Storage Reinvented

Technical Description

IBM XIV Storage System

Storage Reinvented

IBM

Performance

The IBM XIV Storage System offers a new level of high-end disk system performance and reliability. It is a core component of the IBM Information Infrastructure which helps clients address their needs for availability, security, compliance and retention of information. The XIV system provides immunity to hotspots and consistent performance in the face of hardware failure. Several revolutionary practices contribute to this achievement:

Perfect Load Balancing

Each logical volume in the IBM XIV Storage System is divided into multiple stripes of one megabyte (MB). These stripes are distributed evenly across all disks in the system using a sophisticated pseudo-random distribution mechanism.

This unprecedented approach to load balancing ensures:

- The equal use of all disks and modules, regardless of access patterns. Although applications may access some volumes or certain parts of a volume more frequently than others, the load on the disks and modules remains perfectly balanced
- The maintaining of load balance regardless of changes in access patterns, such as the adding, deleting, or resizing of volumes, or the adding or removal of hardware.

Cache and Disk into Every Module

Unlike the design of traditional storage systems, the XIV system embeds the read/write cache in the same hardware module as the disks. This unconventional design aspect produces several advantages:

- **Distributed cache.** The cache is distributed evenly across all modules, enabling all the modules to serve host input/output (I/O) and perform cache-to-disk I/O concurrently. This ensures that the cache never becomes a bottleneck
- **High cache-to-disk bandwidth.** Aggressive pre-fetching is enabled by the fact that cache-to-disk bandwidth is the internal bandwidth of a module, providing dozens of gigabytes per second (GBps) per rack
- **Powerful cache management.** The XIV system reads a large cache slot per disk read, while managing least-recently used statistics in small cache slots. This unique combination is possible due to the system's huge Central Processing Unit (CPU) and high cache-to-disk bandwidth.

Huge CPU Power

Each data module is equipped with its own quad-core processor, giving the XIV system dozens of CPU cores. The system uses this vast processing power to execute advanced caching algorithms that support small cache slots, enable powerful snapshotting and ensure high performance through higher cache hit rates.

High Performance Even During Disk Rebuild

A distributed rebuild mechanism engages all disks in the system during the rebuild process. The end result is a rebuild process that has minimal impact, keeping high performance levels intact.

High performance is enabled by:

– Perfect load balancing at all times
– Utilisation of all system resources
– Innovative cache architecture
– Massive CPU power

High-end Open Storage Based on Serial Advanced Technology Attachment (SATA) Drives

The XIV system innovatively provides high-end open storage based on SATA drives. It leverages their huge capacity, density, power and cost advantages, while ensuring high performance and reliability through several means:

- Attains high-end open performance through comprehensive parallelism, with all drives used concurrently.
- Optimises performance through a sophisticated caching architecture and algorithms
- Ensures a minimal number of disk faults through uniform activity distribution
- Achieves high reliability through a rapid rebuild process that takes just minutes.

