

Management Series

Delivering Holistic Business Continuity with Auto-Replication

From accidental file deletions to fires and floods, all kinds of data disasters threaten business operations and escalate the demands on IT. Better business continuity solutions are needed - to improve backup, speed restore, and reduce lost work. Only remotely replicated, disk-based image copies offer complete data protection and rapid recovery, but historically this function has only been available as an expensive software add-on. Today, EqualLogic, Inc. makes automatic remote replication fast, easy, and affordable as a standard feature of its PS Series array.

Operations Must Be Online

No matter what your organization does or what size it is, business continuity matters to you. Operating around the clock has become extremely important in this internet age, as companies do more and more business over the Web, around the world. In the "old days," there were backup windows at night when customer and supplier interactions ceased. Today, your customers and suppliers may be across the globe—awake and transacting business in a different time zone. To stay in the game, your business needs to operate 24 hours a day, whether serving customers or maintaining internal systems. As you well know, the luxury of nightly downtime is long gone, replaced with demands for high service levels at all hours.

What is a Disaster, Anyway?

Consequently, business continuity is a major strategic initiative for most organizations. You must have a clear "disaster recovery" plan in place to restart business within an acceptable timeframe. The plan should include key categories—buildings, equipment/ infrastructure, people, and data. When most of us hear the word "disaster," we think of a devastating fire, flood, earthquake, or some other dramatic event. These events are infrequent, but they do occur and you must be prepared for them. In terms of data protection, you're safe from these events as long as you have a recent copy of your data stored a safe distance away.

But in the business world, a disaster can be any kind of interruption—such as a software malfunction, virus attack, database corruption, accidental file deletion, network outage, or hardware failure. These occur much more easily and frequently than catastrophic physical phenomena and require you to implement disaster recovery plans such as replacing equipment, recovering data, restarting computer systems, and redeploying staff. For these events, current data is not enough, since it is likely to be corrupted and useless. You need multiple data copies at several different points in time to ensure clean data for business restart.


The challenge facing most organizations is how to minimize the lost revenue and operating expenses these disasters cause. Even without losing buildings and equipment to hurricanes or floods, organizations suffer in many ways, including:

- Loss of employee productivity
- Loss of customer business during downtime
- Loss of customers who, when unable to do business with you, find another vendor—permanently
- Financial penalties from violation of legal requirements

Failure to recover from a disaster has led to the demise of many organizations and threatened the existence of others. That's why companies take business continuity so seriously and why a disaster recovery plan is critical.

Data Protection—One Part of the Insurance Policy

Without business data, very little can be accomplished, so a complete "insurance policy" for your data is of paramount importance. That starts with a regular process for data backup—creating and safely storing copies of data—which is the core of business continuity. Everyone does backup—most of us groan about it, because frankly, backup is a pain—time-consuming, complex, and awkward—but we must do it anyway. Data protection also means finding a way to recover data and restart business if a fire destroys your storage equipment, or a software virus infects your systems, or an application upgrade corrupts data. Improved backup and disaster recovery systems—a better insurance policy—would take significant burden from your IT staff. But it would need



to be affordable and automatic to really make a difference, and guarantee protection from all kinds of physical and digital disasters.

That's what EqualLogic has done. Our PS Series arrays include Auto-Replication as a standard feature, to deliver better backup and complete disaster recovery. Once data protection is solved, you'll have more time and resources to handle other business continuity tasks. And, once you are confident of business continuity, your organization can focus on its key mission—delivering goods and services that build revenue.

But let's go back to the beginning.

A Continuum of Disaster Preparations

Data recovery and business restart depend completely on what you do before the disaster. Do you have the right technologies to deliver the results you need? Do you have the proper plans and processes in place? For most organizations, the key metrics of recovery solutions are "How long does it take to get running again?" and "How much completed work do I lose?" Most organizations develop goals for acceptable amounts of lost time and data, known as recovery time objectives and recovery point objectives. Take a look at a simplified continuum of recovery methods, and their advantages and disadvantages.

