

Enterprise Data Management Solutions
April 2008

IBM **Information Management** software



**Control application data growth
before it controls your business**

Contents

2 Executive summary

3 What drives application data growth?

7 How does data growth impact your business?

9 How can you take control of application data growth?

12 IBM Optim delivers proven enterprise data management capabilities

17 Managing data growth – real world examples

Executive summary

Businesses rely on ERP and CRM applications, which collect information from a variety of sources to support business operations. While this information represents a valuable asset, like any other asset it must be well managed if it is to continue to deliver business value. Without an effective strategy for managing enterprise application data, your organization remains vulnerable to the negative effects of runaway data growth.

Typically, most data growth is organic. That is, as your business grows, so does the amount of application data that is collected, managed and stored. Mergers and acquisitions offer another opportunity for organizations to double or triple the amount of data they possess, as two organizations become one. Data retention requirements add to the capacity constraints as more data must be retained for specific periods of time. Lastly, the “data multiplier” effect accounts for data duplication across an enterprise that significantly increases data growth statistics.

So, how can data growth impact your business? Accumulating vast amounts of application data adds complexity and risk, as well as storage and management costs. Next, managing years of historical data in production databases can impact service levels and disaster recovery initiatives. Expanding online and batch processing windows, as well as routine maintenance tasks, take much more time and can reduce application availability significantly.

Your enterprise applications simply must deliver measurable business value. So how can you manage data growth and capitalize on your investment? This white paper describes some of the ways that organizations can address data growth issues through the power of enterprise data management. Companies need capabilities for assessing

data growth and implementing tiered storage strategies. It is important to classify enterprise data and define service levels. As a recognized best practice, database archiving has proven effective in managing continued application data growth. Look for archiving capabilities that promote a variety of access and storage options and support your data retention requirements.

In particular, the IBM® Optim™ Data Growth Solution offers proven capabilities for managing enterprise data cost-effectively throughout the data lifecycle. With Optim's capabilities to align enterprise data management with your business objectives, your organization will be better positioned to improve service levels, reduce risk and control costs. Enterprise data management helps support your information governance initiatives, simplify your IT infrastructure, support business continuity and increase the business value of your enterprise applications.

What drives application data growth?

Companies across industries worldwide recognize the value of enterprise application data as a critical business asset. But many factors cause enterprise application data to accumulate at a rapid rate. Without an effective strategy for managing enterprise data, your organization remains vulnerable to the negative effects of runaway data growth.

And this growth is expected to continue unabated. With analyst projections forecasting annual compound growth rates for databases at an excess of 125 percent, it is clear that organizations need effective strategies to mitigate the risks associated with managing and storing increasing volumes of enterprise application data. But what are some of the underlying factors that drive application data growth?

Organic business growth. As a natural part of daily business activities, the amount of data in your mission-critical application databases increases over time. It is not unrealistic to assume that in only a few years, a company's eCommerce, CRM and ERP applications will collect millions of new transactions per day! However, as time goes on, the business value of the historical data will decline, creating a backlog of historical information that requires orderly data management and storage policies.

Mergers and acquisitions. Mergers and acquisitions also drive data growth. Businesses expand through merging with other companies or acquiring their competitors. As a result, the IT organization must support the acquired applications – and all associated enterprise data. For industries such as insurance, banking and utilities, which tend to expand through mergers and acquisitions, this type of data growth is commonplace.

Data Multiplier Effect. Data duplication contributes to growth statistics. The challenge of managing large volumes of data is compounded by cloning, or copying, the production database to support various other functions within the organization. It is not uncommon for organizations to maintain several backup copies of critical data or to implement mirrored databases that provide assurance against data loss. Disaster recovery plans often require data duplication in order to store critical data in an alternate location. It is also common to clone an entire production database for use in application development and testing initiatives. What results is known as the “Data Multiplier Effect.”

As data is duplicated, storage and maintenance costs increase proportionally. The Data Multiplier Effect is the result of multiplying every gigabyte (GB) of data in

a production database by the number of replicated copies. The resulting figure represents the company's total data burden – and it is usually a higher number than managers expect.

For example, suppose that a modestly sized production database contains 200 GB of data (see Figure 1). When the production database is copied for backup, disaster recovery, testing, development and quality control, the total data burden increases from 200 GB to 1200 GB, or 1.2 terabytes (TB). The Data Multiplier Effect increases costs and limits the amount of space available for further growth. Unless companies can find ways to streamline the source copy of the data, and control the duplication, this multiplicity of costs will become prohibitive.

Actual Data Burden = Size of production database + all replicated clones




Figure 1. The Data Multiplier Effect is the result of multiplying every gigabyte (GB) of data in a production environment by the number of replicated copies.

