



Tiered information infrastructure: A practical approach to translating strategy into implementation

by Steve Correl

Contents

3 Introduction

5 Practical use cases

- 5 Classify applications according to their information infrastructure requirements
- 6 Understand file system and storage usage to identify opportunities
- 8 Exploit application-aware support for data movement
- 10 Virtualize disk storage to move data seamlessly, improve utilization and reduce management 'touch'
- 12 Deploy a compliance/retention solution that combines disk and tape to utilize lower tiers
- 14 Utilize Content Management to facilitate data movement
- 15 Exploit outside services to help assess maturity, make recommendations and guide deployments

17 Business case

- 18 Costs
- 19 Benefits
- 21 Bottom line

22 IBM capabilities

Introduction

Today, information is created faster than ever before. In some business segments, the information rate is merely multiplying, but in others the increase is literally exponential. In the face of this tidal wave of information, IT is tasked with enabling the lines of business to leverage information, to act quickly, to unlock creativity and innovate, and furthermore to do all these things in a cost-effective manner with control over service levels and risk. In fact, year-to-year IT budgets are not generally increasing, and in some cases they are contracting.

Information lifecycle management (ILM) is often proposed as a strategic response to control cost. In IBM's view, it has a role to play in leveraging information for greater business value as well. However, ILM is multi-faceted and can seem complex to deploy. In working with clients across industries and around the world, IBM is often asked where to start. To help answer this question, we developed the tiered information infrastructure approach, which pulls together straightforward ILM tactics that also make a lot of sense from the business perspective. Based on our experience with clients, we believe this approach comprises achievable steps on the way to achieving the broader strategic goal of implementing ILM.

The tiered information infrastructure (TII) approach responds to several present-day realities:

- *Information is being created at an exponentially increasing rate, and the wide variety of information requires that IT deliver widely differing service levels.*
- *The value of the information to the business shifts over time, sometimes within days or weeks of being created.*
- *The per-gigabyte cost of new storage consistently declines due to technological progress.*
- *While the cost of new storage declines, the overall cost of managing information is increasing. Labor cost increases are driven by the need to maintain service levels despite the surging data volume.*

In order to keep up with these changes over time, IT is continually caught between the pressure to cut spending to stay competitive and the risk of under-servicing and exposing the business to various risks. The risks of under-servicing can include the inability to recover from outages or surges in demand. Organizations are also exposed to audit risks and the costs associated with legal discovery efforts.

A tiered information infrastructure embodies the implementation of a business and IT assessment and analysis. This includes the solutions and people processes that enable organizations to continuously adapt their infrastructure usage to the most cost-effective equipment available. This approach can help organizations:

- *Integrate new purchases as well as leverage existing equipment.*
- *Rebalance resources across operations based on business need and required service levels.*
- *Reduce overall storage labor through improved storage management toolsets.*

We have seen some businesses invest in new storage tiers, yet fail to realize the benefits expected over time. They migrated data to new storage but later ended up with a larger volume of misaligned data. Given the realities described above, we believe an investment in tools and processes is needed to deliver a sustained return on the investment in extra capacity. The IBM TII approach is designed to facilitate the agility necessary to enable businesses to adapt as new capabilities are deployed; to optimize the infrastructure; and to continuously right-tier service levels as new challenges evolve.

This paper outlines practical methods of achieving a tiered information infrastructure along with their business case justification. It covers ways of utilizing a range of tiers to control costs and increase flexibility. We illustrate how a relatively small investment can save hundreds of thousands of dollars even as the amount of managed information continues to grow. We also describe how a TII will synergize with related strategic needs.

Practical use cases

The following sections describe practices we have seen succeed in IT organizations around the world and in a variety of industries. We view them as “best practices” because they appear in multiple successful implementations, although of course they won’t apply in every situation. Depending on the maturity of its IT infrastructure, a business may have implemented only one of these practices, or it may have implemented numerous practices as part of a comprehensive strategic vision and plan.

