December, 2010

Stratus® Continuous Processing® Technology

Automatic 99.999%+ Uptime for Red Hat® Enterprise Linux® Environments

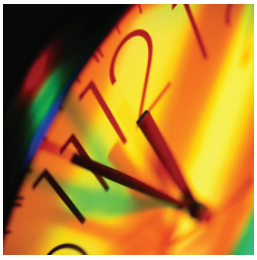

redhat. ADVANCED PARTNER

Stratus® Technologies

Abstract

Stratus Technologies' family of ftServer® systems has been proven to deliver industry-leading uptime of 99.999% and greater for Red Hat® Enterprise Linux® environments. The source of this immediate, transparent availability protection is Stratus® Continuous Processing® technology.

Virtually any software application certified on the Red Hat Enterprise Linux (Release 4) operating environment will run **unchanged and unmodified** on servers that are specifically designed to prevent unplanned downtime. With robustness engineered into the servers' hardware, software, and serviceability, the off-the-shelf fault tolerance that was once exclusive to proprietary systems is becoming a new option in the open source era.

This paper presents the fundamentals of the fault-tolerant technology built into every ftServer system, which enables *The Smarter Approach to Uptime*™. It also explains how the automatic continuous availability protection offered by ftServer systems can yield improved operational simplicity and a significant financial advantage over competing high-availability approaches.

Contents

Fundamentals of Continuous Processing Design

To provide the most complete protection for uptime possible, a comprehensive server solution must address the areas of hardware, software, and service. Only Stratus Technologies, together with its distribution channel partners, and ftServer systems provide this total solution in the Linux® market.

Every Stratus ftServer system includes Continuous Processing features that are the outgrowth of more than 30 years' experience of ensuring uptime for demanding mission-critical and business-critical applications around the world.

Every aspect of this design works in concert to prevent unplanned downtime, not simply minimize it. Preventing downtime is a key design point that differentiates these fault-tolerant servers from "robust" standalone servers and high-availability clusters (which use multiple servers to recover from downtime when one of the servers in the cluster fails). Unlike reliability-enhancing approaches that are not integral to a server's design, built-in continuous availability helps limit exposure to the operator error that industry experts identify as a leading cause of unplanned downtime.

Applications certified on the Red Hat® Enterprise Linux® operating system need not be modified in any way to benefit from these exceptional availability safeguards. This advantage represents a considerable improvement compared with clusters that require failover scripting, ongoing tests (each time the cluster configuration changes), and software modifications to make applications cluster-aware.

Figure 1: Automatic 99.999% Uptime for Red Hat Linux Environments

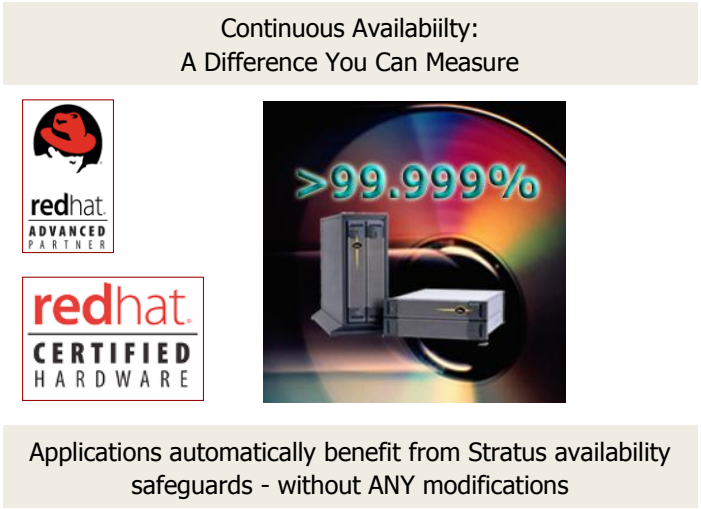

Stratus enables Continuous Processing capabilities in ftServer systems through three fundamental elements:

Figure 2: Core Elements of the Stratus Continuous Processing Design

- **Lockstep Technology —** Lockstep technology uses replicated, fault-tolerant hardware components that process the same instructions at the same time. In the event of a component malfunction, the partner component acts as an active spare that continues normal operation and averts system downtime. The system also detects and corrects transient hardware errors that could cause software failures if left unchecked.

