Facility Cooling Failure: How Much Time do You Have?

A Dell Technical White Paper

Dell │ Data Center Infrastructure

David L. Moss

June 2011

Introduction

Two recent independent data center surveys[1] reported that roughly one-third of data centers experienced a power outage within the prior year. Even though redundancy may be built into your facility, IT equipment can reach critical temperatures during the shift to and from redundant power systems. This paper addresses the variables associated with the speed and extent to which the data center may heat up during a cooling failure. This is often called the *ride-through* time—how much time the facility has to ride through a failure without reaching a critical temperature.

Critical Temperatures

The air temperature at the air inlet is generally the most important temperature for IT equipment. This is usually located on the front surface of the equipment. Most IT systems have an operating maximum temperature of around 35°C (95°F). There are several reasons why you might be concerned about experiencing temperatures above this level:

- Automatic shutdown
- Reliability concern
- Warranty loss
- Performance loss due to component or system throttling

Ride-through Dependencies

The variables associated with ride-through time are numerous. Room density is one of the largest determinants of ride-through time. Additionally, there are dependencies on the following conditions:

- Operating temperature (at the start of failure)
- Amount of facility air over-provisioning relative to IT air consumption
- Room height
- Facility air and coolant movers (fans or pumps) with or without UPS back up
- Some containment systems
- IT power density (watts per square foot)

Test Results

Watts per square foot (W/sq. ft.) is a metric that holds little relevance in most present-day discussions. As varied as today's rack densities can be, this old metric does not reveal much about a data center's ability to cool specific racks. It does, however, paint a reasonable picture about room density and ride-through time. With room density being one of the key factors in ride-through time, the W/sq. ft. metric is a good delineation of behavior during failures. Results included in this paper will show differences in behavior for room densities of 150, 250, 350, and 450 W/sq. ft.

---

[1] 2010 Penn Schoen Berland (32%); IDC 2008 Market Segmentation (37%)

During a cooling outage, rack inlet temperatures trend similarly, regardless of the room density. There is a sharp rise, an inflection point, and then a sustained gradual increase in temperature that is generally linear.

