
Data Domain tells a compelling “tape-less” story

Date: May 4, 2006

Author: Heidi Biggar

Title: Analyst

Abstract: Having shipped more than 1,000 units to more than 300 customers, Data Domain has become an “emerging vendor” success story. The company’s success hinges on its tape-less backup story, which now includes a key vaulting component for disaster recovery (DR) and remote data protection.

The Second Chapter

Anyone can start a business but generally only those who know how to tell a compelling story – and, more importantly, how to execute a technology and business road map – actually make it in the end.

Data Domain is an “emerging” success story. While many start-ups in the industry (and even some that are competitive to Data Domain) struggle with footing and messaging, Data Domain has moved steadily and decisively forward with a clear marketing focus and product portfolio around tape-less backup – a message, or story, that resonates well among end-users today.

Building on this strength, Data Domain this week introduced a new generation of its DD Replicator software – an option that further extends its tape-less story outside the data center. DD Replicator provides off-site vaulting for disaster recovery and remote office data protection, and allows for multi-site tape consolidation.

ESG believes this type of support is very important, especially as end-users begin to address data protection issues with the large – and growing – volumes of data at remote and branch offices (ROBOs). ESG estimates that up to 30% of an organization’s mission-critical data resides outside the data center at ROBOs.

Data Domain fully understands that while primary storage represents a significant data storage challenge for organizations, the larger issue lies with protected, or backup, data. Depending on how the backup process is managed (i.e., is data being backed up to tape or to disk, are full backups being done daily or weekly, etc?), how many copies of the backup data are being made, and whether the data is being de-duplicated (which is very unlikely at this point), backup data volumes are generally many times larger than primary data volumes. Backup volumes are creating a huge capacity and management nightmare for organizations.

The Replication Factor

Admittedly, Data Domain is just one of many vendors (mainly start-ups) over the past couple of years to announce data protection products aimed at improving backup efficiency – and, more recently, ROBO data protection.

Disk-based backup (e.g., Asigra, Avamar, and Signiant), replication (e.g., Double-Take, XOsoft, and Kashya), VTL (e.g., Diligent and Sepaton), CDP (e.g., Revivio), WAFS (e.g., Riverbed and Tacit) and traditional backup vendors (e.g., IBM and Symantec) have all thrown their names into the ROBO data protection ring – and more still are promoting the concept of more efficient (tape-less) backup. However, Data Domain tells a very compelling story for the following (combined) reasons:

- **It is “plug and play” and supports leading backup applications:** This protects end-user investments in traditional backup infrastructure and makes it easy for end-users to phase Data Domain (as a disk device or a virtual tape library) into their existing environments.
- **It includes a data de-duplication component (data is de-duplicated before it is replicated):** Data de-duplication is quickly becoming a household name in IT environments. That said, ESG knows of no other replication product at this time that physically de-duplicates and compresses backup data before it is moved over the WAN. Data Domain claims an average data reduction rate of 20x, compared to traditional backup procedures with weekly full backups and daily incremental backups, depending on the type of data being replicated. ESG Lab has validated the de-duplication feature of the Data Domain product set. What we found most compelling is the thoroughness of Data Domain’s de-duplication process. Data is de-duped two times: 1) remotely at the ROBO and 2) at the central site (a second time) before the ROBO data has been pushed over the WAN. This type of process prevents duplicates locally and at remote sites and minimizes the bandwidth needed for remote vaulting.
- **It is flexible (it allows organizations to set protection levels for different classes of backup data):** DD Replicator allows end-users to prioritize protection for different classes of data. It is not “all or nothing.” End-users, for example, can choose to replicate certain data sets (e.g., Oracle or Exchange) and not others (e.g., Directory A or Directory B). Additionally, end-users can perform bi-directional replication to the same node or many-to-one replication for ROBO protection. Again, data is de-duplicated (across all in-bound data streams) before it is replicated to the ROBO or Data Center Hub. The Data Center Hub checks for non-unique data streams coming from remote locations and only sends unique data streams across the WAN to the main data repository, or hub, for efficient bandwidth utilization.
- **It removes tape from the ROBO/DR equation:** ROBO data is “backed up” via the WAN (using Replicator) to the central Data Domain hub. No tape is necessary, unless the user wants to maintain tape copies in the event of a worst-case doomsday scenario. All active recoveries should occur from disk, either locally, or via the remote store. Doing so can dramatically improve RTOs (recovery time objectives). If tape copies are needed, they can be made at the hub, which also consolidates and shrinks the tape infrastructure.

DD Replicator supports Data Domain’s full line of disk appliances, including the DD410, DD430, and DD460, as well as the DD460g, or gateway device, to which non-Data Domain disk systems (e.g., EMC or Nexsan) can be attached. Pricing for the software starts at \$2,600 (for the DD410).

Bottom Line

By enhancing DD Replicator, Data Domain is taking the right next steps for its product line. ESG is already seeing growing interest in tools, such as DD Replicator, for off-site DR vaulting and in-bound ROBO data protection. In fact, Data Domain is reporting a 40% attach rate for its DD Replicator product among its current customer base (the company has been shipping DD Replicator for system-to-system replication for 1.5 years).

One of the key drivers of Data Domain’s DD Replicator (and other ROBO-focused technologies) going forward will be data de-duplication. Simply put: data de-duplication makes it economically and physically possible to store huge volumes of data locally and to move data across the WAN efficiently and cost-effectively. Data Domain is one of the pioneers of data de-duplication technology and should have somewhat of a competitive advantage due to its comparatively large installed base (1,000+ shipments/ 300+ customers) and the ease of integration of its product line into the existing backup infrastructure (in particular, its backup application support).

From a pure technology standpoint, Data Domain also has a competitive “leg up” on more established players, such as Symantec and, in particular, IBM, which recently entered this market space with their Pure Disk and TSM Express products, respectively. However, they lack the marketing muscle and the channels of these larger companies. CommVault, Computer Associates, and EMC Legato, to date, have not introduced products for this particular market space.

It remains to be seen, however, if a small, but rapidly rising, vendor such as Data Domain can keep up this pace. But for now, things appear favorable for the company. The market needs the type of solution Data Domain offers, its technology works, and the big guys have yet to provide legitimate alternative solutions to the problem. Data Domain needs to continue to execute and accelerate customer adoption. Doing so will make the company more attractive to the market, as well as to would-be acquirers.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. and is intended only for use by Subscribers or by persons who have purchased it directly from ESG. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.