

Data Center Temperature Rise During a Cooling System Outage

White Paper 179

Revision 1

By Paul Lin
Simon Zhang
Jim VanGilder

> Executive summary

The data center architecture and its IT load significantly affect the amount of time available for continued IT operation following a loss of cooling. Some data center trends such as increasing power density, warmer supply temperatures, the “right-sizing” of cooling equipment, and the use of air containment may actually increase the rate at which data center temperatures rise. However, by placing critical cooling equipment on backup power, choosing equipment with shorter restart times, maintaining adequate reserve cooling capacity, and employing thermal storage, power outages can be managed in a predictable manner. This paper discusses the primary factors that affect transient temperature rise and provides practical strategies to manage cooling during power outages.

Introduction

IT equipment is typically backed up by uninterruptible power supplies (UPSs) which supply power until generators come on-line following the loss of utility power to the facility. However, cooling system components such as CRAC (Computer Room Air Conditioner) or CRAH (Computer Room Air Handler) fans, chilled water pumps, and chillers (and associated cooling towers or dry coolers) are typically not connected to UPSs and are generally connected to backup generators. Consequently, the data center supply air temperature may rise quickly following a power failure.

While much attention is devoted to data center cooling system design, most of that effort is aimed at improving the efficiency and reliability of its normal operation under utility power. The lack of attention paid to emergency operating conditions is due partly to a lack of simple tools for data center designers and operators to predict cooling performance under such conditions. However, a recently developed modeling tool¹, makes it easy to estimate data center temperatures following the loss of cooling for various facility architectures, back-up power connectivity choices, and, when applicable, chilled-water (thermal) storage volumes.

Meanwhile, planning for a power-outage is becoming more critical as data center professionals follow industry trends like “right-sizing” cooling capacity, increasing rack power density, adopting air containment, and increasing supply air temperatures. The latter trend is driven in part by ASHRAE’s recently-revised Thermal Guidelines² which allows warmer data center air temperatures than previously recommended. Without other compensating design choices, all of these trends reduce the window of time for safe and reliable operation following a power outage.

Factors affecting the rate of data center heating

Essentially all of the electrical power consumed in a data center is transformed into heat. Under normal, steady, operating conditions this heating rate is balanced by the rate of cooling. During a cooling failure, the heating rate is balanced instead by the rate at which heat is absorbed by the data center air, IT equipment, and the building envelope³. For facilities cooled by chilled-water CRAHs, water in the piping system and any storage tanks also absorbs heat assuming the chilled water can be circulated using backup power.

For a fixed IT load, the larger the volume of the data center, the slower the rate at which temperatures will rise during a cooling outage. However, as more IT equipment is added to a given facility, the air volume and building envelope plays a diminishing role in moderating the rate of heating and the thermal mass⁴ of the IT equipment itself becomes more important. Although perhaps unintuitive, even the “hot” insides of an operating server contributes thermal mass which slows the rate of data-center heating. (As the server inlet air heats up, the effectiveness at which the server can transfer heat to the internal airstream is reduced; this forces the balance of the server heat dissipation to be absorbed by its own solid components such as the chassis and motherboard.)

¹ Zhang, X. and VanGilder, J., 2011, “Real-Time Data Center Transient Analysis”, Proceedings of InterPACK, Portland, Oregon, July 6-8.

² ASHRAE. 2011, Thermal Guidelines for Data Processing Environments, Developed by ASHRAE Technical Committee 9.9.

³ Data center building envelopes are typically well insulated and play only a small role in moderating the rate of heating.

⁴ Thermal mass is the capacity of an object to store thermal energy. In a sense, the servers, walls, etc. act as thermal sponges during a cooling outage and return the heat after cooling is restored.

Chilled-water systems featuring traditional CRAHs or row-based coolers can offer advantages over fluid and air-cooled DX systems when facing a loss of primary power.⁵ First, with no compressor or other refrigeration-cycle components, CRAHs consume less power than DX units and are therefore easier to power by UPS and/or generator. Secondly, any chilled-water storage, if available, can be utilized immediately by CRAHs with only minimal backup power needed. By contrast, glycol-cooled and water-cooled DX systems require more backup power and cannot take advantage of the thermal storage until the DX CRAC (Direct Expansion Computer Room Air Conditioner) units can be restarted which may take several minutes. Air-cooled CRAC DX units offer the least emergency-performance potential as the entire system, including dry coolers (which transfer the heat to the outside air) must be supplied with backup power and, again, may require a lengthy re-start period.

