

Server Buying Guide

How to choose a server for reliability, availability, and serviceability.



The industry's highest measured uptime

Stratus Technologies' uncompromising commitment to uptime is visible every day. We are the first and only server vendor to report the dependability of our installed base of systems worldwide. The Stratus Uptime MeterSM is refreshed daily from actual field data and displayed on our Web site. The results report that ftServer systems surpass five nines of uptime.

After selecting the right application solution, choosing the right hardware is next on your agenda. With so many choices available, won't almost any server do? The answer to that question requires a close look at your operation and its service-level commitments.

The process of selecting a hardware platform is very similar to the application evaluation process you've just completed. Ultimately, your decision should be influenced by the impact the new solution is expected to have on the success of your business. That impact may be measured in increased revenue contribution, lower operational costs, or improved customer satisfaction — or, in the best-case scenario — all three.

As you consider your options, ask yourself the following questions:

- How important is it for your server and application to operate without interruptions?
- What consequences would an unplanned outage have?
- How critical is transaction integrity? Preventing data loss?

- If downtime occurs or there is a need for troubleshooting or repair, how much recovery time is acceptable until your environment is restored to normal operation?
- Do you need to administer and troubleshoot the system remotely?
- How much time can your experienced IT professionals spend administering and maintaining the system?
- How fast do you expect the system's workload to grow? How soon will you need to upgrade?

Considerations such as these are at the core of how consistently a server will support the service level desired, and how economical it will be to own and operate throughout its lifecycle — known as total cost of ownership, or TCO. These factors largely determine whether the server will offer your organization its best return on investment, or ROI.

The characteristics that make a server easy to live with are often described in terms of reliability, availability, and serviceability (RAS). The concept describes a computer system's dependability over time (reliability), its ability to respond to a failure (availability), and its ease, efficiency and effectiveness of service when problems do occur (serviceability).

Better RAS translates into more uptime for servers and applications, which helps ensure that users always have access to business- or mission-critical data and services. (Because the elements of RAS are so inseparable from the levels of uptime they support, the words reliability, availability, and uptime are commonly used to mean the same thing.)



99.9998 %

~ 1 minute yearly downtime

Source: Stratus Uptime MeterSM

Hardware- and software-related incidents are part of the measurement.



Understand your options.

Certain server designs enable RAS more successfully than others. Nearly all come with tradeoffs in terms of system and application uptime, availability and integrity of data, ease of implementation and daily operation, and the cost to buy and own the solution.

Single server

Some of today's commodity servers offer subsets of RAS features once found mainly on proprietary systems: advanced error-correcting code (ECC) memory, mirrored memory, hot-plug disks, hot-plug adapter cards, RAID storage, redundant fans, and redundant power supplies. The extent of included capabilities varies, applies only to certain aspects of the system, and may be found in only the highest-priced enterprise-class servers. The notion is to reduce the probability of failure — but when a standalone server fails, it fails. Such servers aren't designed specifically to prevent downtime or data loss, and they're not designed to recover quickly from problems that may occur. As a result, their downtime can often be measured in hours or even days.



“The biggest benefits we see from the Stratus ftServer system are transaction throughput and 100% availability — that is, the system responds to every transaction received from the network.”

— Banking customer

“We needed to ensure continuous production. As competition and demand increase, we had to eliminate all risks of failure.”

— Manufacturing customer



Clustered systems

Clustering is a technique in which physical connections and software programs link two or more servers so that when a failure occurs on one, its workload fails over to a server that is still operating. Clusters are typically built by assembling hardware and software from multiple vendors.

These server configurations are described as high availability or fault resilient because they are designed to speed recovery from problems. High-availability clusters are not designed to eliminate downtime altogether.

Users are interrupted during a server outage but may be able to log back on quickly. They may notice slower application performance while the faulty server is out of operation, and may have to rerun some transactions using a journal file. As with a standalone server, in-memory data not yet written to disk will be lost if a server crashes. Whenever changes are made to hardware or software within the cluster, its failover capabilities (which depend on a software script) should be updated and retested to ensure that failover will occur as planned.

Fault-tolerant systems

Fault-tolerant systems are often described as providing continuous availability because they are designed to preempt downtime. Traditionally, this approach has provided comprehensive component and functional redundancy within the footprint of a single server. If one of the server's components fails, the server continues to process normally. The redundancy ensures that users don't experience any disruption, no transactions or data are lost, and there is no degradation in performance.

Today's fault-tolerant systems — specifically, the ftServer® family from Stratus Technologies — automatically enable RAS to the fullest extent possible. Replicated components, system software reliability features, and online service features combine to support uninterrupted uptime that has been proven to meet or exceed 99.999%. Applications benefit from the servers' fault-tolerant capabilities immediately; no modifications to software or special configuration changes are necessary.



The current generation of fault-tolerant servers also eliminates a drawback associated with older, proprietary systems: cost. To achieve a competitive initial purchase price, Stratus® ftServer systems use standard Microsoft® Windows® Red Hat® Enterprise Linux® and VMware® vSphere™ operating systems and off-the-shelf hardware components including Intel® processors. Their single-system design permits lower software license costs compared with clusters, because multiple copies of the same software are not required.

