

Overcoming the Challenges of Server Virtualisation

Maximise the benefits by optimising power & cooling in the server room

Server rooms are unknowingly missing a great portion of their benefit entitlement from virtualisation and server consolidation. Optimising the physical infrastructure that supports them, particularly, row-based cooling, correctly sized power and cooling and real-time capacity management are essential elements in realising the full potential in cost reduction, efficiency and reliability. Optimised power and cooling post server consolidation and virtualisation can save up to 50% on your server room energy bill and significantly reduce your carbon footprint while assuring business uptime.

Introduction

There are three key scenarios to understand about server virtualisation and consolidation as it relates to the power and cooling infrastructure in server rooms:

- Power and cooling technology is available today to safeguard availability and meet the challenges of density and dynamic power that often accompany server virtualisation and server consolidation.
- Power consumption will always be less after server virtualisation, as a result of computing consolidation and physical reduction of the amount of IT equipment. With optimised power and cooling to minimise unused capacity (the subject of this paper), power consumption will typically be much less.
- Server room IT energy efficiency will go down after virtualisation, due to fixed losses in unused power and cooling capacity. With optimised power and cooling to minimise unused capacity, power and cooling efficiency can be brought back to nearly pre-virtualisation levels – sometimes even better, depending upon the nature of improvements to the cooling architecture.

Properly designed physical infrastructure will not only provide solutions for the specific power and cooling demands of virtualisation, but – especially if replacing traditional room-based cooling systems – can raise both power density capacity and server room efficiency significantly above what they were before virtualisation.

The changing server room environment

- Consolidation makes the server room a dynamic and unpredictable environment.
- Virtualisation dramatically increases criticality of the physical server availability.
- Higher density blade servers can drive the need for dedicated supplemental cooling.
- IT managers need to deliver higher levels of availability driving the need for greater performance, power, cooling, racks, management and services.
- Corporate governance, sustainable business policy and pending UK government regulations are driving the need for increased energy efficiency requirements.

Challenges to power and cooling infrastructure

Virtualisation creates changes in the server room that present new challenges to the power and cooling infrastructure, with implications to both effectiveness and efficiency. That is how well it performs the job of safeguarding the IT load and how well it conserves power while performing that job. While an upgrade of power and cooling systems is not necessarily required to make virtualisation work, a significant efficiency benefit from virtualisation will be realised with power and cooling that responds to these challenges, which can be characterised as follows:

	Virtualisation challenge to power and cooling infrastructure	Solution
1	Dynamic and migrating high-density loads	Row-based cooling – variable cooling near the load
2	Under loading of power and cooling systems	Scalable power and cooling - matching power and cooling to the load
3	The need to ensure that capacity meets demand at the row, rack and server level	Capacity management tools

While these challenges are not new or unique to virtualisation, the effects are focusing attention on them with a new urgency, especially in light of the expanding interest in energy efficiency and carbon reduction.

Meeting challenge #1 - Dynamic and migrating loads

Row based cooling

While virtualisation may reduce the overall power consumption in the room, virtualised servers tend to be installed and grouped in ways that create localised high-density areas, or “hot spots”. Not only are densities increasing, but virtualisation also allows applications to be dynamically moved, started and stopped – the result of which can be loads that change both over time AND in their physical location in the room. Before virtualisation brought dynamic allocation of server loads, localised high-density hot spots stayed put.

