

<http://www.sparcproductdirectory.com/artic-2002-jan-pb.html>

Unsafe At Any Speed?

Looking under the hood at Sun's recent server engine problems

Sun's cache memory problem:- What did Sun know? When did Sun know it? And what did Sun do about it? - a critical commentary.

Editor's intro:- in 1965 Ralph Nader wrote a book called "Unsafe At Any Speed" to expose design flaws with the Chevrolet Corvair manufactured by GM.

In 2001 Sun's cache memory problems caused millions of dollars of wasted time in corporate America and shattered the myth of Sun server reliability. Sun's reported communications about this subject have tried to pass the buck to a major supplier, IBM. But Sun could have designed around this problem, or detected it earlier.

The recent comments by Sun Microsystems' Chairman and CEO, Scott McNealy, are extremely interesting and shows the absolute importance of truth and candor by senior executives of any company. Often the first signs of deeper problems within the company come from very simple statements by executives that later appear to be misleading and incorrect. Once corporate credibility has been lost it takes years to recover, and sometimes never does.

Lets take a very candid look at the comment (reported in a Computerworld article November 27, 2001)

Q: We reported last year about the problem with the external memory cache on UltraSPARC IIs that was causing a lot of Ultra Enterprise servers to crash. Is that something you're still grappling with, or is it history?

A. We're no longer buying IBM SRAM [static random-access memory]. They were the biggest source of the problem for us. They knew about it before, and they didn't tell us. But we don't have that issue anymore. We designed IBM out of that and put [error checking and correcting logic] across the entire cache architecture.

The Shell Game

As we all know the real purpose of a shell game is to distract the viewer from where the REAL situation is.

What really WAS the defect: ----Where should we look for the REAL problem.

The industry has been rife with rumor and innuendo for three years, not helped by the fact the parts of SUN still today denies there is a problem. No statement has ever been issued by Sun's corporate office to clarify and no clear corporate policy has ever been introduce to rectify the problem.. ScottM seems

be pointing at normal day to day QA problems and outside suppliers. Sun takes a risk by subcontracting the majority of its work to third parties, and in doing so bears the responsibility ensuring compliance. This comment seems to say this is really no big deal and not worth talking about. Classic Shell Game. Watch this hand not THAT one.

The truth is far more serious.

In my view, the real major defect was the complete lack of design of Error Correction in level 1 and 2 of the CPU cache of the UltraSPARC II Sun's flagship CPU processor.

Error Detection and Correction Primer

Transient errors (as opposed to permanent errors due to device failures) can occur in all digital systems. They can be caused by electrical noise, ambient radiation, clock jitter and other causes. As far back as 1948, R.W. Hamming of Bell Labs developed a general theory for error-correcting schemes in which "check-bits" are interspersed with information bits to form binary words in patterns. These are now referred to as Hamming codes.

The rate of failure, and consequences, if uncorrected, are important parameters in system design. Communications systems always use error correction, because the total environment is not under the control of the designer. Memory and storage systems are particularly vulnerable, because they can capture and store a transient error bit which might have no effect in another part of the system. Error detection and correction codes have been in the Digital Design textbooks for over 3 decades, and got a good press in the early 1980's, as a way of fixing soft errors caused by alpha particle radiation in the 64k bit RAM generation. Chip manufacturers later found that alpha particle emitters were contaminants in commonly used chip packaging materials, and removed this prime cause. ECC stayed in large memory systems because it fixed transient problems from many causes including some types of device failures.

As a general rule, in all electronic devices, transfer rates are increasing as time goes on. Also, errors will occur more frequently in any kind of storage device because manufacturers are squeezing more bits into smaller and smaller spaces. As speed and density increase, error correction becomes a necessity because a smaller energy disturbance source can trigger a false bit.

Error checking technologies used in processor boards can be categorised as follows:

BEST: ECC (Error Correction Code) ECC can detect and correct single-bit errors. It is used in high-end PC's and servers.

Medium: Parity This is the most common used method. It can detect errors, but not correct them.

Cheapest: Non-Parity Because there has been an increased quality of memory components and an infrequency of errors, more and more

manufacturers do not include error checking capabilities. This also lowers the cost of the PC. Only used in lower end hardware.

