and ongoing measurement of progress towards efficient facility and IT operations, are also discussed in great detail in this white paper. While this directive affects only government-run facilities, enterprise data centers as well as colocation and cloud computing environments have also become singularly focused on reducing waste in their IT infrastructure and underlying platforms. With the advent of the DCOI and associated PUE mandates, reconfiguration and redesign of the facility may be the only option for government data centers.

According to the National Resources Defense Council (NRDC), data center electricity consumption is projected to increase to roughly 140 billion kilowatt-hours annually by 2020, the equivalent output of 50 power plants, costing American businesses $13 billion in electricity bills and emitting nearly 100 million metric tons of carbon pollution per year. The move toward environmentally conscious, fiscally responsible data center and cloud computing facilities will have a widespread and transformative effect on the industry. Data center managers and facility administrators that adopt sustainability practices, and reduce energy consumption and control costs will survive and thrive — and those that do not will likely see their businesses eventually shutter.

Government Run Data Centers

From a government agency perspective, recent changes in procurement regulations and green IT initiatives are driving federal data centers to the same automation goals as in the private sector. With regulations forcing government-run facilities to replace manual systems with Data Center Infrastructure Management (DCIM) solutions, they will be obligated to measure and monitor progress toward efficiency goals outlined in the recent DCOI initiatives.

While we will discuss recent directives and executive orders around the use of cloud and data centers by the federal government, it's important to recognize that punitive measures are in place, as funding for data center projects that fall outside of complying with the DCOI will not be available. While the highest echelon of government has its eye keenly focused on how generals and other senior agency leadership will address eradicating government-run data center waste, without their support, it is unlikely things will change. Although leadership has turned a blind eye for years, it's apparent that a tipping point is near, as more and more federal data centers are past end-of-life and these inefficient facilities continue to support mission-critical applications. All it will take is one major outage and it is realistic to expect that most of the decision-making leadership will hop on board and put their support behind developing a data center strategy.

That said, let's look at the government data center mandates in detail.

In 2010, the U.S. Office of Management and Budget (U.S. OMB) launched the Federal Data Center Consolidation Initiative (FDCCI), whose goals are the following:

- Promote the use of green IT through reduction of overall energy and real estate footprint of government data centers
- Reduce the cost of data center hardware, software and operations
- Increase the overall IT security posture of the Federal government
- Shift IT investments to more efficient computing platforms and technologies
- Implementation of data center management strategies

In March 2016, the U.S. Government's Executive Order 13693: "Planning for Federal Sustainability in the Next Decade," required all federal agencies to install and monitor advanced energy meters in all government data centers by September 30, 2018. The OMB will monitor the energy efficiency of data centers through PUE metric energy metering tools that will enable the active tracking of PUE for the data center, and which must be installed in all tiered federal data centers by September 30, 2018.

If you have been living in a cave for the past 10 years, PUE is a metric used to manage data center energy efficiency and is determined by dividing the amount of power entering a data center by the amount of power used to run just the computer infrastructure within it.

$$\text{PUE} = \text{Total Facility Energy} \div \text{IT Equipment Energy}$$

It is a measure of how much energy is used by the computing equipment in contrast to cooling and other overhead. PUE is expressed as a ratio, with overall efficiency improving as the quotient decreases toward 1. While this metric may seem to provide a way to compare data center efficiencies, there are some flaws that we would like to point out, including:

- Geographic variations in temperature
- Failing to account for lighting and other non-IT loads, such as the Network Operations Center (NOC)
- Frequently, the IT load values used are the rated power load, as opposed to actual operational consumption

In addition, some subsystems support a mixed-use facility and are shared with other non-data center functions, for example, cooling towers and chiller plants. Hence, fractions of the power attributable to the data center cannot be directly measured.

