White Paper

Intel Information Technology

Computer Manufacturing

Thermal Management

Thermal Storage System Provides Emergency Data Center Cooling

Intel IT implemented a low-cost thermal storage system that maintained cooling at a high-density data center during an electrical power outage. This enabled the data center to survive the outage without costly damage to servers. The system is based on auxiliary thermal storage tanks that feed water into the chilled water supply lines if the main chillers stop working due to an outage. It prevented the servers, which were still running on an uninterruptible power supply (UPS), from overheating by maintaining a chilled water supply to the air handler cooling coils.

Doug Garday and Jens Housley, Intel Corporation

September 2007

IT@Intel

Executive Summary

Intel IT implemented an emergency thermal storage system that enabled a high-density data center to survive a power outage without costly thermal damage to servers. The system, based on auxiliary chilled water storage tanks, kept the data center cool when an outage caused the chillers to shut down.

The system provided a method for surviving a rare power outage that was simple, reliable, and low cost, compared with alternatives such as putting the chillers on a continuous power system.

In data centers with high power and heat densities, a power sag or outage can cause rapid temperature increases. This is because cooling systems temporarily shut down, while servers keep producing heat because they are on an uninterruptible power supply (UPS).

To overcome this problem, Intel IT implemented a thermal storage system at a large regional hub data center. The system, based on two 24,000-gallon tanks containing chilled water at 42° Fahrenheit (F), operated successfully during an outage in late 2006 that lasted several hours.

- When the chillers stopped working, the system added water from the tanks to the chilled water system to maintain cool data center temperatures.
- Chilled water system pumps and air handler fans continued operating because they were on UPS and generator backup power.
- Servers continued to operate for more than 15 minutes due to the light load on the data center at the time of the utility power outage.
- The thermal reserve maintained cooling during this period and for long enough afterwards to allow dissipation of residual heat from the servers.

The system provided a method for surviving a rare power outage that was simple, reliable, and low cost, compared with alternatives such as putting the chillers on a continuous power system (CPS).

Contents

Business Challenge

As data center power and heat density increase, transient losses of electrical utility power potentially have larger impacts and require special design considerations.

A power sag or complete outage can cause cooling systems to temporarily shut down. Meanwhile, servers and other IT equipment keep operating and producing heat for several minutes because they are on UPS.

As a result, data center temperatures can rise rapidly, potentially causing severe damage to IT equipment. Our calculations show that if cooling is interrupted, a high density data center may only take 18 seconds to reach 105° F and 35 seconds to reach 135° F. This is because the power required for these high-performance servers, combined with the high server density, results in considerable concentrated heat, especially when airflow stops.

The response of mechanical cooling equipment to power disturbances varies, depending on the size and duration of the disturbance. Based on our experience, a major disturbance lasting less than a second can cause chillers to shut down, and power sags below about 85 percent of nominal voltage typically are enough to cause problems. Once chillers shut down, it can take several minutes for them to resume cooling after power is restored—too long to prevent damage to IT equipment.

When a data center loses its cooling system, IT equipment on UPS continues to use its cooling fans to draw in air from the data center space in order to cool the hot core of the equipment. The fans transfer this heat to the data center space.

IT equipment also continues to generate heat even after it shuts down. A fully populated cabinet of servers may weigh approximately 2,000 pounds. Most of the mass is metal, with a heat gradient between the hot cores of the IT equipment and the outside of the cabinet at approximately room temperature. This mass stores heat; even after the equipment shuts down, it continues to release heat to its surroundings until it reaches an equilibrium temperature. Solutions to cooling problems caused by power loss need to deal with this residual heat.

Technology Options

There are several methods for increasing the resilience of data center cooling systems to power disturbances. Some data centers requiring very high availability use standby generators for chillers. However, these add significantly to data center cost. Also, generators take several seconds to start up, after which it can take several minutes to restart the chillers.

