

ESG Lab Review

Enterprise-class Backup Performance with Dell DR6000

Date: May 2014 **Author:** Kerry Dolan, Lab Analyst and Vinny Choinski, Senior Lab Analyst


Abstract: This ESG Lab review documents hands-on testing and audits of Dell tests of the DR6000 purpose-built backup appliance. Testing was focused on the throughput performance and accelerated ingest.

The Challenges

A number of realities challenge today's organizations both operationally and financially. Ongoing growth of data volumes drives storage costs up; this is exacerbated by stringent compliance and governance regulations that mandate long-term retention and accessibility. These growing data volumes also make it difficult to complete data protection activities in the time allotted because larger data volumes make the backup process take longer. For many organizations, backups cannot be completed within the backup window and ultimately interfere with production operations.

These trends are borne out in the results of ESG's *2014 IT Spending Intentions Survey*. When asked to identify their organizations' most important IT priorities over the next 12 months, 29% of respondents cited improving data backup and recovery, placing it third behind such crucial initiatives as increasing server virtualization and information security.¹ This ranking indicates that IT organizations are serious about better data protection. In fact, improving backup and recovery has been in the top three most-cited priorities for the past three years. It should also be noted that managing data growth has retained its place, cited among the top four most-cited IT priorities for the past four years.

Figure 1. Top Ten Most Important IT Priorities for 2014



Source: Enterprise Strategy Group, 2014.

¹ Source: ESG Research Report, [2014 IT Spending Intentions Survey](#), February 2014.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by Dell.

The Solution: Dell DR6000

The Dell DR6000 is the latest addition to the DR family of purpose-built data protection and backup appliances. The DR6000 was designed for large enterprises, and offers greater capacity, scalability, and performance than the mid-range-focused DR4100. Key new features are Rapid NFS/Rapid CIFS, the industry's first source-side deduplication for accelerating NFS and CIFS backups, and the Global View interface that enables consolidated management of multiple DR appliances from a single, web-based management console. Dell's all-inclusive licensing includes all features at no additional cost, including future updates when support contracts are maintained.

Like the DR4100, the DR6000 is a 2U backup appliance that includes inline deduplication and compression that reduce storage capacity requirements for protected data. The DR6000 can scale up to 180TB after RAID6 using up to four expansion shelves, offering up to 2.7PB of logical capacity. Many-to-one bi-directional replication is included for disaster recovery. The DR6000 is easy to install and manage, and integrates easily with existing data protection workflows. It supports leading backup applications including Dell NetVault Backup, vRanger, and AppAssure, as well as applications from CommVault, CA, Symantec, IBM, EMC, Oracle, HP, Veeam, Bridgehead, and Amanda.

Target- or Source-side Deduplication

While target-based deduplication is standard in the industry, the DR6000 can accelerate backups using source-side deduplication by leveraging several plug-ins that enable this feature for most backup applications. Dell NetVault Backup software uses the RDA client plug-in. If using Symantec backup applications, the OST protocol is supported. New with the DR6000 is the addition of the Rapid NFS and Rapid CIFS plug-in for source-side dedupe for other backup applications offering a maximum throughput of 22TB/hour. All source-side deduplication chunking and hash computations can be done either at the client server or media server.

Configuration options include:

- Post-RAID capacity points: 9TB (135TB logical); 18TB (270 TB logical); 27TB (405TB logical); 36TB (540TB logical)
- Up to four expansion shelves in 9TB, 18TB, 27TB, or 36TB (after Raid6) units
- Networking: 1GbE, 10GbE, base-T, or SPF+

