# Point-of-sale reliability: A case for proper heat management


Dave Landers
Senior Technical Staff Member
IBM Retail Store Solutions
Raleigh NC



5/13/08
V0.4

## Introduction

Heat is the enemy of all electronic devices….whether the devices are computers, cell phones, or TVs. Disposing of that heat is a key design challenge for product designers, and is not a trivial task. This article will address where heat comes from, and how to remove it from computer devices such as point-of-sale (POS) terminals, kiosks and other similar computer systems. It will include a discussion of fans, and will show how, contrary to popular belief; a properly designed system with a fan can actually enhance long term system reliability.

## Heat sources

In most computer systems, virtually all of the energy consumed is converted into heat, and this heat must go somewhere. (If a system has a fan or other rotating devices such as a disk drive, a very little bit of the energy is converted to rotating motion to turn the device, but for all practical purposes, virtually all of the energy is converted into heat. It should be noted that a disk drive generates considerable heat itself because of the air friction that the spinning disk encounters.)

Most POS systems consume between 75W-100W when sitting idle. That doesn't seem like much until one imagines how hot a 75W incandescent light bulb gets. Imagine the temperatures achieved when this bulb is placed inside of a metal box with no fan? As a simple experiment, a 75W light bulb was placed into an empty point of sale terminal chassis and the maximum temperature inside the box was measured. The interior of the box measured 45ºC (113ºF) in a 24ºC (75ºF) environment…a temperature rise of 21C (38F). This box was vented in a normal manner, but the only air movement provided was by natural circulation. No fan was installed. (A very popular children's toy uses a 75W bulb as the heating element in a small oven that can actually bake a small cake.) Since internal versus external temperature typically scales in a 1:1 fashion, this same experiment performed at a 40ºC (104ºF) ambient temperature would yield an internal box temperature of 61ºC (142ºF)! 40ºC (104F) is indeed hot, but this temperature can easily be found in outdoor environments such as garden centers, amusement parks, gas pump islands, and sidewalk sales areas.

Where does all this heat come? In a word, everywhere. Every electrical component in a typical computer system generates heat. Some components generate a little heat, while others generate a lot. We are all aware of the large heatsink and fan assemblies found on the main processor, but every single component of the motherboard generates heat as well. In figure 1 is a thermal image of a typical motherboard…note all the components that have significantly elevated temperatures and represent significant heat generators

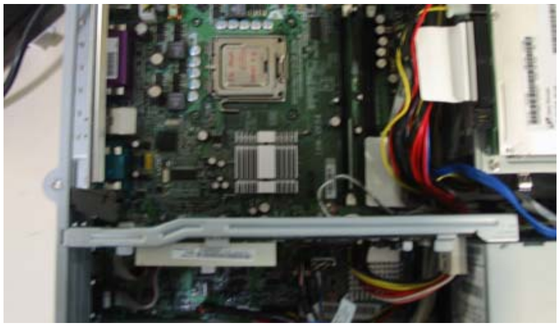

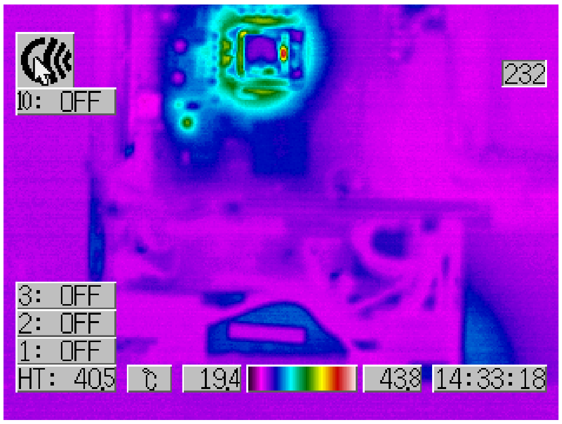


Figure 1: Image (both a normal and thermal view) of a motherboard showing processor with no heatsink installed.  The blue-green glow surrounding the processor is the heat from the CPU radiating from the copper in the motherboard. Note the very red (hot) spot on the right hand side of the CPU.