- **Failsafe Software —** Failsafe software works in concert with lockstep technology to prevent many software errors from escalating into outages. Unlike typical servers or clusters, ftServer hardware and software handles most errors transparently, shielding the operating system, middleware, and application software. Another advantage of the Stratus approach is that it constantly protects and maintains in-memory data.

  Management and diagnostic features capture, analyze, and notify Stratus of any software issues. This allows support personnel to take a proactive approach to correcting software problems before they recur. In addition, hardened device drivers add considerable reliability to the Linux environment on Stratus ftServer systems.

- **ActiveService™ Architecture —** An unmatched combination of ActiveService capabilities enables built-in serviceability not offered by other vendors. Stratus ftServer systems constantly monitor their own operation. When a fault is detected, the server correctly isolates the condition and automatically opens a call that tells the Stratus support center exactly what action to take.

  Remote support capabilities, made possible by the global ActiveService Network, have historically enabled Stratus service engineers to troubleshoot and resolve problems online more than 95% of the time. If necessary, the system automatically orders its own hot-swappable replacement part and ensures the correct part is delivered, within 24 hours, to major locations worldwide. Users can install these components easily while the ftServer system continues to run uninterrupted. In addition, a secure Web-based ActiveService Manager allows Stratus and customer-authorized vendors to collaborate on faster problem resolution.

The following pages provide a closer look at each of these aspects of Continuous Processing technology.

Lockstep Technology

The ftServer family eliminates single points of failure using replicated components that continue uninterrupted processing even in the event of a component malfunction. Hardware faults are handled automatically by the system, without the delay of a failover (as on a cluster) and without loss of data.

While many servers offer duplicated power supplies, fans, and disk drives, Stratus fault-tolerant servers provide protection for core system components that include motherboards, processors, memory, I/O buses, and I/O adapters. Another advantage of this approach is that the server presents a single-system view and runs a single copy of all software, which typically reduces software licensing costs and simplifies administration as compared with multi-node cluster alternatives.

Dual Modular Redundancy (DMR)

The Stratus ftServer line implements dual modular redundancy (DMR), which uses two CPU/memory assemblies (system units). These DMR systems are designed for "five nines" (99.999%) availability that is unmatched by standalone or clustered servers.

Using replicated hardware components that operate in lockstep, the servers maintain multiple CPU/memory units in precise synchronization. These system units run in a lockstep manner from a single system clock source. Fault-detection and isolation logic compares I/O output from the system units; any miscompare indicates an error. The fault-detection logic on each system unit determines which board is in error.

The lockstep processing on DMR components ensure that any errors, even transient errors, are detected and that the system can survive any CPU/memory unit error without interrupting processing and without losing data or state.

Figure 3: ftServer Lockstep Technology

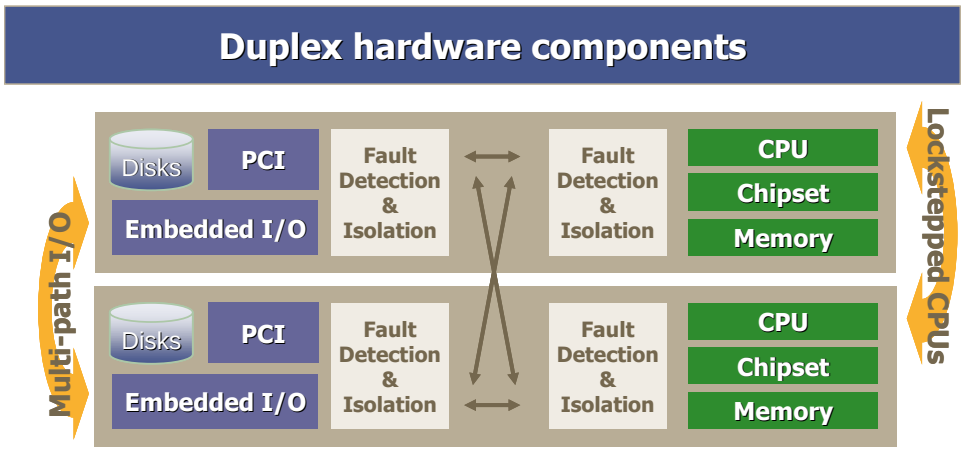

The fault-tolerant I/O subsystem is logically separate from the CPU/memory subsystem. Hardware logic, in the form of custom chipsets, acts as a PCI bridge between the CPU and I/O,

