



IBM XIV[®] Storage System

Optimizing Enterprise Storage Total Cost of Ownership (TCO) with IBM XIV

White Paper | June 2009



The IBM logo, consisting of the letters "IBM" in a white, bold, sans-serif font, set against a blue background with a large, faint, stylized starburst pattern.

Copyright IBM Corporation 2009

IBM, the IBM logo, ibm.com, System Storage, XIV, and the XIV logo are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at ibm.com/legal/copytrade.shtml.

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

This document could include technical inaccuracies or typographical errors. IBM may not offer the products, services or features discussed in this document in other countries, and the product information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Information concerning non-IBM products was obtained from the suppliers of their products, their published announcements or other publicly available sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers. IBM does not warrant that the information offered herein will meet your requirements or those of your distributors or customers. IBM provides this information "AS IS" without warranty. IBM disclaims all warranties, express or implied, including the implied warranties of non-infringement, merchantability and fitness for a particular purpose. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

askxiv@us.ibm.com
www.xivstorage.com

Contents

Introduction.....	1
IBM XIV: Designed for Low TCO	2
Acquisition Costs.....	4
Single-tier architecture advantages	5
Commodity hardware	5
Very High Density Slower Rotation (VHDSR) disk drives.....	6
Built-in thin provisioning	6
Optimal capacity use	7
Simple pricing, with all-inclusive features.....	9
Summary: How to reduce acquisition costs.....	10
Environmental Costs	10
Power and cooling consumption.....	10
Floor space use.....	11
Summary: How to reduce environmental costs	11
Administration and Management Costs.....	12
Less management overhead	12
Summary: How to reduce management costs	13
Downtime Costs	13
Less failure to be managed.....	13
Less human intervention and error.....	14
High availability and non-disruptive changes.....	14
Hot, online, intra- and inter-family code upgrades.....	14
Summary: How to reduce downtime costs	15
Backup/Restore Costs	15
Summary: How to reduce backup costs.....	15
Conclusion.....	15
Appendix: Cost Savings Formulas	17
Floor space savings.....	17
Power and cooling savings.....	18

Introduction

A high-end storage system's total cost of ownership (TCO) is like an iceberg: the vast majority of costs are hidden from view. Comparing systems by looking at obvious factors, like capacity and box price, may be easier. But taking a fuller look at the factors impacting IT budgets over a system's lifetime will yield a more accurate TCO assessment and, ultimately, a more informed decision as to which system is truly more favorable to the company's bottom line.

The chart below shows a fuller set of storage array cost factors and their generally recognized relative impact on a system's overall cost.*¹

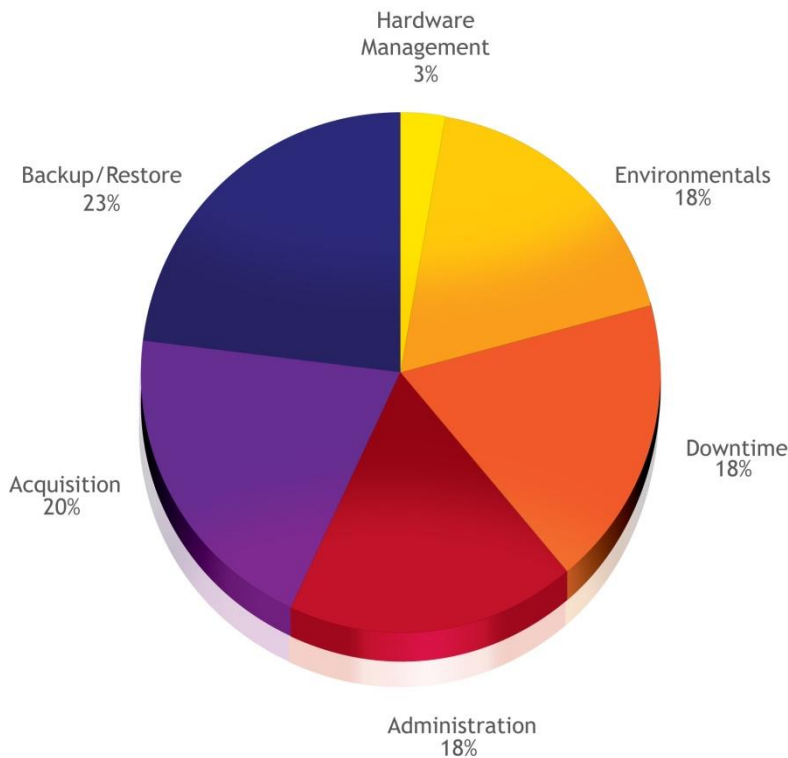


Figure 1: Typical TCO costs of enterprise storage arrays

¹ In "[Governance Dashboard for Enterprise Storage](#)" (*Information Management Magazine*, May 2006), Dick Benton of GlassHouse Technologies Inc. notes that Gartner and Forrester attribute as much as 60-70% of overall storage TCO to the Administration factor. This large figure likely comes from grouping Administration together with Backup/Restore and other cost factor categories. Here, we break them out separately.

