



StorTrends

Continuous Data Technology

These days, the paper-less revolution, corporate governance, business ethics and regulatory compliance have all led to a tremendous influx of digital data. Safekeeping and management of stored data has, therefore, assumed supreme importance. Various new technologies are emerging in an attempt to offer an appropriate solution to address these requirements. Statistics show that data growth is almost doubling every year, yet the Recovery Time Objective (RTO) has remained the same, if not tighter. From the anatomy of data re-usage, it is clear that in overwhelming situations most recent data is recalled.

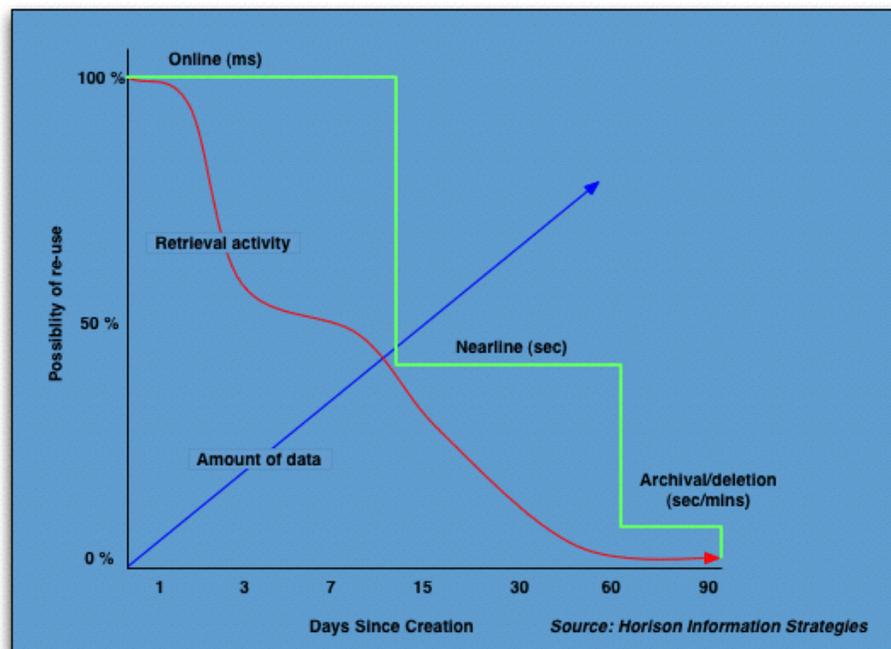


Figure 1: Possibility of reuse as a function of age

This immediacy of need and the corresponding instances to data reference steadily decline as the data ages. Continuous Data Protection (CDP) is now assuming utmost

importance and is no longer an industry buzzword. CDP, which deals with backup and restore, is probably the most important aspect of "Continuous Data Technology" (CDT). More than just recovery, CDT virtually handles any act of granular copy-creation in the enterprise. With this technology, granular instances of chronologically ordered data can be used to serve potentially a wide range of applications and not merely backup and recovery. Some of these are CDR (CDT for Replication) and CDI (CDT for Data Images). According to the Taneja Group: "...an emerging umbrella category, that we term Continuous Data Technologies (CDT), is poised to play a profound role in the coming three years."

StorTrends iTX 2.7, with up to 2,000 snapshots per volume (Near-CDP recoverability) and an efficient I/O journaling architecture (Full CDT), offers the solution, providing a snug hand-and-glove fit for the business need.

