Data Center Physical Infrastructure: Optimizing Business Value

White Paper 117

Revision 1

by Wendy Torell

> Executive summary

To stay competitive in today's rapidly changing business world, companies must update the way they view the value of their investment in data center physical infrastructure (DCPI). No longer are simply availability and upfront cost sufficient to make adequate business decisions. Agility, or business flexibility, and low total cost of ownership have become equally important to companies that will succeed in a changing global marketplace.

Contents

Click on a section to jump to it

P3
Power Protection Products, Inc.

877-393-1223 | www.p3-inc.com

Schneider Electric

Introduction

Data center physical infrastructure (DCPI) is the foundation upon which Information Technology (IT) and telecommunication networks reside. It is the "backbone" of the business, as its elements provide the power, cooling, physical housing, security, fire protection, and cabling which allow the Information Technology to function. **Figure 1** illustrates these critical DCPI elements and their integration into a seamless end-to-end system supported by management systems and services. Viewing DCPI as a whole rather than as individual components is essential to designing and deploying an integrated, understandable system that performs as expected. When individual elements are purchased in isolation from other DCPI elements, the end result is typically a complex and unpredictable DCPI system made up of multiple vendors' products that haven't been designed to work together. Management becomes more complex because a variety of management systems must be used to provide visibility to the entire system, and multiple service contracts become necessary. **Appendix A** of this paper provides more detailed descriptions of each of these DCPI elements.
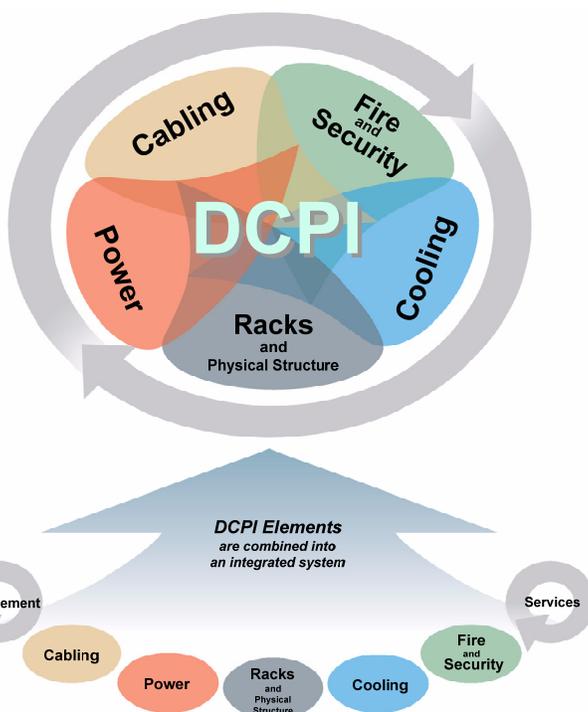
Figure 1

An integrated DCPI system

Most successful businesses today depend on a stable information technology platform. In order to maintain these IT business operations, four layers or building blocks must be present. **Figure 2** illustrates this layered model and the importance of having an integrated DCPI system as a foundation to maintaining business functions. In addition to the DCPI, this includes the *information technology* and the *processes* and *people* to support the operation of these systems. The information technology includes data processing, storage, and communications systems, both hardware and software. Without appropriate planning and design of this technology, the network and ultimately the business cannot function. All processes for operating in this data center or IT environment must be clearly defined, well documented, and standardized in a simple manner for all users to comprehend. When such processes are not effectively implemented, inconsistencies in the operation and maintenance of systems are inevitable leading to unexpected downtime. It is also necessary to have the people to support the operations. This includes having the appropriate staffing level and the right level of skill and training. When proper planning for staffing levels and training / skill levels has not occurred, human error is inevitable.

Figure 2

DCPI is a critical layer of reliable IT business operations
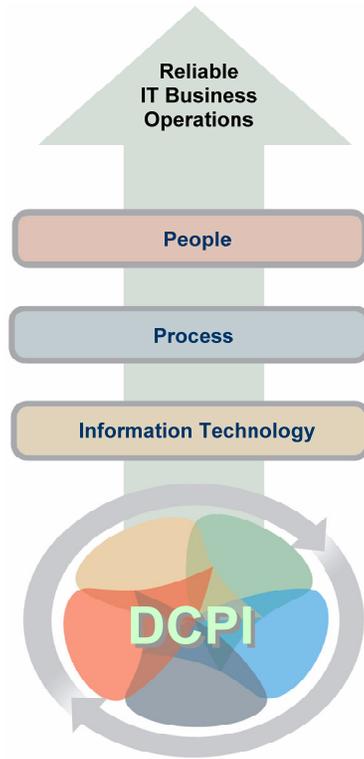
← *With inadequate staffing level or skill set, business success becomes very difficult*

