


Business and IT Requirements for Continuous Data Protection

Protecting Information Assets with Enterprise Rewinder



September
2004



Table of Contents

Executive Summary	3
Protecting Applications Today	4
Requirements for an Effective Solution	5
A High-Level Overview of the Enterprise Rewinder System	7
Normal Operation of the System	8
Managing and Monitoring the System.....	9
Digging Down: What Makes Enterprise Rewinder Distinctive	10
Seamless Integration With Your Environment	10
Ease of Management with Complete Control.....	11
Installation Without Pain	13
Real Application-Aware Protection From Corrupted and Inconsistent Information	14
Simple and Fast Recovery	14
Cost-Effectiveness and Verifiability	15
Summary of Enterprise Rewinder Advantages	16
Conclusion	17
Other XOsoft Products	17
About the Author	18
About XOsoft.....	18



Executive Summary

Software and user errors are estimated to cause 80% of all unplanned downtime. A significant percentage of these errors occur predictably, during database maintenance or batch processing activities. Traditional data protection systems like RAID, tape backups and snapshots provide a critical layer of insurance, but they represent all-or-nothing strategies that cannot prevent the loss of potentially hours of valid, valuable data prior to an error. New continuous data protection systems fill in this gap by allowing data to be restored to an arbitrary point in time, thus providing a solution that is well targeted to the actual problem. An adequate continuous data protection and recovery system must, however, meet several key requirements, including simplicity, cost-effectiveness, application-awareness, flexibility and high performance. We present a detailed description of the Enterprise Rewinder architecture and operation, demonstrating that it has been carefully engineered not just to meet the minimum requirements, but even to set new standards for ease of operation and management. In addition, Enterprise Rewinder offers some unique capabilities, including installation without any downtime and strong application-awareness.



Protecting Applications Today

The days when a daily backup provided adequate protection for a business's critical processes are long gone. Today, downtime, *any* downtime, translates to real and significant costs, in lost productivity because employees can't work, lost business because orders can't be taken, and customer dissatisfaction as your downtime translates to their downtime. A recent study of 80 large organizations by Infonetics Research found that overall downtime costs averaged an astounding 3.6% of annual revenue. Can you afford that kind of bite out of your margins?

Well over half of all downtime is unplanned. Of that, 80% is due to software and user errors: unknown and irreproducible bugs, accidental deletion or configuration changes, batch jobs corrupting important data, and errors from the use of hard-to-understand management tools. More importantly, many of these occur at predictable times, for example, during database maintenance or batch processing runs. What about a solution that actually leverages this knowledge?

There are, of course, a variety of more or less traditional solutions available to protect data, solutions like RAID, tape backup, and snapshots. None of these solutions is well suited to the situation we are interested in, however. RAID is no solution at all – if the database becomes corrupted, RAID merely ensures that the resulting garbage is safe from disk failure. Tape backup and even snapshots are all-or-nothing solutions – they cannot prevent the loss of potentially hours of valid, valuable data prior to an incident. Even a good online replication system requires resynchronization from a backup server and is likely to replicate corruption to the replica server as well. All of these measures have their place in an effective overall application protection strategy, but none address the specific problem as well as they might.

In fact, a November 2003 Gartner report emphasizes that by 2006 the focus of data protection technologies will be on recovery time and that traditional backup technologies and replication-based technologies will merge into a single market by 2005. Continuous data protection technology *does* provide a solution suited to the problem, in particular, a continuous data protection capability that provides the ability to quickly roll back data not just to the last snapshot or backup, but to *any point in time*, so that only operations actually involved in the crash or corruption need be lost. This kind of capability, which can be enabled during operations known to be high-risk, is precisely the kind of targeted, high-leverage solution that is critical to modern IT strategy. Not every system is adequate to the task, however. In the next section, we review the requirements that an effective solution must satisfy. In the remainder of the paper, we first review the Enterprise Rewinder solution in some detail, and then highlight the features that address the requirements. Finally, we briefly reiterate how Enterprise Rewinder fits into a broader risk mitigation strategy for your organization.



Requirements for an Effective Solution

When selecting a continuous data protection product as part of an overall application protection strategy, it is important to remember that a technology is not itself a solution to a business problem – it is only, at best, a component of a solution. The failure to understand this distinction has led to enormous losses of resources invested in IT over the last decade or so – companies have purchased one technology after another without fully understanding how the technology would actually provide a solution to the problems they faced.

