

The hidden costs of storage

*Avoid unforeseen costs by thinking strategically
about data archiving and test data management*



1

Introduction

2

**The true costs
of storage**

3

**Thinking
strategically
about storage**

4

**Finding and
eliminating
hidden costs**

5

Resources

Introduction

Storage is cheap.

That's the myth we've heard many times over. Storage is a commodity, an inexpensive, quick, easy fix to demands for more space and more speed.

The problem is that this “bargain” isn't always the deal it appears to be. As enterprises wrestle with large volumes of fast-moving data—whether they are using analytics to generate business insights, responding to compliance requirements or driving key business processes—they risk being seduced by the siren call of low up-front costs and ignoring long-term, associated costs.

The direct cost of traditional disk storage is indeed lower, thanks to ongoing advances in capacity and technology. But the associated costs of storage—which include staff resources to manage the increasing amount of disk storage, as well as the space to house them and the energy to power them—can be sky high. And those costs are growing.

That's why—as tempting as it may be to just procure more disk storage—buying more of it is not a long-term solution for data growth. However, getting out of the “buy more storage” mindset requires thinking differently about storage.





The challenge goes beyond just handling exponential data growth. It's figuring out how to take advantage of the latest storage and archiving technology for the benefit of both users and IT. It requires giving users fast access to high-priority data, while allowing IT to move older, less-used data to more cost-efficient storage media. It also demands efficient management of the application infrastructure so data is always safe, secure and accessible when it is needed again. At the same time, IT departments have to make test data (minus confidential information)

available when developers are confirming new applications' capabilities. They must also be able to preserve data in accordance with retention and compliance requirements.

The overall goal: be more strategic about storage to accommodate both short-term data needs and long-term retention goals. Done right, making this strategic shift in thinking about storage can reduce costs, improve performance and reduce risks associated with storing data.

The true costs of storage

In a 2011 forecast, research firm IDC estimated that enterprises spent USD1.1 billion on storage in 2011, a number that is anticipated to grow to USD6.5 billion in 2015.¹ That's a staggering number, especially when you factor in the disparities between the direct and indirect costs. Based on IBM calculations, capital expenses for data storage account for only 20 percent of the cost.² The bulk of storage expenditures—the true costs—come from other operational factors (see Figure 1).

This means that, on average, for every dollar enterprises spend on storage, they spend four dollars on the operational elements of managing stored data. And that

may be a low figure; some analysts have calculated this figure as high as 10 times the amount spent on storage.

But it's important to remember that even beyond the capital and operational expenses of storage, there are potential business costs associated with not managing storage appropriately. Enterprise users rely on the ready accessibility of data for their productivity, which is why new data storage strategies are so crucial. Consider the idea of simply buying more storage; this approach presumes that all data is created equal, and it's not. Data has age, value and utility, and it should be stored according to these characteristics.

Industry average fully burdened IT infrastructure cost percentages to support data growth




Figure 1: Direct storage is a small part of the overall budget; operational costs represent the majority of spending.

Not all data is created equal

Enterprises store data for various reasons: it's used regularly, it will be used for testing, it might be used, or it must be kept for compliance requirements. In almost every case, a data archiving strategy can help companies improve application performance and employee productivity, streamline application migrations or ensure business continuity. Consider these scenarios:

- **Call center productivity.** How much data do call center agents need access to when dealing with the vast majority of customers? The last three months of transactions? The last six? With too much data to sift through, database performance suffers, extending the time it

takes for service agents to bring up pertinent data while the customer is waiting—from a few seconds to up to a minute or more. By offloading and archiving older data, call centers can reduce these query times. Older data is still accessible, but call center agents can help most customers faster.

- **Upgrade capability.** Some companies run into trouble when they keep years of production data in online storage—and when it's time to upgrade enterprise applications, enterprises must convert all of the data from the old application to work with the new one. If it's all online, that process can take up to a week—but virtually no company can afford to shut down data access for that long.