To illustrate the heating of a data center following the loss of utility power, we consider examples based on two different predictive models:

- Well-mixed-air model
- Detailed computational fluid dynamics (CFD) model

The well-mixed-air model idealizes the data center as a single CRAH or CRAC and a single IT load with perfectly well-mixed air. So, at any given time, there is only one uniform data center air temperature, one plenum temperature, one cooler supply temperature, etc. The model is only strictly correct in the aggregate – not precise for any particular location – but serves as an adequate representation of IT inlet temperatures for the purposes in this paper. The simple well-mixed model includes design parameters such as chilled-water thermal storage as well as UPS and generator power connectivity for CRAH fans, chilled water pumps, and the chiller plant. The speed and simplicity of this model makes it easy to study and visualize the results of a variety of scenarios quickly. Because we are presently most concerned with understanding the important physics related to and developing design strategies for power-outage scenarios in general (rather than designing any one specific facility), the well-mixed model is the primary tool we use here.

Detailed CFD simulations, on the other hand, require a skilled analyst, and generally take many hours of computational time to produce a useful result. Furthermore, time-varying or transient data center predictions are particularly complex and are, therefore, not routinely performed. However, in return for the considerable complexity, CFD can provide detailed spatial variations in temperature over the course of time following a power outage.

Well-mixed-air model example

Figure 1 shows the response of a typical legacy data center (see sidebar) to the loss of primary power using the well-mixed model. Immediately following the power failure, the room air temperature rises quickly while the plenum temperature remains fairly constant as the CRAH units remain “off”. The generator starts after 1 minute and provides power to the CRAH fans and chilled water pumps. The room temperature drops then rises again – as does the supply and plenum temperatures - since the chilled water in the piping system warms. At 11 minutes, the chiller starts cooling the chilled water and all temperatures trend back down to their normal-operation values. For this example, the room air temperature only marginally exceeds the allowable threshold for about 3 minutes and then remains unacceptably high until about 17 minutes after the power outage.

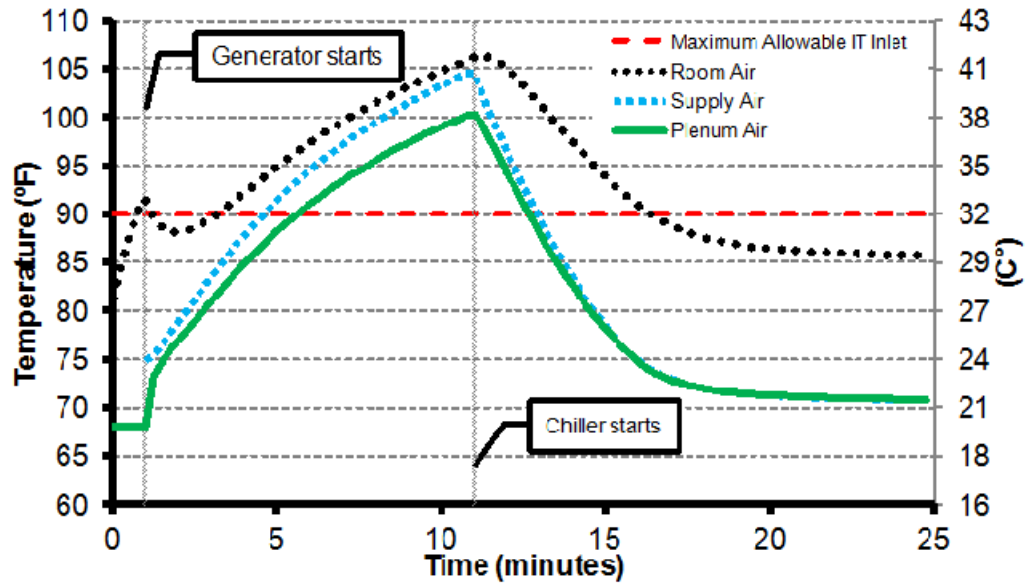
> Assumptions for Example 1

- Well-mixed data center air
- Data center dimensions: 30m (100ft) length x 24m (80ft) width x 4m (12ft) height with 600 mm (24 in) plenum
- Raised-floor cooling, chilled-water CRAHs, no air containment
- Chilled water supply temperature set-point: 7.2°C (45°F)
- Chilled water ΔT : 8.3°C (15°F)
- CRAH airflow: 45 m³/s (96,000 cfm)
- CRAH supply set-point: 20°C (68°F)
- No supplemental thermal storage
- IT load: 400 kW
- Average rack power density: 4.0 kW
- Average rack mass: 545 kg (1200 lb)
- UPS backup for IT equipment only
- All IT and cooling equipment connected to generator
- Generator start time: 60 s
- Chiller restart time: 10 min
- No heat transfer through or thermal mass of walls

⁵ For information on the different types of cooling systems, see [White Paper 59, The Different Technologies for Cooling Data Centers](#)

