



Industry Perspective: N+1 Cooling On Paper Only

By Cary Frame, President and founder of Polargy.

N+1 redundancy is a system design best practice because equipment failure happens; we expect it and we plan around it. In the case of data center cooling, we expect a CRAC to go down at some point, and with an N+1 system design we have a spare CRAC to fall back on. A major problem with this aggregate view of cooling is risk of starving a cold aisle when changing the CRAC lineup or as a result of CRAC failure, and this risk is amplified with cold aisle containment. Fortunately, we can easily and cost-effectively manage this risk with 'air-mover' fan tiles in the cold aisles.

Time and time again we hear operators talk about how changing the CRAC lineup causes cooling airflow problems. Most of these comments come from legacy sites, though we've also heard them from new data centers. Interestingly, many operators fail to connect the dots between this airflow problem and N+1 redundancy. Many operators have a particular CRAC unit they don't dare turn off and yet they assume their spare CRAC unit gives them N+1 redundancy. In our experience, about a quarter of small-to-medium data centers suffer from this false assumption; their N+1 redundancy is on paper only.

The Balancing Act

I can't tell you how often I've heard comments like, "*We need to keep CRAC Unit #3 on all the time or the room overheats*" or "*When we take CRAC #6 down for maintenance we get hotspots on the north side.*" These problems are indicative of typical airflow behavior in raised floor environments:

- CFM into an aisle is highly dependent on under floor pressure and obstructions.
- CFM into an aisle is largely driven by the closest CRAC (the 'CRAC of Influence').
- Changing the CRAC lineup creates large swings in CFM delivered to an aisle.
- Total air supply may be sufficient but local supply may not be (the 'Distribution Problem').

With these understandings, one can easily see how a change in the CRAC lineup can cause under floor pressure to change enough to introduce significant risk of an adverse thermal event.

When we think of cooling sufficiency, thermal safety and preventing problems, it's in terms of normal operating conditions and failure conditions, and both scenarios are highly dynamic. In normal operating conditions we deal with routine changes in cooling demand and supply throughout room while the CRAC lineup remains unchanged. In failure conditions, we deal with a large change in under floor pressure when one CRAC goes offline and a spare unit takes over.

In normal operating conditions, routine changes in airflow demand and supply create risk of falling out of balance and starving a cold aisle. Cooling demand varies at the rack level, aisle level and room level, and can fluctuate either quickly or slowly. For example, a researcher who kicks off a large computational job can quickly heat up one or more racks of number-crunching servers. Or an IT guy swaps a 10kW rack in where a 2kW rack had been, but forgets to mention it to the facilities crew. These changes in demand for cooling create a less obvious change in cooling supply. When cooling demand in one aisle increases, the change in air consumption will affect the supply available to adjacent aisles. Such demand and supply changes during normal operations affect under floor pressure and can result in an aisle with localized low pressure.

In failure conditions, loss of a CRAC unit and replacement with the N+1 redundant unit will cause a change in under floor pressure. Because cooling supply to an aisle is most influenced by the nearest CRAC, and depending on the specifics of the under floor situation, a change in CRACs can result in a low-pressure zone and under-supplied aisles. There may be sufficient cooling supply, but because of the Distribution Problem, there is localized low pressure and even aisle starvation. In this case, even if the N+1 redundant CRAC unit comes online as planned, the best we can say is that the site has only partial redundancy.



Industry Perspective: N+1 Cooling On Paper Only

Fixing With Fans

Fortunately, achieving true N+1 redundancy and mitigating cooling failure risk we've described can easily be achieved with active fan tiles that locally modulate airflow. Raised floor fan tiles, such as the [Frost-Byte™ Raised Floor Fan Tile](#), vary speed to deliver cold air to the aisle based on sensed temperature or pressure differential versus a target setpoint. With several of these “air-mover” tiles in a contained cold aisle, the right amount of cold air is supplied to mitigate thermal risk from inevitable cooling demand and supply changes during both normal operations and failure conditions.

These active fan tiles are built with a matrix of high performance variable speed DC fans in an aluminum enclosure attached to a standard 60% raised floor tile. Commonly, a temperature sensor mounted on the face of server racks controls the fans. Alternatively, sensors that measure pressure differential between inside and outside the contained cold aisle control the fans. This fan tile architecture auto-balances the cold aisles, eliminating starvation risk and improving thermal safety.

Other Fan Tile Benefits

An alternative solution to balancing cooling demand and supply is simply to over-supply an aisle, but with today's emphasis on energy efficiency, the days of oversupplying are largely over. In fact, energy efficiency is the major factor driving the adoption of aisle containment, though even with containment, we sometimes still see oversupply due to balancing challenges. With active fan tiles, these remaining oversupply scenarios can be reduced or eliminated, yielding the full efficiency promise of cold aisle containment.

Additionally, by auto-balancing with active fan tiles, operators achieve labor savings from the elimination of routine manual balancing. The days of walking the room and swapping out perforated tiles are over. Active fan tiles eliminate the need to analyze aisles with a balancing hood (aka: flow balometer) to ensure the CFM in the cold aisle more than matches the IT load in that aisle. Likewise, because conditions in the room and aisles are so dynamic, Computation Fluid Dynamics (CFD) analysis for balancing purposes becomes obsolete since CFD provides a historical ‘snapshot’ of airflow that may no longer be relevant.

Lastly, if active fan tiles are powered through a UPS, they can ensure greater uptime if cooling is completely lost. In a catastrophic cooling failure condition, the under floor plenum holds a cool air reservoir, though without air pressure or air flow. Fan tiles running on UPS backup can continue to deliver and circulate cold air from within this chilled plenum. Testing demonstrates that supply air temperature through the fan tiles remained steady for more than 10 minutes even with all CRACs off.

The benefits of auto-balancing and solving the Distribution Problem that active fan tiles provide allow operators to enjoy true N+1 redundancy. Fan tiles offer significant additional benefits in a data center with cold aisle containment: even greater energy savings, local balancing and thermal safety.

Frost-Byte™ raised floor fan tiles combined with Polargy's cold aisle containment take data center energy efficiency and thermal safety to the next level.

About Cary Frame

Cary is President and founder of Polargy, a provider of hot and cold aisle containment solutions. Cary's prior roles include General manager of SharkRack, a manufacturer of data center infrastructure and vice president of Sales for Converge, a Hewlett-Packard spinoff. Cary earned an MBA from The Wharton School, a BSEE from Michigan State University and is a veteran of the United States Navy.

Don't miss Cary's other Industry Perspectives and posts on the Polargy blog: [The Containment Chronicles](#).