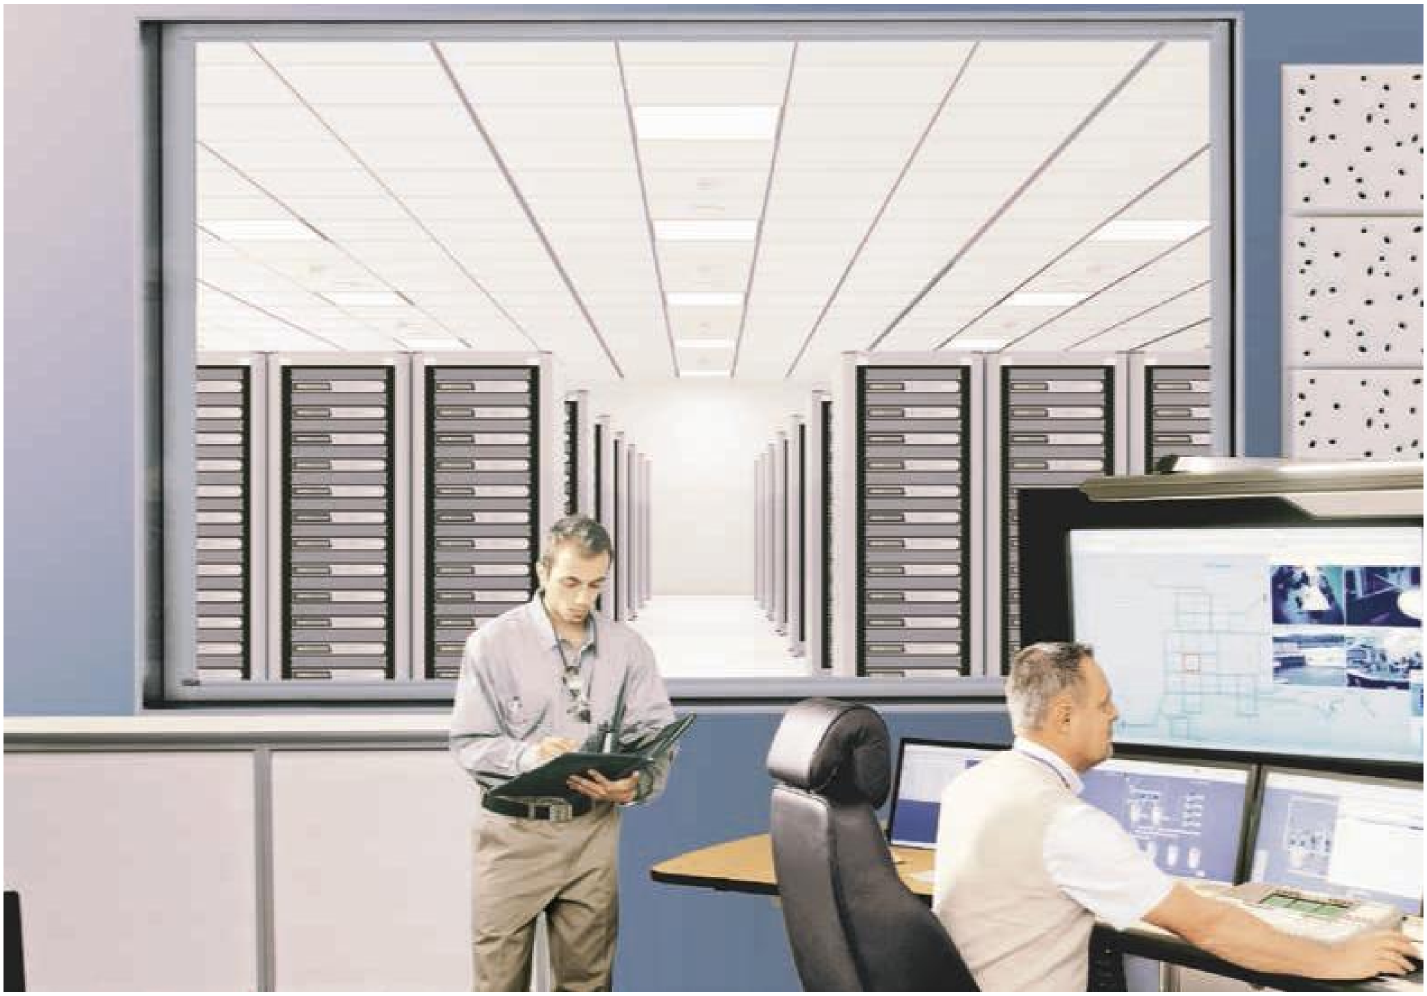

In the crystal ball

Looking ahead at data center design optimization

ZHENYUAN WANG, ALEXANDRE OUDALOV, FRANCISCO CANALES, ERNST SCHOLTZ – "Data center design optimization" is a phrase that rolls easily off the tongue, but actually optimizing the design of a data center is a good deal more difficult than it sounds because the owners, architects and engineers who have a say in the design may all have different priorities. The ability to reconcile the desires of these parties as well as accommodate current and future trends in the industry is a core skill in the art of data center optimization. Energy efficiency is one particularly significant and dynamic trend, and the DC-only, energy self-sufficient feature is one aspect of this trend that is attracting major attention worldwide.

For data centers, a DC-only world would be perfect, especially as DC is native to most renewable energy sources.

Conventional power generators are usually alternating current (AC) based and between the generator and the direct current (DC) electronic loads in, say, a data center, there can be many wasteful AC/DC/AC/DC conversion stages. A DC-only world would be perfect, especially as DC is native to most renewable energy sources. This DC vision has inspired, for example, the DC microgrid-enabled “enernet” ideas of the EMerge Alliance – a not-for-profit, open industry association that is promoting the rapid adoption of safe DC power distribution in commercial buildings through the development of appropriate standards [1]. By reducing the number of AC/DC conversion stages in typical AC-powered electronics, a DC building can be typically five to fifteen percent more efficient. Further, by producing electrical energy locally from biofuel, solar photovoltaic (PV) and wind sources, zero-net-energy buildings (ZEBs) can become a self-sufficient alternative to conventional, externally powered buildings. Data centers are a major application area of this vision.