and provides the core error detection, fault isolation, and synchronization logic for the lockstep architecture. Custom logic within the CPU/memory subsystem contains the primary PCI interfaces, interrupt control functions, and transaction ordering logic. Custom logic within the I/O subsystem contains the voting logic, secondary PCI interfaces, and error registers. The custom chipsets use a passive bus, which the ftServer systems implement in the form of a backplane, to connect the replicated CPU and I/O modules within the server.

Fault-tolerant I/O is implemented through the use of replicated PCI buses, replicated I/O adapters, and replicated devices. All critical PCI adapters are duplicated as well: SAS, Ethernet, remote management, and Fibre Channel. Internal SAS disk storage is mirrored (RAID 1) and connected via two independent storage buses.

Industry-Standard, Modular Components

Leveraging off-the-shelf technology in a modular physical design captures a cost advantage while also reducing development time and improving time to market for new fault-tolerant server models.

The biggest differences from conventional servers are that these fault-tolerant servers separate PCI I/O from the rest of the motherboard and add fault-detection hardware logic — Stratus' custom chipsets — which are essential for lockstep operation and effective fault detection and isolation.

These fault-tolerant servers take full advantage of standard Intel® server components and designs. Red Hat Enterprise Linux on ftServer systems takes advantage of Intel EM64T processor technology that supports both 32-bit and 64-bit applications.

Failsafe Software

Stratus fault-tolerant servers use system software to provide an availability-supporting ecosystem that complements the high reliability of the Red Hat Enterprise Linux operating system.

Stratus' addition of failsafe software capabilities addresses known sources of system and application failures, and minimizes downtime during repair or maintenance:

- Software is shielded from transient hardware errors.
- Hardened drivers prevent software failures.
- Software issues are reliably captured, analyzed, and corrected.
- In-memory data is maintained.
- Extensive integration and error-insertion testing finds and resolves difficult errors.

All ftServer systems running the Red Hat Enterprise Linux operating system provide 100% binary application compatibility. Stratus ftServer systems pass the same rigorous Red Hat Certification tests as other servers, offering further assurance that Linux applications will run compatibly on these systems.

Transient Hardware Errors

The ftServer hardware and system software is designed to detect, isolate, and automatically recover from transient errors as well as hard errors. Because error handling is a known vulnerability in software design, the masking of both transient and hard errors averts many potential software problems. The ftServer hardware and system software trap and handle transient

hardware and software errors that a cluster node or typical server would propagate to the operating system, middleware, or application software.

Hardened Device Drivers

Errant device drivers are acknowledged as a root cause for many operating system crashes. Driver enhancements on Stratus ftServer systems address this significant vulnerability. In the event of a problem, PCI I/O adapters are isolated from the rest of the system.

The following functional enhancements harden device drivers:

- Full support for hot insertion and removal of adapters (also known as surprise insertion and surprise removal)
- Transparent failover (except for tape)
- Ability to run online diagnostics
- Support for online firmware updates

Software Issues Analyzed and Corrected

Software-related issues can occur even with the best of preventive measures. The design of ftServer systems provides a substantial initial advantage in correcting software issues by reliably distinguishing software problems from hardware problems. With conventional servers or high-availability clusters, many problems attributed to software are actually caused by transient hardware errors. Because Stratus fault-tolerant servers automatically detect, isolate, and resolve transient hardware errors, issues are immediately separated into the appropriate category for more effective and timely problem resolution.

More important, the ftServer system design incorporates reliability improvements that help prevent software-induced failures from occurring in the first place. It is worth noting that conventional servers and high-availability clusters do not supply capabilities to prevent software failures. Conventional servers — even those marketed as resilient or robust — do not address prevention of software-induced failures. Clusters address this vulnerability with a restart and recovery mechanism to get software up and running again as quickly as possible; as a result, an application outage occurs.

The ftServer hardware and system software trap and handle transient hardware/software errors that a cluster node or conventional system would propagate, as noted previously. Another advantage of the Stratus Continuous Processing architecture is the ability of the hardware, software enhancements, and service features to assist in isolating and correcting operating system and device driver failures.