Figure 1

Bulk air temperature variation following power outage for example 1



Detailed Computational Fluid Dynamics (CFD) model example

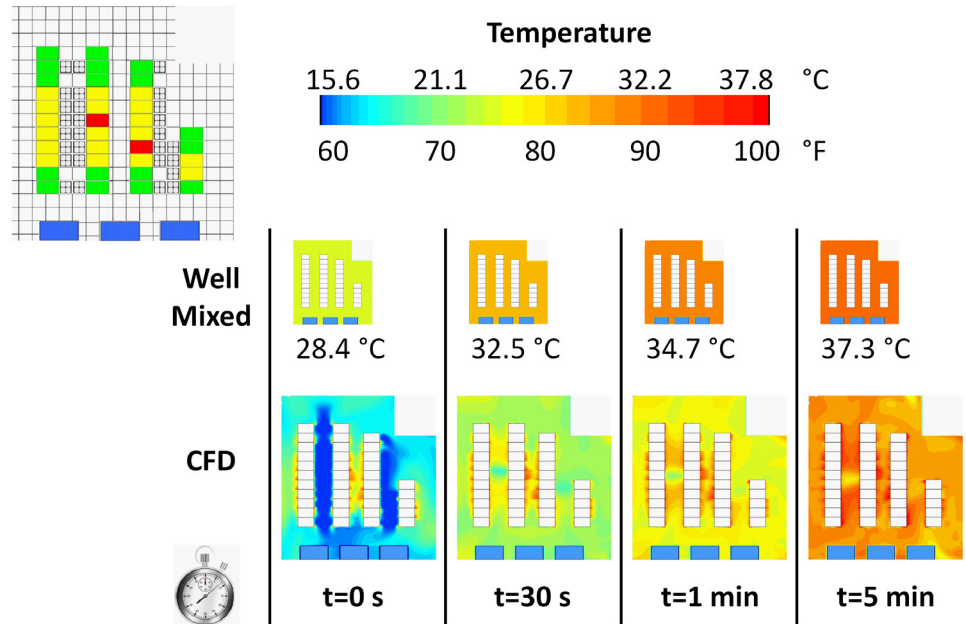
Figure 2 illustrates the response of a different data center (see sidebar) to a loss of cooling power as modeled with CFD. For simplicity, we model only the loss of cooling and assume that, over the period of the simulation, no cooling is restored by backup power or otherwise. For comparison, we show corresponding well-mixed-model results. It is clear from the figure that temperatures become unacceptably hot everywhere after about 5 min. Further, the “cold” aisles quickly become hotter than the “hot” aisles as the racks re-circulate their own airflow. Finally, while the CFD simulation shows substantial spatial temperature variations, the well-mixed model does an effective job at capturing the overall aggregate room temperature variations over time.

Figure 2

Detailed air temperature variation following power outage for example 2

> Assumptions for Figure 2

- Data center dimensions 10m (33ft) length 10.7m (35ft) width 2.6m (8.5ft) height with 610mm (24 in) plenum
- Raised-floor cooling, chilled-water CRAHs, no air containment
- CRAH airflow: 2.1 m³/s (4500cfm) each
- CRAH supply set-point: 20°C (68°F)
- IT load: 400 kW
- Backup for IT Equipment only
- Heat transfer through and thermal mass of walls included
- Average rack mass: 545 kg (1200 lb)
- Rack airflow: 125 cfm/kW
- Rack Power:
 - 1 kW
 - 2.5 kW
 - 6 kW
- Floor Tiles
 - Solid
 - ▣ Perforated



Industry trends negatively impacting emergency cooling performance

Some data center trends and best practices -most aimed at improving performance, efficiency, and manageability under normal operating conditions - may adversely impact operating conditions following a power outage. These trends and practices include:

- Right-sizing cooling capacity
- Increasing power density and virtualization
- Increasing IT inlet and chiller set-point temperatures
- Air containment of racks and rows

Right-sizing cooling capacity

Right-sizing (i.e., aligning capacity to the actual IT load) the capacity of the overall cooling system provides several benefits including increased energy efficiency and lower capital costs. (See White Paper 114, [Implementing Energy Efficient Data Centers](#) for more information.) However, excess cooling capacity is desirable when faced with unacceptably high temperatures following a power outage. In fact, if the total cooling capacity perfectly matched the heat load, the facility theoretically could never be cooled to its original state because after a power outage there would always be heat in excess of the IT load. Just as multiple window air-conditioners effectively cool a bedroom more quickly than a single unit, additional CRAH or CRAC capacity helps return the data center to pre-power-failure conditions quickly. Note that for all architectures, the cooling distribution (airflow) must be such that CRAH or CRAC capacity can actually be utilized (i.e. by use of blanking panels, brush strip, hot / cold aisles, etc.).

