PurpleRage
DYNAMIC BUSINESS SOFTWARE

PeerSync™

Protecting files with synchronisation by combining high availability with disaster tolerance.

Today's businesses have come to rely on data files stored on both server and desktop environments. Those data files often support the core competencies of a business and are the key to conducting operations. Regardless of whether the files are spreadsheets, documents, e-mail or contain other informational elements, many businesses cannot operate without ready access to the data in those files. File synchronisation offers the ability to preserve multiple copies of files across various IT resources; helping to keep up-to-date files available to users under most any circumstance. Implementing file synchronisation can take many forms and can address several shortcomings found with other data protection schemes.

In this paper we define the challenges faced by many businesses and demonstrate how file synchronisation can cost-effectively solve those problems. More importantly, we will take a look at real-world solutions to real-world problems surrounding data protection and availability.

Introduction:

Over the last few years, reliance upon Information Technology (IT) based solutions has increased exponentially. Now businesses of all sizes count on data availability to successfully complete the workday. Of course, different businesses have different requirements. For example, a financial services firm may rely mainly on spreadsheets, while a law office may rely on documents. However, all businesses need to have access to their critical files or they will be severely crippled. Operational differences aside, the common denominator still comes down to keeping data available under most any situation to those who need it, ranging from the remote worker, to the travelling salesman, to the deskbound office worker. All of those employees still have the same need: uninterrupted access to the data necessary to perform their jobs. Therefore, file synchronisation technology can improve productivity and help to reduce the burden placed upon busy IT departments, while still keeping costs under control.

Today's IT professionals face numerous challenges. The complexity of supported data prevents them from just taking a 10,000-foot view of the problem. To guarantee success, IT professionals need to break down data support into manageable elements. This document will cover common data availability situations and can be used as a template for understanding typical business process problems and identifying solutions using a data synchronisation methodology. Research has shown that the vast majority of businesses are concerned with the maximising the data protection and availability under these situations:

- Protecting Data: The Pros and Cons of Various Data Protection Solutions.
- High Availability: The Challenge Associated with High Availability Solutions.
- Geographic Distribution: Protecting Files by Extending Locations.
- Staging Data: Moving Files from Test Environments to Production Systems.
- Data Portability: Keeping Data Files Fresh for the Mobile User.
- Data Synchronisation: Multiple Techniques for File Synchronisation.
- Extending ROI of Synchronisation: Adding Value to the Process.
- Disaster Recovery Solutions: Employing Technology to Improve Business Continuity.

IT professionals will find that combining applicable knowledge with best practices and available solutions will pave the way for resolving the situations created by providing file access in today's complicated IT environments. In many cases, a simple to deploy solution can solve even the most complex of problems.

Protecting data:
The pros and cons of various data protection solutions

Since the dawn of computing, backup has been an essential element of a properly managed computing system. While technologies have evolved, backup has always been a step behind on the evolutionary scale, with capacities and speeds lagging behind modern requirements. Backups exist for a specific reason; properly executed backups allow users to retrieve their data after some external event causes data loss. Those events can be caused by equipment failures, virus attacks, intrusions, or just plain user ignorance.

Traditional backup technology has revolved around event-driven software applications that copy critical files to another storage medium, such as tape or other magnet media, for later data retrieval. While that methodology offers some protection, the very nature of a scheduled or event driven backup means that data can quickly become stale.

Most sites perform backups on a daily basis, which means in the case of a drive failure, as much as a complete day of work could be lost. Multiply that by hundreds of users, and the cost for that day's labour added to down time can quickly grow to astronomical proportions. Another problem with traditional backup software packages is that data stores can become corrupt simply due to the excessive amount of time a backup can take. In many instances, portions of data files, such as indexes or message stores, can change during an active backup, which then can lead to data corruption during a restore.

To counter hardware failures, many businesses have implemented RAID solutions. While mirrored or striped volumes do a good job protecting against data loss due to a drive failure, today's hard drives have increased in reliability and are not usually the first components to fail. In other words, the drive array is no longer the weak link in the chain, most failures occur due to other components. Simply put, when there is a component failure, data still becomes unavailable, regardless of the level of RAID implemented. The simple fact of the matter is that IT departments and knowledge workers can no longer rely on traditional backup technologies to stay competitive. The cost of disaster-proofing networks is escalating due to the inefficiencies of traditional backup solutions, where tape backup has become too slow and disk to disk backups can introduce both storage concerns and unforeseen failures due to separate component failures. The answer to those dilemmas comes in the form of keeping a second copy of the data files continuously available to the user. That can be accomplished with file synchronisation software like PeerSync and FolderMaestro, which excel at synchronising data files across both physical and virtual locations. File synchronisation brings new options to IT departments that are looking to further disaster-proof network storage, while controlling costs.

