

# Achieving Reliable and Efficient Cross-Platform Server Recovery with EMC HomeBase

*A Detailed Review*

---

## **Abstract**

Organizations are seeking a more effective approach to server recovery, which has historically limited their ability to meet aggressive overall system recovery time objectives. Learn about the challenges presented by traditional server recovery methods and how new software offerings like EMC® HomeBase are helping companies overcome those challenges.

January 2008

---

---

Copyright © 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com)

All other trademarks used herein are the property of their respective owners.

Part Number H3501

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	4
<b>Server recovery: A weak link in overall recovery capabilities .....</b>	<b>4</b>
<b>Why traditional server recovery methods fall short .....</b>	<b>5</b>
Server environments change frequently .....	5
Recovery usually happens on dissimilar hardware.....	6
Server configuration details and backup datasets are highly interdependent.....	7
In-house server recovery is labor-intensive and time-consuming .....	7
<b>Assessing new server recovery software tools .....</b>	<b>7</b>
Capture configuration changes daily with minimal production impact.....	8
Enable recovery to different server hardware.....	8
Integrate a server configuration change capture process with backup workflow .....	8
Apply automation carefully to limit errors and speed server recovery.....	8
Support virtualization to add flexibility and increase server recovery efficiencies .....	9
<b>Meeting server recovery objectives with new technology .....</b>	<b>9</b>
Complete, accurate recovery .....	9
Minimized production impact .....	9
Powerful process control capabilities.....	10
Detailed record of configuration changes .....	10
How HomeBase works.....	10
<b>Lower risk, improve efficiency, and increase availability .....</b>	<b>11</b>
Reliable, cost-effective, hardware-independent recovery.....	11
Fast recovery and improved production availability.....	12
Managed server assets and leveraged investments .....	12
<b>Conclusion .....</b>	<b>12</b>

---

## Executive summary

Today's business environment demands aggressive recovery time objective (RTO) and recovery point objective (RPO) targets. For many companies, these targets are approaching zero. However, current server recovery methods are keeping companies from meeting these objectives.

Traditional methods of server recovery have not kept up with the increase in recovery requirements for a number of reasons. First, the manual process of tracking server configuration changes is prone to errors and omissions and produces unpredictable server recovery results. Second, identical recovery hardware is rarely available. The process of reconciling the differences between production and recovery environments takes too long and is unreliable. Third, backup datasets and server configurations are highly interdependent. Manually synchronizing data backups with server configuration information is also a complex, time-consuming process that extends the time to complete a server recovery beyond RTO targets. Finally, traditional server recovery methods consume too many staff resources and do not guarantee successful server recoveries.

Fortunately, new server recovery software offerings such as EMC® HomeBase are helping companies overcome these challenges by eliminating the manual processes that plague current server recovery methods. EMC HomeBase uses an innovative, automated approach to capturing server configuration information that enables daily tracking and monitoring of configuration changes without degrading production system performance. In addition, HomeBase automatically accounts for differences in source and target server environments so that the operating characteristics of the recovery server match those of the production server without a lengthy reconciliation process. Using a flexible API, HomeBase allows server recovery workflows to be easily integrated with data backup workflows, ensuring that server recovery information is always in sync with data recovery information.

Companies that use software offerings such as HomeBase for server recovery are able to successfully recover server environments reliably and efficiently. EMC HomeBase enables you to:

- Meet aggressive overall system RTO objectives
- Virtually eliminate the risk of unsuccessful server recoveries
- Enhance server configuration management capabilities
- Preserve the investments in existing backup systems and utilize production server and storage assets more efficiently

## Introduction

This white paper examines server recovery as a traditional weak link in overall recovery capabilities, and explains why traditional recovery methods fall short. It explains how to evaluate server recovery offerings and details what capabilities a comprehensive solution should provide. The paper also explains the benefits that server recovery solutions provide and details the capabilities of EMC HomeBase.

## Audience

This white paper is intended for technical or IT decision makers and influencers who may be in the process of evaluating backup and recovery solutions. They may also already have backup or recovery software, but are in need of a bare metal recovery and server migration solution across heterogeneous server.

