

EMC DSSD D5 RACK-SCALE FLASH APPLIANCE

Game-changing performance for current and future applications

DO MORE WITH YOUR PERFORMANCE AND DATA-INTENSIVE APPLICATIONS

Trends in social, mobile, and the Internet of Things have led to more users, more devices, and exponentially more data. And this data is increasingly more varied, and no longer primarily transactional in nature nor stored in relational databases.

The results of these changes combined with ever-increasing computing power have been twofold. First, existing applications have had to do more by processing larger data sets faster. In addition, a new generation of applications has arisen to deliver greater business value—by providing new, more innovative ways of looking at data and by deriving more useful and predictive information from greater data sets and more data types.

SUPERIOR PERFORMANCE FOR DATA-INTENSIVE WORKLOADS

Benefits

- Unlock competitive agility by enabling real-time analysis of larger, more complex data sets through unprecedented IOPS, bandwidth, and latency
- Accelerate business innovation through new applications that can take advantage of higher performance capabilities
- Reduce storage TCO with lower \$/IOPS and fewer database and CPU licenses



Thousands to Millions of Users

Siloed

Compute-centric Apps



Billion+ Users
Distributed
Interactive

Structured Data

Batch Processing

Small Working Sets



Structured and Unstructured
Real-time
No Defined Working Set

BREAKING THROUGH THE I/O BOTTLENECK

The increasing demand for faster access to more data by both existing and new applications has put a great deal of pressure on storage, with I/O bottlenecks and bandwidth and capacity issues throttling business-critical applications. These applications include:

- **High-performance database/data warehouse:** These workloads run on legacy databases such as Oracle, as well as massively parallel processing (MPP) databases. Examples include traditional transaction processing at a very high volume as well as more analytics-focused workloads such as fraud detection, risk analysis, and supply chain management.



KEY FEATURES

Density and performance

- 144 TB (raw), 100 TB (usable) flash
- 5U form factor
- Unprecedented IOPS, bandwidth, and latency

High-performance architecture

- Large, high-performance PCIe mesh network
- Separate control and data path
- NVMe technology for shared flash via PCIe
- Redundant connection to up to 48 clients

Enterprise availability

- Multi-dimensional Cubic RAID provides high reliability with low overhead
- Industry-leading durability with flash physics control algorithms
- Dual H/A controllers and redundant components

Native application access

- Block device
- Native direct memory API access
- Application plug-ins

- **High-performance applications on Hadoop:** These workloads require faster processing and data ingestion than traditional Hadoop batch processing. Examples include real-time or near real-time analytics on streaming data for functions such as stock trading, ad serving, or business reporting, as well as predictive modeling based on massive data sets.
- **High-performance custom applications:** These workloads need high-speed storage and run natively via APIs, on traditional file systems like XFS, or on parallel clustered file systems. Examples include analytic applications, such as SAS, as well as many custom applications developed for a wide variety of use cases, including trade simulation, weather forecasting, vehicle design, cancer research, targeted pharmacology, climate analysis, and many more.

EXISTING ARCHITECTURES DON'T MEASURE UP

Current storage architectures have significant constraints and have not been able to satisfy the demands placed by businesses on their mission-critical applications to provide faster, more insightful results derived from larger and more varied data sets.

Flash storage arrays are faster than disk, but are limited by I/O stack bottlenecks, network and fabric latencies, and software and hardware design limitations in moving data.

Direct-attached PCIe flash cards or SSDs provide fast performance at each node but are more complex to manage and lose all of the advantages of shared enterprise storage. Capacity is limited on each card or drive and is tied to individual servers, resulting in stranded IOPS and the need to shuffle and create multiple copies of data to improve performance and reliability. And, with storage and capacity scaling linearly, these architectures often require an over-investment in storage to get the desired computing power.

EMC DSSD D5 RACK-SCALE FLASH APPLIANCE TO THE RESCUE

Rack-scale flash is a new storage platform that overcomes the barriers of existing architectures by bringing shared flash closer to compute. EMC® DSSD™ D5™, EMC's first rack-scale flash appliance, uses the lowest latency transport and combines it with hardware and software innovations that optimize the data path between volatile application memory and persistent storage. DSSD D5 provides the next-generation storage performance demanded by today's data and performance-intensive applications and enables innovative, next-generation applications to solve tomorrow's challenges.