In the unlikely event that an operating system crash occurs, ftServer systems automatically reboot while preserving crash information in one of the replicated CPU/memory units. A kernel memory dump, useful for analyzing the cause of the event, is automatically taken after the system and application are back online.

Simple Network Management Protocol (SNMP) Agent —The ftServer SNMP Agent extends the Red Hat SNMP Agent, allowing third-party enterprise management consoles to remotely monitor ftServer systems. Most enterprise management software — including the Tivoli® Enterprise Console, HP® OpenView®, and CA Unicenter® products — supports SNMP.

The ftServer SNMP Agent sends a notification, in the form of an SNMP trap, any time a system component changes to any of the following states: broken, fixed, removed, or inserted. An ftServer MIB file is provided to allow the enterprise management software packages to understand Stratus alarms.

In-Memory Data Maintained

In-memory data is used extensively in many high-performance, business-critical applications; loss of this data can result in missed transactions or increased downtime. Unfortunately, cluster failover and software crashes both cause the loss of in-memory data. Stratus ftServer systems protect in-memory data from hardware failures using a lockstep architecture that stores memory contents in at least two separate hardware components.

Stratus ftServer systems provide key capabilities that preserve in-memory data:

- Quick dump
- Rapid disk resynchronization (RDR)

Quick Dump

With conventional servers, users have to make an uncomfortable choice after a crash. Either they can keep the application down while they take a system memory dump to be analyzed later, or get the application back up right away — but they lose the information that would help prevent a similar crash in the future. With ftServer systems, users no longer face this dilemma; they can minimize the time applications are offline and capture diagnostic information.

The quick dump capability capitalizes on the replicated hardware in fault-tolerant systems. In the event of an operating system software failure, an ftServer system keeps one duplicated CPU/memory unit offline while restoring the rest of the system to normal operation. As a result, a business-critical server gets back into operation quickly without forfeiting the information required to determine the root cause of the problem.

After the system and application are back in full operation, a standard kernel memory dump is taken using the contents of the offline CPU/memory unit. When the dump is complete, the offline CPU/memory unit is brought back into normal, partnered operation. The system automatically calls the Stratus Customer Assistance Center (CAC) to report the problem. Stratus service professionals can then immediately begin diagnosing the system dump and manage the problem to its resolution.

Rapid Disk Resynchronization (RDR)

RDR delivers higher protection and higher availability through RAID 1+0 for mission critical applications. RDR also delivers improved availability through faster remirroring of disks. If a disk or Customer Replaceable Unit (CRU) is removed for a brief time, only the changed blocks are remirrored. Full remirroring of replacement disks is much faster when using RDR.

Extensive Testing

Stratus employs a rigorous test process that targets fully integrated systems, including all hardware and software options, in a variety of configurations including maximum configurations. Servers are tested under extreme processing and I/O loads; errors are continuously simulated during the test process.

Much of this error-insertion testing is exclusive to Stratus fault-tolerant servers because many of the simulated errors, such as CPU or PCI bus failures, would cause conventional systems to crash. Finding and resolving these integration and error-insertion test issues produces a higher level of software reliability for ftServer systems.

ActiveService Architecture

The guiding design point of Stratus ftServer systems is the ability to detect and resolve problems *before* they cause system downtime. For that reason, the servers' many autonomous and/or proactive service capabilities include recognizing and riding through transient errors without interrupting processing. Uptime is similarly maintained in the event of hard errors, through the use of automatic fault detection, automatic fault isolation, integrated "call-home" remote support, and online component replacement.

The ActiveService architecture begins with the design of the hardware and its technology-enabled Access features and extends to Stratus' global ActiveService Network and Web-based ActiveService Manager.

Figure 4: Active Service Architecture Highlights

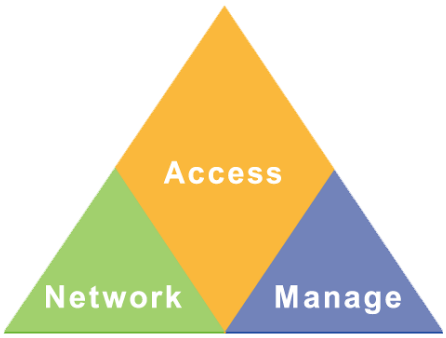

Built-in Serviceability

An example problem resolution scenario (*Figure 4*) demonstrates how built-in serviceability features equip ftServer systems to protect uptime in a way unequaled by other servers.

