Radian MEMORY SYSTEMS

RMS-325

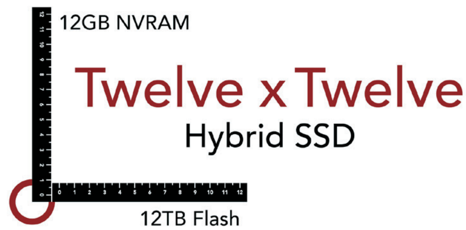

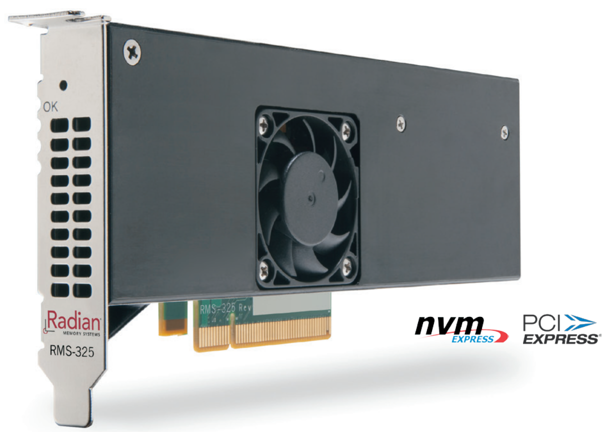

Symphonic CFM or Open-Channel mode

- Flash Capacity: 4TB or 8TB MLC NAND
  3TB, 6TB or 12TB TLC NAND
- User NVRAM: 4GB or 12GB
- NVMe PCIe x8 Gen3 interface
- Half-height, half-length form factor
- On-board Capacitors – no cabling to remote power packs required
- Symphonic™ CFM or Open-Channel* Flash Management mode
- Fault Tolerant Flush-to-Flash™ Backup System
- DuraLife™ Ultra-Capacitor Power Management System
- DiaLog™ OEM Diagnostic Lifecycle Monitoring

The RMS-325 provides host systems with direct control over both its byte addressable User NVRAM, high capacity Flash storage, and transferring data between the two medias. The example below illustrates how the RMS-325 enables a host to write data into its on-board NVRAM, then issue a single Delegated Move command to transfer the data from NVRAM to local Flash storage.

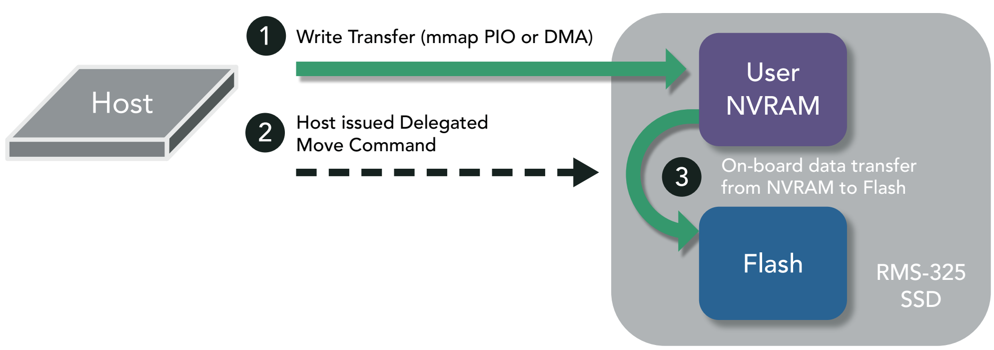

Host Controlled NVRAM combined with High Capacity Flash

The User NVRAM and Flash appear as one PCIe device, but as two distinct block devices. The architecture enables the host to control:

- What data is stored in NVRAM
- What, when and where data should be transferred between NVRAM and Flash
- Issuing a single Delegated Move command that transfers data between NVRAM and Flash
- Avoids overhead of copying data over the system bus, through the host software stack, and into system memory

Radian
MEMORY SYSTEMS

Flash Management: Symphonic CFM or Open-Channel* mode

Today's SSDs are encumbered by a legacy abstraction known as the Flash Translation Layer (FTL) that incurs expensive overhead, unpredictable latency spikes, and suboptimal performance that prematurely wears out the Flash media.

Alternatively, the RMS-325 provides options for Symphonic Cooperative Flash Management or the Open-Channel mode of Flash management that overcome the limitations of FTLs.

Symphonic Cooperative Flash Management (CFM)

The patented, award winning Symphonic CFM technology provides the industry's most advanced Flash management functionality available.

- Magnitude improvement in QoS (latency spikes)
- >80% increase in IOPS & User Bandwidth
- Raw Flash reduced 15% or more
- No device level Write Amplification
- Metrics scale linearly with additional SSDs

Replacing the SSD FTL, Symphonic includes a combination of host-side software libraries and SSD firmware that enables system software to cooperatively perform Flash management processes to realize the full potential of Flash storage.

Symphonic™
Cooperative Flash Management

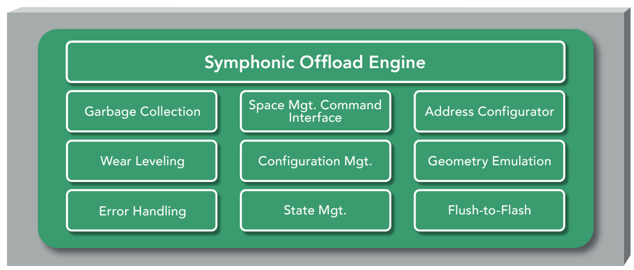

Symphonic is a combination of SSD firmware and software that enables systems to cooperatively perform Flash management

The Symphonic functionality includes configurable address mapping, garbage collection, wear leveling, and reliability features that turn the SSD into an offload engine while operating in host address space. The result is a redistribution of host/device responsibilities that removes an inefficient abstraction layer to dramatically improve Quality-of-Service, performance, cost, and endurance while providing the functionality of a data center class product.