Title picture
Server performance is just one of many factors to be considered when designing a data center.

Other considerations

For data center optimization, however, there are considerations other than energy efficiency.

Capex and opex

Many factors impact the ultimate cost of a particular architecture – for example, mitigating harmonic currents injected to the AC network may require filtering equipment to be inserted between the utility grid and the data center, thus increasing capex.

Reliability and availability

Conventional AC data center designs are classified into different tiers and each tier has its own reliability and availability requirements (see pages 11–15 of this edition of *ABB Review*). Apart from public image damage, outages can also have a huge financial impact on a data center owner/operator. A self-healing function in the power supply network can improve reliability and this is becoming increasingly popular in data centers. On the other hand, reliability and availability improvement often incurs more cost.

Protection and safety

Appropriate protection and safety measures have to be rigorously implemented.

Scalability

To meet growing requirements, some data center owners plan to incrementally expand server capability and power capacity. The latter may involve backup generator type and number considerations, modular UPS converter/battery configuration, etc.

Footprint

A smaller footprint is advantageous where real estate is costly. However, this necessitates higher power density in server racks, UPS converters, etc. and translates into higher cooling system costs.

1 Finding the minimum TCO for a single performance target (eg, reliability, efficiency, environment impact)

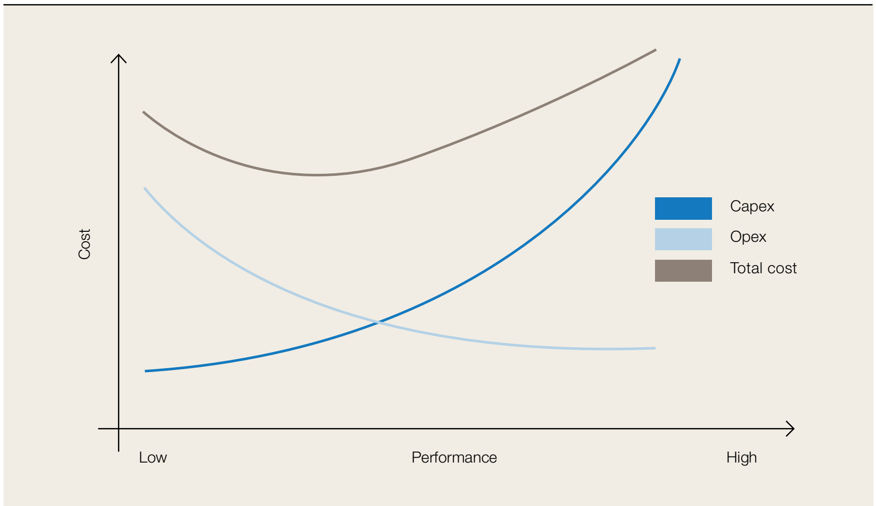

By producing electrical energy locally from biofuel, photovoltaic and wind sources, zero-net-energy buildings can become a self-sufficient alternative to conventional, externally-powered buildings.

