


# *Protect Your Business Against Disruptions and Disasters*

April 2005



**i/+^** International Legal  
Technology Association  
*Peer Powered*

A Publication of ILTA

## About ILTA

---

Providing technology solutions to law firms and legal departments gets more complex every day. Connecting with your peers to exchange ideas with those who have “been there, done that” has never been more valuable.

For over two decades, the International Legal Technology Association (formerly known as LawNet) has led the way in sharing knowledge and experience for those faced with challenges in their firms and legal departments. ILTA members come from firms of all sizes and all areas of practice, all sharing a common need to have access to the latest information about products and support services that impact the legal profession.

*ILTA’s Statement of Purpose: ILTA is the premier peer networking organization providing information resources to members in order to make technology work for the legal profession.*

## Editors’ Note

---

Unfortunately, no law firm is invulnerable to business disruptions or disasters. Whether natural, environmental or man-made, serious or fairly minor, sooner or later an event will happen beyond our control that slows or halts the wheels of business. And when it happens, we will lose precious data, data that is invaluable to the work of our lawyers and the firm or law department itself. Count on it — temporary or permanent data loss is inevitable — and with it comes the loss of money and possibly clients.

But fortunately, it’s equally inevitable we have the ability to weather such an event. Through careful planning and the utilization of technology, we have the power to turn any business disruption into business continuity.

Does your firm have such a plan? Whether your answer is “yes,” “no” or “maybe,” we’re sure these authors — whose articles range from discussing the human factor in business continuity planning to the technical — will help prepare (or better prepare) you for the inevitable.

*Andy Spiegel and Randi Mayes, Editors*


---

## Table of Contents

Business Continuity: What Is the Right Fit for Your Firm? .....	3
<i>by Michael Oh of CISSP</i>	
Writing the Business Continuity Plan: Don't Leave Out the People Factor .....	6
<i>by Atlas Lee of Shook, Hardy &amp; Bacon L.L.P.</i>	
Disruptions Needn't Spell Disaster .....	7
<i>by Gregory Hanna of TOSS Corporation</i>	
Create a Competitive Edge for Your Firm by Ensuring 24/7 Operations .....	10
<i>by Mark Boltz of Stonesoft Corporation</i>	
Data Replication: A Key Element in BCP .....	13
<i>by Jason Buffington of NSI Software</i>	
Never Failover on a Monday: Overcoming The Five Fatal Flaws of Traditional Replication Solutions .....	16
<i>by Russell Sachs of MessageOne, Inc.</i>	

# Business Continuity:

## What Is the Right Fit for Your Firm?



Since that unforgettable day, September 11, 2001, business continuity, disaster recovery and contingency planning have been on everyone's mind. Yet three and a half years later, too few firms have actually implemented any type of business continuity plan. Whatever the size and the scope of a firm, it should have such a plan — one with the “right fit.”

It's easy to think a disruption or a disaster will never happen to your firm. Yet, business interruptions occur all around us all the time, and from three main sources:

*Natural (fire, flood)*

*Human (error, terrorist acts, sabotage, malicious code, electrical power failure)*

*Environmental (equipment failure, software error, power or telecommunication network outage)*

by Michael Oh of CISSP

Natural, human and environmental events cause disruption of operations in widely varying scope, and depending on the firm's size, some of these events can critically impact the ongoing operations. A Data Pro study identifies business disruptions by these percentages:

*Errors and omissions — 50 percent*

*Fire, water, electrical — 25 percent*

*Dishonest employees — 10 percent*

*Disgruntled employees — 10 percent*

*Outside threats — 5 percent*

Some of the more massive disasters such as 9/11 and the Northeast power grid shutdown affect a wide area, while others such as electrical power failures have a smaller impact. Whatever the causes, a business continuity plan is the firm's best defense for mitigating disruptive events and returning to normal operation after an emergency.

### What Is a Business Continuity Plan (BCP)?

According to the National Computer Security Center (NCSC) of the National Security Agency (NSA), a BCP is “A plan for emergency response, backup operations and post-disaster recovery maintained by an activity as a part of its security program that will ensure the availability of critical resources and facilitate the continuity of operations in an emergency situation.”

Such a plan does not have to be overly extensive, but it does have to be implemented; and it should be sized to fit the firm.

As an example, in a one-lawyer firm it could be as simple as having an alternate site such as a home, where operations could continue outside the office. A simple rerouting of the phone number, a backup computer, phone and connection to the Internet may be all that is required.

The needs change for a larger firm in terms of scale, scope and technology — but not in terms of planning. A BCP is a process of planning for disruption so that a firm can continue its critical operations during the disruption and make a smooth recovery when the crisis has abated.

### What Is a Disaster Recovery Plan (DRP)?

A disaster recovery plan (DRP) is a subset of the BCP. Its purpose is to recover essential IT systems, at least to some

degree, in the event of a disruption or a disaster. A DRP has a short-term focus on the continuity of services during the disaster period, whereas the BCP is long-term focused and is designed to normalize operations. The DRP addresses the data center recovery, user operations during a disaster, protection of data and information, IT service capacity and availability during a crisis and the restoration of all systems from the disaster recovery mode.

## What Constitutes a Contingency Plan?

So what is involved in putting together the BCP/DRP? According to the National Institute of Standards and Technology's (NIST) Computer Securities Resource Center (CSRC) documentation on contingency planning, the following procedure should be used to develop the plans:

1. **Develop the contingency planning policy statement.** *A formal policy provides the authority and guidance necessary to develop an effective contingency plan.*
2. **Conduct the business impact analysis (BIA).** *The BIA helps to identify and prioritize critical systems and components.*
3. **Identify preventive controls.** *Measures taken to reduce the effects of system disruptions can increase system availability and reduce contingency life cycle costs.*
4. **Develop recovery strategies.** *Thorough recovery strategies ensure that the system may be recovered quickly and effectively following a disruption.*
5. **Develop an IT contingency plan.** *The contingency plan should contain detailed guidance and procedures for restoring a damaged system.*
6. **Test the plan, and provide training and exercises.** *Testing the plan identifies planning gaps, whereas training prepares recovery personnel for plan activation; both activities improve plan effectiveness and overall preparedness.*
7. **Maintain the plan.** *The plan should be a living document that is updated regularly to remain current with system enhancements.*

These steps should be viewed with three phases that govern the actions to be taken following a system disruption.

**Notification/Activation.** *Describes the process of notifying recovery personnel and performing a damage assessment.*

**Recovery.** *Discusses a suggested course of action for recovery teams and personnel to restore IT operations at an alternate site or using contingency capabilities.*

**Reconstitution.** *Outlines actions that can be taken to return the system to normal operating conditions.*

## How Do You Determine the Best Fit?

To find the right fit for your firm, the BCP as outlined above needs to be completed. These steps are guides and signposts to assist in determining the need so that it can be properly matched to the solutions available. Each step should have the input of the management and the user community to truly determine the critical nature of the resources and their impact on the business.

This risk analysis requires identification of critical systems and weighing the loss of assets versus the cost of implementing mitigating controls. There are four "reactions" to any potential risk:

**Avoidance.** *Don't be involved with the risk.*

**Acceptance.** *Acknowledge and accept the risk.*

**Transfer.** *Shift the risk to someone else.*

**Reduction.** *Take appropriate measures to reduce the impact of the risk.*

*Avoidance* will be impossible since we are dealing with issues outside of our control. *Acceptance* is to accept all results of disruption; this would be unacceptable to most firms. *Transfer* is to transfer the risk to an insurance company, which may be acceptable to a certain extent but may not be sufficient. And *reduction* would be to take measures such as BCP to reduce the impact of any risk.

