Dynamic Processor Deallocation
And
Dynamic Processor Resilience

White Paper

April 16, 2002

Revision 1.07

Revision Information

Revision 1.03: Initial revision reflecting the June 1999 IPR release.

Revision 1.04: Reflects changes release in the September 1999 Support Plus release.

Enhancement of Dynamic Processor Deallocation to keep processors deallocated even if the system is rebooted.

Enhancement of the STM System Information Tool to reflect the current state of all processors on the system.

Revision 1.05: Information added to reflect changes made in support of the Instant Capacity on Demand (iCOD) product.

Revision 1.06: Added information on differentiating Deallocation v/s Deconfiguration.

Revision 1.07: Added information about types of LPMCs and Floating-Point Tests

Introduction

The purpose of this white paper is to provide an overview of an exciting new technology that Hewlett Packard has developed that can significantly reduce system downtime due to processor failures. This technology, called Dynamic Processor Resilience, enables HP-UX systems to monitor the operation of processors, predict failures before they occur, and dynamically deallocate troubled processors before they experience catastrophic errors resulting in system failures. Dynamic Processor Resilience is one of the key technologies that enable HP to deliver industry-leading system availability and is available on all PA8500, PA8600 and future processors.

As the cache sizes incorporated into processors continues to increase, accounting for increasingly higher percentages of processor circuitry, it is critical that correctable cache errors be handled effectively in order to avoid processor-related system failures. On Hewlett Packard systems based on PA8500 or later processors, single-bit cache errors (a single erroneous bit in the data at any given cache memory location) are corrected. However, a double-bit cache error (two erroneous bits in the data) cannot be corrected and will result in a system failure. Statistically, most double-bit cache errors will be preceded by a series of single-bit errors over time as the memory cell begins to degrade. Using Hewlett Packard's Dynamic Processor Resilience and Dynamic Processor Deallocation technology, processor cache can be monitored for correctable errors and the processor dynamically deallocated before correctable errors turn uncorrectable.

Dynamic Processor Resilience also works hand-in-hand with HP's exciting new "instant Capacity On Demand" (iCOD) product. iCOD enables customers to purchase systems that have one or more processors in reserve, which have not yet been purchased. When additional capacity is required, the reserve processors can be purchased and "instantly" enabled. For systems, which have iCOD, enabled, reserve processors automatically replace processors that are deallocated by the EMS CPU monitor (previously named "LPMC monitor") if they exist, thus ensuring that the system continues to run at full capacity. The faulty processor can then be replaced when convenient at which time it will be returned to the reserve pool.

NOTE: Starting HWE 0206 release of Diagnostics, the "LPMC monitor" has been renamed "CPU monitor" in the documentation to reflect the fact that it monitors floating point functionality in addition to LPMCs. The binary name (lpmc_em) remains unchanged.

Types of Errors Addressed

LPMCs

There are four types of Single-Bit Cache Parity Errors that the processor can experience: **I-Cache data error, I-Cache Tag error, D-Cache Data error and D-Cache Tag error.**

NOTE: Starting HWE 0206 release of Diagnostics, the CPU monitor will keep track of each of these types of LPMCs rather than treating them as one type as in earlier version of the monitor.

Floating-Point Errors

Besides monitoring the Cache errors on the processors, the monitor will run tests on the Floating-Point registers to see if they are functioning properly. These tests are run at each Poll Interval (default 1 hr) and if the tests fail on any processor, the monitor will take the Dynamic Processor Resilience (DPR) action – as described later - on it right away. This action is similar to that taken in case of Cache errors. This functionality was added to help isolate CPU failures due to errors in Floating-Point Registers. These tests were available in the CPU Exerciser tool, which does not get executed regularly on a system. Since the monitor is always being executed, it was an ideal solution. The number of test-vectors to be run at each Poll Interval is user-configurable (via a configuration verb in ***lpmc_em.cfg*** file)

Technology Overview

Dynamic Processor Deactivation (DPD)

Incorporated into HP's 11.0 and future versions of HP-UX is the capability to take a processor out of service while the system is running without interruption to applications. This technology is referred to as **Dynamic Processor Deactivation**. Once a processor is deactivated, the HP-UX operating system will migrate all processes that are currently scheduled on that processor to other active processors. HP-UX will then no longer schedule processes to run on the deactivated processor until it is reactivated. Note that if the processor has been assigned to handle interrupts for any I/O drivers, it will continue to do so while it is inactive.

