



# **IBM XIV Storage System: A powerful storage infrastructure for business analytics**

*Including explanation of key aspects of  
XIV technology*

*Rami Elron  
Senior Technical Staff Member  
XIV Product Technical Director*

*IBM Systems and Technology Group  
May 2014*

## Table of contents

<b>Abstract</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>2</b>
<b>The value of analytics</b> .....	<b>2</b>
<b>Storage requirements from analytics have never been more demanding</b> .....	<b>4</b>
An emerging requirement for scalable management.....	4
Management scalability should encompass several key aspects: .....	5
<b>XIV's unique value proposition for analytics</b> .....	<b>6</b>
<b>How XIV's architecture addresses analytics requirements</b> .....	<b>7</b>
XIV is designed to support highly available analytics .....	7
XIV is designed to support high performance analytics.....	9
XIV is designed to support highly scalable analytics .....	11
XIV is designed to support low TCO analytics.....	12
<b>Summary: storage infrastructure matters</b> .....	<b>14</b>
<b>Resources</b> .....	<b>14</b>
<b>Acknowledgements</b> .....	<b>15</b>
<b>About the author</b> .....	<b>15</b>
<b>Trademarks and special notices</b> .....	<b>16</b>

## Abstract

---

*An increasing number of organizations are reporting improved decision-making and risk management following the deployment of analytics solutions. Extraction of business insights from raw data is however a non-trivial challenge, compounded by the extraordinarily large amounts of data that need be processed. An effective storage system represents a cardinal part of any solution set to overcome such a challenge.*

*IBM XIV is a powerful storage system whose technology lends itself extremely well to accommodating analytics requirements. This paper reviews key aspects of XIV's technology and explains how they combine to deliver a robust, high-performing, and feature-rich storage infrastructure for analytics applications.*

# Introduction

---

The current business landscape is subject to a tidal wave of change, driven by an unprecedented growth rate of business-related data, an increasing business reliance on mobile communications and connectivity, and a shift in the way clients and customers interact, notably via social media. Studies and market observations show that companies taking an analytic approach to data are likely to reap noticeable business benefits compared to companies not doing so; yet, the exceptional volume of business-related data represents a major challenge to organizations wishing to stay ahead of the competition.

The field of business analytics - or analytics in short - has been recently enjoying a remarkable surge in popularity. An increasing number of organizations are reporting improved decision making and risk management following the deployment of analytics solutions, enabling higher utilization of resources and better business performance. Extraction of business insights from raw data is however a non-trivial challenge and potentially resembles an excavation project in search of a valuable ore - both calling for a substantial investment in time, cost and labor. The challenge is compounded by the extraordinarily large amounts of data that need be processed, coupled with a need to satisfy high responsiveness demands of competitive business and correlate between numerous data formats. An effective storage system represents a cardinal part of any solution set to overcome such challenge.



**Figure 1:** An evolving business landscape is underpinned by an explosion of data and new technologies

IBM XIV is a powerful storage system for enterprise and cloud environments whose technology lends itself extremely well to accommodating analytics requirements. This paper reviews key aspects of XIV's technology and explains how they combine to deliver a robust, high-performing, and feature-rich storage infrastructure for analytics applications.

## The value of analytics

---

**Analytics can maximize a business' ability to understand how business decisions can impact business success.** Analytics applications operate on raw and aggregated data to produce insights that may help make better business decisions and maximize business benefit. A growing number of

organizations experience pressure to react faster to the dynamic business arena and **adopt a more proactive approach to match increasing customer expectations** – in an effort to gain a competitive advantage. Such companies are increasingly leveraging analytics to identify meaningful patterns in their data, extract new insights and optimize their business processes.

Analytics applications can help businesses become more **agile, allowing them to:**

- Increase **business opportunities** thanks to discovery of trends and patterns of behavior
- Improve **decision making**
- Support better **prioritization and targeting**
- Drive better **development of next generation products** and services
- **Save money**

Analytics can facilitate an ability to make timely, informed, fact-based decisions and predict trends. This can potentially minimize risk and lead to a dramatic improvement in customer satisfaction and operational efficiency.