A business impact analysis (BIA) is critical in determining the appropriate best fit for a firm. It helps you identify the critical IT resources and what tolerance the firm has for interruptions of these critical resources during a disaster or a disruption of business. This is a critical piece of the analysis and will help fully define the requirements, processes and interdependencies of the IT systems to determine the contingency requirements and priorities. The BIA should contain the following actions:

*Identify all critical IT resources such as e-mail, document management, accounting and related interdependencies.*

*Identify the impact of disruption and allowable outage times for the critical resources.*

*Develop recovery priorities for the critical resources.*

Based on the BIA, you can further reduce the risks by looking at preventive measures. This will help mitigate the risks by detecting, deterring or otherwise reducing the effects of a disaster or a disruption. These involve such items as:

*Uninterruptible power supply (UPS) for short-term power loss, spikes and other such disruptions*

*Diesel generator or redundant power feeds for long-term power interruptions*

*Fire suppression systems*

*Water sensors*

*HVAC systems with adequate capacity for redundancy*

*Water/fireproof storage for media*

*Offsite storage of backup media*

*Frequent scheduled backups*

*Master on/off switch for system shutdowns*

*System access and protection securities, such as anti-virus software and firewalls*

## Developing a Recovery Strategy

Priorities and allowable recovery times must be planned. These could also involve a specific decision point by the management to determine when and if a disruptive event requires triggering of further contingency plan steps. For example, if a firm finds that the main office cannot be used for a day due to a short-term issue such as smoke, is that event on Friday at 2:00 p.m. sufficient to trigger a failover to an alternate site? If the process of the failover and the recovery consumes more resources and time and is more disruptive to the normal operations of the firm, should it be triggered anyway? This type of decision point should be outlined and a management resource identified to provide a decision.

Furthermore, how much risk can be transferred? Business continuity insurance can significantly reduce the amount of exposure to the firm. How much does the insurance require in terms of BCP and preparedness? How much do clients require from the firm for the BCP and preparedness? These are other variables that need to be included in the overall strategy for recovery.

IT contingency plans would be the next step in determining the best fit. In the case of a longer-term disruption, alternate site possibilities need to be explored. The needs would be based on the BIA and determined by the business needs of the firm's critical resources. Once triggered at the decision point how much time is required for the alternate site to be operational? These decisions require a review of the various potentials in the chart below.

Site	Cost	Equip.	Telecom	Setup Time	Location
Cold Site	Low	None	None	Long	Fixed
Warm Site	Medium	Partial	Partial/Full	Medium	Fixed
Hot Site	Med/High	Full	Full	Short	Fixed
Mobile Site	High	Dependent	Dependent	Dependent	Not Fixed
Mirrored Site	High	Full	Full	None	Fixed

Additionally, what backup methods would be needed to provide the data at the alternate site? A mirrored site would require a continuous synchronization of the data, where the cold site could do with a 24-hour-old tape backup stored at an offsite facility. How fast should the alternate site ramp up? What does that mean in terms of equipment procurement? Who is providing which equipment? What is the cost? Is this cost-effective?

People are critical in times of a disaster or disruption. Their roles and responsibilities need to be defined, and there also has to be in-depth cross training in case some people are not available during the process. These people should be designated into teams and assigned specific portions of the plan. And don't forget communications. During any disruption, communication is critical, since the natural events will gravitate toward chaos.

Effectiveness is revealed only in regular plan testing intervals and procedures; these should vary from walk-through on paper to actual simulation testing based on the level required by the BIA. There should be testing of all facets, including cross-team coordination and communication, reporting, security and mixed teams in case of unavailability of specific individuals. At each testing interval, update the plan and continue the evolution as the criticality of systems changes. Business requirements change, so, too, the plan should evolve and change.

## Where There's a Need There's a Fit

Anyone can agree on the critical need for business continuity plans, yet for a variety of reasons too many firms fail to implement one — and when they do, mistakes in the initial attempt too often include lack of ownership, lack of testing and lack of planning. To ensure that your organization has and maintains an active BCP/DRP requires you to constantly evaluate your business needs and to size the solution to fit your firm. Such issues as management support, shortage of funds and lack of priority should be addressed not only by technology, but equally important, by examining the business needs and understanding how to find the right fit for your firm.

## Resources:

National Institute of Standards and Technology (NIST): Computer Securities Resource Center (CSRC) — <http://csrc.nist.gov/publications/>

Availability.com — [www.availability.com](http://www.availability.com)

Contingency Planning Management (CPM) — [www.contingencyplanning.com](http://www.contingencyplanning.com)

Disaster Recovery Institute International (DRII) — [www.drii.org](http://www.drii.org)

Iron Mountain — [www.ironmountain.com](http://www.ironmountain.com)

SunGard — [www.sungard.com](http://www.sungard.com)

# Writing the Business Continuity Plan:

*Don't Leave Out the People Factor*

by Atlas Lee of Shook, Hardy & Bacon L.L.P.



**E**mployees are a key factor in recovering a business after a disaster or serious business interruption. Yet, in the hurried quest to return an organization to a state of normalcy or “business as usual,” it’s the people factor that too often gets left out of a business continuity plan that otherwise includes the best laid plans, processes and procedures.

During my career, I have heard too many times about how people were treated badly, misdirected and unappreciated during an effort to recover a business or business function. Such treatment of an employee or group of employees will eventually take its toll emotionally, psychologically and physically. And the ensuing results could potentially include high anxiety, trauma, prolonged illness, isolationism, absenteeism and turnover during and/or after a recovery effort.

In order to avoid or at least minimize the negative impact a recovery effort can have on the employee, the people issue has to be addressed early in the planning stages of developing a business continuity or disaster recovery plan. The initiative must also have the full support and commitment of executive level management to ensure that provisions are implemented and maintained to take care of the people in the event of another disaster or serious business interruption.

Some provisions that organizations should consider having in place for their employees during and after a crisis include:

*Continued pay during a prolonged outage*

*Free meals during long continuous stays in the office or recovery facility*

*Extended childcare*

*Financial assistance*

*Petty cash*

*Utilization of the employee assistance program (EAP)*

*Provisions for temporary shelter*

The people factor should also be part of a firm-wide user awareness program to ensure employees are aware of the existence of the aforementioned provisions. Being proactive in this endeavor will help lessen the anxiety, promote faith in the business, promote *esprit de corps* and act as a morale stabilizer. All this will eventually happen because of the firm’s obvious willingness to put a premium on their people and to lessen the negative impact of any type of catastrophic event that not only affects the business, but the lives and well-being of the employees.

In addition to user awareness, proper training must be given to all levels involved in the recovery planning process. The training must also be applied consistently across all levels within the organization. Training is vital due to the simple fact that everyone should know his or her role in an emergency. An ongoing training program produces continuity, provides tangible and usable communication, espouses confidence and helps prevent chaos, thus reducing the stress, anxiety and other associated people issues during and after an emergency.

## Taking Care of Business — and Your People

Making a conscientious effort to take care of your people is one of the best investments you can make toward normalizing business operations after a business interruption or disaster event. In fact, no other aspect of your business continuity or disaster recovery plan is more important.

# Disruptions Needn't Spell Disaster

by Gregory Hanna of TOSS Corporation



**L**osing data can be disastrous for law firms and legal departments. It is critical to shift the strategic focus from disaster recovery planning to business continuance.

Disasters happen. But with today's technology, law offices can safeguard their data from the many kinds of disasters, even those that bring down the computer system or destroy the main office and data center. The ability to overcome a disaster is critical because much of the lawyer's work product is stored as electronic data. Temporary loss of data can cost anywhere from a few thousand dollars an hour to a few hundred thousand dollars an hour, depending upon the size of the firm. Permanent loss can jeopardize the future of the firm.