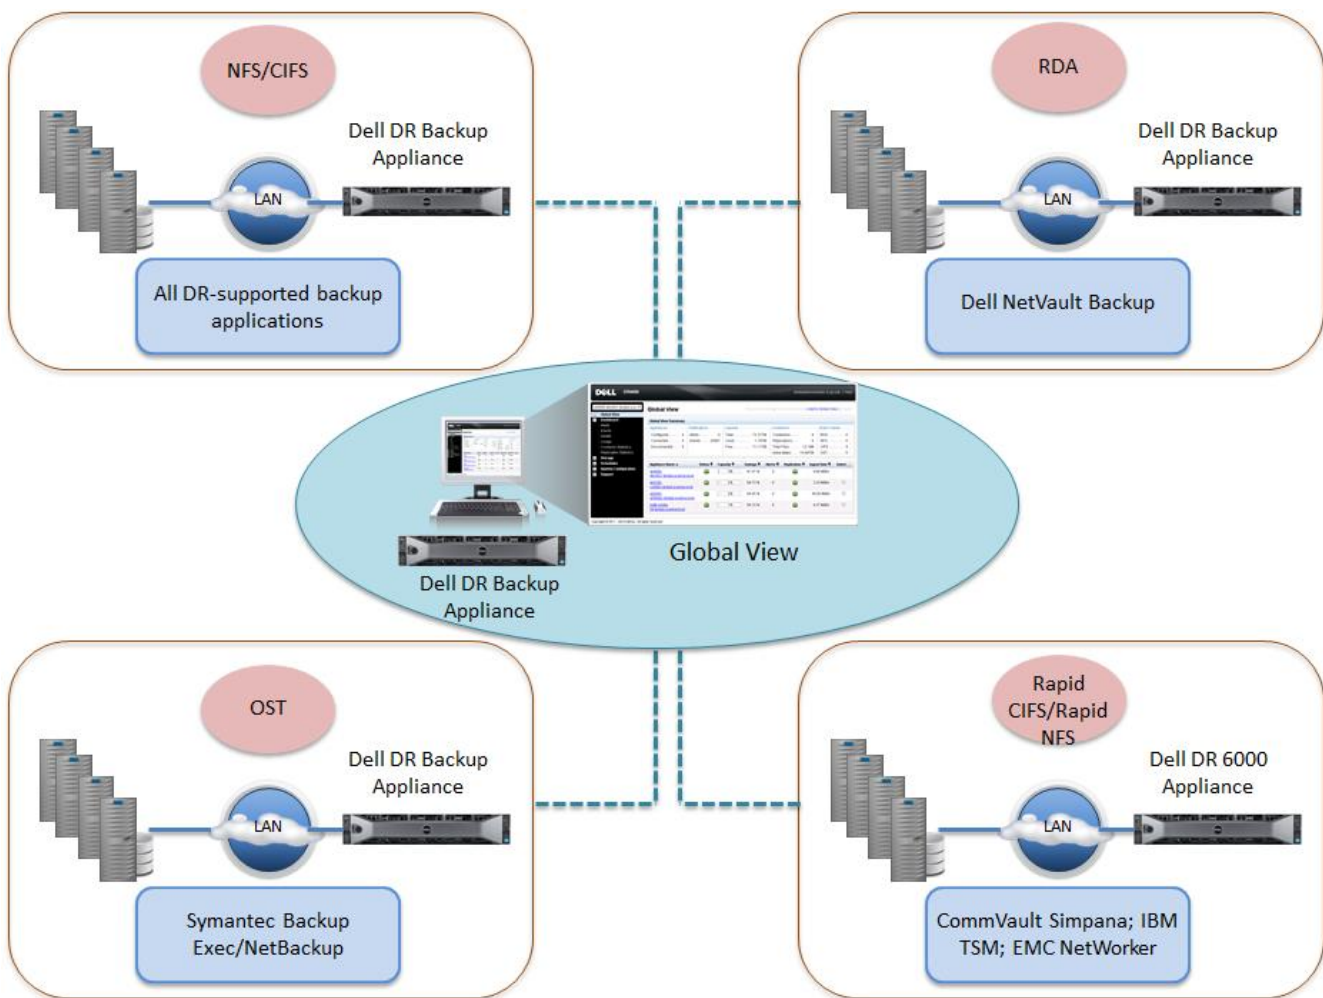
Global View

Global View is the new federated management capability that is included with DR appliances using version 3.0 of the operating system. Designed for multi-DR environments, Global View enables IT administrators to monitor, manage, and perform workflow tasks on up to 64 DR appliances (DR6000, DR4100, and DR4000) from a single GUI instance, as long as they reside in the same Active Directory domain. Global View enables IT to view alerts, events, container usage, and replication statistics.