## Storage requirements from analytics have never been more demanding

---

An **analytic approach to data** blurs boundaries traditionally used to distinguish 'current' from 'historical' data, **calling for more convenient and quick access to both new and old data**. The large (and growing) amounts of data that need be analyzed, together with benefits associated with parallel processing and an ever-increasing requirement to minimize IT costs, compel companies to examine and rethink the approach employed for data storage – with special emphasis on the following needs:

- **High availability and resiliency:** support for high system stability and risk mitigation encompassing effective recovery from disk failures, snapshot capability and enterprise-class disaster recovery
- **Consistent and predictable high performance:** ability to attain consistent high performance across all analytical tasks with minimized need for performance optimization
- **High Management Scalability:** ability to accommodate extended management requirements driven by analytics without requiring lengthy, high-effort data migration projects
- **Minimized total cost of ownership:** effective storage capacity utilization and minimized administration that result in low staffing and training requirements

The significance of these needs is pronounced with storage system management, where administrators are taxed with administration chores concerning system setup, planning, capacity provisioning and performance optimization. Traditional storage systems are often plagued by poor capacity utilization, challenging performance balancing, substantial power and space requirements, high software license costs and excessive administration requirements concerning storage provisioning, monitoring and migration. Compared with a traditional storage solution not designed for analytics, a storage system optimized for analytics would be able to demonstrate higher storage utilization, and support SLAs while increasing utilization. A storage system optimized for analytics would also be able to minimize the need for high-skilled administrators and reduce staffing along with training requirements.

### An emerging requirement for scalable management

Effective analytics calls for scalable storage management. High end storage solutions have compelled customers to consider implementation of a single large storage system in an attempt to accommodate scalability needs warranted by various application profiles. However, large as they may be – traditional storage solutions are unable to support unlimited growth. Traditional storage solutions were also not designed to facilitate management of co-existing storage solutions (even from the same product-line) in an effective or efficient manner. Consequently, the introduction of a new storage system often results in scalability challenges regardless of the maximum capacity limit supported, due to its different management paradigm and a potentially significant investment in data migration and training. Scalable management has become a major challenge in a quest to achieve effective system scalability.

## Management scalability should encompass several key aspects:

- **Administration**: minimized administration and automated operations - are select characteristics expected of a scalable approach to storage system management
- **Security**: segregated administration represents an important element of a scalable management approach, facilitating appropriate and easy policy/role-based assignment of rights to administrators for data across the storage environment
- **Alerts, reporting and troubleshooting**: central control and consolidated management views are two key examples of features that can tremendously ease the management of growing storage environments
- **Setup**: minimal setup and configuration procedures – for both initial setup and expansion use cases - are notable characteristics expected of a scalable approach to storage
- **Training**: the system's ability to effectively accommodate increased demands should not come at the price of challenging management complexity and a steep learning curve.

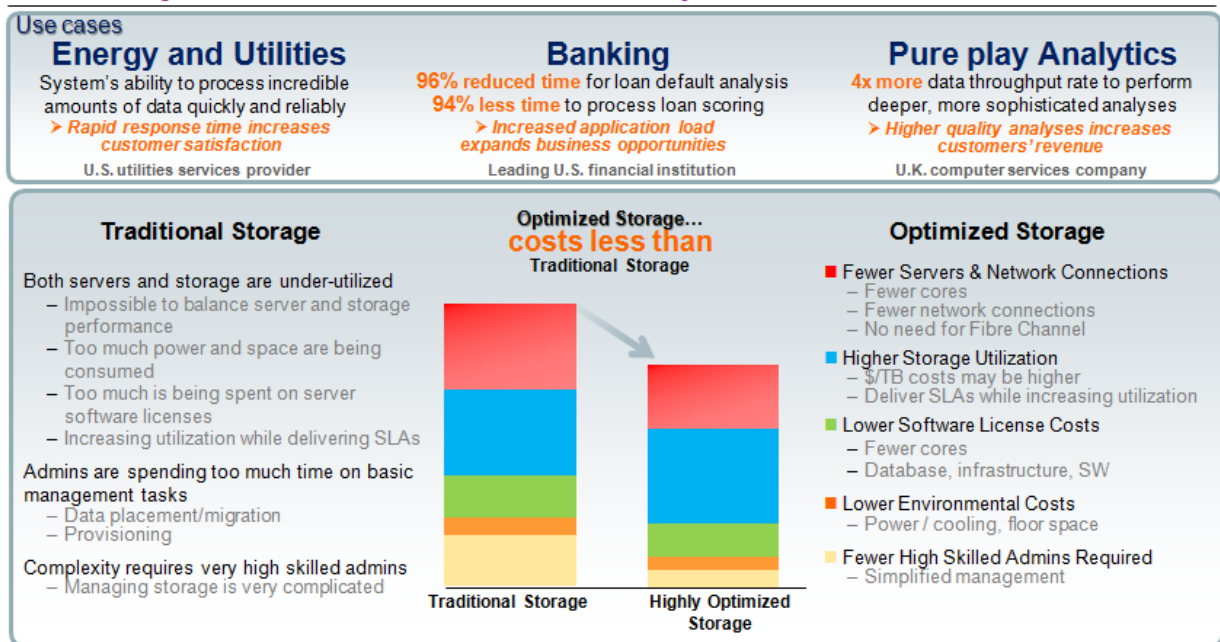
# XIV's unique value proposition for analytics

