Best-practice strategies for designing and deploying modular data centers

Leverage scalable, flexible designs to control costs while enabling business growth

IBM

Contents

Introduction

An explosion in IT demand has put mounting pressure on data center capacity, making it difficult for CIOs to respond to business needs while budgets remain mostly flat. As capacity demands become increasingly unpredictable it is harder to keep up with rapid changes in technology and emerging models such as cloud computing. The challenge is to build a data center that will be functional and cost-effective to operate for 10 to 20 years—while information technology changes every two to three years.

The business requirements for designing a data center have remained fairly consistent around availability (high), capacity (high) and capital cost (low). But these three requirements alone are no longer sufficient. Organizations also need to provide the scalability to upgrade technology and capacity over time and the flexibility to support new technology and computing models while minimizing both capital and operational costs.

All these factors make designing for flexibility critical, which is why modular data centers are becoming the design of choice for many organizations. This white paper examines best practices for designing and building modular data centers, derived from more than 500 IBM implementations for companies of all sizes. It also describes the benefits of using advanced analytics to help establish an effective design strategy.

Expanding the concept of modular data centers

There are a number of offerings in the marketplace today designated as modular data center solutions. One of the most often cited is the data center shipping container, a valuable resource in cases where data center mobility may be crucial to disaster recovery or when temporary resources may be needed to bridge gaps while new construction is being completed. Other modular data center offerings include modular subsystems such as uninterruptible power supplies (UPS) and rack- and row-level cooling that can be added as needed. Some modular data center offerings are prefabricated units that are much like container offerings in that they provide the ability to easily add floor space, power and cooling as demands increase.

While those solutions have their place, IBM has expanded the definition of the modular data center to more effectively address today's needs for flexibility and cost-effectiveness. IBM's approach to modular data center design is not about containerization but about using smaller increments of standardized components, making it possible to add data center capacity when needed and thus more closely align IT with business requirements. It is a concept similar to the design of laptop computers, which have base components like memory chips, processor chips and disc drives—plus they can accommodate additional components when needed, such as additional memory for a new software application or a second hard drive for more data. These new components can be simply "plugged in" to ready-to-use sockets or bays.

IBM approaches data center design in the same way, by enabling organizations to simply "plug in" additional cooling capacity, generator capacity, battery and uninterruptible power capacity when it is needed. In IBM's modular data center design, the electrical and mechanical systems (cable, conduit, piping capacity and so on) are built to accommodate the ultimate capacity of the facility beforehand, rather than retrofitting these systems as capacity needs grow. This enables the data center capacity and configuration to be upgraded seamlessly without disrupting

operations, and it provides the flexibility to easily integrate future technologies such as cloud computing, high-density systems or new water-cooling techniques. Importantly, data centers constructed in this way can be tailored to individual business and IT requirements, with a variety of design choices that can meet the needs of both midsize and large companies.

By building in smaller increments—or modules—organizations pay to build only what is currently needed, with the ability to scale up capacity to meet growing demands as required. This approach has saved each IBM client hundreds of thousands in deferred capital costs for equipment purchases as well as deferred operational expense. Closely matching capital and operating costs to changing business needs helps provide a virtually unheard-of level of fiscal precision in data center planning, one that will likely be appreciated by the CFO.

Making decisions based on needs instead of wants

Every data center design must accommodate business and IT requirements for availability, capacity and cost efficiency. IT leaders are often faced with requests to design a new data center that can provide the highest levels of availability as well as support the most current, highest density IT equipment. The key is to identify what is actually needed by the business and balance that with a cost-effective design.

Availability is at the heart of what every data center must provide. Even for planned outages, organizations have increasing application availability demands and decreasing tolerance for downtime. With application outage windows getting shorter and shorter, mitigating outage risks is critical to business operations. After all, the revenue cost of a single application outage could run from the hundreds of thousands of dollars to millions of dollars per hour, depending on the industry and circumstance. Still, application availability must be maintained in a cost-wise manner.

Organizations can always implement more redundancy in the infrastructure—more power supplies, more batteries, more generators or even more data centers—to reduce the chance of an outage. But at some point the cost of redundancy outweighs the benefit. Companies should continually ask, "Is it worth it? How much outage insurance do we really need?" Not all businesses require the same degree of data center reliability.

