



comtec power
datacentre & server room solutions

Common mistakes in datacentre planning

White Paper

Last edited: March 2013

Document Control

Document Information

Title	Common mistakes in datacentre planning
Author	Nick Ewing

This document is based upon the White Paper “The top 9 mistakes in data center planning”. Originally published by Schneider Electric in 2012*.

Document History

Version	1.1
Date	11 th March 2013

Executive Summary

Whether you call it your datacentre, server room or network closet; why do so many new builds or expansions fail to deliver on the promised benefits? This publication reveals some of the key mistakes organisations make when designing and building new datacentres and examines the benefits of taking a TCO approach to datacentre planning.

Some common mistakes:

- > Failure to take TCO into consideration
- > Underestimating the costs
- > Setting the wrong design and performance criteria
- > Choosing a site before design criteria are in place
- > Space planning before design criteria are in place
- > Designing into a dead-end
- > Not understanding PUE
- > Overcomplicating the design

Introduction

Many businesses are operating a datacentre outside of safe capacity thresholds and with little or no room to expand. The IDC estimates that the average datacentre is 9 years old, yet Gartner claims that any site over 7 years old is obsolete.

Outdated and overcapacity datacentres are a barrier to growth, perform inefficiently and are a potential risk to business productivity and continuity. Sometimes, the only solution is to build a new datacentre. While speed-to-market is critical, organisations that fail to assess their future needs properly will create dead-end datacentres that fail to deliver the expected performance, availability, resilience and efficiency.

So, how can you avoid making some of the more common mistakes when building your new datacentre? The key lies in the approach you take to scoping and designing your datacentre. All too often, businesses base their plans on simple watts-per-square-foot, cost-per-square-foot or tier level. These criteria may be misaligned with your overall objectives or risk profile.

“Datacentre owners have so many problems right now. Their assets are mission critical, but they’re out of control. Power consumption is costing them a fortune. They can’t cool what they have got and cut the risk of a catastrophic outage and if they make an investment, by the time it’s built, it’s already out of date.”

Stanford Group

Poor planning leads to poor resource usage, especially capital expenditure, and can result in an increase in on-going, operational costs. Many organizations get overwhelmed, focusing on speeds and feeds, green initiatives and concurrent maintainability. Whilst these criteria are important, the details occasionally overshadow the big picture during the planning stages. A lot of organisations miss out on the business opportunities presented by adopting a holistic approach to datacentre expansion.

There are a plethora of datacentre consultants and engineers available to help your decision making, but the sheer volume of opinion can make an accurate assessment difficult. Many businesses will lack the internal technical expertise to drive a successful expansion and will need to rely upon external support. Doing so can sometimes result in information overload and lead to flawed decision making.

Failure to take total cost of ownership (TCO) into account

Focusing solely on capital expenditure (CapEx) is an easy mistake to make as the costs to build or expand your datacentre can be significant. Capital cost modelling is, of course, essential; but if you have not included the on-going costs to maintain and operate (OpEx) your business critical infrastructure then you only have half the picture.

There are two critical components required to build a datacentre OpEx cost model:

- > Maintenance Costs
- > Operating Costs

The maintenance costs are the costs associated with the proper maintenance of all critical facility support infrastructure. They include, but are not limited to, OEM equipment maintenance contracts, datacentre cleaning expenses and subcontractor costs for remedial repairs and upgrades.

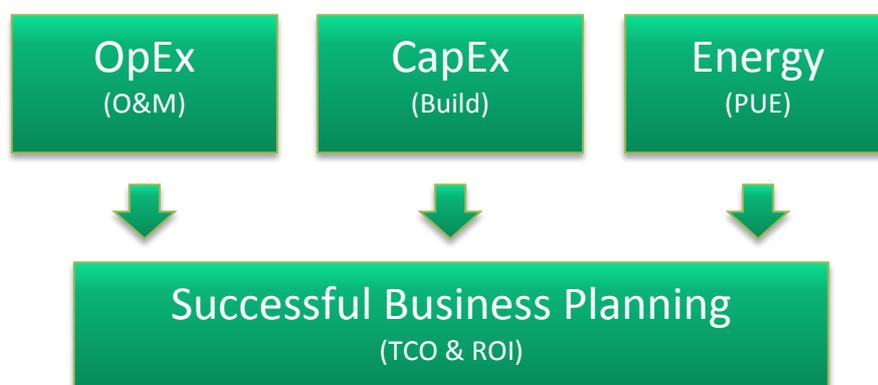
The operating costs are the costs associated with daily operation and on-site personnel. They include, but are not limited to, staffing levels, personnel training & safety programs, the creation of site-specific operations documentation, capacity management and QA/QC policies and procedures.

