

IntelliFlash™ Operating Environment

The IntelliFlash Operating Environment (OE) is flash-optimized and seamlessly integrates multiple grades of storage media to deliver optimal performance (high performance at low latency) and the best economics for a wide range of enterprise applications.

The IntelliFlash OE delivers a single storage platform with a choice of media including NVMe™ flash, performance flash, dense flash, hard disk drives, or a mixture of flash and disk for your storage. You dial up or down the amount of flash to match the performance and economics of your business applications.

We believe that an enduring enterprise storage software and system architecture must accommodate, with ease, many grades of persistent media, today and in the future. Our IntelliFlash OE delivers Total Investment Protection by being able to rapidly incorporate newer forms of persistent media (dense flash, or even persistent memory technologies) seamlessly with full data management software features. The inclusion of the latest technology available on the market drives down the effective cost of storage while driving up performance and capacity.

Media Optimization

The Media Optimization Layer provides the foundation for reliably storing data on different types of storage media, and optimizing the use of media within the system.

Media Management

- IntelliFlash OE is optimized for the underlying geometry of the medium (NVMe flash, performance flash, dense flash, or hard disk) to ensure long life even under high I/O workloads. Data writes are aligned to sector boundaries for disks and to the native page sizes of flash to avoid fragmented I/O and eliminate unnecessary writes to media.
- To accommodate for limited write cycles of flash storage, the IntelliFlash OE tracks flash wear and moves data around to help ensure uniform wear across flash pages.
- To avoid degraded performance resulting from data blocks marked for deletion by operating systems but unclaimed by the flash device, IntelliFlash works in concert with operating systems to continually scrub and free up blocks to minimize write amplification.
- Data is deduplicated and compressed before being committed to media (flash and disk). This minimizes writes and extends the life of the media. To ensure flash endurance, IntelliFlash arrays are built with enterprise quality MLC flash, which has 10X the write endurance of consumer-grade flash.

Storage Pooling

- Provisioning and management of storage in the IntelliFlash OE is optimized and greatly simplified by virtualizing and aggregating physical media into a pool of resources that can be allocated as LUNs or shares. Pool capacity can easily be expanded online with very little management effort.
- IntelliFlash OE employs dynamic stripe widths to eliminate the performance overhead and media wear resulting from the

read-modify-write operations associated with traditional RAID. Reconstruction of failed drives is also faster since unallocated data is not copied.

- The IntelliFlash OE pipelined I/O design ensures performance and optimal writes to media by efficient I/O scheduling and coalescing of write operations.

Resiliency

- All media (flash and hard disks) in IntelliFlash Arrays are dual ported and accessible through a pair of highly available, redundant controllers. The controllers are configured in an active/active manner and can be used simultaneously for data access. IntelliFlash OE seamlessly and non-disruptively fails over I/O from one controller to the other, in the event of a controller failure.
- Regardless of the protocol, IntelliFlash OE enables redundant media fabrics by aggregating I/O ports within and across the two controllers, for high data availability. Online hot spares and quick drive rebuilds minimize exposure to downtime in the unlikely event of multiple failures.
- The IntelliFlash OE enables all hardware components, including media, to be replaced online with zero downtime. Software upgrades to the array are performed with no downtime or loss of access.

Metadata Acceleration

Intelligent metadata handling is a patented, core innovation in IntelliFlash OE and is the key to delivering advanced data services at the speed and scale of I/O that flash offers.

Aggregation

- Traditional storage systems store metadata interspersed with data on disks. Over time, as data is being modified, deleted, and rewritten, metadata and data becomes very fragmented on disk. In addition, traditional data de-duplication implementations also can cause metadata to multiply and grow rapidly. As metadata grows it causes significant deterioration in a system's behavior – especially performance – over time. To avoid this, IntelliFlash OE automatically separates metadata from application data. IntelliFlash's patented Metadata Aggregation and Placement technology organizes, aggregates, and places metadata on low-latency media, enabling the acceleration of advanced data services such as deduplication, compression, snapshots, clones, and thin provisioning.

Placement

- IntelliFlash OE organizes metadata into flexible, multi-layered data structures and distributes them across multiple flash media for full data integrity and protection. Such a multilayered metadata management scheme facilitates storing and managing massive amounts of data at high I/O rates. Additional benefits include high-performance data

management regardless of whether data is stored on flash or disk, inline data reduction at various block sizes, and easy scaling of system capacity and performance.