Renewables

Renewable energy sources, especially PV and wind, should be easily accommodated. Use of renewables polishes the data center's public image and additional capex can often be recouped in renewable-resource-rich locations. Globally, the "green" data center is a growing trend.

Zero-net-energy (ZNE)

ZNE data centers are usually smaller than average and often have access to renewable energy resources. A reliable utility backup and service agreement as well as energy storage will be needed in most cases.

Cooling

Modern data centers have rack power densities over 10 kW/rack and this will continue to increase. Liquid cooling applies over 20 kW/rack – this will translate into higher initial capex.

Location

The geographical location of the data center is a consideration when there are multiple options. The location impacts the real estate cost, cost of electricity, the initial and operating cost of cooling, etc.

The question is: Given so many interconnected factors affecting the final architecture decision, how can one determine an optimal architecture?

Is an optimal design even possible?

The definition of an "optimal design" is important. For some, it is the design that gives the minimum total cost of ownership (TCO) for a given performance target → 1. For others, it is the one with the least environmental impact, or with the smallest footprint, or the highest efficiency, and so on. For a greenfield developer with a strong sense of environmental responsibility and a strong capital position, "optimal" would most likely mean "greenest"; for a small developer who wants a quick return on investment, the smallest initial capex may be his "optimal" and he may not be interested in costly renewable technologies now.

An optimal data center architecture design is always possible for a given data center developer with clear objectives in mind. But some fundamental assumptions and requirements must be discussed:

- The number of years the data center should function before a major makeover.
- The geographical location of the intended data center, as this determines the cost of real estate and energy, alternative energy supply potential, weather (cooling costs) and factors such as contracts with utilities to provide ancillary services, or with other building owners to provide centralized heating services (this can help to offset expenses).
- Average server rack power density for the planned functional lifetime of the data center.
- The reliability and availability targets or, alternatively, the annual outage penalty that can be tolerated.

2 Data center architecture design optimization process. Arrows show dependencies; engineering analysis items are in the light blue blocks.

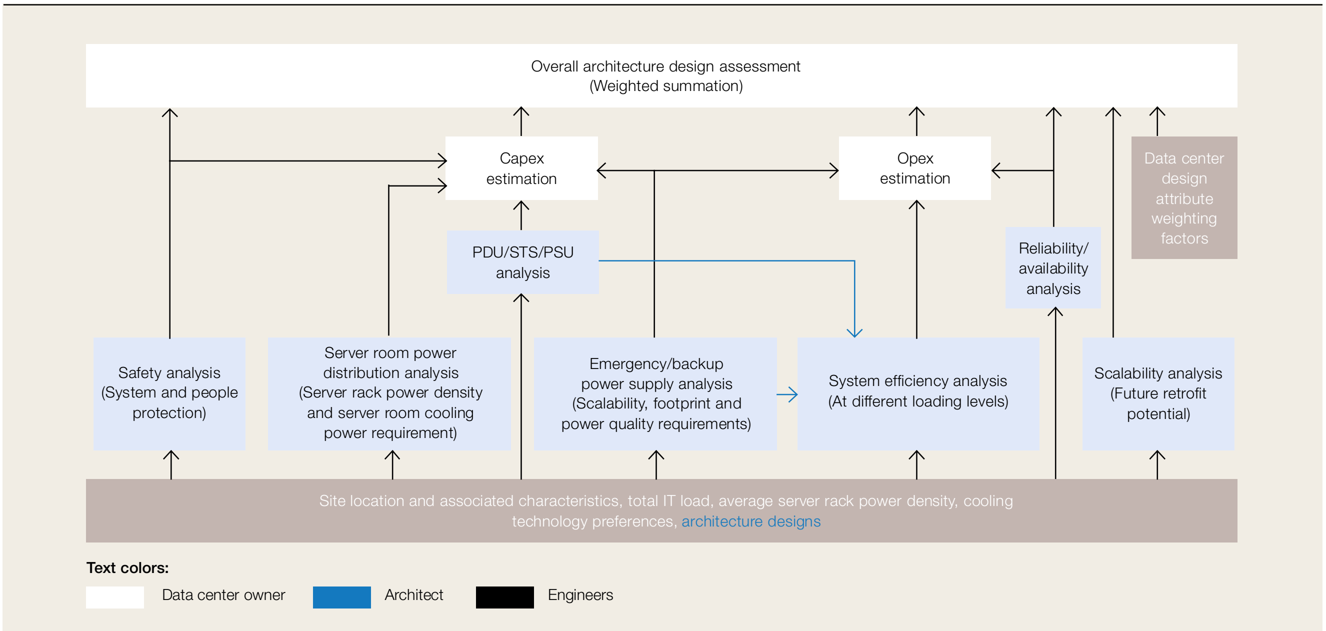

– Site constraints (available space, utility supplies and connection requirements).
– Long-term plan for the site and the data center.

