HALT & HASS Successes on IBM Web Servers

Anthony T. McCray
IBM – Beaverton, Oregon

Biography

Anthony McCray is a Laboratory Technician at IBM's Web Server division in Beaverton, Oregon. His primary responsibilities are performing HALT tests, generating reports based on the test results, and driving failure analysis until root cause is determined. He was first exposed to these techniques in 1995 during an internship supporting the Test Engineering department. Upon completion of his internship he returned to the HALT testing laboratory in 1996 and continues supporting this testing. In addition to approximately fifty HALT tests that have been performed during the last four years he has also been involved in efforts to screen manufacturing product, and perform engineering qualification tests, such as Four-Corners. His Internet address is mccray@us.ibm.com.

Definitions

ASIC	Application Specific Integrated Circuit
BGA	Ball Grid Array
ECC	Error Correction Code
HALT	Highly Accelerated Life Test
HASS	Highly Accelerated Stress Screen
NPI	New Product Introduction
PCI	Peripheral Component Interconnect
PoS	Proof of Screen

Abstract

The purpose of this paper is to communicate the numerous successes that have resulted from HALT, HASS, and other environmental testing used by IBM's Web Server division in Beaverton, Oregon. The engineering and manufacturing facility in Beaverton has been using HALT & HASS techniques for eight years. The first section in this document will familiarize the reader with IBM Beaverton's product so that they will have a better understanding of the product under test. The second section will describe the environmental tests used during product qualification and manufacturing, which include HALT, HASS, Proof of Screen, and Four-Corners. The third section will discuss the issues uncovered and corrected by utilizing these various techniques on the most recent product generations.

Beaverton's Product

In order to understand the test process better it is important to have some knowledge of the product under test. The Beaverton site designs computer servers based on Intel® processors. The current base product is a four-processor system (called a Quad) which can currently be configured with up to 8GB of memory and has 7 PCI expansion slots. Using a technology called the IQ-Link® these Quads can be connected together to form a system with up to 64 processors and 64 GB of memory. The first generation of this product was released in late 1996. The third generation of this architecture, which is based on the Intel® Pentium® III processor, was released in March of 2000.

Beaverton's Process

There are five distinct environmental tests that are performed on Beaverton's products during various phases of the product life cycle. The first is typically Four-Corners testing which is performed by systems engineering (usually the design engineers). The second is HALT testing, which is performed by the NPI Lab. After HALT is complete the HASS profile is defined and PoS is initiated at the contract manufacturer. Once pre-production and production begin, the contract manufacturer performs HASS.

Four-Corners

Beaverton's implementation of Four-Corners utilizes high and low thermal and high and low voltage stresses. The low temperature is typically 10°C and the high temperature is 50°C. Voltage is margined at +10% and -8%. Two system configurations are tested during Four-Corners. The first is an unloaded system, which typically has a minimal memory and PCI configuration. The second is a loaded system, which will have a maximum memory configuration and a greatly increased PCI load. Four tests are run, starting from least stressful (low temperature/high voltage) to most stressful (high temperature/low voltage). Standard diagnostics are used during Four-Corners, and they are run continuously for twelve hours at each corner.

