



Marathon FTvirtual Server™ :
Meeting the Need for Low-Cost, High
Availability for Windows Platforms

Represented By:
Hollis Technologies Inc.
Offices in NH, MA & CT
Phone: 603-577-8958
Fax: 603-880-4827
E-mail: sales@hollis-tech.com
www.hollis-tech.com

Table of Contents

Introduction	1
The Unmet Need for Simple, Affordable High Availability	2
Options for Uptime Protection.....	3
Failover vs Fault Tolerance	
Fault Tolerant Systems	
Proprietary Fault Tolerant Solutions	
Marathon FTvirtual Server Software	5
Automatically Handles Hardware Failure	
Permits Application Operations During Repair	
Provides Full Disaster Tolerance	
Easy Administration, Low TCO	
Appendix 1: Side-by-Side Comparison of High Availability Technologies	

The Marathon logo, SplitSite and Marathon FTvirtual Server are trademarks or registered trademarks of Marathon Technologies Corporation. Microsoft and Windows are registered trademarks of Microsoft Corporation. All other trademarks and registered trademarks are the property of their respective owners. Copyright © 2005 Marathon Technologies Corporation. All rights reserved.

Marathon Technologies Corporation reserves the right to make improvements to this document at any time and without further notice. Marathon Technologies Corporation assumes no responsibility for any errors that may appear in this document.

Introduction

As companies grow increasingly dependent on server-based business applications, such as email and file/print serving, they are turning in greater numbers to entry and mid-range Windows-based servers and blade configurations to run them. These servers have two key advantages—they are cost-effective enough to allow companies to buy them in volume and flexible enough to run a wide range of standard applications. However, they also pose the significant risk and expense of application downtime and data loss.

Until recently, traditional technologies used to protect these servers from downtime have been too costly and complex to justify their widespread use. Less costly alternatives do not protect data or applications and require extensive, labor-intensive backup and recovery procedures in the event of a failure.

For the first time, companies can protect popular applications such as email, collaboration, call-center, database, file/print serving, security, and process automation resident on entry or mid-range servers. A new, highly cost-effective technology delivers the level of availability previously accessible to only the highest end computing platforms.

This paper discusses the various high availability server technologies and introduces Marathon FTvirtual Server™, a software-based fault tolerant operating environment used to protect entry and mid-range servers and blades from downtime due to faults, hardware failures, and physical threats.

The Unmet Need for Simple, Affordable High Availability

Server-based business applications, such as email and file/print serving have become both ubiquitous and essential to the operation of most companies today. Specialized areas such as process control, call center management, and broadcasting are also adopting server-based Windows applications in growing numbers. Because server-based applications are highly flexible and run on cost-effective, entry and mid-range Windows servers, they meet a wide range of customer needs. Large enterprise companies are using these applications to keep productivity high and overhead low in an increasingly competitive economy. Small to mid-size companies are also using server-based applications to achieve the level of efficiency and computing power previously available only to the largest organizations.

However, without a fault and disaster tolerant technology to protect them, entry and mid-range servers are at high risk for costly and politically damaging downtime. The costs of this downtime in end-user productivity, customer service reduction, and IT resources more than offset the savings these servers provide.

Protecting servers from downtime poses additional challenges. The cost and IT resources needed to implement, administer, and maintain downtime protection technologies for these servers is prohibitive. Traditional high availability technologies more than triple the cost of a server and require the adoption of proprietary server hardware and software drivers. These solutions make them impractical for widespread use in enterprise organizations and for any use in small to medium sized businesses. Medium-range technologies such as clusters protect only a limited number of applications and require advanced technical skills to implement and manage. Low-end technologies such as tape backup, load balancing, and remote replication expose end users to downtime, do not protect data and transactions, and require extensive labor-intensive backup and recovery procedures.

The need for a cost-effective, high availability solution for entry-level and mid-range servers is particularly acute in highly decentralized industries with numerous geographically dispersed locations such as retail chains, financial firms, and manufacturing sites. These industries have a large number of commodity servers in locations where there are few or no local IT resources (i.e. retail stores, branch offices, manufacturing plants). Other industries, such as broadcasting and health care systems providers require 24/7 operation of their servers with no tolerance for even brief downtime or minor data loss.