These delays may be acceptable in low-density data centers because temperature increases more slowly after an outage. In high-density data centers, however, delays of even a few seconds cause problems due to the rapid rise in temperature.

To create a CPS capable of delivering power indefinitely without interruption, generators can be combined with technologies such as battery UPS or other types of energy storage systems. These can ride through short power losses or maintain power until standby generators come online. However they also add to cost and complexity, and the resulting system can be less reliable if these systems fail. As the number of UPS systems and generators increases, there is a greater possibility that one of them will not work correctly when needed.

Thermal storage methods offer an alternative to these approaches, providing varying degrees of resilience to power failures at much lower cost, with less complexity.

Thermal Storage Methods

Thermal storage can extend the ability to cool data center IT equipment in the event of a power failure by using thermal reserves to provide temporary cooling during a power outage or sag.

A data center contains various thermal masses at different temperatures. Recirculation air handler (RAH) coils, supply air ducts and raised metal floor (RMF), and much of the air in the data center are masses at a low enough temperature to absorb heat. These are potential thermal reserves. IT equipment and return air paths are heat sources; if airflow stops, heat tends to flow out of these heat sources and into the lower-temperature masses.

To take advantage of this, our data center designs aim to maximize low-temperature masses and minimize heat sources. In Intel data centers, we use enclosures to separate hot aisles from cold aisles, so that all cold aisles are low-temperature air masses. The enclosures maintain temperature gradients and enable control of the hot air path. Hot air leaving the servers travels directly into a hot air return plenum above the drop ceiling. By minimizing the volume of this return air path, the amount of hot air in the system is also minimized when the cooling system is disrupted.

UPS or CPS for Air Handler Fans and Chilled Water Pumps

One simple thermal storage method is to use the cold air already in the air conditioning system. We do this by putting a computer room air conditioning (CRAC) unit or RAH fans on UPS or CPS. This represents a relatively small additional electrical load compared with putting the entire cooling plant on a backup power source. It helps ensure that airflow continues during an electrical outage, providing more time before the data center reaches critical temperature.

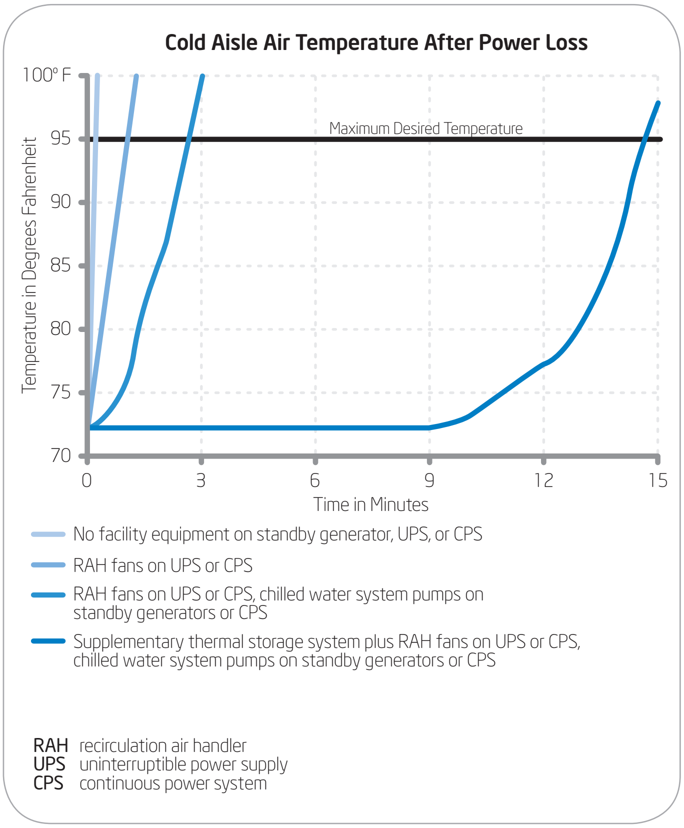

Figure 1. The rate at which data center temperature rises when a utility outage causes a cooling interruption. A supplementary thermal storage system greatly extends the time before a high-density data center exceeds maximum desired operating temperature when a utility outage causes a cooling interruption. Intel analysis based on a 310-watts-per-square-foot high-density data center operating at 72° Fahrenheit when cooling is lost.

