XIV Storage System Product Guide

Bert Dufrasne

Dietmar Dausner

Nancy Kinney

Christopher Moore

Markus Oscheka

Eric Zhang

Storage

IBM

Product Guide

XIV Storage System Product Guide

The IBM® XIV® Storage System is a high-end flash optimized, fully scalable enterprise disk storage system that is based on a grid of standard hardware components.

As shown in Figure 1, the architecture of the system is designed to deliver out-of-the box performance and ease of management while harnessing the high capacity and cost benefits of serial-attached SCSI (SAS) drives.

The XIV Gen3 systems (and the Model 314 in particular) offer powerful real-time compression that is based on the IBM Random Access Compression Engine (RACE) technology.
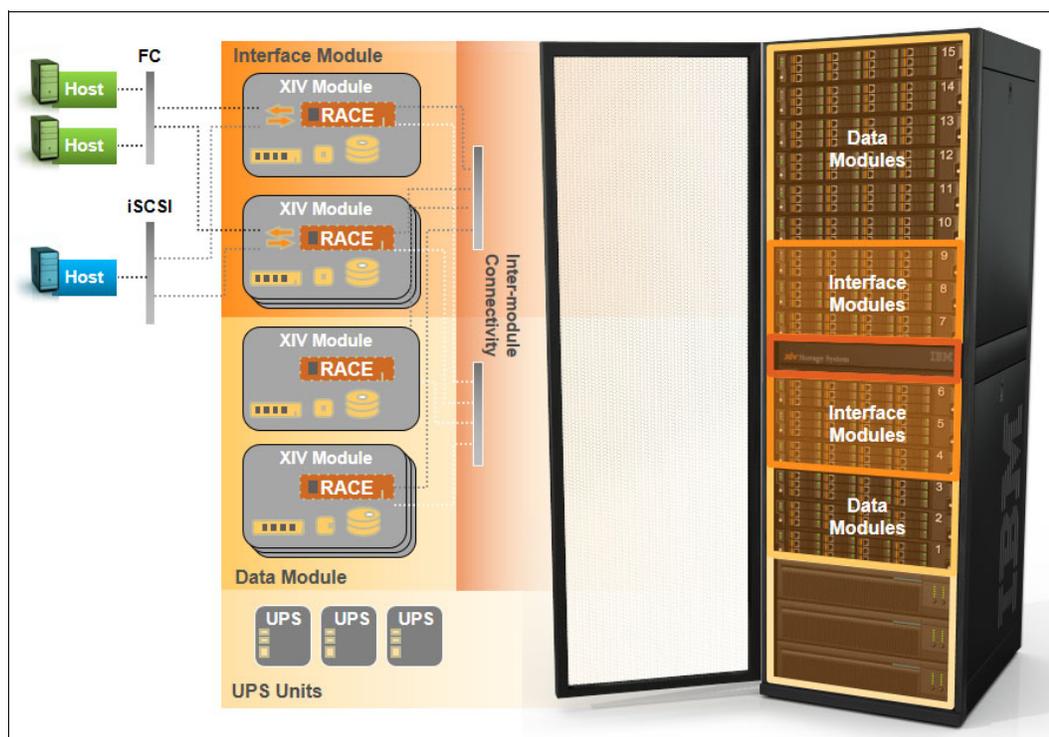

Figure 1   XIV architecture

This IBM Redbooks® Product Guide offers a brief overview of the IBM XIV Gen3 Storage System hardware and software features.

Introducing Model 314

The IBM XIV Storage System Model 314 is the latest addition to the XIV family.

It is a hardware-enhanced XIV Gen3 storage system that is targeted for customers that are requesting high-utilization of Real-time Compression™ (RtC), without performance degradation. Compared to Model 214, its compressed capacity is now up to five times larger. On compressed workloads, performance sees an improvement of 50% or more (with IBM XIV software version 11.6.1). It can also compress smaller volumes.

With Model 314, the Real-time Compression (RtC) feature is included in the IBM XIV license. Therefore, there is no need to obtain other RtC licenses and the feature is enabled by default.

Did you know?

Consider the following points:

► IBM XIV has one of the fastest rebuild times in the storage industry.

► IBM XIV Gen3 offers industry-standard data at rest encryption while avoiding performance affects with self-encrypting drives (SEDs). Encryption can be turned on non-disruptively at any time.

► IBM XIV self-optimizes automatically upon changes in the hardware configuration, such as the addition of modules or replacement of modules upon failure.

► Capacity on Demand (CoD) configurations of XIV Storage Systems are available to provide a method to quickly provision new physical storage and can be ordered at any of the valid partially populated capacities

► IBM XIV Gen3 Model 314 includes the following features:

– Has the compression feature enabled by default, with no separate licensing required

– Dedicates 48 GB of RAM and one 6-core CPU in each module to RtC, as shown in Figure 2 on page 3.