Classify applications according to their information infrastructure requirements

In order to identify opportunities to apply a TII, organizations survey their applications and review each aspect of their existing information service levels to ensure that it is justified. Even where tiers exist, we have often seen organizations that don’t fully understand end-user service objectives (let alone capture the objectives in service level agreements). It is very common that changes in service level over time are not described for applications. Service level comprises attributes such as:

- *Data criticality to business.*
- *Performance needs in terms of read/write characteristics, sequential versus random access, large versus small data transfers. This may also include defining acceptable performance levels when online backup or recovery is taking place.*
- *Availability level needed; for example, maximum minutes of unexpected downtime per year, maximum interval of planned downtime per week.*
- *Recovery point objective and recovery-time objective; i.e., how up-to-date must a recovery be and how long can you afford to wait for a recovery to complete.*
- *Recovery method: by restore or by local replica, by remote replica or some combination.*

Ignoring service level requirements and their change over time can greatly inflate infrastructure costs and result in operational inefficiencies.

For example, both the in-house staff and affiliates of a major television broadcasting company must be able to search and edit digital video sports events. During and very soon after each sports event access frequency is intense as various derivations of the sports event are created and processed. The availability and performance requirements, recovery time and criticality of the raw and derived footage all change very quickly in the days and weeks following the event. Since new sports events occur frequently, there is a huge payback to analyzing the effect of time on service levels. In this example, the client saved an order of magnitude usage of high-level enterprise tiers through a careful time analysis.

Understand file system and storage usage to identify opportunities

Organizations with a big file system infrastructure often have little awareness of the true magnitude of their contents. Similarly, they typically have a limited ability to reorganize and optimize file storage. Sometimes business processes do not use IT resources in an organized way. File systems may be dispersed across departments, and departments may not be resourced to manage their own space usage. Siloed information resources can drive up the cost of providing service. Paradoxically, tiering information resources can lower these costs because lower value tiers are typically managed less intensively.

Storage management tools can survey a diverse set of file systems and summarize file types, characteristics and access patterns, and even identify duplicate files. This analysis is one key to unlocking the use of storage. Many organizations use such tools in some form; however, we believe that certain capabilities are essential to effectively “right-tier” files on a continuing basis and unlock the savings potential in a tiered environment:

- *Look for tools that facilitate data movement across storage tiers. Some storage management tools are integrated with file archiving and deletion tools to enable policy-based automatic movement among tiers. Hierarchical Space Management (HSM) tools can move less critical files to lower-cost storage and optionally bring them back on demand. Some archiving tools are integrated with backup tools to reduce redundant copies of data across multiple tiers, at great savings.*
- *Look for common management capabilities from a few flexible products leveraged to work over a variety of platforms. In a tiered storage environment, administrators will want to monitor and configure many different storage devices and applications from a single interface. Trying to use multiple storage resource management interfaces to juggle different classes or silos of information can consume more administration overhead. The lack of integration makes it harder to leverage insight from one tool into another.*
- *Look for workflow automation support. To the extent you can customize and automate common tasks across different levels of your information infrastructure, you will be able to react to changing conditions with less effort and turn new procedures into routines. Provisioning, data movement, monitoring and performance management are all common tasks that when automated can reduce the labor associated with managing information across multiple tiers. Workflow automation can enable you to use the same staff to get more done with fewer errors.*

Of course, not all these capabilities must be deployed initially. The key point is to work with tools that support the level of integration and automation that you target over time.

Incidentally, one feature we have seen used successfully as a change tool in IT organizations with tiered infrastructures is known as charge-back. The charge-back feature of storage or file system managers lets you report on resource usage such that utilizing departments could be charged for their actual usage. This report will help everyone understand how functional usage decisions result in IT cost. It may not be necessary to actually set up internal billing for resources. The feature is a communication tool that creates incentives for departments to make cost-effective decisions on the service levels they require.

Exploit application-aware support for data movement

Some organizations manage large volumes of information within applications, which presents a challenge to cross-tier movement. For example, applications such as SAP maintain records in interdependent database tables. E-mail servers use a special format for storing messages. And of course, the necessary lifetime of data in such applications can be very long, even if the data is rarely used. Once storage is allocated, the all-too-common (but expensive) choice is to expand it indefinitely and migrate it whole when necessary.

