WHITE PAPER

Designing scalable, modular, digital data centers

“Pay-as-you-grow” design considerations to create sustainable, profitable growth for colocation providers

Table of contents

Meeting the data growth demand

The data center colocation market is growing exponentially, with research estimating that this sector alone was worth $51.8 billion at the close of 2020, registering a CAGR of 12.4% during the forecast period 2015-2020.

In addition to the growing requirement for data in all aspects of industry and society, enterprises large and small are waking up to the benefits of using colocation for their IT needs. Better interconnectivity, improved uptime and flexible resource allocation, as well as additional space on site, reduced utility bills and a lesser need for IT expertise are all appealing benefits for the end user and are further fuelling growth in the market.

For colocation providers, the key to sustainable and profitable growth is to open and lease each phase of the data center as quickly as possible, to start leveraging their investment and generating revenue.

For that reason, many colocation providers are building their data centers in smaller blocks so they can open one section while they begin building the next. This limits the upfront investment and the time before revenue generation begins. It also allows providers to secure tenants earlier, which is important in such a competitive, fast-paced environment.

When using this ‘pay as you grow’ building strategy there are many design considerations and scalable technologies that can minimize the cost of subsequent builds and the disruption to operations.

This whitepaper will provide a holistic view of:
- Data center electrification
- Reproducible and standardized topologies
- State-of-the-art scalable technologies

The advantages of designing scalable data centers

Building for tomorrow
Installing vast data centers from the ground up or adding large scale extensions to existing locations requires considerable upfront costs. Often, the space built is not leased straight away and this ‘building for tomorrow’ mentality can result in delayed revenues and ongoing running and maintenance costs for empty server rooms.

Pay as you grow
An alternative, and an increasingly popular method, is to ‘pay as you grow’ – a strategic growth path which requires less upfront cost and can generate faster revenues.

Using modular systems, replicatable and standardized topologies and scalable digital technologies, colocation providers can expand in smaller blocks, repeating the process many times as demand continues to grow.

There are four main benefits driving a growing preference for scalable data centers:

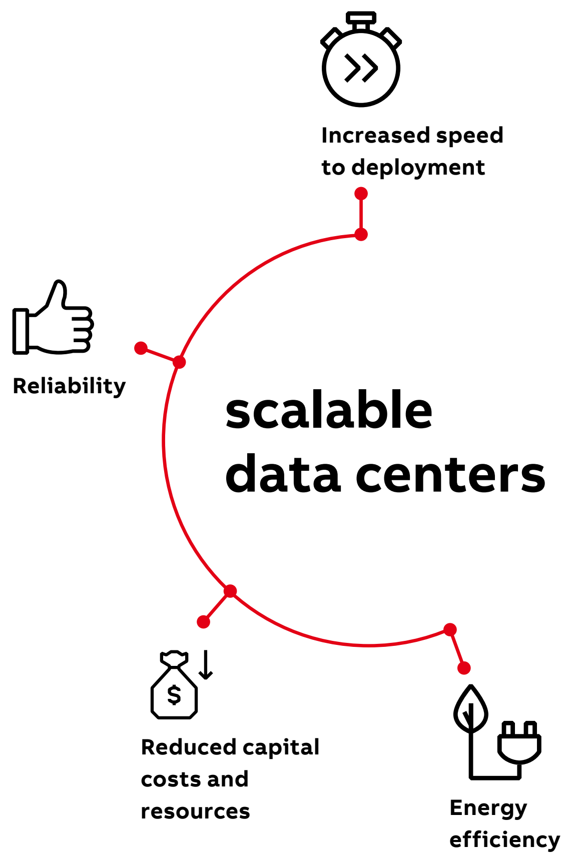

Increased speed to deployment
In a study conducted during the DCD>Building at Scale conference in June of 2020, 39 percent of professionals working in the data center industry said that projects are being deployed in under a year, and 66 percent in under 18 months. To keep up with this demand, data center designers will have to select more scalable solutions.

These lead times are significantly shorter than a decade ago, where data center projects stretched well over 2 years, a build strategy that simply would not meet the data demand of today's rapidly growing market.