StorTrends CDT - a true Event Addressable Storage

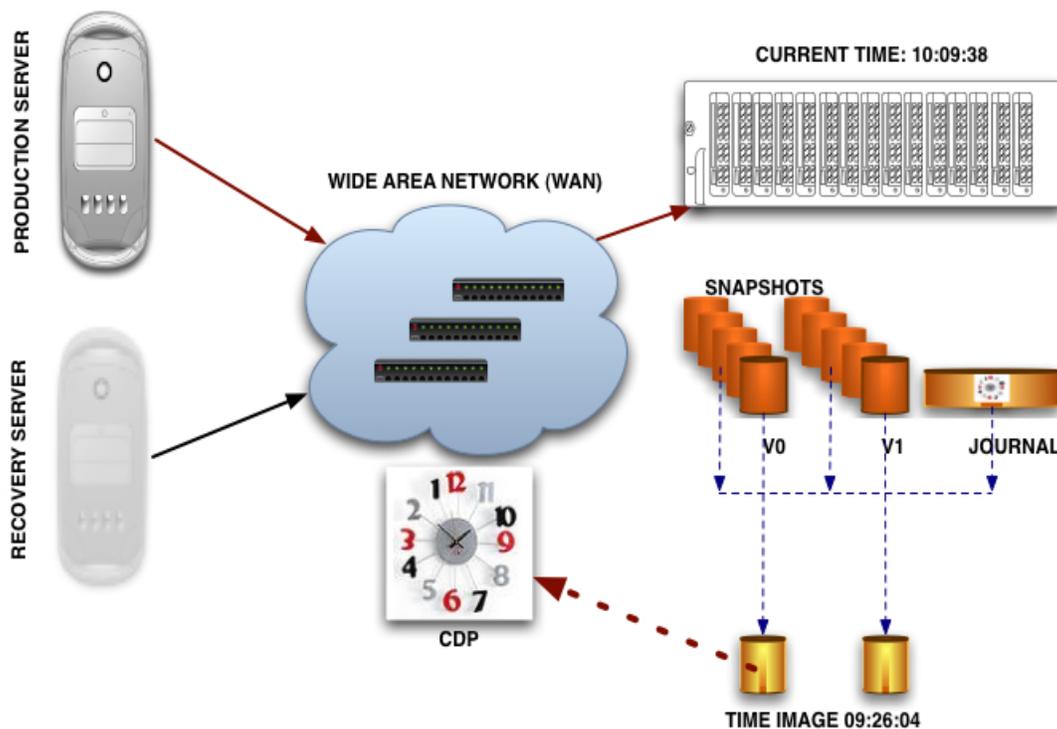


Figure 2: Continuous Data Technology in StorTrends iTX

Today there are various implementations of CDP. A vast majority of vendors implement it on the Host that is serving as the File or Application Server. In Host



based CDP, there is a filter driver or "Data-Tap" residing in the Host Operating Systems that replicates every I/O on the Production Server to the CDP Server. This chews up Host bandwidth and also makes the offering OS specific. In fabric based CDP products, the intervening intelligent switches like MDS 9000 SANtap from Cisco take on this responsibility to replicate I/Os. Though this imparts OS agnostic-ness to the solution, handling transport failures become very messy. On the other hand, StorTrends offer complete target side CDT both at the production and/or recovery sites obliterating the need for any host side agents.

Figure 2 above depicts an iTX box configured to serve Production Server I/Os and at the same time offering EAS. Here two production volumes V0 and V1 are continuously protected by the EAS journal. In the event of disaster or for other data processing needs, a frozen image to any specific point can be recalled, reviewed and rolled back.

AMI StorTrends innovatively combines its low-latency snapshot technology with I/O journaling to offer most efficient CDT repository. Journal files can be configured from up to eight logical drives of any sizes to serve as the event-addressable-storage. To offer the utmost in performance, StorTrends allows configuring separate journal files for different consistency groups. In a typical environment where an average of 10GB production data is changed or generated per hour, a journal file of 1TB can hold four days worth of continuous chronological data. In addition, with eight levels of scheduled snapshots and up to 2000 snapshots per volume StorTrends can deliver any historic data across a very wide time spectrum.

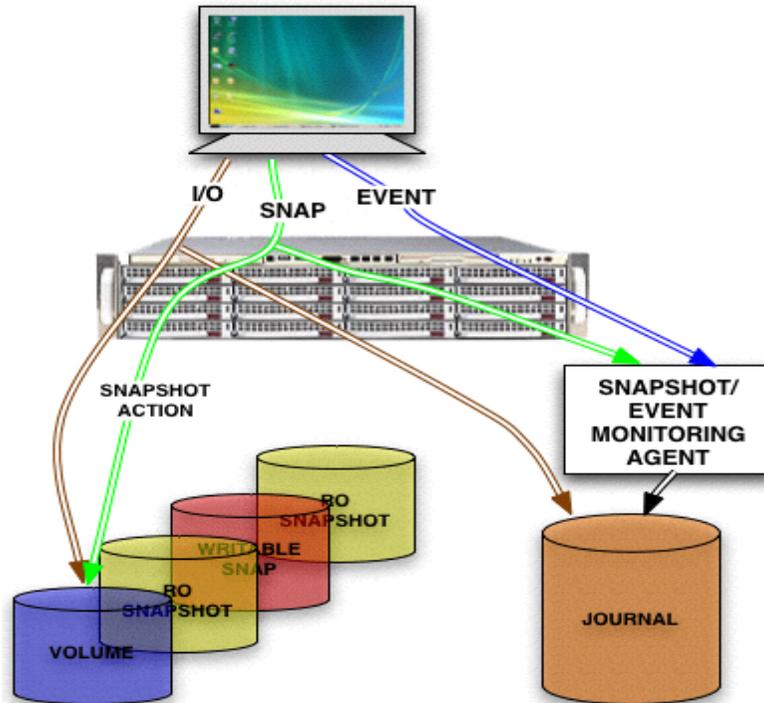


Figure 3: Data and event logging process as a part of Continuous Data Protection

Figure 3 above shows Continuous Data Capturing and the process of time-stamping I/Os and events in the Journal File to form the basis of Time addressable(TAS) and Event Addressable(EAS) storage.

