



# Systems Recovery Guide

## Index

<i>Media based backup and restore</i>	<i>Page 2</i>
<i>Cold Standby System</i>	<i>Page 2</i>
<i>Cold Standby with Data Replication</i>	<i>Page 2</i>
<i>Hot Standby System</i>	<i>Page 3</i>
<i>How long does each method take to recover?</i>	<i>Page 4</i>
<i>How much data will be lost?</i>	<i>Page 5</i>
<i>Implementing a hot standby system using SteelEye LifeKeeper</i>	<i>Page 5</i>
<i>Proactive Protection of applications</i>	<i>Page 6</i>
<i>Shared storage</i>	<i>Page 7</i>
<i>Replicating data where no shared storage is available</i>	<i>Page 7</i>
<i>Failover and Data Replication</i>	<i>Page 8</i>
<i>Backup and Data Replication</i>	<i>Page 8</i>
<i>Disaster Recovery to a remote site</i>	<i>Page 9</i>
<i>Application Recovery with LifeKeeper</i>	<i>Page 9</i>
<i>Recovery of standard applications</i>	<i>Page 10</i>
<i>Bespoke Application Recovery</i>	<i>Page 10</i>
<i>The Generic Application Recovery Kit</i>	<i>Page 11</i>
<i>Customisation Kit</i>	<i>Page 11</i>
<i>Components of an ARK</i>	<i>Page 11</i>
<i>Is the cost of a Hot Backup justified?</i>	<i>Page 11</i>
<i>About Open Minds High Availability Solutions Ltd.</i>	<i>Page 12</i>
<i>About SteelEye Technology Inc.</i>	<i>Page 12</i>
<i>About This Document</i>	<i>Page 12</i>

## Abstract

Due to the pervasive nature of IT in today's corporate infrastructure, and the increase in Intel server deployment, more organisations than ever before consider their business at risk if their servers are not available.

Immediate systems recovery is no longer just required by those with huge budgets and large systems. The demand for systems recovery solutions has come to the Intel server market.

This document is meant as a guide for those considering the next generation of systems recovery solutions for their servers.

## Media based backup and restore

Although media based backup and recovery solutions are very sophisticated today. A backup is rarely done more often than daily and is prone to human error more than most IT operational procedures.

Ultimately, a backup is only as good as the restore it is capable of performing. Many of those who have experienced a full system restore due to hardware failure are aware that it is not always as simple as may first appear.

In spite of this, tape backup does have the advantage of being familiar territory for most people. Tape backups are easy to implement and a cheap solution (if you do not count the cost of downtime while the system is being recovered).

It would be true to say that the pitfalls of tape backup is one reason why organisations are considering alternative methods of restoring users to normal operations in case of downtime. Some of the options available are given below.

## Cold Standby System


A redundant second server is made available to be used as a backup hardware in case of system downtime.

Typically a bare metal restore has to be performed when the server is required and a tape backup is used to perform the system recovery.

## Cold Standby with Data Replication

Data replication is used to make a second copy of data 'real-time'. As the data is changed on one server, it is also changed on the second server. This then acts as an up-to-date copy of the data. The recovery process then involves restarting the application on the second server and having to re-route clients to the backup server. This is a middle ground between a cold standby and a hot standby as data does not need to be recovered from tape and it does not have the disadvantages of having to restore from a tape backup.

### Data Replication Over LAN




## Cold Standby with Data Replication continued

Data Replication is an ideal solution for backing up data over a network. A variety of automated backup procedures can be created that eliminate the need to disrupt user's daytime work activity, or schedule lengthy midnight backup sessions.

The ability to pause the replication process allows a combination of both replication and backup technologies into a single integrated, low latency solution. By performing real-time replication in conjunction with periodic tape backups, users can quickly migrate online backups to media suitable for long-term storage while retaining the immediate availability of data in the event of a local failure.

## Hot Standby System

This solution involves setting up a second server that takes over and running the active application. When the application on the primary server fails, or one of its dependent components, it is recovered on the backup server.



This method offers the fastest recovery method. The backup server is usually up and running in a few minutes. Users may experience data loss of the transaction that they were working on when the system went down. Most applications today are capable of recovering this. The fast recovery means that users can continue working while IT departments diagnose and resolve the cause of the problem. This also means that support does not have to be reactive to external events - reducing out of hours working.

