Sun Microsystems, Inc.

Best Practices Guide

Addressing: E−Cache Parity Errors

October 16, 2001

Revision History

Revision Date	*Approval Needed*	*What Changed*
January 5, 2001	**E−cache Program Office**	**Changes from December 6, 2000 version**
		Page 10: Updated released patch numbers for SSP 3.3 and SSP 3.4.
		Page 14: Updated paragraph under Process to indicate that all kernel update patches are available.
		Page 23: Third bullet changed to indicate that mirrored and non−mirrored modules can be mixed in same system or same domain.
January 23, 2001	**E−cache Program Office**	**Changes from January 5, 2001 version**
		Page 6: Important Notes section added.
		Page 7: First bullet under Item 5 revised to instruct replacement of affected component.
		Page 8: Note added under title. Item 2 revised to instruct replacement of affected component.
		Page 14: Item 2 under 2b revised to instruct replacement of affected component. Note added under Reason.
		Page 25: First bullet under Action revised to instruct replacement of affected component. Second bullet under Action added.
		Page 28: Summary Note under Non−Mirrored SRAM Modules revised to instruct replacement of affected component. Summary Note under Mirrored SRAM Modules revised to instruct replacement of affected component.
February 15, 2001	**E−cache Program Office**	**Changes from January 23, 2001 version**
		Throughout document: References to E3500−E6500 changed to Enterprise 3000−6000 and 3500−6500.
		Throughout document: References to E10K changed to Enterprise 10000.
February 27, 2001	**E−cache Program Office**	**Changes from February 15, 2001 version**

Revision Date	*Approval Needed*	*What Changed*
		Throughout document: References added to include Enterprise 220R−450 – pages 1,3
		Update "Important Notes" on page 6 to include reference to Volume Servers
		Add Volume Server product reference to operational specifications, page 18
		Update "Action" on page 25 to include Volume Servers
		Add notation about front to back cooling for Volume Servers, page 20
May 10, 2001	**E10K and Sunfire PDE**	Page 25: Added PDE approval to mix MSRAM and non−MSRAM modules on the same system board.
October 3, 2001	**Matthew Havelock**	Page 6: Added something about Netra t 1400/1405 failure of modules should be handled by replacing the module.
		Page 9: Added 'Netra t 1400/1405' after Enterprise 220R−450 in second paragraph up from bottom of page.
		Page 18: First paragraph changed to (5−40C/41−104F for Enterprise 220R, Enterprise 420R, Enterprise 450, and Netra t 1400/1405)
		Page 19: Second to last paragraph changed to (This should not be an issue for Volume Server Products (Enterprise 220R−450 and Netra t 1400/5) since they have front to back cooling.)
		Page 24: Third point in Action list changed to (For Enterprise 220R−450, replace with appropriate 400, 450 or 480MHz module. For Netra t 1400/1405, replace with appropriate 400MHz module (#501−6209).
October 8, 2001	**Matthew Havelock**	Page 24: Part number in brackets in bullet point 3 was changed to 501−6209.

Table of Contents

Overview: Best Practices Guide

Sun Microsystems' "Best Practices" document contains field−oriented sets of actions that Sun has found beneficial in maintaining or increasing availability of Enterprise 220R−450, Enterprise servers 3000−6000, Enterprise 3500−6500 and Enterprise 10000. These best practices are targeted at Sun's Service force, as well as customers who maintain their own data center installations. Following Sun's "SunUp" methodology, a holistic approach to availability, this document looks at best practices in all areas of people, process, and product management.

These Best Practices are reviewed and updated as necessary. The topics vary from installation, servicing, and training, to data collection and research. The information reflects the latest lab and field analysis, but is subject to change as further areas of improvement are identified. Sun is committed to delivering continuous quality improvement to our customers, not only on the product side, but through continuing to revise and deliver best practices documents such as this.

These Best Practices are a result of extensive research both inside Sun Microsystems' product and engineering laboratories and partnerships with external customers using Enterprise servers. These Best Practices are currently being communicated to Sun employees and the Sun customer community. Sun Enterprise Services personnel should always consult their Sun Resource Appendix associated with each Best Practices Guide. This appendix serves as a guideline of where to go within Sun for more information regarding the given topic, including internal URLs, discussion groups, or support channels.

This document outlines "best" practices which are not "required" practices, and may not apply in all circumstances. As always, Sun recommends customers work with their service team to identify and analyze unique factors particular to any given environment, system composition, and usage characteristics.

Important Notes

The information contained in this document is provided "AS IS" without warranty of any kind. This content has been derived from a document originally written for Sun Enterprise Services use in maintaining Sun customers' data centers.

When a customer performs its own maintenance, either directly or through a third party maintenance providers, the ultimate decision to use or not use any of these best practices is the responsibility of the customer.

Overview: E−Cache Parity Error

E−cache is an abbreviation for external cache. External cache is sometimes referred to as L2 or Level 2 cache.

Some customers have experienced intermittent e−cache parity errors which can be caused by a faulty component (static random access memory or SRAM) that is overly susceptible to a number of factors. These factors can include temperature, humidity, slot, process running, noise, etc. Another error influencer is ionizing radiation that occurs naturally in the environment. The density of the SRAM memory cells are such that if they encounter one of these high energy particles, the value or bit in the SRAM memory cell can be changed (bit flipped) from a zero to a one and vice versa.

When any of these variables cause the e−cache data and parity to no longer match, the result is an "e−cache parity error" to avoid data corruption from migrating throughout the system. When these errors are transient, which is normally the case, they are classified as "soft errors."

An important note is that e−cache parity errors do not cause permanent hardware damage, nor is there a Solaris bug or behavior of Solaris that induces e−cache parity errors.

