Determining Total Cost of Ownership for Data Center and Network Room Infrastructure

White Paper 6

Revision 4

by Neil Rasmussen

> ## > Executive summary
>
> An improved method for measuring total cost of ownership (TCO) of data center and network room physical infrastructure and relating these costs to the overall Information Technology infrastructure is described, with examples. The cost drivers of TCO are quantified. The largest cost driver is shown to be unnecessary unabsorbed costs resulting from the oversizing of the infrastructure.

Contents
Click on a section to jump to it

Schneider Electric

Introduction

Predicting and measuring total cost of ownership (TCO) for the physical infrastructure for network rooms and data centers is required for return-on-investment (ROI) analysis and other business decision processes. In addition, an understanding of the cost drivers of TCO provides insight into opportunities to control costs. Many users are surprised when they consider that the TCO for physical infrastructure may be comparable to or larger than the TCO of the supported Information Technology equipment.

In this paper, a method for determining TCO of physical data center and network room infrastructure will be described. In this context, the physical infrastructure means all of the facility equipment needed to provide power, cooling, and physical protection of IT equipment, but not the IT equipment itself.

There are no recognized standards for measuring the TCO of the physical infrastructure of data centers. Simple methods of summing various capital and operating expense items do provide insight into total cash outlay, but they do not account for the utilization of the equipment.

Consider the case of two data centers, each 100 kW capacity and built identically; in one case the data center is fully utilized to 100% of the space and power capability; and in the other case the data center has only a single rack with 2 kW of information technology equipment. While the cash cost of operating these two facilities over their lifetime is comparable, the useful return on that investment is totally different. In the totally utilized case, the TCO of the data center is spread over a large amount of information technology equipment providing useful services. In the lightly utilized case, the entire burden of the expense of the physical data center infrastructure must be borne by the single rack. When TCO of physical data center or network room infrastructure is measured from the point of view of the useful work performed, namely how much IT equipment is supported, underutilization can drive extraordinary cost.

This paper will show that when TCO is viewed from the point of view of the useful work performed, the single largest cost driver of TCO for data center and network room physical infrastructure is the unabsorbed overhead cost of underutilized infrastructure. The highest return investment that the average data center operator or specifier can make in relation to physical infrastructure is in rightsizing. The practical and feasible financial return from a rightsizing strategy is quantified.

Relating TCO to the useful work performed

Most attempts to quantify TCO end up expressing TCO per data center, per square foot of data center, or per kW of data center. These metrics are not useful when considering the useful work performed, and are unhelpful to information technology staff in determining project ROI. They don't at all help an individual attempting to determine the TCO associated with the deployment of a new IT application in an existing data center or network room environment.

One measure of physical infrastructure, which is commonly understood by both facility and IT personnel, is the rack itself. The power, cooling, and area requirements of a rack are reasonably well standardized from a facilities standpoint. The amount of IT functionality and equipment that can fit in a rack is reasonably well understood from an IT standpoint. This has lead to the concept of expressing facility infrastructure in terms of "racks" or "rack units", which is now beginning to gain wide acceptance.[1]

---

[1] Snevely, R., *Enterprise Data Center Design and Methodology*, Palo Alto: Prentice Hall PTR, 2002, p. 35

For this purpose, the term "rack" refers to an open frame rack or rack enclosure, and it also refers to proprietary cabinets such as mainframe computers and large disk systems.

Research shows that the typical data center is only utilized to 30% of its capacity. While some data centers are utilized to 90% or more of capacity, there are similar numbers utilized to only 10% of capacity. Furthermore, the utilization of a data center varies during its lifetime according to a relatively consistent pattern. Fraction of utilization, and its variation with time, is an important input into a TCO model. In this paper, the typical model of **Figure 1** is used. For a more complete discussion of this topic, consult White Paper 37, *Avoiding Costs from Oversizing Data Center and Network Room Infrastructure*.

