
Data Domain plans datacenter assault with DDX

Henry Baltazar, Storage Analyst
Simon Robinson, Sector Head
Storage & Systems
IMPACT REPORT
10 October 2006

Data Domain, a self-styled capacity-optimized storage specialist, has managed to grow its customer base and revenue at an impressive rate over the past year or so. Having cemented a solid position in the midmarket, the company now has its eyes on expanding its reach into the datacenters of much larger organizations.

With the launch of its DDX data protection array, the company now feels it can grow beyond its remote-office and department-class market segment and have a bigger impact in the datacenter. To sustain long-term growth, the company plans on extending its technology beyond the data protection market, possibly into the primary and secondary storage markets or archive space.

Impact assessment

The message

Data Domain has expanded its focus up into the datacenter with the release of the DDX array.

Competitive landscape

Data protection startups Diligent, ExaGrid, Sepaton and Avamar – as well as Quantum – all have de-duplication technologies. But NetApp and EMC remain Data Domain's biggest rivals.

The 451 assessment

Data Domain has been growing its revenue and customer count at a strong rate, and looks to be ramping up for an IPO – if a larger player doesn't acquire it first. The release of DDX is a big step, giving it end-to-end data protection covering datacenters and remote offices. The company is taking advantage of customer dissatisfaction with tape-based data protection offerings, but it will eventually need to expand beyond the data protection market into the primary and secondary storage market if it truly hopes to emulate NetApp. We think it remains a highly attractive – if expensive – acquisition target.

Context | Founded in 2001, Data Domain has posted 10 consecutive quarters of growth, a feat not matched since the emergence of **Network Appliance**, it claims. Data Domain has more than 600 customers, and has shipped more than 1,400 systems. In the first six months of 2006, its revenue grew 420% year-on-year. Currently, 35% of its revenue is from outside the US, with particular strength in Japan, and a growing European presence. Company headcount is about 150, with 130 employees split evenly between the engineering and sales teams. With an eye toward future growth, Data Domain recently moved its headquarters to a 46,000-square-foot facility in Santa Clara, California. Data Domain hopes to double revenue next year.

Products | Data Domain's new DDX array clusters up to 16 Data Domain controllers, greatly expanding its existing top-line scalability and performance. Utilizing the third generation of Data Domain's core file system technology, the DDX can back up 6.4TB of data an hour, and can scale to 16PB of capacity. The DDX is available in four-, eight- and 16-controller configurations. Data Domain's top-of-the-line appliance before DDX was the DD460, which could scale up to 220TB and back up 290GB an hour.

The company has also launched the DD560 appliance, featuring 500GB drives (the 400GB drives used in the DD460 have a maximum storage capacity of 1PB and can back up 400GB per hour. It also released a gateway version, the DD560g. This third generation of products is the first to use a 64-bit kernel and code, which should improve its memory utilization and processing efficiency. The new DD560 appliance with 15 500GB drives costs \$95,000, while the DD560g gateway is priced at \$105,000. An entry-class DDX 4 controller array with DD560g controllers is priced at \$420,000.

Like the other products in Data Domain's lineup, the DDX is a backup target device that presents itself as a file share – or as a VTL (virtual tape library) on a fiber channel storage area network – to backup applications. Currently, only 10-15% of Data Domain's customers are using the VTL option, which means that most customer backup applications are writing data to Data Domain appliances and gateways using standard CIFS and NFS file-sharing protocols over Gigabit Ethernet IP networks.

Data Domain is in no rush to push customers to the VTL platform; indeed, it says many prospects initially interested in the VTL option often choose to implement its products using network-attached storage (NAS) since it's simpler. Data Domain products work with most backup software without requiring any changes to be made to existing backup policies, and are typically used to hold a couple of months' worth of backup sets 'online' for quick data recovery. It is not replacing tape-based products for archiving and long-term retention, however.

Before the launch of DDX, Data Domain appliances were typically used to protect departmental class applications and remote offices. Data Domain appliances can efficiently replicate their data over a WAN, and this ability allows customers to aggregate backup data at remote sites to Data Domain units sitting at a central datacenter. With the launch of the larger DDX arrays, Data Domain now has a more legitimate place in the datacenter, consolidating large amounts of data and enhancing backup services for multiple departments. Moreover, Data Domain says the DDX simply reflects the reality of customer demand, which is moving to ever-larger systems.

Data Domain's management interface used to be limited to managing a single appliance at a time, a suboptimal condition for staff working in overcrowded NOCs (network operation centers). With a new version of the Data Domain Enterprise Manager, customers can now manage multiple appliances from a single console.