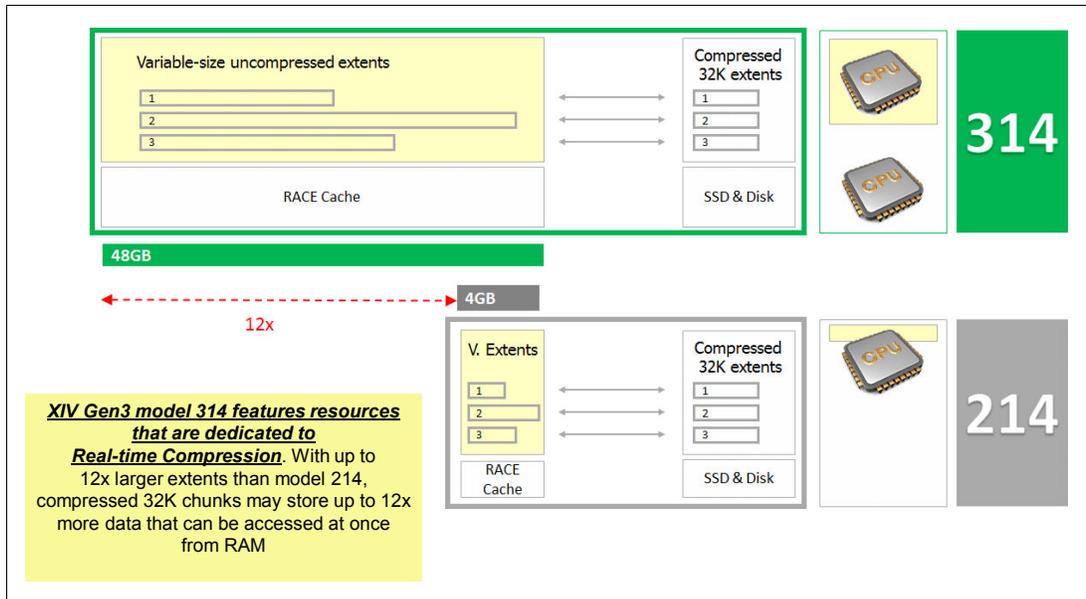– Allows a user configurable maximum soft size of up to 2 PB.

Variable-size uncompressed extents

1
2
3

RACE Cache

Compressed
32K extents

1
2
3

SSD & Disk

CPU
CPU

314

48GB

4GB

12x

XIV Gen3 model 314 features resources that are dedicated to Real-time Compression. With up to 12x larger extents than model 214, compressed 32K chunks may store up to 12x more data that can be accessed at once from RAM

V. Extents

1
2
3

RACE Cache

Compressed
32K extents

1
2
3

SSD & Disk

CPU

214

Figure 2   Model 314: Extended compression power

Product highlights

The IBM XIV Storage System is a scalable enterprise storage system that is based on a grid array of hardware components.

IBM XIV Storage System series is an innovative, high-end, open disk system series that supports business requirements for a highly available information infrastructure. XIV architecture is a grid of standard Intel and Linux components that are connected in any-to-any topology. This architecture supports enterprise-class reliability, performance, scalability, and energy efficiency.

Core features

XIV includes the following features:

► Workload balancing

The workload is evenly distributed over all hardware components always. All disks and modules are used equally, regardless of access patterns. Pseudo-random distribution ensures consistent load-balancing even after adding, deleting, or resizing volumes, and adding or removing hardware. This balancing of all data on all system components eliminates the possibility of a hot spot being created.

► True virtualization

Unlike other system architectures, storage virtualization is inherent to the basic principles of the XIV Storage System design. Physical drives and their locations are hidden from the user, which dramatically simplifies storage configuration.

3

The automatic layout maximizes the system's performance by using system resources for each volume, regardless of the user's access patterns. With the XIV, you no longer decide or plan which type of data protection to use for storing data or plan how many drives you need to dedicate for a specific application. Instead, you allocate the needed storage capacity and the system does the rest. Storage administrators can to respond to growing storage needs in minutes instead of hours or days with traditional storage systems.

► Ease of management

Ease of management that goes below the surface, with a virtualized architecture, highly intuitive Management GUI, and clever implementation of role-based management all add up to create a streamlined user experience with less administration.

► Fast rebuild

The XIV rebuild times are fast because it tracks which blocks contain data and rebuilds only these blocks. This approach results in rebuild times of less than an hour for a 6 TB drive failure when the XIV is 100% used.

Storage efficiency

The following XIV characteristics contribute to storage efficiency:

► Real-time Compression (RtC)

The RtC engine that is used in the XIV is field-proven. For Model 214, a separate license is required for RtC. For Model 314, RtC is enabled by default and does not require an extra license.

The XIV Storage software V11.6 and higher includes Compresstimator that provides specific compression predictions for workloads. The use of compression with XIV results in much higher data density, which reduces the storage system Total Cost of Ownership (TCO).

The compression feature can be applied with a single click to a volume if the system predicts a good compression ratio by using its built-in estimator. There is no need to schedule compression or to schedule conversions between compressed and uncompressed states. For more information, see *Implementing IBM Real-time Compression on the IBM XIV Storage System*, REDP-5261.

► Efficient capacity usage

Capacity utilization functionality, including space-reclamation for virtualized environments, space-efficient snapshots, and advanced thin provisioning, helps support savings in capacity and cost with fewer administration requirements.

► High-density storage

Enterprise-class, high-density SAS disk drives create increased physical space efficiencies. By using 6 TB drives, XIV Gen3 provides up to 485 TB usable capacity in a single rack. The ability to store so much data in one system by using fewer, large capacity drives, and the use of multi-core processors can help to reduce power and cooling expenses for a more energy efficient solution.