Unmet Needs

As end user and customer expectations for continuous service grows, so does the pressure on IT staff to restore service faster and to minimize data loss. Most companies are unwilling to accept the cost and complexity associated with high availability solutions and as a result, there are a growing number of companies running business critical applications on poorly protected servers. To gain the full benefit of entry-level to mid-range servers, companies need a high availability technology that provides:

- **Simple Administration/Low Total Cost of Ownership (TCO).** Technology should be cost-effective to purchase and simple enough for existing IT staff to manage.
- **Continuous Availability.** End users need continuous access to applications without interruption or loss of performance 24 hours a day, 7 days a week.
- **Complete Data Protection.** Technology must ensure that no transactions are dropped and that no data is lost due to faults or hardware failures.
- **Disaster Protection.** Business continuity best practices require that business critical servers will continue to operate through physical damage and disasters.

Options for Uptime Protection

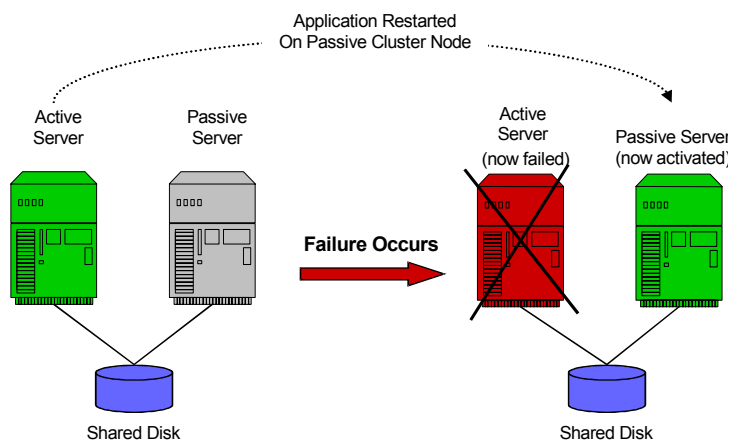
What are the options for providing this level of protection? Industry analysts, such as Forrester Research categorize technologies in terms of the percentage of uptime that is calculated based on Mean Time Between Failures (MTBF) and Mean Time to Repair (MTTR). MTBF is a measure of how long a server is likely to operate before a failure occurs while MTTR measures the time taken to bring the server back online after a failure occurs. High availability solutions, which deliver greater than 99.999% (“Five Nines”) uptime, are at the top of this uptime scale and RAID protected servers, which offer 99% uptime, are at the bottom. While the “nines” criteria is important, this criteria does not capture key differentiators, such as disaster protection, data protection, administrative and maintenance requirements, and TCO.

All high availability server technologies require hardware redundancy to ensure system and application availability. The different high availability technologies are categorized by how each manages redundant hardware, the level of availability they provide, and ease of operations in their design.

Failover vs. Fault Tolerance

A critical distinction among the various technologies used to protect servers from downtime is whether they require some “failover” or recovery procedure or are able to *tolerate* faults. Technologies such as clusters and data replication comprise a class of high availability systems known as **failover** technology. Failover technologies use redundant servers in an active-passive design in which one server actively executes the application operations while the redundant (passive) server is inactive. The passive server is available to assume operations should the active system fail. Using replication technology, data are periodically copied from the active server to a passive replication server. When a hardware failure occurs, application operations cease on the failed system and must be restarted on the replication system.

In a cluster, disk access is shared and accessible by both redundant servers. When a failure occurs in a cluster, the application is migrated from the failed server to the passive server as shown in Figure 1. In some implementations, the failover event is a manual operation that is initiated through administrator action while in others failover is handled automatically. When



using an automated failover technology, scripting, application modification, and testing is typically required to manage the migration of the application from the failed server to the redundant server.

A deficiency of failover designs is that service is interrupted during the failover event resulting in the loss of in-process transactions and a potential loss of application

Figure 1: In a cluster failover, application disks are shared and applications are restarted on the passive server.

data. Furthermore, if the failover event requires manual intervention to reestablish service, significant delays in operations may occur. Additionally, when the failed system is repaired, service must once again be interrupted to return processing to the repaired server.

Other drawbacks to failover and recovery technologies include:

- **Failover/service interruptions can take from a few minutes to days.** On the high end, active/active cluster on Windows 2000 has demonstrated better failover times of a few minutes but can still take more than an hour before service is restored on larger implementations. On the low end, restoring from tape backup can require days.
- Clusters **require more costly SAN configurations** to reach the higher levels of availability
- **More complex and time-consuming fail-back and recovery procedures** require complex custom scripting and application modifications
- **Downtime** may be necessary during repair and maintenance

Fault Tolerant Systems

By definition, fault tolerant systems continue to provide complete service to end-users during and after a fault or failure. With more than 99.999% uptime, fault tolerant technology delivers significantly higher levels of data and application availability than failover and replication technologies. Fault tolerant technologies eliminate the failover and recovery process and provide full service and performance to end-users even during maintenance.