In short, a thorough TCO assessment needs to take into account the cost of:

- ▶ Buying, implementing, running, cooling, and expanding the system
- ▶ Managing, integrating, and testing the system
- ▶ Ensuring the system's reliability and availability
- ▶ Downtime caused by routine maintenance or unexpected equipment failure
- ▶ Savings opportunities lost during the time it takes the system to integrate newly available, cost-reducing technologies
- ▶ Software licensing – at time of purchase and upon each capacity increase

The IBM XIV[®] Storage System has been architected to address all these cost issues. This paper looks at the true total cost of owning an enterprise-class storage system and the XIV system's winning, truly cost-reducing approach.

Important note:

It should be stated that it is *inaccurate* and *even impractical* to compare overall cost of the XIV system with other storage systems using all VHDSR drives or current-generation enterprise-class systems using a significant number of 1 TB drives. This is because **these systems**, unlike IBM XIV, **do not use 1 TB drives for tier 1 or even tier 2 storage**; rather, for archiving and other needs that do not require high performance or high reliability. The XIV system provides enterprise-class performance and reliability for all data at all times. In short, the quality of the storage is different.

It goes without saying that the huge cost advantage resulting from XIV ease of management and software features is relevant in any comparison.

IBM XIV: Designed for Low TCO

Architected to deliver high reliability and performance at exceptionally low TCO, the IBM XIV Storage System is:

- ▶ A fully virtualized, self-managed, self-healing system that balances load perfectly at all times, optimizing capacity use and minimizing human intervention and error
- ▶ Complete with a built-in features set, including differential snapshots, remote mirroring, and thin provisioning
- ▶ Easily managed through a simple graphical user interface (GUI) that minimizes storage management overhead
- ▶ Built entirely of standard components, including low-cost, very high density slower rotation disks (VHDSR)
- ▶ Fully hardware redundant, lowering risk of failure to negligible levels and enabling disks to be added and replaced on the fly without a performance hit or downtime
- ▶ A single-tier platform based on a grid architecture, for easy scaling

The XIV system is designed to comprehensively tackle each storage TCO factor, in multiple ways, as delineated below:

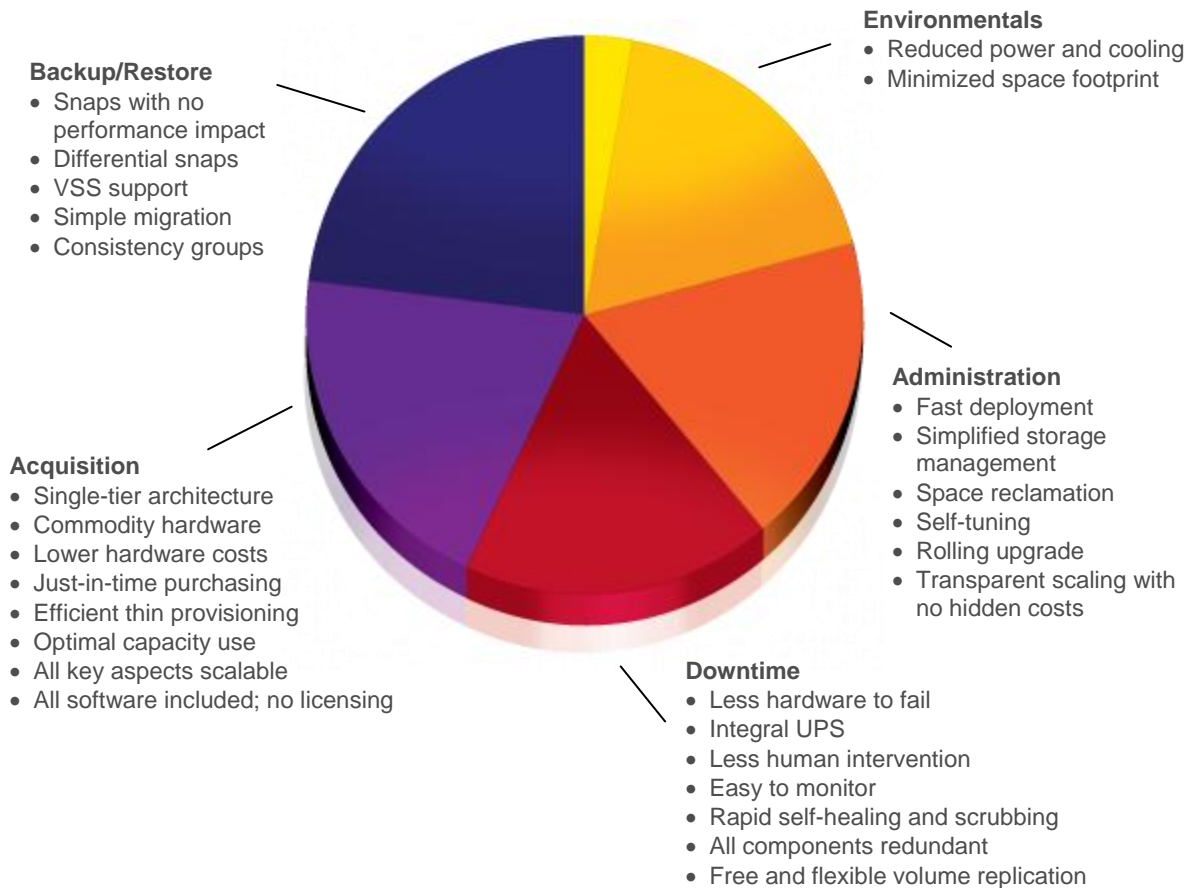


Figure 2: How XIV storage reduces TCO

The result is that the XIV system succeeds in dramatically improving each TCO factor (with the exception of fixed factors, such as hardware maintenance) while other storage systems typically focus on reducing only one or two cost factors. For example, even if a system's upfront cost is largely discounted, its TCO will still be high due to ongoing administration and maintenance costs; a system with improved performance often has higher environmental costs.