Figure 1

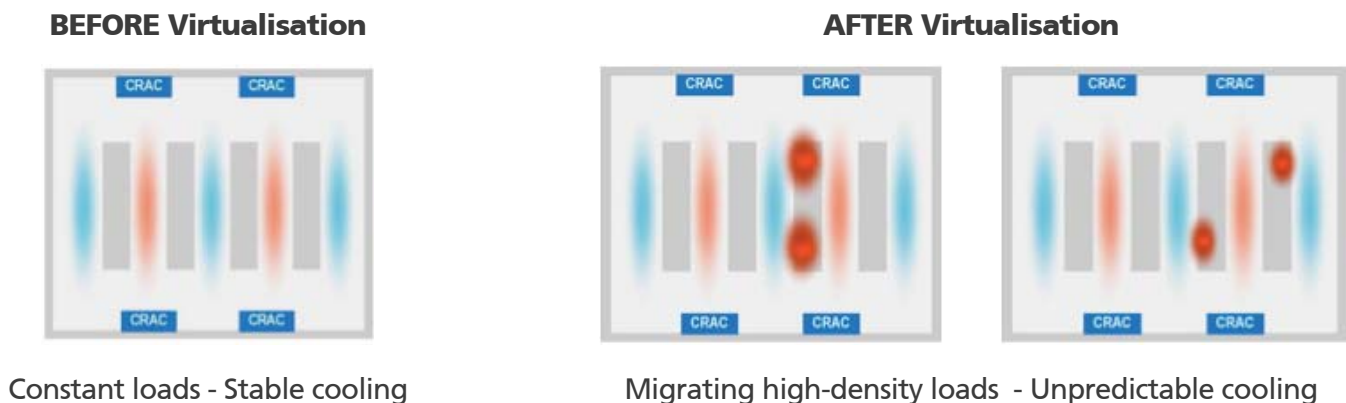


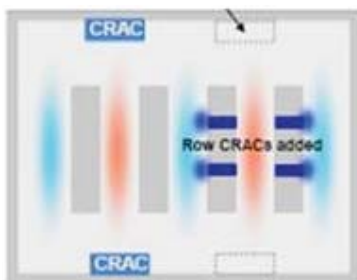
Figure 1 – High-density “hot spots” that vary both in power density AND location can result from virtualised and consolidated IT loads. (Note: CRAC = Computer Room Air Conditioning)

The placement of cooling units close to the servers provides the essential element that is key to efficient cooling - short air paths. A **short air path** between cooling and the load enables a number of efficiency and availability benefits:

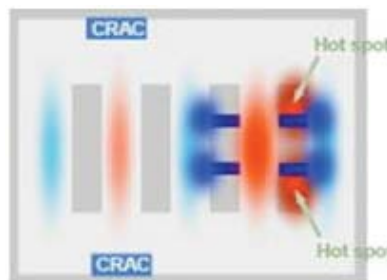
- Targeted cooling that can respond to localised demand.
- Conservation of fan power.
- Reduced mixing of cold supply air with hot return air.
- Increased return temperature (increases rate of heat transfer to coil).
- Reduced – often eliminated – need for make-up humidification (to replace condensation formed on a too-cold coil resulting from a too-low set point).

Figure 2

Room CRAC removed and Row based CRACs added



Row based CRACs sense elevated temperature and increase fan speed to remove extra heat from hot aisle



When temperature decreases, Row based CRACs decrease fan speed to conserve energy

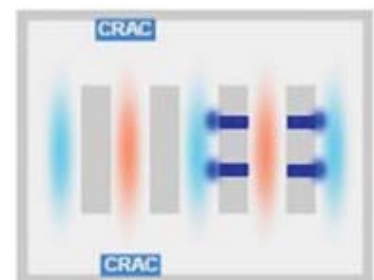


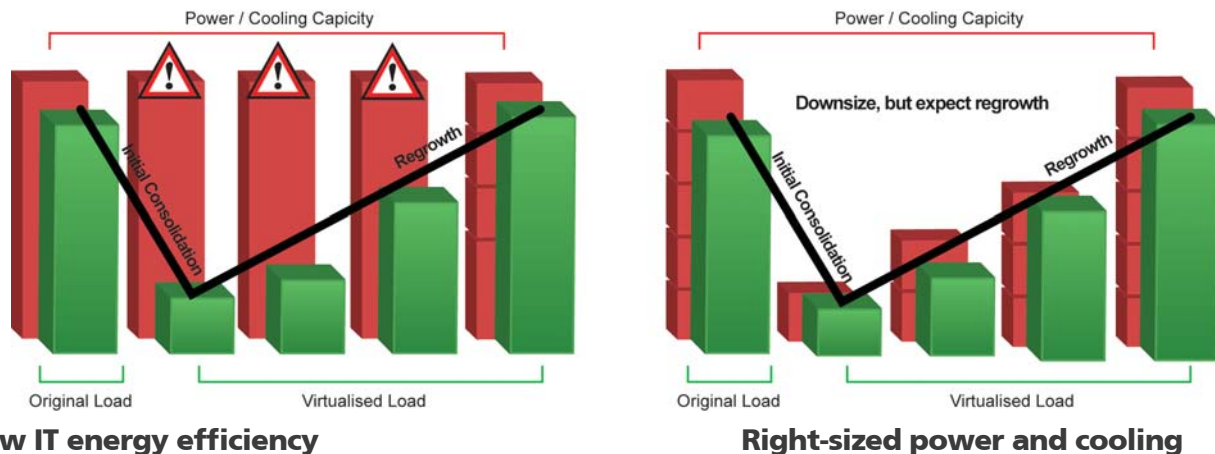
Figure 2 – Row-based CRACs work together to remove extra heat from hot aisle

