

Stratus Raises Hyper-V Virtualization to Fault Tolerance for the Data Center

Analyst: David Reine

Management Summary

There are some functions in our lives, both work and pleasure, where “almost” doesn’t work, and surely “probably” won’t work either. In our leisure life, we may enjoy skydiving. When you ask the instructor *how reliable is the parachute*, you certainly do not want to hear *works right almost all of the time*. In fact, if there were two parachutes available and one cost a little more but was absolutely guaranteed to work every time, there is no question which one you would choose. The same scenario applies when shopping for a new car. If the salesman tells you that the standard ignition works the first time nine times out of ten, and the electronic ignition works every time, but costs more, you might still consider choosing the standard one. However, if he told you that there are two braking systems and one works nine times out of ten, you probably would opt for the option that works every time. When it comes down to a matter of life and death, the decision is easy. It would be wonderful if the “guaranteed” system also was the most cost efficient one.

This same scenario plays out in many mission-critical data centers every day. A choice has to be made between reliability and cost. In a world where Internet access demands that your web-facing servers be available 24x7x365, you cannot afford to have a system that will “almost” never fail. By definition, high-availability systems will fail; however, they are designed with failover mechanisms to keep the outage as brief as possible. A highly available system may be advertised with four “9s” availability. That means that the data center could experience *almost one hour of downtime annually*. That may not seem like a lot, but if you are deploying a 911 system, one hour could mean the difference between life and death. In that situation, a Fault Tolerant (F/T) system may be required, perhaps one approaching six “9s” availability, or downtime of *only 30 seconds per year*. What the data center with F/T applications needs is a low-cost F/T system with outstanding resource utilization, in order to take advantage of that reliability 100% of the time.

Stratus Technologies, an innovator of industry-standard fault tolerance, has extended the capabilities of its innovative x86-based architecture with the availability of Microsoft’s *Windows Server 2008-R2* and *Hyper-V* to their *ftServer* lineup of fault tolerant servers. This provides virtualization capabilities to the most demanding application environments. Whether you are running very critical applications in the core of your data center or other key workloads at the edge, the Stratus *ftServer* can help you achieve full resource utilization to help you lower the total cost of ownership (TCO) of your fault-tolerant infrastructure, and provide you with the maximum flexibility in your choice of hypervisors to deploy that fault tolerance. To learn more about Stratus Technologies and how its fault tolerant line of servers and storage deploys virtualization, please read on.

IN THIS ISSUE

➤ Fault Tolerant Data Center Needs	2
➤ The Stratus Architecture	2
➤ The Stratus ftServer Family	3
➤ Stratus ftScalable Storage	4
➤ Conclusion	5

Fault Tolerant Data Center Needs

The data center with fault tolerant (F/T) application requirements faces the same business constraints as those data centers without the same high level of reliability. These include, but are not limited to, an ever-increasing requirement for more performance and the unrelenting pressure from management to do more with the same or constricted budgets. The IT staff in the F/T data center needs to deploy systems that do not fail, as opposed to those systems in merely high-availability (H/A) environments that can fail but have complex H/A scripts to fail over to another system. These scripts and the clustering associated with them add significant complexity to the server architecture and place an undue burden on the IT staff, adding to the TCO of the IT infrastructure.

More and more, non-F/T data centers are reducing their TCO by migrating their very-critical applications from complex and expensive *UNIX* environments, with traditionally high-performance levels, to industry-standard x86 platforms running *Windows* or *Linux*, configured with multi-core CPUs with similar performance. Many of these data centers are also consolidating multiple older x86 servers with poor resource utilization (less than 20%) onto fewer multi-socket servers with multi-core x86 processors and virtualization hypervisors, in order to improve utilization and reduce TCO. These memory-rich systems enable the data center to reduce the cost for system administration, operating systems, and application licensing, as well as the ancillary cost for floor space and the cost of energy to run and cool the data center environment. Unfortunately, the perceived instability of x86 platforms has limited the use of these platforms for critical environments.

