



# **IBM Hyper-Scale: A powerful new approach to scaling storage management**

*Including explanation of related technologies  
and spotlighting the IBM XIV Storage System  
implementation of IBM Hyper-Scale*

*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>Scalable storage in the 21<sup>st</sup> century</b> .....	<b>2</b>
<b>Scalable storage management: going beyond capacity and performance requirements</b> .....	<b>3</b>
<b>Traditional storage solutions: not a panacea</b> .....	<b>4</b>
<b>IBM Hyper-Scale: a powerful vision for scalable storage in the data center</b> .....	<b>4</b>
<b>The implementation of IBM Hyper-Scale technologies in XIV</b> .....	<b>5</b>
Spotlight: XIV with IBM Hyper-Scale Manager .....	5
Spotlight: XIV with IBM Hyper-Scale Mobility .....	7
Spotlight: XIV with IBM Hyper-Scale Consistency.....	9
<b>Transforming storage management: the special value proposition of XIV featuring IBM Hyper-Scale</b> .....	<b>9</b>
<b>Summary</b> .....	<b>10</b>
<b>Resources</b> .....	<b>11</b>
<b>Acknowledgements</b> .....	<b>11</b>
<b>About the author</b> .....	<b>11</b>
<b>Trademarks and special notices</b> .....	<b>12</b>

## Abstract

---

*Effective scaling of data storage represents a key challenge for all enterprises. IBM Hyper-Scale is a new family of advanced technologies that underlies an innovative approach to storage scalability and offers substantial advantages to customers. The combined power of XIV's groundbreaking architecture with IBM Hyper-Scale brings major benefits to customers and redefines the value proposition associated with scalable storage management.*

*This paper presents IBM Hyper-Scale and explains how XIV's implementation of Hyper-Scale technologies can help customers address storage objectives that challenge traditional storage solutions.*

## Introduction

---

Effective scaling of data storage represents a key challenge for all enterprises. IBM Hyper-Scale is a family of new technologies that underlies an innovative approach to storage scalability and offers substantial advantages to customers. This paper presents the vision behind IBM Hyper-Scale, and explains how the IBM Hyper-Scale technologies implemented by XIV are designed to help customers address storage objectives that challenge traditional storage solutions.

## Scalable storage in the 21<sup>st</sup> century

---

The dynamic nature of business and the burgeoning requirement for quick responsiveness result in extremely dynamic profiles of application criticality and usage which directly affect the application's data storage requirements. In the past, large companies were inclined to employ more than a single storage solution to accommodate their data storage requirements, often dedicating a solution to a specific application – or to a group of applications. In recent years the emergence of storage tiering functionality compelled customers to consider implementation of a single large storage system for different application profiles. However, large as they may be – traditional storage solutions were neither designed to support unlimited growth, nor to support easy co-existence with other storage solutions (even from the same product-line). As a result, the deployment of such solutions could end up becoming a major administrative challenge, warranting massive data migration, calling for considerable investment in training – and coming at a significant cost. Sadly, such endeavor would be destined to repeat itself within a few years, once the system is no longer able to accommodate the ever-growing business requirements.

In recent years an unprecedented surge of user- and machine-generated data has been evidenced, along with an increasing reliance of companies on such data. 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. Such approach arguably blurs some of the traditional boundaries between 'current' and 'historical' data, calling for more convenient and quick access to new and old data, and compelling companies to rethink their original approach employed for scaling storage – with special emphasis on the following:

- Effective and efficient scalability of management (including management policies) as capacity grows
- High responsiveness
- Maximized usage of all storage parts within the overall solution
- Minimized need for performance optimization
- Minimized training requirements
- Minimized TCO

# Scalable storage management: Going beyond capacity and performance requirements

Dictionary.com defines **Scalability** as “the ability of something, esp a computer system, to adapt to increased demands”. In the context of a storage system, the term *scalability* is commonly used as a reference to the maximum capacity and performance supported by such system. Nevertheless, as storage grows in capacity, another scalability-related aspect becomes evident– manageability. With the requirement to accommodate ever-increasing amounts of data stored across multiple physical frames, often extending across multiple sites – the need to “adapt to increased demands” is pronounced.