## Server recovery: A weak link in overall recovery capabilities

Many companies, despite having invested significantly in tools and processes to protect vital data and applications from outages, are unable to meet increasingly aggressive recovery objectives. As tolerance for downtime of any kind—planned or unplanned—has continued to shrink, the need for faster, more reliable

---

disaster and operational recovery has grown. But until recently, server recovery—an important part of overall recovery—has received far less attention than application and data recovery. In what is known as a bare metal recovery (BMR) event, companies must re-create the server configuration from scratch before the application and application data can be recovered and made accessible to the end user.

Now organizations are finding that the traditional process of rebuilding the server environment for a recovery or migration is too complex, involves too many variables, and takes too much time to meet current RTO objectives. These companies are quickly realizing that server recovery is a weak link in their recovery chain.

With RTO and RPO requirements approaching zero, many of these same organizations are now looking for server recovery solutions that provide the same level of operational and disaster recovery capabilities as those delivered by application and data recovery tools. Fortunately, recent software developments are making it possible for companies to achieve reliable and efficient server recoveries even when identical recovery hardware is not available.

## Why traditional server recovery methods fall short

Today, many companies use one of two traditional methods to address server recovery. One is a completely manual process driven by change management documentation. The other relies on capturing server images to record configuration changes. Each approach has limitations. This is largely the result of four primary challenges companies face in dealing with server recovery.

- Server environments change frequently
- Recovery usually happens on dissimilar hardware
- Server configuration details and backup datasets are highly interdependent
- In-house server recovery is labor-intensive and time-consuming

### ***Server environments change frequently***

*The process of manually tracking detailed server configuration changes reliably and efficiently is labor-intensive and prone to errors or omissions.*

Change management as a whole is a long-standing problem for companies. As companies expand with remote locations, a mix of platforms, and additional users, the problem of tracking server configurations becomes a bigger challenge. The number of potential configuration changes has skyrocketed both in the sheer number and variety of changes that can occur as servers and applications are increasingly interconnected and complex dependencies between systems grow. Most companies change their production server environments daily through actions such as applying a system software patch or adjusting application storage and network settings for performance or security purposes. The smallest change, such as a driver upgrade, may be every bit as important to the overall system operation as is the current OS release level.

Aside from being prohibitively time-consuming, the manual process must deliver a comprehensive documentation of server configuration changes in order to support a successful recovery. Any change that is missing from the documentation or incorrectly recorded will likely create errors in the recovered server environment—errors that are especially difficult for companies to resolve in a disaster-stressed scenario.



**Figure 1. Tracking changes manually can be labor-intensive and prone to errors**

Using server images to record configuration changes is a step in the right direction. Unfortunately, it is also a brute-force approach that is ultimately not practical for several reasons. Full server images are large binary files that can occupy up to 15 GB of storage and normally take 30 to 45 minutes to produce. Because of the large operational window needed to capture a server image and the storage capacity each image consumes, an organization cannot reasonably take a server image much more frequently than once a month. When the production environment includes a large number of servers, this issue becomes more significant.

Given the amount of change the typical production environment experiences, companies should take daily snapshots at a minimum to ensure that the latest server configuration information is always available and a reliable server recovery can be achieved. If a week or more elapses between images, there will be no record of system, application, storage, and network changes that have occurred since then and a company will risk not being able to restore business servers back to the most recent state. From a change management perspective, there is also no record of the changes that have occurred prior to the image being taken because only the latest server image is maintained.

Replication of the server images to a remote site, which is a common requirement in disaster recovery, is often impractical as well because of bandwidth limitations. Companies usually do not have enough available network bandwidth capacity to complete the replication of a large server image quickly enough.



**Figure 2. Bandwidth limitations may prevent server images from being replicated quickly enough**

## ***Recovery usually happens on dissimilar hardware***

*Identical hardware is rarely available for the recovery server, which makes it difficult to develop a standardized process that is both reliable and fast enough to meet aggressive RTO objectives. One of the hallmarks of a reliable process is repeatability. Even if a company has accurately recorded all configuration changes, there is a very good chance that they must recover to different server hardware. Resolving the differences requires experts who are familiar with the interdependencies of each operating system and hardware platform involved. These experts must map the configuration settings from one hardware platform and OS combination to another using a specialized process that is difficult to execute accurately and quickly, and even harder to do consistently.*

Because a server image is represented as a binary object, it offers IT no real insight into the actual configuration of the server itself. System images inherently contain hardware specific information, which

---

must be removed in order for the image to be restorable onto different hardware. The success or failure of the system image approach depends directly on IT's ability to locate, remove, and replace this hardware information. As a result, companies that use a server image solution often see "all or nothing" results. In order to be confident of its ability to restore the servers, IT must repeatedly test these images on target hardware. This significantly raises total operating costs.