Clearly, the F/T system needs to have the same potential for benefits as does the H/A system when it comes to driving down the operational costs for the IT infrastructure. Using multi-core open systems microprocessors in an F/T system would enable the consolidation and virtualization of other under-utilized “standard” servers into the F/T environment enabling the data center to eliminate the waste in natural resources, moving to a green environment¹, as well as reducing operating systems and application

costs by migrating from *UNIX* to *Windows* and *Linux*.

Enterprises with very-critical applications are seeking affordable uptime reliability for demanding virtual workloads. They are seeking a way for the *Windows*-based applications most important to their success also to benefit from systems exhibiting fault-tolerant behavior, in order to provide more uptime than they have been experiencing. With more than 30 years of credibility in F/T system development, Stratus has exactly what the CIO seeks. With their *ftServer* systems, Stratus provides an architecture for very critical applications to enjoy continuous operation, while at the same time reducing both the acquisition and operational costs of their infrastructure with Intel x86 processors.

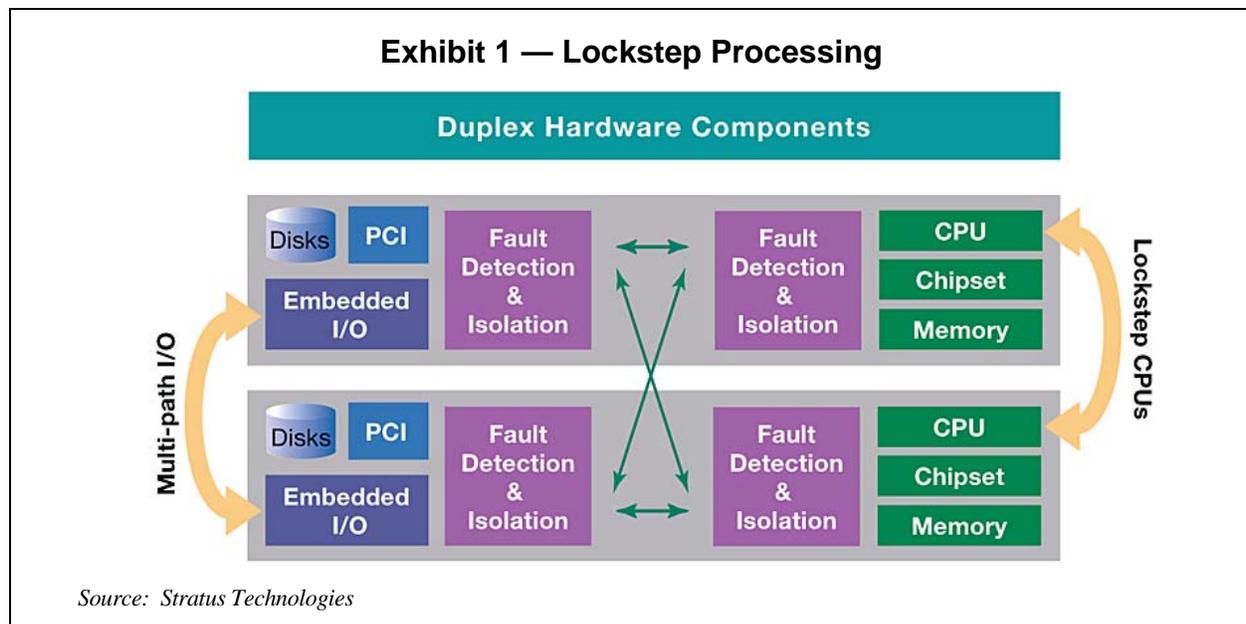
Stratus Architecture

Stratus has taken a unique and innovative approach to delivering almost six “9s” fault tolerance to the data center seeking continuous availability. It has recognized both the need for a low-cost application platform in very-critical environments and the reluctance by many enterprises to deploy an x86 architecture into such an arena because of previous incidents where the hardware failed. It has recognized a perceived reliability issue with *Windows* and x86 commodity platforms and took the steps necessary to deliver the reliability demanded by very critical applications.