Federal Data Center Operators: The Mandate to Transition to Cloud and Data Center Shared Services, and Improve PUE

The ramifications of DCOI and FDCCI and the Federal Information Technology Acquisition Reform Act (FITARA) are clear. The federal government wants to, where possible, exit the data center business and reduce IT operational expenses. According to these policies, government agencies may not allocate funds or resources towards the initiation of a new data center or the expansion of an existing one without receiving Office of the Chief Information Officer (OFCIO) approval. To gain approval, agencies are required to submit a written report on alternatives, including cloud services, shared services or third-party colocation. FITARA enacts and builds upon the requirements of the FDCCI, requiring that agencies submit annual reports that include comprehensive data center inventories; multi-year strategies to consolidate and optimize data centers; performance metrics and a timeline for agency activities; and yearly calculations of investment and cost savings.

Government-run data centers are also being tasked with reducing application and database loads, migrating to Software-as-a-Service (SaaS) and Platform-as-a-Service (PaaS) solutions where possible, as well as increasing the use of virtualization in order to enable pooling of compute, storage and network resources to allocate on-demand. In short, where the private sector data center is migrating, the public sector facility must move as well.

The Presidential Executive Order 13693 relates to the requirement of government-run data centers to collect and report energy usage. Through the use of the PUE metric, the OMB will monitor the energy efficiency of data centers with the requirement that existing data centers achieve and maintain a PUE of less than 1.5 by September 30, 2018. The requirements state that all new data centers must implement advanced energy metering and be designed and operated to maintain a PUE no greater than 1.4, and are encouraged to be designed and operated to achieve a PUE no greater than 1.2.

To achieve these requirements, very efficient electrical and HVAC equipment and distribution are required. Approaches include the use of high-efficiency Uninterruptible Power Supply (UPS) systems and transformers, water- or air-side economizers, high-temperature chilled-water systems, rack- or aisle-level containment systems, direct water-cooled IT equipment, and containerized solutions.

The Order also states that for existing data centers in which a PUE target of less than 1.5 is not cost-effective, agencies must look to the 'Four C's': cloudify, collocate, consolidate, or close to determine a solution. Agencies must include PUE requirements for all new data center contracts or procurement vehicles. Furthermore, any new data center contract or procurement vehicle must require the contractor to report the quarterly average PUE of the contracted facility to the contracting agency, except where that data center's PUE is already being reported directly to the OMB or General Services Administration (GSA) through participation in a multi-agency service program. Interestingly enough, PUE reporting is not required for cloud services.