By taking a balanced and comprehensive approach to reducing all cost factors, the XIV system succeeds in dramatically reducing TCO where other systems fail.

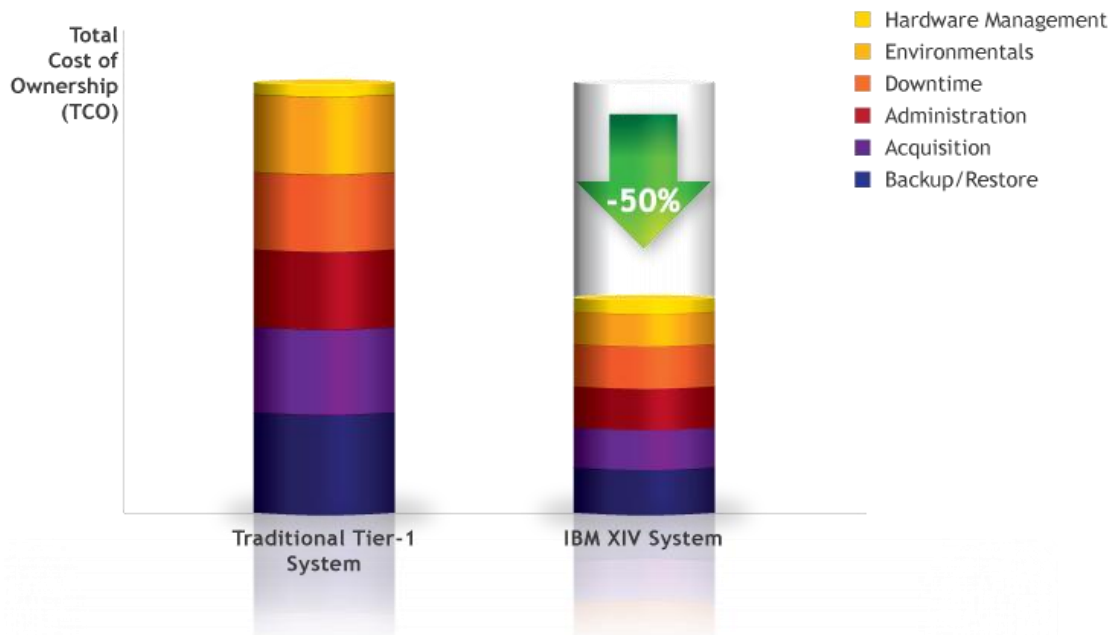


Figure 3: Overall TCO: Traditional Tier-1 storage vs. XIV storage

Acquisition Costs

The IBM XIV Storage System is designed to be cost-efficient in all aspects, while delivering outstanding and consistent performance. The XIV system's:

- ▶ **Single-tier architecture** supports all kinds of storage in one environment and makes scaling easy
- ▶ Use of **off-the-shelf components** makes customized hardware unnecessary and makes it possible to integrate newer, state-of-the-art hardware as soon as it becomes available and without delay
- ▶ Innovative use of **highly economical disk drives** [Very High Density Slower Rotation (VHDSR) drives] offers excellent capacity-for-cost value while meeting enterprise performance standards
- ▶ Built-in **thin-provisioning** allows the acquisition and installation of additional capacity to be delayed until needed, making it possible to defer capital purchases and take full advantage of future hardware and price-point improvements not currently available

- ▶ **Automatic and perfect load-balancing, self-healing capability**, and other architectural features provide tier-1 enterprise-level performance while using fewer disk drives and, ultimately, much less capacity than traditional architectures
- ▶ Simple pricing **includes all features**, making purchasing simple and eliminating software licensing fees. There are no software licensing fees, neither upon initial purchase, nor over time when capacity is added.

The Taneja Group: "Based on what we've seen of its design and architecture, XIV remains unstoppable like the best of enterprise storage, but may be fundamentally more extensible, scalable, and adaptable as technologies evolve. Without a doubt, it brings a new set of economics to the table for enterprise-level storage."

[Taneja Group brief on the XIV system, Nov. 2008](#)

Single-tier architecture advantages

Most IT centers today use Information Lifecycle Management (ILM) and similar practices to cut costs. These efforts involve multi-tiering, which adds a substantial degree of complexity and cost to system management. A multi-tier storage solution typically involves a number of different architectures, each requiring a different kind of know-how to manage. Moving the data between tiers requires effort and planning, and involves hardware and operating expenses for each tier. As each environment grows at its own pace, equipment must be added to one tier or another. Since this growth is typically unpredictable, acquisitions are often not cost efficient.

The single-tier XIV architecture eliminates the complexity and cost of migrating data from tier to tier, and of maintaining different architectures for different tiers. The XIV system offers further cost efficiencies by enabling scaling of all components – capacity, interfaces, cache, CPU power, and internal bandwidth – on the same platform. As such, the XIV system helps organizations curtail and minimize hardware costs while offering the ability to grow at any time, and only as much as needed.