But protecting a law firm's data should not be a life-or-death proposition. Over the past five years, the technology to mitigate computer disasters has improved dramatically. There's really no reason why any law office, even a solo practitioner, cannot be protected against a computer disaster. Follow this basic plan:

*Identify your firm's most critical applications.*

*Know your firm's recovery time objective (RTO) and recovery point objective (RPO).*

*Apply technology to mitigate the failure of critical applications and to meet or exceed your RTO and RPO.*

*Document in detail the steps necessary to resuscitate each application when failure occurs.*

*Practice, at least bi-annually, the recovery process with each critical system.*

A number of technological options exist, running the gamut from those that cost virtually nothing to solutions for megafirms with multiple offices.

Because matching the technology to the organization's needs and budget is the goal, it's critical to ask the right questions and develop the right criteria early on. Mistakes made in the initial planning stages tend to snowball into costly rework. Here are some questions that can get you on the right track:

*What critical applications do we have whose loss for an hour or more would seriously impact our operations?*

*If our offices were destroyed or inaccessible, what is our required recovery time objective — that is, how long will it take to recover our data and get operating again? Is that quick enough?*

*What is our required RPO? RPO tells how much data is at risk. If the required RPO is an hour before the disaster, then you lose an hour's worth of data. If it's a week before the disaster, you lose a week's worth of data, and so on. You can think of RPO as the "freshness" of your recoverable data.*

*How much will system downtime cost in terms of lost productivity and revenue?*

*How long can our systems be down before lack of access to our data begins to jeopardize our operations or causes clients to take their business elsewhere?*

*If key records are lost, how quickly can we reconstruct them?*

*If data from tapes cannot be recovered, what will it cost to re-enter it? And will the required data still be available, or will it have been eradicated in accordance with the organization's retention policies?*

## Basic Security Precautions

Once the organization has a clear sense of its business continuity needs, it makes sense to develop a "call tree" or list of contacts designating who will call whom in the event of a computer disaster. Individuals must know where to go and what to do. For example, the plan may call for employees

with home computers or wireless units to work at home until the disaster is under control. The plan should also include a list of local vendors that can be consulted in an emergency. Some vendors can perform amazing feats of data recovery.

While addressing these critical functions, it's important to observe the kind of mundane precautions that often go unnoticed:

*Check to see if surge protectors are in place.*

*Install an uninterruptible power source.*

*Ensure proper climate control for servers and back-office equipment.*

*Look for obvious risks such as PCs and servers located under sprinklers.*

Law offices can protect software from hackers and viruses by installing and regularly updating firewalls, intrusion detection systems and virus protection software. The key here is to make sure that updates are regularly downloaded to the individual PCs, laptops and servers. Because one person's failure to update can imperil the whole operation, someone should have responsibility for ensuring that the appropriate updates are installed.

If an office is using case management or document management software, all data stored on the network should be synchronized to laptops or desktop PCs so that work can continue even if the network fails. For the same reason, documents should be stored on both the local PC hard drive and server. Once the server is back online, the network can be brought up to date by synchronizing with a desktop or laptop.

Every law firm and corporate legal department should have a written policy requiring the use and frequent changing of passwords. In addition to mandating passwords that are tough to crack, the policy should prohibit the posting of passwords and require users to log off or "lock" the screen of their computers when leaving their desks for an extended period of time.

### Online, Offsite Backup

Many law offices protect their data by backing it up daily to tape. This is a high-risk strategy. If the tapes are stored in the same building as the office itself, a fire or natural disaster could destroy those tapes or render them inaccessible. In the event of a major disaster, it could take days to recover the tapes. Even if the tapes were accessible and undamaged, there's still the risk that the data might not have been properly written to the tape. Most organizations use tapes beyond the manufacturer's recommended specifications, leading to media errors and unrestorable data.

Backup tapes deteriorate over time and eventually become unusable. The general rule is to write to a backup tape no more than 10 times and to use a tape no longer than 12 months. Because of the unreliability of backup tapes, a number of organizations are beginning to turn to IPStor-based VTLs (virtual tape library) and data vaulting.

### Virtual Tape Libraries

Because of escalating data growth, mid-size and large law firms and corporate legal departments face the challenge of moving an increasing amount of data into backup in a shorter period of time. A cost-effective solution to this challenge is the VTL.

In simple terms, this is an appliance with an array of low-cost disks or access to such an array that is connected, via a copper or fiber-channel network at the law office's site, to both its physical tape library and its backup servers. The backup servers "see" the VTL as an actual tape library and automatically back up data onto virtual tapes at a much higher speed than would be possible with physical tapes. The data can then be exported to physical tape for offsite storage. In the event of a data disaster, the VTL can restore the law office's files and directories much more quickly than is possible with physical backup tapes.

For another layer of protection, the onsite VTL can replicate its data virtual tape to a VTL located at an offsite facility. The data virtual tape can then be exported, if desired, to a physical tape library at the remote facility.

The VTL does not replace the law office's tape technology. Rather, it enhances its speed and reliability without requiring changes to existing systems or major backup software investments in new equipment or software. As an added benefit, the deployment of a VTL extends the life of the physical tape library by reducing the wear and tear on drives, gears, robotics and heads.

### Data Vaulting

Moving data back and forth from the organization's main office to a remote storage facility is relatively simple. But what if there are multiple offices, and what if the lawyers and other staff in those offices are using different kinds of computers, including wireless laptops? Data is no longer kept under one roof; instead, it's scattered in pockets across a network that can span the country, possibly even the globe. A hardware failure or software malfunction anywhere in the system could create a major disaster. How then does the multi-location organization protect all of its data and do it in near-real time, with close to zero downtime?

The solution to this problem is called “data vaulting.” With vaulting, all of the organization’s data, including that from laptops and PCs, is compressed, encrypted and periodically transmitted to an offsite “vault.” In the event of a data disaster, the organization can recover its data quickly and efficiently without having to reinstall applications or deal with other delays. Even if the home office is completely destroyed or inaccessible, the legal staff can immediately switch to laptops or home PCs and be right back in business.

### Asynchronous Mirroring

With data vaulting to an offsite location, law offices periodically back up their data — daily, hourly or even more frequently. Such periodic backup still leaves the organization vulnerable to a major loss of data — an entire day’s worth if the office backs up its data daily. For many organizations that’s too great a risk. For them, the next step up is “asynchronous mirroring.” When combined with “replication” and “snapshot” technology, its goal is “zero” downtime.

Asynchronous mirroring continuously monitors the law office’s primary servers and creates in near real time, an offsite duplicate of every bit of data at the block level. If the primary servers go down, offsite secondary servers can stand in, enabling the legal staff to access data with minimal disruption and confusion.

### Foiling Rolling Disasters — Say Cheese!

Despite layers of redundancy, there’s always the danger that data corruption will take place over time rather than instantaneously. Unless the backup data system is protected against such a rolling disaster, it too can be corrupted.

The increased pace of information technology and the growing use of technologies that utilize mirroring techniques have made rolling disasters a cause for concern. Fortunately, there is an antidote. It’s called the “data snapshot.” With this technology, a date/time-stamped, point-in-time snapshot is periodically taken of the data, equating to multiple full backups throughout the day without disruption. In the event of a rolling disaster, the sequence of snapshots enables the recovery experts to quickly search back through the snapshots until they find and select an uncompromised version, or last known “good” point in time, of the data. Once the appropriate snapshot is identified, you can either roll back the production server to it or present it to a recovery server to “mine” through and pull out the data you need.

The frequency of snapshots depends on the organization’s tolerance for data loss or RPO. Snapshots can be as often as once a minute, but the average is hourly.