Scalable management warrants special attention to several notable aspects:

- **Administration:** the ability to minimize administration, automate manual operations, and apply standard policies - are select examples of 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 examples of 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.

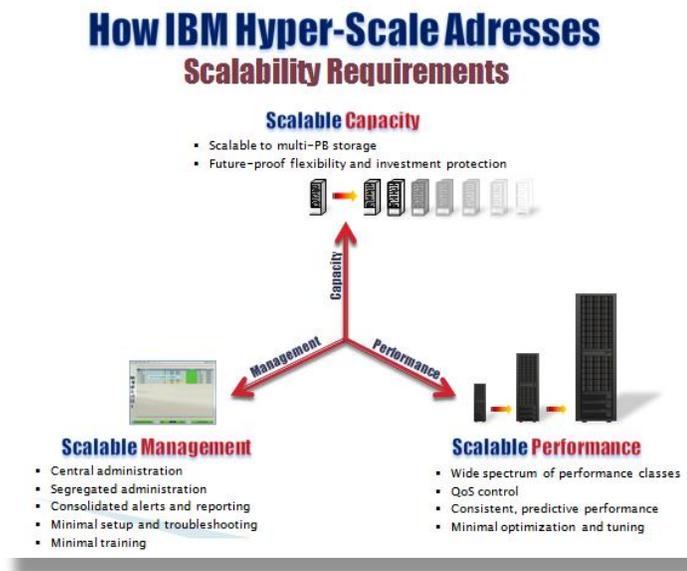


Figure 1: IBM Hyper-Scale addresses three axes of scalability – capacity, performance and management

## Traditional storage solutions: Not a panacea

---

Many companies have more than a single storage solution deployed. In the past, deployment of multiple solutions was warranted to **efficiently** accommodate a range of applications varying in their storage requirements – performance, reliability or others. With the advent of data tiering functionality *within* the storage system, customers have been increasingly considering deployment of storage solutions traditionally known as ‘high-end’ for a variety of application profiles – rather than dedicating such storage to applications strictly considered ‘mission-critical’. Notwithstanding the potential operational value associated with tiering, such value is limited in traditional solutions to the **single system** and may not easily extend across multiple systems – even not across similar systems by the same vendor. Traditional storage systems often regard data movement between systems – a.k.a. data migration - as an *exceptional* use case rather than an ordinary operation occasionally warranted to accommodate quotidian data center requirements, such as workload balancing. Since traditional solutions are not designed to support unlimited scalability, and are also typically not designed to facilitate management of co-existing storage solutions – the introduction of a new storage system often results in scalability challenge regardless of the maximum capacity limit supported:

- **Painful option #1:** If the customer wishes to avoid administration overhead due to the management of multiple systems with different management paradigms, operations and policies, the older storage system will need to be decommissioned, even though it could have continued to support select applications in a satisfactory way. A decision to decommission a system urges migration of data from the old system to the new system – placing further pressure on administration.
- **Painful option #2:** If the customer prefers to not decommission an existing system, the introduction of a new storage system will result in a need to accommodate multiple management policies and operational paradigms, will still warrant lengthy data migration from the existing system[s] to the new system, and will incur an administrative overhead that is likely to increase overall TCO.

Traditional solutions have been forcing customers to make compromises for years. With IBM Hyper-Scale, such compromises are no longer warranted.

