

**The IBM @server pSeries 690**  
***Reliability, Availability, Serviceability (RAS)***

---

*Technical White Paper*

*IBM*

*September 2001*

---

## 1. Advanced RAS Features and Functions

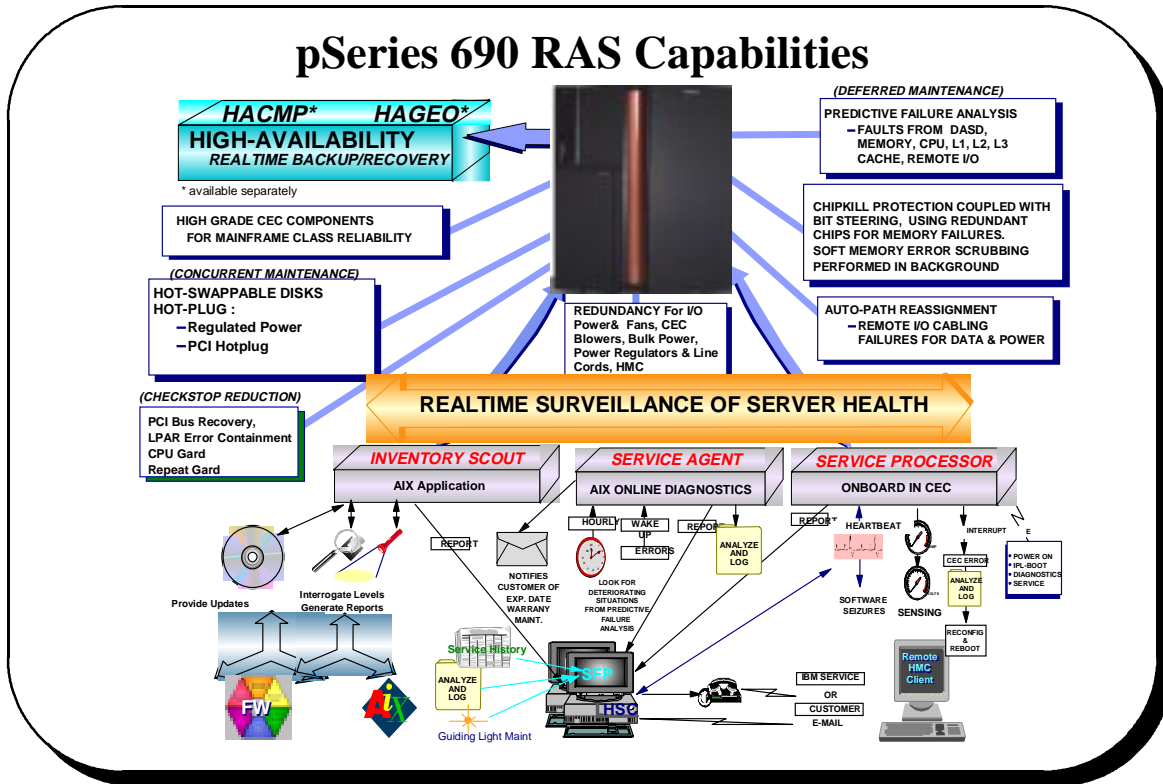
IBM has spent years developing RAS capabilities for mainframes and mission-critical servers. The IBM @server pSeries™ 690 has been able to take advantage of of this knowledge and experience with customer requirements.

The following features provide the pSeries 690 with UNIX industry-leading RAS:

- ◆ Automatic First-Failure Data Capture and diagnostic fault isolation capabilities
- ◆ Self-healing internal POWER4 processor array redundancy
- ◆ Industry-first PCI bus parity error recovery
- ◆ Checkstop-avoiding uncorrectable error handling with partition and process error containment
- ◆ Dynamic error recovery
- ◆ Error checking and correction (ECC) or equivalent protection on main storage, all cache levels 1, 2, and 3 and internal processor arrays
- ◆ Scrubbing and redundant bit-steering for self-healing in main storage
- ◆ Chipkill correction in main storage
- ◆ Fault tolerance with N+1 redundancy, dual line cords, and concurrent maintenance for power and cooling
- ◆ Predictive failure analysis on processors, caches, memory, I/O and DASD
- ◆ Processor run-time and boot-time de-allocation based on run-time errors (Dynamic Processor Deallocation and Persistent Processor Deallocation)
- ◆ Deallocation extended to caches and memory
- ◆ Fault avoidance through highly reliable component selection, component minimization and error mitigation technology internal to chips
- ◆ Concurrent run-time diagnostics based on First-Failure Data Capture for power, cooling, and I/O subsystems

Excellent quality and reliability are inherent in all facets of the p690 product. These measures are designed to ensure that products operate when required, perform reliably, efficiently handle infrequent failures in a nondisruptive fashion, and provide timely and competent repair in many cases either concurrently or on a deferred basis to allow operational resumption with minimal inconvenience. Mainframe diagnostic capability based on internal error checkers, First-Failure Data Capture, and run time analysis of all internal error check states is provided for all CPU, memory, I/O, power and cooling components, eliminating the need for re-creating failures.

## pSeries 690 RAS Capabilities



Reliability is one of the most significant factors in the design of high-end products. RAS (Reliability, Availability, and Serviceability) is an integral part of the p690 and AIX® Version 5 philosophy. It begins with the development of architectures, where RAS innovations are of paramount importance. It flows through design and product development stages, where RAS designs are reviewed, assessed, developed, evaluated, and perfected. It continues through the manufacturing and release processes, where the manufacturing quality is extensively measured and is under continual evaluation. It culminates in service and support; where the reliability is consistently monitored for deviation from the criteria, where warranty and maintenance have high priority, and where significant customer problems are assigned to and addressed by an expert team.

All of the development processes, from the architectural and concept phases of development, through the manufacturing process, and culminating in the provision of service and support are ISO-certified and audited periodically for ISO compliance by representatives of Underwriters Laboratories Inc.

