

**EMBEDDED
DIAGNOSTICS FOR
HIGHLY AVAILABLE
SYSTEMS**

A WHITEPAPER

BY LARRY OSBORN

ALAN SGUIGNA



By Larry Osborn

Larry Osborn, Product Manager Processor-Controlled Test, at ASSET InterTech, has over 25 years of experience in product management, hardware/software product design and development, product delivery to the marketplace and user support. Over the years, Larry has a proven track record for identifying user needs and opportunities in the marketplace, providing innovative solutions and exceeding the expectations of users. At ASSET, Larry is responsible for the profit and loss for a product group. Prior to ASSET, he has held positions with Lockheed Martin, OCD Systems, Windriver, Hewlett-Packard, Ford Aerospace and Intel® Corporation. He holds a Bachelor's Degree in Computer Science from the University of Kansas and various technical and marketing training certifications.



Alan Sguigna – Vice President of Sales & Customer Service

Alan has more than 20 years of experience in senior-level general management, marketing, engineering, sales, manufacturing, finance and customer service positions. Before joining ASSET, he worked in the telecom industry. He has had profit and loss responsibility for a \$150 million division of Spirent Communications, a supplier of test products and services. Prior to his tenure with Spirent, Mr. Sguigna also served in business development positions with Nortel Networks, overseeing the growth of its voice over Internet protocol (VoIP) products.

Table of Contents

Executive Summary	4
Introduction.....	5
Case Studies	8
Telecom	8
Nortel Networks.....	9
Computing	9
IBM.....	10
HP	11
Oracle / Sun Microsystems	12
Semiconductor	13
Intel® Itanium.....	13
PLX Technology PCIe Devices	14
The Economics of High Availability	15
The Future of High Availability	16
Conclusion	17
References.....	18

Table of Figures

Figure 1: The Relationship between quality, reliability and availability.....	5
Figure 2: Nortel Networks DMS-100 Central Control	9
Figure 3: IBM First Failure Data Capture (FFDC).....	11
Figure 4: Cost of Server Outage Downtime	15
Figure 5: Diagnostic and test re-use throughout a product’s lifecycle.	17

© 2010 ASSET InterTech, Inc.

ASSET and ScanWorks are registered trademarks and the ScanWorks logo is a trademark of ASSET InterTech, Inc. All other trade and service marks are the properties of their respective owners.

Executive Summary

At a time when society is becoming more dependent upon electronic devices and consumers expect electronic systems to always be available, these same systems have become increasingly complex. Electronic systems operate many mission-critical applications including public safety systems like 911 emergency calling, police and fire dispatching systems, and aviation control systems. Others are simply mission-critical to individuals or organizations like automotive braking systems, credit card authorization systems and online financial trading systems. Of course, electronics provides essential functionality for medical devices and equipment upon which human life often depends, including heart pacemakers, highly precise diagnostic systems and others. And then there are the electronic systems and devices that consumers formerly looked upon as conveniences but are now considered essential to daily life, such as cellphones, MP3 music players, personal computers, digital television and even gaming consoles. For many of the users of these systems, it is imperative that they are available 24 hours a day, seven days a week and 365 days a year.

When system manufacturers strive to ensure high availability, they first consider the reliability of the system and how this might be improved as much as possible. Of course, this demands that the manufacturer balance what is technically possible with market-driven cost constraints. Finding the root cause of system failures when they occur is essential in this regard. This whitepaper describes the practical and financial benefits of improving system reliability and explores some tools that help manufacturers achieve their goal of highly reliable, high-availability systems.

Introduction

Availability expectations for electronic systems vary with the mission-critical nature of their application. For example, in a commercial airliner, continual availability of certain real-time control systems must be assured, but it is merely inconvenient when a passenger's seat radio fails. In the public switched telephone network (PSTN), dropped cellular calls or the absence of dial tone on a landline are critical to service providers' business results, but not as critical as failing to complete a 911 emergency call that might be a life-or-death situation. In computing environments, a temporary email server outage is merely annoying, while the failure of a stock trading system can cost the brokerage firm and its customers millions of dollars.