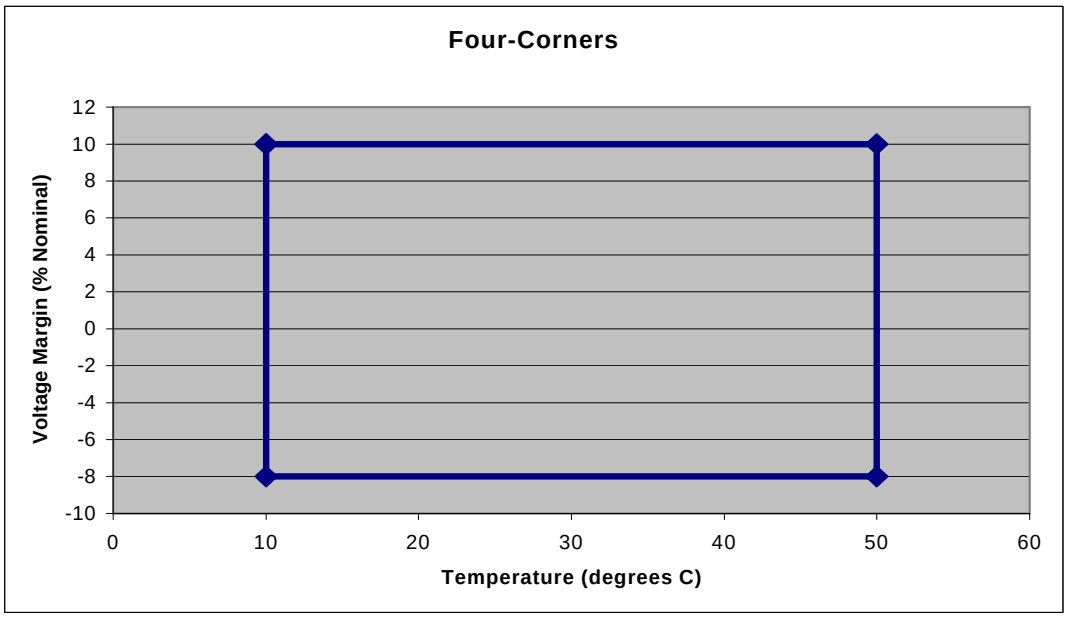

Figure 1: Four-Corners

Highly Accelerated Life Test - HALT

HALT at IBM Beaverton is performed on early pre-production product. The equipment used has the ability to apply thermal, vibration, and voltage stresses, although recently only thermal stresses have been used on most products. Vibration was stopped when evidence was found that relevant failures were not uncovered with it on our particular product. Vibration testing will be resumed as a completely new architecture and new technologies replace the current generation of products. The HALT process used at Beaverton differs from many implementations of HALT in that it is a lengthier process than most. HALT begins with thermal step stress from ambient to the operational limits (from room temperature until an unrecoverable failure, i.e., system crash, occurs). The second stress applied is thermal dwells near the operational limits (the operational limits found during thermal step stress are reduced by 20%), and finally rapid thermal transitions run between the extremes used during dwells. During thermal transitions the temperature changes between these extremes as fast as the chamber can go, or approximately $\pm$ 60°C/minute. Once the transition is complete, the stress chamber will dwell at the extreme for seven to ten minutes to allow the product under test to “catch up” to, and stabilize at, the new ambient air temperature. With this profile it is typical to see temperature transitions on the product of approximately $\pm$ 30°C/minute. Due to the complexity of our product, and the time needed for the diagnostics to execute one full pass, we typically dwell for one to two hours at each stress level during thermal step stress. With this type of time requirement a typical HALT can take from two to six weeks, depending upon the qualification being done (component qual, new design, etc.) and the issues uncovered during the testing. Figures 2 and 3 below are examples of the stress levels and time required to complete thermal step stress, dwells, and thermal transitions.

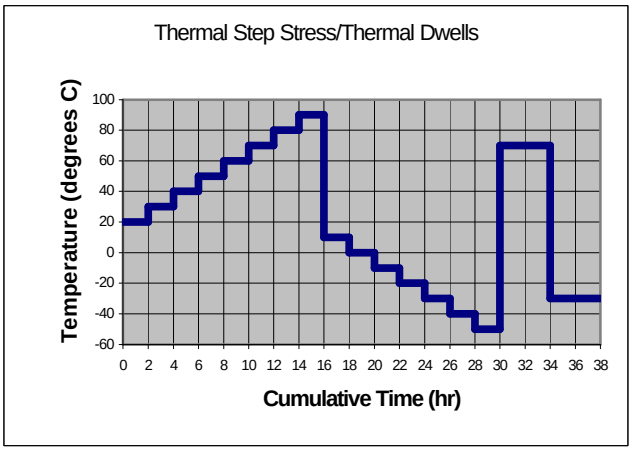

Figure 2: Thermal Step Stress/Dwells

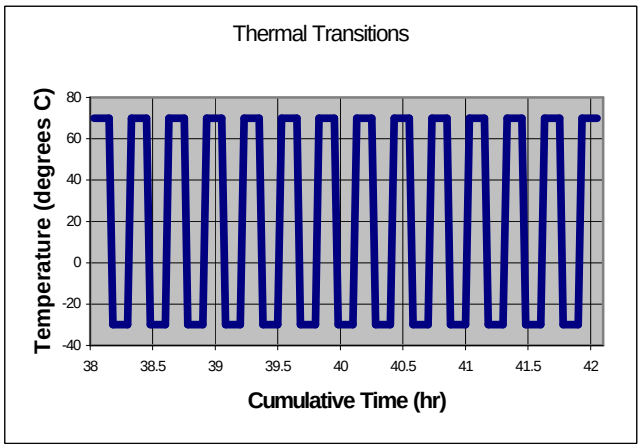

Figure 3: Thermal Transitions

Proof of Screen - PoS

Proof of Screen is performed at the contract manufacturer once HALT is complete and the HASS profile has been designed. PoS is performed whenever production of a new product is about to begin, or when there is a significant change to how the current product is tested (i.e., major hardware revision or a change to the HASS profile). PoS consists of several systems being run through the HASS profile ten times. During PoS the product's response is monitored, and the stress chambers are tuned to ensure that the product in the HASS environment is responding to the input stresses the same as in the HALT environment. Different chambers are used during HALT and HASS, so it is important to make sure the product responds to the input stresses the same in both types of chambers.