There are also other advantages to having a second backup server, such as it can be used to test upgrades and patches before they are applied to the live system. When they are ready to be applied, users can be switched to the backup system and changes applied to the primary system without the pressure of having to return it to service in a short space of time.

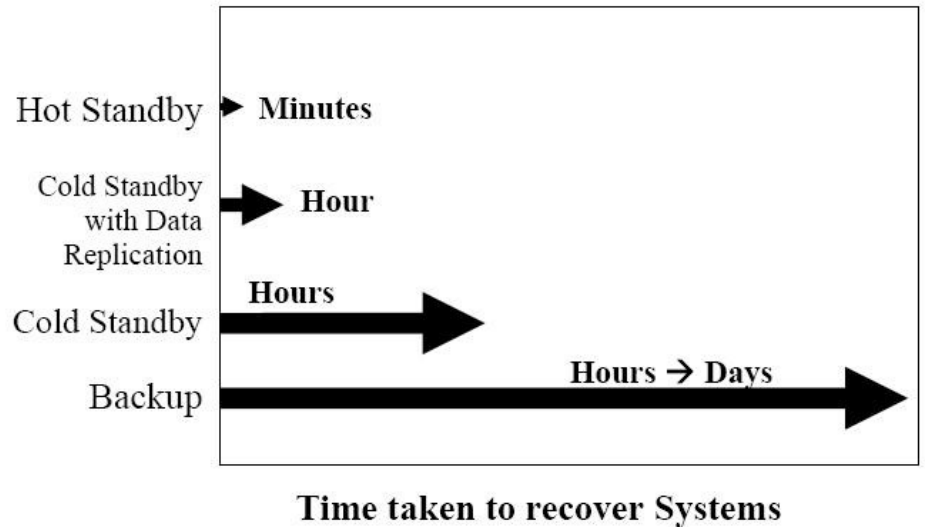
A hot standby is often chosen by those having to guarantee 24/7 availability as it offers the least downtime in the event of not just hardware but application failure.

## Hot Standby System continued

At this point a distinction has to be made between data availability and application availability. Hardware such as SAN's provide availability of data whereas a hot standby system provides availability of applications. A hot standby system is able to perform a tasks intelligently when the system goes down. A standby system will be able to restart applications and ensure that the users regain access to critical applications and data without any manual intervention. Data can be available, but if users cannot access the servers, then it is effectively unavailable. User productivity is determined by application availability.

The combination of Data Replication and a hot standby provide a powerful and resilient recovery solution that incorporates the ability to provide systems recovery and backup.

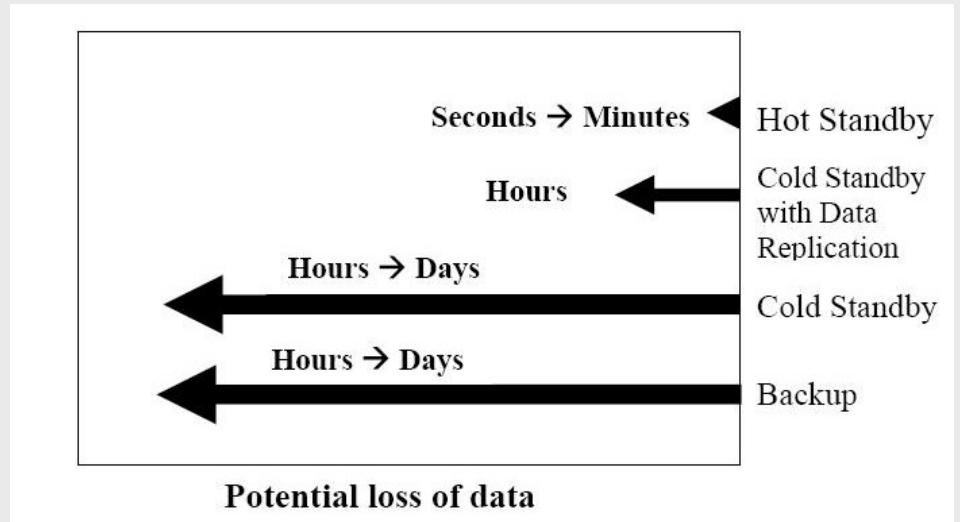
## How long does each method take to recover?