← *The staff's success depends greatly on having well defined, documented processes*

← *Processes and people are dependent on stable IT to achieve business objectives*

← *The foundation layer, upon which everything else resides – without reliable DCPI, the information technology, processes, and people will fail*

This paper discusses a key shift in the way IT planners must think about business value as they make DCPI investment decisions, and presents the drivers that are dictating this new set of value criteria.

Optimizing DCPI to improve its Business Value

Business value for an organization, in general terms, is based on three core objectives:

- Increasing revenue
- Reducing costs
- Better utilizing assets

Regardless of the line of business, these three objectives ultimately lead to improved earnings and cash flow. DCPI investments are made because they both directly and indirectly impact these three business objectives. Managers purchase items such as generators, air conditioners, security systems, and UPS systems to serve as "insurance policies." For any network or data center, there are risks of downtime from power and thermal problems, and investing in DCPI mitigates these and other risks. So how does this impact the three core business objectives above (revenue, cost, and assets)? Revenue streams are slowed or stopped, business costs / expenses are incurred, and assets are underutilized or unproductive *when systems are down*. Therefore, the more effective DCPI is in reducing downtime from any cause, the more value it has to the business in meeting all three objectives.

Historically, assessment of DCPI business value was based on two core criteria: availability and upfront cost. Increasing the **availability** (uptime) of the DCPI system and ultimately of the business processes allows a business to continue to bring in revenues and better optimize the use (or productivity) of assets. Imagine a credit card processing company

whose systems are unavailable – credit card purchases cannot be processed, halting the revenue stream for the duration of the downtime. In addition, employees are not able to be productive without their systems online. Minimizing the **upfront cost** of the DCPI also results in a greater return on that investment. If the DCPI cost is low and the risk / cost of downtime is high, the business case becomes easier to justify.

While these arguments still hold true, today's rapidly changing IT environments are dictating two additional criteria for assessing DCPI business value. First, business plans must be flexible to deal with changing market conditions, opportunities, and environmental factors. Investments that lock resources limit the ability to respond in a flexible manner. And when this flexibility or agility is not present, lost opportunity is the predictable result.

A second additional business value criterion that must be considered is the **Total Cost of Ownership (TCO)**. While upfront cost is still a very relevant factor, it simply does not tell enough of the story. It leaves the decision maker in the dark about long term costs of a solution, including operating and maintenance costs. Upfront cost was often used as the criterion because of a traditional project approach to purchasing capital items. Capital costs, for tax and depreciation reasons, were often separated from on-going expense costs in the accounting justification for a project. Even though upfront cost is only a fraction of the TCO, it was enough to get a project approved and purchased - it allowed the project to commence; And things like the electric bill were not part of the decision making process. Those types of expenses were often viewed as "fact of life" costs that were simply a necessary evil – they just came along with the project and were paid for not with project money, but with operational funding. Business decision makers are now seeing the importance of including these other costs in business value assessment when making critical business choices. A good business investment decision includes consideration of both first costs and on-going operational costs.

Figure 3 illustrates the shift in thinking about the criteria to assess DCPI business value from the slower-paced businesses of the past to the rapidly changing businesses of today. The purpose of the equation is not to provide "quantitative value." Rather, it is intended to highlight the factors that must be considered to achieve high business value. Availability and agility, which are in the numerator or top of the equation, must be *maximized* to increase business value. TCO, which is in the denominator or bottom of the equation, must be *minimized* to increase business value.

Throughout this paper, the data center manager will be challenged to think about availability in a new way, think about cost in a new way, and think about a new performance vector, agility, which simply cannot be ignored in today's business world. All of these performance vectors ultimately translate into dollars, and it is necessary to consider how best to optimize "agile availability per TCO dollar."

Figure 3

Paradigm shift of DCPI business value criteria

$$\text{Value} = \frac{\text{Availability}}{\text{Upfront cost}}$$

The Paradigm Shift

$$\text{Value} = \frac{\text{Availability} \times \text{Agility}}{\text{Total cost of ownership}}$$

When considering this new business value "equation" for DCPI, there are many contributors to each of the three criteria. The following sections discuss those contributors that have the greatest impact on overall business value.

