



# IBM Power Systems: Designed for Reliability

IBM Power™ Platform  
Reliability, Availability and Serviceability (RAS)

Jim Mitchell, Daniel Henderson, George Ahrens,  
and Brian Warner

March 6, 2009

<b>Overview: Designing for Reliability</b> .....	<b>3</b>
<b>Reliability and Highly Available Servers</b> .....	<b>3</b>
<b>Reliability Design Processes at IBM</b> .....	<b>4</b>
<i>Setting Targets</i> .....	5
<i>Evaluating Components, Packaging Accordingly</i> .....	5
<i>Redundancy and Concurrency in All the Right Places</i> .....	5
<i>Continuous Field Monitoring</i> .....	5
<b>IBM Power Systems – Designed for Reliability</b> .....	<b>6</b>

## Overview: Designing for Reliability

The base reliability of a computing system is, at its most fundamental level, dependent upon the intrinsic failure rates of the components that comprise it. Very simply, highly reliable servers are built with highly reliable components. On Power Systems, this basic premise is augmented with a clear “design for reliability” architecture and methodology. The POWER6 reliability strategy evolves from, and improves upon, the reliability design points developed throughout the IBM POWER™ program.

At IBM, trained RAS engineers use a concentrated, systematic, architecture-based approach with the objective to improve overall server reliability with each successive generation of system offerings. At the core of this effort is an intensive focus on sensible, well-managed server design strategies that both stress high system instruction execution performance and require logic circuit implementations that operate consistently and reliably despite potentially wide disparities in manufacturing process variances and operating environments. Intensive critical circuit path modeling and simulation procedures are used to identify critical system timing dependencies, so that time-dependent system operations complete successfully under a wide variety of process tolerances.

This white paper provides an overview of the design points that contribute to a reliable POWER6 processor-based system.

## Reliability and Highly Available Servers

A highly available server (one that avoids end user application outages) incorporates three principles:

1. A “hardened” design: that is, use of reliable components
2. A “redundant” design: a system that duplicates key or critical components
3. An “error handling” design: a system that can actively handle (correct / work around / avoid) faults

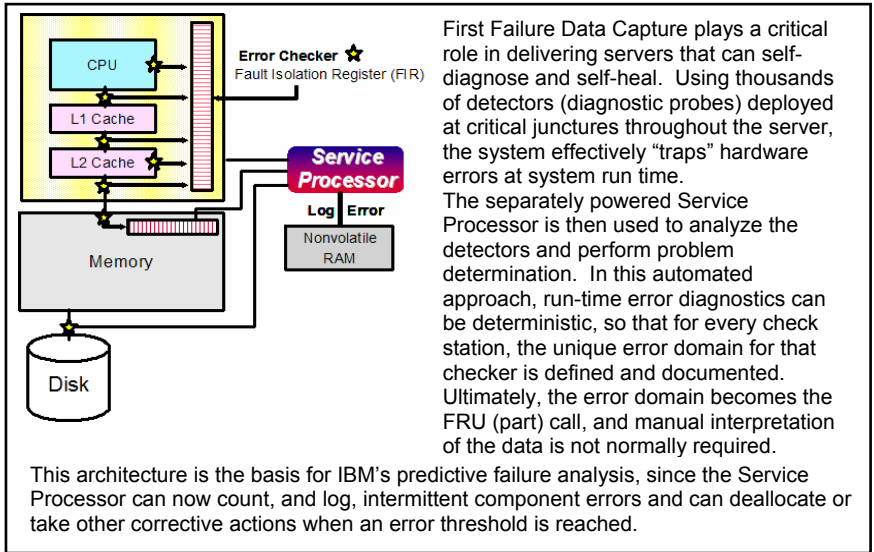
These core principles guide the IBM engineering design team and are reflected in the RAS architecture. The goal of any server design is to:

1. Achieve a highly reliable design through extensive use of highly reliable components built into a system package that supports an environment conducive to their proper operation.
2. Clearly identify, early in the server design process, those components that have the highest opportunity for failure. Employ a server architecture that allows the system to recover from intermittent (random recoverable) errors in these components and/or failover to redundant components when necessary.
  - Employ automated retry for error recovery of: failed operations, failed data transfers in the I/O subsystem, corrupted data.
  - Use sparing (redundancy) strategies including:
    - Duplicated functions
    - N+1 redundancy
    - Fine grained redundancy

IBM engineers draw upon an extensive record of reliability data collected over decades of design and operation of high-end servers. Detailed component failure rate data is combined with data developed through tracking new technology deployments to determine both what redundancy is needed to achieve high levels of system availability, and what level of redundancy provides the most effective balance of reliable operation, server performance, and overall system cost.