One of the largest contributors of heat is the power supply itself.  The power supply's job is to convert the power coming from the wall outlet to a form of power that can be used by the components in the computer.  Heat is generated by the conversion process, because it is not 100% efficient.  A typical  (pre-2008) desktop PC power supply is only 70% efficient at maximum output (meaning that 70% of the input power is converted to output power, and 30% is wasted as heat), but that efficiency can drop to as low as 20% in a scenario where the power output of the power supply is much less than the maximum output value. (Beginning in 2008, one of the rules of the new Energy Star program is that computer power supplies must be 80% efficient across the entire load range).

In summary, the rule about conservation of energy is true with electronic devices. What goes in, must come out, and when as much as 99%+ of the energy going into a POS system comes out as heat… it has to go somewhere.

## Heat, Components, and Reliability

As noted at the beginning of this article, heat is the enemy of all electronic devices.  In our daily lives, we all know how exposure to extended periods of heat can cause materials to change their basic characteristics….whether it be color, flexibility, texture, etc.  In most cases, the function of an electronic device is a consequence of not only its structure, but of the physical properties of the basic materials that it is made of.  If the properties of these materials change over time, then the performance of the part changes as well.

This characteristic is one reason why most electrical components have a tolerance range for a given value.  A resistor with a value of 1000 ohms might have a ±10% tolerance (meaning that the resistor could have a value from 900 to 1100 ohms).  More importantly, it means that the resistor could have one value at the time of assembly, but that value could vary anywhere within that ±10% range thru the lifetime of its use…and this variance is typically driven by material changes within the component.  Heat is one of the main drivers of these material changes.

Virtually all electronic components have minimum and maximum temperature ratings.  These ratings are defined by the component manufacturer and guarantee that the component will meet its specification over the defined temperature range for the defined life at an acceptable failure rate. The most common range is what's called the "commercial range". Devices meeting commercial grade ranges are defined to operate over a 0ºC to +70ºC temperature range. Circuit designers take careful note of these ratings and insure the components that are selected do not exceed these temperature ratings at any operational condition (such as full speed printing or with the maximum amount of memory installed) at the maximum defined temperature limit.

Reliability is typically defined to be something like "no more than X% of components will fail after YY thousand hours if operated within the defined temperature range".   Failure rates increase as temperature increases even in the safe temperature range.  Designers know that better reliability is achieved with cooler components, and seldom is a part allowed to operate at its maximum temperature limit if there is an alternative.   For example, an aspect of IBM retail hardening is to design in temperature guardbands so that components are not operated near their limits in typical operation.

The reliability (or aging) effect of high temperature operation can be used for accelerated life testing. It is possible to significantly accelerate the failure rate of a device by operating it long term at a high temperatures, in many cases beyond the rated temperature limits.  Sophisticated statistical modeling has proven over time that a product can be tested at extended temperatures, and the resulting failures can be used to come up with a reliable indicator of long term performance.

It's important to note that the effect of heat on failure rate is not linear.  The failure rate starts to increase dramatically when the component is operated closer to the maximum temperature rating.   Conversely, if a component is operated significantly below its maximum rating, the failure rate is typically improved over any defined specification value. (A typical failure rate curve plotted vs. temperature would be relatively flat until a certain temperature, and at that point it would start to rise at an ever increasing rate.)

Bottom line, all of the data shows that if heat can be reduced so that components have significant temperature margin from their maximum ratings, we can make systems more reliable over the long term.

## Getting rid of the heat

One of the key design tasks facing an engineering team is how to get rid of the heat. Heat transfer is a well understood phenomena. There are some very simple rules that one needs to be aware of for heat transfer:

- Heat always "flows" from a hotter place to a colder place.
- Different materials transfer heat at different rates.
- Heat flows "faster" when there is a greater temperature difference between the area that is hot and the area that is cold
- Hot air rises
- The heat relationship is fairly linear. If you have a product that is 40ºC on the inside in a 20ºC room, it will be 60ºC in a 40ºC room, assuming all other things are equal.

In electronics, there really is only one place for the heat to go, and that's to the air surrounding the product. (Older mainframes used to pipe in chilled water to carry heat away, but that's not really an option for most modern situations.)