Contemporary CDP Appliance Vs StorTrends CDP

In contemporary CDP appliances, the write I/Os are replicated outside the appliance using some sort of data-tapping agents. This mirrored I/O is then (1) staged in a buffer in the appliance. The old data is (2) read from the recovery volume, time-stamped and (3) journalled out into a "undo" journal device. Subsequent to this, the staged I/O is (4) written on to the recovery volume and finally the (5) operation is completed to the host..

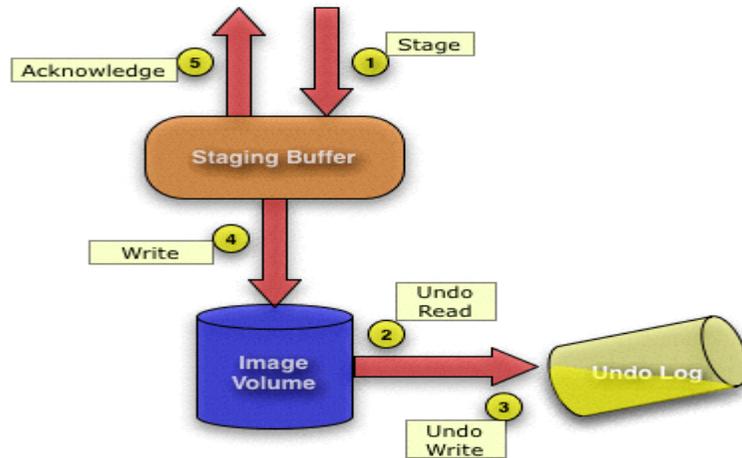


Figure 4: Data logging in a contemporary CDP appliance

During Review or Rollback, a Time image volume is created by replaying the undo journal from the current time to the selected time or event.

StorTrends CDP implementation uses an efficient and innovative “N-Way” I/O router module that can be configured in the run time to offer efficient I/O journaling by eliminating unnecessary data copies and reducing extra I/O overheads. This results in improved I/O latency.

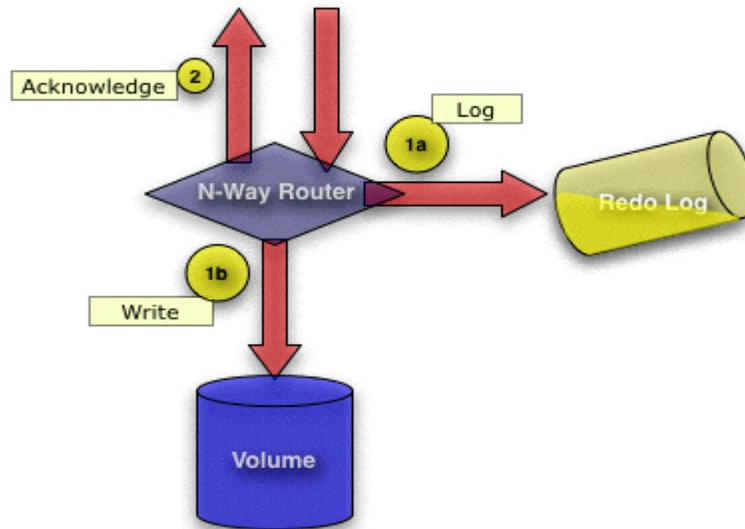


Figure 5: Data logging in StorTrends CDP Journal

During a Review or Rollback operation, initially a rollback to a previous writable snapshot is done and then the redo Journal is played on top of it to the selected point in time.

Under the hood: StorTrends CDT

The Advanced Caching Module of iTX stack serves as the central repository of in-flight data. In addition to many advanced caching functions, this module also implements a "N-Way" data router that allows incoming data to be multiplexed out to various data handling and transporting modules. This helps in effectively channeling data with Zero-Memory-Copy. Therefore, incoming data can be routed out to a remote recovery site (2A) as well as captured by the Journal Module (2B), as shown in Figure 4 below.

The Journalling module timestamps the arriving data and stores them in the journal files giving it the essence of Time-Addressable-Storage (TAS). This module also exposes an event registration interface through which relevant events can be tagged and stored. The figure shows snapshot events (C) generated by Distributed Volume Manager (DVM) And other application specific events e.g. database checkpoints (D) are recorded in the Event-Addressable journal.

At the heart of the technology resides the CDT engine. This works as the client of the Journal and snapshot modules. When historic data is requested, it gives a very fast access to the data using a very dense and efficient in-core index table. Powerful APIs and CLIs are offered to locate, review and rollback the requested image.