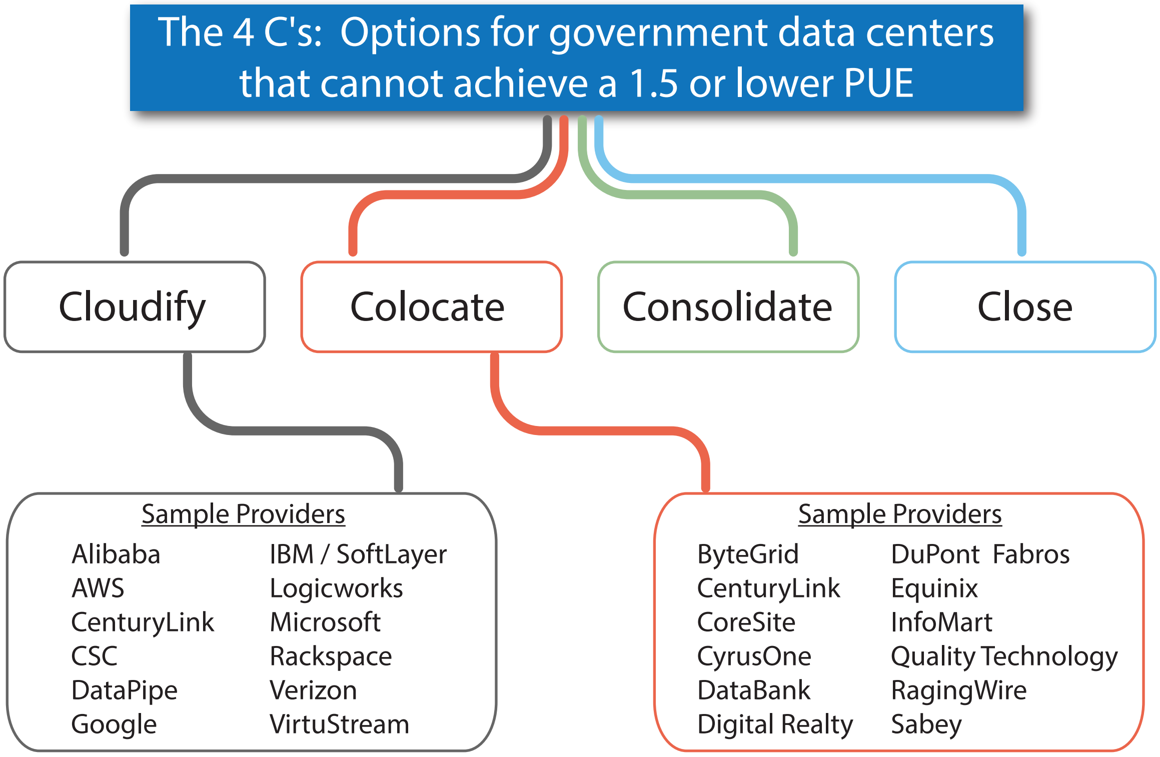

With respect to PUE, clearly a more accurate, uniform method of deriving the PUE metric must be standardized and implemented. A more detailed examination of the issues concerning PUE is included later in this white paper.

> *"We build big boxes that consume extraordinary amounts of power and we aggressively seek out new technologies to manage and deploy that power effectively. That said, if you were to ask me what most efficiency conscious data center managers do? They shut off the lights every time they leave a room. Efficiency always has to start with common sense. Focusing on the basics will get you more energy savings than you think."*
>
> *-Senior Executive at a Global Colocation Company*

Data Center Optimization Essentials

Whether you manage a government facility, enterprise, colocation or cloud computing environment, there are a number of practices and methodologies that are paramount to minimizing risk and capturing the full value of your data center deployment. What follows is a survey of the essential elements of data center optimization.

Airflow Management

Energy is the major driver of cost in today's data center environment, driven by the demands of today's applications that require higher compute levels and by virtualization's capability of extracting the most value from the available processing power. Therefore, today's data centers require a more effective airflow management solution, especially as equipment power densities increase. According to Data Center Knowledge, five years ago, the average rack power density was one to two the kilowatt (kW) per rack, and today, the average power density is four to eight kW per rack, while some data centers that run high-density applications are averaging 10 to 20 kW per rack. Many private and public sector data center managers are combating this increase by following the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) guidelines and running the data center hotter. However, raising the temperature within the facility will only get a data center manager so far. Hence, they must start looking at new and affordable technologies while adhering to best practices and more stringent PUE standards.

Regardless if the data center is on slab or raised floor, the ability to isolate and redirect the hot air from the IT cabinet has transformed the industry. By effectively managing the airflow via some form of containment, data center cooling systems can operate closer to their peak performance capabilities, and thereby, do more work for the amount of energy consumed. Containment has rapidly become the most cost effective technique to address growing energy costs as it can be deployed in both new and legacy facilities and there is a highly competitive marketplace with multiple vendors. Implementing containment within a data center enables cooling equipment to operate closer to its manufactured capabilities.

Hot aisle and cold aisle containment remain widely accepted best practices. The basic concept of aisle containment focuses on the separation of the server inlet cold air and its hot exhaust air. In this architecture, cabinets are arranged into rows in a back-to-back and front-to-front configuration. This design concentrates cool air at the front of the racks, ensuring that air is freely delivered to the server cabinet intakes. As the air moves through the servers, it carries away the heat generated by the equipment back into the hot aisle. This hot exhaust air is then routed back to the cooling equipment where the heat may be transferred to the outside. All of these components must work in concert in order to achieve and maintain the optimal performance of the data center.