Technology | Data Domain's Global Compression de-duplication technology is the company's core asset and allows customers to store vast amounts of data on its appliances. As data is fed into a Data Domain product, it's broken down into, on average, 8KB segments (range from 4-12KB). Each segment is assigned a unique identifier, and a database of these segments is stored within Data Domain's hardware. When a backup application attempts to write a non-unique segment to a Data Domain box, the product acknowledges the segment but does not physically write it to the data set – thus improving the storage efficiency of Data Domain's storage.

The Global Compression process is highly CPU-intensive, which is why it is well suited for backup environments where access to the Data Domain file shares and virtual tapes is limited to a handful of servers. With multicore processors greatly expanding the processing power available for appliances, it is highly possible that Data Domain could launch a family of storage-efficient primary and secondary storage devices in the future. Data Domain will be launching appliances with dual-core processors in 2007.

Since Data Domain devices already present themselves as file shares, it is actually possible for a customer to use current appliances as a NAS device or file server replacement. Because of the CPU overhead associated with Global Compression, a Data Domain product used in this manner would likely perform poorly, especially in transaction-heavy environments. A more likely destination for Data Domain's Global Compression technology is the archive market, where it would be able to provide faster data access relative to current tape-based offerings. A WORM file system – which is a requirement for disk-based archiving – is currently being developed by Data Domain and expected to become available in 2007.

Competition | Data Domain was one of the first vendors to enter the market with a backup appliance using data de-duplication, and it was able to quickly gain market traction. Compared with **Avamar Technologies**, another early market participant, Data Domain has an advantage since it works with current backup software packages and infrastructure, while Avamar forces customers to adopt its own backup software. Data Domain also claims customers prefer its appliance-based approach to software-based alternatives.

Diligent Technologies, ExaGrid Systems and Sepaton are other startups vying with Data Domain. **Hewlett-Packard** is an OEM of Sepaton's, which gives that startup access to client accounts that are typically less receptive to a startup. Sepaton's technology only functions as a VTL at this point, which gives Data Domain an advantage, since it supports NAS and VTL connectivity options. Diligent's technology is also now being offered by **Hitachi Data Systems**. Quantum, armed with de-duplication technology from **ADIC's** acquisition of **Rocksoft**, will release a remote-office backup appliance and an enterprise-class VTL, which will directly compete with Data Domain's products. Besides its highly granular de-duplication capabilities – no longer a unique technology – Data Domain claims its replication and data verification capabilities offer significant differentiation.

Data Domain competes most regularly with NetApp and **EMC**, both of which lack data de-duplication to the same level of granularity as Data Domain. Indeed, Data Domain's much greater efficiency at storing backup data online has allowed it to defeat NetApp staples such as NearStore, as well as EMC's Clarrion Disk Library VTL offering, in multiple accounts. Since Data Domain is now starting to win business worth in excess of \$1m, we feel it is only a matter of time before one – or both – respond by acquiring Data Domain or one of its closest rivals.

SWOT analysis

Strengths	Weaknesses
Data Domain has exhibited strong, consistent growth since its inception and has a solid family of data protection products.	The company is still relatively small compared to more established storage companies, and has many competitors with similar products.
Opportunities	Threats
The data protection products are in high demand now. The archive market, which Data Domain could move into, also has potential.	A number of startups and larger companies are targeting the same data protection space.

About The 451 Group

The 451 Group is a technology industry analyst company focused on the business of enterprise IT innovation. The company's analysts provide critical and timely emerging-technology insight to clients at vendor, investor, services and end-user organizations – insight that aids both strategic and tactical decision making for competitive advantage.

The company's services include the 451 Market Insight Service, which delivers daily insight into emerging enterprise IT markets; 451 TechDealmaker, a weekly analysis service focused on forward-looking M&A within the enterprise IT business; 451 Special Reports, which are produced on a periodic basis to analyze key emerging enterprise IT markets in greater depth; and 451 Strategic Counsel, the company's analyst-inquiry program, which provides clients with direct access to 451 analysts. The company also produces via 451 Events periodic industry summits and investor conferences that address opportunities and obstacles facing emerging enterprise IT markets.

The 451 Group is headquartered in New York, with offices in key locations, including San Francisco, London and Boston. The company also operates Tier 1 Research – an independent division of The 451 Group, headquartered in Minneapolis – which analyzes the financial and industry implications of developments impacting public and private companies within the IT, communications and Internet sectors.

For additional information on the company or to apply for trial access to its services, go to: www.the451group.com