

ESG Lab Review

CommVault Simpana 9 “OnePass” Including Integration with HP X9000 Scale-out NAS

Date: January 2012 **Author:** Jason Buffington, ESG Analyst and Lab Engineer

Abstract: This ESG Lab Review documents hands-on testing of Simpana 9 software from [CommVault](#), specifically its “OnePass” data change gathering and retention mechanisms as well as its integration with [HP X9000 \(IBRIX\)](#) scale-out NAS.

The Challenges

Companies of all sizes continue to struggle with the various aspects of data protection. A great deal of attention is paid to solving not only traditional backup/restore, but also adding archiving and storage resource management to their infrastructures. Along with improving backups of virtualization platforms, laptops, and key workloads, ESG research¹ found that IT end-users planning to implement new data protection initiatives had other goals as well:

- 19% plan to implement data archiving
- 19% plan to implement data deduplication
- 18% plan to re-architect their backup processes
- 13% plan to implement reporting of backup/storage

Users attempting to address diverse backup, archive, and reporting needs often employ technologies from multiple vendors—each with their own agent technologies on individual production servers, as well as their own server back-ends and management interfaces. Each point solution performs its own operations on every production server, including traversing the disk, consuming memory/CPU cycles, and contributing to network traffic.

The Solution: CommVault Simpana 9.0 with “OnePass”

CommVault customers running Simpana software have already learned to appreciate something better than a myriad of point solutions. Simpana software’s common platform delivers backup, archive, search and storage resource management administered from a single console. While built on a single software code base, Simpana software modules have previously utilized separate processes and index databases to run archive jobs, followed by backup and, finally, reporting.

Throughout 2011, CommVault regularly added incremental features to its Simpana 9.0 platform—one of which is a new operating methodology referred to as “OnePass,” which enables backup, archiving, and analytical reporting from a single traversal of the file system. By only reading and/or moving data once, redundant backup, archive, and reporting processes are eliminated to speed operations, reduce storage costs, and simplify management.

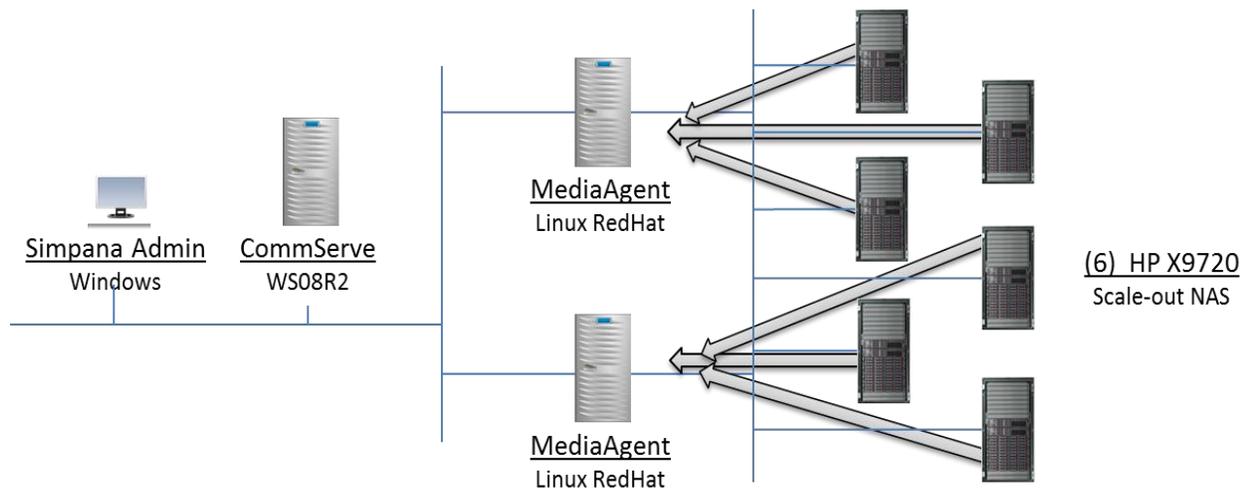
¹ Source: ESG Research Report, [2010 Data Protection Trends](#), April 2010.

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 CommVault.