In this section, we consider the characteristics that a continuous data protection product must have in order to actually solve the business problem, i.e., to achieve targeted, significant risk mitigation for business continuity at an acceptable cost. Throughout, we hold to the overarching principle of *real-world focus*. Real-world focus means that any solution must be well matched to the problems, IT environment, and personnel that your company actually has, not to those that are easiest to address. Without this focus, products are as likely to create as many new complications and problems as they solve.

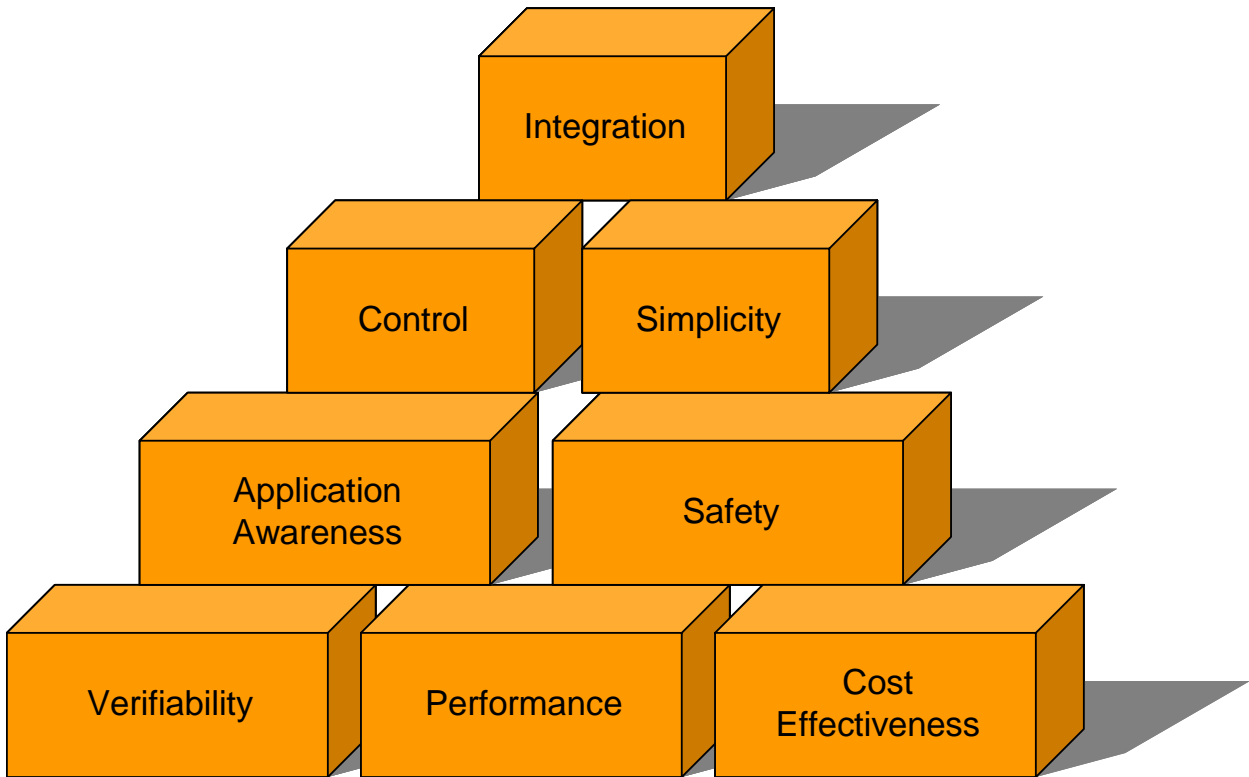


Figure 1. Critical characteristics of an effective data protection solution.

With this broad principle in mind, we suggest that the following represent the most important characteristics that an effective continuous data protection solution must possess:

Integration. Your company should not have to tailor its environment to the needs of the continuous data protection software – a true enterprise solution is the other way around. It must integrate well with existing organizational processes and IT infrastructure. It must



also support the operating systems and configurations that actually exist in your environment.

Control. Critical data protection solutions like continuous data protection must support enterprise-wide management and reporting from a single central point of control, as well as allow distribution of the management and monitoring responsibilities, in order to support the administrative needs of your organization. It must also allow administrators to set bookmarks at times of interest or a known state in order to maintain complete control over the state to be restored.

Simplicity. Today's IT environment, even in organizations of moderate size, is tremendously complex, and no solution should add unnecessarily to that complexity. In practical terms, this means that the product must be simple to configure and install, simple to use and maintain, and it must interact in simple, predictable ways with the rest of the environment. Too often, software products solve a problem that is convenient to solve rather than the problem that the customer has – and the burden of dealing with the misfit falls on the customer. A true solution will *keep the burden where it belongs*.

