



# 5 Questions to Ask About Server Virtualization and High Availability

Part 1

Is Now the Right Time to Virtualize My Critical Tier-1 Applications?

---

The ebook with answers:

- ◆ Is now the right time to virtualize my critical tier-1 applications?
- ◆ Which critical production apps can be virtualized safely?
- ◆ Won't virtualization add more complexity?
- ◆ What happens to my virtual machines if the physical server fails?
- ◆ Will virtualization software alone give me enough high availability?



First of a five-part series from Stratus Technologies

## Is Now the Right Time to Virtualize my Critical Tier-1 Applications?

**More and more often, the answer is yes.**

However, it's obvious that production applications will be more demanding than the non-essential applications where your organization began using x86 server virtualization. And keeping those demands front and center makes all the difference to virtualizing production applications effectively and safely, as we will see.

Let's get going with some context: For years, adding new computing hardware was the path of least resistance when companies wanted to run a new application or support more users. Physical servers proliferated as a result. Power, cooling and facilities costs rose to support all that gear. At the same time, improvements in processor technology produced server hardware

that was increasingly powerful. Too much computing capacity went idle as hardware sprawl met Moore's Law. A tipping point arrived, and server virtualization was in the right place at the right time.

Server virtualization gained traction fast as an accepted way to deliver IT services and applications in the enterprise. The early adopters of x86 virtualization focused on consolidating servers that were lightly utilized. The typical workload was neither performance-sensitive, nor critical in nature.

### Production apps go virtual

Now IT executives — perhaps you're among them — are looking at the next wave of applications to virtualize: extending to business-critical tier-1 applications such as e-mail, messaging and database servers, and business services including online transactions and credit card authorization.

Many such applications that weren't viable candidates for using server virtualization before can now be considered as virtualization software matures. For example, VMware's vSphere™ 4 offers numerous performance and scalability improvements for resource-intensive applications compared with its predecessor, VMware® Infrastructure 3.

### Due diligence first

Still, anywhere that companies recognize a need for high availability or better, IT executives are called on to perform due diligence. Before you move critical production applications to a virtualized environment, you need to know your service level agreements (SLAs) will continue to be met or surpassed.

New concerns come up, for instance, when you concentrate multiple production applications on a physical server. Some IT managers liken the effect to putting many eggs in one basket.

## VIRTUALIZATION, HIGH AVAILABILITY AND ROI



“The bottom line is that with a well-architected virtualization solution, a 6 to 18 month ROI is highly likely.”

- *Chris Wolf*  
Senior Analyst / Burton Group

Explains industry analyst Dan Kuznetsky, vice president of research operations at The 451 Group, “When you think about it — the more workloads that you stack on a single machine — when something does go wrong, the more pain an organization is going to feel because all those applications have dropped on the floor.”

### Burton Group’s Chris Wolf on Virtualization 2009

Economic pressures are prompting more organizations to use x86 server virtualization to support production applications, says senior analyst Chris Wolf of Burton Group in a virtualization podcast for Stratus Technologies.

“Organizations, up until now, have just been starting to virtualize their production applications. With the budget constraints IT is under in 2009, they are going to be under even more pressure to virtualize more production applications. As you virtualize production applications, availability of those virtualized applications becomes an even more critical issue,” says Wolf.

“The thing that’s in virtualization’s favor is that from an ROI perspective, it continues to be an absolute no-brainer. It allows you to remove hardware from your data center, so I’m reducing my hardware maintenance costs. I’m reducing my energy cost associated with that hardware as well. There could be some human resources cost reduction in terms of junior-level IT staff that might be doing some of that hardware maintenance.

“The bottom line is that with a well-architected virtualization solution, a 6 to 18 month ROI is highly likely. Many organizations have already made the initial investments for virtualization — at least in the large enterprises — so it’s an incremental investment.”