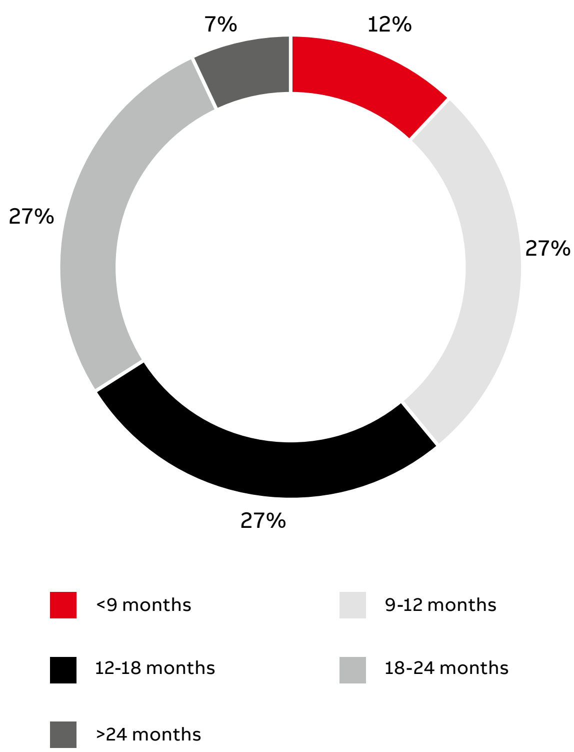

—
Figure 1. Poll results from the DCD>Building at Scale conference suggests 66 percent of data center projects are being deployed in less than 18 months.

Improved reliability
In such a fast pace build environment, using processes that are reliable and predictable delivers many benefits. The consistency and standardization of repeatable designs, reduces risk, maintains a tight build schedule and results in prefabricated systems which are reliable in both performance and in maintenance.

Once scaled, modular designs are reliable having been successfully built, commissioned and used in previous stages. This build method takes the uncertainty out of a campus expansion plan, or the rapid addition of new build centers to an expanding data center campus.

Reduced capital costs and resources
Building for today's requirement using modular build methods allows colocation providers to only invest an amount that will deliver a fast return via rapid uptake of server space and a timely generation of revenue.

By breaking down the upfront investment into smaller chunks, revenues can be generated on each investment before the next modular expansion begins. This allows data centers to be more flexible in their response to market changes. Furthermore, the installation of a modular building block is faster and more efficient onsite, requiring less resources. In turn this translates into lower risk and further time and cost savings.

Energy efficiency
By building in modules, data centers can secure tenants for each phase before building the next. If a data center is built but not occupied, power losses and inefficiencies from equipment such as uninterruptable power supplies will have a negative impact on the site's energy efficiency.

Each modular building block should be around 80 percent capacity to optimize efficiencies. Once a module is around 80 percent utilized ,and at this point planning for the next module can begin.

Introducing the key design considerations for scalability

The task of creating a fortress-like data center that safely and securely stores and manages business-critical data and applications under every conceivable situation, while simultaneously accommodating both short-term and long-term growth, can be daunting.

Overall, this necessitates electrification designs that are more flexible and scalable. These designs should incorporate standard blocks of power, repeated throughout, to allow for future expansion. This approach offers significant improvements over traditional data center designs.

Standardization of design improves operational reliability, yet these designs must still be flexible enough to adapt to a myriad of site requirements.

Design experts recognize these challenges and consider site variables such as utility voltage, total size of the data center, and the optimal design for cooling based on the local climate. Scalable and repeatable designs are based on a standard size IT load that is taken as a building block. Using this tried-and-proven process ensures reliability and maintainability.