### ***Server configuration details and backup datasets are highly interdependent***

*Server configuration settings and backup datasets are highly interdependent and must be brought in-sync before a successful recovery can take place.* Organizations typically execute traditional BMR processes outside the normal data backup workflow. In a recovery event, it is difficult for IT to marry system configuration and backup data information reliably. For example, if the latest system image and backup dataset are taken at different times, any file system changes that occurred since the last system image was taken may create conflicts when the backup data is applied. Previously deleted file systems may reappear. Previously expanded file systems may register full file system errors. Directories that were moved can result in multiple copies of the same file being applied to the recovery server in both the old and new locations. Resolving these inconsistencies is a tedious process that adds significantly to the recovery time and raises the possibility of a system with consistency issues that are not immediately apparent.

### ***In-house server recovery is labor-intensive and time-consuming***

*An in-house server recovery effort consumes significant specialized technical skills and time.* Most companies are ill-advised to take server recovery on as an in-house challenge. There are too many variables resulting from constant vendor changes to system software, firmware, and hardware. Staying on top of it all requires highly specialized skills and experience with each element of the server environment. During a disaster or outage, technical personnel often become quite scarce, which can extend recovery times or even contribute to a failed recovery event. Companies that achieve the most success with server recovery look for help from providers that are focused exclusively on the problem of server recovery.

### **Assessing new server recovery software tools**

When evaluating the server recovery offerings that are available, companies should look for a solution that directly addresses these challenges. They should assess alternatives by comparing each offering's ability to deliver reliable, repeatable server recoveries even for large, mission-critical production server environments that contain a mix of platforms, networks, and storage topologies, and that span multiple locations. The best software tools will provide the capabilities needed for fast and efficient server recoveries to help meet aggressive RTO objectives as well as the growing demands for operational recovery.

---

## ***Capture configuration changes daily with minimal production impact***

Reliability is the first priority for an effective server recovery process. A fast, efficient, and repeatable process is useless if it doesn't deliver reliable results. This requires software tools that track and capture server configuration changes at a level of detail and frequency to allow a complete recovery to take place whenever needed (see box at right). The software should be able to capture server profiles at least once a day without overwhelming storage resources or affecting production system performance. A full accounting of configuration changes must, at a minimum, monitor the elements listed on the right.

## ***Enable recovery to different server hardware***

Because it is nearly impossible, both financially and technically, to guarantee recovery to an absolutely identical hardware environment, the software tools should facilitate the automated translation of the configuration settings from the production environment to the recovery environment. This includes identifying driver substitutions, resolving name differences, and others, as well as more obvious tasks like applying required operating system patches.

## ***Integrate a server configuration change capture process with backup workflow***

Reliability also demands that the configuration information be in sync with the production backup datasets to avoid any inconsistencies in the recovered system. Software tools that facilitate the integration of system recovery and backup workflows, regardless of the backup system used, help ensure that configuration information is captured at the same time backup datasets are captured. When a recovery is required there is no manual intervention needed to bring the system configuration information up to date with the backup dataset. This also helps preserve the investments in existing systems.

## ***Apply automation carefully to limit errors and speed server recovery***

In order to meet lower RTO objectives and increasing operational recovery demands, companies should look for software that allows users to control the automation of the server recovery process. The user should be able to easily control, store, and transmit all aspects of the server recovery process and data securely from a single web-enabled interface. This capability not only improves the accuracy of the overall recovery by reducing errors, but also significantly reduces the time to complete the recovery. Having less human intervention lowers the burden on IT staff and frees specialized, technical skills for higher value

### **A full accounting of configuration changes must monitor, at a minimum, the following elements:**

#### Server hardware:

- CPUs, memory banks and size, motherboard type, network cards, and other adapters

#### Peripherals:

- Internal disk (including SCSI, IDE, Wide-SCSI, and so on)
- External disk (SSA, NAS, and so on)
- Tape
- Removable devices

#### Operating system:

- Variety, version, patches and hot fixes, paging files and paging space, network domain, routing tables, partitions, volumes, file systems, internal and external mounts and mount points, daemon services and service state, network shares, NIS, DNS, YP, name resolution, network tuning, user accounts and groups, ACLs, WINS, and WLM

#### Software applications:

- Version, license, installation location
- Active Directory, Microsoft IIS, NIM, OpenMail (Sendmail), Oracle, PSSP, Sysvol, Terminal Services, TSM, FTP
- Customer-specific application customizations

---

information activities. A side benefit of faster, more efficient recoveries is that companies are able to conduct more recovery tests, which has an additional positive effect on recovery reliability.

