

Protecting Data and Applications from Major Disasters *and* Minor Disruption - A More Effective Approach

INTRODUCTION

In the wake of catastrophic events like Hurricane Katrina and the 9/11 terrorist attacks, enterprises have made disaster recovery a top priority. Over the past several years, most organizations have implemented disaster recovery plans that typically focus on preparedness for a major disaster. After a catastrophic event, these disaster recovery solutions enable companies to regain access to the data, hardware and software necessary to recover and resume critical business operations.

Businesses have started combining their disaster recovery plans with their high availability plans—a strategy that too often results in inadequate availability protection for the most likely and most prevalent disruptions. While enterprises need to be prepared for “the big one,” the reality is that systems usually go down as the result of far less dramatic events, like network failures, disk crashes, or power, fire or telecommunications outages. In these cases, do you really want to restart your system at your disaster recovery location thousands of miles away? Is your organization equipped to deal with business interruptions, potential data loss, and significant reconfiguration efforts every time you encounter a minor system failure?

As enterprises face mounting pressure to keep their systems running 24x7 and provide continuous access to data, applications and transactions, companies need to ask themselves whether it makes good business sense to use traditional disaster recovery solutions to protect against day-to-day disruptions and outages. By its very nature, disaster recovery involves some level of downtime, data loss, and IT intervention that can be extremely costly in today’s fast-paced, competitive business climate. As a result, companies must re-examine their current strategies and consider that the steps they take to protect their business from major disasters are probably not the best way to protect their systems from common failures.

Protecting Data and Applications from Major Disasters *and* Minor Disruptions - A More Effective Approach

DISASTER RECOVERY SOLUTIONS ARE LESS THAN OPTIMAL FOR DISRUPTIONS THAT ARE LESS THAN DISASTERS

Disaster recovery solutions are designed for just that: recovery from true disasters—not minor disruptions. Given this reality, is it in your business' best interest to use your disaster recovery solution to ensure application availability in the event of minor disruptions? To effectively assess the viability of this approach, it's helpful to take a closer look at:

- **Recovery Time Objective (RTO):** How long can you afford to have your system down?
- **Recovery Point Objective (RPO):** How out-of-date can you afford your data to be once the system is up and running again?

RPO and RTO are both measured in units of time, with values ranging from seconds to days or weeks. The closer an application's RPO and RTO values are to zero, the more dependent the organization is on that application. In the case of critical, high-priority applications, the risks associated with recovering from even minor system failures become very significant. As a result, more and more companies are recognizing that using disaster recovery solutions for high availability is simply not an optimal strategy.

When it comes to minor system 'hiccups'—as opposed to serious disasters—the goal is to keep both RTO and RPO to an absolute minimum—if not zero—and to minimize disruption to the IT infrastructure and to the user community. Most conventional disaster recovery solutions, however, fall short in their ability to help companies achieve this goal.

RECONSIDERING DISASTER RECOVERY SOLUTIONS FOR DAY-TO-DAY DISRUPTIONS

Tape vaulting, the process of backing up data to magnetic tape that is stored in an off-site location, continues to enjoy widespread use as a disaster recovery strategy. While still a viable solution for long-term data archiving, tape vaulting offers no application protection and always results in some loss of data and transactions—quite often a substantial loss. Depending on where the backup tapes are stored, RTO can extend to days, resulting in serious business interruptions, potentially accompanied by loss of productivity, customers and revenue. And depending on backup frequency, remote tape vaulting often falls short in achieving all but the most conservative RPO.

ELIMINATING YOUR DOWNTIME COSTS

Gartner estimates that the average cost of downtime, across industries and applications, is \$86,000 per hour. Here are some estimates for applications in mid-sized and large businesses that require high availability:

Application	Downtime Cost per Hour
ERP	\$ 780,000
Supply Chain Mgt.	\$ 660,000
E-Commerce	\$ 600,000
Internet Banking	\$ 420,000
Customer Service Ctr.	\$ 220,000
EFT	\$ 210,000
Messaging/Email	\$ 60,000
Hospital Information System (avg. 3 hospital IDN with 1400 beds)	\$ 60,000
Hospital Information System (avg. single hospital with 500 beds)	\$ 15,840

Protecting Data and Applications from Major Disasters *and* Minor Disruptions - A More Effective Approach

Many enterprises looking to achieve shorter RTO and RPO use disk-based backup, or data replication technologies, which automatically duplicate and update data on another computer across a network—usually to a geographically dispersed location—to safeguard against local disasters. Most replication solutions use asynchronous methods to update the data that reside at a replicated site after the primary data have changed. This approach enables virtually unlimited distances between the production and backup systems, but cannot guarantee data integrity as transactions will likely be lost when a failure occurs. However, data replication typically delivers better RTO and RPO results compared to tape backup and augments it nicely.