System quality, reliability and availability are inextricably linked:



Figure 1: The Relationship between quality, reliability and availability

Quality is a subjective term which relates to a user's perception that a product or service meets or exceeds expectations. The accepted metric for the financial impact of quality on a manufacturer is the Cost of Poor Quality (COPQ), which was developed by IBM quality expert H. James Harrington in his 1987 book *Poor Quality Costs*ⁱ. COPQ defines the costs attributable to poor quality. In a perfect world with perfect products and services, these costs would not exist. COPQ can be tangible and appear on a company's financial statements (warranty returns, for example) or indirect and more difficult to measure (such as disloyal or non-returning customers). COPQ is an essential management tool for various organizational initiatives, including Total Quality Managementⁱⁱ (TQM), Six-Sigmaⁱⁱⁱ, and others.

Reliability is a far more concrete term than COPQ. Reliability engineering is the discipline that deals with the ability of a system to perform its required function for a specific interval of time^{iv}. Often expressed as a mathematical measure such as Mean Time Between Failure (MTBF), reliability is an indirect function of the quality of a system. It also relates to system availability.

In general, component or system reliability combined with Mean Time to Repair (MTTR) make up availability. Reliability is not an accident; it must be designed into systems. This process must be integral to the overall reliability program of an electronics manufacturer. Reliability is accomplished via a threefold process: (1) failure and error prevention functionality is built into the system, (2) redundancy for any single point-of-failure is present in the system, and (3) detailed root cause failure information is recorded when a system fails or is otherwise inoperable. It is a given that all systems will eventually fail; as a result, proactive failure and error diagnostics must be part of a reliability engineering program.

Availability can be readily measured. It also has a direct bearing on the user’s perception of quality for a particular system. Availability is the inverse of system outage downtime and is often expressed in ‘9s’, as in:

Availability	Downtime
90% (1-nine)	36.5 days/year
99% (2-nines)	3.65 days/year
99.9% (3-nines)	8.76 hours/year
99.99% (4-nines)	52 minutes/year
99.999% (5-nines)	5 minutes/year
99.9999% (6-nines)	31 seconds/year

Table 1: System Availability and Downtime

A system that has ‘5-nines’ availability implies carrier-grade service; that is, the system is reliable enough to be deployed in the PSTN with its 911 calling requirements. Mission-critical systems such as brokerage services, aerospace electronics and critical medical systems demand ‘6-nines’ availability or greater. Pacemakers, for example, cannot fail to operate for more than a few seconds in a year, if ever!

Analyzing the root causes of defects can affect quality, reliability and availability, while giving system designers the ability to enhance their products or services by avoiding failures in the future. It is imperative that the underlying source of a software or hardware defect is identified when a system becomes unavailable. Even if the system has returned to full operation, this root cause information is necessary so that corrective action can be taken to improve the design of the system or one of its support processes, and to mitigate the defect so it does not occur again. This

is an iterative process throughout a product's lifecycle. Manufacturing tests cannot identify all possible failure mechanisms. Conditions that trigger failures will occur in the field but these conditions cannot always be duplicated in the manufacturer's test lab.

Components invariably fail in the field. Sometimes a component failure will affect system availability and sometimes it will not. In addition, the source or root cause of the defect may or may not be immediately obvious. In either case, a field service technician or possibly the user will attempt to identify the source of the failure down to the level of a field replaceable unit (FRU). This FRU can then be swapped out of the system and replaced with another that is fully operational. The ability of the technician or user to isolate the failure in a particular FRU is dictated by the granularity of the diagnostics built into the system. That is, the system must be able to isolate faults.

Unfortunately, software failures present a particular challenge for fault isolation. Software problems may in fact be caused by hardware faults. It is common for software to 'lock up' when hardware fails. For example, many consumers are all too familiar with the Windows 'blue screen' whenever a hard fault on a PC occurs. This typically takes place when hardware fails and system software cannot recover.

Assuming that the system's diagnostics isolates a fault to an FRU, the FRU is usually sent back to the manufacturer for test and repair. This is where the common No Trouble Found (NTF) problem may manifest itself. That is, the FRU is re-tested in the manufacturer's lab and it appears fully operational. At this point, it is unknown whether the unit is actually at fault, the system operator mistakenly returned it or whether the lab cannot fully duplicate the conditions experienced in the field at the time of the failure. As a result, the FRU is classified as NTF. Industry practices vary, but often NTF FRUs are refurbished and sent back out as warranty replacements to other customers. And if system diagnostics cannot detect and record the true cause of the problem, the FRU may cause problems for one or more users. Some manufacturers track the number of times an FRU may be returned to the factory and classified as NTF. An FRU may eventually be scrapped after it has been returned for repairs a certain number of times.

