WINCHESTERSYSTEMS®
Purpose-Built Storage

Winchester Systems FlashDisk

Key Storage System Features

Contents

FlashDisk Storage Solutions

Winchester Systems FlashDisk storage systems feature advanced hardware design and provide comprehensive storage services at affordable price points, making them ideal for small and medium businesses (SMBs) as well as enterprises.

The FlashDisk series provides robust data protection to ensure the highest data availability in storage area network (SAN) and direct attached storage (DAS) scenarios. Featuring a modular architecture, thin provisioning, easy and intuitive management, and an attractive price-performance ratio, FlashDisk systems help businesses improve IT efficiency and stay on top of growing storage needs.

Key Features of FlashDisk Storage Systems

Fast Rebuild

As drive capacities continue to grow, the time required for storage to recover from a degraded status to full usability has become more critical. Large storage capacities mean recovery may potentially be very time consuming as there is more data to rebuild, which in turn leads to a greater emphasis on performance. If a rebuild takes too long and the process is not efficient, overall storage system performance becomes affected and the entire array may be compromised.

FlashDisk delivers a fast rebuild to avoid these issues, having been designed with performance as a main priority. Enhanced computing power allows them to perform important tasks such as RAID volume rebuild in far less time than that required by competing products in their price segment. Reduced time requirements for rebuild operations translate into lessened potential for disruption of your productivity, as you regain access to vital data within a shorter time frame, leading to minimized downtime. It also means system resources are devoted to the task for a briefer period, reducing impact on RAID array I/O performance.

FlashDisk systems accelerate rebuild operations through a combination of hardware and specialized firmware-side architecture. To demonstrate the advantages offered to you by FlashDisk fast rebuild, we have conducted tests that compare FlashDisk performance to that of comparable systems from NetApp and Dell. The results clearly show FlashDisk as the faster choice.

Vendor	Model	Drive	Rebuild time
Winchester Systems	**FlashDisk FX**	**3TB NL SAS**	**6.75 hrs**
NetApp	E5460	3TB NL SAS	8 hrs
Dell	MD3260	3TB NL SAS	8 hrs

Modular and Future-Proof Host Board Design

The modular hardware design on FlashDisk systems makes them easy to work with. Parts are cable-free for simple slot-in installation and removal. The ability to quickly swap main components renders FlashDisk products highly future proof. You can purchase a system with 1GbE iSCSI host interfaces and then easily upgrade it to 10GbE iSCSI when you need faster connectivity, without changing any parts beyond the controller module.

To accommodate such flexibility, FlashDisk systems offer a choice of host interface: 16G FC, 8G FC, 10GbE iSCSI (SFP+ or RJ-45), 1GbE iSCSI, and 6Gb/s SAS. For high availability and non-stop operation, FlashDisk systems also feature redundant components. Redundancy allows a RAID array to continue working even if one crucial component experiences failure or requires maintenance. All FlashDisk systems have redundant power supplies. They are also offered with dual controllers, which effectively doubles performance while ensuring high availability – the system stays live even if one controller goes offline.

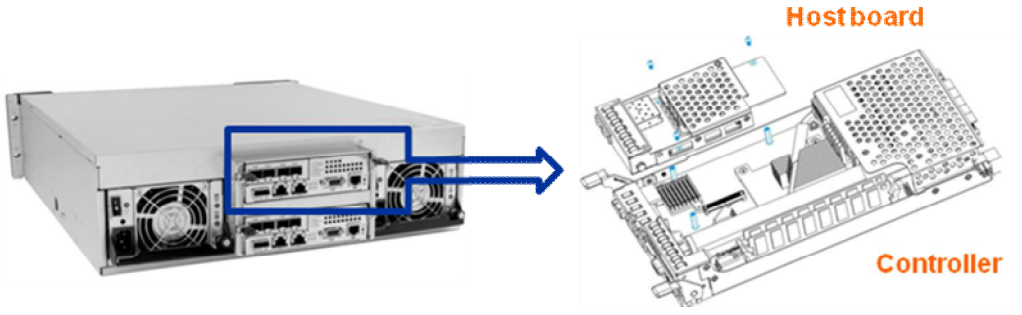

Host board upgrade illustration

Thin Provisioning

Thin provisioning is a technology that allows storage resources to be used in the most flexible manner. By automatically allocating system capacity to applications as needed, thin provisioning helps achieve up to 80% storage utilization while significantly reducing power consumption.