Ready for the Cloud

IBM XIV Cloud Storage for Service Providers for client-hosted cloud solutions require elastic, rapid response, and dynamic demand that is captured within IBM XIV's 485 TB linear scalability.

The immediate and automatic redistribution, without interruption, of added capacity also is important. The grid design stripes data across all modules and disks, which incorporates data redundancy for consistent and predictable I/O performance that is always load balanced.

IBM XIV is ready for the cloud in the following ways:

► IBM XIV integrates with VMware vRealize Suite. This function is available to all IBM XIV releases through the IBM Spectrum™ Control Base Edition version 2.2. It delivers a range of IBM storage integration services in cloud-based architectures that provides provisioning, automation, and monitoring. For more information, see *IBM XIV Storage System and VMware synergy with IBM Spectrum Control Base Edition,* REDP-5131.

► IBM XIV can automate storage discovery and provisioning in OpenStack environments by using OpenStack Cinder integration to manage the cloud in an easier way. We recommend deploying the IBM Cloud Manager. For more information, see *Using the IBM XIV Storage System in OpenStack Cloud Environments*, REDP-4971.

► IBM XIV can automate custom storage operation by using the IBM XIV RESTful API. For more information, see *RESTful API Support in XIV*, REDP-5064.

► IBM XIV multi-tenancy allows secure isolation into domains of IBM XIV storage resources among numerous tenants, with the ability to set different quality of service (QoS) levels per domain. It enables the division of storage system administration tasks into logical domains, by using role-based permissions. It also enables rapid deployments while minimizing the need for extensive planning, tuning, or field-upgrades.

► IBM XIV scales storage capacity and performance linearly and can manage up to 144 systems by using IBM Hyper-Scale Manager.

► The IBM XIV Cloud Storage for Service Providers product further empowers users with the flexibility to combine a robust set of base functions with license-per-need features. The ability to pay for functions when it is required enables cloud providers to tailor services to customer needs even more cost-effectively.

► IBM XIV can be configured to support disaster recovery in the Cloud by using IBM Spectrum Accelerate™ on Cloud as secondary system. For more information, see *IBM Spectrum Accelerate on Cloud*, REDP-5261.

Security and Encryption

The IBM XIV Storage System features the following characteristics:

► Self-encrypting drives

IBM XIV Gen3 Model 214 and Model 314 feature industry-standard data at rest encryption at no extra cost, while avoiding performance affects with SEDs. IBM Security Key Lifecycle Manager is used to manage the security keys.

Also, encryption can be turned on non-disruptively with data initially already on disks. For more information about encryption with IBM XIV, see *XIV Security with Data-at-Rest Encryption*, REDP-5047.

► Payment Card Industry Data Security Standard (PCI DSS)

Security enhancements, including auditing of user actions and user interface locking after a predefined period, allow IBM XIV to meet requirements that are imposed by the PCI-DSS.

Management

The highly intuitive IBM XIV Management GUI and built-in management tools make administrative tasks easy and efficient (with little training or expertise required) from provisioning volumes to monitoring multiple systems. Figure 3 shows an example of a system view with a module pullout in the IBM XIV Management GUI.

Figure 3   IBM XIV Management GUI System view

Consider the following points:

► The IBM XIV Management GUI acts as the management console for the storage system. A simple and intuitive GUI enables storage administrators to manage and monitor all system aspects easily, with almost no learning curve.

► The IBM XIV TOP application allows the user to view and monitor performance information for defined volumes and hosts in real time. It can be started independently or from within the IBM XIV Management GUI.

► The XCLI is a comprehensive command-line interface that allows the user to configure and monitor the IBM XIV Storage System. All of the functions that are available in the IBM XIV Management GUI are also available in the XCLI. It can be used in a shell environment for interactive commands or as part of a script to perform lengthy or complex tasks.

► Performance statistics can be monitored through the IBM XIV Management GUI and XCLI at any time. Monitoring through the IBM XIV Management GUI is easily done by selecting specific filters. The IBM XIV Management GUI then displays the requested data, as shown in Figure 4 on page 7.
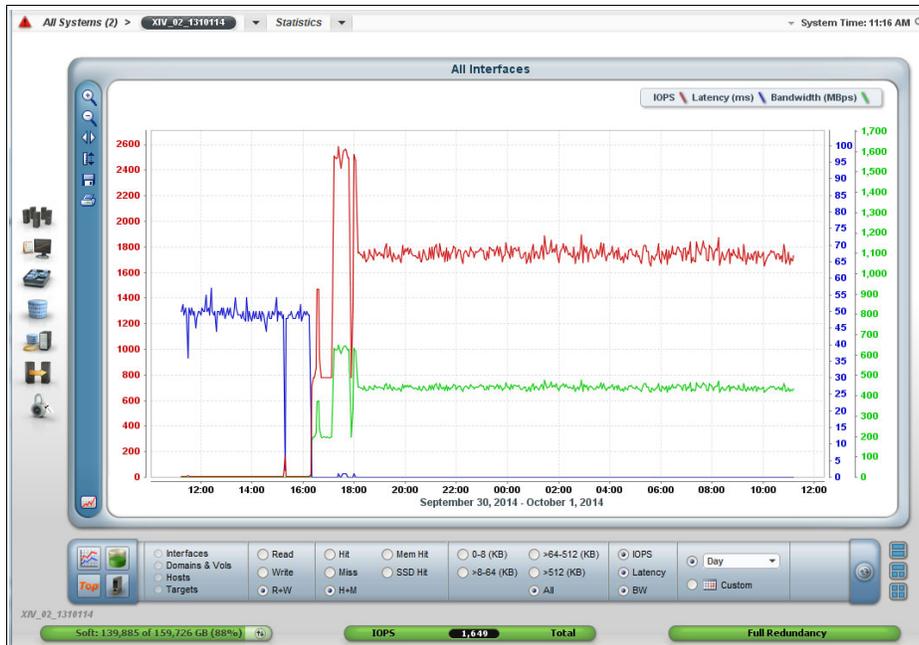
Figure 4 Statistics monitor view in IBM XIV Management GUI

