

Using Thermal Mapping at the Data Center

Gregory Stockton
 CompuScanIR.com, a division of
 Stockton Infrared Thermographic Services, Inc.

INTRODUCTION

For data centers, maintaining a continuous flow of information requires near 100% uptime. As 24/7/365 operations become the norm for business, ensuring that data center infrastructure stays up and running is mission critical. Hard failures can cost a large business millions of dollars in lost productivity and opportunity costs. If this were not a big enough challenge, data centers consume a lot of energy, and in times of rising worldwide demand, high operating expenses can put a serious dent in the information technology (IT) organization's budget.

Because of the need for high uptime rates, the ever-increasing cost of electricity, higher server densities and limits on electrical grid capacity, data center operators are now looking for ways to optimize performance and increase their kW per square foot rating—all while reducing costs. Reducing the consumption of energy in a data center while maintaining high availability is no small task, but the rewards are high. In the US, servers and data centers consumed 61 billion kWh (1.5% of total US electrical consumption) in 2006 and are projected to consume as much as 100 billion kWh by the end of this year¹.

To ensure reliability and economical operating costs, the power distribution and cooling infrastructure must be actively managed. Outages stemming from electrical or mechanical failure can be prevented by physical redundancy practices and predictive/preventive maintenance (P/PM) and are currently being used by most data centers. To reduce electrical consumption and address 'green' operating mandates, data center operators are now exploring ways to raise temperature set points².


What's the solution? While it may not be obvious, a usable 'heat view' provides a big part of the answer. Today, managing a data center involves far more than solid IT operations. Because temperature is directly linked to energy consumption and equipment operation, infrared thermography (thermal imaging and thermal mapping) can be used to monitor power consumption, cooling, and IT operations. In fact, monitoring the electrical equipment, cooling equipment and computing equipment which is at the heart of the data center, is precisely where thermal infrared thermography (IRt) can help the most!

Infrared thermography (IRt) is used to find, diagnose and document problems such as short-cycling of the air conditioning system, loose electrical connections and worn out bearings. After repairs have been made, IRt is used to recheck the equipment to make sure it is operating properly. Two main categories are:

- 1) **Cooling Systems and Heat-Generating Equipment.** Capturing the current condition of the cooling system is the start of the IRt process. Thermal Mapping is a new approach to gather and present that data. Thermal mapping allows IT management, heating, ventilation and air conditioning (HVAC) professionals, consultants and contractors a construct to understand heat-related problems in the data center. It also makes it possible to compare the results to the CFD models used to design them. Because the complete picture is captured in-situ, issues that were not obvious when the room and cooling system were designed become apparent.
- 2) **Electrical Power Distribution and Mechanical Systems.** Performing infrared P/PM on electrical and mechanical equipment is crucial to continuous operation and well accepted in the power industry as standard best practice. In fact, electrical IR has been the most accepted of all IR applications and there are many technical papers on the subject. Infrared predictive maintenance is a must at any data center. The electrical switchgear, motors and motor controls, HVAC equipment, uninterruptible power supplies (UPS), automatic transfer switches (ATS), power distribution units (PDU), batteries and generator equipment and all electrical devices that feed the server systems must be checked with infrared thermography and other testing on a regular basis to assure super-high reliability.

WHY USE IR THERMAL MAPPING IN DATA CENTERS

Stockton Infrared has performed P/PM IR surveys of data centers' electrical distribution systems for over twenty years. On various occasions, our clients have asked us to look at the cooling of the floors and server racks to solve a perceived or real cooling problem. We realized that the method we were using, looking down the aisles (see Figures 1), was satisfactory to look at a certain spot or two on the raised floor or server, but not sufficient to get a picture of the heat distribution in the whole data center. To accomplish that, we had to come up with a better way. So, in recent years, we focused our research to develop a methodology for collecting and post-processing the images into user-friendly, easy to understand 2-D (2-dimensional) and 3-D (3-dimensional) displays which we could present to our clients (see Figure 2). **The purpose of creating a heat or thermal map is two-fold: to confirm efficient design, and to find problems.**



Figures 1a and 1b) "Down the aisles" thermal image of the raised perforated floor and servers of a data center utilizing hot aisle/cold aisle configuration. This is the way that we looked at data centers' cooling systems in the past, but the method is not sufficient to get the overall picture necessary to solve complex cooling problems.