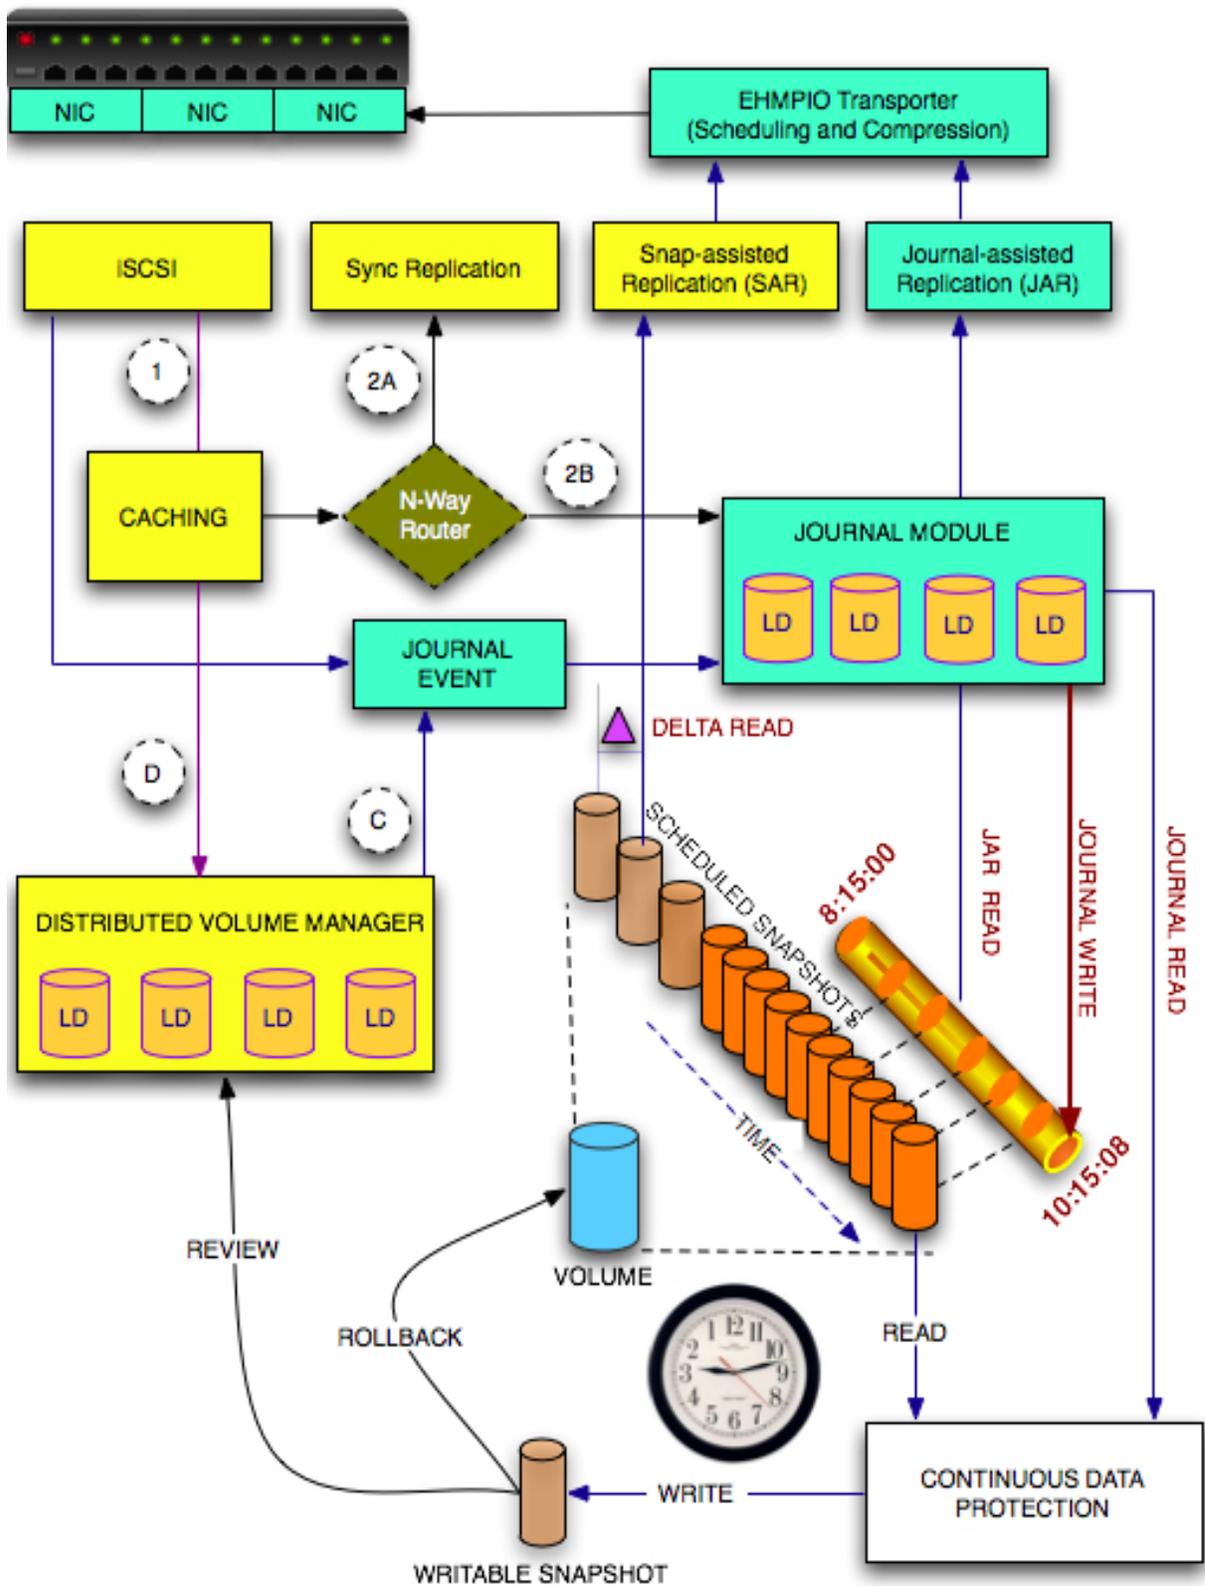


Figure 4: StorTrends Continuous Data Protection architecture

StorTrends Advantage

- **Seamlessly integrated in to Target stack**
- **Supports multiple Journal clients**
- **Consistency group Journals**
- **Multiple performance enhanced Journal Files**
- **Online expansion of Journal Files**
- **Time and Event Addressability**
- **Almost instantaneous creation of APIT image**
- **Creation of multiple APIT events for different types of access**
- **Backward and forward navigation by Events, Seconds and I/Os**
- **Filtering by types of events**



StorTrends

StorTrends iTX 2.7 CDP User Guide

AMI supports Continuous Data Protection (CDP) through a Journaling module that provides Time Addressable Storage (TAS) as well as Event Addressable Storage (EAS). So all volumes / volume groups that have been enabled for Journaling will be available for CDP. Besides, the CDP UI provides a way to consistently review and recover volumes of a Consistency group. For SAN volume / volume groups that are not enabled for Journaling, the UI provides a mechanism to group them logically during a review process and recover data across all the volumes together.

The CDP Utility can be configured and utilized as follows.