Application Awareness. Your company is not seeking to protect a random mass of bits. It is seeking to protect applications – Exchange, Microsoft SQL or Oracle databases, or network file services – and the information associated with them. A true solution will aid the administrator in identifying exactly what must be protected and how, and will be sufficiently aware of each application to ensure that recovery yields a fully consistent state of the application.

Safety. A solution that allows the administrator to restore a database to an inconsistent state is no solution at all. A true solution must provide the ability to easily find and roll back to the latest consistent and uncorrupted state. It must also allow the administrator to bookmark times when the application is in a known state for later recovery. Finally, it must provide adequate internal security and align with the rest of your company's security structure.

Performance. A true solution will strive to minimize overhead at every turn so that the application being protected continues to perform as needed. Equally important, recovery should be fast since the whole point is to improve on other methods of protection.

Cost-Effectiveness. The purpose of a continuous data protection solution is to prevent cost, not to be one. The return on investment in a true solution will be clear. In particular, the total cost of the system, including both the price of purchase and support and the effort required to maintain it, should be clearly related, not to the potential cost of a worst-case scenario, but to the probable cost of the most likely one.

Verifiability. Product claims, even backed by solid testimonials, are not sufficient evidence that the product can effectively protect your applications in your environment – there must be an opportunity to verify the product before purchase.



A High-Level Overview of the Enterprise Rewinder System

Continuous data protection is a relatively new technology, unfamiliar to most, so let us take some time first to explain just what it is and how it differs from other solutions.

Traditional backup freezes the state of a set of data at some specific single point in the past. It provides the ability to return a system to that past state. Replication, on the other hand, tracks the data as it changes. It provides the ability to return a system to the state that existed just prior to a crash.

Continuous data protection offers the best of both worlds. Like replication, it continuously tracks changes in the data. However, instead of simply copying the changes to a replica, it records a local *rewind journal* of the operations required to *undo* the changes that are taking place. Should recovery be required, the continuous data protection and recovery solution simply plays the undo operations backward to any desired point in the past that is recorded in the journal. The rewind ability is analogous to undoing operations in an editor like Microsoft Word.

Enterprise Rewinder provides protection that is tailored and optimized for specific application servers, including Microsoft SQL Server, Exchange 5.5 and 2000/2003, Oracle, network file servers, and others. It is engineered to provide efficient continuous data protection and recovery capability in a way that integrates well with your environment, provides efficient means of recovery, and ensures the integrity of any recovered data.

Imagine, for example, that workgroup-level corruption of your Exchange database occurs during system maintenance. With traditional backup or even replication technology, the entire database would have to be restored. With Enterprise Rewinder, it is necessary only to rewind the individual workgroup-associated files to a time just prior to the corruption. Because of its application-awareness, it is simple to rewind not just to a time before the incident, but also to a time at which the entire database system is known to be in an internally consistent state, not necessarily a trivial thing to do if entire individual files must be restored from a copy.

The architecture of Enterprise Rewinder is quite simple, consisting of three modules. The core of the system is the **XOsoft Enterprise Rewinder Engine**, a Windows service or Unix daemon that runs on each server to control and monitor the work, and communicate with the management interface to administer the system. The Rewinder Engine also installs **XOFS**, Enterprise Rewinder's file system filter driver that captures all changes to data as they occur and prepares the rewind journal. Finally, the **Enterprise Rewinder Manager** may be run from anywhere in the network to administer and monitor the system.

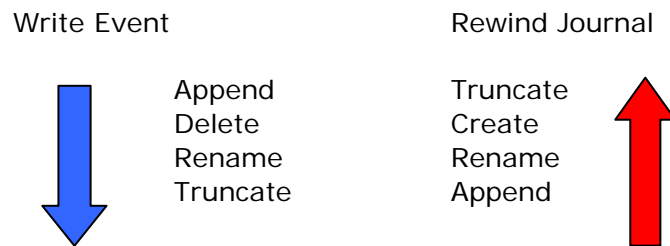
In the remainder of this section, we will cover the operation of Enterprise Rewinder at a fairly high level. We will then turn in the following section to a consideration of some detailed characteristics that distinguish Enterprise Rewinder from other tools on the market.

Normal Operation of the System

In normal operation, the system running on the server to be protected must capture all updates to data of interest and record the operations required to undo them in a rewind journal. It is obviously critical that updates be captured *in precisely the order in which they occurred*, or the rewind journal will likely lead to corruption of the system.