Figure 2) IR research being conducted at an enterprise data center to develop 2-D and 3-D modeling capabilities.

DATA CENTER OPERATIONS AND DATA CENTER TYPES

A data center is a facility which houses computer systems and associated hardware components, such as telecommunications and data storage systems. Most of the installed equipment is in the form of computer systems and servers, mounted in specially-designed cabinets (racks), and dedicated equipment for managing power and temperature and humidity. The equipment on the data center floor is usually grouped in rows with corridors between them allowing IT technicians to access the front and rear of each of the racks.

Enterprise Data Centers

Enterprise data centers (EDCs) provide unified support for corporations with their distributed facilities. They provide integration of many different functions that enable business activities based on internet and intranet services, or both. Corporations highly value the internal control and increased security provided by dedicated enterprise data centers and they are the fastest-growing type of data center. Support of intranet server farms is the primary target of an enterprise data center.

Internet Data Centers

Internet data centers (IDCs) are operated by internet service providers (ISPs). ISPs also operate enterprise data centers. The architecture is very similar to that of the EDCs, but the requirements for scalability are different because the internal user base is different. There are fewer services as compared with those of EDCs, hosting multiple customers.

Colocation Data Centers

Colocation data centers (colos) are owned by companies that house and operate data center equipment for third parties. Corporations use colos for disaster avoidance, offsite data backup and business continuity. Most colo customers are web commerce companies, using these centers as a safe, low-cost environment for redundant connections to the Internet, including “cloud” computing providers. Telecommunication companies use colos to interexchange traffic with other telecommunications companies.

DATA CENTER COOLING SYSTEMS

Traditionally, data centers have been air-cooled. Still today, the typical data center is air-cooled, utilizing the hot aisle/cold aisle layout (see Figure 3). Cooled air is fed from the computer room air conditioning (CRAC) units to the cool aisles under a raised floor through perforated tiles (diffusers) up into the cool aisle, into the equipment and out the hot aisle. The heated air is then returned to the CRAC units.

Data center cooling systems have changed little over the past 25 years, but owing to the issues discussed above, new designs are being developed and tested, two are notable; cold aisle containment and liquid cooling. Cold aisle containment uses a raised floor, but contains the cold air between the cold aisle racks, sending the cold air directly to the server inlets, greatly reducing air mixing and short-cycling. Liquid cooling is used within most CRAC units, but liquid-cooled racks take advantage of the enhanced heat transfer characteristics of liquids. Since the CRAC units can be installed outside the main floor area, this design eliminates short-cycling. These systems are significantly more complex and expensive now, but may become more and more important as server densities increase beyond air cooling capabilities.




Figure 3) Data Center Cooling System (Typical Hot Aisle/Cold Aisle Layout shown). Source: ASHRAE³

DEFINING THE PROBLEM WITH HOT AISLE / COLD AISLE COOLING SYSTEMS

The electronic, electrical and mechanical components within a data center all generate heat. Unless the heat is removed, the ambient temperature will rise, eventually beyond design specifications resulting in electronic equipment malfunction. The temperature and distribution of air within the room is managed by the air conditioning system and influenced by the layout of the server racks. ASHRAE's "Thermal Guidelines for Data Processing Environments" recommends a temperature range of 61-75°F and humidity range of 40-55% with a maximum dew point of 59°F as optimal for data center conditions. Air conditioning systems help control humidity by cooling the return space air but keeping it above dew point. Too much humidity and moisture may begin to condense on internal components. Too little humidity will result in static electricity discharge problems which may damage components. In the case of a very dry atmosphere, ancillary humidification systems may be needed to add water vapor. So, under certain conditions, the control of humidity works against the control of temperature and vice-versa.