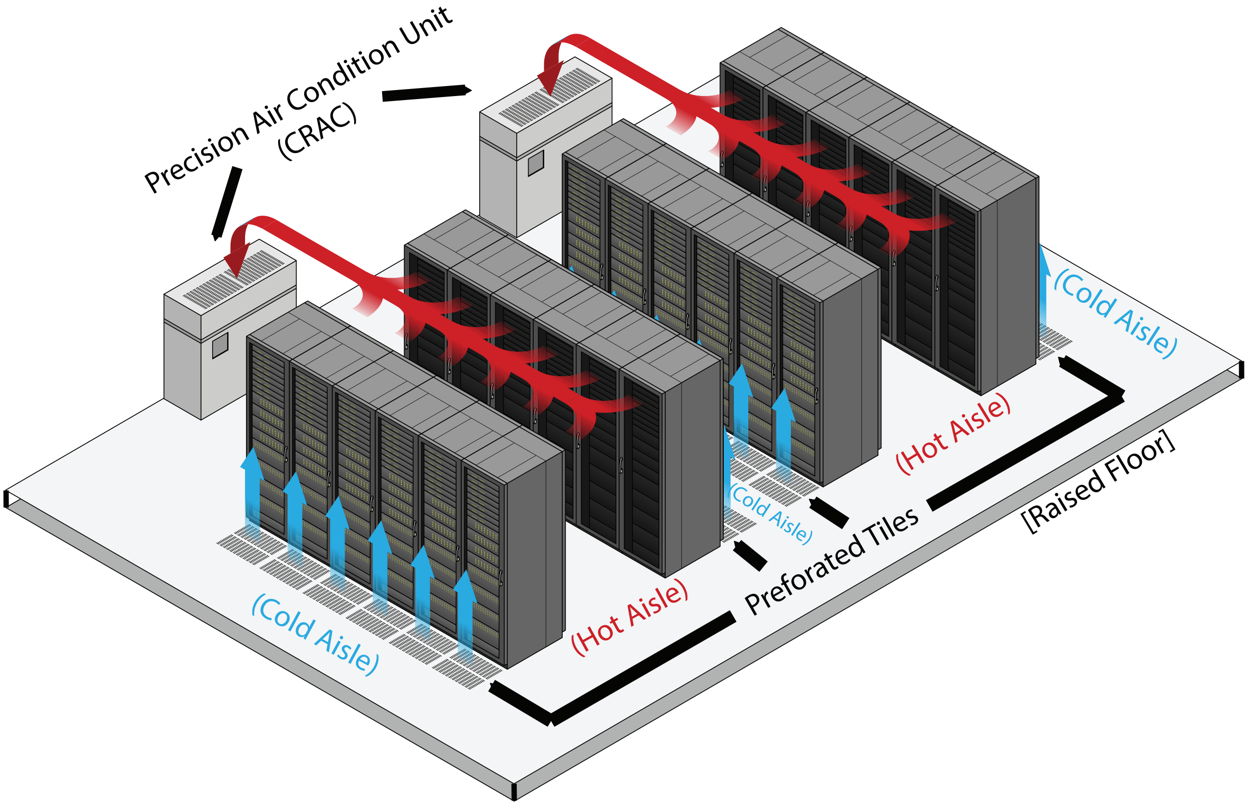

Containment Approach Advantages

Using a containment design within an existing facility, the data center makes increasingly efficient use of the same or less cooling, reducing the cooling portion of the energy bill and thus reducing operating costs. Containment also allows for lower cooling unit fan speeds, higher chilled water temperatures and decommissioning of redundant cooling units, resulting in a lower Total Cost of Ownership (TCO). According to the U.S. Environmental Protection Agency (U.S. EPA), a robust containment solution can reduce fan energy consumption by up to 25 percent and deliver 20 percent energy savings.