4 key design considerations when designing a scalable data center

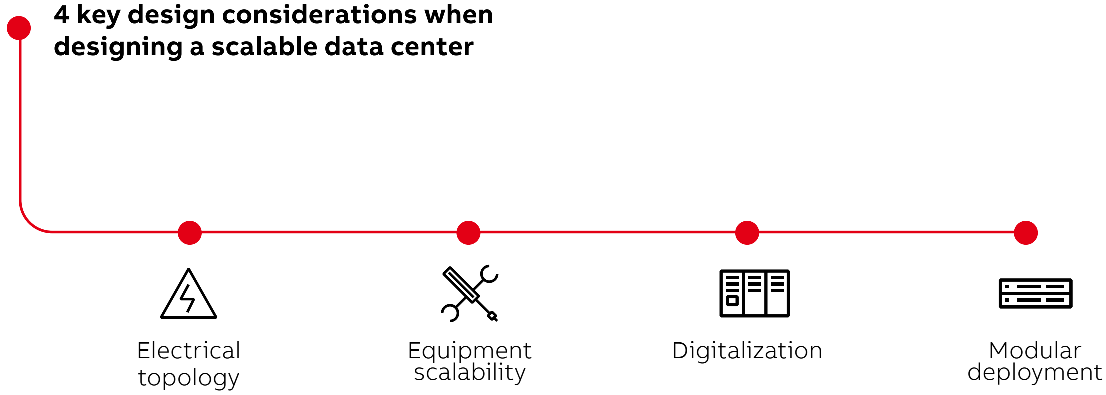

Design consideration 1: Electrical topology

—

Conventional electrical topologies, commonly used in data centers, can be implemented in several different configurations depending on exact project requirements and site conditions. Factors that determine the actual configuration include load kilowatts (kW), available utility service voltages and initial cost.

Scaling electrical topology
Although most data center electrification systems are unique, there are only three main underlying topologies: system plus system, shared redundant and block redundant topologies.

These three basic topologies can be used to create pre-engineered solutions, or standard solution architectures, allowing full integration of equipment for all electrical distribution needs from the utility service to the rack. These sub-systems are factory assembled and tested into power modules (skids or ehouses) so that when the time comes to scale, these pre-engineered, repeatable solutions limit site work and offer a smooth start-up every time you scale.

All pre-engineered solutions can be easily duplicated or modified to serve most data centers. These electrical power solutions typically includes medium- voltage (MV) switchgear, transformers, low-voltage (LV) switchgear, LV switchboards, UPS systems, power distribution units (PDU), remote power panels (RPP), and IT busways – all of which can be combined into integrated skids or eHouses.

System plus system topology
As the name implies, the system plus system topology utilizes two completely independent systems to feed the critical load. The design is based on deploying IT equipment with redundant power supplies (sometimes referred to as dual corded loads). This topology is the basis of design for on-premise data centers, for example, enterprise, financial and government data centers as well as colocation companies.

—
Figure 2. Three most common design topologies for data centers

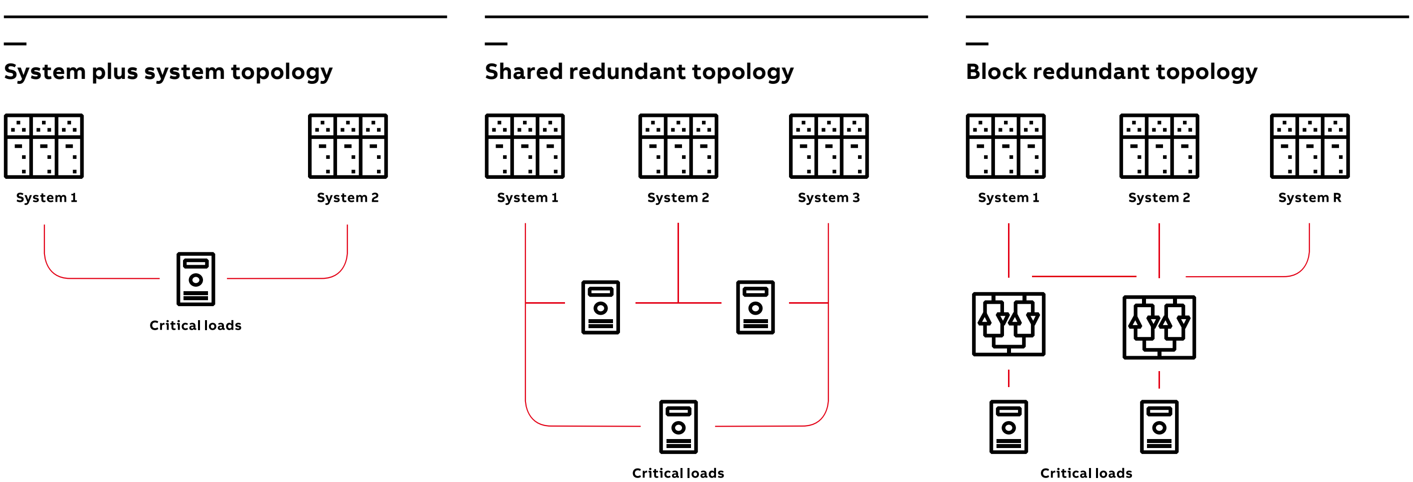

Although the system plus system design has a proven reliability record, costs can be prohibitively high with a maximum utilization of assets of 50 percent and a typical/practical limit of 40 percent.