Looking at the above diagram, it is clear that the tape backup will take the longest to recover. Making it potentially the most expensive solution if the cost of downtime is added to the total cost.

## How much data will be lost?

Looking at the diagram below, the backup and the cold standby stand on a par as they both rely on the integrity of the tape backups. The data replication and hot standby will have data available ready to run so the recovery time is much shorter.

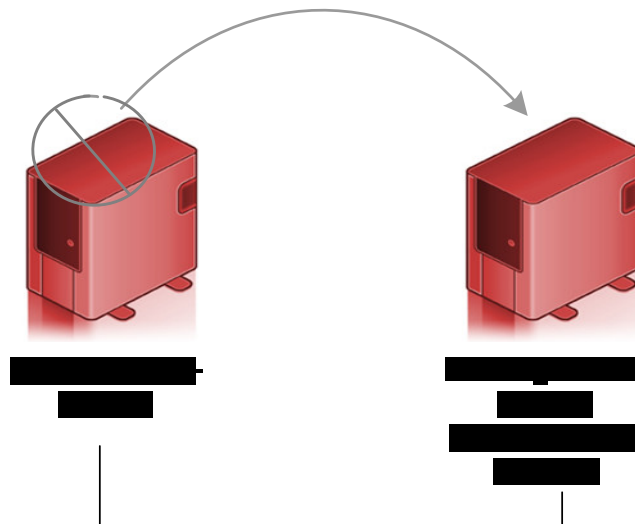


## Implementing a hot standby system using SteelEye LifeKeeper

Creating a hot standby system means planning and implementing two servers that are capable of running the application independently.

It is usual for one server to be the primary server and another one to be the secondary server. Typically, the secondary server is also used as the primary server for another application, this solution give a better ROI as there is no redundant server.

Both servers are then acting as a primary and a secondary for two applications, and both servers should be capable of running both applications.



## Implementing a hot standby system using SteelEye LifeKeeper continued

In the example above, Server1 is the active server for IIS, but the backup server for Windows File and Print server. To achieve this, both servers have to be set up to run both applications independently first. Then LifeKeeper is installed on both servers to monitor the applications health.

SteelEye's LifeKeeper ensures the continuous availability of applications by maintaining system uptime. LifeKeeper maintains the high availability of clustered systems by monitoring system and application health, maintaining client connectivity and providing uninterrupted data access regardless of where clients reside - on the corporate Internet, intranet or extranet.


LifeKeeper provides fault resilience for Linux environments by enabling other servers in a cluster to take over for failed servers or failed applications. Total cost of ownership is reduced, because LifeKeeper supports an active-active server configuration. This model eliminates the need for extra servers dedicated for hot backup and allows clients and applications to failover to other production servers in the cluster.

## Proactive Protection of applications

With LifeKeeper, hardware component or application faults are detected in advance of a full system failure through multiple fault-detection mechanisms. LifeKeeper monitors clusters using intelligent processes and multiple LAN heartbeats. By sending redundant signals between server nodes to determine system and application health, LifeKeeper confirms a system's status before taking action. This reduces the risk of a single point of failure and minimizes false failovers. LifeKeeper also limits unnecessary failovers by recovering failed applications, without a full failover to another server, if the hardware is still active.

LifeKeeper takes an application-centric approach to availability, monitoring the application and all its dependent components. For example, a database such as Oracle or MySQL is dependent on a file system, an IP address being available. These dependencies are configured when LifeKeeper is being set up and a dependency hierarchy is created.

Dependencies between a database application its dependent components



## Proactive Protection of applications continued

When any one of these resources is detected to not respond, it starts a chain of recovery events. The first is to start a recovery on the local server as this would produce the least impact and fastest recovery time for the users.

If a local recovery is not possible, then a recovery on the backup server is initiated.

Virtual IP addresses are used for applications, allowing the application to move from one server to another without any need for reconfiguration or interruption to the client.

The recovery procedure switches over the virtual IP address to the backup server, it then restarts the application, and the whole procedure is automated so it does not require manual intervention.

## Shared Storage

Shared storage handles availability of the data and allows the application to failover while assuming that the data availability is handled by the shared storage.