In addition to working closely with the customer to implement the best practices outlined in this guide, Sun is especially interested in partnering with customers to explore other possible solutions to achieve high availability. The Sun Account Management team and Enterprise Services team will work closely with the customer to develop the best course of action in each situation.

Important Notes:

- **Mirrored SRAM modules will eliminate e−cache parity errors caused by soft errors in the SRAMs. However, the parity protection also plays an important role in detecting other errors on the module to prevent data corruption from occurring.**

  It is still possible to have Solaris report an e−cache parity error on a module with mirrored SRAM, however, we can now be confident that it is not due to soft errors within the cache itself. On the rare chance this does occur, the best practice is to replace the module.

- **For all Enterprise servers (3000−6000, 3500−6500 and Enterprise 10000), routine failures of module or system board will be handled by replacing the affected assembly, i.e., just the module or just the system board. This best practice eliminates unnecessary handling of the components.**
- **For all Volume servers (E220R−450), routine failure of modules will be handled by replacing the module**
- **For Netra t 1400/1405, routine failure of modules should be handled by replacing the module.**

Outline: E−Cache Parity Error Best Practices for Non−Mirrored SRAM Modules

1. Log error according to your standard error−logging procedures.

2. Upon correctly identifying an e−cache parity error, do not replace the module or system board or otherwise alter machine. The system will reboot automatically and continue operations. Leaving the system running at all times eliminates unnecessary handling of modules/system boards and reduces the possibility of early life failures on new components. In addition, this avoids power cycling, which is known to reduce availability of computer hardware.

3. Report error as per your standard service procedure, e.g., contact your Field Engineer (FE), Systems Service Engineer (SSE) or Solution Center.

4. Obtain and install appropriate kernel update patch according to Best Practices Guide.

   Summary of Instructions:

   - Consult Best Practices Guide for specific kernel update patch availability details.
   - Download kernel update patch from Sunsolve.
   - Read "README" file.
   - Install according to your standard installation process.
   - Reboot system.
   - Or, contact your FE, SSE or Solution Center for this process.

   Concurrently, as both a corrective and preventive measure, implement the Data Center Environmental Best Practices outlined in this guide.

The above steps should enable customers to meet or exceed system design goals for availability.

5. If, following implementation of kernel update patch and environmental standards, another error occurs on the same module within 6 months:

   - Replace the affected component, i.e., module or system board.
   - Follow hardware handling best practices (outlined in this guide) when performing any service, including:

     - Proper electrostatic discharge (ESD) processes.
     - Appropriate oil seepage management.
     - Hardware handling should be minimized and should only be performed by trained personnel.
     - Reduce handling of all components to a minimum. Therefore, mass replacement of any system components is strongly discouraged as a technique to improve system availability.

Outline: E−Cache Parity Error Best Practices for Mirrored SRAM Modules

Note: While mirrored SRAM modules will eliminate e−cache parity errors that are caused by soft errors in the SRAMs, it is still possible to have Solaris report an e−cache parity error on a module with mirrored SRAM. On the rare chance this does occur, the best practice is to replace the module.

1. Log error according to your standard error−logging procedures.

2. Upon correctly identifying an e−cache parity error, replace the affected component, i.e., module or system board.

3. Report error as per your standard service procedure, e.g., contact your Field Engineer (FE), Systems Service Engineer (SSE) or Solution Center.

4. Follow hardware handling best practices (outlined in this guide) when performing any service, including:

    - Proper electrostatic discharge (ESD) processes.
    - Appropriate oil seepage management.
    - Hardware handling should be minimized and should only be performed by trained personnel.
    - Reduce handling of all components to a minimum. Therefore, mass replacement of any system components is strongly discouraged as a technique to improve system availability.

5. Sun Enterprise Services personnel please reference item 3.2 in Sun Resource Appendix for information on returning mirrored SRAM modules.

1) Log Error

Overview

The best indication of how well an Enterprise server has run is its historical interrupt data. Verifying that accurate interrupt data is collected, logged and being provided to Sun will shorten the diagnosis and resolution times for problems.

Additional enhancements are being integrated into Solaris that will increase the amount of error reporting when an error occurs. This will allow for more timely and accurate error diagnosis. Incorporating the latest Solaris patch will enable customers to take advantage of this feature.

Process

1. Identify Error

A) Confirm that error is an e−cache parity error:
The data that is required to make the proper diagnosis and define corrective action is, at minimum, the full Solaris panic string. Most important is the Asynchronous Fault Status Register (AFSR).

If you have applied the kernel update patch described in Section 4, the message will appear like this:

WARNING: [AFT1] EDP event on CPU1 Data access at TL=0, errID 0x00000136.134f3a03
AFSR 0x00000000.80408000<PRIV,EDP> AFAR 0x00000000.0f489480
AFSR.PSYND 0x8000(Score 95) AFSR.ETS 0x00 Fault_PC 0x7801aa84
UDBH 0x0000 UDBH.ESYND 0x00 UDBL 0x0000 UDBL.ESYND 0x00
panic[cpu1]/thread=30000d2c580: [AFT1] errID 0x00000136.134f3a03 EDP Error(s)

If you have not applied the kernel update patch, the message will appear like this:

panic[cpu1]/thread=3000225bcc0: CPU1 Ecache SRAM Data Parity Error: AFSR 0x00000000.80408000 AFAR 0x00000000.0bd83bd0

Sun personnel: Consult Sun Resource Appendix 1.0 for information on a web page that compares the specific error conditions.

NOTE: Do not assume that a fatal reset error implies that the system board is defective. Some of the fatal reset errors could be caused by e−cache parity tag errors, which are not e−cache errors. Any fatal reset error in the Enterprise 220R−450, Netra t 1400/1405, Enterprise 3000−6000 and 3500−6500 servers should be evaluated to determine whether it is because of an e−cache error. Sun personnel should refer the Sun Resource Appendix 1.01 for evaluation tool information.

B) If the error is an e−cache parity error, identify the CPU module experiencing the error:
Sun Enterprise Services personnel: Consult Sun Resource Appendix 1.02 for tool information.

b1) In some cases, the error message is not sufficient to diagnose to a single CPU module. If this is so, it is recommended that no CPU module be replaced, as unwarranted module swapping has been proven to degrade machine reliability.

C) A situation where a domain or system has e−cache parity errors that are not specific to a single UltraSPARC−II CPU Module has to be examined closely. Contributing factors such as system load, environmental factors (temperature and humidity), etc., must also be examined to determine if there is an external contributor to these failures.

2. Log error

For Enterprise 3000−6000 and 3500−6500 servers, implement a console logging mechanism to allow all console messages, including the panic messages, to be recorded.

For all Enterprise servers, verify that savecore is implemented by running core_check.sh, which can be run on any SPARC system running Solaris 2.5.1 through 8. This script checks all known configuration deficiencies which would prevent a system from capturing crash dumps. The script also provides recommended corrective actions if deficiencies are found. In addition, it can be used regularly via the system utility cron to monitor the system's ability to capture crash dumps. Sun personnel: see Sun Resource Appendix 1.1 for script info.

For Enterprise 3000−6000 and 3500−6500 servers, the firmware needs to be at least at OBP (Open Boot PROM) version 3.2.25, which corresponds to patch level 103346−25 or later. This will ensure that the panic messages appear in /var/adm/messages after a reboot. The OBP version can be determined by using the command "prtdiag −v".

For any error, gather relevant information to substantiate diagnosis, e.g., core files, Explorer, arbstop/recordstop files. Immediately upload core files and Explorer to SunSolve and inform the Solution Center to move the files to an internal core server. For Enterprise 10000, also ensure that the Explorer is from the system service processor (SSP).

Please note that obtaining the system serial number (S/N) is an important piece of data because it is the "key" that allows us to join all our database tables pertaining to shipped systems and FRUs. If we have this piece of data, we can track systems, FRUs in the system, HOSTID of the system, ship−to address, etc. It is the "key" that ties the corporate databases together.

For all Enterprise servers, verify that the primary swap device (first swap device listed via swap −l and known as the "dump device") is large enough to save a proper vmcore file. This size will vary depending on the amount of main memory and the number of applications running on the system.

In most cases this size of the dump does not need to exceed 2GB. The following table describes the recommended patches for savecore. No specific patches are required for later versions of Solaris.

Table 1 − Solaris savecore Patches

Solaris Version	*savecore Patch*
Solaris 2.5.1	108083−01
Solaris 2.6	107490−01

For all Enterprise systems verify that the appropriate SSP patch or firmware update is applied. These cause panic messages to be properly preserved in the unlikely event of a system crash.

For Enterprise 10000 systems (both mirrored and non−mirrored SRAM modules) the patches are:

SSP 3.1.1: 108638
SSP 3.2: 108676

This will ensure that the panic messages are logged in the netcon console log, where they can be retrieved by tools such as Explorer. Netcon logging must be enabled up to SSP 3.3.

In addition to the above, for Enterprise 10000 systems running mirrored SRAM modules, the following SSP patches need to be current:

	T−Patch #	Released Patch #
SSP 3.1.1	108536−04	108536−05
SSP 3.2	108543−04	108543−05
SSP 3.3	108885−04	108885−06
SSP 3.4	N/A	110304−02

3. Track Error

Overview

Sun's Support organizations should continue to capture the details of every e−cache error as trend analysis and investigation continues.

Process

Sun personnel should immediately track error via Sun Resource Appendix 1.2.

For any error, immediately upload core files and Explorer to SunSolve and inform the Solution Center to move the files to an internal core server. For Enterprise 10000, also ensure that the Explorer is from the system service processor (SSP).

Reason

This data gathering process is the only place from which we will compile the failure information. All other forms of e−cache data gathering will be migrating to the Best Practices data gathering template.

The data is included with the Solaris panic string and contains detailed information regarding many types of errors. This AFSR data, in combination with historic failure information for that system, is used by Sun to recommend the best corrective action to take.

Results

Proper identification of any given error as an e−cache error is the essential first step which needs to be taken in order to know if specific follow−on corrective action needs to be taken. Misdiagnosis of error can result in lost time, effort, and materials used to address it. Proper tracking and analysis has enabled Sun to identify and implement best practices for error correction.

2a) Do Not Replace CPU Module Upon First E−cache Error (for Non−Mirrored SRAM Modules)

Overview

Sun has determined that the majority of e−cache parity errors are transient and will not reoccur on that same CPU module. Therefore, replacement of a module after a single error is unnecessary and will result in unneeded downtime and component handling and other issues which have a negative affect on availability. Should the module experience a second such failure within six months, the error is no longer considered transitory and Sun recommends replacement.

Action

1) Upon first identified e−cache parity error, reboot the system (if needed) and continue normal operations.

2) The error should be recorded for that system (see Part 1, Item 2 − Log Error) and tracked over time.

Reason

E−cache parity errors are triggered by a unique set of circumstances involving a number of different factors, i.e., temperature, humidity, slot, process running, state of registers, etc., that are unlikely to occur in that combination again.

Results

Sun's analysis of modules returned after experiencing a single e−cache parity error found that over 90% showed NTF (No Trouble Found). This is because the modules do not have a design or manufacturing defect. The vast majority of customers who experience a single e−cache failure on a module do not experience it on the same module again.

2b) Replace CPU Module Upon First E−cache Error on Mirrored SRAM Modules

Action

1. Log error according to your standard error−logging procedures.

2. Upon correctly identifying an e−cache parity error, replace the affected component, i.e., module or system board.

3. Report error as per your standard service procedure, e.g., contact your Field Engineer (FE), Systems Service Engineer (SSE) or Solution Center.

4. Sun Enterprise Services personnel please reference item 3.2 in Sun Resource Appendix for information on returning mirrored SRAM modules.

Reason

While mirrored SRAM modules eliminate e−cache parity errors that are caused by soft errors in the SRAMs, it is still possible to have a hard failure generate an e−cache parity error on a module with mirrored SRAM. On the rare chance this does occur, the best practice is to replace the module.

3) Report the Error (All Modules)

Use your standard service procedure as you would for any other service issue (e.g., contact your Field Engineer, SSE or Solution Center). Field Engineers and the Solution Center will be armed with the latest version of this best practices guide, and can help with implementation questions.

As part of Sun's ongoing quality improvement efforts, all errors reported will be consolidated into a central database to give better tracking, feedback, and corrective capabilities.

See section one for more details on electronically logging and tracking errors.

4) Install Appropriate Solaris Kernel Update Patch

Overview

In addition to the upgrade described in the error analysis section, Sun has developed software which is part of the Solaris kernel that can pro−actively inspect data in the cache and in most cases "scrub" any identified soft parity errors before they hit the system.

Process

Since the Solaris improvements have the potential to greatly reduce the incidence of e−cache parity errors, it is Sun Engineering's conclusion that all customers on Solaris 2.5.1, Solaris 2.6, Solaris 2.7 & Solaris 2.8 should install the kernel update patches where possible.

Please note: The Solaris 2.5.1 kernel update patch contains only the cache scrubber. Solaris 2.6, Solaris 2.7 and Solaris 2.8 contain all improvements (cache scrubber, error messaging, handling and reduction).

Sun personnel see Sun Resource Appendix 2.4 for Field Information Notice (FIN) information.

Solaris servers that run Veritas Cluster Server (VCS) below version 1.3.0 should reference FIN information in Sun Resource Appendix 2.4.1 prior to applying the KU patch.

Keep patches up to date:

Sun constantly monitors Solaris OS performance in a variety of areas. When areas of improvement are identified, Sun releases a patch to make this improvement quickly available to the user base. The customer best practice is to work with your account team to ensure all patches are up to date and installed.

For Enterprise 10000:

> All SSP patches are considered mandatory and the SSP must not be allowed to fall behind on patches. Sun personnel should refer to the Enterprise 10000 patch page at Sun Resource Appendix 2.5.
>
> Ensure patches, FINs and FCOs are applied in a consistent and proactive manner in accordance with customer Change Management procedures. Sun personnel can view a query report on the required FINs and FCOs for the machine type at Sun Resource Appendix 2.6.

Two specific enhancements include:

- Error Reporting: Additional enhancements integrated into Solaris increase the amount of error reporting when an error occurs. This will allow for more timely and accurate error diagnosis.

- E−cache scrubber: The implementation of a kernel level e−cache scrubber has been integrated into Solaris versions 2.5.1, Solaris 2.6, Solaris 2.7 and Solaris 2.8 which allows Solaris to pro−actively access or scrub each e−cache location with the goal of reducing the chance of a system panic resulting from a soft error.

Reason

To reduce the likelihood of e−cache data, write−back or copy−out parity errors, a "cache scrubber" has been implemented that periodically flushes modified cache lines out to main memory and invalidates cache lines that have not been modified. By reducing the likelihood that an otherwise nonfatal error in the e−cache will result in a system failure, this procedure improves system reliability.

Results

- Improved Error Handling
  Each error reported by the CPU module is now evaluated to determine whether it is fatal to the operating system, only fatal to a user process, or of no immediate consequence. Fatal errors in the kernel result in a system panic, as they did before. Fatal errors within the user process will now cause the machine to reboot instead of panic, allowing file systems to be fully synched and also preventing the creation of unnecessary kernel core files. Events that do not affect the integrity of either the kernel or user processes are logged, but otherwise ignored. Because UltraSPARC−IIi and UltraSPARC−IIe use simplified error reporting logic (as compared to UltraSPARC−II), the error handling behavior for UltraSPARC−IIi and UltraSPARC−IIe based systems has not been changed. Those systems will still panic on most CPU, e−cache, or uncorrectable memory errors.

- Improved Error Messages
  The CPU, e−cache and memory error messages have been changed to be more accurate and complete. Text descriptions have been rewritten to emphasize the important parameters associated with each event. Also, the logic for reporting hardware errors has changed to ensure that error events are reported accurately, completely, and in the order they occurred. These new error messages will make it easier to determine the CPU that has encountered an error. There are related patches to SunMC will allow it to recognize the improved error messages; without them, the management console will under−report the occurrence of corrected main memory errors.

- Error Reduction
  Lab and beta field testing has shown an improvement factor of about 2x in module reliability for systems using the cache scrubber vs. no scrubber. Individual system results vary widely based on percent of time CPU is idle, environment, configuration, application running and other variables.

Implement Data Center Environmental Standards (Ongoing)

Overview

In order to enhance system availability, Sun has created a series of environmental best practices for operating a data center. These best practices call for tighter tolerances for environmental specifications than those previously described by Sun Microsystems.

For Sun Enterprise Services personnel, additional information regarding optimal environmental practices can be found at Sun Resource Appendix 2.0.

In order to ensure data center readiness for Enterprise 3000−6000, 3500−6500 and Enterprise 10000 servers, a site audit should be completed for verification of acceptable environmental conditions. Below is a listing of the environmental elements to be reviewed.

Process

Follow the instructions in the "Details" section below, including those on temperature and humidity, flooring, system placement, cleanliness, and monitoring.

Where appropriate, schedule audits of Best Practices and environmental reviews for sites experiencing problems. Customers should contact your local ES Account Manager if an audit is required. Sun personnel should follow the escalation path detailed in the Best Practices Sun Resource Appendix.

Reason

Proper implementation of environmental best practices has been confirmed to reduce the occurrence of e−cache errors.

Results

Customers who have implemented environmental best practices following the occurrence of e−cache parity errors have shown a decline in those errors in addition to those seen as a result of any hardware or software improvements

DETAILS:

Temperature & Humidity

The ambient temperature and relative humidity levels where the Enterprise Servers are located should be checked using an electronic thermo−hygrometer. Taking readings from monitoring equipment (mounted somewhere within the data center) or from air conditioning units may provide incorrect data if the data center environment is not correctly maintained.

Optimum ambient temperature: 22 C/72 F.

- Operating specification for the Enterprise Server Range may seem quite wide, (5−40C/41−104F for Enterprise 220R, Enterprise 420R, and Enterprise 450 and Netra t 1400/1405; 5−35C/41−95F for Enterprise 250; 0−40C/32−104F for Enterprise 3000−6000 and 3500−6500, 10−30C/50−86F for Enterprise 10000). This is because most hardware is designed to operate in a wide range of conditions. While these machines may be able to operate within this wide range, more stringent control over environmental factors is necessary if optimal reliability is to be achieved. Therefore, if a high level of reliability is required, then optimal conditions should be maintained.

Maximum rate of change of temperature is 5.5 C/10F in any 60 minute period of operation.

- This optimum level is desirable because it is easier to maintain safe associated relative humidity levels and there is an acceptable wide operational buffer in case of environmental support systems down−time.

Optimum relative humidity: 45% to 50% Rh.

- The operating specification for the Enterprise Server Range is 20% Rh to 80% Rh, but like temperature, the optimum levels of relative humidity must be obtained to improve system reliability.

- This optimum range helps protect hardware from corrosive problems associated with high humidity levels and failures caused by static discharge when humidity is too low. It also provides the greatest operating time buffer in the event of environmental support system failure.

Maximum rate of change of relative humidity is 10% in any 60 minute period of operation.

- When positioning Sun servers in proximity to air handling equipment, ensure that the rate of change of temperature and humidity is consistent with the above.

Subfloor temperature is not to exceed 15.5C/60F.

- Assuming appropriate ambient temperatures are being maintained and there is at least a moderate heat load in the data center, subfloor temperature should generally register near 15.5C/60F.

- This is also the optimal temperature recommended by manufacturers of hardware that utilize direct subfloor cooling.

Table 2 − Optimal Temperature and Humidity Settings

	Optimal Setting	***Maximum Delta***	***Time Period for One Delta***
Temperature	22C or 72F	+/− 5.5C / 10F	60 Minutes
Humidity	45% to 50%	+/− 10%	60 Minutes

The above temperature and humidity levels should be measured and maintained at the **INTAKE** of each system. Simply setting these levels on air conditioner/air handler control panels is normally not sufficient in most larger data center environments. The settings need to be adjusted so that these levels can be achieved at the various system locations.

Flooring

Good subfloor pressure (if used) is 0.05 inch wg/1.3mm wg (water gauge).

- The pressurization of the subfloor void in relation to the above−floor hardware space should be maintained at the optimum level of 0.05 inch wg. If this level is maintained, the conditioned air provided by the environmental support system is delivered to all hardware within the data center. Pressurization should not fall below a level of 0.02 inch wg (0.5mm wg).

Appropriate subfloor tile positioning.

- Air distributed tiles should be placed in such a fashion so as to deliver conditioned air to the intake of each hardware cabinet.
- The Enterprise 10000 utilizes ambient cooling; but it also intakes air from the bottom of the machine, and it is recommended that perforated tiles are placed directly under the machine and in front and at the rear for optimal performance.
- Ensure that large and unobstructed perforations are present under the Enterprise 10000 to enable adequate airflow.

No zinc whiskers present from old subfloors.

- Galvanized tiles are coated with zinc to help protect them from corrosion. If these tiles used the electroplated process, there is a possibility of zinc whisker growth appearing over a period to time.
- Zinc whiskers are conductive contaminants and can actually cause short circuits on components.

System Placement

Enterprise systems can be "rack mounted" inside system racks or cabinets or they can be placed on common shelves. For systems which are not rack mounted, physical placement is important. The critical items to consider when placing any of these stand−alone systems are airflow and ambient temperature. This should not be an issue for Volume Server Products (Enterprise 220R−450 and Netra t 1400/5) since they have front to back cooling.

Inlet and exhaust airflow must be taken into account when these systems are rack mounted or placed on standard shelves. Many of these systems, and some Sun storage devices, are designed to take in cool air on the right side of the system (facing system front) and then exhaust heated air on the left side of the system. By placing these systems or these storage devices at the same vertical level, the normal cooling process can be significantly impacted. The following is a basic drawing of a poor configuration.

Figure 1 − Poor System Placement

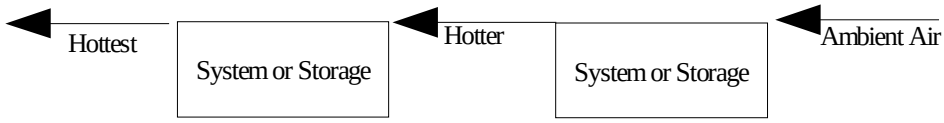

A better solution is to alternate the placement of these side air flow systems. This allows the exhaust of one system to cool to a sufficient level before it is used as incoming air to the second system. As noted, an optional baffle can be installed to help divert heated air if the required distance between the systems is not sufficient.

Figure 2 − Better System Placement

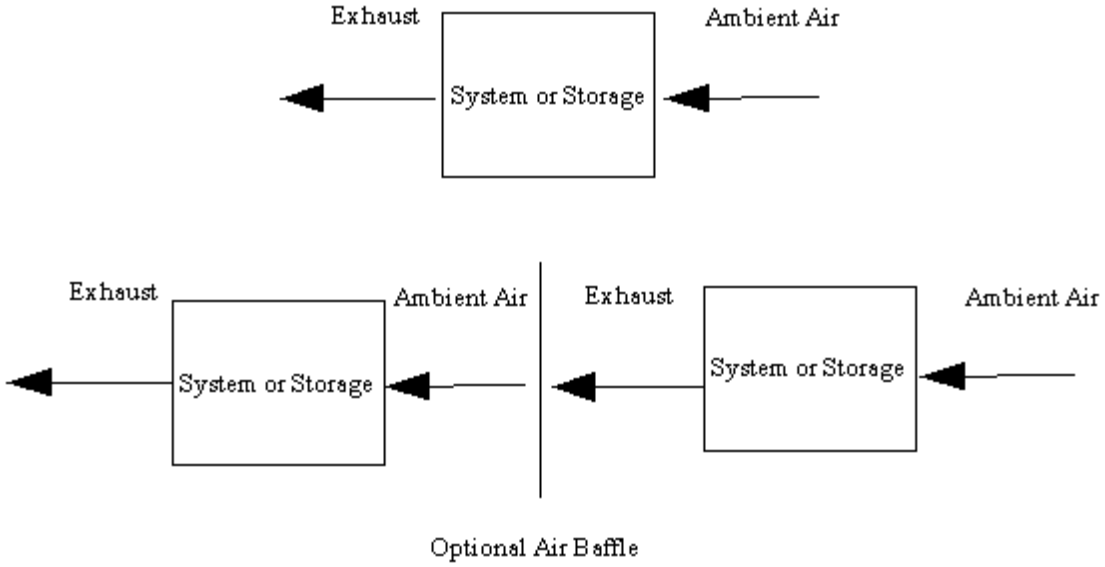

For Enterprise servers that are already rack mounted (i.e. E5x00, E6x00 and Enterprise 10000), verify the following items:

- Do not mount additional hardware (neither Sun nor 3rd party) inside these cabinets without consulting your local Sun Enterprise Services team.
- Verify that perforated/vented floor tiles are placed in the recommended locations for the system type. Again, consult Sun for further information.

The key for any system placement or rack mounting solution is to verify that the **INTAKE** air is at the recommended temperature and humidity levels for **ALL** systems.

Additionally, watch for obvious external influences such as close placement to air conditioners, elevator shafts, industrial equipment or other sources of vibration, particulates or substantial electromagnetic interference (i.e., heavy electrical equipment).

Cleanliness

Good data center cleaning practices minimize air contamination. Even a perfectly designed data center will require continued maintenance. An effective cleaning schedule must consist of specially designed short−term and long−term actions. These can be summarized as follows:

Cleaning Schedule:

Frequency	Task
Daily Actions	Trash removed
Weekly Actions	Access floor maintenance (vacuum and damp mop)
Quarterly Actions	Hardware decontamination and room surface decontamination
Bi−Annual Actions	Subfloor void decontamination and air conditioner decontamination (as necessary)

Filter changing practices and appropriate frequency:

- The Enterprise 10000 has been designed for maximum availability such that the air filters can be changed without the need to power down the system. However, this does introduce some risk to the integrity of the I/O cables, etc., as the cables must be moved in order to access the air filters. In systems heavily configured with many I/O lines, more risk is introduced. As such, the decision to power down the system is left to the discretion of the customer/system service engineer (SSE) whether the Enterprise 10000 is taken off−line or powered down for the filter service.

- The Enterprise 10000 system has four front and four rear air filters that require periodic inspection and replacement. To prevent restricted airflow and possible equipment failure, the Enterprise 10000 air filters should be inspected for debris and trapped particles every three months of operation. Filters should be inspected with a mirror prior to removal. The level of debris found on the filters and surrounding area should be considered in the decision of whether to power the machine down for air filter service. If the machine is left running during a filter change, do not insert any air blocking panel in its place, even for a few seconds. Instead, clean the area around the air filters as best as possible, then make the replacement quickly and carefully. It is considered less risk to pass unfiltered air for a few seconds than it is to stop air flow on an operational system, as individual components can overheat in as little as 10 seconds.

- The actual service life of the air filter is dependent of many variables including the cleanliness of the supply air, location within the data center, etc. The decision when to change filters is left to the judgment of the SSE. When the level of debris is found to be substantial, the filters should be replaced with new filters. Note that it is no longer recommended to clean the old filters as they deteriorate with time and are easily damaged.

- If the air filters are found to collect a substantial amount of debris in less than three months, it is recommended that the air supply system be investigated for sources of contamination and corrective action taken. In addition, Account Management should be notified and updated as required. Sun personnel can access the Enterprise 10000 Service Manual via Sun Resource Appendix 2.1 for additional information.

- It is advisable to have a spare set of filters on−site so that replacements are available when needed.

Ensure a clean AC power source (neutral/earth). The use of Dranetz−BMI analyzers, for example, should only be used by those people who are well−trained on this equipment, and should be part of a more detailed site audit.

Dust and Particulate Control:

- Most data center contamination comes from sources such as unplanned maintenance or construction due to expansion. Also the lack of incoming air quality maintained via air conditioning/air handling units is also a common contributing factor.

- Particulate control is normally achieved via a combination of maintaining general cleanliness in the data center itself, as well as verifying that incoming air to the data center is clean.

- Both above floor and sub−floor cleanliness should be examined and periodic cleaning of these locations performed. Also, verify that all incoming air to the data center environment is filtered or is clean.

- Most air conditioners/air handlers have filters or filtering capabilities. These filtering capabilities should be used and items such as air filters, etc. should be replaced on a periodic basis. This frequency is normally dictated by a combination of the air conditioner/air handlers maintenance manuals as well as the quality of the incoming air.

Monitoring

Optimum Enterprise 10000 CPU temperature is not to exceed 55C/131F.

- This should be achieved if Enterprise 10000 filters are clean, ambient temperature is at the optimum level, there is adequate cool air flow and environmental support systems are maintained.

Optimum Enterprise 3000−6000 and 3500−6500 CPU temperature is not to exceed 50C/122F.

- This should be achieved if ambient temperature is at the optimum level, there is adequate cool air flow and environmental support systems are maintained.

Monitoring and control for temperature and humidity:

- Accurate and comprehensive monitoring of environmental support equipment and in−room conditions is extremely important in an environment as complex and sensitive as a computer data center. The monitoring system should have historical trend capabilities. The data gleaned from analysis of historical psychrometric information can be instrumental in determining seasonal changes or other outside influences.
- A psychrometer is a hygrometer consisting of two similar thermometers with a bulb of one being kept wet so that the cooling that results from evaporation makes it register a lower temperature than one dry one. The difference between the readings constitutes a measure of dryness of the atmosphere.

If the air conditioning unit uses chilled water system (modulated controls), the hardware placement in relation to the air conditioning units is less of an issue since there is typically less dramatic fluctuation caused by air conditioning cycling. A minimum of eight feet is acceptable, although 12 feet is a safer distance.

When compressurized air conditioning systems (water, air, glycon cooled, etc.) with staged cooling (0%, 25%, 50%, 75%, 100%, or 0%, 50%, 100%) are used, these have greater concerns in relation to hardware placement. When these systems are "seeking" between stages, significant (as much as 5 degrees F & 15% Rh) repeated mechanical cycling is observed. In data centers where these types of conditioners are used, place the hardware at least 12 feet away as a minimum distance and 16 feet optimally. It is also recommended that the hardware be located out of the direct airflow of the air conditioners.

Ensure that correct machine grounding procedures are followed. Do not ground Enterprise 3000−6000 and 3500−6500s. For Enterprise 10000, the key component for grounding is solid, stable earth ground at the customer power distribution unit (PDU). For additional information regarding grounding on Enterprise 10000, Sun personnel should go to Sun Resource Appendix 2.2.

Check the earth bonding on the system cabinet and storage expansion cabinets. If a ground strap from the system cabinet and storage expansion cabinets to earth ground is being used, then it is very important to check that the power receptacle (IEC309) ground and the grounding point of the strap (i.e., raised floor ground) are both tied to the same earth ground. Grounds should never be daisy−chained. ESD mats and wrist straps must be used and tested regularly. Sun personnel should check Sun Resource Appendix 2.3 for additional information on ESD test equipment.

Cell Phones in Data Centers

It is suggested that handheld and other portable radio transmitters be kept outside of an 18" area around operating computer systems.

5) Module Type (Mirrored versus Non−Mirrored) Determines when to Replace

Overview

The actions below should be followed upon first occurrence of an e−cache error on a mirrored SRAM module. For non−mirrored SRAM modules, please follow the actions below only if a second e−cache error occurs on the same CPU module within six months. A a second error would suggest that the error is not transient.

Action

- Replace the affected component, i.e., module or system board. If the system is an Enterprise 3000−6000 or 3500−6500, please ensure that the system board is completely locked. Leaving the boards unlocked can lead to system reliability issues. In addition, suspect boards should be examined for sheet metal damage on the front panel−to−board ears. This type of damage can be caused by attempts to extract boards while the locks are still engaged.
- For Enterprise 10000, blacklist the module (if necessary) then replace the affected component, i.e., module or system board.
- For Enterprise 220R−450, replace with appropriate 400, 450, or 480MHz module. For Netra t 1400/1405, replace with appropriate 400MHz module (#501−6209).
- Sun Enterprise Services personnel are encouraged to review system board removal and module installation information provided in Sun Resource Appendix 3.1.
- All 400MHz UltraSPARC processor modules with 8−MByte of mirrored or non−mirrored SRAM are compatible with each other and can be combined in the same system, domain, and on the same system board.
- Do not retorque or reseat modules as an alternative to module replacement.
- Sun Enterprise Services personnel must update the website noted in Sun Resource Appendix 3.0 to reflect any hardware adjustments.
- For module failure analysis information see Sun Resource Appendix 3.2

  Once boards/modules are installed, if possible burn−in the system board and module assembly running Sun VTS and perform engineering recommended tests for one week in the destination system. If one week is not possible, burn−in system board, modules and memory as a whole unit for as long as is practical before putting replacement in production. After the first week, additional burn−in time will return a 10% error reduction per week through the first month.
- For information regarding ordering the CPU Module Installation Procedure training video see Sun Resource Appendix 3.3.
- Sun Enterprise Services personnel, please see Sun Resource Appendix 3.4 for cache size questions.

- For additional availability needs, Sun Personnel see Sun Resource Appendix 3.7.

General Best Practices for Hardware Component Handling

System availability can be enhanced by following proper component handling processes. Although failures caused by handling or installation issues normally appear shortly after the CPU modules are installed, some errors can be latent. Stress to connectors and other components can occur and not be immediately visible via diagnostics or system errors. Over time, temperature changes, power cycling, etc., can cause these stress issues to become either intermittent or solid errors.