A fast-growing market exists in tools to migrate or archive application data such that it can still be easily accessed when needed. These tools enable the movement of application data to manage space and support the implementation of corporate policies and compliance with government regulations. There are typically excellent business justifications for these tools besides infrastructure cost control. CIO-level concerns include complying with regulations, protecting intellectual property or controlling the risk of lawsuits.

Application-aware data movement solutions enable you to define business rules to separate reference, active and inactive enterprise information. They enable you to move records according to their evolving business value. You can keep active records (for example those in the current and most recent past month) in high-bandwidth, high-availability storage tiers. Infrequently accessed records are moved to mid-tier or lower-duty-cycle storage with longer recovery time. Reference data is moved to lower-cost archive or secure non-erasable, non-rewriteable devices.

These solutions may also include features to satisfy other IT stakeholders. For instance, they may support the ability to “scrub” personal data and quickly create mock test data sets. They may include advanced searching capabilities to enable more cost-effective and efficient response to audits and investigations.

As an example, one of the world’s leading shipping lines used an e-mail archiving system to ensure safekeeping of records while freeing up premium network bandwidth and storage space for online applications and the most useful messages. This company wanted a centralized solution to cover more than sixty Lotus® Domino® server locations – as well as lay the foundation for a strategic shift to business content management of information after the e-mail project. Their deployed solution automatically archives e-mail older than 60 days from Domino servers by moving messages to lower-cost disk, yet it allows transparent on-demand retrieval.

After one year, e-mail is transferred to automated near-line tape libraries. E-mail older than four years is stored offline on inexpensive tapes to meet a 10-year retention requirement. Employees can retrieve their messages in original form when needed. It typically takes no more than three seconds for retrieval from disk and six seconds for retrieval from near-line tape libraries. They are moving tens of thousands of messages daily to optimize their infrastructure and realizing significant annual savings on the fully loaded cost of storage, as well as seeing improvements in both e-mail performance and employee productivity.

Virtualize disk storage to move data seamlessly, improve utilization and reduce management touch

Storage virtualization is certainly not required to implement a tiered infrastructure. However, depending on the environment, the features of virtualization may complement and amplify the infrastructure. In particular, virtualization speaks to the key need for data movement and the goal of reduced management touch:

- *Virtualization solutions support the non-disruptive movement of data between disk tiers and movement from legacy storage controllers to newer more cost-effective ones. Non-disruptive data movement can improve availability and reduce off-hours work. It eliminates a major source of downtime and greatly reduces administrator labor in migrations. The easier it is to move data among tiers, the more efficiently you can drive down overall storage and labor cost, and the more flexibility you will have to adapt to future needs.*
- *Due to greater flexibility in moving data to match storage characteristics, you have the opportunity to consolidate like infrastructure workloads and increase storage utilization rates within tiers. We have frequently seen some storage overcommitted in terms of space or bandwidth, while other storage lies underutilized. This imbalance often exists simply because of the administrative cost of changing configurations. Better utilization reduces waste and reduces the cost of incremental storage purchases. Consolidation can simplify management and reduce ancillary overhead such as power consumption and rack space, delivering big benefits in some IT organizations.*
- *Some virtualization solutions also reduce the complexity of managing storage by providing a single point of management for multiple storage devices and infrastructure services, such as business continuance replication. They can even unify services across devices from different vendors.*

These capabilities can control the trend to grow labor cost to provide service levels even while storage capacities rise.

For example, a major online and brick-and-mortar retailer of high-quality outdoor equipment in the United States decided to add virtualization when they purchased a large mid-range storage subsystem. Their existing infrastructure included servers and storage from multiple vendors, but as the company grew, their infrastructure could no longer handle high workloads or meet required service levels. Their goal was to add capacity but also to reorganize their existing data into tiers to accommodate different availability requirements, optimize their use of the new storage, and make better use of their older equipment.