Thin provisioning technology allows users to allocate a large amount of virtual capacity to an application regardless of the physical capacity actually available. At initial setup, thin provisioning does not physically allocate capacity to the prescribed data volume, and the actual space is used only when data writes occur. This on demand method for capacity allocation not only optimizes storage utilization, but also greatly simplifies capacity planning and management tasks. In order to help users easily monitor capacity utilization, FlashDisk systems automatically issue notifications when total capacity utilization is close to reaching user-set thresholds. If you wish to expand capacity, you can do so non-disruptively

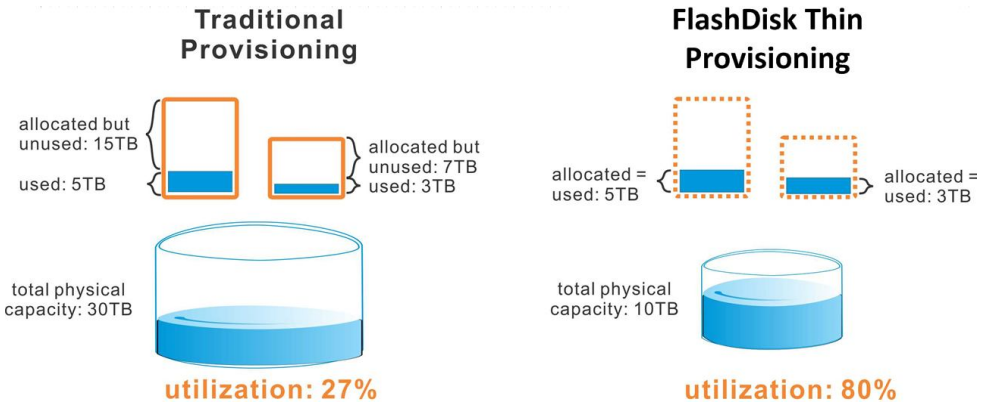

Traditional Provisioning vs. Thin Provisioning

Space Reclamation for Thin Provisioning

Thin provisioning on FlashDisk systems features space reclamation functionality. An often seen phenomenon with thin provisioning is that the space of deleted data on the host system cannot be re-used by the storage system in thin provisioned volumes due to a lack of communication between the host file system and the storage system. This could potentially lead to a situation where space in the storage system is filled to capacity, while on the host a large proportion of the equivalent storage space is still available for use.

With space reclamation on FlashDisk systems, the storage array is made aware of actual space usage in the host and reclaims capacity accordingly, thereby optimizing the utilization of available storage space. Currently, space reclamation functionality is only supported in Microsoft Windows (NTFS) environments.

Capacity Scaling

Capacity scaling involves adding expansion enclosures to storage systems. These expansion enclosures provide large increases in capacity. FlashDisk systems can be scaled to a maximum of 360 drives via enclosures, compared to previous generation models which supported a maximum of 240 drives. They offer users significant flexibility in terms of capacity growth.

Local Replication

Local replication capabilities on FlashDisk systems include snapshot and volume copy/mirror. Snapshots allow users to perform instant, low impact backups with space efficient differential copies. In the instant of snapshot creation, a point-in-time data image is taken without disrupting online applications. Based on the image, data changes copy to the snapshot volume when new writes occur. With copy-on-write design, snapshots protect data from accidental modifications, deletions, and corruption with minimal capacity requirements and performance overhead. By accessing a snapshot copy as the desired recovery point, you can immediately restore system availability from data disruption.

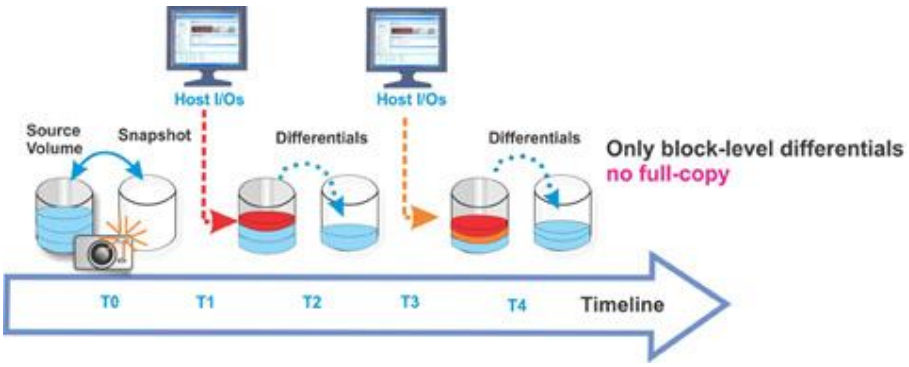

FlashDisk Snapshot Feature

Volume mirror/copy allows users to non-disruptively create full data copies in a single storage pool. Copies created with volume mirror can be constantly synced with source data or split as independent copies. If a volume fails, these copies can be continuously synced and leveraged for quick recovery without any data loss. Split copies can be used to enhance operational resilience by allowing shared access to critical information among production and batch applications, such as testing, data mining, scientific analysis, and tape backup. As business needs require, split copies can be quickly re-synced with the source.

