Auto Diagnosis and Recovery Enhancements for Sun Fire™ Midrange Servers
Updated for Firmware Release 5.19.0

Tricia Wittsack, PTS Engineering Group

Sun BluePrints™ OnLine — August 2005

Part No 819-3729-10
Revision 1.0, 8/10/05
Edition: August 2005

Please
Recycle

Table of Contents

Auto Diagnosis and Recovery Enhancements for Sun Fire™ Midrange Servers

Beginning with firmware release 5.15.0 for the System Controller (SC), several enhancements were made to improve the availability, serviceability, diagnosability, and repair characteristics of Sun Fire™ midrange servers. These enhancements provided in the system controller firmware, combined with enhancements to the Solaris™ Operating System (Solaris OS), implement auto diagnosis and recovery capabilities that can increase system uptime, decrease system outages, improve system resiliency when a hardware fault occurs and minimize service interruptions. These enhancements automate many processes and procedures which required human intervention prior to firmware version 5.15.0. Firmware version 5.19.0 and the appropriate Solaris OS with all relevant kernel updates and patches are required to fully benefit from these enhancements. See Table 1 on page 2 for specific system requirements.

This document is useful for support personnel and assumes a basic technical knowledge of the Sun Fire midrange servers. It addresses the following topics:

- “System Controller Firmware Enhancements” on page 4 describes the firmware release 5.15.0, 5.15.3, 5.17.0, and 5.18.0 enhancements, as well as the recent firmware release 5.19.0 enhancements.
- “Solaris OS Enhancements” on page 15 describes the enhancements to the Solaris Operating System that aid in auto diagnosis and recovery.

Auto Diagnosis and Recovery Overview

The auto diagnosis engine in the system controller performs run-time analysis of faults and provides standardized messages to easily identify the faulty component or Field Replaceable Unit (FRU). Failure data is persistently stored in the component health status (CHS) of the identified unit. This helps ensure that faulty units are not configured into the system during future system reboots and keyswitch operation, and helps prevent faulty components or FRUs from adversely affecting future system operations. The automated recovery features remove faulty components and FRUs from the configuration and allow the system to be placed back into operation quickly without requiring human intervention.

Certain faults can only be diagnosed by the Solaris 8 OS or Solaris 9 OS with the appropriate kernel updates, or by the Solaris 10 OS. Updated properly, the Solaris OS can identify certain CPU and memory events, and can off-line the affected CPU or retire the appropriate page of memory. In most situations, system operation can continue uninterrupted for the specified Solaris OS events. The Solaris OS also updates the CPU’s CHS via the System Controller when attempting to off-line an identified CPU. This helps ensure that these identified CPUs are not configured into the system during future system reboots and keyswitch operations.

Minimum Required Firmware Release and OS Revision

The correct firmware release, Solaris OS revision, and applicable kernel updates and patches are required to fully benefit from the auto diagnosis and recovery enhancements in the Sun Fire midrange servers (see

Table 1). Firmware release 5.19.0 is required. In addition, the Solaris 8, 9, or 10 OS, with the noted kernel updates and patches, is required.

Table 1. Minimum firmware and OS versions required to support auto diagnosis and recovery features in the Sun Fire midrange servers.

Server	Required Firmware Release	Minimum Required Solaris 8 OS	Minimum Required Solaris 9 OS	Minimum Required Solaris 10 OS
Sun Fire™ 3800-6800 servers	5.19.0	Solaris 8 OS, kernel update 117350-24 with patches 116979-05 and 110369-08	Solaris 9 OS, kernel update 118558-06 with patches 117124-08 and 116009-03	Solaris 10 OS, kernel update 118822-03 and patch 119568-01
Sun Fire™ 4900-6900 servers	5.19.0	Solaris 8 OS, kernel update 117350-24 with patches 116979-05 and 110369-08	Solaris 9 OS, kernel update 118558-06 with patches 117124-08 and 116009-03	Solaris 10 OS, kernel update 118822-03 and patch 119568-01
Sun Fire™ V1280 and E2900 servers; Netra 1280 servers	5.19.0	Solaris 8 OS, kernel update 117350-24 with patches 116979-05 and 112162-04	Solaris 9 OS, kernel update 118558-06 with patches 117124-08 and 116009-03	Solaris 10 OS, kernel update 118822-03 and patches 119568-01 and 119570-01

Summary of Auto Diagnosis and Recovery Enhancements in the System Controller Firmware

To properly administer and service a Sun Fire midrange server, it is useful to understand which features are available in the different firmware versions. This document discusses the following features and enhancements in detail to enable proper planning and informed decision-making:

- Sun Fire server firmware release 5.15.0
  - Auto diagnosis — Automated diagnosis of runtime hardware faults.
  - Component health status — Persistent record of information stored in a component.
  - Auto restoration — Automatic restoration of a domain.
  - Domain hang recovery — Detection and recovery from a domain hang.
  - Repeated domain panic recovery — Power on-self test (POST) is run at increasing diagnosis levels to identify and isolate the faulty hardware (if any).
- Sun Fire server firmware release 5.15.3
  - Persistent record of hardware failures identified by the Solaris OS — System Controller receives hardware failure messages from the Solaris OS and stores the component health status (CHS) in the affected FRU.
- Sun Fire server firmware release 5.17.0
  - Persistent logging of certain System Controller error messages and message logs — If the platform is using enhanced-memory System Controllers (SC V2 components), the System Controller platform and domain logs and the error buffer survive a System Controller reboot, thereby ensuring events are available for analysis if a System Controller reboot occurs.
- Sun Fire server firmware release 5.18.0
  - Voltage Core Monitor (VCMON) — Proactive monitoring and analysis of CPU core voltage levels to identify CPU boards which may be at risk for experiencing certain identified types of failures.

- Sun Fire server firmware release 5.19.0
  - CPU Diagnosis Engine added to Auto Diagnosis — Uses information provided by the Solaris 10 OS Predictive Self-Healing to identify the affected component after a CPU fault is detected. For systems utilizing Solaris 8 OS or Solaris 9 OS, the CPU Diagnosis Engine diagnoses specific fatal and non-fatal CPU events without the benefit of additional Solaris OS diagnosis.
  - ECC Diagnosis Engine added to Auto Diagnosis — The ECC (Error Checking and Correcting) Diagnosis Engine correlates collected parity and ECC information from both the System Controller and the Solaris OS to detect bad writers and bad readers, diagnose signaling (Syndrome 0x71) events, and identify components which experience a specified number of ECC events in a specified time period.
  - New domain configuration parameter of `max-panic-diag-limit` — Recovery from repeated domain panics maximum POST diagnosis level is now tunable.
  - Processor CHS updated for voltage core monitoring (VCMON) events.

Table 2 lists the requirements for the various auto diagnosis and recovery features that are available for the Sun Fire midrange servers. For the minimum required Solaris OS kernel updates and patches, see Table 1 on page 2.

Table 2. Firmware requirements for auto diagnosis and recovery features.

	System Controller Firmware Version				
Feature	**5.15.0 - 5.15.2**	**5.15.3 - 5.16.0**	**5.17.0**	**5.18.0**	**5.19.0**
Auto diagnosis	√	√	√	√	√
Component health status	√	√	√	√	√
Auto restoration	√	√	√	√	√
Domain hang recovery	√	√	√	√	√
Repeated domain panic recovery	√	√	√	√	√
Persistent updating of CHS by the Solaris OS for CPU off-line for L2_SRAM conditions		√[a]	√[a]	√[a]	√[a]
Persistent Logging in the SC Platform Log and Domain Logs			√	√	√
Voltage core monitoring (VCMON)				√	√
CPU Diagnosis Engine					√
ECC Diagnosis Engine					√
`max-panic-diag-level` parameter for repeated domain hang recovery					√
CPU CHS updated for VCMON event					√

a. This enhancement requires Solaris 8 OS kernel update 24 (108528-24); or Solaris 9 OS kernel update 09 (112233-09) with patch 116009-01; or Solaris 10 OS.

Summary of Solaris OS Enhancements

Additionally, enhancements have been made in the Solaris Operating System to provide auto diagnosis and recovery capabilities and help improve the availability of the domain.

- *CPU off-lining* — Off-lines a CPU when an L2_SRAM module has an increased probability of experiencing fatal errors. The Solaris 10 OS Predictive Self-Healing provides greater diagnostic capabilities than Solaris 8 or 9 OS, and can isolate a greater number of faults.
- *Virtual memory page retirement* — Retires memory pages that have an increased probability of experiencing fatal errors.
- *Enhanced communication to the System Controller* — The Solaris OS with appropriate kernel updates and patches now sends additional ECC information to the System Controller for use in the firmware version 5.19.0 ECC Diagnosis Engine. Negotiation occurs between the Solaris OS and System Controller to determine fault diagnosis and message handling capabilities.

These enhancements, including the required OS releases and kernel updates to support these enhancements, are discussed in detail in “Solaris OS Enhancements” on page 15.

System Controller Firmware Enhancements

Beginning with firmware release 5.15.0, new features in the system controller firmware provide enhanced auto diagnosis and recovery capabilities. When the system encounters a fatal hardware error that causes a domain to be *error paused*, the hardware fault is automatically diagnosed. The Auto Diagnosis enhancement updates the component health status (CHS) on the affected FRU if the hardware failure can be isolated to a specific FRU or a set of specific FRUs. During the auto-restoration phase, POST consults the CHS and restores the domain with the fault isolated. In addition, if POST encounters a test failure, the CHS is stored in the appropriate FRU.

The System Controller firmware has been further enhanced to detect domain hangs and recover from such situations by resetting and rebooting the domain. Another System Controller firmware enhancement runs POST at increasing diagnostic levels when the domain panics repeatedly, so that the system can identify and isolate any persistent hardware faults.

These system controller firmware enhancements are discussed in more detail in the following sections.

Auto Diagnosis

The System Controller monitors the domains for hardware faults. Auto Diagnosis (AD) is automatically invoked on hardware faults that cause a *domain pause* or *data parity errors.* On Sun Fire midrange servers the data path is protected by parity and ECC. Domain operation is not impacted if data parity errors occur. Domain pauses are fatal errors and stop domain operation. Auto Diagnosis analyzes the following errors:

- Interconnect errors
- Data parity errors
- Internal ASIC errors

As of firmware version 5.19.0, a CPU diagnosis engine and ECC diagnosis engine have been added to Auto Diagnosis to automatically diagnose additional domain *error pause* events. Both the CPU diagnosis engine and the ECC diagnosis engine also have the ability to diagnose certain specific non-fatal events.

Figure 1 shows the Auto Diagnosis phase, Steps 1 through 5. Depending on the fault, three types of diagnosis results are possible:

- Fault diagnosed to a single component
- Fault diagnosed to a set of components
- Unresolved fault diagnosis

When a fault is diagnosed to a set of components, it does not mean that all components in the set are faulty. Rather, it means that the fault is located in a subset of these components (usually just one component).

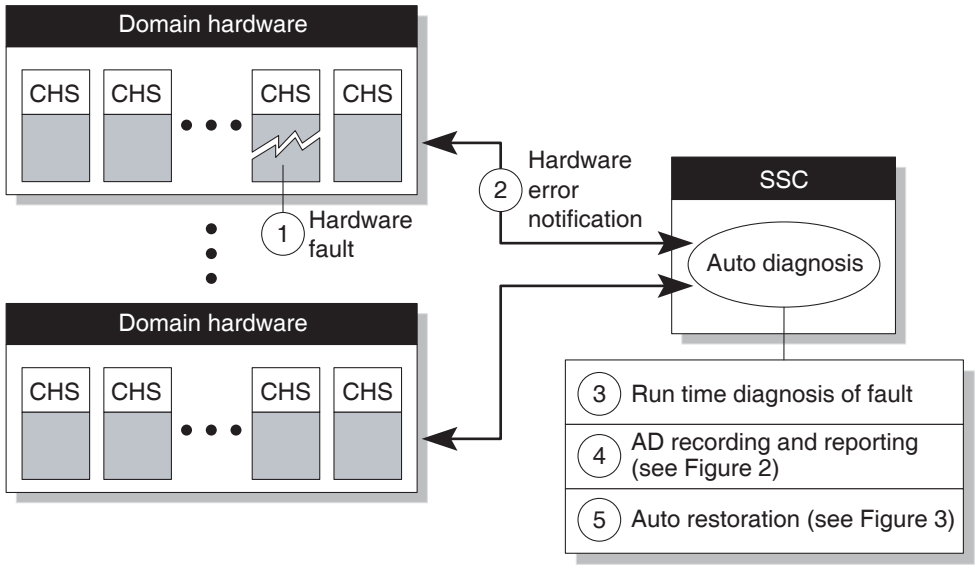

Figure 1. Auto diagnosis process.

CPU Diagnosis Engine

The CPU Diagnosis Engine is an enhancement included in firmware version 5.19.0. The CPU Diagnosis Engine diagnoses fatal and non-fatal errors that can be isolated to a single CPU. Systems running the Solaris 8 or 9 OS utilize the CPU Diagnosis Engine to identify specific fatal and non-fatal CPU events without the benefit of additional Solaris OS diagnosis. The Solaris 10 OS Predictive Self-Healing provides greater diagnostic coverage of CPU events than can be provided by the System Controller CPU Diagnosis Engine.

- *Fatal errors*

  The following events diagnosed by the CPU Diagnosis Engine are fatal errors and stop domain operation:

  – ISAP (system request parity error on incoming address)
  – IERR (internal CPU error)
  – PERR (system interface protocol errors)

- *Non-fatal errors*

  The following events are non-fatal CPU errors and do not stop domain operation:

  – L2-SRAM tag correctable errors (TSCE, THCE)
  – I-cache or D-cache parity errors (IPE, DPE)

  If the Solaris 10 OS is used within the domain, the CPU Diagnosis Engine is neither required nor used for diagnosing non-fatal CPU events. The Solaris 10 OS Predictive Self-Healing sends information to the System Controller for diagnosed non-fatal CPU events and the System Controller updates the appropriate CPU's component health status (CHS) so that future reboots and keyswitch operations do not configure non-optimal components into a domain. The Solaris 10 OS Predictive Self-Healing provides larger diagnosis coverage than is provided in the System Controller firmware.

  Domains which use Solaris 8 or 9 OS do utilize the CPU Diagnosis Engine to analyze certain non-fatal CPU errors. A SERD (Soft Error Rate Discrimination) algorithm detects when a specified number of distinct CPU events have occurred on the same processor in a specified period of time.

Table 3 lists the fatal and non-fatal CPU events diagnosed using firmware release 5.19.0 and indicates whether the CPU Diagnosis Engine or the Solaris 10 OS Predictive Self-Healing is used to identify the failure.

Table 3. Diagnosed CPU events in firmware release 5.19.0.

	CPU Event	**Domains using Firmware Release 5.19.0 and Solaris 8 or 9 OS**	**Domains using Firmware Release 5.19.0 and Solaris 10 OS**
Fatal CPU Events	ISAP	CPU Diagnosis Engine	CPU Diagnosis Engine
	IERR	CPU Diagnosis Engine	CPU Diagnosis Engine
	PERR	CPU Diagnosis Engine	CPU Diagnosis Engine
Non-Fatal CPU Events	TSCE	CPU Diagnosis Engine	Solaris 10 OS Predictive Self-Healing
	THCE	CPU Diagnosis Engine	Solaris 10 OS Predictive Self-Healing
	IPE	CPU Diagnosis Engine	Solaris 10 OS Predictive Self-Healing
	DPE	CPU Diagnosis Engine	Solaris 10 OS Predictive Self-Healing

Note – Table 3 lists only the CPU events that the CPU Diagnosis Engine can isolate. Additional CPU errors are diagnosed by the Solaris 10 OS Predictive Self-Healing. See "CPU Off-lining" on page 16 for specific details.

ECC Diagnosis Engine

The ECC (Error Checking and Correcting) Diagnosis Engine is an enhancement included in firmware version 5.19.0. The ECC Diagnosis Engine gathers and correlates data from both the System Controller and the Solaris OS to determine the proper faulty or suspect component. The ECC Diagnosis Engine can diagnose the following types of ECC errors:

- *Bad writers — A device that is creating data errors.*

  The device could be writing data errors into memory or propagating them onto the system bus. On a write transaction, data with bad ECC is written out to the system bus or into a memory location. Each time that memory location is read it creates an ECC event. This situation often presents itself as many ECC error messages reported from different CPUs with the same ECC Syndrome. This type of ECC error diagnoses to a single CPU.

- *Bad readers — A device which perceives a data error where one does not exist.*

  This type of ECC error Diagnoses to a single CPU.

- *Signaling Events (ECC Syndrome 0x071)*

  A Signalling Syndrome is a purposefully manufactured ECC syndrome that is created by a device to indicate to potential consumers of that data that the error has already been handled by a previous device. If collected data can identify the sourcing CPU and no history in data of prior CPU's memory having errors, the CHS status of the source CPU is updated.

- *Parity errors once a SERD (Soft Error Rate Discrimination) threshold is exceeded*

  Depending on error, may diagnose to a single CPU, two CPUs, or a single board. If a parity error occurs between a repeater board and a system board, or between a repeater board and an I/O assembly, the CHS status of two Boards is updated.

When an ECC error is diagnosed to a memory issue, the System Controller cannot always isolate the fault to a single DIMM. The ECC Diagnosis Engine can identify and isolate memory components in the following manner:

- One DIMM is identified if matching ECC data is received from the Solaris OS.
- Two DIMMs are identified if unable to match Solaris OS ECC data.
- One bank of memory (four DIMMs) is identified if the error is an uncorrectable event and matching Solaris OS ECC data is available.
- Two banks of memory (eight DIMMs) are identified if the error is an uncorrectable error and unable to match Solaris OS ECC data.

If an ECC error is a correctable event, the ECC Diagnosis Engine stores the identified memory component(s) from the event for reference in identifying bad writers, bad readers, and signaling errors. Correctable ECC events do not update the CHS of a DIMM at this time. History of the correctable error is stored to be utilized for auto diagnosis. Correctable ECC errors have been corrected either by the hardware or the Solaris OS. If an ECC error is an uncorrectable error and if the ECC Diagnosis Engine can isolate the error to a single DIMM by referencing the collected ECC history, the CHS of the identified DIMM

is updated and subsequent reboots and keyswitch on operations do not configure this component into the domain. DIMM CHS status can also be updated when DIMMs are identified during POST.

ECC events are primarily handled by the Solaris OS. See “Solaris OS ECC Information to System Controller” on page 20 for further details.

Auto Diagnosis Recording and Reporting

After the fault has been diagnosed, Auto Diagnosis records its diagnosis persistently in the CHS and reports it to the domain console and loghost as shown in Figure 2.

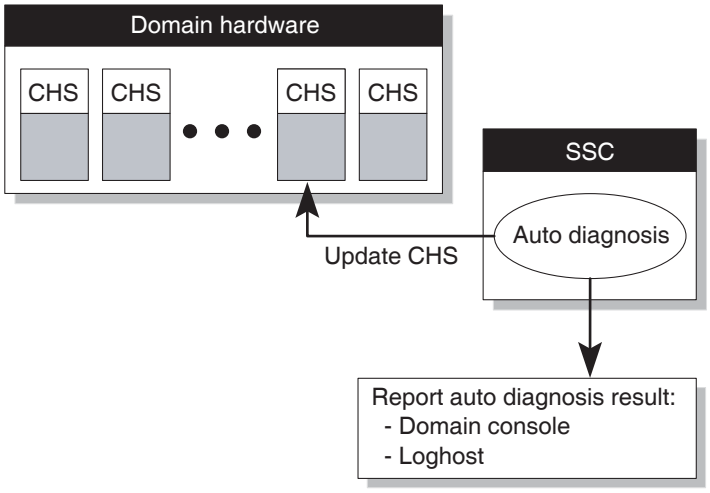

Figure 2. Auto diagnosis recording and reporting performed in the Sun Fire server System Controller (SSC).

The following example shows the Auto Diagnosis result that is output to the domain console for a single FRU diagnosis:

```
[AD] Event: SF3800.ASIC.SDC.PAR_SGL_ERR.60111010
     CSN: 124H58EE  DomainID: A ADInfo: 1.SCAPP.15.0
     Time: Thu Jan 23 20:47:11 PST 2003
     FRU-List-Count: 1;FRU-PN:5014362;FRU-SN: 011600; FRU-LOC:/N0/SB0
     Recommended-Action: Service action required
```

Auto Diagnosis reports a unique event code for the failure type and the diagnostic time. A full description of the Auto Diagnosis output format is in the S*un Fire 6800/4810/4800/ 3800 Systems Platform Administration* manual. In this example, Auto Diagnosis determined that the error is within CPU/Memory board at `FRU-LOC:/N0/SB0`.The reported information enables a service provider to make a quick determination of the problem and initiate corrective service action.

CHS is implemented for the following FRUs and components on Sun Fire midrange servers:

- CPU/Memory boards
- CPUs
- L2_SRAM modules
- DIMMs
- I/O assemblies
- Fireplane switches

Since the CHS and diagnostic information is persistently stored on a component or FRU, it moves with the component. Recurrence of a fault is prevented even if the component is moved to a different location. Service and repair of these systems is therefore easier, and preventing the recurrence of a fault improves the availability characteristics of Sun Fire midrange servers.

Auto Restoration

POST performs the domain auto restoration function. POST runs automatically after Auto Diagnosis, and also can be manually started by issuing the `setkey` command on the System Controller. POST consults the CHS of the domain hardware and reconfigures the domain to isolate the fault (Figure 3).

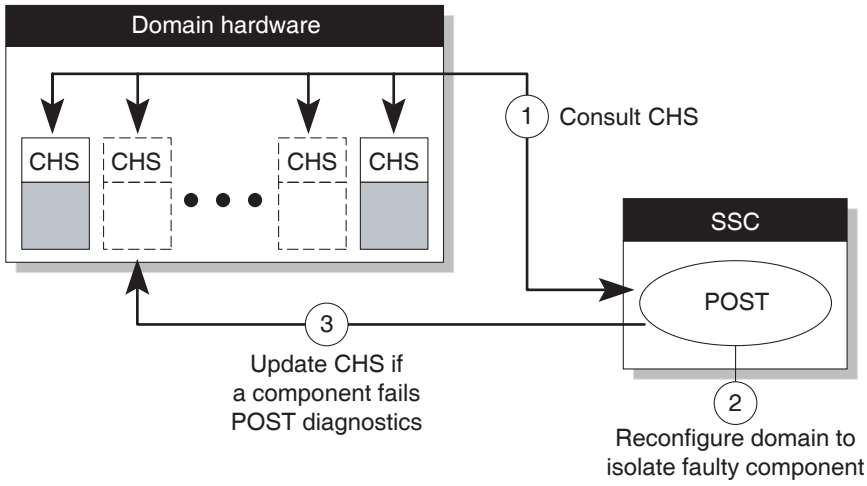

Figure 3. Auto restoration.

After the domain has been restored, administrators can run the `showcomponent` command to check which components have been disabled due to CHS.

If a redundant FRU or component is disabled because of its CHS, immediate replacement is not necessary because the domain is restored with the fault isolated. Utilizing dynamic reconfiguration, the FRU can be replaced at any time with minimal impact to the Solaris OS and user applications. For more information about dynamic reconfiguration, see [1] *Sun BluePrints OnLine* article “Sun Fire™ 3800-6800 Servers Dynamic Reconfiguration.”

Domain Hang Recovery

A domain's heartbeat is a communication mechanism that informs the System Controller that it is alive. If a domain is not updating its heartbeat or is unreachable from the console, this situation is categorized as a *domain hang*. A domain hang condition can occur due to hardware or software issues.

When using firmware version 5.15.0 or above on a Sun Fire midrange server, the System Controller acts as an external monitor for each domain (Figure 4). The System Controller monitors for a domain hang condition and initiates an externally initiated reset (XIR) if the domain heartbeat register is not updated within the maximum time out limit.

The domain heartbeat monitoring is configurable for each domain using the parameter `watchdog_timeout_seconds` in the `/etc/systems` file of each domain. The default time out value for a domain is three minutes. If `watchdog_timeout_seconds` is set to a value less than three minutes, the System Controller defaults to three minutes. For additional details, refer to the `system(4)` man page.

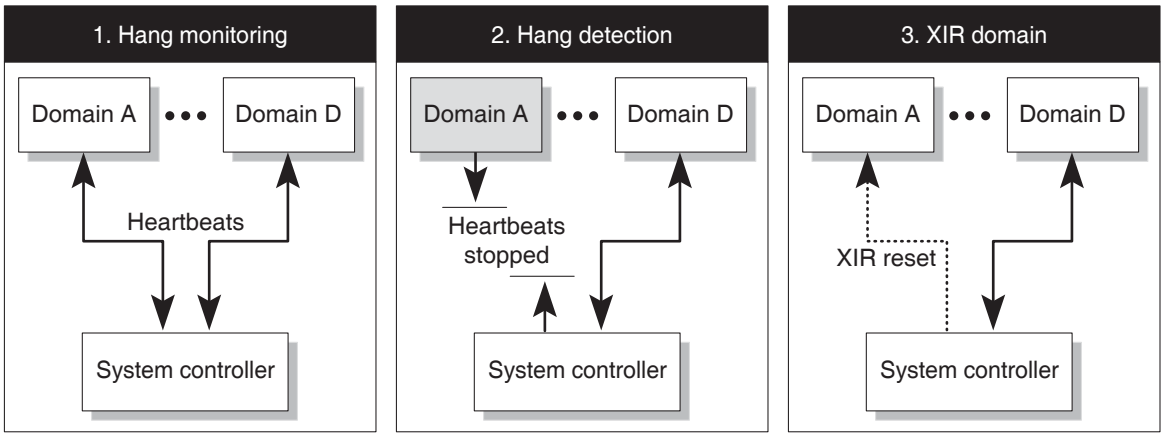

Figure 4. Domain hang restoration.

The following example shows the console output of a domain that was declared hung because the domain heartbeat register was not updated within the maximum time out limit:

Table 4. Example output, domain hung and reset by system controller.

```
Jan 22 17:02:06 sc0 Domain-A.SC: Domain watchdog timer expired.
Jan 22 17:02:06 sc0 Domain-A.SC: Using default hang-policy (RESET).
```

In addition to the heartbeat monitoring, the System Controller also checks if the domain is picking up the interrupts sent to it by the System Controller. The System Controller sends interrupts to the domain when, for example, characters are entered on the domain console. If on a second interrupt the previous one has not been picked up by the domain, the System Controller waits for one minute before declaring the domain hung.

The following example shows the console output of a domain that was declared hung because it did not pick up its interrupts:

Table 5. Example output, domain not picking up interrupts is declared hung.

```
Jan 22 18:09:02 sc0 Domain-A.SC: Domain is not responding to interrupts.
Jan 22 18:09:02 sc0 Domain-A.SC: hang-policy is NOTIFY.  Not resetting domain.
```

The *hang policy* is set by the `setupdomain` command to `notify` or `reset`. If the hang policy is set to `notify`, the System Controller reports the hang condition on the domain console and does not reset the domains (see example output in Table 5). If the hang policy is set to `reset`, the System Controller reports the hang condition on the domain console and initiates a domain reset (see example output in Table 4).

The hang policy is set to `reset` by default. Administrators can use the `error-reset-recovery` variable of the `setupdomain` command to specify whether or not the system creates a core dump when a domain is reset. It is recommended to set the `hang-policy` variable to `reset` and the `error-reset-recovery` variable to `sync` to provide the maximum troubleshooting information for a domain hang. To identify the cause of the domain hang, consult a service provider while referring to the core file.

A system's *hang-policy* can be verified using the `showdomain` command in the domain shell. Each domain's `hang-policy` parameter should be verified to ensure that it is properly set to `reset` for maximum domain availability. For more information about domain setup, refer to the *Sun Fire™ 6800/4810/4800/3800 Systems Platform Administration* manual.

Recovery From Repeated Domain Panics

Domain panics can be caused by software or hardware. To prevent hardware faults from causing panic-reboot loops, the System Controller firmware runs POST at increasing diagnostic levels on recurring panics.

On the first panic, the domain reboots and writes a core file, which can be used to analyze the problem. However, if further panics occur within a short time period, it is desirable to run POST automatically at a higher level as part of domain restoration. POST diagnostics verify the status of the hardware and identify and isolate faulty components (if any). After identifying faulty components, POST updates the appropriate CHS. With firmware release 5.14.0 and higher, the System Controller keeps track of the number of domain panics over time. A panic reboot of a domain has a unique register signature that differs from the normal reboot of a domain. The POST diagnosis level is increased if a panic recurs within four hours of the previous panic. If a panic occurs after four hours from the latest panic, the panic-reboot counter is reset. The panic-reboot counter is also reset if the domain is manually rebooted or a keyswitch operation is performed.

On recurring panics, the domain POST diagnostic level is increased to the next higher level from diag-level `init`. In increasing order, POST levels are `init`, `quick`, `min`, `max/default`, `mem1`, and `mem2`. The

domain is put into standby position if it continues to panic undetected by the user after the highest level of POST is run (Figure 5). For further analysis, consult a service provider while referring to the core file.

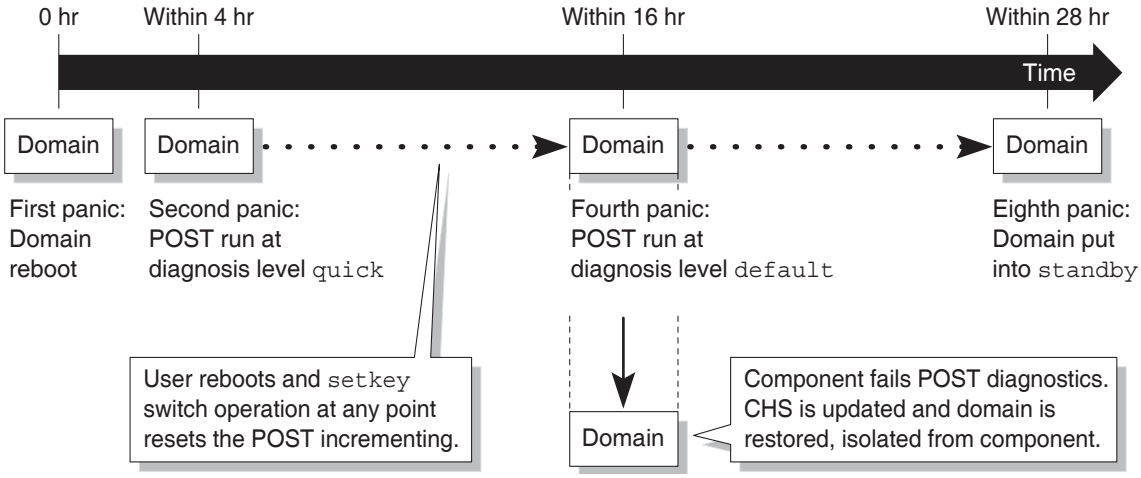

Figure 5. Domain panic restoration.

This domain panic restoration feature prevents a panic-reboot loop of domains. If recurring panics are caused by a software bug, the increased POST level minimizes hardware as a possible cause. Downtime for running further POST diagnostics is not required because the system automatically has already taken the necessary measures.

The standard sequence for recovery from repeated domain panics is shown in Table 6. As of firmware version 5.19.0, a new domain configuration variable of `max-panic-diag-limit` is available. This variable enables each domain to specify the maximum POST diagnosis level to perform for recovery from repeated domain panics. If another panic occurs within four hours after the `max-panic-diag-limit` POST diagnosis level is reached, the domain is put into `standby` mode.

Table 6. Recovery from repeated domain panics.

Sequence for Recovery from Repeated Domain Panics		**Sequence as of Firmware Version 5.19.0 using max-panic-diag-limit set to max**	
Panic #	Diag-Level	Panic #	Diag-Level
1	init	1	init
2	quick	2	quick
3	quick	3	quick
4	min	4	min
5	max/default	5	max/default
6	mem1	6	domain put to standby
7	mem2		
8	domain put to standby		

Persistent Record of Hardware Failures from the Solaris OS

As of Sun Fire server firmware release 5.15.3, certain CPU failures diagnosed by the Solaris OS are persistent. The System Controller receives and stores hardware failure messages from the Solaris OS if the system is using the appropriate Solaris OS release and kernel update. (See Table 1 on page 2 for specific OS and kernel requirements.) Domain reboots and `setkeyswitch on` events no longer configure components that the Solaris OS has previously diagnosed as failed.

Note – The Solaris 10 OS Predictive Self-Healing provides larger diagnosis coverage than either Solaris 8 OS or Solaris 9 OS; see “Solaris OS Enhancements” on page 15 for further details.

When the System Controller receives a hardware event message from the Solaris OS, the System Controller updates the CHS of the affected FRU (Figure 6). POST consults the CHS and does not configure any listed faulty components. System availability is improved as future domain configurations deconfigure the failed components.

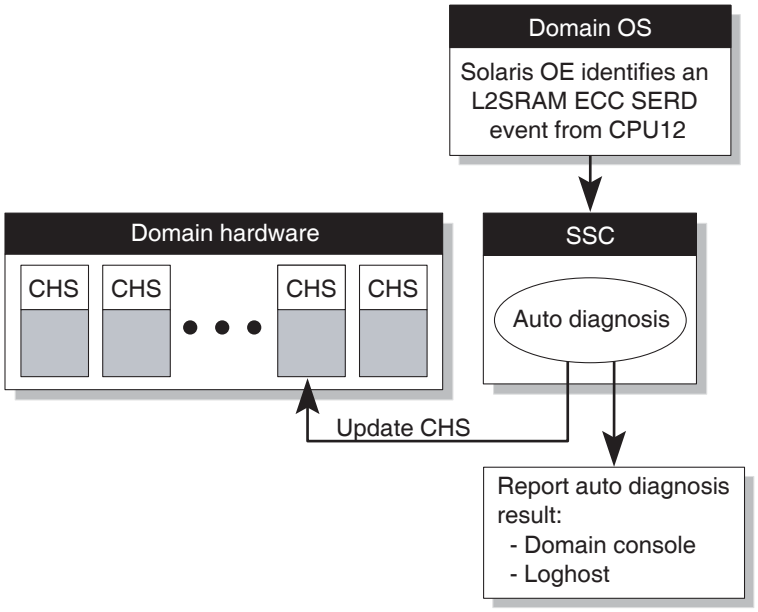

Figure 6. Persistent failure record.

The soft error rate discrimination (SERD) algorithm in the Solaris OS detects when a specified number of distinct CPU events have occurred on the same processor in a 24-hour period. After the specified CPU SERD events, the CPU becomes a candidate for off-lining within the Solaris OS.

The following example Auto Diagnosis failure message specifies the phrase SF-SOLARIS-DE in the ADInfo Domain Log message to identify which hardware events were received from the Solaris OS:

```
[DOM] Event: SF6800.L2SRAM.SERD.f.1b.10040000000091.f4470000
      CSN: 044M347B DomainID: A ADInfo: 1.SF-SOLARIS DE.5_9_GENERIC_112233-09
      Time: Mon Jun 02 23:34:59 PDT 2003
      FRU-List-Count: 1; FRU-PN: 3704125; FRU-SN: 090K01; FRU-LOC: /N0/SB3/P3/E0
      Recommended-Action: Service action required.
```

Components identified for off-lining in the Solaris OS are now communicated to the System Controller so that upon future reboots and keyswitch events these components do not need to be diagnosed again by the Solaris OS. It is important to note that the System Controller updates the affected component's CHS after receiving the ADInfo Domain Log message. This message is sent to the System Controller from the Solaris OS when a component is identified by the Solaris OS as a candidate for off-lining, not when the component is actually off-lined in the Solaris OS.

Persistent Logging in Enhanced-Memory System Controllers (SC V2)

In Sun Fire midrange systems configured with enhanced-memory system controllers (SC V2s) and running firmware version 5.17.0 or higher, system error messages and six levels of messages (alert, critical, emergency, error, fault, warning) in the message logs are retained in persistent storage. This capability enhances service, diagnosis and repair as vital error messages and message logs are retained if the System Controller is rebooted or powered off.

The following commands can be used to view the persistent system error messages and logs:

- The `showerrorbuffer` command — displays the persistent system error messages.
- The `showlogs` command — displays the persistent logs.

The information displayed by these commands can be used by a service provider for troubleshooting purposes. For details on message logs and system error messages, refer to the *Sun Fire Midrange Systems Platform Administration Manual* and the `showlogs` and `showerrorbuffer` command descriptions in the *Sun Fire Midrange System Controller Command Reference Manual*.

The SCV2 is standard for Sun Fire E2900/E4900/E6900 servers. Sun Fire V1280/3800/4800/6800 Servers can be upgraded to use the new SCV2. SCV2 is strongly recommended for sites which do not use a remote loghost to capture and retain platform and domain logs. A remote loghost is still recommended as a best practice in addition to utilizing SCV2. A remote loghost contains a complete history of platform and domain logs which provides the maximum available data for accurate diagnosis of any events. Contact a sales or account representative for details on upgrading to SCV2.

CPU Core Voltage Monitoring (VCMON)

As of firmware version 5.18.0, CPU core voltages are proactively monitored. Voltage on the supply lines to each processor, known as the core voltage (Vcore), is monitored by existing circuits on the System Board. Failures within the CPU socket appear as a connectivity problem between the processor and

Uniboard. This type of failure has been determined to exhibit a very gradual increasing electrical resistance within the socket's conductor columns, between the processor package itself and the printed circuit board (PCB) of the System Board. Only UltraSPARC-III class processors are monitored. Due to a physical design change, the UltraSPARC-IV processors are not vulnerable to this issue. For more information about Vcore Monitoring, see [2] *Sun BluePrints OnLine* article, "Predictive Fault Monitoring in Sun Fire Servers".

The following example shows a VCMON message which is printed to both the platform and domain logs:

```
[VCM] Event: SF6800.VCMON.0.02.1663
     CSN: 036H3005 DomainID: B ADInfo: 1.VMON.18.0
     Time: Thu Sep 16 12:38:59 PDT 2004
     FRU-List-Count: 1; FRU-PN: 5014362; FRU-SN: 004413; FRU-LOC: /N0/SB3/P0
     Recommended-Action: Service action required
```

As of firmware version 5.19.0, processors identified by VCMON have their CHS updated so that future reboots or `keyswitch on` events do not configure the identified processor. Contact an authorized Sun service provider for information on proper actions for boards identified by VCMON.

Solaris OS Enhancements

Enhancements to the Solaris OS provide additional auto diagnosis and recovery capabilities. Kernel updates for Solaris 8 and 9 OS on UltraSPARC™ III and UltraSPARC™ IV systems enhance the error handling capabilities within L2_SRAM modules and pages of virtual memory. Multiple correctable errors on accessing an L2_SRAM module or DIMM indicate a higher probability of experiencing an uncorrectable error. To prevent a fatal uncorrectable error, the Solaris OS attempts to off-line CPUs and affected memory pages are retired. The availability of domains increases because the Solaris OS does not access L2_SRAM modules or retired memory pages that have an increased failure probability.

Additional enhancements enable the Solaris OS to communicate information to the System Controller. For example, updated versions of the Solaris OS communicate information on hardware failures for L2_SRAM modules which the OS attempted to off-line and provide ECC information which is used by the ECC Diagnosis Engine in the System Controller firmware. ECC information from the Solaris OS is only provided if the System Controller is using firmware version 5.19.0 and appropriate Solaris OS kernel updates and patches.

The Solaris 10 OS Predictive Self-Healing provides greater OS diagnostic capabilities than either Solaris 8 or 9 OS. Using the Solaris 10 OS and a minimum system controller firmware version of 5.19.0 enables additional CPU events to be persistent across domain reboots and keyswitch events.

Table 7 lists the auto diagnosis and recovery enhancements that are provided by the Solaris OS and details the required kernel updates and patches. These enhancements are discussed in more detail in the following sections.

Table 7. Minimum requirements for auto diagnosis and recovery enhancements within the Solaris OS.

OS Enhancement	Solaris 8 OS	Solaris 9 OS	Solaris 10 OS
CPU off-Lining for L2_SRAM conditions within Solaris OS	Solaris 8 OS kernel update 108528-20	Solaris 9 OS kernel update 112233-06	Solaris 10 OS (no updates required)
Virtual Memory Page Retirement for DIMM conditions within Solaris OS	Solaris 8 OS kernel update 117000-03	Solaris 9 OS kernel update 112233-12	Solaris 10 OS (no updates required)
Predictive Self-Healing within the OS	*not available*	*not available*	Solaris 10 OS (no updates required)
OS communicates hardware failures to the System Controller for L2_SRAM modules which the OS has attempted to off-line	Solaris 8 OS kernel update 108528-24 with patch 110373-05 and firmware releases 5.15.3 - 5.19.0	Solaris 9 OS kernel update 112233-09 with patch 116009-01 and firmware releases 5.15.3 - 5.19.0	Solaris 10 OS (no updates required) and firmware releases 5.15.3 - 5.19.0
OS communicates additional CPU events from Solaris 10 OS Predictive Self-Healing to the System Controller	*not available*	*not available*	Solaris 10 OS kernel update 118822-03 with patch 119568-01; Sun Fire V1280 and E2900 servers and Netra 1280 servers also require patch 119570-01 and firmware release 5.19.0
OS communicates ECC information to the System Controller for use in the ECC Diagnosis Engine	Solaris 8 kernel update 117350-24 with patches 116979-05 and 110369-08 and firmware release 5.19.0[a] Solaris 8 kernel update 117350-24 with patches 116979-05 and 112162-04 and firmware release 5.19.0[b]	Solaris 9 kernel update 118558-06 with patches 117124-08 and 116009-03 and firmware release 5.19.0	Solaris 10 OS kernel update 118822-03 with patch 119568-01; Sun Fire V1280 and E2900 servers and Netra 1280 servers also require patch 119570-01 and firmware release 5.19.0

a. Patch 110369-08 applies for Sun Fire 3800-6800 servers and Sun Fire 4900-6900 servers only.
b. Patch 112162-04 applies for Sun Fire V1280 and E2900 servers and Netra 1280 servers only.

CPU Off-lining

The Solaris OS attempts to off-line a CPU which is identified as having specific types of ECC events. The Solaris 10 OS Predictive Self-Healing provides a larger diagnostic capability than either Solaris 8 or 9 OS.

The Solaris OS can currently identify and isolate the following types of faults:

- Solaris 8 OS (kernel update 108528-20 or higher) or Solaris 9 OS (kernel update 112233-06 or higher)
  - L2_SRAM ECC correctable events which reach SERD threshold (UCC, WDC, CPC, EDC)
  - L2_SRAM ECC uncorrectable events (UCU, WDU, CPU, EDU)
- Solaris 10 OS
  - L2_SRAM ECC correctable events which reach SERD threshold (UCC, WDC, CPC, EDC)
  - L2_SRAM ECC uncorrectable events (UCU, WDU, CPU, EDU)
  - L2-SRAM tag correctable errors which reach SERD threshold (TSCE, THCE)

– I-cache or D-cache parity errors (IPE, DPE)
– FPU failures
– UltraSPARC-III+ and UltraSPARC-IV ECC correctable events which reach SERD threshold (TO, BERR, DTO)
– UltraSPARC-III+ and UltraSPARC-IV System Bus ECC events (IVU, IVC, EMC)

These diagnosed ECC events are listed in Table 8 for quick reference:

Table 8. ECC events diagnosing to a CPU.

ECC Event	**Solaris 8 OS (KU-108528-20 or higher) or Solaris 9 OS (KU-112233-09 or higher)**	**Solaris 10 OS**
L2_SRAM ECC correctable events which reach SERD threshold (UCC, WDC, CPC, EDC)	√	√
L2_SRAM ECC uncorrectable events (UCU, WDU, CPU, EDU)	√	√
L2-SRAM tag correctable errors which reach SERD threshold (TSCE, THCE)		√
I-cache or D-cache parity errors (IPE, DPE)		√
FPU failures		√
ECC correctable events which reach SERD threshold (TO, BERR, DTO)		UltraSPARC-III+ and UltraSPARC-IV
System Bus ECC events (IVU, IVC, EMC)		UltraSPARC-III+ and UltraSPARC-IV

As of Solaris 8 OS (kernel update 108528-20) and Solaris 9 OS (kernel update 112233-06), the Solaris OS keeps track of the number of ECCs over time on an L2_SRAM module (Figure 7). Two types of ECCs are considered here—non-fatal multibit errors (UCU, CPU, WDU, EDU) and non-fatal single-bit correctable errors (UCC, CPC, WDC, EDC). If an L2_SRAM module experiences one non-fatal multibit error or three single-bit correctable errors in a 24-hour window, the L2_SRAM module is diagnosed with an increased probability of suffering a future fatal failure. In this scenario, the Solaris OS has been enhanced to automatically attempt to off-line the affected CPU module. It is possible that the CPU off-line may not succeed because there might be processes bound to that CPU.

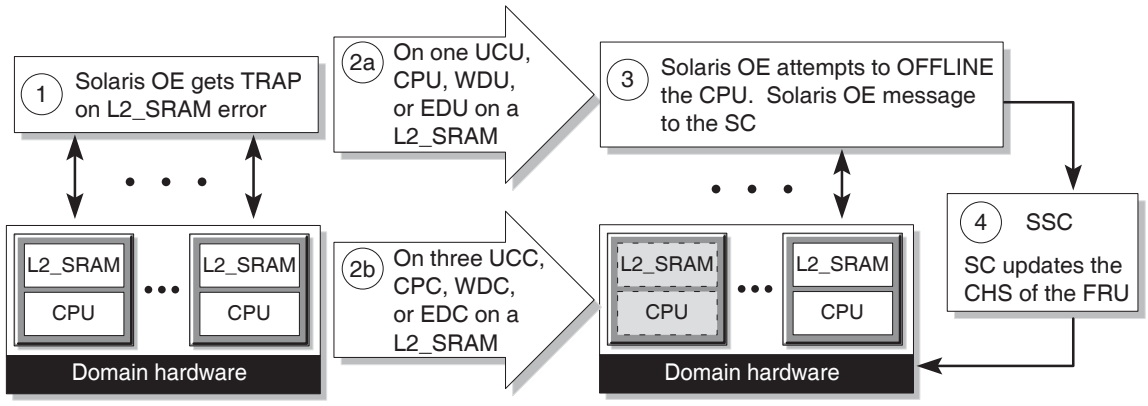

Figure 7. Solaris OS L2_SRAM error handling.

For systems using an updated version of the Solaris OS (Solaris 8 OS kernel update 108528-24 with patch 110373-05 or higher; Solaris 9 OS kernel update 112233-09 with patch 116009-01 or higher; or Solaris 10 OS) and System Controller firmware release 5.15.3 or higher, the Solaris OS sends a message to the System Controller after a CPU is identified as a candidate to be off-lined.

The following example shows the messages on successfully off-lining a CPU that experienced more than two correctable error events in a 24-hour window:

```
Feb  3 06:38:40 doma SUNW,UltraSPARC-III: NOTICE: [AFT1] CPU6 offlined due to more than
2 xxC Events in 24:00:00 (hh:mm:ss)
```

Once a CPU is identified as a candidate to be off-lined, the Solaris OS sends a message to the system controller. The system controller updates the CHS of the affected FRU so that the faulty CPU is not configured into a domain on future reboots or `setkeyswitch on` events. For systems using firmware version 5.15.3 or higher, but lower than 5.19.0, only the L2_SRAM ECC correctable and uncorrectable events are communicated to the System Controller. For systems using firmware version 5.19.0 and an updated version of the Solaris 10 OS (Solaris 10 OS kernel update 118822-03 or higher with patches listed in Table 7 on page 16), the additional Solaris 10 OS CPU off-line events are communicated to the System Controller and the CHS of the appropriate FRU is updated.

Note – The Solaris 10 OS utilizes Predictive Self-Healing and can thereby diagnose more CPU errors than either the Solaris 8 or 9 OS. See the Sun White Paper “Predictive Self-Healing in the Solaris 10 Operating System” for an overview of this new functionality in the Solaris 10 OS.

The following example shows the Solaris 10 OS message upon diagnosing a faulty CPU:

```
SUNW-MSG-ID: SUN4U-6900-6H, TYPE: Fault, VER: 1, SEVERITY: Major
EVENT-TIME: Sun Oct 17 14:15:50 PDT 2004
PLATFORM: SUNW,Sun-Fire-6900, CSN: -, HOSTNAME: 6900a
EVENT-ID: 64fe6c23-12b7-ccd1-f0a7-b531941738f8
DESC: The number of errors associated with this CPU has exceeded acceptable levels.
Refer to http://sun.com/msg/SUN4U-8000-6H for more information.
AUTO-RESPONSE: An attempt will be made to remove the affected CPU from service.
IMPACT: Performance of this system may be affected.
REC-ACTION: Schedule a repair procedure to replace the affected CPU. Use fmdump
-v -u <EVENT_ID> to identify the CPU.
```

Off-lining a CPU with a higher probability of experiencing a fatal error potentially increases the availability of the system. Communication between the Solaris OS and the system controller to persistently store the CHS increases availability and provides easier diagnosis and serviceability of the system. Dynamically reconfigured CPU/Memory boards can be replaced with minimal impact to the Solaris OS and user applications.

Virtual Memory Page Retirement

The Solaris OS (as of Solaris 8 OS kernel update 117000-03, Solaris 9 OS kernel update 112233-12, or Solaris 10 OS) keeps track of the number of correctable errors over time on a DIMM. If more than three errors occur on the same DIMM within a 24-hour period, the domain automatically schedules retirement of the memory page (Figure 8). Memory pages can be retired when all processes have released the page.

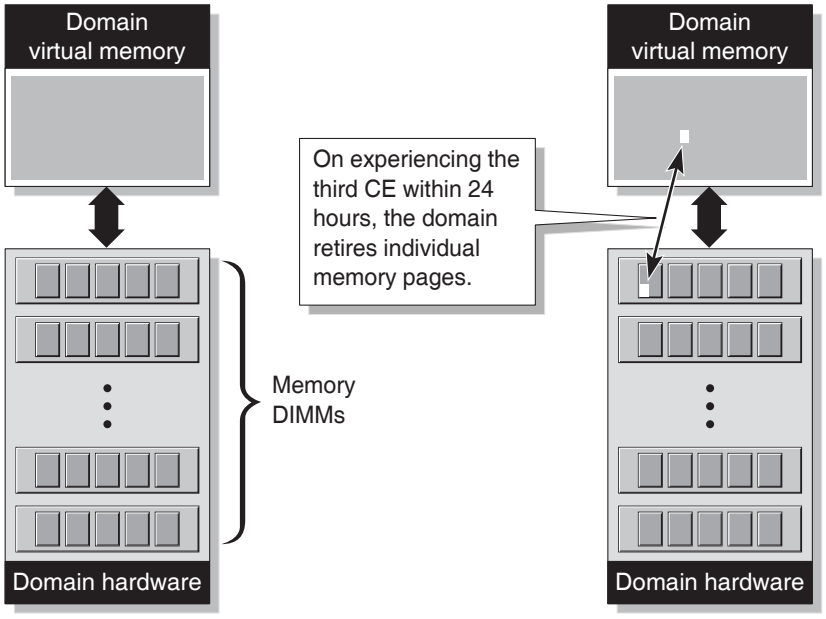

Figure 8. Solaris OS memory ECC handling.

Retired pages are not used by the domain. However, on a reboot, retired pages are accessible again by the domain. For updated releases of the Solaris OS (as of Solaris 8 OS kernel update 117350-24, Solaris 9

OS kernel update 118558-06, or the Solaris 10 OS) that utilize System Controller firmware release 5.19.0, retired memory pages are not persistent across reboots or domain keyswitch events.

The following example shows the messages output when a memory page is retired:

```
Jan  7 04:14:07 doma unix: [ID 596940 kern.warning] WARNING: [AFT0] 3 soft errors
in less than 24:00 (hh:mm) detected from Memory Module Board 4 J3801

Jan  7 04:14:07 doma unix: [ID 618185 kern.notice] NOTICE: Scheduling removal of
page 0x00000001.2bf6c000

Jan  7 04:14:12 doma unix: [ID 693633 kern.notice] NOTICE: Page
0x00000001.2bf6c000 removed from service
```

Solaris OS Enhanced Communication to the System Controller

The Solaris OS (as of the Solaris 8 OS kernel update 117350-24, Solaris 9 OS kernel update 118558-06, or Solaris 10 OS kernel update 118822-03, all with appropriate patches) has the ability to negotiate with the System Controller to identify fault diagnosis and message handling capabilities. The properly updated releases of the Solaris OS also provide a new ECC Information message to the System Controller if the System Controller is utilizing firmware version 5.19.0. The Solaris 10 OS also sends additional CPU off-line messages to the System Controller when the System Controller is using firmware version 5.19.0.

Solaris OS L2_SRAM Failure Information to System Controller

The enhanced Solaris OS kernels have the ability to communicate hardware failures to System Controllers utilizing firmware version 5.15.3 or higher. As of the Solaris 8 OS (kernel update 108528-24 with patch 110373-05), the Solaris 9 OS (kernel update 112233-09 with patch 116009-01) or Solaris 10, a message is sent to the System Controller when the Solaris OS identifies a faulty L2_SRAM module. The failed L2_SRAM module is not reconfigured into a domain on future domain reboots or `setkeyswitch on` operations because the System Controller has recorded the component as faulty in its CHS.

Solaris OS ECC Information to System Controller

The Solaris OS can identify a specific DIMM in ECC correctable events whereas the System Controller firmware can only identify to a granularity of two DIMMs for ECC events. Prior to firmware version 5.19.0, the System Controller had no method to utilize the Solaris OS ECC diagnostic capabilities. A new message format implemented in the Solaris OS (releases listed below) and firmware version 5.19.0 enables the System Controller to utilize the Solaris OS ECC information. Now the ECC Diagnosis Engine (System Controller firmware version 5.19.0) utilizes both the Solaris OS ECC information and System Controller information to diagnose situations which used to require analysis by a service provider.

This new message format for communicating the Solaris OS ECC information is available in the following Solaris OS releases:

- Sun Fire 3800-6800 servers
  - Solaris 8 OS kernel update 117350-24 with patches 116979-05 and 110369-08
  - Solaris 9 OS kernel update 118558-06 with patches 117124-08 and 116009-03
  - Solaris 10 OS kernel update 118822-03 and patch 119568-01
- Sun Fire 4900-6900 servers
  - Solaris 8 OS kernel update 117350-24 with patches 116979-05 and 110369-08
  - Solaris 9 OS kernel update 118558-06 with patches 117124-08 and 116009-03
  - Solaris 10 OS kernel update 118822-03 and patch 119568-01
- Sun Fire V1280 and E2900 servers, and Netra 1280 servers
  - Solaris 8 OS kernel update 117350-24 with patches 116979-05 and 112162-04
  - Solaris 9 OS kernel update 118558-06 with patches 117124-08 and 116009-03
  - Solaris 10 OS kernel update 118822-03 and patch 119568-01 and 119570-01

Negotiation between the Solaris OS and System Controller

Due to the differing capabilities in System Controller firmware versions and Solaris OS releases, negotiation must occur between them to determine the fault diagnosis and message handling capabilities. System Controllers using firmware version higher than 5.15.3, but lower than 5.19.0, can only accept messages from the Solaris OS for CPU off-line attempts when an L2_SRAM correctable or uncorrectable event has occurred. System Controllers using firmware version 5.19.0 can accept the additional CPU off-line messages from the Solaris 10 OS and also can accept additional ECC Information from the enhanced Solaris 8, 9, and 10 OS releases with appropriate kernel updates and patches. Both the System Controller firmware version and each domain's operating system must be considered in determining the fault diagnosis and message handling capabilities of the platform.

For platforms that run multiple domains with differing Solaris OS fault diagnosis and messaging handling capabilities, each domain negotiates its capability independently. As an example, assume a Sun Fire server whose System Controller is using firmware version 5.19.0, which supports receiving the enhanced communication from the Solaris OS. One domain may be running the Solaris 10 OS that has the ability to send additional CPU events to the System Controller and the ability to send additional ECC information, while other domains may be using the Solaris 8 or 9 OS without the capability of sending the additional ECC information. The System Controller appropriately manages each domain based on the capability of each individual domain’s Solaris OS fault diagnosis and message handling capability.

An example of multiple domains communicating with a System Controller running firmware release 5.19.0 is shown in Figure 9. In this example, both Domain A (running the Solaris 10 OS) and Domain C (running the Solaris 8 OS and the required patches) send additional ECC information to the System Controller to be used in the ECC Diagnostic Engine. Domain B is running the Solaris 8 OS without the patches required to support sending additional ECC Information to the System Controller. The Solaris 10 OS, running in

Domain A, also sends additional CPU off-line messages to the System Controller in this example, as the System Controller is using firmware version 5.19.0.

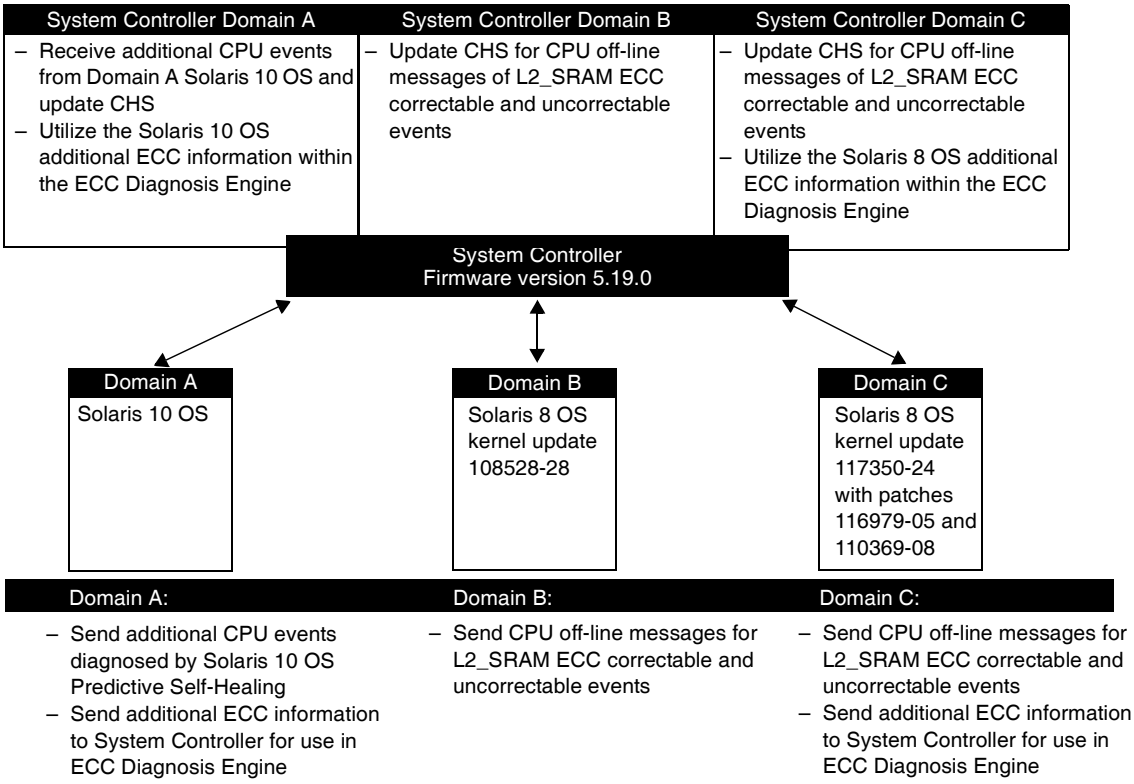

Figure 9. Multiple domains communicating with System Controller running firmware release 5.19.0.

Another example of multiple domains communicating with a System Controller, this one running firmware release 5.15.3, is shown in Figure 10. If the System Controller is using firmware version 5.15.3 or higher (but lower than firmware version 5.19.0), the Solaris OS only sends messages to the System Controller for attempts to off-line CPUs for L2_SRAM correctable and uncorrectable ECC events. The additional ECC information from the Solaris OS is not sent to System Controllers using a firmware version lower than 5.19.0. Furthermore, the Solaris 10 OS does not send the additional CPU off-line messages to the System Controller when the System Controller is using firmware versions lower than 5.19.0.

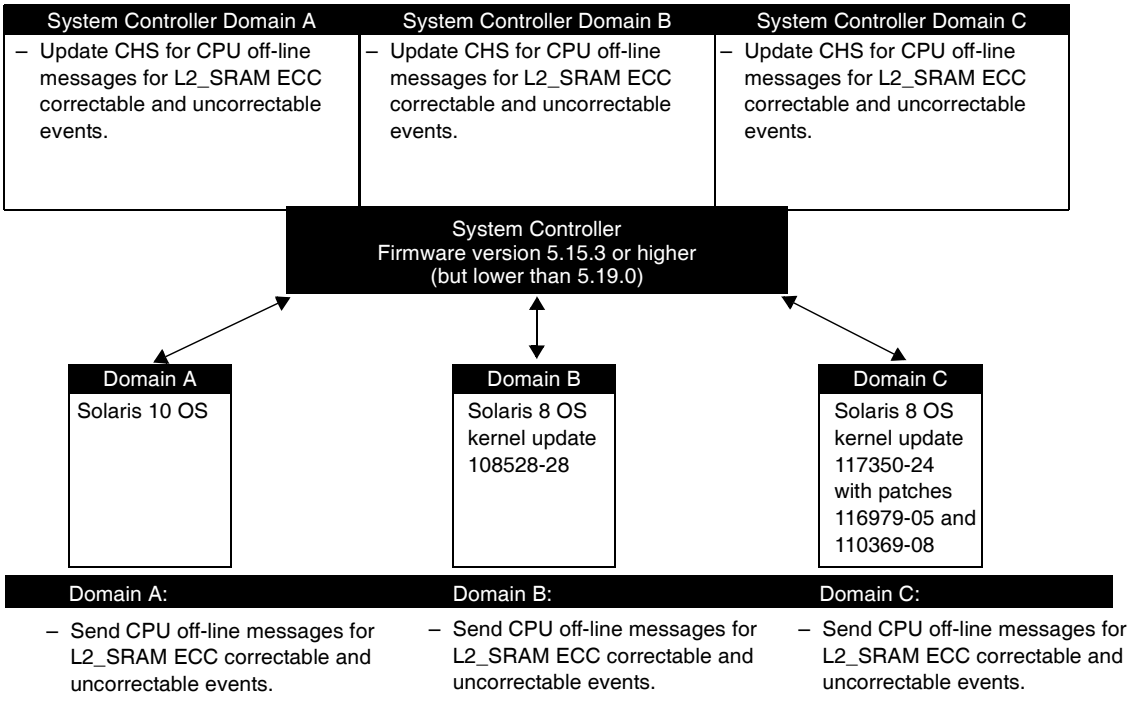

Figure 10. Multiple domains communicating with System Controller running firmware release 5.15.3.

Conclusion

In a mission critical environment, high availability is achieved by system resiliency, the appropriate configuration, serviceability, and efficient and automated restoration processes. The enhancements made in firmware version 5.19.0 address all of these elements and increase availability, serviceability, and diagnosability of Sun Fire midrange servers. To fully benefit from the enhancements discussed within this article, upgrading to firmware version 5.19.0 is required. In addition, upgrading to the Solaris 8, Solaris 9 or Solaris 10 OS with the required minimum kernel updates and patches (see Table 1 on page 2) is also necessary. Use of the Solaris 10 OS, with its Predictive Self-Healing, provides the highest levels of auto diagnosis and recovery.

Standardized messaging, persistent component health status, and auto diagnosis are powerful tools for users and service providers. Auto diagnosis in combination with dynamic reconfiguration on Sun Fire midrange servers greatly increases system availability and decreases scheduled downtime for maintenance. In addition, the proactive actions of the Solaris OS on CPU and DIMM events makes the Solaris OS more resilient and further increases its availability.

About the Author

Tricia Wittsack is a member of Sun's Product Technical Support (PTS) engineering group in Newark, California that is responsible for providing engineering support for Sun Fire midrange servers and Sun Enterprise™ servers. Prior to her current role, Tricia worked in the PTS Global Labs in Newark, California.

Tricia has worked in the computer industry for 15 years. Her technical career began as an electronics technician in the U.S. Navy. After serving in the U.S. Navy, she worked in various technical support positions for a private software applications company and then joined Sun Microsystems in 2000. She received a Bachelor of Science degree in Information Technology from the University of Phoenix.

Acknowledgments

The author would like to recognize the following individuals for their contributions to this article:

- Peter Gonscherowski, Manager of Technical Staff, Product Technical Support
- James Hsieh, Member of Technical Staff, Product Technical Support
- Ken Kambic, Member of Technical Staff, Product Technical Support
- David Re, Member of Technical Staff, Scalable Systems Group

A special thanks to Sundeep Panicker, the Availability ITeam Lead, and the Availability ITeam Developers from the Scalable Systems Group Common Software for their detailed information concerning the new availability features included in firmware version 5.19.0.

References

[1] Gonscherowski, Peter. "Sun Fire™ 3800-6800 Servers Dynamic Reconfiguration," *Sun BluePrints Online*, April 2002.
To access this article, go to: `http://www.sun.com/blueprints/0402/sunfire.pdf`

[2] Re, Dave and Loganathan, Kumar. "Predictive Fault Monitoring in Sun Fire™ Servers," *Sun BluePrints Online*, April 2005.
To access this article, go to: `http://www.sun.com/blueprints/0405/819-2261.pdf`

Related Resources

Sun Fire™ Midrange Systems Firmware 5.16.0 Release Notes, Sun Microsystems, 817-2973.

Sun Fire™ Midrange Systems Platform Administration Manual (5.16.0), Sun Microsystems, 817-2971.

Sun Fire™ Midrange System Controller Command Reference Manual (5.16.0), Sun Microsystems, 817-2972.

Sun Fire™ 6800/4810/4800/3800 Systems Firmware 5.14.X Release Notes, Sun Microsystems, 817-1305.

Sun Fire™ 6800/4810/4800/3800 Systems Firmware 5.15.X Release Notes, Sun Microsystems, 817-1001.

Sun Fire™ 6800/4810/4800/3800 Systems Platform Administration Manual, Sun Microsystems, 817-0999.

Sun Fire™ 6800/4810/4800/3800 System Controller Command Reference Manual, Sun Microsystems, 817-1000

“Predictive Self-Healing in the Solaris 10 Operating System”, Sun Microsystems.To access this article, go to: `http://www.sun.com/bigadmin/content/selfheal/selfheal_overview.pdf`

Ordering Sun Documents

The SunDocs℠ program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is `http://docs.sun.com/`

To reference Sun BluePrints OnLine articles, visit the Sun BluePrints OnLine Web site at: `http://www.sun.com/blueprints/online.html`

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com