For example, when establishing resiliency requirements for its own data centers, IBM found that moving from IBM Level Three (99.982 percent availability) to IBM Level Four (99.995 percent availability) represented a 39 percent increase in capital expenses. IBM adapted its design to achieve the highest levels of availability—at 25 percent less capital cost—by including principles such as concurrent maintenance, using distribution-level static switches to provide additional redundancy, N+1 generators, chilled water storage and 24×7 onsite facilities coverage for maintenance support.

Assessing future IT equipment requirements can be equally challenging with the accelerating pace of change. It requires knowing how much capacity is actually needed while estimating the power density that will be required to support it. In the past decade, industry trends have indicated that the power density required by technology has grown more than 20 times and is now approaching about 40 kilowatts (kW) per high-density server rack. However, Datacenter Dynamics surveys have found that typical data centers only use about 2 to 6 kW per server rack, depending on the geographic region, and they can generally plan to support a maximum of 15 kW per rack.[1]

Using a modular design allows organizations to support this range of expected power densities and to do so with minimal disruption to the operation. Some best-practice modular design principles that enable this level of support include sizing electrical and mechanical systems to meet future capacity needs, provisioning for potential water cooling, using modular UPS and chiller plant designs, and using electrical switchgear that can scale to support growth. Deploying a modular design from the outset is like an insurance policy. As shown in Figure 1, organizations pay a five percent premium in capital costs upfront, allowing them to avoid paying 60 percent more for upgrading power capacity later—while also avoiding a business-impeding retrofit.

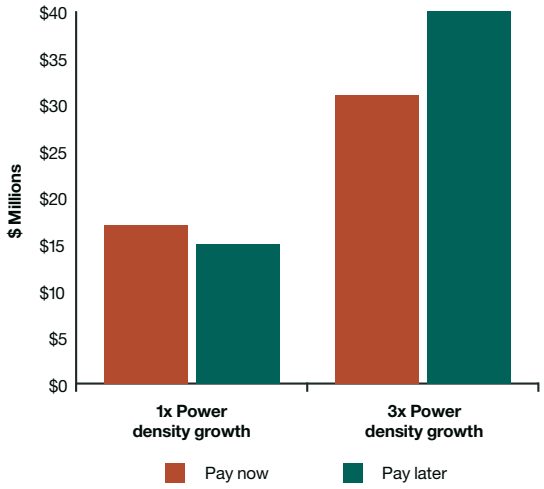

Figure 1: In some cases, an additional five percent upfront investment in modular design can yield a long-term savings of 60 percent for capacity upgrade and can help avoid a disruptive retrofit.

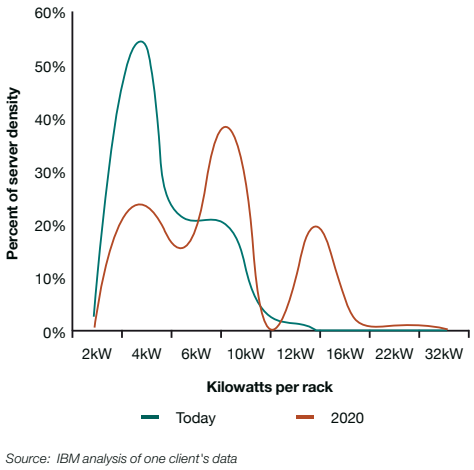

Source: *IBM analysis of one client's data*

Figure 2: Using the output from a data center strategy analytic tool, a large bank was able to determine that by 2020, 60 percent of its server capacity would require more than 10 to 20 kW per rack.

Use analytic tools to determine needs

With technology changing every two to three years, it can be a challenge to design a data center that can support business growth cost-effectively. One best practice is to outline a clear and detailed statement of requirements before data center design begins. Analytic tools geared around data center strategy should also be used. These tools are designed to help companies balance business requirements with objective analysis, enabling them to make sound decisions based on their availability, capacity and total cost of ownership (TCO) requirements.

One such tool used by IBM forecasts IT growth and associated physical infrastructure capacity so that the data center design can accommodate these increases. With input that includes an organization's expected application growth, IT strategy and current data center capabilities, the tool forecasts business and IT demands on the data center capacity thresholds to predict energy and space capacity requirements. Figure 2 shows how output from the tool was used to help a large bank determine required server capacity through 2020.