Many server racks currently in service today are too hot to meet industry standards for maximum IT reliability and performance as reported by Uptime Institute...

"Institute research into computer room cooling indicates 1/3 of all perforated tiles are incorrectly located and 60% of all available cooling capacity is being wasted by bypass airflow. Increasing under-floor static pressure to get air where it needs to go requires permanently blocking all unnecessary air escape routes. This includes sealing cable cutouts behind and underneath products or racks (this unmanaged airflow is what is really cooling most computer rooms) as well as the penetrations in the floor or walls or ceiling and any other openings in the raised floor. Perforated floor tiles with 25% openings can be replaced with 40% and 60% grates to permit a much higher airflow. For sites with unused raised floor space deliberately spreading equipment out to create white space and reduce the averaged gross watts per square foot power consumption will be a viable option."⁴

The weakest link in the system that can lead to a hard failure and loss of availability is lurking in every data center. It is the component that is most susceptible to failure by heat at the lowest temperature. But no one knows exactly where that component is located until it fails. Accurate and even cold aisle cooling is the best practice available to a data center operator. Finding and eliminating 'hot spots' is the goal of any uptime conscious data center manager. Finding and eliminating 'cold spots' is the goal of any energy conscious data center manager. Thermal mapping satisfies both.

TOOLS FOR ANALYZING THE DATA CENTER THERMAL ENVIRONMENT

CFD Thermal Modeling

The data center's cooling system must be designed and engineered to provide cooling to computer components. The objective of the design of the cooling system is to provide a clear path from the source of the cooled air to the intakes of the servers and to return the heated exhaust air to the CRAC efficiently.

Data centers are usually designed and drawn with computer-aided drafting and design (CADD) software and modeled using computational fluid dynamics (CFD) modeling. These tools are available to predict performance for the design of a new data center. CFD is a valuable means of predicting data center thermal performance. Acceptable performance depends on accurate modeling of the energy-consuming components and the heat that they produce. CFD is, however, limited by the granularity of input data and as a result, requires many questionable assumptions.

No matter how complex and well-prepared, CFD modeling is not reality. Simple things like under floor cable or ducting installations have significant impact on theoretical flows of cooled air. Deviations from ideal performance will only show up after physical testing. Also, during and after construction, changes happen. Unforeseen issues like adding servers or increasing server densities are rarely re-modeled after construction. Contractors move equipment, change cabling and conduit routes and HVAC ductwork, inadvertently creating voids and obstructions, reducing or increasing air pressure and diverting the flow of cooled and heated air.

Obviously, these types of unforeseen changes to the thermal dynamics of the cooling system are seldom modeled. **So, IRt is used to validate the CFD model** (in a normal operating condition) and direct HVAC technicians and IT managers to heating problems (hot and cold spots). After repairs have been accomplished, IRt is used to check the repairs.

Thermal Mapping of a Data Center

Thermal mapping is a new approach to capturing the full “in-situ” thermal condition of a data center and all of its equipment. The key advantage is that it is possible to get an overall view of the thermal condition of the entire room for a given point of time while still having the capability to zoom in on specific problems. This is very different from more traditional methods because it allows overall context and viewpoint selection, much like one gets with CFD modeling, but this is actual thermal imagery. For example, reports can demonstrate how a local thermal pattern visible in one aisle is actually the sign of a cooling air blockage across several aisles. When the overall layout of the servers, floor, walls and ceiling is available, what appears to be good thermal performance in one image may actually be wasteful excess cooling when the entire thermal map is analyzed. These problem areas are easy to see only with the overall image.

Temperature sensors have the advantage of monitoring temperatures continuously. Placed in various strategic locations, sensors are a good idea to monitor overall changes, but are certainly no replacement for thermal imaging. Typically, a single thermal image has over 75,000 thousand temperature points and, for instance, a five thousand square foot data center thermal image will have many millions.

