Lost Capacity - The Elephant in the Room

How your short term drivers erode your long term goals

Executive Summary

In the world of economics, short-termism, the policy of concentrating on short-term profit at the expense of long term stability, is often cited as a key factor in the failure of businesses and generally thought of as bad practice. In the data center industry, are current practices driving operators unwittingly into short-termism? The majority of data center operators are making decisions relating to short term gains without the tools necessary to understand the long term effects of those decisions.

Without any visibility into the long term consequences of IT deployment choices and energy saving programs, proper cost/benefit analysis of the proposed changes cannot be performed. The cost of not achieving the full potential of a facility can be massive, increasing the real price paid to provide each kW of processing power by 100%, which often dwarves the short term benefits of many decisions. Without an understanding of the real long term costs, short-term drivers dominate the decision landscape resulting in large swathes of data center capacity lost to productive use, driving up the real price paid for that which is active.

This paper addresses the causes of lost capacity, analyses the real terms cost of under-utilized data center space with a worked example and proposes a way forward for data center operators to start reducing the effects. It is a consensus of opinion between a number of data center professionals in different areas of the industry; owner/operators, consultants and vendors, all of which agree that the real issue facing the modern data centers, the elephant in the room, is lost capacity.

The design of any data center is based on a number of assumptions about the nature of the equipment that will be deployed during the lifecycle. These assumptions determine the power and cooling distribution design, port densities and available U space. They will be based upon the best available data from IT about the future technology plans, but even so these can be proved wrong, even on day 1 of the facility's life. Breaking these assumptions and deviating from the plan envisaged from the designer has a direct impact on how the capacity of the data center is used. For example, a facility designed for a 70/30 mix of blades to rack mount kit will run out of network capacity very quickly if the equipment deployed is 90% rack mount servers.

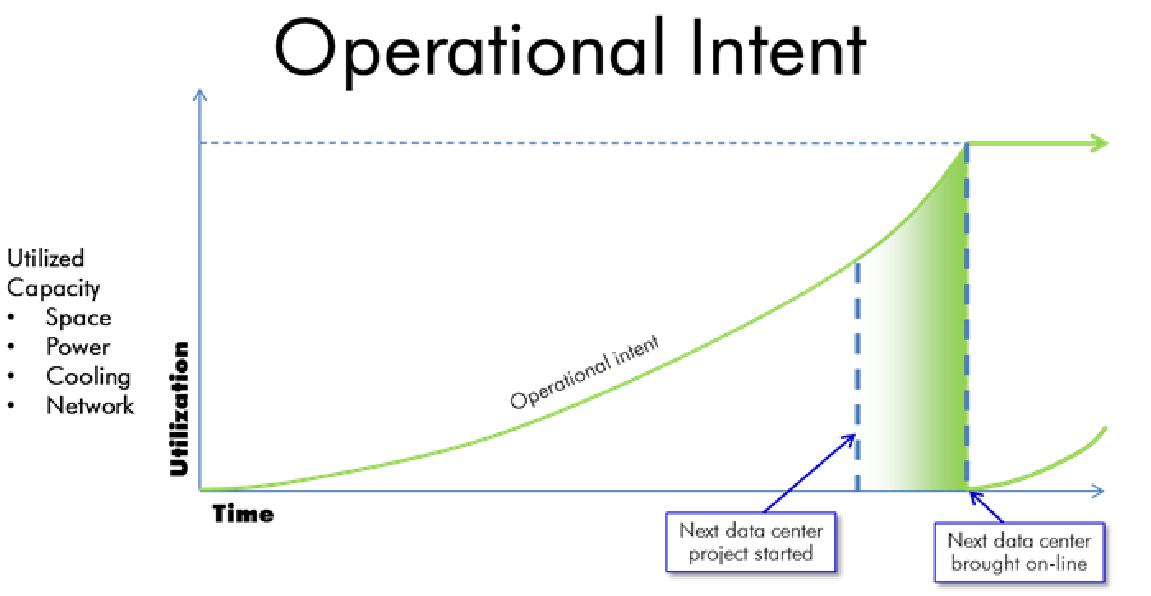

Even though in operations most facilities deviate from the initial plan fairly soon after day 1, it is worth remembering that a data center with a 10-15 year lifespan will experience around 3 complete refreshes. New technology can be very different from that designed for, think of the problems encountered when blades first appeared. This means it is *inevitable* that the original design assumptions will be broken.

