








Special to SearchStorage.com

Data Domain DD400

OVERALL RATINGS

Installation		Documentation does a good job of stepping the user through installation
Manageability/Configuration		The new GUI has dramatically improved manageability
Operation/Performance		Unit performed well in tests, but lacks disk load balancing
Recoverability		Creates a continuous full image
Problem Resolution and Support		Data Domain's tech support was responsive and knowledgeable.
Interoperability		Compression is not compatible with encryption
User Level		
List Price as Tested	\$45,000	

Scale: 1 flame = poor; 2 flames = fair; 3 flames = satisfactory; 4 flames = very good; 5 flames = excellent

User level: ● = novice; ■ = intermediate; ◆ = advanced; ◆◆ = experts only

PRODUCT SUMMARY

Given that Data Domain is a relative newcomer to the storage industry, its products have matured amazingly quickly. Early generations of the company's product were managed by command line only, and it was not until mid-2005 that a GUI was introduced. The GUI was the missing link, however, and the current DD400 generation of product represents a remarkable blend of innovative technology and simple use. The DD 400 family, which includes the DD410, DD430 and DD460 appliance models and the DD460g gateway model, support local backup to disk over file system protocols on Ethernet and simultaneously as a virtual tape library on Fibre Channel.

Data Domain's primary differentiating feature is its "compression." The company refers to its solutions as "Capacity Optimized Storage" (COS). Generically, we would refer to it as single instance storage. But whatever the name, COS can offer significant reductions of used capacity when compared with standard backup, data replication and other VTL technologies.

The confusing element of COS is understanding the phrase "capacity optimization." The company claims a 20:1 compression effect, however storage administrators should not assume that they can backup 20 TB of raw data onto a 1 TB DD unit. Rather, the optimized capacity is achieved over time by avoiding the copying of previously backed up data. When implementing a new DD system, Global Compression or non-redundant storage is used along with standard LZ or optimal GZ compression when transferring data to it. The resulting compression effect for the initial copy of backup data will be between 2:1 and 4:1 in most cases. Subsequent incremental and full backups provide greater capacity optimization due to non-redundant storage. Thus, 20 TB of raw data will be initially compressed to approximately 5 - 10 TB.

When sizing a system, the size of backup data set, backup policy and data retention timeframe are critical parameters. Assuming a typical backup policy and several months of data retention, Data Domain recommends a 1:1 ratio of DD capacity to raw data, but in some cases it may be less. This compares to physical tape libraries where we recommend a 7:1 sizing ratio to accommodate a grandfather/father/son backup strategy. VTLs are typically sized using a 3:1 ratio.

About the Review

Diogenes Analytical Laboratories and SearchStorage.com teamed together to create disk-to-disk buyer's guidelines that help IT buyers differentiate between products and select the best candidates for their own situation. For this report, we took 13 different products into the lab and ran them through their paces. We compared products based on ease-of-use, manageability, features/functions and technical support systems. Our opinions are based on more than a decade of implementing and managing data protection systems.

After the initial data transfer, or backup, COS will continue to look for previously stored character strings and replace them with reference pointers for all types of backup jobs defined in the backup policy. Thus, only new data will be stored. This methodology is very similar to an "incrementals forever" strategy. Therefore, DD systems retain a full backup image after each backup operation even though only new, original data has been added. The extent to which disk space is consumed is governed by two factors: rate of data change plus the rate of original data growth. We qualify this statement with "original," because growth of repetitive data will not result in significant consumption on the DD system. For example, new orders for products in an order entry system would proportionally increase primary disk space, but

would have minimal impact on the DD system because all products, item numbers and descriptions would have already been stored and can be referenced by pointers.

The greatest advantage to COS is its ability to reduce the “multiplier” effect of data growth. The data multiplier effect is characterized by the number of times that original data is replicated. For example, data is often replicated with RAID, in snapshots, mirrors, backup, and so on. Our research indicates that this multiplier effect can range from 7:1 to as much as 15:1, often driven by the large number of backups, such as weekly full, monthly full and annual full. Of course, incremental backups can also contribute to the multiplier effect as well.

The second differentiating feature and advantage of COS is found when replicating or vaulting between systems. Data Domain supports asynchronous replication between units, which can be implemented over a wide area network or local area network. Because replication sends only unique compressed data segments across the network, the amount of data replicated is generally very small. Thus, it is feasible in many situations to replicate backup data between sites using less bandwidth than many other systems.