XIV is a storage system optimized for analytics, featuring characteristics that translate into a powerful infrastructure for addressing analytics-related demands:

- **A groundbreaking architecture** that guarantees high availability and high consistent performance for diverse and dynamic workloads
- **An advanced approach to scalability** that shatters traditional barriers associated with scalability in general and management scalability in particular
- **A compelling management experience and TCO** value proposition

In the past businesses have been compelled to consider dedicated storage solutions for their analytics applications, incurring additional costs, requirements for excessive data migration and a substantial administrative effort overall.

Conversely, XIV's technology maximizes value by enabling businesses to utilize XIV as a general purpose foundation for analytics AND other workloads.



**Figure 2:** XIV's unique capabilities translate to exceptional value for analytics, as demonstrated by customer use cases in various industry sectors. XIV's technology is optimized for analytics, allowing for lower costs, better utilization and fewer staffing requirements than traditional storage solutions. Select use case references are provided at the end of this paper

XIV has been deployed successfully by major organizations in various sectors, including energy and utility, banking, and pure play analytics, among others. Check here to see how a major bank significantly increased performance of SAS® analytics tasks with XIV: <http://www.sas.com/success/bofagrid.html>. (Additional references are provided at the end of the paper).



## How XIV's architecture addresses analytics requirements

---

XIV is not a typical high end storage system. Its unique architecture enables a blend of powerful capabilities that set it apart from traditional storage solutions and make XIV exceptionally well-suited for analytics use cases.

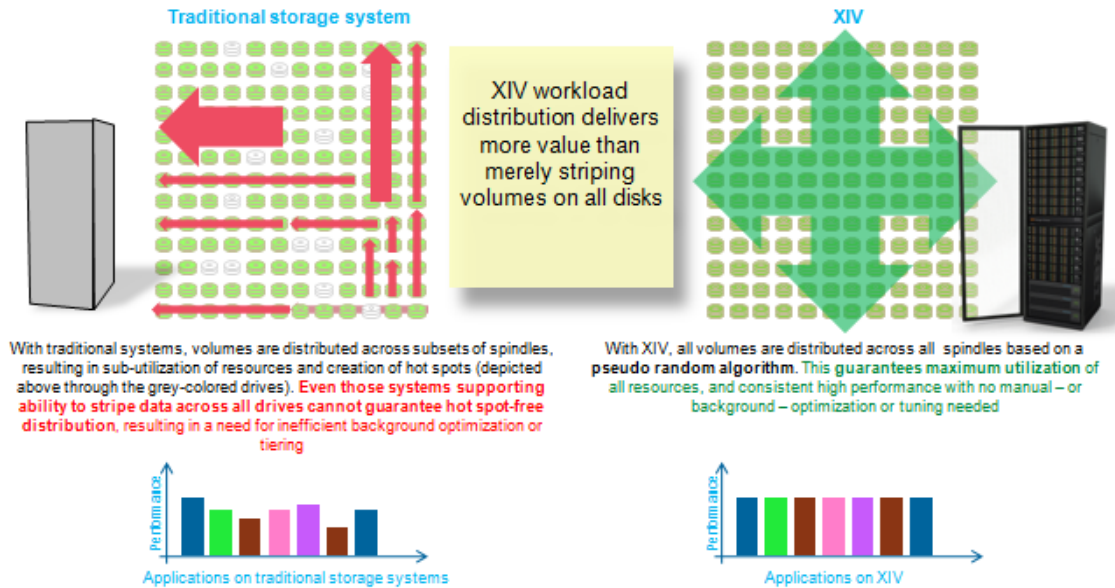
The XIV system features a grid-based architecture. It is supplied in a rack enclosure, and is equipped with modules implemented in an active-active configuration. Two types of XIV modules are featured - *Interface* modules and *Data* modules. Each module type come with multi-core processing, up to 48 gigabytes of DRAM cache and up to 48 terabytes of near-line raw storage (XIV's *Flash Caching* implements an optional SSD drive within each module as extended cache, enabling a potential performance boost for cache-intensive applications). Interface modules also feature external networking connectivity (FC or iSCSI). Modules communicate with each other over an internal InfiniBand network. For a detailed review of XIV's architecture please look here: <http://www.redbooks.ibm.com/abstracts/sg247659.html>.

