DATA CENTERS

SUMMER EDITION

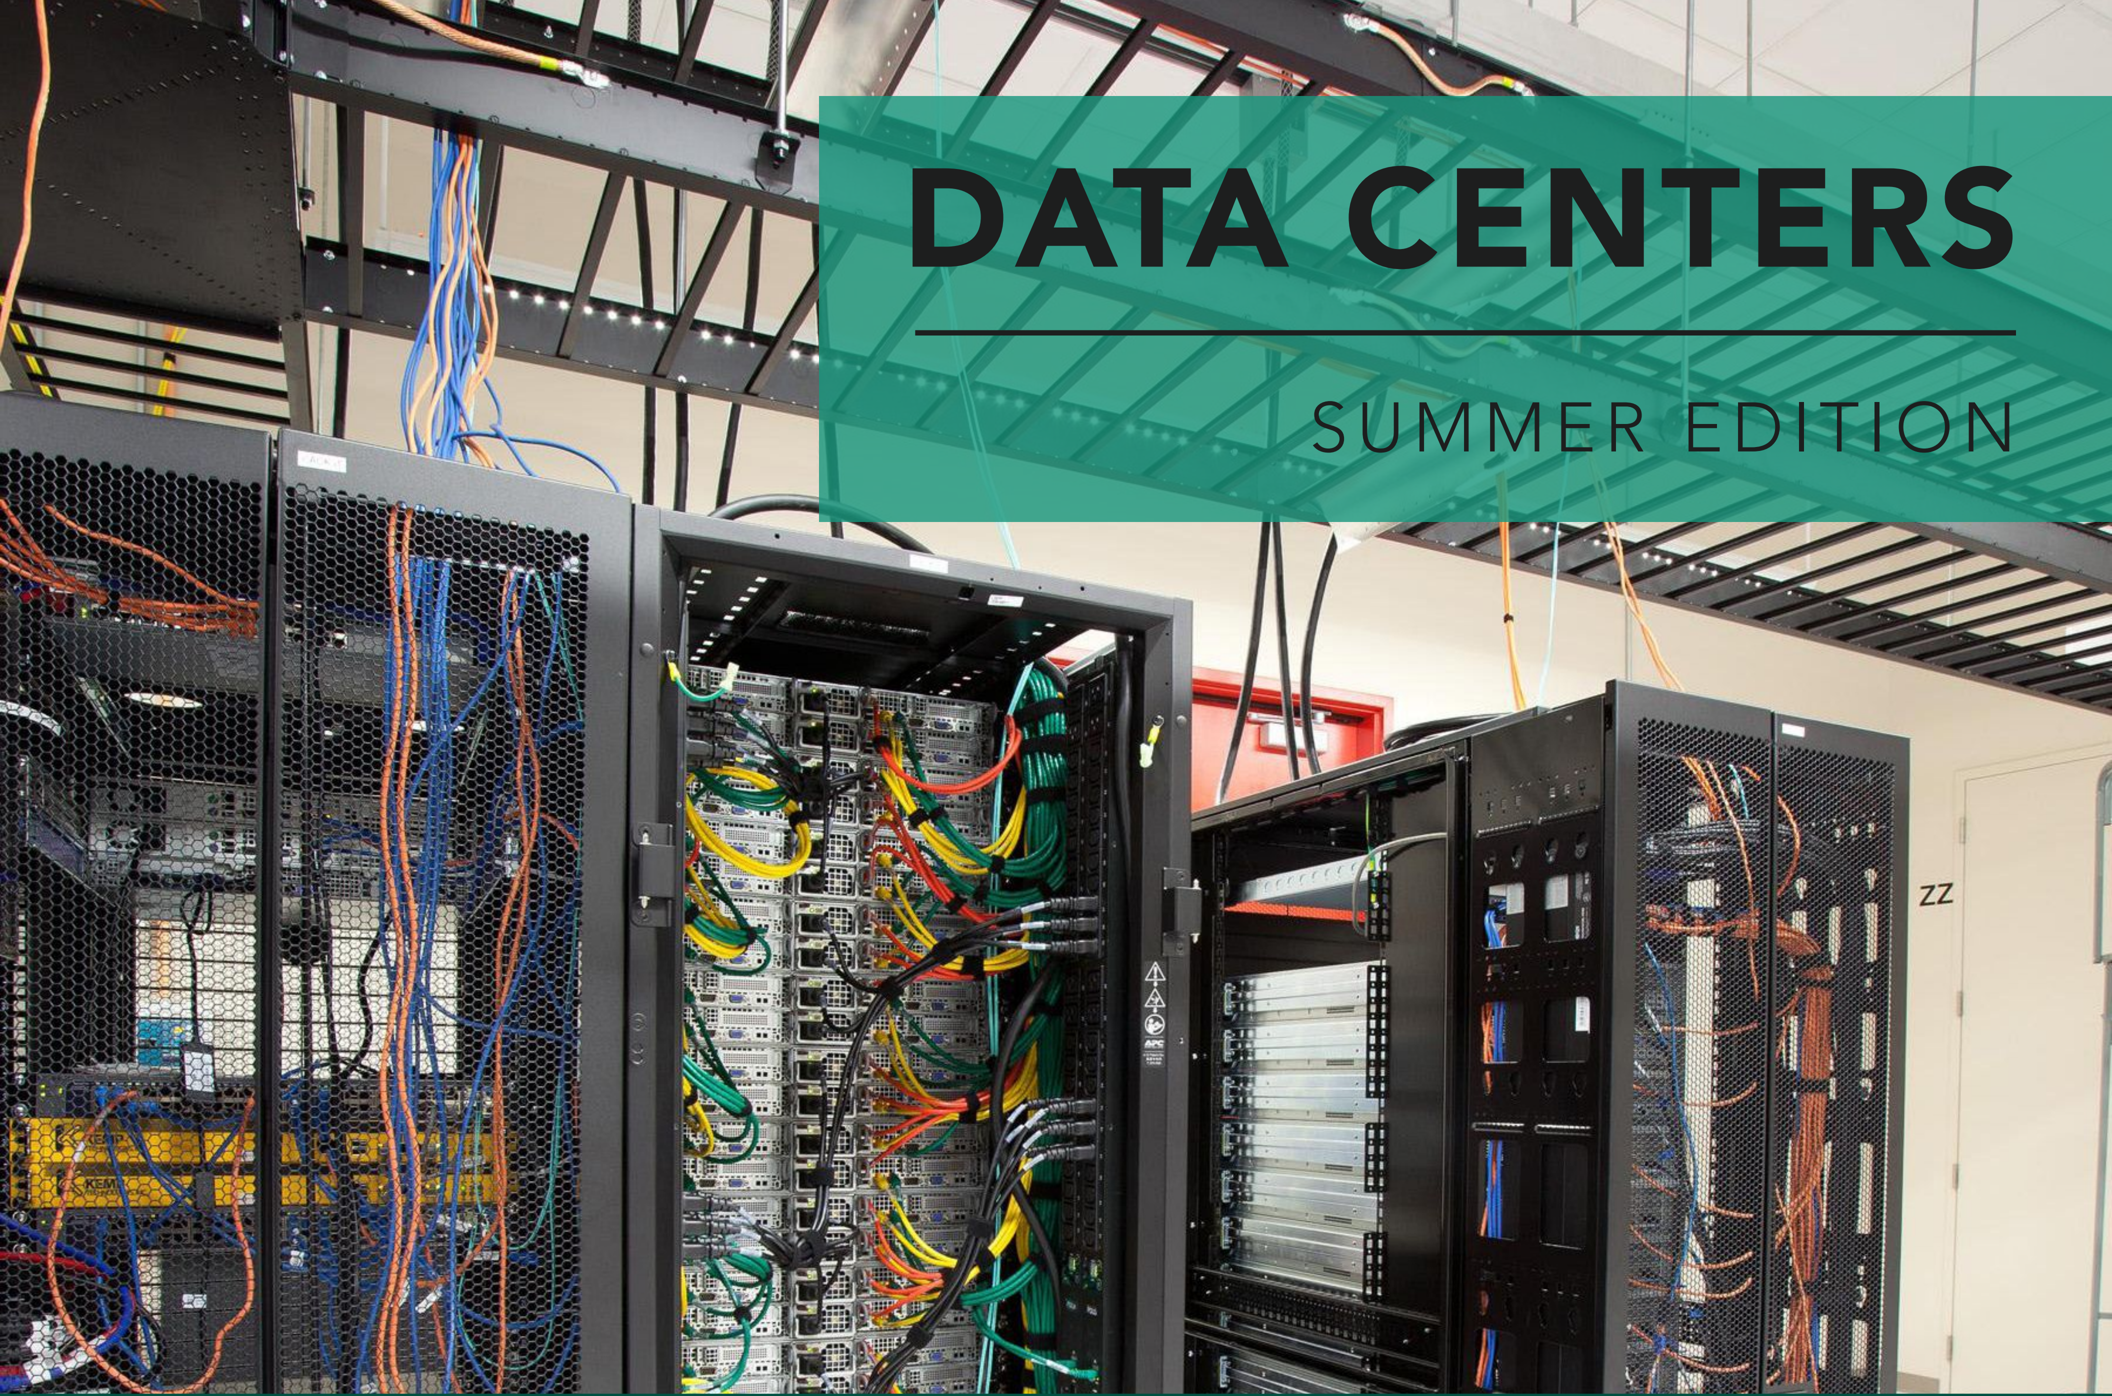

CONSULTING - SPECIFYING engineer eBOOK

Carrier

CHAMPION FIBERGLASS

LEVITON Network Solutions

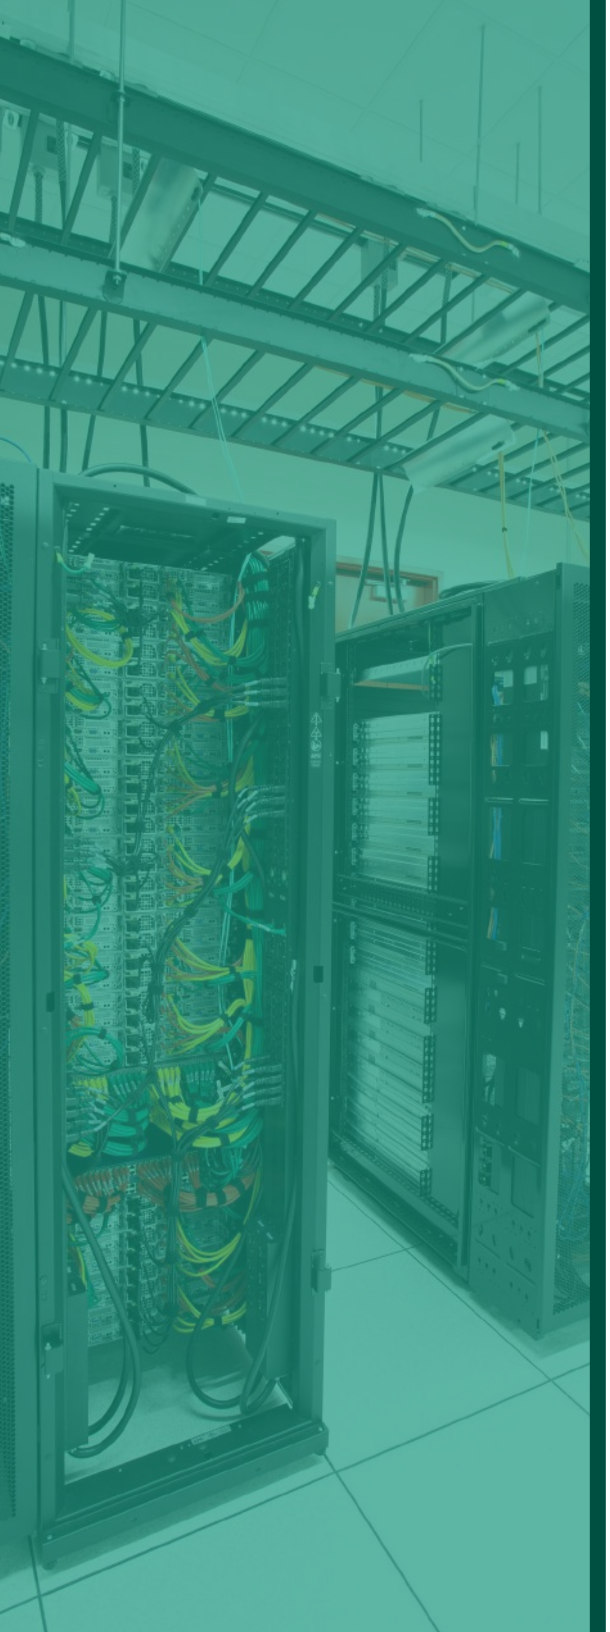

Contents

Designing data center electrical distribution systems

Designing efficient and reliable data center electrical systems requires looking through the eyes of the electrical engineer — and the owner.

Data centers are among the hottest developments in the technology world. The growing needs of the Internet of Things have forced the biggest players in the computing world to spend billions of dollars on new multi-megawatt data centers. This boom in data center construction is largely fueled by the growing use of cloud services, which has put a strain on server capacity (see Figure 1). Additionally, data centers are considered mission critical when their operation is of importance to organizations' economic or functional needs. Even a disruption of a few seconds in the operation of certain types of mission critical data centers could cost millions of dollars.

This article explores data center design through the eyes of both the owner and the electrical engineer. It also discusses the key components of data centers and touches on the codes and standards that apply to data centers and their components.

Preliminary considerations

Data centers, many having servers as their main components, need electrical power to survive. It is, therefore, only natural that any talk about building a data center should begin with figuring out the electrical needs and how to satisfy those needs.

Capacity: Before deciding anything else, the owner must decide the capacity of the data center (in megawatts). In previous planning efforts, it was common to use W/sq ft. However, today it is more common to discuss kW per rack, which may vary from 5 to

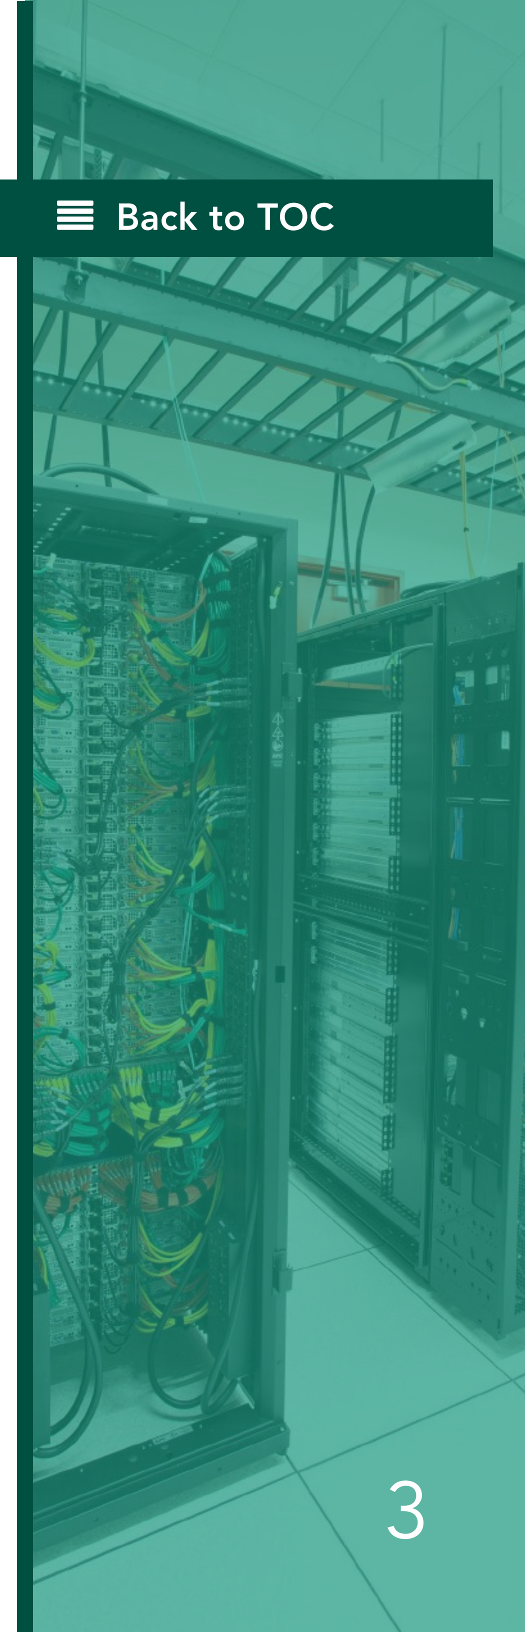

Back to TOC

60 kW. This power concentration per rack can also drive cooling system type and capacity, which must be planned for in the capacity. The owner also needs to consider future capacity.

Figure 1: Increasing demand for cloud services is putting a strain on server capacity. This photo shows data center servers while they are being configured and wired. Courtesy: Jacobs Engineering

Another big decision is to determine the level of redundancy. Reliability is very important for data centers, and disruptions are costly. But the cost of building a data center increases significantly with higher reliability. Therefore, the owner should decide where to draw the line, and determine how much risk is acceptable.

Auxiliary power: After the data center capacity is decided, the facility power must be computed. The facility power includes data center heating and cooling. A focus of recent years is to make the facility (non-data) power as low as possible to improve efficiencies and lower operating costs. To address the efficiency of facility power within a data center, the term "power usage effectiveness" (PUE) was coined. The closer to unity the PUE is, the smaller the nonproducing facility power is.

Years ago, it was normal to account for the facility and cooling load as being half of the total power delivered to the data centers. That means that if a data center had a

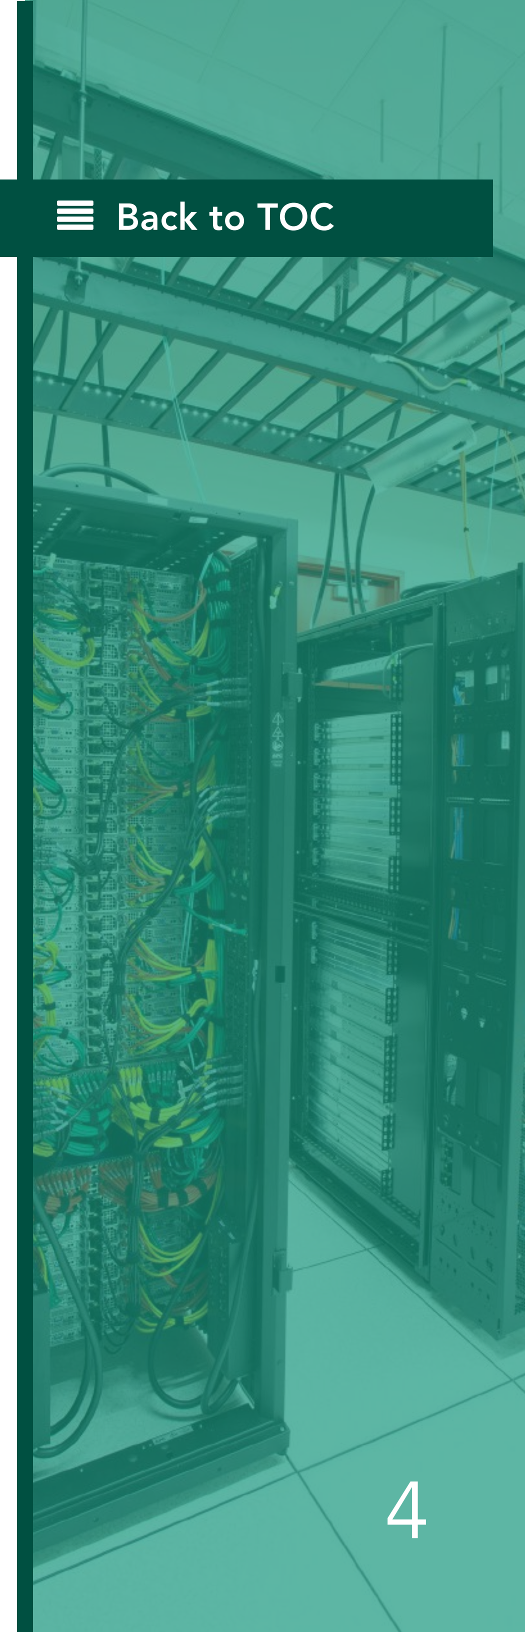

capacity of 10 MW, the facility and cooling load also would be 10 MW, leading to a PUE of 2. PUE equal to 2 is deemed to be average efficiency, but not satisfactory to many data center owners. New technologies have pushed the PUE very close to unity.

Table 1: Uptime Institute tier system

Classification	Description
Tier I	■ Lack a redundant electrical system. ■ Have components without redundant capacity.
Tier II	■ Have components with redundant capacity.
Tier III	■ Maintain duel-powered IT equipment. ■ Have multiple distribution paths to the servers.
Tier IV	■ Have multiple distribution paths to the servers. ■ Have multiple electrical power distribution paths. ■ Have storage capacity and dual-powered HVAC equipment.

Table 1: Data centers are classified into one of four Tiers from lowest to highest reliability. Courtesy: Jacobs Engineering

Reliability and tiers: To classify data centers in terms of reliability, the Uptime Institute created standards referred to as Tiers (see Table 1). Data centers are classified in four Tiers. Tier I data centers don't have a redundant electrical distribution system, and their components don't have redundant capacity. Tier II data centers differ from Tier I data centers in that they have components with redundant capacity. Tier III data centers have dual-powered IT equipment and more than one distribution path to the servers. Tier IV data centers have all the features that Tier III data centers have. In addition, Tier IV data centers are fault tolerant in that they have more than one electrical power distribution path. Tier IV data centers have HVAC equipment that is also dual powered and have storage capacity.

Determining which Tier to select depends on numerous factors. Many organizations used to have large consolidated data centers, which led to choosing a Tier III or Tier IV system. Also, many organizations involved in financial industries choose Tier III

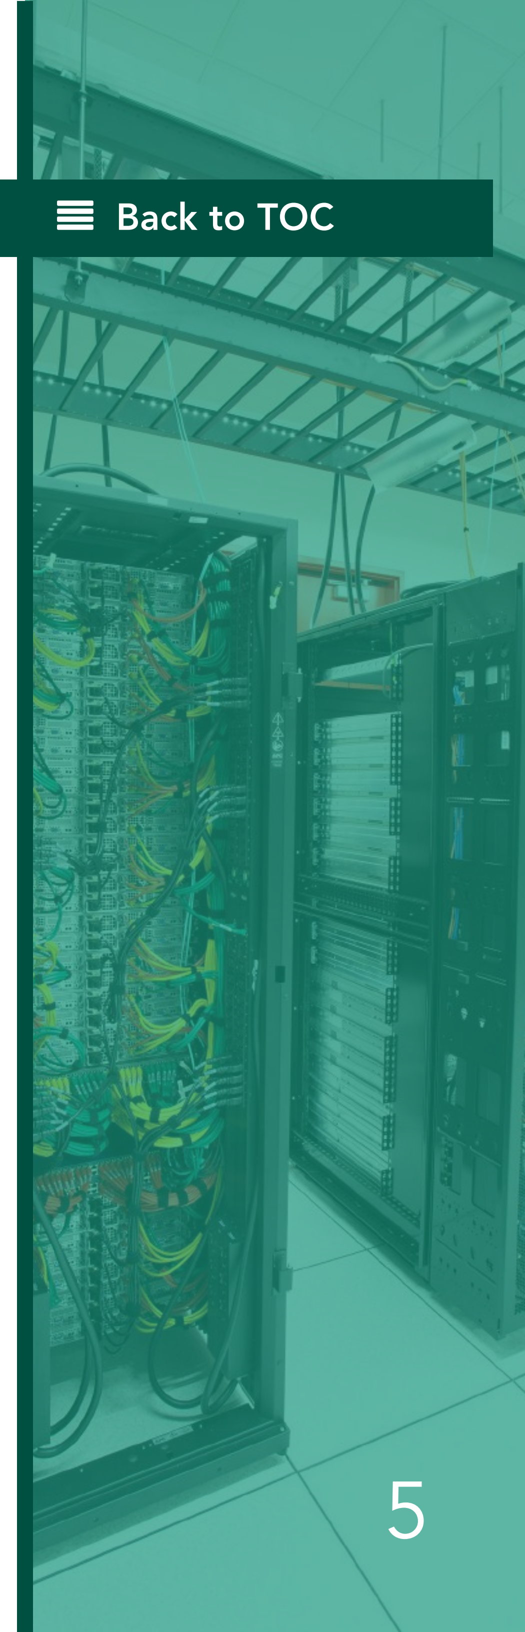

Back to TOC

Back to TOC

and Tier IV systems. Other organizations choose to have multiple data centers that can handle data needs when another center goes down, leading to an ability to use lower Tier systems.

Usage: Data centers are also categorized according to their usage. These include data centers serving a private domain, such as a corporation or a government entity; data centers serving a public domain, such as Internet providers; and multi-user data centers.

Figure 2: For many years, ac has been the dominating form of data center power distribution as shown in this photo of servers powered through overhead busways via busplugs. Courtesy: Jacobs Engineering

Power distribution: Currently, there is debate about what kind of electrical power to use to feed data centers. Should it be ac or dc? Each has merits. Recently, dc power has received increasing consideration because data center computing equipment uses dc power. Having dc power distribution eliminates the need for transformers and ac-to-dc converters on the server floor. Using dc also eliminates harmonics because there is no switching of power. In addition, using dc eliminates conversion steps, which leads to higher efficiency (each conversion step introduces losses), thereby decreasing cost.

However, ac has been the dominating form of power distribution for many years (see Figure 2). The benefits of ac include readily available equipment, lower costs, and

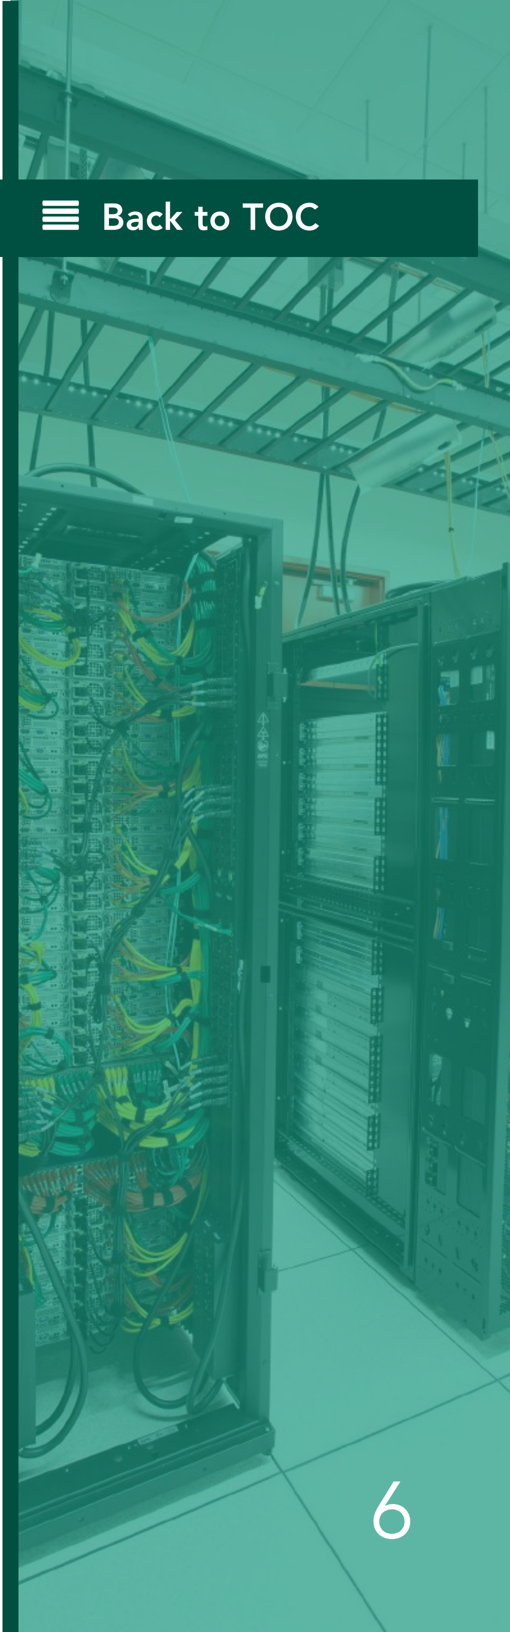

Back to TOC

easier maintenance (because the maintenance crews already know the equipment and the spare parts are readily available). Historically, most ac power distribution systems were designed at 208/120 V. The ever-evolving technologies have helped make the case for using higher ac voltages at 400/415 V, and even 480 V because of the higher power demands and efficiencies delivered by newer electrical equipment.

PUE: Another important factor in data center design and construction is PUE. The closer the PUE is to unity, the better. A data center with PUE of 1.5 is considered the middle line of efficiency. A PUE above that number shows an inefficient data center; a data center with PUE below 1.5 is considered to be efficient. A data center with a PUE of 1.2 is considered to be very efficient. The most important part of a data center is the IT equipment. If there were no supporting (auxiliary) loads, the PUE would be 1. Because the auxiliary loads are necessary, the PUE is always greater than 1. The auxiliary loads include HVAC loads and small electrical loads, such as lighting and receptacles.

Electrical design

After the owner decides on the above considerations, the work of the design professionals begins, especially for the electrical engineers. Electrical engineers have to come up with a design that is efficient, has enough capacity for future growth, and avoids unnecessary frills.

Power distribution elements: There are many parts to electrical power distribution. It starts with utility transformers, which in large data centers are owned by the data center's owner. After the power is stepped down from the utility transmission voltage to the distribution level, it goes through distribution switchgear that redirects power to where it is needed. Typically, the power must be stepped down again, more often than

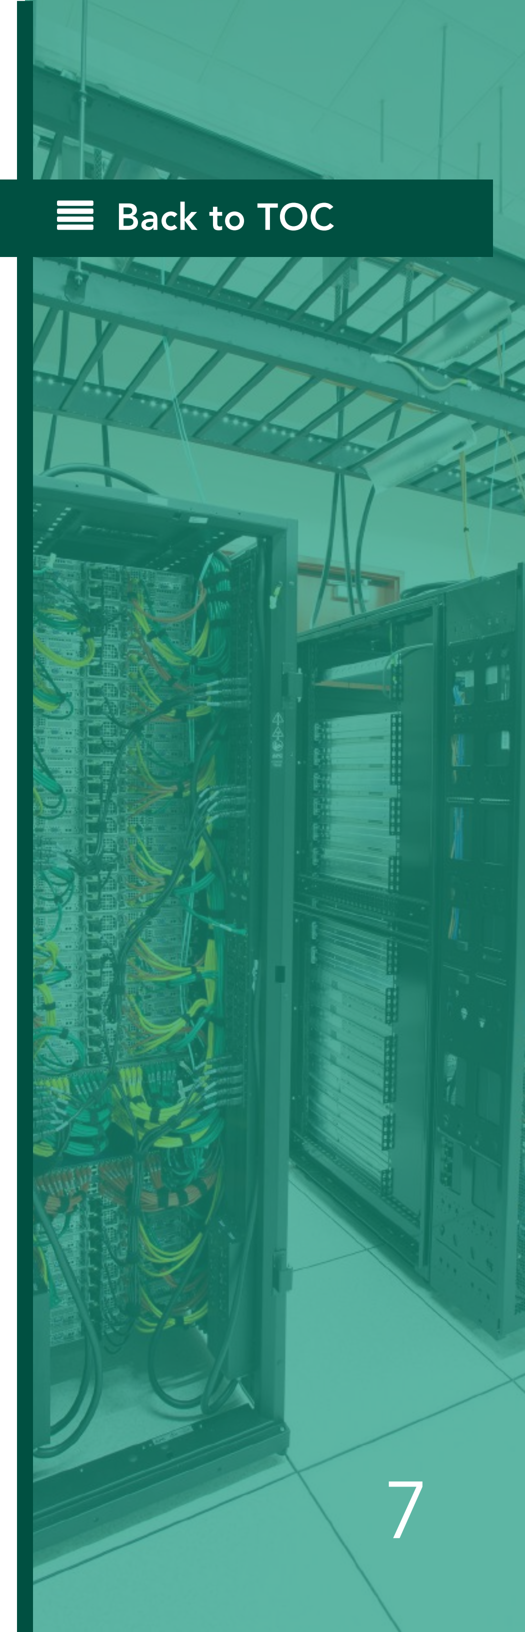

Back to TOC

not, via substation transformers and through more than one path. The standby power, usually present in today's data centers, is often introduced at this level, bringing with it the automatic transfer switch equipment. From the ATS, the power goes to the servers (often via a UPS system), where it switches from ac to dc power to be used by the servers. The next layer of distribution includes switchboards and panelboards that feed the auxiliary load, HVAC loads, and regular house loads. Power monitoring systems could also be employed at this point, which could provide very important information on how different pieces of equipment are working and how power is being used.

