

Save Millions in Storage Costs With These Effective Data Management Best Practices

**FOUNDATIONAL****Refreshed:** 2 February 2016 | **Published:** 25 February 2015 | **ID:** G00269664

Analyst(s): Alan Dayley

Redundant, outdated and trivial (ROT) data is both a waste of IT budget and a potential security risk. Storage architects, similar to those at BAE Systems, need to understand that data reduction requires a combination of data analysis and corporatewide partnerships.



FOUNDATIONAL DOCUMENT

This research is reviewed periodically for accuracy. Last reviewed on **2 February 2016**.

Key Challenges

- According to a Gartner survey taken at the December 2014 Data Center Summit, managing data growth to support business demand is the No. 1 storage concern. ROT data consumes vast amounts of storage, resulting in unnecessarily high storage costs.
- Businesses often have a difficult time determining the value of data, making the decision to categorize, analyze, archive or delete data challenging.
- A general lack of visibility into data poses security and compliance risks.

Recommendations

- Use a file analysis (FA) product to uncover the enterprise data footprint as well as to ascertain the value of data.
- Build a business case for data reduction, focused on quick wins, and establish a data reduction taskforce that includes individuals from executive management, legal, compliance, records management and IT architecture.
- Develop data reduction projects and apply deadlines for decisions that put pressure on both the business and IT to meet data reduction goals.

- Review the results of data reduction projects to set new data reduction goals, and develop (or revise existing) data retention and deletion policies.

Introduction

We would like to thank BAE Systems that generously contributed insights and experiences to the research, including the storage team of BAE Systems Enterprise Shared Services. The contributor to our interviews and case study was Mark Tango (BAE Systems).

Rather than delete unneeded data, many enterprises simply buy more storage year-over-year, bringing added capital and operational costs to IT. Storage managers employ this tactic for two reasons:

- The cost, time and effort for reducing storage may appear to — or in actuality — outweigh the cost of adding more storage.
- In many organizations, acquiring authorization to dispose of data — or even move data to less expensive storage — is nearly impossible.

Despite these challenges, storage is a viable target for cost-reduction in many enterprises. According to "IT Key Metrics Data 2015: Key Infrastructure Measures: Storage Analysis: Multiyear," the annual cost per raw TB of storage in 2014 was \$2,520 (raw capacity is the amount of raw disk formatted for use). This cost includes hardware, software, personnel, connectivity, facilities/occupancy, disaster recovery and unallocated.

To keep storage costs under control, storage managers must effectively manage — and even delete — current data. They must also implement a data management policy to keep future storage growth and related costs at bay.

The five best practices in this research cover both data reduction and data management. These observations are derived from thousands of interactions with Gartner clients. One client, BAE Systems, is featured in this piece as its recent data reduction initiative was enormously successful. Setting out to reduce its enterprise data footprint from 5.5PB to 3.0PB, BAE Systems managed to, over the course of one year, analyze 2.0PB of its data across three data centers. From this work, the company was able to eliminate 881 TB of storage in both Tier 1 and Tier 2 and offer a cost avoidance savings of around \$1.7 million per year through the use of rightsizing overallocated database volumes identified by Active Navigation software. The company also saved capital expenditures by avoiding purchasing additional storage devices. The program also assisted BAE Systems in decommissioning approximately 75 storage units from various vendors, allowing the company to reduce additional support costs as well as avoid risks associated with aging storage devices.

Analysis

Uncover the Enterprise Data Footprint, and Prepare to Identify Its Value

Although storage managers might have a general sense that the opportunity for data reduction is prevalent across the enterprise, they do not always know how to find those opportunities. To uncover the enterprise data footprint, storage managers have, in many cases, tried using various storage resource management (SRM) products or search products. However, both of these product types have their limitations in that they are often constrained by the underlying application and only report basic metadata. In addition, these products produce reports that can be tainted by third-party applications. For example, virus scanners and backup products can change the last access date on files.

Gartner recommends using an FA product to analyze and create a visualization of the file system environment. FA products cost from around \$1,000 per TB for metadata analysis only to around \$5,000 per TB or more for content analysis, including full text indexing.