Data retention regulations. Compounding the data growth challenges are global data retention regulations that require the management and storage of different types of data for extended periods of time.

Corporate policy or government regulations may require that data remains easily accessible for specified periods of time and sometimes specify disposal requirements. Some of these laws are intended to preserve the accuracy and transparency of business records to prevent accounting scandals. Companies may need to protect corporate interests by retrieving historical records to satisfy audit inquiries and resolve claims. Retention requirements carry strict penalties for non-compliance, driving the demand for cost-effective data management and storage solutions.

In the United States alone, there are over 10,000 federal, state and local regulations that govern the preservation of data. For example, HIPAA (the Health Insurance Portability and Accountability Act) requires that healthcare organizations retain original medical records for a minimum of five years and in most cases, decades. The Sarbanes-Oxley Act requires that corporate accountants retain specific records for a minimum of five years. SEC Rule 17a-3 dictates that documents and records must be retained for three to six years, and during the first two years, these records must be stored in an easily accessible place.

By the end of 2006, all global banks had to comply with the requirements of Basel II. Depending on the type of business, national regulatory decisions and risk parameters, banks must retain historical financial data for at least two to seven years. Various other regulations seek to improve the control of obtaining and storing information. Some examples include the Data Protection Act 1998 (UK), the

Anti-Terrorism, Crime and Security Act 2001 (UK), and the Personal Information Protection and Electronic Documents Act (Canada).

These regulations are designed to specify the appropriate standards of care for managing and retaining corporate information. Organizations are required to maintain critical data, keep it accessible and satisfy the demand for documented authenticity and accuracy of their records. It is also important to dispose of data appropriately when the retention period expires. Without the capability to classify, archive, store and ultimately dispose of historical enterprise information, most data is retained in the primary database, increasing the problems associated with long-term data growth.

How does data growth impact your business?

Transaction-intensive, customer-facing applications are collecting and storing more data than ever before. Because you rely on these applications to direct every facet of your business operations, managing continuous database growth is absolutely necessary for controlling costs, improving customer satisfaction and enhancing decision support.

IT professionals are challenged to manage accelerating data growth, while striving to meet service level agreements and maintaining application databases to maintain optimal performance and availability. Runaway data growth can result in a number of outcomes that can be detrimental to your business operations, your profitability and your ability to maintain a competitive advantage.

The potential consequences of continued data growth can affect all areas of your enterprise, making it difficult for you to complete critical business processes

and meet your business objectives. Without an effective strategy for managing enterprise application data, it becomes more difficult to address critical business issues, including meeting information governance requirements, controlling IT infrastructure costs and ensuring business continuity.

Increased IT complexity, risks and costs. Accumulating vast amounts of application data adds complexity, risk and cost to your business. Managing more data means that transaction processing is slower and routine maintenance takes longer to complete. An overburdened IT infrastructure places a strain on capacity, slows operations and jeopardizes customer satisfaction. Boosting processing capacity and tuning the database can provide some relief, but only temporarily.

The cost of storing your data also increases substantially. The Data Multiplier Effect is in full force. Not only must you expand storage capacity for transaction processing, but you must also increase capacity across all cloned environments. Purchasing additional storage leaves many companies in a perpetual game of catch-up. An IT organization can spend millions of dollars annually for additional storage hardware. Still, this short-term, tactical approach does not address the root cause of the problem – rapid data growth.

Diminished service levels. Maintaining databases that contain years of historical data slows response time and impedes access to current information. With too much information to sift through, routine reports take longer to run. Business analysis queries also require more time to complete, limiting your ability to make accurate, timely decisions.

Continued data growth impacts functional processes, such as financial period closeouts or delivery scheduling. Service levels decline, eroding the customer loyalty

you worked so hard to develop. Backup, reorganization and recovery windows are stretched to the point where system availability is seriously threatened. Upgrade and migration projects are more likely to cause costly business disruptions. Without a method for managing ongoing data growth, it becomes increasingly difficult to achieve your business objectives.

Inadequate disaster recovery. In today's environment, disaster recovery is a high priority. In the event of a disaster, the key strategy is to get mission-critical systems operational as quickly as possible. With overloaded databases, all the data (both current and historical) must be restored, just to resume processing for today's transactions. The need to restore historical information simultaneously with current transactions can slow the recovery process by hours or even days.

How can you take control of application data growth?

Your enterprise applications simply must deliver measurable business value. So, how can you manage data growth and capitalize on your investment? The answer rests in your ability to align continuous control of your application data with your business objectives through the power of enterprise data management.