► The IBM XIV Mobile Dashboard, which is supported on Apple iOS and Android smartphones and tablets, is available at no cost from the specific application's store. The mobile dashboard gives the user another way to monitor performance and capacity and receive alert notifications from the IBM XIV Storage System. Figure 5 shows an example of the IBM XIV Mobile Dashboard window.

Figure 5 IBM XIV Mobile Dashboard

► SMI-S 1.6 latest protocol certification helps build dynamic, scalable, secure Microsoft based cloud storage infrastructures with ready-to-use integration with Microsoft System Center Virtual Machine Manager 2012.

7

IBM Hyper-Scale

IBM Hyper-Scale is a family of growing technologies that were developed around an innovative approach to storage scalability. IBM Hyper-Scale includes the following features:

► Hyper-Scale Manager

This feature reduces operational complexity and enhances capacity planning through integrated management for large and multi-site IBM XIV deployments. It enables the IBM XIV Management GUI to access and operate on multiple IBM XIV systems concurrently. It also supports the RESTful application programming interface (API).

The Hyper-Scale Manager runs on a single instance of a virtual machine server or on several servers. Figure 6 shows how the Hyper-Scale Manager allows a storage administrator to work with volumes from multiple systems; in this case to create cross systems snapshots.
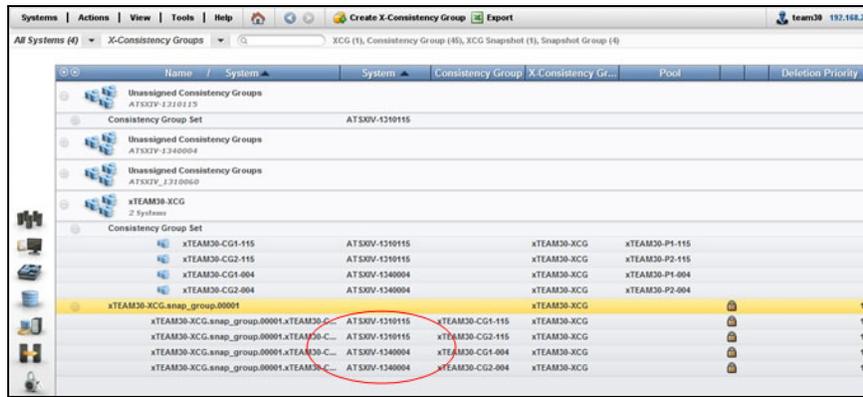

Figure 6   Hyper-Scale Manager

► Hyper-Scale Mobility

By using this powerful function, users can move volumes between storage systems transparently with no disruption to host applications.

► Hyper-Scale Consistency

Cross system consistency (or snapshot) groups enable a coordinated creation of snapshots for inter-dependent consistency groups on multiple systems. This feature is available through the IBM Hyper-Scale Manager only.

Business Continuity

The IBM XIV Storage System provides a rich set of copy services functions that are suited for various data protection scenarios. The following functions also enable clients to enhance their business continuance and data migration capabilities:

► Synchronous mirroring is a data replication solution to use between two storage systems that achieves a recovery point objective (RPO) of zero with a distance of less than 100 km (62.18 miles). In synchronous mirroring, a host write operation is completed on the local and remote sites before an acknowledgment is returned to the host. Synchronous mirroring ensures local and remote always have the same copy of data for a zero RPO.

► Asynchronous mirroring is designed for a nonzero RPO at greater distances. Consistent sets of data are copied to the remote location at predefined intervals while the host writes are acknowledged after they are written on the local site.

- ► Three-site (or three-way) mirroring is a star topology multi-target disaster recovery solution. It uses proven, efficient IBM XIV technology with synchronous and asynchronous mirroring.
- ► IBM XIV data migration is a seamless data transfer tool for migrating data from another source system by simulating host behavior. It synchronizes data between the two storage systems by using transparent copying to the IBM XIV Storage System as a background process. It requires only a short outage on the host side to switch logical unit number (LUN) ownership to the IBM XIV and begin the migration process.

Scalability and performance

IBM XIV Storage System is a scalable enterprise storage system that is based on a grid array of hardware components. The architecture offers the highest performance through maximized and balanced usage of all disks, distributed cache implementation, and exceptional performance characteristics.

The system includes the following scalability and performance features:

- ► Massive parallelism

  With the grid architecture, the system ensures full usage of all system components. All volumes are spread across all spindles in the system. The system harnesses all storage capacity and all internal bandwidth, and it uses all available processing power for host-started I/O activity and system-started activity, such as rebuild processes and snapshot generation.

