



Focus On...

Bare-Metal Recovery: EMC HomeBase

**A White Paper
by Focus Consulting
Anne Skamarock
Barb Goldworm
October 2007**

Focus On... Bare-Metal Recovery: EMC HomeBase

Contents

EXECUTIVE SUMMARY	3
MARKET DRIVERS FOR BARE-METAL RECOVERY	3
BARE-METAL RECOVERY OPTIONS	4
EMC HOMEBASE SOLUTION OVERVIEW	6
AUTOMATED CONFIGURATION CAPTURE	6
FLEXIBLE SUPPORT FOR REMOTE DEPLOYMENT	7
HARDWARE INDEPENDENT RECOVERY	8
INTEGRATES WITH EXISTING DATA PROTECTION ENVIRONMENT	8
END-USER FEEDBACK	8
KEY BENEFITS	10
FOCUS ASSESSMENT	11

Focus On... Bare-Metal Recovery: EMC HomeBase

Executive Summary

CIOs many not be happy about the amount of time and money it takes to protect corporate data assets, but the key question is this - are there solid plans in place for protecting the company in the event of a disaster? In today's computer-based business environment, when business-critical servers are down, what matters is time-to-recovery. When a disaster hits, recovery often starts with bare-metal servers that have to be recovered to full-production levels as fast as possible.

This white paper discusses the drivers for bare-metal recovery solutions and the current options and issues associated with bare-metal recovery, and outlines how EMC's HomeBase solution for bare-metal recovery addresses this hot issue. In addition, both the value and the limitations of the HomeBase solution are reviewed, based on end-user feedback from long-time HomeBase customers. Key benefits of the HomeBase bare-metal recovery solution are summarized, closing with an overall Focus Assessment.

Market Drivers for Bare-Metal Recovery

Disasters come in many forms - both natural (floods, hurricanes, tornados, earthquakes, etc.) and man-made (bombings, terrorist attacks, catastrophic accidents, etc.). Over the last decade, awareness that disasters can destroy a company has risen markedly. On a smaller scale, even though companies employ all sorts of mechanisms to protect their environment against hardware failures, failures happen. In the event of a disaster, large or small, companies of all sizes need to be assured that they have the ability to resume business operations quickly, either locally or at another location, within twenty-four to thirty-six hours of the disaster, or else face closing their doors permanently.

To protect corporate assets from data loss due to human error, most companies have developed backup capabilities. However, many companies have overlooked the need to recreate the servers themselves - the entire server - from bare hardware, in the event the servers fail or are physically destroyed. Some have considered it, but in the past, could not financially justify the expense of additional hardware, planning and testing. A much-cited study in Contingency Planning and Managementⁱ found that 40 percent of businesses that were shut down for only three days, failed within three years. This type of information has given many businesses the financial push they needed to address this issue as part of their standard practices.

However, addressing server recovery has been challenging due to a variety of factors which complicate the process. The complexity of configuration and state information within the operating systems (particularly with Windows) is often compounded by the lack of any automated change management process, resulting in a sea of undocumented changes to OS areas such as the registry. In addition, either the hardware used to recreate the servers at the

“Companies of all sizes need to be assured that they have the ability to resume business operations quickly, either locally or at another location, within twenty-four to thirty-six hours of the disaster, or else face closing their doors permanently.”

Focus On... Bare-Metal Recovery: EMC HomeBase

new site must be exactly the same as the original hardware (which can significantly increase costs), or the re-creation must be able to work across dissimilar hardware. Lastly, the large size of OS images can limit the practicality of server imaging within a reasonable time frame and over limited bandwidth. All of these challenges require careful consideration and planning when evaluating server recovery.

The process of recovering one or more application environments onto a physical server with no operating system or applications pre-loaded (bare-metal), and bringing it back to a production state that exactly mirrors the original production application environment is known as bare-metal recovery (BMR). BMR is often performed using whatever hardware is available at the time, making the recovery that much more complex. To accomplish this level of recovery, simply backing up

application data or using disk-imaging schemes, is not enough. Configuration and state information - operating system, user, and application - must be captured and retained as well. EMC's HomeBase bare-metal recovery solution offers an excellent example of how to deliver this critical capability, facilitating bare-metal recovery of production systems.

