

THE IDEAL FLASH TEMP STORAGE AND GZIP + ERASURE CODING ACCELERATOR




Solve Big Data Processing & Storage
I/O Bottlenecks Simultaneously
(Teragen + Terasort)

60% Increase in
Job Throughput

30% Cost per Job
Reduction

COMPUTATIONAL STORAGE

ScaleFlux CSS Opens Compute and Storage I/O Bottlenecks



Entire Cluster TCO Savings with Easy Integration

Return-on-Investment Analysis: CSS for 9-Node Hadoop Cluster


Baseline of 12 HDDs vs. Baseline + 1.6TB CSS

60% more jobs completed over 3 years

30% dollar-per-job savings

Integrates as easily as an SSD

Dramatically improves server utilization



CSS 1000 Series

Server
12 HDDs * 9 Nodes

Rack

ScaleFlux™ Computational Storage Subsystem (CSS) Product Info

Form Factor

- PCIe AIC & U.2 Drive

Flash Capacity

- 1.6 / 1.92 / 2.0TB
3.2 / 3.84 / 4.0TB
6.4 / 7.68 / 8.0TB

Compute Engines

- GZIP Compression, Erasure Coding (RS), KV-Store
- AES-128/256, SHA-3, ... and others are also available

Adding CSS to big data servers optimizes infrastructure efficiency and maximizes your ROI. CSS integrates as easily as an SSD, but provides much more value with its computational engines, yielding significant savings in TCA & TCO.

Whether your Hadoop deployment is looking to increase job throughput or reduce hardware footprint, ScaleFlux CSS is the ideal solution.

For an in-depth Hadoop Application Note, please request at info@scaleflux.com

UNLEASH HADOOP PERFORMANCE. SCHEDULE A POC