Proprietary Fault Tolerant Solutions

Proprietary fault tolerant designs employ specialized hardware components and redundant server elements to achieve non-stop availability. These solutions process all machine operations on two redundant servers simultaneously and manage redundant disk storage to provide complete protection for both processing and data. The two servers are synchronized, appearing to administrators, the OS, applications, and the network as a single entity. If one server fails, the other continues to operate, delivering continuous service to end-users.

Drawbacks to *proprietary* fault tolerant high availability solutions include:

- High purchase cost due to proprietary hardware requirements
- Required commitment to proprietary brand of server hardware
- Limited choice of supported and qualified peripherals
- The latest generation of CPU / clock speeds are generally not available due to time required for hardware design-in and qualification
- Required use of proprietary drivers

The third main category of high availability technology is non-proprietary fault and disaster tolerant technology. Marathon FTvirtual Server is the only technology in this category. It provides fault tolerance and complete data protection but does so with a cost-effective, software-based solution that works on any standard brand of server or blade server.

Marathon FTvirtual Server Software

Marathon FTvirtual Server™ software is a breakthrough product that delivers true “five-nines plus” (>99.999%) availability, fault tolerance, *and* complete disaster protection to mid-range and entry level servers in a simple, highly affordable software-based design.




Figure 2: Marathon FTvirtual Server software unites two industry standard Intel based servers into a single fault and disaster tolerant configuration. Marathon software creates a perfect environment for applications and OS that is protected from faults and failures.

As shown in Figure 2, Marathon software works in concert with the Microsoft Windows operating system and qualified industry-standard Intel-based server hardware to create a *virtual* server that provides a fault-tolerant operating environment for Windows applications. Users accessing applications running in this fault tolerant environment experience seamless operations and continuous availability even in the event of hardware failures. Marathon software protects any Windows application, providing continuous computing; mirrored data storage; uninterrupted

network access; and service through hardware faults and network failures. These benefits are realized without requiring application modification or custom scripts.

A Marathon solution comprises two standard Intel-based servers that are interconnected with dual gigabit Ethernet networks. These networks are used to coordinate fault tolerant operations. In this configuration, the servers are known as CoServers and the networks are called CoServer links.

The Marathon software creates the FTvirtual Server, a virtual Windows server that is fault tolerant and operates as a guest OS across the two CoServers. The FTvirtual Server is

Characteristics of Marathon Solutions

- **Industry Standard Hardware** – Intel-based servers, unmodified Windows OS and applications, and standard drivers (no scripts, no APIs).
- **Seamless Availability** – Faults have no impact on end users.
- **Data Protection** – Raid 1 mirroring protects data in the event of a failure or site disaster.
- **Failure Identification and Isolation** – Automatically detects, identifies, and isolates errors or failures before any data can be corrupted
- **No Downtime for Repair** – Failed component can be repaired while the system is in operation
- **Automatic Redundancy Restoration** – Repaired/restored components rejoin lockstep operation automatically
- **Disaster Tolerance** – Provide complete protection in the event of disaster

identical to a standard Windows server and appears on the network as a Windows host server. Administration of the FTvirtual Server is like any Windows server. Applications are installed, configured, and run just as they do in a standard Windows configuration. Since the FTvirtual Server environment is fault tolerant, applications operating within the environment are completely protected from hardware and many operating system failures.

Marathon accomplishes fault tolerance by synchronizing data, memory, and program execution in the FTvirtual Server such that every operation executed under the control of one CoServer is identically executed simultaneously on the second CoServer. This technique, known as *lockstepping*, ensures that operations are redundantly processed, synchronized, and executed in a fashion such that when a failure occurs on one CoServer, application operations within the FTvirtual Server can continue under the control of the surviving CoServer.

Marathon FTvirtual Server is not a failover technology. In the event of a failure, applications operating within the FTvirtual Server continue without loss of transactions, data, or network connectivity (see Figure 3). All in-process disk transactions are maintained through the failure and application data is protected by disk data mirrored on redundant drives associated with each server.

