

I D C T E C H N O L O G Y S P O T L I G H T

Efficiently Delivering Enterprise-Class File-Based Storage

June 2010

Adapted from *Novell Delivers a New Way of Intelligently Managing Organizations' File-Based Information* by Noemi Greyzdorf, IDC #216013

Sponsored by Novell

The explosion of file-based unstructured data is creating a number of new challenges for organizations in terms of both storage and data management. In the past two years, the consumption of storage systems for file-based data far outpaced the consumption of storage systems for block-based data. This Technology Spotlight examines the operational challenges that arise as more storage is consumed for file-type data and as new policies for the long-term retention of file-type data emerge. The paper discusses the most visible challenges revolving around controlling the costs associated with data migrations, ongoing data management, and escalating storage facility costs. Within data management, the challenges are around maintaining data protection, improving information management, and improving governance, risk, and compliance. The paper also looks at the role of Novell File Management Suite (NFMS) in this important market.

The Economic and Operation Challenges Posed by the Information Explosion

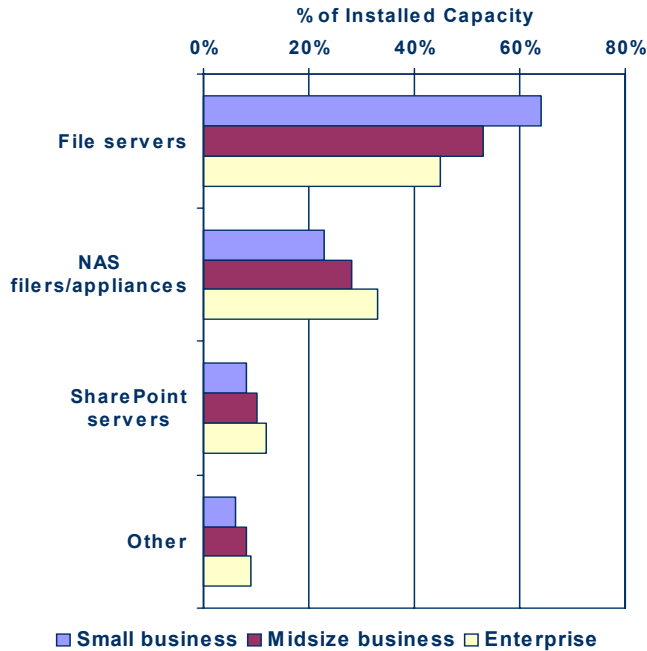
The rapid growth in file-based information within organizations is exposing the weaknesses of existing IT management practices for file-based storage. Many companies continue to deploy large numbers of dispersed, underutilized file servers in support of different workgroups and geographically remote sites. Others have consolidated file-based storage on larger NAS systems but are still struggling to deal with rapid data growth and rapid technology changes (see Figure 1).

The creation, use, and long-term retention of file-based data within organizations also pose significant challenges around optimizing storage utilization for file-based data, reducing the burdens associated with data migration and backup of file-based storage assets, and creating and maintaining an accurate inventory of organizations' file-based information. Some of the more significant challenges are the following:

- **Inefficient storage utilization and backup.** Rapid but variable file growth leads to wide imbalances in file server utilization levels and disrupts backup processes. The actual backups often consist of older, unchanging files (that have already been backed up) or unneeded copies of extraneous files (e.g., employee MP3 files). These low-priority files often account for 80–90% of all files being backed up and can drive backup times beyond acceptable backup windows.
- **Time-consuming and expensive data migration.** The use of dispersed file servers leads to severe underutilization of storage capacity, but efforts to migrate and consolidate data from even a limited number of file servers onto larger, more centralized NAS systems can often take half a year or more.
- **Increased uncertainty and risk in data protection/retention.** The growing diversity of file-based data pools makes it difficult for IT staffs to effectively catalog and organize file-based data. Without knowledge about what is being stored and how it is being used, IT staffs can't respond to performance and cost mismatches or ensure that data protection policies are being followed.