When a company operates a large, heterogeneous production server environment, the automation capabilities become even more valuable. In these environments, companies must execute server recoveries in parallel to meet RTO objectives. Controlled automation in these environments is the only way companies can achieve server recovery in time and at a reasonable cost.

## ***Support virtualization to add flexibility and increase server recovery efficiencies***


Organizations should also look for server recovery software that supports virtualization on both the production and recovery sides. This capability adds another dimension of flexibility and further enhances efficiencies. For example, companies can use the virtualization features to maintain a single physical recovery server that will handle recovery for multiple physical production servers.

## **Meeting server recovery objectives with new technology**

EMC HomeBase is one example of a BMR software solution that is helping companies meet today's more aggressive server recovery demands. HomeBase is a completely new approach to the migration and recovery process for servers, applications, and storage structures. HomeBase is designed to automate the server recovery process and to eliminate the issues associated with current configuration recovery methodologies.

### ***Complete, accurate recovery***

What separates HomeBase from other solutions is the way its profile-based technology breaks down the server stack elements so that enough configuration detail is captured to recover the server completely. At a generic level, the server stack is separated into three layers: the application layer containing all applications and application-created data; the root layer containing information about all hardware drivers and executables; and the configuration layer containing hardware and software settings related to performance and security. HomeBase further refines the configuration layer into functional and support elements. The functional elements include the performance and security settings associated with the applications. The support elements include the server configuration settings associated with the attached storage and network hardware infrastructure, as shown in Figure 3. HomeBase uses this structure to capture all information required to execute a complete system recovery and creates a server profile containing only system-state related data.



**Figure 3. HomeBase defines configurations at the elemental level**

### ***Minimized production impact***

The sophisticated encryption and data de-duplication technology provided by HomeBase results in small, secure Wintel server profiles that consume only 3 to 5 MB of storage. HomeBase-captured UNIX server profiles are even smaller at around 500 KB per profile. The small size of these server profiles makes it practical to capture the server configuration as often as needed without overwhelming storage resources.

---

This also minimizes the bandwidth requirements if the server profiles must be replicated to a remote site, which is a common requirement when recovery must take place in a physically separate facility.

In addition, the operational window required for HomeBase to create a server profile is typically only three to five minutes. When this is compared to the 30 to 45 minutes required to create a server image, it is clear that the HomeBase profile image can be taken more regularly. Most HomeBase customers produce a HomeBase profile on a daily or more frequent basis.

One way HomeBase minimizes the impact on production operations is through its efficient Agent/Server design based on web services principles. A small, unattended Server Agent installed on each protected server captures the server profiles according to a schedule set up by the user on the HomeBase Server. Because the Server Agent executes in a low priority mode, the profile capture process does not impede application performance and makes frequent server profile captures even more practical. The efficient, lightweight design of the Server Agent and the compact server profiles virtually eliminate the chance of missing any configuration changes, which results in greater accuracy of the server recovery process and the ability to meet dramatically lower RTOs.

The design of HomeBase allows it to efficiently support server recovery in a wide variety of environments. For complex production environments with many servers and mixed platforms, HomeBase automates the migration and/or recovery of many servers in parallel. If an organization's production servers span multiple locations, the available HomeBase proxy server helps securely coordinate server recovery across multiple remote locations. As a result, many servers across multiple locations can be recovered cost effectively within the RTO objectives.

## ***Powerful process control capabilities***

The HomeBase server acts as a secure central repository for all server profiles captured by the Server Agents. The web-enabled interface of the HomeBase server allows users to create and manage all policies related to server recovery including how often server profiles are captured, how profiles should be stored, how long profiles should be retained, how configuration changes should be handled, and the specific OS and application software to load onto a recovery server during a recovery or migration event.

At the core of the HomeBase server is the Differential Factoring Engine. This unique technology accurately and efficiently identifies any potential differences between two server profiles, which is a valuable feature during both actual recovery/migration events and ongoing monitoring for configuration compliance purposes.

