

the *Availability Digest*

www.availabilitydigest.com


Stratus Bets \$50,000 That You Won't Be Down

January 2010

Stratus Technologies (www.stratus.com) has been providing fault-tolerant servers for the last three decades. It claims that its current fault-tolerant servers provide in excess of six 9s availability (up 99.9999% of the time, corresponding to 32 seconds of downtime per year). It backs up this claim with the Stratus Uptime Meter found on its home page. This meter is updated daily and is based on all reported service incidents that impacted production during the past six months.



The \$50K Zero Downtime Guarantee



Stratus ftServer Uptime Meter
as of January 4, 2010

Stratus is now putting its money where its mouth is with its Zero Downtime \$50K Guarantee.¹ If you buy a Stratus ftServer® 6300 running Microsoft Windows Server 2008 before February 26, 2010, and if it fails in production during the first six months of deployment, Stratus will pay you \$50,000. In many cases, this is more than the cost of the system.

Outages covered include those caused by Stratus hardware failures, Stratus system software failures, and - yes - even Microsoft Windows Server 2008 operating system failures. If you are unlucky (or lucky) enough to experience a production outage, you can take your compensation in the form of cash or a product credit good for six months.

The Stratus ftServer

Stratus' ftServer was described in some detail in the September, 2007, *Availability Digest* article entitled "Fault-Tolerant Windows and Linux from Stratus."² We briefly review it here.

ftServers provide plug-and-play fault tolerance for Windows and Red Hat Linux applications. Using Intel Xeon chips in a dual modular redundancy architecture, ftServers bring extremely high availability to the industry standard marketplace at affordable prices. The ftServer supports VMware for virtualization and VMware's vSphere for internal cloud computing. ftServer's applicability ranges from use by small businesses with critical applications to large data



¹ Stratus Technologies bets \$50,000 on continuous uptime for ftServer 6300, press release; December 15, 2009.

² http://www.availabilitydigest.com/public_articles/0209/stratus.pdf

centers requiring virtualization support.

In addition to fault-tolerant processors, Stratus also offers the ftScalable array for fault-tolerant storage. These are RAID arrays that can be scaled to 10.8 terabytes of storage.

The Processors

The high availability achieved by the Stratus ftServer product line is achieved by running all applications on x86 dual processors that are lockstepped at the memory access level. Should there be a disagreement between the processors, one of the processors has probably suffered a fault. If the faulty processor has detected its own fault, that processor is taken out of service. Otherwise, processing is paused as each processor enters a self-test mode; and the processor in error is taken out of service. Processing continues with the remaining good processor. If both processors report that they are operating properly, the error is probably a transient error. One processor is declared the winner, and operation continues.

The faulty processor can be replaced and synchronized with the operational processor while the system continues to run.

The ftServer 6300 uses 2.93 GHz X5570 Intel Quad-Core Xeon processors.

The Operating System

Stratus supports Windows and Red Hat Linux operating systems that are commercially available to anyone. Therefore, the ftServers are application binary interface (ABI) compatible with Windows and Linux applications. Any application that can run under Windows or Red Hat Linux on an industry-standard server can run on an ftServer without modification.

Behind the scenes, Stratus spends a great deal of effort ensuring that these operating systems measure up to its high-availability requirements. Working with its partners, Microsoft and Red Hat, Stratus works diligently to harden the operating systems so that operating system faults are minimized.

Stratus engineers have found that the weak links in the software systems are the device drivers. Therefore, Stratus develops its own hardened device drivers that will stand up to the stresses of production environments; and it incorporates them in its standard configurations.

Stratus Call Home

All Stratus servers monitor themselves for faults. If a problem of any kind is detected, the system will automatically call a Stratus support facility (provided customer permission has been granted to do so); and action is immediately taken to diagnose the problem.

If it is decided that a component needs to be replaced, Stratus will immediately send the component to the customer's site. There are many cases in which the first sign of a problem to the customer is when he receives the replacement part in the mail.

