

## Data Domain Introduces Scalable DDX Array; Goes After Enterprise Users

**Date:** November, 2006

**Author:** Heidi Biggar, Analyst

**Abstract:** Data Domain is looking to “replicate” its success in the mid-market in the data center with the introduction of its modular DDX Array. The array can be powered by existing DD460 controllers or, for a significant boost in capacity and throughput, new DD560 controllers.

### Data Domain Goes After the Enterprise

Having made a name for itself in small and mid-size companies - having shipped more than 1,500 systems to nearly 600 customers - Data Domain has now set its sights on the data center. Data Domain believes - and ESG agrees - that its new modular DDX Array equipped with higher-performing/higher-capacity DD560/DD560g controllers is the right vehicle to get them there.

The new DDX Array has all the bells and whistles that Data Domain products have become known for, including data deduplication (a.k.a. Global Compression), WAN vaulting (one-to-one, many-to-one, and bi-directional replication), and RAID-6 data protection.<sup>1</sup> Also, like the existing DD400 Series, the new DDX array works with

**FIGURE ONE**  
**Data Domain Keeps on Scaling**

	DD460	DD560	DDX
<b>Date Introduced</b>	May 2005	Oct 2006	October 2006
<b>Performance</b>	290 GB/hour	400 GB/hour	6.4 TB/hour
<b>Capacity</b>	80-220 TB	400-980TB	7-15PB

all leading backup applications and can be presented to these applications as either a disk or VTL target.

However, from an architectural standpoint, the DDX is different from Data Domain’s DD400 Series in one very important way: Unlike the DD400 Series, the DDX Array scales - and is managed - as an integrated unit, not as a collection of disparate DD400 devices.

Preliminary data from an ESG Snapshot Survey in progress shows that scalability continues to be a key factor in many organizations’ disk-based backup purchasing decisions, from both an initial investment as well as a feature enhancement standpoint. In fact, when asked what types of improvements and/or features they would like to see in their existing VTL solutions, survey respondents (thus far) reported “better/easier management” and “better scalability” as top concerns. Also important were improved performance, and new or better disaster-recovery features.

With the introduction of the DDX Array and DD560 controller, Data Domain can “check off” enterprise scalability and management.

### The DDX Impact

As mentioned, the power of the DDX Array is really two-fold: 1) it allows Data Domain users, for the first time, to scale capacity and performance (by adding controllers) within a single self-contained storage system, addressing ever-increasing end-user performance and capacity requirements, and 2) it is easy to manage.

<sup>1</sup> ESG Lab Validation Report: *Data Domain WAN Vaulting*, August, 2006.

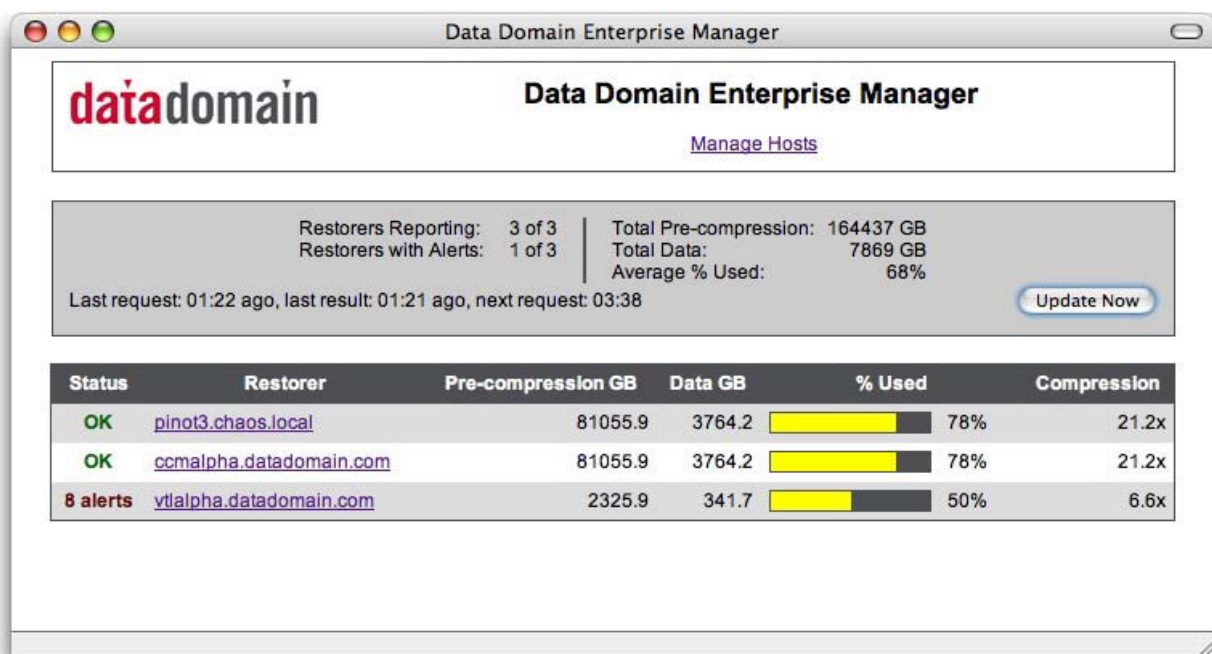
- Scalability:** Each DDX Array can be configured with up to 16 DD460 or DD560/DD560g controllers. The DD560 “g” or gateway supports external (third-party) storage. The company is expected to make shelves available for the DD560 controller (i.e., add its own capacity to the DD560) in Q1 2007. The DDX Array is available in 4-, 8- and 16-controller configurations. Users can add controllers and associated capacity as needed for up to 15PB of “de-duplicated” capacity and up to 6.4TB per hour (using DD560 controllers). At the component level, the new DD560 controller (with internal capacity) scales from 400TB to 960TB and has a data throughput of up to 400GB per hour. The DD460, in turn, scales from 80TB to 220TB and has a maximum 290GB-per-hour throughput (see Figure One). Data Domain’s Global Compression, or data de-duplicating capability, reduces capacity requirements, which help keep costs down from a \$/GB standpoint.