**Hear the entire two-part podcast:**

<http://community.stratus.com/podcasts/virtualization>

# SERVER VIRTUALIZATION AND HIGH AVAILABILITY

Using x86 virtualization for disaster recovery is another situation in which production applications are highly dependent on uptime and performance. Then there are the classic mission-critical applications where any downtime or data loss has the potential for devastating effects.

Wherever “high availability” matters, you’ll want to dive deeper into what that means with respect to your particular production applications and the virtualization solutions you are considering: Will you be protected from downtime and data loss in the manner you thought? What skills and effort will be required of your IT staff to deploy and manage the solution to make sure it all works?

## More about this ebook

In this e-book, we discuss these and other essential considerations to help you decide when and how to move your critical production applications to a virtualized environment.

Upcoming ebook chapters cover:

- Which Critical Production Apps Can Be Virtualized Safely?
- Won’t Virtualization Add More Complexity?
- What Happens to My Virtual Machines If the Physical Server Fails?
- Will Virtualization Software Alone Give Me Enough High Availability?

## About Stratus Technologies

Stratus Technologies focuses exclusively on helping its customers keep critical business operations online without interruption. Business continuity requires resiliency and superior availability throughout the IT infrastructure, including virtual environments.

Stratus delivers a range of solutions that includes software-based high availability, fault-tolerant servers, availability consulting and assessment, and remote systems management services. Based on nearly 30 years of expertise in product and services

technology for total availability, Stratus is a trusted solutions provider to customers in manufacturing, financial services, health care, public safety, transportation & logistics, and other industries.

**For more information, visit [www.stratus.com](http://www.stratus.com)**



COMING NEXT!

## Chapter 2

- ◆ Which critical apps can be virtualized safely?

*Stay tuned to this five-part series from Stratus Technologies*



# 5 Questions to Ask About Server Virtualization and High Availability

## Part 2

### Which Critical Production Apps Can Be Virtualized Safely?

---

The ebook with answers:

- ◆ Is now the right time to virtualize my critical tier-1 applications?
- ◆ Which critical production apps can be virtualized safely?
- ◆ Won't virtualization add more complexity?
- ◆ What happens to my virtual machines if the physical server fails?
- ◆ Will virtualization software alone give me enough high availability?



 *Second of a five-part series from Stratus Technologies*

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

### Which Critical Production Apps Can Be Virtualized Safely?

As we mentioned in the previous chapter, using x86 server virtualization is now possible with many applications that weren't realistic candidates in the past. Attention to a number of guidelines will help you determine which production applications, particularly critical ones, will run in virtual machines (VMs) with the availability and performance needed by your business processes and your users.

### Know your application

First things first: Does your application software vendor or ISV support the virtualization solution you plan to use? The good news is the number of supported applications has grown quickly and continues to climb. For example, some of the global software vendors with support policies in place for VMware® products are listed on the company's site.

Today production application support extends to an increasing number of mission-critical applications that were not appropriate to run in VMs before. Application support is expanding even in rigorous environments that include computer-aided dispatch (CAD) for public safety, and health information systems used in patient care.

It remains a must to consult with your application vendor(s) to be clear about exactly what is supported regarding the specific virtualization technology you plan to use, and to understand the nuances of running your application in VMs.

Issues around application software licensing come to mind first, but they're just one factor. For instance, in certain cases the application vendor may limit support to approved hardware and software configurations that have been pre-tested.

Your software vendor may in addition be a knowledgeable source of advice about specific application characteristics that have an impact on availability and performance.

When you need to go deeper to understand how your application will work with VMs, engaging a professional services provider with expertise in virtualization and availability can help you avoid mistakes that would cost your business dearly.

Among the examples is when you have a reliable legacy application and want to complete a physical-to-virtual migration without destabilizing the environment.

### Identify resource dependencies

In any case, you need to understand which resources your production application depends on, and when.