Figure 1. Inlet Temperature Trend During Facility Cooling Failure

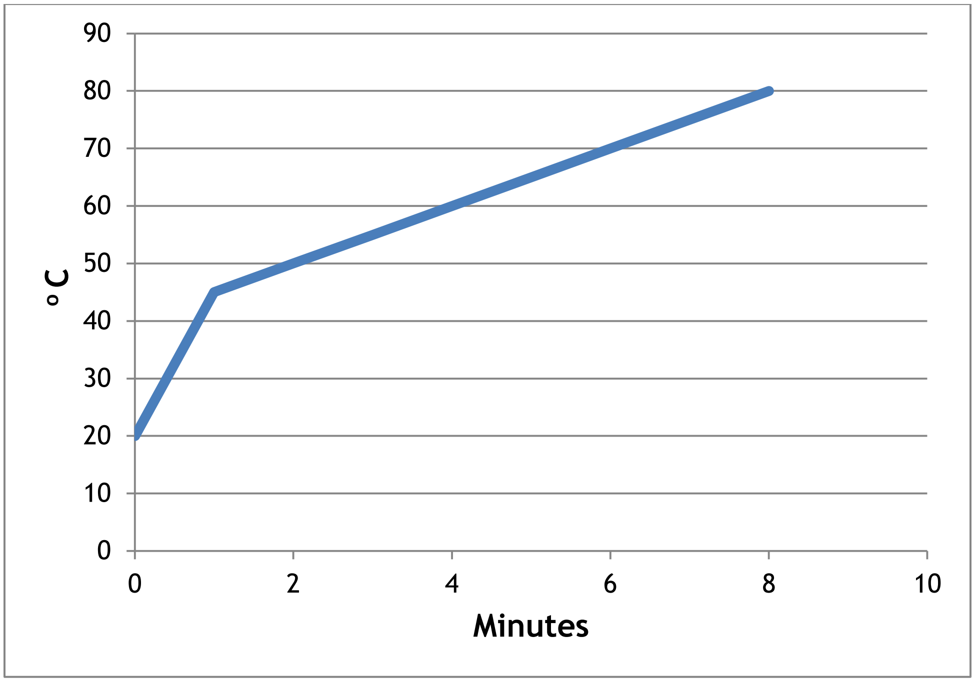

When the cool air production stops, there is a short period during which racks still have access to the cold air. Typically, this is the air that resides in an aisle at the point of cooling failure. The initial steep ramp in Figure 1 depicts the period of consumption of this cool air. If the data center were grossly over-provisioned in terms of volumetric delivery, the steepness of this ramp would lessen and it would take longer to hit the inflection point. During the initial ramp, the room is essentially normalizing on a temperature equivalent to the aggregate exhaust temperature of the IT equipment. The period to the right of the inflection point shows the IT equipment heating up the air and equipment within the room—the room no longer has the ability to remove its own heat. If the room has a taller ceiling, the slope of the secondary trend should decrease. In data centers lacking the back-up of a UPS, chilled water would stop flowing or compressors would stop running, and air handler blowers would stop moving air around the room.

Figure 2. Rack Inlet Temperature During a Total Cooling Failure

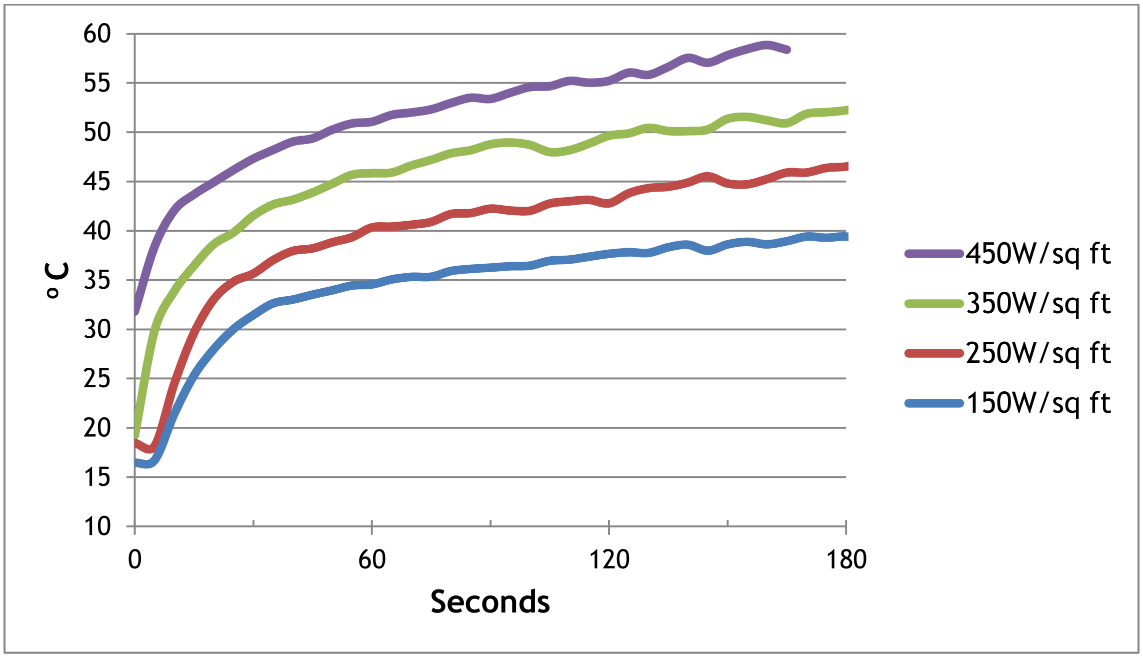

The curves in Figure 2 show the inlet temperature at the top of a rack after the cooling was shut off. In the three lower density experiments, there was enough cooling capacity or delivery capability to maintain a cool temperature at the top of the measured rack. The room's delivery capability was exceeded in the 450 W/sq. ft. test as evidenced by its high starting temperature. It is interesting to note, as seen in Figure 3, that air over-provisioning seemed to only affect the initial rise characteristics.

Figure 3. Inlet Temperatures of Different Racks (450 W/sq. ft.)