Users can also create copies independent from the source with volume copy technology. These copies are ideal for time-critical applications since they are immediately available for use even before the copying process is complete. Similar to split copies created by volume mirror, editing independent copies does not affect the service level or data integrity of primary applications.

Remote Replication

Remote replication allows users to create full data copies across storage systems, either in synchronous or asynchronous mode. If source data fails due to system malfunction or disasters, users can leverage disk-based remote copy to restart services within mere minutes. If the source needs to resume its role, it can be quickly synced with the remote copy to compensate for any differences.

To further ensure the integrity of remote data, remote replication allows users to protect remote copies with snapshot technology. Granular snapshot images can help restore corrupt remote copies in seconds when service resumption is an urgent priority.

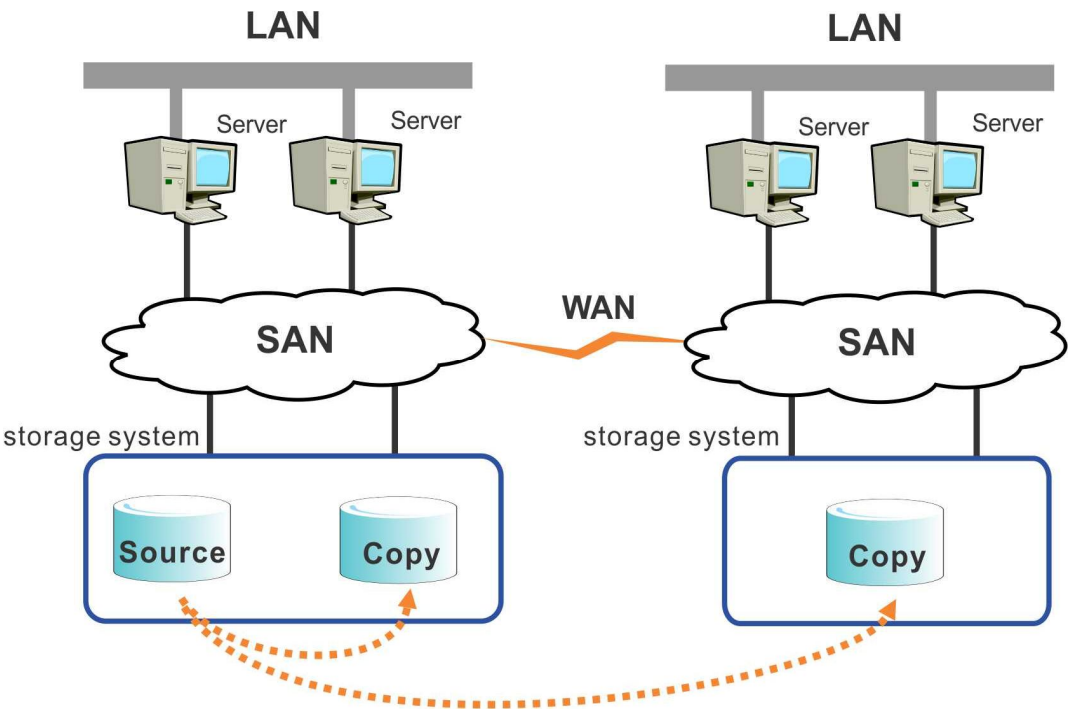

Local Replication and Remote Replication

When using asynchronous remote replication, you have the option to implement data compression, allowing data to be sent more efficiently.

Cache Backup Module

FlashDisk systems improve data protection with CBM (Cache Backup Module). This is achieved with a BBU (Battery Backup Unit) and FBM (Flash Backup Module). In the event of a power outage, the BBU automatically supplies power to allow cached data to be written into the FBM for permanent retention.