Do IT Over

On the very low end, if you have no backup capability, recovering from a failure means simply doing the work over. Lost files, databases, and transaction records will need to be re-created. You have no data protection "insurance," so it costs you nothing prior to the loss—but you get what you paid for! It may be weeks or even months before you are up and running again and you lose an awful lot of work. In essence, you don't pay upfront—but you pay dramatically after a failure. This is an unacceptable situation for business today.

Tape Backup/Restore

The most common approach is tape backup. Disaster protection/recovery is accomplished by creating additional copies of data, saving them to tapes, and transporting these copies away from the primary data regularly. In the event of a disaster, the tapes are brought to functional equipment, and the data is restored by the backup software. By creating physical distance between the primary data and the tape copy, you protect against physical disasters. This dramatically reduces lost work when compared to "do IT over," and operations can be restored much more quickly—typically in hours to days, depending on your procedures. You now have the operating expense of running backup operations regularly and safely storing the backup copies away from the primary systems. Your cost is higher, but you have improved your data protection insurance.

Tape-Based Backup/Restore with Replication

This method involves making disk-based replicas of your data at various intervals. If your system fails, you have data available with which to restart in minutes to hours, instead of the hours to days that backup tapes require. And, you can replicate the system disks—allowing you to recover more quickly by avoiding the typical system re-installation process. Your upfront costs are higher, but you have better insurance and take much less time to recover. It's clear that in selecting a recovery method, you trade off the cost of insurance vs. the cost after failure, and you may get better

insurance if you pay more. But here's a caution: replication technologies vary in costs and in benefits, and don't all offer the same results. Paying more does not necessarily mean you are getting the best protection for your needs. But more on this later.

Replication technologies vary in costs and in benefits, and don't all offer the same results.

insurance if you pay more. But here's a caution: replication technologies vary in costs and in benefits, and don't all offer the same results. Paying more does not necessarily mean you are getting the best protection for your needs. But more on this later.

Backup—The Core of Business Continuity

Few would argue that tape backup is the basic foundation of business continuity. Backup accomplishes several things:

- Creates copies of data
- Makes copies in "backup format" on a transportable tape medium, so it can be moved off site
- Catalogs tape contents, so when you need to recover from the tapes, you can see what data was saved, and when
- Delivers a method for restoring data

Backup offers critically important benefits, as long as backups are done correctly and consistently. These benefits include:

- Disaster recovery—backup operations provide multiple recovery points such as "yesterday," "last week," "last month," "last year." This means that data corruptions or other unexpected changes can be recovered by using backups from a time prior to the data loss.
- A historical record, or archive, of corporate data
- Proof of regulatory compliance
- A legal record of business operations

Backup is the most affordable data protection solution, and covers all types of data loss—the common types (hardware failure, software corruption, virus, human error) as well as the uncommon types (natural disaster/physical destruction of site or equipment). It may not be the most glamorous part of IT—but organizations cannot survive without it.

Recent Backup Operational Improvements

Backups are commonly done locally, across the LAN, or using the SAN. With local backup, each server backs up to its own tape drive. With network backups, the individual per-server tape drives disappear, and the tape library is connected to a single server; backup data is sent over the LAN to that server. This consolidates backup operations, requiring fewer staff and usually fewer tape drives. It is somewhat more complex, but significantly easier and less intrusive to IT.

SAN backups improve operations even more. By moving the tape drives onto the SAN, each application server can send data directly to the tape drive, and use the backup server to manage the catalog and tape library, but not to move data. This offloads the network, offering better performance and faster backup and restore times.

Another recent backup improvement is backup to disk—instead of writing directly to tape, the application server can write to disk, and later the backup server uses the disk backup to create a copy on tape. This method can improve the initial backup performance bottleneck that sometimes occurs when writing directly to tape, and typically shortens backup times—it also provides the opportunity to move tape processing to daytime operations.

Continuous Backups

A step further on this disk to disk to tape methodology is continuous or near continuous backups. This is available today with products like Microsoft® Data Protection Manager, Mimosa NearPoint™, Sonasoft® SonaSafe™, Symantec™ (VERITAS®) Backup Exec™ 10d, and others. These disk-based backups typically provide 30 days of online backups, and reduce the complexity and expense of tape libraries and media.