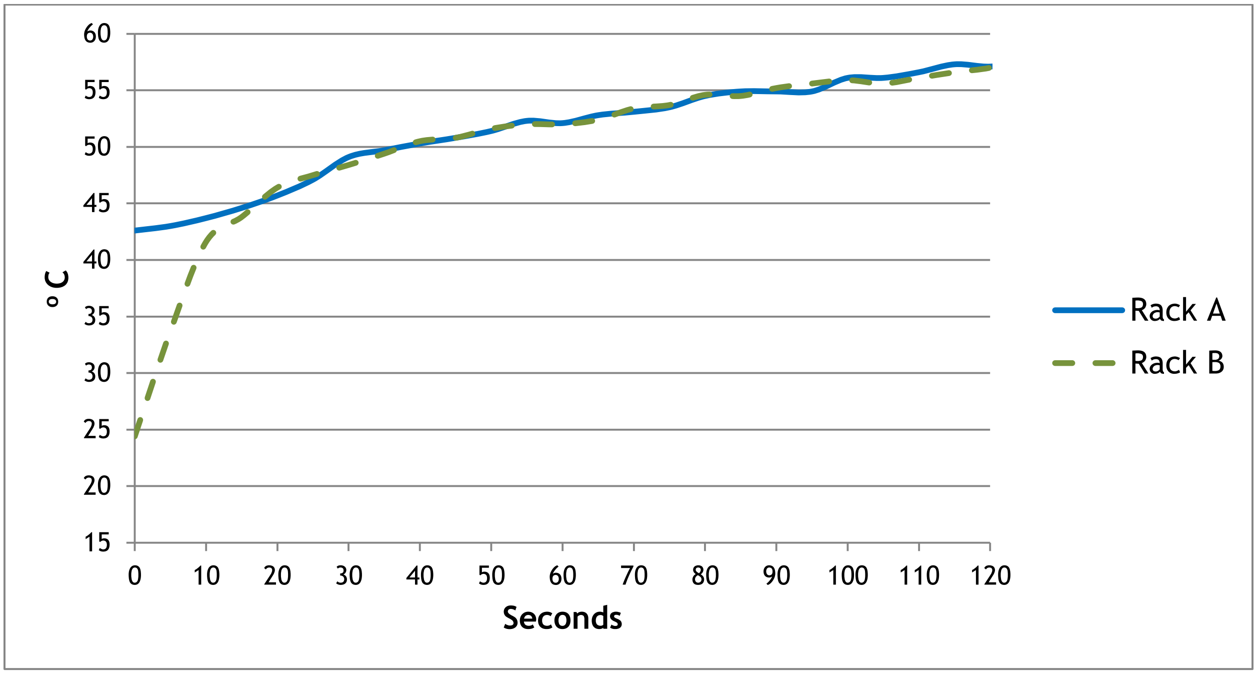

Rack B was located in an aisle of greater air over-provisioning, or to put it in other terms, Rack A was located in an aisle that was under-provisioned. It appears that volumetric over-provisioning of Rack B's aisle did not lengthen overall ride-through time, since the rack that was initially cooler (Rack B) caught up to the hotter rack (Rack A) within a fairly short time frame.

Ride-through tests were orchestrated in an 876-square-foot lab in Dell's Round Rock, Texas facility. This lab has a raised floor with chilled water CRAH (computer room air handler) units that are segregated but connected to the building's chilled water system. The water for the room can either be driven by the pressure from the building's chilled water system or can be rerouted to run using its own pumping system, drawing from the same chilled water feed. This separate control system allows for experimentation with water temperature. More importantly, it enables the water to be turned off to simulate a chiller/pumping failure. For each density case, two tests were run. A shutdown of the chilled water alone simulates a chiller or pumping failure with the CRAH units backed by UPS power. In the second test, the CRAH units were simply turned off to simulate a failure of the chilled water and a failure of the CRAH units as if they were not on UPS power. Figure 4 shows the stark contrast between the gradual temperature increase when the blowers are left running and the harsh increase if air stops moving.