- The CDP utility window can be accessed from the 'Storage Pool Management' window by selecting any container. Click on 'Continuous Data Technology' (CDT) on the top left corner of the window (see Figure 1).



Figure 1: Navigating to CDP Management

- This takes us to the CDP control window. There is an option to either select an entire consistency group / journal file or individual volumes. By default, the 'Start Time' is the time of creation of the container and the 'End Time' is the current time. These can be modified as per the requirement to restore data using CDP.

Note that the above list of CGs / Volumes pertains only to the currently selected container. If you want to perform a CDP on a different CG / Volume, you need to go back and select the corresponding container.

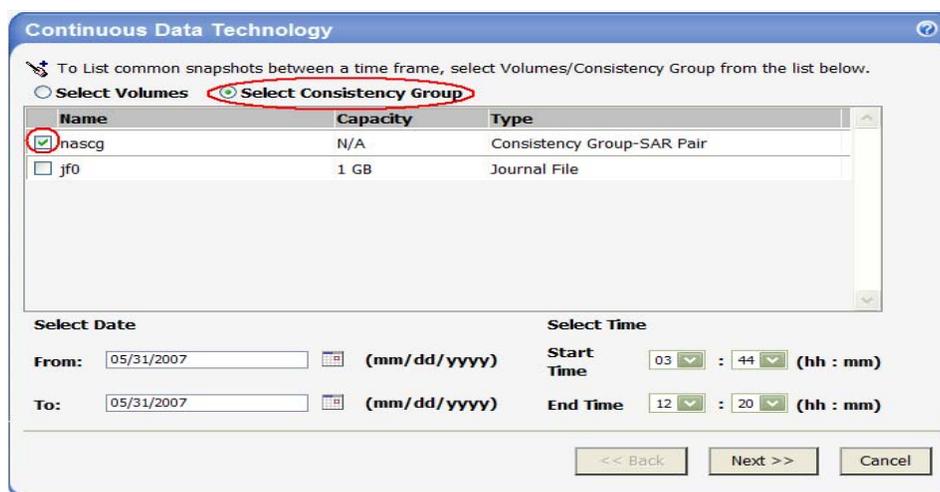


Figure 2: CDP Control Window

There are two interfaces available in ManageTrends to perform the CDP: The Flashback window and the Retrospect window. ManageTrends automatically selects the interface depending on whether a CG/journal or volume is chosen and whether data is available in the selected time zone. For example, if the user chooses a journal file and data is not available in the selected time zone but snapshots are present, then it will take the user to Retrospect (described later in this manual).