By archiving older and infrequently used data, IT departments can reduce their application upgrade windows significantly—allowing them, for example, to schedule application upgrades over a long weekend rather than disrupt the business for a full week.

- **Backup and disaster recovery.** The more data you have, the longer it takes to back it up (or restore, in the case of a business continuity issue). Global enterprises with operations in multiple time zones can't afford an eight-hour backup cycle; they must maintain system availability. That means ensuring that only mission-critical data is backed up, and the rest is archived on a regular basis.

Thinking strategically about storage

When you're at home, the stuff you use most frequently goes in the house, and things you need less frequently go in the garage. But when your garage is full to capacity with objects you rarely use, is it more prudent to buy a bigger house—or simply to find a small rental storage unit?

This logic works for data storage in the enterprise too. Buying more disk storage is like buying a new house to store your holiday decorations, rather than keeping those once-a-year items in an inexpensive offsite storage unit. Ask yourself: How often do you use each type of data? Why? How important is it to keep it handy, and how much does that cost?

Let's look at two strategies for storing data less expensively: archiving and making test data more efficient.

Archiving is for historical data (as opposed to backup, which mirrors production data). Archived data is like your holiday decorations—there's no reason to discard them, but they don't need to be immediately handy. Archived data can be compressed and stored on cheaper, lower-performance media, including WORM drives and tape, but it's still accessible for retention and governance needs. The data lives independently of the application, which supports long-term data retention and referential integrity.





Test data management involves transforming properly structured, relevant production data into right-sized test subsets. Before making this data available in test environments, organizations must protect sensitive information to meet compliance regulations through a process called masking. By using tools designed to automate the creation of right-sized, masked test data sets, organizations can replicate and store

only the test data they truly need—which further helps to reduce storage costs.

Enterprises that take a close look at these two aspects of their data storage strategies often find ways to dramatically increase cost-efficiency in their data centers. For example, rail transportation leader CSX Corporation found that its continued data growth was both degrading its application

performance and increasing its cost and capacity requirements for development and testing across both its mainframe and its distributed platforms. It launched an application test center to provide developers with reliable and repeatable test data, but this precipitated the need to change the way it managed test data. Previously, IT staff used manual “copy and clone” processes to create development and testing environments.

“Sometimes, we would need to deploy many test copies of the same database for different development purposes, each of which would be a full production copy,” says Gary Bachtell, a manager within Enterprise Services at CSX. “For example, the PeopleSoft database has grown to more than 300 GB, and we would have four or five test instances of the production database loaded into our test environment.”

These processes typically took hours to execute—more than 40 hours in some cases—which affected testing schedules. The large test databases also required costly storage capacity.

By implementing IBM® InfoSphere® Optim™ solutions for application testing and managing data growth, CSX improved testing efficiencies, improved the delivery of application enhancements, mitigated risk in nonproduction environments and brought storage costs down over three years.

Virginia Community College System (VCCS) had a more prosaic but equally demanding problem. Faced with the prospect of frequent and expensive server upgrades, it determined that it could archive more than 10 years of inactive student data while still complying with regulations and allowing staff to retrieve data when necessary.

“One of the problems in the past was that one team would consume the data that another team was counting on using. And because we only do a weekly refresh of the databases, the team then might lose development cycles because the test data isn’t available to them. By reducing the size of the databases, we can refresh those databases faster throughout the week, so that developers don’t have to wait.”

—Gary Bachtell,
Manager, Enterprise Services, CSX

VCCS first tried to address the issue by buying more storage, but the team couldn't keep up with the growth rate—and they forgot about the time it takes to implement and tune the database storage appropriately to impact performance. The college then implemented a policy-based archiving solution using IBM InfoSphere Optim solutions and was able to effectively manage data growth, improve service levels and enhance the flexibility with which it stored data.