In the simplest case, the heat inside the box is radiated to the outside air thru the cover of the box. The warm air inside the box warms the covers. The warm covers then radiate the heat into the outside materials (air, tabletop, etc.) surrounding the box. As the air surrounding the box heats up, it slowly rises because of natural convection, and is replaced with cooler air. This cool air is then heated by the covers and a continuing cycle evolves. At some point the temperature inside the box "stabilizes" which indicates that the heat being dissipated into the air by the covers is equal to the amount of heat being delivered to the inside of the box. At this point the temperature on the inside of the box is stable, and the difference in the temperature in the inside vs. the outside is what's known as the "box rise temperature". This box rise temperature typically remains constant no matter what the outside temperature is.

If we want to lower the internal temperature of the box, a first step might be to put large fins on the external covers of the product. These fins provide more surface area for the air to circulate around, and will dissipate more heat than the finless box. The addition of fins increases the amount of heat that can be dissipated; not only do they add surface area, they tend to channel the movement of more air via natural convection processes…similar to a fireplace chimney.  Finned heatsinks are not only common inside of systems, but in many consumer devices (TVs, stereos, etc)

Unfortunately, in many cases, most of these natural cooling solutions do not move enough heat to ensure that the maximum component temperature of components are not exceeded when the system is operated at it's maximum temperature.  This is particularly true in systems that are space constrained or have to get rid of large amounts of heat in a small chassis.  It's also true in cost sensitive situations because of the cost of sophisticated "passive" thermal solutions.  That's where fans become important.

If we were to simply place a large table fan blowing across our product, the temperature reduction would be remarkable, because we are constantly replacing the air flowing across the box with cool fresh air at a far faster rate than would happen if the air was still.  (The same explanation applies to why we blow on our food to cool it off).   No longer is the air heating by itself and rising slowly away.   Adding fins to our box along with the fan would show even more improvement because of the combination of increased surface area and air movement.

Since shipping a table fan is impractical, the most effective choice is to put a small fan (or fans) inside the box, so that it can force cool external air into the box, forcing the hot heated air out of the box. From laptops to servers, this is the typical scenario used in electronic devices today

## Fan Facts

Fans come in a wide variety of shapes and sizes.  A typical desktop computer fan ranges in size from 20mm to 80mm in diameter.  A wide variety of different construction types exist within this family…each fan solution is very specific to its application.  Typical fan specs include the following specific technical attributes:

| | |
|---|---|
| Airflow | volume of air moved per unit time, typically expressed in cubic feet per minute (CFM) |
| Size | Diameter of the fan, most are 40, 60 or 80mm in size |
| Speed | Speed of the fan, at its defined operating voltage. This speed delivers the volume defined in the airflow rating |
| Acoustics | How much noise the fan makes when operating at maximum speed, usually expressed in bels or decibels |
| Bearings | Sleeve bearing, ball bearing or a combination |

The first four characteristics all have relationships with each other. The faster the fan turns, the more air it moves, and typically the louder it becomes. The bigger the fan, the more air it moves for a given speed. The specific physical and environmental conditions (space available, temperature requirements (both outside and inside the box), and how quiet it needs to be (a quiet store, or a loud datacenter) all play in the choices that are available to the designer.

Bearing construction has a dramatic effect on long term fan reliability. A sleeve bearing is simply a tube in the fan body, into which the fan shaft (containing the blade assembly) is placed and rotates within. Sleeve bearings are not sealed and over time any lubrication that is placed within them is lost. It is typical for dirt, dust, etc. to find their way into the sleeve. Eventually the combination of dirt, loss of lubrication, and friction causes wear on both the fan shaft and sleeve. This eventually grinds away some material on these surfaces, and this loose material collects on the fan shaft and sleeve, eventually leading to the shaft not being able to turn freely inside the bearing. Ultimately this causes the shaft to stop turning and the fan fails.

The key advantage of sleeve bearing fans is that they can be made very inexpensively, and their use is prevalent in most cost sensitive applications. This is particularly true in the cost challenged POS space, where they are commonly used in both CPU heatsink fans as well as power supply fans simply because they are the cheapest fan that can be purchased. Almost all consumer grade devices with fans have sleeve bearing fans.

Changing the fan construction one using ball bearings addresses many of the life/reliability concerns. Instead of the fan shaft rotating in a simple tube, the fan shaft is supported at each end by a small ball bearing. Ball bearings are precision mechanical devices and deliver long life in applications that require a rotating shaft such as a fan. The bearing assembly is made of specialized low wear materials, and is lubricated and sealed in the factory. This assembly is designed to keep dirt out, and to keep the factory lubrication in. Consequently, ball bearing fans deliver significantly better reliability than their sleeve bearing cousins. The major negative is that because of the cost of the ball bearings, a ball bearing fan is more expensive to buy than a sleeve bearing fan.