NTFs are often caused by:

- The service technician cannot effectively diagnose a problem, so he quickly opts to return one or more FRUs to the lab. In many cases, multiple FRUs are replaced in a system at the same time, particularly when the technician is under time pressure during a major system outage.
- The inability of the technician to separate software issues from hardware faults.
- Intermittent system problems.

A recent study by Accenture^v found that in the consumer electronics industry returns range from 11 to 20 percent and more than two-thirds of these can be characterized as NTF. Although this data will vary widely by industry and product, there is no question that NTFs and, in general, any system failure for which the root cause cannot be determined represent a huge cost for manufacturers and users alike.

Case Studies

Telecom

Telecom systems have a long history of high reliability/availability requirements. The terms carrier-grade and five-nines (99.999% uptime) are synonymous. The nature of voice calling demands that dial tone must be ever-present on landlines. The absence of dial tone, due to a failure anywhere in the network, would likely result in a complaint call to the service provider. As telephony technology moved, from analog crossbar switches to digital systems, single points of failure diminished and overall system reliability improved. This trend, a number of lawsuits over 911 (emergency communication services) and lower costs per call drove the digitization of voice services in the 1980s and 1990s. And as the Internet gained steam in the 1990s, the same reliability paradigm was applied to data communications as well. Today, as communication technologies have converged on Internet Protocol (IP), much of the voice, data, video, wireline and wireless communication traffic traverses the same PSTN equipment. Single points of failure for any form of communication within the PSTN have become unacceptable to service providers because they can jeopardize emergency calling services and lead to downtime for consumers.

Nortel Networks

When the Nortel Networks Digital Multiplex System DMS-100 was introduced in 1979 it was the first fully-digital voice system designed from scratch with reliability in mind^{vi}. For the first ten years that Nortel built the DMS-100, it used a custom processor (the NT40), a custom operating system (Switch Operating System - SOS) and a custom language (Prototype Type-Enforcing Language - PROTEL). The central control system was a fully-duplicated, matched-sync system with fast hardware failover in the event of a node failure. During normal operations both members of a matched pair ran in software synchronization and a switchover could occur between members with no loss of voice calls.



Figure 2: Nortel Networks DMS-100 Central Control

Various capabilities were added to the DMS-100 family as time passed and competition heated up with the AT&T (later to become Lucent) #5ESS central office voice switch. Both companies, having roughly equal share of the voice switching market in the U.S., highlighted system availability as a competitive advantage. In 1999, reports from the Federal Communications Commission's (FCC) Automatic Reporting Management and Information System (ARMIS) highlighted that Lucent's switch had the best reliability performance by a factor of three or more over its competitors^{vii}. Nortel countered by adding capabilities like DEBUG, which allowed the setting of breakpoints on the DMS-100, and enhancements to the interrupt handler routines when the system would encounter software traps, sanity timeouts and redundant pair mismatches.

Computing

In the high-end computing (HPC) and server market, the quality buzzword phrase is reliability, availability and serviceability (RAS). Originating with IBM which coined this phrase to characterize the robustness of its mainframe computers, RAS is commonly used today to

describe features that keep systems operational for long periods of time by avoiding system outages caused by faults.

- Reliability refers to features that help avoid and detect faults. A reliable system does not continue to deliver corrupted data; instead, it corrects the corruption when possible or else stops and reports the corruption.
- Availability, as previously described, takes into consideration both system downtime and partial system outages. Partial outages are reported when the system remains operational but a fault or faults have occurred. Highly available systems can disable the malfunctioning hardware or software and continue operating, possibly at a reduced capacity.
- Serviceability represents rapid diagnosis of the root causes of failures when problems occur. Early detection of faults can decrease or avoid system downtime, either by the system taking corrective action in real time or by the system providing recommendations for ameliorative actions that will avoid a recurrence of the problem.

It is sometimes mistakenly asserted that the proliferation of virtualization and cloud computing has improved the resiliency of the IT environment such that RAS has become less critical. In reality, nothing could be further from the truth. Even when virtualization or cloud computing has provided for automatic failover, the root cause of the initial failure may simply be propagated to a secondary virtual machine, creating the same threat of downtime and data corruption. In addition, virtualized and cloud computing systems are often not designed to deal with intermittent hardware errors. In fact, the likelihood of downtime increases with the number of virtual machines on a single physical server.