Increasing power density and virtualization

Compaction of IT equipment produces increased rack power densities in the data center. The emergence of equipment like blade servers and certain communications equipment can result in rack power densities exceeding 40kW/rack.

Another technology trend, virtualization, has greatly increased the ability to utilize and scale compute power. For example, virtualization can increase the CPU utilization of a typical non-virtualized server from 5%-10% to 50% or higher. A detailed discussion of the effects of virtualization on the physical infrastructure can be found in White Paper 118, [Virtualization and Cloud Computing: Optimized Power, Cooling, and Management Maximize Benefits](#).

Since both increasing the rack power density and virtualization make it possible to dissipate more heat in a given space, they can also reduce the time available to data center operators before the IT inlet temperatures reach critical levels following a power outage.

Increasing IT inlet and chiller set point temperatures

ASHRAE Technical Committee 9.9 (Mission Critical Facilities, Technology Spaces and Electronic Equipment) developed and expanded the recommended thermal operating envelope for data centers. Increasing the IT inlet and chilled water set point temperature results in an increased number of hours that cooling systems can operate on economizer mode.

It has been estimated that for every 1.8°F (1°C) increase in chiller set point temperature, about 3.5% of the chiller power can be saved⁶. In other words, it gets increasingly expensive

⁶ ASHRAE, 2005, Design Considerations for Datacom Equipment Centers. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.


to cool chilled water the more the set point temperature is reduced below a fixed ambient temperature. (While this applies directly to chilled-water systems, the same trend applies to air-cooled DX systems.) This fact puts pressure on data managers to maintain temperatures at as high a level as possible under normal operating conditions. Consequently, higher IT inlet temperatures leave less time for data center operators to react in a power-failure scenario.

Air containment of racks and rows

Containment can improve the predictability and efficiency of traditional data center cooling systems such as perimeter cooling systems with raised floor or hard floor (i.e. flooded supply). However, containment systems prevent air streams from mixing with the rest of the data center room and this will affect the temperature rise during cooling outages. The temperature rise performance will vary for different containment systems depending on the connectivity of cooling equipment to backup power. For more information about the advantages and deployment of containment, refer to White Paper 135, [Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency](#) and White paper 153, [Implementing Hot and Cold Air Containment in Existing Data Centers](#). **Figure 3a and 3b** show examples of containment with room-based coolers and row-based coolers respectively.

Figure 3a

Containment of rows with room-based coolers



Cold aisle containment system



Ducted hot aisle containment system

Figure 3b

Containment of rows and racks with row-based coolers



Row-cooled hot aisle containment system



Rack air containment system

For hot-aisle containment with row-based chilled-water coolers, if the coolers are not on UPS and containment doors remain shut during a loss of cooling airflow, then there could be a substantial amount of re-circulated hot air into the IT inlets through various leakage paths and IT inlet temperatures will rise quickly. If coolers are on UPS, but the chilled water pumps are not on UPS, then the coolers will pump hot air into the cold aisle without providing active cooling. In this case, only the thermal mass of the cooler (cooling coils, water inside the coil etc.) is utilized. If both coolers and chilled water pumps are on UPS, then the temperature rise depends on the chilled water plant configuration (i.e., storage tank configuration, chiller start time, etc.).

For cold-aisle containment with row-based chilled-water coolers, if the coolers are not on UPS, then the negative pressure in the containment system will draw in hot exhaust through the rack and containment structure leakage paths, thereby raising IT inlet temperatures. If row-based coolers are on UPS, then the temperature rise depends on the chilled water plant configuration (i.e. storage tank configuration, chiller start time, etc.).

Rack-air containment systems, behave similarly to cold-aisle and hot-aisle containment with row-based coolers.

However, for hot-aisle containment with drop ceiling and perimeter chilled-water coolers, or cold-aisle containment with perforated tiles and perimeter chilled-water coolers, whether the coolers and chilled water pumps are on UPS or not, containment has a positive impact on emergency cooling performance because containment provides a cold thermal mass in the open room (cold aisle) or raised-floor plenum associated with the concrete slab, chilled water pipes, etc., to moderate the temperature rise⁷. This is due to the separation of hot and cold air preventing the air streams from mixing, at least initially after the cooling outage.

Strategies for slowing the rate of heating

Despite the challenges provided by recent data center trends, it is possible to design the cooling system for any facility to allow for long runtimes on emergency power. Depending on the mission of the facility, it may be more practical to maximize runtimes within the limits of the current architecture and, at the same time, plan to ultimately power down IT equipment during an extended outage. This section describes four strategies to slow the rate of heating.