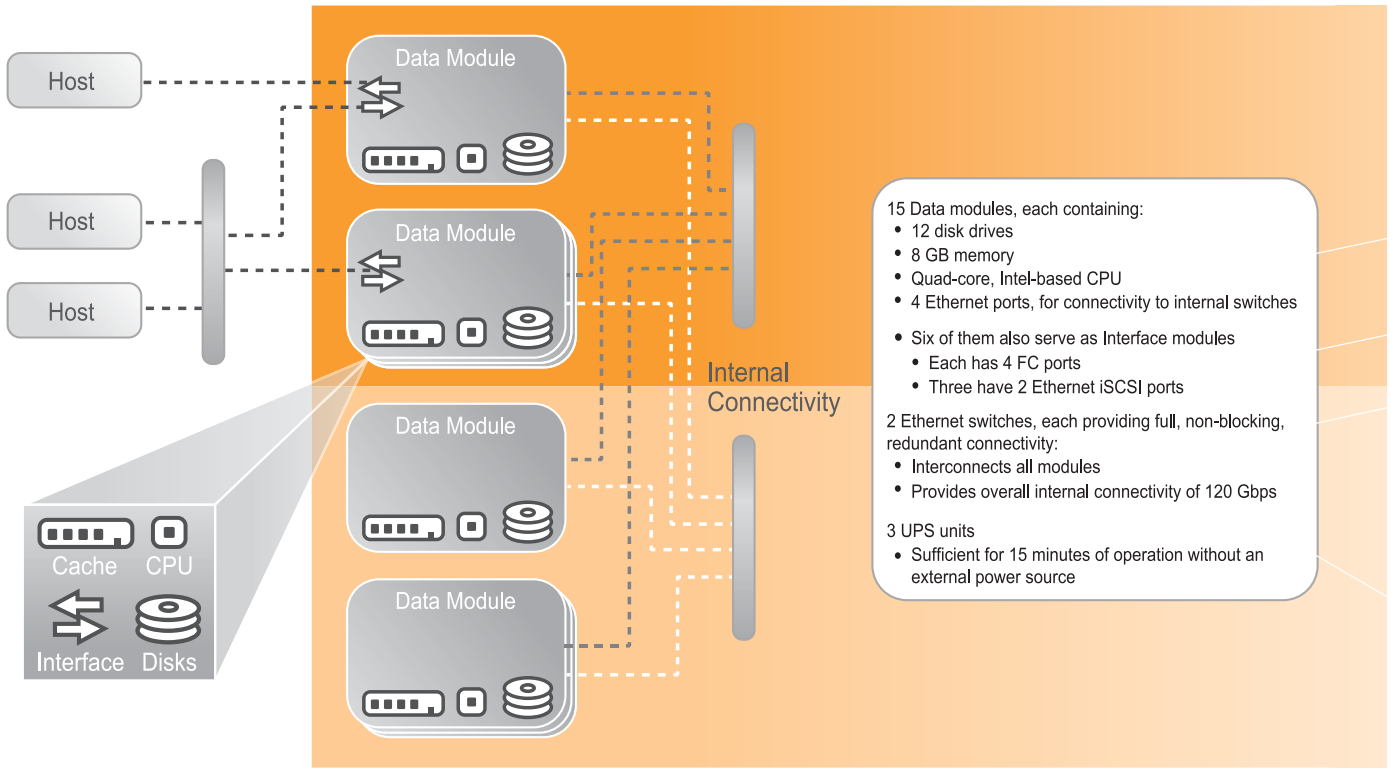

Reliability

The XIV system provides outstanding reliability, enabled by its unique grid architecture and distributed rebuild mechanism. These qualities immunise the XIV system against failures.

Active-Active N+1 Redundancy

The XIV system can survive single failure without affecting host I/O. Every disk, module, switch or Uninterruptible Power Supply (UPS) unit is redundant and protected through an active-active N+1 redundancy scheme. Each of these components is hot-swappable – replaceable without system shutdown.

30-Minute Rebuild Time

The system is designed to minimise the risk of impact from disk failure. It does so by involving every disk in the system in the rebuild process, greatly shortening rebuild time. Using its unique distributed rebuild approach, the XIV system thin-slices all data into 1 MB stripes and distributes and mirrors each stripe on different disks. On a fully utilised XIV system based on 1 TeraByte (TB) disk drives, the system requires less than 30 minutes of rebuild time. The result is a significant reduction in the potential risk of data loss – by orders of magnitude in comparison with other storage systems.

Rebuilding of Only Real Data

Standard storage systems perform disk rebuild on the block level, completely rebuilding the failed disk. The XIV system performs rebuild on only the data that is allocated to volumes and, within volumes, only the data actually written. In actuality, the rebuild time is usually much less than 30 minutes for 1 TB disk drives since, in most cases, not all capacity is allocated and that which is allocated is not necessarily in use.

Self-Healing upon Module Failure

The XIV system employs self-healing even after module failure: the system automatically initiates a rebuild process and returns to full redundancy.

IBM XIV Storage System

Scalability

XIV technology scales in every aspect – capacity, cache size and processing power – and is designed for seamless hardware transitions.