Another positive aspect of thermal mapping is that it allows for trending (comparing data gathered at different times). Information that did not appear to be important at the time of the survey can be used to discover a change in the cooling system. For instance, on one survey an area with no visible problems might be covered by thermal mapping where more conventional surveys might not capture that spot because it showed no apparent problems. When a subsequent survey shows a problem, it is possible to see the change or show that nothing has changed. It is possible to see changes like new cable runs or equipment installed which impacted free air flow.

Every data center has equipment that is important to different groups of people. By methodically capturing all the thermal data and carefully post-processing it into user-friendly displays, these different people can see what is important to them, without having to be on scene during the survey. Also, experts in different disciplines can review the imagery and prepare reports at remote locations.

To create a thermal map, one must collect the thermal and visual imagery in an ordered manner, carefully post-process it into mosaics, and create the construct to display it in 2-D and/or 3-D. To create meaningful reports, the thermal imagery must then be analyzed.





Figure 4) Thermal map mosaic IR image of a set of thirteen server racks in a data center.

2-D Thermal Mapping View of Data Center Floors

2-D thermal mapping provides easy access to large quantities of data in a user-friendly manner. The mosaic IR images provide the ability to analyze patterns that are not clear in single images. Figures 5 show a 2-D thermal map and visible map of the floor of a data center. The upper left corner appears to be warmer in the

thermal map image as compared to the lower end. The warm pattern was the result of an air flow blockage that was caused by cables added after commissioning and the addition of a set of blade servers. With attention being focused on that upper left corner, it is also valuable to look at the 3-D thermal map presented below to see if there are clues to other problems. These problems will be discussed in detail later, but there is an issue with the CRAC unit along the upper wall.



Figures 5a and 5b) 2-D zoomable thermal map image of a 5,000 square foot data center main floor.

2-D Thermal Mapping View of Data Center Server Racks

Thermal mapping of server racks provided below in Figure 6b represent highly detailed front-facing images of the servers. This is different from the looking “down the aisle” shots shown in Figures 1. By combining IR and visible images, it is possible to clearly see the details. By using this presentation it is also possible to show the cabinets with the doors open. In this sample, the details of the top of the servers are shown in a way that allows them to be related to the images of the fronts. Other combinations are available as needed.



Figure 6a) Zoomable thermal map imagery server racks (overview).

Figure 6b is a zoomed-in view of Figure 6a. It allows for a detailed comparison of the easily identifiable hardware and the thermal characteristics as it operates. By taking additional partial or complete surveys at times with different loads, it is possible to get a better idea of the overall thermal performance of any given server under varying loads.




Figure 6b) Zoomable thermal map imagery server racks (details).

3-D Thermal Mapping View of Data Center Floors

Three-dimensional thermal mapping is a new approach to capturing the thermal condition of a data center and all of its equipment and is the most powerful of all tools for presentation to operators, consultants, contractors and HVAC professionals wanting to accomplish adjustments, repairs and redesigns.

A 3-D model can be rotated and viewed from any angle. In this case (see Figure 7), the context of the walls and floor view shows a relationship between a warm floor and an apparent problem with the CRAC unit. By involving an experienced HVAC professional, it was determined that even with a properly functioning thermal control system, the cooling was inadequate. The CRAC unit that is shown to be offline could have taken over for the one that is low on Freon. But because the control system and the operators saw that the defective unit was operating, they did not question that the other one was not operating. Only the IR survey of the entire area showed that something was wrong. The complete 3-D thermal map of the floors and walls provided the big picture that allowed for remedial action.

The under floor blockage was also better appreciated when the condition of the CRAC units was considered. In addition, other views that show the floors and walls in the cooler end, lead one to believe that it is over-cooled. This provides an opportunity to save money. There is also an issue of conflict between the cooling and the humidity control systems. Overcooling below the dew point can trigger [electrical] heating to raise the temperature. This can have significant energy ramifications as the two systems fight one another. Expert HVAC personnel with access to the total picture can perform this kind of analysis without even having to be on site.