Availability – performance vector #1

As discussed, availability is a key driver of business value for DCPI. It is the main reason that DCPI purchases are made. If systems are not available, essential business objectives are sacrificed.

Availability is a term that has taken on a lot of different interpretations by a lot of different people. Some think of it as simply equipment reliability. But there's a lot more to it. The technical definition of availability is:

Availability _is the degree to which a system or component is operational and accessible when required for use [IEEE 90]._

Equipment reliability is certainly one variable that contributes to "systems and components remaining operational," but other factors such as Mean Time to Recover (MTTR) and human factors also play a significant role. Systems go down for many reasons, both planned and unplanned. An example of planned downtime is downtime for regularly scheduled preventive maintenance. Examples of unplanned downtime include human error and natural disasters. Surprisingly, natural disasters account for significantly less downtime than human error caused during day to day operations. Industry experts estimate that human factors cause unplanned downtime 40-60% of the time.

The requirement for availability is expressed in a variety of ways, some more quantitative than others. According to the technical definition, availability is measured as the percentage of time the system(s) are operational, or 99.99%, for example. This measure is often referred to as "9's," with many data center managers striving for the "five nines" goal. Others think of availability in terms of downtime. For instance, 99.999% equates to five minutes of downtime per year. Yet a third way to talk about availability is in terms

$$Value = \frac{Availability \times Agility}{TCO}$$

of "tiers," where each tier has a different set of redundancy and solution requirements. The Uptime Institute is an example of an organization that has defined a set of availability tiers.

With availability as the core reason for DCPI investments, how does one decide what the availability requirement should be? Every data center manager wants the highest availability possible. The key to what drives the requirement, then, is often the cost of the solution. Businesses always have competing projects and limited budgets and by choosing a DCPI solution that is designed for high **reliability**, with emphasis on designing out **human error** and designing in quick **recovery times**, managers can get the greatest availability possible. These three key contributors to a highly available DCPI (equipment reliability, recovery time, and human error) are discussed in more detail below.

Equipment reliability

Reliability is the probability that a device, system, or process will perform its defined function without failure for a given period of time. "Time" is an important component of this definition. Reliability is often a misspoken term, as it is often referred to as "product X is 98.5% reliable." This statement is meaningless without providing a timeframe upon which that reliability applies. So, the proper phrase would be "there is a 98.5% chance that product X will operate

without failure for 3 years." Reliability is also sometimes expressed as unreliability or the probability of failure – basically the opposite variable. For example, the unreliability equivalent for the above statement would be expressed as "there is a 1.5% chance of failure in 3 years of operation."

The physical infrastructure of a data center is comprised of many components. For the system as a whole to be reliable, the individual pieces of equipment must be reliable. Data center managers do not want to be inconvenienced by component failures, regardless of how fast they can be repaired. Every failure is costly in terms of man hours and dollars, and the more reliable the components, the less likely that inconvenience is to occur.

Mean time to recover (MTTR)

MTTR plays an important role in a repairable system's availability. Data centers should be designed to be as reliable as possible in terms of equipment. But with a data center that has a typical life of 10 – 15 years, it is inevitable in that timeframe that systems or subsystems will fail. These failures may or may not result in downtime, depending on the level of redundancy designed into the system. Regardless, when these failures do occur, it is important that the systems have as quick of a recovery as possible. This means the system is diagnosed quickly, parts are readily available, and the system is easily repaired or replaced. Many have either experienced or heard the horror stories of data centers that went down due to a failure, and took days to recover. The impact of such an occurrence on the business is devastating.

Human error

Human factors simply cannot be ignored when considering the availability of DCPI. As previously stated, various industry studies show that human error is the leading cause of unplanned downtime, at up to 60%. IT Equipment is typically swapped out 4 or more times during the life of a data center. When this occurs, the power, cooling, and even the security requirements may change. For example, a new server may require a different receptacle thereby requiring the change of live power circuits. This unstable environment, subject to constant change, creates chaos and invites human error.

Generally vendors provide availability or reliability data to their clients or potential clients, but they exclude human error. When failures of the actual equipment represent such a small proportion of the total and failures from human factors represent such a large proportion, these numbers become very misleading. The more meaningful metric for availability of a DCPI system would include these human-driven failures. Including these failures in the metrics would drive vendors to design complexity out of their systems, design in simple and intuitive interfaces, and design in the ability to accept and effortlessly handle constant change. Doing so would ultimately drive up business value.