However, the cold air in the system can only provide cooling for a limited duration—as little as a few seconds. A much greater cold mass is the cold water in the chilled water mains and branch piping leading to the RAH cooling coils. The chilled water piping can contain thousands of gallons of cold water. If effectively circulated, this can act as a thermal reserve system, providing minutes of precious cooling. This can be exploited by putting the chilled water pumps on a backup power source, such as a standby generator, UPS, or combination of the two to create a CPS, so that water flow continues to cool the data center temporarily even if the chilled water plant shuts down.

These techniques can significantly slow temperature increases, maintaining temperatures within the range required for IT equipment for minutes rather than seconds, as shown in Figure 1.

Auxiliary Cold Water Storage Tanks

If the chilled water piping does not provide enough thermal storage to provide cooling during a loss of power, auxiliary cold-water storage tanks can significantly increase a data center's thermal reserves.

When chillers stop due to a power loss, water from the tanks can supplement the chilled water supply to keep it cold enough to maintain the data center environment near normal operating temperatures.

Cold water tanks have a much lower initial cost than approaches such as putting chillers on UPS and generators, and also have lower operating costs.

Factors that determine the tank capacity required include the data center power and heat density. This is important because it influences the rate at which the temperature rises when cooling is lost.

Cooling Systems

The type of cooling system affects uptime in the event of power loss. For example, different chiller types vary in the speed with which they can restart after an outage. This is a factor in determining the optimum approach to cooling a data center.

Centrifugal and Scroll Chillers

Chilled water systems using centrifugal chillers provide one of the most economical cooling systems, in terms of kilowatts of power used per ton of cooling (kW/Ton). This is an advantage during normal operation. However, centrifugal chillers may not be able to restart themselves as quickly as desired. When power sags of approximately one second or more occur, components making up a chilled water plant can shut down. It may take three to six minutes to restart a centrifugal chiller, or longer in special circumstances.

Another option is to use scroll chillers. These chillers are typically not as efficient as centrifugal chillers, in terms of kW/Ton. However, they offer a quick restart time, and due to their smaller capacity of between 100 and about 350 tons, represent less connected load on CPS or generators. For smaller data centers with lower cooling demands placed in existing buildings, using scroll chillers may be a more attractive option than hardening a larger, shared cooling system based on centrifugal chillers.

Direct Expansion

Many data centers have cooling systems that use CRAC units with chilled water cooling coils. One option is to integrate a direct expansion (DX) Freon* cooling system into the CRAC units and put them on generator power to provide cooling in case the chilled water system or utility power is lost.

In data centers that occupy part of a larger building, it may be more economical to use backup DX systems, rather than providing standby power for the large chilled water plant that supplies the entire building. This also eliminates a single point of failure (SPF); with multiple CRAC DX systems, failures must occur simultaneously in more than one system to result in a significant loss in cooling capacity. When assessing the value of this approach, it is essential to analyze and compare the time it takes to get generators started, CRAC units restarted, their DX systems started, and cooling coils chilled down to cool the air, relative to the rate at which temperature rises in the data center.

Case Study

Intel has designed and built several high-performance data centers. These are high-density, power-efficient data centers housing 20,000 to 50,000 servers at power densities between 10 and 25 kW per cabinet. Our approach has allowed us to support rapidly expanding computing needs at lower cost than other data center designs.

With each new project, we look for ways to reduce costs. Thermal storage tanks have proved to be an effective way to do this.

We implemented a thermal reserve system based on cold water tanks at one of our major regional hub data centers. When designing this facility, we aimed to reduce construction costs. Thermal water storage tanks proved to be an economical way of helping prevent thermal excursions within the server space when an electrical sag or outage occurs. The alternative approach, putting all of the facility chillers on a CPS, would have been more expensive. The utility power supply in the region housing the data center is in general very reliable, and outages are rare events.