Server Recovery Challenges

- **Complex and changing OS configuration and state information**
- **Lack of change management resulting in undocumented changes to OS (e.g. patches and registry changes)**
- **Dissimilar hardware at recovery site**
- **Bandwidth limitations making offsite transfer of large images impractical**

Bare-metal Recovery Options

A number of options are available to businesses for BMR. They span a wide range of technologies, capabilities, and costs. These options range from attempting to use only backup data stored either on tapes or disk, through elaborate remote disk imaging techniques. These techniques are all valid but, in an operational environment, must be thought through and planned carefully.

For example, one may choose to backup everything that is on a system, including the operating system and its associated files. However, because the operating system files change very infrequently, it is difficult to justify using up large amounts of tape or disk space for these files, the bulk of which never change. In an effort to backup only the associated operating system files that do change, some will cherry-pick the files known to be important just to find out, while attempting a recovery, one or two "important files" were missed. Backing up application data is an important process to ensure business-critical data security and availability. It is not recommended that one use backup for collecting all the necessary system configuration and state information required in recovering a system from the ground up, unless the process of understanding what files are needed is automated and done by solutions that are certified by the operating system vendor.

For faster data protection, many companies look to disk imaging techniques to map all data to a twin disk. For single server recovery, even two or three of the most important servers, this is an excellent option. However, if there are hundreds of servers to recover in the event of a disaster,

Focus On... Bare-Metal Recovery: EMC HomeBase

disk-imaging techniques can pose some difficulties. While faster than traditional backups, disk imaging can take 30 or more minutes to transfer data off the system disk, and that is when the second disk is local. Transferring terabytes of data over a wide-area network to a remote facility can be quite time consuming.

Often, BMR solutions require hardware components that match - *exactly*- the hardware configuration of the original server. Duplicating a production data center at a remote location, especially when the remote systems sit idle, can be extremely expensive. Naturally, a similar processor must be used in the recovery system for all the applications to be able to run. However, some bare-metal recovery products require memory sizes, storage interconnect and sizes, and exact file system sizes and mapping to be the same.

This can be due to operating system requirements or BMR solution requirements. For example, if a disk imaging technique is used in a Windows® environment, where the operating system partitions are captured and stored at the remote site, recovery **cannot** be done on dissimilar hardware. When the Windows system is installed, the configuration is baselined. From that time on, when Windows boots, it requires the mapping of the kernel, the disk storage controllers, and the hardware abstraction layer (HAL) to match the original baseline, or the boot will fail.

To enable the use of dissimilar hardware, customers must amend their recovery process to first install the operating system on the bare-metal. When the operating system is installed “fresh” it will create a new baseline with the hardware configuration of that server. Once the operating system has been applied to the hardware, service packs and bug-fixes must be applied. Then the system must be configured for the operational environment. All this must be done before beginning to restore applications and data – taking precious time, sometimes days, and personnel. Once again, the process relies on people logging or remembering all the configuration details needed to bring up a system.

Manual tracking of configuration changes can be fraught with dangers, especially in the case of bare-metal recovery scenarios. Recovering one server may not push the boundaries of the process however recovering even ten servers will. Paper logs can be lost, and electronic files (spreadsheets, the most common configuration change-control tool) can be deleted or not transferred to the remote site or simply not used. In the event of a real disaster, the personnel most familiar with the systems and their associated configurations may not be available to the remote site during recovery. It is necessary to rely on an automated process, without human intervention, which stores all required data at the remote recovery site for automated use in the recovery.

When choosing a bare-metal recovery solution, finding one that integrates with your current processes and addresses your system needs is crucial to creating successful recovery scenarios. The Focus assessment is that EMC’s HomeBase BMR solution has the capabilities and flexibility to meet these criteria.