Checklist of availability considerations

Table 1 provides a list of often overlooked availability considerations that impact business value. In the past many of them were simply not possible, but in today's reality there is no reason to accept any "no" answers to these questions.

Table 1

List of availability considerations

Availability considerations
Are the components of each DCPI element manufactured in high volumes to improve the equipment reliability?
Has redundancy been designed into the DCPI to minimize likelihood of component failure taking down the data center / IT systems?
When a failure occurs, can it be recovered in less than an hour?
Are all elements of the DCPI system designed to integrate together for seamless operations?
Has the system been designed to drive out complexity?
Has the system been designed with intuitive, simple interfaces and proactive management?
Does the system have the ability to accept and effortlessly handle?

Agility – performance vector #2

Agility can be defined as the ability of a system to adapt to change. Adapting to change can mean a lot of things – it means getting systems installed on time, especially for time-constrained projects (**speed of deployment**); it means being able to **scale** in size as needs of a business change while keeping oversizing waste to a minimum; and it means being able to quickly accommodate the ever increasing number of changes required by the IT needs of a business (**ability to reconfigure**). The classic example of a business that lacked agility and failed is many of the big collocation companies. Such companies invested huge amounts of capital to develop solid, high security infrastructure they thought their potential customers would need to host their critical IT equipment. Because their systems couldn't adapt to changing business requirements, they planned for "worse case" in terms of capacity. The amount of guesswork was immense. The result was a huge waste of un-utilized infrastructure, and depleted pocket books.

$$Value = \frac{Availability \times Agility}{TCO}$$
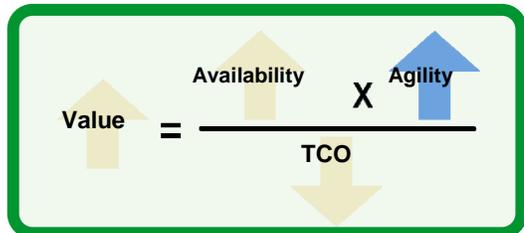
Agility is a key piece of the business value puzzle – as important, in terms of driving final dollars invested, as the other components of business value. When the DCPI investment can be built for present needs and be agile enough to adapt to any future requirement (rather than forecasting ten years into the future), the risk of prediction error is eliminated and the return on the investment is optimized.

Speed of deployment

The speed of deployment is simply how fast the DCPI system can be planned, designed, installed, and commissioned for operation. Historically, this has been on the order of years. **Figure 4** illustrates the typical timeline of such projects. This lengthy timeline is due to the fact that traditional legacy approaches involved a high degree of customization. The process to get more capacity involved several major steps that are custom exercises – one-time engineering, one-time construction design, the construction process itself which is highly

customized, and commissioning that is tailored to integrating a lot of equipment not specifical-ly designed for that integration.

Figure 4

Typical timeline of 400 days for deploying legacy DCPI systems
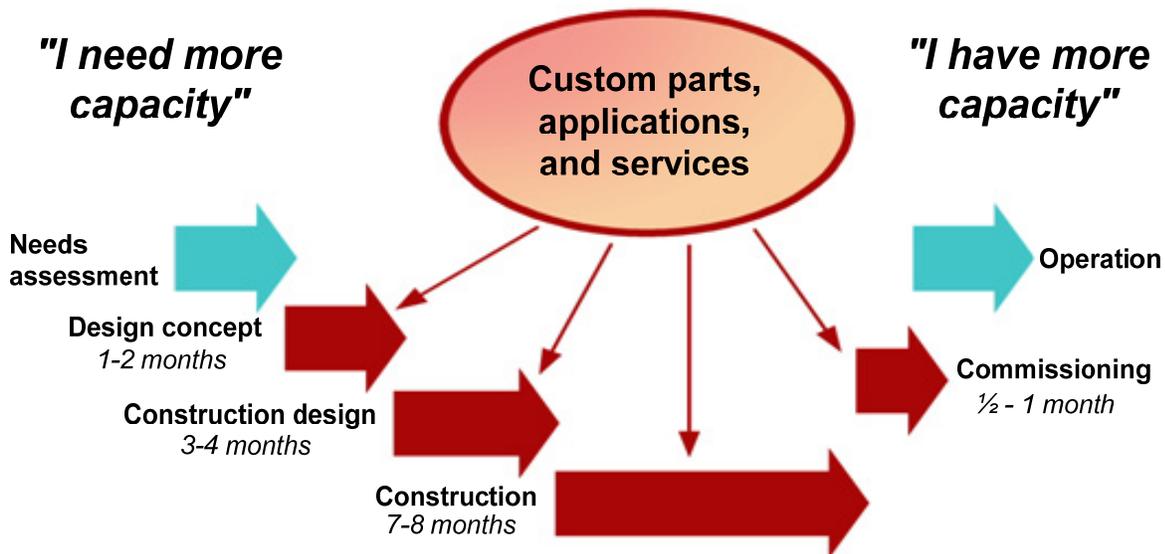
With today's fast-paced business demands, this inflexible timeline is at best inadequate, at worst disastrous. Companies need data centers built in a few months, not years.