- Maintain adequate reserve cooling capacity
- Connect cooling equipment to backup power
- Use equipment with shorter restart times
- Use thermal storage to ride out chiller-restart time

Maintain adequate reserve cooling capacity

As discussed above, the industry trend of “right-sizing” cooling makes sense for normal operating conditions, but having even marginally more cooling capacity than the load can greatly increase the amount of time required to cool a facility that has already become too hot. The key to increased cooling system efficiency is to scale the bulk cooling (i.e. chillers) and cooling distribution (i.e. CRAH units) as the IT load increases. This allows for adequate reserve cooling while improving data center efficiency. For example, a data center designed for a maximum IT load of 1MW may only have a day-one IT load of 100kW. While the chilled water plant piping is sized for the maximum data center load, the installed chillers may support only about 250kW of total load, or about 140kW of IT load. The actual oversizing is dependent on the redundancy requirements and component efficiencies.

Connect cooling equipment to backup power

Referring back to **Figure 1**, the first temperature spike occurs because the CRAH fans and chilled water pumps cannot function until the generator picks up the load one minute after the power outage occurs. The biggest driver for this steep temperature spike is the ratio of IT power to air volume. Immediately after the cooling failure, but before the facility’s thermal mass (i.e., walls, plenums, servers, etc.) can absorb any significant heat, all of the IT power will simply heat the air. Without any cooling, the maximum rate of temperature rise could easily be 5°C/minute (9°F/minute) or more depending on density and room layout. Unless CRAH fans and pumps are on UPS and/or the data center is very lightly loaded, the initial temperature spike will almost certainly violate the temperature gradient values specified by tape vendors or those stated in the ASHRAE Thermal Guidelines.

⁷ White Paper: Data Center 2020: hot aisle and cold aisle containment efficiencies reveal no significant differences. November 2011. Powered by T System and Intel.

In a lightly-loaded facility (i.e. 20% load), placing only CRAH or CRAC fans on UPS until the generator starts, helps maintain proper cooling airflow, limiting recirculation from the IT exhaust to the IT inlet and helps to transfer heat to the pre-cooled thermal mass in the facility. Placing pumps (in addition to CRAH or CRAC fans) on UPS is more effective at reducing the initial temperature spike before the generator starts, particularly in systems utilizing chilled-water CRAH units. In this case, the thermal mass of the chilled water and the piping system alone can significantly extend the window of operation following a power failure. For glycol-cooled DX systems without a free-cooling coil, pumps generally do not benefit from UPS backup because generator power is required to restart the CRAC.

If the chiller plant is located far from the data center or the chilled water system uses double ring pipe loops (used for redundancy and high availability), there can be a large volume of chilled water in the pipes. If the data center is positioned in a large multi-tenant building, a shared chilled water plant would likely be used for the data center which would also provide a large amount of cooling capacity. Note that data center designers and operators should communicate with the facilities personnel to ensure that the data center has the first priority to use chilled water for emergency conditions.

For the two practices above, depending on the type of fan, pump, and backup configuration, a separate and independent UPS may be required to avoid interference with IT equipment. If the fans, pumps, and IT equipment use the same UPS, an isolation transformer may be used for the mechanical load.

Use equipment with shorter restart times

The chiller control system is typically able to ride through a power outage which lasts less than a quarter of a cycle (5 ms for a 50 Hz system and 4 ms for a 60Hz system). Power outages longer than this typically require a restart when power is restored (from the utility or generator). This restart time is typically 10-15 minutes (see sidebar). Advancements in chiller technology have reduced the restart time to 4-5 minutes, a 60% reduction in some cases. A fast chiller restart is not only important for the initial power outage, but also for the momentary brownout (100 ms to 1s) when the ATS (automatic transfer switch) transfers the power supply from generator back to the utility power.

Referring again to **Figure 1**, the second temperature peak occurs because the chiller needs 10 minutes to restart and pick up the cooling load. However, if the restart time had been 5 minutes, data center air temperatures would have only climbed marginally above the acceptable limit of 90°F instead of exceeding 105°F.

Higher-cost quick-start chillers may not be sufficient to prevent unacceptable temperature rise in high-density data centers. However, quick-start chillers are helpful in virtually all cases and, in lower density data centers, may keep temperatures entirely within acceptable limits for the duration of the power outage. Further, with quick start chillers, chilled water and IT temperatures may be maintained at higher levels during normal operating conditions with less risk of exceeding acceptable temperatures under emergency conditions. A balance between first cost and operational cost should be found between the chiller type and the importance of emergency operation after less-expensive options have been investigated.

Use thermal storage to ride out chiller-restart time

For chilled-water systems, additional chilled-water storage can be utilized until the chiller is restarted. If chilled water pumps and CRAH fans are on UPS in a chilled-water system, cooling can be provided with very little departure from normal operating conditions provided the storage tank is adequately sized.