In order for Dynamic Processor Deactivation (DPD) to be available, the system must have at least two active processors. Dynamic Processor Deactivation is not available for the monarch processor--the processor upon which the HP-UX kernel is running. A processor may be deallocated programmatically via an HP-UX system call, via the Support Tools Manager (STM) CPU Expert tool (licensed for use by authorized service personnel).

Dynamic Processor Resilience (DPR)

Beginning with the June 1999 release of the IPR/Diagnostic media, an EMS monitor is provided which monitors the rate of correctable errors in each processor's on-board cache. These errors are manifested as Low Priority Machine Checks (LPMCs). While occasional correctable errors are to be expected in the on-board cache, too many of these errors in a short period of time indicate an increased likelihood that a non-correctable cache error could occur. The EMS CPU monitor will continuously monitor the rate at which LPMCs are occurring and dynamically deactivate a processor, using the Dynamic Processor Deactivation facility, if the factory determined threshold is exceeded. This technology is referred to as **Dynamic Processor Resilience**. For PA8500 processors, for example, the threshold is set at three LPMCs within a 24-hour time period. The monitor sets the threshold for different processors automatically.

NOTE: Starting IPR0009 release, this threshold value is no longer configurable.

NOTE: On N-Class, L-Class and later machines, the processor can be Marked-for-Deconfiguration and so that when the system is rebooted, the processor will be completely removed from system use. This action of removing the processor from the system is known as Processor Deconfiguration. On earlier PA8500-based machines, deconfigured processors will be reconfigured automatically upon reboot. On these machines, it is necessary to deconfigure processors manually via the Boot Console Handler (BCH) if they were Marked-for-Deconfiguration when the machine was rebooted.

NOTE: Starting HWE 0206 release of Diagnostics, the monitor will deactivate the processor *with a special O/S option*, so that it cannot be re-activated *without* rebooting the system. The purpose behind using the new option is to prevent system problems by continued use of the faulty processor in case the user decides to re-activate the processor using the CPU Expert Tool in STM.

The current state of all of the processors on the system can be determined via the STM System Information Tool.

The EMS CPU monitor generates informational EMS events for each correctable cache error that it detects. In order to prevent flooding the administrator with these events in the case where persistent cache errors are occurring, these informational events will cease to be generated once the threshold is met, a serious event is generated, and the processor is deallocated.

The CPU monitor receives immediate notification of LPMCs as they occur. Since no polling delays are involved, the monitor is able to take action the moment the correctable cache error rate exceeds the threshold. When the threshold is exceeded for

same type of cache error (I-Cache Data, I-Cache Tag, D-Cache Data and D-Cache Tag) , one of two actions will be taken:

1. **If the processor IS NOT the monarch**, the Dynamic Processor Deallocation facility will be invoked to deallocate it. The monitor will then generate a serious EMS event indicating that the processor was deallocated and should be scheduled for replacement (see example in figure 3). Subsequent to this, the processor state is checked once every 24 hours and a warning EMS event is generated if the processor is found to still be dallocated. This warning is intended to serve as a reminder that it is essential that the processor be scheduled for replacement.

   If the system is iCOD enabled and there are reserve processors available, a reserve processor will be immediately allocated to ensure that full processing capacity is maintained.

2. **If the processor IS the monarch processor**, it cannot be deallocated. In this case, a serious EMS event is generated, indicating that the processor is experiencing a high cache error rate, that it was not possible to dynamically deallocate it, and that it should be replaced as soon as possible before a catastrophic failure occurs. Subsequent to this, the processor will continue to be monitored. After a 24-hour period has elapsed, the monitor will reset its counters and begin generating informational events again for each correctable cache error that occurs. This time period, referred to as the “repeat frequency”, is configurable. If the cache errors persist and the threshold is met again, another serious event will be generated. Figure 1 depicts the sequence of EMS events that could be generated for a processor that is consistently experiencing cache errors over a period of 48 hours but cannot be deallocated.