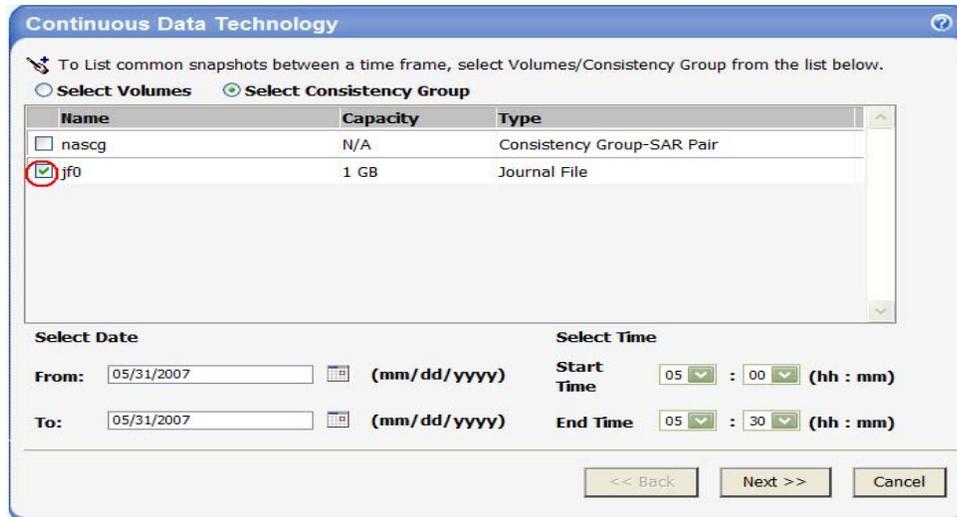


Figure 3: Example of a selected CDP window

We will first look at the Flashback Interface. Figure 3 shows a selected CDP window. Note that a journal file has been chosen and the time zone is chosen such that journal data is available.



Figure 4: CDP Flashback Window – Log view

Figure 4 shows the 'Log View' of the Flashback interface. This interface allows the user to have complete control over the review & rollback process. The user can choose a point or event in time, review the data and rollback if satisfied. It also allows the user to control the review at the IO level. Apart from time, the Flashback interface also has an alternate view called the 'Event View'. This view gives a different perspective for the user to choose an appropriate event to review and

restore. It also gives a feel of the time by indicating the time of the day picturesquely. Most importantly it allows the user to freely review both in the forward and backward directions.

The Log view can be sub-divided into 4 different panels. Figures 5, 6, 7 and 8 describe each panel in detail.



Figure 5: Time Slider Panel

The 'Time Slider Panel' shows graphically, the events and snapshots that are journalled. The slider (1) can be draged to any point on the line with a resolution of 1 second. The red dots represent snapshots and the blue dots represent events that were logged. The orange region prior to the first snapshot indicates a zone where journal data is not available. Similarly, the blue region indicates the zone where data is available. The slider can be moved to any given time at which the system needs to be restored. The Main Control Panel is used to precisely control this recovery time.



Figure 6: Main Control Panel

The Main control panel displays the snapshot (1), event (2) and time (3) controls. The Time control increments the time by a default value of 1 sec. This time step can be changed from the 'Time Step Control Panel'. The number of IO's to be performed can also be specified in this panel.

The time step needed to precisely move the slider is controlled by the 'Time Step Control Panel'. The dial in the panel displays the current time step.

Finally there is the 'Progress & Review Panel' (Figure 8). This panel allows us to review and rollback to the desired time instant. The review will happen to the point at which the slider is pointing to.



Figure 7: Time Step Control Panel

Figure 9 shows an example of using the Snapshot Control. Pressing the left and right button moves the slider from one snapshot to another. Figure 10 and 11 show sliding using time and event control respectively.



Figure 8: Progress & Review Panel



Figure 9: Snapshot based sliding



Figure 10: Time based sliding



Figure 11: Event Based Sliding

Another view for this window is the 'Event View' which is revealed by clicking on 'Event View' on the left corner of the Flashback window (see Figure 11). The CDP utility makes a note of three different types of events taking place in the system: (a) Snapshots (b) Application events and (c) Windows events.

Application events are modifications that happen when an application accesses a database. For example, an event occurs when an application, such Oracle, performs a checksum or a transaction. Windows events are taken from the Windows log to enable the user to correlate the events in the Flashback interface. Clicking on any event gives the information of the process that resulted in that event.