The Stratus *ftServer* architecture uses industry-standard technology and components, making the servers cost-effective to manufacture and deploy, and easy to upgrade as enhancements and improved components become available. Stratus uses a combination of its *Lockstep* processor technology (see Exhibit 1, on the next page.), its *Failsafe*² software, and advanced service technology to help prevent downtime from occurring. Each *ftServer* is fully redundant, integrating fully-replicated hardware components into a unique fault-tolerant design, virtually eliminating any single point of failure and safeguarding data integrity. It detects and isolates errors while presenting a single operating software image. The *ftServer* automatically manages its replicated components, executing all processes in lockstep on each image, delivering

¹ See the issue of *Clipper Notes* dated November 24, 2008, entitled *Capping Energy Demand in the Data Center – “It’s Not Easy Being Green!”*, and available at <http://www.clipper.com/research/TCG2008061.pdf>.

² *Failsafe* software works in conjunction with *Lockstep* technology to prevent many software errors from escalating into an outage. *Failsafe* handles such errors transparently, shielding the O/S, middleware, and application software.



zero switchover time. **Because the replicated components perform the same instructions at the same time, there is no interruption in processing even when a component fails. There is no loss of performance. There is no loss of data integrity. Therefore, with the ftServer architecture, there is no recovery time!** This makes the ftServer the ideal solution for essential applications for financial, public safety, and healthcare institutions.

Each ftServer uses a standard dual-modular (DMR) configuration with two processor/memory modules. If one module fails, or is offline for maintenance or diagnostic purposes, the ftServer continues to operate normally, with the remaining components providing five “9s” reliability or better. This eliminates failover scripting, repeated test procedures, and additional coding required to make enterprise applications cluster-aware.

In addition to the capabilities provided by their innovative design, Stratus has now integrated Microsoft’s *Hyper-V*, along with *Windows Server 2008-R2* into their ftServer platform, complementing the availability of a *VMware solution* from Stratus. However, *Hyper-V* brings its own unique advantages to virtualization on the ftServer.

1. Integrates with Windows Server 2008 to enable your organization to take advantage of the benefits of virtualization in Windows Server environments without adopting a new technology;

2. Leverages the existing device driver support in Windows 2008 with a 64-bit micro-kernel hypervisor;
3. Enables the data center to eliminate the *VCenter Server* and *VMware* licenses;
4. Provides support for SMP in the VM;
5. Eliminates compliance and testing issues by using the Microsoft stack;
6. Eliminates the necessity for a second server, FC switches, and an H/A network (as is required with *VMware* in a high-availability scenario); and
7. Bundles the software cost into the Windows license, especially important for the SMB data center where Microsoft Windows servers and applications are dominant.

With the proven technology of Windows in virtualized environments, Stratus seamlessly and transparently has enabled the ftServer with the virtualization capabilities needed to improve the resource utilization of your fault-tolerant data center or remote office.

The Stratus ftServer Family

The ftServer, now in its fifth generation, comes in three basic configurations: The *ftServer 2600*, *4500*, and *6310*, all supporting Stratus’ *Continuous Processing* for increasingly higher workload environments, including server consolidation, with up to four times the performance of previous generations. **All ftServers run standard applications absolutely unchanged.**

They come as 4U rack-mount platforms with eight internal 2.5" SAS disks (per CRU³) in a RAID-10 format and two embedded PCIe slots, along with an embedded SAS controller and two Ethernet interfaces (10/100/1000). They support 73 and 146GB 15K RPM drives or 500GB drives at 7200 RPM. They also support Stratus' *ftScalable* storage with 72 bays and multiple RAID levels, discussed below. In addition, the ftServer 4500 and the ftServer 6310 have two PCIe expansion slots. Each ftServer comes with two power cords, with a recommendation to plug them in to separate power grids.

The industry-standard ftServer has earned Microsoft's F/T Server Certification in the Windows Server catalog and, now, Windows 2008-R2 Enterprise Edition with Hyper-V is available on all ftServers, as is *Red Hat Enterprise Linux 5*. All Stratus servers also support Active Upgrade⁴ to ensure continuous availability and to reduce complexity, even during software upgrades.

Stratus ftServer 2600

The Stratus ftServer 2600 is a one-socket platform based upon Intel's quad-core *Nehalem* processor running at 2.00 GHz and Intel's *QuickPath Architecture*. It is positioned as an entry-level F/T server and is especially appropriate in a standalone, fixed configuration for the SMB or in a replicated remote office environment. At an entry price of about \$15K, it provides a highly affordable solution for critical applications. The 2600 has 4MB of shared cache and from 4GB up to 16GB of system memory with a memory bandwidth of 19.2GB/s.