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY



*Application support is expanding even in rigorous environments that include message switching in travel and transportation environments, computer-aided dispatch (CAD) for public safety, and health information systems used to support patient care.*

### FAA Virtualizes Critical International Communication Application on Stratus Servers

The Federal Aviation Administration's William J. Hughes Technical Center in Atlantic City is building a virtual computing environment for its operations-critical international message switching system with systems and services from Stratus Technologies.

The project will start with two guest operating systems running on VMware ESX™ software, data center edition. The first guest operating system will run a message switching application for sending flight plans, weather information and other messages to international airports from FAA control centers in Salt Lake City and Atlanta.

The second will run a Windows® application being used as a generic proxy server. Two fault-tolerant Stratus® ftServer® systems and Stratus ftScalable™ storage units will be located in Atlantic City, with another server located in each of the two control centers.

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

Staying clear about resource dependencies can get challenging in a virtual environment, where services and applications are abstracted away from the physical IT infrastructure that supports them. Hitting a bottleneck or a capacity limit will degrade performance and cause your application to become unavailable, provide unacceptable response time or both. You have to be aware of how latency-sensitive the application is, and if it is bound by CPU, memory and/or I/O.

Because peak workload is a key consideration, VM placement matters — especially for resource-intensive production applications. Keep in mind that these applications may benefit from virtualization for purposes of disaster recovery and live migration (see Chapter 3), even when the applications still ought to be hosted on non-consolidated, dedicated physical server(s).

When your application uses specialized hardware services or features, be sure

not to overlook these in your planning. For example, supervisory control and data acquisition (SCADA) systems in manufacturing may depend on legacy or proprietary I/O cards to access specialized plant equipment networks.

### Virtualization software matters

The virtualization software you use makes a difference. Each vendor's server virtualization software has its specific traits; even various editions of the same vendor's software don't necessarily have the same characteristics. Many of those variations have a direct impact on production applications. One straightforward example: VMs that run on “bare metal” — that is, directly on the server — are subject to less latency than those running on top of a host operating system.

You will also want to determine whether it's suitable to run the virtualization software on server hardware you already have. Using the virtualization software's high-availability or fault-tolerant features

in some cases requires a primary and secondary physical server with identical hardware configurations right down to the CPU revision level.

Although high availability and disaster recovery drive many virtualization projects, you also have to determine whether your “high-availability” or “fault-tolerant” virtualization software leaves you exposed to the risks of application downtime and data loss. We'll expand on this point in Chapters 4 and 5.



## COMING NEXT!

### Chapter 3

#### ◆ Won't Virtualization Add More Complexity?

*Stay tuned to this five-part series from Stratus Technologies*



# 5 Questions to Ask About Server Virtualization and High Availability

Part 3

Won't Virtualization Add More Complexity?

---

The ebook with answers:

- ◆ Is now the right time to virtualize my critical tier-1 applications?
- ◆ Which critical production apps can be virtualized safely?
- ◆ Won't virtualization add more complexity?
- ◆ What happens to my virtual machines if the physical server fails?
- ◆ Will virtualization software alone give me enough high availability?



 Third of a five-part series from Stratus Technologies

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

### Won't Virtualization Add More Complexity?

When any game-changing technology changes the rules of the game, nobody ought to be surprised. Virtualization follows the pattern established by earlier game-changers: Some things get easier. Other things get more complex.

Physical server sprawl may be replaced by virtual server sprawl — and fast, because provisioning new VMs is close to effortless compared with getting budget approved to buy more hardware. In addition, applications, networking, storage and backup are all affected by the physical-to-virtual abstraction of x86 server virtualization. IT responsibilities, policies and tools also have to evolve now that there is a physical infrastructure, a virtual infrastructure and myriad interdependencies between them.

### Disaster recovery eased

x86 server virtualization does make it easier to provide aspects of the high availability warranted by critical

production applications. Disaster recovery (DR) and business continuity is one area where server virtualization is making resilient IT easier and more affordable, by enabling shorter recovery times using VMs and relaxing the requirement for identical hardware at the DR site.