Going through so many pieces of equipment requires meticulous work. The design professional must be mindful of the cost of equipment and cables and also the losses introduced by each piece of equipment. Having so many pieces of electrical and mechanical equipment means that the engineer also must be mindful of many codes and regulations associated with these designs.

Relevant codes: The relevant codes for data center design professionals include ANSI/TIA -942-2005: Telecommunications Infrastructure Standard for Data Centers, NFPA 70: National Electrical Code, and ASHRAE: Standard 90.1: Energy Standard for Buildings Except Low-Rise Residential Buildings. Other very important codes include International Building Code, International Mechanical Code, International Plumbing Code, International Fire Code, International Fuel Gas Code, International Energy Conservation Code, NFPA 72: National Fire Alarm and Signaling Code, and NFPA 90A: Standard for the Installation of Air-Conditioning and Ventilation Systems.

Depending on the size of the data center and the type of building hosting it, other codes such as NFPA 13: Standard for Installation of Sprinkler Systems, NFPA 30: Flam-

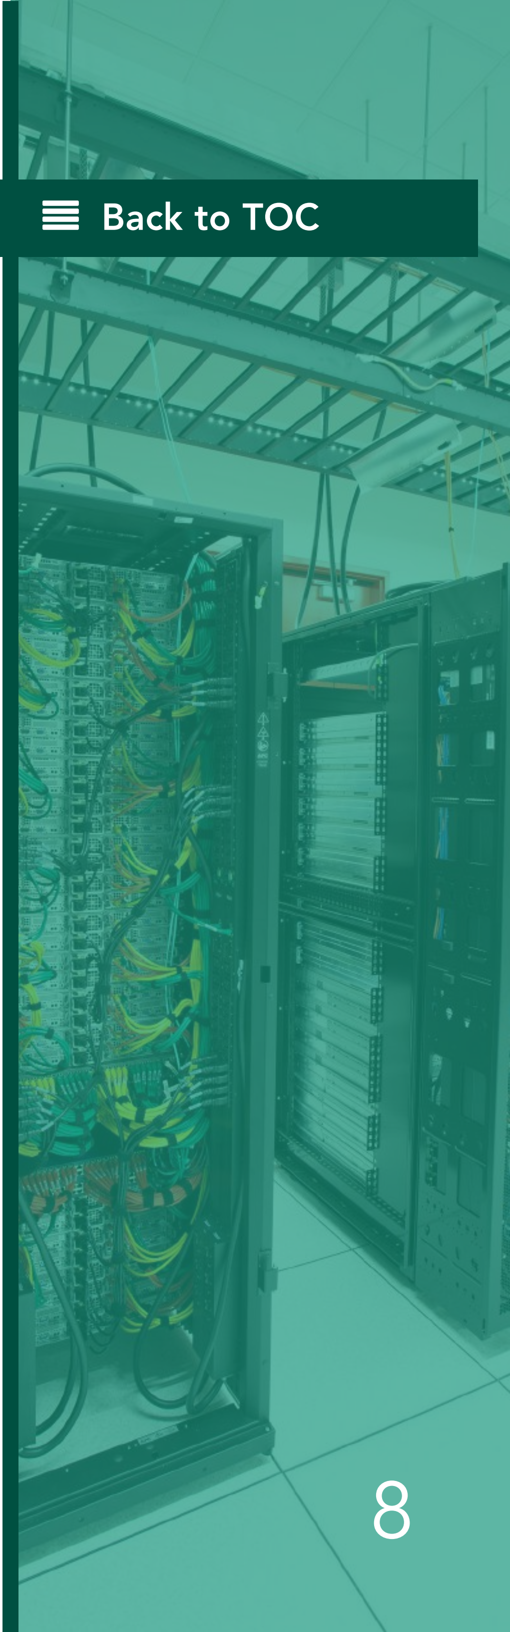

Back to TOC

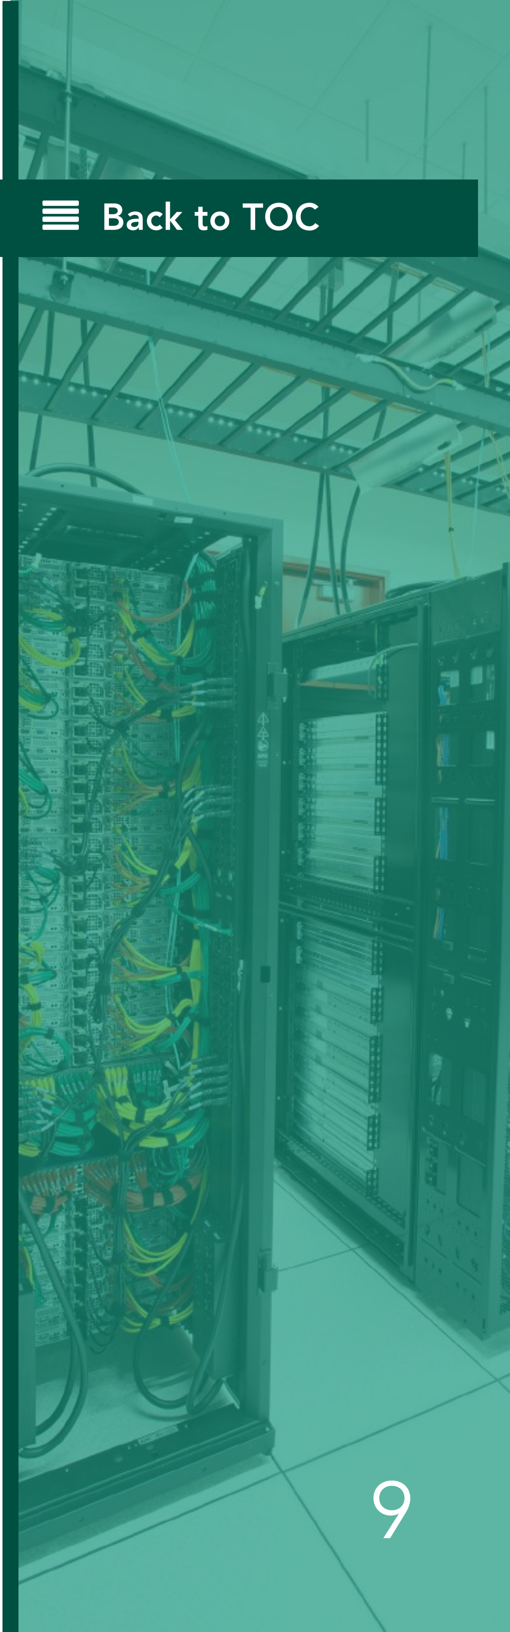

mable and Combustible Liquids Code, NFPA 10: Standard for Portable Fire Extinguishers, NFPA 101: Life Safety Code, NFPA 110: Standard for Emergency and Standby Systems, NFPA 780: Standard for Installation of Lightning Systems, and NFPA 20: Standard for Installation of Stationary Pumps for Fire Protection may apply.

Utility service: As with any other project, designers start by considering the utility service. Because of the importance of reliability, owners must engage early on with the utility company to discuss the service. Depending on the size of the data center, the service options include a separate dedicated utility line or an existing, very reliable line.

The electrical designer, in close collaboration with the owner, must decide how many layers of equipment will be there. The more equipment introduced, the more points of failure are present. In mission critical facilities, it is important to avoid single points of failure.

The utility service will most likely be medium voltage. Depending on the size and location of the data center, the service could be between 13.8 and 345 kV. The next step is to step down the voltage to a level usable for the servers. Most data center IT equipment works with dual voltage, 100 to 120 V ac and 200 to 240 V ac. The higher voltage—208 or 240 V—increases efficiency, thereby lowering losses. Having servers powered at 415 V ac further increases data center efficiency, making for a better PUE. If the designer decides to use the higher voltage, 415 V, the auxiliary mechanical load would then be at 480 V. This means that autotransformers must be used to take the power from 415 V to 480 V.

At what point does one decide to convert the medium voltage to low voltage (below 600 V)? The answer to this question depends on the size of the data center and the distance from the service drop. If the data center is part of a campus, the data center can be quite far from the service drop. If that is the case, it is preferable to distribute the electrical power at a voltage level as high as possible, typically 13.8 kV. If the service voltage is higher than 13.8 kV, the first transformation will be at the service entrance, stepping down the voltage from whatever the utility voltage is to 13.8 kV. This power is delivered to the data center where the second transformation takes place, stepping the voltage down to 480 V or 415 V.

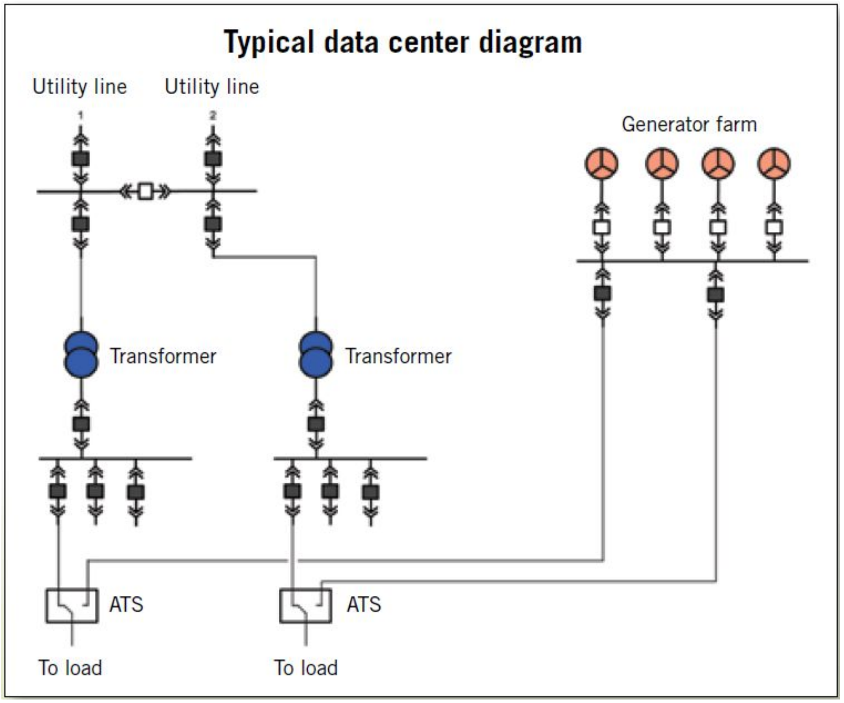

Figure 3: This one-line diagram of a typical data center shows the tie breaker on the primary side of the transformers. However, locating the tie breakers on the secondary side is just as effective. The tie breaker makes it possible to have two sources of normal power. Courtesy: Jacobs Engineering

Redundancy: What sets data centers apart is the level of redundancy. But everything comes at a price. The more layers of redundancy that are added, the more expensive construction of the data center becomes. Granted, having a data center blackout (or brownout) is very expensive as well.

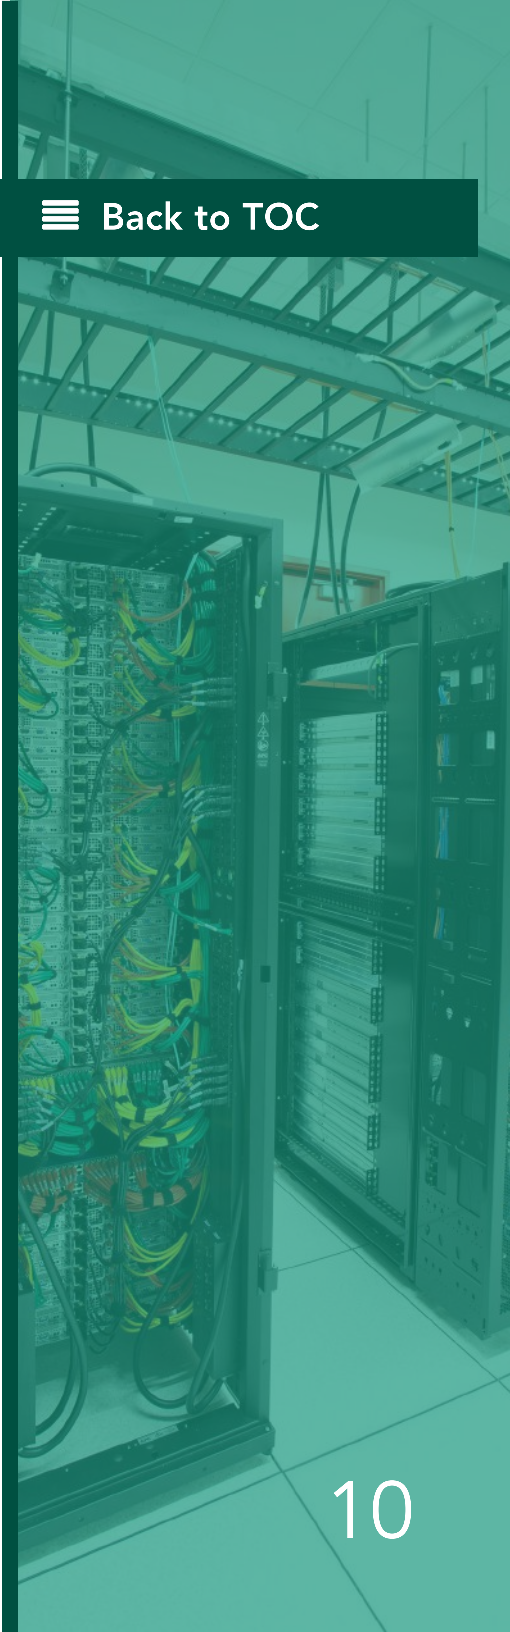

Back to TOC

Back to TOC

The servers, by design, come with two power supply options. In addition, they are backed up by batteries. Therefore, there are two different normal power supplies to each server. That means that the servers would be served from two different substations. To be fully redundant, the substations need to be fed from two different utility lines. In the best-case scenario, the utility lines have a tie between them at some point in the electrical distribution system, and each utility line has enough capacity to carry the entire load of the data center. This scenario describes a fully redundant, normal power data center (see Figure 3).

Figure 4: Although data center design can be complex, the completed project can be efficient, reliable, and robust if designed well. Courtesy: Jacobs Engineering

The normal power redundancy is very important, but it is not enough by itself. The normal power is often backed up by a standby system. The standby system is generally composed of generators, which could be diesel, natural gas, or a hybrid. Diesel generators are the preferred type of generation because they are reliable machines and can be easily maintained. Depending on the type of building the data center is housed in, the generators may or may not be part of the life safety system. Nevertheless, the generators are usually set to be ready to back up the power system very quickly, usually in 10 to 30 sec. The time depends on how long the server backup batteries can last.

Back to TOC

Final thoughts

Although designing a data center's electrical distribution system may seem straightforward, there are inherent challenges. The electrical engineer must:

- Work closely with the owner to determine current and future data center capacity.
- Work with the owner to decide which data center Tier would be appropriate for the client's needs.
- Work closely with the owner to determine the level of redundancy.
- Design a system simple enough to be easy to operate, but one that is also robust.
- Eliminate single points of failure.
- Design a very efficient system with the goal of achieving a PUE under 1.5.
- Apply the relevant industry codes and regulations.

Designing data centers is complex (see Figure 4). Building data centers is very expensive, as is their operation and maintenance. Continuous collaboration with the owner is extremely important—more so than in any other type of project. The successful completion and implementation of the design depends on that collaboration.

Eduard Pacuku, PE, Jacobs, Philadelphia

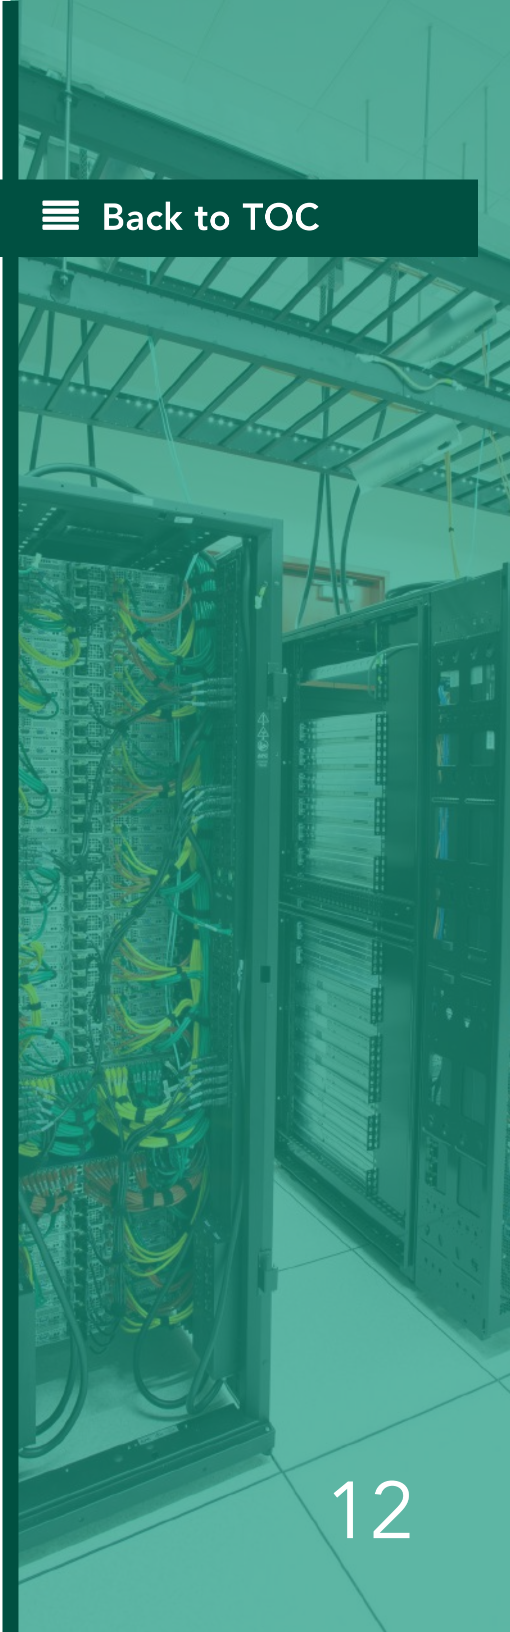

Automated Logic's WebCTRL® building automation system and Nlyte's Asset Optimizer®

Maximize Uptime

Meet Sustainability Goals

Support Increased Workloads

Reduce Costs

Improve Efficiency

CARRIER HVAC SOLUTIONS
DESIGNED FOR THE DATA CENTER

Turn to Carrier for purpose-built solutions and expertise to help you face all your data center challenges with confidence.

Data centers have never been more strategic or critical than they are today. They have also never faced greater scrutiny regarding efficiency, availability and flexibility. As your data center evolves, we're here to help you operate with optimized performance at every stage of the lifecycle.

Management Systems
Control building and IT infrastructure for greater efficiency and optimization with Automated Logic's WebCTRL system and Nlyte's DCIM software solutions

Equipment
Delivers essential cooling while reducing both your energy consumption and carbon footprint

Service and Support
Ensures peace of mind and helps enable continuous operations around the clock, and around the globe

Connect with us at **carrier.com/datacenters** to learn more about these and other lifecycle solutions designed for the data center:

MGMT SOFTWARE

EQUIPMENT

SERVICE & RENTALS

Carrier

carrier.com/datacenters

Data Center Sustainability… Searching for the Oasis

"When the well's dry, we know the worth of water."

– Benjamin Franklin

Integrating Dry Cooler Technology to Address Water Usage Effectiveness (WUE) Concerns

Throughout time—regardless of geographic location— water has always been held as a precious resource. However, the current combination of global warming, drought and ever-increasing consumption are not only creating previously non-existent water-stressed areas but have made uninterrupted access to water a community-based 'call-to-arms' to restrict many commercial developments. And nowhere has this been more apparent than in data centers that provide shared access to applications and data using complex networking, computing, and storage infrastructures.

There is little argument that data centers play a critical role in the future of our world's economic growth. Of the approximately 8,000 data centers currently in the world, 33% are located in the United States.[1] Whether it's an Enterprise, Managed Services, Cloud-based, Colocation, Edge or Hyperscale data center, two of the major operational certainties they all share are extremely high demand for power and cooling. Carrier is providing both advanced-technology equipment and digital lifecycle solutions to support the unprecedented growth and criticality of data centers. More than 300 data center

Back to TOC

owners and operators with over one million racks, spanning enterprise, colocation and edge benefit from Carrier's optimization solutions across their portfolios.

Data centers are one of the most energy-intensive building types, consuming 10 to 50 times the energy per floor space of a typical commercial office building. Collectively, data centers account for approximately 2% of the total U.S. electricity usage.[2] The current trend towards server consolidation (rack vs. blade servers) and virtualization may help reduce hardware within data centers but does not always reduce energy consumption. As for water usage, a large data center can consume anywhere between 1 and 5 million gallons a day—as much as a town of 10,000 to 50,000 people.[3] Water is used for a data center's cooling system, which ensures that the heat produced by thousands of servers within these massive facilities is precisely controlled so that their internal servers can run uninterrupted 24 hours a day, 7 days a week. Additionally, data centers indirectly consume large amounts of water off-site at a power generation plant to supply the data center with electricity.

As data networking, computing and storage demands continue to surge—and as water becomes scarcer due to climate change—the development of data centers is attracting greater scrutiny and resistance by both local residents, municipalities and governments throughout the US. Additionally, addressing environmental issues such as Power Usage Effectiveness (PUE), Carbon Usage Effectiveness (CUE) and Water Usage Effectiveness (WUE) are becoming mission-critical for all data center owners and operators.

Technical References

Historical

- **PUE** has been used extensively in data center research, development and evaluation as one of the most crucial indicators to gauge system effectiveness. It is defined as

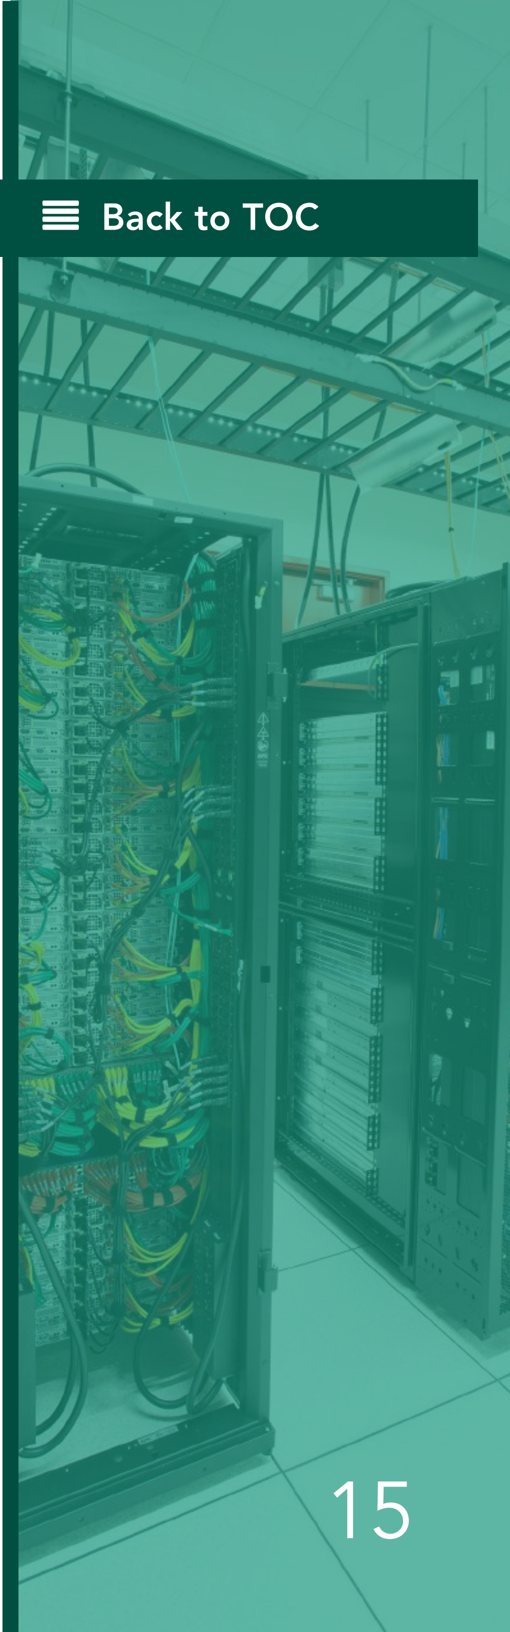

the ratio of a data center's overall energy consumption to the energy used just by its IT equipment.[4]

Driven by New Sustainable Efforts

- **WUE** in recent years has become an important focus for data center operations due to the continued advancement of sustainability efforts. It is a value identified as the ratio between the use of water in a data center system that consists of water loops, adiabatic towers, humidification, water-driven energy production, etc. and the energy usage of the IT equipment.[4]

- **CUE** aims to gauge its sustainability by means of pollutant emissions as an outcome. It is the relationship between the CO_2 emissions generated by the data center and the energy usage of IT components.[4]

Traditional Data Center Cooling

Maintaining and managing precise temperatures within data centers is crucial to maintaining the functionality and life of IT equipment. IT equipment temperatures which rise above set parameters can result in expensive downtime or equipment failure.

Traditionally, **water-cooled chillers** are paired with cooling towers that evaporate water to reject the heat generated by servers, IT equipment and mechanical infrastructure found within the data center. This type of cooling system removes and releases heat produced inside a data center into the outside environment, through a cooling tower that uses a water evaporative process. Cooling towers are heat exchangers that use water and air to transfer heat from a data center's cooling system to the outdoor environment. Most commonly, they are used to remove heat from the condenser water leaving a chiller.

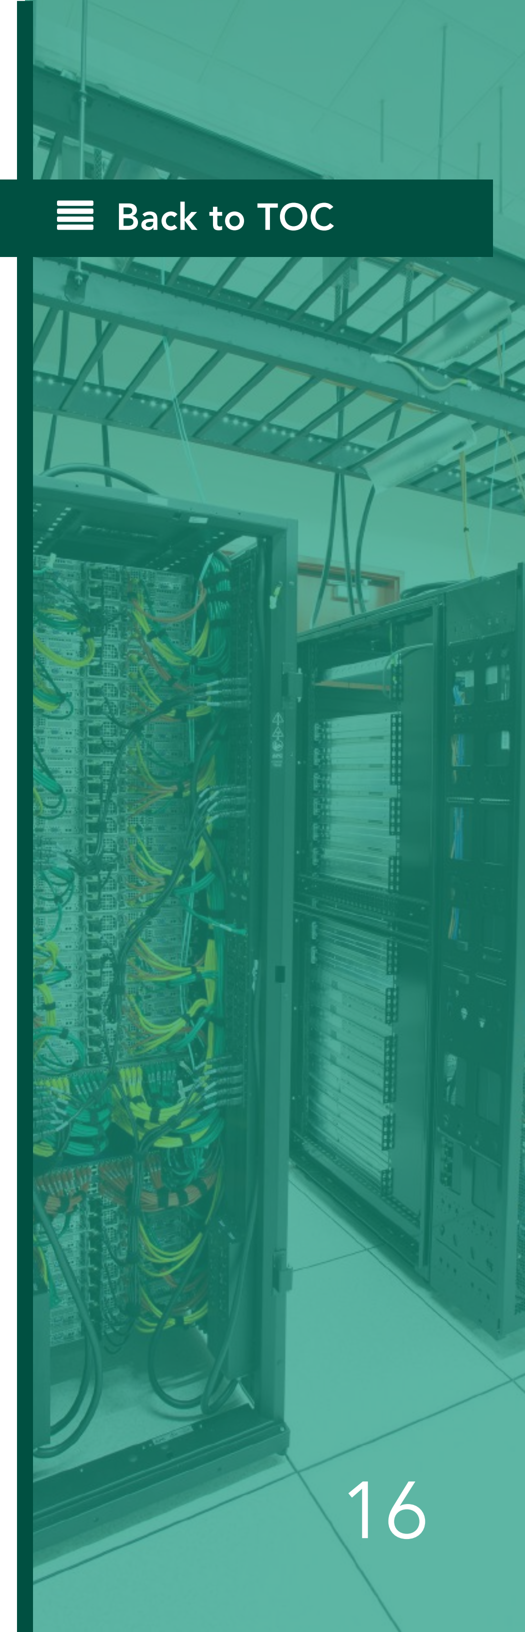

Back to TOC

Back to TOC

As an alternative to water-cooled chillers, stand-alone **air-cooled chillers** are often used to eliminate water loss from cooling towers. With limited cooling tonnage capacity, this equipment is often found in either smaller data center facilities or large data centers that utilize 100+ air-cooled chillers to meet the load. Air-cooled chillers are a lower-cost packaged product where as water-cooled chillers are available in larger capacities (4000+ tons with a single compressor), resulting in fewer units to maintain.

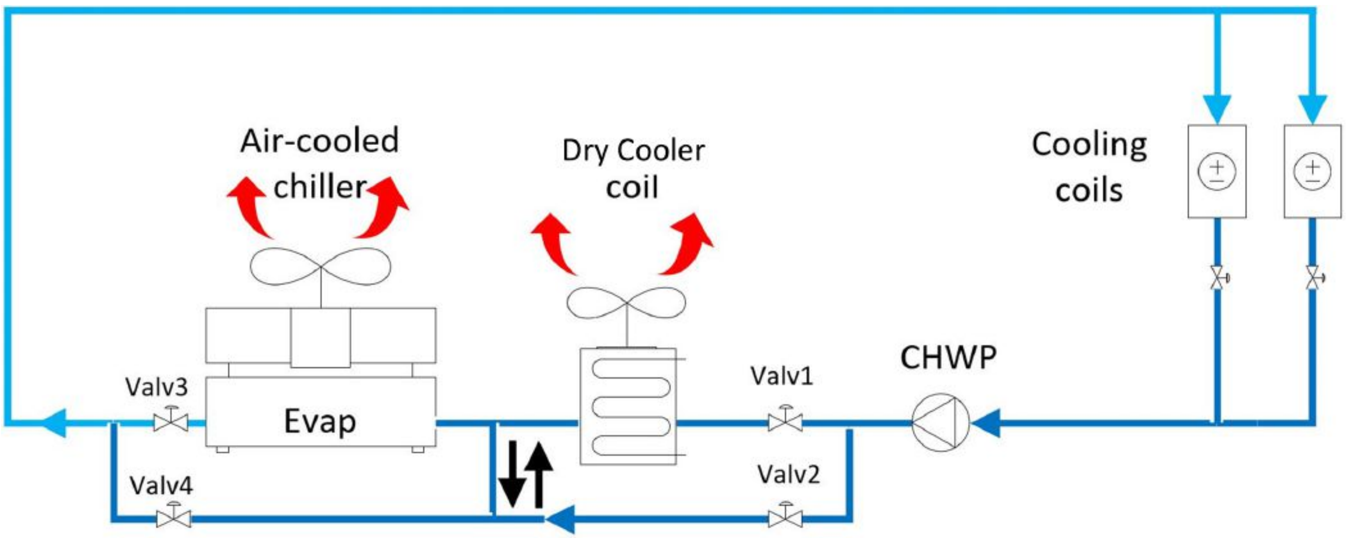

Dry Cooler Technology – Helping Meet ESG and Sustainability Goals

Environmental, social and corporate governance (ESG) is a business framework for considering environmental issues and social issues in the context of corporate governance. Data center owners and operators are acutely aware of the importance of creating and meeting ESG/Sustainability goals within their facilities, and continually explore innovative technologies to accomplish them. As climate change and diminishing natural resources gain importance, solutions to help reduce water and power consumption are of paramount concern.