A fully-equipped XIV rack features 15 modules (6 Interface modules, 9 Data modules), up to 325TB of usable capacity using 180 Near-Line 4 TB SAS drives, 720 gigabytes of DRAM, and up to 12 terabytes of Flash Caching. Notwithstanding its formidable processing power, what makes XIV especially suited for analytics is its ability to effectively harness the performance potential, guarantee consistent high performance with zero tuning for a diverse mix of applications, and maximize performance for analytics applications – which benefit from parallel processing. But XIV goes even farther and combines such power with an unprecedented management experience yielding TCO benefits that are may not be realized using traditional storage architectures. Running analytics applications in parallel with other applications using a single storage solution can significantly reduce storage capacity requirements, data migration needs, and administrative effort – contributing to an already-compelling value proposition for analytics.

### **XIV is designed to support highly available analytics**


XIV's advanced data distribution enables high system availability featuring characteristics that are not offered by traditional architectures.

Traditional storage solutions commonly place all volume capacity within a subset of the system's data modules (often within a single hardware module). Such approach exhibits several limitations that can adversely affect availability and performance, notably - a small spindle-count per volume, a long rebuild duration following drive failure (hours, and even days) and a higher risk for data unavailability due to environmental factors affecting certain modules rather than the system in general (e.g., heat), among others. Traditional storage solutions may also feature support for standby spare drives that are employed if failed system drives need be replaced and their data needs be rebuilt. Such drives are, however, not often subject to health monitoring examination before their actual usage, thereby increasing likelihood for drive problems to surface when the drive's health is critical for recovery. Also, once staged, spare drives are subject throughout the data rebuild process to high workload spikes that might adversely impact the drive's performance and health.



**Figure 3:** Traditional data distribution vs. XIV's advanced data distribution. XIV's technology implements an advanced approach for data distribution which obviates hot spots – rather than requiring manual or background tuning processes to be run, as is commonly done with traditional storage solutions.

On XIV, the volume data is stored on one-megabyte-sized data 'partitions' that are distributed across all system drives and modules in a pseudo-random fashion, with each partition being mirrored to another drive on a separate module. Such approach results in a large spindle-count per volume (180 drives with a 15-module XIV system), an extremely short data rebuild duration (<15 minutes per 1TB of data) and a lower risk for impact by environmental factors affecting a subset of system modules. As for spare storage, XIV features the concept of hot spare **capacity**, pre-allocating free capacity on all drives rather than dedicating standby drives for data rebuild scenarios only. Such approach allows for timely health inspection and monitoring of all spare capacity, and minimizes the likelihood of drive failure due to extreme workload spikes. It also facilitates higher storage performance as the system can leverage a larger number of spindles.