The facility uses both centrifugal and scroll chillers. The three 1,200-ton centrifugal chillers, which are more efficient in terms of kW/Ton, supply the main cooling system: a 55° F chilled water system providing sensible cooling of the areas housing the IT equipment and power supplies. The two smaller 175-ton scroll chillers supply a smaller-capacity system with chilled water at 42° F. This is used for latent cooling (dehumidification) and non-critical loads. The 42° F system also trickles through to cool the water in the thermal reserve tanks.

We hardened the 55° F system to continue operating through as many short-term sags as possible, but this did not completely eliminate the risk that a sag or outage could cause the chillers to be offline while heat was still being generated in the data center.

We identified two worst-case scenarios:

- A utility power sag shuts the chiller plant down and the UPSs continue to power the IT equipment in the data center, producing heat. Without any supplemental thermal storage, the servers in the data center will suffer thermal damage and shut down due to high ambient temperatures before the UPS battery reserve is exhausted.
- Utility power is lost and the chiller plant shuts down. Then, late in the UPS battery drain period, the utility power comes back on line. As a result, the IT equipment is now running as it normally does on utility power to the UPS, but the chiller plant has been offline for nearly five minutes and will take several more minutes to resume normal cooling.

In both of these situations, the IT equipment may be producing heat while the cooling system is off. The IT equipment may not sense the loss of cooling until it is too late to avoid thermal shutdown. The UPS for the IT equipment was designed to provide five minutes of power when the equipment is fully loaded, and if cooling stops while the IT equipment is running, it takes only a few seconds for the temperature to rise dramatically. This could damage servers that would potentially cost millions of U.S. dollars (USD) to replace.

Our solution was to install a large supplemental thermal reserve system. Our system is based on two 24,000-gallon cold water tanks, sized to provide enough capacity to cool the data center for seven minutes longer than the UPS battery life when the IT equipment is running on full load. Our calculations indicated that this would represent a huge improvement over using only the normal data center thermal reserves, comprised of the cold air in the data center and water in the main chilled water system.

Our supplemental thermal reserve system promised to solve many potential issues:

- In the case of a short power sag, the reserve would provide cooling until the centrifugal chillers restarted.
- If there is a complete loss of power, the reserve would provide cooling for the UPS battery drain period, plus the time needed to remove the residual heat from the data center after the IT equipment shuts down.
- If utility power returns toward the end of the UPS battery drain period, the thermal reserve would still have the capacity to provide cooling long enough to get the chiller plant back on line.
- If, in the future, we decide to add standby generators, the thermal storage will still serve a useful purpose by providing continuous cooling during the period when the power source is switched from utility power to standby generator.