> Chiller restart time

Chiller restart time is the time period from when the chiller receives power after undergoing a power outage to the moment it can provide its rated cooling capacity. All chillers have a program to protect themselves from being damaged under abnormal conditions.

Once the power returns, the chiller checks the status of components like the control board, compressor, oil system, and water system. If the system check results are normal, then the chillers will start up and re-establish the chilled water temperature as soon as possible. This total process takes several minutes, and is dependent on the type of chiller. Some factors that influence chiller restart time include chiller configuration, lubrication oil pressure, controller reboot time, ability to manage diagnostics, and type of power loss.

Low-pressure thermal storage tanks for chilled water systems have a much lower initial cost than placing chillers on a UPS and can even be made of plastic. The volume and type of tank depends on various factors including space constraints and load-bearing capacity if installed on a roof or an elevated floor. Their use is particularly recommended for high-density data centers when even a very short period without cooling can be problematic.

Thermal storage requirements should also consider the temperature stratification inside the tank. For large diameter tanks, the height of the mixed layer can be reduced by using a water distributor which controls the speed at which the warm return water enters the tank. Additionally, the piping and controls should be configured so that the storage tank can be bypassed after the chiller is restarted. This prioritizes the delivery of the coldest water immediately to the data center over the re-cooling of the storage tank water.

Comparison of strategies for slowing rate of heating

We now return to the data center of Example 1 and consider the relative merits of the strategies for slowing the rate of heating following a loss of primary power, again, using the simple “well-mixed-air” model. Referring to **Figure 4**, the “baseline” plot shows the same room-air-temperature curve as that shown in **Figure 1** which assumes that CRAH fans and chilled-water pumps are connected only to generator. By also connecting the CRAH fans to UPS, there is a modest improvement in the initial period before the generator starts as the additional thermal mass of the pre-cooled plenum can be accessed. (Note that the well-mixed model may under-predict the value of placing CRAH or CRAC fans on UPS as it does not include the additional benefit of proper airflow management which can keep hot air from easily re-circulating into rack inlets.)

If both CRAH fans and chilled-water pumps are placed on UPS, the initial temperature spike is eliminated as the thermal mass in the piping system is immediately accessible. By reducing chiller restart time from 10 minutes (11 minutes from the start of the outage) to 5 minutes (6 minutes from the start of the outage), acceptable temperatures are reestablished much more quickly and the maximum room air temperature is reduced from 106°F to 98°F. Adding thermal mass storage by itself does nothing in the initial period before the generator starts since CRAH fans and chilled-water pumps are not on UPS. However, once this equipment comes back online, the chilled-water storage effectively keeps data center temperatures close to acceptable limits until the (standard) chiller can be restarted. Finally, if all of the above strategies are implemented simultaneously, there is only one modest temperature ramp before the (fast-start) chiller restarts and temperatures never exceed acceptable limits.

Figure 4

Room Air temperature variation following power outage for data center of Example 1 after implementing strategies for slowing rate of heating

