

ESG Lab Review

Dell FluidFS v5: Flexible and Secure Scale-out File System

Date: February 2016 **Author:** Aviv Kaufmann, Senior ESG Lab Analyst and Mike Leone, Senior ESG Lab Analyst

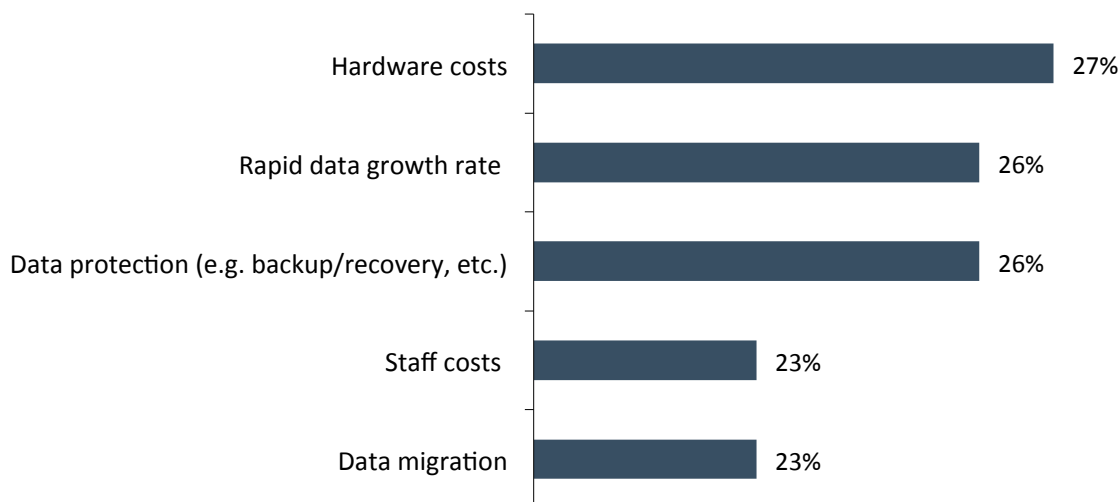
Abstract: This ESG Lab Review documents hands-on testing of Dell's FluidFS v5 scale-out file system with a focus on validating the latest enhancements in the areas of simplicity, scalability, and data governance.

The Challenges

Due to their simplicity and flexibility, NAS storage solutions have proven invaluable to enterprise organizations as in-house solutions for file sharing, backup, and archive needs. The emergence of server virtualization, private and public clouds, and the need to support an ever-increasing mobile workforce has led to increased demands in the performance, capacity, security, and connectivity of NAS storage solutions. With an ever-increasing need to manage and protect a growing amount of unstructured data, it is no small surprise that respondents to ESG research have identified hardware costs, rapid data growth rate, data protection and migration, and staff costs as top storage challenges faced by enterprise organizations.¹ Storing, protecting, and managing access to such a large amount of unstructured data is costly both in terms of capital expenses (hardware and software) and operational expenses (management and expertise).

Figure 1. Top Five Storage Challenges of Enterprise Organizations

In general, what would you say are your organization's biggest challenges in terms of its storage environment? (Percent of respondents, N=373, top five shown)



Source: Enterprise Strategy Group, 2016.

With so much unstructured data to manage and protect, traditional scale-up NAS storage solutions are no longer sufficient. Organizations require a NAS solution that can scale out, grow capacity and performance on demand as needed, bridge the gap between SMB and NFS with a single global namespace, and remain in compliance with security, government, and organizational regulations.

¹ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October, 2015

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 FluidFS v5 and the FS8600 NAS Solution

Dell FluidFS v5 is the fifth generation of Dell’s high performance scale-out file system. FluidFS provides a single, global namespace for NFS and SMB clients, greatly simplifying the management of unstructured data, while offering enhanced data governance and application integration capabilities. FluidFS provides enterprise-class features with simple all-inclusive licensing, backed by Dell’s world-class Co-Pilot support.