Figure 7) 3-D Thermal Map of the actual heat in a data center. Thermal mapping discovered a CRAC unit that was low on Freon. At the same time, it showed that another CRAC unit which could have handled the load was not operating because the control program showed that the other CRAC unit was operating, even though it was not cooling that area the room.

THERMAL IMAGING OF THE POWER DISTRIBUTION SYSTEMS IN DATA CENTERS

Electrical and mechanical predictive maintenance is the second part of the equation. Infrared predictive maintenance (IR/PM) is a must in a mission critical facility like a data center. The main electrical switchgear, HVAC equipment, UPSs, ATs, PDUs, batteries, battery panels, generator panels and switches, motors and motor controls and the electrical panels that feed the server systems must all be checked with infrared thermography on a regular basis to assure super-high reliability (see Table 1).

Tier	Rating	Availability
Tier 1	Basic	99.671%
Tier 2	Redundant Components	99.741%
Tier 3	Concurrent Maintainable	99.982%
Tier 4	Fault Tolerant	99.995%

Table 1) TIA-942 standard for reliability in data centers. This standard was developed by the Telecommunications Industry Association (TIA) to define guidelines for data center reliability.

The stakes are high, so there must be total accountability of all survey results, especially all of the equipment associated with the cooling, UPS, and server systems. Documentation is very important. This can be accomplished by recording the entire survey on digital videotape and/or capturing fully-radiometric IR images of all equipment, whether problems exist or not. In either case, a data log of all equipment surveyed must be created including a time/date stamp reference for all equipment.

Safety is of the utmost importance during any infrared survey of electrical and mechanical gear and no different in data centers (see Figures 8a & 8b).



Figures 8a & 8b) Safety around electrical equipment is the #1 consideration during infrared surveying. Some of the options: 8a shows an IRISS infrared window installed in electrical cabinet⁵. The thermographer safely holds an infrared camera up against a viewing port or window to see inside the switchgear. These windows can be custom fit to any data center equipment. 8b) Thermographer wearing proper personal protective equipment (PPE) inspecting a battery bank during a full battery discharge test.

Specialized Infrared P/PM Inspections


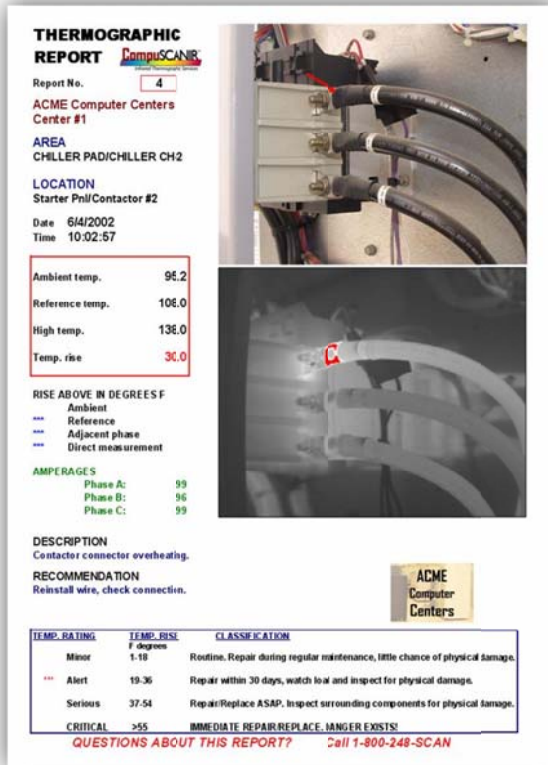
IR surveys are performed on the systems that support the uninterrupted operation of data centers on the electrical and mechanical equipment. This IR method is predicated on the fact that there is a locked relationship between heat and resistance on electrical systems (heat and friction on mechanical systems). Figures 9 shows a sample report of an electric problem on a chiller control.