There are also 2 factors that dramatically affect fan reliability but are not commonly known: fan speed and operating temperature.

- Average fan speed has a direct effect on reliability. There are numerous equations that attempt to define the effect of speed on fan reliability. A conservative generalization indicates that if fan speed is cut to 50% of the maximum rated speed, fan reliability triples. Most ball bearing electronic fans are rated to have 50,000 hour lives at full rated speed. If the speed is cut in half, the reliability would then triple to 150,000 hours on average (17 years at 24 hours/day)….at least equal to, if not far longer than, the life of the other components found in the product.

  A valuable side benefit of reducing fan speed is that it minimizes the acoustic impact of the fan.

- Temperature is the other reliability factor. As average temperature is reduced, fan reliability increases. One study indicates an almost 2X improvement in reliability if the fan is operated at 25C vs. 70C.

Bottom line, with thoughtful design (ball bearings, fan speed control, etc.), fans can easily be made to operate with a lifetime equal to or greater than the life of other, non moving components in the overall product, and the presence of the fan will dramatically improve the failure rate characteristics all the other components in the product during it’s lifetime.

## Processors and fans

Most modern CPU chips today dissipate anywhere from 30 watts (for mobile processors) to 100 watts (for high performance desktop processors). With the introduction of Intel’s Core family of CPUs, a 30W to 60W range is typical. Most processors of this type have a requirement that the case temperature of the processor not exceed 70C under any condition. Keeping the case of a 30W processor installed in a desktop PC below 70C in a 25C ambient room, while not trivial, can sometimes be done without a fan. Keeping that same processor cool at 40C is almost impossible without a fan, as is cooling a 60W CPU in even a 25C room.

X86 CPUs from the major suppliers (Intel, AMD, and VIA) are classified into 4 basic categories: server, desktop, mobile, and ultra low power. Server class CPUs are outside the scope of this discussion.

The desktop category of processors encompasses the family of parts used in the typical point of sale system. (CoreDuo, Pentium4, Celeron from Intel, Athlon and Sempron from AMD, and the C3/C7 family from VIA Technologies are typical brands). Speed and price are king in this space, with power consumption/heat generation typically not considered important. Power ratings of these processors have recently been reduced from 100W in the last generation of Pentium4 and Athlon processors to 60W in the current crop of dual core processors. One key fact driving this power reduction was the practical difficulties in cooling these 100W processors. They required massive heatsinks and high speed fans to cool them, and both OEM manufacturers and consumers complained about the size and noise of the heatsinks. It's important to note that a major advantage of the desktop category of CPUs is price; of the three categories, these are always priced much less than the processors found in other categories.

Power consumption (because of battery life) is the defining feature in mobile processors, such as the Intel Pentium-M, AMD Turion, VIA C7-M, etc. Most of these processors are still rated in the 30W range (although some are as low as 10W), and almost all current laptops still require heatsink fans in order to keep the processors within their thermal design limits. (Part of this requirement is due to the fact that there is little room for a large finned heatsink in a laptop.) Mobile processors are also significantly more expensive (as much as 3X) than their desktop counterparts, making their use in cost sensitive applications that don't require battery power very difficult to justify.

Finally, there are very low power X86 compatible processors available (sometimes called ULV processors). ULV stands for ultra low voltage…these processors typically run on 1.2V or less. Common characteristics of these processors are power ratings of 5W or less, and it's not difficult to design fanless solutions to keep them cool. Unfortunately, they offer much lower performance levels (typically 400MHz – 1.2 GHz) AND are higher in cost than either their mobile or desktop cousins. Usage of these devices is typically restricted to very compact point of sale where it's not possible to install a fan at all, or in laptops that are exceptionally thin and light. (Lenovo X series laptops use ULV CPUs).

