



## Effective cooling means careful planning

**Cooling the datacentre and reducing power costs is a critical part of any managers' job. From initial design to the day before decommissioning, it is a constant balancing act to keep power and cooling in balance. Unfortunately, too many see this as a one off exercise.**

When a datacentre is commissioned, significant amounts of time and effort are invested in planning the location of walls, racks, cables, pipes, power and cooling. The result of all of this is that the construction crews create a facility with a known infrastructure. This is incredibly important as, once created, there is little you can do to change those parts of the infrastructure you do not own.

The power available to datacentres is no longer an endless pot. Energy companies are not only capping the power, they are seeking to reduce it where possible. The days of deploying more hardware and asking for another power supply are gone. So, whatever power budget you start with has to reflect the lifecycle of the datacentre - the average being 15 years.

Over time, hardware becomes more power efficient. However, the demands to deploy more hardware is outstripping the ability of the hardware vendors to reduce power and cooling footprints. This means a delicate balancing act between reuse, and buying new hardware.

Many datacentres start with a thermal map showing how much heat can be supported across the datacentre. It provides a view of where hardware can be placed, powered and cooled, meaning that there comes a point in time where everything is in balance. From this point on, however, everything starts to change.

As you add more hardware, you need to plan power and cooling. Power, in many sites, is a fixed element. When the power and racks were

installed, there was an expectation of the maximum amount of power that could be delivered to a rack. While it is possible to make changes, it is expensive to do so and immediately changes the whole thermal model of the facility.

One technology, more than any other, has created this change – blade systems. They are more power efficient than racks of 1U servers and create less heat. The challenge for deployment, is their dense computing nature. A full rack of blade systems can replace up to 64 1U servers. Looked at in rack terms, you have almost halved the number of racks you needed, and reduced the amount of heat generated.

However, simply reducing the footprint doesn't mean you can install blade systems without understanding their impact on the overall facility. Re-routing power to provide dense power grids to racks of blade systems changes the airflow. The previous airflow model was designed to work around obstructions such as racks. Removing unnecessary racks has a significant impact on the effectiveness of the airflow, meaning hot and cold air may lose containment.

The fact that the heat is overall lower doesn't compensate for the fact that it is ejected into a smaller space, which ends up being hotter than it used to be. In addition, the heat isn't spread as evenly from floor to ceiling, and the cooling demands of devices higher in the rack are much higher than was originally planned for.

There are other significant changes, especially in the way the air is drawn in and expelled. A standard rack server tends to be cool air at the front, hot air at the back. This works well when planning hot and cold aisles. But not everything behaves like a server.

Switches tend to draw air from the side, as do many other appliances such as firewalls. This means there has to be sufficient cool air for them to use. They can be extremely hot on the top of the case – much hotter than servers, which creates localised hotspots within the rack enclosure. A lot

of switches are designed to sit at the back of the rack, rather than the front. This means you need to consider how they will access cool air and where the hot air will be expelled.

Storage systems are an even bigger challenge. They can draw air up and out, so you need to contain the air coming from the top of the system. This has a significant impact on how air is being managed and the containment of hot and cold aisles, as the air consumed and expelled is outside that aisle containment scheme.

To cool exceptionally hot racks, there is a range of water-cooling options that attach to the side of a rack. They do an excellent job of removing the heat from the dense computing solutions, but air that is being pushed to the front of those racks through the floor is now wasted. The hot air at the back of nearby racks ends up spilling out into the space where the contained racks would normally put air. The impact of this is a mixed airflow, and poor temperature management over time.

Installers will often look at the thermal model to see areas where there is spare cooling capacity. What they don't do is remodel the room based on the demands of the new hardware they are about to install. Again, this unbalances the cooling of the room, which leads to unexpected hotspots and a lot of reactive, remedial action. The result is a viscous cycle of time and money spent chasing, rather than planning a solution.

If you want to really control your power and cooling costs in the datacentre, you need to constantly review what is happening, and start using simulations to assess the impact of change. This will improve the behaviour of the datacentre, save money and over time allow a greater planned density of computing. With the demands for increased computing from business units, it makes sense to start planning now.