Solution Overview

Figure 2 provides a conceptual view displaying multiple DR implementations using multiple backup protocols, all centrally managed using Global View. Four remote locations are displayed, each including servers that are LAN-connected to a DR appliance, and each replicating to a central location for disaster recovery. The upper left shows a site executing NFS/CIFS backups that can leverage all backup applications supported by the DR series. The upper right shows the RDA plug-in that Dell NetVault leverages to speed backups. At the lower left is an implementation using the OST plug-in to accelerate backups using Symantec applications. Finally, the lower right shows a DR6000 implementation that leverages Rapid CIFS/Rapid NFS for faster backups using Simpana, TSM, or NetWorker. In the center, the Global View management interface can monitor and manage all locations from a single console.

Figure 2. Dell DR6000 and Global View



ESG Lab Tested

ESG Lab performed remote testing and audited Dell tests performed in Hyderabad, India. Testing focused on accelerating throughput performance for a range of backup applications using the Rapid NFS/Rapid CIFS, OST, and RDA plug-ins. Peak performance was also tested.

Global View

ESG Lab first explored the Global View console. With Global View, administrators can monitor and manage up to 64 DR6000, DR4100, or DR4000 appliances. Figure 3 shows a Global View instance running on a DR4000 and tracking data on four appliances. The top half of the screen provides a summary of all the appliances being monitored, showing the number of appliances configured and connected, the events and alerts, capacity information, and details regarding containers, replications, total files, and active bytes. In addition, active clients are listed by the protocol being used for backup. The bottom half shows details for each appliance being monitored, including status, capacity, savings, alerts, replication status, and ingest rate.

The blue link at the top enables the administrator to add appliances to the Global View monitoring screen. In the left navigation pane, the administrator can select the **Dashboard** link for specific details (including usage, capacity, and throughput charts) for the appliance hosting the instance of Global View (in this case the DR4000).

Figure 3. DR6000 Global View

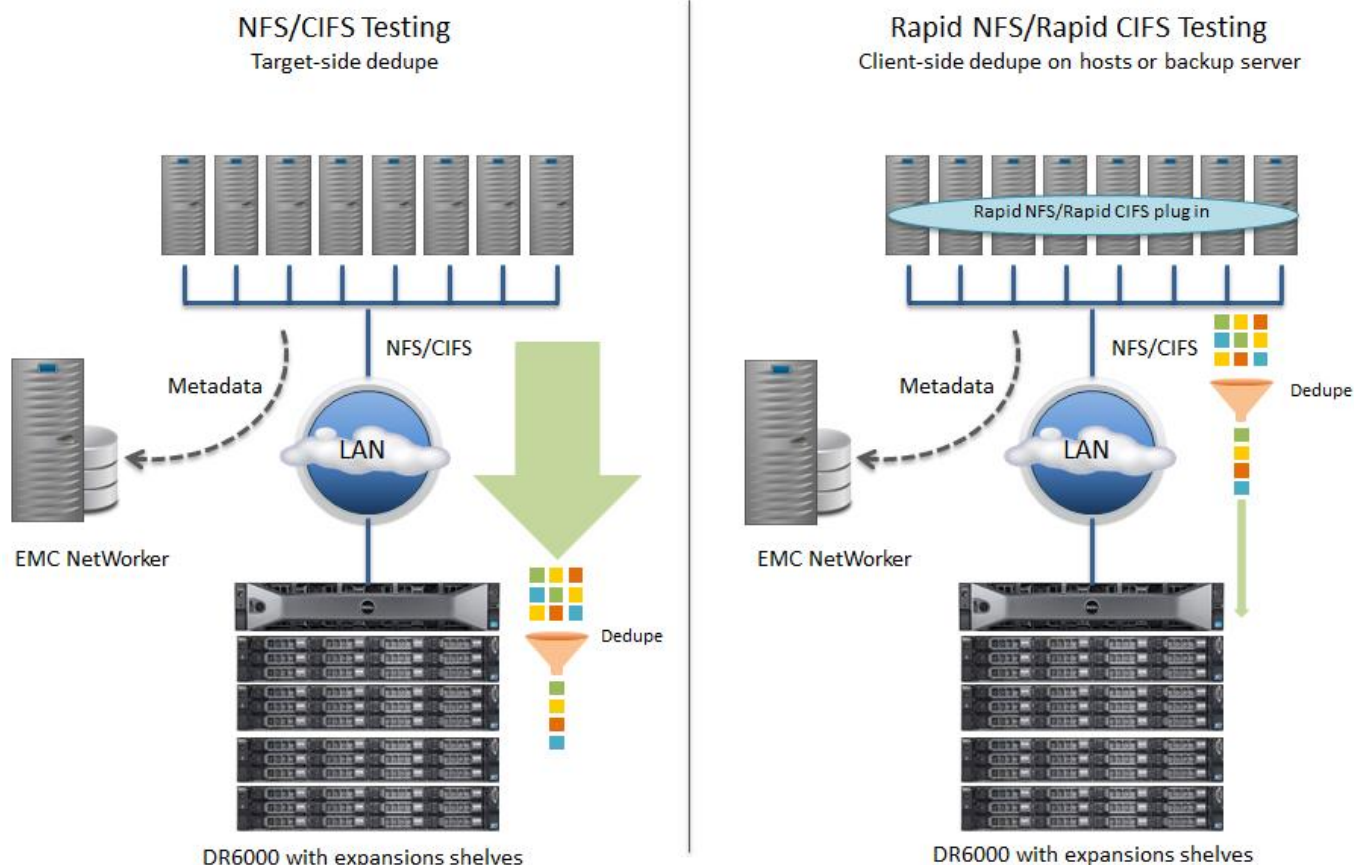


Performance Test Bed Set Up

A significant feature delivered with the DR6000 is new technology for accelerating ingest rates for NFS/CIFS backups. Figure 4 shows an overview of the test environment that was used to compare throughput performance for NFS vs. Rapid NFS, and CIFS vs. Rapid CIFS. The test bed was the same for both traditional NFS/CIFS backups and Rapid NFS/Rapid CIFS backups. Each test bed included eight Dell R310, R320, or R620 servers, all with eight CPU cores and 16GB RAM, connected to the LAN using 10GbE; testing involved scaling from one to eight clients simultaneously.

Backups were run using EMC NetWorker 8.1, and the NetWorker server collected metadata and managed all backups. The Dell DR6000 backup appliance was configured with 18TB of storage, plus four expansion shelves of 18TB each, for a total of 90TB.


Figure 4. NFS/CIFS versus Rapid NFS/Rapid CIFS Test



As shown on the left of Figure 4, backups were performed for NFS/CIFS using DR6000 inline deduplication and compression, so all data was transferred from the clients over the network for every backup. On the right, the Rapid NFS/Rapid CIFS client plug-ins enabled client-side deduplication, so that after the initial backup, only unique data was transferred over the wire.

Figure 5 shows NetWorker screens for the four-client NFS backup test. The background screen shot shows the EMC NetWorker **Devices** tab, indicating the NFS device that the backup clients were writing to. The foreground screen shows the NetWorker **Configuration** tab. Four Linux servers (*r310-perf-01*, *r310-perf-02*, *r310-perf-04*, and *r310-perf-05*) are scheduled for second ingest backups.

Figure 5. NFS Testing, Four Clients, Second Ingest



Why This Matters

Many enterprise organizations require 24x7 uptime and have little tolerance for interruption, even for critical data protection tasks. In these cases, IT looks for ways to speed backups, but that can be complicated by the ongoing growth of data volumes. In addition, many enterprises need to retain data longer, particularly to comply with industry regulations and corporate governance mandates, and as a result, they continuously spend more on backup storage.

The Dell DR6000 provides deduplication and compression to reduce data volumes and save on storage capacity regardless of lengthy retention requirements. In addition, the DR6000 offers accelerated data ingest for a wide range of backup applications using Rapid NFS/Rapid CIFS, RDA, and OST. With the DR6000, enterprise organizations gain the opportunity to speed backups for many backup applications, ensuring maximum uptime for production operations.

For organizations utilizing multiple remote or disaster recovery locations, or many DR appliances at a single location, Global View reduces multi-DR environment management, monitoring, and configuration by offering controls and status from a single pane of glass.

Performance


ESG Lab audited performance tests designed to compare the performance of CIFS vs. Rapid CIFS, and NFS vs. Rapid NFS. In addition, peak throughput testing was conducted for NFS, CIFS, Rapid NFS, Rapid CIFS, RDA, and OST.

NFS/Rapid NFS and CIFS/Rapid CIFS

For the NFS/Rapid NFS and CIFS/Rapid CIFS testing, data was generated using internal Dell tools. Individual 128GB data sets containing file data were created and placed on each client; next, for each client a second 128GB data set containing 90% duplicate data was added. Separate EMC NetWorker 8.1 servers were used to execute backups of 64-bit Windows servers using CIFS, and 64-bit RHEL 5 servers using NFS. Using NetWorker’s “Client Direct” feature, data was sent directly to the DR6000, bypassing the NetWorker server; only metadata was sent to the NetWorker server. To measure throughput with 90% duplicate data, a backup of the 128GB original data set was performed. Next, a second backup containing 90% duplicates of the first backup was done, and throughput was measured for that second backup.

The first test involved a single client backup. Throughput was subsequently scaled by increasing to two, four, six, and finally eight clients. For tests involving multiple clients, throughput was measured during the time all of the clients were sending data. Figure 6 demonstrates backup application performance scalability for NFS compared with Rapid NFS, as the number of clients and backup streams scaled from one to eight.

Figure 6. NFS vs versus Rapid NFS Performance




What the Numbers Mean

- Throughput for NFS traffic scaled from 246 MB/sec for a single client to 1,244 MB/sec for eight clients.
- Throughput for Rapid NFS traffic scaled from 310 MB/sec for a single client to 2,753 MB/sec for eight clients.
- The Rapid NFS throughput improvement increased as clients/backup streams were added.
- With eight backup streams, Rapid NFS delivered greater than 2x throughput compared with NFS.
- Because disk latency for one to two clients limits performance testing, this testing involved adding clients in parallel to push DR6000 throughput performance.

Figure 7 demonstrates backup application performance scalability for CIFS compared with Rapid CIFS backups, using the same procedures.

Figure 7. CIFS versus Rapid CIFS Performance



What the Numbers Mean

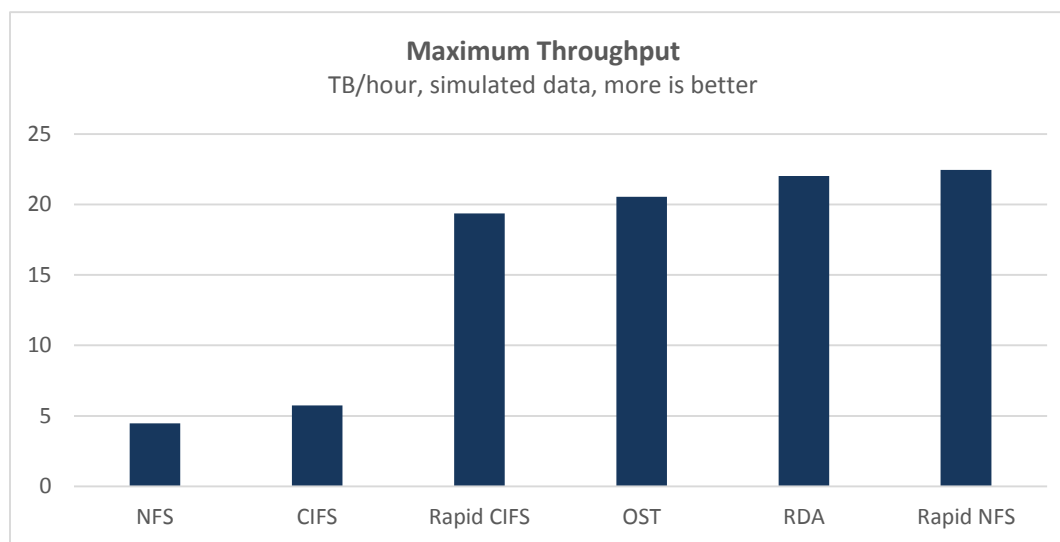
- Throughput for CIFS traffic scaled from 277 MB/sec for a single client to 1,245 MB/sec for eight clients.
- Throughput for Rapid CIFS traffic scaled from 275 MB/sec for a single client to 2,619 MB/sec for eight clients.
- The Rapid CIFS throughput improvement increased as clients/backup streams were added.
- With eight backup streams, Rapid CIFS delivered greater than 2x throughput compared with CIFS.
- Because disk latency for one to two clients limits performance testing, this testing involved adding clients in parallel to push DR6000 throughput performance.