Another timeline anecdote many data center managers can relate to is the project delays caused by repeated attempts to forecast the future. By eliminating the time and resources necessary to plan for capacity and densities before they are needed, project timelines can be shortened.

Ability to scale

When capital funds are not free-flowing (which is generally the case), the ability to scale is key to being agile. The average data center is significantly overbuilt. Research shows that typical data centers today are utilized to less than 50% of their infrastructure capacities. In fact, according to "Data Center Power Requirements: Measurements from Silicon Valley," a thesis written by J.D Mitchell-Jackson, et al., typical data centers are approximately only *one-third full.* 66% is not a small amount of waste. Imagine that the ROI for that project could have been three times greater. Two-thirds of the investment could have been applied to other company projects with greater return potential.

Data center managers didn't overbuild because they wanted to waste. They overbuilt because they had to assume the worse case in terms of capacity, based on their analysis of future business needs, while not having visibility beyond 3-4 years. Until recently, the alternative was even more negative – to underbuild and not have the ability to increase capacity without disruption and extensive cost to the business. White Paper 37, *Avoiding Costs from Oversizing Data Center and Network Room Infrastructure*, discusses the signific-ance of rightsizing in greater detail.

Link to resource
White Paper 37

Avoiding Costs from Oversizing Data Center and Network Room Infrastructure

Ability to reconfigure

Consider the collocation company that built out 50% of their facility with DC power and 50% with AC power, only to later learn that 99% of their customers would require AC power. The

DC investment was nearly a total loss due to this lack of agility. Now imagine the present – what if this same collocation company could have built only what their current customers required, and reconfigure only as needed for different power densities, different levels of redundancy (targeted availability), and for different voltages and plug types, without ripping apart the data center.

Imagine another company that designed and built a data center, only to find out after it was commissioned that the facility needed to be relocated to another state. With legacy systems, the majority of the resources invested cannot be relocated and ultimately results in huge financial impacts to the business, since a substantial second investment has to be made for the new location. Systems that are portable and have the ability to relocate can add substantially to the bottom line.

Checklist of agility considerations

Table 2 provides a list of often overlooked agility considerations that impact business value. As with the availability considerations, these are items that should be expected in a successful DCPI deployment. With the fast paced and fast changing business conditions of today, business success means answering "yes" to all of these agility considerations.

Table 2

List of agility considerations

Agility considerations
Can a new data center be built in days instead of months or years?
Can additional capacity be added quickly to the existing DCPI system? Could it be doubled in the next month?
Can the data center's DCPI be installed with minimal site work?
Can a significant percentage of the infrastructure be moved to a new site if relocation becomes necessary?
Can relocation of the data center occur in months rather than years?
Can a portion of the data center be supplied with redundant DCPI (targeted availability)
Can a plug type be changed in a matter of minutes during IT refreshes?
Can backup runtime be scaled as requirements change?

TCO – performance vector #3

Total Cost of Ownership or TCO is clear to most managers – it is dollars spent and is very tangible. The TCO for a particular data center will depend on the size of the facility. However, it is possible to express total cost in a way that is useful for any size data center. A suggested way of expressing the TCO requirement is in terms of the total cost per rack over the life of the data center. A typical value for lifetime cost of DCPI per rack is $100,000. This is a substantial investment (comparable to the cost of the IT equipment investment) and it is therefore crucial that the manager understand the value received for the cost.