The current gold standard in disaster avoidance is a completely virtual office that utilizes asynchronous mirroring along with point-in-time snapshots to replicate data, applications, paper documents and telecommunications. The result is nearly instantaneous access to data from any location 24/7. Then, even if a disaster destroys the organization’s offices or denies access to them, the staff can continue to function with minimal disruption over secure connections utilizing thin client technology.

The use of redundant data circuits, whether they are in a fully meshed MPLS network, point-to-point private lines or Internet IP, provides an additional layer of data continuity. The secondary bandwidth provider should be one whose network functions independently of the local telephone company.

### Don’t Forget Phones and Paper Files

Securing the law office’s electronic data is just part of the job. Lawyers spend much of their time on the phone, and they frequently have to consult paper documents. To keep telephone communications up and running through a crisis, some law offices are using a virtual PBX (private branch exchange) service that stands in if the local system is unavailable. Such a virtual PBX can instantly replicate all of the features of the organization’s own phone system (voicemail, call forwarding, call “follow me,” which transfers calls to a predefined call list (*i.e.*, cell phone, home phone, etc.) and can also route calls to an Internet phone or notebook, if VoIP (voice over IP) is enabled.

Many law offices store paper documents in files or warehouses, where they are difficult to access and where they can be lost, mislaid, destroyed by fire or ruined by water from sprinklers. Protecting paper files is relatively inexpensive and easy. The office simply scans the documents into a digital format and then automatically indexes them onto online digital storage. Besides protecting the documents, such digital storage makes documents immediately accessible. There’s no need for someone to waste time going to the file room or warehouse to locate them or to suffer the frustration of having to search manually for lost or misplaced documents.

### Good Advice

A truism in IT is that by the time a new development reaches the marketplace, it’s already obsolete. What was leading edge just a few years ago is likely to be commonplace today. This means that if law offices are not alert to advances in technology, technology will pass them by. To make sure that they get the best technology for their budget, law firms and legal departments should periodically review their systems, bringing in outside experts when necessary. After all, “an ounce of prevention is worth a pound of cure.”

# Create a Competitive Edge for Your Firm

by Ensuring 24/7 Operations

by Mark Boltz of Stonesoft Corporation



**T**oday's law firms face many challenges to business continuity. From a network security perspective alone, many forms of threats exist that can lead to disruption in communication, loss of productivity, loss of client confidence and diminution of the firm's reputation. In order to be competitive in today's world, firms must be able to ensure both network security and business continuity and provide the safe, secure, confidential sharing of data at all hours of the day or night. Two of the most difficult challenges to business continuity in the area of network security are the availability of the firewall/VPN gateways and the Internet Service Provider (ISP).

## A Need for the Always-On, Secure Network

According to a recent report by Infonetics Research, large companies lose two to 16 percent of their annual revenue due to network downtime. Service providers can account for up to 30 percent of the downtime costs. In the legal profession the failure of a firewall/VPN gateway or an Internet provider can mean delays in filings, ineffective collaboration with partners and co-counsel, e-mail delays and the inability to access crucial online research, such as Lexis-Nexis.

In an industry that relies heavily on reputation and word-of-mouth endorsement, a failure or breach of security or continuity leads to credibility loss, violation of client privilege and other severe consequences. Additionally, today's cases frequently require the establishment of remote offices for onsite counsel, sometimes for weeks or months on end. These satellite offices require secure, uninterrupted connectivity to the firm's other locations.

Even though the importance of reliable, secure continuity in computer networks is growing, many firms still rely on outdated technologies and single points of failure for their connectivity. Firewalls are rarely made redundant, either as an oversight or because the overworked IT staff is not aware of new advances in high availability. Frequently VPNs are not deployed, because the perceived unreliability of the Internet connectivity has pushed organizations towards more costly frame relay circuits and other leased line options. Today, better ways exist to meld the continuity of the network security components; and firms that deploy these new technologies often gain competitive advantage.

## The Importance of Firewall/VPN Redundancy

Firewalls are a control point between two or more networks, controlling traffic to and from each place. Whether the firewall is a software-based solution or a hardware appliance, a failure of the firewall means the disruption of all communications. Mission-critical business processes come to a standstill when the firewall is down. Whether the outage is planned, such as an early morning service window to perform an upgrade, or unplanned, such as the failure of a hard drive, the overall cost to a firm can be massive. Even small firms can expect a minimum of six hours of unplanned downtime a year, in addition to the planned downtime for maintenance operations.

VPN gateways, often contained in the same device as the firewall, are often used for communications between networks. VPNs, or virtual private networks, allow for the secured extension of the firm's network to remote offices over the Internet through encryption technology. The failure of the VPN gateway, through either planned or unplanned downtime, also leads to disruption of business continuity. With the VPN down, offices can no longer exchange information, communicate via e-mail, share files or access research among many other important tasks. Many organizations opt for frame relay or other leased lines with specific service level agreements instead of VPNs to avoid this problem. The dedicated communication lines allow them to connect offices, for a price. Firms that have addressed the continuity of their network security are more competitive, better able to share knowledge and data globally while ensuring the confidentiality and integrity of their clients' information.

### Methods of Firewall/VPN Availability

Multiple methods now exist to achieve firewall/VPN availability: hot standby, active-active pairs and clustering.

#### Hot Standby Firewall/VPN Gateways

The most basic solution to gateway availability is to have a second box as a "hot standby," ready to intercede at a moment's notice. Cisco's PIX firewall operates in such a configuration, where the standby firewall is linked to the primary via a serial cable. Configuration changes are synchronized between the two devices, and if something happens to the primary, the secondary will take over. For stateful inspection firewalls such as the PIX, most traffic should failover seamlessly, unnoticed by the end user. Organizations using the slower, yet more secure, application proxy firewalls, such as those made by Gauntlet or Raptor, can also have hot standby configurations, but the transparency of the failover is lost.

Although this active-passive approach improves the situation, the second device sits idle most of the time. Additionally, the solution doesn't provide for scalability. If additional performance or features are required in the future, both boxes must be replaced entirely by new equipment. Because the hot-standby approach allows for only two boxes, additional boxes cannot be added to accommodate growth later. Though the redundancy solves the continuity problem, basic hot-standby configurations also add additional layers of complexity. Both boxes must be configured identically, with each change now being entered twice instead of just once.

#### Active-Active Firewall/VPN Gateways

An improved approach is an active-active configuration, such as the type used in most Juniper NetScreen firewalls.

Juniper's NetScreens can have two devices share the load at the time. In this configuration they are often still connected by a dedicated serial cable, though using network links for synchronizing data between the boxes is becoming more common. If one of the two devices fails or needs to be taken out for maintenance, the work of that gateway is transferred to the operational one. As with the hot-standby configuration, traffic should failover transparently. With an active-active configuration, the investment in hardware and any necessary licenses is put to better use.

Active-active configurations are still limited, however. If one of the devices is down for maintenance, for example, the other is once again a single point of failure. Both devices must typically be identical in hardware power, memory, configuration, software version, etc. And both must be configured so that they will operate with peak loads below 50 percent of the machine capacity, otherwise a failover will create an overload condition in the remaining box. As with the hot-standby solution, both devices must be configured identically, with the same routing and network information, same security policy and other data. An active-active solution is also not scalable, as new machines cannot be added to the group. The only option for upgrading in the future is to replace the existing equipment with new gear.