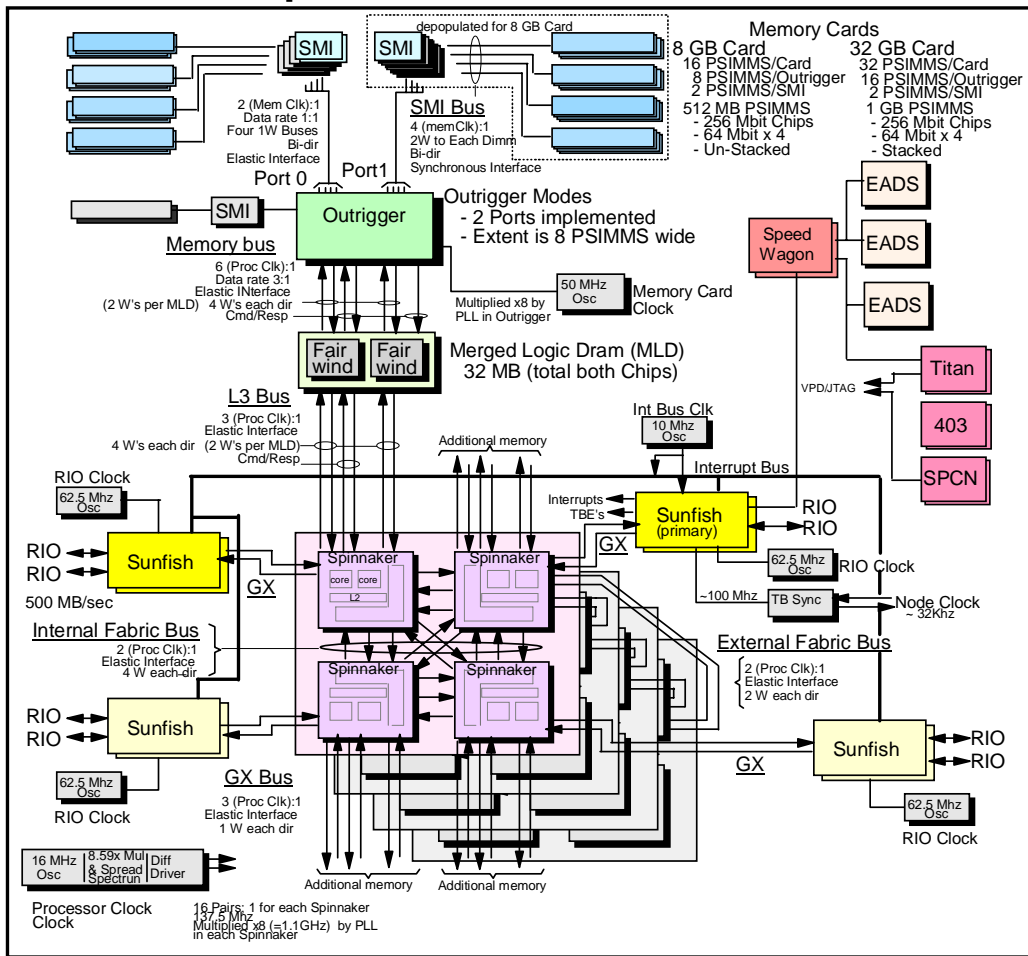
## 2. Reliability - Fault Avoidance

Major design efforts have contributed to the development of the p690 to analyze single points-of-failure within the Central Electronic Complex (CEC) to either eliminate them or to provide hardening capabilities to significantly reduce their probability of failure. The best way to harden a system is to prevent the errors from occurring in the first place.

Components within the CEC are designed to provide "mainframe" levels of reliability. These components provide the superior levels of reliability which are available and undergo additional stress testing and screening above and beyond the industry-standard components that are used in several UNIX-based systems today.

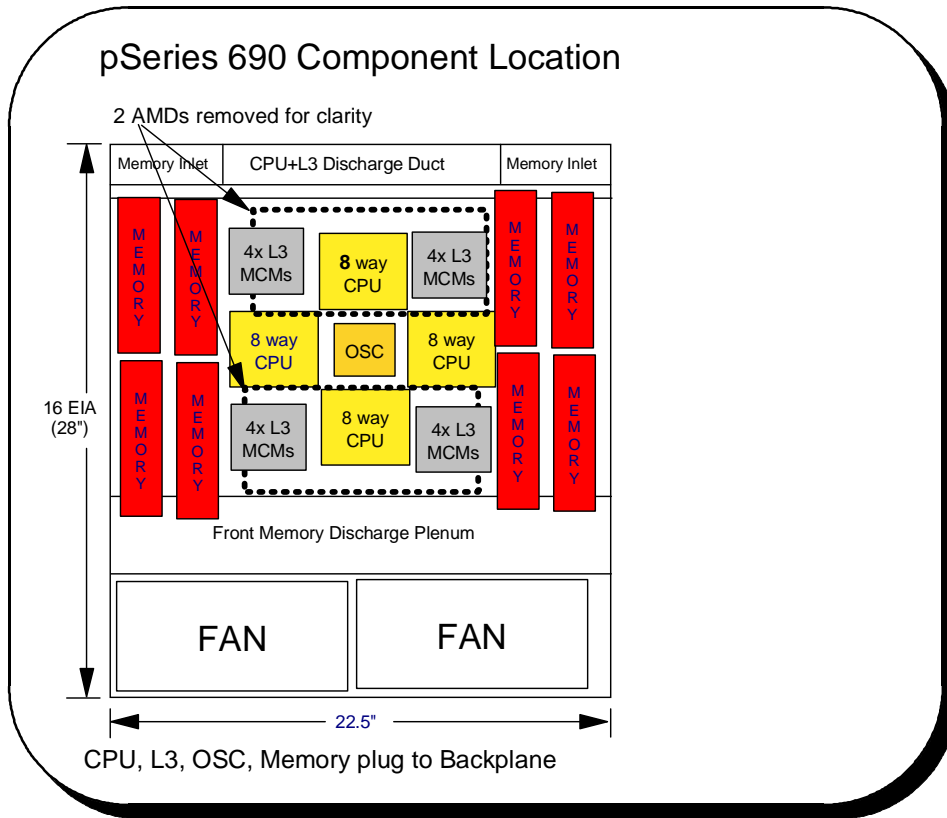
Fault avoidance is also served by minimizing the total number of components, and this is inherent in POWER4 chip technology, with 2 processors per chip. In addition, internal array soft errors throughout the POWER4 chipset have been systematically masked using internal ECC and recovery techniques. Going beyond ECC in the memory subsystem, the basic memory DIMM technology has been significantly improved in reliability through the use of more reliable soldered connection to the memory card.

### p690 CEC Structure



As seen on the architectural diagram above, the processors are mounted on a Multi-Chip Module (MCM) which is connected directly to the p690 backplane. This packaging methodology was carefully designed to provide increased levels of reliability by eliminating the tiered packaging levels of separate processor modules mounted on processor cards mounted on the backplane than some Unix servers utilize. In addition, the MCMs utilize the high density, copper technology mounted on a "Silicon-On-Insulator"

substrate to provide the densest, highest performance chips running at reduced temperatures which also increases reliability.



We use the term Air Moving Device (AMD) as a generic term for either a fan or a blower. Variable speed AMDs in the system unit along with variable speed AMDs in the I/O units allow for increased air flow to maintain proper cooling levels in case of an AMD fault.

This packaging provides for electromagnetic compatibility (EMC) shielding to minimize errors induced by electrical noise, and provides positive retention seating to help prevent shocks or vibrations from loosening critical system connections. These packaging features eliminate many of the intermittent errors experienced in UNIX servers containing less robust packaging.

During the design and development process, sub-systems go through rigorous verification and integration testing processes. During system manufacturing, the p690 system goes through a thorough testing process to ensure high product quality level. Extensive error detecting and checking circuitry helps maintain the integrity of data stored and transported in the system. The system design facilitates the recognition of component errors that are either corrected dynamically, or properly reported for isolation and repair. Parity on the system bus, cyclic redundancy checking (CRC) on the Remote I/O (RIO) bus, and the extensive use of ECC on memory and arrays provide some of these capabilities.

BIST (Built-In Self-Test) and POST (Power-On Self-Test) are designed to check the processors, caches, and associated hardware that are required for proper booting of the operating system every time the server is powered on. Additional testing can be selected at power-up time to fully verify the system memory and the chip interconnect wiring as an added reliability measure.

The system reboots in the extended test mode following a failure to check that all components are thoroughly tested and verified. If a non-critical error is detected, or if an error occurs in resources that can be deconfigured from the system, or if a processor has been marked for deconfiguration by Persistent Processor Deallocation, the boot process will attempt to proceed to completion with the faulty device automatically deconfigured. Detected errors are logged in the system non-volatile RAM (NVRAM). Run-time algorithms then gather the information from the NVRAM, perform analysis on it, and log it to the AIX log facility.

The AIX log facility is where hardware and software failures are recorded and analyzed by Diagnostic Error Log Analysis (ELA) routines to provide warnings to the system administrator on the possible causes of system problems. If the system concludes that service is required, an update is made to the Diagnostic Event Log which is monitored by Service Agent. Service agent analyzes the Diagnostic event log entry and determines whether to forward a service action event to the Service Focal Point application running on the Hardware System Console. If service is required, the Service Focal Point (SFP) will provide any filtering of duplicate service requests which may have been received from multiple sources or operating system partitions and initiate the call home for service. Data on the nature of the failure, the parts to be replaced and additional data on the specific machine configuration is gathered and transmitted to the service provider. This enables the service representatives to bring along needed replacement hardware components when a service call is placed, thus minimizing system repair time.

