REPORT EXCERPT

PREVIEW

2017 Trends in Storage

NOV 2016

Henry Baltazar, Research Director, Storage

Tim Stammers, Senior Analyst, Storage

Steven Hill, Senior Analyst, Storage

Though storage is still a lucrative market, change is coming from disruptive technologies such as flash and object storage. Established vendors are being pressured to adapt to new customer requirements, including cloud computing. Provisioning speed and efficiency improvements are top of mind for many stakeholders.

ABOUT 451 RESEARCH

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to more than 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
20 West 37th Street
3rd Floor
New York, NY 10018
P 212-505-3030
F 212-505-2630

SAN FRANCISCO
140 Geary Street
9th Floor
San Francisco, CA 94108
P 415-989-1555
F 415-989-1558

LONDON
37-41 Gower Street
London, UK WC1E 6HH
P +44 (0)20 7299 7765
F +44 (0)20 7299 7799

BOSTON
75-101 Federal Street
5th Floor
Boston, MA 02110
P 617-261-0699
F 617-261-0688

ABOUT THE AUTHOR

HENRY BALTAZAR

RESEARCH DIRECTOR, STORAGE

Henry Baltazar is a Research Director for the Storage Channel at 451 Research. Henry returned to 451 Research after spending nearly three years at Forrester Research as a senior analyst serving Infrastructure & Operations Professionals and advising Forrester clients on datacenter infrastructure technologies. Henry has evaluated and tested storage hardware and software offerings for more than 15 years as an industry analyst and as a journalist.

Executive Summary

INTRODUCTION

To survive in the new reality of flat storage budgets and intense data growth, storage professionals must get more value out of every dollar they pour into their infrastructures. Not long ago, all-flash arrays (AFAs) were viewed as luxury items for IT datacenters and were earmarked for the most performance-sensitive workloads. With the steadily declining cost of NAND flash and ubiquitous deployment of efficiency-boosting deduplication and compression technologies, we are at a point where an organization cannot afford to pass up the superior performance, reliability and consistency of AFAs. Disk is far from dead, but if a company isn't leveraging flash, storage will likely remain the biggest bottleneck in its infrastructure.

Object storage is another sector that will move deeper into the mainstream IT landscape in the coming year. As organizations struggle with the requirements of hybrid cloud and long-term data management and distribution, metadata management and manipulation will be a major driver that pushes object storage out of the 'cheap and deep' archive space into more primary storage use cases.

Hyperconverged infrastructure (HCI) has blossomed into a legitimate infrastructure architecture over the last few years. It will also disrupt the storage landscape, since it provides a simplified means for delivering compute and storage resources to customers and service providers without requiring deep storage administrator expertise. We expect to see explosive growth in this market segment, though at this point it is still a relatively small part of IT infrastructure spending, and a substantial segment of customers do not have it in their near-term plans.

451 Research's 2017 Storage Trends

Source: 451 Research, 2016

	WINNERS	LOSERS
Organizations Must Modernize Storage To Do More With Less	Enterprise end users that focus on reducing management costs; vendors that provide flexibility; service providers that embrace hybrid cloud principles	Vendors that stick with the status quo; customers unwilling to make changes; service providers that rely too heavily on traditional storage players
HCI Will Disrupt Traditional Infrastructure	HCI vendors that embrace cloud; enterprises and midsized organizations that find ways to integrate HCI into their infrastructures	HCI vendors that don't address networking concerns; enterprise storage professionals that focus on derailing HCI projects
Flash Prices Will Continue Falling, Driving More Datacenter Use	Vendors that protect their installed bases; providers that move quickly to use all-flash storage to offer higher levels of application performance; enterprises that identify the applications that will deliver the biggest business benefit from a performance boost, and also consider the potential for reduced operational costs	Vendors that lose sight of the big picture; service providers that under-invest in infrastructure; enterprise IT organizations that move too slowly
Flash Disruption Will Enter a Second Phase	Vendors that enjoy the first-mover advantage; service providers that move quickly to host analytics and other applications that the new generation of AFAs target; enterprise IT organizations that get input from the businesses they serve	Vendors that are slow to embrace new technologies; service providers that fail to provide a range of performance tiers; enterprise IT organizations that don't establish a give-and-take of information with business units
IT Will Begin Recognizing Metadata as the Key to Long-Term Data Management for Production and DR/BC Platforms	Storage vendors that are most capable of addressing the challenges of data categorization and metadata creation; service providers that can simplify and accelerate the use of hybrid storage; enterprise customers that already integrate metadata as part of their storage infrastructures	Storage vendors that don't recognize the growing importance of metadata as part of a long-term data management strategy; service providers that don't explore the growing potential of cloud storage beyond basic archive and file sync/share capabilities; enterprise customers that address their data growth challenges by simply adding more storage capacity