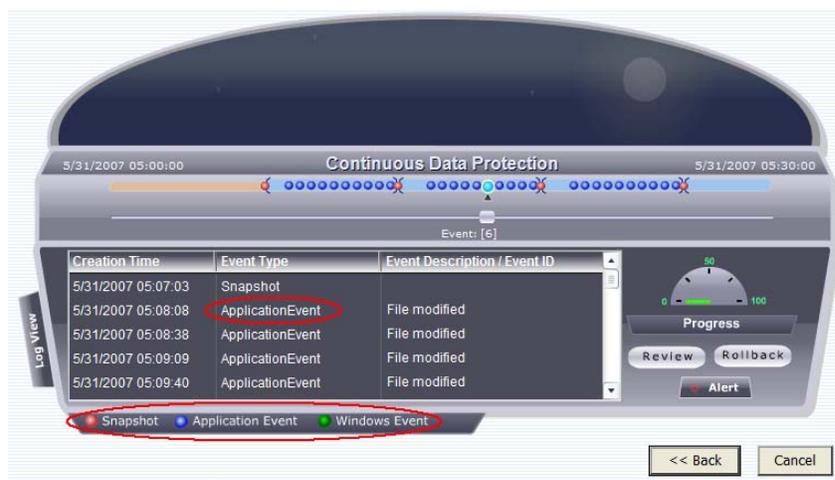


Figure 12: CDP Flasback Window - Event View

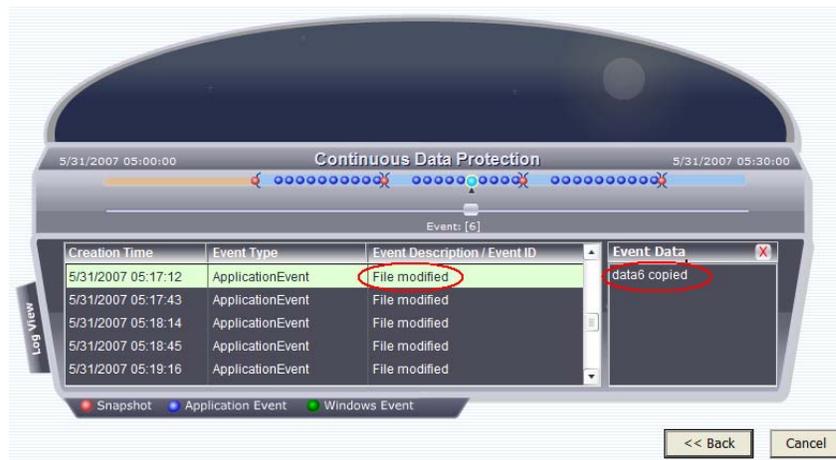


Figure 13: Event Description

This view enables the user to choose an event to review and rollback. We can switch between the Log and Event Views by clicking on the appropriate button on the left-hand corner of the window.

After choosing the recovery time or event, click on 'Review' to start the review process. This can also be done from the 'Event View' (Figure 12). The 'Progress' bar starts blinking and stops blinking when an indicator appears on the 'Time Slider panel'. The dial shows the progress of the review process. Once the progress is complete, the 'Alert' button flashes and the status can be known by clicking it (Figure 15).



Figure 14: Progress of Review



Figure 15: Completion of Review

When a review is successfully completed, a writeable image of the volume at the given time instant is provided for review. If a CG is being reviewed, a single target with all consistent volumes as LUN is created.

Connect your application server (MS Exchange, SQL etc) to this image and perform a consistency check to make sure that this is the image you wish to rollback. Once you are satisfied, logout of the target and perform a rollback operation. If you are not satisfied, logout and choose another review point and review again. This operation can be performed as many times as you want until you get a satisfied image (rollback point).

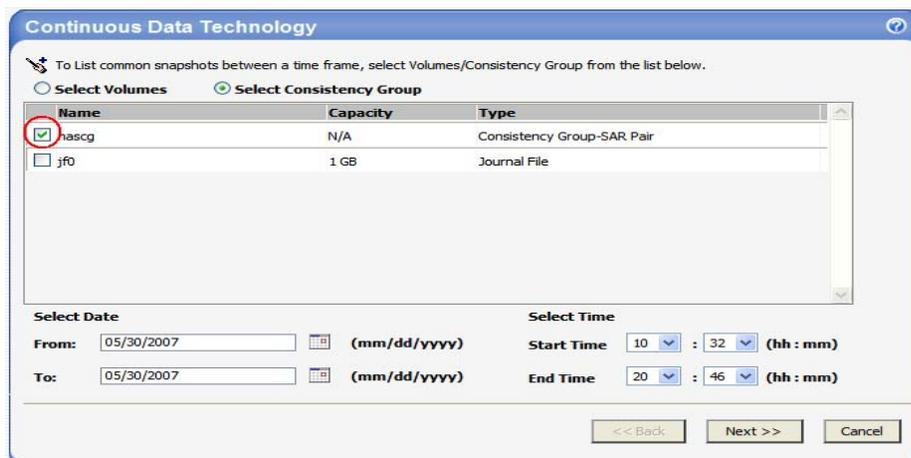


Figure 16: CDP Window to access the Retrospect window

Another CDP GUI is the retrospect window, which is used to restore individual volume data. Choosing a consistency group rather than a journal file from the CDP Window accesses this GUI.