The virtualization software uses a shared resource pool of physical servers and storage to provide the base for DR. The process works as follows: When a physical server shuts down, virtualization software can automatically (if supported by the implementation) use a storage area network (SAN) to retrieve image(s) of the affected VM(s), including configuration state, disk state and so on. The replicated VM images can then be restarted at the disaster recovery site. Data can also be replicated to storage elsewhere on the network.

### Maintenance minus the downtime


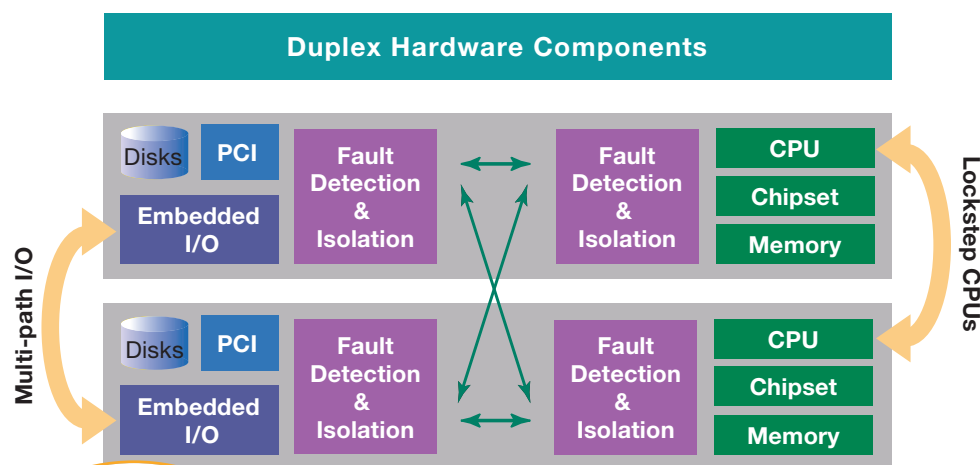
Some server virtualization also offers live migration, which makes accommodating

planned maintenance much easier than in the past. Live migration lets you move VMs from one physical server to another, while users perceive little or no impact on application availability or performance (provided there are enough system resources to handle the load). Planned maintenance can be handled without the usual application downtime.

A word to the wise, though: Hardware compatibility remains a prerequisite for live migration of VMs between physical servers. Otherwise applications are liable to crash due to differences in the instruction sets between the host CPUs. Of course, technology suppliers are aware of this and are making strides to reduce constraints on live migration. As an enterprise-ready virtualization solution, VMware provides compatibility checks in VMotion™ to help prevent attempts at live migration between non-compatible CPUs. In addition, microprocessor leader Intel is enabling enhanced compatibility in newer processors to relax the restrictions.



## Lockstep Hardware Architecture: Simpler Availability for Critical Production Apps



Redundant components within ftServer® systems operate in lockstep — processing the same instructions at precisely the same time. In the event of a malfunction, the partner component simply continues processing without any interruption or degradation in performance. The system's built-in fault-detection features automatically capture, analyze and report any issues. Even in-memory data is fully protected from loss or corruption and, in virtualized environments, full multi-core SMP support is made available to **ALL** VMs.

Stratus® ftServer® systems running the vSphere™ platform provide out-of-the-box continuous availability for virtualized workloads. The ftServer architecture provides component and functional redundancy within the footprint of a single server to keep deployment and support simple. Stratus' hardware-based fault tolerance eliminates the overhead of software emulation along with the I/O limitations and scalability constraints that software lockstepping imposes. Being able to use true symmetric multiprocessing (SMP) allows applications running in virtual machines to scale across multiple CPU cores.

While leveraging SMP performance to its fullest, your compute-intensive production workloads are assured of an industry-leading six nines or greater (99.9999%+) availability across all VMs. What's more, you can use VMware vMotion to migrate workloads between conventional x86 servers and ftServer systems to open up a new availability-on-demand option.