Figure 4. 350 W/sq. ft. Test

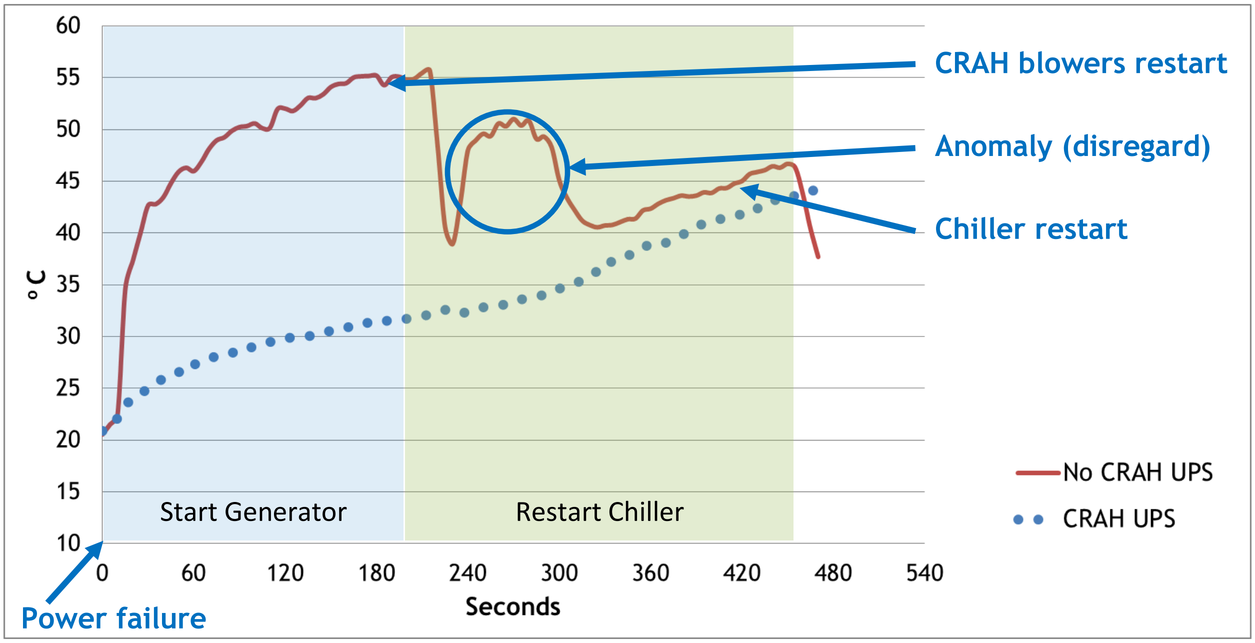

The generator should be started as soon as possible to get air movement around the room, especially when CRAH units are not backed by UPS. Normally, the generator comes on very quickly, but it could take longer than desirable. The three minutes represented in Figure 4 is considered an extremely long time, but not out of the realm of possibility. The chiller startup procedure is more involved. It typically happens well before seven minutes have passed, but it could take that long or even longer if it doesn't start up the first time. Figure 4 also contains an anomaly which is highlighted in the figure; this can be ignored since it is probably not reflective of a true failure.[2]

The air temperature just prior to the simulated failure in this test was about 21°C (70°F). For greater facility efficiency, Dell recommends operating at a slightly higher temperature of 25°C (77°F).[3] Had this test been conducted four degrees higher as suggested, both curves would have shifted up four degrees. In that case, at 350 W/sq. ft. and a nominal operating temperature of 25°C, this test would have resulted in an inlet temperature approaching 60°C (140°F) with no UPS backing of the air handlers. With reasonably dense rack spacing, 350 W/sq. ft. density is no more than 14 kW per rack on average. This is a fairly high density, but not extreme. One key takeaway from this test is the value added by keeping CRAH units on UPS.

Further Discussion on Variables

UPS-Backed Facility Fans

Figure 4 shows the dramatic difference between a still data center and one with backup power to keep air circulating. If you are deployed with this level of density, it is only a matter of seconds until you are over spec in the still data center, but if the air is moving, this extends your time to minutes. The incremental cost to add UPS capability to cover the facility fans should be no more than an additional 5%-15% above what you are spending to back up the IT. If at all possible, you do not want your systems

---

[2] See Appendix A for a greater explanation of test anomalies.
[3] "Data Center Operating Temperature: The Sweet Spot," 2011, David Moss

hitting the 50°C to 60°C (122°F to 140°F) temperatures seen in Figure 4. Even system with relatively high graceful shutdown limits will probably have shut down by the time they reach these levels. By design, Dell does not force a shutdown at high temperatures. There are, however, over-temperature protection limits built into all power supplies that would likely trigger by the time inlet temperatures reach the 50°C-60°C (122°F to 140°F) levels. Other vendors may have default shutdown limits that are lower than this.