ESG Lab Testing

ESG Lab tested the new OnePass functionality at a shared CommVault and HP test facility located in Denver, Colorado. The ESG Lab test bed consisted of a typical Simpana software configuration of one CommServe and two MediaAgents, each configured to protect three HP X9720 scale-out NAS nodes sharing a single file system, as seen in Figure 1.

Figure 1. The ESG Lab Test Bed: CommVault and HP Scale-out NAS



The test bed was provided by HP to assess Simpana 9.0’s ability to protect a high-volume of unstructured data.

ESG Lab investigated how CommVault consolidated data protection methodologies using the OnePass architecture. The left side of Figure 2 shows the typical IO patterns of three related data management workflows, including traditional backup, file-archival for reducing disk consumption, and reporting services. The right side of Figure 2 shows the combined workflow of the OnePass-enabled agent in Simpana 9.0.

Figure 2. Comparing Three Traditional Data Protection Workflows to "OnePass" within CommVault Simpana 9.0

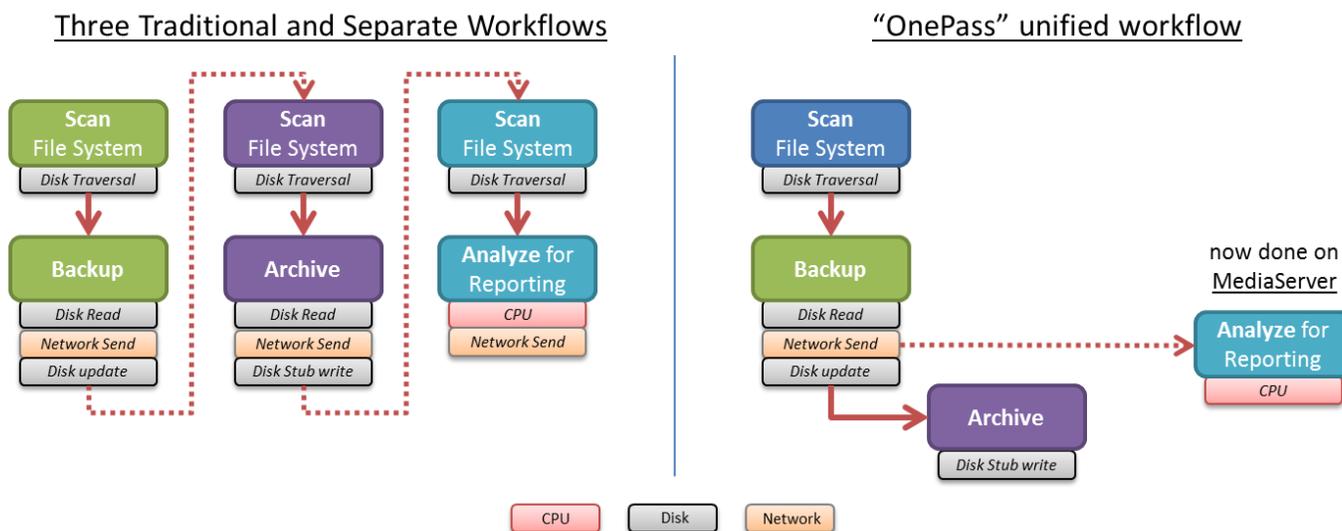


Figure 2 shows how "OnePass" traverses the production storage only once, thereby eliminating significant IO redundancies on the primary server, which should dramatically reduce backup windows and the IO penalties associated with data protection and management tasks.

In a traditional environment using three data management tools, ideally with some level of integration or at least reporting, one might:

1. Perform a traditional backup for data recoverability using traditional incremental methods.
2. After the backup is complete and therefore recoverable (just in case), determine if any files are candidates for archive (hierarchical) management. These files should be "stubbed" to save space, meaning that the original file is replaced with a "stub" pointer referring back to the original file held in near-line storage. This ensures that the actual contents are able to be retrieved transparently when the file is accessed.
3. With the backup finished and the appropriate files migrated to near-line storage, update the reporting system for usage and capacity.