FIGURE 1

File-Based Storage Options



Only 26% of small businesses and 14% of midsize businesses rely solely on file servers.

Main file-based storage uses are:

- File/print (44%)
- Backup/archive (32%)
- Virtual machine (17%)

Main challenges today are:

- Large: remote file servers
- Medium: NAS proliferation
- Small: FS sprawl

n = 195

Source: IDC's 2009 Trends in File-Based Storage Survey

The growth in file-based data and the data management challenges it poses are becoming major issues for IT organizations planning future datacenter investments:

- More, inefficiently managed file-based storage requires datacenter space and energy to power and cool systems. In recent years, this has become a major problem in some geographic areas because the availability of energy has been limited by the delivery network. No matter what an enterprise is willing to pay, there is just no more energy to be had.
- Though datacenter space may not seem like an issue, adding more file-based storage to an existing datacenter is an incremental cost; when a datacenter becomes full, the cost of adding more storage becomes extremely high. There is great incentive to store as much as possible in the available space (e.g., leveraging denser SATA base storage systems), avoiding having to add space or move to a bigger facility.
- Growth in storage means adding more capacity that needs to be properly allocated and provisioned. In many organizations, the sprawl of file-based storage systems has created the challenge of managing it for capacity and performance. With some systems sitting idle while others are at capacity, organizations need a more efficient way to load balance users and their data.
- Traditional architectures require significant investments in the migration of data at times of storage systems refresh or organizational restructuring. A lot of time may be spent on planning and migrating data across storage devices, costing organizations time and money.

Achieving Storage Efficiency Through Intelligent Data Management

Today, organizations use many tools to manage storage devices, but these tools are much less useful when it comes to managing information, especially when stored as files. The industry is actively designing solutions that help achieve greater operational efficiencies when it comes to managing exploding file-based data assets. The challenges on the data management side are more extensive and complex, with significant implications for an organization's infrastructure costs and business responsiveness:

- Protecting data is critical for organizations. As the amount of data that needs to be protected and hours of operations increase, the window for data protection continues to shrink.
- Organizations rely on their information assets to stay competitive in the market. With the onslaught of new data, it is often very challenging for organizations to identify valuable data sets and ensure they are available, accessible, and retained for the duration of their value life cycle.
- The proper retention of data addresses the need to reduce the cost associated with older less frequently accessed data while also mitigating risks of data loss and compromise, lack of compliance with regulations, and failure to employ best practices around data governance — in other words, the right data at the right place at the right time.

In the world of file-based unstructured data, storage administrators want to leverage intelligent file management solutions that drive efficiency from the hardware investments and improve the effectiveness and efficiency of storage and data management. Such a solution provides the following:

- Allows organizations to deploy (and consistently manage files across) a wide range of disk storage tiers with different performance, capacity, availability, and cost characteristics, without requiring wholesale displacement of existing assets
- Uses a common set of scalable and highly available data migration, data life-cycle management, and data protection services across these multiple storage tiers
- Provides a data/file discovery/indexing capability so that IT administrators can inventory existing file-based storage assets, identify mismatches in usage and/or value, and enable proactive policies for intelligent storage tiering and data retention

Considering Novell File Management Suite

Novell File Management Suite has been designed to address the challenges administrators face in managing growing data and expanding storage. NFMS' approach to helping administrators drive greater efficiencies is to manage data and storage with the understanding of what that data represents and who owns it.

Through identity-based policies and integration with directory services, Novell File Management Suite automates data management, ensuring that user data policies and storage management standards are enforced.

Novell File Reporter

It's impossible to design policies without understanding an organization's data assets. Novell File Reporter helps managers understand what data is stored and the value it potentially represents to the organization: by owner, by location, by date, by server, or across all network shares. The information gained using Novell File Reporter is critical in the process of taking control of file-based data. Once an organization better understands its data assets and how they are being used, Novell Storage Manager and Novell Dynamic File Services can be used to set and execute policies that will create an efficient and effective file management framework.