Commodity hardware

Traditional enterprise-class storage systems are usually built from proprietary hardware. This typically makes them costly due to the huge investment in their development and version testing – but, more importantly, creates a built-in rigidity that makes such systems unable or slow to integrate new technologies.

The XIV system minimizes costs and provides exceptional flexibility by treating hardware as an interchangeable commodity. Built entirely of off-the-shelf hardware components, the XIV system is designed to integrate the latest disks, modules, interconnects, and

other components, enabling organizations to capitalize on market advances in capacity and cost while meeting their dynamic performance and budget requirements.

The XIV system allows integration of multiple generations of hardware – new alongside older – on an in-production system. Component upgrades and replacements are non-disruptive and require little management. Adding capacity with the XIV system is similarly seamless, with the new capacity available immediately, without the need to reconfigure and without performance degradation.

Very High Density Slower Rotation (VHDSR) disk drives

Enterprise storage systems are traditionally configured with Fiber Channel (FC) drives generally known for speed and reliability and, consequently, high cost. The XIV system uses only Very High Density Slower Rotation (VHDSR) disk drives – known for their very high densities as well as comparatively low cost. The XIV system applies innovative load balancing and other advanced algorithms to provide tier 1 performance, reliability, and availability, coupled with the dramatic per terabyte savings, and power, cooling, and maintenance efficiencies that result from using the capacity-rich VHDSR drives.

Gartner Inc. notes that: "...higher-capacity SATA disks generally cost less than smaller FC or Serial Attached SCSI (SAS) disks on a \$-per-GB basis, and consume much less power and cooling. When the use of these high-capacity 1 TB disks reduces the overall disk count in the storage system, relative to a system configured with FC or SAS disks, it also reduces the frequency of disk-related repair activities."²

For more information on XIV's high-end performance using VHDSR drives, see the [IBM XIV Performance Reinvented](#) white paper.

Built-in thin provisioning

Most enterprise systems now offer thin provisioning – the ability to define a system's logical capacity as larger than its physical capacity, so as to defer physical capacity purchases (see below). However, in many systems, this feature has been added to the existing architecture, making it harder to manage and of limited scope. The IBM XIV system provides thin provisioning as a core feature of its design, handily managed at the click of a button.

Organizations can leverage XIV's easy to manage thin provisioning capability to significantly reduce capital and operating expenses and postpone capacity purchases by acquiring physical capacity for only the total space actually written rather than the total space allocated. For more information, see the [IBM XIV Thin Provisioning](#) white paper.

² "IBM's XIV: Next-Generation Scale-Out Disk Storage," Stanley Zaffos, Valdis Filks, Nov. 20, 2008, Gartner, Inc. and/or its Affiliates

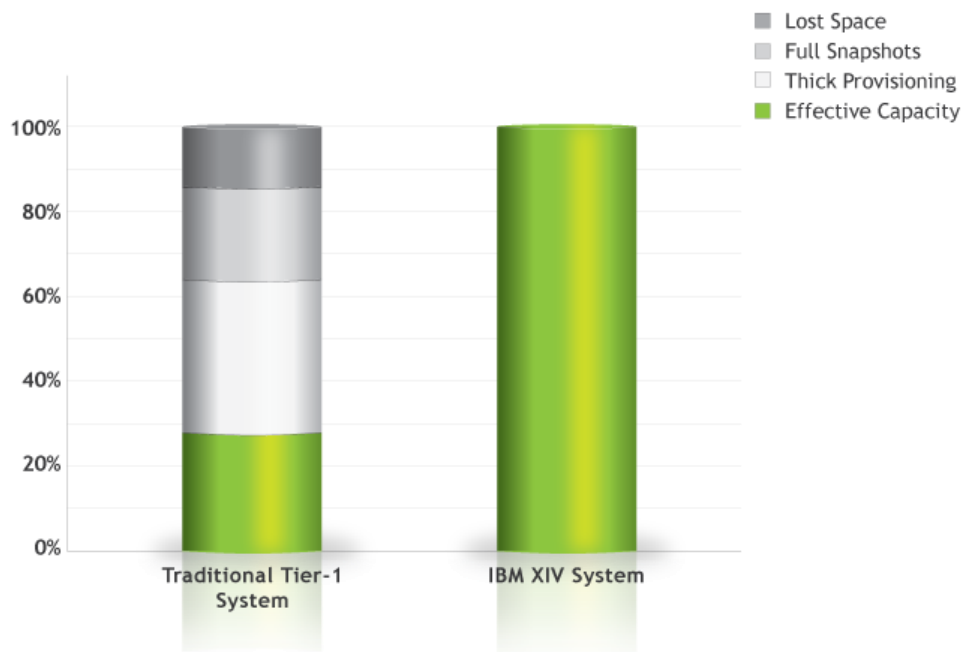


Figure 4: Comparison of typical space utilization

Optimal capacity use

As every storage manager knows, as much as 70% (or more) of a storage system’s total storage capacity goes unused. Whether economic times are good or bad, that’s a troublesome scenario of capacity waste.

Why is it so hard for systems based on current-generation technology to use disk space better?