Focus On... Bare-Metal Recovery: EMC HomeBase

EMC HomeBase solution overview

Recovery of any server must encompass the entire server ecosystem. This ecosystem includes applications and their associated data, the operating system, and the hardware. Each of these layers has associated configuration information such as software levels, users, security, performance settings and more (see Figure 1).

Many IT shops do an excellent job of securing the application data in backups and stop there. EMC's HomeBase solution works in conjunction with traditional backup, capturing point-in-time images of critical operating system and application configuration states to assure total system recovery from bare-metal hardware.

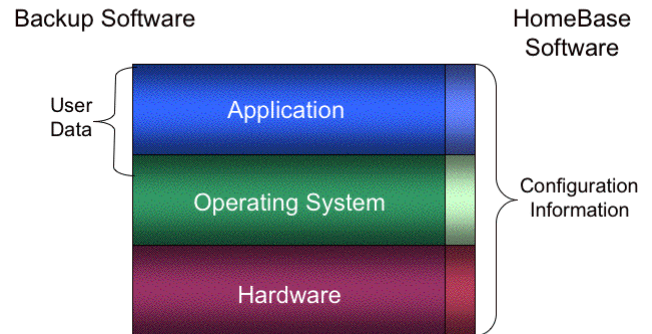


Figure 1: System configuration layers

These images, called HomeBase profiles, are created according to user-defined schedules. Once a profile of the Production Server is created, it is sent to the HomeBase Server for use during a recovery. The information contained in the profile contains only configuration information so its small size makes it easy to transmit to a remote HomeBase Server.

When servers need to be recovered, these profiles are used by the EMC HomeBase Agent, also installed on the recovery system, to bring the Recovery Server (Server B in Figure 2) to a production state. The HomeBase Agent is responsible both for creating and restoring profiles, to maximize ease of use.

Automated configuration capture

The EMC HomeBase bare-metal recovery solution, once configured, automatically snapshots all configuration and state information on business critical servers at the scheduled intervals. The profile created from the snapshot is sent to the EMC HomeBase Server and is both compared to the previous profile and tested against user-defined business rules. If the comparisons show changes in the hardware or software configurations, such as the addition of a new application or storage device, the EMC HomeBase Server will generate an alert reflecting these changes. Improving significantly over manual change management, EMC HomeBase automates the tracking of configuration changes, allowing the administrator to verify that change management procedures have been followed and, if

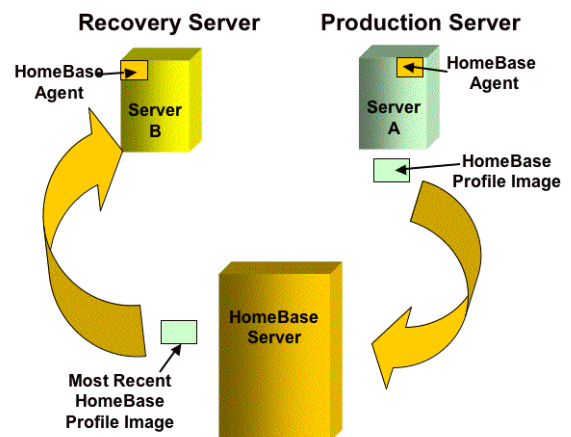


Figure 2: EMC HomeBase profile and recovery cycle

Focus On... Bare-Metal Recovery: EMC HomeBase

necessary, to adjust the disaster recovery plan to accommodate the change. It also allows the administrator to back out unwanted configuration changes. If, for some reason, the EMC HomeBase Server does not receive an updated profile at the scheduled time, an alert will be generated notifying the administrator of the missing profile, to avoid gaps in the profiles.

Flexible support for remote deployment

Unlike a full server snapshot which can take awhile to create and requires large amounts of storage capacity and possibly network bandwidth, EMC HomeBase profiles take less than five minutes to execute and require only kilobytes of storage. This makes it easy to integrate into a production application environment unobtrusively. The small size of the profile makes trivial sending the data across the network to a remote location, rather than having to schedule data transfers across the network at off-peak hours or risk saturating the entire network. This small footprint allows IT to build a DR plan that coincides with the business objectives rather than one that is dictated by the recovery software. EMC HomeBase can be configured to run as often as necessary, e.g., once an hour, or if changes to the environment rarely take place, it can be run less frequently.

The small profile size also allows many servers to utilize a single EMC HomeBase Server, gaining efficiencies and lowering total cost of ownership (see Figure 3). One large pharmaceutical company uses EMC HomeBase to profile 18 servers. Their EMC HomeBase server is located at their remote site, hundreds of miles away. According to the Storage Director, "The profiles take up less than 100 MB's for all 18 systems."

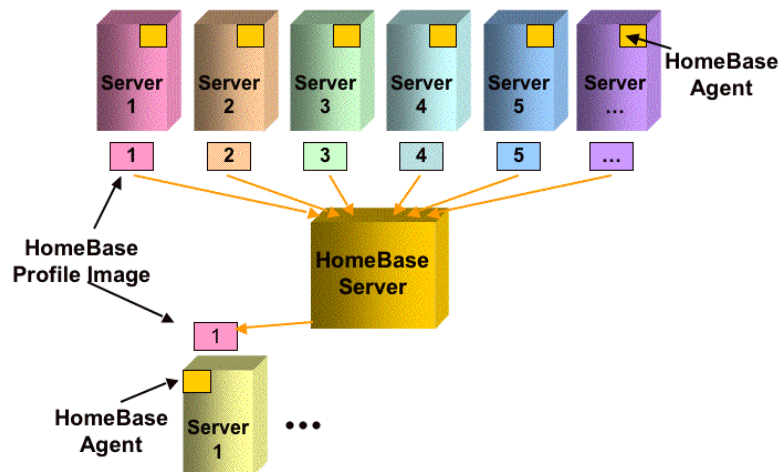


Figure 3: One EMC HomeBase server supports many production servers

The EMC HomeBase solution supports recovery of a single server through a point-and-click graphical user interface. Multiple servers can be recovered simultaneously through batch interface to the EMC HomeBase Server software providing flexible recovery alternatives to align with business needs.

Focus On... Bare-Metal Recovery: EMC HomeBase

Hardware independent recovery

Along with flexibility of the runtime environment, EMC HomeBase supports recovery to servers with dissimilar hardware configurations. While the processors must be of the same family, e.g., x86, memory, networking and storage do not have to mirror the original configuration. To bypass operating system issues, this solution requires a clean version of the operating system to be loaded on the server first. Then EMC HomeBase uses the production server configuration information to intelligently and automatically configure the recovery environment. This saves cost for the customer as they can populate their remote site with a less expensive or newer/faster failover infrastructure. In contrast, relying completely on image technology for disaster recovery requires an identical infrastructure at the recovery site, making it a far less flexible solution.

Integrates with existing data protection environment

EMC HomeBase operates in conjunction with the application backup software, such as EMC NetWorker, EMC Avamar and EMC Retrospect, as well as others. The backup software typically collects and restores user, application, and some operating system data while EMC HomeBase profiles, captures and restores the associated server configuration. EMC HomeBase's short runtime window and small data size makes it affordable to run many times during the week to capture all configuration changes without having to rely on manual configuration control policies. This, together with the associated backup data, provides an efficient, effective and affordable disaster recovery solution.

End-user feedback

End-user feedback is very positive on HomeBase. For example, according to one Storage Director from a large pharmaceutical company, and a Systems Engineer from a large insurance agency, their firms have been extremely happy with the HomeBase solution. As the Storage Director put it, "There is nothing else on the market like it."

