

0 1

0 2

0 3

0 4

# NoDup<sup>®</sup>

redu»io

# NODUP®

Maximum capacity savings with in-memory block-level dedupe and compression

## KEY BENEFITS

### MORE THAN 4X SAVINGS

Global dedupe and compression across cache, volumes, clones and history regularly achieves high savings ratio

### NO CONFIGURATION, NO MANAGEMENT

Always-on reduction, with no best practices or any tuning required

### BUILT FOR HIGH PERFORMANCE

Optimized resource use improves performance of flash, disk and cloud storage

### SAVINGS FOR 3RD-PARTY STORAGE

Reduces the overall required capacity for data migrated into a Reduxio system, freeing up for other purposes

## WHAT IS NODUP

NoDup is a patented real-time data reduction technology that brings unparalleled storage efficiency and density to Reduxio storage systems. The NoDup engine removes duplicated data and compresses the unique blocks. The data optimization process occurs at ingest as data is received from a host. This means that deduplication is global and no duplicate data is stored on the system across all volumes, clones, and historical data retained for BackDating. Unlike other solutions, both deduplication and compression are combined for maximum savings for practically any data – databases, virtual machines, file shares and more.

## WHY NODUP?

Data is a critical resource that requires 24x7 fast access. More and more data is being generated every day, and it is required to keep it for longer duration. On top of that, multiple copies are created for backups, analytics and development, which means even more data needs to be stored and managed. This growth requires more storage and more floor space, but IT budget is usually limited and supporting such growth is a challenge.

Data reduction technologies have been added to some legacy storage systems to try and relief some of those issues but they typically causes performance degradation. As a result, data reduction, when used, is enabled only




on a few select volumes or used only on archive, less active data. Most of the solutions do not support global deduplication which means lower savings.

Reduxio's TimeOS was designed with NoDup, an in-line, in-memory data reduction engine as an integral part of the architecture. Maximum reduction is achieved as all data - primary, clones and data history are deduplicated together. With NoDup, history for BackDating is deduplicated with the primary data for minimal space consumption. NoDup reduces the bandwidth requirement for data sent to a NoRestore repository and minimizes the data footprint on the repository - irrespective of whether the repository is another storage system or is in the public cloud.

### DON'T DEDUP – NODUP!

NoDup ensures that only unique blocks are stored in the system. Since the process is done prior to data being placed in cache, NoDup provides a significant boost to performance. More data can be served from memory cache and since less data is moved between tiers it maximizes the performance of the tiers and minimizes resource utilization. For example in the case if VDI all IO could be served from memory significantly reducing boot times. NoDup dramatically reduces data footprint, boosts storage performance and supports advanced data management capabilities - a perfect solution for one of today's top challenges.



Always-on – No need to give up reduction for certain workloads

In-line in-memory reduction increases the effective size of cache

Global reduction - A single large dedupe domain across volumes, clones and history

Maximum savings with dedupe & compression combined

Entire VM environments are collapsed to a single copy that is typically served right from cache

Tiering data movements are minimized since data is already reduced