Since deploying the tiered system with a storage virtualization appliance more than a year ago, they have had zero unscheduled downtime, can perform most upgrades and maintenance without scheduled downtime, and have more flexibility to expand and move data around. If business demands change, such as moving an application to a different class of service, the application need not be taken offline. These capabilities, together with simplified management, have been important for handling their explosive growth. An added bonus of the new solution is higher performance for the company's business-critical applications. The added caching present in the virtualization appliance has helped legacy as well as new disk performance.

Deploy a compliance/retention solution that combines disk and tape to utilize lower tiers

Compliance, archive and data retention solutions respond to critical corporation-level risks and pressures, although they also offer a great opportunity to meet additional IT goals through a tiered information infrastructure.

Compliance and retention solutions help reduce the risk of penalties or brand damage associated with not meeting regulatory requirements. They are designed to preserve authentic copies of all kinds of records for the legally prescribed period. They can also help you easily produce records when needed, and help you eliminate records once they are no longer absolutely necessary.

From the IT perspective, these solutions can be a powerful tool to deal with exploding information volumes in areas such as e-mail, digital images or content, medical records and database transactions. Innovative archive and retention solutions can move information to lower-cost infrastructure tiers under automation or policy control. They can enable IT to exploit less expensive infrastructure, with less manual effort than would be needed otherwise. As a side effect, offloading older information to an archive can let you optimize the utilization of more expensive tiers, controlling cost and improving service levels.

Consider the following example of an IT subsidiary of a leading European automotive company that manages trading, import, financing and insurance systems for several car brands. The company is subject to stringent government requirements for data retention. While it had used optical jukeboxes to archive data, according to its projections the company was looking at tripling the amount of jukebox hardware over three years, with a substantially increasing requirement for management effort. After an assessment, the company decided to move to a data retention system with an integrated backup storage manager capable of utilizing write once, read many (WORM) tape.

In the new system, information is initially archived to the 3.5TB data retention device, where it resides on lower-cost disk in immutable form, yet can be accessed very quickly as needed. Migration policies control the automatic movement of information to near-line WORM tape, where it resides for four years, at much lower cost, yet can still be accessed transparently for contract activity by the car brand companies as needed. After four years the tapes are moved offline, where their retention continues to be automatically managed for the required time. The total system currently uses 100TB of online and near-line storage and 200TB of offline storage. Because of the automatic, policy-managed use of lower-cost media, the company has reduced its incremental labor cost as well as saved by avoiding continued hardware purchases as its data grows.

Combining policy-based migration tools with near-line tape automation is a highly cost-effective TII approach. Automated tape libraries are available in a wide variety of configurations, from a single 2U rack form factor holding 8TB all the way up to systems with thousands of tape slots holding petabytes of data. They can return data from tape in seconds or minutes depending on configuration. As with disk, the fully loaded cost of near-line tape continues to decline due to technological innovation. Tape media cost is a fraction of disk and can be converted to offline storage as needed. The solution can automatically manage the location of each critical document and help you optimize usage of your higher tiers.

Utilize content management to facilitate data movement

Content management (CM) solutions have evolved to handle the growing degree of digital recordkeeping and the need for integration of business processes. This evolution has resulted in powerful capabilities to move information within the infrastructure, and for this reason content management is often aligned with a TII. There are a number of areas within CM that can offer excellent payback.

We have seen clients start small but phase in CM over the course of years as more and more types of business records and objects are brought under management. Documents can be captured and indexed in multiple formats. Multiple existing repositories can be federated to appear as one to the end user, even as repositories are migrated across different tiers of storage over time. Documents can be structured into defined workflows and routed into different storage tiers as appropriate to their life phase. Objects can be managed through their lifecycle, while linked to records in a central repository, and archived to reduce compliance risk.

Such projects typically support other IT goals or mandates; for example, content within output from host applications (such as ERP/CRM) can be brought under management. Documents not previously captured can be imaged and electronically managed. Presentation through Web and e-mail portals can deliver managed content to users via easy-to-use interfaces.

For example, an IT division of a large Austrian banking group has been phasing in content management for several years. They have focused on their high volume and critical information such as customer and transaction records, invoices, trading reports and human resources documents. Over time they have moved large amounts of information to lower-cost, high-capacity tape and optical storage, reducing loading time and freeing up higher-performance tiers to accommodate growth. In addition to reducing storage cost, the CM solution has put an end to time-consuming manual and fragile scripted searches across multiple databases. Business records within e-mail attachments are automatically captured into the system.