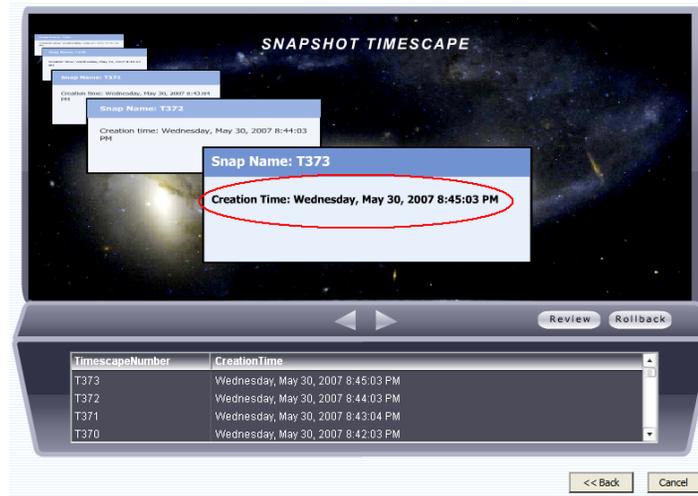


Figure 17: Retrospect Window

The 'Retrospect Window' GUI looks as shown in Figure 17. The startup screen shows latest snapshot that was created within the given time frame.

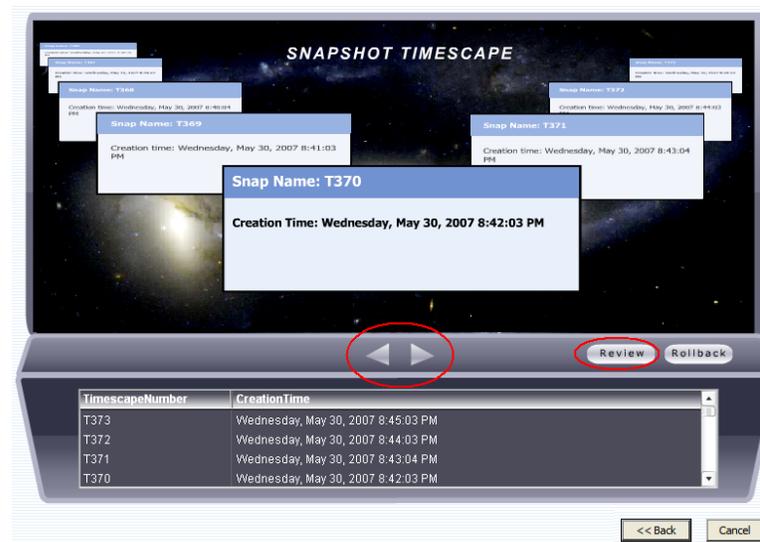


Figure 18: Scrolling Snapshots

Each flying window represents a consistent snap group taken across all the volumes of the CG. For a SAN CG, this is achieved using the VSS Scheduler and for a NAS CG, the NAS scheduler does the job. The scroll buttons can be used to access any snapshot. Clicking the 'Review' button starts the reviewing process.

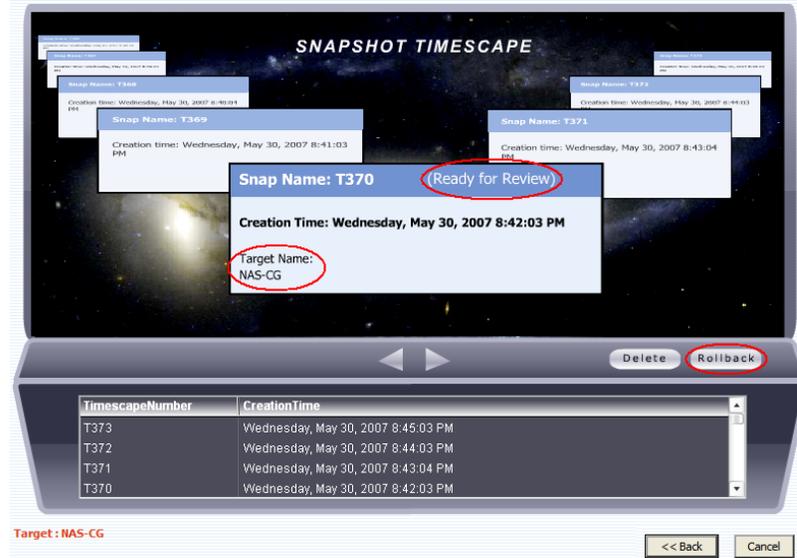


Figure 19: Reviewing Snapshots

Once the review process is completed, a writeable image of the volume is created and appears on the bottom left corner of the window and the status changes to 'Ready for Review'.

Connect your application server (MS Exchange, SQL etc) to this image and perform a consistency check to make sure that this is the image you wish to rollback. Once you are satisfied, logout of the target and perform a rollback operation. If you are not satisfied, logout and choose another review snapshot and review again. ManageTrends automatically deletes the previous target and creates a new one. This operation can be performed as many times as you want until you get a rollback point.