**Figure 4:** XIV's core technology provides a compelling value proposition for analytics workloads. Featuring a grid-based, massively parallel, and distributed architecture, XIV enables exceptional performance, availability and ease of use that is not attained with traditional storage solutions.

XIV features a powerful implementation of snapshots featuring advanced functions such as writable snapshots, and the ability to take a snapshot of another snapshot, restore a snapshot from another snapshot and even take a snapshot of a consistency group. XIV additionally features special snapshot functionality with mirroring (*Snapshot Mirrors*) and support for snapshot coordination across multiple XIV racks (*IBM Hyper-Scale Consistency*).

XIV also features support for advanced synchronous and asynchronous replication functionality, including volume or consistency group replication, independent replication settings per volume/consistency-group, independent RPO and interval settings, minimal RPO of 30s (the minimal replication interval is 20s), snapshot mirroring, and support for streamlined failover and failback. Importantly, XIV features an extremely easy setup and management replication paradigm that minimizes training and administration requirements.

## **XIV is designed to support high performance analytics**

Analytics applications call for high performance storage that can accommodate both sequential and non-sequential workloads equally well. XIV's technology and parallelism enable extremely high performance for analytics applications. XIV deployments show great performance across multiple types of analytical workloads, and have demonstrated the ability to attain high data throughput rate with no contention between analytics tasks.

One of the most critical factors responsible for XIV's exceptional performance is its special data distribution. Each volume is distributed on XIV across all system drives in a manner that obviates hotspots, resulting in **absolutely no need to ever tune the system – either manually or through a background system process**. To portray this from another perspective - customers will experience with XIV high performance right from the moment the system is deployed and will continue to enjoy high performance later on, with no tuning needed. XIV's architecture enables the system to recover from a disk failure significantly faster than traditional systems and impressively – doing so with a negligible impact on the performance of running applications. These remarkable capabilities probably explain best why XIV is able to satisfy so effectively performance requirements of multiple, mixed, and diverse workloads that are characteristic of analytics profiles.

XIV's architecture manifests special value whenever new modules are added as well: adding a module adds storage capacity along with a carefully matched allotment of processing power and cache. Critically, perfect load balancing is attained regardless of the addition, deletion or resizing of volumes; XIV will only redistribute data upon changes to hardware availability – e.g., addition of a new module.

XIV's patented cache management is one of the salient factors behind XIV's stellar performance. XIV systems feature an advanced caching implementation and are equipped with large amounts of DRAM cache (a 15-module rack is equipped with 720GB). XIV's *Flash Caching* enables flash drives to be installed as an extension for read cache (up to 12TB with a 15-module rack). XIV's Flash Caching implementation is designed to effectively and efficiently accommodate diverse workloads potentially demanding extremely high performance without a need for physical tiering. With XIV Flash Caching, flash drives are placed in the data path between the system's near-line drives and the DRAM cache. For a detailed review of XIV's Flash Caching implementation look here:

<http://www.redbooks.ibm.com/abstracts/redp4842.html>.

XIV has demonstrated notable performance benefits using flash for analytics use cases. There are nevertheless analytics workloads which are predominantly large sequential I/O in nature, where the advantage of flash is arguably less pronounced. Since XIV is designed to effectively accommodate high sequential performance requirements using near-line drives alone, sequential reads that are larger than 64K will bypass Flash Caching, allowing the system to dedicate flash resources to non-sequential workloads. Administrators can also specify the volumes whose data will be cached through Flash Caching.