In the case of Simpana OnePass functionality, the operating methodology is similar ... yet optimized:

1. The agent conducts a backup of changed files.
2. With the backup changes successfully committed on the media server, the same agent then assesses the files as candidates for archival, and, if so, stubs the file.
 - No additional file system traversal is necessary because it was done during the backup.
 - No additional disk "read" or network "send" operations are performed during stubbing, as would be required by a separate archival product. The archival process knows that the backup process already read the file and sent it during the backup operation—so it already exists within the Simpana unified storage pool.
 - Either way, the archival routines within the OnePass agent simply perform the stubbing operation of replacing the actual file with a stub—after which the file-system driver will handle retrieval requests in case the file is accessed.
3. With the backup complete and the appropriate files archived, the reporting mechanism updates its information. Again, this occurs without any incremental disk traversal or network operations because Simpana OnePass uses a common index and reporting mechanism from a single collection.

ESG Lab tested the unified OnePass operating model by first conducting separate backups, archives, and report generation using Simpana 9.0 without the OnePass methodology at work. The files were spread across six nodes of an HP NAS and were backed up in parallel by one of the two Simpana media server nodes seen in Figure 1. After the initial testing, ESG Lab audited the results of a similar prolonged test provided by CommVault.

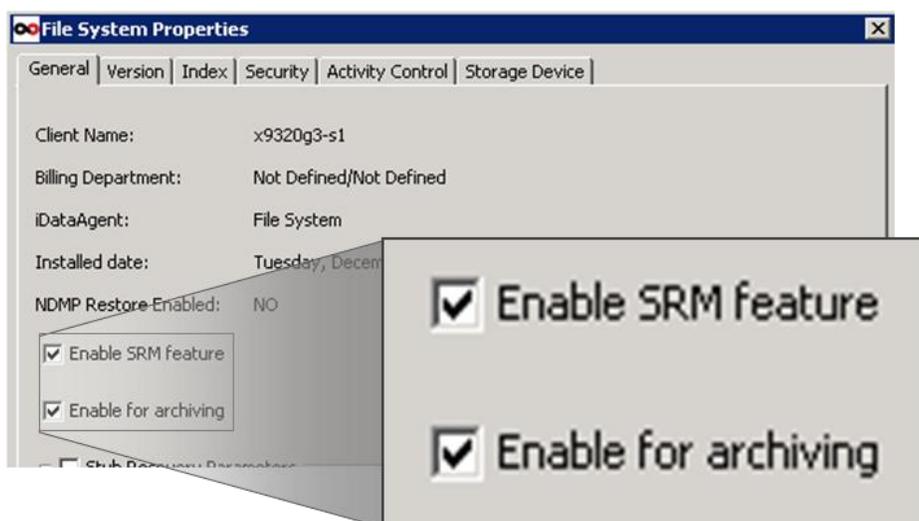
ESG Lab found that the overall backup time was reduced anywhere from 30% to 200% based on three key factors: data types and sizes, amount of redundancy among stored files (e.g., versioning), and archival retention settings that will vary by company. At the low end, even a 30% time savings may mean the difference between compliance with backup window SLAs or not. At the high end, the incremental nature of these backup processes, coupled with nearly transparent archival and SRM functionality, may make the entire backup tax nearly vanish for some production environments.

While less quantifiable, ESG Lab noted that by 1) only traversing the file system once, and 2) offloading the analysis processes to the Simpana MediaServer seen in Figure 2, an appreciable amount of disk IO and CPU processing should be relieved from the production server(s). This means that the production platforms should spend far fewer resources on data protection/management, reserving resulting in more IO and CPU for production purposes.

ESG Lab was impressed by how simple the process was to enable OnePass for Simpana customers. As is typical, the actual agent software components are upgradable through either a push from the Simpana administration console or an .MSI through the customer's typical software deployment tool. The software can be deployed at any point even if the OnePass functionality is not immediately enabled.

Figure 3 shows how enabling the Archive or SRM reporting functions within the unified agent (i.e., enabling "OnePass") is simply a matter of two checkboxes within the backup configuration in the Simpana administration console.

Figure 3. Enabling "OnePass" via Two Checkboxes within the Simpana File System Agent



Why This Matters

Most IT professionals instinctively hope for a unified data protection approach. Historically, they looked for a single backup solution that protected the range of devices in their environments. With continually growing data sets, systems are often becoming "too big" to back up with traditional methods, so solutions for archival and reporting are becoming equally sought after. And while those are good goals, the reality of running at least three different data protection, retention, and analysis agents and processes on a production server is highly undesirable if it means managing multiple tools, supporting many agents, and continually switching between tools due to various financial, environmental, or workload-specific constraints.