#### Clustering Firewall/VPN Gateways

While active-active meets the needs of many organizations, this method of redundancy is not scalable. As the needs of the firm grow or change, the system is inflexible and will most likely need to be replaced with newer hardware or a completely new system, which can often be a disruptive and costly process. For these reasons, clustering of the gateways becomes the most appropriate answer for most firms. Check Point's VPN-1 and Stonesoft's StoneGate both support clustering the firewall and VPN gateway — Check Point up to four and Stonesoft up to 16 machines. In either case, the firewall/VPN gateway still looks like a single device to those on the network, but in reality the traffic is load balanced across all of the available devices. The failure or maintenance of a device causes the transparent redistribution of the load to the remaining available machines. With three or more machines in the cluster, downtime of a node no longer means reducing the firewall/VPN gateway to a single point of failure. With clustering technology, the system fully utilizes the investment in hardware and licenses, while ensuring a scalable system. Additionally, new nodes can be added to the cluster without disrupting or altering the existing configuration. The cluster usually communicates with a "heartbeat" between the nodes, often on a dedicated network. Backup heartbeats ensure that the failure of a network card or switch doesn't disrupt the communication between the nodes for the remaining networks.

## ISP Redundancy

A reliable link to the Internet is critical to issues of business continuity and security. Even with full clustering of the firewall/VPN devices, the failure of the Internet connection is disruptive. From basic problems like the hungry “copper seeking” backhoes at the nearby construction site, to misconfigurations of equipment, to distributed denial of service attacks, access to the Internet can fail. DSL and cable modems are inexpensive bandwidth but are difficult to get service level agreements (SLA) for — even with an SLA they are not as reliable as the more expensive frame circuits many have come to rely on.

Remote offices often connect with frame relay and are limited in what services can be provided to them. High-bandwidth items like voice over IP, centralized virus scanning and centralized databases are not possible. Distributed architectures with each office receiving its own copy of databases, virus scanners, URL filters, etc. create enormous complexity for the IT staff and inevitably lead to failures in continuity. These problems can be solved with reliable, redundant Internet connections combined with redundant firewall/VPN gateways.

Today, firms have several options to achieve ISP redundancy. Each of these can improve connectivity, potentially allowing expensive leased lines or frame relay circuits to be dropped in favor of Internet connectivity.

### Border Gateway Protocol (BGP)

The most common approach to redundant links is BGP. This protocol allows routers to share information about traffic between sites. BGP requires an organization to attain an autonomous system number (ASN). A peering arrangement must be established between each ISP to allow the BGP information to flow between sites. Due to the increase in routing information, border routers must be robust systems with lots of memory. Additionally, the firm should have a BGP expert on the IT staff to program and troubleshoot any problems that may arise. Although BGP is robust and common for high availability, it is often limited to the headquarters and one or two other key sites in an organization because of its complexity. It also tends to scale poorly, because additional ISPs complicate the peering arrangements and the routing tables. BGP is typically not an option for low-cost lines like DSL or cable-modems. For firms with the staff and equipment, it may be an option.

### External Load Balancers

Another approach is to use external load balancers. These are often dedicated equipment that sits outside the firewall/VPN

gateways and redirect traffic through the appropriate ISP lines. The load balancers also need to be redundant, often in an active-passive configuration; otherwise the load balancer itself is a single point of failure. Again, due to cost and complexity, companies using external load balancers tend to use them only for the main site and key locations, so the benefits do not extend to the entire organization. External load balancers also require additional IT skills, as well as maintenance and support agreements, and added management software.

### Multi-Link Technology

The third option is to use multi-homing or multi-link technology in the firewall/VPN gateway itself. With this option, the firewall/VPN cluster load balances traffic among all the available ISPs, regardless of their type. It is not necessary for the ISPs to be aware of each other or for additional hardware and software to be used. In other words, it is a simple method to achieve very robust Internet reliability. With this technology, VPN failover is transparent, and the bandwidth of the ISPs is aggregated. Multi-link or multi-homing technology enables more robust IP services at branch offices and other remote locations, including voice over IP (VoIP), video conferencing and more.

For example, a firm can set up a remote, temporary “office” to serve a client, providing reliable but inexpensive connectivity to the firm’s headquarters. Each attorney at the remote office can still securely use his same phone extension through VoIP and access all the other tools, including research databases and collaboration, with coworkers.

Increased reliability of the ISPs, combined with the transparent VPN failover capability of multi-link, also allows firms to displace costly leased lines or frame circuits. T1s, DSL, cable modems and other data services can be combined to create an effective, reliable link even if the individual lines are not.

### Prepare for the Unexpected

Although firewall and VPN technologies provide for good security, they can be a bottleneck and a problem for business continuity. ISPs are also a point of potential failure when the unpredictable happens. But by using available technology in network security products, it is possible to prepare for the unexpected and create a robust, redundant system. Firms should explore the different high availability options in the market today. Additionally, technologies like BGP and multi-link increase the reliability of the Internet connections. The increased reliability and bandwidth enables firms to provide more robust services to remote locations, increasing the competitive edge of firms who leverage these technologies.

# Data Replication: A Key Element in BCP

by Jason Buffington of NSI Software



**T**he need for business continuity planning is as vital in the legal profession as it is in any major industry — in some ways, more so. Traditional backup systems, commonplace in many firms, do not offer the level of data protection that is critical to professional productivity and practice continuity. This article will explore the key aspects of good planning and good backups.

## Business Continuity Planning — a Must-Have for Law Firms

Many benefits accrue from business continuity planning, including high availability of systems, good disaster recovery procedures and better backup systems — all of which mitigate unplanned outages. Additionally, Sarbanes-Oxley and other compliance regulations have impacted data and record retention for many business sectors served by lawyers, including financial (SEC & FDIC), government (DOD 5015.2 and CO-OP) and healthcare (HIPAA and JCAHO).

Two notable factors make business continuity planning in the legal segment especially important:

*The average legal practice has the highest percentage of revenue generating workers per company of any major industry in the world. Whereas in a typical company, one might have up to 30 percent in a sales or revenue generating capacity (with the remainder in manufacturing, support and operational roles), the average law practice comprises a broad range of the billable resources, supported by a disproportionately smaller administrative, operational and support structure. This results in an even greater need for high productivity and guaranteed uptime due to the hourly billable nature of the legal workplace. Arguably, a small law firm needs business continuity planning to the same degree that a large international legal practice does. And the need is more prevalent in the legal profession than in similarly-sized organizations across any other industry.*

*The second characteristic of the legal profession exacerbating the need for business continuity stems from attorney-client privilege. Because of the level of disclosure that clients must give to their own attorneys as part of most major litigations, clients are often forced to trust significant amounts of their own data to their attorneys. By extension, this means that continuity regulations around the viability and security of data arguably transfers to the practice representing those organizations.*

Imagine a situation in which a client company has survived a regional crisis through its own business continuity preparations, only to find that the law practice representing them has gone out of business due to loss of client information, billing records or even key discovery contents. A law firm's IT department should focus on the three core aspects of business continuity planning:

*High Availability — Protecting Productivity*

*Disaster Recovery — Protecting the Practice*

*Better Backups — Protecting the Data*

## High Availability — Protecting Productivity

In business continuity planning, one of the primary goals most often pursued is that of ensuring uninterrupted productivity.

It's rather obvious that a large firm requires significant uptime for the hundreds of users who depend on the firm's IT resources — what may not be as clear is the small firm's equal dependence on its data. Good data protection is important to firms or law departments of any size.

### Small Firm, Large Need for Protection

Consider a small practice, a 25-person office (10 attorneys and 15 support personnel). If the production server goes down in the middle of the day, data loss for both the first part of the day and productivity loss for the entire day will be incurred. We can measure productivity loss as the entire amount of manpower, which either is incapacitated or will need to be reapplied while repeating a task. In this case, we have data loss, plus a productivity loss for the time it takes to replicate lost work, plus the hours where users cannot access their data and may be completely idle.