- ► Processing power

  The IBM XIV Storage System open architecture uses the latest processor technologies and is more scalable than solutions that are based on a closed architecture. The IBM XIV Storage System avoids sacrificing the performance of one volume over another; therefore, it requires little to no tuning.

- ► Innovative cache memory

  Up to 720 GB of total system cache with flexible and powerful cache implementation allows the IBM XIV system to use large slots for reads and pre-fetched data while enabling it to manage a smaller slot size for improved cache management and better performance. In the Model 314, another 720 GB of RAM (48 GB per module) is dedicated to the compression process.

- ► Flash Cache (optional in Model 214)

  For ultra-high performance needs, IBM XIV Gen3 optionally offers up to 12 TB of management-free Flash Caching available to all system data. The flash caching option can be installed non-disruptively and provide up to a 90% reduction in I/O latency for random read workloads with block size up to 64 KB. Operating with advanced flash algorithms, the latest IBM XIV Gen3 components help meet requirements for high-performance workloads.

- ► Enhanced connectivity

  IBM XIV Gen3 offers multiple active/active IO interfaces with improved host connectivity with up to 24 8 Gb FC ports and up to 12 10 Gb Ethernet ports or 22 1 Gb Ethernet ports for connecting to iSCSI-attached hosts.

Availability and serviceability

The IBM XIV Storage System maximizes continuous operation and minimizes the performance degradation that is associated with nondisruptive planned and unplanned events, while providing the capability to preserve the data in a disaster.

The system includes the following availability and serviceability features:

► Disk failure is managed by an efficient rebuild process that brings the system back to full redundancy in minutes. In addition, the IBM XIV Storage System extends the self-healing concept, resuming redundancy even after failures in components other than disks, or upon changes in the hardware configuration, such as the addition of modules.

► The rapid restoration of redundant data across all available drives and modules in the system during hardware failures and the automatic redistribution of data across all newly installed hardware are fundamental characteristics of the IBM XIV Storage System architecture.

► IBM XIV Storage System dynamically maintains the pseudo-random distribution of data across all modules and disks while ensuring that two copies of data are available always when the system reports *Full Redundancy*. When a disk drive or a module fails, the IBM XIV data distribution algorithms automatically identify the non-redundant partitions and then begins to make copies of only those non-redundant partitions. It then places this data in reserved areas of the disk drives that are for exactly this circumstance.

► Non-Disruptive Code Load (NDCL) enables upgrades to the IBM XIV Storage System software from a current version to a later version without disrupting the application service.

► The provision within the IBM XIV Storage System to efficiently and flexibly create snapshots, which are coupled with the ability to define consistency groups of volumes, constitutes integral elements of the data preservation strategy. In addition, the IBM XIV Storage System data mirroring functions facilitates excellent potential recovery point and recovery time objectives as a central element of the full disaster recovery plan.

Architecture and key components

The IBM XIV Storage System is composed of the following physical components, as shown in Figure 7 on page 12:

► IBM XIV Storage System is delivered in a standard IBM T42 rack, fully cabled for 15 modules, even if a partially populated rack is ordered.

► The hardware of the Interface Modules **[1]** and Data Modules **[2]** in the IBM XIV Storage SystemGen3 is based on an Intel server platform optimized for data storage

► A patch panel **[3]** at the rear of the rack is used for connecting the IBM XIV Storage System to the customer's network. The Fibre Channel ports on the Interface Modules are connected to the patch panel by using 50-micron cables.

All external connections must be made through the patch panel. In addition to the host connections and to the network connections, more ports are available on the patch panel for service connections.**[4]**

► An Automatic Transfer Switch (ATS) **[5]** allows the IBM XIV Storage System to attach to two independent power sources to protect against loss of customer utility power. If there is an external power failure, the ATS automatically transfers the load to the redundant power supply.

► The IBM XIV Storage System has three internal Uninterruptible Power Supplies (UPSs) **[6],** which are used as cache batteries so the data in cache always is written down to disk, if there is a power outage.

Two of these UPSs are enough to maintain cache battery power availability. This redundant design is n+1 because the IBM XIV Storage System has one redundant UPS. The UPS complex features enough battery reserves to sustain two emergency power shutdowns.

► The 1U Maintenance Module **[7]** and the modem, which are installed in the middle of the rack, are used for IBM XIV Storage System support and for the IBM personnel to maintain and repair the system. This device is used only to gain remote access to the IBM XIV System through the modem for support personnel.

► The IBM XIV Storage System uses redundant InfiniBand switches **[8]** for communications between the interface and data modules.