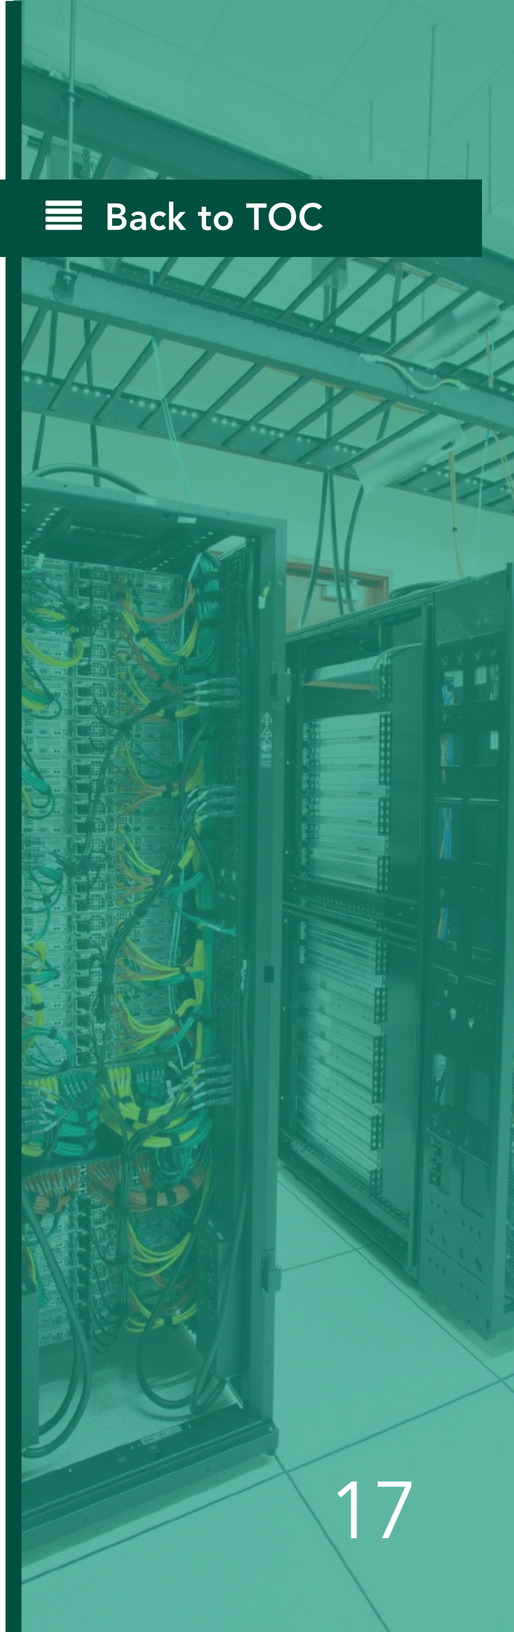

Back to TOC

"Data center operators have made great strides in power usage effectiveness over the past 15 years," said Michel Grabon, Director of Data Center Solutions, Carrier. "Continual technology advances with higher powered server processors present power-consumption and cooling challenges requiring specialized solutions that Carrier provides."

To help meet the cooling challenges relative to water usage reduction, **Carrier Dry Coolers** are a closed-loop system that does not evaporate water. They remove excessive heat from the cooling liquid and reject it directly into the atmosphere. Fans blow air over a heat exchanger which then cools down the liquid. Additionally, Carrier Dry Coolers eliminate water treatment requirements and Legionella concerns since they don't use any water or evaporative cooling. Hence, water conservation and lower maintenance are the key considerations for using Carrier Dry Coolers.

Unlike cooling towers, Carrier Dry Coolers do not evaporate any water which is key for data center owners and operators due to public scrutiny. They also positively impact corporate sustainability goals such as Water Usage Effectiveness (WUE) which data centers are aiming to get as low as possible. Dry Coolers, compared to air-cooled chillers, have the same WUE but better PUE. Data centers have a constant heat load from the servers and evaporate enormous amounts of water with cooling towers. Installing Carrier Dry Coolers in conjunction with a water-cooled chiller is more efficient and doesn't utilize water.

A very real concern in many parts of the US today is the future availability of water. Data center owners and operators who are considering water-cooled chillers with evaporative cooling towers in dryer areas may run the risk of water availability. For example, if the life of a data center's cooling plant is 20 years, the concern for data center operators in

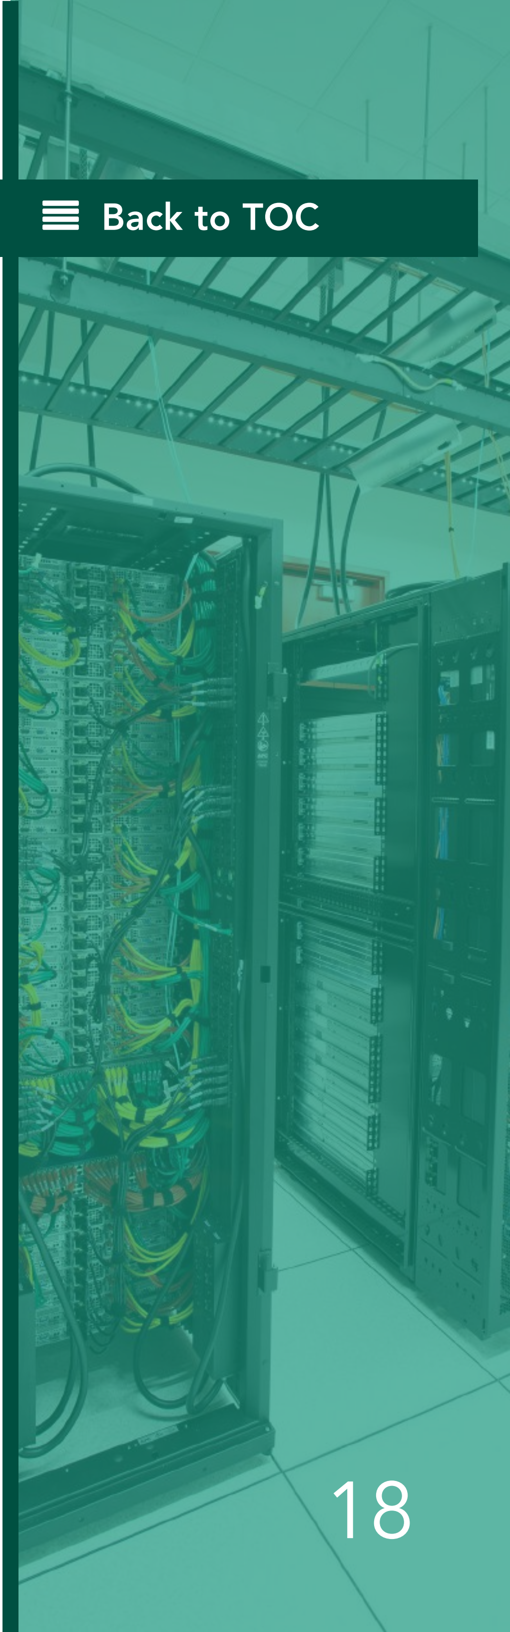

Back to TOC

those areas is if they'll be able to get water, and in 10 years' time, what the price will be for that water to be able to continue to operate their type of system.

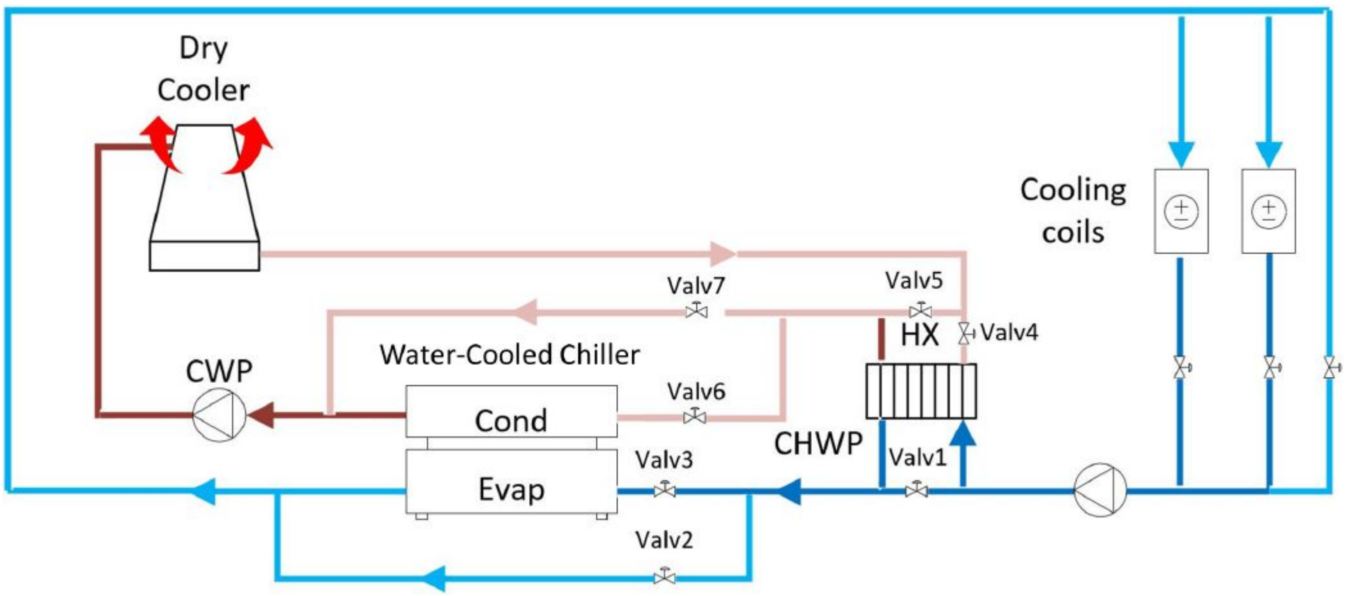

Sustainability / ESG Issues

Cooling/evaporative towers use chemical treatment for water and have 'blow-down,' which is the result of evaporating so much water that you get sludge which is then blown down the drain. Over time, correcting this requires time-consuming maintenance costs. The treated water is also emitting contaminants into the air when the evaporator is running.

In a cooling/evaporative tower, the condenser water loop is open to ambient air, so any airborne contaminants can ultimately end up back in chiller tubes. Over time, this condition can reduce a chiller's efficiency. Add to this the downtime, maintenance and cleaning costs to remove the built-up scale and grime on the tubes. Conversely, a dry

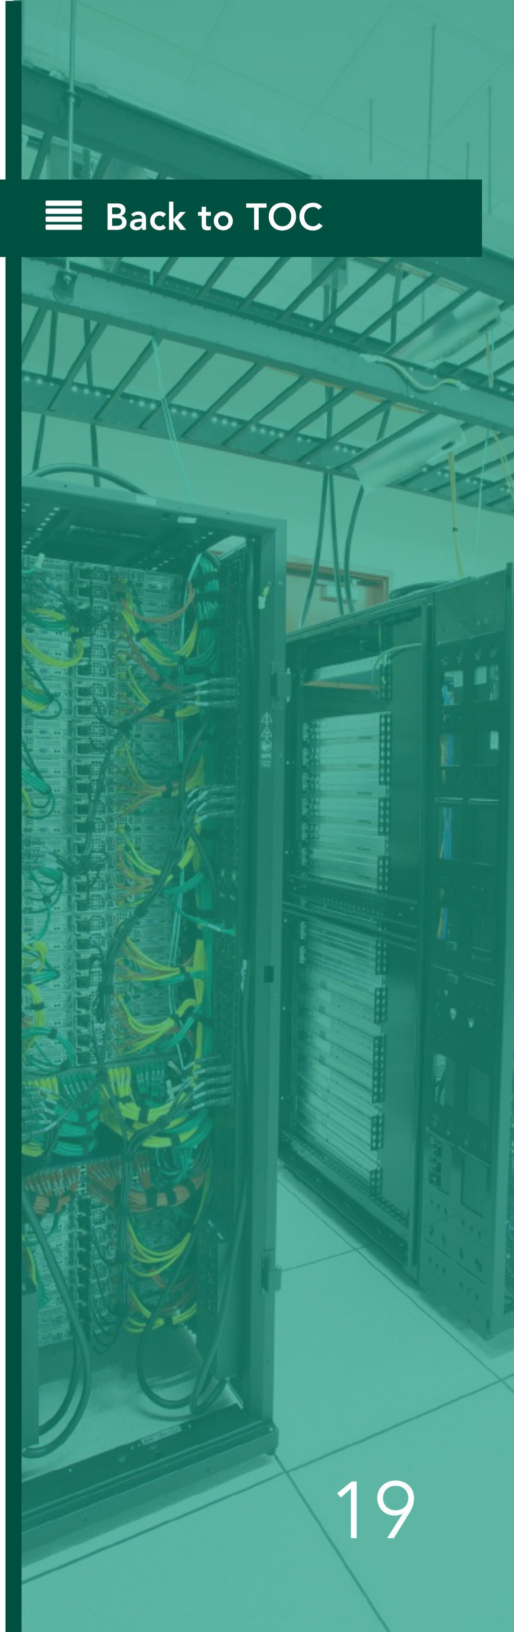

Back to TOC

cooler is a closed loop system designed to eliminate water loss. This feature proves critical when addressing concerns often found in water-stressed areas. Additionally, a dry cooler's fresh water remains clean which contributes to less chiller maintenance required.

Features and Benefits of Carrier Dry Coolers

Dry Coolers remove excessive heat from the liquid and reject it directly to the atmosphere. Unlike evaporative cooling towers, Carrier Dry Coolers do not require water. Instead, fans blow air over heat exchangers and thus, cool down the liquid.

- As a closed loop system, Dry Coolers eliminate water treatment requirements and Legionella concerns since they don't use any water for evaporative cooling. Hence, water conservation and lower maintenance are the key advantages of using Dry Coolers.

- Dry Coolers installed in conjunction with water-cooled chillers allow for fewer chillers due to capacities up to 4000+ tons being available in a single compressor versus an average of 300 tons being available in a single compressor of an air-cooled chiller.

- Carrier Dry Coolers utilize a unique fan design that allows for a smaller footprint that delivers significant cooling tons per square foot when compared to competing technologies.

- When using a Dry Cooler, the owner is able to operate in both full and partial economization mode without the efficiency loss of a secondary heat exchanger due to both loops being closed loops versus an open evaporation cooling tower loop.

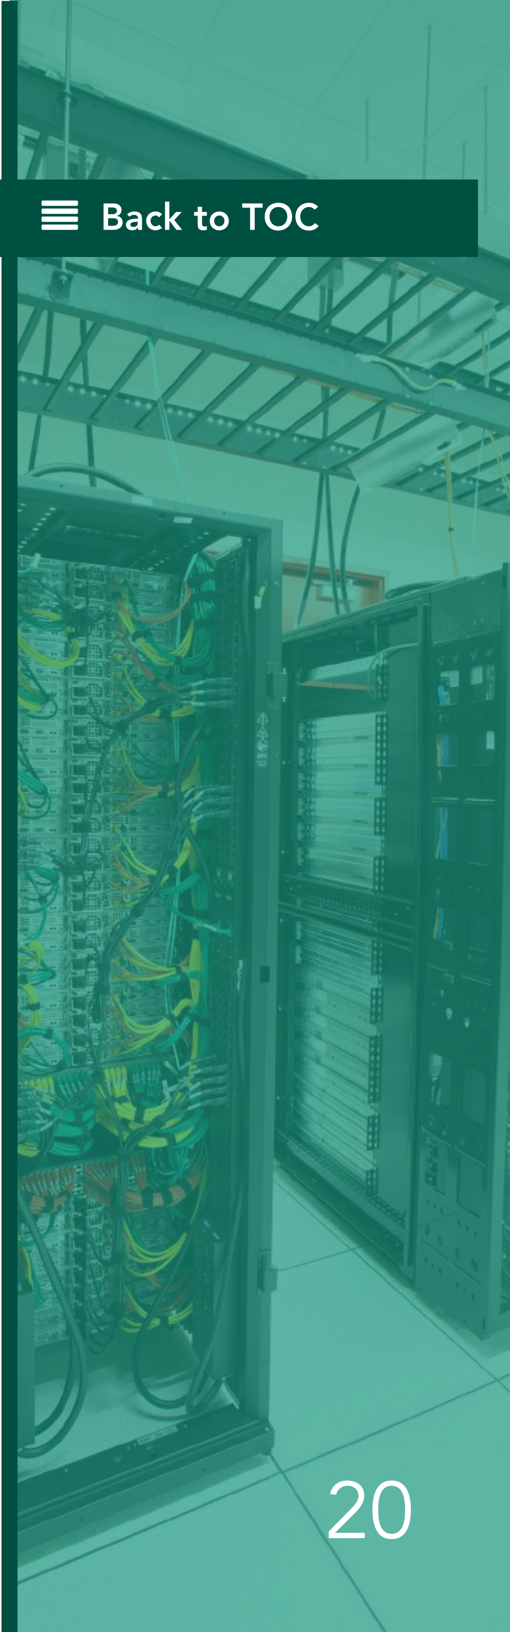

Back to TOC

Smaller Footprint

With available square footage on data center roofs at a premium, Carrier Dry Coolers incorporate a unique modular design to facilitate stacking over water-cooled chillers to better utilize space. This convenient separation of components allows operators to save vital roof space for other equipment.

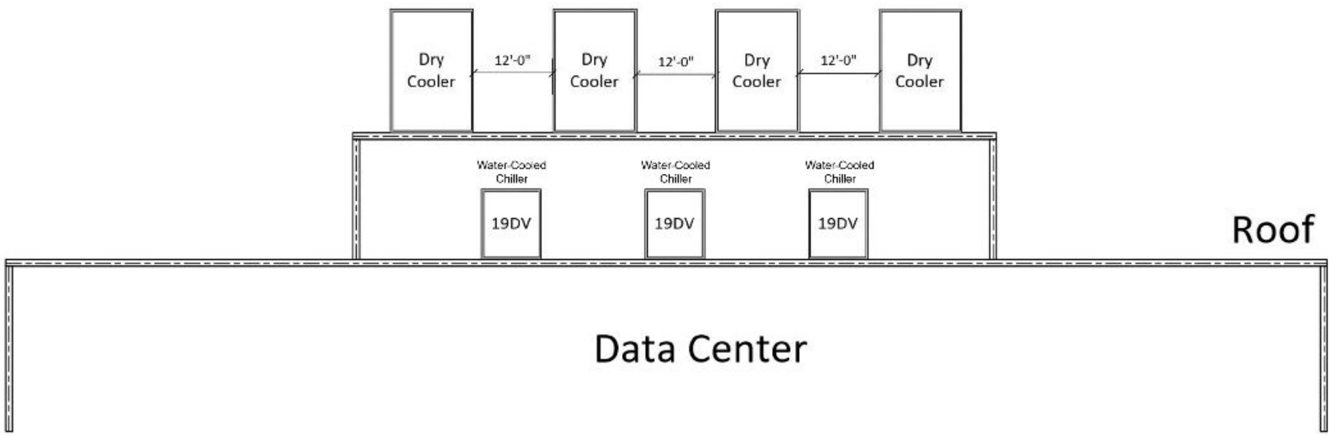

Supports Water-Positive Initiatives

Carrier Dry Coolers do not require a constant water supply as opposed to evaporative cooling towers. This feature eliminates water supply, disposal and maintenance issues. When compared to air-cooled chillers the water-cooled chiller with Dry Coolers can support a reduction in PUE that can reduce the amount of water utilized in the production of electricity at the power plant serving the data center.

Design Flexibility

Carrier Dry Coolers can be custom designed for any size data center. In addition, due to the custom design, additional sound attenuation is available to meet the overall community sound pollution levels.

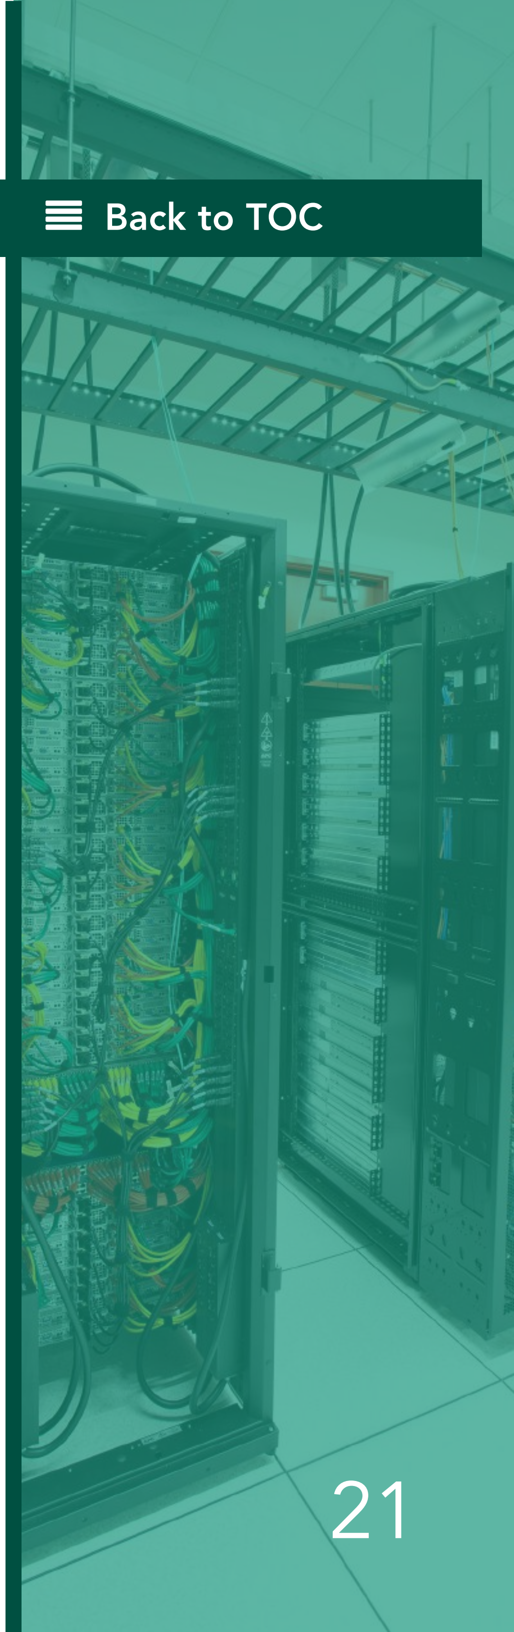

Back to TOC

Easy Installation and Start-up

Carrier Dry Coolers arrive on site 'installation ready' and are wired with product-integrated controls for quick connection to any building automation system.

Controls

Carrier Dry Coolers can be easily programmed to operate efficiently year-round, even in geographical regions that experience lower temperatures.

Efficiency

The fan technology utilized allows for a smaller footprint and delivers the most industry-leading cooling tons per square foot. This feature also contributes to low operating costs after initial installation. In addition, the heat sent to the Dry Cooler can easily be used as a low-grade heat source for building heating, district heating and other process needs to improve the overall system efficiency.

LEED and Local Utility Rebates

In certain installations, Carrier Dry Coolers positively impact data centers that are applying for LEED certification. Additionally, there may be opportunities to receive local utility rebates.

Life Cycle

As demonstrated in the above discussion utilizing Dry Coolers with water-cooled chillers can have distinct benefits when compared to the other common design methods used today.

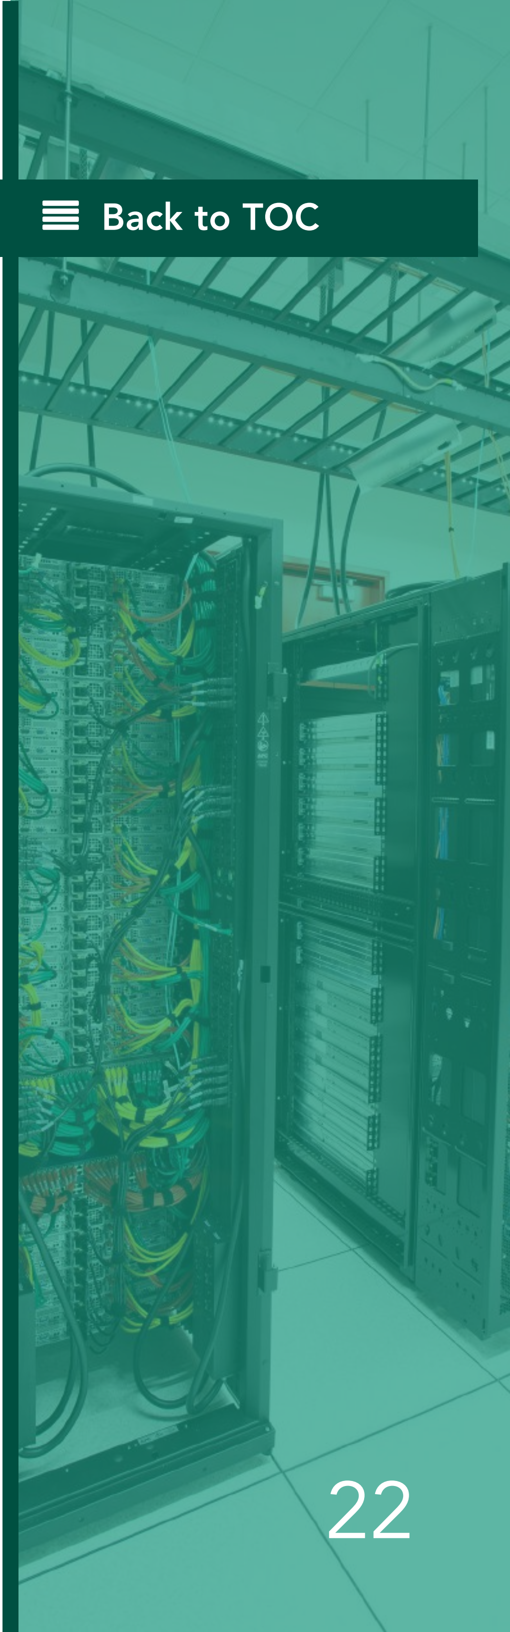

- When compared to using direct evaporative air handling units, this solution eliminates water usage and the introduction of particulate in the airstream serving the data hall.

- Larger capacity water cooled centrifugal compressors available up to 4000+ tons can significantly reduce the number of compressors utilized, compared to air-cooled chillers, and can also reduce the PUE.

- The elimination of evaporative cooling towers in a traditional water-cooled chiller design eliminates the water usage (improved WUE), chemical treatment and tube cleaning.

For future data center designs, the design team should investigate whether utilizing Dry Coolers with water-cooled chillers would provide a preferable solution versus the traditional designs used in the past. Newer, centrifugal compressor designs that allow for the higher condensing temperatures required for dry-cooler operation, have only recently become common in the market, allowing for this new design concept to become a reality.

Lifecycle Solutions for Complete Data Center Optimization

Data centers have never been more strategic or critical than they are today. They have also never faced greater scrutiny regarding sustainability, efficiency and flexibility. To help address these issues, Carrier delivers purpose-built solutions and expertise to help face all data center challenges with confidence.

Carrier Dry Coolers and our full breadth of HVAC equipment deliver essential cooling while reducing both energy consumption and carbon footprint. Our

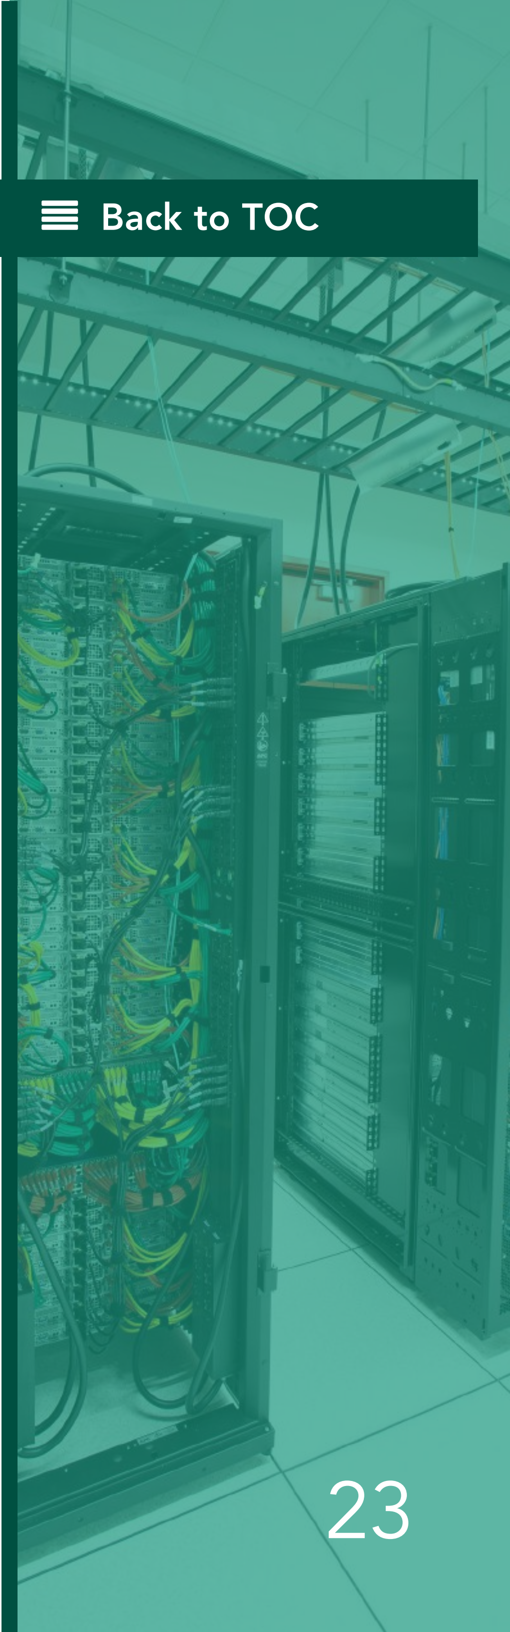

Back to TOC

building management systems integrate building and IT infrastructure to provide greater visibility and optimization. Our service and support ensure peace of mind and enable continuous operations. And as a truly global partner, we're there wherever and whenever you need us. It all adds up to the confidence needed to operate with optimized performance at every stage of the data center lifecycle—and at Carrier, it's delivered by design.

Carrier's range of smart and connected solutions deliver upstream data from the data center ecosystem to cool, monitor, maintain, analyze and protect the facility to meet green building standards, sustainability goals and comply with local greenhouse gas emission regulations. Carrier's Nlyte data center infrastructure management (DCIM) tools share detailed information between the HVAC equipment, power systems and servers/workloads that run within data centers, providing unprecedented transparency and control of the infrastructure for improved uptime.

Carrier's purpose-built solutions are integrated across its solutions portfolio with efficient and high-performing HVAC equipment, DCIM tools and building management system to help data center operators use less power and improve operating costs and profitability for many years.