Stratus ftServer systems themselves include capabilities that provide the first level of customer support. Maintenance software within the server runs transparently to constantly monitor, diagnose, and report system events, accurately isolating faults to the component level. In the event of a hardware error or failure, the system automatically isolates the condition while continuing operation on a second, replicated component. In contrast with a cluster, there is no failover time and users do not experience a loss of performance.

Figure 5: Problem Resolution Scenario

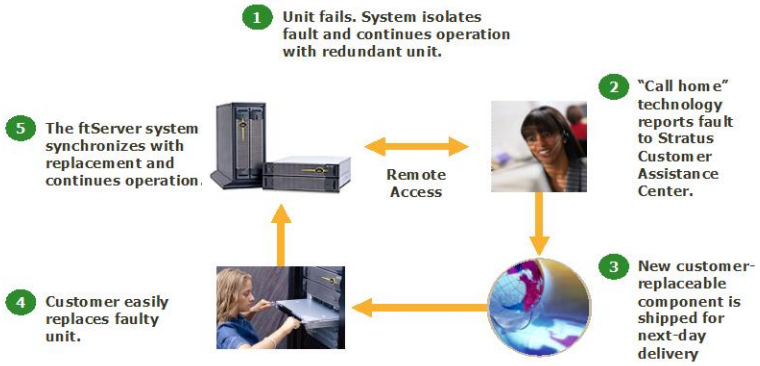

The server automatically reports any problem condition to a Stratus Customer Assistance Center (CAC) via a secure, dialed connection. The global ActiveService Network provides a worldwide infrastructure that enables remote access to every customer system. Authorized support professionals are able to remotely investigate critical problems 24/7, without the need to visit the customer site. In practice, Stratus service engineers have historically been able to troubleshoot and resolve problems online in more than 95% of cases.

In service situations where a replacement component is needed, the ftServer system automatically orders the correct replacement part. Replacement parts ship for next-day delivery to most locations. Because most ftServer components are customer replaceable units (CRUs), the replacement part can be installed easily by a user, without requiring special tools or assistance from a field service engineer.

The system brings the newly installed hardware into operation automatically, synchronizing it in lockstep with its partner. The system and the application continue to run normally throughout the entire process.

Reduced Exposure to Operator Error

The preceding section illustrates the point that ftServer systems, from the very first models, are designed for online component replacement to simplify service and reduce exposure to operator error. ftServer systems take this concept to the next level by reducing the number of system components and by tying components into a common chassis design that includes a blindmate backplane.

Extensive use of status indicator LEDs and keyed components eliminate potential operator errors during service operations. And because no operator commands are required to initiate component replacement or system reconfiguration, chances for error are even further reduced.

ActiveService Network
Like the systems it supports, Stratus' 24/7 service infrastructure was created with the express purpose of maximizing uptime for critical applications.

Every ftServer system is built to take advantage of the ActiveService Network, which provides a secure, continuous link between the servers and Stratus' technical experts and CACs worldwide. Connectivity to the Stratus ActiveService Network is provided via a dial-up modem, an Internet connection or the combination of a dial-up modem and Stratus ftGateway™ software. The ftGateway software feature allows several ftServer systems to share a common dial-up connection to the Stratus ActiveService network. This capability limits the need for phone lines and makes it easier to manage service connections for multiple ftServer systems located at a single site. The ActiveService Network enables online, around-the-clock monitoring and remote troubleshooting of systems regardless of their location, which virtually eliminates the delays and costs associated with on-site service.

Authorized service engineers use the ActiveService Network to access, investigate, and configure customers' ftServer systems. The network's powerful remote service management tools include remote reset, remote console capabilities, information capture and storage, and security features.

Diagnostic and analysis technologies allow the ActiveService Network to be used for determining the root cause of an event, and for uploading error logs and system dumps. Stratus service engineers can likewise use the network to install software patches, diagnostic routines, and hot fixes as needed.