Over the ten-year useful life of the bank's data center, the projected business growth and the adoption of new technologies such as virtualization and cloud computing would result in a shift in the expected server density—necessitating a potential two-to-four times increase in power density. By using these calculations, the bank avoided under-building its power and cooling capacity and the expense of redesigning the data center at years five, seven and ten. Instead, the bank's modular data center is predesigned to seamlessly expand capacity by three-to-five times its current power density to support projected IT changes and technology adoptions. What's more, by integrating dual-zone power and cooling capacity into the design, the bank will save US $1 million in annual energy costs.

Establishing control over how decisions are made

Establishing governance over data center investment decisions is essential to a well-designed and well-run data center. Building a new data center—large or small—is a significant capital expense. However, while capital costs typically garner 100 percent of the

attention, they are only about 15 percent of the total costs of building and operating the facility, making it imperative for IT leaders to understand who the stakeholders are for managing both kinds of costs. They need to know who is responsible for funding the data center build, and who is responsible for maintaining it. The CFO and board of directors typically oversee capital investments; data center operating expenses usually fall under the real estate organization or under the CIO.

IBM has found that designing to optimize the overall square footage of the data center is no longer appropriate, as the cost of the physical "shell" only makes up ten percent of the total capital costs (not including the cost of real estate). By contrast, about 60 percent of the costs of building a new data center are derived from the electrical and mechanical equipment, including UPS, chillers, air-conditioning systems, battery backup systems and generators (see Figure 3).

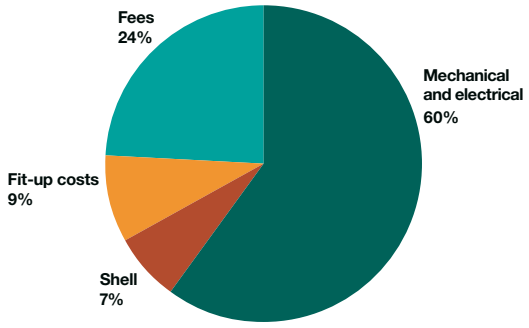

Source: IBM estimates

Figure 3: When building a data center, 60 percent of the capital outlay is used for mechanical and electrical systems while only seven percent goes to constructing the building shell.

IBM has also found that over the 20-year expected life of a new data center, operational expenses typically are three-to-five times the upfront capital investment. This means that a data center that costs roughly $50 million to build will incur $250 million in operating expenses over its 20-year life cycle (see Figure 4). Certainly, the importance of understanding these cost ratios cannot be overstated. It is critical to optimizing both capital and operating expenditures over the long term.

These findings also make clear the advantage of designing and building in smaller modules and putting off major investments in the mechanical and electrical equipment until they are needed. Organizations can defer upfront capital expenditures by as much as 40 percent and operational expenses by as much as 50 percent.

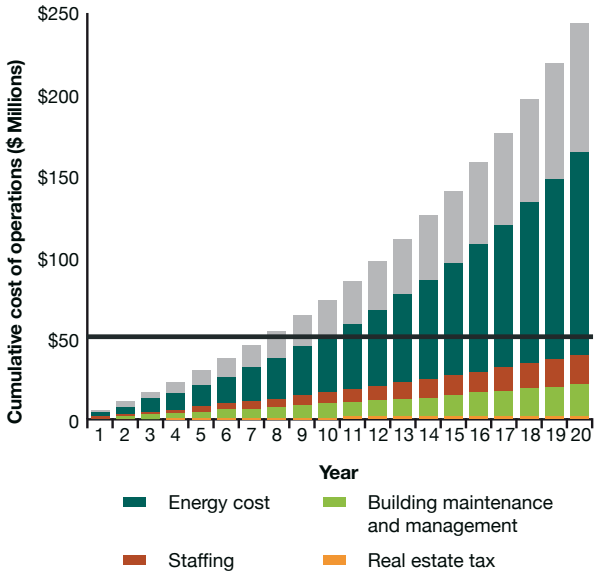

Figure 4: Operating costs can be as much as five times the capital costs; a $50 million data center (black line) will cost about $250 million in operating expenses over its lifetime.