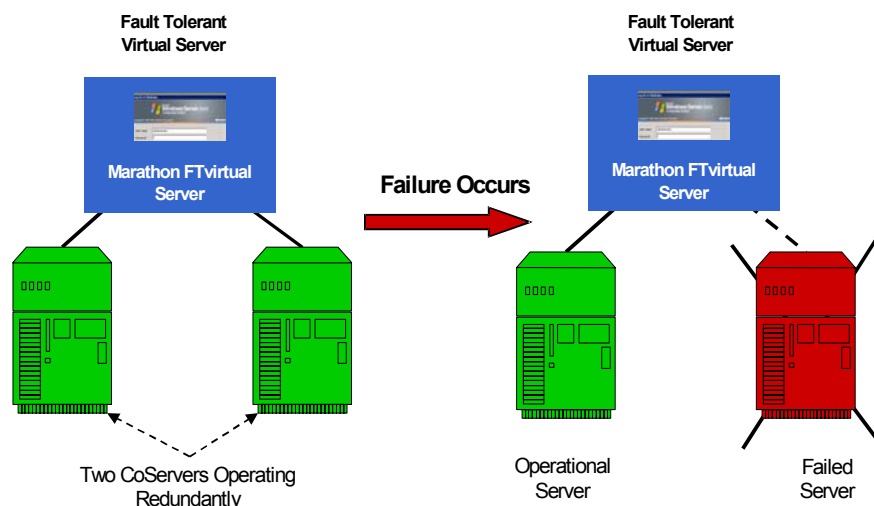


Figure 3: Marathon fault tolerant application environment survives server failures.

Marathon Automatically Handles Hardware Faults

Marathon continually monitors all system devices to determine which are operable and which are not. Marathon automatically discovers failing or faulted hardware and reconfigures hardware resources to maintain system operations in an optimal fashion and without requiring operator intervention. Errors detected in server memory, motherboards, or CPUs result in the complete removal of the faulty server. Errors discovered in disk or network adapters will result in device reconfiguration in a manner necessary to continue on-going operations. The reconfiguration process is not detectable by the user, application, or the operating system.

Marathon Permits Application Operations During Repair

Marathon allows you to shutdown one CoServer to repair, replace, or upgrade hardware components and install software upgrades while application operations continue on the

remaining CoServer. After a failed component is repaired or restored, it is simply booted and returned to service. The Marathon FTvirtual Server software automatically reconfigures and resynchronizes operations and data between the two CoServers and returns the Marathon configuration to full redundancy.

Marathon Provides Disaster Tolerance

The Marathon optional SplitSite® feature maintains the synchronous virtual fault tolerant environment even when the CoServers are physically separated. If one CoServer is destroyed, the surviving CoServer continues to compute through the disaster without failover, hesitation, or end-user disruption.

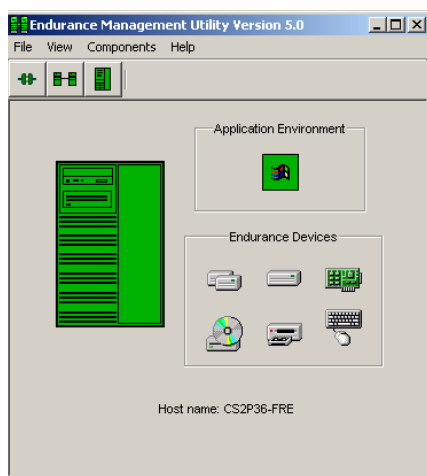


Figure 4. Marathon System Management Utility allows administrators to manage the virtual fault tolerant environment as a single entity.

Easy Administration, Low TCO

The Marathon management utility provides clear, detailed views of system status and allows administrators to manage the Marathon as a single entity. Management using industry standard SNMP frameworks is also supported using the Marathon management MIB.

Because Marathon FTvirtual Server was designed for low TCO and easy administration, it enables companies of all sizes to take advantage of the benefits of high availability and continuous operation. Marathon eliminates the administrative complexity and advanced skill set required for alternative solutions by using any brand of industry-standard, unmodified Intel-

based server hardware, unmodified applications and any Windows OS.

By using Marathon, companies can be assured of nearly 100% uptime while reducing IT resources needs to maintain availability.