Scalability in Every Dimension

- Thanks to perfect load balancing, the XIV system fully uses all disks on all racks, without administrative intervention
- Cache is embedded in each data module enabling cache size to grow linearly with disk capacity. Cache bandwidth and, as such, performance levels, grow linearly with cache size
- The number of host interfaces and their bandwidth grow proportionally to system capacity, increasing linearly the number of applications supported and total throughput
- Internal switching capacity stays unchanged regardless of system size, avoiding bottlenecks and ensuring that throughput grows proportionally to capacity
- Processing power is embedded in each module, ensuring that system size has no impact on performance during snapshots, caching and self-healing.

Easy Migration and Upgradeability

- Any XIV system can scale from the minimal to maximal configuration
- The system automates the transfer of data to new hardware, automatically load-balancing all data across new and existing hardware
- The use of standard components enables the system to capitalise on the latest industry technologies
- Hardware upgrade activities can take place without interruptions to production.

> *The XIV system's virtualisation and simplicity, together with its powerful GUI management, significantly reduce the time and effort required for managing storage resources.*

Thin Provisioning

The XIV system was built to provide thin provisioning. It powerfully allows organisations to use thin provisioning to trim physical capacity buying to an as-needed-only basis:

- Provides the flexibility to define logical volume sizes as larger than physical capacity, while physical capacity need be just larger than the data actually written
- Enables the periodic reclaiming of disk space no longer in use. Instant space reclamation is also enabled for supported applications
- Automatically shrinks volumes upon migrating data from a non-XIV system, offering great power and space savings.

Snapshots and Mirroring

- XIV snapshot enables a near-infinite number of snapshots and with virtually no performance overhead, opening the door to new storage paradigms based on snapshotting
- XIV also supports business continuity through synchronous and asynchronous differential mirroring that offers flexible backup and restore options between remote sites for rapid recovery
- Unlike other implementations, XIV asynchronous mirroring supports multiple consistency groups with different recovery point objectives, as well as multiple targets and mirrored pairs, flexible scheduling, event reporting and statistics collection.

Total Cost of Ownership (TCO)

The XIV system offers numerous direct and indirect cost advantages:

- Power, floor space and cooling expenses are significantly reduced due to the use of SATA drives – without compromising performance or reliability
- Ease of management greatly reduces administrator activities and related costs
- Differential snapshots, elimination of orphaned space and thin provisioning capabilities enable data centres to use less system capacity to achieve the same goals.

Ease of Management

Workload-Reducing Virtualisation

The XIV system is fully virtualised, freeing the user from the need to plan how logical volumes are spread over physical disks. The system assigns physical resources automatically, ensuring optimal resource utilisation. The user sees newly added hardware components as additional free space.

Powerful Management Tools

Graphical User Interface (GUI) – A simple and intuitive GUI allows all administrative operations to be carried out quickly and easily, with minimal training and knowledge.

Command Line Interface (CLI) – A powerful CLI tool enables the writing of complex scripts for high-level system administration and integration with hosts and applications.

Single Tier Solution

- High-end open performance, reliability and features
- Mass storage affordability
- Avoids the cost of software and administration overhead required for some tiered storage solutions.

For more information

To learn more about IBM XIV Storage System, contact your IBM representative or IBM Business Partner, or visit: www.xivstorage.com OR: **ibm.com**/systems/storage/news/centre/xiv/

Additionally, IBM Global Financing can tailor financing solutions to your specific IT needs. For more information on great rates, flexible payment plans and loans as well as asset buyback and disposal, visit: **ibm.com**/financing

Recyclable, please recycle.

IBM XIV Storage System

IBM United Kingdom Limited
PO Box 41
North Harbour
Portsmouth
Hampshire
PO6 3AU
United Kingdom
IBM Ireland Limited
Oldbrook House
24-32 Pembroke Road
Dublin 4
IBM Ireland Limited registered in Ireland under company number 16226.

TSF03012-GBEN-03