Disk predictive failure analysis provides the capability to detect an imminent disk failure and report the findings through the process described above. The disk can then be replaced on a deferred or concurrent basis, depending on the configuration of the disk subsystem. RAID controllers and hard disk hot-plug design often allow normal operation to continue in the presence of a failed disk, and a concurrent repair can be performed while the applications continue to run.

Surveillance of system operation is one of many functions provided by the service processor, which is a separate microprocessor subsystem whose many additional functions will be described in more detail in a later section. During boot time, a surveillance monitor in the service processor is automatically enabled to check for "heartbeats" from the boot firmware. If a heartbeat is not detected within a default period, the service processor is designed to cycle the system power and attempt to reboot until the system either boots successfully, or a retry threshold is reached. If the threshold is reached, the service processor logs the error, leaves the system powered on, and provides the user with various options to assist in diagnosing the error. The error logs can be interrogated and various

options exist for attempting to reboot the system. The service processor is designed to report the error to the Service Focal Point. This capability will be described in more detail in the later section on Serviceability.

In SMP mode, the service processor can also be optionally configured to monitor for operating system hangs or failures while providing a “heartbeat” indication to the operating system for which AIX is designed to detect a failed service processor. If enabled, the service processor can log operating system hangs or fails and report them the next time that the system is successfully booted. This function is disabled in the Logical Partitioning (LPAR) mode of operation.

Similar to the firmware surveillance scenario, the service processor can be enabled to notify the Service Focal Point and report the operating system surveillance failure condition.

---

### **3. Availability - Redundancy and Recovery**

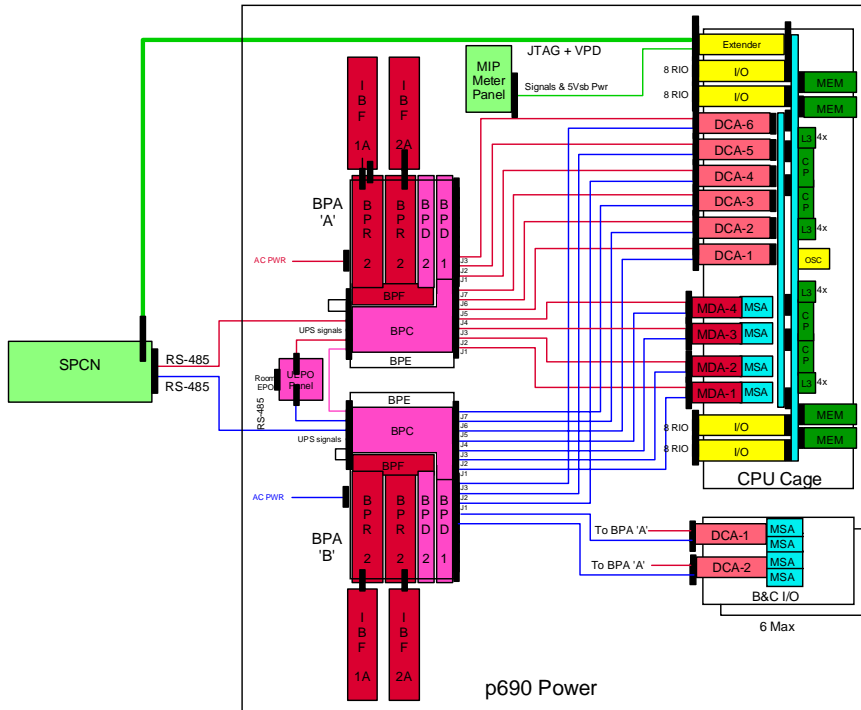
The p690 brings new mainframe levels of availability features and functions to the enterprise server arena which are described in the following sections.

#### **Power and Cooling Subsystem**

Because of the converged design point with the mainframe IBM @server zSeries™ for the CEC Rack power and cooling, the p690 brings new mainframe levels of availability features and functions to the pSeries enterprise servers. Within the CEC Rack, the N+1 power and cooling subsystem provides complete redundancy in case of failures in the bulk or regulated power supplies, the power controllers and the cooling units as well as the power distribution cables. As on the zSeries mainframe, concurrent repair is supported on

all of the power and cooling components.

## pSeries 690 Uses zSeries Power Subsystem



The p690 also features redundant Internal Battery Features (IBF) designed to maintain system operation during brown-out conditions. For full power loss protection, the p690 supports optional uninterruptible power supply (UPS) systems in addition to, or in place of, the IBF features.

The Journaled File System (JFS) in the AIX operating system is designed to maintain file system consistency and helps prevent data loss when the system is halted abnormally due to an interruption of power. In addition, power monitoring hardware is designed to provide early detection of loss of source power (i.e., black-out / brown-out conditions). This same power monitoring hardware also detects the loss of redundant bulk power supplies, regulators, fans, and blowers and internal battery features and reports them to the operating system for system error logging and maintenance notification. Components can be replaced either on a deferred or concurrent basis as previously discussed.

The I/O drawer extends availability by providing N+1 power supplies and fans. The power supplies can be repaired concurrently, while the fans can be repaired on a deferred basis.

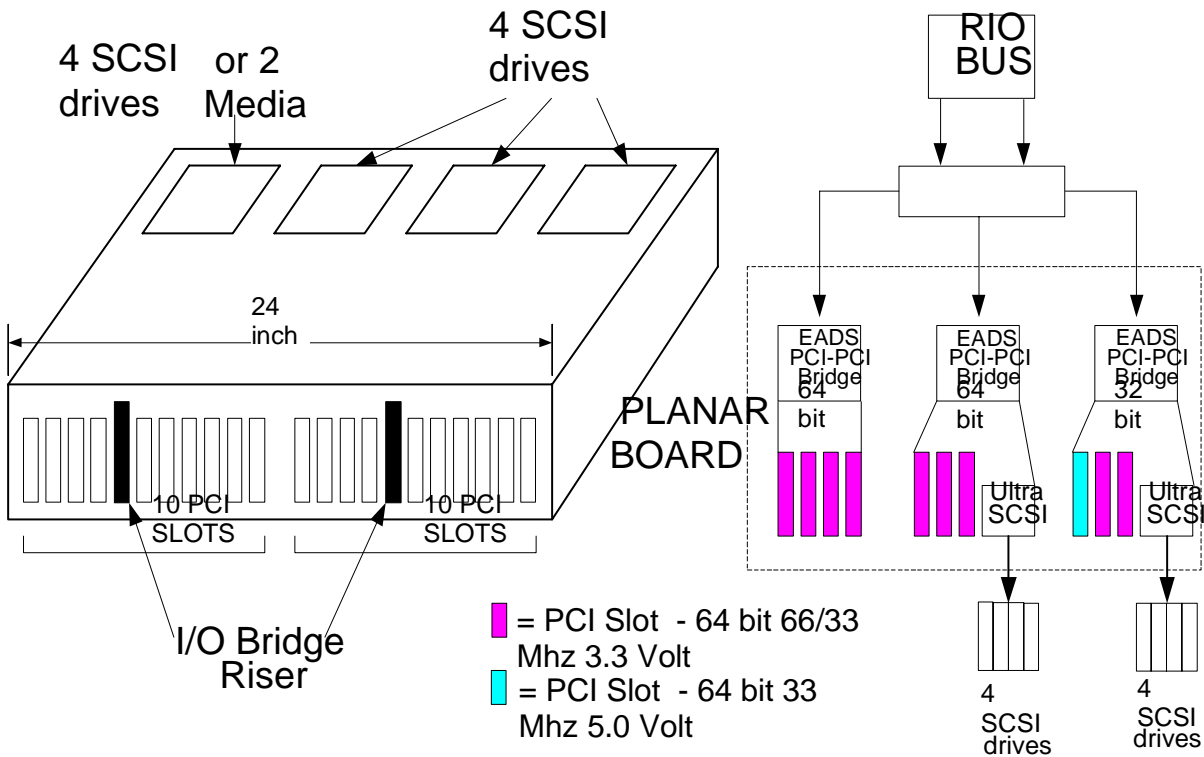
The interface from the processor to the I/O is through the Remote I/O (RIO) link. This link, in a similar method to the SSA interface, uses a loop interconnect technology to provide redundant paths to I/O drawers. RIO is a proven and robust interconnect technology. RIO availability features include CRC checking on the RIO bus with packet



retry on bus timeouts. In addition, if a RIO link fails, the hardware is designed to automatically initiate a RIO bus reassignment to route the data through the alternate path to its intended destination.