High availability:
The challenge of high availability solutions

Businesses have spent countless amounts of money attempting to live the high availability dream, only to find that there are always situations where high availability solutions fail. In the past the answer has been to throw more software, hardware and money at the problem, while costs steadily spiral out of control. Businesses deploy high availability solutions because they want to avoid down time at all costs. But, when you consider the real goal of high availability, you will find that the technology focuses on keeping users connected reliably to their data.

Traditional high availability solutions are built around complex technologies, ranging from mirrored server farms to automatic fail over devices. Before trying to build a high availability solution, businesses need to consider other options. First off, is the proposed solution overkill for businesses needs? Secondly, what services, software and data files need to be constantly available? After determining the answers to those questions, many businesses may be able to consider other options, namely file synchronisation.

File synchronisation offers a robust, tuneable solution for those looking to build a high availability environment. The advantages offered by a file synchronisation based solution are numerous, especially when one factors in the requirements of high availability. First off, file synchronisation products, such as PeerSync Server and FolderMaestro from Peer Software, offer the ability to keep two or more exact copies of data files in different locations. A simple example of how that capability applies to high availability would be a desktop computer attached to a network, with the "my documents" folder replicated and synchronised on a network share. If the network share becomes unavailable, then the user still has access to up to date critical files. The same concept applies if the user's PC fails, the user can switch to another PC and access the up-to-date data files on the network share.

While that example is an over-simplification of the capabilities of both PeerSync and FolderMaestro, it still illustrates how file synchronisation can complement or even in some cases replace expensive high availability solutions. Complex environments can also benefit from these application's file replication and synchronisation capabilities. In many cases, high availably data stores may be in constant use, which can present an open file problem, in other words, if the file remains open, it may not be able to be backed up. PeerSync overcomes that issue by implementing an open file manager, which allows open files to be replicated or synchronised. High availability is usually associated with high bandwidth; which impacts network performance. Both products solve that problem by offering bandwidth throttling, a technology that minimises synchronisations impact on network overhead.

Geographic distribution:
Protecting files by extending locations

One of the best ways to protect valuable data is to keep a fresh copy at another location. Traditional backup solutions handled this by rotating tapes offsite, but that created additional overhead and added expense in the form of both media and time. What's more, even if tapes are rotated offsite daily, the recoverable data will be at least a day old. But most importantly, the sad fact of the matter is that someone can forget to rotate tapes offsite on a daily basis, leaving a company out of business if a flood, fire or other disaster causes a facility loss.

To counter those problems, some organisations have turned to offsite/online backup services. Those services tend to be expensive, and put a company's valuable data in the hands of another party. Another problem with online backup services is that they are usually based upon scheduled backups, which means data is only as fresh as the last backup. File synchronisation offers a viable solution to the offsite backup conundrum. With the low cost of broadband Internet connections and storage products, backups based upon file replication and synchronisation become an affordable alternative to traditional offsite solutions.

When deploying a geographically dispersed backup solution, businesses need to consider several elements to guarantee success. One of the most important elements is the location of the backed up data. Luckily with PeerSync and FolderMaestro any location that can offer a viewable shared folder is suitable. For example, a small business owner could use his home computer as a replication point. For a larger business, a remote office with a NAS device may be more suitable. Keeping a current copy of data files at another location is not without its challenges. Those responsible will have to address several factors that can complicate the viability of a remote backup. First off, handling open files is of the utmost importance. Secondly, managing bandwidth usage merits serious consideration. Thirdly, automatic reconnection and validation need to be considered. There are other less critical considerations, which may have an impact on implementing a remote backup strategy, such as data compression or byte level replication, reporting and e-mail notifications.

All of those needs can be met by PeerSync and FolderMaestro which offer an extensive feature set. Those working with large files will appreciate byte level replication, which only synchronises the changes in data files, speeding up the process, while conserving bandwidth. The open file manager capabilities in PeerSync helps to further close the backup loop, by backing up in-use files, such as databases and documents. In short, file synchronisation makes remote backup an affordable reality for most any business.

Staging data:
Moving files to production systems

One of the biggest challenges faced in IT production environments is validating and testing solutions before deployment. Whether it's modifying log-on scripts, or updating web pages, best practices dictate that those changes should be tested on a non-production system before being rolled out to active systems. Most IT departments have the testing methodology down to a science, but still struggle with the deployment of changes across dozens, if not hundreds, of systems. Those deployment issues consume valuable man-hours and can interrupt daily operations, while creating a risk that some system updates will be overlooked.