The backup server executes a full backup to disk of the all server data, initializing itself with a complete data set. After that, the backup server will execute frequent incremental backups at designated intervals—every 5

minutes, every hour, whenever a log changes in your email or database system—whenever you want. These new technologies have redefined "incremental backup" to mean only the actual changes in each file—instead of the entire changed file—are backed up. Less data means backups run much faster, and you can backup more frequently because the overhead of each backup is significantly reduced.

If you lose a file at 4 pm, you don't have to go back to yesterday's version—you may have a 3 pm version, significantly reducing lost work and lost productivity. You gain dramatic improvement in backup frequency, increasing fidelity without hurting performance or increasing administrator burden.

Continuous backups can help improve catastrophic system loss, however the continuous backup server must first be recovered, before any other recovery can begin.

Backup Challenges

This is not to say that backup is easy and convenient, because it's not. There are numerous challenges, some operational and some logistical.

- Data is saved in backup format, and is not usable until converted back to its original format. Because the data is in backup format, recovery is slow. The software must copy data back and reprocess it on new storage to convert it into data that your applications can use. This can be a lengthy restore process—hours or even days. This process is further complicated when entire systems, including the backup server, need to be restored.
- Backups are an ongoing operational burden. They must be run regularly in order to be useful, occupying staff time and impacting production systems. Operations—and service levels—are impacted while backups are running. Think of backup like a large application workload that runs on your infrastructure daily—it puts demands on your servers, networks, and storage that must be factored into planning for each system. Performance is often degraded during backups, and application disruptions often occur.
- You must ensure compatibility of backup applications with heterogeneous infrastructure components. Backup is not easy to execute properly, and errors can corrupt your data. For live data, if applications are not properly prepared before backup begins, you cannot guarantee successful data recovery.
- After a failure there can be significant loss of work, because you will lose any data since the last backup. Most organizations aim for a nightly backup, commonly resulting in 12 or more hours of lost work if a disaster strikes.

Backups Suffer from Lack of Application Coordination

The most common problem for executing backups is coordination with running applications. With open applications, open files, and data being spread across multiple disks, it is easy to make errors. For instance, imagine that you begin backing up Disk 1, then Disk 2. You have a database with a number of files that span Disk 1 and Disk 2. If one of those records is being changed while backup is executed, chances are good that the backup will not save the correct data. But you won't know for certain.

How do you solve this problem? First, you can stop applications while you run backups. This means downtime, impacts service levels, and is antithetical to business continuity. What about setting up your backup software to skip any files that are open? With databases running around the clock, that would mean they never get backed up at all.

What Does Recovery Involve?

The overall goal of business continuity solution is to keep business running without interruption if possible, and with minimal interruption if not. But what actually happens when a disaster occurs?

Step 1: Identify that a disaster has occurred.

A physical disaster is likely to be noticed immediately, but the more common corruptions and failures may not be noticed until some time has passed. You eventually realize that you haven't received expected files, or a virus is spreading—and now you have to take action fast.

Step 2: Decision to Implement Recovery Scheme.

Once a problem is discovered, senior management must be alerted and decisions made about how to react to the disaster. The recovery plan must be reviewed and initiated.

Step 3: Set Up Infrastructure.

You may need to shut down systems, find the right data (tape or replicas), alter networking schemes, and change how employees work.

Step 4: Restart Servers.

Once alterations have been made and you are ready to get back to business, IT must restart the servers.

Step 5: Restart Applications.

Since your applications were not shut down cleanly due to the disaster, restarting applications can require some work. For instance, if your database was running at the time, it now must roll forward and roll back. If your database is large, that will delay business operations just that much longer—unless you have replication, which makes restart faster.

While the first two steps are likely to be about the same regardless of your replication method, Auto-Replication offers an advantage in steps three and five. Multiple replicas give you fast data restore. Also, since you have a replica that was coordinated with the application, application restart is fast. You can restart using the last Point in Time copy, and the database will start again without delay. If multiple failures occurred coincident with the disaster, recovery is ensured because of the multiple recovery points available with Auto-Replication. Other replication technologies cannot offer that, since they don't offer a previous "known good" data copy.