- ▶ **Burden of buying capacity early.** Pre-allocating storage requires the presence of the defined physical capacity in the system, even if years go by before its use. Upfront purchases commit capital unnecessarily and waste IT resources in managing the unused capacity. Upfront capacity purchases also deprive the organization of benefiting from advances in technology, cost improvements and other advantages at the time the capacity is used.
- ▶ **Management overhead.** Storage teams often over-allocate volumes to avoid the time-consuming task of resizing them later. Reclaiming unused space is cumbersome, requiring time and skill – purchasing more capacity is often considered the lesser evil. (Of course, extra capacity also requires management overhead.)

- ▶ **Inappropriate architecture.** Techniques that help reduce storage pre-allocation, such as thin provisioning, are typically layered on top of the existing architecture and, as such, are limited in their impact.
- ▶ **Capacity-devouring copying.** Many applications, such as testing and data mining, require copies of the original data at the application level (file copy, etc.) or storage level (snapshots and full copies). Current-generation architectures copy both used and unused blocks, thus increasing the amount of unused space.

In contrast, the XIV system is designed to optimize capacity use in many ways, including through self-management, ongoing space reclamation, and powerful native features, such as thin provisioning and snapshots, that save space as they are applied. The XIV system is able to meet the same storage needs as traditional systems with many fewer net TBs. This translates to less spending on physical capacity and fewer related costs by virtue of a system that operates without hotspots and without system down-time or performance degradation during maintenance or scaling.

Self-management

Core to the XIV system's optimization of capacity use is its self-management of data. XIV's automated volume distribution mechanism optimizes the use of capacity across all system disks at all times, including at peak times and regardless of IT activities, application size, or user demand. Its automated approach to volume distribution, including automated load balancing, avoids onerous and error-prone manual allocation tasks. The system also automatically performs space reclamation (details below), further maximizing capacity use on a per terabyte basis.

A global leader in media and entertainment: "Replicating our production was taking a week. With XIV, it takes us 150 minutes."

Differential snapshots (full backup volumes)

XIV's snapshot technology uses a differential implementation that copies only the data that differs between source and snapshot – dramatically and efficiently minimizing the space otherwise consumed by copies and keeping system performance unaffected.

In addition, XIV's snapshot and replication mechanisms take the required storage space from a single, general pool. Traditional systems typically require the allocation of separate storage pools for snapshots and replication. The downside to this approach is that when the pools are from different tiers, spare space is taken from all over the system, increasing the overhead needed to track where the data resides.

Hotspot-free

As described, the XIV system is fully virtualized, distributing data automatically and maintaining perfect load balance at all times, without hotspots. The lack of hotspots eliminates the need to pre-allocate volumes to achieve specific performance requirements.

Hands-free scaling

Whenever new physical capacity is added, the XIV system automatically integrates the new capacity across all the disks and modules and swiftly redistributes the load to maintain perfect balance. For more information on IBM XIV volume distribution and load balancing, see the [IBM XIV Performance Reinvented](#) white paper.

Space reclamation

XIV performs space reclamation automatically and offers easy management tools for leveraging this capability to the fullest. The XIV system recognizes “zeroed out” space as not being in use; it releases any space marked by zeros into the general pool, thus reclaiming it. XIV management tools enable administrators to easily zero out space no longer in use so that the XIV system can automatically reclaim it and allow its reuse. Traditional systems have a different approach to space reclamation; yet, most of these approaches are not currently functioning and will not be available for a number of years. This lack in traditional systems is one of the sources for inflated capacity requirements.

Simple pricing, with all-inclusive features

Simplicity is a key virtue of the XIV system; this simplicity applies to purchasing, as well. Buying an XIV system is a simple matter of specifying the desired gross/net capacity, with no other costs involved.

The XIV system takes an all-inclusive approach to features: every XIV rack comes with all software built-in, as itemized below, at no extra cost and with no license fees whatsoever:

- ▶ Snapshots
- ▶ Consistency groups
- ▶ Remote mirroring
- ▶ Storage management
- ▶ Performance monitoring
- ▶ Thin provisioning
- ▶ Self-tuning
- ▶ Host connectivity packages
- ▶ Data migration
- ▶ Special host features, like MS VSS support
- ▶ Native multi-path support of the host operating system

The lack of software license fees differentiates the XIV system from the vast majority of systems on the market. Most traditional tier-1 storage systems have complex pricing models that, by nature, inflate system cost and complicate attempts to assess true TCO. Such enterprise storage pricing schemes commonly involve: itemized hardware and infrastructure costs, itemized license fees for software features, including for their initial purchase, upon capacity increase, and annual license renewals, as well as other software fees, such as for version compatibility.

The XIV system's native self-healing and load-balancing mechanisms offer the added cost-benefit of sparing or minimizing any purchase of peripheral software enhancers, such as performance optimization tools.

Industrial manufacturer: "The all-in-one software [of XIV] is great—with our previous vendor, retroactive version control was costing us more than the upgrade charges on extra disk space."

Summary: How to reduce acquisition costs

The IBM XIV Storage System helps organizations minimize capital outlay by means of:

- ▶ A single-tier architecture
- ▶ Hardware as an interchangeable, upgradable commodity*
- ▶ Inexpensive VHDSR disks
- ▶ Reduced capacity needs, up front and ongoing
- ▶ Easy scaling – and on the same platform
- ▶ Built-in features and performance optimization

** Also enables organizations to exploit declining storage costs and efficiency improvements*

Environmental Costs

The IBM XIV Storage System is an extremely "green," environmentally advantageous system, offering up front and ongoing savings in power, cooling, and floor space. These savings are so substantial that they alone can cause the XIV system's TCO to be much lower than that of a comparable system.