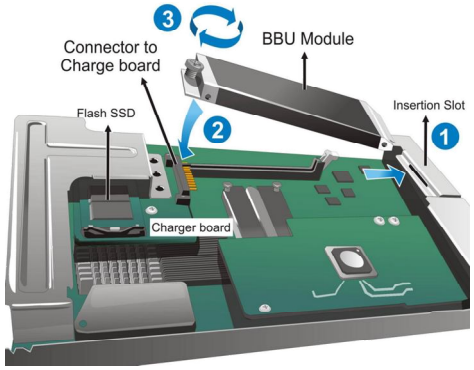

CBM (Cache Backup Module)

Hybrid Connectivity: Cost Effective Remote Replication and SAN Tiering

FlashDisk series systems feature hybrid host connectivity, including Fibre Channel (FC) and iSCSI ports. With hybrid connectivity, you are able to implement highly cost effective remote data replication via iSCSI. Traditionally, protocol converters or FCIP routers are required to replicate data sets from a FC-host storage system onto another FC-host storage system at a remote site over Ethernet. One router needs to be installed at the local site, while a second needs to be installed at the remote site. These FCIP routers require significant investment.

Remote replication on hybrid systems can be implemented via iSCSI, which can be directly connected to IP networks. This allows users to skip protocol conversion and directly send data over Ethernet to remote sites. In this way, you do not need to separately buy two FCIP routers, leading to considerable savings.

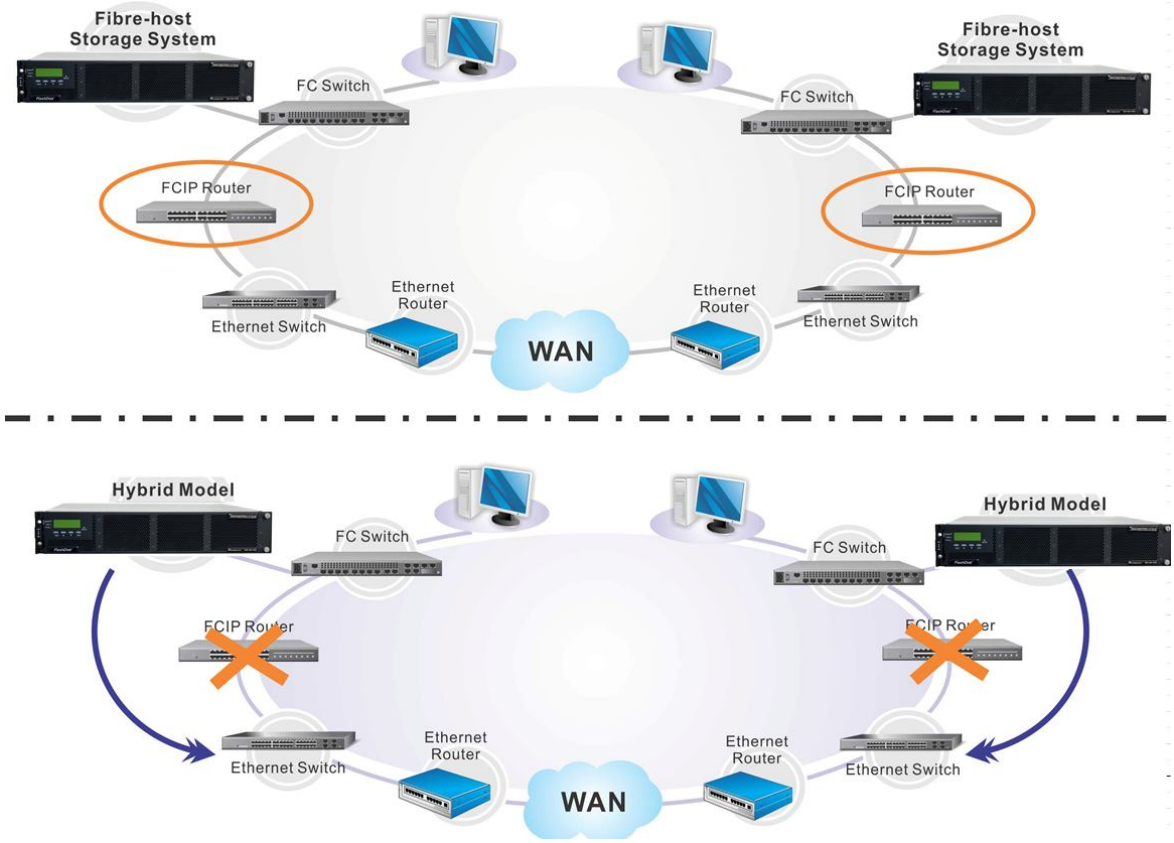

Remote Replication without and with Hybrid Connectivity (upper/lower)

In addition, hybrid connectivity enables users to deploy FC and IP SAN at the same time to take advantage of consolidated SAN tiering. SAN tiering means that a single storage system can meet service level requirements in FC SAN and IP SAN configuration concurrently, extending its value and usefulness thanks to better utilization.