As noted above, there are two system components are involved on the server, XOFS and the XOsoft Enterprise Rewinder Engine.

XOFS is a low-level driver that is installed through a standard mechanism to act in concert with the file system on the server. Once running, XOFS is able to see every I/O operation on the file system when it occurs and to copy information on any operation that changes a file or directory included in a scenario for which it is responsible – all such information is copied in precisely the order in which it occurs for recording in a journal. Unlike normal replication, however, Enterprise Rewinder does not record the operation itself, but rather the *inverse operation*. In other words, it generates the operation required to undo what was just done, as shown in the figure below. Thus, for example, if an operation appends 8K bytes to the end of a file, XOFS will record an operation to truncate the file by 8K bytes, and similarly for every other kind of operation.



When the journal is replayed during a recovery, from back to front, the effect is to cancel each of the events, one by one, in the proper order, in order to restore the file to the state at any previous time.

The XOsoft Engine controls the process – it is responsible for starting and stopping XOFS, providing it with lists of files and directories to be monitored, and passing on bookmarks defined by the administrator.

The entire process is both simple and highly efficient in order to maximize the performance and reliability of the recording with minimum impact on the performance of the application.



Managing and Monitoring the System

As the primary means of interaction between the system administrator and the Enterprise Rewinder system, the Enterprise Rewinder Manager is a critical part of the overall solution. The Manager is a Microsoft-certified management interface that is intuitive, simple, and straightforward to use. Indeed, most Enterprise Rewinder administrators find that they do not actually ever need to use the user's manual.

At the same time, the Manager is extremely powerful, giving the administrator complete control over all continuous data protection processes across the enterprise and a clear view into the working of the system.




Figure 2. All data center rewind journals and activity is monitored from a central console.

The fundamental basis for operation of the system is the *scenario*, a structure describing what information is to be protected (the source directories and files), where it is located (the application server to be protected), and how it is to be protected (the parameters controlling recording of the rewind journal). Multiple independent scenarios may run on a single server.

The Enterprise Rewinder Manager provides centralized administrative control to set up scenarios, load them onto the appropriate servers, and set them running. If desired, multiple instances of the Manager may be run independently by different administrators, although only a single administrator may control any given scenario at one time. System security can be provided either through Enterprise Rewinder's internal password-based mechanisms or integrated with your site's normal Windows security.



The Manager is highly application-aware and able to aid the administrator in locating appropriate files and setting up optimized parameters based on application characteristics, as described in detail later in this paper.

The Enterprise Rewinder Manager, together with the XOssoft Engine and XOFS, also provides monitoring capabilities, including logging within the Manager, alerts through email, and writing information to standard system logs.

Digging Down: What Makes Enterprise Rewinder Distinctive

Now that we have a big-picture view of the system – what it does and how it works – let us turn now to consider some of the features that make Enterprise Rewinder stand out from other solutions.




Figure 3. Enterprise Rewinder management console.

Seamless Integration With Your Environment

Any company of even moderate size already has more than enough to do maintaining its IT systems – and it wants to arrange protection of the systems it actually has, not provide systems the solution vendor would prefer that it have.



Enterprise Rewinder was engineered with this as a central design principle. As a consequence, it integrates with your environment in several important ways.

First, for Exchange 2000/2003 environments, Enterprise Rewinder is already integrated with Veritas Backup Exec, and integration with additional backup and snapshot solutions is underway. With this integration, the behavior of Enterprise Rewinder may be tied directly to the beginning or end of a backup on a particular storage group or set of storage groups.

Next, as further detailed below, Enterprise Rewinder requires no reconfiguration of the servers to be protected whatsoever. In addition, Enterprise Rewinder is available for a variety of applications, including Exchange 5.5, Exchange 2000/2003, Microsoft SQL, Oracle and file servers.

Integration with internal security is obviously a critical issue as well. Enterprise Rewinder supports and integrates fully with the standard Windows security models.

Finally, alerts generated in the system may be handled in a variety of ways. They may be routed to the Enterprise Rewinder Manager for display, recorded in standard system logs, or routed externally via email. In addition, they may be passed to an arbitrary administrator-defined script to be integrated with other corporate system alert systems.

Ease of Management with Complete Control

We have already described the inherent simplicity and intuitive nature of the Enterprise Rewinder Manager. However, the ease of management and control provided by the system goes well beyond this.

The Enterprise Rewinder manager allows centralized administration of every continuous data protection scenario across the entire enterprise, while also allowing administration and monitoring to be distributed however is appropriate for your organization.