Appendix 1 Side-by-Side Comparison of Active and Passive Redundant Technologies

	Marathon FTvirtual Server	Proprietary Active Redundant	Traditional Two-Node Clusters
Fault Tolerance Method	Active Redundant, Software <ul style="list-style-type: none"> Completely software based for flexibility, ease of use, and cost savings. Two redundant servers operating in lockstep Failure on one server has no impact on second. 	Active Redundant, Hardware <ul style="list-style-type: none"> Hardware, motherboard, and chipset based solution, resulting in higher costs, low flexibility Two redundant servers operating in lockstep Failure on one server has no impact on second 	Passive Redundant <ul style="list-style-type: none"> Active server backed up by secondary passive server Hardware faults trigger failover events requiring start up of application and transfer of control to passive server
Availability	Excellent <ul style="list-style-type: none"> No shared resources representing single points of failure All Windows applications supported OS and application protected from hardware failures and transient OS faults >99.999% uptime 	Excellent <ul style="list-style-type: none"> No shared resources representing single points of failure OS and application protected from hardware failures and transient OS faults >99.999% uptime 	Moderate <ul style="list-style-type: none"> Failures incur service interruption and loss of in-process transactions Repairs require that service is interrupted to return repaired hardware into service Fail-over times depend on failure class and configuration More costly SAN-based hardware configuration for optimal availability
Choice of Hardware Server Brand	Open <ul style="list-style-type: none"> Qualified leading server brands and hardware (HP, Compaq, Dell, IBM) Standard VGA No PCI card limitations All Windows device drivers supported Latest CPU, clock, chipset 	None. Proprietary-Only <ul style="list-style-type: none"> Proprietary motherboard I/O in proprietary section of system Proprietary ASICs Proprietary storage subsystem Disks only available from vendor Only vendor-certified PCI cards supported Only proprietary-hardened device drivers supported 	Open <ul style="list-style-type: none"> All major server brands
Support	Open/Cost Effective <ul style="list-style-type: none"> Open architecture on standard hardware and software minimize training and support needs Phone support available 24/7/365 worldwide 	Vendor direct-only <ul style="list-style-type: none"> Mandatory vendor support contract Service via vendor proprietary remote console card using vendor diagnostics. Software generates a call to vendor service center. Spares are vendor proprietary 	Complex <ul style="list-style-type: none"> Requires specialized training and advanced skill set Familiarity with application in cluster environment Custom failover scripting may not be supported
Administration & Management	Simple, cost-effective <ul style="list-style-type: none"> Unmodified OS, server hardware, and application Managed like a reference server 	Specialized <ul style="list-style-type: none"> IT skills specific to vendor hardware and environment required to manage proprietary hardware. 	Complex, costly <ul style="list-style-type: none"> Specialized IT skills required to manage and install applications in a cluster environment

	Marathon FTvirtual Server	Proprietary Active Redundant	Traditional Two-Node Clusters
Disaster Tolerance	Complete <ul style="list-style-type: none"> SplitSite® optional capability allows two CoServers to be physically separated for disaster tolerance Continuous computing through catastrophic damage to one system 	None <ul style="list-style-type: none"> Remote data replication only. 	Moderate <ul style="list-style-type: none"> Cluster over a distance capability available In-flight data lost during failover Operations pause during failover
Data Protection	Excellent <ul style="list-style-type: none"> No loss of in-flight transaction during hardware or OS fault No failover process. Continuous computing through hardware and OS failure Continuous data access 	Excellent <ul style="list-style-type: none"> No loss of in-flight transaction during hardware or OS fault No failover process. Continuous computing through hardware and OS failure Continuous data access 	Poor <ul style="list-style-type: none"> In-flight transactions lost during a failover Memory fragmentation may prevent failover unless MS guidelines are followed Application context lost during fault Access to data interrupted during failover Data recovery required
End User Protection	Excellent <ul style="list-style-type: none"> No failover state. Continuous operation at normal performance levels throughout hardware and OS faults User sessions unaffected Application context maintained In-flight transactions completed 	Excellent <ul style="list-style-type: none"> No failover state. Continuous operation at normal performance levels throughout hardware and OS faults User sessions unaffected Application context maintained In-flight transactions completed 	Poor <ul style="list-style-type: none"> Degraded performance for users during failover System reduced to 50% capacity of the original Active/Active cluster User sessions & context dropped during failure
Upgrades	Excellent <ul style="list-style-type: none"> Software-based technology upgraded with new chip introduction 	Poor <ul style="list-style-type: none"> Chip/motherboard-based technology Slow to adapt to new technology Proprietary ASICs to enable the technology must be redesigned with each technology turn 	Complex, costly <ul style="list-style-type: none"> Specialized IT skills required to do basic upgrades for a cluster environment
Blade Server Protection	Excellent <ul style="list-style-type: none"> Provides full fault and disaster protection for blade servers 	None	None

Marathon Technologies Corporation
295 Foster Street
Littleton, MA 01460 U.S.A.
Phone: 978.489.1100
Fax: 978.489.1101
E-mail: info@marathontechnologies.com
www.marathontechnologies.com