Now, assuming you are the leader in your industry, Marketing says that you make the Rolls Royce, DELUXE product, that's why you command the top \$\$\$\$ prices. Which one do you use: Sun's answer on the UltraSPARC II was the cheapest: NONE, no ECC on your primary L1 cache and cheap non-recoverable parity on L2. In my view this is not a minor issue it's a MAJOR design problem.

Sparc ULTRA II

	L1 Cache (Inst.)		L1 Cache (Data)		L2 Cache	
	Tag	Data	Tag	Data	Tag	Data
UltraSPARC-II	None	None	None	None	Parity	Parity

ECC on L1 and L2 Cache reduces error rate drastically

As everyone knows most companies today subcontract the majority of the actual product work to outside original equipment manufacturers and Sun is no exception, BUT, the final quality of the product is the foremost responsibility of the brand manufacturer, in this case SUN. There is even a rumor that extensive email went from both the SUN CPU suppliers to Sun about this design and its potential failure problem..

Let's say you are Ford motor company you assemble most of your products, Dana Assembly sends you complete auto chassis for your top of the line model, but forgets or intentionally does not, because they were directed not to, put brakes on the chassis, Ford assembles and distributes, again without the brakes. Customer has an accident, no brakes. Who is liable. Doesn't take a genius to figure it would be Ford..

But you say " Who would be nuts enough to leave out the brakes?"

My point with ECC " Who would..."

NO other enterprise vendor uses any of the cheaper options ALL use ECC, it's a cost issue plain and simple, ECC costs 30% more to include and so Sun's bean counters killed it in the design phase.

OK let's move on,

You know you have a problem with the design of your ENTIRE premium product range, what do you do?

Be entirely honest and tell your customers what's up and how you are going to fix it.

Try and blame someone else.

Go into Denial, that anything is wrong.

SUN still today is oscillating between 2 and 3 . hoping that the problem will disappear, but as we know it's a SUN engineering design problem and they don't vanish, they are like pesky relatives, they hang around forever.

The first people Sun blamed for the problem was its customers and now its suppliers.

What would make them do this?

Maybe it's the fear of the potential liability? That's speculation of course, but let's look at the kind of impact this problem can have for a customer who encounters the problem.

The following is a true story, and like every exciting tale, names and places have been eliminated to protect the innocent....

Take a typical \$100 million dollar a year financial company whose entire company operation runs on SUN hardware running Solaris software with Server client desktop Sunrays. The ideal perfect utopia SUN customer.

SUN's market pitch when the customer bought the system was, Large, fast rock solid reliable, easy to maintain, central servers running thin desktop clients. The best of the best, the premier UNIX based business system. SUN's marketing argument is that desktop PCs attached to smaller servers have a high maintenance and failure rate, therefore bigger more reliable servers with less intelligent desktop units is much better. Pay big bucks up front and get the best of bread. True argument as far as it goes except when the big SUN server fails more often than its humble PC desktop. At least with the cheaper PC, even when a server fails some form of local work can still be done. Primary SUN server is down and so is your entire company.

One day a central tier one primary application server just quits and brings the entire companies business to a grinding halt. The corporate Sysad's are rapidly getting older and white haired trying to trace the problem without success, screams are now being heard from the corporate wing "What the hell have you clowns done" and other juicy comments. . Luckily enough they have also bought the best support package possible " GOLD", they call the support hotline and a local support SWAT team is dispatched, they too cannot find the problem and so they take the entire server apart and reassemble it and guess what, it starts up and runs with no problem.

Big sigh of relief and our MIS/IT system heroes go of to fight another battle.

A week later, same problem, this repeats itself again and again and again and again. After rebuilding the server several times SUN tech support points fingers at the network, the operators, the cabling, the VAR, the corporate

sysad's, the environment even outer space, that's right OUTER SPACE. At one time cosmic rays got the blame. The reader will not be faulted if you think here SUN is creating Science Fiction. The company client thought so to.

Let's highlight the Value added reseller here. The VAR, who originally introduced the client company to the SUN product range. The entity that is most trusted by the corporate client. Right in the middle of this HURRICANE of a problem "SUN fires the VAR" Incredibly stupid because the VAR now goes to represent another UNIX vendor. Who do you think the corporate client REALLY trusts.