Complete Lifecycle Service and Support Specifically Designed for Data Centers

BluEdge® Service Solutions

The BluEdge service platform delivers expertise, service, and greater peace of mind throughout the lifecycle of your HVAC equipment. Our service tiers offer a range of

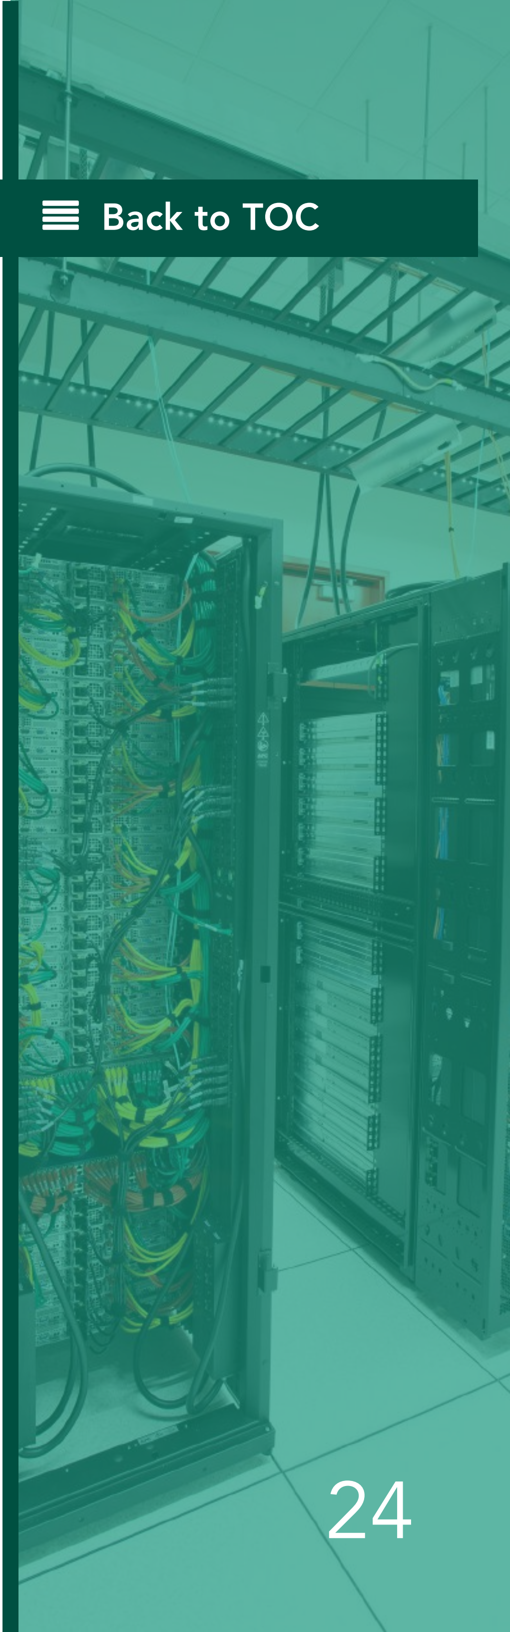

Back to TOC

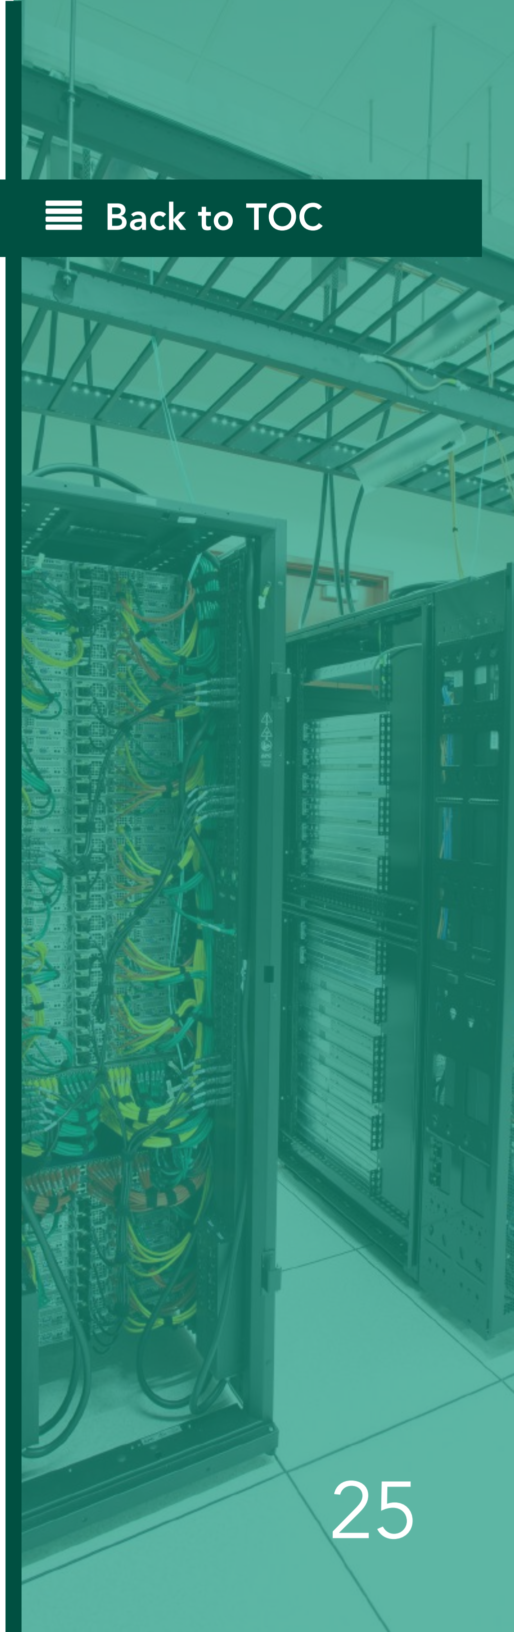

options, all powered by Abound™ HVAC Performance—your source for actionable insights that can reduce operating costs, increase efficiency and improve uptime.

Advanced Analytics Solution for Predictive Maintenance

Abound Predictive Insights™ solution provides improved asset management across your building portfolio. This enables a simpler, smarter way to analyze your building operations and optimize efficiency, comfort and performance. The Abound technology is based our award-winning CORTIX™ artificial intelligence (AI) platform. This advanced AI platform can predict failures and maintenance issues long before they manifest into breakdowns, while also helping pinpoint efficiency improvements.

Data Center Infrastructure Management (DCIM)

Carrier's Nlyte® Asset Lifecycle Management and Capacity Planning software provides automation and efficiency to asset lifecycle management, capacity planning, audit and compliance tracking. It simplifies space and energy planning, easily connecting to an IT service management system and all types of business intelligence applications.

Carrier Rental Systems

Carrier Rental Systems can help you maintain data center uptime when temperature control is interrupted by scheduled upgrades or unforeseen events. With customized plans according to your timeline, budget and available space, you can use our readily available rental units to ensure continuous operations while you select, order, and receive new equipment.

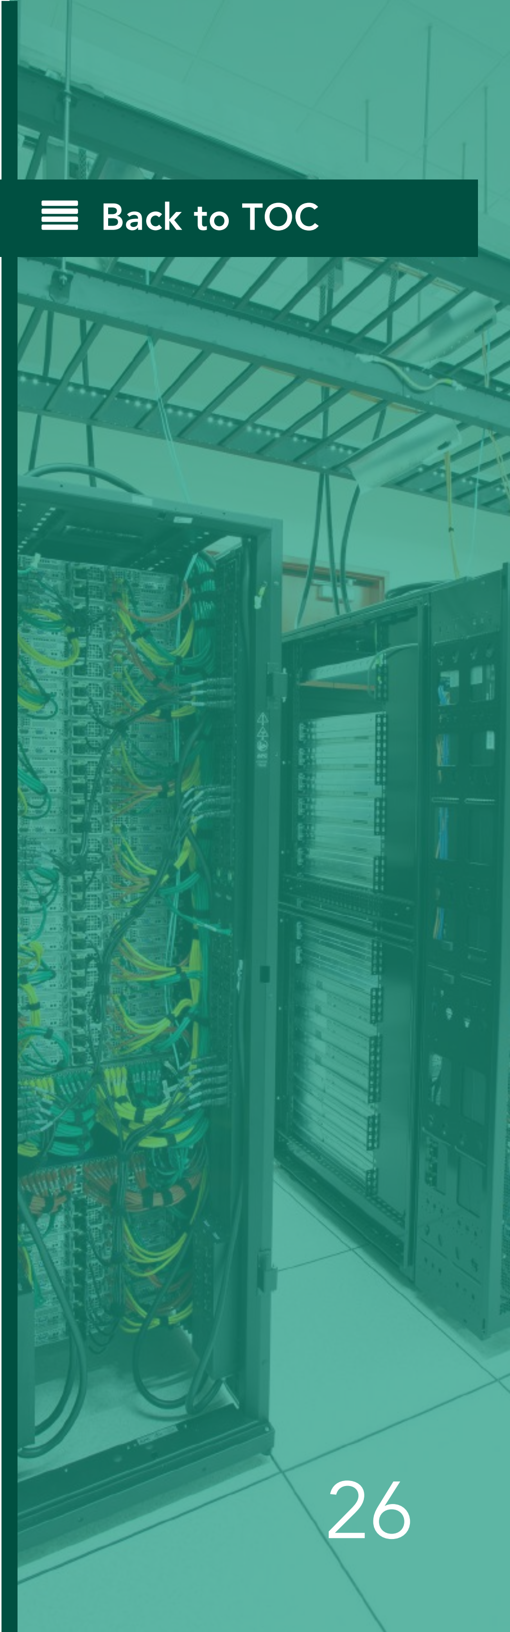

Back to TOC

Citations:

1. https://techjury.net/blog/data-center-statistics/

2. https://www.energy.gov/eere/buildings/data-centers-and-servers

3. https://www.datacenterknowledge.com/sustainability/new-front-water-wars-your-internet-use https://submer.com/blog/pue-cue-and-wue-what-do-these-three-metrics-represent-and-which-is-one-is-the-most-important/

4. https://submer.com/blog/pue-cue-and-wue-what-do-these-three-metrics-represent-and-which-is-one-is-the-most-important/

The development of ASHRAE 90.4: Energy Standard for Data Centers

Back to TOC

ASHRAE 90.4 is a new energy standard for data centers that calls for meeting minimum efficiency requirements.

When addendum "bu" was added to ASHRAE 90.1-2007: Energy Standard for Buildings Except Low-Rise Residential Buildings in 2010, it caught the data center industry off-guard. While data center operators had a considerable amount of representation on the ASHRAE TC9.9: Mission Critical Facilities, Data Centers, Technology Spaces and Electronic Equipment technical committee that shapes the environmental standards for data centers, the ASHRAE 90.1 Standing Standards Project Committee (SSPC 90.1) was something else altogether. This addendum added significant prescriptive requirements to ASHRAE 90.1 for air- and water-side economizers in data center HVAC systems. Up until that point, data center HVAC systems were effectively exempt from energy code requirements. The major players in the data center industry reacted strongly to these changes and the resulting firestorm served to help shape future standards development, namely ASHRAE 90.4-2016 Energy Standard for Data Centers.

What was 90.1-2007 addendum bu?

Addendum bu added a new definition to ASHRAE 90.1: "computer room." This definition was as follows:

A room whose primary function is to house equipment for the processing and storage of electronic data and that has a design electronic data equipment power density exceeding 20 W/sq ft of conditioned floor area.

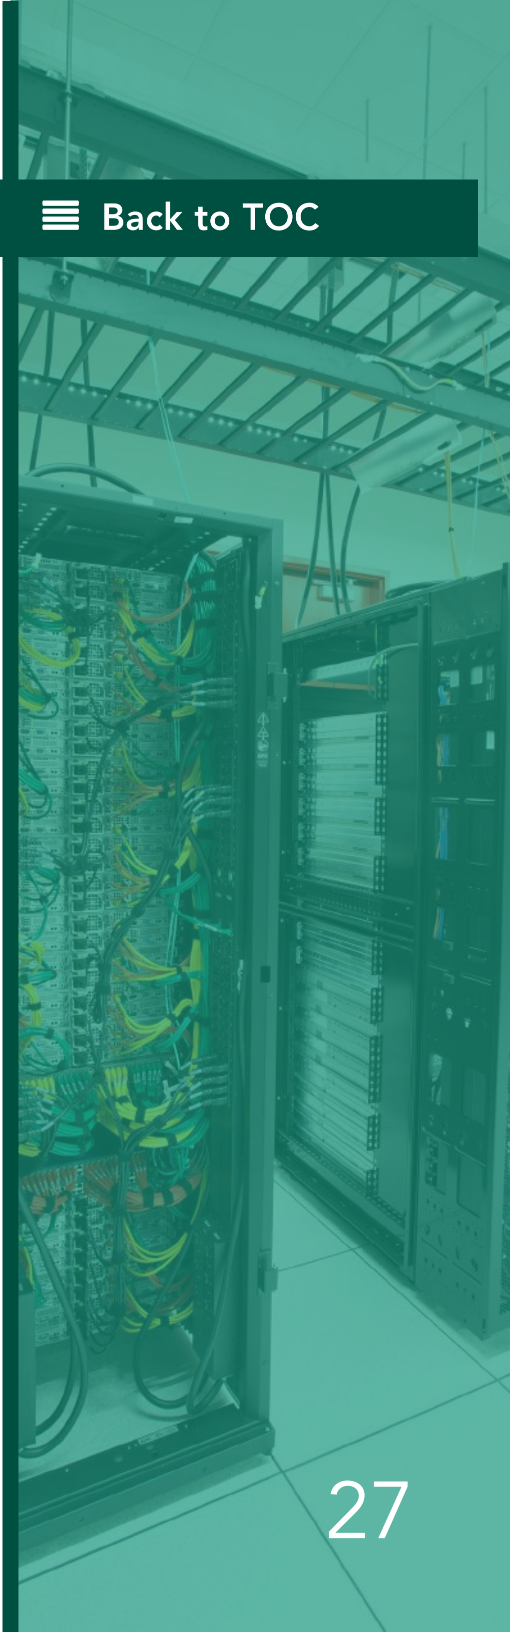

Back to TOC

This definition has been altered substantially in recent versions of the code, which will be discussed.

The major change attributed to addendum bu was the addition of economizer requirements for cooling systems with fans that serve computer rooms. This more or less aligned data centers with requirements for HVAC systems in other types of buildings and affected sections 6.4.1.1, 6.5.1, and Table 6.8.1H within 90.1-2007. There were several exceptions to the economizer requirement including when:

- The total combined design load of all computer rooms in a building is less than 3,000,000 Btu/h (250 tons) and not chilled-water cooled.
- If chilled-water cooled, the room design load is less than 600,000 Btu/h (50 tons).
- Less than 600,000 Btu/h (50 tons) of computer room cooling is being added to an existing building.
- The authority having jurisdiction (AHJ) does not allow cooling towers.
- Where at least 75% of the design load is critical in nature (i.e., NEC Article 708: Critical Operations Power Systems, Tier IV data centers, data centers that process financial transactions, etc.).
- The economizer provision affected every climate zone except for zones 1A and 1B (in the United States, typically only tropical regions, such as the southern tip of Florida, Puerto Rico, Guam, and Hawaii are in these climates zones).

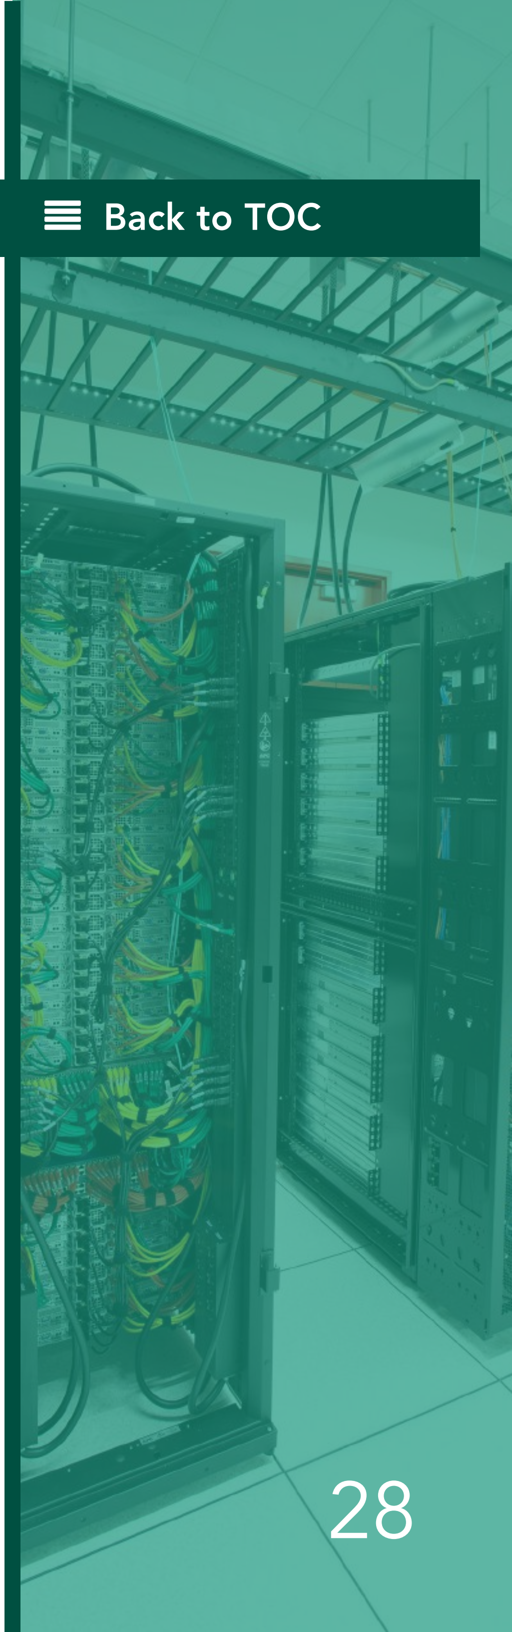

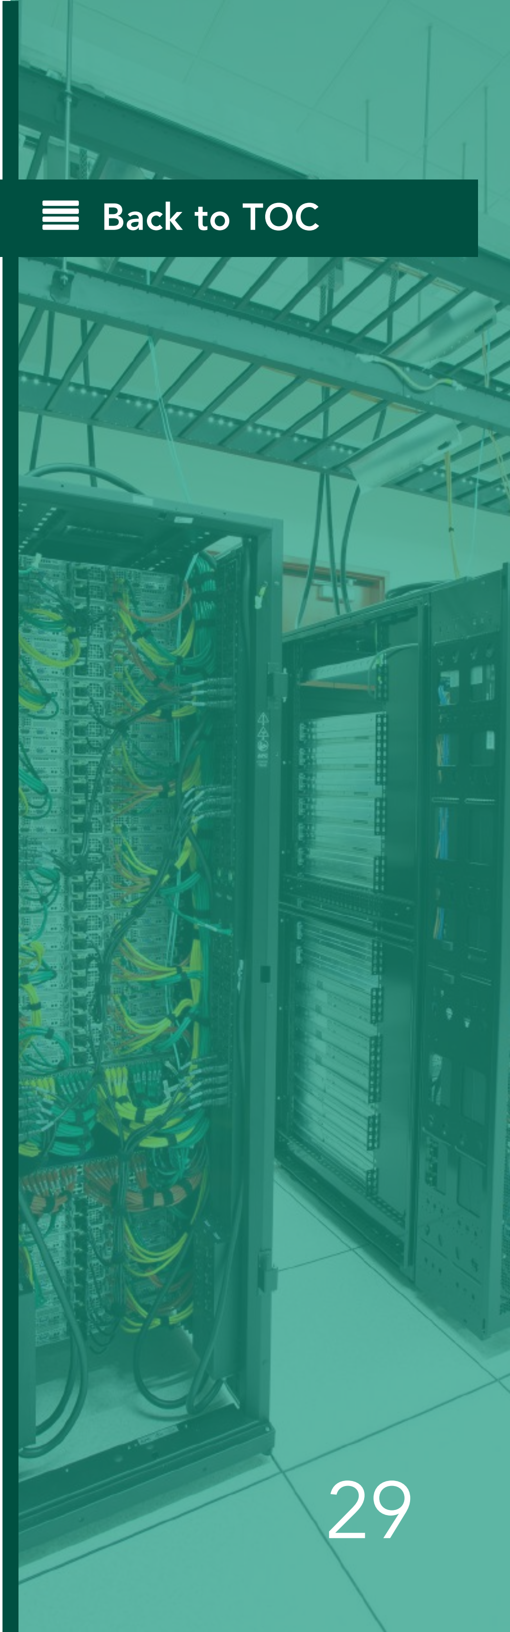

Back to TOC

Why did the data center industry object to bu?

While the foreword to the addendum clearly outlined ASHRAE's justification for the new requirements and exceptions contained in it, the data center industry still had unusually strong objections. The primary issue was that the requirements were "prescriptive," mandating specific design solutions. The data center industry's argument was that in a product sector experiencing explosive growth/change, they should not be limited with prescriptive requirements that stifle technical innovation. Rather, the preference was that the requirements be performance-based and focus solely on quantifying energy efficiency, not the exact methods used to achieve that efficiency.

The data center industry's most public response to addendum bu was published on Google's public policy blog in April 2010. Google traditionally reserves use of this widely read forum for official commentary on privacy, net neutrality, anti-trust, and similar major regulatory topics. This open-letter response was signed by the who's who of the data center industry. ASHRAE did issue a formal response to the open letter, maintaining that the new exceptions in combination with alternative compliance paths, such as the energy cost budget method, reasonably addressed those concerns. By this point, however, the lines had already been drawn between the two groups.

Why did efficiency requirements suddenly apply to data centers?

Like all new codes and standards, nothing happens quickly. Changes are generally the result of a long, deliberate sequence of events. ASHRAE has long considered addressing data center energy efficiency in ASHRAE 90.1. In August 2007, the EPA issued a report on data center energy efficiency to the U.S. Congress. The key takeaway from this report was that the nation's data centers were responsible for about 1.5%

Back to TOC

(61 billion kWh) of total US electrical consumption in 2006. The report further forecasted that this use would double by 2012. Now with the issue quantified, the wheels of change could be set in motion. Given this dramatic projected increase in energy usage, adoption of new energy efficiency requirements would be justifiable per the key provisions of federal energy efficiency legislation in effect at that time, the Energy Policy and Conservation Act of 1975 (EPCA) and the Energy Policy Act of 1992 (EPAct).

Those projections ended up being wrong. Based on a Lawrence Berkeley National Laboratory report issued 9 years later in 2016, usage did not increase nearly as dramatically as predicted. Instead of doubling, data center electrical usage was estimated at about 1.8% (70 billion kWh) of total 2014 U.S. electrical consumption. This represents only a 15% increase from 2006 levels. Ironically, the forecasting error is mostly attributed to the emergence of hyperscale/cloud data centers—which include many of the signatories to Google's open letter who have since embraced forms of air and water economizers in a significant percentage of their data centers. Google, in fact, has managed to reduce their fleetwide trailing 12-month power-usage effectiveness (TTM PUE) in 2016 to 1.12, even when using fairly conservative metrics.

The emergence of ASHRAE 90.4

Traditionally, energy codes have addressed only HVAC systems used for comfort cooling and heating and left out requirements for systems related to process cooling/heating. With the aforementioned conflicts over new data center efficiency requirements starting with the addendum to ASHRAE 90.1-2007, members of the technical committee for TC9.9 initiated a new standard in response to those perceived shortcomings. After almost 5 years of work, the result was ASHRAE 90.4-2016: Energy Standard for Data Centers.

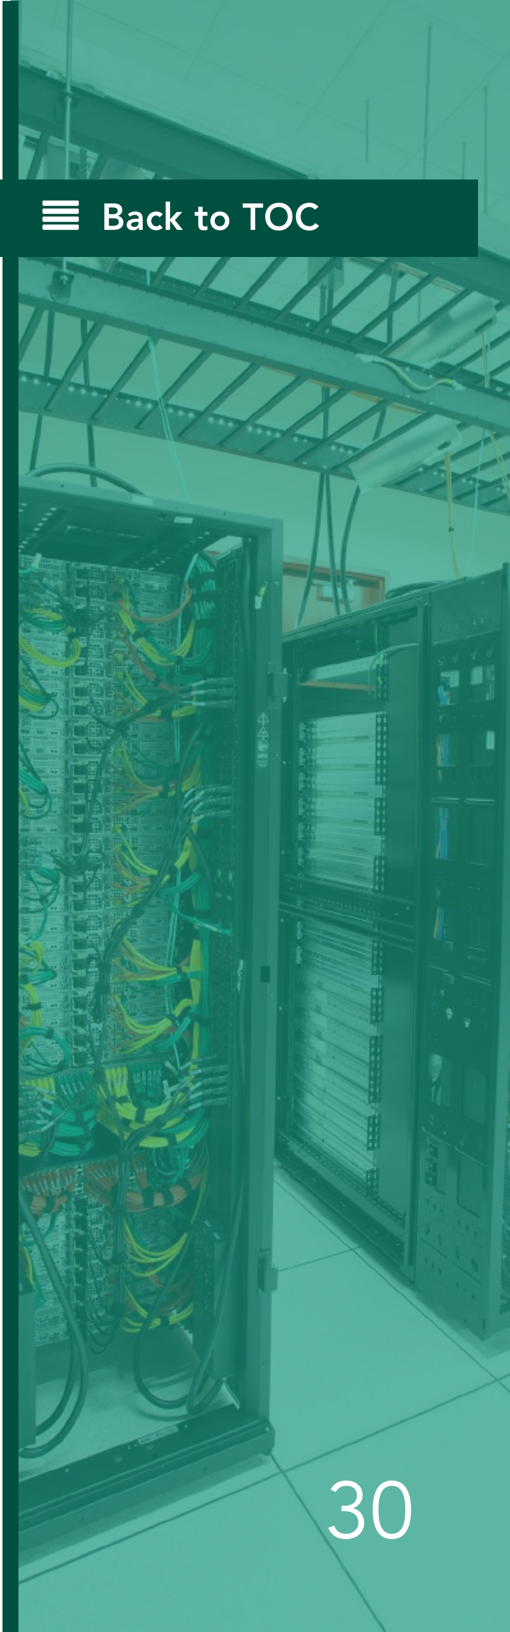

Back to TOC

Even though ASHRAE 90.4 was a consensus standard with input from numerous industry groups, there was still significant heated debate, specifically over minimum acceptable efficiency levels and associated measurement metrics. One of the first hurdles was agreeing on a formal definition for a "data center". As mentioned previously, addendum bu had added an unusually broad definition for a "computer room." Based on that earlier definition, there would be no useful distinction between an intermediate distribution frame (IDF) closet in an office building and a large, hyperscale data center. The proposed addendum cs to ASHRAE 90.1-2010 laid the groundwork for ASHRAE 90.4 by attempting to revise the definition of computer rooms. That addendum also introduced a new definition, "data center," to which ASHRAE 90.4 now applies. Those two definitions have since been revised in subsequent addenda to ASHRAE 90.1. The final definitions as incorporated into ASHRAE 90.4 make a distinction between these two occupancy types by adding power usage thresholds for each. The revised definition in ASHRAE 90.4 for a computer room is as follows:

A room or portions of a building serving an information technology equipment (ITE) load less than or equal to 10 kW or 20 W/sq ft or less of conditioned floor area.

A data center is defined by ASHRAE 90.4 as follows:

A room, building, or portions thereof, including computer rooms being served by the data center systems, serving a total ITE load greater than 10 kW or 20 W/sq ft of conditioned floor area.

This distinction is important. ASHRAE 90.1 still applies to computer rooms but does not apply to the mechanical and electrical distribution systems in data centers. How-

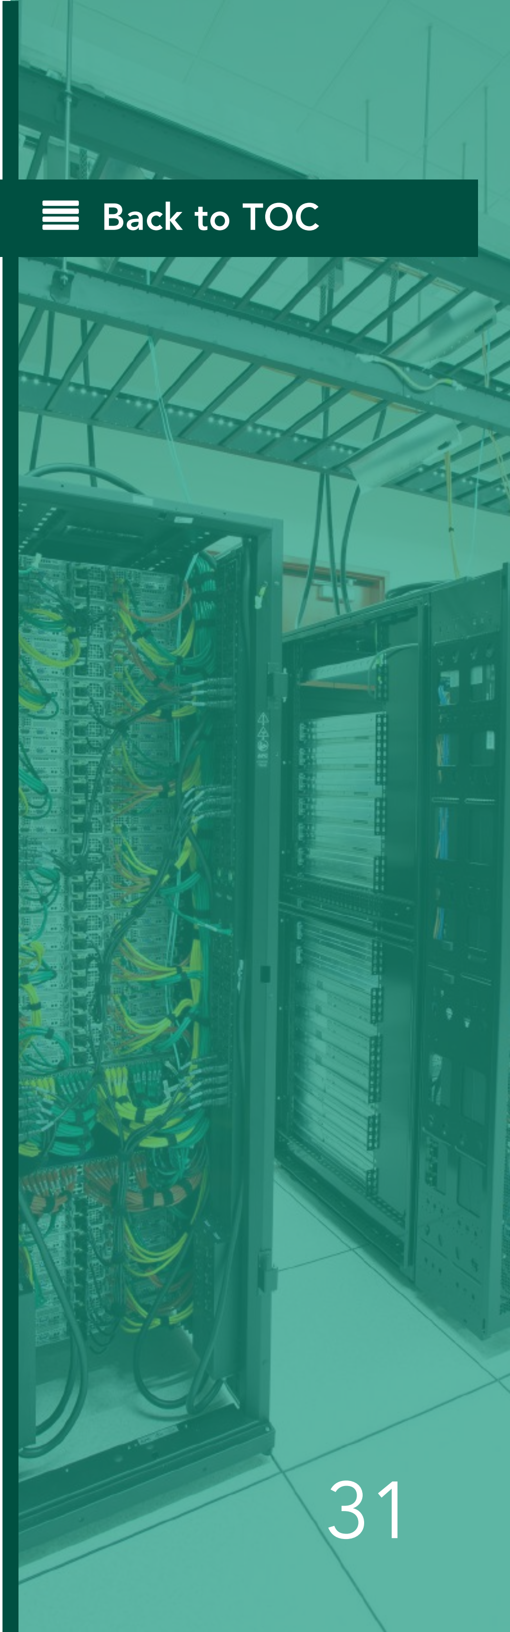

Back to TOC

ever, ASHRAE 90.1 still applies to other portions of data centers, namely envelope, service-water heating, and lighting. Neither standard applies to mechanical and electrical equipment in telephone exchanges or essential facilities (i.e., NEC Article 708 Critical Operations Power Systems, Tier IV data centers, data centers that process financial transactions, etc.).

Efficiency is meaningless if you can't calculate it

A primary point of contention during the writing of ASHRAE 90.4 was how to quantify minimum efficiency requirements. The Green Grid Association, a consortium of data center industry companies that works to improve the efficiency of data centers worldwide, developed a metric known as power usage effectiveness. PUE is a ratio that quantifies the relationship between the energy specifically used by ITE and the total energy used by a data center. It is the most widely accepted efficiency metric in the industry. However, it has a few shortcomings:

- It indicates nothing about the efficiency of the IT equipment itself. Since it does not consider productivity (how much compute processing capability per unit of data center input power or percentage usage of computer server equipment), underused IT equipment can skew the energy profile of the overall facility.

- While it is useful in providing insight to how energy usage changes in response to changes in a data center's infrastructure (deployment of new servers, changes to how air conditioning equipment is operated, etc.), it is based on actual measured energy usage of an active data center. It can change dramatically based on changes in uncontrollable factors (i.e., server use, weather, etc.). It is not especially useful for theoretical baseline calculations performed during the initial design of the data center.