DSSD D5 delivers ultra-dense, high-performance, highly available, and very low latency shared flash storage to up to 48 clients. D5 is connected to each node through PCIe Gen3 and leverages NVMe™ technology, which delivers the performance of PCI-attached flash. At the same time, D5 is a standalone appliance that is disaggregated from compute, delivering the benefits of shared storage.

The result is next-generation performance with **average latency as low as 100µs, throughput as high as 100 GB/s, and IOPS of up to 10 million.**

DSSD D5 not only meets the demands of the high-performance, data-intensive workloads of today, but also delivers a crucial link to emerging next-generation applications that are processing extremely large, fast-growing working sets with 100 percent hot and active data.

DENSE AND SHARED FLASH

DSSD D5 reduces your data center footprint and costs with ultra-dense, shared flash. The appliance provides up to 36 flash modules, with 144 TB RAW (100 TB usable) capacity, in a five rack-unit chassis that can be accessed redundantly by up to 48 direct-attached clients.

HIGH-PERFORMANCE ARCHITECTURE

DSSD D5 provides the storage performance demanded by today's and tomorrow's applications by leveraging NVMe technology to connect shared flash storage directly with clients through PCIe Gen3, eliminating many layers in the traditional I/O stack.

DSSD D5 also separates the data and control paths and utilizes the world's largest and highest performance PCIe network, which provides each compute node direct access to each flash module, providing faster performance.

ENTERPRISE AVAILABILITY

DSSD D5 enterprise-class availability and serviceability features provide increased application uptime. These features include dual ported client cards, dual H/A controllers, redundant components, and industry-leading flash reliability and resiliency with Cubic RAID™, dynamic wear leveling, flash physics control, and space-time garbage collection.

NATIVE APPLICATION ACCESS

DSSD D5 accelerates existing and future applications with multiple access options to its flash media. Legacy applications, such as those running on Oracle, can communicate with DSSD D5 using its high-performance block driver without any modification. You can also modify applications to communicate directly with DSSD D5 using its Flood Direct Memory APIs for maximum performance. And, DSSD has also created an API interface plug-in for applications running on Hadoop Distributed File System.

UNLOCK THE POWER OF FLASH STORAGE FOR YOUR MISSION-CRITICAL WORKLOADS

DSSD D5 enables you to tap into all of the performance and capacity that flash storage can provide, maximize the capabilities of your existing data and performance-intensive applications, and easily transition to next-generation applications that can take better advantage of the ever-increasing growth in data and computing power.

With DSSD D5, you can:

- Unlock competitive agility
 - Enable real-time analytics for accelerated decision-making with no production downtime
 - Analyze larger data sets to deliver more comprehensive results and improved decision support
 - Analyze more complex problems and run multiple complex analyses and workloads simultaneously
 - Increase the performance of traditional and next-generation databases

- Accelerate business innovation
 - Build new real-time applications and services
 - Enable applications to be real-time versus batch to improve business results

- Reduce TCO
 - Reduce costs and data center footprint with industry leading IOPS and throughput in a shared direct flash platform
 - Deliver up to 10 times IOPS/\$ and reduce TCO by up to 30 percent compared to all-flash arrays
 - Deliver game-changing performance and operational efficiency

CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit www.emc.com, or explore and compare products in the [EMC Store](#).

EMC², EMC, the EMC logo, Cubic RAID, DSSD, and D5 are registered trademarks or trademarks of EMC Corporation in the United States and other countries. The NVMe word mark is a trademark of NVM Express, Inc. © Copyright 2016 EMC Corporation. All rights reserved. Published in the USA. 02/16 Data Sheet H14828

EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Color of flash modules on front of DSSD D5 may vary from depicted image.

The EMC logo is located in the bottom right corner of the page. It consists of the letters "EMC" in a white, bold, sans-serif font, with a small superscript "2" to the right of the "C". The logo is set against a solid blue rectangular background.