Power to the drawers is controlled from the power controller in the CEC through the SPCN link. This link is implemented in a point-to-point technology. Any break in the loop is recoverable via alternate routing through the other linkpath and can be reported to the service provider for a deferred repair.

## pSeries 690 I/O Drawer - Packaging for RAS



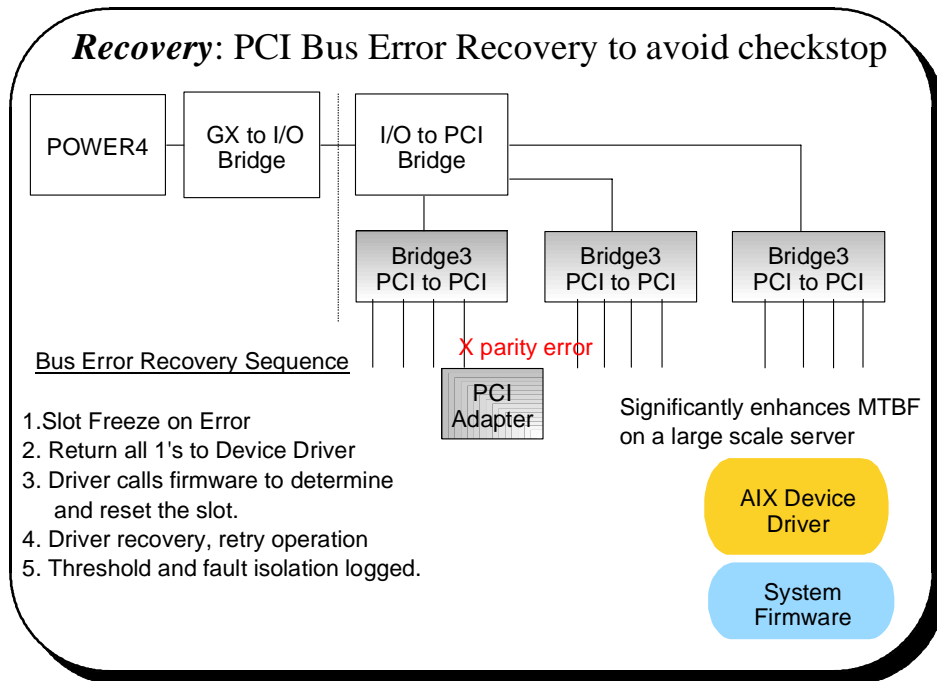
### ◆ p690 I/O Drawer:

In addition to the redundant power, cooling and cabling for the I/O drawer previously described, each of the four internal four-pack disk subsystems has its own backplane and is controlled by its own disk controller. These controllers are physically located on two independent planars and on reside on separate PCI busses thereby enabling mirroring across pairs of the four packs and reducing single points-of-failure within the I/O drawer. For additional levels of redundancy, mirroring can be configured across I/O drawers to further reduce exposure to disk failures due to a midplane failure within a drawer.

The disk drawers provide the ability for a disk to be hot plugged thus enabling the disks to be reconfigured or replaced dynamically. The p690 systems use AIX on-line (concurrent) diagnostics with Error Log Analysis and Service Aids that allow administrators or IBM service representatives to diagnose potential system malfunctions without interrupting end-user operations. Auto-restart (reboot) options, when enabled, can automatically reboot the system following an unrecoverable software error, software hang, hardware failure or environmentally-induced (AC power) failure as described in the reliability section. These standard high-availability features, coupled with the IBM High Availability Cluster Multiprocessing (HACMP for AIX) program product offering, offer outstanding server availability.

### PCI Bus Parity Error Recovery

In predecessor systems, PCI Bus parity errors cause a global machine check interrupt, which eventually requires a system reboot to continue. In the p690 system, a new I/O drawer hardware, system firmware, and AIX interaction has been designed to allow transparent recovery of intermittent PCI bus parity errors, and graceful transition to the I/O device unavailable state in the case of a permanent parity error in the PCI bus. This mechanism, the PCI Extended Error Handling (EEH), and is illustrated in the following diagram.



To support this, AIX device drivers must be written to a standard called EEH enabled, which requires that they respond to a special data packet generated from the affected PCI slot hardware by calling system firmware which will examine the affected bus, and allow the device driver to reset it, and continue. For maximum system availability and recovery, IBM recommends that only EEH enabled PCI adapters be used.

Each of the hot plug PCI adapter cards can be serviced from the rear of the system without having to move the drawer into a service position. This in-place repair provides greater reliability by eliminating the need to place the I/O drawer in a service position.

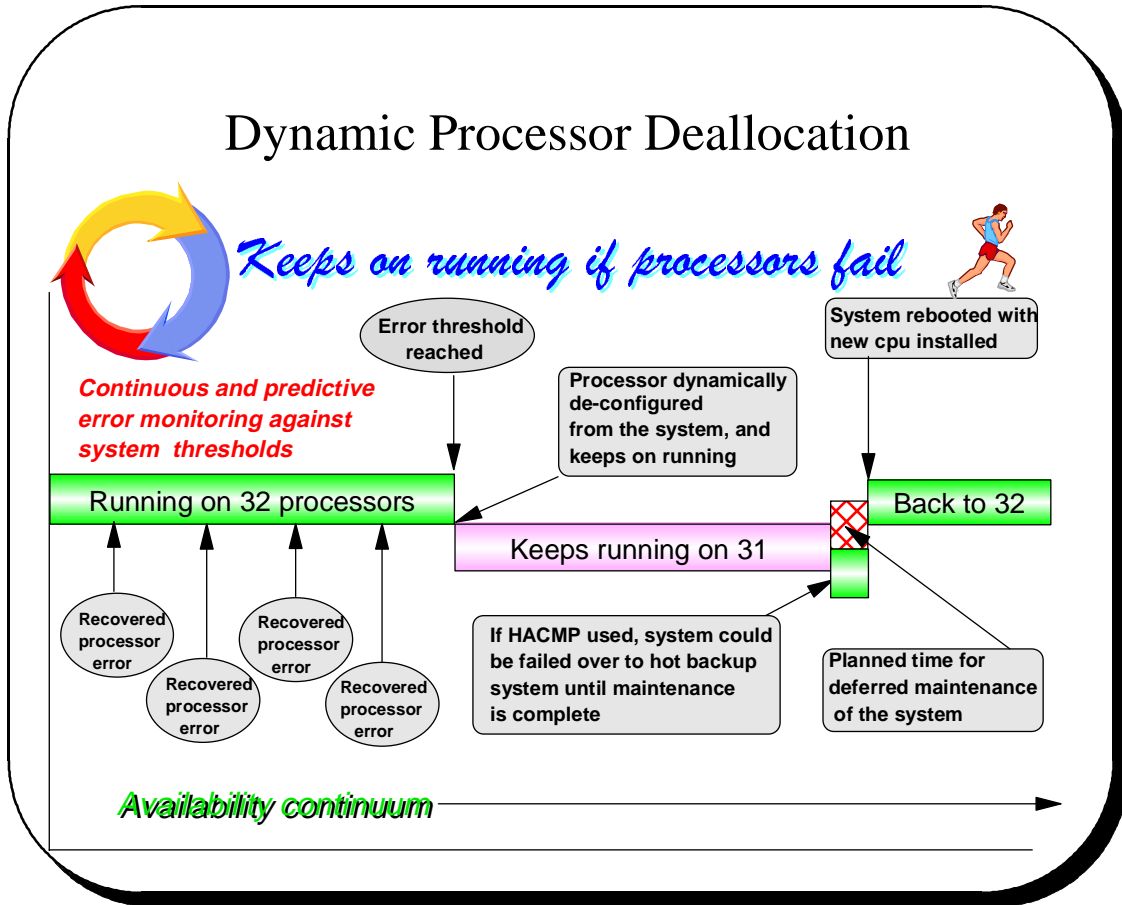
This involves sliding out the drawer and all of the attached cables in order to be able to physically remove the PCI adapter from the top of the drawer. This movement of the drawer in and out of service position can easily lead to inducing additional errors in PCI connectors and/or the associated cables plugged into the adapters.