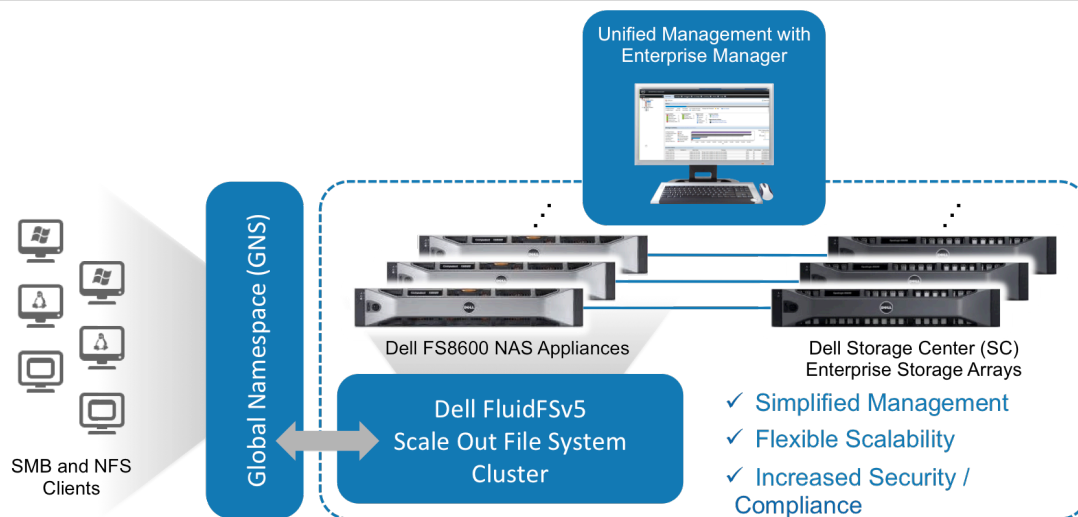
The Dell FS8600 NAS appliance leverages custom Dell hardware that was purpose-built to best support Dell FluidFS technology. The FS8600 consists of two highly available active-active hot swappable NAS controllers connected through a high performance 40Gbps PCIe midplane. Each controller contains two quad-core Intel Xeon processors, 24-48 GB of memory (depending on client network attach option), and an internal HDD for local boot and cache consistency. The client network can be connected using 1GbE or 10GbE connections, and the Storage Center SAN can be connected via 8Gb Fibre Channel or 10GbE iSCSI.

Dell FluidFS v5 running on FS8600 hardware can be deployed as a scale-out clustered file system that leverages Dell Storage Center (SC) storage arrays for storage capacity and enterprise storage features. The FS8600 includes all major file system features such as snapshots, replication, thin provisioning, and data reduction. Storage capacity and performance can be scaled non-disruptively and independently of each other by simply adding disks to SC storage systems, adding additional SC storage systems, or by adding additional FS8600 appliances and/or additional SC storage systems to the cluster. Host connectivity and data management are greatly simplified since all capacity is provided and managed under a single, global namespace (GNS).

New features supported in Dell FluidFS v5 Include:

- **Protocol Enhancements:** Support for user data via FTP, improved SMB performance (through support of BranchCache and sparse files), and improved NFS interoperability (support for using the AD Global Catalog as an LDAP repository).
- **VMware Enhancements:** Consistent snapshots for VMs running on NFS datastores on FluidFS.
- **Access Control Enhancements:** Support for AD Group policy-based file security policies, data governance audit support through Dell Change Auditor and other third-party auditing solutions, and the ability to tighten security by disabling self-restore by end-users.
- **Architectural Enhancements:** Introduction of FluidFS Global Namespace that can provide access to 10s of PBs of data spread across discrete devices under a single SMB or NFS mount point.
- **Management Enhancements:** Support for PowerShell and REST system management APIs for seamless integration with existing customer applications.