They have also digitized paper and microfiche records, and consolidated storage of human resources documents, saving valuable office space. Access by thousands of bank employees across branch offices is centralized and improved via tools they are already comfortable with – e-mail software, Web browsers, and desktop office PCs. In addition to the other benefits, the CM and optical storage solution helps the company meet legislative requirements for the banking industry, even in the face of increasing data volumes.

Exploit outside services to help assess maturity, make recommendations and guide deployments

Consulting services can be extremely helpful to IT organizations that seek to implement a tiered infrastructure. Some IT consulting services specialize in strategy. Others specialize in planning or technology project deployment, and some IT consulting services focus on the challenges special to an industry. Of course, many IT organizations already have these functions in house. If so, why would it make sense to bring in an outsider?

Engaging an outside consulting service has a variety of advantages that are often ideally suited to effect a change to “business as usual.”

Outside services can:

- *Bring broad experience gained from working with many other organizations, and by listening carefully, translate the best of those practices into your specific situation.*
- *Apply valuable expertise or tools that are not cost-effective for your business to maintain in-house. Services teams typically have subject matter experts with specialized tools to use your time as efficiently as possible.*
- *Integrate conditions, issues and opportunities from across the diverse parts of your organization. The independent perspective can see patterns and trends, or hear problems that may be hidden by everyday operations, and act as a foil to ease potential internal political issues.*
- *Serve as a useful, independent point of comparison with your internal strategy or architecture team, and thus bring in fresh ideas. They can be a lever, if necessary, to help you challenge the existing status quo.*

For these reasons the services advantage can be a very cost-effective way to optimize your environment and help your IT organization innovate.

As an example, a major public research university in the central United States needed to respond to rapid increases in demand for storage of student, faculty research and medical center data. In fact, the university projected growth rates of 50–70% over the next three years above their existing 45TB of storage.

The university believed its management tools/automation would not handle that growth within current headcount, and did not think it could meet availability and performance objectives.

The university engaged a business partner and a small separate services team to analyze the environment and conduct a tiered storage assessment. A key success factor in this project was to give the assessment executive level support. The team interviewed and gathered data spanning several functional areas within the IT organization (as well as some areas outside IT). They identified strengths, issues and opportunities. They made prioritized, detailed recommendations linked to the client's strategic goals and discussed the analysis with the IT leadership team. The leadership team used or adapted the most practical and achievable parts of the studies into their planning cycle.

As a result, the university implemented a three-tier infrastructure leveraging their existing multi-vendor storage and adding some new equipment. Most importantly, they invested in tools to centralize storage administration, and tools to classify and migrate information in applications and file systems within the tiered, centralized infrastructure. Since deployment, the university reports successfully using the tools to manage their storage growth with no increase in the number of administrators. Quarterly costs of storage purchases have decreased, while uptime has improved dramatically.

Business case

The financial business case for a tiered information infrastructure is based primarily on a simple principle: The spiraling costs associated with information depend on the service level that is provided. Second, savings come from pushing information to the lowest-cost tier that will support the business need.

Of course the full picture of benefits will vary according to which approach is taken. This analysis focuses on the financial advantage to being able to react efficiently to the reality of the growing volume and shifting value of information. We are concerned with exploiting existing and added resources efficiently rather than moving wholesale to a new, larger set of equipment.

Our analysis considers a client with 50TB of storage with a projected annual average growth rate of 50%. This client chooses a TII approach to increase their control over their environment. We estimate the costs and benefits of implementing the solution over three years – even though many investments would have a longer lifetime. Frequently a TII project will be part of a larger strategy and set of investments, with other costs, and with other synergistic benefits.

Costs

We will assume that the chosen approach will have costs in the form of hardware or software, maintenance, and cost of labor to implement it over time (see Table 1). These are due to an initial acquisition of \$300K of new infrastructure tools and capabilities, plus the cost of exploiting them over a three-year period. Even though we're spending \$300K exclusive of buying more storage capacity, it is worthwhile because we get a significant payback leveraged over the growing infrastructure over time. This analysis focuses on how efficiently the growing storage needs are managed.