The IBM Optim Data Growth Solution is a single, scalable enterprise data management solution designed to meet your evolving business needs. From small to large organizations, from single applications to global business centers, data management is streamlined using a consistent proven strategy.

With Optim, you can simplify enterprise data management to accelerate business-critical projects. Provide open access to current and archived data, complete easier upgrades, implement cost-effective tiered storage strategies, and profit from superior application performance and availability. With capabilities that allow you to manage

application data growth across your enterprise, you can improve application service levels, mitigate risk and control costs.

The following capabilities are necessary for effective enterprise data management:

- **Assess application data growth and tiered storage strategies.** With proven capabilities for assessing your data, you can easily identify the applications where data is accumulating the fastest. Gaining a complete understanding of which areas are accumulating the most information allows you to apply the most effective storage strategy. Assessing your data allows you to be proactive with enterprise data management, enabling you to identify and address potential problems before they impair business results.
- **Classify data and define service levels.** Once you have assessed your application data, it is useful to classify it according to its business value. Classifying each type of data record or business object allows you to define the appropriate level of performance. For example, current transactions are the highest priority, often requiring sub-second throughput. In contrast, reporting data, used mainly to generate management information reports, has a lower service requirement. It can be safely moved to a somewhat slower, but less expensive environment. Finally, inactive data that is being retained for compliance purposes can be maintained in an offline environment.
- **Initiate archiving as a best practice.** Archiving capabilities are an essential part of an enterprise data management solution for controlling data growth. Database archiving allows you to segregate and safely remove historical data from the production environment, freeing valuable capacity for priority business needs.

Implementing a policy-driven archiving strategy, based on the age and status of the data, enables you to manage each class of data according to its unique service requirements.

- **Store enterprise data cost-effectively.** Storing archived data according to its evolving business value is a logical component of an enterprise data management strategy. Continuing the three-tier classification example, current transactions are maintained in high-speed, primary storage. Reporting data is relocated to mid-tier storage. Reference data is retained on a secure WORM (Write Once, Read Many) device, keeping it available if an audit request should arise. Implementing a tiered storage strategy allows you to reclaim capacity and maximize the value of your existing storage infrastructure.
- **Promote data accessibility.** Decision makers must have access to data, whether it is current or historical. A comprehensive enterprise data management solution will allow those decision makers to access the right information at the right time. Authorized business users must be able to query and browse all active, inactive and reference data. Reliable access makes it possible to generate reports and respond quickly to audit and discovery requests. If additional business processing becomes a requirement, you must be able to restore archived transactions.
- **Manage data disposal.** It is imperative that organizations control the disposal process to help ensure that data is eliminated at the end of the required retention period. Effective enterprise data management makes it possible to automate both retention and disposal, while still offering the capability to manually remove data. Applying suitable and secure methods for data disposal allows you to prevent your information assets from becoming liabilities.

IBM Optim delivers proven enterprise data management capabilities

You invest millions of dollars in your enterprise applications and the supporting infrastructure to promote optimal operating performance, improve decision-making and gain a competitive advantage. IBM Optim provides the power of enterprise data management, so you can derive the most business value across your enterprise.

Proven enterprise data management capabilities allow you to optimize the performance of your mission-critical applications and manage database growth. By streamlining your application databases, Optim enables you to provide exceptional service to business units and profit from the superior performance and availability. Key business processes that are imperative to your organization are completed on time. And with continuous access to your data, decision makers get the information they need, when and how they need it.

Optim also helps organizations control the costs of data growth management (see Figure 2). By streamlining critical databases and simplifying complex infrastructures, you can leverage your existing IT investments. Key IT processes require fewer resources. With Optim, you are able to reclaim storage capacity and take advantage of a cost-effective data growth management strategy.