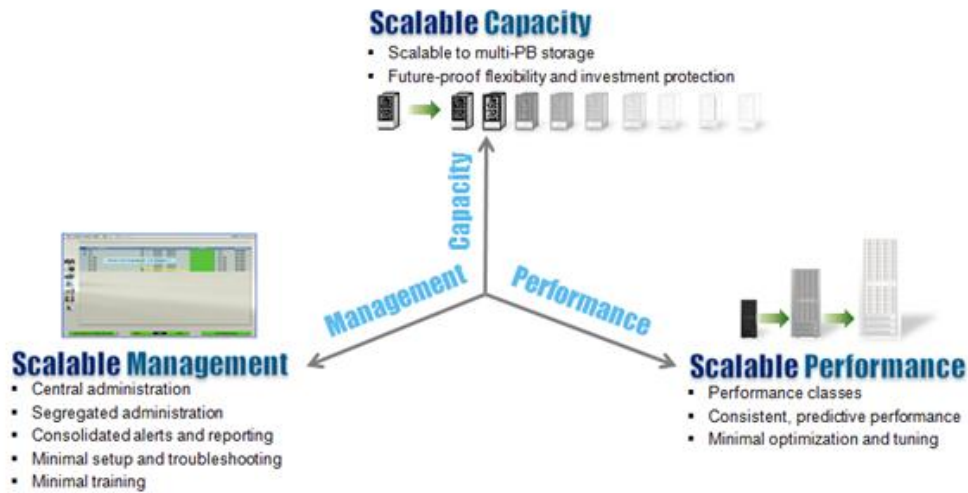
The term "Powerful" alone doesn't really do justice to XIV's remarkable architecture advantages, such as its caching. XIV's advanced caching scheme is effective, flexible, and smart:

- **Effective:** XIV's advanced caching approach enables concurrent service by all cache units. Rather than being managed as a single pool, cache is distributed across all modules, with each module's cache serving IO related to the module's storage. Each module caches data that is stored on its own disks, and the high cache-to-disk bandwidth enables aggressive pre-fetching and cache management. Thanks to XIV's effective data distribution the system does not require special planning to match the 'right' cache capacity with storage, and to realize linear performance scaling as modules are added.
- **Flexible:** XIV implements a flexible caching method that supports both large and small cache slots, unlike other caching methods than compromise on a single slot size. Employing only large cache slots can help increase performance, but can also make it more difficult to manage the precious cache area efficiently as only some part of the cached data ends up being used. Conversely, employing only small cache slots sacrifices performance for management flexibility. XIV's caching implementation offers the combined value of large and small cache slots: the system employs a large slot size when reading data from disk (which is dynamically increased depending on the I/O pattern), yet manages cached data in small slots for efficient, flexible cache capacity utilization that ultimately leads to better performance.
- **Smart** – Each XIV module maintains knowledge concerning the system's data distribution; modules can therefore send messages to each other concerning data partitions that are projected to be accessed in the near future (based on tracked workload patterns) and should be preemptively fetched from disk to cache; this is highly beneficial for sequential workloads common to analytics use cases. In addition, XIV's patented implementation of Flash Caching also minimizes wear leveling effects associated with flash drives, ensuring high performance, reliable operation and longevity.

XIV's caching implementation optimally leverages XIV's data distribution and parallelism, resulting in a storage system that is able to maximize the effectiveness of caching and maintain extremely high performance and low latency for analytics applications.

## XIV is designed to support highly scalable analytics

*Storage scalability* is commonly used as a reference to the maximum *capacity* and *performance* supported by a storage system. That said, the burgeoning significance of analytics and an increased pressure to accommodate analytics applications at minimal cost accentuate the significance of another scalability-related aspect – *manageability*. Management scalability encompasses administration, security, alerts, alerting, reporting, troubleshooting, setup and training. With the requirement to accommodate ever-increasing amounts of data stored across multiple physical frames (often extending across multiple sites) and an emerging need to apply analytics on such data– management scalability requirements are pronounced.



**Figure 5:** IBM Hyper-Scale technology addresses three major pillars of scalability required by analytics applications – Capacity, Performance and Management. XIV’s implementation of Hyper-Scale technology allows for unprecedented power and flexibility when managing extremely large capacity

IBM Hyper-Scale is a family of new technologies that underlies an innovative approach to storage scalability and offers substantial advantages to customers. At its core, **Hyper-Scale is designed to treat all storage as a single, large hyper store**. Featuring IBM Hyper-Scale technology, XIV enables a novel paradigm of management scalability that addresses challenges that storage administrators have been dealing with for years. Effective scaling of data storage for analytics represents a key challenge for all enterprises.