Each InfiniBand switch contains 36 ports that include 40-Gbps full bidirectional bandwidth per port. The switches are powered by redundant power supplies and fan modules to eliminate any single point of failure.
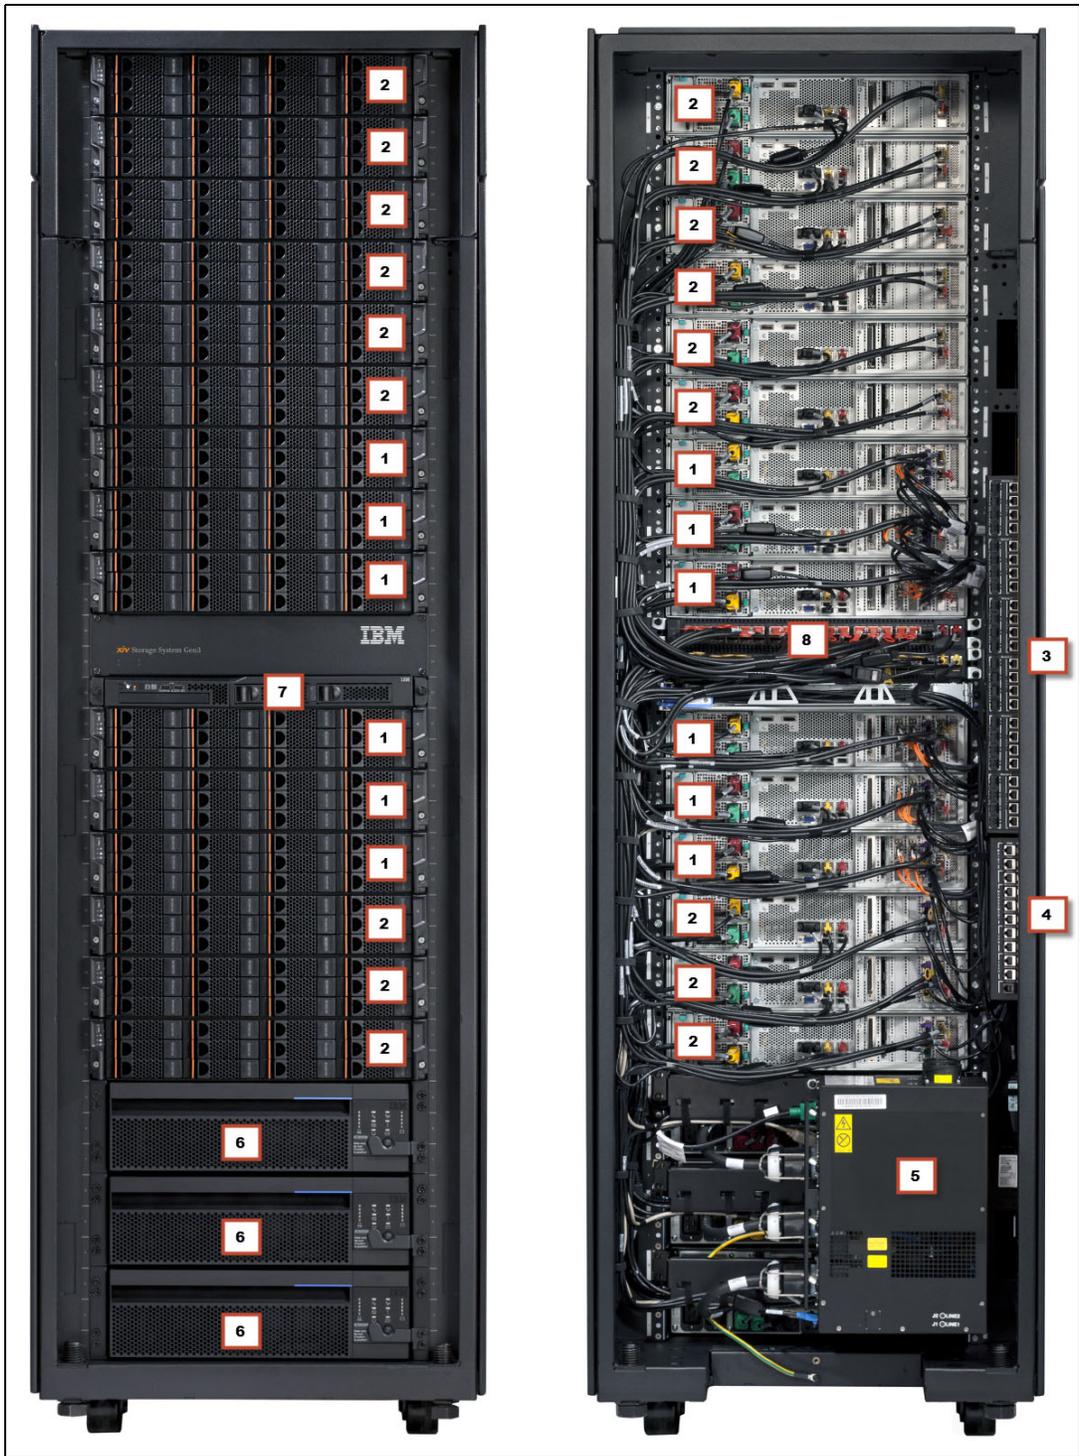
Figure 7   Physical components of an IBM XIV Storage System

Models

The rack specifications for capacity, connectivity, system resources, and power usage are shown in Figure 8. The specifications differ based on the number of modules in the system, and the size of the disk drives it contains.