Using industry statistics, this 25-person office has a manpower cost of \$1,500 per hour (considering average salaries and benefits for partners, lawyers, paralegals, and administrative staff positions). Assuming half a day of data loss and a full day of productivity loss, a single one-day outage will cost this small organization \$23,000 (exclusive of any lost revenue). If we include the average billing rate (assume 60 percent of revenue generators' hours are billable and multiply that by industry-average rates), the firm will lose an additional \$26,000 of revenue. While this small office may have considered itself "too small" for business continuity planning, a single outage per year carries a not-so-small pricetag of \$49,000.

Downtime happens, inevitably. Business continuity planning is about reducing its costs. If \$7,500 can be spent to mitigate a \$49,000 loss and any future \$49,000 losses for the entire time that the business continuity technology is in place, there is a huge return on investment even for the smallest firm.

### Economies of Scale

In a large legal practice, HR costs tend to scale linearly, while billing costs increase at higher rate due to productivity gains from leveraging advanced tools and often higher billable rates. E-discovery technologies, document imaging, online research libraries — all these things improve the productivity of the large practice. But they make the practice even more dependent on its systems, which therefore causes an even higher loss of productivity during any kind of service disruption.

As a final example, if this outage happened to a medium-sized practice of 100 employees, the single outage would cost \$135,000 in manpower plus \$184,000 in lost billing.

For a 250-300 user law firm, the greatest business continuity requirement may not necessarily be about immediate

resumption of activity (high availability described above) as much as about the survivability of the company itself. By definition, an SMB (small or medium business) comprises up to 500 employees. Various surveys, including Gartner, indicate that a medium-sized business (500 users or less) has a 50 percent chance of going out of business after a disaster if they cannot gain access to its data within the first 24 hours following the crisis. There are many factors that contribute to this grim prediction, but a notable point is this: if data cannot be accessed within the first 24 hours after a crisis, it's highly likely that the company cannot begin the rest of its business recovery efforts fast enough to avoid eventually going out of business.

### Disaster Recovery — Protecting the Practice

Ironically, it wasn't so long ago that the terms "disaster recovery" and "business continuity" were interchangeable. Today, most industry professionals recognize that business continuity is a broader strategy around uptime, data protection and crisis resilience. We'll focus on the last of these goals, which is still referred to as "disaster recovery." Simply put, disaster recovery means "get the data out of the building" and then plan for how you will access and utilize it.

### Insure Yourself

Many disaster recovery plans never get off the ground due to lack of executive sponsorship, which is reflected in an annually deferred budget. Put another way, disaster recovery becomes "something we'll try to do next year." In our earlier example of high-availability, we saw the financial implications of just a single day of downtime. One could extrapolate from those same numbers, stretched over several days but now possibly without the hope of recovering the data, what the cost of a disaster might be. Chances are, that those costs are not in the budget either.

Disaster recovery planning costs ought to be viewed like any type of insurance. Everyone knows that they need it; no one is necessarily excited about spending for it — but the entire motivation is to spend pennies now instead of dollars later. For disaster recovery, a business needs to understand the financial ramifications of a crisis. After looking at the big picture, a firm can typically justify the expense of preparing for a disaster. The advantage of spending for business continuity planning over insurance is that there are benefits that are reaped now, not just after a crisis.

### Don't Let Documentation Delay Your Plan

Firms can become myopically focused on grandiose "disaster recovery plans" or binders full of documentation. And while full-scale DR plans should include documentation, the fault lies in delaying implementation pending the completion of grandiose documentation.

## Better Backups — Protecting the Data

### Why Isn't "Tape" Good Enough?

Imagine having to go to your senior partners or management committee to report that a server crashed in the late afternoon. All of the data for the attorneys, paralegals and some information provided by clientele were all lost. In addition, imagine notifying them that the same server will be down for most of tomorrow while the pieces are being ordered and the server repaired.

In traditional data backup, one should be prepared for the fact that the organization will experience at least half a day of data loss and one day of downtime in what is typically a "best-case scenario." This is not specific to the legal sector, but rather, specific to tape backup in general. If data is backed up nightly, data loss will be measured in "days." In addition, if spare hardware resources aren't readily available, most of the next business day following a crisis will be spent getting the parts to repair the downed server. Many times, an additional day of data loss may occur due to the likelihood that 30 percent of all tapes are not restorable.

And while the above tape scenarios are applicable to all tape backup environments, the effects are more strongly felt in the legal community than in many other industries because of the legal community's dependence on hourly productivity. Law firms and legal departments must protect their data more often than nightly. This takes us from a nightly tape process and into the realm of continuous data protection or real-time replication.

### Real-Time Data Replication

With this new scenario, as data is changed, the changes are being transmitted from the production server(s) to redundant server(s) at alternate sites. Instead of having tape-protected data from last night, replicated data on the target server is a virtual twin of that on the production server.

### A Boon for the Large Firm

A large multi-location law practice is much like a bank with many branches, a chain of retail stores or any other enterprise with distributed offices. Remote offices tend not to have dedicated IT personnel. While this is obviously notable during a crisis, it also requires a routine burden for remote office personnel to manage system backups. Perhaps an administrative assistant or office manager is tasked with swapping tapes, cleaning cartridges and even validating last night's backup. This process comes with an appreciable manpower cost and introduces the possibility of human error into one's data protection strategy.

One reason that many firms still use this approach is because bandwidth is not available (or cost-effective) to back up the

remote offices to the core data center. However, low bandwidth lines can be used to replicate the data from the remote sites to the IT data center because only the byte level changes are transmitted.

### Advantages of Data Replication

From a budget perspective, no other data protection technology is as cost-efficient. Leveraging host-based replication, one need only put a simple software license on each production server and then point it at an offsite location.

Replication will transparently, automatically and without routine manual intervention, send the data to a remote site, which by definition is the beginning of disaster preparedness.

Instead of each site handling its own tape backups (and changing tapes, cleaning cartridges, monitoring jobs, etc.), all of the data on the production servers have a consistent copy at the core data facility, and the centralized IT team can perform centralized backups. Tape backups can be done during the day from the replicated copy of data.

Backups of the remote offices can occur, even though the remote office production data is actually still in use on their real servers. This results in fewer backup jobs to manage and maintain regardless of the size of the environment(s).

### In Summary

A major difference in the business continuity needs of law firms from those other types of business is their notably higher dependence on hourly productivity and data protection. This is compounded by the regulatory requirements that many law practices must comply with as part of supporting their clientele in various specific industries such as healthcare, financial and government.

Unlike other industries in which only larger organizations place a priority on business continuity, we have seen how even the smallest of practices have significant uptime requirements and data dependencies. And of course, as the firms get larger, the needs are further exacerbated. Traditional methods of data protection simply do not fit the vast majority of law firms. Because of the manpower associated with the creation of data, its loss is intolerable; and nightly tape backups simply do not provide sufficient protection.

Data replication technology may be the solution that addresses your firm's needs to protect its data, its productivity and therefore, its practice.

# Never Failover on a Monday:

## Overcoming the Five Fatal Flaws of Traditional Replication Solutions



As data and applications have become critical to every law firm's ability to serve clients and remain in business, firms of all sizes are investing more heavily to protect these assets. Replication solutions protect data and enable and deliver high availability to mission-critical applications, but the solutions are vulnerable to failure. We'll examine the inherent flaws and provide insight on how to avoid them.

### Protecting Critical Applications: A Growing Need Within Law Firms

Law firm technology departments have experienced tremendous growth in the past 20 years. This is due in part to the advent of the Web browser, which has enabled IT solutions to evolve from proprietary, back-office systems to Web-enabled, front-office applications. Productivity has surged, and the number of users within a firm relying on mission-critical applications has exploded. The data that powers these applications has become the lifeblood of all law firms.