Meeting challenge #2 – Under-loading of power & cooling systems

Providing scalable power and cooling - Matching power and cooling to the load

The reduction in IT load as a result of server consolidation offers a new opportunity to take advantage of modular, scalable architecture for power and cooling. Until now, the usual argument in favour of scalable architecture has been the ability to start small and grow as needed, to avoid over-investment and wasted operational cost from infrastructure that may never be used. With virtualisation, scalable architecture now allows scaling down to remove unneeded capacity at the time of initial virtualisation, with the later option to re-grow as the new virtualised environment re-populates. Whether scaling up or scaling down, the idea is the same – power and cooling devices are less efficient at lower loading, making it wasteful to be running more power or cooling than you need. “Right-Sized” infrastructure keeps capacity at a level that is appropriate for the actual demand, taking into account whatever redundancy and safety margins are desired.

Running more power and cooling equipment than needed is like leaving your car running at idle. When not in use energy is consumed, but no useful work is accomplished. Since virtualisation can significantly reduce load, over-sizing becomes an important efficiency issue. While the electricity bill will go down due to the lower IT load and less air conditioning, the proportion of utility power that reaches the IT loads will drop, signifying wasted power that could be conserved to further reduce energy consumption.

Figure 3

Low IT energy efficiency

Power and cooling not downsized after virtualisation.

Right-sized power and cooling

Figure 3 – Using scalable power and cooling to minimise the inefficiency of unused capacity during consolidation and re-growth.

Meeting challenge #3 - Ensure that capacity meets demand at the row, rack and server level

Capacity management - Knowing what's going on in real time

The dynamic nature of virtualised computing demands accurate, timely and actionable information about power and cooling capacities. Ensuring that power and cooling are keeping up with a changing load profile that can shift from one day to the next.

Capacity management provides instrumentation and software for the real-time monitoring and analysis of information about the three essential capacities in the server room: Power, Cooling and Physical space. Capacity management enables these resources to be utilised effectively and efficiently throughout the server room, through continuous real-time visibility to capacities at the rack and server level. With this data, management software can identify locations where there is available capacity of one or more resources, where a capacity is dangerously low, or where there is unusable (stranded) capacity of one or more resources.

An effective capacity management system uses automated intelligence and modelling to monitor power, cooling and physical space capacities at the room, row, rack and server level. It also suggests the best place for adding equipment, predict the effect of proposed changes and to recognise conditions or trends in time for corrective action to be taken. Capacity management that comprehends server location and loads, power and cooling capacity available to servers, temperature fluctuations and power consumption, not only protects against downtime from localised shortages of power or cooling, but also increases efficiency by optimising the use of available resources. A holistic system like this can:

- Model the system-wide effects of proposed server changes.
- Compare alternative layouts using detailed design analysis.
- Confirm that proposed changes will not cause a power or cooling overload before deployment.
- Verify that a change was made as planned.
- Reserve power, cooling and rack space so new equipment can be installed quickly.

The need for capacity management is greatest when there is change – the hallmark of today’s virtualised server room such as a changing server population, varying power density, load migration, the steady advance of new technologies and increasing pressure to conserve energy. Unmanaged change in the server room can compromise availability, thwart planning and waste resources. Effective capacity management comprehends the mechanics and wide-ranging effects of change, allowing the server room to utilise its power, cooling and physical space to maximum advantage. With this intelligence, virtualisation can fulfill its potential for efficiency and business value.