Shared redundant topology
A variant of system plus system topology is often employed to reduce overall cost. This is called shared redundant or system + utility. Here, the "system" has N+1 uninterrupted power supplies (UPS) while the "utility" has no UPS.

Colocation, hyperscale and cloud data centers generally use the shared redundant topology, which comes in a variety of configurations. The design is normally designated by the number of systems over the number of loads, such as 3N/2 for three systems and for two loads, or 4N/3 for four systems and for three loads.

For example, using 1 MW blocks of IT load, a 3N/2 system would have 3 MW of capacity feeding 2 MW of IT load.

With this configuration, customers can improve utilization to 66 percent or even to 75 percent with 4N/3. Consequently, this topology serves as the basis for the design of many colocation and hyperscale data centers. While the use of the shared redundant topology does improve asset utilization, it nonetheless requires operators to monitor loads.

Block redundant topology
Block redundant topology, also known as catcher topology, utilizes a Static Transfer Switch (STS) to transfer the critical load from the primary or active system to the reserve or catcher system. Such a topology is used by both hyperscale and colocation data centers.

For data centers with single corded loads, this topology is usually the most cost-efficient design. With the block redundant topology, an asset utilization of 80 percent is possible and there is no need to constantly monitor loads to maintain redundancy. The primary disadvantage of this topology is the reliance on static transfer switches; this increases the cost and complexity of the design.

Block redundant topology depends on the ability of the catcher or reserve UPS module (or modules) to handle a step load. The active UPS can be loaded to full capacity and the reserve UPS has no load in normal operation. The reserve system can be larger than the active systems. Block redundant can be applied with single cord or dual corded IT loads. As such, if single corded IT loads are used, the static transfer switch is a single point of failure.

Design consideration 2: Equipment scalability

To ensure your data center design is modular and scalable, it is essential to select equipment that is scalable. Switchgear, uninterruptible power supplies (UPS), power distribution units (PDU) and remote power panels (RPP) are all examples of scalable equipment. Get this right and specifying future expansions will be time and cost efficient.

MV GIS Switchgear
Selecting medium voltage Gas Insulated Switchgear (GIS) that can be scaled up without the need to empty the gas from the system makes scaling quicker and easier without disrupting service.

Market-leading flexible switchgear designs can include a three-phase encapsulated, arc-resistant switchgear for single and double busbar applications with separate gas-filled busbar and circuit breaker compartments.

Specifying equipment with plug-in technology enables safe, fast and easy installation without the need for special tools.

UPS built for flexibility
The UPS plays a vital role in ensuring clean and continuous power for the critical infrastructure. A modular UPS based on a decentralized parallel architecture (DPA) allows data center designers to scale the power on demand and lower upfront investment.

By using the DPA design, every power module has its own fully independent entity, with its own control and all required hardware for autonomous operation. With DPA architecture, single point of failure is removed providing continuous uptime even during maintenance or system expansion.

Power distribution units (PDU) and remote power panels (RPP)
Power distribution units (PDUs) and remote power panels (RPPs) help to meet the demands of power-intensive applications, delivering unsurpassed power monitoring and distribution in a safe, reliable, space-saving footprint. PDUs/RRPs can be configured as needed and they are scalable regarding the number of outputs and required output power, ensuring continuous power to critical applications.

Furthermore, market leading systems of this nature can help panel builders to dramatically reduce the cost of certification and facilitate easy post-installation alterations such as the ability to change or extend the number of servers connected.