With the growing dependency on mission-critical applications and the need for data integrity comes the challenge of how to protect them. New threats such as viruses, data corruption and SAN failures have materialized. While there are many options for protecting data — tape backups, disk mirroring, electronic vaulting, etc. — data replication is increasingly common for ensuring availability of critical applications and data.

### Replication 101: Solution Overview **rep·li·ca·tion** *n.*

Simply defined, data replication is creating a copy of data. That copy becomes insurance against anything happening to the primary set of data. In the event of an emergency such as hardware failure or data corruption, the backup data can be

by Russell Sachs of MessageOne, Inc.

used to recover important information, ensure data integrity and/or continue business operations.

### Data Protection vs. High Availability

When evaluating replication options, it is important to understand the specific reasons you need or want a solution. Two primary needs benefit from replication solutions: data protection and high availability.

Data protection is primarily focused on data integrity, ensuring that the data and all changes to it are captured in a backup copy. Immediate access and continued use of the data are not top priorities. High availability, on the other hand, is concerned with three attributes:

*Data integrity*

*Immediate access*

*Continued application use*

High availability solutions must hedge against regional disasters and outages. This requires the deployment of the high availability replication solution in a geographically remote backup environment, ideally more than 20 miles away. Distance adds the complexity of bandwidth latency and cost constraints, factors that are negligible over a LAN.

**Using replication to deliver high availability is more complex than data protection.**

### Using Replication to Provide High Availability

Understanding replication and its prominent place in delivering high availability for mission-critical applications requires a deeper look at the key elements that enable high availability, particularly across a geographically diverse or complex IT environment.

**Data Replication** — The set of processes that collect and move data from one source to a second source. Data replication includes a process to watch a data repository for new and changed data, which it then duplicates and places in a replication queue. The replication queue stores and manages the pending data transactions. These are replicated to the secondary data repository. A process monitors the transaction to ensure its safe arrival at the destination. The backup data provides a source safe from many issues that can affect the primary data source. This component provides sufficient coverage to protect data sources when the data is not supplied to applications that require near continuous uptime.

**Failover** — The set of processes that switch a secondary data source from its backup role to a primary role. Often, this is tightly integrated with other processes that redirect users to an alternative set of application servers. Depending on the application, further security processes and control logic are required. This assumes that all systems are synchronized, all software versions and revision levels are consistent, the systems have been tested and personnel are standing by at both data centers. Failover can then proceed smoothly, during which time users have near-continuous access to their critical applications with minimal interruption.

**Failback** — Returning operational control back to the original system. In many instances, a company’s primary environment is more robust or better situated to deliver superior end user performance, benefiting from greater bandwidth or a concentration of local users. For these performance or security reasons, a systems administrator will seek to failback to the primary system as soon as possible once the primary system has been restored to a fully functional state. Failback is similar in many aspects to failover, but adds additional steps and complexity to account for resynchronizing the two systems. This involves recovering the primary system to a working state and reestablishing data integrity, as the two data source have become out of synch since the initial failover.

**Replication is easy; failover and failback are difficult and risky.**

## Replication Shortcomings

Though replication solutions offer promises of risk mitigation, higher availability and data protection, the attempted use of these solutions can result in frustration and wasted time and money.

Installation alone can be a deterrent. Firms frequently have such numerous issues associated with installation or management of a complex solution that they abort their implementation and “shelve” the replication software. Once installed, many firms find replication manageable. But then a bigger challenge may be encountered when attempts are made to activate replication solutions to maintain application availability and recover in the event of a large-scale systems failure or disaster. Failover processes are difficult to test and don’t always work as expected during tests. Accounting for all permutations of failover requires extensive process mapping across the entire IT environment that supports the application. The expertise required to perform this mapping rarely exists within the company. Failure to adequately address all possible failover scenarios creates risk and possibly a false sense of security.

Even when failover does work, reverse replication and failback are extremely difficult. Because of the challenges

associated with failover and failback, many organizations will failover only at the end of the week, when they can use the weekend for damage control and restoration of production. This lack of confidence means that many system problems will result in unnecessary downtime, as administrators will choose downtime over complexity and untrusted systems.

**Replication and failover complexities can prevent many firms from fully realizing the true benefits.**

## Top Five Causes of Replication Failure

An examination of the following causes of replication failure provides both insight into the challenges a firm faces when seeking to successfully implement a replication solution and a useful checklist to benchmark your efforts.

### 1. Secondary Environment Not Ready for Failover

A firm establishing a secondary backup of its most important applications and data must address many challenges, one of which is maintaining two nearly identical environments. All software, patches and access levels need to be consistent. Over time, many changes that are made to one system fail to get implemented on the other system. These differences often result in a secondary environment that is out of synch with the primary environment to the point where failover cannot occur.

Another factor impacting failover readiness in the secondary environment can be that critical processes are not functioning properly. Most organizations do not have clear visibility into the readiness of their systems for failover and are surprised by one of the following problems:

*Replication is not performing normally.*

*The replication queue is too large.*

*The secondary environment is not healthy.*

*Primary/secondary environment software and configurations are out of sync.*

*Dependent systems are not designed for failover.*

It is critical to take the time necessary to develop processes to control the introduction and distribution of changes and updates to both environments. One undetected change can cause delays in service resumption spanning hours or days. It is equally important to monitor all critical processes that impact the readiness of the secondary environment, as early problem detection insures system readiness when a situation demands use of the secondary system.

### 2. Manual Error in Failover Process

The weak link in a failover process is often people. Manual errors introduced during a failover sequence can corrupt the entire failover. The more complex and step-intensive the

failover process (for example, it takes 350 steps to failover a 10-server Microsoft Exchange environment), the more likely mistakes will occur that may result in failure. Examples of manual problems include:

*Missed process steps*

*Steps executed out of sequence*

*Process initiated before dependent steps fully executed*

*Misjudgment of state of primary or secondary environment*

*Typing error(s)*

*Steps out-of-date for current software versions*

### 3. Experts Not Available During Crisis

Failover processes are dependent on technical experts who may not be available during a crisis. The failover process touches upon all technical disciplines — from hardware and operating systems, to applications and databases, to networking and security. In a large organization, these disciplines are highly specialized with different personnel responsible for each. If even one person is unavailable — because, for example, he is occupied with other crisis efforts or has no Internet access — the process can break down.

As part of the process to develop and deploy a failover solution, it is important to establish a list of skills and resources available for each skill. Full contact information and predetermined communications protocols need to be created, continuously updated and readily available to all team members. These steps will aid recovery efforts and mitigate some personnel risks.

### 4. Failover Process Unable to Scale

For large firms, the scale of a failover or recovery effort can become a critical bottleneck. The technical staff is limited, often managing large numbers of systems. A critical application failure or a facility problem can result in dozens or more systems that require failover. Multi-server failover is a serial, manual process. The technical staff comes under tremendous pressure to rapidly restore service. Under these conditions, it's easy to identify scaling issues such as:

*One administrator can only failover one server at a time.*

*A 25-server environment will take 10+ hours for one administrator to failover or fallback.*

*Many mutually dependent systems will not work until the entire environment has failed over.*

### 5. Untested Failover Assumptions Don't Work

Complex multi-server failover is often too sensitive and complex a process to fully test. Therefore, it's impossible to know exactly what will happen during a real crisis. Failover breakdown issues include:

*Large, complex environments have many failure types and scenarios.*

*Multi-server failover involves many moving parts.*

*Server-by-server failover is very different from a holistic failover of the entire environment.*

*Different failures result in different behaviors.*

Incorporating “what-if” scenario-planning sessions and “pre-mortems” is one way to mitigate risk. These role-playing efforts allow a technical staff to identify untested failover scenarios and potential bottlenecks.

