Transient Error Protection

The Smarter Approach to Uptime™

The industry's highest measured uptime

Stratus Technologies' uncompromising commitment to uptime is visible every day. We are the first and only server vendor to report the dependability of our installed base of systems worldwide. The Stratus Uptime Meter[SM] is refreshed daily from actual field data and displayed on our Web site. The results report that Stratus systems surpass five nines of uptime.

Since day one, Stratus' engineering and development efforts have been dedicated to making its hardware and software solutions the most reliable and available in the industry.

Every Stratus server comes equipped with Continuous Processing® features that are the result of more than two decades of availability achievements. Through its *Smarter Approach to Uptime,* Stratus addresses the many different factors that negatively impact solutions availability. One of the least understood of these is the transient error and the risk it poses to maintaining solutions uptime and data integrity.

Two classes of hardware errors

Computer hardware crashes can generally be attributed to two classes of errors: hard errors and transient errors. While both hard and transient errors usually result in downtime for a standard server and initiate a failover recovery procedure in a cluster — the similarities end there.

Hard errors are usually reproducible, consistent and easy to isolate. In contrast, transient errors are unpredictable random events that are virtually impossible to isolate on a conventional server.

Compounding the problem of transient errors is that they can cause silent data corruption that results in the system generating false outputs. The consequences can be severe. Irretrievable loss of critical data, costly solutions downtime, and failure to meet regulatory compliance may all be at stake when silent data corruption goes unchecked.

99.9998%

THE **SMARTER** APPROACH TO **UPTIME** ™

Hardware- and software-related incidents, including the Microsoft Windows operating system, are part of the measurement.

What causes transient errors?

Transient hardware errors can occur as a result of many different factors:

Technology factors: The huge performance gains exhibited by today's systems can be attributed directly to the increasing complexity of integrated circuits. But, the very characteristics that enable such gains also increase the likelihood of transient errors. An Intel research and development paper states, "Circuit susceptibility to transient error mechanisms is increasing with each process generation. Some are increasing at an exponential rate."[1]

Additional factors include the increase in transistor densities, lower operating voltages, and increasing clock frequencies. The Intel paper also cites a paper given at the International Conference on Dependable Systems and Networks in 2002. "It is predicted that by the year 2011, the error rates in combinational logic will reach the levels at which we had to have protection in memory."[2]

Environmental conditions: Natural or man-made radiation and changes in temperature, altitude, and humidity can all cause transient errors.

Design issues: The typical industry-standard server is designed with price/performance as its primary goal. When availability is viewed as a secondary objective, minimal design margins are all too often the outcome. Such systems are prone to transient errors when subjected to just the right combination of system load, component manufacturing process, and environmental conditions.

Manufacturing issues: Process control and monitoring issues within the manufacturing process can result in marginal components that lead to transient errors. In their quest to drive down production costs even further, vendors may opt for shorter test cycles or rely on extensive testing of samples — practices that also result in the release of components that cause transient errors.

Over time, these factors can cause affected components to move from a fully functioning state to an intermittent state and, finally, to a hard failed state. Depending on the defect, the component may be in an intermittent state for a relatively long period of time during which transient errors may occur more frequently.
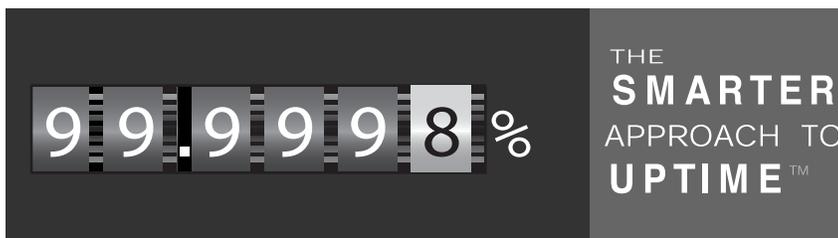
[1]"Firmware-based Platform Reliability", Intel Corporation 2004.
[2]P. Shivakumar, M.Kistler. "Modeling the Effect of Technology Trends on the Soft Error Rate of Combinational Logic."

Addressing the Growing Threat of Transient Errors.

While no computer system can prevent a transient error from occurring, Stratus' line of fault-tolerant systems has been uniquely engineered to detect, isolate, and withstand transient hardware errors.

Engineered for Error Handling

Stratus servers are designed using replicated core system components, including motherboards, processors, memory, I/O buses, and I/O adapters. At the heart of the system design is the use of multiple CPU-memory units operating in synchronized operation. Lockstep processing ensures that any errors, including transient errors, are detected and that the system can survive any CPU-memory unit error without interrupting processing and without loss of data or state information.

While no computer system can prevent a transient error from occurring, Stratus line of fault-tolerant systems has been uniquely engineered to detect, isolate, and withstand transient hardware errors.

In addition to riding through the error condition, Stratus systems capture and log information about the transient occurrence and will automatically take the affected component out of service if it reaches a threshold beyond which it has been determined that a failure is likely to occur. In the event a component is taken out of service, its partner component simply continues to operate as normal.

Extended Software Protection

The Stratus approach to availability extends to the system software as well. Because software is particularly vulnerable to hardware errors, proper error handling can avert many potential software problems. With conventional servers, many problems attributed to software are actually caused by transient hardware errors. Stratus failsafe software reliably distinguishes software issues from hardware issues — greatly contributing to effective and timely problem resolution.

Stratus fault-tolerant hardware and driver hardening technology shield the operating system, middleware, and application software from the impact of transient errors. As part of the driver hardening process, Stratus conducts extensive transient, as well as hardware, error injection testing. This means potential problems are identified and resolved before the system is ever installed at a customer site.

The Continuous Processing Advantage

Stratus' approach to availability is based on a design philosophy that detects, isolates, and corrects errors before they cause system downtime or corruption of valuable business data. Preventing downtime is a key design point that differentiates Stratus servers from conventional servers and high-availability clusters.

The result is uninterrupted uptime that has been proven to meet or exceed 99.999%. Your applications benefit from the design innovations of Stratus servers from the time you load them on the system; no software modification or special configuration is necessary.

Find out how valuable and simple it is to pair your application solution with fault-tolerant systems from Stratus Technologies.

Stratus®
Technologies

www.stratus.com