Lower Normal Operating Temperature

It could be tempting to run a data center as cold as possible in an effort to have more time to switch over to redundant systems. The effect on ride-through time of the temperature set point is linear. If starting 5°C (9°F) lower, you will be about 5°C lower at all times during a failure. These tests were run with normal operating temperatures of 21°C (70°F). A typical chilled-water facility can get the air down to as low as about 13°C (55°F), and the ideal operating temperature is closer to 25°C (77°F). That 12-degree difference will still only buy you seconds on the non-UPS trajectory. If those seconds ensure you do not hit a critical temperature causing a server shutdown, you might estimate it to be worth it. You will be incurring a continuous OpEx penalty for running very cold nominal temperatures. Judging from the trajectory of the UPS-backed temperature increase in Figure 4, Dell recommends avoiding the OpEx penalty and investing in UPS for your CRAH units. You should ascertain whether or not you can get the chillers back on quickly and how many more minutes a colder nominal operating temperature affords you. It is important to balance this consideration against the long-term cost savings gained by running at 25°C (77°F) inlet temperatures instead of 13°C (55°F).

Even if you have an excursion of only a few minutes past allowable IT temperature limits, it would have a negligible effect on the equipment as long as the temperature does not reach automatic shutdown levels. Dell equipment will not automatically shut down until a much higher temperature than the allowable maximum. Additionally, in instances of short-term minor temperature excursions, your Dell warranty would remain effective. As you consider all factors, you should check the warranty language from your other equipment providers to better understand their specific warranties and limitations.

Room Height

Room height should lessen the slope of any of the curves presented in this paper. If the room height is doubled, it should approximately double the ride-through time. It might be useful to know size of the test room. Since room heating is directly related to volume, these results could be extrapolated to rooms with a different ceiling height. Where heat density is implicated in watts per square foot, the room size was 876 square feet (81 square meters). The volume above the raised floor was 11,400 cubic feet (323 cubic meters), and the volume below the raised floor was 1,314 cubic feet (37 cubic meters).

Containment

Containment can also play a role in ride-through time, especially if it is some form of tight containment like a chimney product. In August, Dell will begin shipping the PowerEdge Energy Smart Containment Rack.[4] Any containment product associated with this test leveraged these specific units, which have a sealed front door enclosing a deep front plenum for air to travel from under the rack to feed the equipment. Because of its tight containment path, the distance from IT exhaust to intake is quite lengthened. It helps to dampen the extreme temperatures that the IT systems see during the

---

[4] "Managing Data Center Costs with Dell PowerEdge Energy Smart Containment Rack Enclosures," David Moss, 2011, http://content.dell.com/us/en/enterprise/d/business~large-business~en/Documents~energy-smart-containment-rack.pdf.aspx

cooling outage (see test results in Figure 5). The dampening may be enough to lessen the risk for facilities not leveraging a UPS.

