



IBM XIV[®] Storage System

Host-based Concurrent Multipath Software and the IBM XIV Advantage

White Paper | June 2009





Storage Reinvented

Concurrent Multipath Software & the XIV Advantage

Copyright IBM Corporation 2009

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

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

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

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





Contents

What is Concurrent Multipath Software?	1
Multipath I/O in Practice	1
Channel Failover	1
Concurrent Multipath I/O	2
Load Balancing	2
Active/Active Storage Arrays	2
Multipath I/O Architectures	2
TCO: Host-Native vs. Storage-Vendor Multipath Software	3
Software Cost	3
Quality Issues	3
Deployment Costs	3
XIV's Approach to Multipath Software	4
Summary	5



What is Concurrent Multipath Software?

When only one physical path exists between a computer system's CPU and mass storage devices, a failure in any component along the I/O path will shut down the entire storage system and the applications that rely on it. To prevent this kind of catastrophic and costly event, hardware vendors have provided more than one physical path for I/O, and storage vendors have provided multipath software to take advantage of the alternate routes available. If hardware, such as a controller, port, or switch, should fail, the multipath software immediately and transparently reroutes I/O to the application.

Over time, multipath software began to use idle redundant hardware to provide greater throughput at all times, not only upon hardware failure. This feature became known as "concurrent multipath software."

Recently, operating system (OS) vendors have begun offering concurrent multipath modules that are gaining increasingly wide acceptance. As Wikipedia* notes, "native multipathing technologies in operating systems like Windows, Linux and HP-UX have progressively advanced and [are] preferred by customers. Most of these software packages come free of charge thereby increasing the momentum of customer adoption in the market."

An added advantage of OS vendor-created (host-based) concurrent multipath software is that it is backed by the OS development teams, receiving not only ongoing support but also the benefits of a continuous stream of new features.

The XIV system strongly recommends the use of standard OS concurrent multipath software over proprietary multipath software provided by storage vendors. This paper delineates why.

** Wikipedia entry for "multipathing" as it appeared on February 5, 2009*

Multipath I/O in Practice

Innovations tend to create new opportunities — as well as unanticipated new issues that need to be addressed. In this section, we review the development of multipath software and its pursuit by the various players.

Channel Failover

Redundancy in the connection between a computer server and a storage system was first provided to ensure that failure of equipment along the path would not be catastrophic. The aim was for any failure of a component along this path to automatically and seamlessly activate an alternate path, enhancing storage system reliability. This capability, known as "channel failover," has been part of the operating system and/or specialized device drivers on open servers for the past 20 years.



Concurrent Multipath I/O

With idle redundant I/O equipment now present, it was inevitable that vendors would eventually develop software to leverage it. Over time, features were added to storage-vendor multipath software to allow I/O operations to run **concurrently** on multiple paths, potentially increasing overall I/O throughput and bandwidth. This improvement gave mainframe-like multipath I/O functionality to open system storage.

Load Balancing

Some implementations even attempted to optimize overall I/O performance between the server and storage system, by providing dynamic load balancing between the I/O paths.

Active/Active Storage Arrays

Not all storage systems are capable of providing active/active concurrent multipath access to volumes. Most mid-range storage systems provide only active/passive access (channel failover capability with no concurrency). Some mid-range storage systems provide failover of the control device driver on the host server, but this is not to be confused with real, concurrent multipath capability. Rather, these software solutions serve to activate channel failover or high availability path failover.

The IBM XIV storage system is an active/active storage array, architected to enable a volume to be presented for concurrent multiple connectivity through the FC and iSCSI ports. On the host side, multipath software would typically require an I/O device driver to be installed (or a layer closely coupled with the device driver). All I/O operations on the server need to pass through such a driver to facilitate multipath access. Concurrent multipath software is provided by the operating system vendor (e.g., MPIO) or by a specific storage vendor.

Multipath I/O Architectures

Multipath I/O software architectures can be classified into four main technology groups:

- ▶ **Host Bus Adapter (HBA) level** – This type of software primarily supports failover (transfer) of control between “heads” of a mid-range storage system (mostly host adapter failover) or internal storage controllers within storage arrays.
- ▶ **Storage vendor-specific multipath software** – This is first-generation multipath software; typically provided by a specific storage vendor (e.g. IBM® System Storage™ Multipath Subsystem Device Driver (SDD), EMC PowerPath). Most multipath software built by a storage vendor has been created to support only the storage of that vendor. In addition, in most cases, the software is difficult to remove or change. The host software code is at the device driver level, tightly integrated with the operating system, and requires patches to it. This aspect has been known to cause major problems in compatibility that can lead to loss of data and unavailability.



- ▶ **Symantec's multipath driver** – Storage Foundations from Symantec provides concurrent multi-path software that works with selected qualified storage vendors on specific operating systems. This is a chargeable item from Symantec, which provides one product across various storage and operating system vendors.
- ▶ **Host-based, open, native multipath I/O software** – This is second-generation multipath software, provided by the OS vendor. Offerings available on the market today include Microsoft MPIO framework, Solaris MPxIO, and Linux Mpath, among others.