—
Figure 3. Typical data center one-line

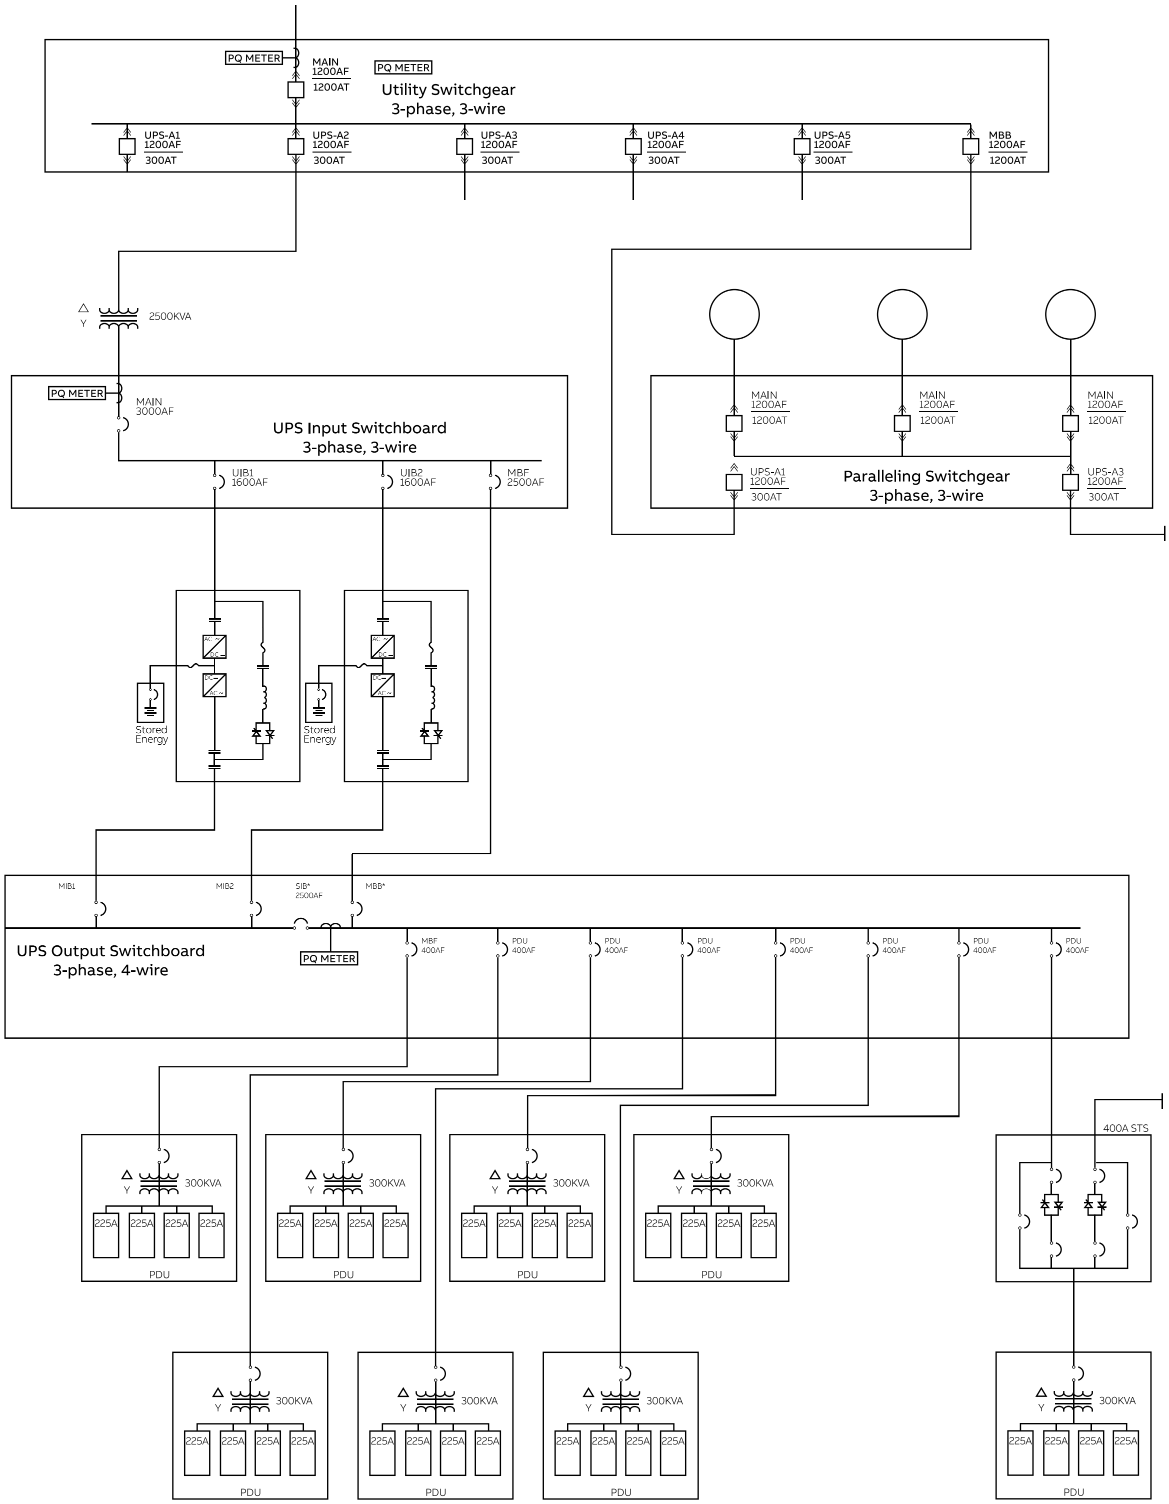

Case study: The importance of scalable equipment

The challenge
Scalability across all equipment specification is critical in implementing a pay-as-you-grow data center design strategy. If some of the equipment selected is not scalable, it can cause considerable cost and time delays in future expansions, as US Colocation Data Center, Volico discovered:

Volico, like many progressive colocation providers, had a pay-as-you-grow strategy in place, but the UPS specified in the initial design was not scalable and this posed a major expansion roadblock. Even though all other equipment was scalable, this one specification error had put a stop to Volico's pay-as-you-grow plans.

The solution
Needing a UPS that was scalable and flexible, Volico opted to replace its current UPS equipment with ABB's modular DPA UPS system, comprising two frames of 10 module sets delivering 1MW of power.

The Conceptpower DPA 500 can scale vertically, up to 500 kW in a single frame, and horizontally by adding up to six parallel frames, to a total of 3 MW of power.

The result
This replacement UPS allows full system scalability, allowing designers to add more power every time Volico expands simply by inserting another module. The system is based on unique slide-in UPS module sets, each rated at 100 kW. Modules can be inserted or removed from the frame while the system is running securely in double conversion.

Using online-swap modularity, combined with the system's full redundancy helps data centers achieve six-nines availability (99.9999 percent). The DPA 500 provides total reliability with more than 96 percent energy efficiency, reducing total cost of ownership compared to other UPS systems.