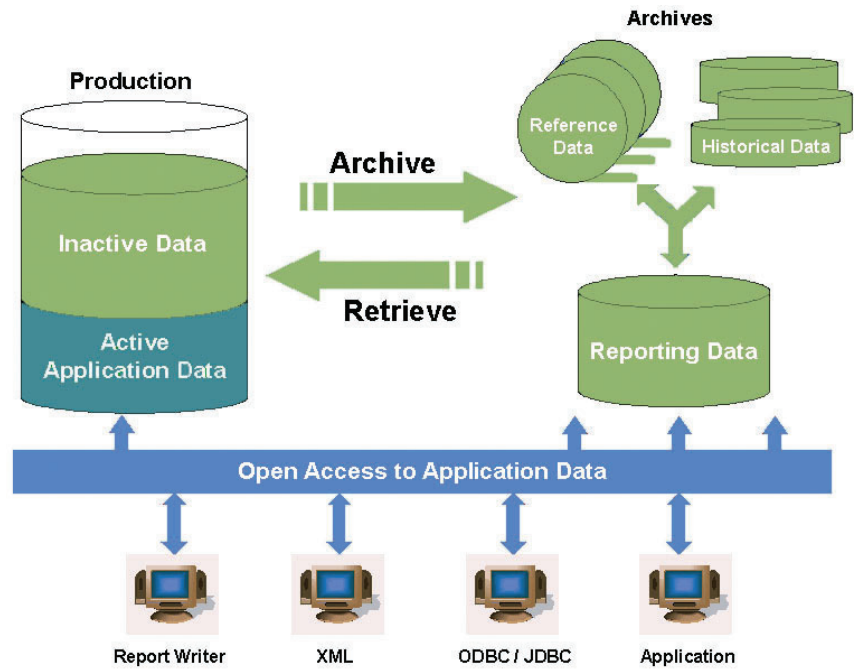


Figure 2. Optim lets you segregate inactive application data from current activity, safely remove it to a secure archive and keep it accessible.

Protecting your company from legal and other liability is critical. Optim enables you to mitigate the risks encountered in business today. Optim’s data management capabilities allow you apply functional policies to govern data retention. You can automate data retention to support compliance initiatives and respond quickly and accurately to audit and discovery requests. And in the event of a disaster, employing a staged recovery strategy helps ensure the continuity of your business.

Optim is a single, scalable solution that supports all major enterprise databases and operating systems – IBM DB2®, Oracle®, Sybase®, Microsoft® SQL Server®, IBM Informix®, IBM IMS™, IBM VSAM®, Microsoft Windows®, UNIX®, Linux® and IBM z/OS®. It also supports the key business applications in use today – Oracle® E-business Suite, PeopleSoft® Enterprise, JD Edwards® EnterpriseOne, Siebel®, Amdocs® CRM and all of your custom and packaged applications.

Manage enterprise data throughout its lifecycle. Optim provides capabilities for controlling your application data from creation to disposal. By implementing a proven enterprise data management strategy, you take command of your mission-critical data throughout its entire lifecycle and realize measurable benefits across your organization.

- **Archive.** Database archiving is a recognized best practice. Once separated from current activity, inactive application data is safely moved to a secure archive. Optim manages your data at the transaction, or business object level – such as service orders, payments or shipments – allowing you to define the policies and criteria for moving these transactions into an archive. Ongoing archiving allows you to manage continued data growth, simplify maintenance, speed disaster recovery and ultimately reverse the Data Multiplier Effect.
- **Store.** With Optim, you determine the appropriate storage location for each class of application data, based on its business value and access requirements. Utilizing a tiered storage strategy, you reclaim capacity by storing only current transactions in the high-performance production environment. Your reference and reporting data can be stored safely in nearline or offline storage. Tiered storage strategies organize your data cost-effectively, based on its business value, and help you manage your rapid data growth.

- **Access.** Optim delivers access to the information you need, when and how you need it. With capabilities to query, browse and generate standard or custom reports, you can respond quickly and accurately to audit or discovery requests. Optim supports a broad range of methods to access archived data. Application-based access offers a consolidated view of current and historical information through the existing application interface. “Self-help” data access permits end-users to leverage existing skills and tools (such as Crystal Reports™ or Microsoft® Excel®) with no additional training or assistance from the IT group. Application-independent access utilizes industry standards methods, such as ODBC/JDBC, XML or SQL, and reporting tools, such as Cognos® or Business Objects™, to access archived business transactions without impairing online transaction processing (OLTP) performance.
- **Dispose.** A comprehensive retention compliance strategy must include a means for disposal. Optim allows you to control and automate data disposal when retention periods expire. You minimize the risks associated with retaining records longer than the stipulated time period. Whether you choose to automate the disposal process for increased efficiency or manually select archived transactions for disposal, Optim enables you to retain only necessary information in primary databases and archives.

Align data management with your business objectives. Your ability to meet your business objectives is critical. Optim offers measurable benefit across the enterprise, enabling you to achieve your goals:

- **Satisfy information governance requirements.** Companies that deploy effective enterprise data management strategies for Information Governance can achieve significant benefits. Corporations must protect all sensitive data, no

matter where it is stored (online in an archive database, near-line on a file server, offline to tape, disk-based WORM device or long-term permanent storage). Adapt and respect security measures present in your existing database, application or network. Control access to data processing functions, objects and data from within Optim. Set and automate retention policies. Optim also maintains easy access to current and archived data allowing you to mitigate the risk of non-compliance with data retention regulations and respond quickly to audit and discovery requests.