TCO: Host-Native vs. Storage-Vendor Multipath Software

Increasingly, for both technical and commercial reasons, enterprises are adopting host-native multipath rather than storage-vendor-specific products such as those cited above.

Software Cost

On the commercial side, the primary motivation for choosing host-based (native) multipath software – particularly now, when one of the main IT challenges is to reduce expenses and overall IT operating costs – is that the software is free. It is included as part of the operating system.

Vendor-specific multipath software, in most cases, has a separate, per server license charge (SDD is an exception), and involves additional software management licenses.

Quality Issues

Host-based multipath is also widely believed to be of higher quality. The software is tightly integrated into the operating system, supported by OS developers, and benefits from a continuous stream of new features.

Storage-vendor-specific multipath must constantly catch up with operating system upgrades and updates. Since it consists mostly of device drivers, storage-vendor-specific multipath software is inherently sensitive to changes within the operating system and has a track record of frequently failing in production. The cost of such failure must be factored into a computer system's TCO.

Deployment Costs

Many industry anecdotes recount lengthy deployment delays as customers are forced to wait for the latest supported device driver to arrive from a storage vendor. The costs to an organization of not being able to use storage that has already been purchased include not only the capital costs, but the costs of overloaded storage systems, slow response times, and delay in the deployment of new applications.

Such hardware deployment delays and their related costs do not occur when using host-native multipath software.



XIV's Approach to Multipath Software

XIV strongly recommends the use of standard OS concurrent multipath software, such as MPIO, rather than storage vendor proprietary multipath software, for many reasons:

- ▶ **Native concurrent multipath software is available and mature.** As noted, the software is well integrated within the operating system, released with the operating system, and fully supported by the operating system vendor. In sharp contrast, non-native multipath software is at the device driver level or slightly above, making its development and certification tricky, especially for outside storage vendors. The constant need for third-party vendors to “catch up” often delays release and deployment of new storage platforms.
- ▶ **XIV is an active/active array, uniquely suited to native multipath.** As an active/active array, XIV architecture allows concurrent I/O processing from multiple paths. A common claim of storage vendors is that the native multipath software provided by some OS vendors lacks dynamic multipath or other advanced scheduling. In fact, the “round robin” provided by all the native multipath offerings works especially well in active/active storage arrays such as the IBM XIV system. As the adoption of multipath software grows – and this is already the trend – so, too, will the bundle of real-time and other features and capabilities that native multipath software provides.
- ▶ **Native multipath allows multi-vendor storage support.** Storage-vendor-specific multipath software takes full control of the I/O paths and, as such, excludes other storage vendors. It is well known in the industry that facilities to remove storage-vendor-specific multipath software from host systems either do not exist or do not work. A main concern of enterprise IT departments is the loss of control over their assets and being locked in to a single supplier.
- ▶ **Native multipath is features-rich and rapidly enhanced by operating system vendors as the main path to I/O devices.** It also allows customization (e.g. IBM XIV, Windows, DSM, or Linux multipath-tools), for smooth integration by the storage vendor of special features and statistics.
- ▶ **It's free!** Why pay for capabilities that are already part of the operating system? Most storage vendor proprietary multipath software is very expensive. It includes general license fees as well as “per host” client charges. This is true for all vendors other than IBM's SDD (although that software, too, comes at a cost). Host-based native multiple path software does not cost extra.



- ▶ **It helps meet customer standardization requirements.** IT shops sometimes seek to standardize on one supplier of multipath software for all servers and multiple-vendor storage solutions. Only Symantec provides such software, Symantec DMP, which is part of the Symantec Storage Foundations package. The IBM XIV Storage System has been certified and qualified by Symantec for the Storage Foundations suite. For specific support platforms, please refer to IBM XIV interoperability on SSIC and to Symantec support documentation. An ASL module distributed by Symantec is required for full support. The Symantec product meets customer requirements for multipath software common to all server hosts and storage vendors.

The IBM XIV system does not generally support proprietary multipath software such as IBM SDD and EMC PowerPath. The one exception is Symantec's Storage Foundations DMP. The XIV system supports this one offering because it itself supports all storage and all servers. Unfortunately, this product is expensive.

In the future, IBM XIV seeks to create customer packages inclusive of the XIV GUI, to simplify deployment of host-native multi-path environments.

Summary

The XIV system's active/active storage array is best able to exploit the benefits offered by host-based concurrent multipath I/O. In our opinion, host-based native concurrent multipath software is superior to storage-vendor-specific multipath software, for many reasons:

- ▶ better integrated with the operating system
- ▶ more reliable in operation
- ▶ receives the attention of OS vendors, with constant upgrades and new features
- ▶ no delays waiting for storage-vendor-specific device drivers
- ▶ keeps control in the customers' hands by allowing simultaneous support of multiple storage vendors
- ▶ is free

With these clear advantages, host-based native concurrent multipath software will continue to widen its lead over storage vendor-specific multipath software.