Rack Configuration								
	Model 214 Only							
Total number of modules (Configuration type)	6 Partial	9 partial	10 partial	11 partial	12 partial	13 partial	14 partial	15 full
Total number of data modules	3	3	4	5	6	7	8	9
Total number of interface modules	3	6	6	6	6	6	6	6
Number of active interface modules	2	4	4	5	5	6	6	6
Interface module 9 state		Disabled	Disabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 8 state		Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 7 state		Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 6 state	Disabled	Disabled	Disabled	Disabled	Disabled	Enabled	Enabled	Enabled
Interface module 5 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 4 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
FC ports	8	16	16	20	20	24	24	24
iSCSI ports (1 Gbps – mod 114)	6	14	14	18	18	22	22	22
iSCSI ports (10 Gbps – mod 214)	4	8	8	10	10	12	12	12
Number of disks	72	108	120	132	144	156	168	180
Usable capacity (1 / 2 / 3 TB) (Mod 214 only)	28 TB 55 TB 84 TB	44 TB 88 TB 132 TB	51 TB 102 TB 154 TB	56 TB 111 TB 168 TB	63 TB 125 TB 190 TB	67 TB 134 TB 203 TB	75 TB 149 TB 225 TB	81 TB 161 TB 243 TB
Usable capacity (4 / 6 TB)	112 TB 169 TB	177 TB 267 TB	207 TB 311 TB	225 TB 338 TB	254 TB 382 TB	272 TB 409 TB	301 TB 453 TB	325 TB 485 TB
# of CPUs (one per Module) Mod 214 # of CPUs (two per Module) **Mod 314**	6 n/a	9 **18**	10 **20**	11 **22**	12 **24**	13 **26**	14 **28**	15 **30**
Memory (24 GB per module w 1/2/3 TB) Memory (48 GB per module w 4/6 TB) **Mod 314: Memory (96 GB per module w 4/6 TB)**	144 GB 288 GB **n/a**	216 GB 432 GB **864 GB**	240 GB 480 GB **960 GB**	264 GB 528 GB **1056 GB**	288 GB 576 GB **1152 GB**	312 GB 624 GB **1248 GB**	336 GB 672 GB **1344 GB**	360 GB 720 GB **1440 GB**
{Optional for 1, 2 ,3, 4, 6 TB XIVs} 400 GB Flash Cache {Optional for 4, 6 TB XIVs} 800 GB Flash Cache	2.4 TB 4.8 TB	3.6 TB 7.2 TB	4.0 TB 8.0 TB	4.4 TB 8.8 TB	4.8 TB 9.2 TB	5.2 TB 10.4 TB	5.6 TB 11.2 TB	6.0 TB 12.0 TB

Figure 8   IBM XIV Gen3 Rack Configuration

Some of the configurations that are shown Figure 8 include the following components:

► Partially populated configurations of 6, 9, 10, 11, 12, 13, or 14 modules, which allows for more granularity of capacity options. Modules can be added as capacity needs increase, up to 15 modules. For these configurations, some of the interface modules are not available. Those models are indicated as "disabled" in Figure 8.

> **Note:** The first column of configuration in Figure 8 is applicable to Model 214, not to Model 314. Model 314 starts with a nine-module configuration. The Model 314 is available with 4 TB or 6 TB drives only.

► A fully populated rack contains nine data modules and six interface modules for a total of 15 modules. The total usable capacity is 161 TB for a complete system when equipped with 2 TB drives, 243 TB of usable capacity when fully populated with 3 TB drives, 325 TB of usable capacity when fully populated with 4 TB drives and 485 TB of usable capacity when fully populated with 6 TB.

► Fully and partially populated systems can be configured with 1 or 10 Gbps iSCSI ports. The number of ports that is available for the 1 Gbps option is 6 - 22. The number of 10 Gbps ports is 4 - 12 per system.

► For Model 214, the memory options that are available for fully and partially populated systems in 24 GB per module for systems that are configured with 1, 2, or 3 TB drives and 48 GB per module for 4 or 6 TB drives.

For Model 314, the memory for fully and partially populated systems is 96 GB per module for systems that are configured with 4 or 6 TB drives (48 GB of the 96 GB RAM is dedicated to RtC).

Specifications

Figure 9 shows other physical characteristics of the Model 314.

Physical features	
Temperature	10°C – 35°C (50°F – 95°F)
Altitude (max)	2,134 m/7,000 ft
Humidity	25% – 80% noncondensing
Noise Level	56 dBA
Dimensions (height × width × depth)	202 cm × 66 cm × 120 cm (79.53 in. × 25.98 in. × 47.24 in.)
Maximum weight	1044.5kg (2303.1 lb)
Clearance front/rear	120 cm/120 cm (47.24 in./47.24 in.)
Redundant power feed	Yes
Input voltage	180 – 264 V ac at 60 A or 30 A (±10%)
Power usage	7.5 kVA max. consumption (15 modules with 4TB or 6TB disks)
Host connectivity	
Fibre Channel rates	Up to 24 x 8 Gb FC ports
iSCSI rates	Up to 12 x 10 GbE ports or 22 x 1 Gbps or 1GbE ports
Warranty	1 and 3 year limited warranty, onsite service, same day 24×7

Figure 9   IBM XIV Specifications

Options

The IBM XIV Storage System can be purchased with different options, including the following options:

► Flash Cache is a feature that is available for IBM XIV Gen3 systems. The system can be equipped with 400 GB of flash cache for every module. In addition, the IBM XIV Gen3 systems with 4 TB or 6 TB drives can be equipped with 400 GB or 800 GB flash cache for every module. For Model 314, 800 GB flash cache is used regardless of the capacity of the hard disk drives that are used.