ESG Lab found that, with its most recent innovations in the 9.0 Simpana platform (which could arguably be called R2), CommVault seems to have achieved something that most suite-based or pseudo-integrated platform products strive for and that so many backup administrators with multiple products have longed for: not just interoperability across data protection and management processes, but actual unification with a single agent per production platform, running truly combined processes to reduce its disk/network/CPU footprint while still accomplishing multiple protection and management goals.

Simpana Archive Integration with Scale-out NAS

"OnePass" is not the only innovation recently delivered for the Simpana 9.0 customer base. Along with backing up large file systems, CommVault now also offers its archival capabilities as the near-line extension of scale-out NAS platforms, including the HP X9000 (IBRIX) product family.

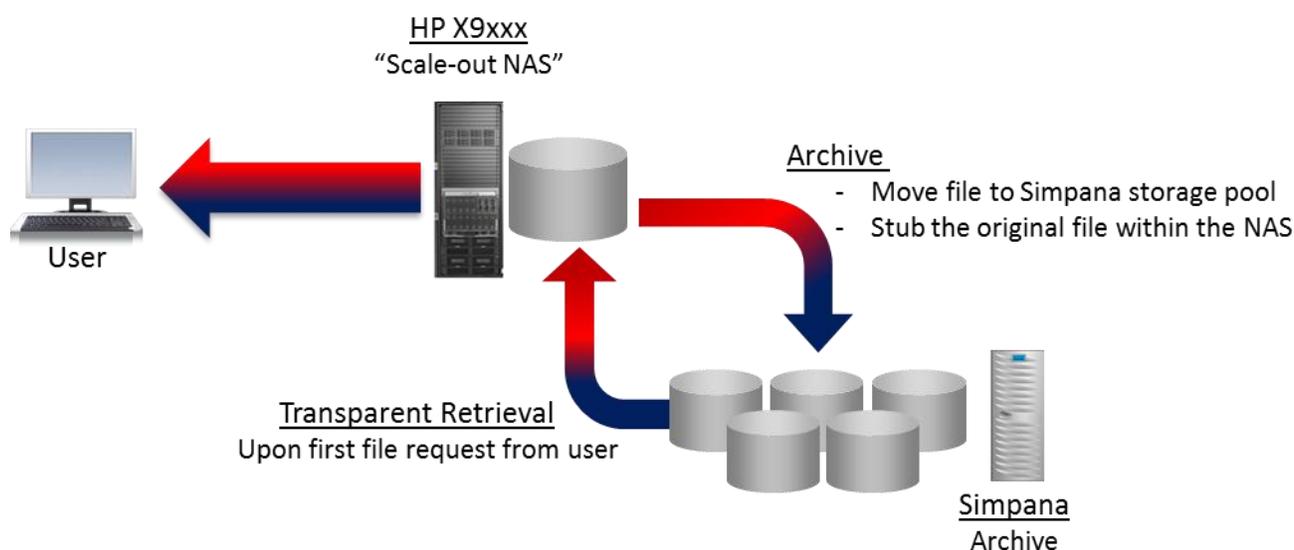
By integrating the Simpana software's archival ability with scale-out NAS, CommVault software is able to offer an additional tier of near-line storage, enabling organizations to leverage a wider range of storage options at a better price point.

ESG Lab Testing

ESG Lab initially treated the HP X9720 platforms as the production server farm being backed up by Simpana. By reconfiguring the test environment, ESG Lab was also able to test Simpana archival storage as a near-line expansion of a scale-out NAS appliance.

Figure 4 shows the reconfigured test bed with the production NAS being archived by the recent enhancements in Simpana 9.0, using the HP X9000 (IBRIX) platforms as a recent example.

Figure 4. Using Simpana software's Archive as Near-Line Extended Storage for Scale-out NAS



In Figure 4, files accessed from the X9000 platforms can be either taken from their own storage pool or transparently retrieved from the Simpana Archive. While the user "sees" all files just as they would expect to within the NAS, those files may be within the primary storage of the scale-out file system or within the Simpana archival storage pool (using any storage that Simpana software supports).