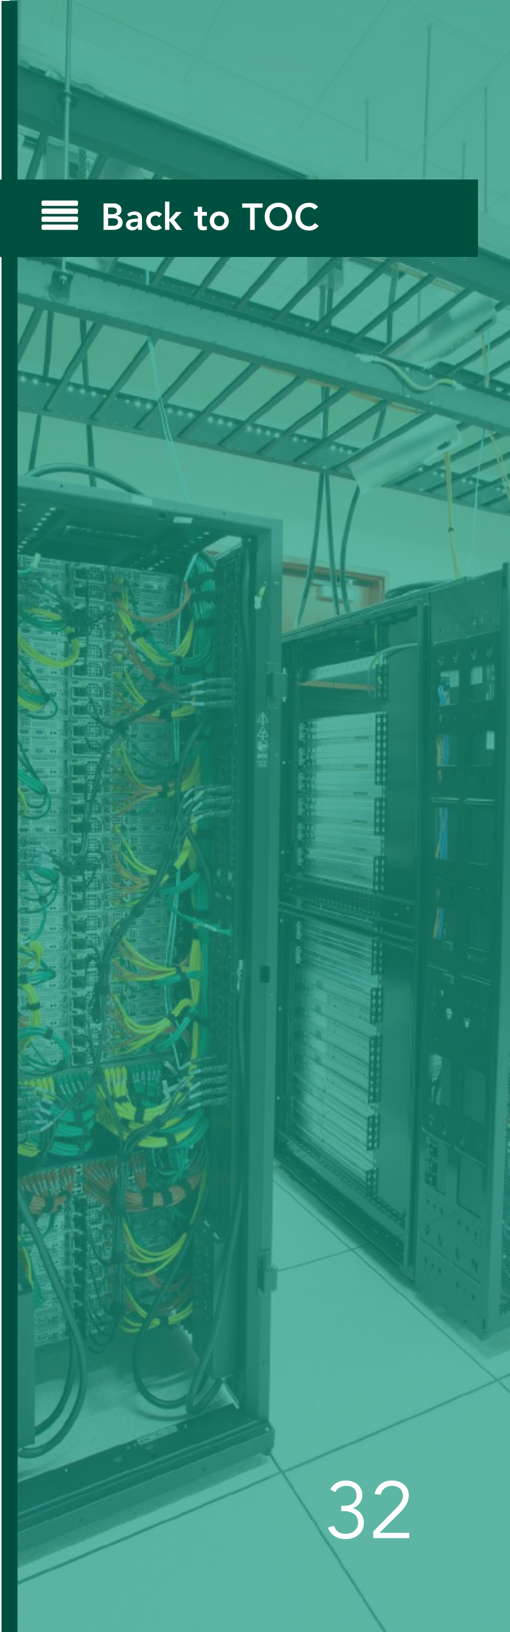

Back to TOC

- Depending on the Green Grid level of metering (where or how it is measured), you may not capture all measured energy consumption in a way that is directly comparable to otherwise similar facilities.

The ASHRAE 90.4 committee initially tried, during early draft revisions of the standard, to apply a concept that they also called PUE, but potential confusion over comparisons to the Green Grid version of PUE eventually led to the elimination of PUE nomenclature from ASHRAE 90.4.

In place of PUE, ASHRAE 90.4 brought forth two new metrics, electrical loss component (ELC) and mechanical load component (MLC). A conscientious decision was made to address mechanical system and electrical distribution system efficiencies separately. The ultimate goal of separating the two was to provide more design flexibility. If either the mechanical or electrical systems didn't meet the minimum efficiency level required by the standard, an alternative compliance path could be used where trade-offs would be allowed. For example, if replacement of a major portion of the electrical distribution system (i.e., deployment of a new larger uninterruptible power supply; UPS) triggered a compliance requirement for an existing facility, replacement of an otherwise functional existing mechanical system that didn't meet the minimum required efficiency could potentially be avoided by an offsetting increase in the efficiency of that new electrical distribution equipment. The associated minimum-efficiency threshold levels were structured to effectively create an 80/20 policy where, after trade-offs are considered, generally only the bottom 20% of existing facilities will be forced into upgrades.

Electrical loss component

ELC quantifies the inefficiencies/losses of different parts of the electrical distribution

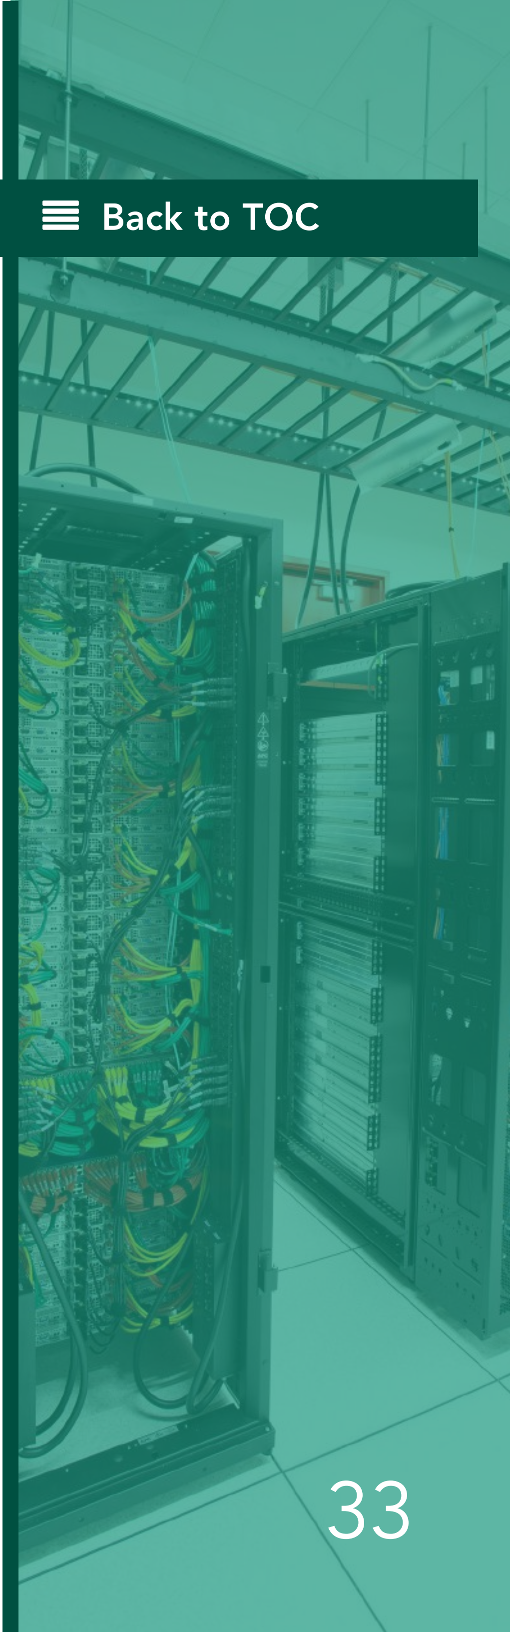

Back to TOC

system, from the utility service entrance all the way through to the receptacle at the ITE cabinet. The ELC value is the percent loss (i.e., a 75% overall efficiency equates to an ELC of 0.25). There are three distinct segments of this power path defined by ASHRAE 90.4 as follows:

- Incoming electrical-service segment (from the utility service disconnect/demarcation to the UPS input)
- UPS segment (limited to the UPS equipment and any associated paralleling gear for multimodule designs)
- ITE distribution segment (from the UPS output to the end of the branch circuit at the point of use/receptacle including all transformers/power distribution units, remote power panels, busduct, branch conductors, etc.)

The focus here is solely on the portion of the electrical distribution system that delivers power to the data center ITE load. If multiple paths to the ITE exist, the pathway with the greatest losses needs to be used in the calculation. However, losses attributed to the portion of the electrical distribution system serving associated supporting systems (computer room air conditioning units, chillers, lighting, etc.) are not included in ELC. Also, emergency/standby generator systems that are normally “off” are not considered as part of ELC calculations.

The ELC calculation for the UPS segment also recognizes that efficiency is affected by the physical size of the UPS system and by the level of redundancy. For example, it is expected that in a fully redundant 2N UPS system, neither UPS will be loaded to more

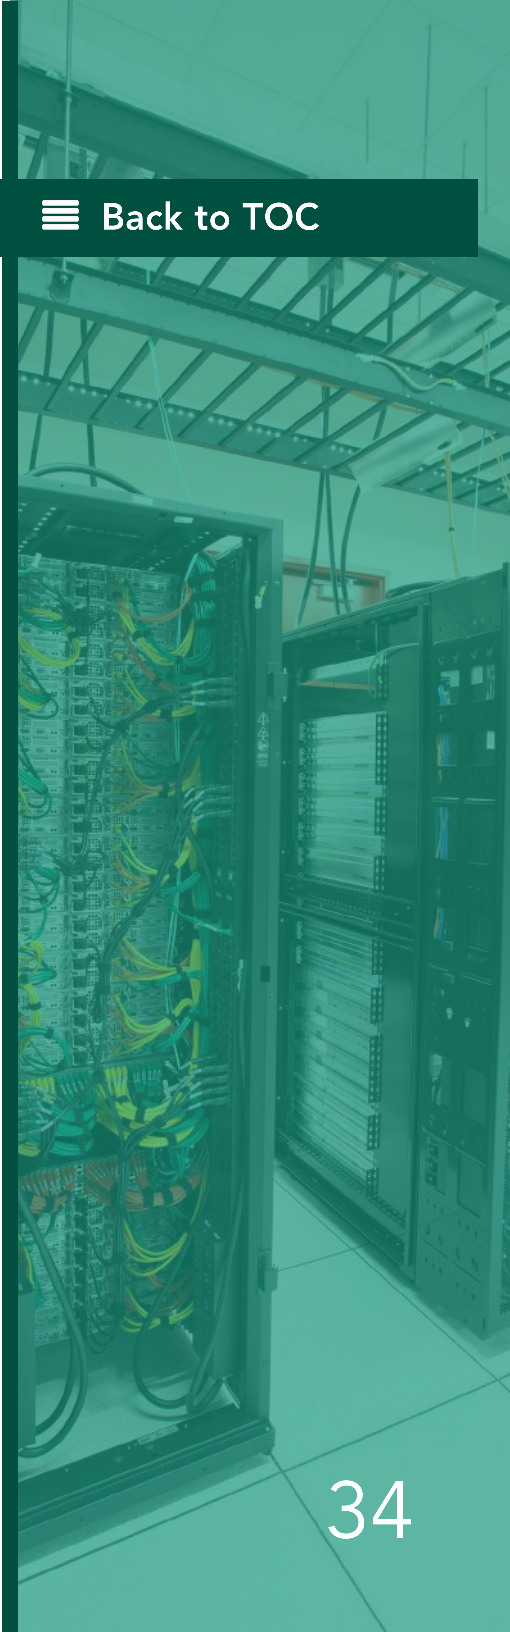

Back to TOC

than 50% under normal operating conditions. A lightly loaded UPS will generally be less efficient than a more heavily loaded UPS with less redundancy. As such, maximum allowable ELC values as detailed in tables 8.2.1.1 and 8.2.1.2 include the following considerations:

- Is total ITE design load either less than or greater than 100 kW?
- Is UPS system configuration single-feed (N, N+1, etc.) or dual-feed with two distinct output busses (2N, 2N+1, etc.)?

ELC calculations must be made at two different load points (100% and 50% of the load expected at each UPS). For example, ELC calculations for a 1N UPS would be made at 100% and 50% and those for a 2N UPS would be made at 50% and 25%.

Mechanical load component

MLC takes a slightly different calculation approach than ELC. MLC is the sum of all data center HVAC equipment power usage (including humidification, if present) divided by the baseline ITE design power. Unlike ELC, where the overall system is subdivided into three parts, the distinction is de-emphasized between individual data center HVAC system components. Rather, the overall power usage of the data center's HVAC system is the key metric here. The only consideration is that the maximum acceptable MLC value changes depending on the climate zone in which the facility is located. This may make certain HVAC technologies that are more effective in certain climate zones more attractive. Regardless, this focus on the system rather than the individual components effectively makes MLC performance-based and not prescriptive, which addresses some of the intense criticism directed at the addendum bu for ASHRAE 90.1.

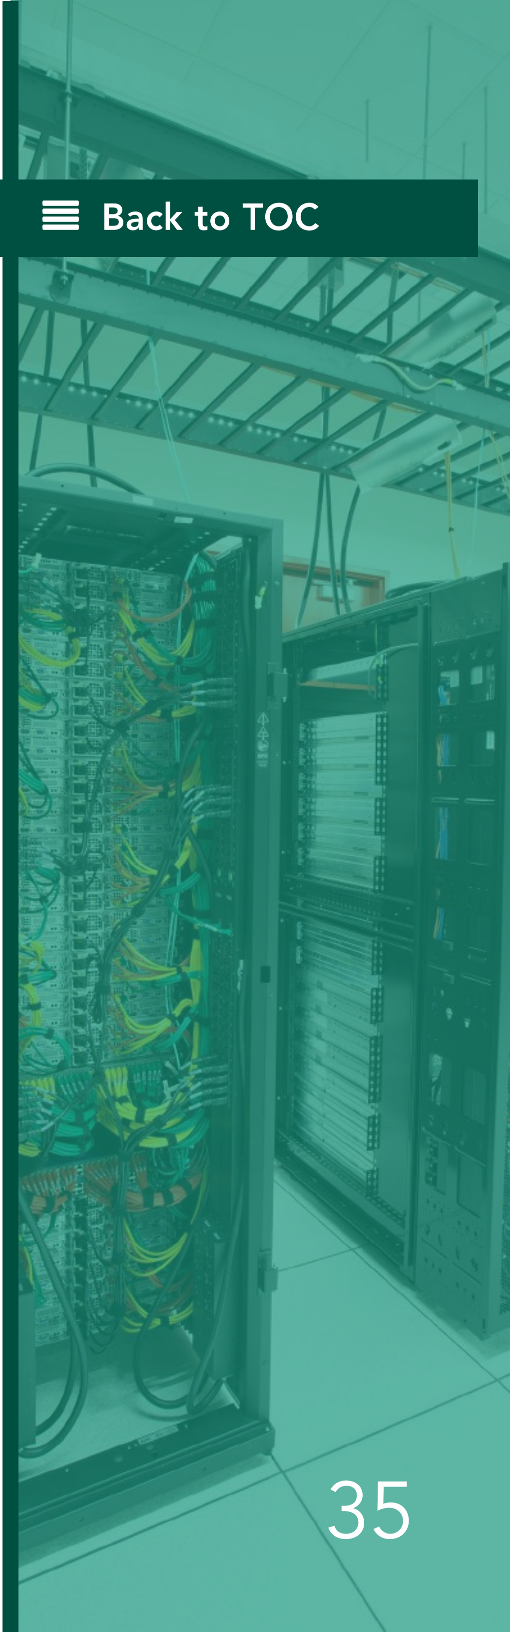

To give the design team more flexibility, there are two MLC compliance paths—power and energy. ASHRAE 90.4 does not mandate that one or the other be used in specific situations. Rather, the decision concerning which path is more appropriate is left to the design team. The power-compliance path calculates peak MLC at both 100% and 50% ITE design loads (kW). The energy path calculates annualized MLC at both 100% and 50% ITE design loads (kWh). Ultimately, the determining factor for which path is more appropriate lies in whether seasonal variability in weather at the data center location will hurt or help overall HVAC system efficiency.

How existing facilities are addressed

The very first question that many clients will nervously ask is if

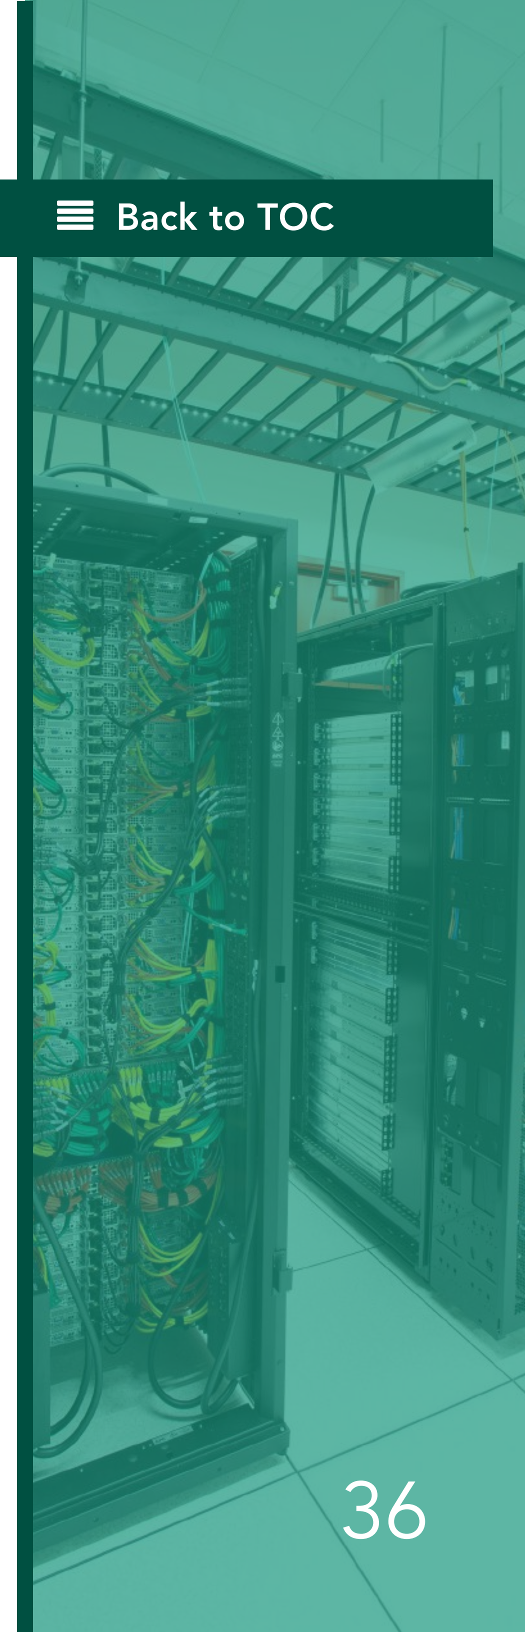

Back to TOC

Back to TOC

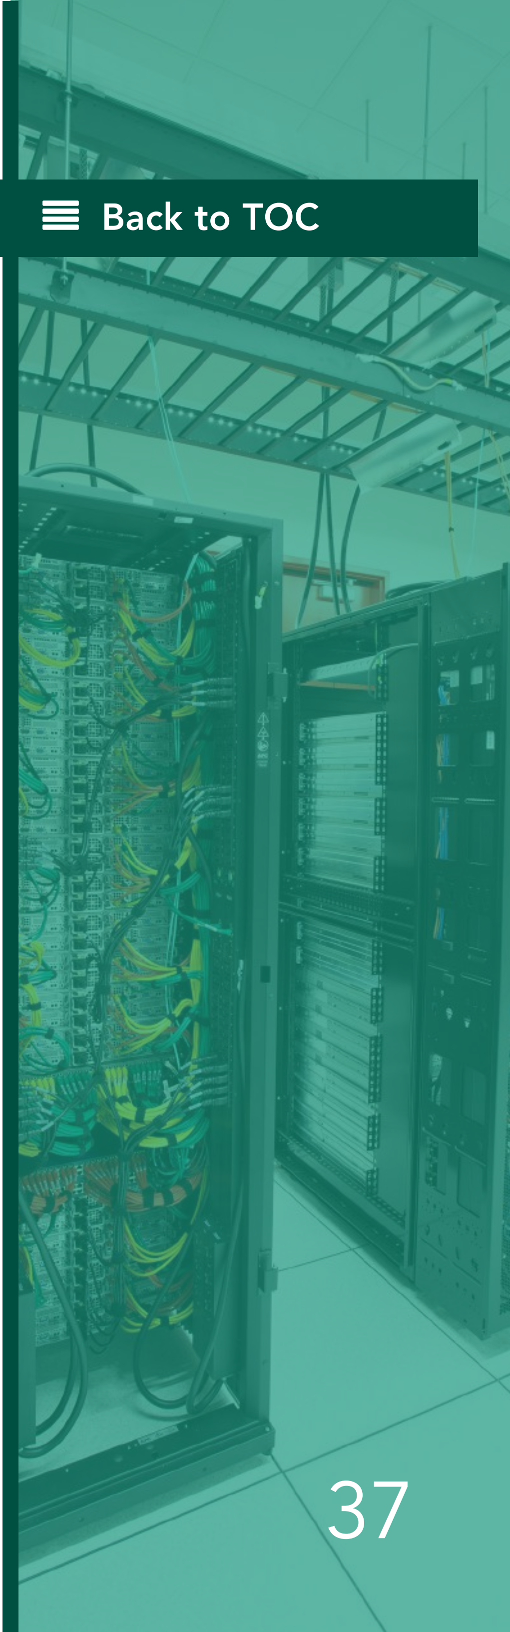

ASHRAE 90.4 applies to their existing data center. The answer is a definite "maybe." Since this is a new standard, it is unclear exactly how individual AHJs will interpret and apply it. Some of the language, especially as it pertains to alterations, is somewhat confusing. However, ASHRAE 90.4 does attempt to make a distinction between new data centers, additions to existing data centers, and finally, alterations to existing data centers. The applicability is as follows:

- All provisions of ASHRAE 90.4 apply to new data centers.
- ASHRAE 90.4 applies to an addition only if it increases area or connected load by 10% or more.
- Alterations shall comply, provided that compliance will not result in the increase of energy consumption for the building. An alteration is defined as "replacement not in kind."

When will ASHRAE 90.4 take effect?

SSPC 90.4's original goal was to have the issuance of ASHRAE 90.4 to coincide with the release of ASHRAE 90.1-2016. However, after the first public review of ASHRAE 90.4, more than 600 comments were received in 45 days. While the sheer volume of comments was not totally unexpected, it did slow down the adoption process. Per the ANSI requirements for a consensus-based standard-making process, any issues raised during the public commentary period must be formally addressed prior to acceptance of the standard. Ultimately, this meant that final approval was delayed until mid-2016.

Back to TOC

While placeholders referencing ASHRAE 90.4 were incorporated into ASHRAE 90.1-2016, it is unclear if ASHRAE 90.4 will be adopted by the International Code Council (ICC) during the current code-development cycle for inclusion in the 2018 International Energy Conservation Code (IECC). Proposals for inclusion by references were brought before the ICC 2016 Group B Committee last year, but were rejected. The reason for rejection wasn't necessarily because of any inherent flaw in ASHRAE 90.4. Rather, the standard was new and had not been reviewed. Regardless, if the prevailing energy code in a particular jurisdiction is IECC, it is questionable that ASHRAE 90.4 will be enforced any time in the immediate future for those locations.

John Yoon, PE, LEED AP

***John Yoon** is a lead electrical engineer at McGuire Engineers Inc.*

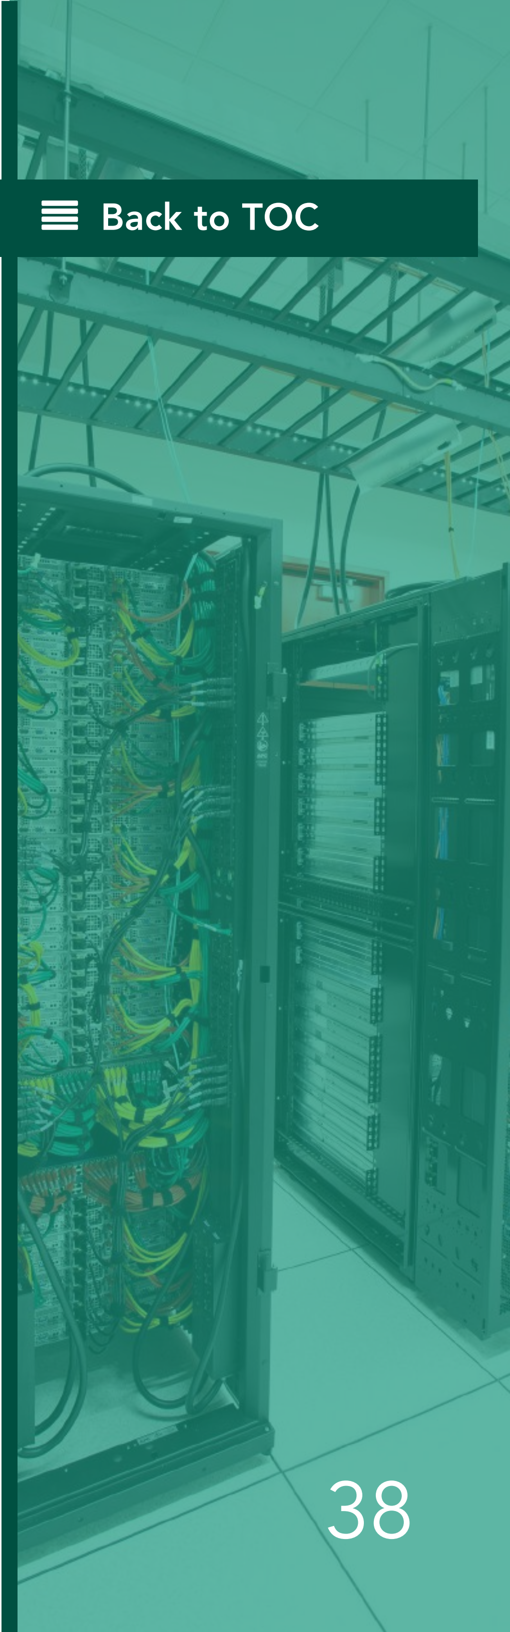

Back to TOC

CHAMPION FIBERGLASS

No Burn-Through Elbow eGuide

Learn how data center and utility projects benefit from no burn-through fiberglass elbows.

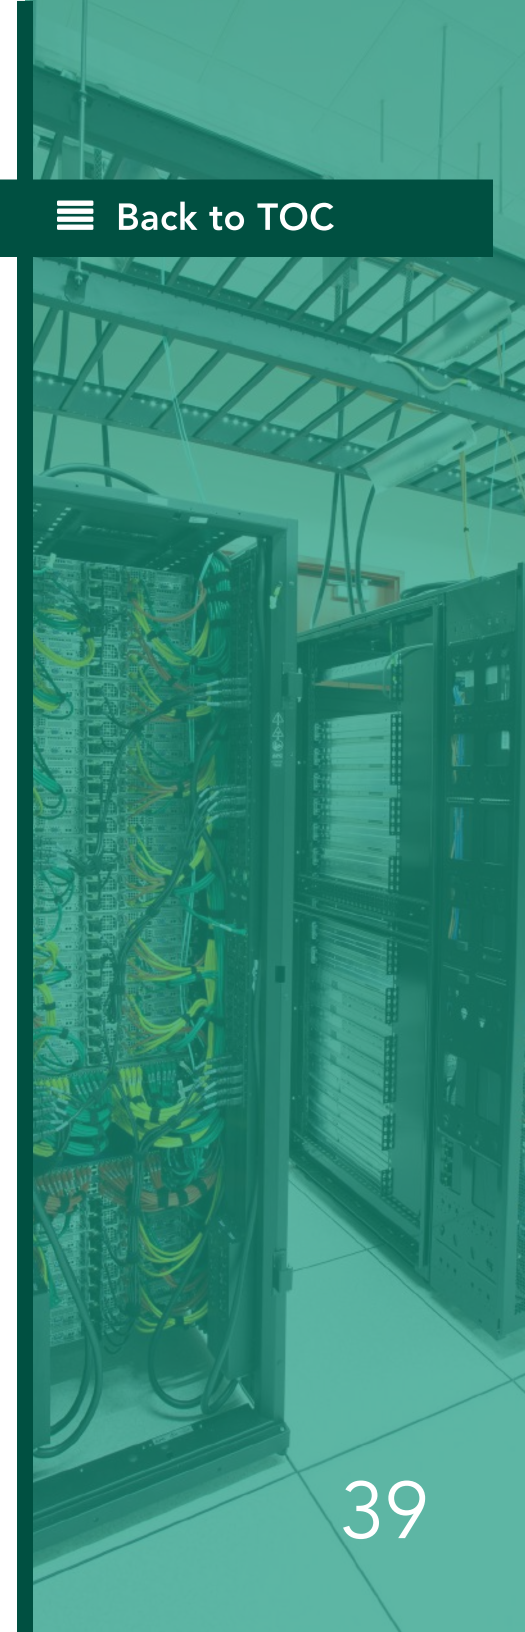

Conduit Sweeps and Elbows Offer Solutions in Data Center Applications

Fueled by legislation and new technology, data center demand is **forecasted to continue to grow by close to 10% per year until 2030**, according to a report by consulting firm McKinsey.

For the electroindustry, this demand will drive the need for materials within data centers such as cables, conduit and elbows, especially the need for a large number of electrical conduit sweeps and elbows in varying sizes. As a leading supplier of electrical conduit to data centers, Champion Fiberglass has developed **product lines** that deliver solutions for common project challenges, including:

- Repairs due to burn-through

Back to TOC

- Installation challenges
- Fault concerns
- Project longevity

Let's review these solutions in more detail and highlight real data center project examples.

Eliminate Repairs in Data Center Electrical Systems

Cable-dense data center projects are at higher risk of costly repairs resulting from burn-through. A single circuit winding process provides flexural strength and the ability

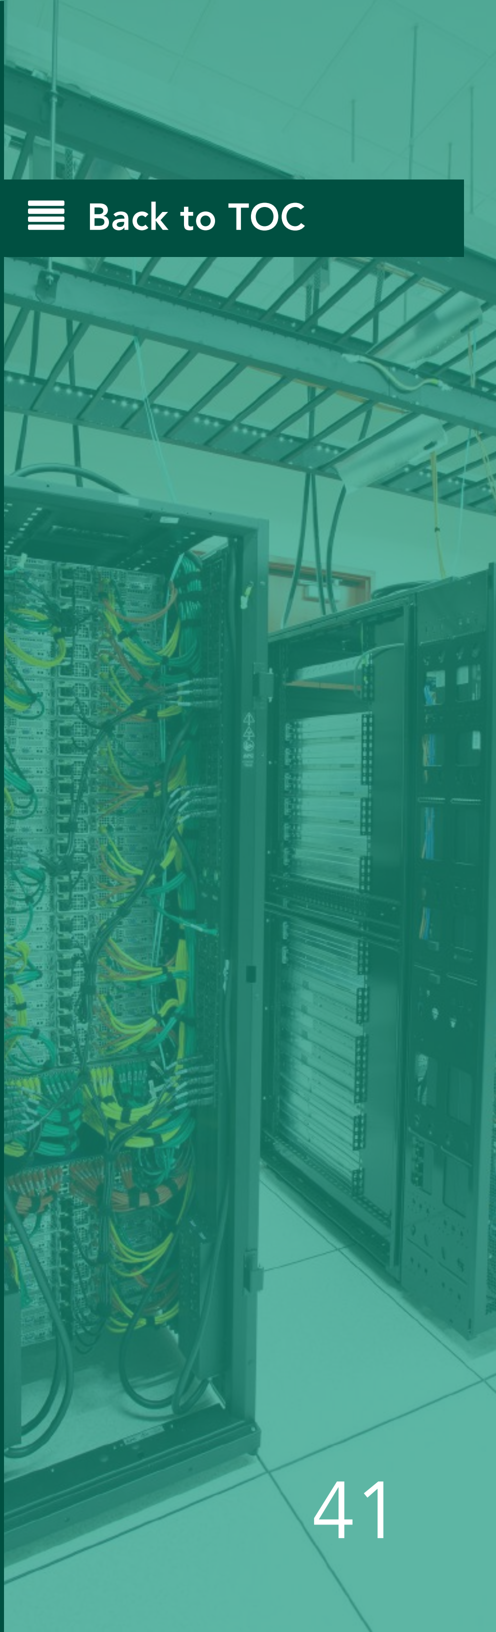

Back to TOC

to absorb mechanical impact, keeping cables from burning through RTRC fiberglass conduit elbows and preventing the need for costly repairs. Fiberglass is also the only electrical conduit substrate where cables won't melt or weld to the inside of the conduit, so cable repairs can be made without damaging the conduit.

Realize Safer, Easier Data Center Installations That Save Money

RTRC conduit sweeps have a .38 coefficient of friction, much lower than .55 for GRC and .90 for PVC SCH 40 and SCH 80, which allows them to be used in longer conduit runs with fewer pull boxes. This allows for a faster, smoother pull with less labor time, fewer materials, cost and risk. Fiberglass conduit elbows also offer no conductivity, resulting in safer installations with reduced risk of electrical injuries.

Fiberglass conduit is lightweight with ease of portability that contributes to easier, faster installation. In fact, for jobs with fiberglass conduit sweeps at larger diameters significant labor savings can be found through faster installations. According to **NECA**, fiberglass installation labor time is a fraction of what it is for galvanized rigid steel (GRC) or PVC-coated steel. You can estimate your potential project savings with the **Champion Fiberglass Elbow Calculator**.

Experience Product Strength and Longevity

Electrical contractors and specifying engineers continue to choose Champion Fiberglass elbows as part of a complete electrical conduit system with a long life and low maintenance, thanks to the conduit's strength, durability and corrosion resistance.