Figure 5. 350 W/sq. ft. Test, Including Dell Energy Smart Containment Rack

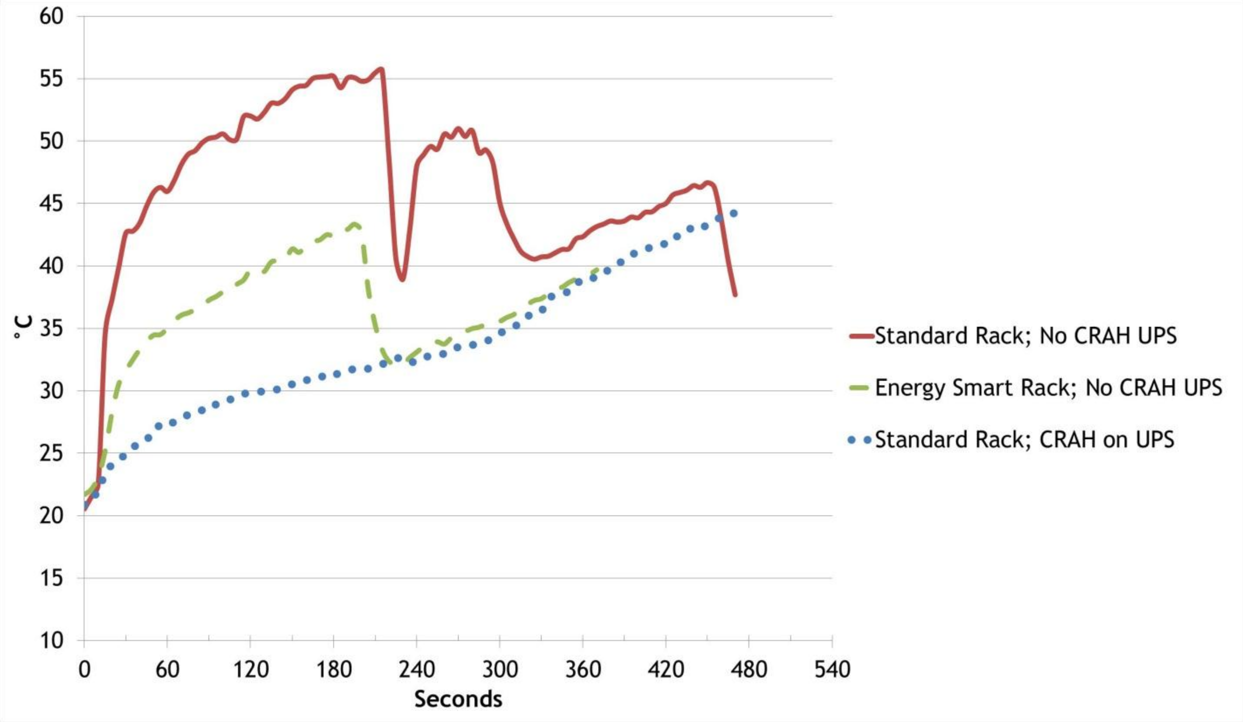

Summary

The higher the room density, the greater the chance that equipment could experience critical temperatures. The lag of component temperatures during an increase of system inlet temperature is quite small. Within seconds of surpassing the system maximum inlet temperature, you would expect to see some component approaching its maximum specification. This is a number, however, that is derated and intended to guard against long-term exposure at high levels. The short-term exposure of components to temperatures past their maximum specification will have no measurable effect on long-term reliability. Reliability fears are therefore misplaced when considering the high temperatures resulting from a facility cooling outage lasting only minutes.

The concern, then, becomes one of warranty, component throttling, or system shutdown. You may see a shutdown at the non-UPS temperature levels seen in Figure 4; you might also see some throttling prior to that shutdown. Dell's warranty remains valid in instances where equipment has been exposed to temperatures beyond allowable limits for short excursions. It is worth understanding whether this is the case with other vendors you leverage in your data center. Throttling typically does not happen until several degrees past maximum allowable temperatures, but throttling is likely the least of your worries during a cooling failure. Dell systems typically will not hit a shutdown trigger point until at least 10°C (18°F) past the maximum allowable temperature.

Conclusions

The addition of UPS support to data center air handlers turns seconds of ride-through time into minutes. At a fairly small cost above what you are spending to back up your IT equipment, you could also back up your air handlers. You should understand how long it takes to start your generator and restart your chiller. Use the data in this paper to calculate how quickly your facility might heat up and at what risk you might be. Considering the large energy differences associated with running very cold data center temperatures, it is not a recommended strategy to extend ride-through time. Tight containment might help extend ride-through time if you are unwilling or unable to put the CRAHs on UPS. Complacency could lead you to believe you have more time than required to get your data center running again, but this can be risky. Even if you have previously managed to get a facility back up and running in plenty of time, it is possible your density has progressed to the point that you no longer have as much time.

Appendix A: Graphic Anomalies Explained

Initial Temperature Dip

The initial temperature dip seen in Figure 2, and repeated here in Figure 6 for clarity, suggests that the temperature actually went down for a short period of time when the cooling was lost.