Improve Data Protection and Recovery by Leveraging Disk Image Copies (Replication)

Replication can produce tremendous benefits, but the various replication technologies offer quite different features, protections, and costs. There are three key replication methods: Snapshots, which are local image-based copies; Continuous Replication, including two kinds of remote mirrors, and Remote Point in Time Replication, which is essentially a Snapshot transported to a remote location.

Snapshots

A Snapshot is an instant copy of application data taken at a particular point in time. The copy is typically collocated with the primary data and shares storage with it. Snapshots can include complete copies (generated by breaking off a disk from mirrored disk sets), or just the data that has changed since the last Snapshot (using "copy on write" technology). There are several key benefits of Snapshots. First, they are created during the normal operation of production volumes, and when well implemented, do not impact performance—they require no downtime, and do not hinder operations. Second, you can easily maintain multiple Snapshots, offering you multiple recovery points—for example, an hour ago, 4 hours ago, yesterday, or last week. Third, Snapshots can be instantly restored back to the volumes from which they were created. This image of disks is in application-usable format, so no data conversion is necessary, resulting in fast restores. Fourth, snapshots can be "assigned" to another server and be used for backup, testing, and other operations that can be done in parallel, offloading processing from application servers.

Snapshots Cover Many Failures

Because Snapshots offer you a "known good" copy of your data, they are good protection against any outage or failure. You might take a Snapshot copy once a day or once an hour, depending on your needs, and use it to run backups from another server. Since you are using the Snapshot and not the production volume, it doesn't matter if backup takes 5 minutes or 5 hours—your production systems will not be involved or impacted. You'll still need to coordinate the Snapshot with the application, but Snapshots can be created quickly—typically in under a few seconds. In addition, many applications are Snapshot aware, so this can be done online.

Let's say you make Snapshots every hour. Now, you are protected from every circumstance—if you accidentally delete a file, you have a copy from no more than an hour ago. You can mount the Snapshot and recover that file immediately. What about a software corruption or virus? When you discover that your data was corrupted four hours ago, you can go back to the Snapshot from five or 12 hours ago and know that you have a good copy of data. You've lost some data, but you can restore quickly. Another way to take advantage of this technology is to make a Snapshot of your data prior to installing a new system disk or software patch. Then perform the installation—if any problem arises, you can simply revert to the Snapshot and be back online instantaneously.

You can see that using Snapshots in conjunction with typical tape backup can reduce the amount of data you lose in a failure, dramatically reduce recovery time, and offload application servers. Snapshots can also provide you with history; although you are not likely to keep a year's worth of Snapshots like you would with backup tapes, you can keep multiple copies for a week or more. The one vulnerability of Snapshots is that they are collocated with the primary data—any physical disaster that destroys the primary data will destroy the Snapshots as well.

Continuous Replication Methods Have Limited Use

Continuous Replication is basically RAID, or disk mirroring, extended over distance. There are two kinds of Continuous Replication, synchronous and asynchronous; both execute "continuous writes" to remote storage. With Continuous Replication, every time a write is done to the production disk, a write is done to the remote site. Therefore, the copy is identical to production data at every moment.

Continuous Replication methods minimize data loss in the event of a physical disaster. However, they cannot be used for protection from common disasters since they provide no previous "known good" copy of data, just one copy of the production data. So if your production database is corrupted, you can't look to the replica because it is corrupted too. Clearly, you must do backups to give yourself a historical copy and truly protect data—but since these replicas were not coordinated with applications, they cannot be used to run backups. Only Snapshots are known to be good copies of data at a given point in time. Other weaknesses of synchronous or asynchronous Continuous Replication include:

- They require low latency network links to support continuous writes. Typically, your remote site must be within about 100 miles/150km of the production site to meet the latency and application performance requirements.
- Continuous Replication requires high bandwidth network links, in order to keep both sides up to date. For both methods, if you write over the same block 100 times in an hour, you sent 100 writes over the network to the remote storage. This is different from a Snapshot, which will write only the changed data—one write, instead of 100 or more.
- They offer no history or other recovery points—so if you have corrupted data at the production site, you also have corrupted data at the remote site.
- Application recovery takes longer than with Snapshots, because there is no coordination with the application.
- These weaknesses combine to create an expensive solution that only protects you from some of your business continuity threats (physical disasters).