Related resource
White Paper 37

Avoiding Costs from Oversizing Data Center and Network Room Infrastructure

Figure 1

Utilization fraction of data center power infrastructure over lifetime
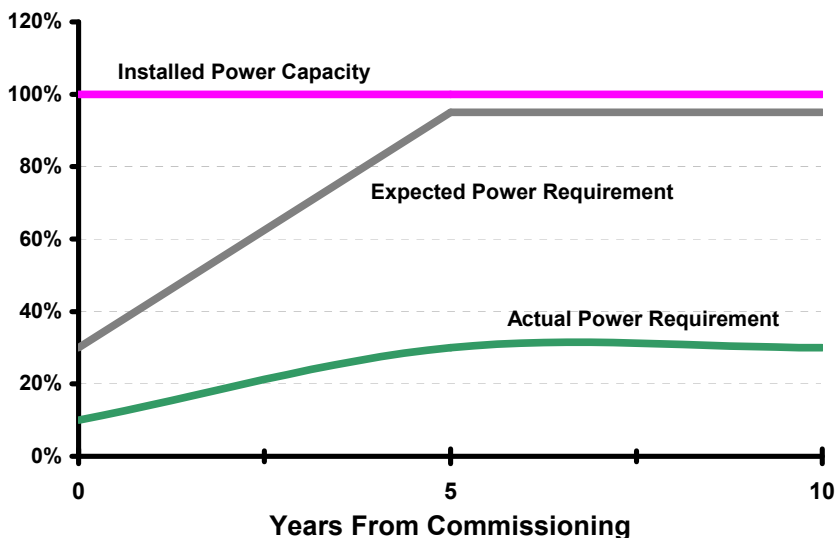
When TCO is expressed in a per-rack basis, the total cost of the data center or network room is allocated across the utilized racks. In this way, no unallocated overhead costs are created, and costs associated with data center or network room physical infrastructure can be more accurately and directly related to the IT infrastructure.

Typical TCO example

To determine TCO and express it on a per-rack basis requires a significant amount of data, including capital, engineering, installation, and operating cost data for the various elements of physical data center or network room infrastructure, as well as design-related parameters such as sq. feet per rack, watts per rack, utilization schedule, expected lifetime, redundancy options, etc.

For the data presented in the subsequent sections, TCO calculations were performed using the Schneider Electric TCO Calculator software application, developed by the Schneider Electric Data Center Science Center (accessible at http://www.apcc.com/tools/isx/tco).The TCO Calculator was set up to use average costs for capital equipment, installation, engineering, and operating costs. The results presented reflect values based on industry and site averages. Utilization data is assumed to exhibit the profile of **Figure 1**.

To determine TCO data for a typical data center, the typical data center was defined to be one exhibiting the following characteristics:

- Power rating: 100 kW
- Power density: 50 W/sq ft
- Life Cycle: 10 years
- Average rack power: 1,500 W
- Redundancy: 2N

Tests confirm that changing these parameters over typical ranges does not materially affect the results or conclusions of this paper.

Findings

The total cost of ownership of a rack in a data center is approximately $120K over the data center lifetime. In many cases, this cost is comparable to the cost of IT equipment that a rack is likely to contain over the data center lifetime. Approximately half of the lifetime per rack TCO of $120K is capital expense, and half is operating expense. These costs break down into categories as shown in **Figure 2**.
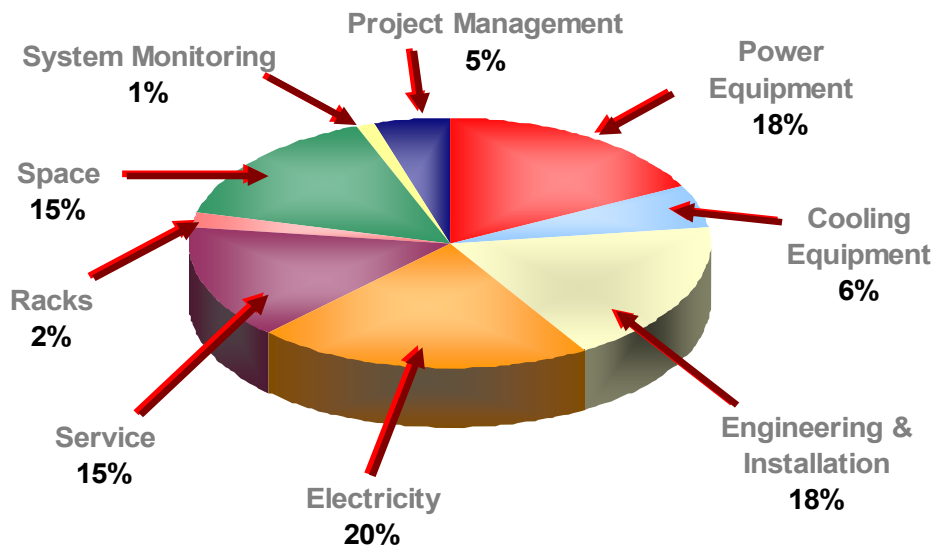
Figure 2