### Software lockstep has drawbacks

Complexity increases when virtualization depends on software lockstep in order to achieve high availability or fault tolerance. Although reliability can be improved by using software lockstep, this technique imposes additional overhead on CPU, network and I/O processing.

Furthermore, software lockstepping depends on primary and secondary physical systems: at least one duplicate physical server, a duplicate copy of any software and the planning to ensure that failover is going to work properly. That amounts to adding servers — not consolidating them. In this sense, software lockstep resembles the server cluster approach that's been a high availability alternative for many years.

Software lockstep also uses replication and a dedicated Ethernet connection to provide a heartbeat that keeps the primary and secondary virtual machines in sync. The replication and heartbeat

### Software lockstep: important considerations

- ◆ Does not provide sophisticated data protection or diagnostic features: faults may propagate to other servers
- ◆ Hardware still fails: blackouts occur during failover, inflight data may be lost
- ◆ May need to over-configure system, network and storage to compensate for latency and limited CPU resource usage
- ◆ Implementation requires two (or more) of everything: adds complexity; impacts purchase, deployment, management of hardware and software infrastructure
- ◆ Runs on a variety of hardware but complete system and CPU compatibility is required to provision and migrate partner VMs from one system to another

can slow down application response time due to latency and processor overhead, particularly for applications that have high transaction volumes. Where does that lead? Complaints from users and a drop in service level performance.

A lesser-known drawback of software lockstep is that it limits a virtual machine to a single processor core. Then VMs can't use the symmetric multiprocessing (SMP) or multicore capacity of a CPU, which restricts the application from scaling up in performance.



## COMING NEXT!

### Chapter 4

- ◆ What Happens to My Virtual Machines if the Physical Server Fails?

*Stay tuned to this five-part series from Stratus Technologies*



# 5 Questions to Ask About Server Virtualization and High Availability

## Part 4

### What Happens to My Virtual Machines If the Physical Server Fails?

---

The ebook with answers:

- ◆ Is now the right time to virtualize my critical tier-1 applications?
- ◆ Which critical production apps can be virtualized safely?
- ◆ Won't virtualization add more complexity?
- ◆ What happens to my virtual machines if the physical server fails?
- ◆ Will virtualization software alone give me enough high availability?



## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

## What Happens to My Virtual Machines If the Physical Server Fails?

For production applications, examining exactly what would occur during failure and recovery is vital. Advises industry analyst Dan Kuznetsky, “Even when people are thinking about availability and reliability, they don’t often think about how long does it take for a function to move from a failed environment to one that’s working.

“Next, when they choose an availability regimen — a set of tools and processes — they often don’t think about how long will it take for that resource to become available again once the failure scenario has started,” he says.

### Failure isn’t consequence free

Because a single server typically plays host to multiple VMs in a virtualized environment, the implications of a physical server failure are anything but trivial in high-performance, high-density and/or mission-critical settings. The

principle is fundamental: Any single physical server is a potential single point of failure for every VM it runs.

Failing to observe the interdependencies between the services and applications in your IT environment is another risky omission. Forrester Research explains, “Applications are now so interdependent that one highly resilient application may utilize data from multiple applications with lower levels of resiliency. In the event of a disruption, highly resilient applications are at the mercy of their less resilient application partners, causing potential critical business processes to fail.”<sup>1</sup>

And although we touched on this point in an earlier chapter, it bears repeating: Virtual machines are almost effortless to provision. That tends to keep x86 virtualized environments in a state of perpetual motion, with the potentially harmful side-effect of introducing single points of failure where none existed before. You have to remain aware of

which critical production VMs are placed on which physical server(s).

### Failover: resiliency with downtime

The failover capabilities provided by virtualization software are a useful initial defense against failure for production applications where some downtime is acceptable. For these more forgiving apps, the failover and restart — which may extend three minutes, four minutes or more — is an improvement over a prolonged outage. It’s necessary to evaluate which users and applications can tolerate the pause while failover/restart takes place.