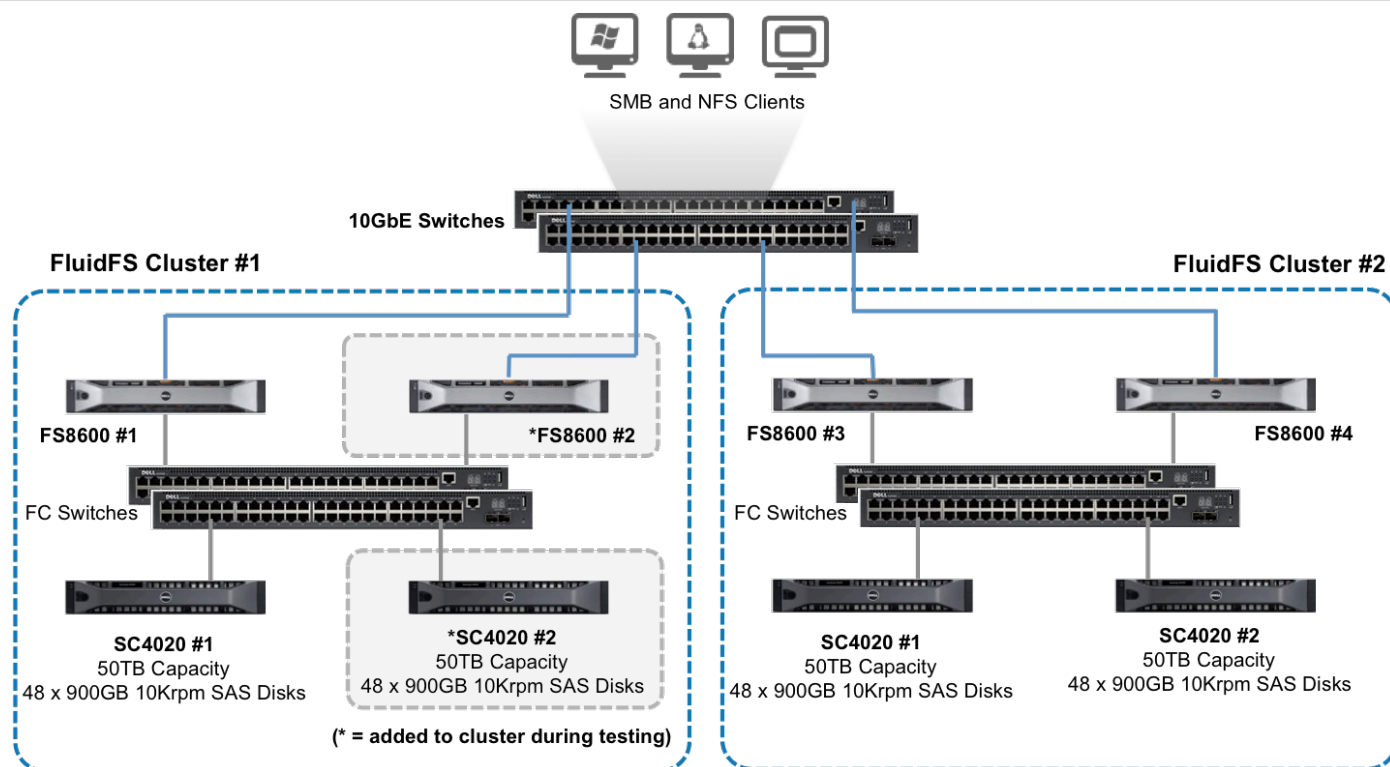
Figure 2. Dell FluidFS v5 Deployed on the FS8600 NAS Appliance



Simple Flexibility

ESG Lab performed remote, hands-on testing of the Dell FluidFS file system running on two separate FS8600 clusters with block FC storage provided by up to four Dell/Compellent SC4020 storage systems. ESG Lab validated the simplicity and flexibility built into the Dell FluidFS scale-out file system with a goal of seeing how Dell FluidFS can help to simplify the management of an ever-growing quantity of unstructured data for an organization while keeping that data available to clients on a global scale under a single namespace. Of particular interest to ESG were some of the new protocol and architectural enhancements included in version 5, making cross-protocol client access and file system scalability simpler, with greater flexibility in deployment options. Figure 3 shows the ESG Lab validation test configuration.

Figure 3. ESG Lab Test Configuration Used in Hands-on Validation



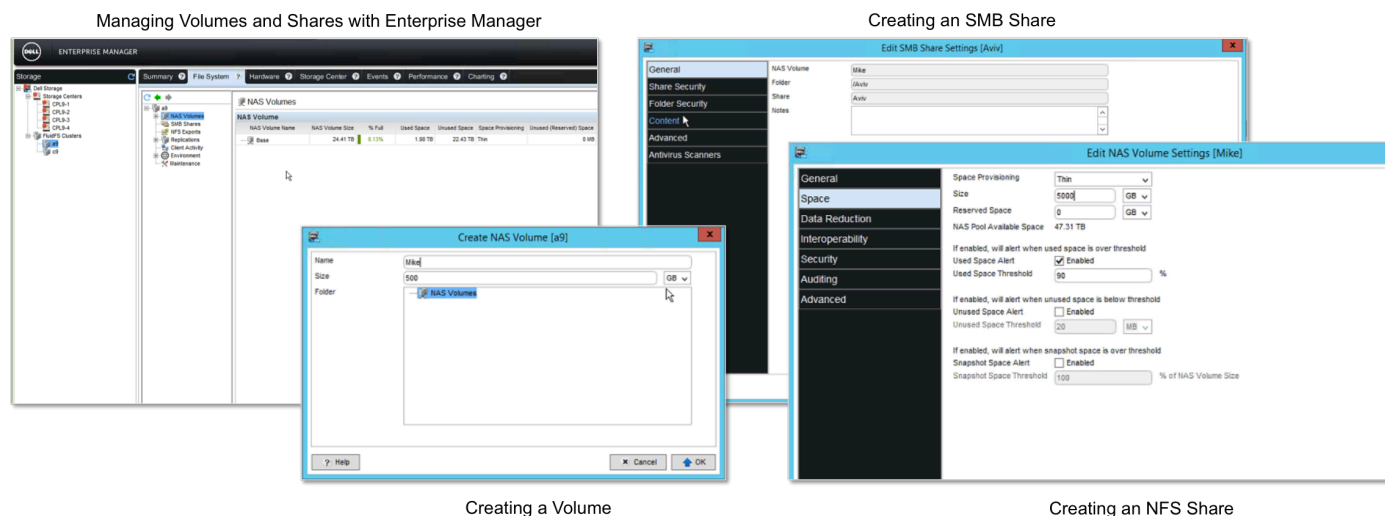
To get started, ESG Lab reviewed the steps required to deploy an FS8600 cluster for the first time. One of the key value propositions of the FS8600 cluster is that the initial deployment, configuration, installation, and testing are performed by Dell professional services and included in the price of the system. Dell service professionals deploy the FS8600 appliances and network connections, and then use Dell Enterprise Manager to connect to the appliance; provide the necessary IP addresses; set up the network, clustering, and SC storage configuration; and join the Active Directory before creating the first NAS storage pool. The entire process is performed in under an hour, removing the burden of installation from the end-user and ensuring that the deployment is protected from any unforeseen issues down the road.