Breakdown of TCO cost components for a typical rack in a high availability 2N data center

A review of the breakdown of TCO cost components provides insights into opportunities to control or reduce costs in various areas.

Opportunities to control TCO

A variety of strategies to control lifetime total cost of ownership are apparent. These include efficiency improvement, improved planning, rightsizing the system, negotiating costs, self service, etc. Using the TCO calculator, it is possible to examine the effect that various scenarios have on TCO, in order to identify fruitful areas for investigation and investment. The per rack TCO savings for a typical 2N data center or network room resulting from a variety of scenarios are summarized below in **Table 1**:

Table 1

Data center physical infrastructure scenarios for reducing data center TCO

Scenario	TCO $ saved per rack	% of TCO saved
Purchasing power equipment with 2% higher electrical efficiency	*$1,472*	*1.1%*
Reducing the electricity rate by 1 cent per kW–Hr	*$3,100*	*2.4%*
Eliminating the raised floor	*$4,200*	*3.3%*
Increasing the cooling performance coefficient by 100%	*$5,500*	*4.3%*
Obtaining space at no cost	*$12,000*	*9.4%*
Obtaining all capital equipment at 50% discount from standard	*$15,700*	*12.3%*
Rightsizing the system to the actual requirement over time	*$76,400*	*60.1%*

The savings in the table are per rack; therefore the savings for any size data center or network room can be determined by multiplying by the number of racks. Note that in each case substantial savings are possible when compared with a typical legacy design, but the feasibility of obtaining the benefits is uncertain.

Rightsizing of the system provides the largest potential benefit among the above scenarios. This is due to the oversizing effect described earlier and summarized in the previous **Figure 1.** Large savings are obtained by rightsizing because A) data center or network room infrastructure that is never needed is never deployed, and B) data center or network room infrastructure that is needed is not built out until it is needed.

Practical benefits of rightsizing

A deployment strategy that could prevent oversizing was shown in the previous section to offer the potential of reducing the cost of infrastructure by 60%. Ideal data center or network room architecture would be "rightsized" and only incur the infrastructure costs that are actually required at a given time. To achieve the theoretically available cost savings, the ideal data center or network room architecture would only have the power and cooling infrastructure needed at the moment; it would only take the space that it needed at the moment, and it would only incur service costs on capital infrastructure capacity that was actually being used. It would be perfectly scalable. While such an ideal architecture is not currently feasible, it is practical and cost effective to consider an approach to data center and network room deployment that is scalable and modular in order to obtain a significant fraction of the cost savings available by rightsizing.

With technology available today, many components of a data center and network room, such as Uninterruptible Power Supply (UPS), power distribution, and air conditioning, can be deployed in a modular and scalable method. One example of such a scalable architecture is the InfraStruXure architecture from Schneider Electric. When such components are deployed over time to match the actual requirement, cost savings are realized not only for the equipment itself, but also for service contracts and electricity. There are many contributors to the total cost of a data center that are difficult to scale with time and are typically incurred upfront, such as facility space improvements, facility switchgear, and engineering costs. When the modular, scalable technology is used to the extent which is currently practical, the TCO Calculator estimates that approximately 50% of the theoretical savings from rightsizing can be obtained. This result is show in **Figure 3**.

