

## Disaster Recovery and High Availability Solutions for Microsoft Exchange using SteelEye LifeKeeper

### Index

<i>High Availability and Disaster Recovery for Exchange</i>	<i>Page 2</i>
<i>Protection for the Exchange Environment</i>	<i>Page 3</i>
<i>Configuration Considerations</i>	<i>Page 4</i>
<i>Active/Standby Configuration using Data Replication</i>	<i>Page 5</i>
<i>Active/Standby Configuration using Shared Storage</i>	<i>Page 6</i>
<i>N+1 Configuration</i>	<i>Page 7</i>
<i>Off-site Disaster Recovery Solution</i>	<i>Page 8</i>
<i>Combining Disaster Recovery and High Availability</i>	<i>Page 9</i>
<i>Network Considerations</i>	<i>Page 10</i>
<i>Data Replication Considerations</i>	<i>Page 10</i>
<i>Virus Checking and other third party Exchange Add-ons</i>	<i>Page 11</i>
<i>Appendix - LifeKeeper Functionality Matrix</i>	<i>Page 11</i>

### Abstract

Messaging has become *the* critical business application. This white paper describes how to use LifeKeeper to provide high availability protection and Disaster Recovery for Microsoft Exchange environments.

All solutions mentioned in this white paper are available today from Open Minds High Availability Solutions.

### Introduction

Businesses are increasingly relying upon Microsoft Exchange to perform business critical messaging and day to day tasks. Because of this reliance, it is becoming necessary for the corporate messaging infrastructure to be available all the time, with minimal or zero visible disruption. Traditionally only larger enterprises have been able to justify high availability solutions, as they required high end custom hardware or proprietary operating systems, or a mixture of both. Our solution is able to escape from

both of these constraints, and allows for any IBM PC compatible hardware to be used, along with the Microsoft Windows family of operating systems (specifically NT, 2000, 2003). It works on MS Exchange 5.5, Exchange 2000 and Exchange 2003.

## High Availability and Disaster Recovery for Exchange

- High Availability is the ability for a service to recover from failure, whether this is hardware or software related, in order to provide continuous service to the user. High Availability systems often allow for a small window of downtime. After a failure has been detected, the host locally recovers from the failure, or a fail over takes place to another server. The window of downtime can vary from a few seconds to minutes, depending upon the application and failure. In the case of Exchange servers, recovery time, depending on the number of users, can be as low as 1 minute.
- Disaster recovery is the ability of a server or service to recover from a failure where the servers are located across 2 separate sites. This distance is used to protect critical servers against site disasters.

Both high availability and disaster recovery are a part of the SteelEye LifeKeeper solution. As failover can take place on a local or remote server. There is also a combined 3-node solution that incorporates a local and a remote failover.

The LifeKeeper Microsoft Exchange Protection Suite for Windows recovers MS Exchange services from hardware or software failure. After a failure has been detected, the Exchange services are recovered on a backup server. The backup server can be located on the same LAN as the active Exchange server, or on a WAN to provide disaster recovery. There is a small window of downtime while the software establishes the failure of the active server to provide the service, and recovers it on the second server.

It can work equally well with and without shared storage. Features of SteelEye's implementation for an Exchange Fault Resilient solution include.

- LifeKeeper works on all versions of Windows, there is no requirement for the high end, Enterprise or Data Centre versions of Windows.
- LifeKeeper works on all versions of Exchange including Exchange 5.5, Exchange 2000 and Exchange 2003, it does not require Exchange Enterprise Edition.
- LifeKeeper does not require shared storage (e.g. SAN, Shared SCSI or a NAS)
- LifeKeeper provides for a disaster recovery solution, giving ultimate protection to important data.
- LifeKeeper lets you run applications other than Exchange on the protected server and gives the ability to support any application running on that server. For a smaller business, this allows for existing servers to be utilized.

## High Availability and Disaster Recovery for Exchange Continued

- LifeKeeper returns the failed server back into the cluster seamlessly by re-synchronising data and making the failed server the backup server.
- LifeKeeper modifies Exchange server information for all domain users in the Active Directory (applicable for Exchange 2000/2003)

LifeKeeper protects, not just from hardware failure but also the Exchange services. Therefore, hardware resilience and software resilience is achieved via the LifeKeeper Exchange solution. The solution has been deployed worldwide and in the UK we have customers running the Exchange 5.5, 2000 and Exchange 2003 solution.