Table 1

Infrastructure Investment Costs	Initial	Year 1	Year 2	Year 3	Total	Present Value
Hardware and Software	(\$300,000)				(\$300,000)	(\$300,000)
Maintenance			(\$75,000)	(\$75,000)	(\$150,000)	(\$118,332)
Administrative or Services	(\$40,000)	(\$60,000)	(\$60,000)	(\$60,000)	(\$220,000)	(\$189,211)
Total Costs	(\$340,000)	(\$60,000)	(\$135,000)	(\$135,000)	(\$670,000)	(\$607,543)

We assume an initial infrastructure investment of \$300K in hardware and/or software, which includes one year of maintenance and support. After the first year, yearly maintenance is \$75K. Administrative labor associated with initial deployment is \$40K (approximately 0.3FTE). After initial deployment, the effort needed to exploit the solution is \$60K (approximately 0.5FTE) per year.

Benefits

The analysis will assume only two general benefits of the solution:

- (1) We expect to save money on future storage purchases and associated cost because a TII will enable us to shift data onto less expensive tiers; and,
- (2) We expect to save money on administration labor because automation or tooling assistance with data movement will reduce the work needed to adapt to data growth. Table 2 shows how we assume how these benefits occur over three years.

Table 2

Tiered Data Distribution Changes	BAU	New Solution		
	(fixed)	Year 1	Year 2	Year 3
Enterprise	70%	60%	47%	23%
Midrange	20%	20%	23%	27%
Low-cost	10%	20%	20%	30%
Near-line tape			10%	20%

Administration Cost Changes (FTE per Managed TB)	Year 1	Year 2	Year 3
BAU	0.14	0.13	0.12
New Solution	0.13	0.11	0.10

In the Business As Usual (BAU) case, 70% of data is stored on enterprise-class storage, and the distribution ratio is maintained as storage grows. Labor of 0.14 FTE is needed to manage each terabyte (TB) of information infrastructure in Year 1, and this labor ratio decreases over Years 2-3 due to economy of scale. We assume the New Solution tooling enables staff to migrate data to lower-cost tiers over Years 1-3, and enables staff to manage about 6% more data in Year 1, and 18% more data in Years 2-3 than would otherwise be needed to accommodate storage growth.

To translate these benefits into dollars, we also make certain minimal assumptions about the cost of storage, cost of administrative labor and cost of money¹.

Depending on the approach chosen, a tiered infrastructure is likely to have other benefits besides the two listed above. A careful and realistic assessment of a specific approach to a real-life business situation will include other factors. For example, following are some benefits that we have seen from a TII solution:

- *We don't assume storage utilization improves, although in some approaches that is a valuable benefit.*
- *Improved end-user productivity, although fairly frequently seen, is not assumed here.*
- *Reduction in risk when the solution also addresses related IT goals such as business continuity, compliance, etc. is not measured.*
- *We are not measuring strategic benefit – the various ways that having a more flexible infrastructure enables the business to respond to competitive pressures.*

The latter two benefits, while hard to quantify, can be extremely valuable across the broader organization, and indeed can drive the whole investment provided the other factors are positive. However, to limit the scope of the analysis given in this paper, we use the more conservative set of benefits.

Table 3 summarizes the value of the benefits under the assumptions as stated above.

Table 3

Total Infrastructure Benefits	Initial	Year 1	Year 2	Year 3	Total	Present value
Reduced cost of storage			\$231,625	\$598,772	\$830,397	\$641,292
Reduced administration cost		\$50,400	\$219,024	\$315,395	\$584,819	\$463,790
Total benefits		\$50,400	\$450,649	\$914,166	\$1,415,215	\$1,105,082

After the first year, the TII enables us to shift purchases of new storage to a less expensive tier than we otherwise could. Reduced cost of storage reflects hardware as well as service-level costs associated with higher-tier storage for information. Reduced administration cost shows the migration effort saved by exploiting TII tools to move data to accommodate storage increases.