Caching

- To deliver consistent I/O, data services, and performance at scale, IntelliFlash OE employs a Caching and Scaling engine. This engine caches the most frequently-accessed metadata and application data in high performance storage layers (DRAM, persistent memory, and flash) enabling predictable low latency for data reads and writes. The caching algorithms are adaptive and optimized for various I/O patterns to ensure high cache hit rates. These innovations in the IntelliFlash OE provide the foundation for our flash arrays to deliver consistent and predictable performance for a wide range of enterprise applications.

Scaling

- The Caching and Scaling engine is designed to seamlessly adapt to media with differing latencies by using multiple levels of cache.
- The caches employ DRAM, persistent memory, and dedicated high-performance flash media for metadata and application data. These devices are protected against device errors or failures. As data and consequently metadata grow over time, the IntelliFlash OE enables automated expansion of metadata storage space. This in turn ensures dynamic scaling of metadata management and advanced data services to handle massive data sets.

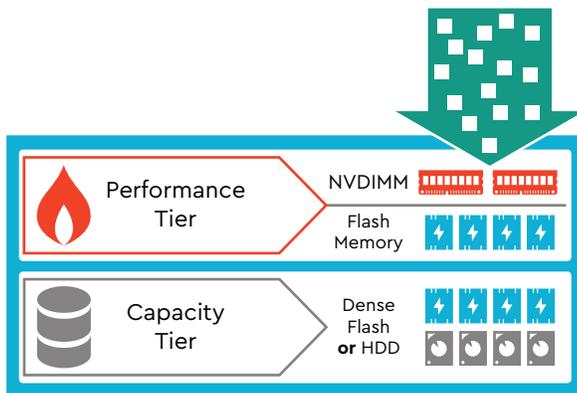


Figure 1. IntelliFlash metadata acceleration and caching architecture

Data Services

The IntelliFlash OE data services reduce the effective cost of owning and operating storage, and enable continuous data availability, data protection, and recovery features.

Data Reduction

- The IntelliFlash OE data reduction service includes inline deduplication, inline compression, and thin provisioning.
- Deduplication can be enabled at the storage pool level, or at a per LUN or share level. Moreover, each LUN or share can be configured with different block sizes ranging from 4KB to 128KB.

- A choice of compression algorithms is provided to meet the needs of various workloads. Compression, like inline deduplication, can be turned on or off at the storage pool level or at the individual LUN or share level.

- Inline deduplication and compression serve a multifold purpose. By compressing data and eliminating redundant blocks, they reduce the overall storage consumption and minimize media wear. By enabling more data to be served out of persistent, high performance, storage media, they also accelerate performance.

- IntelliFlash OE innovations allow the read cache to be persistent across controller failures resulting in uniform, consistent performance even after a controller failure event.

- Thin provisioning optimizes capacity utilization through just-in-time consumption of actual storage. Integration with VMware® through the VMware APIs for Array Integration (VAAI) enables graceful pause-and-resume of virtual machines when a thin provisioned LUN runs out of space.

Data Protection

- The IntelliFlash OE data protection services help keep data safe, secure and available with no impact on performance. The services include end-to-end checksum data integrity validation, auto-healing, at-rest encryption, snapshots, clones, full volume copies, and replication.

- For every data block written to the array, a checksum is computed and stored. Complete data integrity is ensured by storing the checksum and data in separate data structures. Multiple copies of the data are stored if needed. To protect against silent data corruption every read operation of a block is verified against the block's checksum and in the event of a mismatch the data is recomputed, verified, and returned to the application.

- Security features include 256-bit AES encryption of data at rest, helping protect against theft or loss of a drive during maintenance or transit. IntelliFlash delivers in-line encryption of data on flash and hard disk with no loss of performance. Key management is performed natively in the system without user intervention.

- Instantaneous, read-only point-in-time images of a single data volume or a set of volumes are taken with snapshots. The snapshots require no space reservation and little to almost no space. Subsequent snapshots allocate just enough space to accommodate new or changed blocks in the data volumes.