In the past, busy IT departments attempting to institute updates via manual or batch-file based processes encountered numerous barriers to success, due to issues involving open files, system down time, scheduling and verifying updates, and reporting on update failures or successes. Those issues can be resolved by switching over to a process that based on file synchronisation.

File synchronisation products, such as PeerSync, offer the ability to work with open files, use persistent technologies to make sure updates occur, and offer robust error and success reporting via comprehensive logs. File replication and synchronisation offer the ability to move files across platforms based upon date and time stamps or other specific file parameters. That becomes an important consideration when rolling out changes in an organisation. For example, a web development team could update some static HTML pages and then need to move those changed pages out to production web servers. The easiest way to accomplish this on a large scale would be to setup a central shared directory to house the updates and deploy a file synchronisation solution. Developers would then copy the updated files into that directory. The many-to-one relationship supported by PeerSync would allow the remote web servers to detect the updated files and then pull those files down, making what was once a time-consuming manual process an automatic instantly viable update.

File synchronisation capabilities can greatly improve updating production environments and the same techniques can be applied to most any file updating or roll out process. What's more, file synchronisation opens the door to other data protection solutions. There is no reason why one can't leverage the capabilities offered by synchronisation to solve problems beyond data rollouts. A typical IT department can implement synchronisation to solve an immediate problem, such as data file push/pull, and then advance the solution to support offsite backup or real-time backup needs. In other words, file synchronisation offers a solution to many problems.

Data portability:
Keeping data files fresh for the mobile user

In today's connected world, almost every business has the wandering or remote user. Whether it's the business owner, director of sales, or a field technician, one element remains the same: those users need to have access to their data files. IT departments have struggled to keep mobile users happy, but past solutions have failed miserably. IT support personal have attempted to train users to copy or e-mail files, or have instituted log-on scripts that copy files from a shared drive to the workers' PC.

Those solutions have created even bigger problems. In many cases, travelling workers would inadvertently copy an older version of the files and destroy the updated versions, or become confused by folder structures and scatter different file revisions around their local hard drive and network shares. Those problems became support issues, wasting help desk time, or even worse, losing important data files.

Leveraging file synchronisation can solve those problems, without users needing to worry about where a file is and if it is up-to-date. Furthermore, an IT department can benefit from the "set and forget" capabilities offered by file synchronisation products, which translates to reduced help desk calls. What's more, a file synchronisation platform can be extended to include backup capabilities, which means critical files, such as a user's local mailbox, can be included in the synchronisation process.

A simple synchronisation solution would be to use PeerSync Workstation or FolderMaestro Laptop, both low-cost synchronisation products that support background synchronisation and replication. The product could be used in a multitude of fashions to keep files up-to-date. An IT department could script the synchronisation process to run via a log-on script whenever the user attaches to a network share, that could occur over a remote VPN connection or whenever the user is in the office.

That style of implementation will help to keep critical files backed up, while keeping data files fresh. It is important to understand that file synchronisation must be a two-way process, especially in the case where a user has more than one system (i.e.: a desktop at work and a notebook for travel). In that situation a share can be setup to store all critical data files and then synchronised or replicated files to the local hard drives of each system.

The synchronisation process will insure that only the latest files take precedence. Not only is that a good solution for the travelling user, it allows users who work from multiple locations (a home PC and an office PC) to switch between systems with no complications. When you add up the requirements, file synchronisation becomes the near perfect solution to address the needs of the mobile and remote worker.

Data synchronisation:
Multiple techniques for file synchronisation

Data file synchronisation and replication offers a vast array of solutions to the typical IT department. What once started out as a simple approach to copying files between storage locations has evolved into a robust process that can be used for a range of services, from directory maintenance, to backup, to deploying software.

That may make some wonder as to what has changed with file synchronisation products to make all of that possible. The answer comes from innovations, such as those from Peer Software. PeerSync, PeerLock and FolderMaestro bring important features to the market, which help to extend the viability of the technology by adding capabilities. A major enhancement comes in the form of open file handling which allow locked files or files that are in use to be synchronised. That helps solve the real time backup dilemma for many organisations.

Another key feature comes in the form of directory pruning; administrators can now setup synchronisation software to remove files from target directories, which helps to keep disk space under control, while eliminating suspicious or un-authorised files. Byte level replication helps to extend file synchronisation to control the replication of massively large files, without sacrificing bandwidth or performance. In short, byte level replication only synchronises the portions of a file that has changed, greatly reducing overhead.