The ABB modular DPA UPS
ABB offers a best-in-class modular UPS with an efficiency rating of 96 percent, a compact footprint and low TCO.

There are a range of options for scaling the ABB modular DPA UPS. These include the 250 S4, which is ideal for mid-size data centers with single power capacity of 300kW based on power modules of 50kW. In addition, the new MegaFlex DPA offers power capacity up to 1500kW based on power modules of 250kW.

For Volico, this modular design perfectly resolved the short-term need to replace a system that was a barrier to expansion, while giving scope for future growth.

Design consideration 3: Digitalization

—

There are many benefits to adding intelligence to electrical equipment using open protocols (such as IEC 61850), but for scalability, the most important benefit is that it reduces wiring, minimizing installation and commissioning time.

IEC 61850 provides a standardized framework for substation integration that specifies things like communications requirements, functional characteristics, structure of data in devices, naming conventions for the data and more.

IEC 61850 is now being increasingly used across smart grid, renewable, and process industry applications and is benefiting equipment such as protection relays, circuit breakers, communication gateways, programmable logic controllers (PLCs) and supervisory control and data acquisition (SCADA) architectures.

Together, these devices make it possible to design and operate a fully integrated protection and supervision system that spans the voltage ranges required. This comprehensive approach is made possible by exploiting IEC 61850 to create features such as advanced logic selectivity based on device-to-device communication, real-time diagnostics and integrated engineering.