IBM

IBM has a long heritage of RAS and systems engineered for reliability. Its Power Systems platform, including the IBM eServer™ p5, pSeries® and iSeries™ families, have been designed for increased availability and to support new levels of virtualization^{viii}. All IBM Power Systems employ a First Failure Data Capture (FFDC) architecture. The FFDC methodology employs hardware-based fault detectors (diagnostic probes) which act as instruments embedded in the system to report fault details to a service processor.

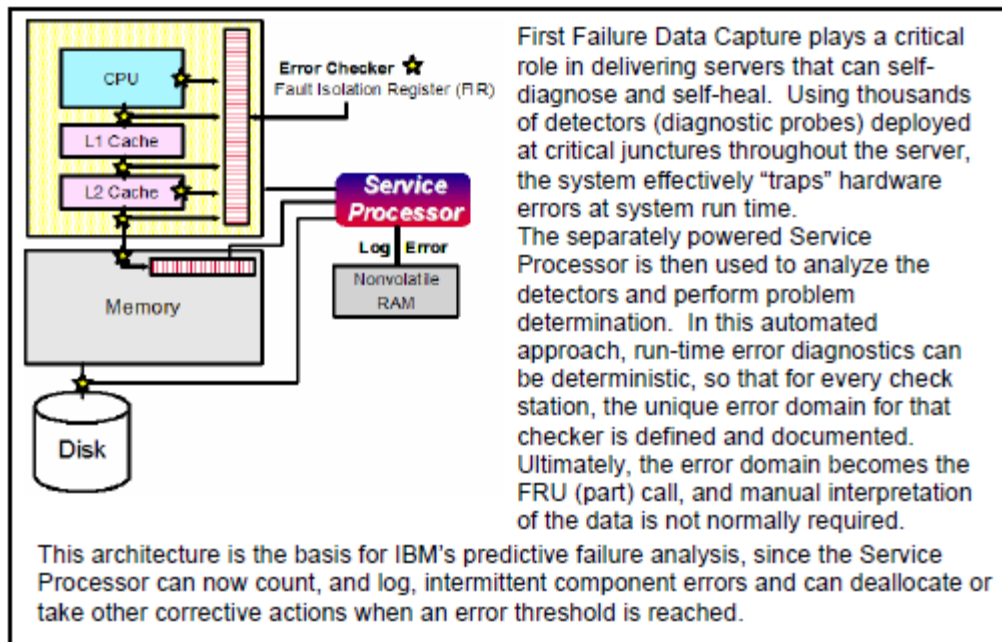


Figure 3: IBM First Failure Data Capture (FFDC)

Every server in the POWER6 and POWER5 processor-based product families includes advanced availability features like Dynamic Processor Deallocation, PCI bus error recovery, 'chip-kill' memory, memory bit-steering, L3 cache line delete, dynamic firmware update, and redundant hot-plug cooling fans N+1 power, and power cords (optional in some configurations). POWER6 processor-based servers add dynamic recovery features like Processor Instruction Retry, L2 cache line delete and L3 hardware-assisted memory scrubbing. Many of these functions rely on IBM FFDC technology, which allows the server to efficiently capture, diagnose and respond to hardware errors the first time that they occur.

HP

HP's Integrity and NonStop systems feature high levels of reliability and availability. Enhanced RAS capabilities on these platforms include Dynamic Processor Resiliency, double-chip spare, cache-safe technology and enhanced PCIe I/O error recovery. These platforms provide an Offline Diagnostics Engine (ODE) environment for field technicians. The ODE troubleshoots the platform which is running without an operating system. A partial list of the diagnostics and tools^{ix} provided under ODE are:

CIODIAG2	Core I/O diagnostic tool
CPUDIAG	CPU diagnostic tool
DFDUTIL	Firmware update utilities for SCSI disks
FCFUPDATE	Fibre Channel firmware update tool
IODIAG	Diagnostic tool for HP Tachyon Fibre Channel board
MAPPER	Utility for mapping physical layout of the SPU and its peripherals
MEMDIAG	Memory diagnostic tool
PCIUTIL	PCI HBA firmware update utility
PERFVER	Utility to run supported SCSI/IDE self-tests on devices
PLUTODIAG	Diagnostic tool for Core Electronics Complex (CEC)

HP NonStop systems power approximately 75 of the 100 largest fund transfer networks around the world^x.