Other advantages of containment include:

- The ability to support racks running at higher power densities
- A reduction in the cooling capacity-to-power ratio closer to a one-to-one match by minimizing the latent cooling requirements
- Achievement of savings through annual utility bill reduction without significant additional CapEx
- Reduced repair costs due to the reduction in workload on the existing cooling equipment

In order to implement an effective containment architecture, many of the following components are often utilized in concert to contribute to more efficient flow of cold and hot air within the environment:

- Rack-based chimneys and CRAH / CRAC Plenums
- End-of-row aisle doors
- Aisle ceilings or overhead vertical wall systems
- Perforated floor tiles in the appropriate locations
- Floor grommets to create a tight seal around through-floor cables and fittings
- Blanking panels to close empty rack slots
- Rearrangement and redistribution of racks or servers to a hot aisle / cold aisle configuration
- Removal of unneeded obstructions hindering underfloor air flow

Challenges of Hot Aisle and Cold Aisle Containment

There are multiple ways to contain hot and cold air and each containment method has its own benefits and challenges. When hot aisles are contained, the rest of the data center is left to serve as the cold aisle. Depending on the supply air temperature, this may create a data center with an expensive, uncomfortably cold environment.

One issue to keep in mind regarding containment is the operating temperature of the hot aisle. While it can be a significant impact on cost savings, requiring trained technicians to work in a space that is greater than 140 degrees Fahrenheit eventually becomes a workplace safety issue that should be addressed prior to any deployment.

When cold aisle containment is utilized, the balance of the room outside of the cold aisles may be far too warm to be comfortable. Operators entering a data center may get the mistaken impression that a cooling problem exists, when in reality, the servers are and cool in their protected space.

Data Center Infrastructure Management (DCIM)

DCIM tools monitor, measure, manage and / or control data center utilization and energy consumption of all IT-related equipment, including servers, storage and network switches, as well as facility infrastructure components, including Power Distribution Units (PDUs) and CRAC units.

There are several factors driving the demand and adoption of DCIM, primarily those around power management, asset management and technology advancements. In tandem, these facets of operation contribute to the profitable and efficient management of the data center along a number of dimensions, such as:

- Capacity Planning and Management: The ability to quickly model and allocate space for new servers, and manage power and network connectivity
- Asset Lifecycle Management: Centralized databases allow more accurate record keeping and more efficient processes
- Change Management: Fully integrated workflow management, including automation of work orders and workflow activities
- Measurement of uptime and availability
- With the advent of virtualization and cloud computing, it provides the ability to manage the physical machines containing the virtualized environments
- Power Management: Constant monitoring with alerts before circuits fail; this includes the ability to locate stranded capacity to avoid costly build outs, and intelligent PUE analytics and reporting tools

DCIM is highly relevant to the DCOI initiative given the requirement to provide the OMB with a strategic plan detailing energy consumption and forecasted savings. According to Gartner, utilizing a DCIM can lead to energy savings that reduce a data center's total operating expenses by up to 20 percent. DCIM tools have the ability to integrate live data to monitor power consumption and temperature, allowing the manager to see the current state of the data center. Further, since cooling is based on design capacity of a rack, and this varies greatly from rack to rack, there is frequently an over-prediction of cooling needs resulting in a worst-case design and therefore overbuilt HVAC gear and the expense that goes with it.

As a word of caution, DCIM tools are expensive and while it is recognized as a powerful and transformative technology, more often than not deployments are difficult and the business benefits are challenging to measure and ultimately quantify. Given that calculating an ROI is a challenge, it remains a technology platform that is not as commonly deployed as one would expect. For example, in smaller federal and enterprise data centers, the ROI on a DCIM solution is extremely challenging given the prohibitive cost of a DCIM.

Reconfigure and Redesign

Most federal and enterprise data center managers have to go beyond containment and deploy a DCIM to capture a sub-1.5 PUE. When that occurs, firms are left with few options and they all fall under the umbrella of redesign or reconfigure. The challenge any data center manager faces when reconfiguring is that budgets are not unlimited. Therefore, in considering the investment, not only should there be a desired ROI but annual maintenance costs must also be calculated into the equation. There are three primary categories of a redesign project: equipment, floor plan zoning, and critical power capacity matching.

Efficient and Agile Equipment

As we move towards new technologies to identify efficiencies, the basic building blocks of the data center remain the same: racks, server / storage devices, power, cooling and connectivity. We will focus on two of the five core elements — power and cooling.