Given the definition of "optimal" and the fundamental assumptions and requirements, multiple data center architecture designs can be developed and analyzed to determine the best candidate. This process, however, requires the involvement of all parties: the owner, the architect and the engineers (for IT, network, electrical, cooling, etc.) ➔ 2.

The data center owner's role

The data center owner (or recipient of the optimized architecture solution) plays a pivotal role in the optimization process as he is well acquainted with many aspects relevant to the data center design. These include but are not limited to:

– The geographical location, with the associated information mentioned above.
– Planned load capacity (in MW) in the short-term and in the future. This impacts the oversizing and reliability considerations of the power equipment.
– Average server rack power density (kW/rack). This will influence the cooling system design and the dimensioning of the power equipment.
– Preferred cooling technologies.

Most importantly, the owner gives major input to the process due to the fact that it is he who will weight the different attributes in the overall assessment.

The architect's role

Based on the owner's inputs, the data center architect will come up with several designs. These designs can be based on DC, conventional AC or a mixture of the two. A design can also incorporate multiple emergency/backup energy sources, protection schemes, etc. In principle, the architecture will roughly determine the cost and performance attributes of a data center – more exact figures will be determined later by rigorous engineering calculations and evaluations.

The engineer's role

Engineering analysis takes center stage after the owner's requirements and the architecture have been clarified. Provision of the power supply alone involves numerous analyses ➔ 2:

– Power distribution unit (PDU), static transfer switch (STS) and power supply unit (PSU) analysis. Depending on the architecture and total IT load, the type, rating, footprint, power density, efficiency, reliability, cost and number required must be determined.
– Server room power distribution analysis. Depending on the server rack power density and the selected cooling technology, this analysis

Scalability is important as, to meet growing requirements, some data center owners plan to incrementally expand server capability and power capacity.

Cost minimization and energy efficiency maximization can be treated as the two most important data center design objectives – reliability is a given requirement and cannot be compromised.

3 A ZNE data center must consume no net energy from the utility grid over a specified time period.

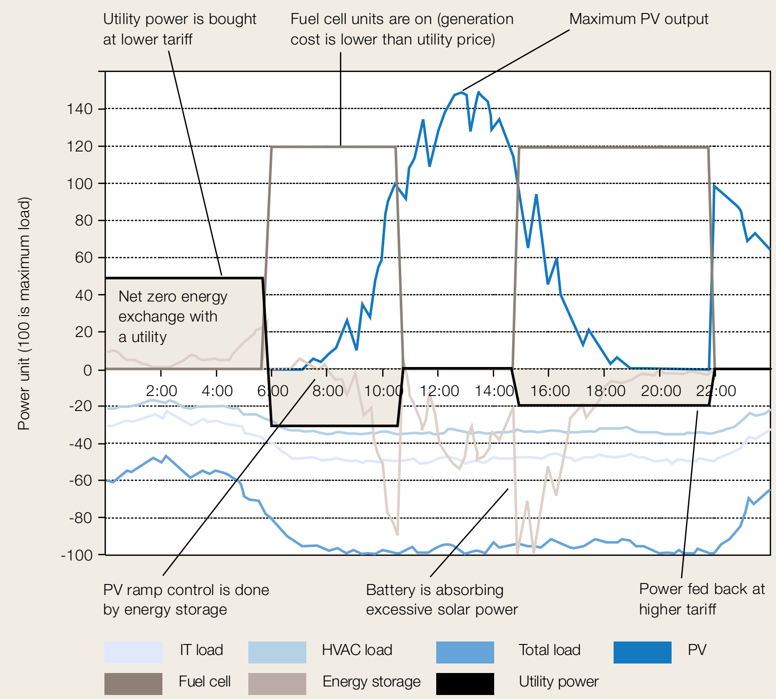

determines the size, length and safety grounding for the power distribution bus and feeder.