## IBM Hyper-Scale: A powerful vision for scalable storage in the data center

---

IBM Hyper-Scale is a family of new technologies from IBM that enables customers to treat storage in a novel way that results in substantial benefits. IBM Hyper-Scale introduces the following concepts and accompanying advantages:

- A single ‘hyper store’ metaphor for all managed storage, that is able to support unlimited growth without incurring the management overhead of a new system. The hyper store is comprised of storage containers – which do not have to feature the same configuration or version
- Ability to add a new system to the hyper store without compromising system functionality

- Ability to preserve investment in an existing system generation: no need to decommission old systems in the hyper store
- Ability to move data between systems in the hyper store without disrupting host applications to accommodate business and administrative needs – e.g., performance optimization, maintenance windows, system repurposing, capacity allocation preferences, etc.
- Ability to maintain a standard, uniform management approach in the hyper store
  - Single administrative login, and central hyper store management
  - Consolidated view of all hyper store aspects
- Ability to manage a volume grouping spanning a large number of frames
  - Ability to group and manage volumes residing on different XIV frames as a single cross-system consistency group
  - Ability to coordinate the creation of snapshots across XIV frames
- Ability to minimize TCO and training requirements
  - Ability to add new systems to the hyper store in a streamlined manner
  - Ability to enforce automation policies across systems within the hyper store
  - Ability to minimize the need for performance optimization – manual or automatic
  - Ability to maximize self-healing capabilities that minimize the need for immediate administration intervention

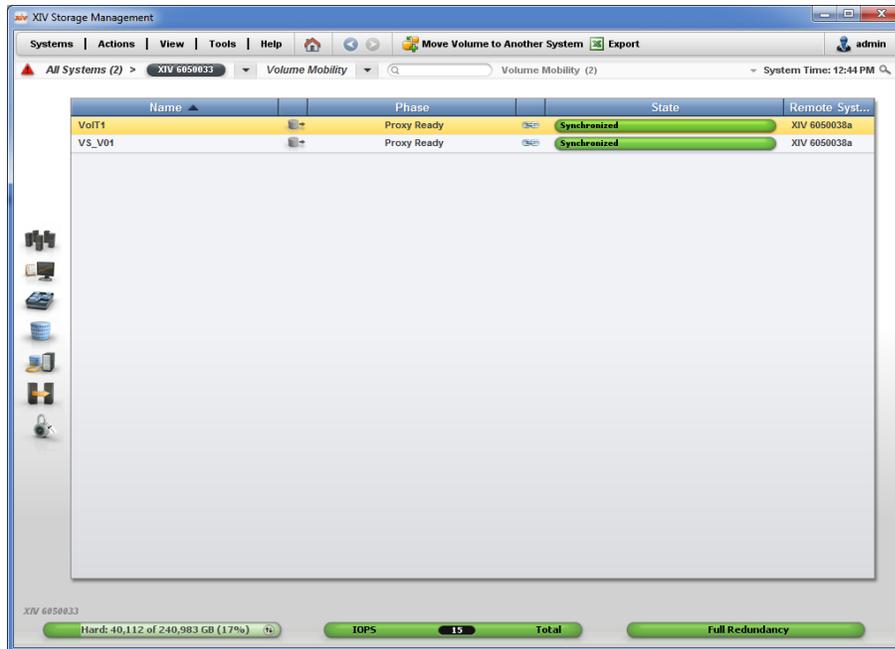
## Implementation of IBM Hyper-Scale in XIV

IBM Hyper-Scale features three major technologies: **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; an online migration facility for volumes implementing **IBM Hyper-Scale Mobility**; and support for cross-system consistency groups implementing **IBM Hyper-Scale Consistency**.

### Spotlight: XIV with IBM Hyper-Scale Manager

*IBM Hyper-Scale Manager* offers a powerful management paradigm for multiple storage containers that realizes unified administration of multiple frames as a single hyper store. It is designed to help customers accommodate capacity growth requirements, along with the benefits of minimized setup, management and training.