After deploying the initial configuration, consisting of two FS8600 clusters (each connected via FC to a single SC4020 storage array), ESG Lab used Dell's familiar Enterprise Manager software to manage the entire deployment, including the FluidFS clusters, FS8600 NAS appliances, and SC4020 storage arrays. Using the simple and intuitive interface, ESG Lab created a 50TB NAS storage pool that used capacity provided by the SC4020 storage array. The allocation of capacity on the SC4020 was completely automated and transparent to the administrator, with no additional steps needed to create or map volumes on the SC4020. This, combined with the proven automated tiering capabilities of the SC4020, greatly simplified the management of storage, saving the administrator valuable time and reducing complexity.

ESG Lab then created new thin-provisioned NAS volumes as well as new SMB and NFS shares. It was easy to control the shares access permissions, modify file-based policies, and allow access to clients in a variety of ways (IP address, IP

address range, subnet, or clients in a Netgroup). Mounting the volumes on clients was just as simple and required no extra steps because the clients were already members of the Active Directory (AD) domain. Even the NFS volumes were able to leverage the Global AD catalog as an LDAP repository, saving administrators the burden of having to maintain separate repositories for NAS shares. The mounted volumes simply appeared in the client's file system automatically following creation of the volume. The simple volume creation process and management interface are shown in Figure 4.

Figure 4. FluidFS Simplified Management of NAS Volumes and Client Connectivity



Next, ESG Lab validated the ability of Dell FluidFS v5 to non-disruptively scale in a variety of flexible ways. The design of the FluidFS architecture allows organizations to scale the capacity and/or performance characteristics of the global file system independently of each other. A FluidFS file system can be scaled by adding drives, storage systems or channels, NAS appliances to the cluster, or by adding ports at any level. Not having the capacity tied to the performance means that organizations can make better timed choices, enabling them to take advantage of the latest technologies in drives, connectivity options, storage systems, and NAS appliances, if and when they make the most sense to the business.