If you have not calculated a 5 year operations and maintenance (O&M) budget you cannot build an effective return on investment (ROI) model that will support smart business decision making.

If you are planning to build or expand a business-critical datacentre, your best approach is to focus on three basic TCO parameters:

- > Capital Expenditure
- > Operating Costs
- > Energy Costs

Leave any component out, and you run the risk of creating a model that does not correctly align your organisation's risk profile and business expenditure profile. If you are making a decision about whether to build your own or to utilise a co-location or hosted solution, the impact of failure to take TCO into account could be even greater.



Underestimating the costs

Another common mistake is the estimate itself. Financial requests made to boards of directors for capital to expand or build a datacentre are often too low.

The decision making process usually goes something like this:

- > The capital request is made and tentatively approved.
- > Financial resources are allocated to investigate, capture information and create a true budget.
- > Time is spent driving the budget process.
- > The findings reveal that the initial budget request was too low.
- > The project is delayed, careers are impacted and the ability to deliver a quality service to internal and external clients is impacted.

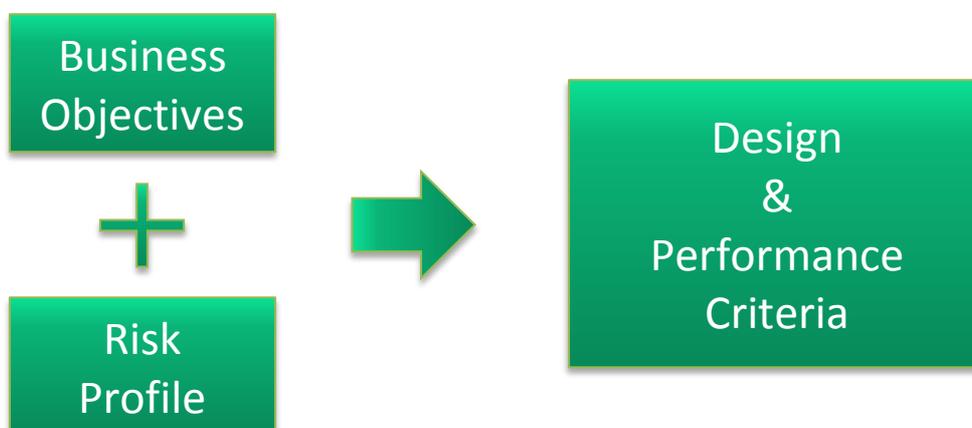
This takes you back to our previous common mistake – failure to take TCO into account when building a financial model. In reality, mistakes associated with underestimation can be easily avoided, but are likely to result in major issues if you fall into the next trap.

Setting the wrong design and performance criteria

There are two major missteps that can send your organisation into an overspend spiral. First of all, everybody wants high levels of redundancy and resilience, but this is not always supported by the business imperative. Secondly, most visions of kilowatt per rack position are not supported by IT requirements.

Too often, the “we must have so many kilowatts per rack” approach is not justified. Don’t fall into the trap of over-building; it is a waste of capital and higher tier facilities also lead to higher O&M and energy costs.

Establishing the right design and performance characteristics before building your CapEx and OpEx models is the foundation for a sound business model that will deliver ROI.



Selecting a site before design criteria are in place

Companies often start searching for the perfect space in which to build before having their design and performance criteria in place. Without this vital information, it doesn't make sense to spend time visiting or reviewing multiple sites.

This “cart before the horse” scenario is common where the imperative is to build or renovate a facility in the immediate vicinity (often the same building as other business functions). Businesses that are more experienced in datacentre planning will usually consider power, cost, connectivity and even geographical risk factors before selecting a site.

The problem with selecting a site prematurely, or based on a narrow set of criteria, is that the site may not meet your design criteria. For example: having a datacentre located on the second floor of your office building may be convenient, but business-critical datacentres are often subject to a long list of criteria that cannot be met by a multi-tenant space. Occupying the second floor of your office building may result in significantly higher build costs, will limit the potential for further office expansion and will not provide a robust business continuity solution.