Oracle / Sun Microsystems

Oracle's RAS approach uses an iterative lifecycle approach, which encompasses system architecture, system design, PCB fabrication and assembly, system bring-up, and field data and telemetry^{xi}. In particular, Oracle/Sun has pioneered remote telemetry as a key advantage for troubleshooting NTF issues^{xii}. Using multivariate state estimation technique (MSET), which the company pioneered for work in the nuclear power industry (where breakdowns that result in NTF diagnoses are simply not an option), Oracle remotely monitors temperatures, voltages, currents and a variety of other performance metrics.

One often-cited^{xiii} application of MSET is the detection of cooling air flow perturbations in large enterprise servers. For example, a scrap of paper may partially block an air inlet. The higher temperatures that result from this are not sufficient to trip the system threshold alarms, but they may contribute to accelerated reliability issues over the longer term. Without MSET, the systems might detect the higher temperatures, but would be unable to diagnose that the blocked airflow was the root cause.

Semiconductor

Semiconductors in many ways are analogous to printed circuit boards. An increasing number of package options, such as MCM (multi-chip module), SoC (system-on-chip) and SiP (system-in-package), present opportunities to incorporate processor, memory, DSP, SerDes, and other functionalities in one device. These devices begin to look like systems themselves; as such, they are subject to the same quality, reliability and availability constraints that platforms, which might consist of multiple circuit boards as well as these complex chips, are.

Most semiconductors contain built-in embedded instruments which can be accessed to perform a number of functions that relate to reliability and availability. These embedded instruments can be used to validate the chip itself or the entire systems. Such instruments are often referred to as BIST (Built-In Self Test). Examples of BIST are memory BIST (MBIST), logic BIST (LBIST), I/O BIST, power/temperature monitors and others.

Intel® Itanium

Itanium processors are at the heart of many mission-critical high-end server systems that support continuous operations such as brokerage and stock-trading services. The HP Integrity line of servers has incorporated many of the RAS features in the Intel Itanium processor to provide non-stop service. At the processor level, some of these are documented in the Itanium Error Handling Guide^{xiv}. A convenient comparison between Intel’s Itanium and Xeon chips has been compiled by Hoffman Labs^{xv}, a third-party support provider:

Feature	Platform Alternatives	Description
Cache ECC coverage	Xeon and other x86	Processor cache error detection and correction.
Single-bit Memory Error Correction	Xeon and other x86	Detects and corrects single-bit memory errors.
Double-bit Memory Error Detection and Retry	Xeon	Detects double-bit errors and retries a memory read.
Memory ECC on the data bus	Xeon MP	ECC protection on bus traffic.
Internal Logic Soft Error Checking	claimed unique	Memory parity for large arrays.
Bad or Poisoned Data Containment	claimed unique	Host flags bad memory so the processor does not re-use it.
Cache Reliability, “Intel Cache Safe” or “Pellston”	Xeon MP	ECC that detects and distinguishes between hard and soft cache errors during initialization and run-time. Subsequently disables the cache line with the hard error.
Memory Sparing	Xeon (requires OS assistance and is platform specific)	Makes extra or “spare” memory space available. Found in some Northbridges, including the Intel 5000P and 5000X chipsets.

Memory Mirroring	Xeon (requires OS assistance and is platform specific)	Configures redundant memory arrays (RAID 1 for memory). Found in some Northbridges, including the Intel 5000P and 5000X chipsets.
Hot Plug I/O	All platforms (requires OS assistance and is platform specific)	Inserts or removes PCI-X or PCIx controller while the system is powered and booted.
Memory Hot Swap	Xeon MP (requires OS assistance and is platform specific)	Inserts and removes memory while the system is powered and booted.
Processor Lock Step	Xeon MP on specific platforms (requires OS assistance and is platform specific)	Cross-checks processor output per clock cycle.

Table 2: Intel’s Itanium vs. Xeon Error Handling

Of course, RAS features and capabilities are much more extensive than simply the set of features supported by the processors. The platform, operating system, system software and application software all contribute to the system’s RAS. One example of this are the RAS features supported by Intel’s Interconnect Built-In Self Test (IBIST) embedded instrumentation capabilities^{xvi}, which are present on Itanium and upper-end Xeon platforms. This instrumentation, which is embedded in Intel’s chips, validates the high-speed I/O signal integrity of printed circuit board designs through physical-layer bit error rate and margin testing. In the factory and in the field, IBIST provides the ability to detect manufacturing variances such as solder voids, micro-cracks, incompletely plated vias, annular ring breakout, flaws introduced during the imaging process (i.e. pinholes, nicks, cuts), plating thickness variations and delamination. ASSET’s ScanWorks® Platform for Embedded Instruments is the only tool platform in the industry that supports Intel IBIST embedded instrumentation.