## Protection for the Exchange Environment

### Exchange Services Monitored

Not only does LifeKeeper perform a failover in case of a catastrophic hardware failure, but it can also easily monitor optional exchange server services. LifeKeeper increases the uptime of the Exchange Server by constantly monitoring key Exchange Server Components :

Message Transfer Agent (MTA)  
System Attendant  
Directory and Information Store  
World Wide Web Publishing Service  
SMTP  
Information Store and  
Routing Engine

The periodic monitoring of the services is called the “QuickCheck” and “DeepCheck”. QuickCheck's are generally simple tests to ensure that the service is running, while DeepCheck's provide a more rigorous test that validates the functionality of the services. Should either test fail, then LifeKeeper will attempt to restart the service. Should this restart fail, or the service fails again, fail over to another system will take place. The interval at which the services are checked by either approach can be customised.

### No Need for 24/7 Staffing of IT Departments

Reduced downtime and enhanced data protection result in lower staffing costs. In the event of a failure, there is no urgent rush for an administrator to repair the service, as the backup node will have taken over. This helps maximize the investment in Exchange.

### No Need for downtime during upgrades and maintenance

As well as being useful in preventing system disasters, LifeKeeper can also be used to assist system upgrades or for installing new hardware and/or applications. This is possible as the service can be manually switched over from one node to another at will, allowing the inactive node to be upgraded/ repaired. The administrator can then check that the newly upgraded node works, and if not switch back to using the other system. This reduces the need to schedule downtime for common maintenance tasks and upgrades, helping to enhance the availability and reliability of the service as a whole.

## Configuration Considerations

The Exchange Server Application Recovery Kit (ARK) provides for the installation and operation of an Exchange Server in either a shared storage environment or one using data replication. Clients connect to the computer name through Active Directory. The recovery kit for Exchange manipulates the Active Directory, allowing clients to be unaware of a failover – there is no need for reconfiguration.

The main configuration issue is what sort of storage is to be used (e.g. SAN, NAS, SCSI, Internal disk etc) and then is the cluster going to be spanning a wide distance to provide a disaster recovery solution.

If some form of shared storage is used, then it is not necessary to use the data replication software. The advantage of this is that there is no need to do a resynchronization between disks after a fail over, however shared storage is often more expensive and difficult to use in a disaster recovery solution as there is often a limit on the distance it can span.


Without shared storage, data replication software is used. This software mirrors at the block level on one disk (the source) to another (the target). At any one time while the mirror is in operation, only one side of the mirror is accessible – this prevents corruption of the data. Because the replication is occurring below the file system level, it is not affected by open or locked files. The use of data replication software, which runs over TCP/IP, allows for a true disaster recovery solution, where the distance between nodes is limited to the budget for bandwidth.

*Active/Standby Configuration.* Microsoft Exchange Server itself is designed so that a given computer can run only one instance, and the name of the Exchange Server is fixed and cannot be changed. This limits configuration options to active/standby for Exchange servers. However, an Exchange server can be a part of an active/active and N+1 environment with other application servers.

*Exchange Server Installation.* This is relatively straight forward and automatic. The installation details and instructions are detailed in the “MS Exchange Server Recovery Kit Administration Guide” which is on the Life-Keeper Documentation CD.

## Active/Standby Configuration using Data Replication

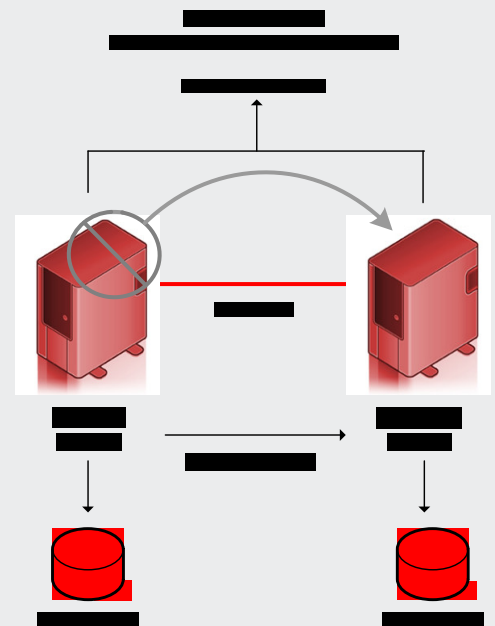
For a two node cluster in an active/standby configuration. Initially the active server is on the left (Primary Exchange). There is a heartbeat between the two servers, and data replication taking place between an internal volume on the primary server and another internal volume located in the secondary server. While replication is taking place, only the source machine has access to the data to prevent data corruption.