The scenarios used as the basis for configuring the continuous data protection were specifically designed to support protection of the application, rather than the server. For this reason, a key early parameter in configuring a scenario is the selection of the type of application being protected. The Enterprise Rewinder Manager's application-aware design provides auto-discovery of all types of database files, thus offering a powerful facility to aid the administrator in selecting *all* the necessary files for protection, while still providing complete flexibility in defining what files are to be protected through easy-to-use file browsers and filters. An example of Exchange database auto-discovery is shown below.

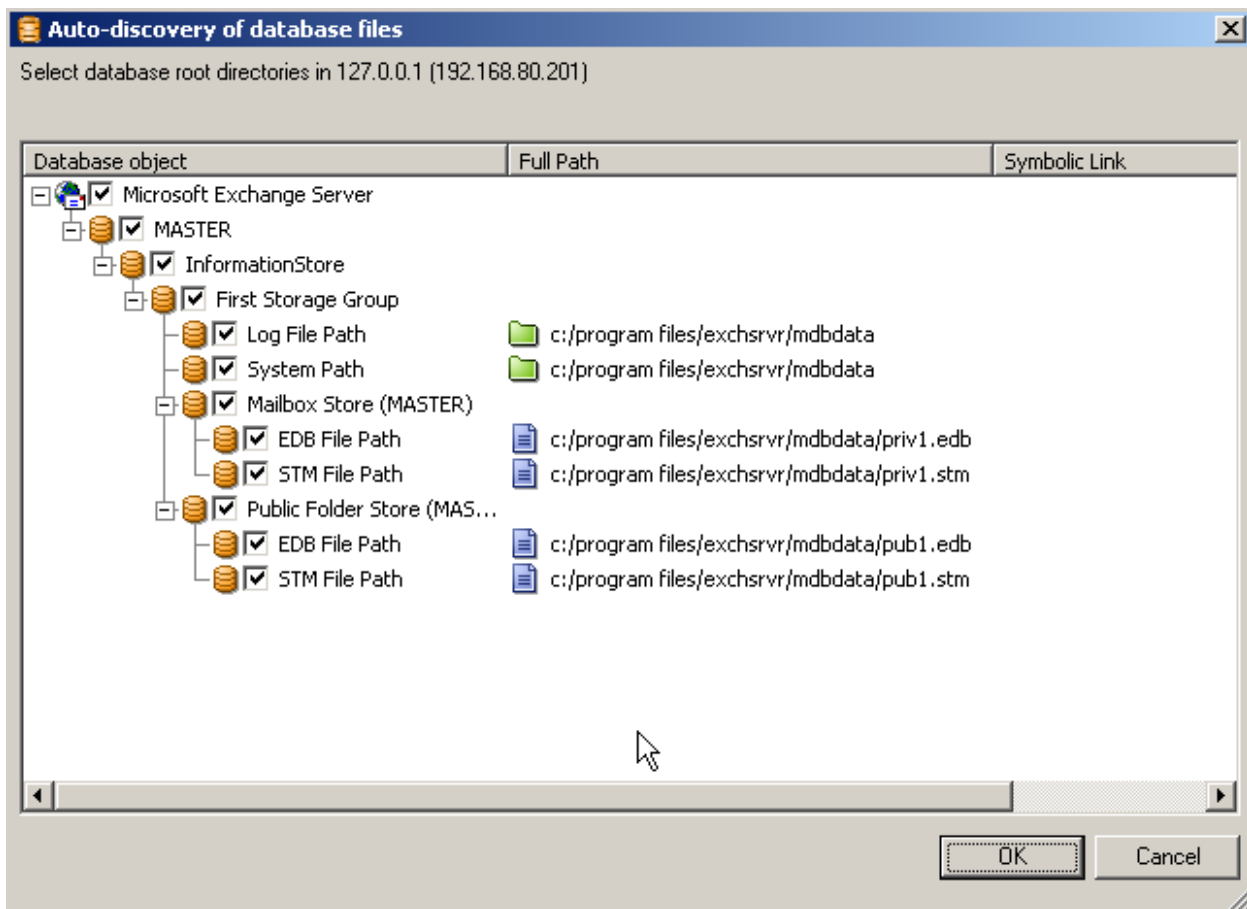


Figure 4. Auto-discovery of database files for ease of set-up.