Continuous Replication is complicated to set up and complex to manage. Some large firms use Continuous Replication for a few applications that cannot afford to have certain transactions—financial dealings, for example—be even moments out of date. These firms also count on Snapshots, Remote Point in Time Replication, and tape backups because they know that Continuous Replication serves only a small fraction of their data protection needs.

Auto-Replication Covers All the Bases

Snapshots improve backup/restore and protect you from common disasters. Using network and storage technology, you can create Remote Point in Time copies, which are essentially remote Snapshots. The Snapshot technology gives you a replica from which to backup/restore and protects you from common corruption and human error failures, and the addition of physical distance protects from physical disasters. Usable data can be restored at the remote site, without having to convert backup data. This form of data protection is slightly more expensive than tape backups alone, because you need network links and storage at the remote site. But the advantage is clear—total protection and faster recovery. This method has all the benefits of Snapshots, without the weakness of collocation with primary data.

	Multiple Recovery Points	Distance from Primary Data	Recoverability	Improves Backup	Protects Against
Snapshot	✓		✓	✓	Common corruptions only
Continuous Replication		(some)	(some)		Physical disasters only
Remote Point in Time Replication	✓	✓	✓	✓	Common corruptions and physical disasters

This chart shows how the key business continuity requirements (number of recovery point copies, distance from production data, easy recoverability of data, backup improvement) are addressed by each replication method, as well as what protections each offers.

EqualLogic Builds Automatic Remote Replication into the Array

Most organizations need a holistic business continuity solution that delivers better, faster, more storage- and network-efficient backup and disaster recovery, regardless of the kind of disaster. EqualLogic, Inc. has taken the bold step of including Auto-Replication, which is automatic, Remote Point in Time Replication as a standard feature of the PS Series arrays—working from the premise that including more standard tools will vastly improve storage management for organizations of all sizes. EqualLogic has engineered Auto-Replication into its array for rapid, reliable backup and restore that ensures business continuity.

Remote Point in Time Replication in the PS Series array delivers the advantages of Snapshots with the additional protection of physical distance. Auto-replication can be performed in both directions—Site A can replicate to Site B, Site B can replicate to Site A, and they can replicate to each other. However, there is no need for a one-to-one correspondence—replication is between groups, not specifically between arrays. As a result, Site A may have five arrays and Site B may have three, and each can grow independently without impacting on the other.

Auto-replication uses "copy on write" technology to replicate data between PS Series groups and provide efficient operation over any type of network. With Auto-Replication, you create an image copy quickly at a remote location, and gain a stable data copy to use for backup or any other activities. These copies are network and storage efficient, don't hinder operations or performance, and offer multiple recovery points and quick restore. They are a practical, affordable, easy-to-use answer to the need for data protection that covers all kinds of functional failures.

Auto-Replication: Easy, Fast, and Affordable

A number of features make this solution extremely useful to any organization seeking backup and disaster recovery improvements. These features focus on three underlying benefits: making replication easy, fast, and affordable.

Network and Storage Efficiency

Because Auto-Replication makes Remote Point in Time copies, replications after the first one can include only changed data, making the process both network—and storage-efficient. You won't be clogging the network with multiple writes or needing massive amounts of storage capacity in order to retain data copies.

Multiple Recovery Points

As with Snapshots, you can keep many differential copies for multiple recovery points. If you replicate a volume called Database (DB) at noon, 3 pm, and 6 pm, then DB(noon), DB(3pm), and DB(6pm) will store the differences from each other. They will use very little disk space, while looking to the administrator like three full copies.

Instant Recovery

Auto-Replication creates image copies on disk that are ready to use immediately—no lengthy conversions from backup format. Like a Snapshot, the replica can be brought back online instantly—just at a different location.