Stratus ftServer 4500

The ftServer 4500 is a one- or two-socket platform based upon Intel's quad-core *Nehalem* processor running at 2.00 GHz and the *QuickPath Architecture*. It is ideally positioned as a scalable, versatile engine for departmental applications that simply cannot fail, and has proven to be Stratus' most popular model. At an entry point of about \$25K, the 4500 provides headroom for growing or unpredictable workloads. The 4500 has 4MB of shared cache and from 4GB up to 48GB of system memory with a memory bandwidth of 38.4GB/s.

Stratus ftServer 6310

The ftServer 6310 provides a high performance aspect to the high end of Stratus' fault-tolerant server family. It is an enterprise-class system optimized for transaction-intensive Tier-1 applications with the capability to deliver maximum performance and reliability for critical computing. It is a processing intensive platform, ideal for database and transaction-intensive applications, as well as consolidation, yet carries a price tag of only about \$45K.

It is a two-socket fault-tolerant system, replacing the *Nehalem*-based ftServer 6300 with 33% more performance, using Intel's hex-core *Westmere* processor running at 2.93 GHz including Intel's *Hyper-Threading* technology, *QuickPath Architecture*, along with a 10Gb Ethernet interface. The 6310 has 12MB of shared cache per processor and up to 96MB of system memory with a maximum memory bandwidth of 64GB/s.

Stratus ftScalable Storage

The ftScalable storage system combines economics with scalability in a densely configured 2U drawer. It is a high-performance, modular array designed for continuous availability with redundant components, integrated automatic controller failover, and hot standby features to ensure maximum data integrity and protection, including *EnviroStor* battery-free cache protection. It can address dedicated, shared, and networked storage environments. It is an open, fault-tolerant storage platform qualified to support both ftServer systems and x86 platforms from other suppliers. It allows the IT staff to configure dynamically and expand the platform as quickly as the needs of your enterprise demand.

Optimized for performance-oriented applications with I/O intensive or high-bandwidth traffic, generation 2 of the ftScalable array supports 24 2.5" SAS drives in a 2U rack-mountable chassis, with Stratus' patented *DupliCache* redundant cache optimization⁵. The data center can configure up to three chassis modules, with a maximum capacity of 72 hot-swappable drives, 10.5TB of 15K SAS drives (73GB or 146GB) or 36TB of 7.2K SAS drives (500GB). In addition

³ Customer Replaceable Unit.

⁴ Enables the IT staff to "break the mirror" between the two replicated systems for continuous operation during maintenance periods and restore F/T behavior when actions are complete.

⁵ Write data sent to the first RAID controller is concurrently mirrored to the second controller in an active-active pair. The memory bus is used just once to write the data as opposed to multiple passes in conventional systems. This increases bandwidth and results in extremely low write-back latencies.

to the drives, all RAID I/O modules, power supplies, and fans are hot swappable. The ftScalable system supports 8Gb Fibre Channel SAN or DAS host connect, with up to 1000 snapshots optional. It supports RAID levels 0, 1, 10, 3, 5, 6, 50.

The ftScalable array is an ideal complement to an ftServer in a financial, public safety, or healthcare data center (or other similar environment), where it can play a significant role in keeping applications running when hardware components fail or serving in a disaster recovery, business continuity or regulatory compliance strategy.

Conclusion

Stratus has taken the risk out of moving to an open systems architecture for the data center with fault-tolerant application requirements. By implementing Continuous Processing in hardware, they have simplified the IT infrastructure, and lowered the TCO, for the data center, as opposed to running a proprietary implementation with high-availability networking and software, adding to the cost and complexity of the environment. By adding Hyper-V to their virtualization options, Stratus provides a seamless and transparent integration of virtualization for F/T Windows applications. By using commodity components, Stratus has reduced the initial acquisition cost as well, taking the sting out of the decision making process.

When push comes to shove, an ftServer enables the IT staff to sleep at night, knowing that their very-critical applications receive the proper level of care. If you are not sleeping well and want to take the risk out of your data center operations, take a close look at Stratus' fault tolerant product family.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.