VMware High Availability (HA) is a capable example of a virtualization solution with failover and restart capabilities that protect applications against unplanned downtime caused by guest operating system as well as x86 server failures. Upon detecting a guest OS failure, VMware HA will automatically restart VMs after a pre-determined interval. In cases where VMware HA

<sup>1</sup> Forrester Research, Inc., X86 Server Virtualization for High Availability And Disaster Recovery, October 24, 2007

## THE IMPLICATIONS OF PHYSICAL SERVER FAILURE



The principle is fundamental:  
Any single physical server is a  
potential single point of failure  
for every VM it runs.

detects an x86 server hardware failure, it automatically restarts the affected VMs on another server in the resource pool. Applications that would otherwise be left without a failover option can benefit from a more timely recovery.

### Consider the worst case

While automatic failover provides resiliency for VMs, tier-1 applications

### More Workloads Become Critical

According to industry veteran and vice president Dan Kuznetsky of The 451 Group, when enterprises plan their virtualization strategies, assessing application availability requirements is often not getting the attention it deserves.

“I’ve come to the conclusion that an awful lot of people aren’t really thinking about availability, security or management when they embark on a virtualization strategy. They never think what happens if something goes wrong,” he says in a podcast recorded for Stratus Technologies.

“They’ve not even gone the next step of, ‘Let’s look through our application portfolio and determine: Is this painful but not critical? Is this critical? Is this a showstopper?’ And I think that if you do look at it, it suddenly becomes clear that some application workloads that were ancillary before have become critical. How many people rely so much on e-mail and their calendar management system that if they fail, no one can do any work?”

“People need to stop, and take time, and think. Because just [putting] things into a virtual machine, or [putting] things into a virtual machine and adding some other layers of virtualization around it, doesn’t solve some of the inherent problems of availability and reliability.”

Hear the complete podcast:

<http://community.stratus.com/podcasts/virtualization>

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

are subject to data loss and additional downtime when a physical server does fail. No matter which virtualization software you choose, best practice calls for assessing worst-case scenarios for your production applications when VMs must restart to recover from an outage.

An unplanned server outage will require restarting VMs, as just discussed. But when a VM reboot is necessary, the information about application state as well as in-flight data — data not yet written to disk — will be lost.

In addition, error(s) are able to propagate to the secondary VMs upon restart. The reason is that virtualization software solutions on the market today do not isolate the cause of a failure, whether in software or hardware. The failure that required the VM reboot could be repeated/replicated along with the same threats to uptime and data integrity.

Software-based high-availability or fault-tolerant solutions are also not designed to deal with transient (temporary) hardware errors including device driver malfunctions, which may spawn downtime and data corruption when left uncorrected.

Your contingency planning ought to include the time and skills that your IT staff will expend on problem resolution. For instance, when application state is lost during a VM restart, determining where the error originated and finding the root cause gets more difficult because valuable diagnostic information disappears as a result of the reboot.

Complementing the resiliency of your virtualization solution with fault-tolerant server hardware — the case in point being VMware® vSphere™ 4 and Stratus® ftServer® systems — offers a means to harden your virtualized environment against unplanned downtime and data loss.

In Chapter 5, we explain how adding full-function Stratus fault-tolerance to your VMware vSphere environment can enable you to meet the continuous availability needs of key tier-1 applications and equip you to provide availability on demand.



COMING NEXT!

### Chapter 5

- ◆ Will Virtualization Software Alone Give Me Enough High Availability?

*Stay tuned to this five-part series from Stratus Technologies*



# 5 Questions to Ask About Server Virtualization and High Availability

Part 5

Will Virtualization Software Alone Give Me Enough High Availability?