Each time a new profile is received from any Server Agent, the Differential Factoring Engine in the HomeBase server compares the new profile to the old profile from the same server to determine any configuration differences. Any task that is mandated by the configuration changes as defined in the user-created policies stored on the HomeBase server will be executed.

During a migration or recovery event, the Differential Factoring Engine automatically differentiates between the production and recovery server's profiles and identifies the steps necessary to configure the target server and storage to the required performance and security characteristics regardless of the differences between the two hardware platforms.

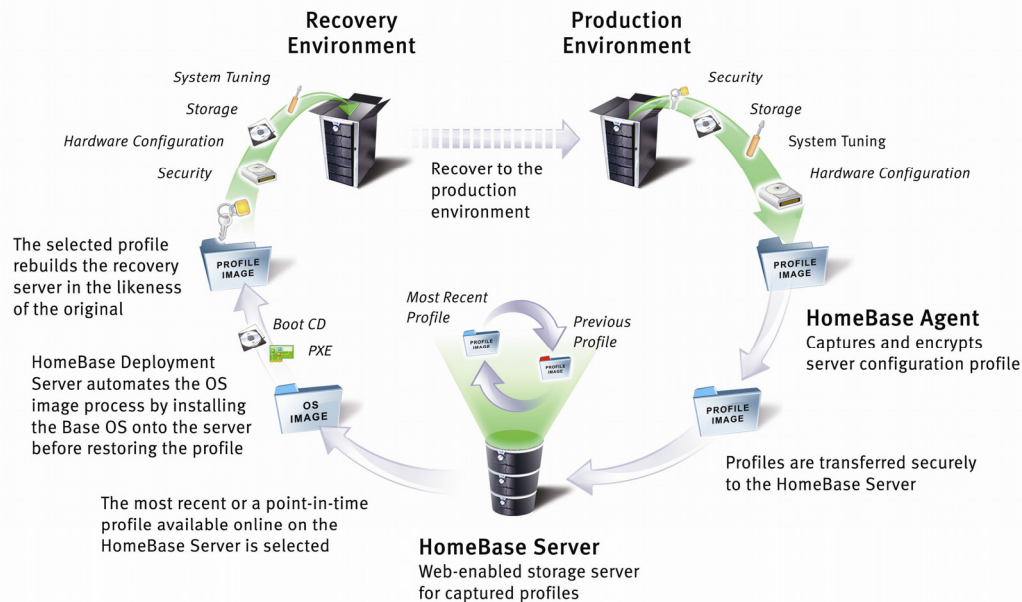
## ***Detailed record of configuration changes***

HomeBase can also help meet compliance requirements related to system change and auditability. Since HomeBase stores as many of the previous server profiles as defined in the retention policies, HomeBase is able to provide a secure, auditable trail of the detailed changes made to any server configuration that is being monitored.

## ***How HomeBase works***

HomeBase helps a company achieve a flexible, fully automated server recovery process that monitors changes made to a protected production server's configuration and executes a reliable server recovery or migration as disaster or operational recovery requirements dictate.

The HomeBase process works as shown in Figure 4.



**Figure 4. HomeBase's automated server recovery process**

On an ongoing basis, the HomeBase Server Agent installed on the production server automatically captures the server profile according to a schedule set up by the HomeBase user. The Server Agent automatically transmits the encrypted, compressed profile securely via a user-selected protocol to the HomeBase Server where it is stored and retained according to the policies defined by the user.

Whenever a new profile is received by the HomeBase Server, the Differential Factoring Engine compares the new profile to the previous profile from the same server and determines how the configuration has changed. If there are changes, HomeBase checks the configuration against the conditions set up in the user-created policies to see if any action is required. If so, HomeBase will automatically execute the action, such as sending an e-mail alert notifying stakeholders that a configuration element does not comply with corporate standards, or generating an SNMP trap to transmit to the network management service in place.

When a production server must be migrated or recovered, the HomeBase user simply selects the source production server and the target recovery server via the web interface. The Differential Factoring Engine automatically identifies the differences between the current profiles of the production and recovery servers and identifies the appropriate configuration settings for the target server and storage environment. Following a deployment template predefined by the user, the HomeBase Server begins an automated recovery of the production server environment. HomeBase first installs the required OS on the recovery server, followed by the necessary hardware drivers and executables, file system, and third-party software, and then applies the new configuration settings for the recovery environment so that application data recovery can begin immediately. At this point, the backup system applies the backup dataset for a fully recovered and operational system.