From a power consumption perspective, it's very important to note that processors are constantly changing how much power they consume, and when they aren't doing anything, most consume very little power. This is true no matter what type of CPU it is. Even a 60W rated CPU only uses 10-15 watts when it's just idling waiting for work to be done (such as a sitting at a Windows desktop). Only when a CPU has CPU intensive tasks (such as media encoding/decoding, database analysis, simulations, etc.) does the CPU actually consume anywhere close to its rated power. **This fact is extremely important and has a significant impact in design decisions.**

## Fans are GOOD!

A properly designed system with carefully selected fans is just as reliable (if not more so), has a much wider range of operating temperatures, and likely delivers far more performance than a fanless one. More performance translates to more flexibility for the end customer. In most cases the price to the customer is less than a fanless solution as well.

A properly designed fan system includes the following attributes:

- implementing multi level fan speed control for all fans contained in the system;
- all fans must be of the ball bearing type
- Careful fan selection and vent design to ensure “flow thru” cooling of the entire product.

Fan speed control combined with ball bearing fans is the key to delivering extended fan reliability. (Fan speed control is simply a system that controls fan speed based on the internal temperature. As the temperature goes up, internal circuitry detects this increase and speeds up the fan.) Average fan speed of systems with fan speed control is 20-30% of rated speed, which extends the fan far beyond the life of most of the other components.

Since fan speed is so low most of the time, systems with fan speed control deliver acoustics performance almost equal to fanless systems in many application environments, particularly those in normal office temperature environments. A system with fan speed control typically keeps the fan speed to a level where the fan noise isn’t even noticed.

Finally, systems with fan speed control have little to worry about from a dirt and dust perspective. The total volume of air that moves thru the box is low, and contamination is of negligible importance. Only when large quantities of air are moved at high velocity, do dirt, dust, and grease contamination become an issue.

The presence of fans will also provide a much lower internal operating temperature for the life of the product, extending the lifetimes of all components in the product. In the simple experiment described at the beginning of this article, a small 60mm fan was installed in the metal box with the light bulb. This fan was operated at 1/3 of its rated 3600RPM speed and was virtually silent. The temperature rise of the box went from 22C (40F) to 9C (16F), with the surface of the box going from very hot to cool to the touch.

Some suppliers in the industry are now trying to convince buyers that fanless solutions deliver solutions with higher levels of reliability, simply because there is no longer a fan that can fail. As these solutions are examined more closely, it becomes apparent that significant sacrifices or limitations have been made that compromise performance, reliability, or cost. Even worse, some deliver the potential for increased failure rates.

Specific examples:

- reduced temperature ranges for approved operation
- much slower processors
- power supplies packaged separately from the main system. (This approach adds significant cost because of the cost of the 2 vs. 1 box approach, and the cost of the cabling and connectors required to support this approach.) Even worse, the reliability of this solution is compromised because of the additional complexity and failure points associated with the two box solution. (Connectors are the most problematic part of any system…the more connectors, the more exposed one becomes to failures.) The space required for 2 box solutions is typically much larger as well.
- In many cases, suppliers put fans in the external power supplies of these so called "fanless" systems, negating many of the claims that a fanless advantage provides. In most situations, these fans are easily blocked by other things that are buried in the cabinetry along with the power supply.
- Higher average temperatures for the system, which can lead to higher failure rates, particularly over a long period of time.
- Requiring separate power supplies for high power devices such as printers, displays, etc. In general, it's very difficult to make a power supply with an output in excess of 100W or so without an internal fan. The complexity/reliability issues associated with 3 different power supplies (one of the system, one for the printer, and one for a flat panel display) are far worse than any reliability concern about a single, "higher power" power supply with a fan, particularly one with fan speed control.

## Conclusion

The benefits of POS systems with properly engineered fan solutions are many:

- Lowers the internal temperature (and enhances reliability) by moving more heat out of the box quickly
- Allows for the product to be used in much higher operating temperature environments
- allows the use of higher power CPU's
- allows the use of lower cost CPU's
- allows the use of higher performance CPU's

Properly engineered fan solutions have excellent reliability characteristics and do not show failure rates out of proportion to other components used in the system. The key phrase is "properly engineered"…careful selection of fan type/bearing construction, air flow characteristics, acoustic impact and fan speed control all must be considered in the solution. IBM’s field experience with millions of POS systems over many years confirms this conclusion.

The ultimate benefit of a system with a well designed fan solution is improved system reliability and better performance at a lower cost.