In the event that a failure occurs, on the primary exchange server, that is not recoverable locally, then this will occur :

Because the server on the left is no longer accessible, no form of replication is currently taking place. Exchange will come into service on the secondary server (on the right), hence it now has access to the local storage. The heartbeat channel is flat – i.e. there is no heartbeat occurring as one system is out of service.

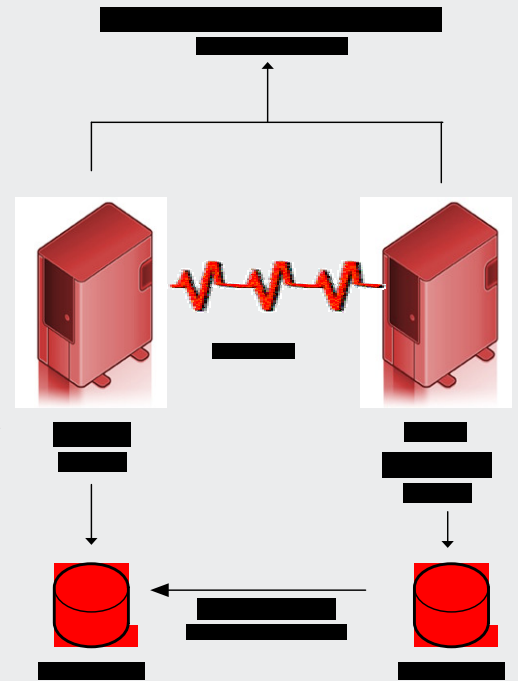
On failover, the Active Directory is manipulated so that the Exchange Server is fully recovered.



## Active/Standby Configuration using Data Replication Continued

Then, when the server on the left hand side is repaired from the failure that took place – this may be hardware related or routine maintenance (e.g. Application of security updates or service packs). The data replication direction is now reversed, and, depending on the type of failure, either a partial or a full resynchronization will take place from the now active server (on the right) to the standby server (on the left). A partial resynchronization would have taken place unless there is a disk failure, here, only the newly modified blocks will be copied across.

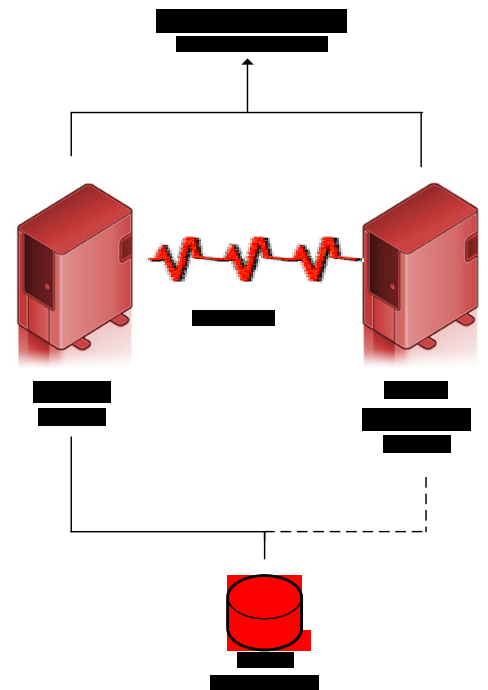
An automatic switchback, though possible, is not generally recommended. This is to guard against a service “bouncing” while repairs are underway (the machine may be rebooted multiple times while repairs are undertaken).



## Active/Standby Configuration using Shared Storage


LifeKeeper manages access to the shared storage, ensuring that only one computer at a time has access. It also automatically detects a failure of either the node, or a service a node is running allowing it to perform a predetermined recovery.

In the diagram to the side, initially the left hand node (marked ACTIVE) has access to the shared storage, while the other node is locked out. Exchange is running on the active server, and the heartbeat between the nodes is functioning without problem.



## Active/Standby Configuration using Shared Storage Continued

In the second picture, the power fails (or the server crashes). This causes the heartbeat communication between nodes to die, and results in the standby server taking over. The standby server gains access to the data, and starts running the Exchange service. Clients are able to connect without change, due to the migration of a floating IP address from the failed server to the active server.