Peak Throughput Testing

ESG audited testing designed to show peak performance of the DR6000 using multiple protocols. Testing included NFS, Rapid NFS, CIFS, Rapid CIFS, OST, and RDA, as shown in Figure 8. The DR6000 was configured with 18TB of storage and four 18TB expansion shelves for 90TB total.²

² The test bed for this project was focused on Rapid NFS and Rapid CIFS. As a result, it was not possible to provide a test bed and number of backup clients capable of maxing out backup throughput for all protocols. The test bed used EMC Networker for peak throughput testing for NFS, Rapid NFS, and Rapid CIFS, and used Dell NetVault Backup for peak throughput testing for RDA. For peak throughput testing of CIFS and OST, ESG audited results of simulated testing using Dell’s internal tool that wrote directly to the DR6000 over 10GbE. The same data set size and duplicate data parameters, Dell clients, and network parameters were used for all peak throughput tests.

Figure 8. Maximum Throughput for Various Protocols



What the Numbers Mean

- Peak throughput for NFS was 4.47 TB/hour.
- Peak throughput for CIFS was 5.74 TB/hour.
- Peak throughput for Rapid CIFS was 19.37 TB/hour, approximately 3.4x the throughput for CIFS.
- Peak throughput for OST was 20.53 TB/hour, approximately 4.6x the throughput for NFS.
- Peak throughput for RDA was 22.01 TB/hour, approximately 4.9x the throughput for NFS.
- Peak throughput for Rapid NFS was 22.44 TB/hour, approximately 5x the throughput for NFS.

Why This Matters

Many large enterprises are accustomed to running hundreds of Windows and Linux clients in a single environment that all need to be backed up. The ability to speed backups enables enterprises to not only save time and get back to production faster, but to also reduce the amount of network bandwidth consumed. This ensures minimal impact to network resources so that other operations don't suffer.

ESG Lab validated that the DR6000 can accelerate backups for NFS and CIFS workloads for multiple backup applications. Scale testing confirmed that as the number of backup clients increased, Rapid NFS and Rapid CIFS enabled greater than 3x faster ingest using source-side deduplication than standard NFS and CIFS. Peak throughput testing confirmed that Rapid CIFS, OST, RDA, and Rapid NFS enabled throughput of 19.37 TB/hour to 22.44 TB/hour, delivering 3x-5x greater throughput than NFS and CIFS.

The Bigger Truth

Through both acquisition and development, Dell has established a powerful and versatile stable of data protection solutions that solve challenges for customers of all sizes. The most recent entrant is the DR6000, an enterprise-focused backup appliance that is a robust follow-on to the midmarket-focused DR4100. The DR6000 adds more capacity, scalability, and performance than the DR4100. Both appliances support a wide range of backup applications, and offer inline target-based deduplication as well as the choice for source-side deduplication through RDA and OST. The DR6000 expands to the list of source-side deduplication options, offering Rapid NFS and Rapid CIFS for source-side dedupe for many more backup applications.

ESG testing focused on two key features of the DR6000: management and performance. The simplified, consolidated management and faster performance of the DR6000 enable customers to maximize their return on investment. The new Global View GUI that can be used with all Dell DR Series appliances using OS 3.0 or greater consolidates monitoring and management of many DR appliances in a single screen. This is just the type of solution that large enterprises are looking for to simplify management and, as a result, reduce costs. In addition, the DR6000's faster performance and additional source-side deduplication opportunities ensure maximum uptime for production operations, regardless of the selected backup application. These feature improvements are right in line with ESG research, indicating that return on investment is the top-cited consideration reported by respondents for justifying IT investments in 2014.³

ESG Lab validated the ability to monitor and manage multiple DR appliances from Global View. Administrators can see an overview of system statistics and health across the enterprise, with dashboards showing storage savings and throughput performance. In terms of performance, ESG Lab confirmed that Rapid NFS and Rapid CIFS improved throughput by more than 3x over NFS and CIFS. In addition, ESG Lab confirmed that the protocols used for source-side deduplication (Rapid NFS, Rapid CIFS, RDS, and OST) enabled 3x-5x the performance of standard NFS and CIFS.

Complexity and inefficiency are two of the most insidious threats to business operations, often draining human and financial resources. The DR6000 builds on the previous foundation of deduplication for capacity efficiency, adding simplicity and performance that can benefit larger enterprises.

Will Dell's conversion to a private company impact its ability to deliver solutions that customers want? Will the company effectively leverage the channel for strategic benefit? It's too soon to make those calls, but if the DR6000 is any indication, its data protection solutions seem to be right on track.

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. 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.

³ Source: ESG Research Report, [2014 IT Spending Intentions Survey](#), February 2014.