Standard Feature

Because Auto-Replication is a standard, easy-to-use feature of EqualLogic's affordable iSCSI storage solution, you can start slowly and increase usage over time. No additional licenses are required.

Standard Networking

Auto-Replication uses standard TCP/IP and iSCSI so you can replicate over any IP network, over any distance, including T1, T3, DS3, or OC networks. No network devices or special equipment are required, just standard TCP/IP networking. As a result, with Auto-Replication you can continue to manage your network traffic in the same way you always have—no new management skills or tools are necessary. The flexibility, use of standard networks and procedures, and familiar network management make Auto-Replication simple to use.

Delegated Management Model

This solution is easy to set up, and uses a delegated management model that enables cooperation between administrators at different sites, while ensuring complete independence and security. Divisions of one organization that don't normally work together can share storage capacity for data protection; similarly, an ISP that serves as a disaster recovery site for many organizations can more easily and efficiently delegate capacity with assurance of data protection.

No Host Software

Auto-Replication works in the storage array, leaving the servers to handle their own jobs.

Simple Management

Like other PS Series capabilities, Auto-Replication is managed and monitored through a single interface, preserving the architecture's ability to operate as a single unit.

Improves Backup

You can keep multiple replications for as long as you like, so they can act like backups—only much easier to access. If you use them like a backup, you make backup more efficient, offload backups from a server at the production site, and make restore much faster and simpler.

Automatic or On Demand

Replication can be done on demand or automatically according to a customer-defined schedule, whether that is every 10 minutes, every hour, three times a day, weekly, or whatever suits your IT environment.

No Disk Reconfiguration

Administrators can turn replication on or off at any time without reconfiguring disks. You may decide to replicate some volumes that have contained data for two years—you need only turn the feature on and select the volumes to replicate.

Flexible Site Configuration

Each location can configure its PS Series Group as needed by that location. There is no requirement that the locations be the same size, configuration, or RAID level. Auto-Replication supports RAID 10, RAID 50, and RAID 5.

What Can You Do with Auto-Replication?

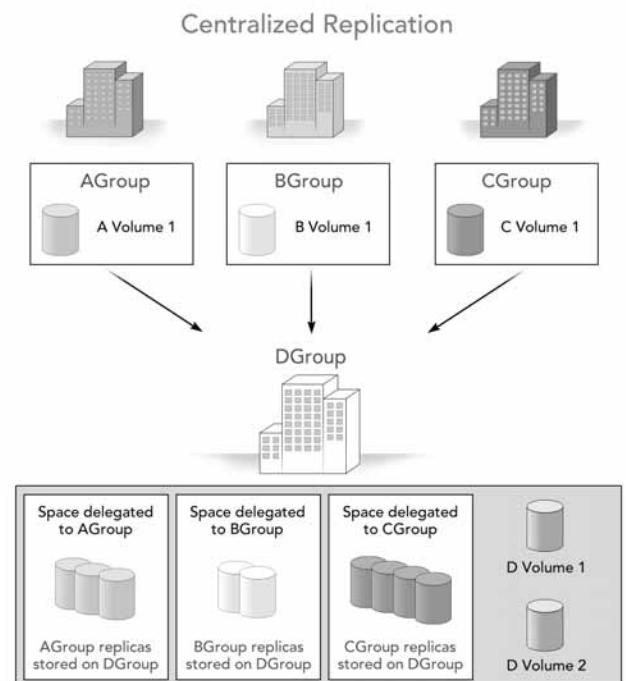
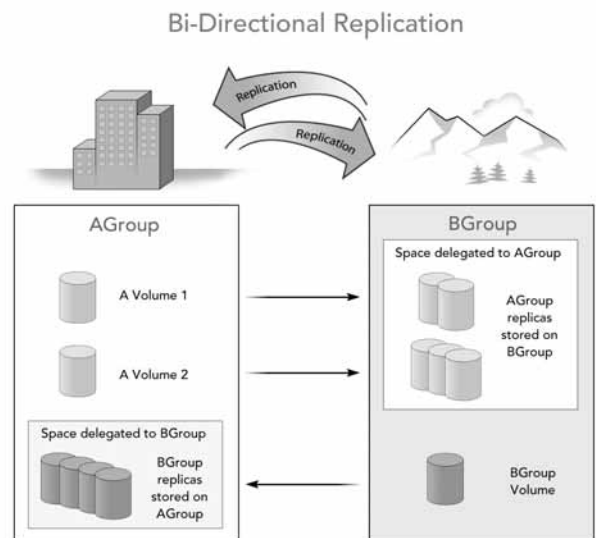
With EqualLogic's PS Series arrays, remote replication becomes accessible and affordable for most organizations, whether their remote locations are across campus, across the country, or around the globe. There are many uses of this functionality, but a few are described below.