METHODOLOGY

Reports such as this one represent a holistic perspective on key emerging markets in the enterprise IT space. These markets evolve quickly, though, so 451 Research offers additional services that provide critical marketplace updates. These updated reports and perspectives are presented on a daily basis via the company's core intelligence service, 451 Research Market Insight. Forward-looking M&A analysis and perspectives on strategic acquisitions and the liquidity environment for technology companies are also updated regularly via Market Insight, which is backed by the industry-leading 451 Research M&A KnowledgeBase.

Emerging technologies and markets are covered in 451 Research channels including Business Applications; Cloud Transformation; Data Platforms and Analytics; Datacenter Technologies; Development, DevOps and IT Ops; Enterprise Mobility; European Services; Information Security; Internet of Things; Mobile Telecom; Multi-Tenant Datacenters; Networking; Service Providers; Storage; and Systems and Software Infrastructure.

Beyond that, 451 Research has a robust set of quantitative insights covered in products such as Voice of the Enterprise, Voice of the Connected User Landscape, Cloud Price Index, Market Monitor, the M&A KnowledgeBase and the Datacenter KnowledgeBase.

All of these 451 Research services, which are accessible via the web, provide critical and timely analysis specifically focused on the business of enterprise IT innovation.

For more information about 451 Research, please go to: www.451research.com.

Table of Contents

TREND 5: IT WILL BEGIN RECOGNIZING METADATA AS THE KEY TO LONG-TERM DATA MANAGEMENT FOR PRODUCTION AND DR/BC PLATFORMS

Implication: *Businesses are rapidly becoming aware of the value of information that becomes increasingly invisible as it moves through traditional storage platforms, making the metadata-rich capabilities of modern object storage one of the only ways to provide extended data visibility and automated governance of company data across tiers, whether on-premises or in the cloud.*

Impact to the Market

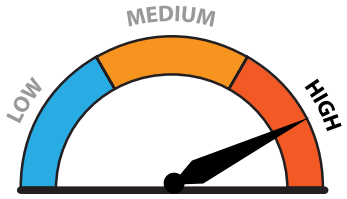

Since the inception of SAN and NAS technology more than two decades ago, the primary focus of the storage industry has been on centralizing data storage in purpose-built systems. The idea was, and remains, sound for a number of enterprise storage applications; however, these systems were built on block and/or file-based systems that relied on file naming with extensions, limited file attributes and a hierarchical directory system as the sole identifiers of the nature of the data itself. Unfortunately, this model still remains dependent on extremely limited metadata. It provides few if any reference points that storage administrators and their systems can use to identify, tag and programmatically manage the difficult-to-identify unstructured data that makes up the lion's share of enterprise data growth (see Figure 7).

Figure 7: Unstructured Data Is the Most Common Type in Enterprise Data Growth

Source: 451 Research's Q1 2016 Voice of the Enterprise: Storage

Q. Approximately how is your organization's total storage capacity, including primary and backup/archive storage, distributed across the following applications/workloads?