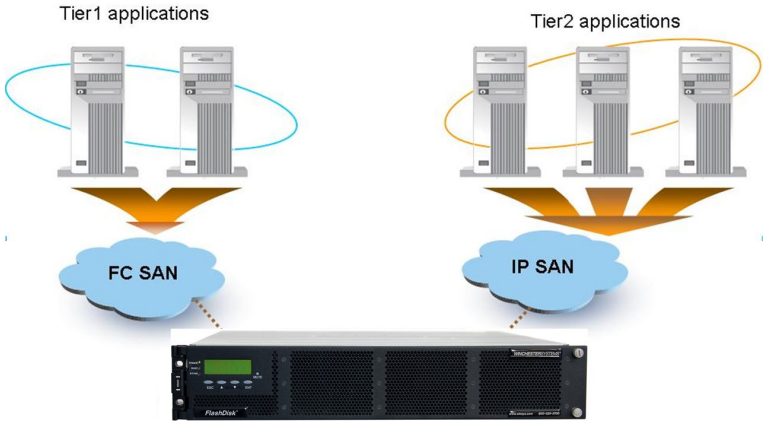

SAN Tiering

Automated Storage Tiering

SAN improves data management efficiency, but is not the most efficient approach to serving multiple datastores since by default a SAN treats all data equally.

Based on frequency of access and tenure, data can be divided into “hot” and “cold” categories. While standard non-tiered SAN designs utilize resources equally for both categories, the addition of automated storage tiering leverages different drive types to offer better service and make the most of investment in storage hardware. Automated storage tiering intelligently analyses data hotness (or access frequency) and duration of retention (how long data has been stored), and moves application workloads to the best-suited tier.

For example, very hot data can be tiered to SSD storage for the quickest access available, while cooler data and applications can be tasked to more affordable and plentiful mechanical disk drives on other tiers. Taking two identical investments in storage, the one with automated storage tiering brings dramatically better performance and efficiency.

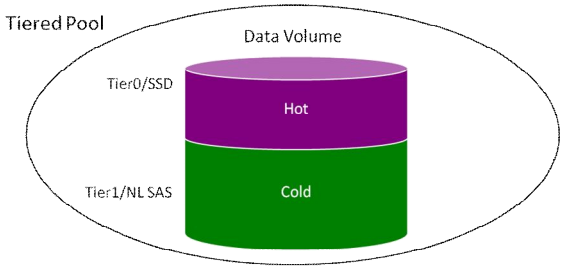

Tiering illustration

SSD Read Cache Pool

Current FlashDisk FX and VX Series storage systems support SSD Read Cache to dramatically improve read performance in the most essential services and applications.

How SSD Cache works

Cache is a component that transparently stores data so that future requests for that data can be served more efficiently. It is fundamental to storage, especially for read-intensive applications. SSD Read Cache allows fast SSD to be used to extend the cache pool capacity of storage systems and can store large amounts of frequently-accessed data. SSD Read Cache increases more than just cache capacity. It also improves the cache hit rate, in effect serving more data with considerably greater efficiency than non-SSD caches.

FlashDisk SSD Read Cache is an intelligent solution that dramatically improves read performance, especially for read-intensive applications, and greatly minimizes latency while supporting large cache pool capacities. It is also simple to set up, manage, and maintain via the FlashDisk Global Manager interface. Therefore, it is recommended to use SSD Read Cache if application workload has a high percentage of read behavior and users require specific amounts of vital data that consistently depends on intensive read requests.

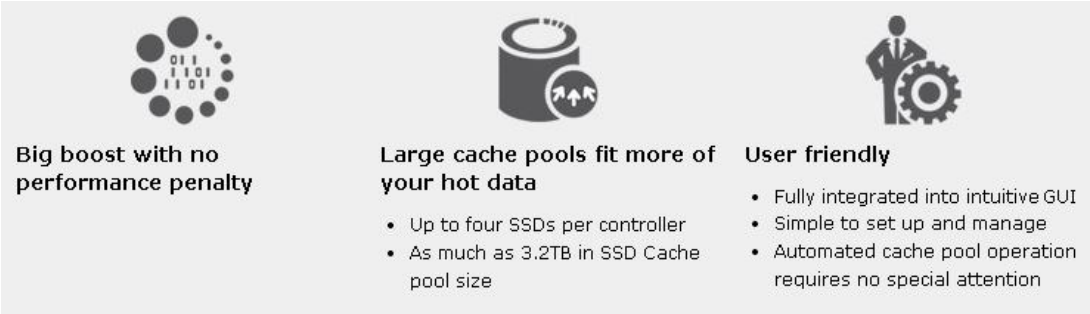

Summary of SSD Read Cache features and benefits