Traditional Disaster Recovery

In addition to improving backup, Auto-Replication delivers traditional disaster recovery. With PS Series groups at your production site and your remote site configured as replication partners, you can schedule regular replication at your disaster recovery site. This provides you with multiple recovery points that protect you from site failures as well as corruptions and viruses. If a virus strikes, you have several recovery points from which to restore data and quickly restart business operations.

Centralized Replication & Backup

Organizations with several locations can centralize disaster replication and backup. Instead of building a complete backup/restore infrastructure at each location (including backup servers, tape drives, backup software, and staff), you can use one location for backup and replication. For example, if Sites A, B, and C are production sites, you can eliminate backup equipment at all of them and only build that architecture at Site D. Site A can schedule replication to Site D at its convenience, and so can Sites B and C. Tape backups can be done at Site D from the Remote Point in Time Replicas, which are also available for quick restore should a disaster occur. Auto-Replication also improves backups by enabling multiple recovery point replicas, offloading backup away from production operations, and running backups from replicas.



"Buddy System"

The PS Series array's delegated management model enables independent divisions of an organization to share storage capacity for replication while remaining autonomous. For example, a large manufacturing company has two distinct divisions—one builds helicopters, the other builds elevators. These divisions function completely independently—they have their own headquarters, data centers, staff, etc. With Auto-Replication between PS Series groups, the divisions can agree up front to designate 500 GB of their storage capacity to the other division for replication. Once passwords are exchanged, each can replicate freely, using the other site's data center as its remote site. The Helicopter Division storage administrator has complete control of his/her replicas in the Elevator Division's data center, but no access or privileges on the Elevator division's arrays. After the initial agreement, the staffs at each site do not need to communicate or work together—they function with complete independence and total security. Each administrator continues to manage his/her own SAN without interruption, and each remotely manages the replicas.

This approach offers greater security and cost efficiency. Now, organizations can avoid building complete data centers for disaster recovery—instead, they can leverage infrastructure at another location within the company. They avoid large expenditures for buildings and equipment, data center infrastructure, staffing, and initial and ongoing expenses of facility management. Auto-Replication's security and delegated management model make this possible, allowing storage sharing without requiring interaction between storage administrators.

Moving Data Online

IT departments can also use Auto-Replication to move data or business applications between data centers more safely and efficiently. By including Auto-Replication as a standard feature, even occasional use of replication is both easy and affordable.

Summary

Every organization needs a holistic business continuity solution for data protection. Auto-Replication protects you from all disasters, as well as improving backup. You cannot afford to skimp on this critical business insurance—but you don't have to buy and manage add-on software that only solves niche problems, either. With PS Series arrays from EqualLogic, you have a complete replication solution that is efficient, affordable, easy to manage, and delivers total protection.

**For more information regarding EqualLogic and the PS Series,
please visit www.equallogic.com or contact us at 888-579-9762 ext 7792**



9 Townsend West, Nashua, NH 03063
Tel 603.579.9762 / Fax 603.579.6910 / www.equallogic.com

Copyright © 2004–2006 EqualLogic, Inc. All rights reserved. EqualLogic, PeerStorage, and Simplifying Networked Storage are trademarks or registered trademarks of EqualLogic, Inc. All other brand or product names mentioned are the trademarks or registered trademarks owned by their respective companies or organizations.

WP902_02.06.06