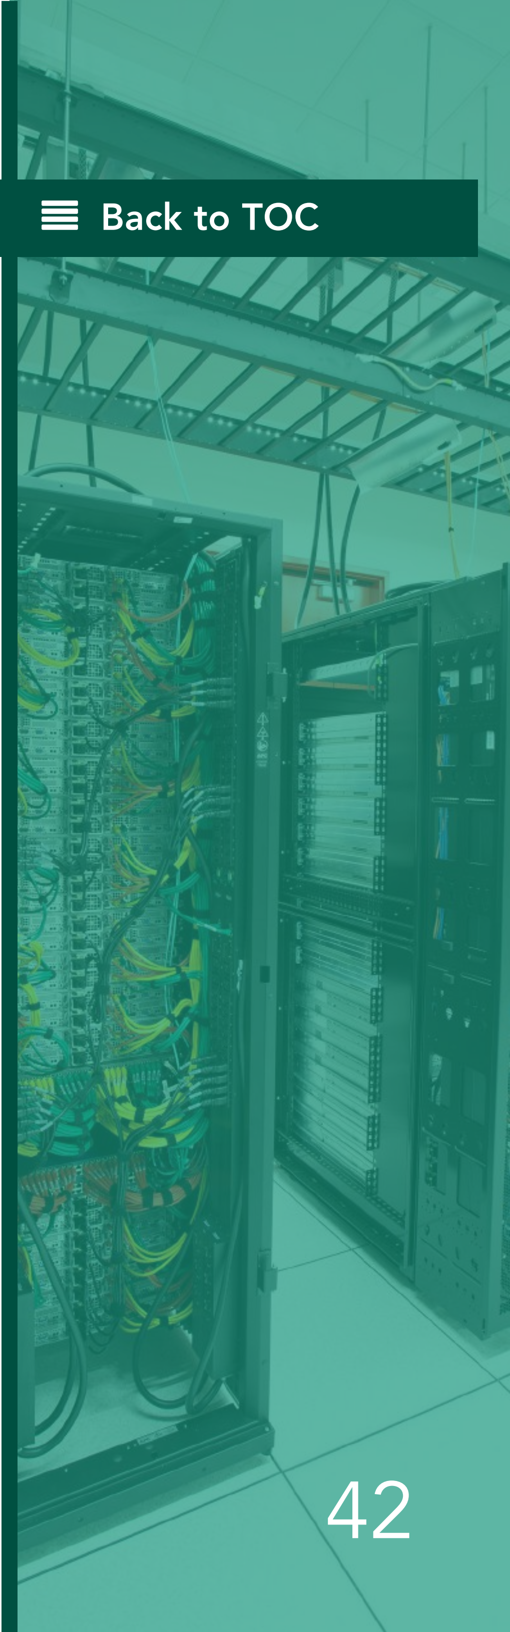

Back to TOC

Discover Data Center Project Successes

Champion Fiberglass has a long history of success in data center projects. In a recent **project**, fiberglass conduit offered the following wins:

- Fiberglass conduit served as a reliable substitute for PVC throughout a product shortage
- Fiberglass conduit saved material costs over other conduit types
- Low coefficient of friction, no burn-through and fault resistance provided ultimate protection for important data center electrical
- On-time deliveries and speedy installation (with guidance from the Champion Fiberglass team) allowed the project to finish on time to hit goals

Get the details of this case study **here**.

For another data center **project**, fiberglass conduit proved a better alternative to galvanized rigid steel elbows. Material costs were less and the lighter weight of fiberglass

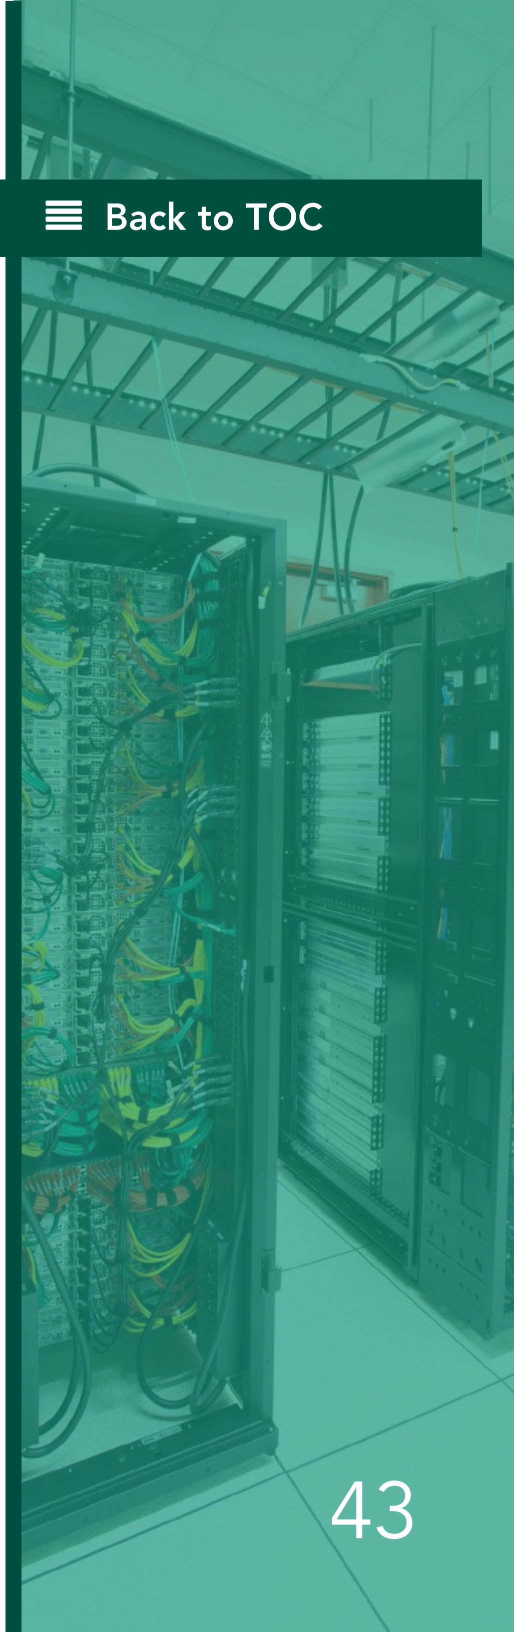

conduit elbows saved substantial time in installation. Results include:

- Fault resistance and elimination of cable burn through to support this data center for the entire lifecycle of the facility

- Immediate savings in terms of easier field handling, lower freight cost and a lower total installation cost

- For this data center, a savings of over $500,000 in materials and installation was realized by choosing fiberglass conduit instead of galvanized rigid steel

Find the full case study **here**.

Learn more about **Champion Fiberglass Elbows**.
Download the complete No Burn-Through Elbow Eguide.
To purchase fiberglass conduit elbows today, reach out to **your local Rep**.

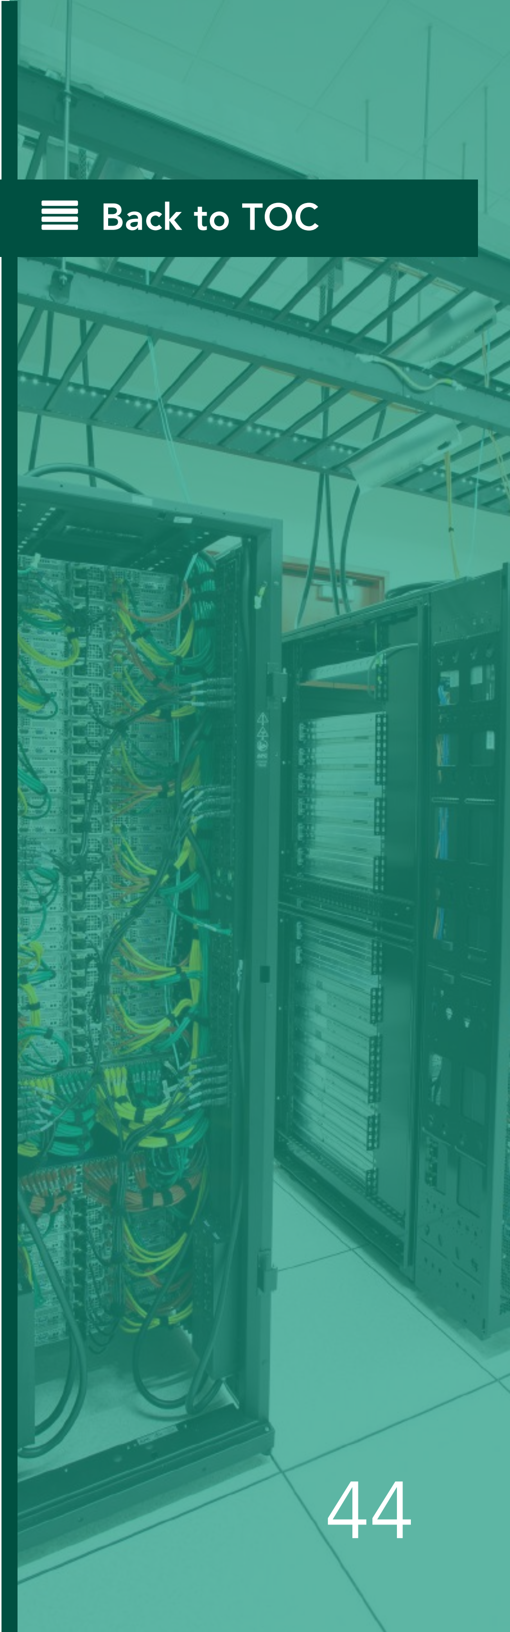

Codes and standards drive trends, changes in data center design

Several codes and standards are pushing the engineered systems in data centers in different directions

Bill Kosik, PE, CEM, BEMP, Senior Energy Engineer, DNV, Oak Brook, Illinois – Matt Koukl, DCEP, Principal, Market Leader Mission Critical, Affiliated Engineers Inc., Madison, Wisconsin – Kenneth Kutsmeda, PE, LEED AP, Global Technology Leader – Mission Critical, Jacobs, Philadelphia – Ben Olejniczak, PE, Senior Project Mechanical Engineer, Environmental Systems Design Inc., Chicago – Brian Rener, PE, LEED AP, Mission Critical Leader, Smith Group, Chicago – Jonathan Sajdak, PE, Senior Associate/Fire Protection Engineer, Page, Houston

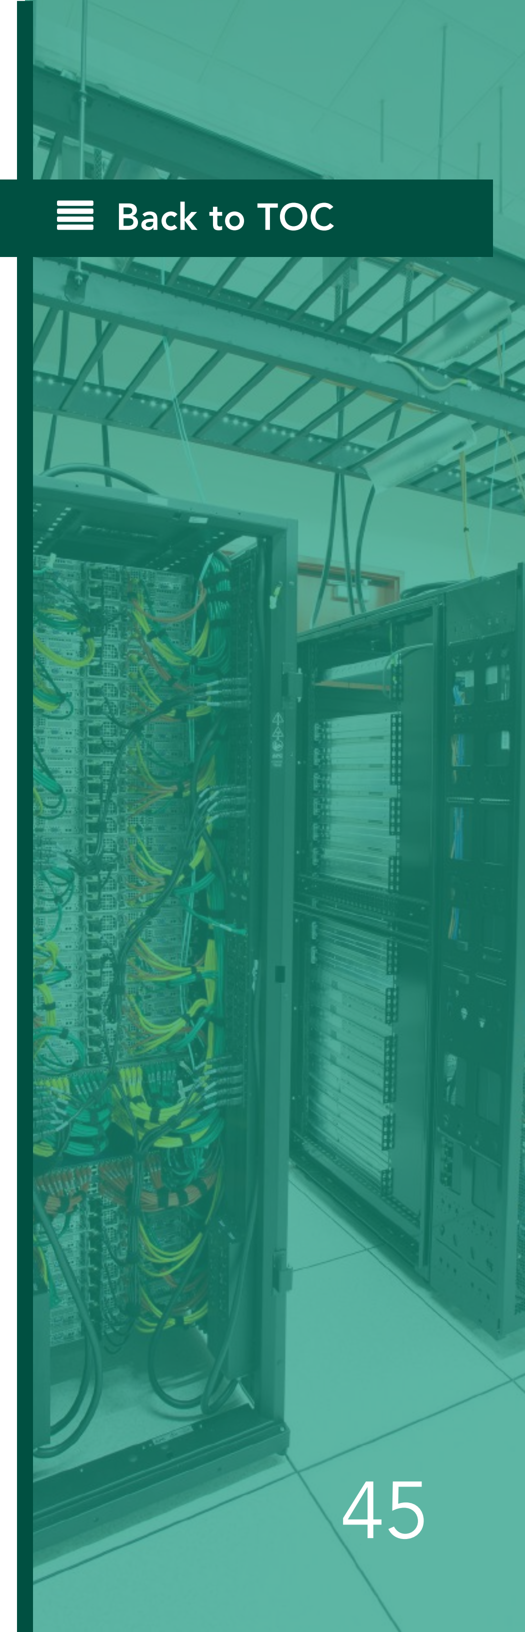

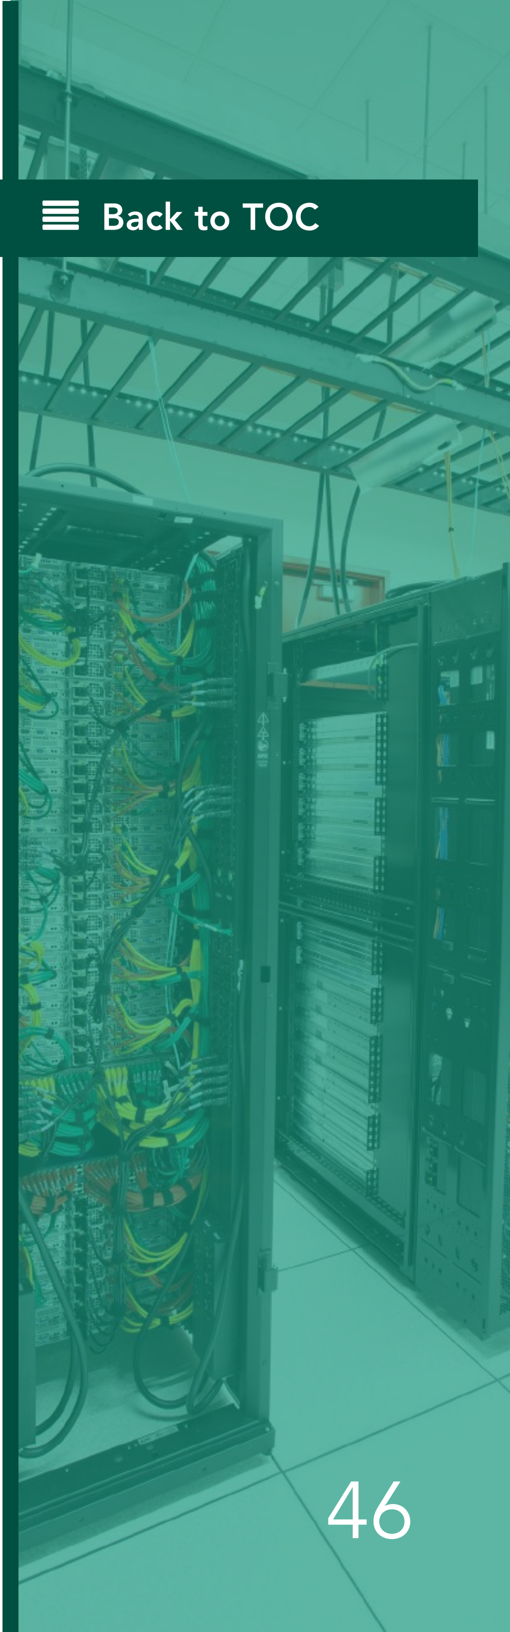

Back to TOC

What are some best practices to ensure that such buildings meet and exceed codes and standards?

Jonathan Sajdak: At the beginning of each project, it is key to identify all applicable codes and standards. Coordination with the authority having jurisdiction during the early stages of the design can also be very helpful. During these conversations, the applicable codes and standards can be reviewed and a high-level overview of the design can take place to ensure both parties are in agreement with the requirements. One of the biggest challenges that can pop-up on a project is when an owner requirement, insurer requirement or local amendment was overlooked as the design stage nears completion. It may be difficult to implement the required system or design feature due of the lack of available space in the facility or its impact on other disciplines.

One additional clarification to highlight is that in most jurisdictions, clean agent systems cannot be used in lieu of automatic sprinkler systems. Clean agent systems are not intended to be used as a replacement for sprinkler systems, but instead be used in addition to as an extra layer of protection with the goal of rapid suppression/containment of the fire and mitigating damage to property and assets.

Ben Olejniczak: One of the most important best practices that I could recommend is to simply stay on top of the latest revisions of all the documentation that may influence your design. It sounds simple, but it actually requires a fair amount of time investment to understand and keep track of subtle changes that occur in each revision. Also, it's very important to stay on top of the latest technology that equipment vendors are working on to ensure they are keeping up with the changes as well. With all the changes happening within codes and at the legislative level, sometimes it is difficult for all parties to stay current.

Back to TOC

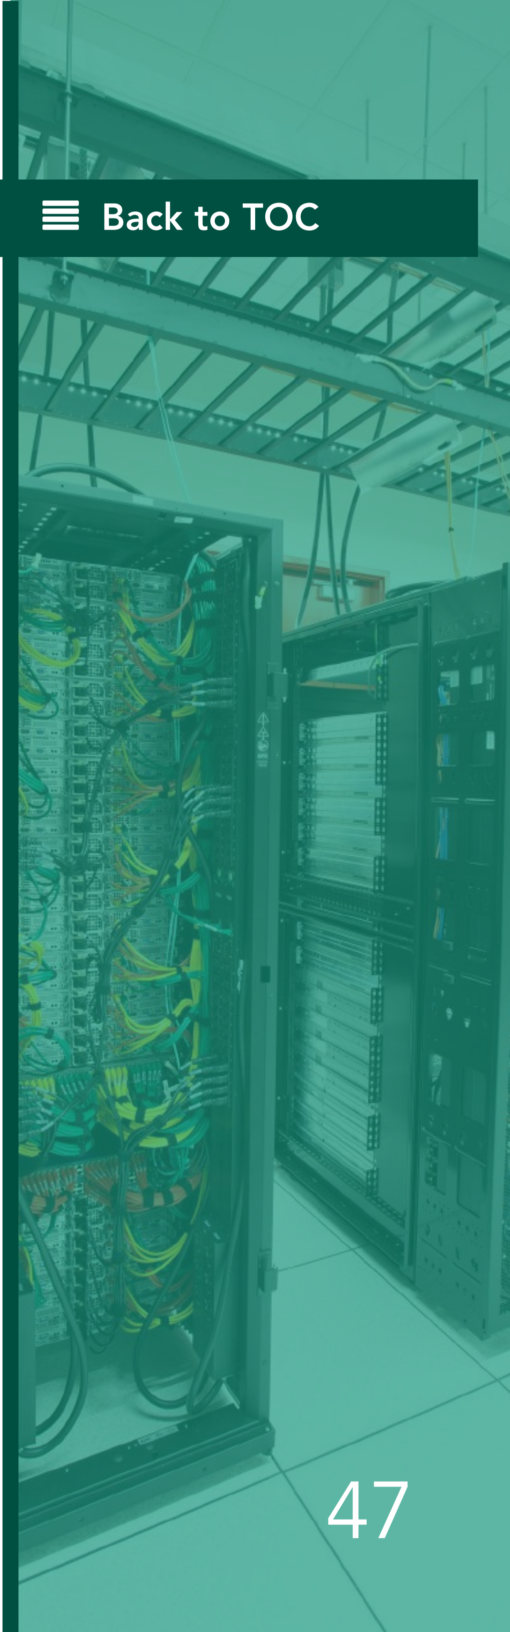

How has ASHRAE Standard 90.4-2019: Energy Standard for Data Centers affected your work?

Bill Kosik: I think one of the main ways ASHRAE 90.4 has affected my work is by the standardization of certain design elements. An example of this: the code has instructions on how to calculate the electrical losses based on the reliability of the electrical system. Seemingly minor, the electrical losses could total up to nearly what it takes to power the fans. This also is conducive to modeling what-if scenarios to calculate the power loss and associated air conditioning energy consumption.

How are codes, standards or guidelines for energy efficiency impacting the design of data centers?

Matt Koukl: Data centers are unique compared to traditional building facilities that accommodate creature comforts while balancing the facility's energy use. With high sensible loads and computing power involved, data centers have a significantly higher base load and energy use intensity than most other facility types. As designers of these facilities, we must ensure the infrastructure and heat rejection systems operate at the highest levels of energy and water efficiency. A common approach to achieving the code energy requirements is by evaluating the supply water temperatures to gain the highest number of economizing hours to reduce the annual energy use of mechanical cooling systems.

Ben Olejniczak: As codes, standards and guidelines become more restrictive, it will require the designer to think critically about their approach to designing a data center. Enhanced complexity generally accompanies systems that integrate economization/

free cooling from a controls perspective. To be clear, this is not a downside to the design — it just requires deeper thought and analysis. Systems that may have been assumed as the status quo may no longer be feasible and alternate solutions may be required to be investigated.

What new or updated code, standard, guideline organization or association do you feel will change the way such projects are designed, bid out or built?

Kenneth Kutsmeda: Two new codes that engineers should be aware of are International Fire Code – Section 1206 Electrical Energy Storage Systems and NFPA 855: Standard for the Installation of Stationary Energy Storage Systems. Both codes contain important new requirements (location, separation, quantities, etc.) that relate to the installation of battery systems in particular lithium-ion batteries. UL 9540A is a testing standard that can provide compliance to the new IFC and NFPA code regulations, but it must be approved by the local AHJ.

What are some of the biggest challenges when considering code compliance and designing or working with existing buildings?

Brian Rener: On the electrical side, it is critical to understand exiting fault and arc flash levels and models in order to properly specific rated and safe additions and modifications.

Jonathan Sajdak: With existing buildings, one challenge we have seen is the use of existing sprinkler piping in facilities that are planned to be converted into data centers.

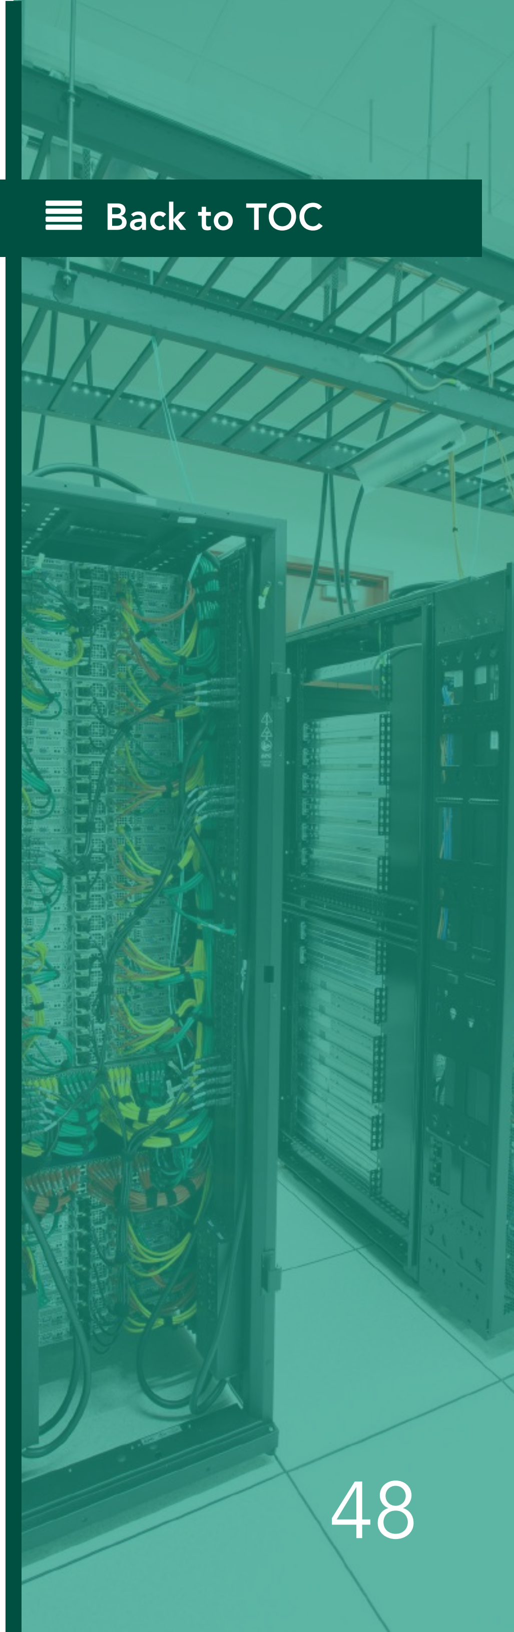

Back to TOC

When an existing building and system is inherited, attention should be given to the condition of the existing piping (both internal and external) and determining whether or not it needs to be replaced. Corrosion mitigation is key in data centers and the integrity of the sprinkler pipe is vital. Corrosion is one of the leading causes of systems leaks, which could result in unwanted discharge of sprinkler systems in equipment areas. Supplementary systems or features to help with this can include nitrogen generation, automatic air venting and pipe thickness/material selection.

Kenneth Kutsmeda: The biggest challenge I faced with code compliance is the interpretation of the code by different jurisdictions. Codes can be interpreted differently based on the experience and expertise of the person reviewing the code. Early in my career a challenge I faced repeatedly with AHJ was the elimination of emergency power off, or EPO, system for the data center. EPO systems were potential points of failure that could take down the data center, so most clients wanted them eliminated. Data center designs were adjusted so that they no longer governed by NFPA 70: National Electrical Code Article 645 and therefore no longer required an EPO system. Most of the AHJs had experience with the requirements of NEC 645 IT equipment that required EPO buttons in data centers. Convincing AHJ that the data center was no longer governed by NEC 645 was a big challenge.

Consulting-Specifying Engineer

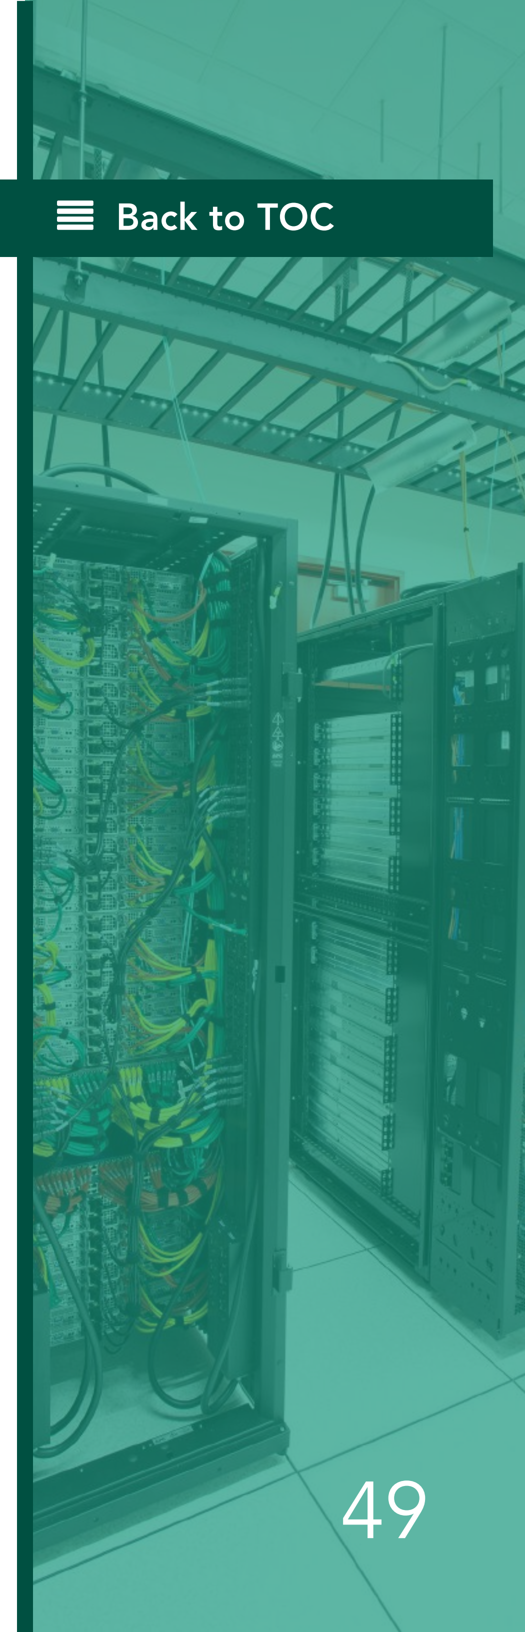

Back to TOC

OPT-X™ High-Performance Fiber Systems

Engineered for ease of deployment and manageability, while ensuring migration options for future bandwidth growth

Available Worldwide

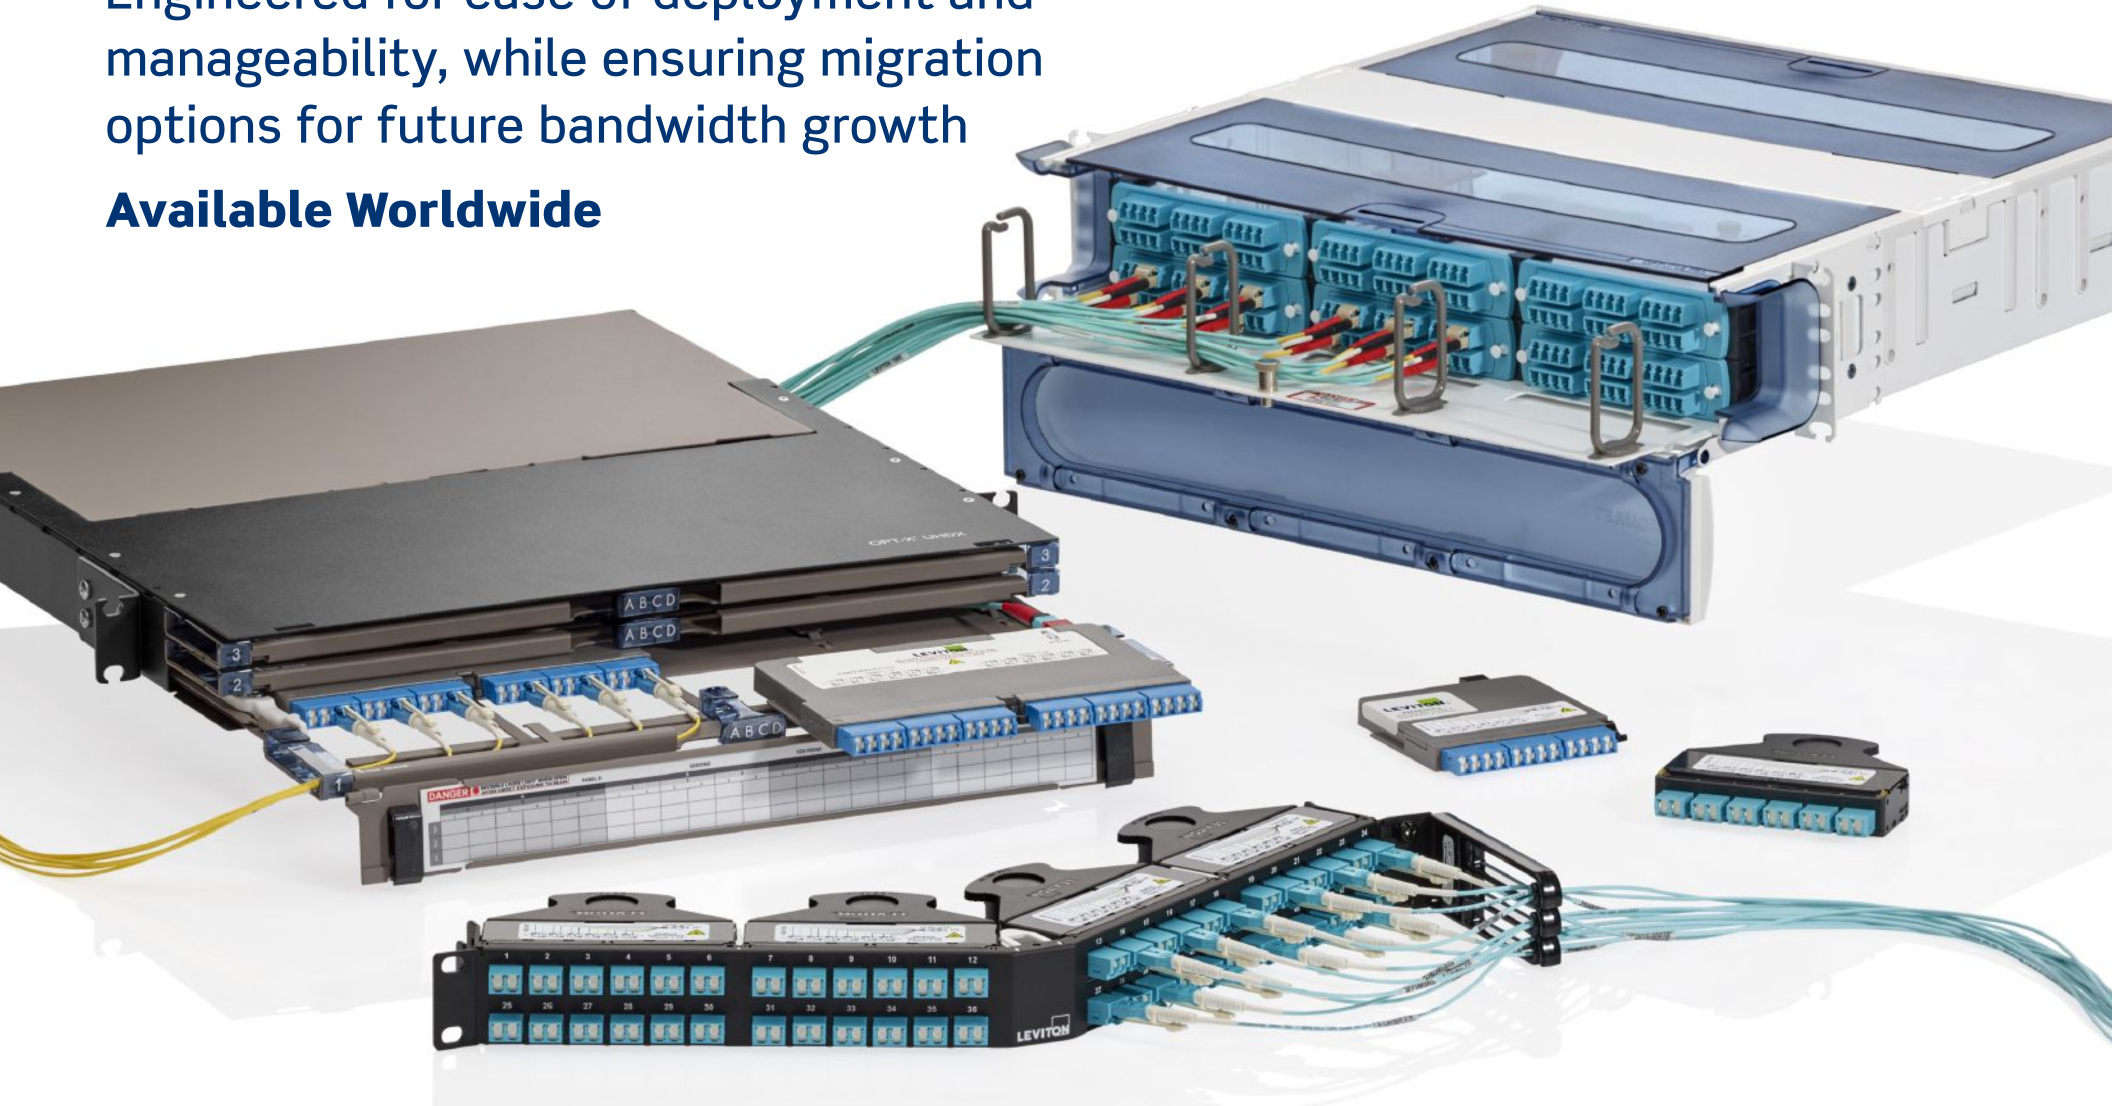

Back to TOC

LEVITON
Network Solutions

Leviton Network Solutions — Solving Data Center Challenges

Every data center has a unique set of network priorities, from consolidation and cost control to quick deployment and zero downtime. You can count on Leviton to ensure that your data center's infrastructure meets your specific requirements. Leviton network connectivity is installed in some of the most advanced data centers and applications in the world. And our data center specialists have extensive experience with the most popular active equipment, network topologies, and transmission methods.

