Tier-X™

reduxio

Tier-X™

Continuous autonomous data relocation engine integrating flash and high-capacity media in a single storage pool

KEY BENEFITS

FLASH FOR CAPACITY, NOT JUST PERFORMANCE
Optimizes the use of flash by storing actual data on it, not just cache

....

OPTIMIZED DATA RELOCATION
Fine-grain tiering of deduplicated and compressed data

....

FULLY-AUTOMATED, PERIOD
No configuration, no policies, no schedules

....

FAST & ADAPTIVE
Automatically reacts to changing workloads in real-time
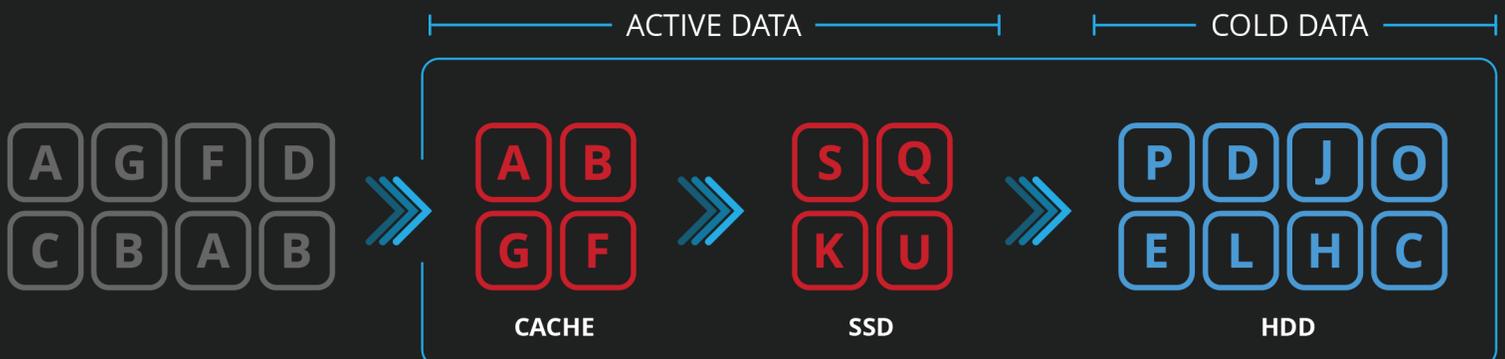
TIERING - NOT CACHING

Tier-X is an active storage tiering technology integrating high performance flash and high capacity media to form a single, multi-tier storage pool. Data is deduplicated and compressed in memory on ingest to store only unique blocks across cache, volumes and history. Tier-X provides consistent low latency for both reads and writes from or to any storage tier.

MORE PERFORMANCE PLEASE

As data centers get consolidated and virtualized, the required performance from storage systems has increased sharply. While flash storage promises much higher performance, the cost premium over high-capacity disks demands that a combination of flash and disk be used to ensure affordability across all applications and use cases. But vendors either avoid disks altogether, focusing only on high-end, low-capacity use cases, or add a limited flash-based cache, with marginal gain in system performance and no gain in capacity since cached data means duplicated data - one copy in flash and one on disk.

THE SOLUTION: CONTINUOUS TIERING

Tier-X provides the cost benefits of tiering and the performance benefits of caching. Fine-grained, continuous and autonomous tiering of optimized data ensures that the system is reactive to the changing needs of applications.