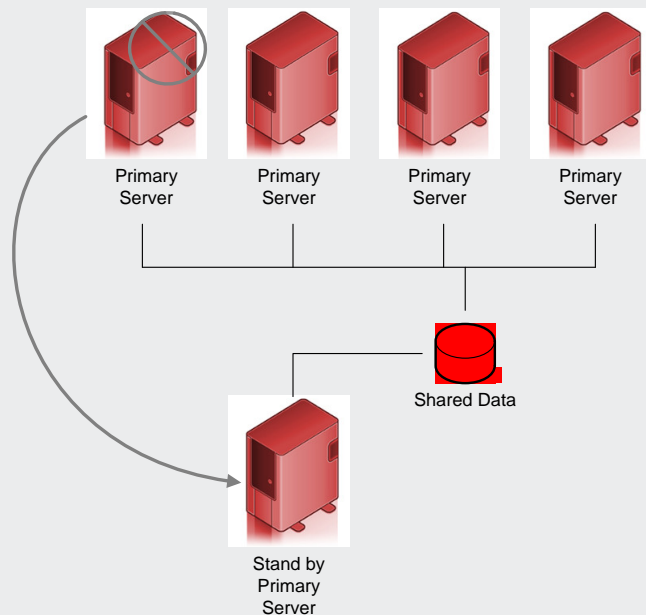


## N + 1 Configuration

Available only for shared storage configurations, it is possible to have a number of servers all being protected by one standby server. The cost advantages of this solution are significant as it reduces licence costs and hardware costs.

For example, if 4 servers are protected by one backup server (as shown in the diagram), there will only be 5 servers in total, rather than 8 servers if each server had its own separate backup server. This results in a reduction in costs of hardware, licences, support and maintenance.