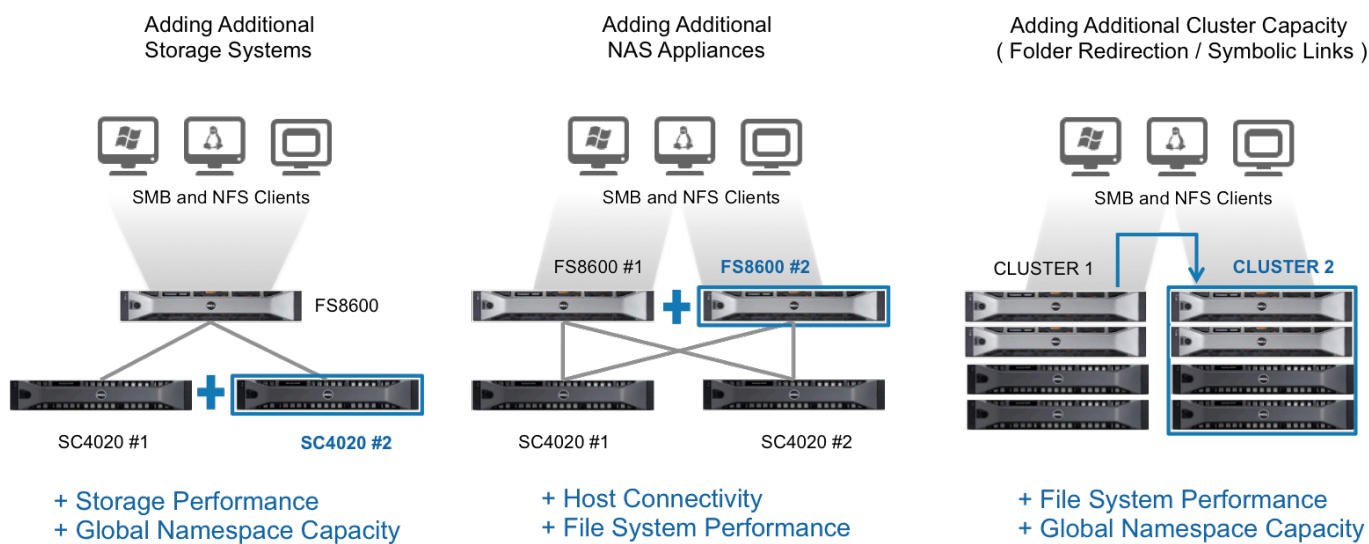
To create a load on the system, ESG Lab used the open source tool, *frametest*, to simulate a 1.3GB/Sec workload that consisted of heavy 12MB video write streams (a workload that might be used by a video surveillance organization, for instance). It should be noted that this write-intensive workload was used to place a heavy load on the system in order to gauge the ability of the solution to provide high levels of performance while scaling capacity and does not represent the maximum performance of the FS8600 NAS scale-out solution. ESG Lab used the performance charting capabilities of Enterprise Manager to view the performance of the system over time as it was scaled to grow capacity and/or performance capabilities by adding new components to the configuration.

While the workload was running, a second SC4020 was added to the cluster and the NAS pool was extended to a total of 100 TB to make use of the new capacity provided by the second array. ESG Lab noted that there was no impact to the running workload, and that writes (including redirected re-writes) were sent to the newly added SC4020 until system capacity was balanced. Up to eight SC arrays could provide storage for a single FluidFS cluster (up to a 4PB single file system without the need to use Microsoft DFS), and the SC4020 storage arrays could be used in a unified manner to provide storage to block-based hosts and the FS8600 cluster simultaneously. ESG Lab was then shown how Dell professional services could add a second FS8600 appliance to the cluster with no impact to the running workload, providing greater parallelism and potential performance in both host and storage connectivity.

Finally, ESG Lab grew the namespace by giving a mounted host access to a volume located on the second FS8600 cluster through redirection. Folder redirection can be used to merge portions of two separate file systems located on different clusters into a single Global Namespace or mount point for both SMB and NFS clients. This allows organizations to grow a namespace without adding any additional hardware to a cluster and may be useful in cases in which organizations or datasets must be merged while providing a single view to all clients.

To validate this ability, ESG Lab used two identical FluidFS clusters (each consisting of two FS8600 appliances and two SC4020 arrays). On the first cluster, we created a file system and a base volume, and added both an NFS export and SMB share to the root of the volume. On the second cluster, we created a 5TB folder named “/AdditionalData” (with an SMB share and NFS export that was added to the root of the folder) that would be used to non-disruptively grow the file system on the first cluster through redirection of existing NFS and SMB clients to the shares on the second cluster. Under the Global Namespace tab of the “base” NAS volume on the first cluster, we created a redirection folder named “/AddedData” that pointed clients to the virtual IP address of the second cluster and the SMB share and NFS export associated with the “/AdditionalData” folder. The folder and additional capacity immediately showed up as a sub-folder under the base file system on both SMB and NFS clients (as a symbolic link or redirected folder). This redirection ability gives organizations the ability to provide all clients with a single consistent global namespace for SMB and NFS clients while eliminating interoperability issues and complex management procedures. Figure 5 shows the three flexible methods validated by ESG Lab to non-disruptively grow the global file system.

Figure 5. ESG Lab Validated Flexible Options to Grow Capacity and Performance with Dell FluidFS



Why This Matters

NAS storage solutions have been widely adapted in the enterprise, thanks mainly to the ability to effectively store and protect large amounts of unstructured data. In fact, 25% of ESG research respondents identified NAS as their primary on-premises, disk-based storage technology, making it the most-cited response.² While scale-out NAS solutions have solved some of the historical problems related to scale, many scale-out NAS storage solutions remain difficult to deploy and complex to manage. Administrators still struggle with managing access for clients across devices and protocols while providing a single unified view of a global namespace.

