Building from SCRATCH

Building a data center rather than leasing one is a bold decision. Understand why an organization might do it and what approaches work best.

BY DANNY BRADBURY

If a job is worth doing, then sometimes, it's worth doing yourself. For some companies, that may also include data center construction. Building your own data center from scratch can be an expensive and time consuming task. The most mission-critical ones can take 15-20 months or more to get up and running. But there are advantages, too.

Companies that decide to build their own data centers are going against the market, warns Scott Stein, managing broker for consulting firm Global Data Center Solutions. He says that there is plenty of empty real estate well-suited for data center operations. "Prices have plummeted," he says.

Most experts agree that the threshold for cost effectiveness is around 5Mw in size. Build one below this size, and economies of scale will work to a company's disadvantage. The likelihood that it will be more expensive in the long term than leasing or colocating inside a wholesale or retail colocation grows dramatically.

The decision to lease or build your own data center doesn't always come down to cost-effectiveness, though. Those who choose to build must have very specific and compelling reasons for doing so. Some companies may have a corporate policy that requires them to control all of their own IT assets, including the facility. That requirement may come from customers. Some companies may prefer for accounting reasons to keep real estate as an asset on the balance sheet, rather than on the profit and loss statement.

The Uptime Institute has historically catalogued the price of an enterprise-class mission critical data center up to $22,000 per Kw.

Another reason may be that a company's physical requirements are so exacting that it is difficult to find a physical property to match them. An organization that deals with sensitive information must have a 48-meter setback from a main highway, with a Sensitive Compartmented Information Facility (SKIF). That will help rule out the majority of generic colocation or leased data center facilities, requiring the company to build its own, despite requirements for a relatively low computing capacity.

Some companies may try to find existing facilities to satisfy such requirements, but the added benefit can be mitigated by potential risks. An underground bunker may satisfy security goals, for example, but it is difficult to protect against water penetration in such facilities.

Typically, customers building data centers will define costs on a per-kilowatt basis. Data centers can range wildly in the cost per kilowatt. The Uptime Institute has historically catalogued the price of an enterprise-class mission critical data center up to $22,000 per Kw. There's a delta between a purpose-built enterprise facility and a wholesale or retail facility, and the key is economies of scale.

Many Fortune 500 companies whose core revenues don't come from IT services won't be purchasing high-capacity generators and high-capacity switch equipment in the same volumes as a Google or a Microsoft, which can get down to lower rates with data centers that are 50 Mw or more in size. "You're buying 1-4 generators. They're buying hundreds at a time," says Stein.

Building in Stages

Any data center construction project will include a number of stages. The project team will begin by identifying broad requirements of the project, including the location of the facility, its size and capacity, and the plan for expansion over the coming years.

Location is one of the most important factors in building a new data center. Often, companies will site them in close proximity to their headquarters and field offices, but this is often not the most efficient location from an operations perspective, says Bryan Loewen, senior managing director for real estate consulting firm Newmark Grubb Knight Frank.

Some of the biggest markets for data centers are in New York and New Jersey (far from the cheapest locations for energy). "A trend is that the financial industry drove these markets initially," he says, adding London, Hong Kong and Sydney internationally to that mix.

Other markets were driven by other industries. Santa Clara, another prime spot for computing facilities, has its own power company (Silicon Valley Power), with reasonable rates and high reliability rankings for power. The region is also culturally sympathetic to technology developments, given the number of tech firms there. But it also sits on a fault line.

> "I think the in-row stuff has largely been driven by vendors, and then you start to buy gear from them and it locks you in,"
>
> — Steve Miano, Managing Principal, PlanNet

The precise location of the building will also take some other considerations. Organizations will need access to good power and fiber links (which can limit the use of more remote, colder regions that can take advantage of air cooling). Often, a facility will need solid access to water, too. "Those big ticket items are important," says Loewen. "Water is sometimes mitigated with airside cooling, but power and fiber are crucial."

A data center will also be graded at various tiers of reliability, according to the ANSI/TIA-942 standard. The first tier (effectively a server room) involves a single electrical path and single points of failure all the way up the line. There will be redundancy in the second tier, and concurrent maintenance capability (hot swapping) at the third tier. This includes sufficient capacity and distribution to carry data and power loads on one path, while performing planned maintenance on another. A fourth-tier data center is fault tolerant to 99.995 percent availability, and includes multiple power and cooling paths. It may even include connections to multiple different power substations for extra resilience. The Uptime Institute maintains a similar tier ranking.

That said, Michael Fluegeman, a principal with data center design and construction firm PlanNet, argues that highly fault-tolerant tiers may not be as significant today. "Data center IT strategists are doing a much better job in spreading out risk among multiple facilities," he says. "Therefore we don't have to make a particular facility fail-proof and robust. The risk is getting spread between facilities via geography."

Security is another consideration when choosing the location of the facility, and the underlying physical design, says Lawrence Smith, founder of data center design and construction consultancy ABR Consulting Group. "I have done tier four buildings where there are bollards around the building, and secure fencing," he says. "The receiving areas are blocked off from view. There is tremendous security between different environments inside the building."