3. Develop server hardware that can detect and report on failures and impending failures.

All IBM Power Systems employ a First Failure Data Capture (FFDC) architecture. This methodology uses hardware-based fault detectors to extensively instrument internal system components. Each detector is a diagnostic probe capable of reporting fault details to a dedicated Service Processor. FFDC, when coupled with automated firmware analysis, is used to quickly and accurately determine the root cause of a fault, regardless of phase of system operation and without the need to run “recreate” diagnostics. The overriding imperative is to identify *which component* caused a fault — *on the first occurrence of the fault* — and to prevent any reoccurrence of



- The FFDC methodology is also used to predictively vary-off (deallocate) components for future scheduled repair. The system will continue to operate, perhaps in a degraded mode, avoiding potentially expensive unscheduled server outages
- In those rare cases where a fault causes a partition or system outage, FFDC information can be used upon restart to deconfigure (remove from operation) a failing component, allowing the system or

the error.

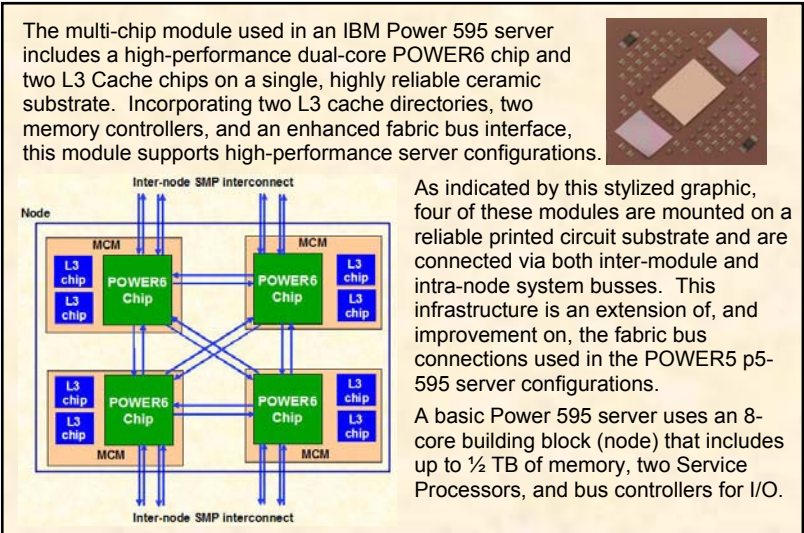
This technology is used to create server hardware that is self-healing, that automatically initiates actions to effect error correction, repair, or component replacement.

- The goal of self-healing/sparing is to avoid faults by employing sparing where it can effectively prevent a scheduled or unscheduled outage. Self-healing concepts can be used to fix faults without having to remove or replace a part.
- IBM’s FFDC methodology is used to accurately capture intermittent errors — allowing a Service Processor to diagnose potentially faulty components. These unique design characteristics allow POWER6 processor-based servers to *capture and isolate* potential processor failures when they occur. Then, using saved system state information, a POWER6 processor-based server<sup>1</sup> can use Processor Instruction Retry and Alternate Processor Recovery mechanisms to transparently (to applications) recover from errors on the original processor core or on an available spare processor core.

partition to continue operation, perhaps in a degraded mode, while waiting for a scheduled repair.

**Reliability Design Processes at IBM**

In the “big picture” view, servers with fewer components and fewer interconnects have fewer chances to fail. Seemingly simple design choices such as integrating a POWER6 processor chip and two L3 caches chip on a single module in a Power 595 server can deliver excellent performance and reduce the “opportunity” for server failure. In this case, an 8-core processor book is built with four highly reliable modules, allowing improved modularity and concurrent book repair and upgrade. Chip to cache interfaces are contained on module — improving reliability and performance. Not only does



<sup>1</sup> Processor Instruction Retry and Alternate Processor Recovery are available on all Power Servers although Alternate Processor Recovery is not available on the BladeCenter® JS12 and JS22.

this reduce the total number of replacement parts, it enables a system with an improved cooling design — delivering better performance per watt than previous high-end POWER processor-based models.

### *Setting Targets*

During the system definition phase of the server design process, well before any detailed logic design is initiated, the IBM RAS team thoroughly evaluates system reliability attributes and calculates a server “reliability target.” This target is primarily established by a careful analysis of the potentially attainable reliability (based on available components), and by comparison with current IBM server field reliability statistics.

Over the past decade, IBM RAS engineers have been systematically adding mainframe-inspired RAS technologies to POWER products, resulting in dramatically improved system designs.

### *Evaluating Components, Packaging Accordingly*

From the smallest to the largest server, system packaging is designed to deliver both high performance and high reliability. In each case, IBM engineers perform an extensive “bottoms-up” reliability analysis using part-level failure rate calculations for every part in the server. These calculations assist the system designers in selecting a package that best supports a reliable design.