This type of solution is used by cluster hardware such as the HP ProLiant G3 cluster. The LifeKeeper software is used to failover to the second node. As LifeKeeper sits on top of the software, the software itself does not need to be cluster-aware. LifeKeeper does the work of monitoring the software, and all of its dependent components down to hardware level.

This solution is suitable for any shared storage such as SCSI, NAS, Fibre, iSCSI or SAN.

## Replicating data where no shared storage is available

Where no shared storage is available, a failover scenario needs to plan how the data availability is to be maintained in the event of a failure.

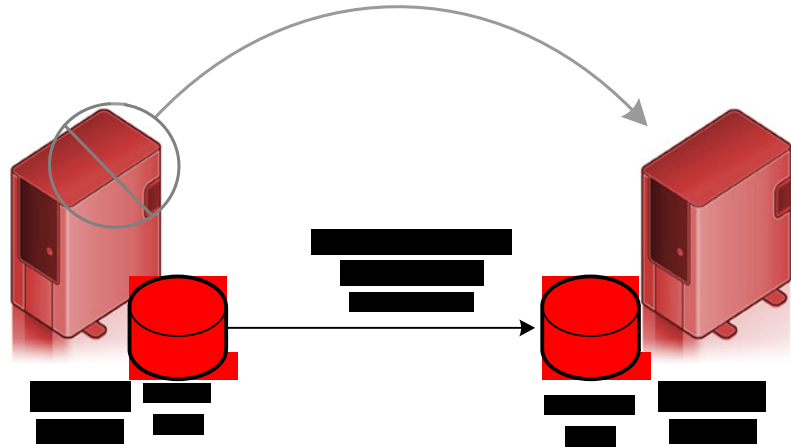
In this case, when the failover occurs, the data has to be as up-to-date as if it were on the primary server. This is achieved through data replication.

LifeKeeper replicates data at the block level and allows those blocks to be user defined. It can mirror in increments as small as a single byte, and can select to replicate only changed data in order to minimize the impact of systems and network bandwidth.

Data Replication enables users to define how and when data is mirrored, with facilities for continuous, periodic and scheduled replication, as well as synchronous or asynchronous replication. In addition, with change logging, synchronization of disks is fast and dependable.

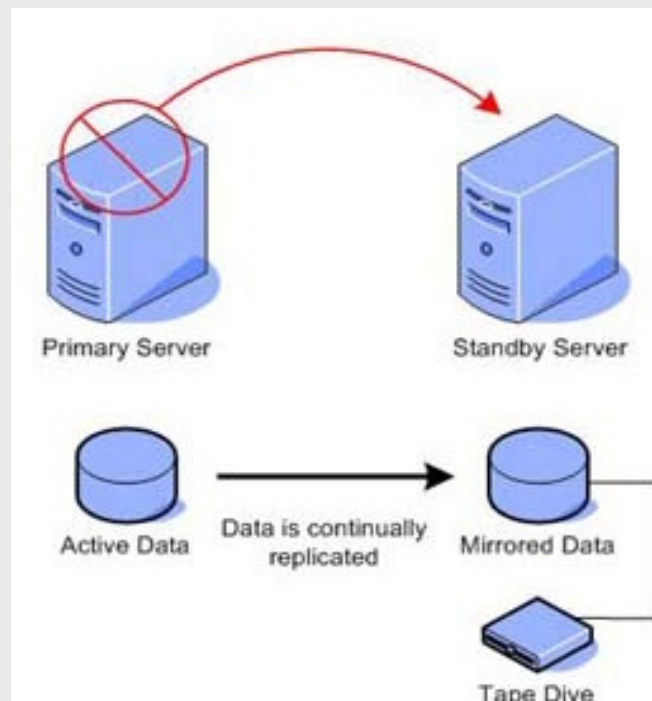
## Failover and Data Replication

Data replication ensures that the data on the backup server remains current. In the event of server or application downtime, LifeKeeper initiates a recovery of the application on the backup server. When the application restarts, the data is up-to-date due to the replication of the data. When the failed server is repaired, the data is replicated back.



## Backup and Data Replication

Data Replication is an ideal solution for a company's backup procedures. An IT department can create a variety of automated backup procedures that eliminate the need to disrupt user's daytime work activity, or schedule lengthy midnight backup sessions.