This system provides carriers for the PCI adapters along with a “second generation” hot plug connection mechanism to allow individual hot plug PCI replacement without affecting surrounding adapters or moving the drawer into and out of a service position.

### **Dynamic Processor Deallocation**

Central Electronics Complex: Building on the robust features of the processor complex, the p690 system has enhanced error detection and prevention capabilities by using unique RAS features called Dynamic Processor Deallocation and Persistent Processor Deallocation.

Both utilize the pervasive error detection and isolation capabilities of the processor card along with the error recording facilities of the service processor to provide a method to record key fault information. Information is recorded when a processor has reached a threshold of recoverable errors.



If the failure of a part can be predicted before the failure results in a system outage, this enhances the reliability of the system. Predictive failure analysis is performed on recoverable processor errors during run-time. If a processor exceeds a threshold of recoverable errors such as on an L2 cache access, the event is logged. More than just logging the error, the processor will actually be marked and deconfigured from the system while the operating system continues to run. This feature allows the processor to be repaired on a deferred basis while helping prevent an unscheduled system outage.

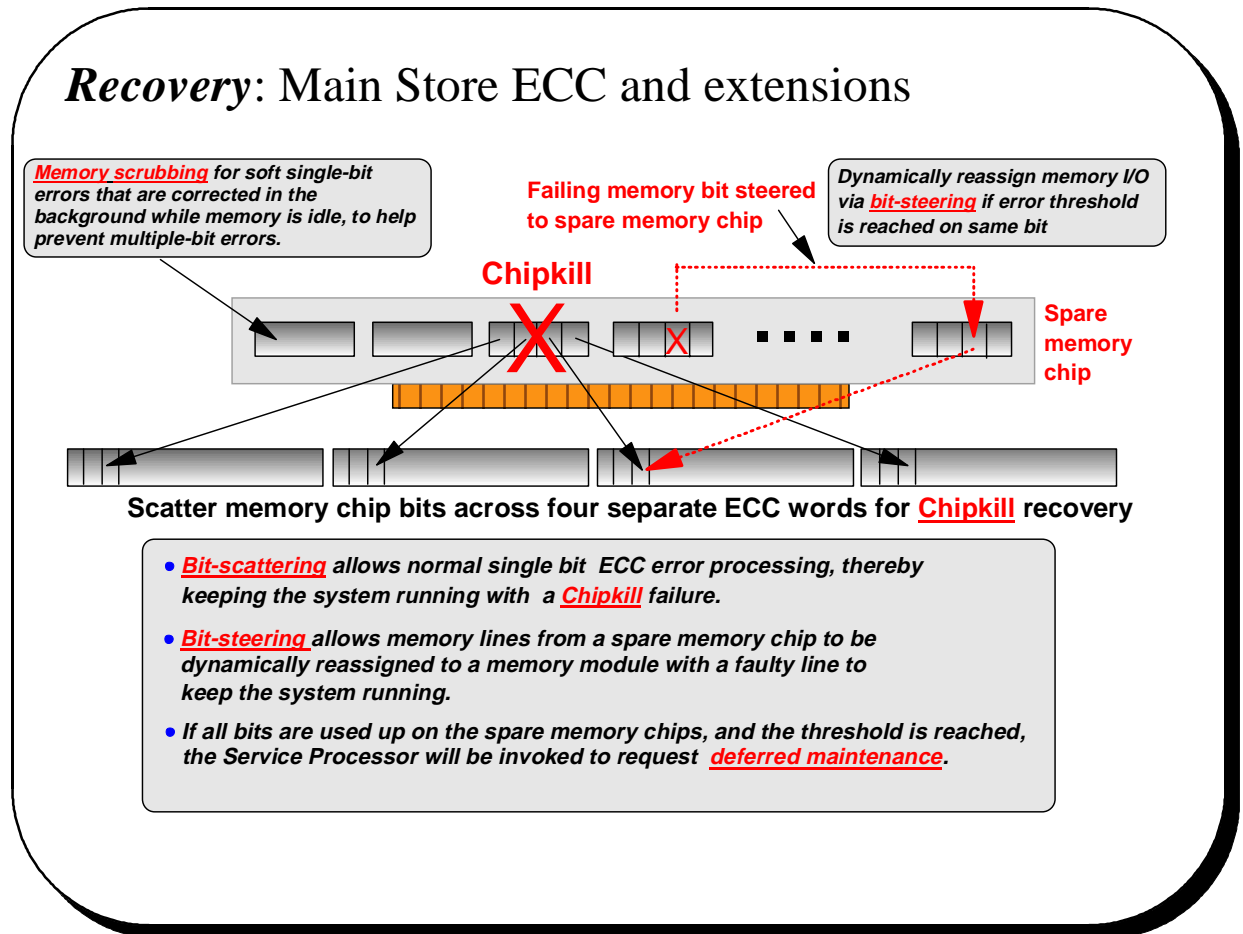
Persistent Processor Deallocation ensures that on subsequent reboots, the processor remains out of the system configuration until replaced.

Enabling of Dynamic Processor Deallocation and Persistent Processor Deallocation are customer options.

### Memory Error Correction Extensions

The standard memory card has single error-correct and double-error detect ECC circuitry to correct single-bit memory failures. The double-bit detection helps maintain data integrity by detecting and reporting multiple errors beyond what the ECC circuitry can correct. The memory chips are organized such that the failure of any specific memory

module only affects a single bit within an ECC word (bit scattering) thus allowing for error correction and continued operation in the presence of a complete chip failure (Chipkill recovery).



This memory card also utilizes memory scrubbing and thresholding to determine when spare memory modules within each bank of memory should be used to replace ones that have exceeded their threshold value (dynamic bit steering). Memory scrubbing is the process of reading the contents of memory during idle time and checking and correcting any single-bit errors that have accumulated. These single-bit failures could be either solid (technology failures) or soft failures (intermittent errors caused by noise or other cosmic effects). If an error is detected, the system hardware is designed to correct it by passing the data through the ECC logic that corrects the fault and then writing the corrected contents back to its memory address location. This function is also used to restore correct memory data after bit steering takes place. Scrubbing is a hardware function on the memory controller chips that takes place during memory idle time, and does not influence normal system memory reference performance.

In order to prevent an uncorrectable memory error from causing a system outage, the service processor is designed to initiate a deferred maintenance request on memory cards

that have used their spare bits and are experiencing additional correctable errors (memory predictive failure analysis).