Current processes allow little visibility in the short term effects and no visibility into the medium and long term effects of deviating from the original plan. This leaves both facilities and IT making decisions with no understanding of the future impact of those decisions. In the very best scenario, the changes will have zero long term impact, but this is rare. More commonly, decisions taken today will have some negative impact on the total usable capacity available in the facility. Indeed, according to Gartner [1] the majority of data centers fail to reach their intended capacity. The real cost of this lost capacity is huge. Consider a data center that was intended to house 1MW worth of equipment for 15 years is considered to be full at 50% load after 6 years. The real-terms cost of the data center will have increased significantly as you are getting 500kW

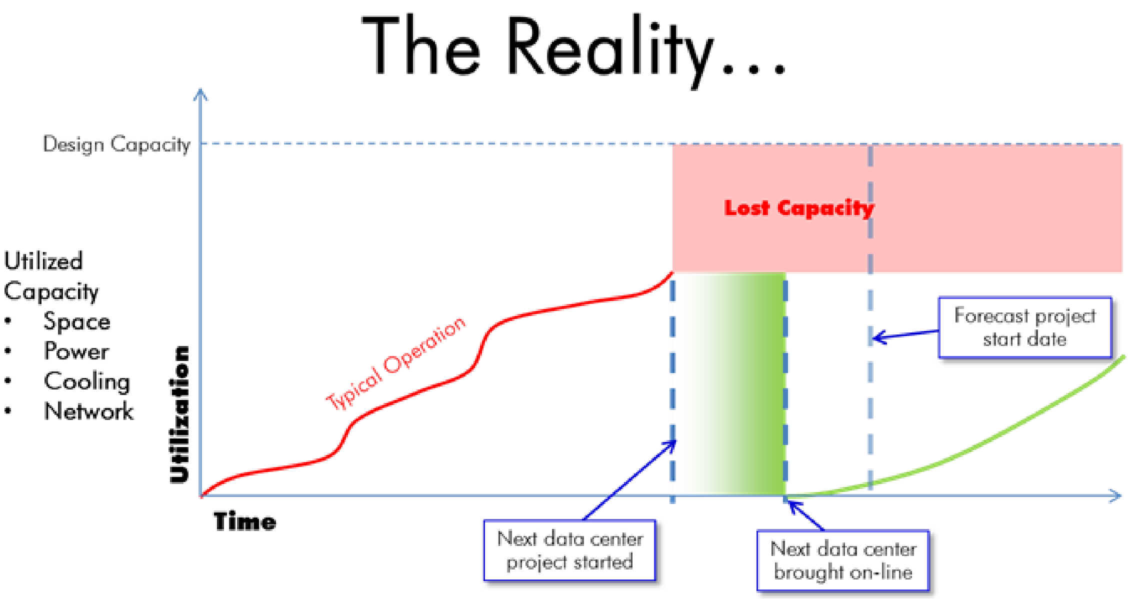

for the price of 1000kW. You will need to build another data center, at great capital expenditure, 7 years early and 400kW of capacity you will effectively pay for twice. There will be a lag between realising you need more space and that space coming on line, which could result in delays to the deployment of critical new systems that could cost the company a commercial advantage. Unfortunately, the industry is yet to really acknowledge these issues, preferring to focus on reducing operational expenditure.

