

## The Case for Continuous Data Protection

This paper will discuss the benefits of using CDP-based solutions, particularly in place of traditional, tape-based backup products that currently dominate the SMB market..

### CONTENTS

Abstract	2
Snapshots vs. Continuous Data Protection	2
— Snapshot	2
— SonicWALL CDP Solution	3
— Recovery time is important	3
Tape Rotation vs. Archiving vs. CDP	4
— Tape rotation	4
— CDP, how it works	5
Conclusion	7



**Abstract:** Today small- to –medium businesses (SMBs) are learning what large enterprises learned years ago – that as the value, complexity and volume of their data increases, conventional fixed-schedule single point in time (SPIT) backup to tape or CD no longer provides adequate protection. New customers are added, transactions are processed, customers are billed. And mounting concerns for security, compliance and business continuity make protecting this data a top business priority.

So, why do most companies still risk this critical asset by only backing it up once a night, thus putting their reputation and business at risk? The answer you'll hear is, "That's how we've always done it." Nightly tape-based backup was invented in the early days of computing, when data was entered through batch processing and there were duplicate paper records for everything.

Since then, businesses large and small have moved to a highly interactive, fully computer-dependent, networked world. Business is no longer 9-to-5. Due to the very nature of this new "always-on" world, the risk of exposure to data loss is constant. Viruses, worms and hackers join human error, data corruption, system failure and physical loss that can destroy your data at any time.

Despite the passing of time, conventional tape-based backup hasn't changed much in 30 years. It remains the one batch process we still use in our computing environments. With today's constant exposure to loss and constant dependence on data, a new approach is needed.

Although SMBs are beginning to require enterprise IT functionality, they also require a simple-to-install and easy-to-use solution. In order to acquire these functionalities, most small businesses hire IT consultants, then purchase backup software and backup hardware. These solutions are low quality, pose integration risks and require support from multiple vendors.

Seventy-five percent of SMBs backup via antiquated tape systems. Until now, these businesses have had little choice but to live with the inherent inadequacies of tape. According to a recent Storage Magazine survey, 77% of respondents found failed tapes while testing their backup system. Tape is prone to human error, requires manual intervention and does not provide for instant recovery. By the year 2006, over 70% of SMB data – e-mail and attachments, contact, patient, and customer records and business documents will be stored on a PC (source: IDC). For this reason, instant recovery of data is a top of mind issue for these business owners.

## Snapshots versus Continuous Data Protection (CDP)

### Snapshot

Until recently, storage technology has been one-dimensional—a set of blocks stored in specific locations with no history or sense of time. When the application overwrites the disk, the data simply changes. If the change is mistakenly caused by a faulty upgrade, administrator misstep, application error or any of the many ways corporate data can be corrupted, bad data is written to the disk blocks and good data is lost forever. The only solution available to date has been to perform tape-based data backups. Tape backups represent a single point in time (SPIT) copy of the application data stored on secondary storage media.<sup>1</sup>

In SPIT the recovery process is divided into two parts: data restoration and application recovery. Data restoration is purely data movement – copying the (hopefully) correct data from one place to another. When basing their recovery on SPIT images, companies faced with data corruption must rely on the most recent uncorrupted SPIT image. This SPIT image must be treated with kid gloves because it is the company's only copy of the previous data image. Application recovery should never be attempted directly with this image. Instead, the contents of the image must be copied to another location and operated on there, insulating the pristine copy from mistakes or problems in the recovery process. If the SPIT image were to be corrupted during the recovery process, the company would be forced to use an older SPIT image (which means more data loss), or worse, restore from tape. As a result, recovery time is dominated by data movement. There is nothing dynamic about SPIT images, and as such they are treated as very fragile things, indeed. Further, the reliability of tape systems have been questioned for years.<sup>2</sup>

