The Availability Corner

Fault Tolerance vs. High Availability

September/October 2005

Dr. Bill Highleyman
Dr. Bruce Holenstein
Paul J. Holenstein

Why is it that we see industry-standard servers advertising five 9s of availability while NonStop servers acknowledge four 9s? Are these high-availability industry-standard servers really ten times more reliable than fault-tolerant NonStop servers? Of course not. To understand this marketing discrepancy, let's take a look at the factors which differentiate fault-tolerant systems from high-availability systems.

To start with, there is no reason to assume that a single NonStop processor is any more or less reliable than an industry-standard processor. In fact, a reasonable assumption is that a processor will be up about 99.5% of the time (that is, it will have almost three 9s availability) whether it be a NonStop processor or an industry-standard processor.

So how do we get four or five 9s out of components that offer less than three 9s of availability? Through redundancy, of course. NonStop servers are inherently redundant and are fault tolerant (FT) in that they can survive any single fault. In the high-availability (HA) world, industry-standard servers are configured in clusters of two or more processors that allow for re-configuration around faults. FT systems tolerate faults; HA clusters re-configure around faults.

If you provide a backup, you double your 9s.[1] Thus, in a two-processor configuration, each of which has an availability of .995, you can be dreaming of five 9s of hardware availability. But dreams they are. True, you will have at least one processor up 99.999% of the time; but that does not mean that your system will be available for that proportion of time. This is because most system outages are not caused by hardware failures!

The causes of outages have been studied by many (Standish Group, IEEE Computer, Grey, among others), and they all come up with amazingly similar breakdowns:

Hardware	10% – 20%
Software	30% – 40 %
People	20% – 40%
Environment	10% – 20%
Planned	20% – 30%

[1] See Highleyman, "*Availability Part 1 – The 9s Game,"* The Connection; September/October, 2002.

These results are for single processor systems. However, we are considering redundant systems which will suffer a hardware failure only if both systems fail. Given a 10-20% chance that a single system will fail due to a hardware failure, an outage due to a dual hardware failure is only 1% to 4%. Thus, we can pretty much ignore hardware failures as a source of failure in redundant systems. (This is a gross understatement for the new NonStop Advanced Architecture, which is reaching toward six or seven 9s for hardware availability.)

So what is left that can be an FT/HA differentiator? Environmental factors (air conditioning, earthquakes, etc.) and people factors (assuming good system management tools) are pretty much independent of the system. Planned downtime is a millstone around everyone's neck, and much is being done about this across all systems. This leaves software as the differentiator.

Software faults are going to happen, no matter what. In a single system, 30-40% of all single-system outages will be caused by software faults. The resultant availability of a redundant system is going to depend on how software faults are handled. Here is the distinction between fault-tolerant systems and high-availability systems. A fault-tolerant system will automatically recover from a software fault almost instantly (typically in seconds) as failed processes switch over to their synchronized backups. The state of incomplete transactions remains in the backup disk process and processing goes on with virtually no delay. On the other hand, a high-availability (HA) cluster will typically require that the applications be restarted on a surviving system and that in-doubt transactions in process be recovered from the transaction log. Furthermore, users must be switched over before the applications are once again available to the users. This can all take several minutes. In addition, an HA switchover must often be managed manually.

If an FT system and an HA cluster have the same fault rate, but the FT system can recover in 3 seconds and the HA cluster takes 5 minutes (300 seconds) to recover from the same fault, then the HA cluster will be down 100 times as long as the FT system and will have an availability which is two 9s less. That glorious five 9s claim becomes three 9s (as reported in several industry studies), at least so far as software faults are concerned.

So the secret to high availability is in the recovery time. This is what the Tandem folks worked so hard on for two decades before becoming the NonStop people. Nobody else has done it. Today, NonStop servers are the only fault-tolerant systems out-of-the-box in the marketplace, and they hold the high ground for availability.