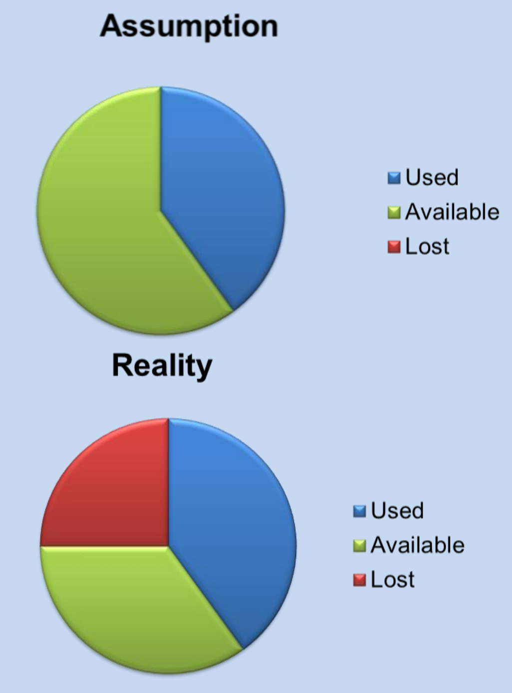

Capacity planning and forecasting tools work on the assumption that all unused capacity is available, this is never the case

What do we mean by capacity?

When commissioning a new data center build, a figure for the total amount of IT equipment that the new facility will hold is defined. This DESIGN CAPACITY will be specified in either kW or MW. The facility is then designed such that the AVAILABLE CAPACITY = DESIGN CAPACITY + REDUNDANT CAPACITY (if it did not it would be a pretty poor design!). The cooling system will then be sized to remove that amount of heat at the desired resilience. A desired power density will calculated from the DESIGN CAPACITY and the size of the white space and the cabinet layout designed to deliver this. The power system and network will be designed to provide the correct power and desired connectivity to all of the cabinets.

So, the AVAILABLE CAPACITY in the facility is made up of four resource components: *SPACE, POWER, NETWORK* & *COOLING*. As IT equipment is deployed, it uses up these resources and the AVIALBLE CAPACITY reduces as the INSTALLED LOAD increases. If a lack of any of these components prevents the deployment of IT equipment then the AVAILABLE CAPACITY effectively becomes 0 or Resilience is conceded either knowingly or unknowingly. If this happens before the INSTALLED LOAD reaches the DESIGN CAPACITY then there is lost capacity or a reduction in the Tier rating in the data center

How does capacity get lost?

The four resource components are interlinked and each affects all the others, but we can try to look at each one individually to understand how they run out and prevent IT deployments before the design capacity of the data center is reached.

Space

The total U space available in a facility does not tell the whole story about how much equipment can be deployed. The other important factor is the amount of contiguous U space. If all of the U space in the data center is in blocks no bigger than a single U, then any deployment requiring more than a single U cannot go ahead! Obviously this is an extreme example, but as more processes move towards cluster based computing requiring a large number of rack mount servers in close proximity, the demands for large blocks of contiguous U space are increasing.

Current rack state shows slot 28 next available for deployment.

Server in slot 10 scheduled for decommission before deployment takes place.

Result is the fragmentation of available space.

So how does contiguous U space get lost? Just like the hard drive in a PC gets fragmented as programs and files are added and removed, so does data center space as IT equipment is deployed and decommissioned.

Deployment plans are generally made using the current state of each rack, so any future decommissions are not taken into account when choosing a location to place a new piece of IT. For example a new 4U server needs to be installed next week, the current state of the cabinet suggests that slot 28 is the first available U but this is not actually the case. One of the 4U servers at the bottom of the cabinet is planned to be decommissioned tomorrow, so by the time of the installation the first available U slot will be slot 10. The installer does not know this, and the plan is finalised to install in slot 28. The result is that the remaining available U is now fragmented into two sections, rather than one contiguous block.

Power

At design, as a general rule each PDU will be assigned to a number of cabinets and then the total power from that distribution unit will be divided evenly to each cabinet. Some designs may specify high density areas or network areas with higher power availability, but in the designated server cabinets an even distribution of power is going to be the case. However, the relationship between the amount of space required and

What you think you have vs What you really have

This is a simple example of how a lack of one capacity component can prevent the others from being used. Figure (a) shows 5 cabinets with differing loads.