Cooling

Free cooling and economization are approaches to realizing significant efficiencies and a rapid shift in PUE. Most all new data centers are designed to leverage free cooling but legacy facilities face an uphill battle, since installing and commissioning new equipment could cause outages detrimental to the business. As an additional risk, justification of capital investment is a challenge given that legacy facilities are difficult to monitor with precision, resulting in poor ROI estimates.

One category of efficiency improvement solutions that falls outside of free cooling are those that bring water to the cabinet. Firms such as CoolIT Systems, now Stulz, and LiquidCool Solutions are manufacturing solutions that allow for a concentrated and targeted area where water can absorb the exceptional thermal temperatures produced by the IT hardware. The OpEx and CapEx savings can be significant as much of the electricity to run the CRAC or CRAH units can be saved. Water-to-the-rack has been slow to be adopted because of the legacy anxiety of bringing water into the data center. That being said, the behavior of industry leading colocation firms like Digital Realty (NYSE: DLR), RagingWire (a subsidiary of NTT Communications), Equinix (NYSE: EQIX), and DuPont Fabros (NYSE: DFT) to name a few, all have the capability to bring water to the cabinet at their new purpose-built facilities. As global leaders in managing and maintaining data center facilities, their decision to offer this option to their Fortune 1000 clients reflects their confidence that this a direction towards which the industry is moving, although this is a solution that has been slow to be adopted.

Power

Minimal innovation is coming in the form of the physical equipment that brings power into the data center, as we all have to work within the laws of physics. The area where this segment of the industry is transforming is in the solutions that Gartner, the information technology research and advisory company, classifies as Software Defined Power (SDP). Software can script tasks and activities from analyzing power quality, to switching loads, to legacy consumption for forecasting purposes. Ideally, the software would also be able to track energy sources to identify if the power is coming from a renewable energy source and then quantify tax credits or apply for one of the many federal energy rebate programs. What is most appealing about a software solution of this sort is that it has the potential to be cloud-based and interdisciplinary, collecting facility data along with data from the IT environment.

The end game of SDP is that controls can be put in place to monitor electrical usage and alert the data center facility manager when waste begins to occur. The beauty of this approach is that, should the systems be monitored by a discipline and/or function, alerts can be mapped to departments and accountability can be assigned to both people and budgets. This type of transparency will give the data center manager the information necessary to extract waste from the facility's total consumption.

Efficient and Agile IT

Efficient and agile IT is just as much about efficiency as it is about risk. The data center manager, while incentivized to deploy efficient technologies, must also educate their CFO as to the financial and operational benefits of the technology to justify the investment.

The illustration represents an ineffective usage of zones, with hot spots indicated in red and cool spots indicated in blue, scattered throughout the data center and resulting in an inefficient facility.

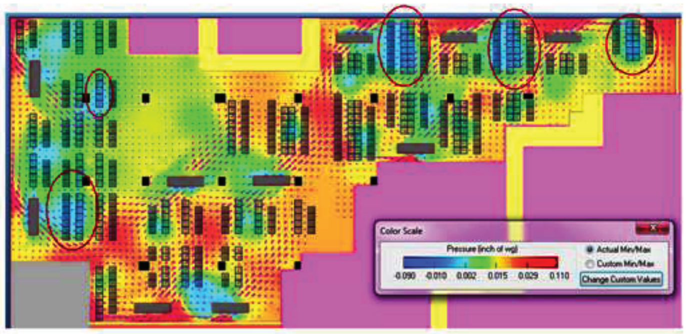

Today's data center managers can rejoice in that for several decades facility professionals have pleaded for manufacturers to provide solutions that extend the life of equipment while also improving efficiency. As a result, there is a robust market of electrical and mechanical solutions to leverage as a vehicle to reduce PUE.

Floor Plan Layout and Data Center Zoning

The same way local municipalities have a zoning board to control all matters relative to community planning, a data center should be planned with the same discipline. Given that once racks are deployed and applications are live supporting active internal and external clients along with partners, it becomes a virtual impossibility to reconfigure the space and recover losses.