► IBM XIV Cloud Storage for Service Providers offers an innovative pay-per-need scalability matrix. The offer empowers cloud providers to cater to tenant requirements with ease and efficiency through flexible software licensing. For more information, see this website:

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&appname=gpateam&supplier=897&letternum=ENUS214-173&pdf=yes

► Advanced system placement is a pay-as-you-go program with which an organization can purchase the IBM XIV Storage System for a fraction of the price upon installation. When the system reaches a predetermined capacity threshold, the client is charged for the full balance of the system and can acquire the next system for $1 (US). This program is designed for rapidly growing environments.

- Capacity on demand is a program that helps organizations with less aggressive growth projections. A Gen3 system can be ordered with a specified amount of authorized storage capacity along with extra capacity that is not intended for initial use. The extra capacity is purchased as needed when allocated to a storage pool.

Warranty information and upgrades

IBM XIV Storage System Model 314 Device Type 2812 and Device Type 2810 offer a flexible warranty choice. Device Type 2812 supports a 3-year warranty to complement the 1-year warranty offered by the functionally equivalent, Device Type 2810.

Technical advisors

Technical advisors (TAs) are available with the initial warranty and for any other fee after the warranty expires. The TA supports the initial installation and updates of the IBM XIV Storage System. The TA also enhances end-to-end support by being a client advocate to proactively manage problem resolution through the support process for hardware.

Call home and remote support

Call home is the capability of the system (when enabled) to send event notifications to the IBM XIV Remote Support Center. This feature enables proactive and failure notifications to be sent directly to IBM for analysis.

Remote support enables trained IBM service personnel to connect to the IBM XIV Storage System to analyze a problem, repair it remotely if possible, or assist the onsite IBM service support representative.

Related information

For more information, see the following resources:

- *IBM XIV Storage System Gen3 Architecture, Implementation, and Usage*, SG24-7659-04:

  http://www.redbooks.ibm.com/abstracts/sg247659.html

- *IBM XIV Storage System Business Continuity Functions*, SG24-7759:

  http://www.redbooks.ibm.com/abstracts/sg247759.html

- IBM XIV Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/STJTAG/product_welcome/xiv_kcwelcome page.html

- *IBM XIV Gen3 with IBM System Storage® SAN Volume Controller and Storwize® V7000*, REDP-5063-00:

  http://www.redbooks.ibm.com/abstracts/redp5063.html

- *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904-02:

  http://www.redbooks.ibm.com/abstracts/sg247904.html

- *Solid-State Drive Caching in the IBM XIV Storage System*, REDP-4842-00:

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

- *IBM XIV and VMware Synergy with the IBM Storage Integration Server*, REDP-5131-00

  http://www.redbooks.ibm.com/redpieces/abstracts/redp5131.html

► *XIV Storage System in a VMware Environment*, REDP-4965-00:

  http://www.redbooks.ibm.com/redpieces/abstracts/redp4965.html

► *IBM XIV Security with Data-at-Rest Encryption*, REDP-5047-00:

  http://www.redbooks.ibm.com/abstracts/redp5047.html

► *IBM Hyper-Scale in XIV Storage*, REDP-5053-01:

  http://www.redbooks.ibm.com/abstracts/redp5053.html

► *IBM XIV Storage System: IBM Hyper-Scale Mobility Overview and Usage*,
  REDP-5007-00:

  http://www.redbooks.ibm.com/abstracts/redp5007.html

► *IBM XIV Storage System Thin Provisioning and Space Reclamation*, REDP-5001-00:

  http://www.redbooks.ibm.com/abstracts/redp5001.html

► *Using the IBM XIV Storage System in OpenStack Cloud Environment*, REDP-4971-00:

  http://www.redbooks.ibm.com/abstracts/redp4971.html

► *RESTful API Support in IBM XIV*, REDP-5064-00:

  http://www.redbooks.ibm.com/abstracts/redp5064.html

► *IBM XIV Storage System Multi-Site Mirroring*, REDP-5129-00:

  http://www.redbooks.ibm.com/abstracts/redp5129.html

► *Space Reclamation in IBM XIV with Windows 2012*, TIPS1011:

  http://www.redbooks.ibm.com/abstracts/tips1011.html

► IBM Hyper-Scale: A powerful new approach to scaling storage management (XIV):

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=T
  SW03281USEN#loaded

► IBM Offering Information page (announcement letters and sales manuals):

  http://www.ibm.com/common/ssi/index.wss?request_locale=en

Authors

This Product Guide was produced by a team of specialists from around the world working at
the International Technical Support Organization, San Jose Center, including the following
authors:

► Bert Dufrasne
► Dietmar Dausner
► Christopher Moore
► Markus Oscheka
► Nancy Kinney

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from 2 - 6 weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ► Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- ► Follow us on Twitter:

  http://twitter.com/ibmredbooks
- ► Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- ► Explore new Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- ► Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-5272-00, was created or updated on January 5, 2016.

Trademarks

IBM, the IBM logo, and ibm.com 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®	Redbooks®	System Storage®
IBM Spectrum™	Redpaper™	XIV®
IBM Spectrum Accelerate™	Redbooks (logo) ®	
Real-time Compression™	Storwize®	

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

IBM®

Printed in U.S.A.

Get connected

Redbooks

ibm.com/redbooks