- Writeable, point-in-time images are created with the cloning feature. Clones, like snapshots, are thin and perform sparse allocation as needed for new data. Both snapshots and clones are application and virtual machine consistent.

- The IntelliFlash OE supports creating full copies of a data volume within the same pool, across pools in a single array or across pools in multiple arrays.

- Asynchronous snapshot-based replication of data volumes can be performed across pools and across arrays – including replication between an all-flash and hybrid flash array, for a cost effective disaster recovery solution. Various replication topologies including one-to-many and many-to-one are supported.

Data Recovery

- In case of data loss or corruption, snapshots enable near-instantaneous recovery.
- Data loss due to media failures is mitigated through multiple data resiliency options. The intelligent data reconstruction capabilities of IntelliFlash ensure quick recovery and minimizes the window of exposure to another storage device failure.
- Site to site remote replication facilitates disaster recovery and business continuity in case of a data center outage. Periodic asynchronous replication and data compression ensure optimal utilization of the WAN across datacenters by minimizing the amount of data transmitted for replication.
- The ability to take point-in-time VM consistent and application-consistent snapshots, and replicate data in a coordinated fashion using aggregated volumes in a consistency group enables seamless application data recovery.

Data Services

- The IntelliFlash OE is architected to natively provide block and file protocol access. Supported block protocols include iSCSI and Fibre Channel; file protocols include NFS, CIFS, and SMB 3.0. The IntelliFlash OE is extensible, enabling other protocols to be added seamlessly in the future.
- All protocols can be used simultaneously over a variety of storage ports.
- For all protocols, redundant fabric capabilities allow data access over multiple paths.

Management Flexibility

- The IntelliFlash OE provides flexibility through multiple management options.
- The Web UI is a simple and easy-to-use interface for comprehensive management of IntelliFlash arrays.
- Storage management and data protection tasks can be scripted and automated through a programmable, task-oriented REST API.
- The vCenter web client and desktop client plugin allows VMware datastores running on IntelliFlash Arrays to be seamlessly managed through VMware vCenter®.
- Microsoft® Hyper-V virtual machines can be managed through Microsoft Systems Center Virtual Machine Manager (SCVMM).
- The IntelliFlash Call-Home capabilities provide various alerts that are sent to administrators and Customer Support to ensure proactive and timely intervention.

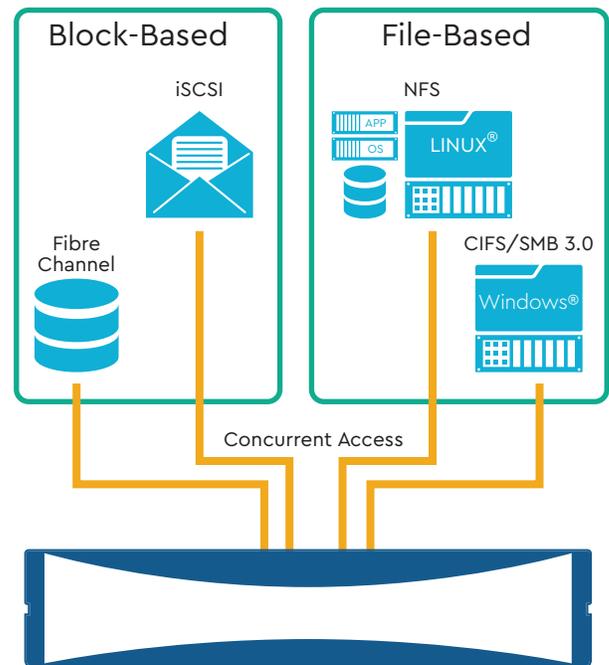


Figure 2. IntelliFlash concurrent multi-protocol access

IntelliFlash Arrays deliver incredibly high performance while maximizing efficiency not only for IT, but for your business, by keeping your storage costs in check and making your data come alive. For more information on how IntelliFlash can help make your data come alive, visit westerndigital.com/intelliflash

Western Digital.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000
www.westerndigital.com

© 2018 Western Digital Corporation or its affiliates. All rights reserved. Produced 8/18. Western Digital, the Western Digital logo, and IntelliFlash are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. All other marks are the property of their respective owners. Before that, please also add: The NVMe word mark is a trademark of NVM Express, Inc. VMware vCenter is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks are the property of their respective owners.