At each step of the way large \$\$\$ sums are spent by the company to comply with SUN's recommended directives. Events go on for a period of three years until with luck, string and bailing wire, relative stability is created but no real solution has been found.. Today the company can only get basic stability running an older version of Solaris 2.6 mixed with 2.7 that has been superseded by 2.8 and so support is iffy at best. With three versions of Solaris active in the corporate MIS/IT systems, NEW problems arise. Sun advises the client to upgrade all systems to the latest Solaris OS 2.8 but cannot tell the client HOW. Several times the company wants to trash the entire system and rebuild it with another vendor's, but exactly how do you rebuild an operation system used on a daily basis. It's like changing and rebuilding a supertanker fully loaded in the middle of an ocean storm.

Company costs for , Employee time, both technical management and operational, physical, structural, environmental , incremental expenses can easily be in the \$millions and lets not forget it's a financial company who, when electronic transaction die in mid process has to manually audit them, add the inconvenience issue and its several \$\$\$ Million.

Now comes a bombshell! The corporate client finds out from an independent source that the problem really is a design defect in the UltraSPARC II CPU of the SUN hardware. SUN has known this all along but followed a disinformation, smoke and mirrors strategy to cloud the issue. That means SUN was aware of the potential problem in the design phase of the UltraSPARC II several years earlier, its actions forward where pure posturing.

A smart liability attorney will tell you that potential lawsuits could be \$\$\$ MULTI MILLION.

How widespread is the problem : Potentially every UltraSPARC II system that Sun makes from the humble Ultra to the E250 / E450 / E4500 all the way to the E10000K. Sun has had thousands of unresolved complaints. So potential liability could be in the \$\$\$\$ Billions.

As any good attorney will tell you, in a liability issue, the quicker you confess a known problem and take steps to rectify it, the better you look when you eventually end up in court, as you surely will. The more you try and cover up, the worse your crime becomes Even US Presidents have fallen, not for the offense itself, but for the cover up..

Why on earth would a premier supplier even risk this type of stupidity?

Well the deep dark secret of the MIS/IT world is that Sun knew that a rampant disease had spread through the industry. Its name "WANBIGAWUN". Probably 60% of Enterprise systems are over designed originally at the tier one level. Many of the early Sun installations hardly ever broke a sweat in the normal day to day operations. The hardware and software vendors also helped with the pre- design capacity planning configurations with no more tools than great WAGS.

At budget time the vengeance of the nerds predominated. The same humble corporate been-counter who the week before had been giving the MIS/IT Sysad grief for buying magic markers that cost an extra 50c now was being told that the new email server would HAVE to be an EWUNKAY 32 processor box running the world. "Yup we gotta havit – Yup we will grow into it" (maybe in 50 years) ha ha. Several companies (many .com) and scientific research institutions used the main enterprise servers as nothing more than boxes with lights on to impress visitors, investors and donors. Install the EWUNKAY put it in self test and leave it, don't even connect it to the network. Many large .coms and research start ups did not even have a single UNIX Sysad on its staff.

Now all this is well and good. You know that the client probably will never take your deluxe buggy (server) that's supposed to be able to go 1000 miles an hour and carry 200 people past 20mph with 4 up, BUT, and this is a big BUT, what happens when the customer / client starts to crank up the load as technology advances beyond anything you imagined when you designed it. Woops the wheels fall off. -----Not Good.

That's what happened and is still happening to the Sun UltraSPARC II Servers.

OK: now we have established, What Scott M or SUN corporate knew and when he/they knew about it. I happen to know that Scott personally got a detailed email about this one.

Now lets see what SUN did to rectify it

Without letting on to the customer what the problem really was.

As with every story there are villains and heroes, some of the true heroes are within Sun itself. The Sun local technical support team in particular did everything possible to placate and work with the client many times with total non support even being directly attacked from the rest of the Sun organization. These guys and gals stand tall, they tried and tried, Heroes every one of them.