IEC 61850 is ideally suited to the automation of data center power infrastructure. For example, in traditional analog switchgear, each component is physically connected with copper wiring, while digital switchgear requires a few optical fibre wires between cabinets. This reduces the wiring by up to 90 percent and saves time on assembling, testing, installing and commissioning the system.

Systems of this nature are also much more scalable because changes to the configuration can be done remotely using software, as opposed to changing out hardware or reassembling wiring.

There are further resources on the importance of digitalization in scalable data center design here.

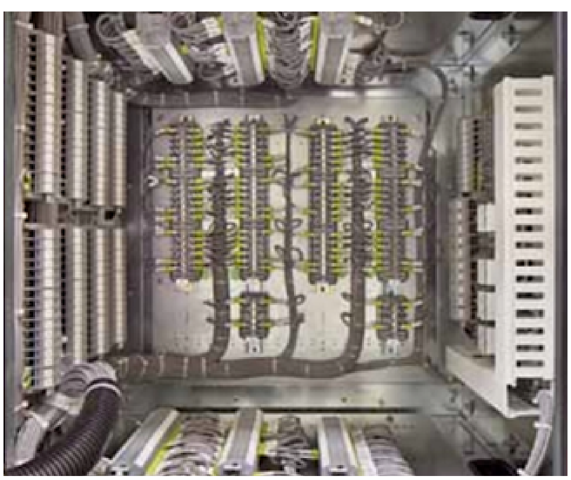

—
Manual switchgear

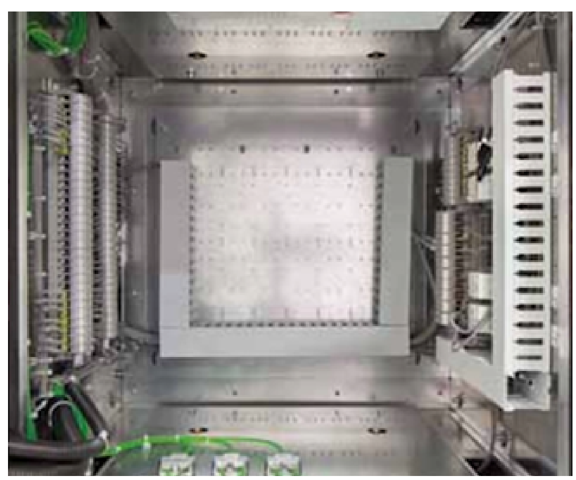

—
Digital switchgear

Design consideration 4: Modular deployment

—

With topology and scalable equipment specified, the fourth design consideration is modular deployment. Pulling the system together in a modular way will improve its scalability and have less impact on operations. In particular, modular deployment is centered around skid units and eHouses.

The poll taken during the DCD Building at Scale conference also showed that 74 percent of data center professionals are already dealing with some type of modularization:

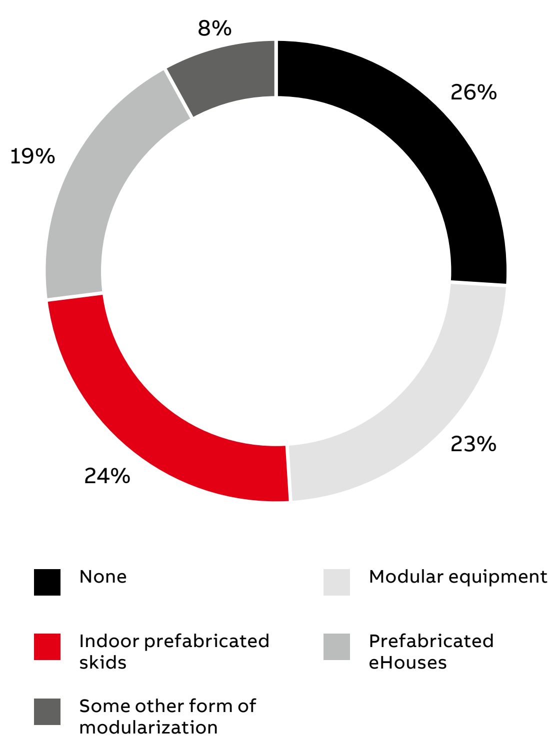

—
Figure 5. Poll results from the DCD>Building at Scale conference suggests 74 percent of data center professionals already have some experience with working with modularization.