Intel has recognized that high availability is a function not only of the quality and reliability of its own silicon, but also of its ability to interoperate with and verify the operations of other components in the system.

PLX Technology PCIe Devices

PLX has more than 65 percent of the PCI Express (PCIe) switch chip market and has embedded BIST capabilities into its third generation PCIe devices. One such capability is its I/O BIST solution called *visionPAK*™. This software-based tool which takes advantage of an embedded instrument within the silicon accesses internal data paths and state machines for debugging purposes^{xvii}. It can measure signal eye-width inside the chip at the receiver, inject errors into a circuit board to check system behaviors, apply loopbacks to transmitters to debug data paths, and monitor PCIe packet activity and performance.

These capabilities provide a value proposition similar to RAS and the Intel IBIST embedded instrumentation which is being placed in the Itanium, Xeon 5500 and future processor families.

The Economics of High Availability

The cost of system downtime is industry- and situation-dependent. Several studies have attempted to quantify this cost. For datacenters, Hennessy and Patterson^{xviii} have attempted to quantify lost business and productivity that result from system downtime. (Figure 4)




Figure 4: Cost of Server Outage Downtime

The \$6.45 million-per-hour cost for downtime of financial services/brokerage servers is often cited in industry reports to highlight COPQ. Since this study was compiled in 2002, its costs figures are likely conservative.

In addition to the effects downtime has on IT environments, system outages in the telecommunications industry also receive intense scrutiny. It is estimated that the difference between ‘four-nines’ and ‘five-nines’ availability can mean a savings of \$4 million per year for large service providers.

Stories of system outages seem to appear almost daily. Two outages for Telecom Corporation, New Zealand’s largest carrier, in December 2009 and January 2010 shut down service for over 200,000 customers for several days. The service provider blamed Alcatel-Lucent for its troubles.

Research In Motion has had several extended system failures, most recently in December 2009. And, of course, Toyota is now under intense scrutiny for failures in its braking, acceleration and power steering subsystems.

Reliability engineering attempts to mathematically model the effects on availability that enhancements to a product's design will yield. A combination of these models and business models for a given application can project cost savings attributable to enhanced diagnostics.

The Future of High Availability

At the time of writing this paper (March/April 2010), it is estimated that Toyota will spend in excess of \$2 billion for recalls of its cars due to poor performance in its brake, acceleration, power steering and other vehicular systems. At a fundamental level, reliability issues likely have caused these systems to become unavailable for brief periods of time, leading to intermittent or momentary failures and, in some cases, resulting in loss of life. The root cause of some of these issues is as yet unknown, but speculation includes software bugs, hardware wear-and-tear, electro-magnet interference (EMI) and others. It is possible that additional embedded diagnostics capabilities in some form could have provided an early-warning for Toyota and allowed the company to take proactive steps to eliminate faults. The total impact of this, in terms of COPQ, lost productivity and human lives has yet to be determined.

The accelerating complexity of electronic systems and the increased extent to which chips, circuit boards and systems are networked will result in a greater emphasis on reliability in the future. For example, a brokerage services server could go down because of an intermittent logic failure in a single chip in the system, which arose from high temperatures, and which the system software could not recover from. Over time it will become apparent that more embedded instruments in silicon and systems should be made available to diagnose these failures and take preemptive corrective action. In addition, standards such as IEEE P1687 IJTAG and others will be critical to the widespread adoption of embedded instrumentation technologies. The re-use of diagnostic and test technologies throughout a chip/board/system lifecycle will lead to optimized solutions for system availability:

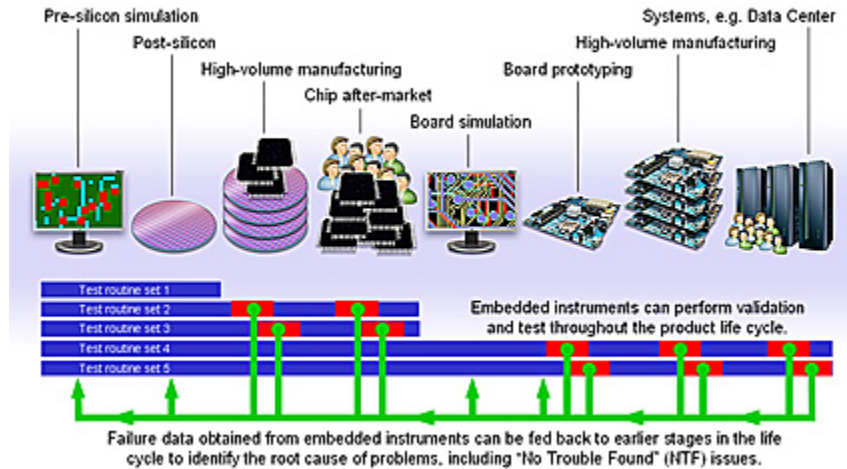


Figure 5: Diagnostic and test re-use throughout a product’s lifecycle.

Conclusion

The electronics industry is entering an era of renewed emphasis on high availability. ‘Five-nines’ (99.999%) availability has been a staple of the telecom industry for many years, but performance above and beyond this level is needed for mission-critical applications such as critical medical systems, lifeline communications, brokerage and public safety systems within the medical, financial, aerospace, automotive and other industries. In particular, server, storage and telecom equipment will be subject to increasing demands for higher levels of reliability as a way to help ensure availability, especially in virtualized and cloud computing environments. Systems are certainly becoming more intelligent and interconnected. To achieve high-availability, they must have embedded instrumentation as well. This instrumentation and a focus on embedded diagnostics will ensure competitive advantage.

References

- ⁱ Harrington, *Poor-Quality Cost*, American Society for Quality, ISBN 9780824777432, OCLC 14965331, 1987.
- ⁱⁱ Peratec Limited, *Total Quality Management*, Springer, ISBN 0412586401, September 2009.
- ⁱⁱⁱ Pyzdek and Keller, *The Six Sigma Handbook*, McGraw Hill, September 2009.
- ^{iv} Pham, *Springer Series in Reliability Engineering*, ISSN 1614-7839, www.springer.com/series/6917, 2006.
- ^v Accenture, *Big Trouble with “No Trouble Found” Returns*, 2008.
- ^{vi} Nortel Networks, *DMS-100 Family System Description*, 1988.
- ^{vii} Nortel Networks, *DMS-100 Family System Description*, 1988.
- ^{viii} Mitchell et al, *IBM Power Platform Reliability, Availability, and Serviceability (RAS)*, <http://www-03.ibm.com/systems/power/hardware/whitepapers/ras.html>, June 2009.
- ^{ix} HP, *Offline Diagnostics Environment Administrator’s and User’s Guide*, <http://docs.hp.com/en/5992-6605/5992-6605.pdf>, April 2009.
- ^x HP, *Offline Diagnostics Environment Administrator’s and User’s Guide*, <http://docs.hp.com/en/5992-6605/5992-6605.pdf>, April 2009.
- ^{xi} Freda et al, Love, *IPC/iNEMI Reliability Summit*, http://thor.inemi.org/webdownload/newsroom/Presentations/APEX_2007/IPC-iNEMI_Reliability_Summit/Sun_Freda_Sauter.pdf, February 2007.
- ^{xii} Riske, *Contrarian Minds*, <http://research.sun.com/minds/2004-1007/>, 2004.
- ^{xiii} Gross, *Continuous System Telemetry Harness, Achieving True Autonomic Computing, Self Healing Qualities, and Proactive Fault Avoidance Through Telemetry*, http://research.sun.com/sunlabsday/docs.2004/talks/1.03_Gross.pdf, 2004.
- ^{xiv} Intel, *Intel Itanium Processor Family Error Handling Guide*, <http://download.intel.com/design/Itanium/Downloads/24927802.pdf>, April 2004.
- ^{xv} HoffmanLabs, *Itanium Reliability, Availability and Serviceability (RAS) Features*, <http://labs.hoffmanlabs.com/node/95>, February 2007.
- ^{xvi} Nejedlo and Khanna, *Intel IBIST: The Full Vision Realized*, International Test Conference, October 2009.
- ^{xvii} PLX Technology, *PEX 8648 – 48-lane, 12-port PCIe Gen 2 Switch*, http://www.plxtech.com/files/pdf/apps/ExpApp68_PEX8648.pdf, July 2009.
- ^{xviii} Hennessy and Patterson, *Computer Architecture*, Morgan and Kaufmann, 2002.