Figure 3

Per rack lifetime cost for three types of data center and network room infrastructure types
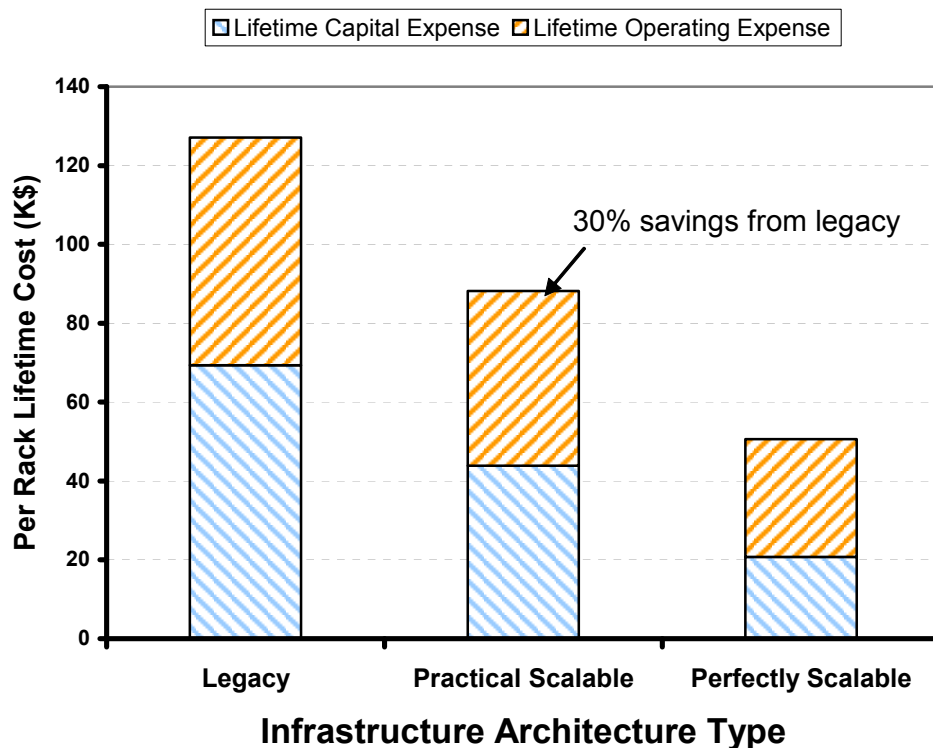
The savings in the "practical scalable" example in **Figure 3** do not reach the theoretical or "perfectly scalable" level because no practical scalable technology yet exists to allow certain data center components like fire suppression, raised floor, facility space, or switchgear, to be deployed in a modular scalable fashion. Furthermore, some components like UPS cannot be minimally deployed to exactly match the load requirement, but must be deployed in steps incorporating a safety margin. Nevertheless, a very significant savings in lifetime TCO are practical.

65% of the savings of the "practical scalable" example of **Figure 3** were capital expense reductions, and 35% were operating expense reductions. The cash-flow requirement is reduced every year by using the scalable approach, with the largest savings in year 1. Over 90% of capital is typically spent in year 1 using legacy design approaches, yet this is when the utilization of the facility may be lowest and the knowledge regarding future requirements is poor. This can make the investment difficult to justify based on ROI.

The savings obtained in a specific situation will vary based on assumptions and constraints of the particular project; for a more exact estimate the TCO Calculator should be run with site-specific inputs.

Conclusion

Expressing TCO for data center and network room physical infrastructure on a per-rack basis normalizes the measurement of TCO, providing a metric that can be used to compare data centers / network rooms and to compare different approaches to design.

The use of a TCO calculation tool and method was described. This tool permits the evaluation of cost control strategies, and can estimate TCO for specific installations.

The per rack lifetime TCO of a high availability data center is approximately $120K. The oversizing of infrastructure is a major contributor to this cost, and on average 30% cost savings can be obtained by implementing practical design techniques using modular scalable data center infrastructure.

✎ About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 19 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC InfraStruXure system.

Prior to founding APC in 1981, Neil received his bachelors and masters degrees from MIT in electrical engineering, where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.

Resources

Avoiding Costs from Oversizing Data Center and Network Room Infrastructure
White Paper 37

White Paper Library
whitepapers.apc.com

TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative