



## Data Management for the Mid-Sized Enterprise — Enigma Data Solutions' *SmartMove*

Analyst: Anne MacFarland

### Management Summary

The pervasive clutter of enterprise data is getting worse. There are more data capture points. More people are hitting on the same data sets for more kinds of analysis to wring out costs and extend the appeal of everything you do. If you have virtuously merged your stovepipes of data, the problem is exacerbated.

What the enterprise needs for its data is the technology-equivalent of an intelligent closet – a closet that is infinitely expansive, so everything you cram into it fits; infinitely knowledgeable, so everything you want is easily at hand; and also able to provide all your sports coats or dress shirts on demand from which you can choose. Of course, it would also keep all your possessions safe for as long as you mandated, and would prevent other people from borrowing or throwing away your stuff. However, **when the mid-sized enterprise looks for something with which to manage its data, they are often faced with choices that are too big, too skimpy, or with features that just don't fit their needs. Many technology solutions address the problem from a storage-centric, not business-centric, point of view.** They embody the closet and its needs, but don't address for what purpose a business is using the closet. Moreover, for the smaller enterprise, many solutions involve IT expertise beyond what your enterprise can support on an ongoing basis

Instead of thinking about what you want to buy, think in terms of what you want to do, beyond the basic of replication for data protection (also known as *backup* and, over long distance, *business continuance*). As a smaller enterprise, you have data being generated in a number of places (servers, laptops, appliances) and stored both on the networks and on individual devices. This data needs to be accessible to everyone who has a business need for it. Data access times should be reasonable, but also should match the priority of the need for information, particularly when some resources may have to be redirected to urgent, unexpected business demands. And business data needs to be much, much easier to find.

To address this situation, Enigma Data Solutions of West Sussex, England, and Houston, Texas, offers *SmartMove*, a product that helps reasonably intelligent person characterize and create policies for data stored in CIFS or NFS files. It then takes those policies and creates an intelligent enterprise information environment that can respond to organizational change as well as to spikes in demand while keeping business priorities front and center. Read on for the details.

### Three Business Priorities for Data Management

#### *1. Manage the Data by Its Business Value, Yet Be Responsive to Sudden Business Demands*

The policies governing the data management systems must be flexible enough to allow for contingencies of drastic change in demand. If a surge in demand requires access to some files be degraded, it must be access to less important data that is affected. Laissez-faire or hands-on storage administration is increasingly inadequate, but simplistic policies and

### IN THIS ISSUE

> Three Business Priorities for Data Management.....	1
> How Enigma's <i>SmartMove</i> Works.....	2
> Conclusion .....	3

ruthless automation will not suffice, either. This is an area where business-level controls must prevail.

## **2. Migrate Data When it is Necessary or Useful, But Leave Enterprise Data Structures Unified and Access Permissions Intact**

Migrating data to improve data center performance and reduce data center costs is only acceptable to the business side of the house if their data appears as if it was where it always was, and if the access permissions remain intact. The key to accomplishing both these requirements is to leave a link or stub, so that, from the point of view of the application and the permissions embedded in it, the data is still in the same place. The link or stub redirects the fetch, and it takes a little longer, but not much, particularly since the data moved is, by policy, inactive. This is like a dimensional warp, which is how everybody would like their closets to work, if there were a dimensional pocket.

## **3. Keep the Data Safe, Secure, and Conformant to Government/Governance Regulations**

One copy of electronic data is never enough. Data must be replicated for redundancy, and replicated remotely for business continuity. This replication must be risk-free. The original data should not be reduced to a stub until the data copied to the new location has been verified. SmartMove does this. Retention periods can be written into the rules. ACLs<sup>1</sup> are maintained, so restrictions on access to data are not compromised. In addition, if needed for full compliance, SmartMove can write to unalterable media or a Compliance storage appliance as a secondary target.

## **How Enigma's SmartMove Works**

SmartMove is a data groomer. Used in an environment with heterogeneous storage devices of different quality and support costs, SmartMove leaves the control of the business policies to business people, and automates the transparent migration of data of lesser value (as defined by those policies) to less-expensive storage. It is business-facing, lightweight, data management – a less draconian, labor-intensive form of the old Hierarchical Storage Management (HSM). Evoking secondary targets, it can be a lightweight, opportunistic form of data replication. It is a full Information Lifecycle Management solution. And, it also can be used for enterprise data archiving.

SmartMove scans the file system from the server where it is installed, not with agents. It is in the data path during this scan, which is often done during early morning hours or on a weekend, but is otherwise out of the data path, including during data migrations.

<sup>1</sup> Access Control Lists

This reclusive behavior is appropriate to a non-production utility process, and minimizes any effect on infrastructure performance. As an opportunistic utility function, SmartMove it is not a single point of system failure. If a job fails, it is retried later; just as open files are moved later, when they are closed. Data is never at risk.

## **SmartMove's Infrastructure Role**

The SmartMove Software runs on *Windows (XP Pro, 2000 and 2003)*, *Linux (SuSe and RedHat)*, *Sun Solaris*, or *EMC Celerra Dart 5.3*. The server capacity it demands will depend on the volume of data to be moved, especially on how much has to be moved concurrently. SmartMove comes with a Hypersonic Database to store migration information, filters, and rules. If the volumes and the transfer rate of data are high enough, the software can be configured to use an Oracle database instead<sup>2</sup>. Administration of the software can be done from anywhere on the network (LAN).

## **What Storage Devices and Applications SmartMove can Use**

SmartMove can migrate any CIFS/NFS file from any storage device - from Mainframe to JBOD to EMC's *Centera* to tape. For both source and destination, it works with SAN, NAS, or DAS. This lets an enterprise get more business value out of the kinds of storage it already has, particularly if the enterprise topography includes two sites that can be the disaster-recovery remote sites for each other.

The targets of a migration can also be Tivoli *Storage Manager*, Legato *DiskXender*, Veritas *StorageMigrator*, or other HSM management software. The data can be migrated to any NFS or CIFS storage. Since SmartMove was developed as a migration tool for Network Appliance, it can use a Network Appliance Filer with *SnapLock* as a secondary target to create a guaranteed retention archive that will comply with government regulations. SmartMove also works with EMC *Celerra* and, on *Solaris*, as a gateway to EMC *Centera*. Another Enigma product, *Linear Disk*, presents *Centera* as a file system, so SmartMove can migrate files to *Centera* directly, without having to negotiate *Centera*'s API.

## **How to Use SmartMove**

The first step is for business managers to set policies, perhaps using the modules described in Exhibit 1. These policies can then be aggregated into rules that will govern the IT data migration jobs.

The *Archiving Module*, using the Bundles

<sup>2</sup> 1 TB data sets are the limit of what the Hypersonic database can handle.

described in Exhibit 1, below, can be used to move completed files and other inactive data to a separate structure when and if that becomes desirable. This function is part of the basic SmartMove product.

Through an optional *Reporting Module*, an administrator can preview the data to which the rules apply to make sure that the files should be moved. The preview also indicates how much storage will be freed up – whether it is enough may be just as important as or more important than what files are moved. The reporting tool can also scan file systems on a regular basis, so the administrator can instantly

#### Exhibit 1 - *SmartMove* Vocabulary

- **Projects** - The File folders in which files are kept become “projects.” Projects are containers characterized with useful metadata elements including some that can be added by the business, as it desires. By specifying Project as a unit of aggregation, capacity planning and costing of this aggregation become possible, which gives the business an idea of what its business processes and initiatives cost.
- **Modules** - SmartMove provides pre-defined rules and policies for specific industries and their regulatory needs, known as Smart Modules.
- **Filters** - Like people, data is not stale just due to type, age, business use or lack of access, but usually due to some combination thereof. You can set up multiple parameters by setting up separate filters for age, file type, business priority, etc. You could set up separate filters for end of month or end of year, when some otherwise inactive data becomes suddenly more valuable.
- **Rules** - Rules are aggregations of filters. They can thus be complex where necessary, but easy to decompose and recompose when change is necessary.
- **Jobs** - These rules can be applied to migration “jobs.” By these tools, a businessperson can characterize how and with what frequency files should be migrated.
- **Bundles** - Smart Move’s Archive feature, an option for licenses under 10 Terabyte, will take all the data from a project (pulling back data that has been migrated elsewhere) and will *bundle* it into a TAR file on secondary storage, preserving all the original dates and ACLs (access permissions). More metadata can be added at this point, if desired.

model various what-if scenarios for data migration.

The optional *Unmigrate Module* allows an administrator to restore previously migrated files to their original location for faster data access

#### **Pricing and Go-to-Market**

The Enigma Data Solutions software is priced by the amount of storage it has migrated. The original entry-level capability was 10 Terabytes for \$29,000. At customer request, licenses now start at 0.5 TB and come stripped of the additional functionality inherent in the 10 TB license, for \$6,000. License keys to upgraded with Reporting and Unmigrate Modules can be received by email, and deployed, using an overwrite/refresh, by the administrator.

#### **Other Uses for Smart Move**

In addition, there are some creative uses for Smart Mover.

- It could be used to migrate selective data components, from old devices (servers or laptops) to new devices.
- It could be used to replicate files locally to branch offices taking on new organizational roles.
- It could be used to locate all the files relevant to a project, product, customer, or partner. Here the use stops short of migration and uses the characterization at the project level to find the business equivalent of “all the sports coats in your closet.”

The key is being able to filter the data to suit the occasion, which is seldom a matter of volumes, disks, or LUNs. Once you start managing data by its business value - and add the ability to sift out the data you want or don’t want from the great mass of what you have - the technology ramifications of enterprise reconfiguration and events become easier to manage.

#### **Conclusion**

SmartMove is a useful set of tools and capabilities. If you can your characterize enterprise files, and rationalize them, it is not just your storage environment that will benefit. Enterprise processes will run smoother. Enterprise search will be a rational process, gleaning a higher percentage of useful results. The data ramifications of the prospect of enterprise change will be less daunting - and all this is achieved by an unassuming utility application that runs un-noticed in the background.

Think about what SmartMove could do for your enterprise. It may be just what you need.



### ***About The Clipper Group, Inc.***

***The Clipper Group, Inc.***, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

***Anne MacFarland is Director of Infrastructure Architectures and Solutions for The Clipper Group.*** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at [Anne.MacFarland@clipper.com](mailto:Anne.MacFarland@clipper.com) or at 781-235-0085 Ext. 128. (Please dial “128” when you hear the automated attendant.)***

### ***Regarding Trademarks and Service Marks***

**The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager,** and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosure***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.