- **Simplify your IT infrastructure.** A complicated infrastructure is risky for your business. Optim's enterprise data management capabilities allow you to simplify your IT infrastructure and make the best use of your resources. Approach the challenges of data growth management proactively, through routine archiving to segregate historical from current data and maintaining databases at a manageable size. The consequences of the Data Multiplier Effect are reversed, mitigating the data growth rate and providing ongoing control of your expanding storage platforms and spending. Implementing lifecycle storage strategies enable you to optimize utilization and defer the costs of high-speed resources, while still maintaining data access. Database archiving capabilities also provide the means to decommission or retire redundant and obsolete applications, allowing you to reduce infrastructure complexity and reclaim IT assets.
- **Support business continuity initiatives.** Enterprise data management strategies allow you to minimize the downtime associated with routine business operations and help you facilitate maintenance and backup. Controlling data growth allows you to reduce the possibility of batch overruns, lengthy backup windows or other impediments in meeting service level agreements. Your data

is available when you need it. With a tiered recovery strategy, active information from the production database is recovered first, enabling operations to resume as quickly as possible in the event of a disaster. Reporting and reference data can be recovered later, without impacting the continuity of your business. Optim provides ongoing control for managing data growth so your applications and databases remain operational when you need them.

- **Increase the business value of your enterprise applications.** Optim empowers you to manage, retain and control mission-critical application data so you can achieve the greatest business value. By aligning application data management with your business objectives, you consistently reach performance targets, provide superior service and promote customer satisfaction. Classify data based on its business value and access requirements, and apply the appropriate level of resources to achieve desired service levels. Optim helps you organize and manage your application data so that you can provide decision makers with timely access to the data they need. Reporting on historical information takes less time and effort with Optim's broad range of access methods; from application-based access (through your current interface) to self-help access and application independent access, leveraging industry standard methods and tools.

Managing data growth – real world examples

Many clients have realized measurable benefits from implementing Optim to effectively manage enterprise data. For one client in the retail industry, data growth was negatively affecting online response time for its business-critical retail management application.

Business users and customers alike were experiencing slow response time, while expanding windows for batch processing, backup, and recovery processes were interfering with open-for-business hours. These issues were having a serious impact on the success of the business. With Optim, this client was able to archive historical application data and streamline its production database. As a result, a significant amount of high-priced storage capacity was reclaimed and application response time improved by more than 50 percent.

Similarly, a client in the telecommunications industry felt the effects of rapid database growth on its ability to complete marketing campaigns and other business initiatives. Compliance with data retention regulations resulted in a complex production database containing years of historical data as well as current transactions. As business expanded, the client began to experience increasing storage requirements and slowing response times. Continual maintenance and tuning provided only temporary relief. After implementing Optim and gaining control of rapid data growth, the company increased its ability to launch frequent and aggressive marketing campaigns to generate new business opportunities, while continuing to provide superior customer service.

Another telecommunications client was seeking to resolve the issue of database overload in its mission-critical CRM application. With tens of thousands of new cases being added daily, and a backlog of over one million closed cases in their primary database, it was crucial to implement a solution that would help manage this growth. Optim enabled the client to archive several million closed cases per year, storing the data on a more cost-effective medium. The performance, availability and reliability of its customer-facing application greatly improved, and the client was able to meet and exceed rigorous service level agreements.

About IBM Optim

IBM® Optim™ enterprise data management solutions focus on critical business issues, such as data growth management, data privacy compliance, test data management, e-discovery, application upgrades, migrations and retirements. Optim aligns application data management with business objectives to help optimize performance, mitigate risk and control costs, while delivering capabilities that scale across enterprise applications, databases and platforms. Today, Optim helps companies across industries worldwide capitalize on the business value of their enterprise applications and databases, with the power to manage enterprise application data through every stage of its lifecycle.

For more information

To learn more about IBM Optim enterprise data management solutions, contact your IBM sales representative or visit: www.optimsolution.com.



© Copyright IBM Corporation 2008

IBM Software Group
111 Campus Drive
Princeton, NJ
USA, 08540-6400
www.optimsolution.com

Produced in the United States of America
04-08
All Rights Reserved.

DB2, IBM, the IBM logo, IMS, Informix, Optim, VSAM and z/OS are trademarks or registered trademarks of the IBM Corporation in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. UNIX is a registered trademark of The Open Group in the United States and other countries. Excel, SQL Server and Windows are registered trademarks of Microsoft Corporation in the United States and other countries.

All other company or product names are trademarks or registered trademarks of their respective owners.

References in this publication to IBM products, programs or services do not imply that IBM intends to make them available in all countries in which IBM operates or does business.