The pharmaceutical company is a large multi-national corporation with over 74,000 employees worldwide, selling products into over 150 countries. The Storage Director is located at the U.S. corporate offices but has worldwide responsibilities. The corporate offices started out using a service provider that utilized the HomeBase solution (then called URM, Universal Recovery Methodology) in conjunction with tape backup. They were happy with HomeBase but after two years of not being able to have a smooth DR test, were unhappy with the group providing the service. A year ago, they decided to take control of their own DR process and personnel. To do this, they had to lease space for the DR site hundreds of miles away, purchase recovery servers, and software licenses and hire personnel to staff the site. "Even with the cost of going out and finding a remote site, purchasing and maintaining the equipment, software and licenses, and hiring personnel to work at the site, we are spending about one-third of what we would have paid the service provider."

***"We are spending
a third of what we
would have paid a
service provider."***

Global Storage Director
Multi-national
Pharmaceutical Company

Focus On... Bare-Metal Recovery: EMC HomeBase

Today this company is able to successfully test and restore their production servers reliably. They have one HomeBase Server at their remote (DR) site that captures profiles, because the amount of configuration data collected is fairly small. They have the HomeBase client loaded on 20 production servers (all systems are in one Data Center). They capture a profile once a week and keep all profiles for one year. This process creates less than 100 MBs for all the systems. "We currently have a manual process for recording changes done during the week but we may switch to performing the HomeBase profiles more often to capture the changes." HomeBase is now part of this company's test cycle when bringing up a new production application.

The insurance agency is a subsidiary of a larger corporation which uses a service provider for its DR process. The insurance agency has roughly 100 business-critical servers with roughly another 100 from the parent company all running Windows 2000 or 2003 with a recovery time objective (RTO) of 72 hours. The insurance agency and parent company are located in the mid-west of the United States with the service provider's DR location in the New York City area. "The HomeBase Agent installation and configuration [on production servers] is very easy."

"HomeBase allows us to recover in less than half the time we allotted [for our recovery time objective]."

Systems Engineer
Large Insurance Agency

This company uses SRDF (Symmetrix Remote Data Facility) to backup their application data to the remote site and HomeBase to profile their system configuration data. They test their DR process once a year and have performed two successful tests. "HomeBase allows us to recover those systems in less than half the time we allotted" thereby exceeding RTO.

Both customers agreed that once up and running, they have never seen a problem with HomeBase itself. However, the customers described a couple of areas for consideration and improvement with the product.

Users agreed emphatically that, "It's important to understand that every application environment is different and therefore needs different procedures for recovery. *Every application is unique!*" For example, the pharmaceutical company found that in their environment Active Directory and SQL Server can be tricky because both these environments add complexity. And according to the SE with the insurance agency, "60% of our servers we can recover in our sleep. Our challenge comes with the 20 systems running MSCS," Again, because of added complexity, the insurance agency has difficulties fully recovering their Microsoft Cluster Server (MSCS) configurations as a cluster. To solve this, they restore each system as a standalone then rebuild the cluster.

With any network-based solution, proper sizing of the network is critical to success. The need for incorporating this into best practices was emphasized by a comment from the insurance agency using the remote service, "Even though the profiles only take 1½ to 3 minutes to run, we can only send five profiles at a time." The SE doesn't know why this is but suspects it might be the server provider's network. While this is not a HomeBase-specific issue, it provides insight for those who are in the process of or about to implement their own DR strategy.

Focus On... Bare-Metal Recovery: EMC HomeBase

The only complaint they both had with the product was that the interface to the HomeBase Server is not as friendly as they would like. As it turns out, both customers interviewed are running older versions of the HomeBase product. Having listened to their customers, EMC's most recent release of HomeBase targets this very issue with significant interface and usability upgrades.

On a final note, both customers are pleased that EMC recognized the value of this product enough to purchase the Indigo Stone Company and hope EMC puts some money into making it easier to use.

Key Benefits

EMC's HomeBase solution works in conjunction with existing data protection processes to allow bare-metal recovery of a single server or hundreds of servers. It has the ability to automatically determine configuration updates made from the base operating system and to use that information to restore remote systems to a production level quickly. This has the added benefit of giving IT the ability to track and validate change management automatically, rather than relying on manual methods.