TCO is not simply the initial investment of DCPI, but includes all other costs associated with operating and maintaining that investment for its life (typically 10-15 years). It should be clear that in today's competitive business world, it is TCO, and not upfront cost alone, that drives real business value. **Figure 5** illustrates a typical breakdown of costs for data centers implemented using traditional systems.
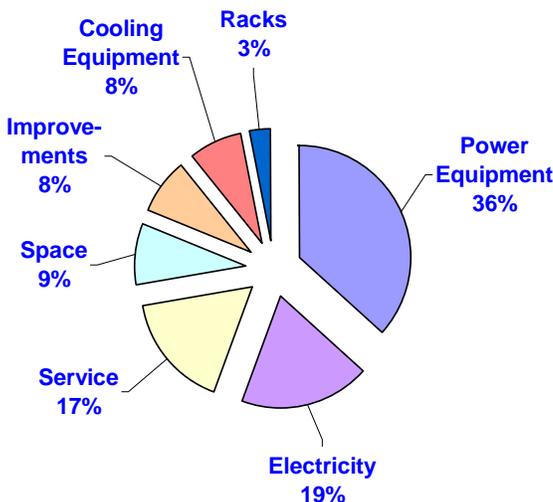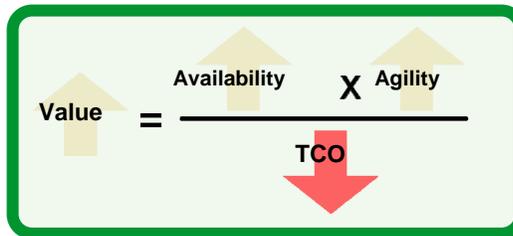
$$Value = \frac{Availability \;\; X \;\; Agility}{TCO}$$

Figure 5

TCO breakdown for traditional systems

Of these costs, approximately 50% is operating expense and 50% is capital expense. And of this total, a substantial percentage (roughly 30%) is typically wasted due to poor planning and design decisions.

Oversizing is the dominant contributor to excess TCO. This results in considerable waste, including excess **capital cost**, **operating cost**, and specifically, **energy cost**. On average, over 50% of installed DCPI capacity is wasted. While this waste is caused primarily by uncertainty, the answer cannot be simply to reduce plans – there certainly are cases where the actual load ends up equaling or exceeding the expected load, and underbuilding data center capacity can have even greater consequences than overbuilding The real key to achieving an optimized TCO is to implement an DCPI solution that is designed with the ability to easily scale so that the infrastructure can be rightsized to match the needs at any particular point in the data center's life. No one wants to be responsible for the data center that built out a megawatt of capacity only to learn that the actual load never achieved more than 300 kW. White Paper 6, *Determining Total Cost of Ownership for Data Center and Network Room Infrastructure*, discusses a method for measuring TCO and suggests strategies for minimizing these costs.

Link to resource
White Paper 6

Determining Total Cost of Ownership for Data Center and Network Room Infrastructure

Capital cost

Capital expense, also referred to as CAPEX, is the money spent by a company to acquire or upgrade property, plant, and equipment assets, with the expectation that they will benefit the company over a long period of time (more than one year). These costs are depreciated for the expected life of the equipment and have special tax treatments. Capital cost for a DCPI project is a very real and measurable expense. This includes the cost of the DCPI equipment, as well as the costs associated with designing and deploying that equipment. The single biggest opportunity to optimize (reduce) this capital cost is through the ability to scale

the infrastructure. Matching the capacity to the load avoids significant (up to 3 times) over-investment in capital. Another opportunity to reduce CAPEX is through the reduction of technical labor needed. Implementing a solution that minimizes site work during installation drives down this upfront labor cost.

Operating cost

Operating expense, also referred to as OPEX, is all of the costs associated with maintaining a business. Operating cost for a DCPI deployment includes the DCPI operating staff, training expenses, and maintenance and repair costs. Energy cost is also an operating cost but has been separated to emphasize its significant impact on business value.

It is important to consider the level of predictive maintenance available in the system, as this can reduce the need for programmed maintenance (which not only reduces expenses but also reduces the risk of introducing defects). The level of services needed also greatly depends on the complexity and customization of the system. The amount and type of repairs also is determined by the nature of the system. Modular systems allow for more self-servicing and replacing of failed components, whereas legacy customized systems require more onsite invasive repair procedures.

The more complex and customized a system is, the greater the amount of design effort and site preparation time, and ultimately the greater the amount of highly technical staff needed to support the system. This generally means greater spending on training to bring the skills of the employees to the level needed to operate the systems successfully.