Use analytic tools to balance capital and operating costs

One way to improve investment governance is to assess which technology components will cost less to purchase versus which will cost less to operate over the 15- to 20-year life of the data center. However, the least expensive equipment purchase does not always equate to the most cost-effective choice. Making technology and vendor tradeoffs is critical, but if the decision makers are only taking capital costs into consideration, then the long-term operating costs may be overlooked.

IBM uses a data center life cycle cost analytic tool to provide insights into both capital and operating costs and help evaluate the right vendor and technology tradeoffs to balance those costs. Consider the findings for one 20,000-square-foot data center, where IBM analyzed the cooling systems of 20 possible vendors grouped in three technology categories (see Figure 5).

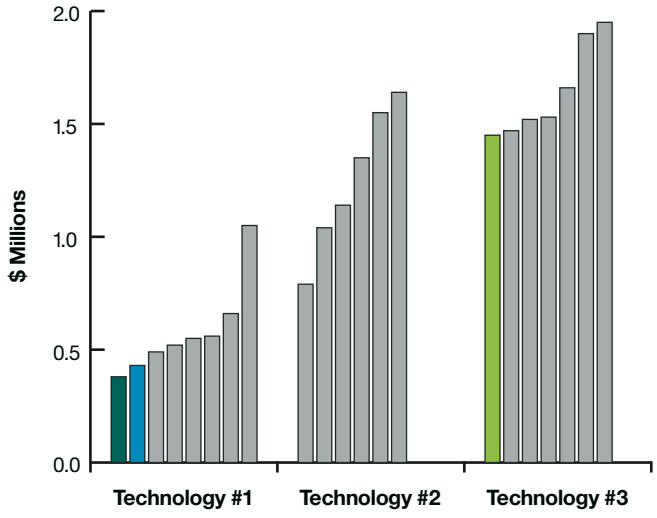

Source: IBM estimates

Figure 5: Capital costs among technology vendors evaluated for one IBM client varied by as much as five times for one 20,000-square-foot data center.

IBM found that equipment costs varied widely, and the size of the capital investment did not always equate to the operating costs over a ten-year period (see Figure 6). In fact, the analysis determined that the technology that cost the least (indicated in teal in Figures 5 and 6) was the most expensive to operate in the long run. However, the cooling system that cost $20,000 more to deploy (indicated in blue) than the lowest cost solution resulted in a $1.5 million operating cost savings over ten years. Conversely, equipment from the low-cost vendor in the highest-priced technology category (indicated in light green) cost the least to operate, but even those savings did not cover the higher initial capital investment.

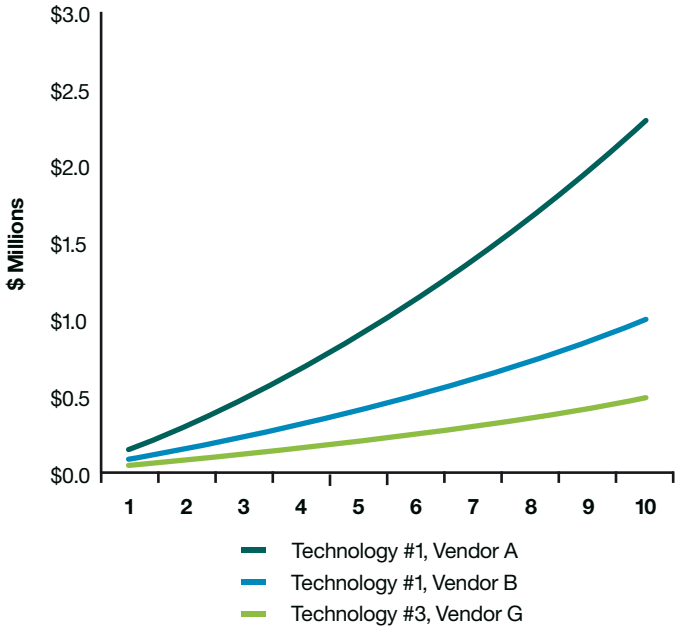

Source: IBM estimates

Figure 6: IBM's analysis shows that the technology that required a $20,000 higher capital investment (shown in blue) produced a $1.5 million savings in operating costs over ten years, producing the better value for the client.