---

The ebook with answers:

- ◆ Is now the right time to virtualize my critical tier-1 applications?
- ◆ Which critical production apps can be virtualized safely?
- ◆ Won't virtualization add more complexity?
- ◆ What happens to my virtual machines if the physical server fails?
- ◆ Will virtualization software alone give me enough high availability?



 Last of a five-part series from Stratus Technologies

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

## Will Virtualization Software Alone Give Me Enough High Availability?

The answer to this question for enterprise IT is: probably not everywhere, especially for production apps that depend on higher availability and simplicity of operation. Factor in the following when assessing your availability needs:

- Your business' recovery time objective (RTO)
- Your recovery point objective (RPO)
- The amount of planning, testing and ongoing IT management needed to enable that RTO and RPO

## Full-function fault tolerance

When used with software-based x86 virtualization, fault-tolerant server hardware can be used to make a good thing better by addressing virtualization concerns that relate to availability and performance: including scalability limitations, single points of failure, unplanned downtime, fault isolation, data integrity, problem resolution and more.

The combination of the VMware® vSphere™ cloud operating system and Stratus® ftServer® technology lets you maximize uptime while optimizing your resource utilization. The full-function fault tolerance of the ftServer family offers an effective means to harden virtualized — and cloud — environments without adding complexity.

Eliminating single points of failure is inherent to the ftServer system hardware design, which is described as providing continuous availability. If one of the server's components fails, its duplicate keeps the server processing with zero interruption and no degradation in performance. Failover is eliminated rather than minimized. Because protecting uptime and availability happens on a single physical server rather than requiring multiple servers, hardware cost of ownership and software license fees are less expensive and less complicated.

Automatic fault isolation and diagnostic tools are integral to the server architecture. In addition, the internal redundancy within an ftServer system protects integrity of committed (completed) transactions and preserves in-flight data in the event of a component failure or a transient error.

## Hardened for 24/7/365 availability

Software, including your VMs, need not be modified in any way to benefit from the ftServer system's reliability and availability features. Stratus ftServer systems run the VMware ESX hypervisor and are fully API-compatible with applications supported by VMware ESX software. Leveraging many years of availability expertise, engineering groups at VMware and Stratus have worked together to let off-the-shelf VMware software transparently take advantage of fault-tolerant ftServer features such as hot-plug replacement of PCI cards.

Availability protection continues with Stratus proactive remote service for

## SERVER VIRTUALIZATION AND HIGH AVAILABILITY

## Stratus ftServer Systems:

## Full-Function Fault Tolerance for VMware vSphere HA Environments

Full-Function Fault Tolerance	Performance	Simplicity of Operation	Financial Advantage
<p>Hardware &amp; Software</p> <ul style="list-style-type: none"> <li>• Purpose built</li> <li>• Predictive failure analysis</li> <li>• Industry-standard components</li> <li>• Transient error protection</li> </ul> <p>Service</p> <ul style="list-style-type: none"> <li>• 24/7 support</li> <li>• Extensive diagnostics</li> <li>• Online repair</li> <li>• Single point-of-contact</li> <li>• VM root-cause analysis</li> </ul>	<ul style="list-style-type: none"> <li>• Full virtual SMP for up to 8 cores per guest</li> <li>• Hardware fault-tolerance ideal for latency-sensitive applications</li> <li>• No additional system or network overhead incurred with hardware fault tolerance</li> </ul>	<ul style="list-style-type: none"> <li>• Drop-in, transparent fault tolerance for vSphere™ environments</li> <li>• Fault tolerance for off-the-shelf ESX™ hypervisor — without any modifications</li> <li>• No additional network, storage or management requirements</li> <li>• Single-system, fault-tolerant vCenter™ Server management platform</li> </ul>	<ul style="list-style-type: none"> <li>• Single copy of VMware®, operating system and application software</li> <li>• Single-system fault tolerance</li> <li>• Availability on demand extends fault tolerance to entire cloud</li> <li>• Fault tolerance extended to all VMware Editions</li> <li>• Full utilization of system capacity</li> </ul>


VMware vSphere ESX environments that extends to 24/7 support coverage. We uniquely offer a support service that performs root-cause determination for Red Hat® Enterprise Linux® and Microsoft® Windows® Server 2003 and 2008 virtual machine environments on ftServer systems.