While some NAS vendors provide their own "archival" capabilities through storage tiering and near-line capacity, it doesn't always align with the "unified" data protection benefits described above unless 1) backup and reporting are also performed within the NAS/SAN and 2) the NAS/SAN platform is common across the entire corporate environment. Using a software-based approach, customers may be able to leverage the unified data protection/management capabilities of CommVault software across a wide variety of production servers and NAS platforms consistently—and as a complement to any data management functions that may be offered by the NAS itself.

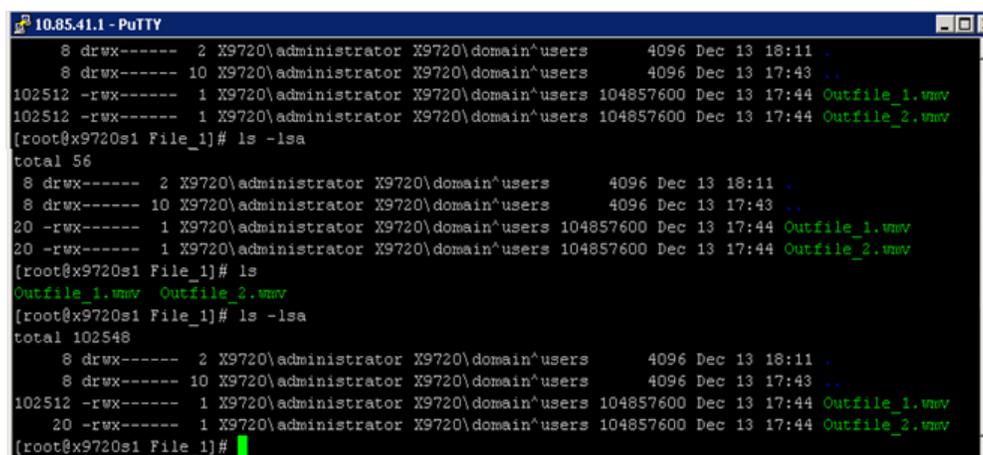
ESG Lab tested this by enabling the Simpana Linux file server agent on each of the HP X9000 NAS nodes. While many data management products purportedly present challenges when integrating with IBRIX platforms, CommVault is able to use its standard agent with the addition of a registry key on each IBRIX node.

After enabling the agent, ESG Lab tested the user experience by defining archival policies within Simpana software for various files and then retrieving them from an NFS client workstation.

Figure 5 shows two files used during testing of the archive integration with scale-out NAS:

- The top file listing shows the files were originally 100 MB.
 - The left statistic reveals each file consumes 102,512 KB.
 - The right statistic reports each file's size as 104,857,600 bytes in the directory listing.
- The middle of the screen reveals that the files were stubbed after archive—consuming only 20 KB each within the NAS, while still displaying 100 MB in the directory listing.
- The last file listing shows that after accessing one of the files, it has been retrieved and thus consumes its regular capacity within the NAS while the other file remains archived until first access.

Figure 5. NFS Client's Experience in Retrieving Files from an Archive-Enabled NAS



```
10.85.41.1 - PuTTY
8 drwx----- 2 X9720\administrator X9720\domain^users 4096 Dec 13 18:11 .
8 drwx----- 10 X9720\administrator X9720\domain^users 4096 Dec 13 17:43 ..
102512 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_1.wmv
102512 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_2.wmv
[root@x9720s1 File_1]# ls -lsa
total 56
8 drwx----- 2 X9720\administrator X9720\domain^users 4096 Dec 13 18:11 .
8 drwx----- 10 X9720\administrator X9720\domain^users 4096 Dec 13 17:43 ..
20 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_1.wmv
20 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_2.wmv
[root@x9720s1 File_1]# ls
Outfile_1.wmv Outfile_2.wmv
[root@x9720s1 File_1]# ls -lsa
total 102548
8 drwx----- 2 X9720\administrator X9720\domain^users 4096 Dec 13 18:11 .
8 drwx----- 10 X9720\administrator X9720\domain^users 4096 Dec 13 17:43 ..
102512 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_1.wmv
20 -rwx----- 1 X9720\administrator X9720\domain^users 104857600 Dec 13 17:44 Outfile_2.wmv
[root@x9720s1 File_1]#
```

Note, that while Figure 5 shows the attributes from an NFS perspective, Windows (CIFS) users would have a similar experience where the actual consumption size is masked and the user perception is all files being offered and stored on the HP NAS.