By using an FA product, a storage manager can analyze unstructured data at the metadata level. With insight into metadata, a storage manager can determine if data is orphaned, belonged to an ex-employee or is classified, for example. Allowing visualization of the data to the line of business (LOB) owners is instrumental in determining the value of data. Because of the linkage between identifying LOB owners of data and determining that data's value, storage managers should select an FA product that works with their directory service.

Using the FA capabilities of Active Navigation, BAE Systems crawled 1.5PB of data in two phases:

1. **File-level analysis** — BAE Systems began by looking for easy opportunities for data reduction by scanning for redundant, outdated and trivial (ROT) data. An example of trivial data would be the 4,000 copies of the invite to the company's 2007 annual picnic. In addition to ROT, BAE Systems also scanned for log files, pictures and disallowed files types, such as PST files and music files. Phase one did not, however, include integration with full records management, as the business was not yet ready to tackle this.

BAE Systems found that this first phase was critical, as the resulting analysis enabled legal counsel to create and approve policies for defensible deletion. In the past, when legal asked for information about the data, respondents provided Excel spreadsheets with thousands of lines of information. Taking this pain away, BAE Systems' FA product provides visualizations and drill-downs.

2. **Content-level analysis** — Although BAE Systems has not yet started this phase, phase two will involve scanning at the content level for sensitive data, including data subject to export control, data containing personally identifiable information and spillage, which is data that is stored on a network with a lower security classification than required.

Although BAE Systems set out to reduce its enterprise data footprint from 5.5PB to 3.0PB, the company chose not to analyze all 5.5PB of data. Instead, it began by analyzing specific data that added up to 2.0PB after determining that was a manageable workload for its data management

team. Moreover, while the team could have scanned all 5.5PB, the time to address those files would exceed the relevance of the collected data.

Get Executive Buy-In, and Establish a Data Reduction Taskforce

A data reduction initiative must start with executive buy-in, because it will result in the deletion of data. However, we still recommend performing FA before approaching the business. The FA product provides easy-to-digest, interactive visual graphs and data that help the business user make informed decisions. FA outputs the metrics needed for establishing an airtight business case.

To get executive buy-in, storage managers should develop a business case with the following:

- **A focus on quick wins.** This includes identifying content that is clearly out of policy and requires minimal user intervention to make decisions. Example quick-win content includes ROT, orphaned data and any storage that is allocated but unused.
- **Adequate quantitative information for executives to make a decision regarding beginning the data reduction initiative.** This means attributing a cost savings to a particular data reduction action.

Once the business case is created, storage managers must present it to those individuals who can approve the initiative, can make decisions about data deletion and retention, and understand the company's policies surrounding data deletion and retention. Typically, these individuals include:

- A legal and compliance manager.
- A records manager.
- The CIO or an appropriate representative from the office of the CIO.
- Two to three representatives from each LOB area that coordinate activities with LOB management.
- Enterprise architects.

Together, these individuals should form a data reduction taskforce. Storage managers can bring a project manager into this team to act as a liaison to the lines of business (LOBs) and track progress.

Through its data reduction initiative, BAE Systems found that the reporting capabilities of its FA product were, by far, the most important motivator in getting business leaders to act on data reduction. This was because BAE Systems was able to use the reports to show business leaders the potential cost savings associated with its data reduction projects, which included:

- Rightsizing its storage. This meant deleting data and maximizing the usage of storage solutions to enable much less expensive maintenance contracts with NetApp and Dell, and avoid the need to maintain the same year-over-year storage buying habits. Much of this was accomplished through the appropriate allocation of storage to each business unit.
- Repurposing aging storage arrays.

Once the storage team was able to display cost savings coupled with a clear picture of the data, the acceptance rate to assist increased by a factor of five.

Set Executive-Level Policies and Goals for the Data Reduction Initiative

Together, the individuals on the taskforce should start with selected data sources and set simple and straightforward policies for retention and deletion. Where feasible, these policies should be extensible across the organization as more projects are added. While it requires more work upfront, a single set of policies removes the cost and inefficiencies of managing specific policies that apply only to certain business groups, for example, and enables better data governance.

In creating a set of policies, storage managers will need to:

- Gain executive approval to own the data deletion process.
- Get legal approval for the defensible deletion of data.