ESG Lab validated that Dell FluidFS was simple to deploy and manage and could be scaled in a variety of flexible ways to match the needs of the business. The file system grew in both capacity and performance capabilities as we added SC storage systems and FS8600 NAS appliances with no disruption to client activity. FluidFS eliminated most of the complexity of managing unstructured data spread across discrete NAS devices by providing a global namespace to SMB and NFS clients while eliminating interoperability and management issues, providing a better end-user experience, and reducing both operational and capital expenses.

² Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

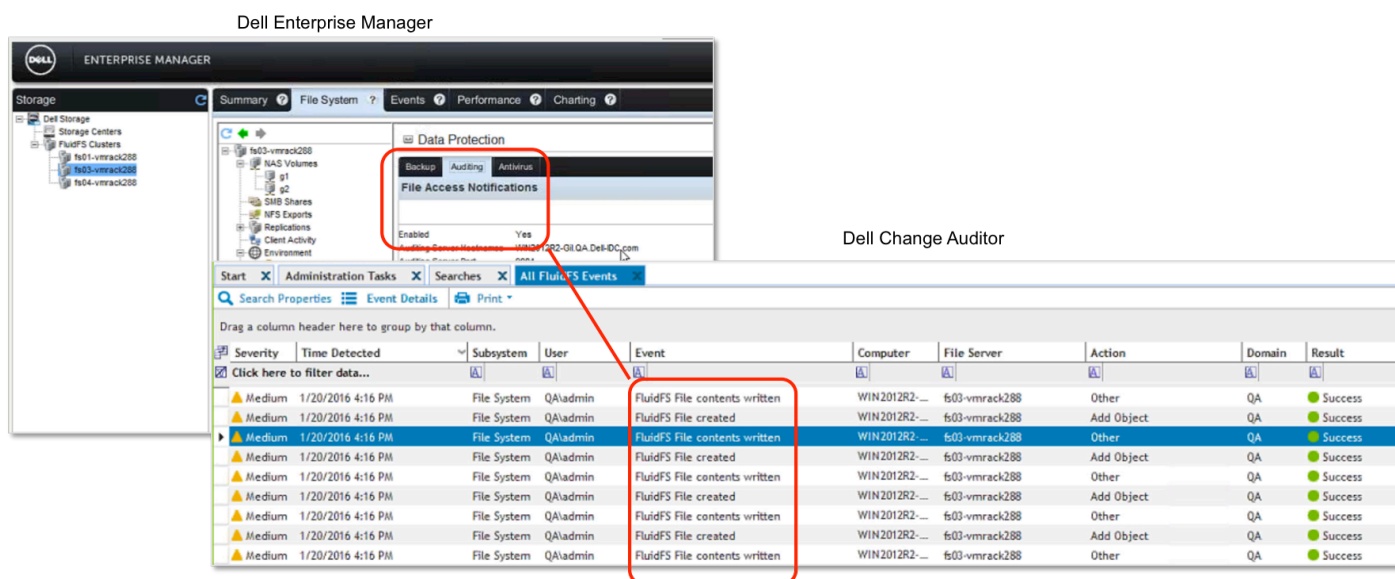
Security and Data Governance

Next, ESG Lab validated some of the ways in which a Dell FluidFS v5 deployment can be used to provide more secure access to data for clients and remain in compliance through integration with Dell Change Auditor. With the added pressure to remain in regulatory compliance and protect against data breaches, it is critical for organizations to know who is accessing their data, and how they are doing so. File systems typically rely on system access control lists (ACLs) to control access to data. If an end-user has permissions to modify an ACL, then she can easily get to any file she wants. To help organizations understand and better control data access, Dell FluidFS v5 added support for integration with leading data governance tools such as Dell's own Change Auditor and other third-party auditing solutions. To get a better understanding of this tight integration, ESG Lab performed hands-on validation of some typical data governance scenarios using FluidFS with Dell Change Auditor.

Dell Change Auditor can be used to securely audit all activities related to ACLs and file access to the FluidFS file system as well as to Active Directory, Exchange, Windows file servers, SQL databases, and other applications. Change Auditor can be installed on a physical or virtual server and uses a SQL database to store events, provide searchable reports, and generate alerts on suspicious activity. Agents are deployed as close as possible to monitored devices and a fully customizable amount of information is reported back from each monitored device to the Change Auditor database.