Novell Storage Manager

Novell Storage Manager (NSM) allows administrators to manage file resources based on users' identities and customizable policies. This creates a management framework that simplifies and automates many provisioning and allocation tasks that are costly and prone to human error. Administrators can create user templates and assign permissions based on predefined roles associated with the template. The result is greater data security through policy adherence and a timely and accurate assignment and deprovisioning of access rights, cost reduction, and productivity gains by eliminating complexity and automating administrative tasks. The templates also facilitate governance best practices by providing a way to track and audit access rights to information.

Most powerfully, this approach allows administrators to manage all user identities and information across all file servers as a single pool rather than separate independently managed domains. Policy can be applied consistently, and any change in an organization's or a user's location doesn't require major reconfigurations and transfers to network drives.

Novell Dynamic File Services

Not all data has the same value to an organization, and this value may change over the life cycle of that data. In an effort to reduce storage costs, organizations seek dynamic ways to move data to the most appropriate storage tier based on data requirements. Novell Dynamic File Services executes on the policy that requires data to reside on different tiers of storage at different times in the life cycle. The migration of data to a different storage tier reduces the overall cost of storage by leveraging more cost-efficient storage media. Additionally, migration of data may reduce the amount of data that needs to be backed up. Moving data across storage tiers dynamically without disrupting the end user creates operational efficiencies by reducing complexity and administrative resources and lowers overall capital expenditures.

Organizations must be able to respond to the market; this means that all resources must be aligned with business objectives. The Novell File Management Suite has been designed to align storage resources with the demands of the organization, making collaboration simpler, managing file data-based identity, and reducing complexity and costs in how an organization creates, tracks, and disposes of file data. The results are achieved without impact to the end users or a significant rearchitect of the environment.

Challenges

The average storage administrator is bombarded with a variety of messages, and it is challenging to stand out among all the noise. Novell has a well-recognized and well-respected brand associated with directory services, security, and operating platforms, but it is not well known for storage management.

Novell File Management Suite presents an opportunity for Novell to deliver a storage and data management solution and establish a brand in the storage market. This will require an investment in a go-to-market channel that will be able to articulate the value proposition of this solution and deliver services to help end users realize the benefits of such technology.

In addition, Novell must extend the value of NFMS to encompass emerging cloud-based storage services. While Dynamic File Services is able to utilize off-premises storage providers as second-tier storage, Novell must develop equivalent capability for Storage Manager and File Reporter to ensure policies are retained in the cloud. Companies must be able to take advantage of these new services to add layers of tiering and reach new locations without jeopardizing existing policies and assets.

Conclusion

The exponential growth in file-based storage presents an opportunity for solutions providers to design products with greater efficiency and more intelligent storage and data management to get at the heart of the problem this growth creates. Having intelligence on data stored can help storage managers make smarter decisions about data storage, protection, archiving, retention, and disposition. Managing storage based on predefined roles streamlines the process of managing data placement and storage.

Novell File Management Suite makes it easier for hard-pressed administrators to leverage the policy engine to automate many of the tasks associated with storage management, thus delivering lower cost of operating the environment with the confidence that the data stored possesses value and is relevant to the continuing operation of the business. Administrators can:

- Automate movement of data to lower-cost storage tiers and defer spending on high-cost storage
- Accelerate migrations of data from file servers or older NAS systems to more scalable and cost-effective NAS systems
- Reduce the time and resources set aside for file backup (improving information availability and enabling cost-effective disaster recovery for file-based information)

The key to attaining these benefits, of course, is effective implementation of the solution, including the adjustment of related processes (e.g., backup processes) and the leveraging of new efficient storage options such as data deduplication. A well-thought-out implementation (e.g., tiered storage) also makes it easier to expand use of NFMS. Finally, a sound implementation makes it easier for an enterprise to react to changing business conditions and new information needs. Working with a partner that understands the full capabilities of NFMS will ensure that rapid information growth doesn't translate into uncontrolled storage growth.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC. For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com