The Availability Corner

Let's Measure System Reliability in Centuries

November/December 2004

Dr. Bill Highleyman
Dr. Bruce Holenstein
Paul Holenstein

NonStop server reliability currently leads the computing industry by a factor of ten in most cases, according to studies by several different analysts. Even those that approach NonStop reliability, such as IBM Simplex, miss the mark by a significant amount. Today, the mean time between failures of fully configured NonStop systems from any cause is in the order of five to ten years. Some customers report even higher MTBFs.

But should we stop here? No, says the NonStop Enterprise Division which has a major effort underway for Indestructible Scalable Computing led by Wendy Bartlett. No, say we, in our recently published book "Breaking the Availability Barrier: Survivable Systems for Enterprise Computing." We have in our hands today the technology to increase MTBFs five or ten fold. And when a failure does occur, we would lose just a portion of our computing capacity and deny service to no user. Furthermore, the probability of losing more than a single increment of capacity would be measured in centuries. Most important, this level of availability includes all causes of system failures from hardware failures to software failures, operator errors, environmental faults, and scheduled downtime.

How to achieve such a high level of availability at little additional cost is the subject of this series of Availability Columns. There is so much to say on this topic that this series is already scheduled for three years, and undoubtedly there is enough material to perpetuate this topic for at the least the anticipated achievable uptime of a NonStop system.

This month's column serves as an introduction of that which is yet to come. The continuing columns will build on a series of concepts necessary to achieve availabilities measured in centuries. Some of the topics that we will cover include:

- The difference between high availability systems (UNIX, etc.) and continuous availability systems (NonStop, of course)?

- Various measurements of availability and reliability important to understand where we are at and where we want to go.

- User perceptions of the availability they need in mission critical applications.

- Achieving century reliability by distributing applications in active/active configurations using NonStop's Application Clustering Services (ACS) over ServerNet or by distributing applications over a large geographical range.

- Methods for synchronizing databases which are distributed geographically.

- Total cost of ownership (TCO) and why NonStop servers optimize this important measure.

- Zero downtime migration to eliminate planned downtime.

- Grid computing and NonStop's role in local and wide area grids.

- Optimizing your Recovery Time Objective (RTO – the time to recover from a fault), from hot standbys to synchronized distributed databases.

- Optimizing your Recovery Point Objective (RPO- the amount of data lost due to a fault) by minimizing database synchronization latency.

- The NonStop Indestructible Computing (ISC) Initiative.

- The NonStop Application Cluster Services, a major part of ISC.

- The NonStop Advanced Architecture to be introduced with the new Itanium systems.

- The role of enterprise storage systems (ESS).

- The people in NED who are driving the quest for ever higher availabilities.

- Availability best practices as reported by winners of the annual NonStop Availability Award.

- Case studies of successful high availability systems, from configuration to management.

Much of this material has been published in our recent book, "Breaking the Availability Barrier: Survivable Systems for Enterprise Computing." The second volume of this book is now being written and is scheduled for publication late this year. Some of these topics will be discussed in detail in that book. Until then, if there is some topic on the above list that you would rather hear about sooner rather than later, or a topic which you would like to see added, let Bill Highleyman know at billh@sombers.com.

We have a lot of material to cover on a topic which is not only complex but is of utmost importance to all mission-critical applications. NonStop will continue to lead the pack when it comes to systems that hardly ever fail. So follow us in our on-going columns as we take you along this exciting path.

NonStop spells the difference between systems that are mostly go and those that never stop.