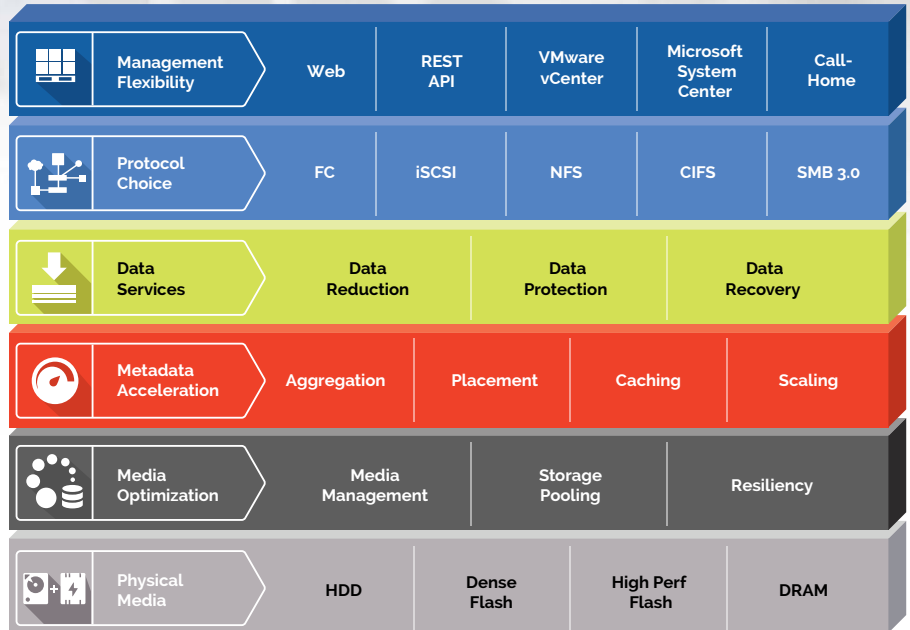


IntelliFlash™ Architecture



IntelliFlash™ is our flash-optimized storage software architecture. It seamlessly integrates multiple grades of storage media to deliver optimal performance (throughput and latency) and best economics for a wide range of enterprise applications.

It is designed to give you a single platform with the flexibility of using all flash, a mixture of flash and disk, or a mixture of high-performance flash and high-density flash for your storage. You dial up or down the amount of flash to match the performance and economic demands 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 architecture 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



IntelliFlash 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 storage system.

Media Management:

- IntelliFlash is optimized for the underlying geometry of the medium (flash or hard disk) to ensure long life even under high I/O workload. 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, IntelliFlash tracks flash wear and moves data around to 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 and minimize write amplification.
- Data is deduplicated and compressed before being written to media (flash and disk), minimizing writes to the media and extending the life of the media. To ensure flash endurance, Tegile arrays are built with enterprise quality eMLC flash which has 10 X the write endurance of consumer grade flash.

Storage Pooling:

- Provisioning and management of storage in the IntelliFlash architecture is optimized and greatly simplified by virtualizing physical media and aggregating them into a pool of resources that can be allocated as LUNs or file shares. Pool capacity can easily be expanded online with very little management effort.
- IntelliFlash uses 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 pipelined I/O architecture 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 driven Tegile 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.
- Regardless of the protocol, IntelliFlash 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 architecture allows for 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.



IntelliFlash Metadata Acceleration

Intelligent metadata handling is a patented, core innovation in IntelliFlash 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 data and metadata together, with metadata being interspersed with data on disks. Over time, with data being modified, deleted, and rewritten, metadata becomes very fragmented on disk. In addition, traditional data de-duplication also can cause metadata to multiply and grow rapidly. As metadata grows it causes significant deterioration in a system's behavior over time. To avoid such behavior, IntelliFlash automatically separates metadata from application data. IntelliFlash's Metadata Aggregation and Placement Engine 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 organizes metadata into flexible, multi-layered data structures and distributes them across multiple flash media for full integrity and protection. Such a multi-layered 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 across flash and hard disks, and easy scaling of system capacity and performance.

Caching:

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

Scaling:

- The Caching and Scaling Caching engine is designed to seamlessly adapt to media with differing latencies by using multiple levels of cache.
- The caches use DRAM and dedicated high performance flash storage media for metadata and application data. These devices are mirrored for protection against device errors or failures. As data and consequently metadata grow over time, IntelliFlash 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.



IntelliFlash Data Services

The IntelliFlash data services reduce the effective cost of owning and operating storage through data footprint reduction, and enable continuous data availability through a comprehensive set of local and remote data protection and recovery features.

Data Reduction:

- The IntelliFlash data reduction services include inline deduplication, inline compression, and thin provisioning.
- Deduplication can be enabled at the storage pool level, or even at a per LUN or file share level. Moreover, each LUN can be configured with different block sizes ranging from 4KB to 128KB.
- A choice of compression algorithms are 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 file 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 (flash and hard disks) wear. By enabling more data to be served out of flash and DRAM, they also accelerate performance.
- Tegile 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 provisioning of 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 data protection services keep data safe, secure and always available with no impact on performance. The services include end-to-end checksum protection, auto-healing, at-rest encryption, snapshots and clones.
- 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 tree structures in the block storage hierarchy. Multiple copies of the data are stored if needed. To protect against silent data corruption every read operation of a data block is verified against the block's checksum and in the event of a mismatch a copy of the data block is read, verified and returned to the application. The system auto-heals the incorrect block using RAID/Mirror copies.
- Data security is provided using 256-bit AES encryption of data at rest, protecting 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 required for encryption 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 using snapshots. The initial snapshot requires no space reservation and takes little to 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 using 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.

Data Recovery:

- In case of data loss or corruption, snapshots enable instantaneous recovery.
- Data loss due to media failures is mitigated through multiple RAID/ Mirroring options. The intelligent data reconstruction capabilities of IntelliFlash ensure quick recovery and minimize the window of exposure to another media 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.



Protocol Choice:

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



Management Flexibility:

- The IntelliFlash architecture provides management flexibility by providing multiple management options.
- The Web UI is a simple and easy to use management interface for comprehensive management of Tegile 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 data stores running on Tegile 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 Tegile Customer Support to ensure pro-active and timely intervention.



Management Flexibility



Protocol Choice



Data Services



Metadata Acceleration



Media Optimization



Physical Media