Virtualisation’s effect on energy efficiency

Virtualisation will always reduce power consumption due to the optimisation and consolidation of computing onto a smaller number of physical devices. However, if no concurrent downsizing or efficiency improvement is done to the power and cooling infrastructure, then the server room’s efficiency curve will remain the same and efficiency (DCiE) will be reduced (Figure 4).

To improve post-virtualisation energy efficiency, the server room infrastructure efficiency curve must be raised by optimising power and cooling systems to reduce the waste of over-sizing and align capacity with the new, lower load (Figure 4). The greatest impact on the efficiency curve can be made by going from room-based to row-based cooling and by “right-sizing” the power and cooling systems. In addition to improving efficiency, optimised power and cooling will directly impact the electricity bill by reducing the power consumed by unused power and cooling capacity.

Figure 4

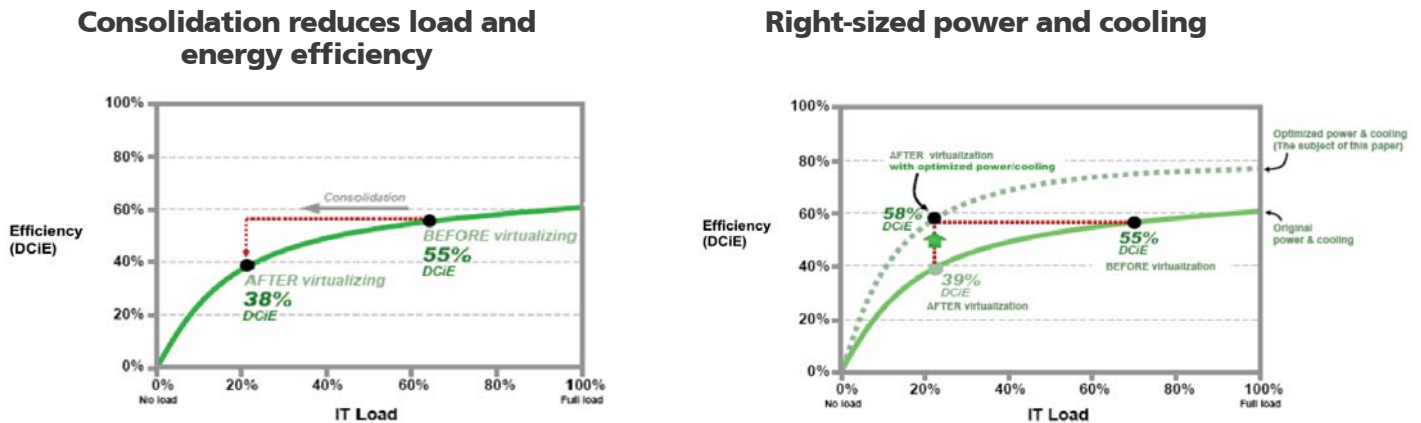


Figure 4 – Effects on energy efficiency

Reduced power consumption and reduced electricity bills

Power consumption will always be less (better) after virtualisation due to reduced server population. However, efficiency will be lower (worse) if power and cooling systems are not downsized and optimised to align with the new, smaller IT load. In other words, physical infrastructure that doesn’t “slim down” to match the lower IT load will still be consuming power to maintain excess or misdirected capacity. Figure 5 illustrates this typical outcome – reduced power consumption combined with lower infrastructure efficiency.

To compensate for this problem and realise the full energy-saving benefits of virtualisation, an optimised power and cooling infrastructure will reduce power consumption and maximise energy efficiency, as illustrated in figure 5. This optimised infrastructure will incorporate design elements such as the following

to maximise the electrical efficiency of a server consolidation and virtualisation project:

- Power and cooling capacity scaled down to match the load.
- Fans (and pumps) that slow down when demand goes down.
- Equipment with better device efficiency to consume less power in doing the job.
- Cooling architecture with shorter air paths (e.g. move from room-based to row-based).
- Capacity management system to balance capacity with demand.
- Blanking panels to reduce in-rack air mixing.