Having Enough Elbow Room

Data center capacity is one of the other major considerations when designing a facility, says Smith. "The electrical design engineers will want to know what you're putting in, and what you're taking from your existing facilities, and how much it will grow in the next 5 years,'" he says.

The idea is to plan that growth and build accordingly. This can make a self-build project economically more viable, because building for future capacity rather than limiting a development to current capacity can bring economies of scale.

Capacity planning has changed over the years because of increased power density. It is possible to fit far more computing capability into a cabinet than 10 years ago, but that capability brings additional power requirements. "The industry has advanced considerably in the last 30 years," says Loewen, adding that even in the last decade, things have altered. "The power densities in data centers were dramatically less than they were today."

He estimates that data centers were getting 10-50 square watts per foot 10 years ago, compared to 100-150 watts per square foot now. "There is a lot of product that's being placed into a small footprint."

Even so, things can vary dramatically depending on an individual company's requirements and capabilities. Today, companies can vary in their power density from 3-4Kw per cabinet, through to 25-30Kw per cabinet, say experts.

This density will affect the design of the data center. An aggressive power density may warrant two zones: a high-density zone with a different cooling technology to a lower density area.

Cooling techniques are changing, too, says Steve Miano, managing principal of PlanNet. He is moving away from in-row cooling. "I think the in-row stuff has largely been driven by vendors, and then you start to buy gear from them and it locks you in," he says. "I believe that you want to keep the mechanics and the batteries and the material away from the IT equipment."

Instead, outside air cooling is becoming a design consideration. Companies can use either direct or indirect outside air. The first involves bringing in outside air, filtering it, and circulating it. This carries potential cleanliness issues, because filtering it may not always remove all of the pollutants. "A lot of people are realizing that this isn't as much of a problem. Circuit boards might get dirty, but they'll change them every few years anyway," says Miano.

The alternative is indirect outside air, in which a heat exchanger is cooled by the outside air. The air never enters the data center. This can be less efficient than conventional cooling, say experts, but can still minimize the mechanical cooling required. These air cooling technologies—especially the indirect ones—explain why companies are building facilities close to the edge of the Arctic Circle.

ASHRAE's Efforts

However, facilities projects don't need to move to Finland to take advantage of free cooling. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) is helping here. In 2011, it updated its guidelines for allowable temperature variations in IT equipment. First established in 2004, ASHRAE's environmental classes for data center temperature and humidity were formalized in its TC 9.9 standard, which vendors agreed would not invalidate warranties.

In 2008, it widened the recommended temperature ranges for the two classes it had published that were specifically for IT equipment used in data centers. In 2011, it recommended two new classes (A3 and A4) for IT equipment in data centers that once again had expanded envelopes. This was designed for different classes and priorities of equipment.

In 2012, the Green Grid (an organization dedicated to green computing) published a selection of global "free cooling maps," which showed the free cooling capabilities offered in different parts of the world. This will be a boon to any company that wants to build a free cooling data center from scratch, and isn't bound to its headquarters' city of residence.

Clearly Define Specifications

Along with compute density, networking and storage are other major capacity considerations in a data center. Requirements here, too, are accelerating at a high rate. As MIPS, network throughput and storage capacity requirements increase, it becomes more important than ever to get IT and construction teams working together. Unfortunately, this doesn't always happen.

ABR's Smith recalls projects where IT teams have failed to contribute appropriate specifications until too late, creating significant problems for the project, and leaving the facilities team to work on their own, making unwarranted assumptions. "If they don't get the material from the IT people, they will move ahead with their project," he says. The results can be disastrous.

"Make sure the specs doc is extremely complete," he continues. Power and cooling systems, electrical and cabling infrastructures should all be outlined in high detail. "You will pay money to get documentation like this put together so that you can bid on this properly."

A well-managed project with lots of communication from the beginning can result in detailed operations information when the facility goes live. Data center infrastructure management (DCIM) systems have enabled data center designers to gather information about current and projected IT hardware assets and populate the data center model with them.

Building information modeling (BIM) enables designers to map out the data center facility in 3D using CAD techniques, which can help to design airflow, network cabling, and electrical distribution. Bringing the two together helps designers to understand the effect that IT equipment will have on the facility's overall operation. Ultimately, information about the IT equipment can then be brought over into asset management systems when the time is right.

In a well-managed project, there are several levels of commissioning that run hand-in-hand with design and construction. Level one involves the test of the design and specifications of the equipment provided for a project, ensuring that it is designed properly for the proposed data center environment. The level two test involves checking equipment onsite when it arrives. Installation of the equipment is verified, and then systems such as backup generators, UPS units and other components are tested individually. Finally, an integrated system test runs through all operational scenarios.

Ultimately, the decision to build a data center rather than lease one will depend on a number of factors, from the economical, to the technical. Company culture may even play a part. But those who do choose a greenfield project will find themselves faced with a wide variety of choices that could make their projects more efficient and productive in the long run.

Danny Bradbury has 20 years of experience as a technology journalist. He writes regularly about enterprise technology issues including data center management, security, software development and networking. He blogs at www.wordherder.net.