Into these cabinets there is demand to deploy a new HPC cluster which requires 11U of rack space and 1.4kW of power. Because of the way the cluster components communicate they have to be in a contiguous block of space.

This constraint means that none of the cabinets can accommodate the new equipment and the remaining capacity is lost.

Figure (a)

Cab	Available Space	Largest Block (U)	Available Power
1	13	8	1.2
2	15	8	1.7
3	26	26	0
4	17	8	2
5	17	10	1.7

power draw is different for different technologies. For example, rack mounted storage systems use large numbers of low powered disk shelves and can fill an entire cabinet without reaching the power limit whereas servers and blades can use up an entire cabinets power in less than half the available U space.

This mismatch between different technology leads to a number of issues. Cabinets containing lots of low power equipment will run out of space before they run out of power, leaving the remaining power unavailable for use. Conversely, cabinets with more dense processing equipment run out of power before the space in the cabinet is used up (you can start to see how the 4 components are interlinked here).

The facility manager will try to avoid this unbalance as much as possible, often at the expense of usable capacity. A 90kw PDU will have 30kw available on each of the three phases, but if the cabinets on phase 1 are filled but only with 10kW of power (say a large number of disk shelves), then to keep the design capacity the 20kw remaining from phase 1 gets migrated to the other phases or the 20kW of capacity is lost. This leaves 20kw of capacity unusable without increasing the risk of equipment failure.

Network

The port density in the data center day 1 is decided during the design based upon a best guess of the technology mix that is going to be deployed. Just as the ratio of power consumed to space used varies for different types of equipment, so does the ratio of ports to power. Rack mount server equipment can use as many as 10 copper ports for around 500W of processing power whereas a blade chassis can provide up to 5kW of processing power with the same number of ports. If a data center is provisioned with a port density of 48 network ports and 5kW of power per cabinet then deploying a blade chassis would utilise all of the power but leave 80% of the ports unused whereas rack mount servers could use all of the ports but only 50% of the power.

Cooling

The recurring theme in the previous sections has been the fact that in design an even spread of demand is assumed but in practice the non-uniformity of IT equipment breaks this assumption. Cooling is no different. Different pieces of IT equipment require different amounts of air, in different directions, to cool the chips inside. This variation happens not just between different types of technology, but between manufacturers and even between different generations of the same model!

Again design is for even spread of airflow demand, but this does not happen in operation. In fact, it is often not even delivered! Other changes to the DC have effects on the cooling system as well (rerouting over power cabling, extra networking capacity

Different technology is cooled in different ways

etc), changing the cooling profile of the facility. However, unlike the other components, air, the delivery method for cooling is invisible, adding an extra level of complexity to the problem.

Imagine a data center where all of the power and data cables were invisible. For any given piece of equipment in the facility you could not see how it was getting power and networking, or whether it was resilient, you would just know that it was powered on and connected to the LAN. Now imagine trying to deploy a server into this data center. You would put the server into the rack, press the power button and wait to see whether it got power and network connections and again, you would have no idea from where they were coming or whether it was resilient. This is how the cooling system works currently.

The problem is often compounded by neither Facilities nor IT departments taking responsibility for managing the airflow within the data center. The variation in cooling profile between different pieces of IT equipment means that the ability of the cooling system to deliver the correct amount of air to the inlets of all the IT devices in the facility is totally dependent on the choice of equipment and where it is located. In many organisations Facilities Management will have responsibility for ensuring that the cooling system can remove the total amount of heat, but have no control over IT deployment. The IT teams will control IT deployment, but the cooling system will be outside of their remit. This allows airflow management to fall through the gap between the two departments, with no-one taking responsibility, creating hotspots and ultimately losing capacity.