Highly Accelerated Stress Screen - HASS

All new production material is tested in a HASS screen at the contract manufacturer. The first part of the current HASS screen is called pre-stress, which consists of dwells at -30°C and 60°C. Once pre-stress is complete HASS begins with two hours of thermal transitions between -30°C and 60°C with ±20°C/minute transition rates. The profile finishes with more dwells at -30°C and 60°C for one hour each. The current profile is five hours long, including pre-stress.

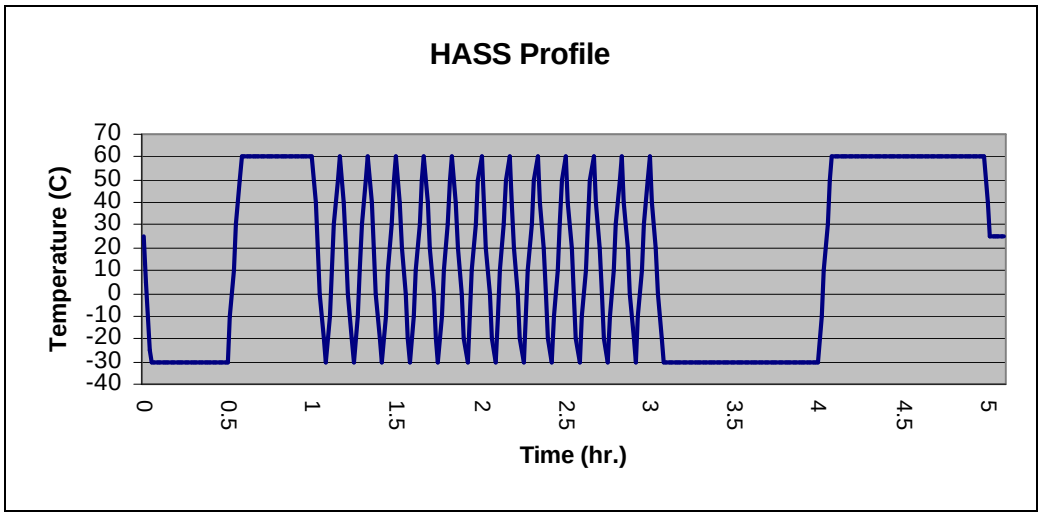

Figure 4: Current HASS Profile

Successes

Over the years IBM Beaverton has had great success using these techniques. The sections below describe some of the defects found in the current product line, starting with the second generation of the architecture.

Second Generation

The second generation of Beaverton's product was based on the first Intel® Pentium® II Xeon™ processor. While functionally the same as the previous generation, this was a major hardware redesign. Several issues were uncovered during HALT, including ASIC problems and a documentation issue.

HALT

The first issue uncovered with HALT was an ASIC problem. This particular ASIC was only used in a multiple Quad configuration, and as such, this failure only occurred in a system with three or more Quads. The failure occurred repeatedly above 60°C, which caused a lot of concern since all of the previous products that had been tested had much higher upper operational limits (90°C and above). Systems Engineering became heavily involved in debugging this failure, spending about three weeks reproducing the failure and capturing traces on several different samples of the product. During this time, a timing issue was uncovered in the ASIC. Failure analysis was driven back to the designer and manufacturer of the part. The manufacturer discovered a flaw in their simulation tool, which resulted in the parts running slower than expected. As it was already too late to affect the design of this product, a tougher screen was implemented at the supplier to screen out the slow parts. The operational limit of this system increased from 60°C to beyond 90°C once good parts were reworked onto the HALT samples.

A second problem that was found involved another ASIC in this system. This problem also caused the system to crash above 60°C. This ASIC was manufactured by the same company as in the previous example, and so a similar problem was suspected. Failure analysis showed that there was a slow path in the ASIC, which caused a parity signal to take longer to reach its destination than the data. This issue resulted in a timeout and subsequent system crash. As with the previous case, it was already too late in the design cycle to affect the ASIC design. A software change was implemented to correct this issue.