Bottom line

Table 4 shows the final bottom-line result of the analysis. The project pays for itself in roughly two years. Over three years, the internal rate of return is a healthy 54%. In all likelihood, the benefits from the investment will continue to accrue well beyond three years.

Table 4

Cash Flow Summary	Initial	Year 1	Year 2	Year 3	Total	Present value
Total costs	(\$340,000)	(\$60,000)	(\$135,000)	(\$135,000)	(\$670,000)	(\$607,543)
Total benefits		\$50,400	\$450,649	\$914,166	\$1,415,215	\$1,105,082
Yearly net savings	(\$340,000)	(\$9,600)	\$315,649	\$779,166	\$745,215	\$497,539
ROI	82%					
Payback period (years)	2.04					
IRR	54%					

Combining the costs from Table 1 and the benefits from Table 3 gives the yearly net savings shown. Return On Investment (ROI) on this cash flow is 82% over three years, with the project paying for itself in just over two years. Internal Rate of Return (IRR) is 54%.

In spite of using fairly conservative assumptions, this simple analysis illustrates the latent opportunity in a tiered infrastructure. The business conditions of a quickly growing volume of information, and the need to control the cost of providing the appropriate level of service, create a financial incentive to invest in agility. Under these conditions, there is a substantial opportunity cost to business as usual.

IBM capabilities

IBM offers a complete portfolio of products, solutions and services. We create innovative solutions and work extensively with business partners in all the areas mentioned in this paper. We have a proven track record in ILM and tiered infrastructure projects.

Our solutions are independently recognized as leading the industry. For example, IBM Federated Records Management for DB2® Content Manager OnDemand was named the AIIM Enterprise Content Management “Best of Show” 2006 in the records management category – for the second year in a row. Likewise, the IBM System Storage™ DR550 was named AIIM’s “Best of Show” 2006 in the hardware storage category, also for the second year in a row.² The IBM System Storage SAN Volume Controller has more than 2,000 clients and is entering its fourth year of market acceptance. The SVC holds leading performance results for disk virtualization, and in fact, leading results for all storage controllers that have published benchmarks with the Storage Performance Council.³

Equally important, IBM has deep experience helping clients succeed with infrastructures – with our products as well as other companies' products. We offer skilled professionals worldwide to assist with assessment, design, and deployment, and to help reduce the complexity of your existing environment. Gartner acknowledges IBM as the market share leader in storage services and ranks us as #1 in ability to execute.⁴ For example, IBM Global Services offers patent-pending, proven and repeatable techniques for facilitating data classification, service level definition, policies and hardware tiers.

For more information

For more information on how your organization can implement a tiered information infrastructure, please contact your local IBM representative or IBM Business Partner, or visit our Web site at:

ibm.com/servers/storage/solutions/ilm.



© Copyright IBM Corporation 2007

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A.

Produced in the United States of America
03-07
All Rights Reserved

¹We also make the following assumptions:

Discount rate: 10%; Fully loaded cost of an administrative FTE: \$120K; Annual increase in FTE cost: 4%; Fully loaded cost per TB of managed storage in Year 1: Enterprise: \$55K, Midrange: \$25K, Low-cost: \$10K, Near-line tape: \$2K; Annual reduction in fully loaded cost of new purchases of storage after the first year: 15%.

²AIIM:

<http://www.aiim.org/article-pr.asp?ID=31452> (2006)
<http://www.aiim.org/article-docrep.asp?ID=29867> (2005)
<http://www-306.ibm.com/software/swnews/swnews.nsf/n/hhal6q3s7v?OpenDocument&Site=data>

³SVC

<http://www-03.ibm.com/systems/storage/news/press/20061206.html>
<http://www-03.ibm.com/systems/storage/news/press/20060525.html>

(This paper does not use the claim of 53% of storage virtualized)

⁴Gartner's Magic Quadrant 2Q'06: <http://mediaproducts.gartner.com/reprints/emc/vol2/article2/article2.html>

DB2, Domino, IBM, the IBM logo, Lotus and System Storage are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.