Total cost of ownership has been a staple for evaluating servers for decades but rarely is considered for the data center facility. Using analytics to forecast TCO prior to making data center investments, organizations can reduce total data center operating costs by 15 to 30 percent over a data center's useful life—and can potentially save enough in operating costs to pay for the data center build.

Finding a trusted partner from concept to completion

Two out of three organizations turn to a trusted service provider for help in designing a new data center. When choosing a vendor, look for one that can provide data- and fact-based analysis, derive actionable recommendations and assist in deployment. Steer away from companies offering a fixed architectural approach, with limited ability to address unique IT and business requirements. It is also important to choose a data center service provider that can model alternative designs, recommend vendor and technology tradeoffs, help justify expenses for higher initial investments and suggest ways to reduce operating expenses. Always take into account your business goals, and make sure the provider will assist in meeting those. Lastly, understand available skill sets (and limitations) within your own facilities and IT organizations, and seek a provider that can offer complementary expertise.

Additionally, look for a provider that can:

- Coordinate specialists across IT and facilities
- Employ a variety of advanced analytic tools to determine data center strategy and provide cost insights
- Demonstrate full transparency on capital versus operational costs to help make strategic vendor trade-offs
- Draw on extensive experience with newer technologies—such as backup power solutions and bringing water cooling into the data center—as well as global capabilities for those organizations that are located worldwide
- Provide research-based guidance and recommendations in addition to well-engineered, cost-effective data center design choices
- Employ local contracting licenses and working contacts to keep costs at a minimum and maintain project control

Choosing the right modular design

From its experience with hundreds of modular data center implementations, IBM has found that no two organizations' business and IT requirements are alike. Fixed architecture-based data center designs such as a single-story layout assume "one size fits all clients" and often fail because they are rarely able to address each organization's full scope of requirements.

IBM offers a comprehensive set of modular data center offerings to satisfy a range of cost, scalability and modernization requirements, including:

- Scalable modular data centers—ideal for midsized, growing businesses that need to replace or consolidate servers to reduce sprawl or that need to modernize aging infrastructures. IBM has found that a scalable modular data center design can help reduce TCO by 15 to 25 percent over traditional data center models.
- Enterprise modular data centers—best suited for large companies that are likely to expand their data center capacity and need to defer capital and operating expenses by between 40 and 50 percent.
- Portable modular data centers—suitable for organizations that require a temporary solution to provide capacity while they take 18 to 24 months to build a new data center. They provide a fully functional data center that brings computing capacity closer to customers, manufacturing sites or remote operations.
- Higher-density-zone data centers—a good choice for older data centers that need to support heat-generating blade servers. Using this design approach, IBM clients have realized a 35 percent cost reduction over a site retrofit.

Summary

Modular data centers—with their inherent flexibility and scalability—provide a significant opportunity for CIOs to adapt to their existing IT environments for change. They can help companies adopt new technologies, cost-effectively align capacity to business needs and begin to factor in energy efficiency as a key IT operating metric. These innovative approaches to data center design and build enable CIOs to play a pivotal role in supporting the CEO's business objectives, specifically, faster response to change, reduced costs, and better alignment between business and IT.

IBM can help CIOs seize this opportunity. In fact, IBM owns or manages more than 8 million square feet of data center space and has designed and built more than 30 million square feet of data center space for clients in more than 40 countries. Building on that know-how and experience, IBM uses advanced analytics to help clients establish an effective design strategy. IBM provides a range of choices ideal for both midsized and large clients. Today, with the flexibility to choose and the potential for significant cost savings, 80 percent of IBM clients are selecting a modular data center design over a traditional design.

To meet the business and IT demands of today and the unpredictability of tomorrow, organizations need more flexible ways to manage capacity growth and integrate new technology into the data center environment. Modular data center designs hold the key to such flexibility, enabling continuous IT capacity growth while cutting operating costs and helping to minimize environmental impact.

For more information

To learn more about how your organization can plan and successfully implement a modular data center, please contact your IBM marketing representative or IBM Business Partner, or visit the following website: **ibm.com**/services/smarterdatacenter

IBM Global Technology Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
September 2011

[1] Datacenter Dynamics Datacenter Research Group, January 2011.

Please Recycle

SFW03015-USEN-00