Figure 6. Initial Temperature Dip

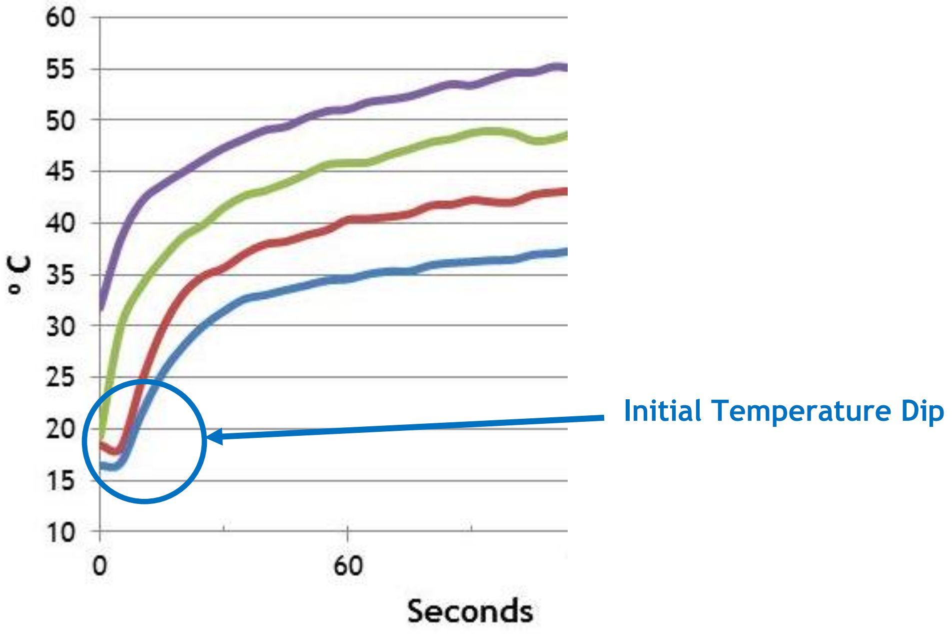

There are a couple of explanations for the dip or for a delay in the increase. To reiterate, this was a graph of the complete shutdown of cooling and airflow. To achieve this in testing, each of three CRAH units was manually shut down. The amount of time it takes to shut down the three CRAH units is several seconds walking from unit to unit, and then it takes some time for the unit to power down and for the fan blades to slowly come to a halt. To timestamp the data, as the first CRAH unit was being shut off, a spare thermocouple (one of twenty) was dipped into an ice bath. Data points were being logged every five seconds. Therefore, during the first official seconds of the test (when the ice bath is initiated), between one and five seconds elapse until the first reading is taken. Approximately three to five seconds elapse before all CRAH units are turned off. It takes four to five seconds for each CRAH unit to shut down, and then it takes five to ten seconds for the fan blades to stop turning. Although diminishing in quantity during this period, some cool air is still being driven by the blowers. When the unit is de-energized but still spinning down, the air temperature should actually drop. In a CRAH unit, the air is cooled down by the coil, drawn through the blowers, and then out the bottom of the unit (in a down-flow air handler). As it passes by the blower motor, it absorbs the same amount of power as the motor uses to turn the fan blades. This can result in a 1°C to 2°C (1.8F to 3.6°F) increase in exhaust temperature after the air has been cooled by the coil. During the spin down, it is no longer dissipating this heat. This is the likely reason for the dip seen in some of the test results. Even without the dip, some of the sideways movement in the graph could be due to timing issues between the data-point start with the ice bath and the pushing of the off button on the CRAH unit.

Temperature Bounce after Generator Start

Without backup power for the CRAH units, temperature increases were substantially higher. In our experiment, we manually turned the CRAH units back on after about three minutes to simulate the generator start up and reenergizing of the CRAH units. Temperatures dropped substantially and quickly, but they bounced back up for a period as seen in Figure 4 and Figure 5, and repeated in Figure 7.