**Not only for mission-critical apps**

Classic mission-critical applications are a natural fit for pairing the VMware vSphere operating system and ftServer technology. As the underlying server hardware used with offerings including VMware HA, Stratus ftServer systems deliver continuous availability in a virtualized infrastructure ready for your tier-1 enterprise applications.

Additional use cases include disaster recovery and application-dense server consolidation as noted in earlier chapters. In a virtualized world, the duo of the VMware vSphere platform and ftServer technology benefits a growing set of cases that wouldn't have ranked

## Stratus Value Proposition



Total Availability



Operational Simplicity



Financial Advantage

as critical before enterprises began running production applications in VMs.

**Management/provisioning platform:** An entry-level ftServer 2600 offers a robust, single-server host platform you can use to support VMware vCenter™ Server and manage the entire VMware vSphere 4 environment. You are able to stop a problem at the host level from interrupting access to management capabilities. An ftServer system can also host golden image(s) of the VM(s) you use to provision your production apps, to continuously safeguard the availability and integrity of the golden image.

**Availability on demand:** Stratus ftServer systems can act as the critical resource in a server pool, providing availability on demand during load balancing, volume spikes, time-constrained processing windows, etc. You are positioned to support tiered levels of service, with the flexibility to bump up VMs to a higher level of availability protection whenever your business needs call for it.

### Virtualize with Confidence

ftService<sup>SM</sup> offerings extend the value of Stratus servers with proactive remote support that sustains the industry's highest levels of uptime for VMware vSphere environments. Our total solutions approach leverages powerful ActiveService™ capabilities built into every ftServer system, linking you to a global service network for detecting, troubleshooting and resolving problems fast:

- Industry-leading 99.9999%+ availability across all VMs
- Downtime and data loss prevented, not just minimized
- Automated 24/7 monitoring, error detection, problem isolation, fault management
- Global incident management system
- Issues resolved online more than 98% of the time
- Server orders own customer-replaceable components
- Engineering response to critical issues in minutes\*
- Single-source accountability includes root-cause problem determination for VMware ESX hypervisor\*

\* with Assured Availability Plus service coverage

Learn more about our Worldwide 24/7 Services: <http://www.stratus.com/services>




TAKE THE NEXT STEP

**Contact Stratus at:**

**[www.stratus.com/services/virtualization.htm](http://www.stratus.com/services/virtualization.htm)**



## 5 Questions to Ask About Server Virtualization and High Availability



### **The Availability Company®. For nearly three decades.**

Stratus, the innovator of industry-standard fault tolerance, now extends VMware virtualization capabilities to the most demanding application environments. From the heart of the data center to key workloads at the edge, Stratus ftServer systems help you achieve full resource utilization and maximum flexibility.

**Learn more:** [www.stratus.com](http://www.stratus.com)



- ◆ Last of a five-part series from Stratus Technologies



THE AVAILABILITY COMPANY®

[www.stratus.com](http://www.stratus.com)

Stratus, ftServer and the Availability Company are registered trademarks; ftScalable, ActiveService and the Stratus Technologies logo are trademarks; and ftService and Assured Availability Plus are service marks of Stratus Technologies Bermuda Ltd. VMware and the VMware logo are registered trademarks and vSphere, VMotion, ESX and vCenter are trademarks of VMware, Inc. in the United States and/or other jurisdictions. Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and/or other countries/regions. Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries/regions. Red Hat and Enterprise Linux are registered trademarks of Red Hat, Inc. in the United States and other countries. The registered trademark Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. All other trademarks and registered trademarks are the property of their respective holders.

© 2010 Stratus Technologies Bermuda Ltd. All Rights reserved.

X1001-A