After installing the SQL database, ESG Lab deployed Change Auditor on a VM by simply defining a name for the instance and joining it to the Active Directory. The Change Auditor agents were then installed on the Windows servers and automatically connected to the appropriate Change Auditor coordinator running in the AD. Once installed, we could use Enterprise Manager to see that the FluidFS cluster was actively being audited by Change Auditor. Using the Change Auditor interface, we could review any activity on the Dell FluidFS file system. No additional interaction was necessary on FluidFS or Change Auditor to begin the auditing process. In Change Auditor, it was easy to select which volumes, directories, and file types to monitor and to fully customize the amount of information that was logged to create an optimal balance between logging overhead and the amount of information gathered. File data on the FluidFS device intended for wider consumption could be audited in a quick and simple manner, while access to more sensitive information on the FluidFS system, such as financial data, can be scrutinized much deeper. Figure 6 shows the tight integration between FluidFS v5 and the Dell Change Auditor interfaces.

Figure 6. Integrating Dell Change Auditor with Dell FluidFS v5

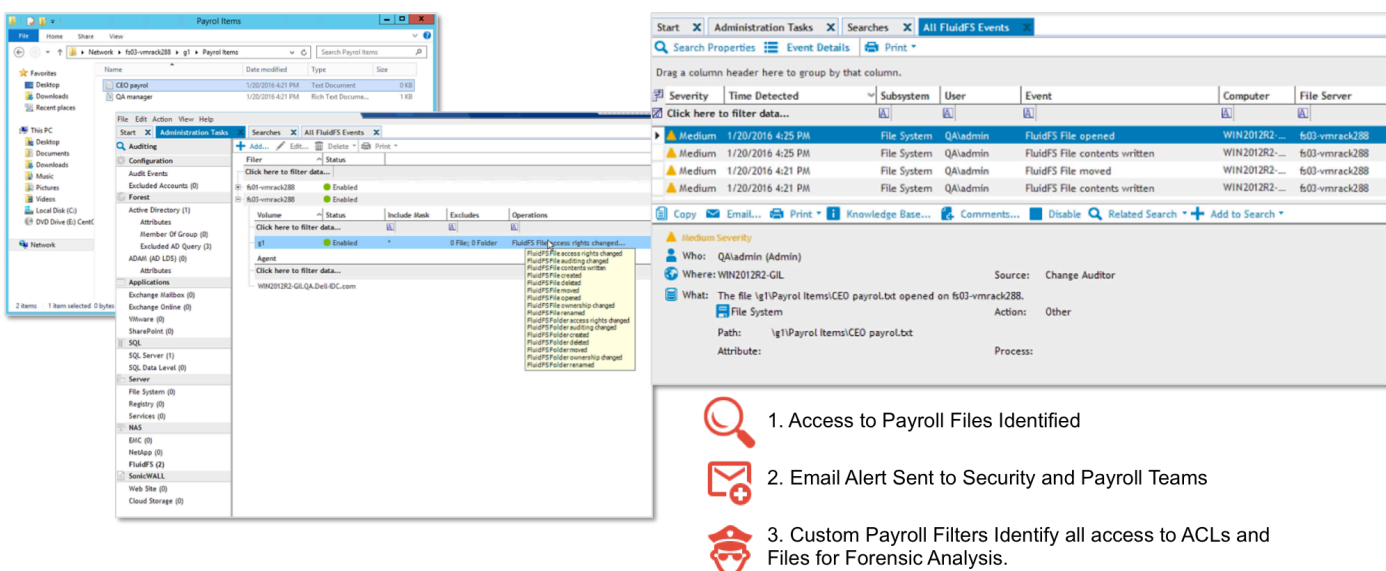


Next, a set of files was copied to the FluidFS file system and in under ten seconds, the newly written files were visible in the Change Auditor user interface. Rather than update each file in real time, caches on the FS8600 help to combine and optimize reported information in consolidated bursts to limit network traffic. ESG Lab then performed some typical file activities to the copied files. The Change Auditor interface was then used to easily create filtered searches to query

events and see when files were opened, modified, deleted, renamed, or copied. File events were categorized and color-coded by severity for quick visual analysis with simple file operations categorized as medium severity, and modifications to the Active Directory ACL categorized as high severity.

Finally ESG Lab walked through a simulated scenario to see how Change Auditor can be used to better protect an organization by quickly identifying the possibility of an internal or external data breach. In a directory named “Payroll,” we placed a few simple text files named “CEO Payroll” and “QA Manager” to represent organizationally sensitive information. Using Dell Change Auditor, we then created custom filtered searches and alerts to monitor all actions related to activity on the payroll files and ACLs and then send alerts to both the payroll and security teams should anyone access the files. After simply opening the payroll file, ESG Lab verified that the alerts had been sent and a large amount of forensic information was available to help identify the methods used to gain access to the file. It was quite apparent that Dell FluidFS file system’s integration with Change Auditor can help organizations identify and react to possible breaches quicker, while giving them the information necessary to better prevent future breach attempts.