It's a bit more complicated than that…

Unfortunately the four components cannot really be considered individually. Each affects the other which makes the problem considerably more complicated. The availability of one component can drive deployment which has a direct impact on others. For example, the availability of network ports or power may push a deployment into a certain area which is detrimental to the performance of the cooling system. It is this delicate balancing act between the different parts of the data center eco-system that has to be performed to get the most out of the investment made in the facility.

The Cost of Lost Capacity

The real cost of lost capacity is a reduction in the Return on Investment (ROI) for the whole data center project and an increase in the cost per kW of processing power. Using the True TCO Calculator provided by the Uptime Institute [2], we were able to make an estimate of the Total Cost of Ownership (TCO) of a 1.3MW data center over a 15 year lifespan, based on a real case study, as an example.

The total cost of the data center is made up of four distinct parts:

a) The **Capital Expenditure** required for the construction and fit out of the facility.

b) The **Fixed Operational Expenditure** required to run the facility (this includes elements like staffing costs and taxes)

c) The **Load Dependant Operational Expenditure** which is mostly made up of the cost of the utility bills required to power and cool the installed load.

d) The cost of buying the IT equipment itself.

Input Data

Total Area	10,000 ft^2
Number of Cabinets	400
kW per Cabinet	3

The processing requirements of the business will not change if the data center does not reach capacity, so the cost of the IT equipment can be ignored from these calculations (the servers will still be bought; they will just end up in another facility!).

The Capital Expenditure laid out to build the facility is a fixed cost paid up front, this will not change over the life of the data center. In this case the capital expenditure is $50M, so annualised over the 15 year lifespan of the facility this cost amounts to $6.3M a year*.

The Fixed Operational Costs take into account all of the running costs of the facility that are not directly related to running the IT load. This includes the cost of employing various members of staff (e.g. site management and security), general maintenance and cleaning of the building and property taxes. Unlike Capital Expenditure, these costs may vary over time although they are generally out of the control of the data center owner. In this case, the fixed operational costs amount to $3.1M a year.

*This value is taken from the Uptime Institute True TCO calculator and is calculated using the static annuity method. For further explanation please see [2].