---

<sup>1</sup> Rowan, Revivio Software, 2004

<sup>2</sup> Rowan, Revivio Software, 2004

## SonicWALL CDP Solution Overview

The SonicWALL CDP solution shifts the data protection focus from backup to recovery, enabling users to virtually recover rapidly to any point in time by continuously capturing modifications to data. Any computer or server connected to a SonicWALL CDP appliance is continuously monitored to determine if files have been added, deleted or modified. If a file status or database changes, the difference in the file is automatically saved on the CDP appliance (for instant data recovery) and then to a secure offsite storage facility (for disaster recovery). The CDP software also continuously updates the SonicWALL Web infrastructure providing it both health status, and an instant directory update so remote administrators have current file information. When a file has changed, only the changed data is backed up to the appliance, not the entire copy of the revised file—saving valuable space even before compression of the file. Because the SonicWALL CDP solution runs on Linux OS and the files are compressed (and therefore dormant), it is not susceptible to Windows-based viruses. Finally, SonicWALL CDP software updates are seamlessly downloaded to the host appliance, and distributed to client PCs and servers without user intervention. Therefore, interoperability and upgrade issues with operating systems and other software are eliminated.

SonicWALL CDP backs up files as they are saved by users or at 30 minute increments for database transactions. Therefore, a user can go back to a minute, an hour, a day, a week, a month or a year before an interruption occurred. Users can perform this operation from their desktop without calling the help desk. SonicWALL CDP works on three types of files:

1. Unique Individual Files (i.e., Word, Excel, PowerPoint). Each time a user saves a document it is automatically saved to the SonicWALL CDP appliance, and then to an offsite location (optional). Each time the user saves the files only the difference between the previous version and the current version is saved on the SonicWALL CDP appliance.
2. Outlook PST Files: Outlook.pst files are backed up using SonicWALL's proprietary file system filter driver. This technology allows Outlook to be backed up even when the application is open.
3. SQL Server, MS Exchange Databases: SonicWALL CDP provides seamless support for databases and applications without the need to integrate third party software modules. IT Administrators can backup logs, differentials, and fulls. SQL Server and Exchange are time based backups, and can be set to the ½ hour level.

## Recovery Time is Important


Recovery Time Objective (RTO) defines how long the business can tolerate being offline or down from a failure. It typically includes the three phases of application recovery: analysis (typically 5% of the RTO); data restoration (typically 90% of the RTO); and recovery (typically 5% of the RTO). The largest portion of RTO is data restoration. SonicWALL CDP reduces time required for data restoration by over 95% and reduces RTO by over 80%.

## Tape Rotation Versus Archiving Versus CDP

Implementing a tape rotation process is not considered archiving. Archiving is simply taking a snapshot of data and saving it for a period of time. Government regulations such as Sarbanes / Oxley and HIPAA require users to maintain data for up to seven years. **A user can archive data at any time via SonicWALL's CDP offsite archive feature or via local archiving.**

### Tape Rotation

A tape rotation process overwrites tapes depending on the process used. Although there are many approaches to tape rotation, a popular implementation, and simple example, is the "Grandfather-Father-Son" approach. In this approach a company earmarks a tape for each day of the week and rotates these tapes on a weekly basis – almost like a treadmill. See the illustration below.



At the end of each week, a weekly tape is taken. **Therefore, by the end of the 2<sup>nd</sup> week, all daily tapes are overwritten from the first week.** All daily file changes from the week before are lost.

Depending on the weekly schedule, typically the 5<sup>th</sup> weekly will overwrite the first weekly, similar to the daily approach illustrated above. At this point, the most granular backup resolution is at the weekly level, except for the current week where dailies are available. Finally, at the end of the first month a monthly snapshot is taken.

The problem with this approach is the gap in the backup window (backup resolution). As tapes are overwritten (i.e., dailies after the second week, weeklies after the fourth week), the ability to recover data to a specific point in time becomes harder and harder. After two months, the best backup resolution will be at the previous month. After three weeks, the best resolution will be at the weekly level. Of course, at any point in time an archive can be taken, however, this is outside of the tape rotation process.

## Continuous Data Protection, How it Works

The SonicWALL Continuous Data Protection (CDP) Appliance replaces the tape rotation process AND allows users to archive their data.