The Enterprise Rewinder application also provides significant flexibility to tailor parameters to a particular environment and application. In most tools, access to dozens or more configurable parameters is hardly a favor to the administrator – at best, they enable tailoring of the system to a particular environment by specialists from the vendor. In Enterprise Rewinder, however, the most important parameters depend primarily on the characteristics of the application being protected, including Microsoft Exchange, Oracle and Microsoft SQL Server. Here again, the Manager’s awareness of these characteristics enables it to provide administrators with application-tailored defaults for all parameters in the system automatically.

For system monitoring, the Enterprise Rewinder Manager provides complete visibility into the working of the Enterprise Rewinder system. It can act as a collector for a configurable set of events that can occur anywhere in any scenario running on the system, from fatal errors to suspicious events to simple progress reporting. In addition, the status of any system running a scenario under the Manager’s control is visible in the interface.

The Manager provides a simple-to-use facility for browsing rewind journals either in raw form, should it be necessary, or in a form that condenses the raw I/O event data into a more intuitive format that displays the names of files affected and collapses long sequences of operations into simple open-write-close sequences. Most importantly, the Manager displays special events for administrator-defined bookmarks and for points at

which a database is automatically determined to be in a fully consistent state for use in performing recovery.

Finally, a comprehensive command-line interface (CLI) is also available for administering the system.

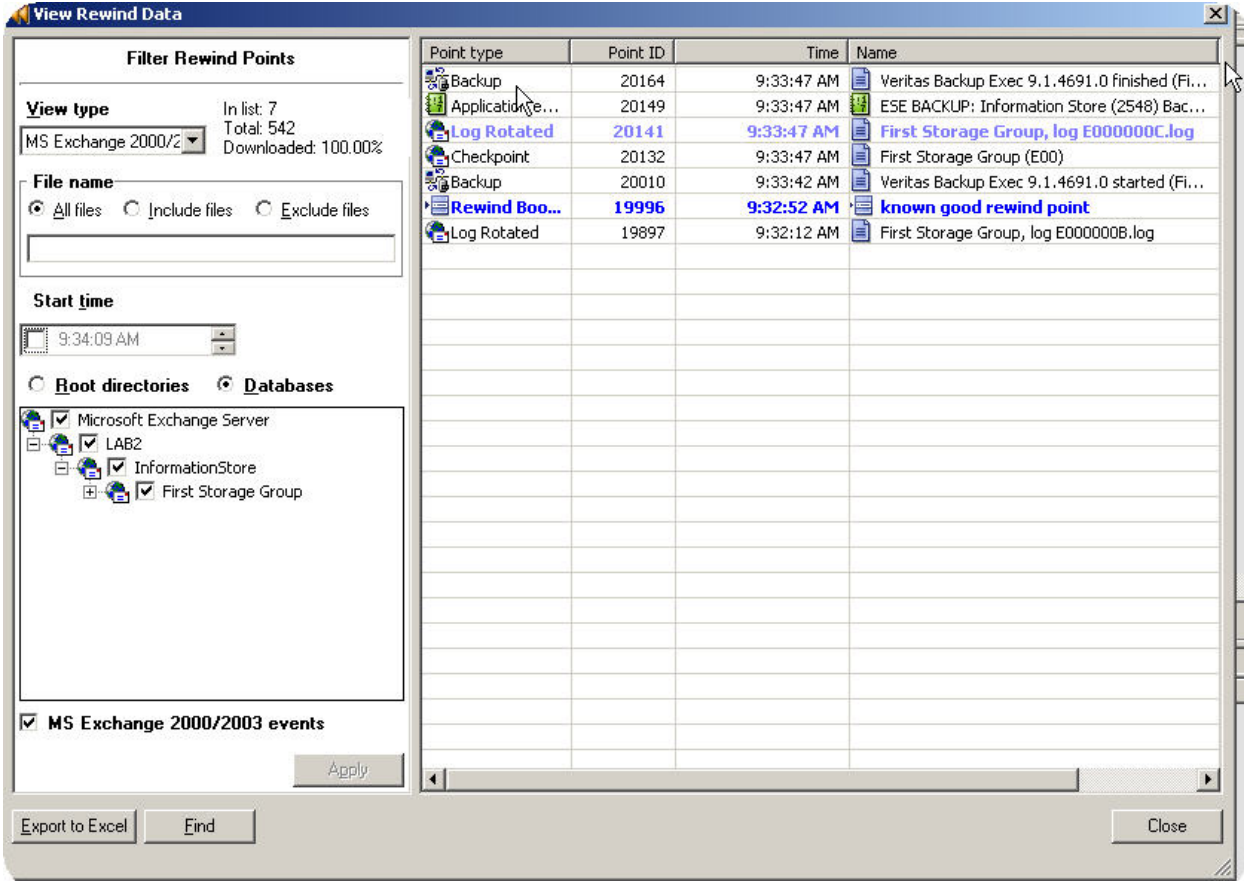


Figure 5. List of database- and environment-specific events for rewind.