HomeBase's flexible runtime environment allows profiles to be captured as often as required to ensure accurate server recovery. The quick snapshot and small profile size lend themselves well to remote transfer of profile data many times per day, if needed, providing a non-intrusive easy to deploy DR solution

Using HomeBase allows customers to achieve true hardware independence at their DR site, with the possibility of significant savings in hardware costs. The small profiles allow current networking infrastructure to be used to transfer to the remote site rather than requiring the purchase of large WAN pipes.

Key Benefits

- **True hardware independence**
- **Non-intrusive, easy to deploy**
- **Cost-effective protection**
- **Remote server recovery**
- **Asset tracking/reporting**

Due to its efficient recovery model, the HomeBase BMR solutions saves costs associated with server downtime. With an estimated cross-industry average cost of downtime between one and 1.6 million dollars per hour, reducing recovery time from to 72 hours to 24 hours, as experienced by the insurance agency, equates to between 48 and 76 million dollars in savings. HomeBase provides cost-effective protection for all servers in the enterprise, no matter how big or small.

Focus On... Bare-Metal Recovery: EMC HomeBase

Focus Assessment

Focus has long supported the inclusion of DR process, planning and testing for all companies, regardless of size. We find great value in a solution that easily integrates with existing data protection and IT processes, saving customers implementation costs. When disaster hits, any solution that can cut time-to-recovery in half not only saves the company money but may simply save the company.

Focus believes that the HomeBase solution, especially with the resources that EMC can bring to bear on continuing to improve the interface and ease of use, will help more companies incorporate bare-metal recovery capabilities into their disaster recovery/business continuance plans and processes.

¹ "2001 Cost of Downtime – Eagle Rock Alliance." Contingency Planning and Management

About the Focus Analysts

Anne Skamarock, Research Director with Focus, has spent nearly 30 years in software engineering and technical marketing, as an end-user, vendor, analyst, and author, with Sun, SRI, Solbourne, StorageTek, and Enterprise Management Associates (EMA). For the past 10 years, she has focused on systems, software, storage, storage networking, and storage management solutions, and the intersection points between systems, software and storage. A frequent speaker at conferences, she was co-creator and Program Manager for the ROI Planning Lab at the 2007 [Server Blade Summit on Blades and Virtualization](#), and chaired *Interop's* Network Storage Track for several years. She has been a regular expert columnist for *Network World* and *TechTarget* and is co-author of the book, [Blade Servers and Virtualization: Transforming Enterprise Computing While Cutting Costs](#).

Barb Goldworm, president and chief analyst of Focus, has spent 30 years in technical, development, marketing, sales, senior management, and industry analyst positions with IBM, Novell, StorageTek, EMA, and multiple startups. Barb chaired the [2007 Server Blade Summit on Blades and Virtualization](#), is virtualization chair of the [2008 Blade Systems Insight](#) conference, created and chaired the Network Storage Track of *Interop*, and was one of the top 3 ranked analyst/ knowledge expert speakers at SNW. Barb is on the advisory board for several TechTarget sites, and has been a regular expert columnist and speaker for *TechTarget*, *Computerworld Storage Networking World Online*, *Network World* and *Virtual Strategy Magazine*.

Focus delivers research, analysis and consulting, focused on systems, software and storage in the enterprise, SMB, and government markets. Focus areas include: Systems, Storage and Enterprise Management (Physical and Virtual); Server, Desktop and Application Virtualization/Streaming; High Availability, Disaster Recovery, Business Continuity, Backup, Data Protection; Storage Networking (NAS, SAN, Fibre Channel, iSCSI); Storage, Networks, I/O and File Virtualization; Storage Technologies (Clustered File Systems, Data deduplication, VTL, etc.); Blade Systems (Server, Workstation and PC Blades); and Business Benefits of Technology (ROI, TCO). www.focusonsystems.com