– Emergency/backup power supply analysis. Emergency power supply refers to uninterruptable power supply (UPS) systems, which can be based on batteries, ultracapacitors or flywheels; backup power supply refers to diesel generator sets or other types of generation devices that can provide power for hours to days. The architect may have considered the type and redundancy, but this analysis details the ratings, auxiliaries (protection and control), footprint, efficiency, reliability, cost and number of these power supplies. Construction cost differences between alternative technologies are considered in the layout of the emergency/backup power supply rooms (eg, converter/battery rooms).

– Safety analysis determines appropriate protection devices and grounding practices – including the type, rating and number of the protection devices, and the size/length of the grounding conductors. The fault-current limiting function of converters is considered in the protection device dimensioning.

– System efficiency analysis will be done for at least three loading levels: 20, 50 and 100 percent. Efficiency curves of PDU/STS/PSU and UPS converters are the main inputs to this analysis. Server room distribution feeders are quite short and their efficiency can be assumed to be 100 percent, when they are considered.

– The system efficiency analysis result is the major input for the opex estimation, as are data center emergency/backup operation cost and outage revenue loss or penalty. Capex is estimated based on data center IT power supply/distribution equipment and protection equipment costs. Other types of opex and capex

Multiple data center architecture designs can be developed and analyzed to determine the best candidate.

are considered to be the same for all architecture designs and are not considered in the optimization process.
– Reliability/availability analysis is important to ensure the architecture meets certain requirements [2]. A DC data center design may be on a par with a higher-tier AC design in this respect due to savings in power conversion stages.
– Scalability analysis looks at potential power equipment modularity and hot-swapping benefits, or the integrated data center.

Overall architecture design assessment

The overall assessment is usually straightforward – it can be a simple weighted sum calculation. However, as explained earlier, the data center owner bears the ultimate responsibility in assigning ranking weights to different design attributes.

Trends

In general, cost minimization (both capex and opex) and energy efficiency maximization can be treated as the two most important data center design objectives (reliability is a given requirement and cannot be compromised). In addition, data center design optimization has to consider the major industry trends:
– Greener: Designs using renewable or reduced carbon energy resources are of growing interest. A zero net energy (ZNE) data center is a goal.
– Modular: Data centers can be quickly constructed and maintained by using standardized and plug-and-play-capable server racks, power modules, battery packs, cooling equipment and generator modules.
– Cloudier: Economies of scale can be exploited by colocating the IT services of several organizations, especially cloud service providers, in one data center.
– Hotter: With the advent of blade servers, the power density of server racks has increased significantly, posing cooling challenges.

Further, the ZNE building concept is attracting interest due to several drivers [3] that are also relevant to data centers:
– Rapid price drop of local generation technologies (mainly PV panels)
– Controllable loads – heating, ventilation, air conditioning and lighting; the IT load can be shifted, especially in cloud computing data centers
– Progress in energy-efficient building construction technologies (in the case of data centers, more efficient architectures, too)

By definition, a ZNE data center must consume no net energy from the utility grid over a specified time period ➔ 3. Since data centers are characterized by a very high consumption density (100 times that of an average office building) with relatively low daily/seasonal variation, several key factors must be considered in ZNE data center architecture designs:
– Availability of energy supply for local generation
– Type, operation mode and size of local generation
– IT load balancing
– Near-term IT load and local generation forecasting

The design can be AC or DC. However, a DC design will be more efficient, making it easier to achieve ZNE operation.

For fuel-cell and PV-powered ZNE data centers, microgrid operation, ie, self-powered and isolated from the grid, is a real possibility. However, design optimization to accommodate microgrid operation as well as all the other requirements mentioned above is a whole other story.

The owner gives major input to the process due to the fact that it is he who will weight the different attributes in the overall assessment.

Zhenyuan Wang
Ernst Scholtz
ABB Corporate Research
Raleigh-Durham, NC, United States
zhenyuan.wang@us.abb.com
ernst.scholtz@us.abb.com

Alexandre Oudalov
Francisco Canales
ABB Corporate Research
Baden-Dättwil, Switzerland
alexandre.oudalov@ch.abb.com
francisco.canales@ch.abb.com

References

[1] B. T. Patterson, "DC, Come Home," IEEE Power & Energy Magazine, November/December, 2012.
[2] F. Bodi, "DC-grade" reliability for UPS in telecommunications data centers, in 29th International Telecommunications Energy Conference, Rome, Italy, 2007, pp. 595–602.
[3] S. Pless, P. Torcellini, "Net-Zero Energy Buildings: A Classification System Based on Renewable Energy Supply Options," National Renewable Energy Laboratory, Golden, CO, Technical Report TP-550-44586, 2010, Available: http://www.nrel.gov/sustainable_nrel/pdfs/44586.pdf