Installation Without Pain

Installation of software that must integrate closely with the operating system is not normally a topic to be emphasized in a white paper. Installation is typically one of the messier parts of the system, often involving complex procedures for configuring the machines and software, and usually requiring that all applications be stopped and that the machines be rebooted a time or two. Put briefly, a fair amount of pain and downtime, precisely the opposite of what the solution is intended to provide.

In the case of Enterprise Rewinder, in contrast, installation is a part of the system to be highlighted since it illustrates XOssoft's commitment to keeping the burden for the solution on the software, where it belongs.

Actually installing the software is, of course, no more complicated than installing any application program – insert a CD, follow a couple steps in a wizard, and you are done.



Unique to Enterprise Rewinder, however, is its patented ability to set up and start a continuous data protection immediately on a live, actively running system. Specifically, Enterprise Rewinder installation is **soft** and **completely live**: there is *no need to reconfigure* any of the servers, there is *no need to reboot* any of the servers, and there is *no need to stop the applications* whose data is to be protected. Enterprise Rewinder automatically plugs itself into the running system and even finds and links to files that are open and in active use in order to begin replication.

Real Application-Aware Protection From Corrupted and Inconsistent Information

With most important applications, particularly databases like Exchange, Microsoft SQL Server, and Oracle, the application may not be valid at just any arbitrary point in time – the image on disk can easily be in an inconsistent state because the application is in the middle of a transaction. Similarly, rewinding to a completely arbitrary point in time may well leave the application data in an inconsistent state.

For this reason it is critical that a continuous data protection solution take this into account – it must be able to distinguish between valid and invalid states of the data and to ensure that the administrator is able to recover a valid one.

Application-awareness is a central principle of Enterprise Rewinder's design – as a result, it is able to track when the database is consistent and when it might not be. During recording, Enterprise Rewinder makes special entries in the rewind journal at points where the data is known to be consistent – entries that are visible to the administrator through the administrative interface when it is necessary to perform recovery.

In addition, Enterprise Rewinder allows the administrator to set a labeled bookmark with the simple push of a button. Should something go wrong, the bookmark is easily accessible through the Enterprise Rewinder Manager as a point for recovery.

With these capabilities, recovery to a valid, consistent state of the database becomes straightforward – the administrator need only search in the administrative interface backward to the latest valid state before the crash or corruption took place.

Simple and Fast Recovery

While it is never desirable that things go wrong, dealing with recovery from a system crash or corruption is, of course, the whole point of the Enterprise Rewinder system. The questions that must be asked are: when disaster does strike, how difficult will it be to recover and how long will it take?

With Enterprise Rewinder, the procedure is completely straightforward and quite fast. The administrator need only stop the application and then use the Enterprise Rewinder Manager to rewind the data to one of the points of known consistency displayed in the rewind journal, either an automatically labeled consistent state or one of the bookmarks set earlier by the administrator, a matter of no more than a few minutes. Once this is done, the application can be restarted and operations continued normally.



Cost-Effectiveness and Verifiability

The costs of downtime or data loss obviously vary from one company to another, however, studies indicate that they are typically higher than many are aware. As noted in the introduction, Infonetics Research found that overall downtime costs reach 3.5% of revenue. Research by Creative Networks, Inc. has shown that Microsoft Exchange downtime can easily cost thousands of dollars per hour in a mid-sized organization. Thus, Enterprise Rewinder's price is actually *less* than the likely cost of a single incident. Given the high likelihood of at least one incident in a year, the cost will almost certainly be recouped within the year. In fact, because Enterprise Rewinder can be installed on a live, active system, it immediately avoids one downtime incident.

In addition to the cost benefits of soft installation, Enterprise Rewinder provides several other sources of potentially significant cost advantage. These include its ease of use and the fact that no configuration changes are required, two features which translate directly into savings on maintenance overhead. Further, Enterprise Rewinder requires no investment in third-party software, has minimal overhead, and, of course, provides easy and immediate recovery after an incident.

Finally, Enterprise Rewinder, like all XOsoft products, is available for download with a 14-day license for testing purposes, allowing customers to verify that the solution will actually work for them *before* committing to purchase.