IBM Hyper-Scale consists of 3 main pillars: IBM Hyper-Scale Manager – a comprehensive management application; IBM Hyper-Scale Mobility – a powerful functionality to move volumes between storage containers with no disruption to host applications; and IBM Hyper-Scale Consistency – a function enabling the creation and management of a volume group spanning multiple frames. XIV sports features leveraging all IBM Hyper-Scale technologies: a management UI implementing the IBM Hyper-Scale Manager and featuring XIV’s acclaimed UI; support for cross-system consistency groups implementing IBM Hyper-Scale Consistency; and an online migration facility for volumes implementing IBM Hyper-Scale Mobility.

Hyper-Scale Mobility is designed to provide online, non-disruptive data mobility between storage containers in a manner that is transparent to host applications. Online data mobility is increasingly becoming an effective approach for solving key customer needs in the modern data center and cloud environment, including challenging over-provisioning scenarios, optimal re-location of data to improve applications performance, and hardware refresh requirements. Online data mobility extends benefits associated with tiering beyond a single system to multiple systems. Hyper-Scale Mobility facilitates data migration by providing access to a volume through multiple storage controllers as a single, multipath volume on the host and controlling which storage controller accepts I/O from the host. This helps maximize utilization and effective performance balancing across multiple frames, facilitates aggressive thin-provisioning, and helps customers streamline migration between system versions. For additional information on IBM Hyper-Scale and its XIV implementation look here:

<http://public.dhe.ibm.com/common/ssi/ecm/en/tsl03121usen/TSL03121USEN.PDF>.

## **XIV is designed to support low TCO analytics**

In an ever-challenging quest to minimize IT costs, organizations are expecting storage systems to demonstrate a compelling TCO value proposition. An ability to accommodate analytics' requirements for high storage performance and scalability at a low TCO – let alone doing so for both production and analytics environments using a single storage solution - enables reduced administrative overhead and response time – as well as maximized agility and cost savings. XIV is unique among storage systems in its ability to deliver performance associated with high-end solutions at cost levels typically associated with mid-range solutions. Such value is made possible thanks to XIV's grid technology, parallelism, caching and special data distribution. XIV's architecture was designed to attain high performance using low cost near-line storage, realizing an unprecedented low TCO for high-end systems.

XIV is exceptionally efficient and features a powerful implementation of thin provisioning. The feature is supported at the storage pool level, allowing for the creation of storage pools that are thinly provisioned along with pools that are not. The system features a comprehensive UI for managing and monitoring capacity utilization as well as related alerts and also features support for Thick-to-Thin migration, allowing for potentially large savings in capacity when migrating volumes from legacy storage to XIV. In addition to XIV's efficient support for thin provisioning, XIV also supports space reclamation with various environments (VMware, Microsoft and Symantec)

XIV's remarkable ease of use contributes greatly to management efficiency, enabling effective management of a multitude of XIV system at a fraction of the staffing and training requirements warranted by traditional storage systems. XIV's architecture and single tier approach obviate planning and administration operations traditionally required for RAID group setup and management, volume and snapshot capacity allocation, and tiering configuration and management. Notably, XIV obviates any manual or background tuning.

✓ **Exceptional PERFORMANCE**

- Scalable grid architecture
- Tuning-free
- Any workload combination anytime
- Flash on board

*“The main thing that attracted us about XIV was the performance... we would be able to rely on **consistent performance**... which would give us more time to focus on running more **sophisticated analytics** that would deliver **increased value** for our clients.”*

✓ **Exceptional AVAILABILITY**

- Sophisticated data distribution
- Self healing
- Unparalleled disk rebuild times
- Enterprise-class availability

✓ **Exceptional SIMPLICITY**

- Simple planning, provisioning, and agility
- Connect and forget

**Figure 6:** XIV's key value points for analytics include exceptional performance, availability and simplicity

XIV's warrants minimal training to effectively set up and manage, and provides automated processes to facilitate data entry procedures. XIV's Graphical User Interface presents a simple yet powerful view of the storage system and outstanding operational status. XIV's UI consolidates data from multiple systems, simplifying the detection of issues calling for attention, and abetting task prioritization. The system also features support for mobile platforms that allows administrators to monitor system performance on the go using mobile devices. The XIV GUI has been critically acclaimed as an industry benchmark for storage management ease of use, and is a major reason behind the system's low training and staffing requirements.