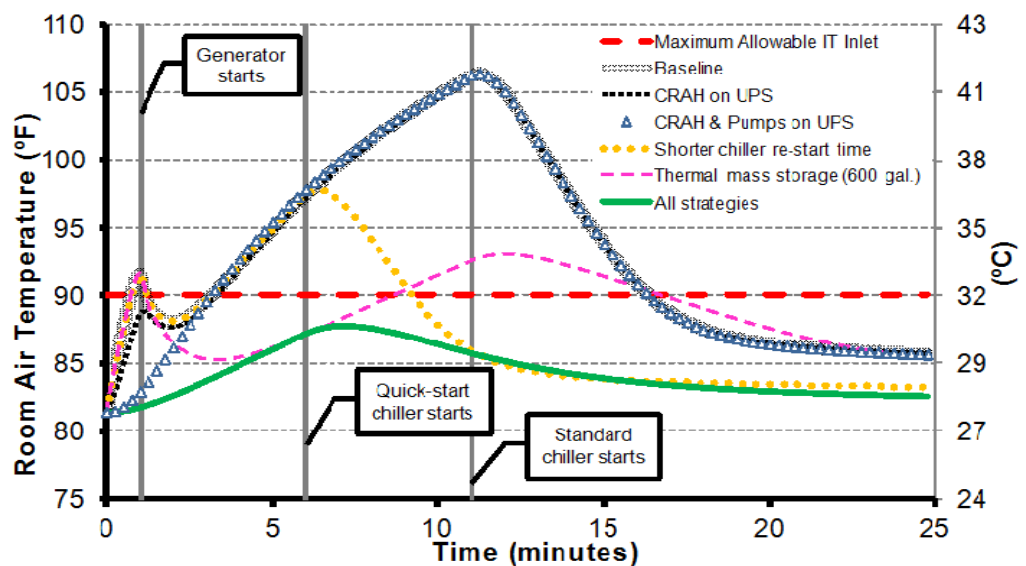


Table 1 compares the four basic strategies discussed above in the context of our example data center. For chilled-water CRAH systems, the best option may be to simply first ensure that CRAH fans and chilled water pumps are connected to generator (“baseline” in **Figure 4**) and then add thermal storage (Strategy 4) to ride out the chiller re-start time. For high-density facilities (with a steep initial temperature rise), Strategy 2, placing CRAH fans and pumps on UPS, may be necessary to avoid unacceptably-high temperatures before the generator starts. The use of a fast-start chiller (Strategy 3) may make sense for a new facility but other options are more economical when attempting to improve the emergency response of an existing facility.

For DX CRAC systems, the first step is again to ensure that all components are connected to generator. With air-cooled, glycol-cooled, and water-cooled units, it may be possible to place only the CRAC fans on UPS (Strategy 1) to realize some benefit in the early period before the generator starts. (Note, however, that for some CRAC units, placing fans on UPS may adversely impact the DX system start-up time when power is restored.) For glycol-cooled or water-cooled DX units without a free-cooling coil, there is no additional benefit to placing pumps on the UPS because the cooling fluid cannot be utilized until the CRAC units are restarted. For glycol-cooled or water-cooled DX units with a separate free-cooling coil, pumps and fans on UPS (Strategy 2) is beneficial. Of course, it is also possible to power entire DX CRACs by UPS; however, the large units required would be expensive and inefficient under normal operating conditions.

Some DX CRAC systems include a “multi-cool” option that allows for redundant cooling with the addition of a chilled water coil. Cooling is achieved with both the internal compressor (via the DX coil) and an external chiller (via the chilled water coil). With these systems it is more effective to place CRAC fans and chilled water pumps on the UPS. Provided CRACs can be restarted substantially faster than the chiller, thermal storage (Strategy 4) is also beneficial.

Table 1

Comparison of strategies for slowing rate of heating

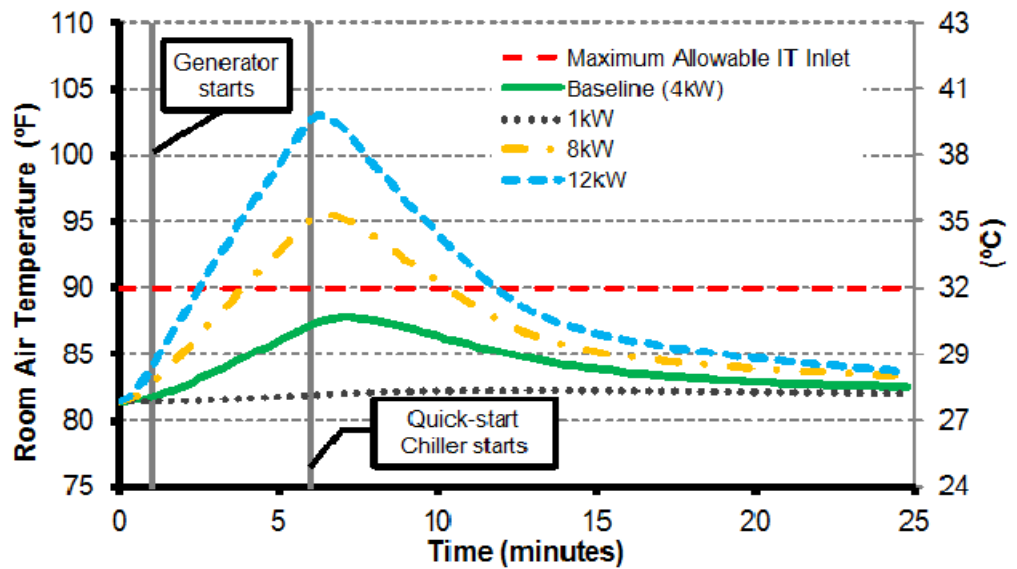
Solutions		Goal	Assessment by Application
1	Place CRAH or CRAC fans on UPS	Reduce initial temperature spike before generator starts	<ul style="list-style-type: none"> Only strategy in table applicable to air-cooled DX A reasonable strategy for water-cooled and glycol-cooled DX Better to implement #2 below for chilled-water CRAH applications.
2	Place CRAH or CRAC fans and chilled water pumps on UPS	Eliminate initial temperature spike before generator starts	<ul style="list-style-type: none"> No additional benefit over #1 above for water-cooled and glycol-cooled DX applications Better than #1 for chilled-water CRAH applications and “multi-cool” DX CRAC systems
3	Use fast-start chiller	Reduce secondary temperature spike after generator but before chiller starts	<ul style="list-style-type: none"> Effective for chilled-water CRAH and “multi-cool” DX CRAC systems utilizing chillers
4	Use chilled-water storage	Reduce secondary temperature spike after generator but before chiller starts	<ul style="list-style-type: none"> Effective for chilled-water CRAH and “multi-cool” DX CRAC systems utilizing chillers For chilled-water CRAH systems, even more effective when used with #2