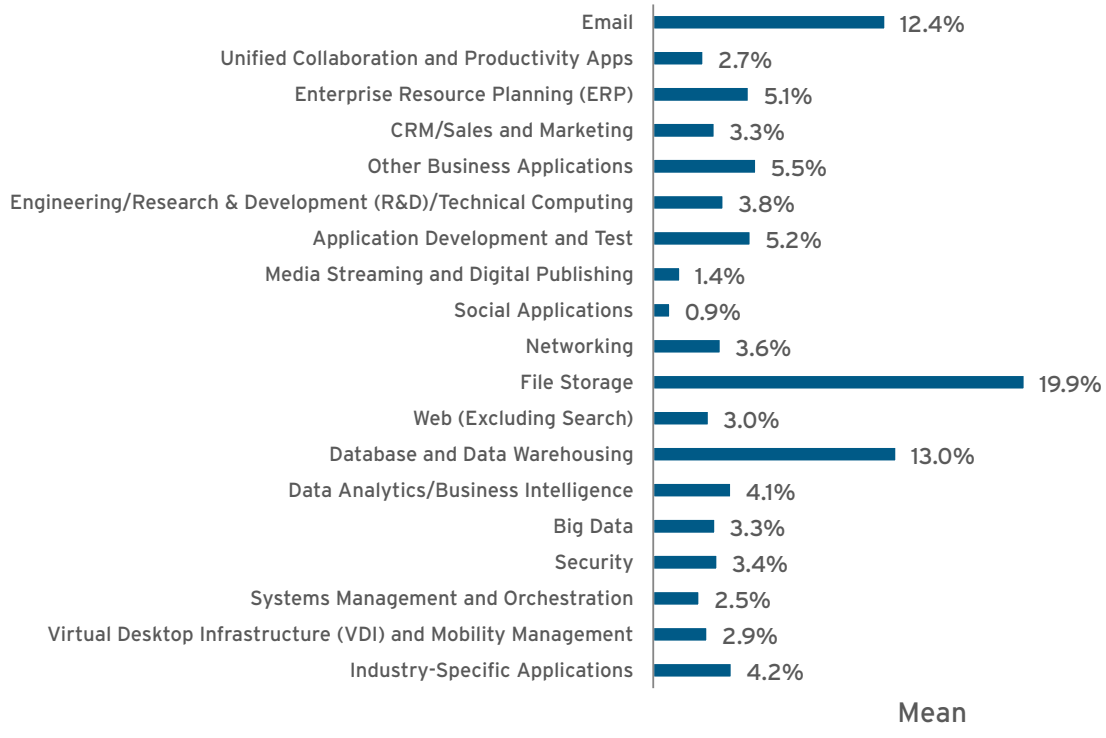

n = 721

Unstructured data consists of a broad group of data types, ranging from basic text documents to every form of audio, video and imaging data imaginable. Much of this media data is represented in the 'File Storage' category shown in Figure 7, but is increasingly being found in a number of other categories in the form of items such as email attachments, security videos and other work products that incorporate media. Media files are typically larger than text files and difficult to compress; they can also be extremely iterative, existing in multiple formats and versions with only nominal differences, if any.

This is not a new problem, and many of the applications for the media creation industry have been addressing the challenges of data management within the context of their own platforms. Media file types such as JPEG, MP3 and even Microsoft Office documents have embedded metadata as part of their file formats. Metadata is commonly used as an important part of the workflow platform for vertical applications such as video editing, where many versions of a work in progress may exist until the final product is put together. The problem with this form of embedded metadata lies in the fact that it doesn't extend to current block and file data management technologies: Even though much of the metadata can be displayed as details in a file's properties, file-based systems contain no mechanism to utilize that metadata for advanced data management purposes.

Conversely, object storage was specifically built around a metadata-based framework that supports data management programmatically, based on whatever extended metadata information may be attached to each individual data element. Along with basic information like filename, creation date and access control permissions, object data can be associated with any number of metadata fields. These fields in turn can specify a broad combination of global and/or granular data management policies. For example, a given piece of data can be automated to only be available in the US, backed up at several locations, targeted for extra-strong encryption, access-monitored as part of an audit trail for a legal department and then deleted on all systems on a given date seven years in the future. Even though this may seem like a relatively extreme example today, it only scratches the surface of what may be expected of storage systems in the future.

Aside from the obvious benefits of being able to automate highly granular, policy-based storage management, there is a growing interest in improving visibility into all of the data we have been collecting for years. In the current model, data moves in a path that goes hot, warm, cool and cold, and is usually tiered based on most recent access. This is the case because traditional storage architectures don't really take into account anything other than last access, and to be fair, this remains a relatively efficient model for generic tiering. But as data moves further and further down the traditional storage path, it becomes less and less visible, eventually ending up in an offsite cold storage tier – the functional equivalent of a dungeon where it may never see the light of day again. This is a case where 'out of sight, out of mind' can be counterproductive to business, because the increasing use of analytics is giving historical data a new lease on life. Good metadata can be the difference between looking something up on the internet and searching for a needle in a haystack.