Last but not least, all XIV software features – advanced snapshots, replication, QoS, security, thin-provisioning, host connectivity, virtualization and others – come included with each XIV system. Considering that Direct Attached Storage systems and traditional SAN solutions typically offer similar functions at a considerable high extra charge – XIV's packaging represents an attractive offering further contributing to a significant value proposition for analytics.

## Summary: storage infrastructure matters

---

While analytics is rising in popularity, customers may still find the acquisition of a specialized infrastructure for analytics to be prohibitively high and difficult to justify. XIV represents a powerful storage infrastructure for analytics applications, featuring an architecture that enables the system to demonstrate high availability, performance, and scalability.

Importantly, XIV's architecture is designed to effectively accommodate a mixed application environment, allowing customers to consider whether they really need to dedicate a storage system for analytics purposes at extra costs and considerable administration overhead. XIV's architecture lends itself extremely well to supporting diverse and dynamic application workloads running in parallel - which is very challenging to accomplish by a traditional storage solution. XIV deployments supporting analytics applications highlight the ability to successfully support analytics workloads and other application data using the same storage infrastructure. As a result, XIV customers are compelled to place multiple diverse workloads on systems and reap the benefits associated with its advanced management paradigm.

XIV's ability to sustain consistent high performance for mixed workloads (e.g., large sequential, small random and others) with no manual or background optimization is a major advantage for all applications – and especially for analytics. XIV's administration simplicity results in a substantial reduction of administrative setup, management, and other operations compared to traditional high-end storage systems. Finally, XIV's efficient storage capacity utilization within a small physical footprint, low power requirements and minimal management requirements, which keep operational, staffing and training costs extremely low – translate into a low total cost of ownership – contributing to an exceptional value proposition.

## Resources

---

The following links provide useful references to supplement the information contained in the paper.

- **Video: Bank of America with XIV and SAS -- by IBM Business Partner Destiny:**  
<https://www.youtube.com/watch?v=KSTJLwJk1Y8>
- **Bank of America avoids gridlock in credit-risk scoring, forecasts using SAS®**  
<http://www.sas.com/success/bofagrid.html>
- **Sumerian accelerates the delivery of sophisticated analytics**  
<http://ibm.co/1azLqI7>
- **Statistics Finland counts on IBM XIV for large-scale decision support**  
<http://ibm.co/1jxmKTl>
- **IBM XIV Storage System Gen3 Architecture, Implementation, and Usage**  
[ibm.com/redbooks/abstracts/sg247659.html](http://ibm.com/redbooks/abstracts/sg247659.html)
- **Solid-State Drive Caching in the IBM XIV Storage System**  
[ibm.com/redbooks/redpieces/abstracts/redp4842.html](http://ibm.com/redbooks/redpieces/abstracts/redp4842.html)
- **Analyst paper: In&Out on XIV Gen3 Grid Storage**  
<http://ibm.co/1fnfBjp>



## Acknowledgements

---

Author would like to thank the following team members for their help, guidance and technical review of the paper:

Alain Azagury, Director of SWG Technical Strategy, IBM Academy, Member

Vicente Moranta, Business Line Manager - Cloud, Object and Big Data Storage, IBM Systems & Technology Group

Ed Segura, WW Sales Executive, XIV, IBM

Eyal Abraham, Manager, Storage Solution Sales, IBM Systems & Technology Group

Eitan Hadayo, Business Solution Architect, IBM Systems & Technology Group Big Data & Analytics

Rivki Matosevich, XIV Architect, IBM Systems & Technology Group

Diane Benjuya, Market Segment Manager, XIV, IBM Systems & Technology Group

Elan Freedberg, Market Segment Manager, XIV Storage, IBM Systems and Technology Group

## About the author

---

**Rami Elron** is an IBM Senior Technical Staff Member and the XIV Product Technical Director within the IBM Systems and Technology Group. He is a subject matter expert in storage architecture, mirroring, and security. He has been helping drive technical strategy and innovation for the IBM XIV Storage System series since 2008. He can be contacted at [relron@il.ibm.com](mailto:relron@il.ibm.com)

## Trademarks and special notices

---

© Copyright IBM Corporation 2014.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. 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 "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.