Now lets point to a BIGGIE villain, a whole group. Sun Sales/Marketing/Professional Services and Leasing (SSMPSL) . Now as you recall the VAR has been terminated and SSMPSL, has direct access to the customer. SSMPSL used ever trick in the book to turn this issue in to a revenue generating situation to benefit SUN. Blame the network, no problem give Sun professional service \$100K/250K and they will pretend to fix the problem. Blame the company Sysad , same SSMPPL response. Blame the local

environment, you need Sun offsite hosting. Hey your system is unreliable you need SUN Disaster Recovery, Sun monitoring. Need true High availability because the customers system is failing, SSMPPL says why buy more of the same, just double your budget, You need back ups for the back ups High Availability (HA) for the HA for the HA. Yes at one stage that was a suggestion, and not tongue in cheek either. Heard right hear at a great Sun Sales pitch.. Extend your lease spend more \$.

At NO time has Sun Corporate announced to the customer what the TRUE problem was or even offered restitution for the \$\$\$\$ wasted.

Multiply this thousands of times and it would be safe to say, Sun/ ScottM/Sun Corp did very little to fix the issue, in fact most of Sun made a lot of money out of it, notably SSMPSL.

Now we move to the infamous NDA (Non Disclosure agreement)

that Scott says Sun has had for many years.

What the customer was screaming about was the clause in the NDA that seemed to give Sun total immunity for any alleged or proven wrongdoing prior to telling the customer what the alleged repair fix was. No company in there right mind would do this and this customer in particular told Sun to stick it where the Sun don't shine, pun intended.

This problem was revisited in a follow up interview of Scott McNealy, Chairman and CEO at Sun Microsystems in another Computerworld article - November 28, 2001, from which the following Q and A's are quoted.

Q. Are you fully confident that your new Sun Fire 15K server is free of this whole memory cache problem?

A. We designed all of that stuff out, yeah. In fact, all of our old products we've upgraded to mirrored SRAM. It handles it on the fly, and the problem went away. We're exceeding all of our design specs on all of our servers right now.

Sparc ULTRA III E15K

	L1 Cache (Inst.)		L1 Cache (Data)		L2 Cache	
	Tag	Data	Tag	Data	Tag	Data
UltraSPARC-III	Parity	Parity	Parity	Parity	Parity	ECC

ECC on L1 and L2 Cache reduces error rate drastically

OK> We KNOW that the UltraSPARC II does not have the design fix so .lets look at the architecture of SparcIII on the Sun Fire 15K

Whoops : Looks like ScottM made a mistake it's only got ECC on the data at level 2, still the cheap route.

Some major companies / universities / government labs that use enterprise servers and specify ECC All level 1 @ 2 and now 3 have stopped looking at an testing the UltraSPARC III because it does not meet their basic QA requirements.

Ok lets move on again and say that----- Sun tomorrow sees the light and falls on its sword and confesses

"Mea Culpa" yes we did wrong, we bad , we bad, but we will fix it.

Question is HOW besides opening up the corporate coffers REALLY wide and sending some ill-gotten gains back to the customer..

Sun has a whole new hardware line based on the UltraSPARC III, that is eventually intended to totally replace the infamous UltraSPARC II.

Question is ?? -- Is it any better than the UltraSPARC II.

While most enterprise vendors have to ECC on both level 1 and 2 and is introducing level 3, the UltraSPARC III only has EEC on level 2 data and parity on level 1.

What you say "But this is silly don't they ever learn"

Guess not

About the author: Peter Baston.

A seasoned professional with more than 20 years experience in the field of advanced product development and implementation for the high tech industry, He is the founder of IDEAS, which works with a broad spectrum of Fortune 500 companies, transitioning and implementing new development concepts from theory to use and operation.

He is known as the Logic Man, someone who brings perfect clarity to complex systems His experience ranges from advanced MIS/IT systems, software to hardware to networks and how they interrelate to business core practices and rules. He has also spent many years as a turn around specialist for major US financial and VC companies, with problem investments, who needed rapid understanding of complex technological companies.

He grew up in England and Rhodesia in Africa, now Zimbabwe, which as he often says is the perfect environment to study totally opposite systems and beliefs in how to get "Nowhere with Everything" and "Everywhere with Nothing".

He holds an MBA in Business with specialization in Marketing from URHO and is a frequent speaker and lecturer, communicator on his favorite subject " How complex systems can be made to IMPROVE the quality of life" Peter avidly champions the philosophies of Don Norman, Allan Cooper, Richard Feynman, Seth Goodin, George Gilder, Oliver Sacks, David Macauly and Lewis Mumford and believes no University Technical or Business graduate should be released into the real world until they have mastered. Terry Prachett's Discworld series