IBM POWER6 Processor-based Systems: Designed for Reliability

IBM System p Platform

Reliability, Availability and Serviceability (RAS)

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

June 11, 2007

1 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 POWER6™ processor-based 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.

2 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 two processor cores on a single POWER chip can dramatically reduce the “opportunity” for server failure. In this case, a 32-core server will include half as many processor chips (and chip socket interfaces) as with a single CPU per processor design. Not only does this reduce the total number of system components, it reduces the total amount of heat generated in the design, resulting in an additional reduction in required power and cooling components.

2.1 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 reliability statistics.

In general, RAS targets are set to exceed the reliability of currently available servers. Over the past decade, IBM RAS engineers have been systematically adding mainframe-inspired RAS technologies to POWER products, resulting in dramatically improved system designs.

2.2 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.

Maintaining full binary compatibility with IBM's POWER5™ processor, the POWER6 chip offers a number of improvements including enhanced simultaneous multithreading, allowing simultaneous, priority-based dispatch from two threads (up to seven instructions) on the same CPU core at the same time (for increased performance), enhanced virtualization features, and improved data movement (reduced cache latencies and faster memory access). Each POWER6 core includes support for a set of 162 vector-processing instructions. These floating-point and integer SIMD (Single Instruction, Multiple Data) instructions allow parallel execution of many operations and can be useful in numeric intensive high performance computing operations for simulations, modeling, or numeric analysis.

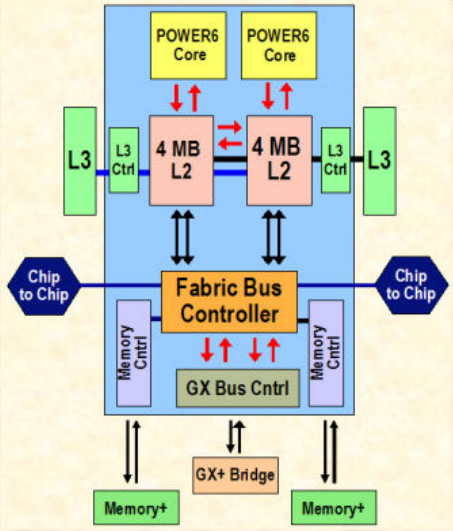

Restructuring the server inter-processor "fabric" bus, the System p™ 570 supports additional interconnection paths between 4-core processor building blocks, allowing "point-to-point" connect between every building block. Fabric busses are protected with ECC, enabling the system to correct many data transmission errors. This system topology supports greater system bandwidths and new "ease of repair" options.

Since 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 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, POWER6 processor-based systems are built with redundant, variable speed fans that can automatically increase their output to compensate for increased heat in the central electronic complex.

2.3 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.

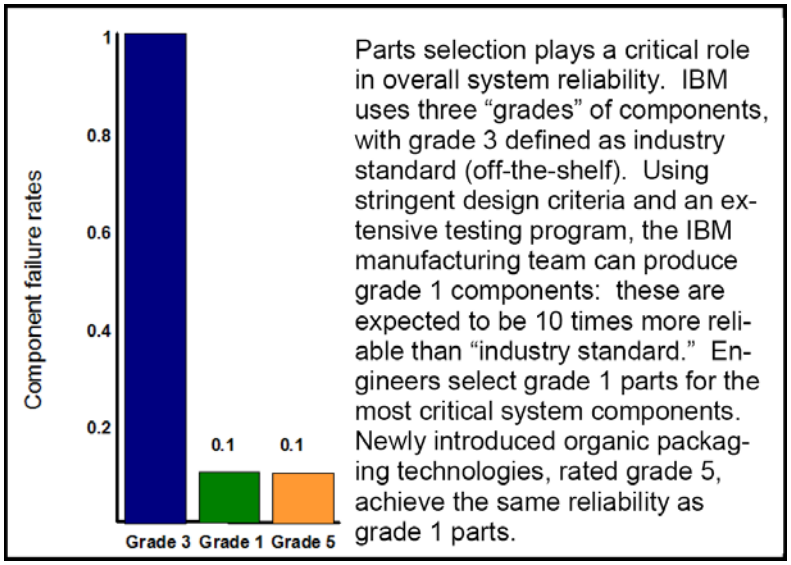

POWER6 processor-based systems include measures that compensate for, or correct, errors received from components comprised of less extensively tested parts. For example, industry grade PCI adapters are protected by industry-first IBM PCI bus enhanced error recovery (for dynamic recovery of PCI bus errors)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

2.4 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, testing to determine if system firmware and maintenance procedures and tools 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, rather than depending only upon projections. IBM engineers then undertake "in-flight" corrections to improve current products being deployed, and use this information

when designing future server products.

3 IBM POWER6 Processor-based 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
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
June 2007

PSW03018-USEN-00