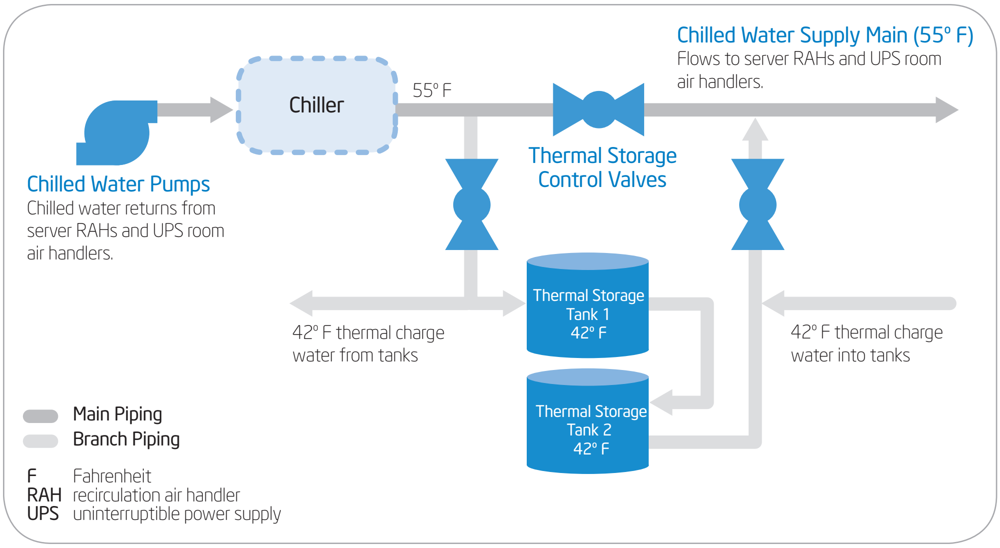

Figure 2. Diagram of supplemental thermal storage system loop. The valve on the main chilled water supply line is open during normal operation. In the event of an outage, this valve closes and the two valves on the thermal storage branch piping open to allow the water in the storage tanks to flow into the chilled water supply, providing cooling to the air handlers.

Thermal Storage System Operation

Our data center and thermal storage system are shown in Figures 2 and 3. During normal operations, the centrifugal chillers supply chilled water at 55° F to the RAHs that cool the IT equipment. Meanwhile, the low-temperature scroll chillers maintain a trickle of cold water to the thermal water storage tanks, which keeps the tanks at about 42° F. Storing water at this low temperature reduces the cost and size of the thermal storage tanks needed. It also means that in the event of an outage, the colder water from the tanks can pass through the system more times before it becomes too hot to provide data center cooling.

In the event of a power outage, the centrifugal chillers stop. The IT equipment keeps operating because it is on UPS. The thermal storage tank valves open and add water from the tanks at 42° F into the main 55° F supply feeder line, helping to keep the main chilled water supply at a low temperature.

The chilled water pumps are on a separate facility CPS, so they keep the cold water moving through the RAH cooling coils. The RAH fans are also on this CPS, so they continue to move air through the cooling coils and deliver cold air into the data center space.

If utility power does not return before the UPS batteries drain completely, the servers shut down. Two 2-megawatt standby generators support critical facility equipment and loads. Part of this generator capacity continues to power the water pumps that keep chilled water running to the coils as well as the RAH fans that keep a nominal amount of cool air flowing through the servers. This removes the residual heat that has built up in the server cabinets and helps prevent thermal damage.

The thermal storage system contains additional safeguards. In normal operation, the thermal storage tank valves are closed, isolating the 42° F water in the tanks from the main 55° F chilled water system. However, if a control failure should occur, the valves fail to the safe position of bleeding thermal reserve water into the chilled water system.

Response to Outage

A power outage one evening in late 2006 put the thermal storage system to the test. The outage lasted several hours.

Like other Intel high-performance data centers, the facility is highly automated, with no need for operations staff to be present 24x7. As a result, operations staff were not onsite when the outage occurred.

When utility power was lost, the IT equipment continued running and generating heat because it was supplied by UPS power.

The thermal storage system worked as designed. The building control system sensed the power outage and activated the thermal storage system. Valves in the thermal storage tanks opened, feeding cold water into the chilled water supply line. The facility UPS system kept the chilled water system pumps and RAH fans running until the standby generators started and picked up the load.

Because the servers were lightly loaded at the time, UPS batteries lasted much longer than the five minutes for which we designed our original system, providing power for more than 15 minutes. However, the light server loads also meant that less cooling was needed. As a result, the system provided adequate cooling during the entire period that IT equipment was on UPS, and long enough afterwards to remove residual heat.