There are a number of challenges in adopting metadata as part of a long-term data management strategy, with perhaps the greatest being the fact that most companies haven't even thought about collecting metadata. Building metadata after the fact can be difficult, especially for unstructured media files that don't lend themselves well to categorization. If metadata-based storage management takes off like we believe it will, there will be a growing interest in products and services that can best simplify the collection and categorization of post-creation metadata, as well as a new focus on building useful and functional metadata at the time of data creation.

THE EVOLUTION OF BACKUP, DISASTER RECOVERY AND BUSINESS CONTINUITY

In our VoTE Storage survey from Q4 2015, improvements to data backup and disaster recovery ranked as the top storage objective for 2016 (see Figure 8). Metadata plays a key role in the disaster recovery and business continuity markets due to the changing nature of data backup in general, in part due to the increasing use of cloud-based workloads but also because of the growing use of cloud-based storage platforms for backup and recovery purposes. While traditional backup and recovery remains an extremely important part of protecting business data, there is a growing need to better protect production workloads as well as their data for both on-premises and cloud-based applications. Metadata offers the ability to deliver granular recovery time objective (RTO) and recovery point objective (RPO) capabilities for disaster recovery and business continuity applications. It also provides the programmatic control of data needed to enable the automation, tracking, movement and synchronization necessary to identify and spin up interrupted workloads as well as restore access to critical data services regardless of physical location.

Figure 8: Improved Data Backup and Disaster Recovery Leads Storage Objectives

Source: 451 Research's Q4 2015 Voice of the Enterprise: Storage
Q. What is your organization's top storage objective for the next year?