SonicWALL CDP works on an intelligent time and version decay algorithm. **It is not based on a schedule.** CDP is based on the premise that recovery from any point in time is far better than an arbitrary SPIT chosen by an IT administrator. Each time a user saves a document the difference between the previous version and the current version is automatically replicated to the CDP appliance. The CDP solution saves 15 versions of unique individual files for an unlimited time period. Future releases of SonicWALL CDP will allow administrators to set version limits per file and the decay time period. Once a 16<sup>th</sup> version is created the system will delete a previous version of a file (but never the first version) based on where the most incremental amount of changes occurred within the smallest time interval for that file (10 minutes, one hour, one day, one month, three months, etc). This technique is far superior to tape rotation systems where a file is only saved once a day (and potentially lost after the first week). In addition, CDP allows a user to recover multiple versions from one month ago where a tape rotation gives only one copy from a specific month. Again tape falls down against CDP.


**Deleted Files:** If a file is created and then accidentally deleted between backup cycles, it's very probable that the tape rotation process will overwrite tapes with the last good version of the file. CDP eliminates this problem, as the file is **never** deleted from the SonicWALL CDP appliance. Only the administrator can purge the file from the SonicWALL CDP appliance. This guarantees that at least one version of the file is always available. This is a significant advantage to revision-based CDP over tapes.

**File Changes:** Evaluating changes to a file after three days, and then again after three months with CDP.


SonicWALL's CDP data protection algorithm is based on time blocks and file versions. The system monitors the following time blocks: 1 ) 30 minutes, 2) 2 hours, 3) one day, 4) one week, 5) two weeks, 6) one month, and 7) n months. Time blocks are measured and work backwards from the most recent saved version of the file. Therefore, the user is always working in the first time block (one hour). The current version limit is 15, however, a future release of SonicWALL CDP will allow administrators to set the file version limit and the time blocks.

For example. A user creates a file and saves it 16 times the first half hour. Because the user is inside of the first time block (one hour), the 16<sup>th</sup> save will delete the previous revision with the smallest incremental change (or delta). When the next time block passes (eight hours) from the most recent saved version of the file, the algorithm will guarantee that at least one revision is available from time block two (eight hours). Therefore, if the user saves the file 15 times quickly, the algorithm will delete a revision within the first time block (the most recent hour), guarantying a version from the second time block (eight hours ago). This process goes on for months to ensure that the user has at least one version from each time block. If a revision exists in each time block, the algorithm will delete the version with the smallest delta in the time block

As the diagram below illustrates, CDP allows users to go back to any saved version of the file. Therefore, there is much less chance of losing data. With tape, a user can only go back to the previous month. Therefore, incremental changes made to the file between months will be lost.



The scenario in reverse is illustrated below. With tape, all daily file changes lose the previous version of the file. With CDP, all versions are saved and can be retrieved



\* Based on configurable file version limits. To be released and patched automatically to SonicWALL CDP appliances

Once again, this changes the paradigm from one of backup to one of recovery. Users can still take a snapshot of data and create a local archive. The CDP component of the solution replaces the tape rotation process while minimizing the backup window.

## Conclusion

Disk-based continuous data protection is a better solution because it:

- Allows users to recover files themselves
- Allows users to recover files from any previous version of the file
- Eliminates the laborious time-consuming process of notoriously unreliable single point in time (SPIT) tape systems
- Allows users to archive their data both locally and offsite
- One Click Recovery: Reduces recovery time objective (RTO) by over 80%.
- Simple, affordable, hands and hassle free

Disk-based continuous data protection virtually eliminates the backup window while reducing recovery time from hours to minutes. The unprecedented data growth over the past years, coupled with unrelenting competitive pressures to remain operational around the clock, demand more than daily backups that have high-failure rates. While conventional tape-based backup was a strong IT tool 15 years ago, it clearly does not protect the business from the challenges of today's dynamic threats.

And any IT manager would agree that reducing recovery time and tightening the backup window can only be beneficial. SonicWALL CDP makes them possible.