Power and cooling consumption

The XIV system offers a power and cooling savings of at least 63% over comparable high-end enterprise systems. Simple math shows how:

High-density disks typically feature 7200 RPM, in contrast to the speedier 10K/15K RPM of Medium Density Fast Rotation (MDFR) disks. A VHDSR drive-based system, then,

consumes half the power of an MDFR drive-based system, with at least 25% more savings per drive due to lower energy consumption. A system that uses VHDSR drives of 1 TB capacity per drive, rather than MDFR drives of 146 GB capacity per drive, consumes 10% less power.

See the Appendix for formulas that calculate the comparative energy saved by the XIV system in a given scenario. For more information on how the XIV system reduces power usage, see the [IBM XIV Power Consumption Reinvented](#) white paper.

Floor space use

Today, VHDSR drives offer capacities that are twice as much or more than those of MDFR drives. Replacing an existing MDFR storage system with the XIV system enables customers to use at least 50% fewer disk shelves. For existing storage, this can translate to a space savings of more than 80% (to clarify: VHDSR drives take up the same space as the MDFR drives, but provide 1 TB of storage capacity rather than 146 GB).

Award-winning regional and academic medical center: “We’ve shrunk our floor footprint from 10 tiles down to two. With the more efficient XIV rack, we’re seeing an 80% reduction in our cooling and power use.”

*Greg Johnson, chief technology officer
Virginia Commonwealth University Health System*

Additional floor space is saved as a result of proportionately less related infrastructure, such as power, cooling, and network hardware.

The [Appendix](#) provides formulas for calculating the comparative floor space savings provided by the XIV system in a given scenario. Savings may also accrue by avoiding costs associated with IT facility renovations required to accommodate more physical storage and the related organizational disruption that would take place.

While these indirect savings are often overlooked and/or underestimated, the simple fact is that if physical space is not available, a data center must be expanded or relocated. The space-efficient IBM XIV Storage System may well help protect an organization’s existing infrastructure investment.

Summary: How to reduce environmental costs

The XIV system is inherently “green”: high storage densities and ongoing efficient use of capacity create a minimized footprint – including excellent floor use per capacity ratios, and related power and cooling efficiencies – from the point of deployment and onward.

Administration and Management Costs

Less management overhead

The IBM XIV Storage System is architected to manage itself as much as possible and be managed with minimal effort. Its easy and intuitive user interface is like no other in the industry, and an outward reflection of an efficient, well-structured architecture.

Ease of use

The XIV system's highly intuitive GUI greatly simplifies day-to-day storage administration, enabling most tasks to be executed in just a few clicks. The result is a dramatic reduction in the time required to, among other things, provision storage, resize volumes, create storage pools, and even take snapshots – in most cases, a reduction from hours and days to a few minutes or less. Training on the XIV system is also quick. Exceptional ease of use helps organizations reduce costs by enabling them to assign fewer IT individuals to storage management tasks, even data migration tasks.

In fact, this ease of use combines with XIV role-based access and authorization to offer a new IT paradigm: any staff member, including one without special storage skills, can provision and manage storage. It is conceivable, for example, that DBAs handle provisioning, without external assistance, controlling only those volumes they are authorized to handle based on role. The XIV system's ability to enable lesser skilled personnel to manage ongoing growth translates to substantial potential savings in overall TCO.

Easy monitoring

The XIV system hugely simplifies the monitoring and maintenance of system health, including through alerts sent to select personnel via automated e-mail, SMS, and SNMP that clearly identify a problem and its source. The result is a highly transparent system with speedy mechanisms for easily and efficiently handling maintenance and issues as they arise.

Reduced deployment time

The XIV system obviates the traditional need to intensively plan the layout of application volumes on physical storage units. Its automated, optimized approach to volume distribution and load balancing significantly reduces storage deployment time; many XIV customers report that deployment took less than a day.

Self-tuning

The XIV system automates performance tuning, with the complexities of load balance, the location of mirrored data, and data distribution hidden from the user. With the XIV system, consistent high performance is achieved without the need to tune, reconfigure or rebalance to disperse hot spots. It maintains optimal performance, with no need for performance optimizing software products.

“Thinner” system

The XIV system’s use of high capacity drives and ability to optimize every terabyte results in fewer disks for the same volume needs; this translates to administrative cost savings through fewer disks to manage.

Single-tier platform

The XIV system spares IT teams the administrative efforts otherwise required to support additional architectures or migrate data from one tier to another in an attempt to keep up with changing needs or trim storage costs.

Summary: How to reduce management costs

The IBM XIV system dramatically reduces storage management effort and related costs by automating performance tuning and other traditionally manual management functions, simplifying daily administrative tasks, and offering capacity efficiencies that reduce the overall amount of physical capacity that needs to be managed. The XIV system further reduces management costs by reducing the level of administrator skill needed.