ACTIVE DATA — COLD DATA

| A | G | F | D |
| C | B | A | B |

| A | B |
| G | F |

CACHE

| S | Q |
| K | U |

SSD
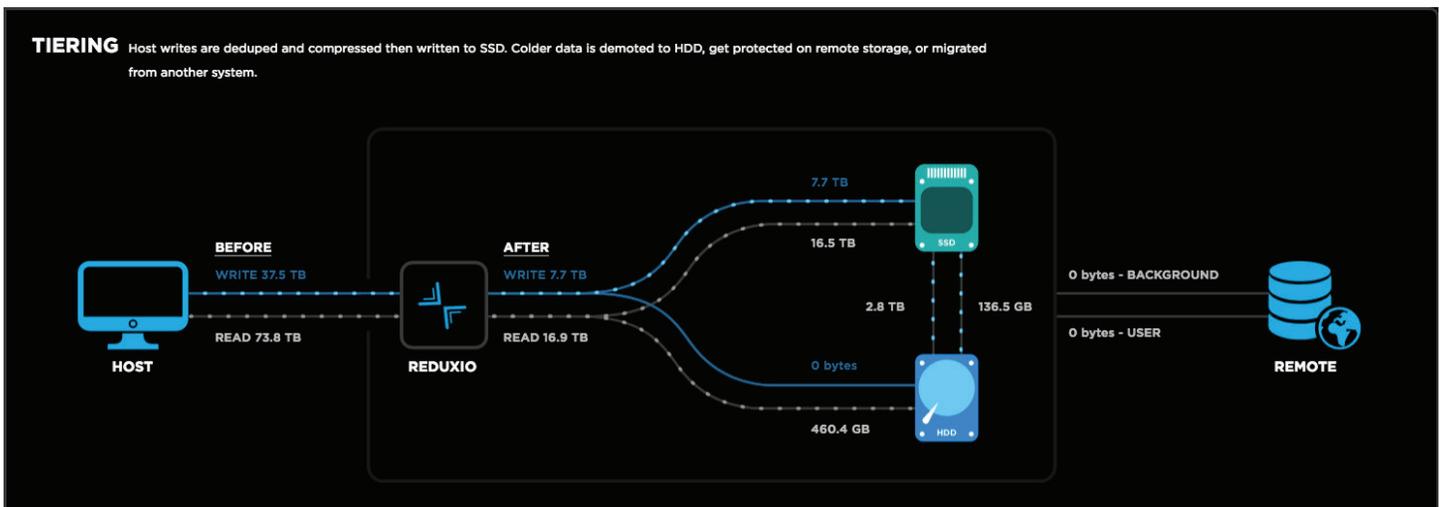
| P | D | J | O |
| E | L | H | C |

HDD

HOW DOES IT WORK?

All incoming writes are deduplicated, compressed and stored in protected memory cache, and then destaged to the SSD tier configured with enough capacity (typically more than 10TB) to hold the entire active working set of enterprise customers - and even more. Effectively operating similar to all-flash arrays, Tier-X enables high performance and low latency for both reads and writes. As data blocks on the SSD become cold, they are immediately moved to the HDD tier. Writes are coalesced using an algorithm that considers the relationship between data blocks and the unique characteristics of the media to ensure that performance of the slower disk tier is not affected when the data is moved. Reads are fast as they are usually done from cache or SSD. When a block is read from the HDD, Tier-X algorithm will automatically read other related blocks to improve read performance.

TIER-X - FAST AND EFFICIENT

Tier-X is the key component for achieving flash performance from Reduxio's multi-tier system. Data reduction that includes both deduplication and compression, allows Tier-X to store more data in the flash and HDD tiers. Data is stored on both SSD and HDD using a special algorithm that continuously and seamlessly moves blocks between the tiers according to their usage. The Tier-X engine improves the system performance by grouping blocks and then promoting the entire group once one block was accessed, while enabling the usage of the high capacity tier to lower the cost of the system.

TIERING Host writes are deduped and compressed then written to SSD. Colder data is demoted to HDD, get protected on remote storage, or migrated from another system.

Large flash tier services all writes and most reads providing low-latency read and write performance

Adaptive to workloads and media supporting highly random workloads that would otherwise require all-flash arrays

---

Optimized use of capacity since all data is written after dedupe and compression

Continuous – Tier-X is active at all times and dynamically and instantaneously adapts to changing workload conditions

---

Operating at the block level, Tier-X relocates only the necessary data, unlike legacy sub-LUN tiering

SSDs are optimally used to store the hot blocks of data, in addition to the cache layer

---