



Tools & Services
for
Safety-Critical
&
High-Reliability
Embedded Systems

Tools from
Requirements
Capture
to
Product Life Cycle
Management

Services including
Consultancy
Compiler Validation
Contractor Register

0808 1800 358

info@phaedrus.com

www.phaedrus.com

Test & Measurement: Embedded diagnostics are critical for today's always-on systems

Reg Waller looks at the problems in obtaining high availability for embedded systems

Users of many of today's electronic systems expect them to work the first time and every time. In fact, from life-saving emergency communications to mission-critical financial systems and a slew of embedded applications in between, a case could be made that society has become totally dependent upon high-availability systems. The cost of downtime is staggering and unacceptable in many cases.

System availability depends on system reliability, which in turn relates to product quality. There are several common metrics for measuring reliability and quality. Quality metrics include IBM's cost of poor quality (COPQ). Reliability is a function of a systems' mean time between failure (MTBF) and, by inference, availability is MTBF weighted by a systems' mean time to repair (MTTR). Expressed another way, overall system uptime can be very high when either (a) MTBF is very long or (b) MTTR is very short. The challenge system designers face is anticipating all possible failure scenarios, alone and in combination, while ensuring (a) and/or (b).

In the telecoms industry, the term 'five-nines' is heard often, referring to mission-critical systems that must be available 99.999% of the time (see Table 1). In the final analysis though, performance is also critical. A system can be available, but its performance may be degraded. The source of system outages or performance degradations may be in either software or hardware.

Reliability is accomplished by 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/or
3. Detailed root cause failure information is recorded when a failure occurs and performance degrades or when the system is rendered 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 programme.

Analysing the root causes of failures and the defects that caused them can positively affect the quality, reliability and availability of systems. 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 and to mitigate the effects of defects. Comprehensive design validation, quality assurance and manufacturing tests, although essential to a product's success, cannot possibly anticipate all possible failure mechanisms in the field.

Embedding diagnostics

There are several ways to diagnose defects and failures, and subsequently provide this data for further analysis and improvement to the system. For example, the telecoms industry, where five-nines availability in the network is an imperative, has deployed several methods. Hardware redundancy is an obvious technique, but embedding certain diagnostic capabilities is just as important. Long before the days of the internet, voice switching systems had diagnostic capabilities such as breakpoint-setting, enhanced interrupt handlers for software traps, onboard failure logging and others embedded into them. Recently, other industries have followed telecoms' example.

Similar to five-nines in telecoms, the current phrase for high-availability in the high-performance computing and server marketplace is reliability, availability and serviceability or RAS. Coined by IBM to characterise the robustness of its mainframe computers, RAS today describes a set of features for diagnosing failures and poor performance while ensuring system availability. RAS is defined as:

- 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 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.

IBM has followed the RAS philosophy and has embedded diagnostic capabilities into its Power Systems series of servers to let the server efficiently capture, diagnose and respond to hardware errors the first time that they occur. HP's Integrity and Nonstop systems also feature high levels of reliability and availability via embedded diagnostics and an offline diagnostics engine for field engineers. High-availability servers from Oracle and Sun have also deployed a range of embedded diagnostics functionality based on the multivariate state estimation technique.

Delving a level deeper, the chips themselves that go into high-availability systems are being equipped with diagnostic capabilities based on embedded instruments. As complex devices such as system-on-chip, system-in-package and 3D multi-die packages look more like a printed circuit board every day, chip designers have realised that their devices are subject to the same quality, reliability and availability constraints that govern circuit boards and systems.

As a result, many of today's complex semiconductors contain built-in embedded instruments that 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), IO bist, power and temperature monitors, and others.

One example of chip-level RAS and embedded diagnostic features is Intel's IBist interconnect bist, which is that company's embedded instrumentation technology. IBist is present on Itanium and upper-end Xeon platforms for the validation of high-speed IO signal integrity on printed circuit board designs through physical-layer bit error rate and margin testing.

PLX Technologies is another example of a chip company that is embedding diagnostic capabilities. Its Visionpak is an IO bist engine that accesses internal data paths and state machines for debugging purposes. It can measure signal eye-width inside the chip at the receiver, inject errors into a circuit board to check system behaviours, apply loopbacks to transmitters to debug data paths, and monitor packet activity and performance on PCI Express.

No trouble found

In the real world of installed high-availability system, the need for embedded diagnostics is accentuated by the no-trouble-found (NTF) problem.

Assuming that the system's diagnostics or an offline diagnostic tool isolates a fault to a field replaceable unit (FRU), the FRU, which is often a circuit board, is usually sent back to the manufacturer for test and repair. This is where the common NTF problem may be manifested. That is, the FRU is tested in the manufacturer's lab and it appears to be fully operational. As a result, it is classified as NTF.

Industry practices vary, but often NTF FRUs are refurbished and sent back into the field as warranty replacements for other customers, where it could cause problems for another user. Some manufacturers track the number of times an FRU may be returned to the factory, classified as NTF and shipped out again to another user. An FRU may eventually be scrapped after it has been returned for repairs a certain number of times, but by this time, several users may have become dissatisfied with the system's reliability. With embedded diagnostics in such a system, the system itself would be able to detect and record the true cause of the problem, eliminating the NTF designation and providing data to improve the product design for greater reliability and quality.

High-availability future

The accelerating complexity of systems and the increased extent to which chips, circuit boards and systems are networked will result in a greater emphasis on reliability and availability in the future. For example, a stock brokerage firm's server could go down because of an intermittent logic failure in a single chip that has overheated. The system's software would not be unable to recover from this hardware condition.

Many manufacturers are coming to the realisation that embedded instruments in silicon and systems can be employed to diagnose failures such as this one and take pre-emptive corrective action. In addition, standards such as IEEE 1149.7 enhanced boundary scan and IEEE P1687 internal JTag, which standardises the interface to embedded instruments, and other such specifications will be critical to the widespread adoption of embedded instrumentation technologies. Ultimately, the significant economic benefits of automatically feeding back failure data generated by embedded instruments at every phase in a product's life cycle will be enlightening to system designers and lead to optimised methods for system availability (see Fig. 2).

Five-nines availability has been a staple of the telecoms industry for many years, but performance at or above this level is needed for many embedded mission-critical applications. In particular, server, storage and telecoms equipment will be subject to increasing demands for higher levels of reliability as a way to help ensure availability, especially in virtualised and cloud computing environments. Systems are certainly becoming more intelligent and interconnected. To achieve high-availability, they must have embedded instrumentation for self-diagnostic purposes. This instrumentation and a focus on embedded diagnostics will ensure competitive advantage.

Reg Waller is European director for Asset Intertech

28 September 2010, [Asset Intertech](#)