Figure 7. Using Dell FluidFS v5 and Dell Change Auditor to Quickly Identify a Simulated Payroll Data Breach



1. Access to Payroll Files Identified
2. Email Alert Sent to Security and Payroll Teams
3. Custom Payroll Filters Identify all access to ACLs and Files for Forensic Analysis.

Why This Matters

Never before has public awareness of data breaches been more apparent. Reported data breaches at Sony, TJX, and eBay have helped to spread fear into the minds of consumers, which can ultimately impact the bottom line of the business. ESG research revealed that, as the result of a security incident within the past 24 months, 32% of organizations have suffered a breach of confidential data and 26% have experienced a violation of regulatory compliance or audit failure.³ With the added pressure to remain in regulatory compliance and protect against data breaches, it is critical for organizations to know who is accessing their data, and how they are doing so. But security is expensive, and a lack of skills certainly makes the job of preventing data breaches more difficult.

ESG Lab validated that the built-in security features and tight integration with data governance tools simplifies the process of keeping data safe on the Dell FluidFS v5 file system. ESG Lab was able to quickly and easily configure FluidFS and the Change Auditor to log any access to simulated payroll data, send notifications and alerts to the security and payroll administrators when files were accessed, and perform forensics on the access log to better identify vulnerabilities and help prevent future attacks.

³ Source: ESG Research Report, [Cyber Supply Chain Security Revisited](#), September 2015.

The Bigger Truth

Unstructured data continues to grow at an amazing rate. ESG research revealed that 88% of organizations believe that their storage capacity is growing at double-digit annual growth rates.⁴ Despite carrying the general perception of being complex to manage, NAS storage solutions continue to be the technology of choice when it comes to primary enterprise storage. While scale-out NAS storage solutions have taken a step in the right direction in simplifying the historically complex management and deployment procedures, most of these solutions offer little flexibility beyond adding another identical brick and have done little to simplify the task of supporting both SMB and NFS clients simultaneously.

Dell FluidFS is a fifth generation scale-out file system that is field proven to simplify the management of unstructured data while delivering the scalability and performance required for traditional NAS file workloads (like file sharing, backup/archive, VMware, and Oracle) and more demanding workloads (such as video surveillance, medical imaging, life sciences, media and entertainment, and higher education). Dell FluidFS v5 running on FS8600 hardware can be deployed as a scale-out clustered file system leveraging Dell Storage Center (SC) storage arrays for storage capacity and enterprise storage features.

ESG Lab validated that Dell FluidFS was simple to deploy and manage and could be scaled in a variety of flexible ways to match the needs of the business. Because the system is installed, configured, and tuned by Dell services, the global namespace grew in both capacity and performance capabilities as we added SC storage systems, FS8600 NAS appliances, and even FluidFS clusters with no disruption to client activity. FluidFS eliminated most of the complexity of providing a global namespace to SMB and NFS clients while eliminating interoperability and management issues, providing a better end-user experience, and reducing both operational and capital expenses.

ESG research revealed that information security was both the IT priority most-cited by respondents and also the area in which their organizations were most often identified as having a shortage of skills for 2015.⁵ ESG Lab also validated that the built-in security features and tight integration with data governance tools simplifies the process of keeping data safe on the Dell FluidFS v5 file system. ESG Lab was able to quickly and easily configure FluidFS and the Change Auditor to log any access to simulated payroll data, send notifications and alerts to the security and payroll administrators when files were accessed, and perform forensics on the access log to better identify vulnerabilities and help prevent future attacks. Tight integration between the FluidFS and Dell Change Auditor interfaces simplified security and governance so that an IT generalist could perform all of the required tasks, lowering operating expense and freeing up valuable security resources for more pressing issues.

While a scale-out file system that leverages modular building blocks certainly makes growing a global file system easy and predictable, there is still something to be said about having the flexibility to make the optimal choices in technologies to meet the ever-changing needs of the business. Dell FluidFS v5 offers the best of all worlds by making deployment and management of the global file system and client connectivity simple, without losing the flexibility to choose how to best grow the environment. Not having capacity tied to performance means that organizations can make better timed choices, enabling them to take advantage of the latest technologies in drives, connectivity options, storage systems, and NAS appliances, if and when they make the most sense to the business. If your organization is looking to simplify the management of your scale-out NAS environment while better meeting the challenges imposed by security, data governance, and client connectivity, ESG Lab recommends that you consider some of the new capabilities of Dell's FluidFS v5 and the FS8600 NAS appliance.

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.

⁴ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

⁵ Source: ESG Research Report: [2015 IT Spending Intentions Survey](#), February 2015.