Figure 5

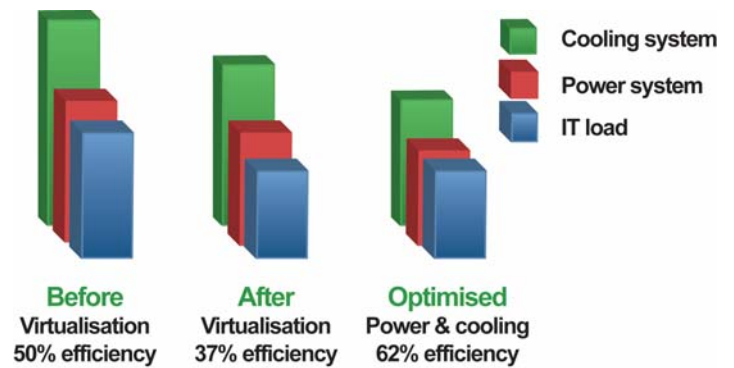



Figure 5. Example: Reduced power consumption post virtualisation and with optimised power & cooling

Conclusion

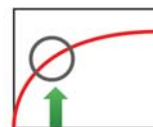
Virtualisation is an undisputed leap forward in server room and data centre evolution – it saves energy and increases computing throughput, frees up floor space, facilitates load migration and disaster recovery. Less well known is the extent to which virtualisation’s entitlement can be multiplied if the power and cooling infrastructure is optimised to align with the new, leaner IT profile. In addition to the financial savings obtainable, these same power and cooling solutions answer a number of functionality and availability challenges presented by virtualisation. Figure 6 summarises the effects of the optimised power and cooling infrastructure described in this paper, both as it meets the technical challenges of virtualisation and as it delivers general performance benefits.

Figure 6

Challenges of virtualisation



Dynamic and migrating high-density loads



Underloading from consolidation



The need to ensure capacities down to rack level



Meet technical Challenges

Optimised power and cooling



Deliver entitlements

Entitlements from virtualisation



Reduced Overall power use beyond virtualisation alone



Increased power/cooling efficiency (DCiE)



Figure 6: server consolidation and virtualisation challenges and their solutions

The three major challenges that server consolidation and virtualisation pose to physical infrastructure are dynamic high density, under-loading of power and cooling systems and the need for rack-level, real-time management of capacities (power, cooling and physical space). These challenges are met by row-based cooling, scalable power and cooling and capacity management tools, respectively. All three of these solutions are based on design principles that simultaneously resolve functional challenges, reduce power consumption and increase efficiency.

With server consolidation and virtualisation, a parallel upgrade of power and cooling will optimise both the architecture and operation in a number of ways that safeguard availability, enhance manageability, lower power consumption and increase efficiency. Properly designed physical infrastructure will not only provide solutions for the specific needs of server consolidation and virtualisation, but can also raise both power density capacity and IT energy efficiency significantly above previous levels.

Comtec provide a range of assessment services to guide organisations toward efficiency savings that offer the greatest net reward for server consolidation and virtualisation. For more information, contact the Comtec Power team on **0845 899 1428** or email power@comtec.com.

About Comtec

Comtec Power specialises in the design, build, maintenance and ongoing management of comms rooms, server rooms and datacentres. Our work encompasses new builds as well as the refurbishment and upgrade of live facilities. Our goal is to create technically excellent facilities which are scalable and efficient to meet the power and cooling demands of latest generation IT equipment. Comtec has a thorough and experienced knowledge of the IT industry and is perfectly positioned to deliver highly resilient solutions for high density IT deployments including blade servers and HPC installations. Comtec's maintenance and inventory management services help to reduce the threat of downtime mitigate the risk of thermal shutdown and keep facilities operating optimally.

About APC by Schneider Electric

APC by Schneider Electric provides protection against some of the leading causes of downtime, data loss and hardware damage: power problems and temperature. As a global leader in network-critical physical infrastructure (NCPI) solutions, APC sets the standard in its industry for quality, innovation and support. Its comprehensive power and cooling solutions, improve the manageability, availability and performance of sensitive electronic, network, communications and industrial equipment of all sizes. Visit www.apc.com for further information.



Visit www.comtec.com/power
or call **0845 899 1428**
for further information.