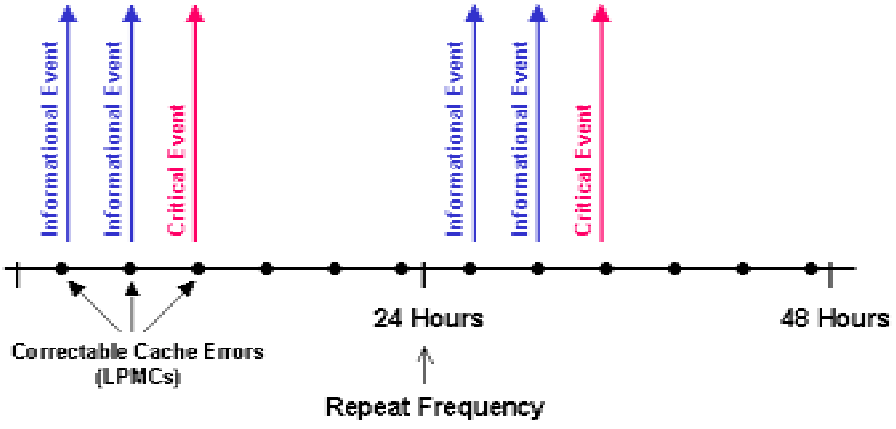

Figure 1 - EMS LPMC Event Sequence

NOTE: On N-Class, L-Class and later machines, the monitor will also try to mark the processor for Deconfiguration whether the processor in question is a Monarch CPU or not.

NOTE: The EMS CPU monitor will detect and prevent most processor failures that are related to cache errors. However, although cache errors account for the majority of processor failures, it is not possible to detect and prevent all processor related system failures.

For a complete listing on each event that can be generated by the EMS CPU monitor, refer to the event listings for the CPU monitor provided at:
http://docs.hp.com/hpux/onlinedocs/diag/ems/lpmc_em.htm

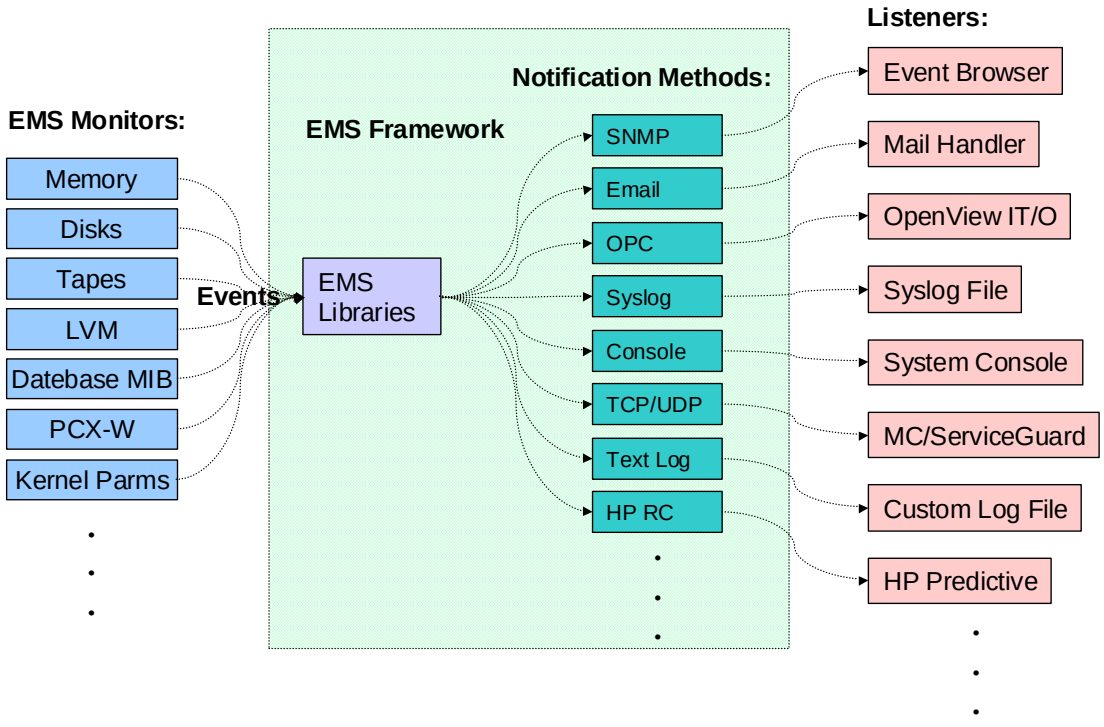

Figure 2 – The HP Event Monitoring Service

HP Event Monitoring Service (EMS)

Dynamic Processor Deallocation and Resilience (DPD and DPR) occur in context of the HP Event Monitoring Service. This is a comprehensive monitoring facility that

provides an array of monitors that continuously assess the health of various hardware and software components on HP-UX systems. These monitors will generate EMS events whenever a failure condition is detected. These events are then reported through various customer configurable reporting mechanisms such as e-mail, syslog, console messages, text logs, SNMP, etc. Figure 2 shows a high-level diagram of EMS.