Figure 7. Temperature Bounce

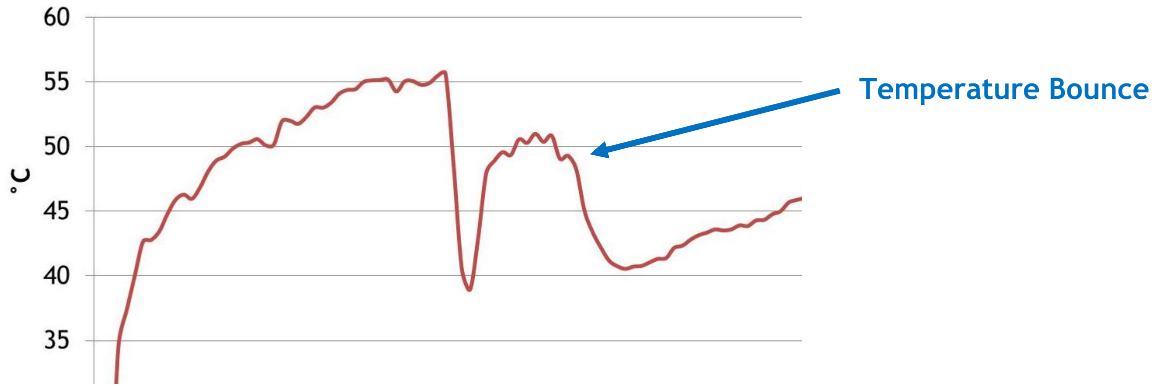

In our test scenario, the CRAH units were shut down in standard fashion, which caused the chilled water valves to slowly reduce flow down to completely closed. With a power failure, however, they would be stuck open without backup. When reenergized, the CRAH unit blowers ramp quickly to 100% flow which results in the initial steep drop in temperature. After reaching 100%, they look to their control algorithm. By default, they seek the same percentage of maximum speed as the chilled water valve is open. Since the chilled water valve is slowly creeping up from a completely closed state, the blowers are quickly adjusting to match the valve (but only down to their minimum setting of 60% flow). They do this much more quickly than the water valve, and the temperature climbs back up as the flow rate goes down. The temperature flattens out as the flow rate reaches 60%. As the chilled water valve passes through 60% on its way to fully open, the blowers follow.

In an actual failure, the water valve would have stopped partially or fully open. Upon reenergizing, the water valve has a relatively small distance to travel to get to fully open (and may already be there). If it were not fully open by the time the blowers start to seek their operating point, there would be a smaller bounce of shorter duration.

Appendix B: Additional Data

Alternate Room Density Tests

The graphs in Figure 8 and Figure 9 represent the other room densities tested during cooling failure.

Figure 8. 150 W/sq. ft. Test

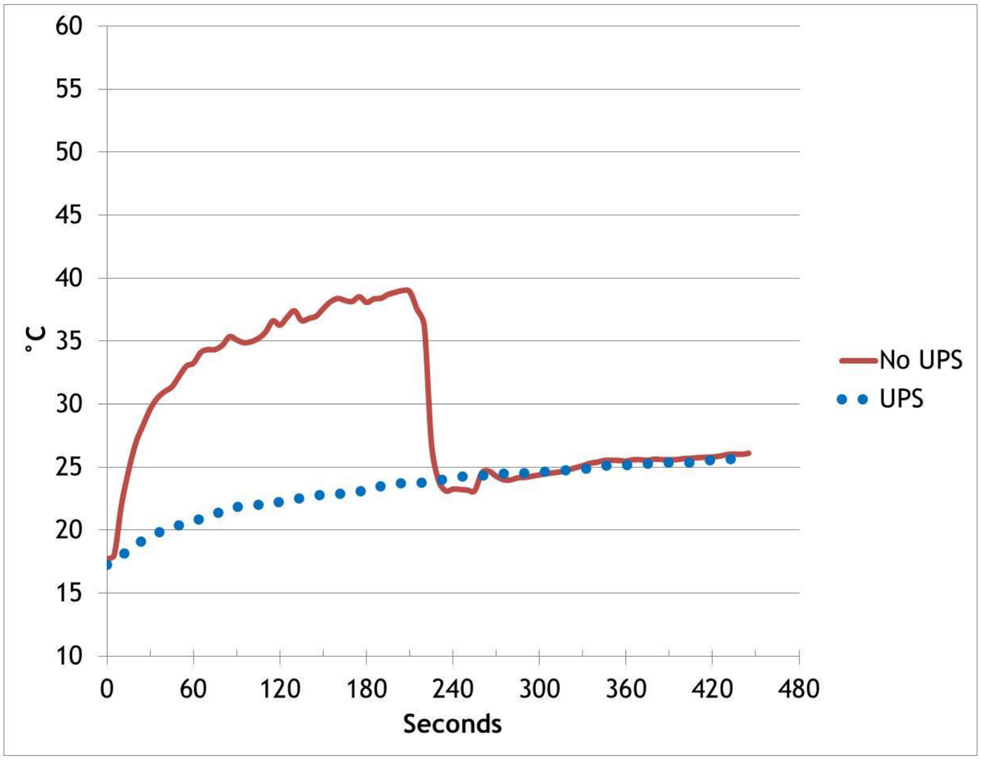

Figure 9. 250 W/sq. ft. Test