## Lower risk, improve efficiency, and increase availability

Companies that utilize a server recovery software solution like HomeBase realize benefits that extend beyond an improved server recovery process.

### **Reliable, cost-effective, hardware-independent recovery**

The innovative HomeBase approach to server recovery allows servers to be recovered accurately and completely even when identical hardware is not available. HomeBase easily handles the volume of detailed

---

configuration changes that occur in the most complex, mixed production server environments. Its efficient web services design and compact server profiles allow detailed server changes to be captured as often as needed without affecting production system performance. For companies with both Windows and open system servers spanning multiple locations, the workflow and profile replication capabilities provided by HomeBase facilitate parallel recovery to remote locations. This results in reliable protection for hundreds of servers at a fraction of the cost of traditional methods.

Powerful HomeBase administration features help significantly lower the operating costs associated with the server recovery process. The web-enabled interface improves ease of use, allowing users to preset OS install and deployment templates to simplify the recovery process. With HomeBase, companies need fewer technical resources to recover a greater number of production servers. In addition, HomeBase helps companies use existing resources more efficiently. Minimal storage and bandwidth requirements and support for virtualization technologies like VMware<sup>®</sup> let organizations extend the economic life of existing server and storage assets and reduce overall cost.

### ***Fast recovery and improved production availability***

HomeBase uses its design efficiency and process repeatability to automate the steps of server recovery and eliminate the time-consuming, error-prone manual processes associated with traditional BMR approaches. With HomeBase, companies are able to meet RTO objectives as low as minutes with a significantly lower risk of recovery failure. A faster, less intrusive recovery process makes it practical to conduct more comprehensive and frequent testing of the recovery process. Production system availability is improved because the downtime associated with server-profile capture and migration or recovery testing is minimized.

### ***Managed server assets and leveraged investments***

Companies using HomeBase improve more than just overall system recovery capabilities and efficiencies. The software's support of open interfaces and design standards allows organizations to easily integrate HomeBase capabilities with existing IT investments including data backup and replication systems. The ability to integrate with an existing backup system is a critical feature required for successful overall recovery. HomeBase does not force a company into a specific backup solution. Instead, it ensures that a company's previous backup investments are fully leveraged. The product's open API lets companies easily integrate the server recovery workflow with the backup workflow produced by any of the leading backup and replication systems including EMC NetWorker<sup>®</sup> or Symantec NetBackup. HomeBase also complements EMC backup and replication products.

From a configuration management standpoint, HomeBase can also help manage large server farms. Using the software's powerful configuration monitoring and control capabilities together with installed CMDB solutions, companies can efficiently propagate required configuration changes to a single server or group of servers, or globally to all servers through the intuitive web-based HomeBase interface. IT can distribute and install standard configurations, OS builds, and third-party software across Windows platforms as part of general server provisioning, migration, and repurposing. With just a few clicks, HomeBase can generate detailed, component-level reports to export into standard formats for asset reporting.

## **Conclusion**

As recovery demands increase, it is clear to many companies that server recovery is an essential element of overall system recovery. New software solutions like HomeBase are helping organizations finally achieve the automated, hardware-independent server recoveries needed to meet today's shrinking RTO objectives. HomeBase delivers the reliability levels needed to deploy server recovery in even the most complex production environments that undergo significant daily change. HomeBase easily accommodates that change to deliver fast, efficient server recoveries without degrading production operations in the process.

The open, flexible design of HomeBase helps extend its benefits by allowing organizations to combine its capabilities with other installed systems for more efficient server provisioning, migration, and

---

configuration management. Through a standardized API and SNMP interface, companies can integrate HomeBase with existing systems like backup and CMDB solutions to preserve and expand the value of previous IT investments.

Solutions like HomeBase allow companies to properly address the challenges of server recovery and achieve higher levels of protection. Organizations that use HomeBase are able to not only lower the risks of recovery failure, but also eliminate significant operating costs and in the process extend the value of existing infrastructure investments.

HomeBase is just one element of EMC's comprehensive information protection solutions that address everything from data backup and replication to highly available mirrored operations.

To find out more about EMC HomeBase and other EMC information protection solutions, please visit [EMC.com](http://EMC.com).