Examples of BAE Systems' policies include:

- Data retention time frames associated with overall information life cycle management (ILM) policies.
- Quotas or folder size limits.
- Locking down the data so that personal mail folders cannot be created and data cannot be copied onto personal devices or to other locations, such as enterprise file sync and share (EFSS), unless there is a corporate-approved EFSS solution.
- A personal clean-up period before making folders read-only. This gives users a grace period of 90 days to clean or move needed data. After this point, the data can either be made read-only or deleted.
- BAE implemented a quarantine period. If the LOB, compliance or legal have concerns about data deletion or even deep archive, grant a grace period of, perhaps, one year in which data access is restricted to an IT request for retrieval. If no requests are made within the quarantine period, then the data can more comfortably be disposed of.
- Chargeback to the LOB. Chargeback rate can be based on either used storage or allocated storage. Note that by going by allocated storage, as BAE did, the LOB will feel greater fiscal pressure to evaluate its true storage needs.
- BAE implemented a rewards for success program for storage. Rewards must include compiling an overall cost-savings report per quarter for presentation to executive management. This allows the data manager to remain employed and gives management the confidence that they have the right person for the job. Rewards can be a redirection of funds to that manager for additional productivity products or additional personnel.

With a policy in place, the storage manager should then work to set goals for the data reduction initiative. This includes setting monthly goals per LOB and delegating tasks to the related LOB team for completion within a set time frame.

In general, data reduction goals fall into four categories. These are:

- **Infrastructure optimization and cost reduction.** This can be done through appropriate allocation of storage and chargeback.
- **Compliance adherence.** This is accomplished by retaining only required data for extended periods of time and disposing of nonregulated data in a much shorter time frame.
- **Risk mitigation.** This is important because as most corporate data ages, its value dramatically decreases while risk due to litigation exposure can increase.
- **Cost avoidance.** This is achieved through allocation, chargeback, maintenance cancellations, power reduction and a smaller data center footprint. BAE Systems chose, at this time, not to consider the cloud for aging data, but this should be a consideration for many organizations.

Develop Data Reduction Projects That Put the Onus on Both IT and the Business

Data reduction projects should initially focus on quick wins. This is the typical 80/20 rule — 80% of the benefits often can be derived by concentrating on the obvious 20% storage problems. To identify areas for data reduction, and to see if your file storage growth constitutes one of the biggest opportunities for improvement, begin with the metadata from the FA. By carrying out an initial file inventory and reporting on the findings, storage managers can determine the size and scale of the problem. This allows the business to prioritize areas that are likely to realize the quickest and largest cost savings.

Examples of quick wins at BAE Systems include:

- Removing redundant database and application volumes.
- Rightsizing volumes.
- Removing ROT.
- Setting chargebacks based on allocated storage, not on used storage, to control future data growth.

Importantly, although some quick-win projects will already have been identified in the business case, they are not necessarily approved. The CIO or an executive under the CIO ultimately must set and approve clear, measurable goals for the initiative. Acquiring this permission early in the process will eliminate ongoing requests to the C level for each disposal request.

To help ensure the success of data reduction projects, storage managers should tie their performance rating to the success of those projects. Policies, such as chargebacks to the LOBs and personal clean-up periods, as mentioned in the previous section, also can put pressure on the business to constantly evaluate its storage usage. If low-touch data is deemed nondisposable, it

can be moved to lower cost storage such as BAE Systems did by utilizing aging arrays as recycled units.