Each EMS event contains information identifying the system that generated it, time stamp, and failure details including a description and cause/action text. Figure 3 shows an example of a serious event generated by the CPU monitor when it deallocated a processor.

More information on EMS hardware monitors can be found on Hewlett Packard's documentation web site for Diagnostics and Monitors at:

http://docs.hp.com/hpux/diag

```
-------------------------------------------------------------------------------
Event Time.........: Wed Feb 20 08:55:50 2002
Severity...........: SERIOUS
Monitor............: lpmc_em
Event #............: 100624
System.............: hptest17

Summary:

     Module at Hard Physical Address = 0xfffffffffc478000 : Cache Error(s)
     detected on Processor 16.

Description of Error:

     2 parity errors have been detected in either Instruction or Data memory
     (I-Cache or D-Cache) in 1 Day(s). These are indicated by 100611 100612
     EMS Event(s) generated prior to this event. The operating system has
     recovered from the errors, but this is an abnormally high failure rate.
     Contact your HP support representative to have the processor checked
     before a catastrophic failure occurs. The monitor tried to deactivate the
     processor for the current session but was not successful. However, the
     processor is deconfigured for the next boot session. The system should
     be rebooted as soon as possible.

Probable Cause / Recommended Action:

     Contact your HP support representative to have the processor checked.

Additional Event Data:
     System IP Address...: 15.8.147.100
     Event Id............: 0x3c73d51600000000
     Monitor Version.....: B.01.00
     Event Class.........: LPMC
     Client Configuration File...........:
     /var/stm/config/tools/monitor/default_lpmc_em.clcfg
     Client Configuration File Version...: A.01.00
          Qualification criteria met.
               Number of events..: 1
     Associated OS error log entry id(s):
          None
     Additional System Data:
          System Model Number.............: HP9000/HALFDOME
          EMS Version.....................: A.03.20
          STM Version.....................: A.30.00
     Latest information on this event:
          http://docs.hp.com/hpux/onlinedocs/diag/ems/lpmc_em.htm#100624

v-v-v-v-v-v-v-v-v-v-v-v-v    D  E  T  A  I  L  S    v-v-v-v-v-v-v-v-v-v-v-v-v

Component Data:
     HPA.....................: 0xfffffffffc478000
     Processor Number........: 16
     Physical Device Path....: 4/10
     Serial Number...........: N/A

Dynamic Processor Resilience action threshold met.
     Number of parity errors : 2
       Received in : 24 hour(s)
```

Figure 3 – Example Dynamic Processor Resilience Event

Dynamic Processor Resilience and HP Predictive Support

HP Predictive Support is a facility provided to HP Support customers that detects and predicts certain failures on customer systems and reports them directly to HP. Dynamic Processor Resilience is indirectly supported by this facility. While HP Predictive does not currently relay the events generated by the LMPC monitor back to HP, it does monitor the status of all processors on the system, senses when a processor has been deallocated, and reports it back to HP.

For further details on HP Predictive Support, contact your local HP Support office.

Availability of Dynamic Processor Resilience

EMS is provided free of charge as part of the IPR/Diagnostic media, starting with the June 1999 release. When installed, EMS is enabled by default, including the CPU monitor, providing immediate protection against avoidable processor failures. This monitor is available on HP-UX 11.0 for V-Class and N-Class systems.

NOTE: It is highly recommended that EMS remain enabled at all times. Disabling EMS will result in loss of the protection offered by Dynamic Processor Resilience.

Future Enhancements

Several enhancements to Dynamic Processor Deallocation and Dynamic Processor Resilience are under investigation. These include:

- Enhance Dynamic Processor Deallocation to migrate I/O interrupts off of deallocated processors.

- Enhance Dynamic Processor Deallocation to enable deallocation of the monarch processor.