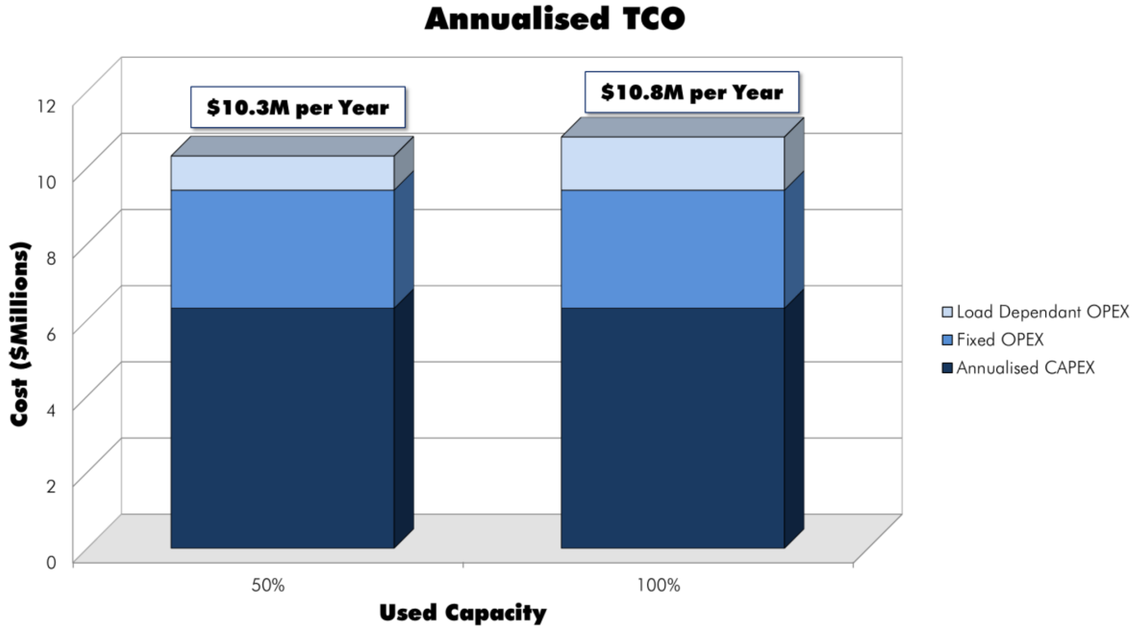

The Load Dependant Operational Expenditure is the cost of powering and cooling the IT systems installed in the data center. These costs are the only real variable component as they are dependent on the amount of load installed in the facility. They are also somewhat under the control of the data center owner as they can be reduced by energy efficiency initiatives. In this case for 50% load the cost is $0.9M per year and at 100% load the cost is $1.4M per year.

The graph above compares the total yearly cost of our example data center at both 50% load and 100% load. Because the only component of the total cost that varies with the installed load is also the smallest, the total figures are almost exactly the same.

This is fairly alarming, however, the real situation is much worse than this. As mentioned earlier, the IT demand from the business remains the same so there is still a requirement for another 0.65MW of processing power (to give the full 1.3MW of the design). This means another identical data center will be built to satisfy the business need, but the same problems arise and the second facility is also considered full at 50%. So, the total cost of 1.3MW of processing power is actually *double* the original estimate.

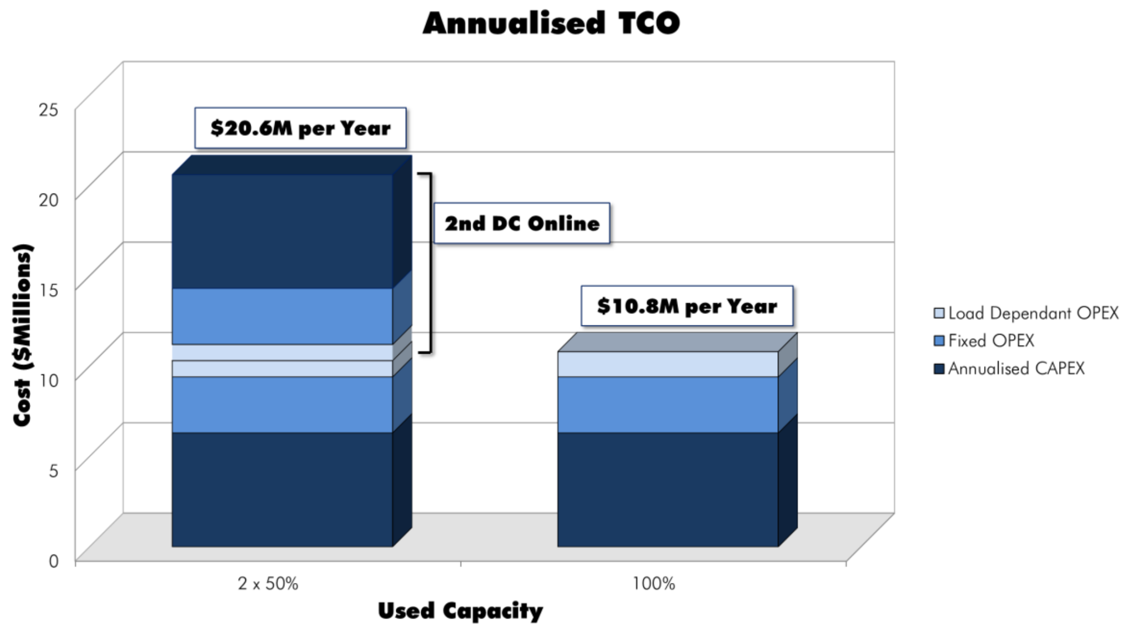

This increases the $/kW processing up even further, to twice that of a single, fully utilized facility. Because each data center is not fully utilized you are paying the extra cost for the underutilization twice, which is significantly higher than the cost of a single, fully utilized facility:

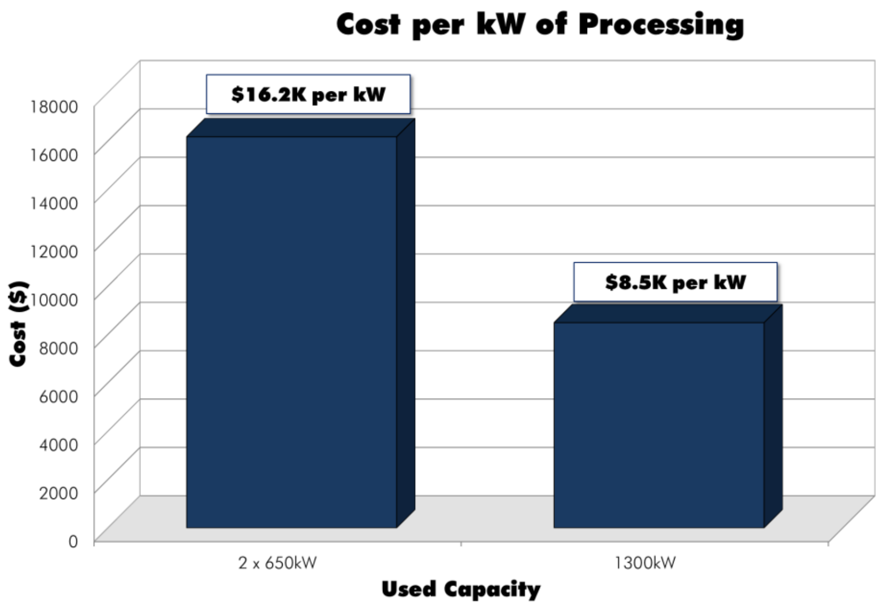

How to avoid lost Capacity

In any data center that needs to keep up to date with changes in technology, avoiding lost capacity completely is impossible, but it can be minimised.

To avoid lost capacity, every time there is a change that deviates from the design assumptions **the question should be asked: will this change sacrifice capacity in the long term?**

To answer this question information needs to be gathered from all stakeholders in the data center space. More and more data center management systems are converging information between IT and facilities. In the past, simulating the future in an operational environment was impractical because of the amount of data required and the immediacy of the decisions. However, **modern simulation techniques can now take the collated data from DCIM tools and show the effects of future plans within the timescales required for operations.** So combining modern data center management systems and data center simulation software is ideal for answering the difficult question posed above.

Once the question is asked, and an answer provided, then lost capacity will fall to a minimum. It may be that some deployments that result in losing capacity cannot be avoided, but at least the consequences will be understood. This means proper cost benefit analysis can be performed and signed off and the resulting reduction in available capacity built in to the forecasting for the next data center build.

Conclusion

In the ideal data center, all four components of capacity would reach 100% utilization at the same time. The way to achieve this is to design the facility with a specific plan in mind and never deviate from this. However, in today's mission critical facility this is rarely an option. The requirement to have the agility to keep pace with improvements in technology and demands from the business to remain competitive is paramount and means that plans constantly evolve. To ensure that this evolution does not sacrifice capacity, and hence greatly reduce the return on the investment in the facility, they need to be checked before they are implemented. Simulation techniques are the only way for data center operators to gain this foresight into the future state of their facility.

Authors:

David King MEng - Consultant Engineer, Future Facilities

Dr Matt Warner - Product Manager, Future Facilities

Contributors:

Gary Cane – Associate Director, CBRE

Roy Freedman – International Corporate Real Estate Services Manager, BNP Paribas

John Lane – Partner, Cundall LLP

Craig Felton – Senior Infrastructure Consultant, PTS Consulting

Ken Baker – Senior Consultant, HP Critical Facilities Services EMEA

References

[1] Cappuccio, D.J. "Data Center Efficiency and Capacity: A Metric to Calculate Both", Gartner, September 2009

[2] Koomey, J et al. "A Simple Model for Determining True Total Cost of Ownership for Data Centers", Uptime Institute, Santa Fe, NM, USA, March 2008