The third issue uncovered during this process had a strange root cause. This issue involved the ASIC discussed in the previous paragraph. In this case the system was crashing after a multiple bit ECC error was detected by the ASIC. Again, Systems Engineering came to the HALT Lab to aid in debugging efforts. They soon discovered that the multiple bit ECC error that was being reported, and crashing the system, was actually a correctable single bit ECC error. In addition, multiple bit uncorrectable ECC errors were being treated as single bit ECC errors, which would have resulted in silent data corruption at a customer site. Further research found that the software was coded to treat the errors in this manner. It was then found that the documentation written for this particular code was incorrect. When the documentation was originally created the ECC error handling had been done incorrectly, and the software had been written according to this documentation. This error was fixed by a new software release that was coded correctly, and of course, updated documentation.

Third Generation

The third generation of this product was very similar to the second. Changes were made to support faster processors, and increase the bus speed. These changes required another major hardware redesign, although the basic architecture remained the same.

Four-Corners

Four-Corners was the first environmental test applied to this product. One major failure mode was found during testing. In a loaded (over 4 GB of memory) multiple Quad configuration, single bit memory errors were streaming from a special physical memory location reserved for communication between the Quads. Due to the increased bus speed, the reads from the SRAMs were too slow. The fix for this issue involved changing the resistors on the lines between the bus and the SRAMs to speed up this transaction. The 'slow-corner' of this test (high temperature/low voltage) detected this failure mode.

HALT & HASS

During HALT several issues were uncovered. The resistor problem found during Four-Corners as discussed above, was one of them. In addition to this, two other major problems were found, both of which led to the same failure mode. The first involved a small BGA terminator board that was used to terminate the signals on several ASICs in the new product. During testing at cold

temperatures the system was crashing repeatedly. It was discovered that the manufacturing process for the terminator boards was faulty. The layers of the terminator boards were delaminating, and during periods of low temperature testing when the materials were contracting, buried resistors and vias were opening. After repeatedly having problems, production of this part was moved to a different manufacturer. This issue was actually known about before it was discovered in HALT, but HALT was able to detect it much faster and reproduce it more consistently than other test methods (room temperature tests). Parts from the new manufacturer were not immediately available for upgrading the HALT samples. The HALT samples were reworked with better parts from the original manufacturer. This failure mode continued to occur, however with the new parts the margins and operational limits had been increased.

Shortly after parts became available from the new manufacturer our contract manufacturer began its initial builds for production. Even though the new builds had the new parts, they were still failing at low temperature with the same failure mode. This particular failure mode cut the yield of one card to almost nothing. With the persistence of this problem systems engineering became involved in the debug effort. After several experiments were performed a very close look was taken at the schematics. The result of this effort was the discovery of a design flaw. A resistor that supplied a reference voltage to a key ASIC (the ASIC terminated by the terminator board) appeared to be incorrect, causing a reference voltage to be 50 mV too high. An experiment was conducted where this resistor was changed on several boards and the samples were run for extended periods at low temperatures (below where the failures were occurring). The resistor modifications corrected the failure mode on all of the boards. Next the resistors on the HALT samples were reworked. At this time the HALT samples still had suspect terminator boards installed. These boards were also tested at low temperature extremes. The resistor modification fixed the HALT samples as well. Although the original terminator boards were defective they were not the only issue that manifested itself in this failure mode. This experience provides a great example of the importance of driving all failures to root cause, implementing corrective action, and verifying the fix. If samples of terminator boards from the new manufacturer had been installed on the HALT samples it is likely the resistor flaw would have discovered long before the contract manufacturer began building boards for general availability. Fortunately, this problem was detected early in the production cycle, and with the rework being relatively simple, this issue did not adversely affect the release of this product.

Summary

The intent of this paper was to describe some of the faults that have been discovered and corrected through the use of HALT and HASS type testing. While the specific implementation of HALT and HASS is atypical, the results are not. The failures described in this paper are significant. Had any of these issues been allowed to manifest themselves in the field the costs would likely have been

significant in terms of customer downtime, warranty repair costs, and loss of business. As more of these tests are completed the designers and manufacturers of these products learn from their mistakes, and the lessons learned are rolled into the design or manufacturing process for the next product. While past mistakes are generally not repeated, new technologies and new manufacturing processes will likely result in new failure modes. HALT, HASS, Four-Corners, and other environmental tests will continue to play a key role in uncovering and correcting these issues.