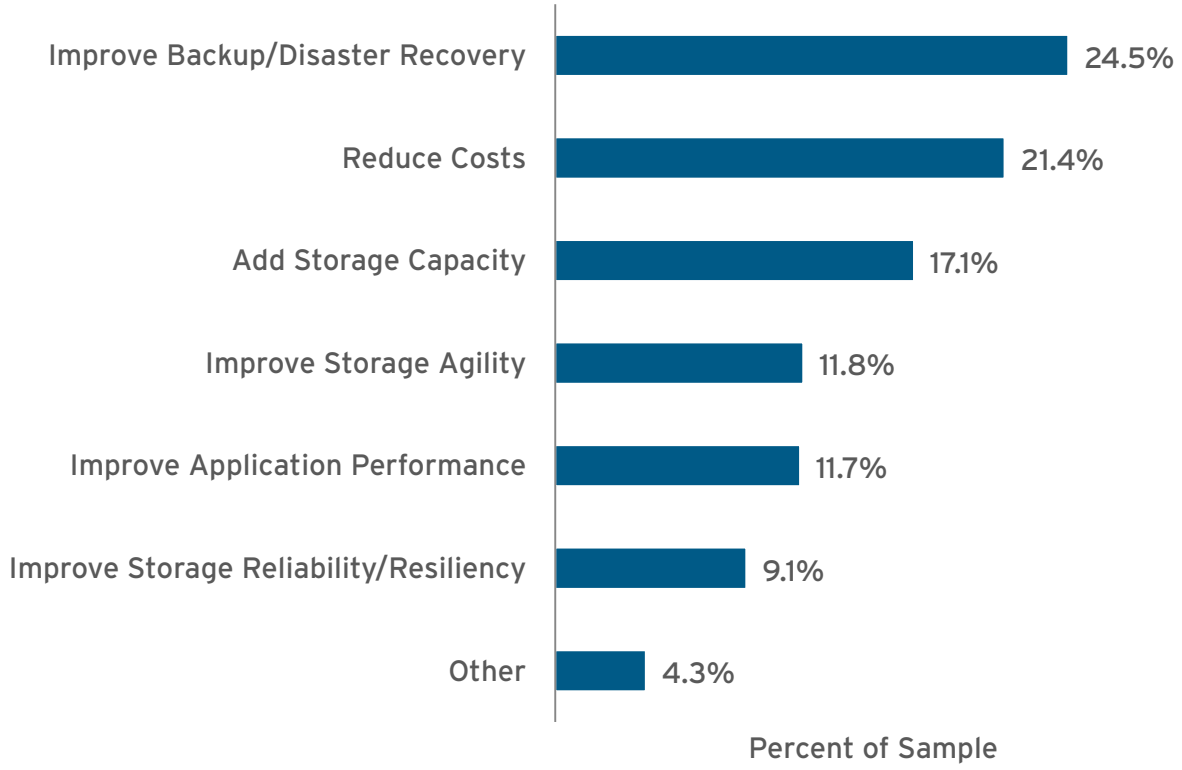

n = 625

As an added benefit, metadata already being generated as part of the backup process is playing a new role in better identifying and utilizing information that's already present. As data progresses down the traditional 'hot, warm, cool, cold' storage path, it can become practically invisible, especially if cold storage only consists of large, consolidated backup sets carefully stashed away in tape vaults. But a number of backup vendors have already been collecting metadata as data passes through the backup system as a necessary part of their platforms. Companies such as Veritas and Commvault have already started laying the groundwork for metadata services that actually start from their respective backup platforms and move back upstream to populate active archive applications and enable deep analytics of existing data. In addition, traditional storage vendors like NetApp are already speaking in terms of an 'information fabric' that utilizes metadata to provide end-to-end data visibility and management regardless of storage tier or location. We also believe that metadata will become increasingly important as IoT data starts entering the storage mainstream, in part because IoT looks to be an application capable of gathering massive quantities of data and one where the metadata about that information can be nearly as important as the data itself.

Traditional storage architectures without rich metadata capabilities will not be able to provide the scale and automation necessary for the massive and granular storage management needs of the future. Object-based storage has been around for decades, but has been pigeonholed as a slow, 'cheap and deep' storage market. However, we believe that beyond the 'traditional' object storage archive model, there is a renaissance underway that utilizes the rich metadata capabilities of object storage to provide the next generation of SDS-based distributed storage, not to mention massively scalable cloud storage platforms such as Amazon S3, Microsoft Azure and Google Unified Object Storage. This form of object technology is necessary to support the detailed and granular behind- the-scenes storage management capabilities necessary to provide the same degree of performance, data protection and management on commodity hardware that was expected of proprietary storage platforms, yet make it scalable to exabytes and beyond.

RECOMMENDATIONS

- **Storage vendors should provide a cohesive metadata strategy for long-term storage management** that better supports common policy management, informational visibility and automation for data regardless of where it resides. They should also provide assistance in helping customers determine a long-term vision for metadata creation that supports good data governance, industry-specific metadata customization and flexible hybrid cloud capabilities.
- **Public cloud storage service providers should look to support common metadata capabilities** that will allow customers to seamlessly integrate public storage services with on-premises storage capacity, supporting the same management tools and policy structures.
- **Enterprise customers should recognize the importance of metadata collection** and its role as part of a long-term data governance and management strategy capable of addressing the next-generation business, technical and legal expectations of corporate data.

WINNERS

- **Storage vendors that are most capable of addressing the challenges of data categorization and metadata creation** for existing data that provides unified management and visibility regardless of data format and across all forms of on- and off-premises storage environments.
- **Service providers that can simplify and accelerate the use of hybrid storage,** especially in the case of DR/BC applications where public cloud services are used as an alternative to on-premises secondary storage platforms for data, as well as production workload protection.
- **Enterprise customers that already integrate metadata as part of their storage infrastructures.** Metadata will allow organizations to more efficiently manage their data and workloads throughout their lifecycle and react to changes.

LOSERS

- **Storage vendors that don't recognize the growing importance of metadata as part of a long-term data management strategy.** The role and capabilities of object storage are evolving, and vendors who only look at object as a slow 'cheap and deep' archive market will be left behind.
- **Service providers that don't explore the growing potential of cloud storage beyond basic archive and file sync/share capabilities.** Poor and inefficient metadata and data management will prevent service providers from optimizing their services and will hamper their ability to plan for the future.
- **Enterprise customers that address their data growth challenges by simply adding more storage capacity.** The 'save everything' solution of increasing capacity without informational visibility and automated management is simply throwing good money after bad.