The reliability of electronic components is directly related to their thermal environment – large decreases in component reliability are directly correlated with relatively small increases in temperature. POWER6 processor-based systems are therefore carefully packaged to insure adequate cooling. Critical system components such as the POWER6 processor chips are positioned on printed circuit cards so that they receive “upstream” or “fresh” air, while less sensitive or lower power components like memory DIMMs are positioned “downstream”. In addition, Power Systems are built with redundant, variable speed fans that can automatically increase their output to compensate for increased heat in the central electronics complex

(CEC).

### *Redundancy and Concurrency in All the Right Places*

Ongoing detailed RAS analyses help the design team pinpoint server features and design improvements that will have a significant impact upon overall server availability. This enables IBM engineers to differentiate between “high opportunity” items that most affect server availability and that need to be protected with redundancy and fixed via concurrent repair, versus “low opportunity” components which seldom fail or have low impact on system operation and that can be deconfigured and scheduled for deferred, planned repair.

During system development, components that have the highest failure rate and/or potential impact upon availability are identified and the system is designed to contain their impact to overall server RAS. For example, most POWER6 processor-based systems will include redundant, “hot-plug” fans and provisions for N+1 power supplies. Many components in the Central Electronic Complex (the CEC includes processor, memory, and core I/O) are built using IBM *grade 1* or *grade 5* components, parts that are designed and tested to be up to 10 times more reliable than their “industry standard” counterparts.

POWER6 processor-based systems include measures that compensate for, or correct, errors received from components comprised of less extensively tested parts. For example, industry grade PCI adapters are protected by industry-first IBM PCI bus enhanced error recovery (for dynamic recovery of PCI bus errors) and, in most cases, support “hot-plug” replacement if necessary.

### *Continuous Field Monitoring*

Of course, setting failure rate reliability targets for component performance will help create a reliable server design. However, simply setting targets is not sufficient. For this reason, IBM dedicates considerable resources to field tracking, to ensure that hardware components are performing as predicted and to fine-tune existing predictive models.

IBM field engineering teams track and record repairs of system components covered under warranty or maintenance agreement. Failure rate information is gathered and analyzed for each part by IBM commodity managers, who track replacement rates. Should a component not be achieving its reliability targets, an IBM commodity manager will create an action plan and take appropriate corrective measures to remedy the situation. Aided by IBM's First Failure Data Capture (FFDC) methodology and the associated error reporting strategy,

commodity managers build an accurate profile of the types of field failures that occur and initiate programs to enable corrective actions. In many cases, these corrections can be initiated without waiting for parts to be returned for failure analysis.


The IBM field support team also continually analyzes critical system faults, assessing system firmware and maintenance procedures and tools to see if they are effectively handling and recording faults. This continuous field monitoring and improvement structure allows IBM engineers to ascertain, with some degree of certainty, how systems are performing in client environments. IBM engineers then undertake "in-flight" corrections to improve current products being deployed, and use this

IBM's POWER6 chip was designed to save energy and cooling costs.

Innovations include:

- A dramatic improvement in the way instructions are executed inside the chip. Performance was increased by keeping static the number of pipeline stages but making each stage faster, removing unnecessary work and doing more in parallel. As a result, execution time is cut in half or energy consumption is reduced.
- Separating circuits that can't support low voltage operation onto their own power supply "rails," dramatically reducing power for the rest of the chip.
- Voltage/frequency "slewing," enabling the chip to lower electricity consumption by up to 50 percent, with minimal performance impact.
- Innovative and pioneering techniques allow the POWER6 chip to turn off its processor clocks when there's no useful work to be done, then turn them on when needed, reducing both system power consumption and cooling requirements.
- Power saving is also realized when the memory is not fully utilized, as power to parts of the memory not being utilized is dynamically turned off and then turned back on when needed.

When coupled with other RAS improvements, these features can deliver a significant improvement in overall system availability.



information when designing future server products.

## IBM Power Systems – Designed for Reliability

POWER6 processor-based systems continue a long heritage of a balanced design for reliability. IBM's RAS organization carefully models systems for years before a product is sold, enabling designers to make educated choices when deciding which components require additional measures such as redundancy and concurrent repair. By tracking the evolution and performance of components after a product is released, IBM engineers can continue to fine-tune system configurations, models, and provide input for the design of future systems.



© IBM Corporation 2009  
IBM Corporation  
Systems and Technology Group  
Route 100  
Somers, New York 10589

Produced in the United States of America  
March 2009  
All Rights Reserved

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features, and services available in your area.

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

IBM, the IBM logo, POWER, POWER5, POWER6, System p are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

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

The IBM home page on the Internet can be found at <http://www.ibm.com>.

The IBM Power Systems page can be found at <http://www.ibm.com/systems/power/>.

The AIX home page on the Internet can be found at <http://www.ibm.com/servers/aix>.

A more detailed white paper "[IBM Power Platform Reliability, Availability, and Serviceability \(RAS\) - Highly Available IBM Power Systems Servers for Business-Critical Applications](#)" can be found on the Internet by selecting "white papers" at <http://www-03.ibm.com/systems/power/resources/>.

POW03019-USEN-00