Disaster Tolerance

The ftServer is a single system that will tolerate any single fault. To provide disaster tolerance in the event of a data-center disaster, a remote backup system can be kept synchronized with the primary system via Double-Take asynchronous data replication.³

³ [Replicating Windows and Linux Environments with Double-Take](http://www.availabilitydigest.com/public_articles/0408/doubletake.pdf), *Availability Digest*, August 2009.
http://www.availabilitydigest.com/public_articles/0408/doubletake.pdf

The Fine Print

On the surface, this offer seems pretty impressive. But what about the fine print? Upon reading the terms and conditions, one does not find any gotchas – only clarifications that are appropriate.

Included in these terms are the following definitions and requirements:

- Orders must be accepted by Stratus by 2/26 and scheduled for delivery within 90 days.
- The Microsoft Windows Server 2008 Enterprise Edition Operating System must be purchased, but it can be downgraded to Windows Server 2003 if application compatibility requires this..
- The system must be installed within one year of delivery using Stratus' installation services.
- The Assured Availability Plus Service contract must be activated by the date of installation.
- The system cannot be used for development within the guarantee period.
- The customer must maintain an active modem or Internet connection between the system and Stratus Active Service Network (referred to above as Stratus Call Home).
- The customer must replace customer-replaceable units within one business day of receipt of the unit.
- The determination of the failure cause is determined by Stratus' root cause analysis.
- Failures not covered include:
 - known defects or bugs as published in product release notes.
 - failure of non-fault-tolerant components such as USB ports, keyboards, or monitors.
- The guarantee period begins on the first day of production deployment by the customer.
- The guarantee period ends on the earlier of:
 - six months from the customer's first use of the system in a production environment, or
 - one year from the date of installation, or
 - eighteen months from the shipment of the system.

Protecting the Edge

Stratus has recently announced a new system, Avance, that is tailored for what they refer to as “edge computing” – everything outside of the corporate data center.⁴ Avance runs on a pair of standard x86 servers connected by an Ethernet link. The servers may be up to 0.5 kilometers apart. The servers run as a primary/backup pair, and Avance creates a single-system image for the operators and users.

⁴ [Stratus' Avance Brings Availability to the Edge](http://www.availabilitydigest.com/public_articles/0402/avance.pdf), *Availability Digest*, February, 2009.
http://www.availabilitydigest.com/public_articles/0402/avance.pdf

The Ethernet connection is used for two purposes. One is to exchange health information between the two servers. The other is to keep the backup disk synchronized with the primary disk via synchronous replication.

Using the open source Xen hypervisor, Avance supports both Windows and Linux virtual machines. The backup server seamlessly takes over the functions of a failed virtual machine or the primary processor should it fail. In the worst case, a catastrophic server crash will be recovered in less than two minutes.

Planned downtime can be eliminated by rolling upgrades through the system one node at a time.

Stratus claims that Avance can achieve an availability greater than four nines. It is not covered by an uptime guarantee – yet.

Summary

The impact on system availability of vendor-supplied hardware and software is rapidly diminishing to the point that it is no longer a significant factor. Systems such as Stratus' ft Server, HP NonStop systems, IBM's Parallel Sysplex, and HP's OpenVMS clusters that exhibit availabilities greater than six nines effectively take these systems out of the availability equation. Rather, availability is determined today by application bugs, operator errors, and environmental faults such as power, cooling, and data-center destruction.

Six nines of availability can be extended beyond the data center by active/active systems,⁵ which provide distributed processing by geographically-separated nodes synchronized via data replication. A node failure is recovered by simply rerouting users to surviving nodes.

Continuous availability is no longer a technological problem. It is an exercise in balancing system cost with downtime cost. Stratus' ftServer is an affordable starting point to achieve extreme availabilities. Stratus says so – with its wallet.

⁵ [What is Active/Active?](http://www.availabilitydigest.com/public_articles/0101/what_is_active-active.pdf), *Availability Digest*, October 2006.
http://www.availabilitydigest.com/public_articles/0101/what_is_active-active.pdf