In a utopian world, the data center floorplan is designed in concert when the data center facility itself is also being designed. In doing so, the data center team is essentially forced to elevate their understanding of the IT platform and what compute the facility will be housing. By understanding the end-user's requirements, the devices the IT group will need are selected and thus give the data center manager and designer the baseline tools they need to create zones. Zones are ideally designed to create a concentration of racks that are running at similar densities, thereby improving the efficiency of the cooling platform.

The essentials of laying out an efficient floor plan include:

- Understanding the compute platform that the IT environment will be housing and the storage and server infrastructure being used
- Controlling the airflow using a hot-aisle / cold-aisle rack layout
- Aligning all of the racks with the floor and ceiling tile systems providing future airflow optionality
- Minimizing areas of stranded cabinets housing no active compute
- Updating the floor plan annually and, each year, creating a two-year plan.

In working with the head of IT capacity planning at a Fortune 50 pharmaceutical firm, their discipline of tracking and forecasting, without constraining the business, became a game-changer for their organization. s they began to deploy 2.4MW across 20,000-square feet, the capacity team was able to model out a floor

CompuDynamics™
UPTIME IN NO TIME

plan that took into account the day one requirement, while also leveraging historical data to forecast what their layout will look like two-to-five years down the road. This level of attention to detail has enabled this enterprise to optimize their investment on day one and also mitigate the risk of future inefficiencies.

Critical Power Capacity Matching

While the electrical and mechanical technologies are becoming more efficient, the game is changing because of the availability of modular solutions. 451 Research, the technology research and advisory company, classifies these pre-engineered and pre-built technologies as Prefabricated Modular (PFM) designs. Whether built in a factory or in the field, modular designs will drive down build costs, making investments in data center solutions easier to rationalize on a balance sheet. With the technology more in line with the optimal buying process, modular building blocks simplify the scaling of power and cooling for the modern data center.

Ultimately, the capacity of the power / cooling systems should be matched to the load. This is more complicated than it sounds when the load is somewhat hard to predict from one technology shift to another. Nevertheless, if one can utilize scalable chunks of capacity in a just-in-time manner, one can realize energy savings and improved PUE performance.

Wrapping Up

Federal, enterprise and collocation data center managers are all dealing with rules and regulations intended to remove inefficiencies from their respective processes. That should come as no surprise. It is refreshing to see teeth behind these rules and regulations in the form of reducing and eliminating budgets to full-on closure of the data center. These measures are extreme but we need behavior to change. As the British inventor and industrial designer James Dyson says, "Fear is always a good motivator."

That being said, the economic benefits of optimization are something Compu Dynamics would like to see, and one vehicle to achieve that is for the capacity planning function to perform a greater role in the organizational decision-making process. As applications become more robust and require unique physical architecture, we would also like to see a capacity planning function become more valuable to the business strategy, given that this function will be a liaison between the IT and facilities teams. While this white paper stresses the importance of the data center manager understanding the applications that reside in the data center, there is also a burden on the IT manager to understand the principles of data center design and management. The beauty of these roles evolving is that there is not just the single benefit to the agency or enterprise, but the employee also becomes more versatile, opening a multitude of career opportunities.

With mechanical and electrical elements of the data center so intertwined, leveraging a design, installation and ongoing maintenance partner with both mechanical and electrical skill sets is optimal. Engaging a contractor with competencies in both data center electrical and mechanical work reduces risk because the complexity of day-to-day operations.

Hundreds, if not thousands, of government and enterprise facilities built over 10 years ago in the U.S. require attention and the prioritization of activities to receive the optimal ROI of any investment. While these facilities are dated, Compu Dynamics' methodology around optimization will extend the lifetime value of these investments.

About Compu Dynamics

Compu Dynamics is the leading provider of facility maintenance and data center design, deployment, optimization, and ongoing operational support. Compu Dynamics delivers solutions to help businesses achieve their optimal goals by leveraging specialized technicians and technologies to minimize risk of data center and facility outages. Headquartered in Northern Virginia, with operations in both Virginia and Maryland, Compu Dynamics is positioned to drive value for government agencies, enterprises, and third-party colocation providers of any size. Additional information can be found at www.compu-dynamics.com, or by following Compu Dynamics on Twitter and LinkedIn.