### **Uncorrectable Error Handling**

All levels of storage protected by ECC are capable of failures whose resulting errors are uncorrectable because the failure corrupted more than one bit of an ECC word. While Chipkill ECC formatting eliminates these in mainstore, for internal arrays this technique can't be used. In the p690, checkstops due to uncorrectable errors are mitigated by new hardware behavior which marks data at the first instance it is detected, notifies the service processor of the exact physical element causing the error, and moves the marked data throughout the storage hierarchy without further error notification until it is referenced in a CPU. The CPU hardware presents a synchronous machine check interrupt to AIX, with hardware state indicating the address of the referring instruction. In the case of LPAR, the effect will be a software reboot of the partition. The ultimate capability of this error handling will be software process terminate rather than partition reboot.

---

## **4. Availability - First Failure Diagnostics and Reconfiguration**

The ability to correctly diagnose problems in a computer is the bedrock capability upon which availability is based, and without this pervasive capability, even simple problems which behave intermittently can be a cause for serious and prolonged outages. p690 provides un-matched capability in both IPL and initialization diagnostics, based on internal test procedures, and in run-time first failure diagnostics based on strategic error checkers operating full time to detect and capture precise error signatures with pre-determined hardware fault domains.

### **Value of run-time diagnostics to repeat outages and reconfiguration capability**

Based on experience with servers implemented without the run time first failure diagnostic capability, it is possible to project that high impact outages would occur 2 to 3 times more frequently without this capability. This is especially true for a design with any form of deconfiguration as a fundamental availability feature. It is not possible for a system to automatically deconfigure an element that can't be diagnosed. Experience and field data comparing products with and without this capability show a marked reduction in high impact outages for servers with run time first failure error capture diagnostics as implemented in the p690 family.

The diagnostics goal for the p690 system is to isolate 95% of the failures to a single Field Replaceable Unit (FRU). For 5% of the failures, two FRUs plus any boards or wires that interconnect the FRUs are candidates for fault identification. In these 5% of cases, manual isolation procedures may be employed by the service person. In order to attain these isolation goals based on error checkers alone, the entire system logic design must contain strategically placed error checkers.

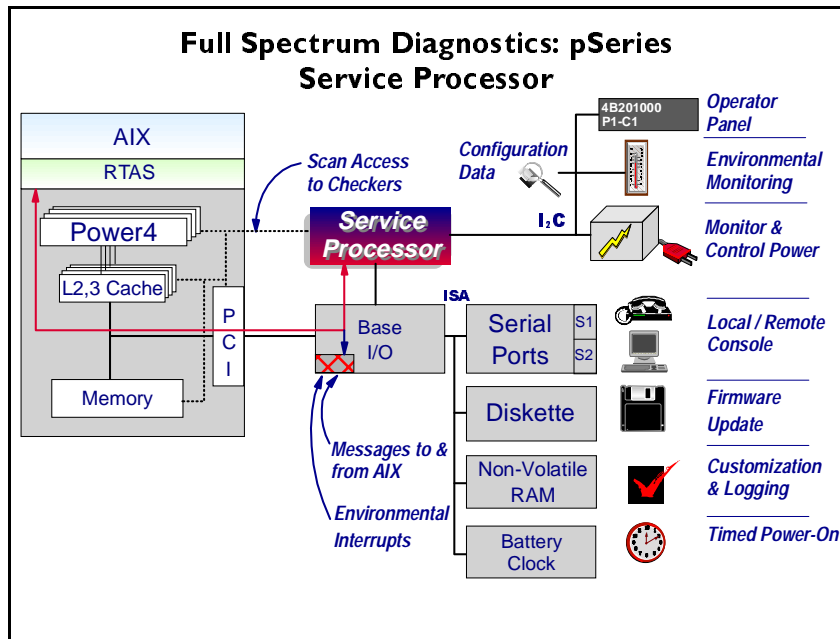
All p690 error checking mechanisms, including parity, ECC, and control checks, have three distinct but related attributes. First, checkers provide data integrity. Second, checkers initiate appropriate recovery mechanisms, from bus retry based on parity error detection, to ECC correction based on hardware detection of a non-zero syndrome in the ECC logic, to firmware executing recovery routine based on parity detection. Third, and equally important, all error check stations have been placed in the data and control paths of p690 systems to deterministically isolate physical faults based on run-time detection of each unique failure that may occur.

All error checkers are instrumented with software readable error capture Fault Isolation Registers (FIRs) and blocking logic so that for every detected error the error is recorded only by the first checker that encounters it. This form of instantaneous run-time diagnostics greatly enhances other forms of diagnostic testing, such as BIST, which relies on reproducible defects, rather than intermittent ones often present or evident only at run-time. Run-time error diagnostics are deterministic, in that for every check station, the unique error domain for that checker is defined and documented. Diagnostic validation consists of dynamic run-time injection of intermittent error conditions, to determine that the correct physical component is called out by the diagnostic.

To accomplish the error detection and fault isolation objectives stated above, there are over 5600 fault isolation register bits representing over 15000 internal error checkers.

### Role of Service Processor in run time diagnostics

The role of the Service Processor in FRU isolation is similar to that in IBM 308X, 3090, and 9021 machines.



The Service Processor is a separate, independent processor that provides hardware initialization during system IPL, operation monitoring of environmental and error events, and maintenance support for the p690. For run time diagnostic purposes, the

communication between service processor and the p690 consists of (1) Attention signals from the p690 hardware and (2) Read/Write communication between the service processor and all hardware internal FIRs, using specialized JTAG ports between the service processor and all p690 chips. This diagnostic Read/Write capability of hardware error registers is simultaneous, asynchronous, and transparent to any system activity running on p690 H. Said another way, these FIRs are known only to the service processor, and are not accessible by system software.

The system is designed to generate an appropriate Attention signal to the service processor when an error is detected in hardware. The ultimate response of the service processor is to Read the appropriate FIR, based on analysis of the “Who’s On First” (WOF) structure, and to examine the active FIR bits, and post the FRU callout in the p690 NVRAM. The NVRAM acts as a mailbox between service processor, system firmware and the AIX running on the p690 H. The architected FRU callout is moved by system firmware to AIX and into the AIX system Error Log, along with notification about the nature of the event, usually a deferred repair, based on the p690 internal element availability mode ( one of CPU, L2, L3, memory, or PCI adapter offline). Following the analysis of a recovered event, the service processor resets the FIRs so they can accurately record any future error events.

### **Run time diagnostics for PCI adapters**

From the above, it is clear that one of the critical corequisites of run time diagnostic isolation is run-time access to internal error state. For the p690 unique hardware, this access is via special service processor data ports or hardware implementation specific memory mapped error register for runtime system firmware access. For industry standard PCI adapters, this scan path access is not currently available. The alternate is to use the existing adapter device driver access to error sense information contained within the adapter hardware to indicate the failure state. IBM encourages device drivers to be written so that they respond to any adapter error indication by reading and logging into the AIX error log all sense data from the adapter. This sense data is examined by an AIX function called Diagnostic Error Log Analysis, which creates the appropriate FRU callout. Before supporting any PCI adapter in an AIX release, the adapter specification is examined to be sure that all available internal sense data is (1) logged and (2) defined so that Diagnostic Error Log analysis will call the correct PCI adapter FRU. This behavior is also tested by IBM during Maintenance Package Verification for the adapter.