**Figure 2:** IBM Hyper-Scale Manager

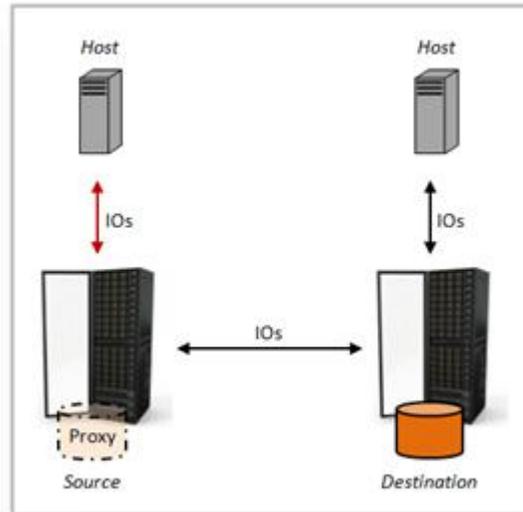
XIV's GUI implements IBM Hyper-Scale Manager technology, making all administration tasks of IBM Hyper-Scale Manager accessible from the XIV GUI, including central configuration of user access rights, hosts and event rules. One of Hyper Scale Manager's key features is its support for extensive capacity reporting. The system collects capacity-related data over time, and enables creation of capacity usage reports for storage pools that help detect the potential risk of reaching designated usage thresholds. Such data can be processed later by analytics applications to reveal insights that could be used to improve capacity planning (see example chart below).

Hyper-Scale Manager's GUI automatically opens the relevant view and highlights new objects after their creation (pool, volume, host, etc). The value behind the GUI's presentation approach is easily evidenced with Hyper-Scale Mobility (see below), where the GUI facilitates movement of volumes between XIV frames. Administrators can view volumes in mobility and their status, along with a raft of GUI enhancements designed to ease and streamline Mobility management.

Hyper-Scale Manager also features support for REST API, providing XIV users with a new and advanced management API that allows retrieval of information from all frames in one result set and enables application of configuration tasks to multiple frames in the hyper store. Hyper-Scale Manager provides an API for Query requests (requests that return object entities managed by the IBM Hyper Scale Manager), and Update requests (requests that create, delete, or update object entity state). It's also worth noting that the IBM Hyper-Scale Manager can be installed either as a virtual appliance or on a Linux (RHEL 6.0 and above) physical server, resulting in increased administration flexibility.

## Spotlight: XIV with IBM Hyper-Scale Mobility

Online data mobility is increasingly becoming an attractive approach to solve 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.



**Figure 3:** In Proxy Mode volume data migrated to a Destination can be accessed indirectly by the related host's applications through the Source acting as a proxy. Before the customer moves out of Proxy mode, a direct connection between the host and Destination should be established – and at that time the host's applications may access the volume data either directly or indirectly. Once Proxy mode is exited the host's applications will only access the volume data directly.

*IBM Hyper-Scale Mobility* is designed to provide online, non-disruptive data mobility between storage containers in a manner that is transparent to host applications. *IBM 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/Os 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.