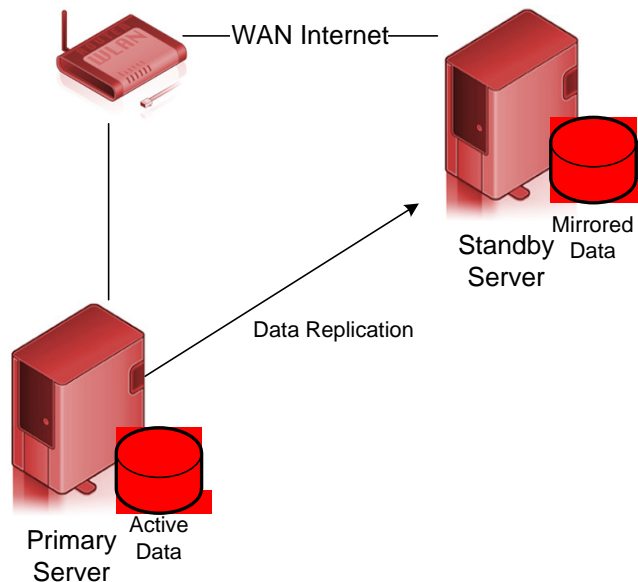


## Backup and Data Replication continued

The ability to pause the replication process allows IT departments to combine both replication and backup technologies into a single integrated, low latency solution. By performing real-time replication in conjunction with periodic tape backups, users can quickly migrate online backups to media suitable for long-term storage while retaining the immediate availability of data in the event of a local failure.

## Disaster Recovery To a Remote Site

Recovery to a remote site is achieved by replicating the data over a Wide Area Network. A failover to the remote site is initiated when there are problems with the application or hardware.



In the example above, the data is replicated between the two servers. When the failover occurs, the data is current so the users can continue working on a recent copy of the data.

To replicate data over a WAN, sufficient bandwidth has to be available, especially when there is a large amount of data being updated continually.

## Application Recovery with LifeKeeper

One of the key advantages of SteelEye's LifeKeeper architecture is that neither system kernels nor target applications need to be changed in order to provide application and data protection. Instead of requiring changes to users' compute environments, LifeKeeper's Application Recovery Kits (ARKs) are used as an efficient and flexible mechanism to provide integration between LifeKeeper, its GUI and any target application whose reliability and uptime require protection.

These ARKs are a set of administration and recovery components that provide LifeKeeper with the ability to manage and control a specific application. A template recovery kit is included with the LifeKeeper core to help users get started with developing their custom ARKs, if any are needed. This recovery kit template may be easily modified to allow protection of a generic or "customer-defined" application.

## Recovery of Standard Applications

There are standard ARKs available for most of the commercially available databases, web servers and applications.

Following, is a brief description and overview of some of these optional ARKs. Contact us for information on specific application recovery scenarios.

**Web Server Recovery Kit** - provides protection and recovery for automatic switchover of the Apache Web Server and Secure Web Server (SSL) processes on Linux as well as protection of MS-IIS on Windows.

**Database Recovery Kits** - These LifeKeeper ARKs provide resource definition, failover and recovery software for Oracle 8i and 9i, Informix, and DB2 WE, EE and EEE, respectively, running on Windows or Linux platforms. A kit for SQL Server is available on Windows 2000. A kit for MySQL server is available on Linux.

**IP Recovery Kit** (included with core) - Provides switchover software for automatic recovery of IP addresses. These IP addresses may be migrated to another NIC within the same server or to a NIC on an alternate server.

**Network Attached Storage Recovery Kit** - provides resource definition and recovery software for cluster systems that have mounted an exported file system from an NFS server or Network Attached Storage (NAS) device in the cluster. This kit allows the use of NAS as an alternative to directly-attached shared storage.

**NFS Server Recovery Kit** - provides resource definition and recovery software for automatic switchover of NFS exported file systems.

**Print Services Recovery Kit** - provides resource definition and recovery software for automatic switchover of printers and print queues.

**Email Messaging Systems Recovery Kit** - provides resource definition and recovery software for processes, mailboxes, and mail queues in either Sendmail or MS-Exchange environments.