Virtual Technician Module
The ftServer system provides out-of-band management capabilities through the Virtual Technician Module (VTM). The VTM allows for remote communication between the Stratus ActiveService Network and the customer's system, regardless of the server's state and is replicated for fault-tolerance. The VTM allows operations staff or service engineers to remotely power on/off or reset/reboot the system, and manage the security of incoming and outgoing communications through the ActiveService Network. To ensure system access, the VTM is an intelligent system that operates independently of the host computer.

The Virtual Technician Module is Stratus' second-generation remote access technology. The Virtual Technician Module introduces remote service capabilities including full remote keyboard, video and mouse, remote floppy/CD and out-of-band alerts.

ActiveService Manager
Complementing the ActiveService Network is the ActiveService Manager. This Web-based service tool supports online call management for ftServer systems. Designed to provide 24/7, real-time interaction with Stratus CACs, the ActiveService Manager allows users to review call tickets that have been created automatically by the system, as well as create and update support calls that are instantly routed to the appropriate support professionals within Stratus. In addition, the ActiveService Manager displays a complete incident history of Stratus systems throughout the customer's enterprise.

Stratus ftService^{SM} Options
Stratus provides a single source of accountability for complex inter-related platform, system software and Linux operating system issues. Customers may choose from several levels of proactive ftService^{SM} support options that make use of ActiveService technologies.

- **Assured Availability Plus[SM] for Linux:** Our premium service offering provides our most comprehensive level of 24/7 uptime support. It's our most popular service coverage among customers with demanding business-critical systems. Assured Availability Plus[SM] coverage is sold in conjunction with the Red Hat Enterprise Linux Premium Edition annual subscription services.

- **System Availability Plus[SM] for Linux:** Our basic ftService offering provides complete platform support during local business hours. System Availability Plus[SM] coverage coverage is sold in conjunction with the Red Hat Enterprise Linux Standard Edition annual subscription services.

Third-party collaboration is also a standard part of Stratus ActiveService Network and Web support center. Stratus ftService customers may designate non-Stratus vendors to view the status of calls through the ActiveService Manager and the ActiveService Network, building a virtual call queue for collaborative problem solving from anywhere in the world.

Focus on Mission-Critical Services
Because large-scale deployments place greater demands upon IT departments, Stratus 24/7 Worldwide Services provides a range of professional services offerings that focus on achieving maximum uptime and performance for ftServer computing solutions. Services include design, implementation, and management of availability solutions, as well as training and education.

Stratus' start-of-the-art Center for Fault-Tolerant Computing, located at its U.S. headquarters in Maynard, Massachusetts, provides compliance and availability testing of third-party products and customer applications, in addition to benchmarking and porting services.

These options help customers implement reliable, cost-effective solutions without the extra overhead of bringing technical skills in-house, diverting their internal staff from other projects, or contracting with several firms to design and install the application and infrastructure.

Conclusion

By distributing and supporting Red Hat Enterprise Linux on the ftServer product line, Stratus Technologies is taking fault-tolerant servers into today's Enterprise Linux market. Because the servers' reliability features operate transparently and automatically, applications certified on Red Hat Enterprise Linux stand to benefit from 99.999% or better server uptime protection without the need for human intervention, nor additional programming and testing.

Outstanding operational simplicity makes deploying and managing these servers easy and cost-effective. Every aspect of these continuously available servers is engineered to prevent unplanned downtime, not simply allow for quick recovery as high-availability clusters and "robust" standalone conventional servers are designed to do. Compared with reliability-enhancing approaches that are not intrinsic to a server's design, enterprises can reduce their exposure to the operator error that industry experts cite as a leading cause of unplanned downtime.

All of this produces a tangible financial advantage over competing approaches by reducing costs associated with complicated deployment, unplanned downtime, and ongoing support expenses.

Additional Resources

Stratus publishes a series of white papers that describe other technologies, features, and benefits offered by ftServer systems. **Learn more at www.stratus.com**.

Abbreviations and Acronyms

CAC	(Stratus) Customer Assistance Center
CPU	central processing unit
CRU	customer replaceable unit
DMR	dual modular redundancy
I/O	input/output
PCI	Peripheral Component Interconnect
RAID	Redundant Array of Independent Disks
RDR	Rapid Disk Resynchronization
SAS	Serial Attached SCSI

X934-C