Scheduling and scripting capabilities add flexibility to the synchronisation process; administrators can launch synchronisation jobs via a script or via a schedule, a handy solution for files that do not need to be synchronised in real time. For example, a script can be created that runs only after a business closes a reporting quarter, allowing effortless storage of archival data.

What can be accomplished using PeerSync's options is almost unlimited. By combining several features of a synchronisation package, a single script could be created to backup, archive and then deploy files to specific targets. That helps to save time by automating manual tasks. File synchronisation and replication utilities offer solutions to many common IT problems. IT staffers need to focus on the possibilities and shift their thinking towards data's meaning in the work environment, then deploy the most cost-effective tools to accomplish their goals.

The real key to deploying synchronisation comes in the form of understanding how an organisation's data is used and how that data should be protected or preserved. The flexibility of synchronisation solutions allows IT departments to address most of the day to day concerns by keeping data both safe and available under most any condition.

The ROI of synchronisation:
Adding value to the process.

When evaluating the cost of deploying a synchronisation package, IT managers need to consider the benefits offered by the technology. In the backup realm, there are two costs: materials and time. On the materials side, tape has proven to have a high cost per gigabyte. On the time front, traditional backup to tape technologies often reduce system performance while taking several hours to complete. Those deploying file synchronisation can reduce costs significantly. Synchronisation can be used to move backup files to low-cost online storage, which then becomes the target for backup tape operations. That helps to reduce time concerns; tape-based backups can be setup to occur at any time, because the tape backup process now has no impact on live systems. What's more, multiple daily backups may be contained on a single high capacity tape, reducing the number of tapes needed.

The cost benefits of synchronisation continue when IT managers consider how the technology can be used to replace or supplement cost-intensive high availability systems. The major costs associated with high availability are software, hardware and bandwidth. A synchronisation process can be used to keep most data files available during network problems, reducing the need for a robust high availability solution. For example, a user's documents or other data files can be synchronised to the user's local hard drive, leaving important files available if there is a network failure. Then once the network is restored, the updated files will be automatically returned to the network. On the network side, critical files can be replicated from a high cost RAID array to low cost NAS unit, which can still provide access to files if the raid system fails or needs to be taken down for maintenance.

One of the biggest returns on investment scenarios for a file synchronisation solution comes in the form of supporting mobile users. File synchronisation can solve many of the high cost problems associated with remote or mobile users. For instance, when a travelling user leaves the office, they are often without access to their data files. Synchronisation allows that user to maintain access to their files, without needing assistance from a help desk or working with contrived manual processes. The savings come in the form of efficiency; users will always have access to their data, while reducing the burden on IT departments. IT managers can maximise their ROI by combining all of those features into a single solution built around file synchronisation technology. No other software product promises to solve as many problems as a well implemented synchronisation package, while significantly reducing overall operational costs. It simply comes down to leveraging all of the capabilities offered by the technology.

Disaster Recovery:
Using new technologies to build business continuity.

Business continuity and disaster recovery planning are the latest business buzzwords. To understand the disaster recovery process and business continuity environment, one has to break the concepts down into the smallest and most critical elements, namely intellectual property. By following best practices associated with a file synchronisation solution, businesses can speed disaster recovery and improve business continuity.

A disaster recovery plan outlines the procedures necessary to preserve business continuity. The concept of business continuity is best described as keeping a business functioning under any condition. File synchronisation lends itself well to those concepts.

The first area that a disaster plan addresses is a loss of facility, in other words, what can be accomplished if the office is unavailable. File synchronisation solves that problem by geographically distributing data, or for smaller locations, backing up data to a portable computer that is removed from the site daily. That data can be quickly retrieved for use at an alternate location by simply synchronising the backed up data directories with the replacement hardware, or even use the data directly off of the backup location. Those techniques prove to be faster then traditional tape restoration.

Virus attacks or computer worms also affect business continuity. Including multiple event-drive replications as part of a backup strategy can reduce the impact of those issues. After the problem is resolved, IT managers will find restoring data from a synchronised share much quicker then relying on tape.

Business continuity concepts also extend to major upgrades or equipment replacement – events that can disrupt data access, resulting in lost productivity. Synchronisation can solve those continuity issues by keeping a fresh copy of the data stored in an alternate location, i.e. the user's local hard drive. This allows users to keep working with their data files, even if the network is down for maintenance, and then automatically restores the files back to the network, once it is restored.

Disaster recovery also extends down to the desktop. In many cases, companies put a plan in place to address the failure of a key PC, such as a payroll or sales system. File synchronisation can help keep that PC backed up in real time, without any end use intervention. Using technology such as open file management, all critical local files, ranging from e-mail PST files to accounting software files, can be replicated to another location. That speeds replacement of a failed PC and minimises down time. Simply setup a new system and then replicate the backed up data down to the new system. That ideology extends down to small businesses, which normally have no disaster recovery plan in effect. When it comes to businesses continuity, file synchronisation makes a lot of sense, while still going easy on the budget.