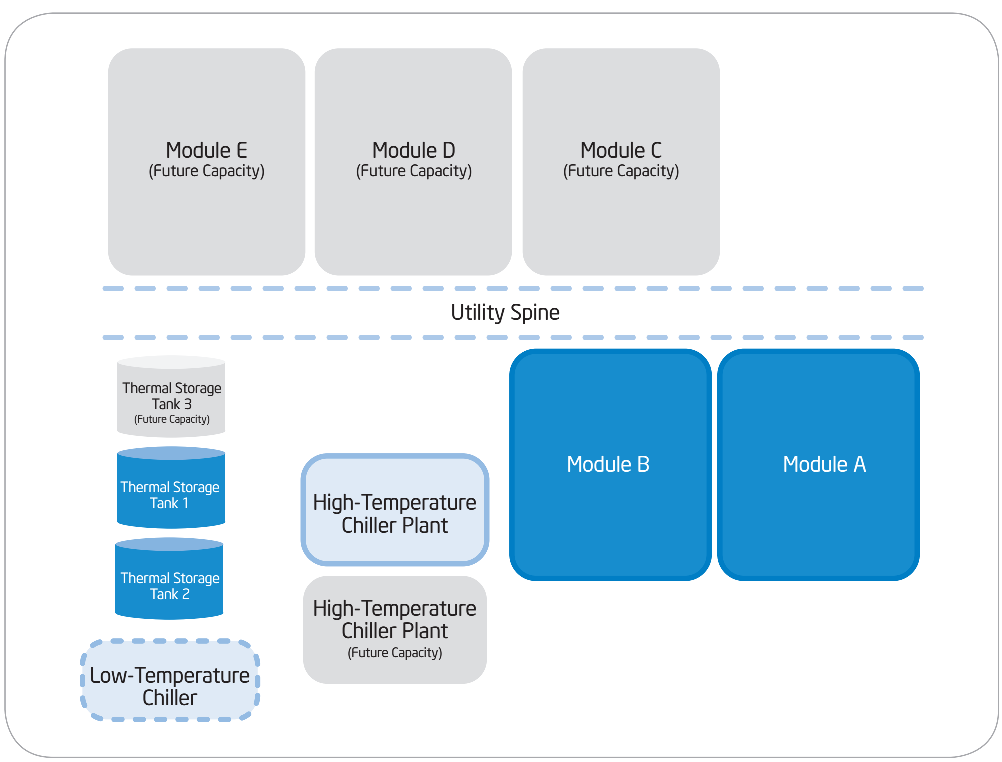

Figure 3. Schematic plan of data center, showing thermal storage tanks and chillers.

As the UPS batteries drained, the IT equipment shut down.

Operations staff were automatically paged when the outage occurred. When they arrived approximately 45 minutes later, the data center was at normal operating temperature.

Cost Analysis

The system provided a relatively low-cost solution, compared with the alternative approach of putting the entire cooling system on CPS. The cost of the tanks was approximately USD 300,000, and the cost of UPS and generators for the RAH fans and pumps was approximately USD 500,000.

The thermal storage system offered additional advantages. It is simple and reliable, and does not require the extensive periodic testing that a full UPS- and generator-based system would require.

We designed the system so that standby generators could be added for the centrifugal chillers, cooling tower fans, and condenser water pumps in future to provide a further level of reliability. The system is designed so that these generators can be installed without having to shut down the data center.

Conclusion

Our thermal storage system enabled a high-density data center to survive a power outage without damage to IT equipment. The cost of providing this simple, reliable system was far lower than alternative methods. Based on our experience, thermal storage tanks are a cost-effective way to provide temporary cooling in high- and medium-density data centers, preventing millions of USD in potential damage to IT equipment. We have since used a similar approach to design thermal storage systems at other data centers.

Authors

Doug Garday is a senior mechanical engineer with Intel Information Technology.

Jens Housley is a senior project manager and architect at Intel Corporation.

Acronyms

CPS	continuous power system
CRAC	computer room air conditioning
DX	direct expansion
F	Fahrenheit
kW/Ton	kilowatts of power used per ton of cooling
RAH	recirculation air handler
RMF	raised metal floor
SPF	single point of failure
UPS	uninterruptible power supply
USD	U.S. dollars

www.intel.com/IT

Printed in USA
0907/SEP/RDA/PDF

♻ Please Recycle
ITAI Number: 07-1205w