Modulization with skid units and eHouses
Skid units and electrical house (eHouse) packages enable data centers to address speed to market challenges. Common to traditional construction methods, packaged solutions are typically supplied as individual components with installation and interconnections provided by third parties.
Such a solution is an open frame mounted, compact unit with factory installed equipment and interconnections. Because they are pre-engineered, prefabricated and pretested, indoor skid solutions can be built off-site parallel to other construction efforts, thereby accelerating the construction schedule.

The eHouse is a prefabricated, pre-engineered and pretested system; it is an environmentally controlled building with factory installed equipment and interconnections. Testing all components prior to shipping ensures that risks in the field are minimized. These units are easy and cost-effective to install and can be placed close to the main loads. Although, usually permanently installed, they can be relocated as a colocation center grows or a customer's IT equipment is changed.

Skids units and eHouses can both be used to accelerate construction schedules and easily expand existing capacity. These one-piece systems can be installed with limited impact on existing operations, and reduce on-site work limiting the number of people on site at an operational data center.

Case study: Solution architectures for scalable data center design

The challenge
USA-based GIGA data centers had acquired a long, rectangular building in North Carolina suitable for conversion into a data center. Working with ABB, the brief was to create a critical power distribution design that was fully scalable to allow for further pay-as-you-grow expansions.

The design needed to be scalable, so that GIGA could offer flexible and modular data center technology at a competitive cost to all its customers.

The solution
GIGA selected ABB as its design partner based on the expertise of its team in providing flexible, innovative technical electrification solutions within a challenging time frame.

Together, GIGA and ABB worked to design a system that would support 60 MW of IT load. Amazingly, the first phase of the data center conversion was completed in less than six months.

Subsequently, ABB collaborated with GIGA's mechanical, electrical and plumbing consultants to design a flexible, scalable and efficient packaged solution based on a system plus system topology. This design is scalable and has the capability of expanding in increments of 2 MW of IT load.

The result
ABB's electrification solution for GIGA Data Centers included TLE UPS modules (to support the IT server load) and a decentralized parallel architecture (DPA) UPS for the site server load. In addition to uninterruptable power supplies, the solution included UPS maintenance bypass cabinets, lithium-ion battery systems, lighting panel boards, and dry-type transformers. The DPA UPS was chosen for its effectiveness for lower power requirements and scalability.

ABB's packaged solution included all project management and field services from the start-up and commissioning of all equipment to the commissioning of all systems. ABB's successful end-to-end solution is a compact and efficient data center that delivers the power distribution and protection performance GIGA needs today with the potential to expand in the future.

Conclusion and key takeaways

The increasing need, and desire, for data shows no sign of slowing. Data center owners and operators need to design data centers that are capable of responding, not only to the needs of today's data demand, but to the data requirements mid to long-term too. The potential for future expansion should be at the forefront of design, in the selection of technologies and the construction of new centers.

Adding standardization and modularity to planning, together with digitalization strategies will ensure future-proof data centers are designed to meet standards for global expansion, in a way which is cost and time efficient.

Ultimately, it is about using clever design and technology to ensure scalability and speed of build, without compromising on future performance.

Key takeaways

- Using prefabricated and modular construction techniques in data center design and the use of skids and eHouses allows a pay-as-you-grow approach to expansion which is faster, energy efficient and more cost effective, with minimal disruption to operation
- By scaling in line with demand, upfront cost is kept low and revenues are generated faster
- There are a number of important considerations when planning a scalable data center, they are topology, scalable technology, digitalization and modular deployment
- Solutions architectures can provide a blueprint for each expansion, maximising investment, energy efficiency and speed of build
- ABB Data Center's team are highly skilled in the design of scalable data centers and offer solutions architectures and a range of scalable technologies

To discuss your pay-as-you-grow expansion plan, or how your data center can benefit from scalable data center design, click here to request to speak to a technical expert.

ABB Ltd.
ABB Inc.
305 Gregson Drive
Cary, NC 27539

www.abb.com/datacenters

9AKK107992A4185 REV A JUN 2021