### **Redundancy for array self-healing**

While the most likely failure event in a processor is a soft single bit error in one of its caches, there are other events which can occur, and which need to be distinguished from one another. For the L1, L2, L3 caches and their directories, hardware and firmware keeps track of whether permanent errors are being corrected beyond a threshold. If exceeded, a deferred repair error log is created. Additional run-time availability actions, such as CPU vary off, or L3 cache line delete, are also initiated. L1, L2 caches and L2, L3 directories on the POWER4 chip are manufactured with spare bits in their arrays which can be accessed via programmable steering logic to replace faulty bits in the respective arrays. This is analogous to the redundant bit steering employed in main store as a



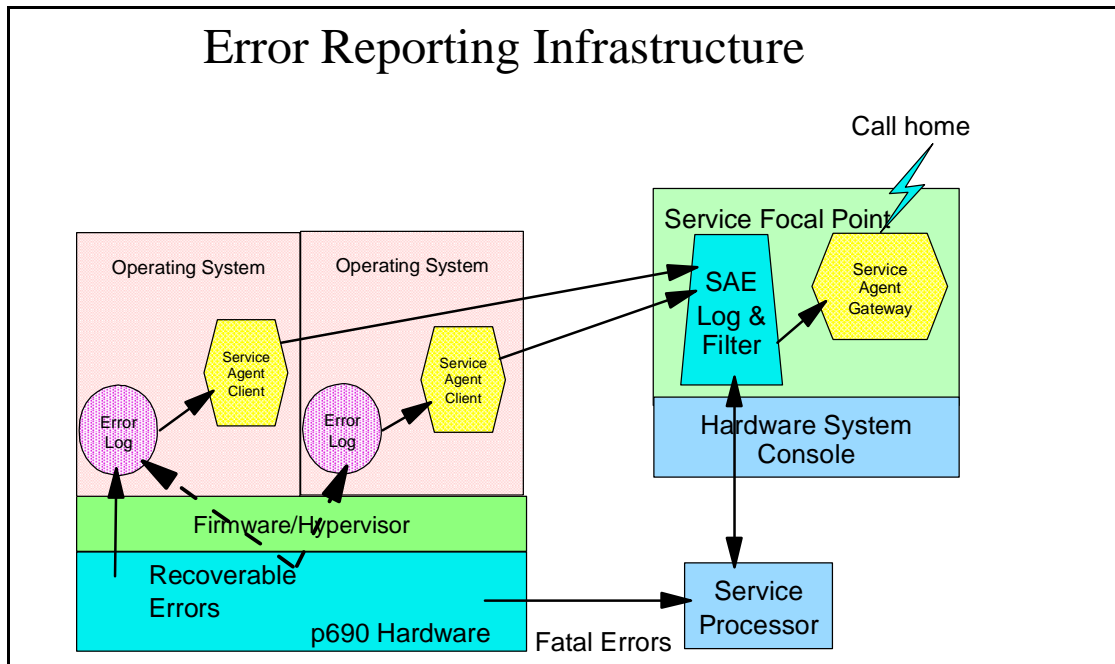
mechanism to avoid physical repair that is also implemented in POWER4 systems. The steering logic is activated during processor initialization and is initiated by the Built-in System Test (BIST) at Power On time. L3 cache redundancy is implemented at the cache line granularity level. Exceeding correctable error thresholds while running causes invocation of a dynamic L3 cache line delete function, capable of up to 2 deletes per cache. In the rare event of solid bit errors exceeding this quantity, the cache continues to run, but a message calling for deferred repair is issued. If the p690 system is rebooted without such repair, the L3 cache is placed in bypass mode, and the system comes up with this cache deconfigured.

---

## 5. Serviceability - Effective Problem Resolution

The pSeries 690 is designed to be installed and maintained by a trained service representative.

There are three main components to the p690 error reporting strategy. The first component is the service processor and the second is Service Agent. These two components provide reporting capabilities on unexpected changes in the system environment to the Service Focal Point (SFP) application running on the Hardware System Console (HSC). The SFP is the third component in the error reporting strategy and provides the error filtering, extended error data gathering and call home capabilities to report the service action request to the service provider as well as notifying the customer. A high level overview will be provided first with more detailed specifics on each component to follow.



System failures that prevent the system from coming back to an operational state (ie. Operating System inoperative) will be reported to the SFP application by the service processor. System failures that do not prevent the system from coming back to an operational state (recovered through hardware, firmware or software techniques) will be reported by the Service Agent application since the operating system is operational.

## **Service Processor**

The service processor provides for excellent RAS service features such as first failure data capture analysis explained in the prior availability section and surveillance monitoring described previously. It also provides functions such as; power-on/off of the system, reading the service processor and POST error logs, reading vital product data (VPD), changing the bootlist, viewing boot sequence history, and changing service processor configuration parameters, all of which can be performed remotely. Customers can enable console mirroring on the system console so they can monitor all remote console activity. For this option to work, a modem must be attached to one of the serial ports and configured appropriately.

## **Service Agent**

The Service Agent application is shipped standard with the p690 system. There is no additional charge for its usage as long as it is covered by an IBM Warranty or an IBM Maintenance Agreement. There are two main components of the service agent application. The client version resides on each copy of the operating system and monitors the system while AIX is running. The Service Agent monitors and analyzes all recoverable system failures, and, if needed, can automatically forward a service action event request to the service focal point (SFP) application running on the Hardware Management Console (HMC).

The second component of Service Agent is the gateway function and resides on the HMC as part of the SFP application and provides the consolidated focal point to place a service call to the service provider (without customer intervention). By doing so, the product can reduce the amount of downtime experienced in the event of a system component failure by giving the service provider the ability to view the error report entry, and if needed, order any necessary replacement parts, prior to arriving on-site. The opportunity for human misinterpretation or miscommunication in problem determination is therefore mitigated.

Service Agent is designed to automatically report problems based on default settings. The customer may modify the default values to prevent service agent from placing a service call during hardware upgrades, testing, or in the event that the failed component is not covered by an IBM Service Agreement (e.g., a third-party disk subsystem). The customer can also configure the product to only alert personnel within the customer's IT department (via e-mail). This function can be configured instead of, or in combination with, the ability to automatically place a service call to the IBM Service Center.

Service Agent is driven on a timed cycle (by default, set to one hour). At the completion of the cycle, the service director awakens to check diagnostic results, and failed or pending transmission events. In the event of an entry in the AIX Error Report, however,

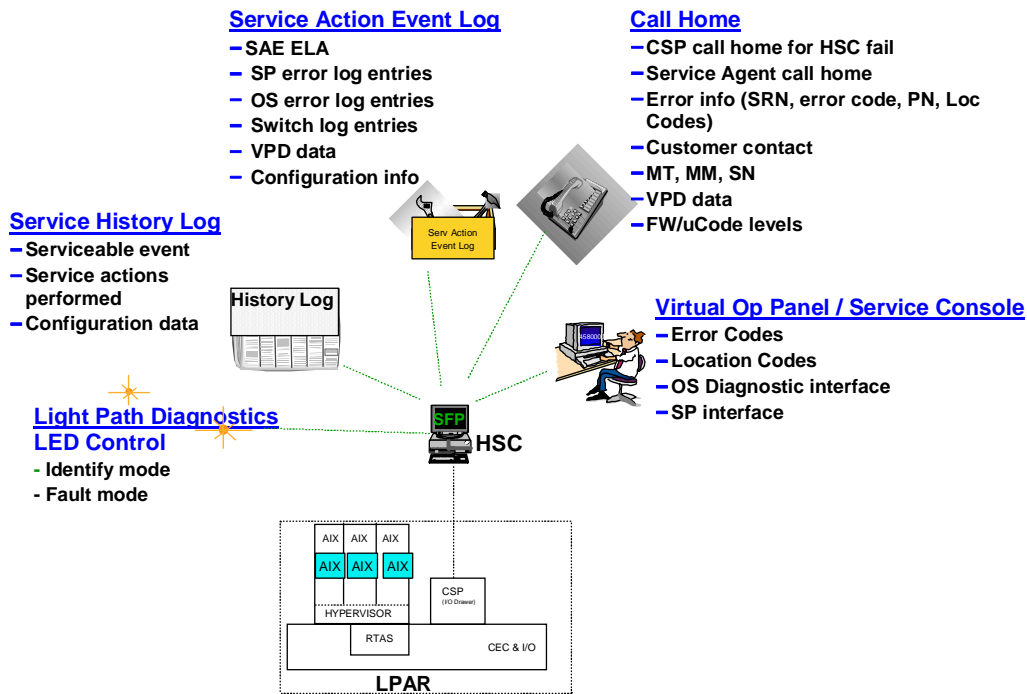
the Service Agent is designed to automatically start actions to prepare and send a request for service.