After enabling archival, ESG Lab configured recurring jobs to enable migration of data from the shared file system within the six IBRIX nodes to the Simpana ContentStore. Files that have been migrated will be returned to make file requests from a client workstation accessing the NFS shares on the X9000. ESG Lab observed no appreciable lag in performance or changes in the users' experience as file requests were routed to the Simpana platform and transparently retrieved from the CommVault software-powered archive.

Why This Matters

ESG research² shows that scale-out NAS is no longer just for extreme usage scenarios; it is becoming more and more mainstream. And while platforms like the HP X9000 (IBRIX) offer significant storage performance, they sometimes require proprietary data protection methods and often lack the extensibility to be protected by more typical third-party software solutions. CommVault and HP/IBRIX have partnered in such a way that a simple registry key enable the Simpana archive capability.

By combining the archive (and backup) capabilities of Simpana with the scale-out NAS functionality of HP's X9000 series, CommVault customers can not only achieve their performance goals for NAS, but do so while managing costs and capacity through Simpana software's archive ability.

² See: ESG Research Report, [Scale-Out Storage Market Forecast 2010-2015](#), December 2010.

ESG Lab Validation Highlights

- ESG Lab examined and tested the combined methodology of "OnePass" with appreciably reduced overall data protection jobs, as well as reduced impact to the production servers due to the consolidated network and disk operations of "OnePass."
- ESG Lab observed how easy it was to enable Simpana software as an archive to a scale-out NAS, without perceivable changes to the end-users' experience.

Issues to Consider

- ESG Lab found that while it would be easy for an experienced Simpana operator to add the OnePass functionality to their environment, the Simpana administration console may appear complex to someone new. This is a reasonable result of a very mature ninth-generation codebase that continually adds new features and options based on feedback from over 15,000 customers.³ Those considering converting to Simpana for its OnePass functionality, its other workload-specific capabilities, or its ability to provide an archival store for scale-out NAS should be prepared for a learning curve which can be offset by training.
- While the HP X9000 is just one of the scale-out NAS platforms supported by the Simpana software archival function, customers will want to ensure that their specific platform is currently covered. With CommVault routinely producing updates and incremental functionality, those not directly supported today may be supported later in 2012.

The Bigger Truth

Most environments struggle with a myriad of data protection and management technologies, perhaps because of workload-specific requirements, data center solutions that are less ideal in remote offices, or simply different data management goals (e.g., backup, archive, and reporting). For many, the sentiment has often been *"If there was a unified solution that did everything well, then we would all own it already."* For others, the potential interoperability of suite-based software or simply complementary products from the same vendor have left customers disappointed as they discovered that each product operates as if it were the only tool that matters.

By simply enabling the "OnePass" capabilities within Simpana 9.0, CommVault customers can enjoy something that many others should find very enviable: a single agent that backs up, archives, and reports on each production server, with only one network stream and significantly optimized disk-I/O impact. The result is something that appears so intuitive that it should be the measure by which other unified products aspire—where functions/technologies may have originally been developed or even acquired separately, but eventually become folded into a single agent talking to a unified back end.

Along with observing the before and after effects of "OnePass," ESG Lab also tested integration of the archival capabilities of Simpana software with scale-out NAS, showing an appreciable benefit to customers with applicable platforms. Without changing the client experience or installing client-side software, even the most advanced NAS platforms can take advantage of an additional tier of storage through the near-line capabilities of Simpana.

If you are currently using a variety of data and management technologies for different purposes and have been disappointed by the lack of integration or coexistence supportability, then Simpana may be exactly what you have been looking for. While individual test results will vary, the fact that common disk reads and network operations are unified should be a valuable optimization method that all environments can take advantage of. Looking at the unified workflow of Simpana software's OnePass methodology should make you ask, *"Why doesn't everyone do it like that?"*

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 CommVault [press release](#), November 2011