Conclusion

File synchronisation, replication and backup solutions can cost-effectively solve many of today's problems with keeping data available and up-to-date. At Purple Rage we've been providing the business world with what we think is the best file synchronisation, replication and backup software solutions available. The PeerSync, PeerLock and FolderMaestro products from Peer Software Inc represent the latest generation of this mature technology and offer you a full range of data sharing and protection options for both LAN and WAN environments.

We have fully functioning software product trials that can help you evaluate your needs and make a more informed choice. If you'd like to get information please visit purplerage.com/peer or email enquiries@purplerage.com.

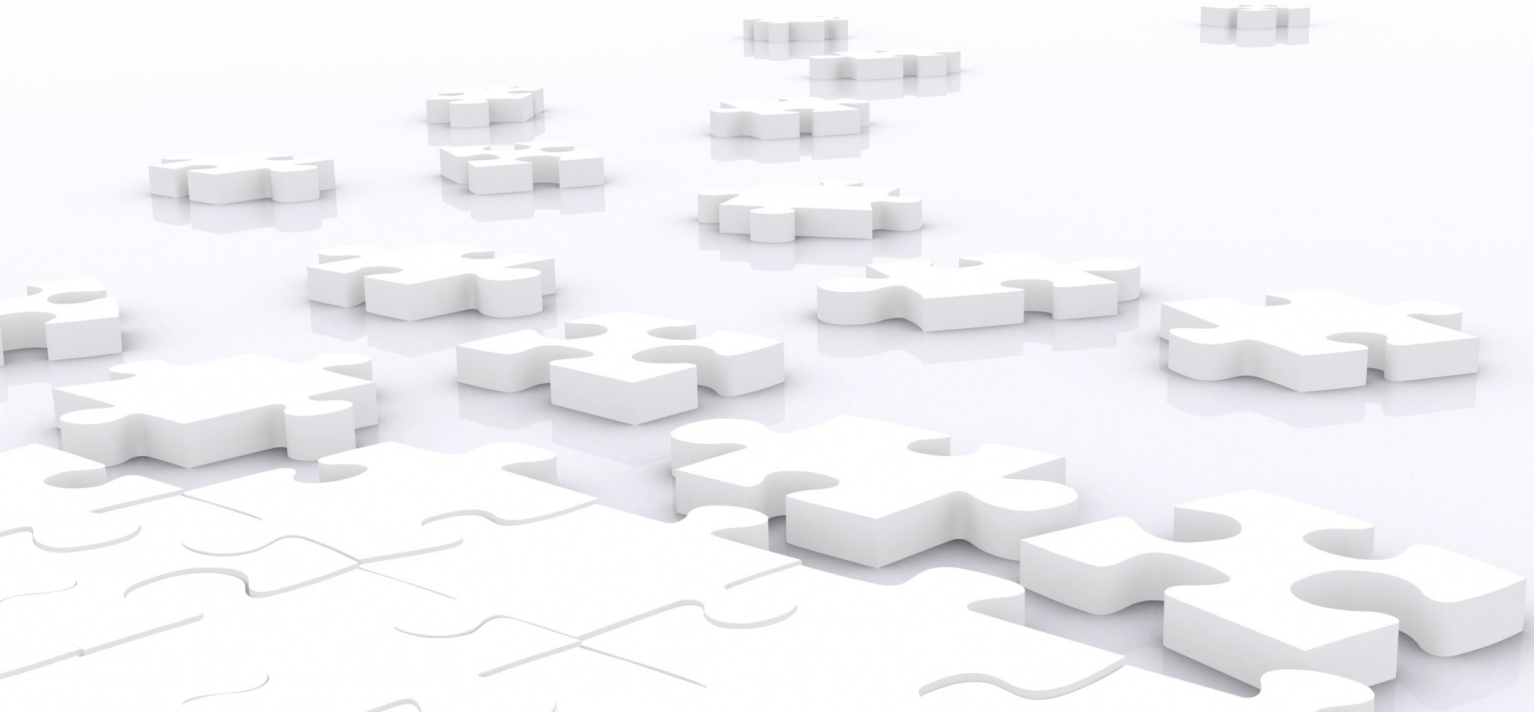

PurpleRage
DYNAMIC BUSINESS SOFTWARE

PO Box 243, Malvern WR14 9AX,
United Kingdom
+44 (0)1684 576343
enquiries@purplerage.com
www.purplerage.com