### The Impact of Replication Failure

Understanding the impact of replication failure is essential to making informed business decisions about replication investments. The following provides a summary of the key issues associated with replication failure.

**Replication adds complexity.** Developing a replication solution requires doubling the amount of hardware and software while requiring additional bandwidth to handle replication traffic to the secondary environment. If the secondary becomes the primary, more bandwidth is needed to deliver service to end-users.

The additional equipment and systems bring additional demands on the administrative staff. The new environment significantly increases complexity, providing further challenges for the staff.

**Replication failures have long recovery times.** Restoring an entire environment requires significant skills and complex recovery processes take time to implement. The availability of critical resources and personnel becomes a bottleneck. Often, there are periodic delays during the failover process as key skills are not available. If large amounts of data must be moved as part of a recovery effort, bandwidth constraints can further prolong the recovery time.

**Replication failures can cause other problems.** Replication failures can have widespread impact. Database corruption can occur, requiring substantial efforts and time to restore a database to a functional state. Even when the databases are not corrupted, effort is required to determine how much data was lost during the failure and to recover the missing transactions.

*Replication failures have a high cost.* Replication is almost exclusively used for critical systems. Data in these systems is important enough to protect and therefore recover. The time and effort required to recover the data, returning it to a usable state, is substantial. Replication failure forces application downtime, resulting in negative economic consequences: lost revenue, missed opportunities, degraded client satisfaction and possible loss of business. Any replication failure will prove costly.

## To Succeed, Eliminate the Risks

The key to successfully implementing a replication solution is to understand all the risks and eliminate as many as possible. Start with the risks identified above in the top five reasons for failure. Through careful evaluation of approaches, you can select a solution that is the best approach for your company.

### Plan a Parallel Approach

Large firms have complex computing environments and little tolerance for downtime or lost data. When investigating the top causes for replication failure, manual recovery of a large environment can prove to be a bottleneck. An administrator can only address problems serially, requiring lengthy recovery periods to return an entire enterprise to full operations. As the recovery process progresses, fatigue further increases human errors and downtime. Firms should seek solutions that allow the overall recovery process to be conducted in parallel rather than in more traditional serial approaches.

### Monitor the Replication Environment

Successful replication and failover starts with a solid understanding of the environment and immediate knowledge of problems so that minor issues get fixed before they turn into major problems. Firms should seek to deploy an integrated and comprehensive monitoring solution that watches all replication processes and the secondary environment, ensuring that the secondary environment is always ready to assume the role of primary if the need should arise. As complexity and human error introduced during a crisis are responsible for many failover failures, a monitoring solution can simplify complex information on a single console, providing status and confidence in failover readiness.

### Automate the Solution

Human performance is strongly and negatively impacted by crisis situations. People respond and act more slowly and with more errors. These issues arise because a crisis situation is rarely encountered and often requires unusual steps to remedy. The crisis creates pressure, which affects human judgment. This introduces errors that can ultimately prevent a successful failover and resumption of normal operations.

Business process automation of the entire process is the only way to eliminate the human risk and dependency, as well as speed the time-to-failover. Automation of the complex sequence of steps necessary to successfully failover an application environment is the key to mitigating the risk.

Automated application failover and failback capabilities not only eliminate the risks associated with these actions, but enable on-demand testing of the system. This allows a firm to build confidence and trust in their solution, while identifying minor system discrepancies before they create recovery problems.

### Embed World-Class Expertise

Automating the entire process only solves the problems that have known solutions. Accounting for all possible failover scenarios requires assembling a team of experts from all technical disciplines including applications, network, hardware, operating system and other systems experts — and then systematically creating a massive decision tree to capture all “what if” conditions. This requires tremendous teamwork and access to acknowledged world experts, people rarely found under one roof.

Law firm IT departments should look for automated solutions that have effectively assembled these experts and have productized their knowledge into a fully automated system, representing the combined expertise of the best of every technical discipline across the entire IT infrastructure.

### Seek a Service Model

Realizing that even with the robust monitoring, management, automation and embedded expertise, many firms will not want to manage these types of solutions themselves. Firms should begin to look at an emerging class of service delivery models, which provide centralized management of the monitoring and failover infrastructure. A number of qualified vendors possess experienced operations staff that can compliment in-house experts a firm may have.

## Summary

Replication solutions represent a powerful way to protect your firm’s most important data and applications. But replication solutions, particularly those that are used to provide failover and failback for high availability environments, are fraught with risks — risks that have resulted in a relatively low success rate for the use of replication solutions in law firms. Increasingly, law firms are seeking enterprise-class solutions that deliver automated replication, failover and failback for their most mission-critical applications and data — solutions that are designed to eliminate the traditional risks that have plagued these types of solutions.

# About the Authors

---

**Mark Boltz** is a senior security consultant with Stonesoft Corporation, an innovator in business continuity and network security solutions. He has over eight years of experience in network security and over 14 years in network and system administration. Mark has been cited by various magazines and given presentations at events such as SANS, RSA Conference and SHARE. He can be reached at [Mark.Boltz@stonesoft.com](mailto:Mark.Boltz@stonesoft.com).

**Jason Buffington** is the Director of Business Continuity for NSI Software, which enables high availability and disaster recovery via replication software. Jason has worked in the networking industry since 1989, with a majority of that time being focused on data protection. He is a Certified Business Continuity Planner and a Microsoft MVP in Storage. Jason can be reached at [jbuffington@nsisoftware.com](mailto:jbuffington@nsisoftware.com).

**Gregory Hanna** is President and CEO of TOSS Corporation ([www.DisasterAvoidance.com](http://www.DisasterAvoidance.com)), which has been providing security, network and business continuity services since 1992. He can be reached at 1.888.884-TOSS (8677) or [alData@DisasterAvoidance.com](mailto:alData@DisasterAvoidance.com).

**Atlas Lee** is Director of Business Continuity for Shook, Hardy & Bacon L.L.P. in Kansas City, Missouri. He has been with the firm for 17 years and has 21 years' experience in the information technology field. Atlas is a frequent speaker on business continuity planning, disaster preparedness/recovery and information systems security, and he has authored several white papers on those subjects. He can be reached at [ALEE@shb.com](mailto:ALEE@shb.com).

**Michael Oh** is the Founder and CEO of Heavy Water Ltd., based in New York, and has been in the technology field for more than 25 years. He founded the company in 1992 and has specialized in the infrastructure engineering, architecture and security. Michael is widely recognized in the New York tri-state area as a leading expert in the field. He can be reached at [moh@heavywaterltd.com](mailto:moh@heavywaterltd.com).

**Russell Sachs** is Co-Founder and Vice President, Legal Solutions, of MessageOne, Inc. Previously, he practiced law for many years at a major New York City-based litigation firm. Russell can be contacted at 212.812.5017 or [russell\\_sachs@messageone.com](mailto:russell_sachs@messageone.com).

---

## Disclaimer

This report is designed for use as a general guide and is not intended to serve as a recommendation or to replace the advice of experienced professionals. If expert assistance is desired, the services of a competent professional should be sought. Neither ILTA nor any author or contributor shall have liability for any person's reliance on the content of or any errors or omissions in this publication.

## Copyright Notice

Copyright © ILTA 2005. All rights reserved. Printed in the United States of America. No part of this report may be reproduced in any manner or medium whatsoever without the prior written permission of ILTA. Published by ILTA, c/o Editor, 2110 Slaughter Lane, #115, PMB 149, Austin, TX 78748.



2110 Slaughter Lane, #115  
PMB 149  
Austin, Texas 78748

• Address Service Requested

PRST FIRST CLASS  
MAIL  
U.S. POSTAGE PAID  
PERMIT NO. 1149  
AUSTIN, TEXAS

[www.iltanet.org](http://www.iltanet.org)