Summary of Enterprise Rewinder Advantages

The previous sections have described in quite significant detail how Enterprise Rewinder works and what distinctive advantages it offers as a data replication solution. Here we revisit our earlier list of characteristics that any effective data replication solution must possess and see how Enterprise Rewinder's features measure up.

Integration and Flexibility

- Use of scenarios to perform replication only on desired data.
- Integration with Veritas Backup Exec
- Optimized support for a variety of critical applications on all server configurations, including clusters.
- Ability for administrator to set bookmarks at points of known state and to rewind to them, if a problem occurs [see Safety as well].
- Ability to configure spool size for storage of rewind journals at any time, even during recording.

Control

- Centralized management of all scenarios across the enterprise.
- Ability to delegate management and monitoring.
- Straightforward and powerful management interface (a comprehensive CLI is also available).

Simplicity

- Soft installation, without any need to reconfigure the servers involved.
- Patented capability for live installation, without any need to reboot the servers or even to temporarily shut down the applications to be protected.
- Seamless integration with the environment, including security, bandwidth management, and alarm-reporting mechanisms.
- Intuitive interface with powerful application-specific aids.

Application Awareness

- Ability to recognize and record points at which databases are in a consistent state.
- Ability to configure scenario parameters to the application targeted by the scenario.
- Auto-discovery of database files to relieve administrators of the need for detailed knowledge of database storage architecture.

Safety

- Ability to recognize/record points at which databases are in a consistent state.
- Ability for administrator to set bookmarks at points of known state and to rewind to them, if a problem occurs.
- Ability to fully integrate with security environment.



Performance

- Careful design to support rapid recovery.
- Engineered to minimize overhead on the server.

Cost effectiveness

- Clear return on investment with price keyed to typical cost of a single recovery incident.
- Ease of use and setup reduces administrative overhead cost.
- No investment needed in third-party software.
- Initial installation requires no downtime.
- Easy and immediate recovery after an incident.

Verifiability

- Ability to download for free 14-day trial.

Conclusion

Effective use of scarce IT resources demands targeted, high-leverage solutions that mitigate real, identified risks.

In this paper, we have outlined the characteristics that such a solution for protecting your organization's most critical applications must have – any product that lacks these characteristics cannot provide the level of protection that today's pace of business requires. Enterprise Rewinder not only possesses the necessary characteristics, but also far exceeds the minimum requirements in several ways.

We invite you to download Enterprise Rewinder for trial at www.XOsoft.com, or to contact a XOsoft representative at 1-866-WANSync. Recorded technical demonstrations and webcast are also available on the XOsoft website. We look forward to hearing from you.

Other XOsoft Products

XOsoft offers several other products to protect access to your critical data, as well as to add value through fast and completely flexible content delivery. Please check out our website or contact a XOsoft representative for more information.

WANSync HA
WANSync for Disaster Recovery
WANSync for Content Delivery



About the Author

Eric Jackson is Vice President of Products at XOsoft. Mr. Jackson has nearly 20 years experience in the development and commercialization of advanced software technology.

In 1986, he joined the staff of the Applied Math program at Princeton University. His work there focused on the design of software to solve complex turbulent flow problems on advanced supercomputer and parallel supercomputer architectures and research into turbulent flow phenomena. Mr. Jackson has numerous publications in scientific journals, including *Nature*, *Physical Review*, *Physical Review Letters*, *Proceedings of the Royal Society of London*, *Proceedings of the SPIE*.

In 1994, he moved to a private consulting firm in Princeton. There, collaborating closely with researchers at leading computer chip manufacturers such as LSI Logic, Toshiba, DEC, IBM and others, he designed and implemented one of the first and most sophisticated software applications to automatically analyze and correct computer chip designs to improve their manufacturability.

In 2000, he co-founded Ibrix, Inc. He is the inventor of the Ibrix distributed file system, a parallel file storage system able to scale in size and performance to millions of terabytes. Mr. Jackson served as chief software architect and led the product development effort through the initial phase of the company's existence.

In 2003, Mr. Jackson founded DeepWeave, a management consulting practice helping leading-edge software companies with technology commercialization and strategic product development.

About XOsoft

XOsoft's continuous application availability solutions keep enterprises working. The company's patented business continuity, disaster recovery and continuous data protection software platforms significantly lower the risks of decreased productivity and lost revenue by minimizing both planned and unplanned critical application downtime. Many of the world's largest corporations and government agencies have chosen XOsoft to ensure uninterrupted access to their most valuable information assets. For more information, please visit www.XOsoft.com.

A 14 day trial evaluation of the software is available for download at:
<http://www.xosoft.com/download/index.shtml>