Effect of power density on thermal storage capacity

As discussed previously, increasing rack power density can reduce the time available to data center operators before the IT inlet temperatures reach critical levels following a power outage. **Figure 5** shows the effect of rack power density on the data center room air temperature following the loss of primary power using the well-mixed model (see assumptions in side bar). Note that the 8kW/rack and 12kW/rack scenarios exceed the maximum allowable IT inlet temperature because the IT rack quantity and thermal storage capacity are held constant. This means that, as rack density increases, the IT load increase. Increased IT load is addressed by increasing the thermal storage capacity.

Figure 5

Room air temperature variation following power outage for varying rack power densities



> Assumptions for Figure 5

- Data center of **Figure 4** with all strategies implemented (CRAH and pumps on UPS, thermal mass storage connected)
- Server/rack IT population (thermal mass) the same for all scenarios
- CRAH and chiller capacities scale up with the IT power
- Thermal storage capacity fixed at 600 gallons (2.3 m³)
- IT load increases with density - 400kW, 800kW, and 1.2MW

Table 2 provides the thermal storage capacity required to not exceed ASHRAE allowable of 90°F (32°C) for four different rack densities with both chiller types. The table assumes the piping size is same for all scenarios, however, in practice the piping diameter would increase with IT load which slightly decreases the required chilled water storage capacity.

Table 2

Thermal storage capacity required to not exceed ASHRAE allowable IT inlet temperature of 90°F (32°C)

Rack density	Total IT load	CRAH & chiller capacity	Standard chiller (12 minutes)	Quick-start chiller (5 minutes)
1kW/rack	100 kW	150 kW	0 gal (0 m ³)	0 gal (0 m ³)
4kW/rack	400 kW	600 kW	1,300 gal (4.9 m ³)	350 gal (1.3 m ³)
8kW/rack	800 kW	1,200 kW	3,800 gal (14.4 m ³)	1,300 gal (4.9 m ³)
12kW/rack	1,200 kW	1,800 kW	7,000 gal (26.5 m ³)	2,500 gal (9.5 m ³)

Conclusion

Today, with new data center design trends including increasing power density, warmer supply temperatures, right-sizing of cooling equipment, and the use of containment, data center temperatures can rise very quickly during a cooling outage. Predictive models and design strategies make it possible to ensure continued reliable operation of, or ample time to power-down, IT equipment following a power outage.



About the authors

Paul Lin is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operation research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environment. Before joining Schneider Electric, Paul worked as the R&D Project Leader in LG Electronics for several years. He is now designated as a "Data Center Certified Associate", an internationally recognized validation of the knowledge and skills required for a data center professional. He is also a registered HVAC professional engineer. Paul holds a master's degree in mechanical engineering from Jilin University with a background in HVAC and Thermodynamic Engineering.

Simon Zhang is a Sr. Research Engineer with APC by Schneider Electric working on data center design, operation, and management software platforms. He has extensive experience with real-time cooling predictions & indoor airflow simulations, and has author/co-authored 12 patents (granted or pending) and over a dozen peer-reviewed technical papers on data center cooling and energy assessment techniques. He is actively involved in data center communities and has chaired and organized many technical sessions of ASME and IEEE conferences. He received his M.S. in Mechanical Engineering at Syracuse University in 2006, and is finishing an MBA degree from Boston University in 2012.

Jim VanGilder is responsible for Schneider Electric's data-center-cooling software encompassing both software development and related research. He has authored or co-authored more than 30 published papers and holds more than 10 patents in the field of data center cooling. He is currently the Chair of the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Technical Committee 4.10, Indoor Environmental Modeling, and is a registered professional engineer in the state of Massachusetts. He has master's degree in mechanical engineering from Duke University with a background in Computational Fluid Dynamics (CFD).



Resources



[The Different Technologies for Cooling Data Centers](#)

White Paper 59



[Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency](#)

White Paper 135



[Implementing Energy Efficient Data Centers](#)

White Paper 114



[Virtualization and Cloud Computing: Optimized Power, Cooling, and Management Maximizes Benefits](#)

White Paper 118



[Browse all white papers](#)

whitepapers.apc.com



[Browse all 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 at
www.apc.com/support/contact/index.cfm