Using the correct hardware replacement processes is critical. Customers should **NOT** handle Enterprise 3000−6000, 3500−6500 and Enterprise 10000 system hardware. Please work with your Enterprise Services team for all hardware replacement or augmentation issues.

Electro−Static Discharge (ESD)

System availability can also be enhanced by avoiding the ill effects of ESD. The Enterprise servers use state−of−the−art components and, as such, are susceptible to the effects of ESD. Although most of the system hardware is modular and easy to physically handle, the components are still sensitive to ESD. Ensure that items such as personal grounding straps and a grounded work environment via the use of an approved ESD mat are used for **ALL** component replacements. Check the earth bonding on the system cabinet and storage expansion cabinets. If a ground strap from the system cabinet and storage expansion cabinets to earth ground is being used, then it is very important to check that the power receptacle (IEC309) ground and the grounding point of the strap (i.e., raised floor ground) are both tied to the same earth ground. Correct ESD and handling practices must be used when handling spares. ESD mats and wrist straps must be used and tested regularly.

Preventive action: Oil seepage

Any possible oil seepage should be cleaned only if modules are replaced.

If a CPU module replacement is necessary, check the system board for oil contamination seeping from the thermal pads that cool the underside of the CPU module. For additional information regarding oil seepage, Sun personnel see Sun Resource Appendix 3.5.

For Enterprise 10000, remove the oil contamination with solvent and lint−free wipes as detailed in Sun Resource Appendix 3.6 and install the new CPU assembly or system board. Only clean the area under the CPU module experiencing the error. Do not remove and clean under other hardware.

This oil is non−conductive and will not interfere with machine operation if the module is not removed. However, if a module is replaced oil residues may lead to particulate collection at connection sites which could cause errors.

Torquing Guidelines

Do not routinely retorque or reseat CPU modules or reseat memory DIMMS.
Do not retorque or reseat CPU modules on new or pre−configured (CTO) system boards.
Do not retorque or reseat CPU modules as an alternative to module replacement−−use error messages and procedures as guidelines.

Minimizing replacement−related interrupt

Steps can be taken to minimize the interrupt when a component requires replacement. For those cases when the Power On Self Test (POST) does not fail a suspect CPU Module, the CPU Module can be disabled and the system restored to operation quickly. A service action can be scheduled days or weeks in advance during a scheduled maintenance window to minimize a second interrupt for the service action.

The following are some actions that can be taken to ensure system stability without the need for an immediate service action:

- Use the following psradm(1M) command from Solaris to disable the CPU Module. For restrictions, read the manpage of psradm.

 # psradm −f <cpu_number>

For Solaris 2.7 systems and higher, also disable I/O interrupts.

 # psradm −i <cpu_number>

Note these commands can be scripted as a reboot will put the CPU module back on−line.

- For Enterprise 10000 systems, the blacklist(1M) process on the SSP can also be used. Note that a reboot using the bringup(1M) command must be issued after any changes are made to the blacklist file for those changes to become effective.
- If the Enterprise 10000 configuration will support this, consider the use of Dynamic Reconfiguration (DR) to remove suspect hardware for replacement while the domain continues operation.

General

Minimizing the amount of hardware that is replaced in a system also reduces the chance for early life failures, which are always a possibility, despite the rigorous testing at Sun production and assembly centers.

Reason

If the error has been identified as non−transitory (or "hard"), replacement of the module/board should improve availability.

Results

For those customers experiencing a "hard" failure updates, module/board replacement has resulted in substantial improvements of module reliability (when combined with environmental and software upgrades).

Summary of Best Practices Guide

Overview

The following is a summary of the key components of Sun Microsystems Best Practices Guide Addressing E−Cache Parity Errors.

Table 3 − Summary of Sun Microsystems E−Cache Best Practices

Best Practice	***Best Practices Summary***
Environmental	
Maintain consistent temperature	Maintain consistent temperature at system locations.
Maintain consistent humidity	Maintain consistent humidity at system locations.
Minimize dust/particulates	Maintain a clean data center environment.
Perform regular filter maintenance	Clean/replace filters on air handling units and on Enterprise 10000 systems.
Hardware Handling	**NON−Mirrored SRAM Modules**
Replace CPU module only if second e−cache error occurs	Upon first occurrence of e−cache error, DO NOT replace module. Implement environmental Best Practices and install cache scrubber software. If a second ecache error occurs, replace the affected component.
Hardware Handling	**Mirrored SRAM Modules**
Replace CPU module upon first occurrence of e−cache parity error	Upon first occurrence of e−cache error, replace the affected component. Implementation of environmental Best Practices is still appropriate.
Hardware Handling	**All Modules**
Do not perform mass replacements	Reduce early life failures and hardware handling.
Do not handle hardware	Consult Enterprise Services for hardware maintenance and augmentation.
Observe ESD procedures	Verify that ground straps and mats are used.
Pro−active FIN, FCO & patch process	Consult Sun for current FIN/FCO/Patch recommendation.
Disable suspect hardware on−line	Use Solaris/SSP commands to disable suspect hardware.
System Placement and Positioning	
Multiple server & storage	Observe air flow and cooling considerations.

Best Practice	***Best Practices Summary***
Rack mounting servers or storage	Consult Sun when rack mounting hardware.
Stand−alone server placement	Correctly position systems in relationship to vents/tiles.
Error Logging	
Saving system core dump	Verify that system is correctly configured.
Develop a log for all interrupts	Track all interrupts to validate system stability.
Capture CPU AFSR information	Capture AFSR information to ensure correct diagnosis.
System console logging	If possible, implement a console logging solution.
Report the error	Sun Personnel should immediately track error via Sun Resource Appendix 1.2.
Ensure all patches are up to date	Ensure that all savecore patches as well as SSP patches are up to date.