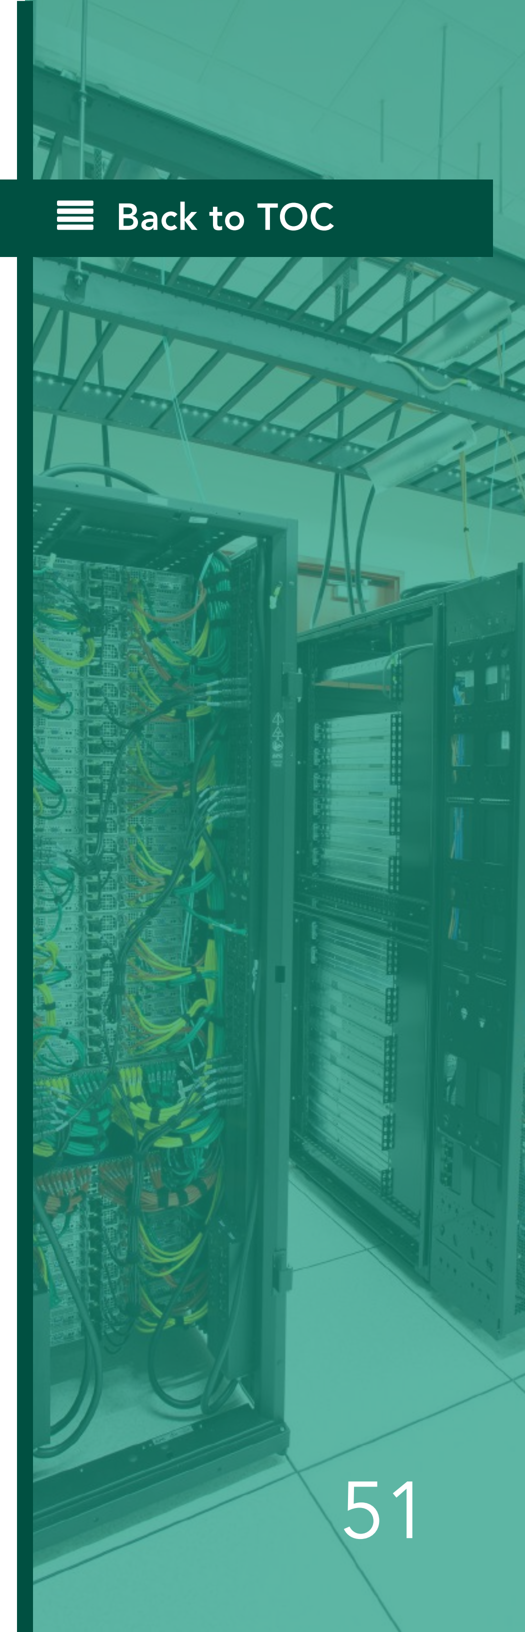

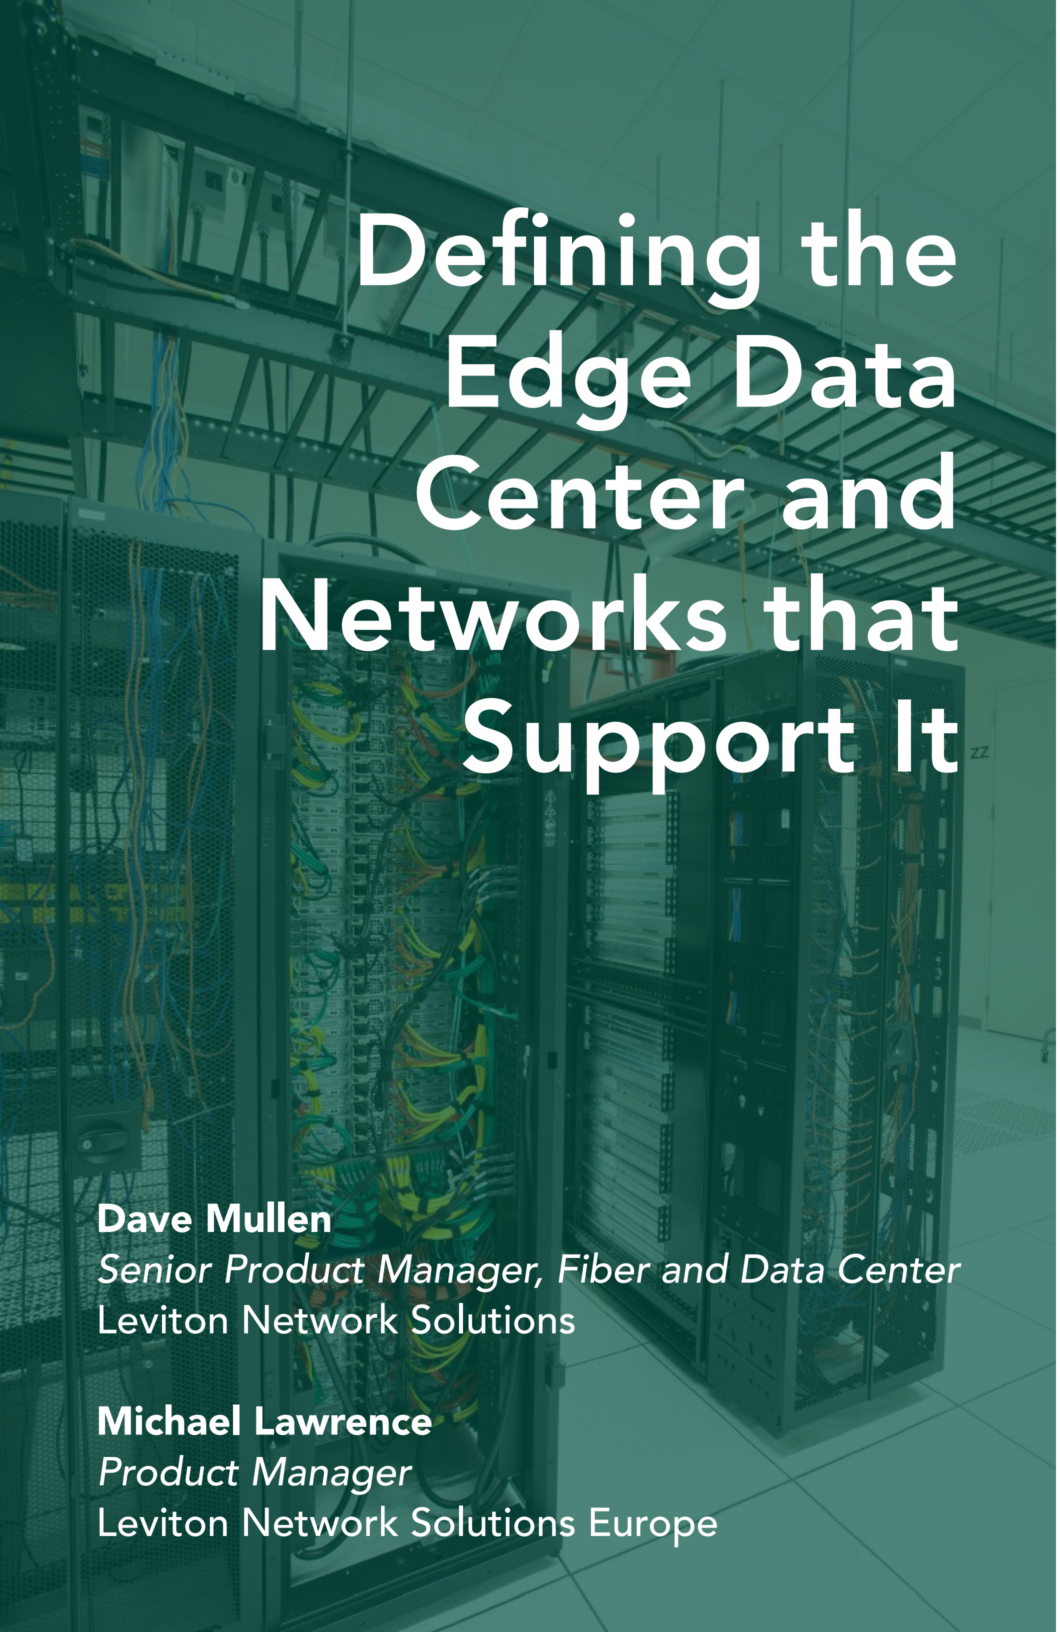

Defining the Edge Data Center and Networks that Support It

Dave Mullen
Senior Product Manager, Fiber and Data Center
Leviton Network Solutions

Michael Lawrence
Product Manager
Leviton Network Solutions Europe

Looking back 8-10 years ago, the "cloud" was the big topic when it came to data centers. Today, it is the "edge" garnering hype and press. While it may seem like a buzzword, edge computing is a very important development. As more and more devices around us become connected and critical business applications need to be on-site — from smart factories to smart cities with IoT connected devices — the latency or time delay for data to transfer needs to be much shorter. This simply can't be addressed without bringing the computing power of a data center to highly distributed edge computing, closer in proximity to where the connected devices are located.

The biggest catalyst for the growth of edge data centers comes from emerging 5G cellular technology. 5G is opening up opportunities for new IoT applications and smart city technologies that rely on real-time data, such as improved automation in factory and buildings, flow of pedestrian traffic in dense urban areas or sporting events, and even responsive autonomous vehicles in the future. 5G is creating a complex digital transformation that will converge telecom and IT networks, and edge data centers will be needed to address the lower latency and higher bandwidth requirements that 5G brings.

Back to TOC

In general terms, a medium latency for data transmission from an end device to a centralized or hyperscale cloud data center can be around 20 milliseconds or longer. When moving the data storage and processing to the edge, latency can drop to 10-15 milliseconds. This is considered low latency. This may seem fast: for reference, our brains need about 13 milliseconds to recognize what our eyes see. However, some emerging IoT applications for on-premises networks such as factory assembly lines will require ultra-low latencies that drop down to 5 milliseconds or less. Low and ultra-low latency performance can only be accomplished with an architecture where edge and traditional cloud data centers work together by sharing processing power and reducing latency for when applications require it.

Understanding the Edge

The definition of "edge" can be a little fuzzy, but a good general definition of edge computing comes from James Stranger, the chief technology evangelist at CompTIA. According to Stranger, the edge is "the practice of capturing, storing, processing and analyzing data near the client, where the data is generated, instead of in a centralized data-processing warehouse." Thus the data is stored at the edge of the network, rather than always with a traditional data center.

Not all edge data centers look the same. Service edge data centers could take the form of a "cloudlet," or small-scale cloud data center that moves some of the resource-intensive computing closer to the edge. Similarly, small data centers used by colocation and service providers could serve that edge location role. These are dedicated facilities that even have the power, cooling, and security commonly associated with a traditional data center, and might host 10 to 100+ cabinets. These might be typically found in second tier or mid-size cities with less than a million people, and offer low or medium latency.

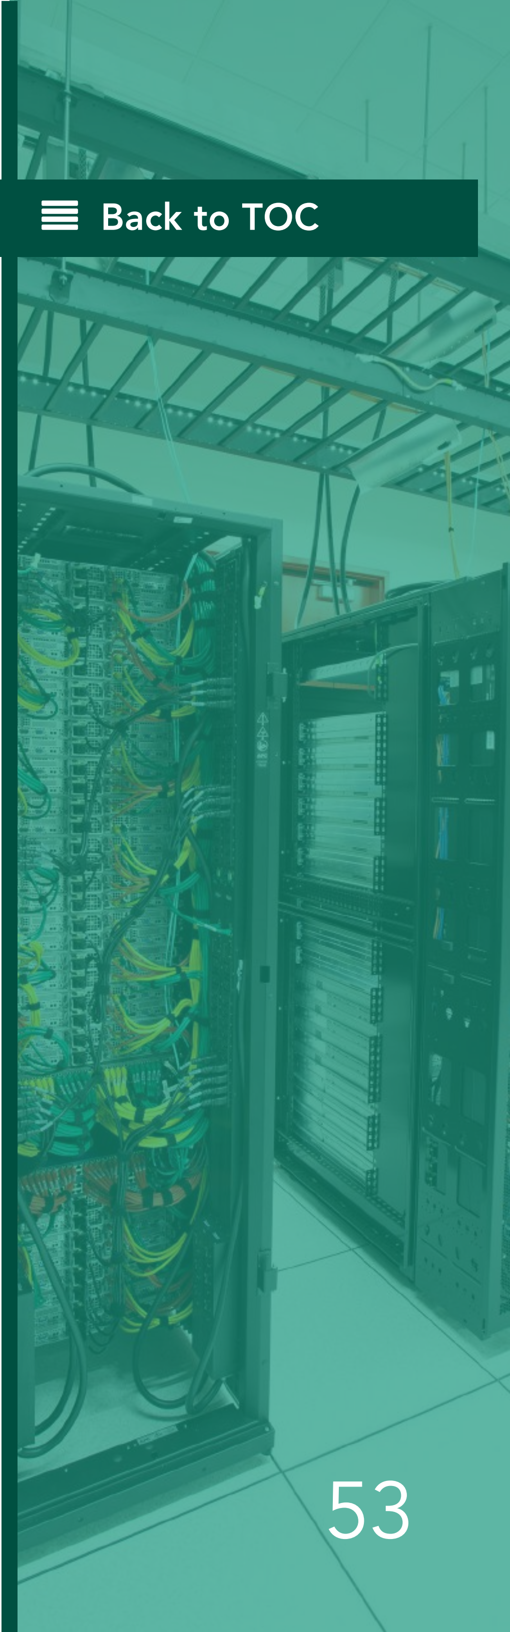

Back to TOC

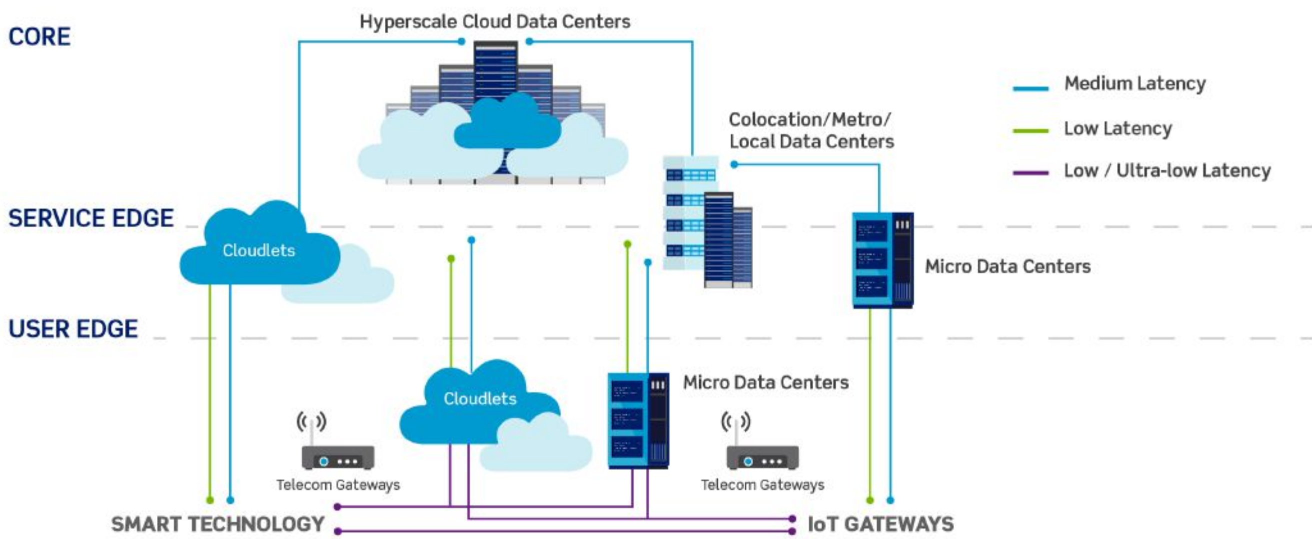

User edge data centers are located in even closer proximity to users and end devices, with the goal of providing low to ultra-low latencies at 5 milliseconds or less. These tend to take the form of micro data centers or very small data centers ranging in size from a half-rack height up to five cabinets. They are versatile solutions that could be located in a warehouse, wiring closets or remote sites — anywhere on premises to support workloads that are critical to a business.

Data rates and cabling infrastructure

It is important to recognize that the effect of 5G deployments goes beyond adding more edge computing — it will also place more strain on the core cloud computing performed in centralized data centers. Roughly 90% of data is still processed in these data centers today, and 5G will speed up the introduction of 400 Gb/s and 800 Gb/s optics and switches in hyperscale and cloud data centers as a way to move data faster.

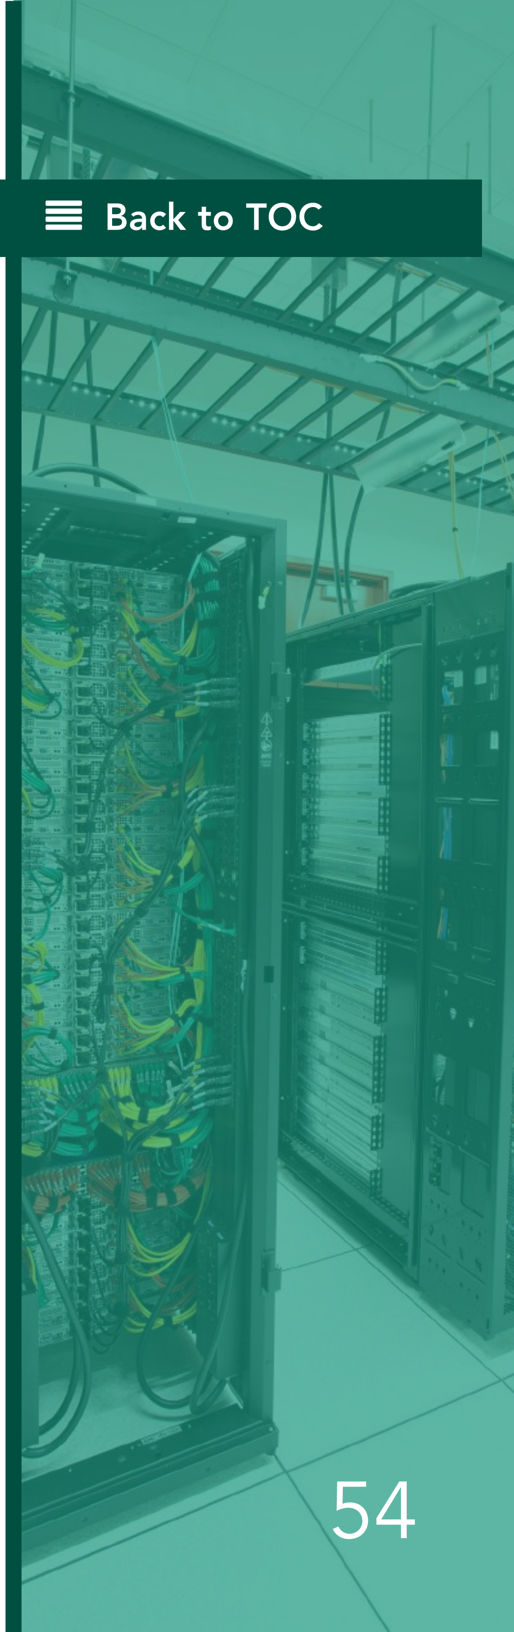

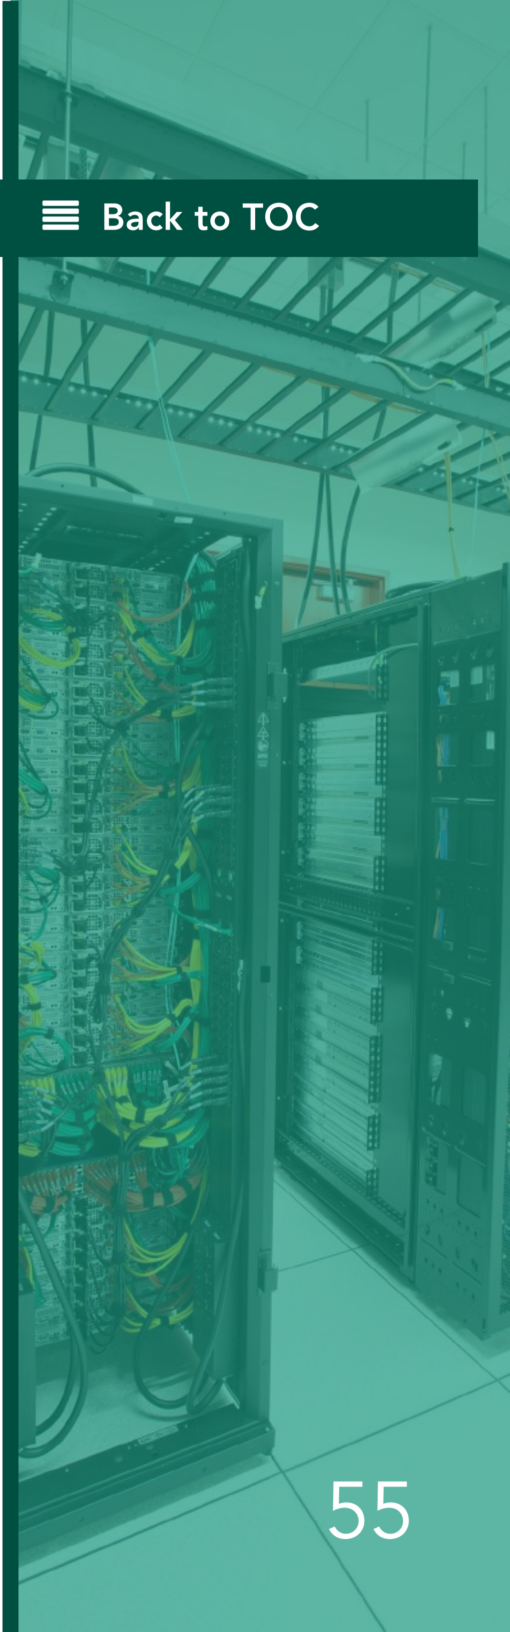

Back to TOC

400 Gb/s switch options entered the market fairly recently, introduced by manufacturers in late 2018 and early 2019, and adoption of these new switches took off in 2020. The new 400 Gb/s switches, based on 12.8 Tb/s chips, provide much faster speeds and greater network density. Today, 200 Gb/s and 400 Gb/s switches make up 20 percent of total data center switch shipments, according to a Dell'Oro 1Q 2023 Ethernet Switch Data Center Report.

However, as we now see more 200 Gb/s, 400 Gb/s, and eventually 800 Gb/s in hyperscale cloud data centers, we will also see 100 Gb/s optics deployed at servers and in edge data centers. 100 Gb/s is now a fundamental building block in data centers, and there will remain a strong demand for 100 Gb/s in the near future.

It is also important to note that the majority of 100, 200 and 400 Gb/s transceiver options are for single-mode networks because of bandwidth and distance capabilities. This trend is also partially a result of the decreasing cost of single-mode optics, prompted by their adoption by cloud companies with major purchasing power, and recent standards committee activities that specify more single-mode options for higher speeds. As this trend continues, centralized and near edge data centers may find single-mode solutions to be more enticing.

Network Security at the Edge

Networks for micro data centers use a combination of copper and fiber cabling, usually with copper connectivity and fiber uplinks. As micro data centers are sometimes located in more exposed or even ruggedized environments, cabling and connectivity should be robust and protected within the rack. These data centers may not have back up power or the high levels of security, fire detection or cooling that are all part of a

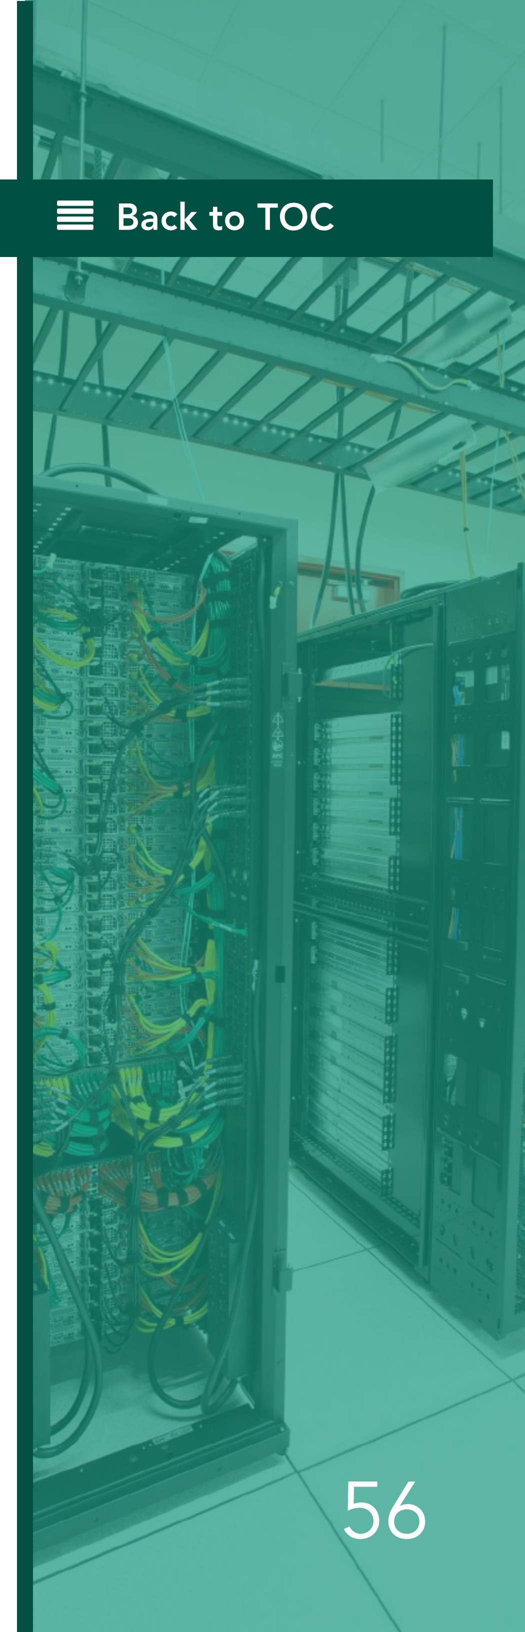

Back to TOC

centralized data center. Secure connectors and assemblies that lock into ports might be a greater priority, along with locking cabinets and additional security.

As edge data centers are often distributed in more remote locations, intelligent port management may be helpful in monitoring networks from a central location. Intelligent patching allows for real-time visibility of the status of network connections. Managing and tracking connections is a simple means to identify unauthorized access and to control disruptions caused by human error. As the number of edge data centers grows, intelligent patching is scalable to accommodate increased demand, offering continual network security as the network grows.

Ultimately, edge and micro data centers can end up having very different requirements than large data centers. These differences include the physical infrastructure, and there are numerous considerations you can encounter. Is there existing structured cabling that you are integrating into? Is it single-mode or multimode? Is trunking required? Guidance from structured cabling manufacturers and consultants can help you navigate this new edge data center frontier.