Downtime Costs

Downtime is expensive. The loss or lack of access to mission-critical data can harm an enterprise directly and indirectly. The XIV system keeps data availability continuous and minimizes the extent and cost of downtime through redundant hardware, automated and pro-active maintenance, and efficient hardware repair with minimal human intervention.

Less failure to be managed

Through ongoing automatic monitoring of hardware and data, the XIV system enables – in most cases – detection of imminent failure in time to prevent data loss or data unavailability. The system performs data scrubbing continuously, comparing data copies, identifying and correcting inconsistencies quickly and efficiently. All system components are redundant and constantly monitored to minimize downtime from physical failure. Upon detecting signs of imminent failure, the system creates a third copy of the data on the at-risk disk immediately and automatically alerts operations personnel. The system is designed for component replacement on the fly, without downtime or even a perceptible degradation in system performance.

Should a component fail without warning, the XIV system's rapid self-healing features restore complete system redundancy in minutes, without disruption to vital enterprise functions.

For more information on how the XIV system prevents and handles failure, see the [IBM XIV Reliability Reinvented](#) white paper.

Less human intervention and error

The XIV system further prevents system downtime by avoiding human error that leads to equipment or system failure. Its self-management of volumes avoids the traditional need for manual configuration planning – a task that not only consumes much time, but often increases the risk of data unavailability or data loss (DUDL) due to human errors made when the plans are carried out. Its automatic self-healing, which returns the system to full data redundancy upon hardware failure, avoids the need to rush to repair a failed component, avoiding the kind of human error that can take place when parts are replaced under pressure and at odd hours of the day and night. Furthermore, the XIV system calls for maintenance tasks to be performed only when the system is fully redundant, thus ensuring against data loss caused by human error during maintenance activities, such as removal of a working disk rather than the faulty disk.

High availability and non-disruptive changes

In most enterprise storage systems, planned downtime for data migration and maintenance is costly – and a major inconvenience for the IT team and end users. Minimizing downtime involves much planning. Multiple teams are involved to represent the interests of end users, management, vendors, and service personnel. Negotiating downtime for maintenance is another inevitable source of friction, as maintenance must compete with other, higher priority demands. In traditional systems, rolling out a business application may also require system downtime, with similar costs and conflicts.

The IBM XIV Storage System is designed and implemented to grow, shrink, or change without the need to reconfigure or shut down the system at any time. The administrative headaches enumerated above have been eliminated through XIV's single-tier, constantly balanced architecture.

The XIV architecture allows hardware components to be added to or removed from a new or existing, operational XIV system without downtime. The XIV system provides an easy and seamless non-disruptive procedure for component changes with little management effort and no disruption. Upon any physical capacity change, the system automatically redistributes volume data, maintaining full system equilibrium.

The XIV system's built-in, redundant UPS units provide ongoing protection against data loss while increasing reliability by avoiding the wear and tear to electronic and electromechanical parts caused by inevitable data surges. By sparing its hardware parts undue exposure to fluctuating power, the XIV system essentially lengthens the lifespan of these parts, saving costs by bringing their shelf life closer to their theoretical lifespan.

Hot, online, intra- and inter-family code upgrades

The XIV system allows hot code upgrades and patches within and outside a code family. This eliminates the need to schedule plant downtime, while ensuring ongoing access to

application data. The capability also allows storage teams to proactively implement critical fixes rather than waiting for scheduled downtime windows, not only ensuring continuous availability but also preventing costs resulting from escalation of issues left untreated.

Summary: How to reduce downtime costs

The IBM XIV system is exceptional in its ability to avoid most of the traditional causes of unplanned downtime. It achieves this through reduced hardware requirements, automated and simplified management, self-healing, and hot-code fixes. In addition, the XIV system architecture provides for ongoing changes in storage capacity without incurring degradation in performance or requiring scheduled downtime.

Backup/Restore Costs

With traditional enterprise systems, even backup and restore processes are costly. Traditional backup processes involve the writing of full volumes of data onto additional storage space, including the copying of unused space that has never been written to. The XIV storage system handles backup by copying only the blocks of data that have been actually written to – it does not copy zeroes. In contrast to the high rates of unused space in legacy systems, the XIV approach greatly reduces the time and space that backups require and, consequently, the costs involved.

In addition, the XIV storage system offers logical backup and low-cost recovery through differential snapshots. This innovative mechanism allows organizations to expand and fortify their backup practices without necessarily increasing their physical capacity needs.

IBM XIV offers up to 16,000 differential snapshots, on demand, with no administrative overhead or perceptible degradation in system performance. Differential snapshots can save 15% to 30% of storage requirements, contributing further to the XIV system's low overall TCO. For more information on the XIV snapshot capability, see the [IBM XIV Snapshots Reinvented](#) white paper.

Summary: How to reduce backup costs

The XIV system's easy-to-use differential snapshots take up much less space and fewer system resources to back up data in the system. This approach translates into overall backup and recovery savings through reduced physical capacity and related administration.

Conclusion

A storage system's total cost of operation is comprised of a complex blend of obvious and less obvious factors. It is critical to look beyond easy to compare factors, such as initial cost and power, to consider the indirect factors that are costing the organization money over time. The XIV system is designed to reduce both direct and indirect costs, including

by making storage simple to deploy while keeping data protected from risk and highly available.