Enterprises can derive significant benefits from being more strategic about storage. For instance, as CSX found, better strategies for test data management helped it easily refresh and maintain test

environments, as well as conduct tests faster with more accurate information. Boosting database performance boosts employee productivity.

As VCCS discovered, proper archiving helps to reduce total cost of ownership for databases and data warehouses by compressing historical data. With less data remaining on fast, expensive production-level storage systems, IT staff can manage these systems more efficiently and conduct backups more rapidly. Archiving dormant data helps to increase application performance and improve employee productivity—while still supporting data retention and compliance requirements.



Finding and eliminating hidden costs

Is your enterprise spending more time and money on storage than it should? Start by asking yourself these questions:

- How much time do you spend each week on hardware capacity-related performance issues?
- How long are you keeping data in your production systems? How long are you supposed to keep it—if at all?
- Are you using any shortcuts to keep your storage costs down, like using backup copies for data retention? That's a lot of data to sift through if you have an e-discovery request—not to mention a significant disaster recovery risk.

- Is your information trustworthy? Are you satisfied that errors and duplicates are being found and fixed, and that data privacy is protected wherever the data is being used?

Once you've answered these questions, IBM InfoSphere Optim solutions can help you find and eliminate hidden storage costs. Based on information collected in interviews with current InfoSphere Optim solutions customers, Forrester found that organizations can realize benefits in the form of improved operational and capital cost savings, improved IT and end-user efficiency, as well as higher levels of data protection and application performance.³



In addition to helping enterprises manage data growth and improve application efficiency, InfoSphere Optim software can help them prepare for application retirement and consolidation, manage the data lifecycle in big data environments and deliver data lifecycle management to cloud and virtual environments.

On the archiving side, InfoSphere Optim can help enterprises control rapidly growing data volumes and enhance application performance by archiving data, and capturing and storing historical data in its original business context.

This approach helps support long-term retention by archiving data in such a way that it can be accessed independently of the original application. No matter what access methods the applications use—such as ODBC/JDBC, XML, IBM InfoSphere Data Explorer or third-party reporting solutions—the data is still accessible when needed. In addition, InfoSphere Optim helps enterprises define and maintain data retention policies consistently across the enterprise, supporting not only packaged ERP applications but also custom applications across heterogeneous environments.

On the test data management side, InfoSphere Optim helps enterprise build realistic test data environments by automating the creation of right-sized, masked subsets of accurate data. This helps reduce the overall size of test environments so developers can run them quickly, yet with a high level of properly structured accuracy. Just as with the archiving capability, the software supports both custom and packaged ERP applications in heterogeneous environments, so that organizations can manage data no matter what kind of application it resides in.

Resources



The benefits of archiving and test data management strategies are palpable: better production database performance, more accurate retention of data, and better alignment between data and appropriate storage mechanisms. Reduce the associated costs of storage—the expensive ones—through better management, and you can dispense with myths of cheap storage that only limit your organization’s capabilities.

To learn more about strategies to address data growth and cut storage costs, check out these resources:

- [InfoSphere Optim Self-Service Business Value Assessment for Managing Data Growth: Calculate your ROI](#)
Answer five quick questions in this online assessment and get an unbiased third-party view of exactly how much IBM can help you save.
- [Whitepaper: Control application data growth before it controls your business](#)
- [InfoSphere Optim Data Lifecycle Management solutions](#)



© Copyright IBM Corporation 2013

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
March 2013

IBM, the IBM logo, ibm.com, Informix and Optim are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

¹ IDC. "Worldwide Archival Storage Solutions 2011–2015 Forecast: Archiving Needs Thrive in an Information-Thirsty World." October 2011. Doc # 230762.

² IBM Business Value Assessments.

³ Forrester Research. "Your Enterprise Data Archiving Strategy." February 2011. https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US&source=sw-infomgt&S_PKG=500006080&S_CMP=Optim_forrester_data_archiving_strategy_analyst



Please Recycle