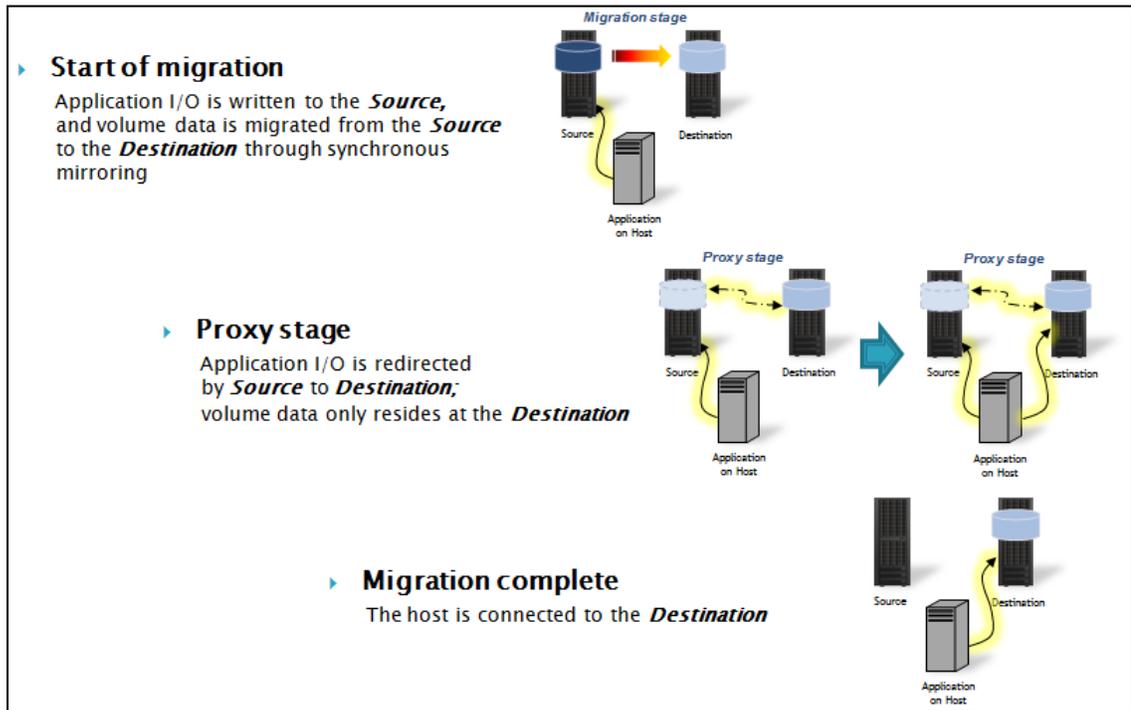
XIV's *Online Volume Mobility feature (OLVM)* implements *IBM Hyper-Scale Mobility* technology, with support for hosts running AIX, Linux, and a variety of Windows Server versions. The feature offers robust functionality with streamlined and automated operation. Volume mobility between the original frame (Source) to another frame (Destination) does not interrupt host activity, and is not disrupted by Rebuild, Redist, Phaseout, or failover on the Destination. **A notable stand-out feature of XIV's support is Proxy mode.** In this mode, volume data migrated from the Source to the Destination is accessed indirectly by the related host's applications through the original frame acting as a *proxy*. The host can remain connected to the Source without a need to zone and move it to the Destination; the host can be moved after the migrated volume data is located on the

Destination. Different volumes can have different Source and Destination controllers. A controller can be either a Source or a Destination per volume. Hosts can access the volume from either controller – and a single host can be connected to both controllers.

XIV's implementation of IBM Hyper-Scale Mobility features special advantages:

- Requires no special hardware
- Requires no special licensing
- Responsibility for data integrity is not placed upon the user
- Streamlined migration of mirroring with Offline Initialization
- Streamlined cluster volume migration
- Proxy mode

XIV's powerful implementation of Hyper-Scale is further bolstered by XIV's new management UI, based on *IBM Hyper-Scale Manager* technology.



**Figure 4:** Migration phases

## Spotlight: XIV with IBM Hyper-Scale Consistency

A growing number of organizations have multiple storage systems hosting inter-related data stores, whose snapshots are required to be consistent across the pertinent systems.

*IBM Hyper-Scale Consistency* enables an administrator to establish a volume consistency group spanning multiple XIV frames – called a *Cross-system Consistency Group*, or an *XCG*. Hyper-Scale Consistency also enables administrators to coordinate the creation of snapshots across XIV frames. XIV's implementation of Hyper-Scale Consistency features a powerful command set for XCG administration, including the ability to stop and resume I/O to an XCG at will. This functionality may also be leveraged by IT developers, ISVs and IBM partners to develop custom data protection solutions.

As demonstrated by the previously-mentioned IBM Hyper-Scale technologies, Hyper-Scale Consistency enables an administrator to extend the value of proven XIV features (in this case – XIV's Consistency Group) across a multitude of frames.

## Transforming storage management: The special value proposition of XIV with IBM Hyper-Scale

---

The power of IBM Hyper-Scale technologies combined with XIV's unique architecture translates to a substantial value proposition for customers. The XIV system features predictable and consistent high performance that requires no performance tuning – manual or automatic. Adding the system's support for self-healing, XIV represents the optimal storage container within a hyper store, enabling minimized TCO. Following are notable advantages that the XIV system featuring IBM Hyper-Scale brings to customers, compared to traditional high-end solutions:

- Traditional solutions have an upper capacity limit which cannot satisfy alone the requirements of most large enterprise customers in years to come. XIV with IBM Hyper-Scale has no such architectural scalability limits.
- Traditional solutions represent an island of storage and are not architected to seamlessly co-exist with other storage solutions from the same product line, or different versions of the same solution. XIV with IBM Hyper-Scale encourages usage of different solution generations as different tiers.
- Traditional solutions warrant a complex migration project. XIV with IBM Hyper-Scale features streamlined online migration functionality.
- Traditional solutions do not guarantee consistent, predictive storage performance for applications. The integration of IBM Hyper-Scale with XIV scales XIV's staple benefits to a multitude of systems.
- Traditional solutions warrant considerable planning for volume setup, snapshots, thin-provisioning and QoS. The integration of IBM Hyper-Scale with XIV enables application of XIV's groundbreaking technology to extremely large storage environments.

## Summary

IBM Hyper-Scale is a family of novel technologies by IBM designed to dramatically increase the customer's ability to scale storage. IBM Hyper-Scale features technologies that can transform the way scalable storage is managed, facilitating greater scalability without compromising the functional aspects of storage management. XIV's implementation of IBM Hyper-Scale is reflected in three cardinal features – IBM Hyper-Scale Manager, IBM Hyper-Scale Mobility and IBM Hyper-Scale Consistency. XIV's implementation of IBM Hyper-Scale Manager offers a powerful consolidated management view designed to accommodate large XIV deployments, and to ultimately realize a vision for a single hyper store. It already extends the vast benefits associated with XIV's UI to the management of very large environments. XIV's implementation of IBM Hyper-Scale Mobility provides customers with a means to move volumes between XIV systems without disrupting host applications. In contrast with traditional approaches treating data migration as an exception – IBM Hyper-Scale Mobility is designed to be used a standard means to accommodate business and administrative requirements as diverse as workload balancing and system repurposing. Finally, XIV's implementation of IBM Hyper-Scale Consistency empowers customers with the ability to create consistency groups spanning multiple XIV frames. The combined power of XIV's groundbreaking architecture with IBM Hyper-Scale brings major benefits to customers and redefines the value proposition associated with scalable storage management.

Category	Aspect	XIV feat. IBM Hyper-Scale
SCALABILITY	<b>Practically unlimited</b> number of frames supported	✓
	Capacity managed can grow to <b>multi-PB</b>	✓
RELIABILITY	<b>Reliable, robust</b> data migration between frames	✓
PERFORMANCE	<b>Consistent, high performance</b> with no tuning	✓
	Data migration has <b>minimal performance impact</b>	✓
	Support for <b>concurrent data migrations</b>	✓
MANAGEMENT & EASE OF USE	<b>Breakthrough UI</b> - streamlined setup and administration for operations	✓
	Multiple systems are <b>administered as a single system</b>	✓
	Data migration is done in a manner <b>transparent to hosts</b>	✓
	<b>Minimal training</b> requirements	✓
TCO	<b>Low TCO</b> with enterprise-class performance and reliability	✓
	<b>Minimal training costs</b>	✓
	<b>Full functionality</b> - included at no extra charge	✓

Figure 5: IBM Hyper-Scale brings notable benefits to storage management

## Resources

---

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

- IBM Redpaper: IBM XIV Storage System Gen3 Architecture, Implementation, and Usage  
[ibm.com/redbooks/abstracts/sg247659.html](http://www.ibm.com/redbooks/abstracts/sg247659.html)
- IBM Redpaper: IBM Hyper-Scale in XIV Storage  
<http://www.redbooks.ibm.com/abstracts/redp5053.html?Open>
- Data Sheet: IBM XIV Storage System  
[ibm.co/1chEuOB](http://www.ibm.co/1chEuOB)
- Performance white paper: IBM XIV Storage System  
<http://ibm.co/1aXejbT>
- Asynchronous white paper: IBM XIV Storage System  
<http://ibm.co/1jRS9xq>
- 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:

Rony Shapiro, XIV Architect, IBM Systems & Technology Group;

Moshe Weiss, Manager, XIV Storage Management, IBM Systems & Technology Group

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

## About the author

---

**Rami Elron** is an IBM STSM 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.