The main disadvantage to COS is that its benefits are lost when it receives encrypted or compressed data stream. Compression engines (including COS) look for duplicate strings of characters and when found, are referenced by pointer rather than stored again. Encrypted data appears to compression engines as random data, and therefore no matching strings are found. This limitation only applies if the data is encrypted on the source disk or is encrypted over the network to the DD unit. For IT organizations that use encryption on physical tapes or do not encrypt data within the data center, the issue will not be a problem.

DIFFERENTIATING FEATURES AND FUNCTIONS

FEATURE/FUNCTION	DESCRIPTION
<p><i>Compression</i></p>	<p><i>Although the DD400 products use LZ or GZ compression, this actual compression accounts for only a fraction of the amount of compression that the product actually provides. Instead, the product uses a pattern-matching algorithm to avoid storing the same data more than once. After the first baseline full backup, each subsequent backup stores only unique compressed data segments and related data pointers.</i></p>
<p><i>Continuous full image and multiple backup versions available for rapid restore</i></p>	<p><i>DD systems have a inherent snapshot-like capability and retain each backup image after each backup operation even though only new, original data has been added. Thus, restoring either a full system image or a individual file is a simple matter.</i></p>
<p><i>Replication</i></p>	<p><i>The DD systems include optional software replication between units. This replication can be either local for data redundancy or remote for disaster recovery; both methods are asynchronous and work in WAN and LAN environments. The replication technology is optimized in the same way as the disk storage in that only completely new and unique compressed data patterns are replicated between DD systems.</i></p>

BUYER GUIDANCE

Kudos:

The DD400 series compression, or non-redundant storage, can save a significant amount of disk space (and therefore cost). Whereas a VTL would typically be sized by at least 3:1 VTL space versus primary disk space, each model in the DD400 series can safely be sized at 1:1. Because the rules of non-redundant storage apply to replication between units, it is feasible to keep a remote site or disaster recovery facility up-to-date with lower bandwidth connections than other solutions. The DD400 also maintain , so all recoveries can be from the DD unit rather than a combination of disk and tape. The conceptual aspects of the DD400 are more challenging than the reality of the system; most IT professionals will have no problem adjusting.

Caveat:

The benefits of COS will be lost when sending encrypted data to the DD400 systems. (Encrypted data appears random to a compression engine and cannot be compressed.) This may be an issue for organizations that encrypt data on disk in the data center, such as classified information or highly sensitive regulatory data. Setting up the DD400 takes a little more work than a simple VTL, but not much.

Who should consider this product:

DD400 systems will appeal to organizations looking for leading edge technology and optimized storage capacity and are willing to invest a little time in upfront planning.

ABOUT DIOGENES ANALYTICAL LABORATORIES, INC

Diogenes Analytical Laboratories, Inc. is an independent organization dedicated to helping Information Technology buyers reduce the inherent risk and uncertainty associated with technology purchases. Our goal is to create an informed I.T. consumer and provide the complete information needed to make smart purchase decisions. This report is based on Diogenes Analytical Laboratories' actual lab testing experiences and was not funded, sponsored or commissioned by any vendor. The opinions expressed in this report are those of Diogenes Analytical laboratories, Inc.

Diogenes offers a full range of services for IT buyers. These include:

- In-depth product evaluations and comparisons
- Comparative Buyer's Guides
- RFP generation and review
- Head-to-head in-lab "bake-offs" for product finalists
- ROI and TCO analysis
- Advisory services

For more information, go to www.diogeneslab.com or call 303-402-9900.

This report was not funded, sponsored or commissioned by any vendor. Diogenes is an independent testing lab that does not endorse products. Nothing in this report should be construed as an endorsement of any product or technology for any purpose.

© 2004-2006 Diogenes Analytical Laboratories, Inc. All rights reserved. Neither this report nor any portion of it may be reproduced in any form without the express written consent of Diogenes Analytical Laboratories, Inc. The use of this report is governed by a separate Product and Services Use Agreement which may be viewed at www.diogeneslab.com.

Trademark notices:

The Diogenes logo, Spotlight symbol and flame are trademarks of Diogenes Analytical Laboratories, Inc. SearchStorage.com is a trademark of TechTarget, Inc. All other trademarks are the property of their respective owners.