**File Share Recovery Kit** - built into the LifeKeeper core product, it provides support for the Windows 2000 and Windows NT file share kit. Similar functionality is also available on Linux with an ARK for Samba.

## Bespoke Application Recovery

Applications that do not have a pre-packaged recovery kit can be easily recovered by using the templates provided within LifeKeeper (Generic application recovery). Or by using the Customisation Kit.

## Generic Application Recovery Kit

The Generic Application Recovery Kit included with the LifeKeeper core allows protection of a user-defined application that has no associated recovery kit to define the resource. This recovery kit provides the user with a generic administration interface to create a LifeKeeper resource, and define monitoring and recovery scripts that are customized for a specific application. SDK template files are provided with the Generic Application Recovery Kit for these user-supplied scripts. Any dependent resources such as file systems, raw devices, or IP addresses are created separately using the appropriate recovery kits. These are then linked to the application resource with dependencies to create the application resource hierarchy. Refer to the LifeKeeper Online Product Manual for additional information on using the Generic Application Recovery Kit.

## Customisation Kit

The LifeKeeper Customisation Kit provides a user-friendly way to develop integration between the LifeKeeper core and any other application, enabling users or SteelEye partners to write their own custom recovery kits.

Usually a custom recovery kit leverages as many of the existing LifeKeeper resource instances as it can. For example, the Apache Web Server Recovery Kit supplies its own webserver/apache resource instance, but its job is merely to start and stop the web server. It creates dependencies to the LifeKeeper supplied comm/ip and gen/filesys resources, to protect the IP addresses and file systems it needs to function correctly.

## Components of an ARK

Each ARK consists of two separate components: an action component (for recovery and monitoring) and an administration component (for creation, extension, and removal). The action component performs start, stop, and monitoring operations for the application associated with a resource instance. The administration component is usually attached to a top-level resource and performs administration operations, such as creating the entire resource hierarchy, including the associated dependencies.

The administration component also includes a properties file that enables easy integration of the new ARK throughout the rest of the GUI.

## Is the cost of a Hot Backup justified?

The recovery solutions chosen by organisations depend on the value placed on the data and the uptime of the system. In terms of recovery time, a hot standby system offers the shortest recovery time with the least pain when a system goes down. However, until recently the price has been prohibitive for most organisations. Today, with the cost of hardware reducing dramatically and the combination of software tools that have become more affordable over the years, it is an option being considered by more and more organisations. The cost of additional hardware and software is rarely greater than the cost to the organisation of lost productivity and operations resulting from systems downtime.

## About Open Minds High Availability Solutions Ltd

Open Minds High Availability Solutions Limited provides support, training and consultancy for High Availability IT deployments.

Founded in 1990, Open Minds focus is in using their experience of delivering high availability and server recovery to serve their customers. Open Minds solution architects have over 10 years experience of LifeKeeper and other High availability products on various platforms, including NT, Window 2000, Unix and Linux.

Using a unique network of partners and products, Open Minds can plan and implement systems availability on a simple two-node cluster on a LAN, to a sophisticated multi node, multi-application disaster-recovery across a WAN.

Open Minds works with its customers and partners to create systems and application availability scenarios in line with their business objectives.

## About SteelEye Technology Inc

SteelEye Technology Inc (<http://www.steeleye.com>), is a leading provider of IT solutions for business continuity and disaster recovery.

The SteelEye LifeKeeper family of software products and services offer unique scalability in terms of integrated solutions for data protection, high availability clustering, and wide-area disaster recovery on Windows and Linux.

SteelEye LifeKeeper enables enterprises of all sizes to ensure continuous uptime of business-critical systems. With 'out of the box' support for the widest range of applications, databases, and storage subsystems running on Intel-based Windows and Linux servers, SteelEye LifeKeeper ensures enterprise-grade reliability at a fraction of the cost and complexity of traditional solutions.

SteelEye LifeKeeper is proven in the world's most demanding business environments. Today, Global 1000 companies rely on SteelEye LifeKeeper to keep their systems running and their people productive.

## About This Document

Original author - Open Minds High Availability Solutions.  
Revision - 5.0  
Last updated - September 20th 2006.

Portions of this document are used under permission from SteelEye Technology Inc.

Please send any feedback, questions or criticism to [info@openminds.co.uk](mailto:info@openminds.co.uk)