The ftServer series keeps ownership costs low as well. They are maintained like a standalone industry-standard server — but with self-monitoring capabilities, remote troubleshooting, and true hot-swap components all built in. Furthermore, Stratus’ Active Upgrade™ technology enables customers to perform online software upgrades and critical operating system hot fixes while the application continues to run uninterrupted.*

Stratus ftServer systems don’t miss a beat in the face of component failures and system software faults that cause other servers to crash. Their combination of remote-access serviceability and automatic problem reporting allows service engineers to resolve issues online in more than 95% of cases. And when needed, ftServer systems automatically order their own hot-swappable, customer-replaceable parts.

* Windows only

“After-sale support is key. The ‘call-home’ feature and the way Stratus Technologies supports its servers make all the difference.”

— Public safety customer

Award-winning ftServer systems

The built-in continuous availability and operational simplicity of Stratus’ industry-standard servers add up to long-term value. Drawn from the hundreds of IT products evaluated each year by InfoWorld, ftServer systems have been twice honored as “Best Fault-Tolerant Server”.



	2600 Model	4500 Model	6300 Model
Positioning	Highly affordable, entry-level, fault-tolerant server	• Scalable, versatile, expandable price/performance leader for applications that cannot fail	Enterprise-class server delivering maximum performance and availability for mission-critical computing
Deployment model	Replicated remote locations	Departmental business processing	Critical enterprise and business/operations processing
Workload	<ul style="list-style-type: none"> • Stable, standalone • Fixed configurations 	<ul style="list-style-type: none"> • Growing or unpredictable • Multi-tasking applications 	<ul style="list-style-type: none"> • Processing-intensive • Database server • Transaction-intensive • Server consolidation

Take a closer look.

Features	Single system	Clustered systems	Stratus ftServer system
Key uptime features	Duplication of select components; varies by vendor and hardware platform	Software-based failover; duplication of select components	<ul style="list-style-type: none"> • Lockstepped processors; redundant components • Failsafe software technology • Integrated, proactive serviceability
Uptime philosophy	Minimize downtime due to failure	Fast recovery from failure	Failure prevention
Recovery time	Minutes to days	Minutes to hours	Zero
Fault detection and isolation	Limited	Limited	Yes: to the component level
Transaction integrity	<ul style="list-style-type: none"> • Integrity of database transactions dependent on hardware configuration • All in-flight data lost on failure 	<ul style="list-style-type: none"> • Integrity for database transactions only • All in-flight data lost on failure 	Protection of in-flight data ensured by mirrored memory and redundant components
Deployment	No application modification or special programming required	<ul style="list-style-type: none"> • Deployment requires careful planning to eliminate single points of failure and to properly size servers • Failover scripts must be written and manually tested 	No application modification or special programming required
Application modification	None	May need cluster APIs to ensure proper failover	No application modification or special programming required
Systems management	No special skills required; single-system view	More complex: administration of multiple servers and ongoing change management procedures require skilled personnel	No special skills required; single-system view
Dependence on IT staff	None	High: effectiveness of cluster usually dependent on implementation and management skills of personnel	None: set up and administered like any industry-standard server
Service philosophy	System monitoring, reporting, and diagnostic capability varies greatly by vendor and platform; typically does not approach capability of Stratus	System monitoring, reporting, and diagnostic capability varies greatly by vendor and platform; typically does not approach capability of Stratus	Integrated system monitoring with automatic problem reporting. More than 95% of issues are resolved online
100% Availability Guarantee: Service Program	Not offered	Not offered	Standard with premium service*
Availability statistics	Not available	Not available	Published daily
Cost of ownership	Extremely aggressive price nullified when cost of downtime is high	Aggressive acquisition costs eclipsed by cost of downtime, duplicate software licenses, and personnel costs associated with ongoing change management	TCO advantage over clusters when cost of downtime and personnel are included; significantly less expensive than single server when cost of downtime is high
Online upgrade	Not available	Complex to manage	Active Upgrade technology**

* Terms and conditions apply

** Windows only

Specifications and descriptions are summary in nature and subject to change without notice.

Stratus, ftServer, the Stratus logo, and the ftServer logo are registered trademarks; the Stratus Technologies logo, Active Upgrade, and The Smarter Approach to Uptime are trademarks; and Uptime Meter is a service mark of Stratus Technologies Bermuda Ltd.

Microsoft, Windows and Windows Server are registered trademarks of Microsoft Corporation in the United States and other countries. Red Hat, Enterprise Linux, and the Red Hat Shadowman logo are registered trademarks of Red Hat, Inc. in the United States and other countries. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. VMware, the VMware logo and vSphere are trademarks or registered trademarks of VMware, Inc. All other trademarks and registered trademarks are the property of their respective holders.

© 2010 Stratus Technologies Bermuda Ltd. All rights reserved.

X832-F