Energy cost

The largest potential for TCO savings comes in the form of efficiency cost. A UPS system that is oversized is less efficient and requires substantially more energy to run the than a system that is sized to match the load. This is because operating efficiencies improve dramatically as the actual load approaches the design load.

In redundant UPS systems, the size of the redundant module(s) also impacts efficiency. This is because in a typical N+1 configuration, the load is shared across all modules. Consider a data center load of 80 kW. If the UPS system is designed with four 20 kW systems to support the load and one extra 20 kW system for redundancy, the 80 kW load would be shared across the 5 systems. In other words, each UPS would support 16 kW, or 80% of capacity. This system could also be designed with two 80 kW systems, where each runs at 50% capacity. In this example, the first design would provide greater efficiency, and therefore reduced energy cost.

The technologies chosen for DCPI can have a huge impact on energy costs. Traditional UPS systems, for example, are only about 85% efficient at 50% load levels. Newer on-line technologies, such as delta-conversion UPS systems, maintain an efficiency of 93% or better at these low load levels. With electrical costs well over $0.10 per kW-hour in most areas of the United States, this efficiency edge contributes significantly to a low total cost of owner-ship.

Checklist of TCO considerations

Table 3 provides a list of often overlooked TCO considerations that impact business value. As with the availability and agility considerations, these items are essential to making wise financial business decisions regarding DCPI.

Table 3

List of TCO considerations

TCO considerations
Is the project being assessed on TCO rather than upfront cost alone?
Is the capacity of the DCPI rightsized to optimize the capital investment?
Is the electrical efficiency optimized for the data center, by scaling appropriately?
Are high efficiency DCPI components included in the design?
Are long-term maintenance costs minimized for the data center?

Strategy for optimizing business value

In today's business environment, it is clear that an updated set of criteria must be reviewed to optimize the business value. **Figure 6** re-emphasizes what these new criteria should be. This translates into real cost that affects the business in many ways – not only in terms of dollars, but also strategy and planning.

Figure 6

The new business value equation

$$Value = \frac{Availability \times Agility}{TCO}$$

In legacy DCPI systems – with their unique one-time engineering and inflexible, typically oversized, design –it was not possible to achieve a system that could optimize all three performance vectors of this equation. It was not feasible to get a system of high quality (availability), in a reasonable amount of time (agility), at a reasonable cost (TCO). There was generally a sacrifice – give up one of these variables to get the other two. If you wanted it fast and cheap, you would sacrifice quality. If you wanted it high quality and fast, you'd certainly pay the price. That old paradigm, happily, no longer exists.

There is a proven strategy for optimizing all three business needs – through standardization, and more specifically, **modular standardization**. Applying the same approaches as in designing cars to IT systems, for instance, can substantially improve the business value of DCPI. Consider servers of the past – each server used to fit in a different type of rack, which resulted in incompatibility – but when the server industry became standardized, servers could fit in nearly any rack. This is just one of many examples of the direction the DCPI industry is

heading. White Paper 116, *Standardization and Modularity in Data Center Physical Infrastructure*, describes the wide-ranging benefits that standardization and modularity provide to a DCPI investment.

Conclusion

DCPI is foundational to a data center. Changes in today's business environment, particularly the shortened IT refresh cycle, are driving a need for change in thinking about how to assess the business value of DCPI. The new DCPI business value equation described in this paper is summed up in this. Availability must be thought about in a new way, accounting for key drivers including human factors. Cost must be thought about in a new way: no longer is upfront cost alone sufficient for making business decisions – the cost over the infrastructure's life is what really drives value. And agility is a performance vector that now must be considered, as it directly measures the ability of the DCPI to meet unforeseeable demands and opportunities.

The key to achieving high business value under this new paradigm is standardization, modularity, and scalability. The DCPI industry must be taken out of its old "craft industry" mindset where two different facilities have two totally different designs and two totally different sets of unique problems. Systems must be modular and scalable to match ever-changing needs of data centers, improve the reliability and availability of the infrastructure, and optimize TCO.