The DDX’s operating system (DD OS 4.0) gives each controller its own namespace, which can be used by one backup application/media manager or shared among multiple backup applications/media managers by targeting different subdirectories (in the case of NAS) and different virtual libraries (in the case of VTL). Of course, the number of applications/media managers supported is limited to the throughput and available capacity of each namespace and, ultimately, the sizes of the data groups for each namespace are limited by the speed of the media manager. For this reason, while each namespace can use up to 16TB of raw space, they are generally much smaller. Users establish backup policies for each namespace, which are adjusted as controllers are added to ensure that the backup policy is targeted to the right namespace. ESG Lab will validate this claim, as well as the scalability of the DDX Array, in a follow-up Lab.

- Management:** The array is managed via Data Domain’s Enterprise Manager GUI. This tool gives high-level information about the array, including capacity, compression and general component “health” information, as well as component-specific information about each individual controller.

ESG Lab recently used the Enterprise Manager to monitor multiple DD400 appliances during the replication process.<sup>2</sup> (The Lab was conducted prior to the DDX announcement. ESG Lab is scheduled to perform DDX testing in the coming months.) The Manager gave us a view of the amount of data stored on each device, compression ratios and the physical capacity consumed by that data (see Figure Two).

Figure Two: Data Domain Enterprise Manager



<sup>2</sup> ESG Lab Report: *Data Domain: The DD400 Series: Capacity and Optimized Bandwidth Protection*, June, 2005.

There are two main applications for the DDX: It can be used as a scalable high-performance disk-to-disk (D2D) backup appliance (as either a disk target or a VTL) or as the central site target in a remote backup configuration for consolidation and disaster recovery purposes. In this situation, a company would replicate data using Data Domain Replicator (after it has been de-duplicated via Data Domain Global Compression) from remote sites to the central site. Data Domain multi-site remote replication and best-in-class data de-duplication make this a compelling and valuable way of protecting remote office data.

### The Bottom Line

With the introduction of the DDX Array -- and companion DD560 controller -- Data Domain has made its intentions to go after the enterprise loud and clear. Whether or not the company will be able to replicate its early success in the mid-market at the enterprise level remains to be seen. Without an OEM partner to help open doors in the enterprise, the road before Data Domain is challenging at best, impenetrable at worst. But ESG believes Data Domain has compelling solutions whose economics are easily quantified. The company addresses end-user requirements and technology issues head on, and it stays with or ahead of the pack in terms of feature roll-out (e.g., data de-duplication). (That said, its GUI could use a little work.)

Certainly, there is opportunity for Data Domain to live in the Enterprise data center supporting local backup to disk requirements, and the company has a great story for remote office backup with its multi-site replication support and data de-duplication technology. In fact, the remote backup use is justification alone (The DDX would act as an aggregator for a large number of remote offices backing up to a central site). The VTL guys don't really have an answer for capacity-optimized remote office backups yet - especially where multiple sites are concerned.

So, while we understand why Data Domain is going after larger enterprise environments (the big dollars are in the larger configurations), we'd be remiss in not pointing out the potential consequences of doing so. By going after the enterprise, Data Domain runs the risk of diluting its core mid-market focus, and missing opportunity there. Nonetheless, Data Domain has a window of opportunity here, which could get them to a whole new level. Perhaps there is even an IPO on the "remote" horizon.