If the customer uses IBM, the IBM Service Center receives the machine type, serial number, host name, Service Request Number (SRN), extended error data such as configuration information and a problem description (taken directly from the failing machine's error report). The Service Center analyzes the problem report and determines whether service action is necessary. The Service Center will also determine if any hardware components need to be ordered to complete the service action.

The Service Agent Gateway code on the SFP application also gives the customer the option to establish a particular SFP and HMC system as the problem reporting server. A single SFP/HSC, accessible over the customer network, can be used as the central server for all other machines on the local LAN who are running the Service Agent client application (a second system can be configured as a backup Service Director to the central server). If the Service Agent application on a remote client decides a service request needs to be placed, it forwards the pertinent information to the Service Agent gateway server who dials the service provider from its locally attached modem. In this scenario, the customer only needs to maintain a single analog line for providing call-out capabilities for a large set of servers. This capability is a significant enhancement which provides unified reporting of both fatal and recoverable errors for all lan attached servers utilizing a single phone line. Communication between systems is handled through the Remote Procedure Call functionality (RPC) provided in AIX.

## Service Focal Point

# Service Focal Point Overview



LPAR environments usually add complexity to servicing, but the p690 ships with software to reduce this complexity. In order to accommodate error reporting, analysis and repair in the LPAR environment, a new application was developed to run on the Hardware Management Console (HMC). This application is called the Service Focal Point (SFP) and leverages the design capabilities of the HMC to provide equivalent "virtual" function to the current capabilities presented by physical op panels, service processor TTY menu interfaces and system firmware interfaces as well as capability for configuring/reconfiguring building block hardware into partitions.

The Service Focal Point is a system infrastructure which manages serviceable event information for the system building blocks. It consists of resource managers that monitor and record information about different objects in the system. It is designed to filter and correlate events from the resource managers and initiate a call to the service provider when appropriate. It also provides a user interface which allows a user to view the events and perform problem determination. When a problem is corrected the user can record actions that have been taken to resolve the hardware problem. This stored data can then be accessed by service representatives on future calls to determine what actions have already been taken on the system and adjust the service action plan accordingly. These features of SFP support the overall problem management strategy in a complex system.

The SFP application receives service action events from the service processor for critical system down situations, and from the Service Agent application programs running on the individual logical partitions for system recoverable or predictive events as well as Operating System or device driver detected events.

### **Service Action Event Log**

The SFP is designed to collect the serviceable events from different building blocks together in a Service Action Event (SAE) log. The log entries are generated by analysis routines that run on an error that has occurred in a building block. The resource manager for the building block forwards information about the event to the service focal point and the information is placed in the SAE log. The particular content of the error data depends upon the type of the error and on the system configuration itself.

The SAE log on the SFP also contains pointers to extended information that may have been recorded at the time of a serviceable event by the building block. Extended error collection includes not only the collection of first failure data capture, but also vital product data, partition information, operating system error logs, service processor error logs, error register data, etc.

When the SFP receives a new log entry, filtering is done to determine if this is a unique event. The filtering is done because sometimes an event notification can come from more than one resource manager for the same event or a resource manager may forward a notification for an event which previously occurred but has not yet been corrected.


### **Service Agent Component**

When a service action event is logged in the SFP, the system needs to communicate the failure back to the service provider. During this "call-home" function, particular error data and system configuration information needs to be sent to the service provider to drive the service delivery infrastructure. The SFP utilizes the Service Agent focal point application residing on the HSC along with the HSC modem to initiate the call home and transfer the pertinent error information to the service provider. When a call home is required, Service Agent manages the connection to the service provider which is used to open a problem record. The problem record is used by the service delivery team to determine whether or not to dispatch a customer engineer (CE) with the appropriate service parts to the system to perform a repair.

When a CE does perform a repair on the system, the SFP is used to identify the source of the problem and record information relating to the repair. When the CE has performed a repair, the SAE log entry will be updated with FRU replacement information and any comments that the CE has. The information stored by the SFP represents the system's service history and is used to ensure proper maintenance over the life of the system.

## Light Path Diagnostics

### Frame/Rack Identifier Light



- ▶ Frame/Rack LED mounted visible along w/drawer LED
- ▶ Set to Flash when in Identify mode
- ▶ Reset as part of system verification process

To assist the IBM service representative in locating the correct system unit and drawer that contain the fault requiring repair, the SFP will enable the capability to flash LEDs on the respective system unit and drawer that contain the fault

### Microcode Discovery Service

Microcode Discovery Service provides the capability to determine if the p690 is at the latest microcode and firmware levels. Using a secure Internet connection and a Web browser, Microcode Discovery Service captures the machine data and generates a real-time comparison report showing subsystems that may need to be updated.

The Inventory Scout application will run as a daemon on the server to accomplish this function. The tool will create a file containing the current level of all microcode (adapters, devices, system, and support processor) levels in the system. This file will be used to compare the system level codes against the latest available levels on the IBM Web site. A

report is then generated identifying any new updates available along with a link from which the updates can be downloaded.

Flash updates for firmware are performed by trained service personnel. The microcode updates are available from a support page on the internet. From there, the code is downloaded to the server and installed. There are several options for downloading the updates, which include downloading to AIX workstations as well as DOS, OS/2®, or Windows®-based PC workstations. Also, the update can be downloaded directly to the p690 and then installed. This capability allows the update to be performed remotely. Service aids or AIX command line options can be used to install the updates.

---

## 6. Service Support

The IBM @server pSeries 690 is a new generation of 64-bit SMP enterprise servers designed for mission-critical applications. The service and technical support structure for this system reflect the importance that this product will play in your business.

Hardware service requests will go to IBM's remote support center for initial problem diagnosis. This approach provides more direct access to skilled specialists. These specialists can either solve the problem over the phone or help get it resolved as quickly as possible by identifying the failing part or component and the specific skills required to resolve the problem. Service specialists are backed up by a Product Engineering team that has been highly trained and provided with additional tools to assist in problem identification and resolution.

---

## 7. RAS Conclusion

The RAS features and functions designed into the base system, extended by features such as RAID controllers, the HACMP program product, and remote service capabilities, combine to make the p690 a system that will meet the needs of a mission-critical marketplace, and provide the growth, expansion, and performance required by this market.

The p690 has been designed to be the most reliable high end UNIX server ever produced by IBM, delivering more than 5X better MTBF when compared to already highly reliable IBM high end UNIX servers.



---

## 8. Notices

© International Business Machines Corporation 2001

IBM Corporation  
Marketing Communications  
Server Group  
Route 100  
Somers, NY 10589

Produced in the United States of America  
10-01 All Rights Reserved

More details on IBM UNIX hardware, software and solutions may be found at:  
[ibm.com/servers/eserver/pseries](http://ibm.com/servers/eserver/pseries).

You can find notices, including applicable legal information, trademark attribution, and notes on benchmark and performance at [ibm.com/rs6000/hardware/specnote.html](http://ibm.com/rs6000/hardware/specnote.html)

IBM, the IBM logo, AIX, pSeries, OS/2, zSeries, Chipkill are registered trademarks or trademarks of the International Business Machines Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Limited.

Windows is a registered trademark of the Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM may not offer the products, programs, services or features discussed herein in other countries, and the information may be subject to change without notice.

General availability may vary by geography.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Any performance data contained in this document was determined in a controlled environment. Results obtained in other operating environments may vary significantly.

IBM may have patents or pending patent applications covering subject matter in this paper. The furnishing of this presentation does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.