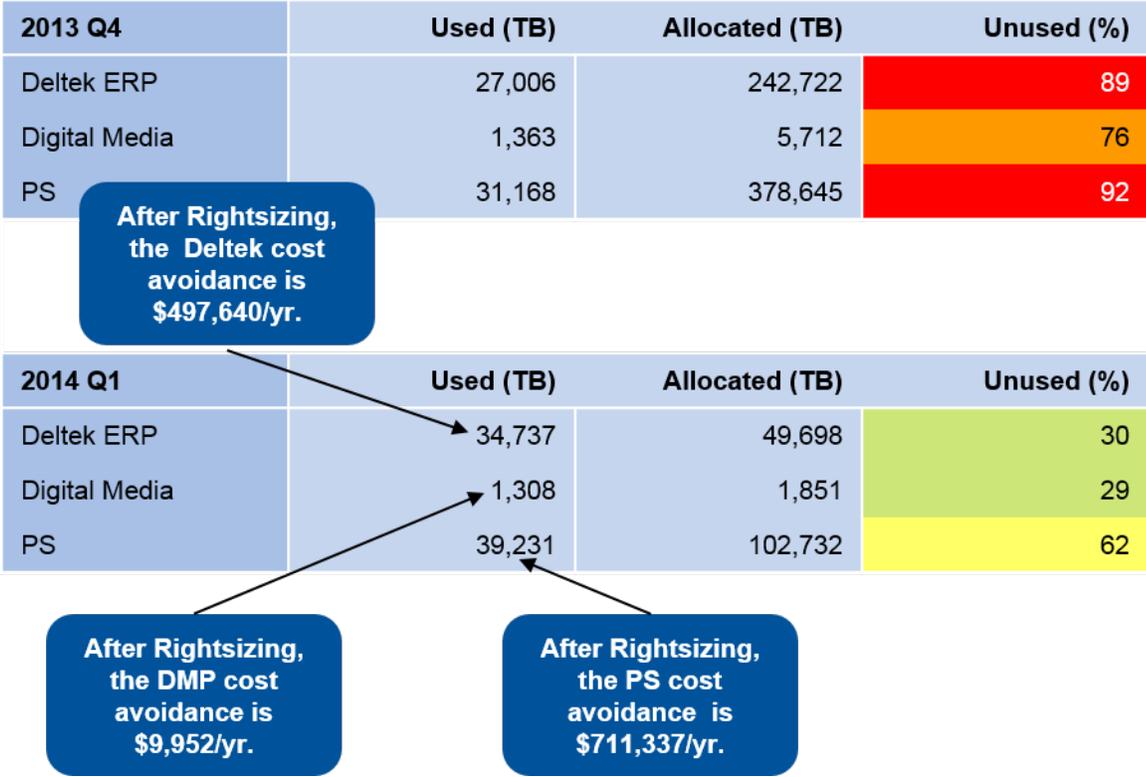
Measure the Results of the Data Reduction Projects, and Use the Results to Inform Future Projects and Data Management Policies

Storage managers must track the results of data reduction projects and regularly report the results to stakeholders. In addition to keeping stakeholders abreast of a project's status and success (or failings so corrective action can be taken), results can help:

- Lay groundwork for a broader data management practice.
- Set realistic expectations for future data management projects.
- Form new data management and data governance policies or revise existing policies.

To report on its storage reduction progress, BAE Systems developed a report to detail what was being stored, what could still be reduced, what had been reduced and the cost savings achieved (see Figure 1).

Figure 1. BAE Systems Report



Source: Gartner (February 2015)

Currently, the BAE Systems team presents this report to executive management on a monthly basis to determine and evaluate sectors lagging in the cost reduction effort. Incentives are available to those LOBs that meet or exceed goals, and pressure to improve performance around storage management is placed on teams that are falling behind. The net effect is that LOBs that have bought into storage management, and thus reduced costs, are enabled to have more resources available to do additional business-value-based projects as budgets are not reduced and the dollar savings can be repurposed.

Gartner Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"Market Guide for File Analysis Software"

"Hype Cycle for Storage Technologies, 2014"

"Does File Analysis Have a Role in Your Data Management Strategy?"

"Storage and Data Management Are Not the Same Thing"

"Selecting the Best Archival Storage Architecture for Your Needs"

"Best Practices for Data Retention and Policy Creation Will Lower Costs and Reduce Risks"

"Cool Vendors in Information Governance and MDM, 2014"

GARTNER HEADQUARTERS**Corporate Headquarters**

56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters

AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit <http://www.gartner.com/technology/about.jsp>

© 2015 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. or its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. If you are authorized to access this publication, your use of it is subject to the [Usage Guidelines for Gartner Services](#) posted on gartner.com. The information contained in this publication has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This publication consists of the opinions of Gartner's research organization and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "[Guiding Principles on Independence and Objectivity](#)."