

## Continuous application availability needed to prevent losses

HP's NonStop servers effectively help clients minimise losses when a system failure occurs

By Ron LaPedis, CBCP, CISSP, HP NonStop Enterprise Division

DISASTER TOLERANCE IS BECOMING a "must have" for an increasing number of applications. But disaster tolerance is not an absolute. There are levels of tolerance to system and site failures, just as there are levels of tolerance to any other unwanted stimuli - noise, heat, you name it. With many applications, it is enough to react quickly to a system component or indeed an entire site failure and recover from it. But for others, recovery is not nearly good enough. Recovery, no matter how fast and how complete, implies that some damage has already been done - services disrupted, money lost.

Customers with applications that are unable to tolerate any type of loss - in time, data, transactional integrity, etc. - require continuous application availability. This article presents the approach taken by the HP NonStop server; according to the Gartner Group, the only out of the box ultra high-availability system on the market today.

### Good company

According to a 1999 Research Note from D. H. Brown Associates, Inc., NonStop servers process 66 per cent of credit card transactions, 95 per cent of securities transactions, and 80 per cent of automated teller machine (ATM) transactions. They also participate in 75 per cent of electronic funds transfer (EFT) networks. As the vision is realized of "taking everything to the Internet," there is no question that the flagship NonStop platform will continue to occupy a proud position at the center of HP's enterprise strategy.

### Redundant but not idle

The NonStop Server was designed as a linearly scalable, fully integrated platform made up of redundant hardware, system software, middleware, and database, all based upon open standards such as SQL, Unix, Java, CORBA and so on. Unlike some other

**Avoiding the need for recovery is the fundamental principle behind the design of NonStop systems, and it is the basis for the software approach to disaster protection that HP recommends for critical applications running on these same servers.**

fault-tolerant systems, every piece of a NonStop server is active. When a failure occurs, the load is picked up by the remaining components within seconds with no apparent visibility to the user.

Surveys show that planned downtime for hardware and software maintenance or upgrades is at least



With ATMs an essential part of peoples' lives, any system disruptions or failures cannot be tolerated

twice as prevalent than unplanned downtime and therefore reduces application availability even more than unplanned downtime.

The NonStop server was designed so that most maintenance can be done online, without application downtime. Disks and processors can be added or upgraded, disk cache can be altered, and the SQL database supports alteration of indices, partition split/merge/move, and disk or tape backup, all without application downtime.

A hardware or software fault causes a takeover, not a failover, and application processing continues with an unnoticeable delay and with no loss of data. Unlike other systems, a failure does not force hardware or software reconfiguration, rebuilding of memory data structures, or database lock reinstatement.

But even the most fault-tolerant system cannot keep your application running if an external event makes the system unavailable. A failure or planned upgrade of your communications, power, or physical infrastructure, or a local or regional disaster can take down your application unless you have

distributed it across multiple physical locations. But how complicated is it to implement, and how long does it take for a remote site to resume processing?

**The optimal approach to true business continuity is to create a disaster-tolerant environment that distributes the processing across multiple sites, removing the need for recovery entirely.**

### Spanning site outages

The optimal approach to continuous application availability is to create an application environment that distributes the processing workload across multiple physical sites, removing the need for recovery entirely. When disaster strikes, surviving portions of the environment immediately take over processing for the failed parts, maintaining database consistency and keeping business-critical services online without the need for a lengthy recovery process. When the missing computing power is restored, parts of the application are migrated back to it. Other vendors talk about five-9's server uptime. We talk about five-9's application uptime, even across a total site failure. This is not magic nor the future. This is simply an extension of the architecture of a single fault-tolerant NonStop server.

The important thing is that sharing processing in real time across sites provides a much greater measure of protection against disaster-induced downtime than

merely replicating bits as Enterprise Storage Systems (ESS) do. With continuous processing, any node can take over the work of any other node in a matter of

seconds. There is no need for system reboots, application restarts, or database reconfigurations. Nor do local workloads on the backup node have to be jettisoned.

### Summary

Companies need to look at each business process and each system one at a time to determine what is acceptable in terms of application availability. The cost of downtime must be finely weighted, both in and of itself, and versus the costs of protecting against it. But it's a fact that many applications are becoming increasingly time-critical, and that lost time is equating to ever larger amounts of lost money. Money which comes off of your company's bottom line. Money which can be saved by evaluating an HP NonStop server for your application.

In any event, technology should never be selected before the a thorough risk analysis and business impact analysis are done on your key business processes. And then, whatever technology or

combination of technologies is employed for disaster tolerance, they should fit within a larger continuity planning and process framework for ensuring business survival.

### About the Author

Ron LaPedis has been with Hewlett-Packard's NonStop Enterprise Division for 22 years, and is the senior product manager for platform security and business continuity products. He has been a Certified Business Continuity Professional since 1990 and is a Certified Information Systems Security Professional. He has published several articles and has taught and consulted in these fields around the world. Ron is a visiting scholar at the East China College of Computer Technology in Shanghai, and is also a licensed radio amateur, instructor, and volunteer examiner.



Ron LaPedis