N (activity) +1 (backup) Failover

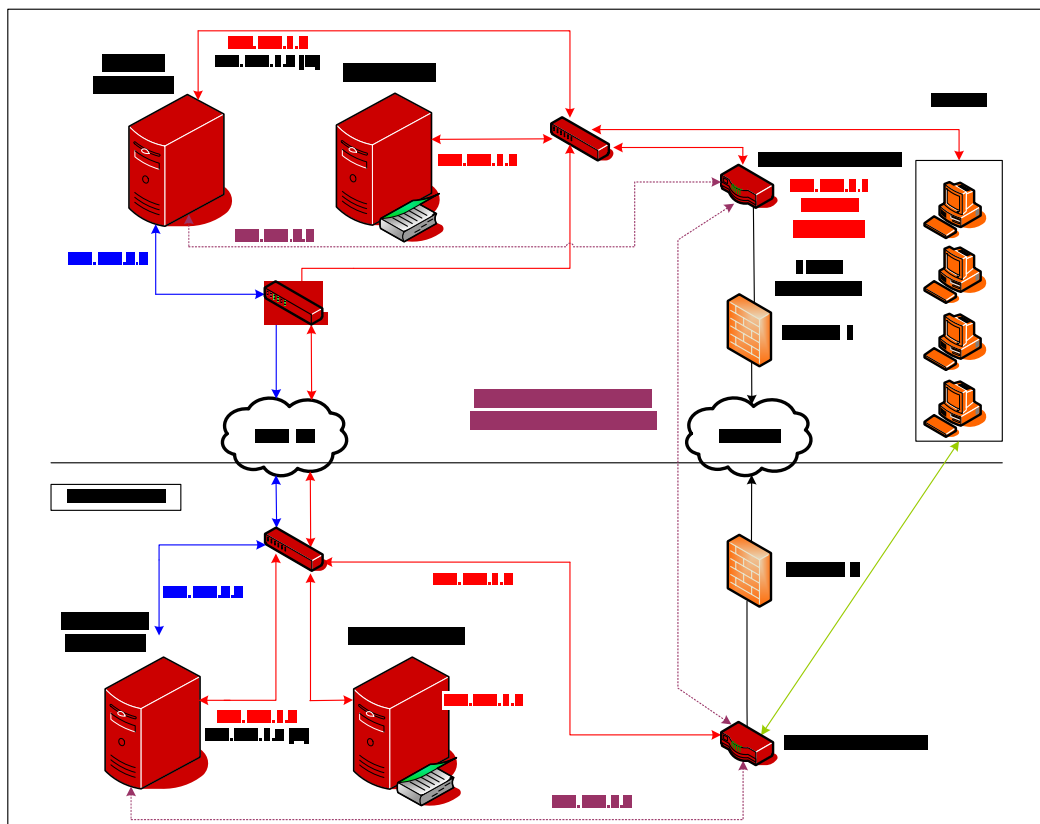


## Off-site Disaster Recovery Solution

It is also possible to have a 2-node failover over a Wide Area Network. The beauty of this is that it allows organizations to deploy a disaster recovery system relatively inexpensively. The following is a two node disaster recovery solution, designed by Open Minds for a London based client. The solution consists of four nodes – two Exchange servers and two Active Directory servers. The Exchange servers are configured in an active/standby configuration, while the Active Directory servers are configured to replicate between themselves at all times. As you will see from the diagram below, one pair of servers (standby Exchange, and Active Directory) are located in an off site data centre, while the primary Exchange server and the other Active Directory server are located on the customers premises. Should the local Exchange server fail, LifeKeeper will move the service to the data centre. This fail over will be seamless to clients, who will continue connecting as normal.

In the event of a disaster occurring, clients will continue to connect to the server by connecting to the data centre via a VPN link over the Internet. This is under the assumption that they are unable to physically access the office, due to the disaster, and will be for example working from home or in a remote office. Replication of data between the two sites takes place over a dedicated link such as a LES10 (10mbit) link.

In order to reduce the chances of a false fail over taking place (or a split brain scenario), two heartbeat channels are created between the LifeKeeper protected Exchange servers – the first is over the 192.168.2.X network (over the LES10). The second heartbeat channel is over the 192.168.3.X network, which is connected to the remote site via a VPN between the two routers. As you can see the network is configured to allow access to the DR site over the LES10 should the local exchange server fail.

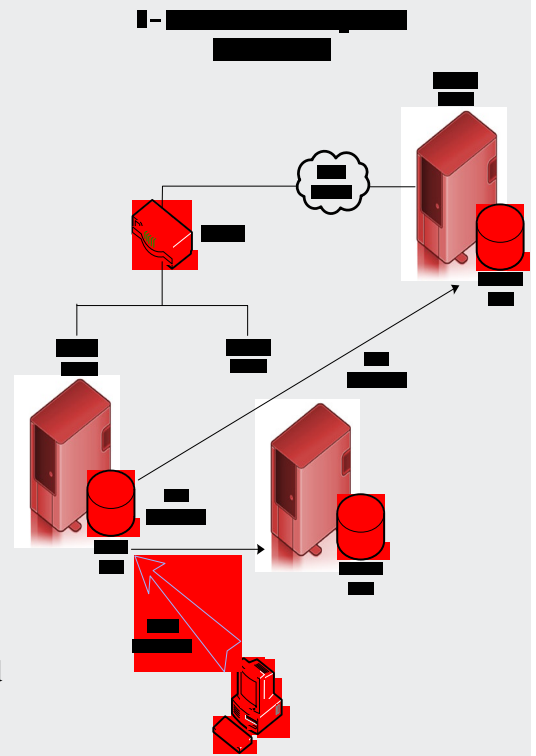


## Combining Disaster Recovery and High Availability

This involves setting up both a local and a remote backup server, both of which are capable of running Exchange. It combines the benefits of high availability and disaster recovery. The local server is used as the main backup server, and the remote backup server is only used if there is a site disaster.

Therefore, the expectation is that the local backup will be the server most commonly used. Use of this solution often reduces the requirement for expensive WAN links as the local server is the main backup

This is a 3-node Exchange solution. The Exchange data is replicated 2 ways. Firstly to the local server, and secondly to the backup server over the WAN.



The 3-node solution offers the advantages of a local failover for those not-quite-so-disastrous events such as network failure or Exchange unavailability.

This solution offers the benefits of :

- Least disruption is experienced by users when a server or application failure is experienced.
- Local failover can be used administration purposes such as upgrades and maintenance
- Local failover is in place for 'everyday' disruptions such as network glitches .
- The remote server is only used for site disasters.
- Wan replication can be stopped during peak traffic and restarted as required.
- If the WAN link is unavailable, the local copy of the data is still available and Exchange is still protected.

## Network Considerations

In all cases it is recommended that there are at least two communication paths between the LifeKeeper protected servers. This guards against false failovers, split brain scenarios and data corruption.

A split brain scenario is when neither node can see the other, this results in both nodes thinking the other is out of service, and bringing the services on line locally. In many situations there will be no data corruption from this, as clients normally only be able to see one particular server at a time (if they could see both, then you need to ask why there was not a backup heartbeat path over the public LAN).