Space planning before design criteria are in place

The amount of space required to house your datacentre infrastructure components can be significant. In the most robust of systems, the ratio of raised floor to support gear can be as high as 1 to 1; leading some organisations to base their space planning on the quantity of IT equipment alone.

Critical IT and communications equipment requires a lot of space, as does the essential power and cooling equipment required to keep your datacentre running efficiently. Add to this the requirement for office space for support personnel, storage space and IT equipment staging areas and it is easy to see why you have established your design criteria before you begin space planning.

“While the physical design of a datacentre is critical, how a site is operated and maintained plays a more significant role in achieving site availability”.

The Uptime Institute

Designing into a dead-end

Modular datacentre solutions are based on the principle of being able to add chunks of additional infrastructure on a just-in-time basis, to minimise unnecessary capital expenditure. However, businesses can still dead-end themselves by incorrectly predicting their future requirements.

Everything can, and will, change. A flexible, modular design can hold the key to long-term success. Even the best kilowatt/rack planning can be obsolete due to consolidation, exponential business growth via acquisition, or a sharp turn to an unplanned high density footprint.

Electrically, you should make sure that your design includes the ability to add UPS capacity to existing modules without an outage. Design your input and output distribution systems to accommodate any future change in your base build criteria. The cost to oversize distribution for future capacity needs is not significant in your overall TCO model.

Mechanically, most users can meet their cooling requirements via conventional perimeter cooling with the proper floor height and hot/ cold aisle planning. However, one high-density rollout can change everything. Make sure your core design allows for the flexible (uninterrupted) implementation of custom in-rack/in-row cooling solutions.

Misunderstanding PUE

Power Usage Effectiveness (PUE) is an effective tool to drive and measure datacentre infrastructure efficiency. However, broad energy efficiency claims can lead to a degree of misunderstanding; there is almost always a capital cost associated with achieving lower PUE.

Organisations often set PUE goals without taking into account all of the contributing factors. It is essential to understand what return on your capital investment is required to achieve your objectives. Ask yourself, what is the TCO relative to the target PUE?

Here are two cautionary examples of misunderstanding the relationship between PUE, ROI and TCO:

- > When calculating PUE during the design phase, was it calculated based on the “perfect” day or on a yearly average?
- > Was the calculation based on a fully loaded or partially loaded datacentre? All equipment efficiency curves fluctuate according to load profiles. PUE changes daily, if not hourly, under normal operating conditions.

Use PUE to your advantage to meet your overall business goals, but be cautious. Try not to get trapped into misusing the calculation formula to justify overall capital expense and operating expense budgets.

Overcomplicated designs

Simple is best; whatever you target tier rating. Redundancy objectives often drive an increase in complexity, as do the multiple approaches to modular design. Whether you engage with an internal resource or an external consultant, try and keep things simple. Why?

- > Complexity often means more equipment and components. More parts mean more failure points.
- > Human error. Whilst the precise numbers vary, statistics show that most datacentre downtime is down to human error. Complex systems increase operational risk.
- > Capital investment. Simple systems are less costly to build.
- > Operations and maintenance costs. More equipment means an incremental rise in O&M costs.

Design with the end in mind. Whilst the design might look good on paper and it can be all too easy to justify a configuration or predicted uptime; if you haven’t considered maintainability during operations or servicing, your systems performance may become compromised.

Conclusion

Many datacentre builds and expansions fail to deliver on their promises; yours doesn't have to. By avoiding the mistakes outlined in this paper you are well on your way to delivering a successful datacentre project.

In summary:

- > Start with a Total Cost of Ownership approach
 - Evaluate your risk profile against your business expense profile
 - Create a model that incorporates CapEx, OpEx and energy costs
- > Determine your design and performance criteria
 - Base your criteria on your risk profile and business objectives
 - Ensure those criteria determine the design; including tier level, location and space planning - not the other way around
- > Design with simplicity and flexibility
 - Design a system that will meet your uptime objectives but keep build and operation costs down – simplicity is the key
 - Accommodate unplanned expansion by incorporating flexibility into the design
- > If PUE is a part of your criteria, understand the implications and costs associated with it

Through proper planning, using the TCO approach, you can create a datacentre facility that meets your organisation's performance goals and business objectives for today and tomorrow.

Contact us

For further information on the datacentre services available from Comtec Power, contact us:



Tel: 0800 008 7599



Fax: 01737 336105



Email: sales@comtecpower.com