Experience has shown that enterprises that deploy the XIV system have achieved:

- ▶ up to 70% savings in tier-1 capacity costs
- ▶ up to 80% savings in power, space and cooling costs
- ▶ up to 80% savings in ongoing system management costs

The simplicity of buying and configuring an XIV system further reduces its overall TCO, by eliminating hidden costs and reducing the overhead typically associated with bringing new technology and systems into a data center. The XIV system's use of density-to-cost-efficient VHDSR drives slashes capacity costs by requiring fewer disks and associated operating costs: each disk requires less power, cooling, and space, as well as less supporting infrastructure and add-on software.

The calculations provided here assume that MDFR drives are replaced by the equivalent amount of VHDSR storage. But this assumption actually grossly overstates the requirements. If a storage project has a significant non-utilization rate – and this is not at all unusual – then eliminating the unused storage can slash up to 70% of the cost right off the top.

XIV's self-balancing, single-tier architecture, with advanced monitoring and failure alerts, requires less training, management, and planning than the typical enterprise storage system – and much less unplanned downtime.

In a nutshell, the IBM XIV Storage System takes a breakthrough approach on multiple levels to deliver a total cost of ownership that is dramatically lower than that typically associated with high-end storage solutions offering comparable functionality and services.

Appendix: Cost Savings Formulas

The formulas below allow you to compare the environmental savings of your current storage solution with the savings you can realize using the XIV storage solution. We have provided examples for Medium Density Fast Rotation (MDFR) drives in 146 GB, 300 GB, and 450 GB sizes, and in RAID-1 and RAID-5 (5+1) configurations. To make similar calculations for a RAID-6 (4+2) installation, simply replace the 0.8 coefficient by 0.75.

Floor space savings

With higher density drives; the XIV system can meet the same needs using fewer drives. To determine the savings in floor space from the XIV system's use of high density drives, use values relevant to your storage operation in the formulas below.

- ▶ To compare RAID-1 drives to XIV's VHDSR drives, use the following formula:

$$\% \text{ Savings in Floor Space} = \{1 - (\text{Your FC Drive Size}) / (\text{the XIV system High Density Drive Size})\} * 100$$

For example: Assuming you are comparing your existing 450 GB drives to XIV's 1 TB (1000 GB) drives, the floor space savings is calculated as follows:

$$\% \text{ Savings in Floor Space} = \{1 - (450) / (1000)\} * 100 = 55\%$$

- ▶ To compare RAID-5 drives to XIV's VHDSR drives, use the following formula:

$$\% \text{ Savings in Floor Space} = \{1 - 0.8(\text{Your FC Drive Size}) / 0.5(\text{the XIV system High Density Drive Size})\} * 100$$

For example: Assuming you are comparing your existing RAID-5 300 GB drives to XIV's 1 TB (1000 GB) drives, the floor space savings is calculated as follows:

$$\% \text{ Savings in Floor Space} = \{1 - 0.8*(300) / 0.5*(1000)\} * 100 = 52\%$$

To calculate your actual savings, multiply the resulting percentage by the real cost of the floor space currently in use by these disk drives at your data center.

Your potential savings may be even greater should your data center be at its physical limits. Switching to the XIV system may allow renovations or a move to a larger facility to be delayed or avoided, saving the expense and disruption such a change would entail.

Power and cooling savings

High density drives are also more efficient to operate. The XIV system's ability to meet the same needs with fewer drives and through greater per drive efficiency creates corresponding cost savings in power and cooling.

To determine the savings in power and cooling from replacing your Medium Density Fast Rotation (MDFR) drives with Very High Density Slower Rotation (VHDSR) drives, use values relevant to your storage operation in the formulas below.

- ▶ To compare RAID-1 drives to XIV's VHDSR drives, use the following formula:

% Savings in Power & Cooling = {1 - 0.75 (Your MDFR Drive Size) / (the XIV system VHDSR Drive Size)} * 100

For example: Assuming you are comparing your existing 146 GB drives to XIV's 1 TB (1000 GB) drives, the power and cooling savings are calculated as follows:

% Savings in Power & Cooling = {1 - 0.75 (146) / (1000)} * 100 = 89%

- ▶ To compare RAID-5 drives to XIV's VHDSR drives, use the following formula:

% Savings in Power & Cooling = {1 - 0.75 [0.8 (Your MDFR Drive Size)] / [0.5 (the XIV system VHDSR Drive Size)]} * 100

For example: Assuming you are comparing your existing 300 GB drives in a RAID-5 configuration to XIV's 1 TB (1000 GB) drives, the power and cooling savings are calculated as follows:

% Savings in Power & Cooling = {1 - 0.75 * [0.8 * (300)] / [0.5 * 1000]} * 100 = 52%

- ▶ To compare 450 GB RAID-5 drives to XIV's 1 TB (1000 GB) drives, use the following formula:

% Savings in Power & Cooling = {1 - 0.75 * [0.8 * (450)] / [0.5 * 1000]} * 100 = 46%

As noted, to make similar calculations for a RAID-6 (4+2) array, simply replace the 0.8 coefficient in the formulas by 0.75.

Use your own figures to calculate your potential savings. Apply the percentage that you calculate to your data center's current utility bills to truly grasp the significance that these savings can have on your annual budget. The results may astound you.

We invite you to contact your IBM/XIV representative to learn more.