About the author

Wendy Torell is a Strategic Research Analyst at Schneider Electric in West Kingston, RI. She consults with clients on availability science approaches and design practices to optimize the availability of their data center environments. She received her Bachelors of Mechanical Engineering degree from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

Resources
Click on icon to link to resource

Avoiding Costs from Oversizing Data Center and Network Room Infrastructure
White Paper 37

Determining Total Cost of Ownership for Data Center and Network Room Infrastructure
White Paper 6

Standardization and Modularity in Data Center Physical Infrastructure
White Paper 116

Browse all white papers
whitepapers.apc.com

Essential Power System Requirements for Next Generation Data Centers
White Paper 4

Re-examining the Suitability of the Raised Floor for Data Center Applications
White Paper 19

Physical Security in Mission Critical Facilities
White Paper 82

Browse all TradeOff Tools™
tools.apc.com

Mitigating Fire Risks in Mission Critical Facilities
White Paper 83

P3
www.p3-inc.com

OMAHA
1207 S. 75th Street
Omaha, NE 68124
402 393 1223 P
402 393 4779 F

KANSAS CITY
6603 A Royal
Pleasant Valley,
MO 64068
816 792 4641 P
816 792 4171 F

DES MOINES
4725 Merle Hay Rd.
Suite 105
Des Moines, IA 50322
515 277 5770 P
515 277 5771 F

ST. LOUIS
62 Soccer Park Rd.
Fenton, MO 63026
636 717 3422 P
636 717 3430 F

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative

Appendix: elements of DCPI

DCPI includes power, cooling, racks & physical structure, security & fire protection, cabling, and the management and servicing of these elements. DCPI is the foundation of reliability of IT operations for a business. As the foundation, it must be dependable, and ideally, "out of sight, out of mind." Below are brief descriptions of each of the DCPI elements.

Power

There are many elements that make up a solid power infrastructure, including all systems from the electrical service entrance of the building to the loads of the data center or network. This is comprised of the main distribution, the generator(s), the UPS systems and batteries, surge protection, transformers, distribution panels, and circuit breakers.

Link to resource
White Paper 4
Essential Power System Requirements for Next Generation Data Centers

White Paper 4, *Essential Power System Requirements for Next Generation Data Centers*, discusses the problems and challenges with traditional power systems and presents a set of requirements moving forward.

Cooling

Cooling systems that are required to successfully remove heat from a data center include computer room air conditioners (CRAC), any associated subsystems that allow the CRAC to operate, including chillers, cooling towers, condensers, ductwork, pump packages, piping, and any rack-level distribution devices.

Racks and physical infrastructures

There are many systems that could be considered part of the physical structure of a data center. The most critical elements are the IT racks which house the IT equipment and physical room elements such as dropped ceilings and floors (both raised floors and concrete "slab" floors).

Link to resource
White Paper 19
Essential Power System Requirements for Next Generation Data Centers

White Paper 19, *Re-examining the Suitability of the Raised Floor for Data Center Applications*, takes a close look at the past reasons for raised floors and suggests that their widespread use is no longer justified in data center environments.

Security and fire protection

Security systems and fire protection systems are essential to maintaining the integrity, safety, and availability of the data center. Subsystems included here are physical security devices at the room and rack level and fire detection & suppression systems. Some examples of physical security devices include biometric devices, keys, codes, & cards. Examples of fire detection & suppression systems include intelligent type smoke detectors, clean agent fire suppression systems, and linear heat detection.

Link to resource
White Paper 82
Essential Power System Requirements for Next Generation Data Centers

Link to resource
White Paper 83
Essential Power System Requirements for Next Generation Data Centers

White Paper 82, *Physical Security in Mission Critical Facilities*, discusses physical security elements in greater detail and suggests procedures for optimizing the security of the facility. White Paper 83, *Mitigating Fire Risks in Mission Critical Facilities*, describes the detection, suppression and prevention of fire in greater detail, and presents best practices for mitigating fire risks.

Cabling

The cabling infrastructure encompasses all data cables that are part of the data center as well as all of the power cables necessary to ensure power to all of the loads. Cable trays and cable management devices are also critical to support the IT infrastructure as they help to reduce the likelihood of downtime due to human error and overheating.

Management

Management is an element of DCPI that spans all of the elements discussed above. In order to have a reliable DCPI, it is important to have visibility to all of the components of the physical infrastructure. Management includes systems such as Building Management Systems (BMS), Network Management Systems (NMS), Element Managers (such as InfraStruXure manager), and other monitoring hardware and software.

Services

There is a broad range of services that is necessary to support DCPI systems throughout its life cycle. These services can be broken into five categories: (1) Consulting and design services; (2) Installation services; (3) Maintenance and repair services; (4) Monitoring services; and (5) Decommissioning services.

877-393-1223 | www.p3-inc.com