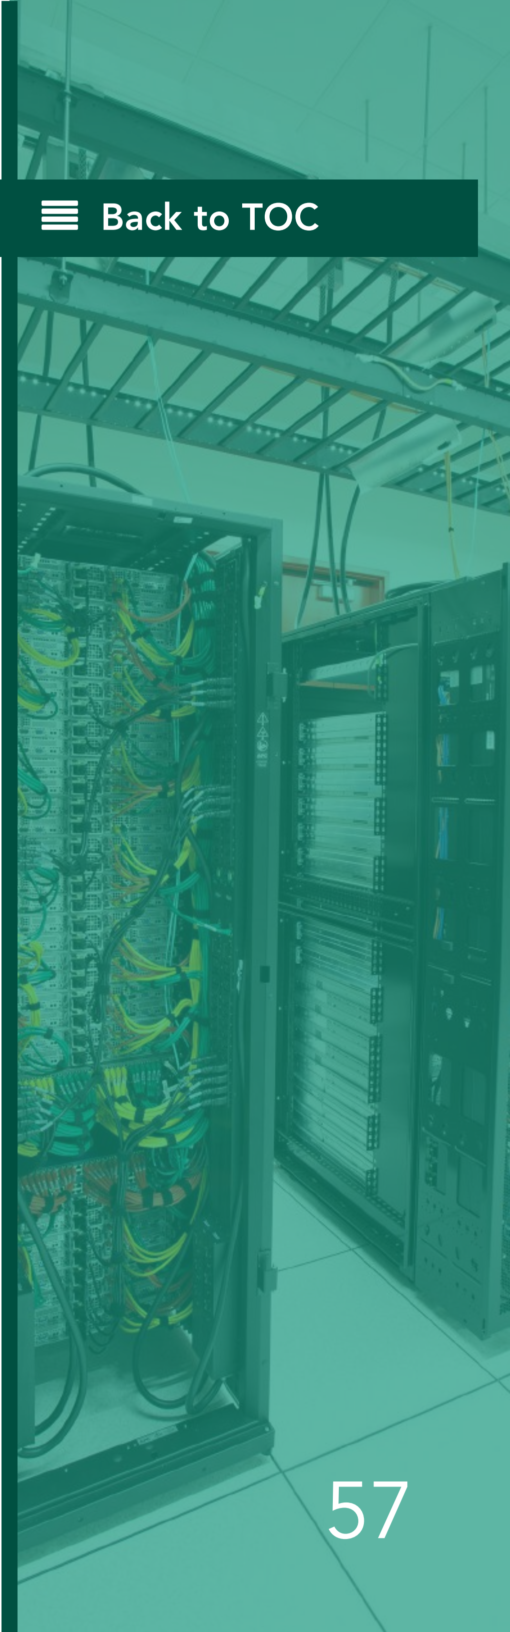

Back to TOC

Dave Mullen
Senior Product Manager, Fiber and Data Center
Leviton Network Solutions
Dave *works closely with Leviton customers to understand their applications. As a result he's had a hand in bringing many of Leviton's recent cutting-edge fiber optic solutions to the market, such as the OPT-XTM Unity Fiber System, OPT-X SJP pre-terminated cabling, and e2XHD Patching System. He brings a wealth of experience to the role, with nearly 20 years working in fiber optics and more than 10 years in electronics. He currently supports two TIA fiber optic committees.*

Michael Lawrence
Product Manager
Leviton Network Solutions Europe
Michael *is a highly skilled professional with over 15 years of expertise in the manufacturing industry. With Leviton, he has provided invaluable support for a range of copper, fiber, and pre-terminated solutions. Presently, he specializes in fiber networks and data center applications, and is responsible for implementing Leviton's pre-terminated fiber assemblies and Intelligent Management Solution SMARTPATCH™ across Europe and the Middle East.*

Data centers at the front line of future motor technology

Data centers are becoming more powerful as they gather more data and need high-efficiency motors to match their data consumption, which takes a lot of energy.

Big Data has become the foundation upon which industrial automation providers are building new solutions and services. This comes as the rest of the economy also embraces Big Data as a mechanism for generating previously unattainable insight. According to the United States International Trade Commission, the amount of new data expected to be created in 2025 will be 175 ZB. This would represent a 146 x increase from 2010 levels.

To support this increase in data, infrastructure in the form of data centers is continuing to be built and leased out to companies in need of server space. Research group Cloudscene estimates there are 8,382 data centers globally and the continued growth in demand for data is pushing this number even higher. Data centers are very energy intensive operations and, as such, face constant pressure to reduce consumption. Equinix, one of the largest data center operators, consumed 7,140 GWh of electricity in 2021. Based on wholesale energy rates per MWh published by the International Energy Agency (IEA), this would put Equinix’s yearly energy costs at approximately $378mn which is over 10% of its cost of revenues.

Much of this energy consumption stems from the motors used throughout the facilities. Within data centers, motors are used in fan, pump, and compressor applications to help cool the servers which they house. Cooling is exceptionally important to data

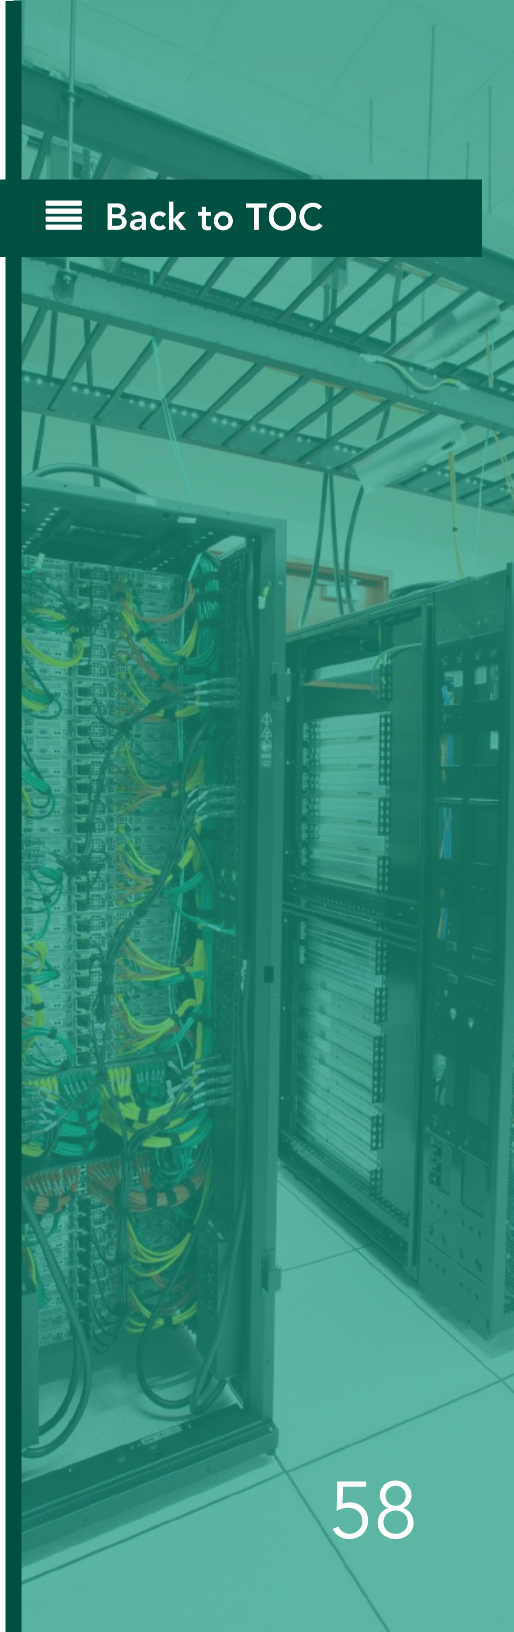

Back to TOC

Back to TOC

centers given the heat generating nature of servers, therefore each facility is equipped with robust cooling systems which can utilize air or liquid-circulation. Both methods require many motors to support the cooling process. The number of motors needed per data center can in some cases reach into the thousands. By some estimations, the motors used in these cooling systems can account for 30 to 40% of total energy use within the facility. Data centers have been hyper-cognizant of the efficiency levels of their motor systems.

High-efficiency motors need sector-like data centers

Data center operators have long hunted increased efficiencies in their motor systems. Historically, this meant utilizing IE3 or NEMA Premium class motors. These motor efficiency classes are defined by their performance across different powers and until recently IE3 and NEMA Premium represented the state of the art. This has changed, however, with improvements in motor technology and increased availability of IE4 and IE5 class motors, which are now offered by most motor vendors.

Unfortunately, when it comes to high-efficiency motors (IE4 & IE5 or NEMA Super Premium), demand has almost always been legislation-driven due to their associated premium price. For the data center sector, it's different; it represents one of the few sectors where there is significant demand for high-efficiency motors coming from the end-users themselves. This has led the sector to become a testing ground for many new motor technologies as data center operators are especially motivated to reduce energy costs from their operations.

We estimate the number of low voltage ac motors sold per year within data center applications to be somewhere in the 450-500k range. When you compare this to the

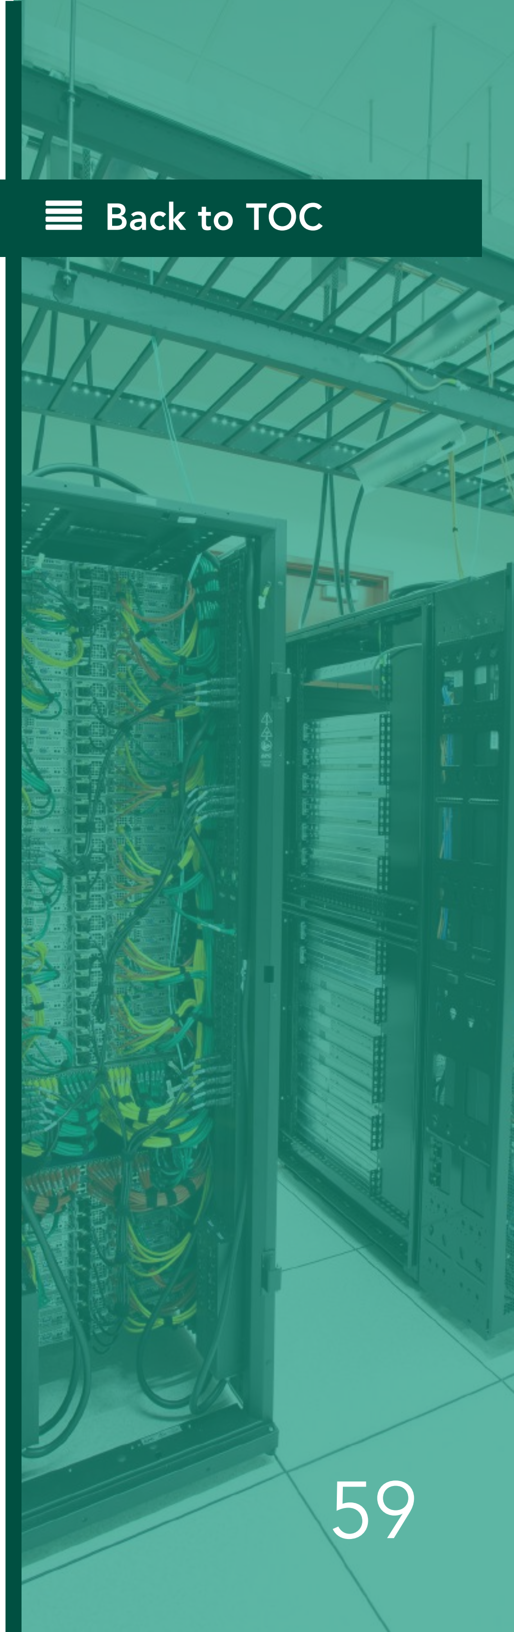

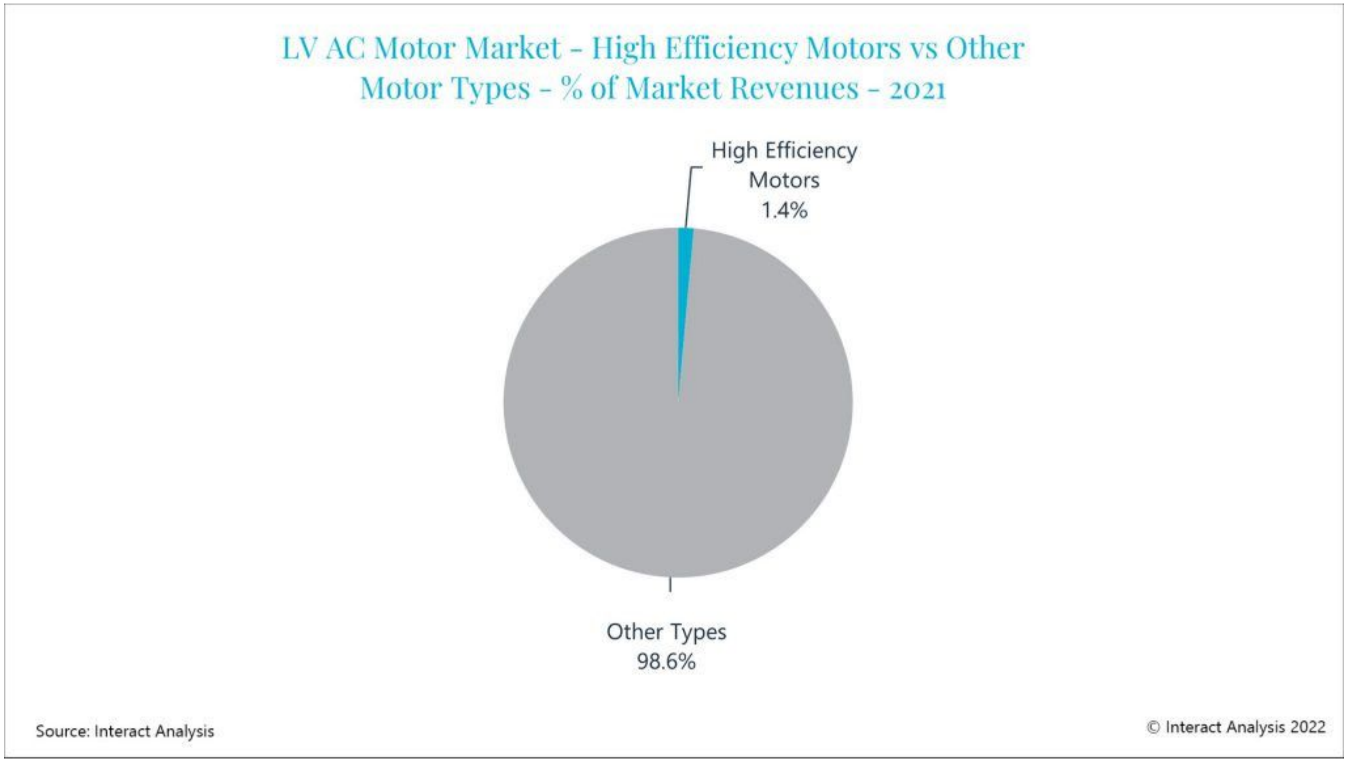

High-efficiency motors are defined being IE4 (NEMA Super Premium) or higher. Data is sourced through reported data from motor suppliers. Courtesy: Interact Analysis

nearly 20mn low voltage ac motors sold globally in 2021, this sector doesn't appear to be particularly special. However, while the sector is small relative to the entire low voltage ac motor market, it is a sizable opportunity when compared to the market for high-efficiency motors.

We estimate the number of high-efficiency motors sold globally in 2021 to be approximately 300,000. When examined in this context, the number of motors sold to data centers becomes much more significant. As established motor vendors and start-ups

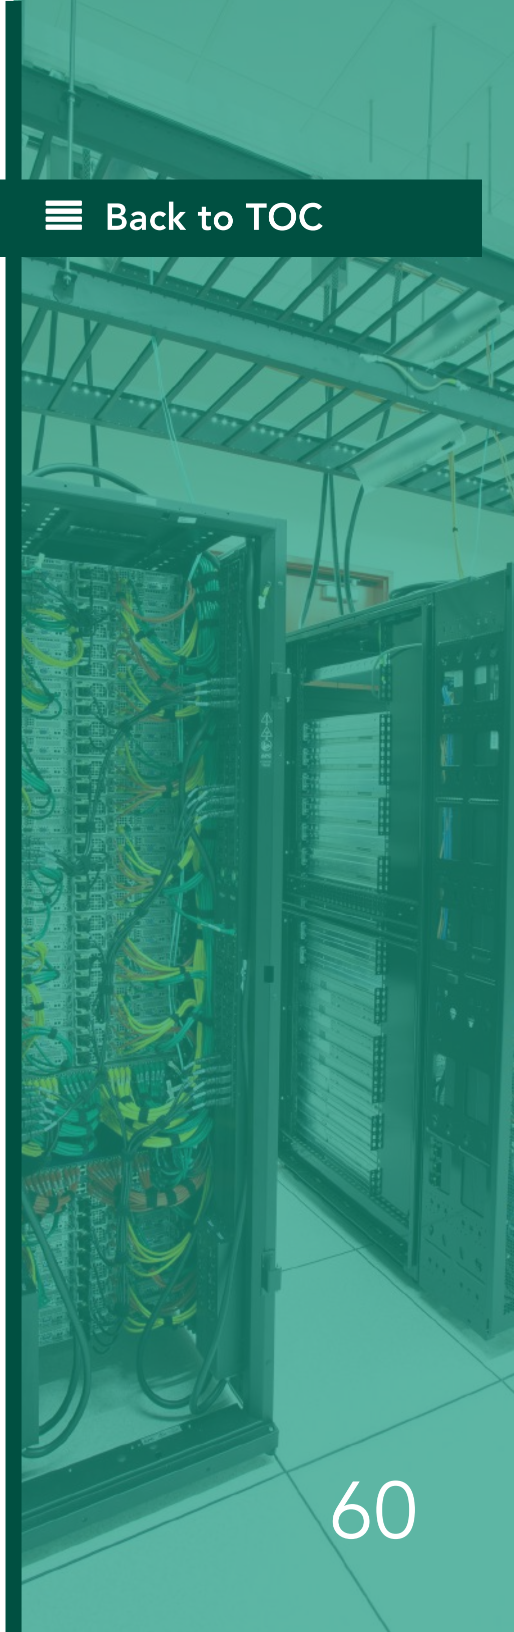

Back to TOC

Back to TOC

alike look to sell and refine their high-efficiency motor tech, it will be data centers who are among the earliest customers.

Motor vendors are benefiting from data centers' willingness to adopt

Most high-efficiency motors are designed for variable torque applications like fans or pumps. There have been several startups which have developed novel technologies aimed at displacing the status quo within these applications. For example, Infinitum, a motor startup who recently raised $80m in Series D funding, has developed an axial-flux motor which it claims addresses many of the shortfalls of traditional motors. In addition to improvements in efficiency brought on in part by its integrated variable frequency drive, Infinitum's motor is significantly lighter and quieter. Both qualities are important in the context of cooling system applications, particularly when they are installed in the numbers often associated with data centers.

Established motor vendors are also developing new technologies to bring higher efficiencies to these applications. In 2020, the largest industrial/commercial motor vendor in the world, ABB, released its EC Titanium motor. This motor is a prime example of the emerging trend of greater integration between the drive and motor. Like Infinitum's product, ABB's EC Titanium contains an integrated variable frequency drive which helps achieve an efficiency level which surpasses the IE5 efficiency class threshold. While IE5 has yet to be officially defined by the International Electrotechnical Commission (IEC), it is commonly understood as achieving a 20% reduction in losses when compared with the officially defined IE4 class.

With any new product release, there needs to be early adopters utilizing the technolo-

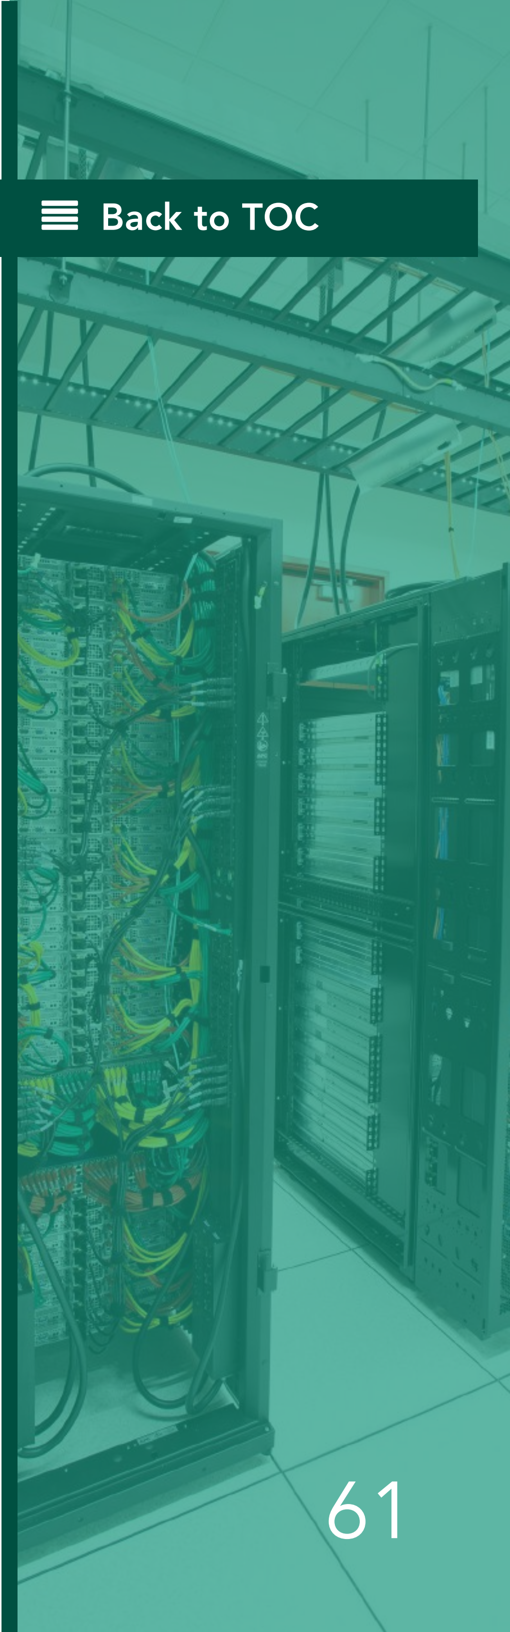

gy to work out the kinks before adoption begins en masse. Data centers appear to be one of these key early adopters for high-efficiency motor technology. As new data centers are constructed, and existing ones look to retrofit new motor technology, expect the sale of high-efficiency motors to grow. This offers motor vendors a sizable application to refine their technology and achieve some level of scalability ahead of the next round of motor efficiency regulations.

Blake Griffin

Graduating Magna Cum Laude with a degree in business and economics, Blake's area of coverage includes various industrial automation technologies such as low voltage AC motor drives and industrial/mobile hydraulics for Interact Analysis.

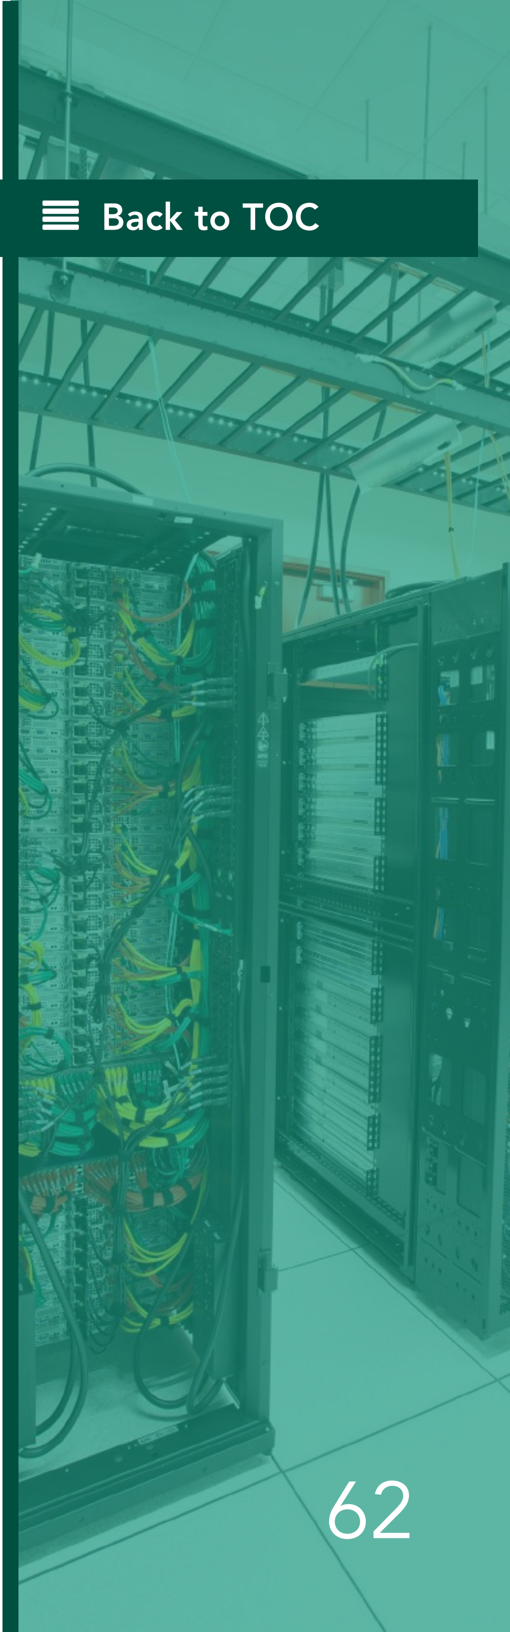

Back to TOC

Grant awarded to make data centers more sustainable

Back to TOC

Penn State was awarded a four-year grant from the National Science Foundation (NSF) to make data centers more sustainable.

Syed Rafiul Hussain, assistant professor of computer science and engineering in the Penn State College of Engineering's School of Electrical Engineering and Computer Science, was awarded a four-year, $206,598 grant by the National Science Foundation (NSF) as a collaborator on a project to make data centers more sustainable.

Data centers, which are the banks of servers collecting and processing massive amounts of information, take an enormous amount of energy, according to Hussain.

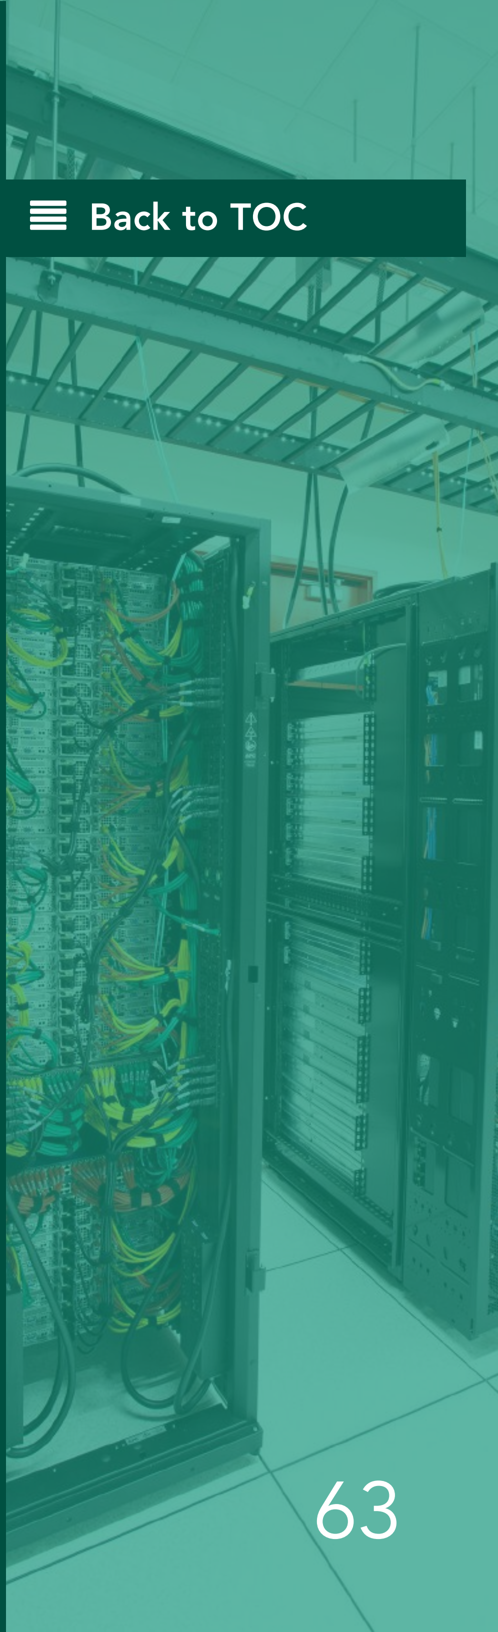

"Data centers already contribute significantly to the global carbon footprint, and the rise in popularity of resource-intensive big data and machine learning workloads is poised to make data center operations unsustainable," Hussain said. "This project designs a suite of Sustainability Aware Software Systems (SASSY) to enable 'sustainable-by-design' data centers."

According to Hussain, SASSY is a measurement framework that focuses on sustainability through a holistic lens — energy source, carbon footprint of computing equipment and device reliability — to measure per-job, end-to-end sustainability costs. With the grant, Hussain and his team plan to develop new programming models, security architectures and tools to enable users to specify their sustainability and performance objectives and empower third-party regulatory agencies and users to verify the sustainability costs. These parameters will guide SASSY to make data-center-wide sustainable management choices.

Ashley WennersHerron

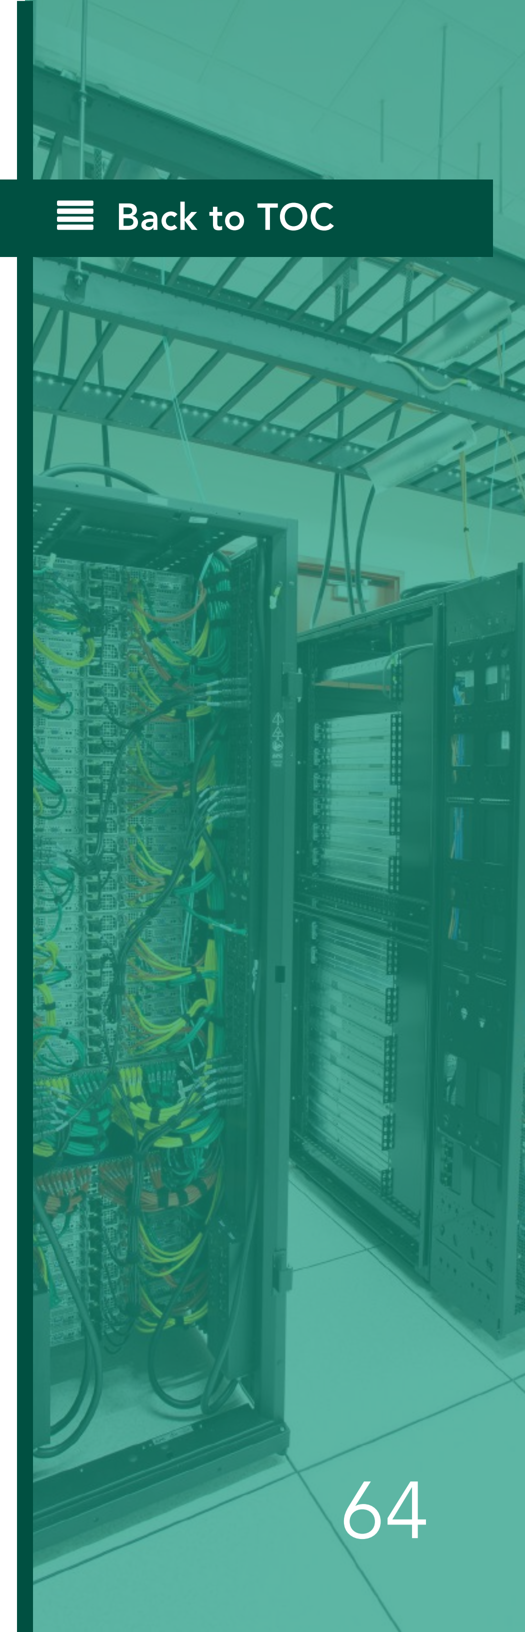

Back to TOC

Content Archive

2023 Winter Edition

2022 Fall Edition

2022 Summer Edition

2022 Winter Edition

Data Centers

Thank you for visiting the Data Centers eBook!

If you have any questions or feedback about the contents in this eBook, please contact CFE Media at *customerservice@cfemedia.com*

We would love to hear from you!