If undertaking a disaster recovery configuration, it is necessary to also have an active directory controller on the remote site. Otherwise clients will not be able to authenticate with the Exchange server(s).

## Data Replication Considerations

The data replication software, copies at the block level on a disk. It is therefore unaware of files and how much of the disk is actually being used. There is no replication of white space and only changed data is replicated. This minimises the traffic on the network. In addition, if the network speed is an issue at certain times of the day, it is possible to pause the mirror at peak traffic times, and restart it. On restart, it will automatically resynchronise.

In the event of a disaster recovery type solution being deployed, where the nodes are separated by some distance, an Asynchronous mirror can be used, which improves performance on the active node (but does increase the chance of data loss). The asynchronous mirroring can be undertaken over lower speed links if necessary – unfortunately it is hard to estimate what speed of link is required for each customer due to differing requirements and usage patterns. As mentioned above, only write requests are mirrored to the remote server, so if a large amount of the activity is reading then a low speed link may be feasible.

Open Minds have customers using links from as low speed as 2mbit/s to 100mbit/s for disaster recovery solutions.

In some configurations DR situations it may therefore be necessary / desirable to move the off site server to the same room as the active server in order for a full resynchronization to take place quickly over a local gigabit link, after which the mirror can be paused on the target node, the node powered down and moved back to the remote data centre. The mirror can then be resumed / un-paused and a partial resynchronization allowed to take place (where only the changed data is copied across).

When a failed server is returned to service, the mirrors will be automatically resynchronised.

## Virus Checking and Other Third Party Exchange Add-ons

LifeKeeper is capable of controlling other services at the same time as Exchange, allowing for e.g. A virus checker to be running on the same node as the Exchange server at all times. This is possible due to LifeKeeper's inherent flexibility, and the lack of a requirement to recode the application to a cluster API.

Open Minds have Exchange customers using third party virus checkers, and back up solutions integrated with the Exchange recovery kit, so in the event of a failure taking place virus filtering and backups continue unimpeded.

## APPENDIX LifeKeeper Functionality Matrix

<i>Functionality</i>	<i>Description</i>
Maximum Nodes Supported	Only one Exchange node can be active at one time. But LifeKeeper supports up to 32 active nodes.
Mixed Platforms	LifeKeeper does not require that each cluster node is identical. The only requirements are that an Intel based platform is used running Linux or Windows. Each node must be configured to run the service, and able to handle the workload.
Data Replication	Block level replication White space not mirrored Only changes tracked Automatic resynchronisation
Exchange Administration	Administration and maintenance are aided by the ability to move the Exchange service between nodes during maintenance.
Environment Monitoring	The LifeKeeper Exchange Recovery Kit has the ability to monitor the Exchange services, as well as any other additional services (E.g. 3 <sup>rd</sup> party virus checkers).
Client Reconnect	Clients do not need to reconfigure, and will often be unaware of an exchange fail over.
Resource Monitoring	LifeKeeper monitors dependent resources, and should one fail, and not be recoverable locally, fail over to the backup system will take place.

**APPENDIX**  
**LifeKeeper**  
**Functionality Matrix**  
**Continued**

<i>Functionality</i>	<i>Description</i>
Protected Services	Message Transfer Agent (MTA) System Attendant Directory and Information Store WWW Publishing services SMTP Information Store Routing Engine +Additional third party services.
Connectors	Lotus Connector X400 Connector MS-Mail connector cc:Mail
Remote Administration	LifeKeeper protected cluster can be administered remotely through a web based Java GUI.
Multiple Network Communication Paths (Heartbeats)	It is possible (and recommended) to use multiple network paths between a given pair of servers. This enhances reliability and reduces the likelihood of a false fail over, or split brain scenario.
IP Local Recovery	In the event of a network card failing, it is possible for LifeKeeper to initiate recovery locally using an appropriately configured network card – this removes the requirement to purchase network card bonding products.
External Interface Support	LifeKeeper is capable of sending SNMP traps to external monitoring hosts, allowing for real time problem notification.

**Further Information**

For more information or a demonstration of the Microsoft Exchange Failover contact Open Minds High Availability Solutions

Phone 0845 3453943  
+44 (0) 121 313 3943  
email: [sales@openminds.co.uk](mailto:sales@openminds.co.uk)  
Web <http://www.openminds.co.uk>