While highly effective for disaster recovery after major disruptions, asynchronous data replication presents significant drawbacks when used as a high availability solution for common day-to-day failures. Because many disaster recovery solutions, such as asynchronous data replication, offer some type of failover to allow application re-start on the secondary backup system, many businesses have chosen to use these solutions for application availability as well. However, the manual failover processes, data loss and minimal system and failure analysis associated with most disaster recovery solutions can lead to false and failed failovers. In addition, with disaster recovery solutions, businesses often face significant effort to get users reconnected once the application is restarted at a very remote location. As a result, these solutions too often prove less than optimal for ensuring application availability.

APPLICATION AVAILABILITY + DISASTER RECOVERY = EFFICIENT PROTECTION

For companies looking to stay ahead in today's competitive business climate, it's important to ensure efficient protection of critical data and applications across virtually any set of circumstances. That's why companies would be well advised to reconsider their reliance on disaster recovery solutions for responding to minor disruptions and system failures. No longer willing to risk data loss and costly business interruption every time a system goes down, these companies are recognizing the need for local-area availability solutions that enable applications to operate through virtually any circumstance—whether a failed network connection, power outage, storm, building fire or even a more widespread natural or man-made disaster. With local-area availability approaches, the goal is to protect the systems and applications so as to avoid having the disaster recovery solution kick in until absolutely necessary—if at all. This way, businesses avoid the typical processes incurred by a disaster recovery failover when dealing with minor disruptions.

How can an enterprise efficiently protect its data and applications from both major disasters and minor disruptions? Optimal protection comes from combining true

Protecting Data and Applications from Major Disasters *and* Minor Disruptions - A More Effective Approach

availability solutions that prevent disruptions and ensure immediate and seamless recovery from the majority of failures with conventional disaster recovery solutions that provide remote application and data backup to protect the business in the event of a widespread catastrophic event. Consider the case of Hurricane Katrina, in which floods and storms ravaged large areas of both Louisiana and Mississippi. While this area encounters many storms, one of this magnitude is a rare occurrence. Under most circumstances, a true local-area availability solution between Louisiana and Mississippi, for example, would likely be sufficient to protect the business from most failures and outages. In the rare case of Katrina-like devastation, a long-distant disaster recovery solution would ensure that the business can recover when it is safe and practical to do so.

This comprehensive approach provides true, non-disruptive availability protection for the majority of failures and outages without requiring massive infrastructure and client-side changes and without major disruptions. At the same time, it ensures your business is protected from the lesser chance of a major outage and that you meet any regulatory or legal requirements regarding data protection.

THE MARATHON SOLUTION

Marathon Technologies provides a complete range of automated, fault-tolerant-class availability solutions for all essential Windows applications. Our patented everRun™ software completely synchronizes two standard Windows® servers, including the operating system, application, network interfaces, storage, and data. These servers can be separated geographically for optimal disaster tolerance—not just disaster recovery. Unlike cluster or failover solutions that require two fully configured and managed systems, everRun creates a very simple and automated availability solution that requires little setup and configuration while providing the most stable and reliable platform for running all of your critical Windows applications.

everRun prevents interruptions and downtime by fully automating fault management. Whether a component failure, a complete system failure, or even a site-wide outage, everRun can ensure the Windows application continues to operate uninterrupted. Best of all, the process is completely transparent to the user, to the application, and to the operating system.

Protecting Data and Applications from Major Disasters *and* Minor Disruptions - A More Effective Approach

WORLDWIDE HEADQUARTERS

Marathon Technologies Corporation
295 Foster Street, Littleton, MA 01460
Tel 1.800.884.6425 / 1.978.489.1100
Fax 1.978.489.1101

EMEA HEADQUARTERS

Marathon Technologies UK Ltd
Regus House, Trinity Court
Wokingham Road, Bracknell
Berkshire, RG42 1PL
Tel +44 (0) 1344.706.241
Fax +44 (0) 1344.706.242
Email:
emea@marathontechnologies.com
web: www.marathontechnologies.com

And to provide optimal protection—even in the event of “the big one”—consider disaster recovery solutions that perform data replication to remote sites located hundreds or even thousands of miles away. Should your region suffer such a major outage that it brings down the entire everRun disaster tolerant configuration, you can easily recovery your business once you are able to get systems back online.

What does this mean for your enterprise? With Marathon, your organization can protect its data and applications from both major disaster and minor disruptions—in the most cost-effective and people-effective manner possible.

Want to learn how more about how Marathon's availability and disaster recovery solutions can provide continuity of your operations across virtually any set of circumstances? Contact us for more information or to take a test drive.

MARATHON
Run to Infinity