Since the thermographer can see all the heat, he/she can immediately see excess heat on electrical and mechanical apparatus and make judgments on whether a problem exists and how to report potential faults such as loose connections, damaged bearings, overloaded or imbalanced circuits, faulty fuses and the like. The power distribution system in data centers is often checked during 'maintenance windows' or times designated by IT and facilities management when testing can be performed and maintenance activities can be safely carried out, without chance of interrupting critical IT activities.

Dual-power technology requires two completely independent electrical sources, tied together with switchgear. When [not if] the normal source of power fails, these dual-path power supply systems quickly switch to a back-up source. A UPS system keeps the power flowing on batteries until the normal source is restored or another source is brought on-line and synchronized. Generators are set at the ready to provide power while the batteries are providing power and before a second source of utility power cannot be utilized. Usually, the UPS, through a PDU (power distribution unit) takes the AC power and *converts* it to DC. There, the bank of batteries is tied in. The DC power is then *inverted* back to AC to feed the computer hardware.

Utility main power supplies are typically owned by the local power company but are sometimes owned by the user. Often, a looped system feeds power from two different power company substations and can be "back fed" if the power stays out on the primary. No matter who the technical owner of the utility equipment is, it must be checked with IR just like all other components.

Commissioning newly added power equipment with IR has become standard. Since some systems often cannot be tested on-line, they must be examined at times when the impact of testing is low, so that simulations can be run. By pulling power from a load bank, resistive load testing is used to fully simulate and test all equipment by increasing loads incrementally. Any problems that are encountered during this "burn-in" are repaired immediately and the system is rechecked before putting the equipment in operation.



Figures 9a & 9b) Electrical infrared predictive maintenance sample report. This problem was reported and repaired during a planned outage, causing zero downtime.

SUMMARY

Since the heart of the data center is its power and cooling system, data center managers are now using infrared technology to monitor and repair these systems. Through the use of IR thermography; predictive maintenance and thermal mapping, data centers of the future will be more reliable and use less energy.

FURTHER READING

For more information on data center maintenance using thermal imaging...
<http://www.compuscanir.com/papers.html>

REFERENCES

- 1 Report to Congress on Server and Data Center Energy Efficiency (Public Law 109-431) U.S. Environmental Protection Agency ENERGY STAR Program, August 2, 2007
- 2 Data Center Knowledge website Miller Webworks LLC.
<http://www.datacenterknowledge.com/archives/2008/10/14/google-raise-your-data-center-temperature/>
- 3 ASHRAE. 2004b. *Thermal Guidelines for Data Processing Environments*. Atlanta, GA: American Society of Heating, Refrigerating and Air Conditioning Engineers
- 4 UpTime Institute, Procedures and Guidelines for Safely Working in an Active Data Center page 9., UpTime Institute (12/18/06).
- 5 10 Things You Need To Know About Infrared Windows Copyright ©2009 by IRISS, Inc. All rights reserved

ABOUT THE AUTHOR



Author Biography – Gregory R. Stockton

Gregory R. Stockton is a principal in three infrared companies; Stockton Infrared Thermographic Services, Inc. (www.stocktoninfrared.com), United Infrared, Inc. (www.UnitedInfrared.com) and RecoverIR, Inc. (www.RecoverIR.com). Stockton Infrared is a nationwide multi-disciplined infrared service contractor. United Infrared is a nationwide network of infrared thermographers providing training on a variety of applications and the business of infrared thermography. RecoverIR is an aerial thermal mapping company primarily focused on power utility issues such as improving energy efficiencies, weatherization, and identification of lost energy.

Greg is a certified infrared thermographer with thirty years of experience in the construction industry, specializing in maintenance and energy-related technologies. He has published fifteen technical papers on the subject of infrared thermography and written numerous articles about applications for infrared thermography in trade publications. He is a member of the Program Committee of SPIE (Society of Photo-Optical Instrumentation Engineers) Thermosense and co-chairman of the Buildings & Infrastructures Session at the Defense and Security Symposium.

Copyright 2010

Infrared Training Center / Stockton Infrared Thermographic Services, Inc.