Open-Channel*

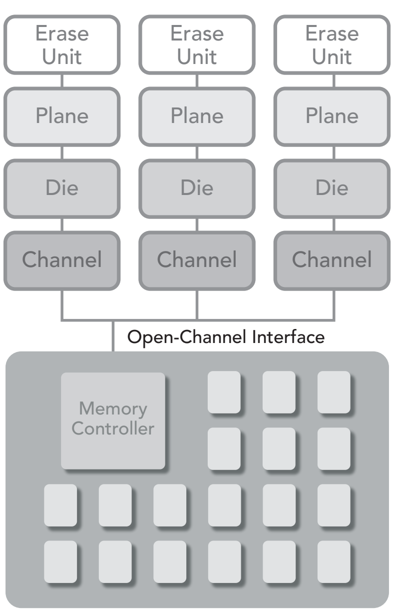

The Open-Channel interface exports the SSD's physical geometry directly to the host system

Open-Channel is a standard hardware interface for SSDs that enables either application-specific host FTLs or tight integration of Flash Management functions into host system software. Based on a paradigm of moving Flash Management functionality from the SSD to the host, an Open-Channel device exports mappings of the actual physical geometry of the SSD to the host system. This provides developers with a platform to design custom tailored host FTLs or integrate system software stacks to achieve levels of parallelism (bandwidth) and deterministic latency not possible with conventional FTL SSDs.

**Open-Channel Solid State Drives NVMe Specification Revision 1.2, April 2016*

The NVRAM is visible as a standard block device that supports block DMA or is byte addressable via mmap PIO. With exceptional, consistent performance for small random writes and unlimited write endurance, host controlled NVRAM can minimize latency spikes and updates to Flash, improving wear out.

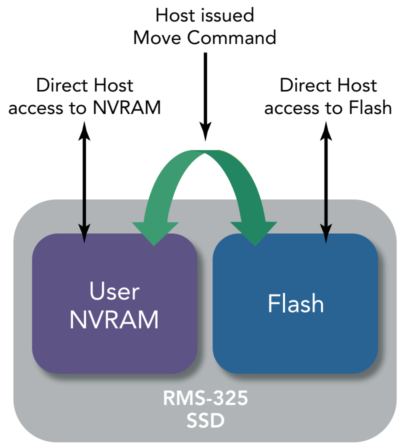

NVRAM/Flash Hybrid Applications:

- Coalescing and buffering writes in NVRAM, then deterministically scheduling data transfers to Flash
- Write-ahead logging and caching to NVRAM for low latency response times
- Persistent NVRAM store for metadata involving high frequency or random updates
- Ideal for efficient check pointing of NVRAM data to larger capacity Flash storage
- Device Read cache can be populated by host, with deterministic, low latency response times

Architectural Advantages

- Optimal for tiering different classes of memory storage media
- DMA engine improves system performance and offloads host resources
- Scale NVRAM and Flash tiers proportionately as SSDs are added to the system
- Delegated Move capability reduces traffic across the system bus and host overhead
- Device appears as a block device that supports DMA or byte addressable mmap PIO
- Simplifies low latency support of PCIe PeerDirect, RDMA, or NVMe-over-Fabric implementations

---

Cost Effective System-Level Memory Tiering

While many FTL based SSDs include NVRAM in addition to Flash, their NVRAM can only be utilized for internal functions and is not visible to host systems. But host systems possess more intelligence about data, including how they want different data prioritized and consequently what media tier different segments should be stored upon. By providing hosts with the ability to select the media and control over transfers between the media, the RMS-325 enables a more deterministic approach to system-level memory tiering. It also enables a more balanced, efficient utilization of capacity for each tier that translates into overall cost reductions.

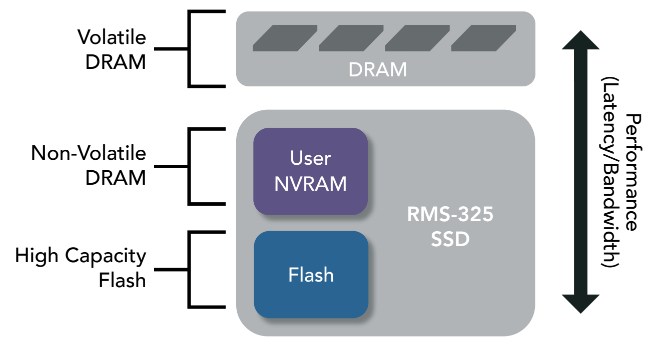

NVMe DMA→High Performance NVRAM

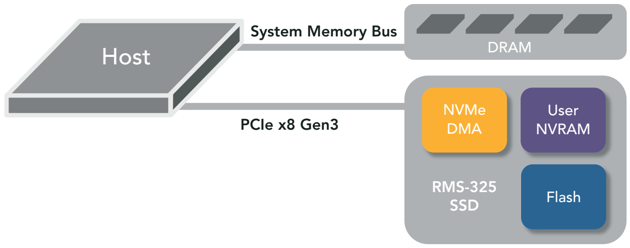

• DMA engine with NVMe command set often outperforms a RAM disk on the memory bus without a DMA engine

• Doesn't require introducing new persistent storage concepts to OS or software applications

• No requirement to cable to a remote power pack

Simple, Reliable, Persistent

Operating systems and host software applications often encounter significant challenges when trying to establish how to handle persistent storage if it is located on the system memory bus, as is the case with NV-DIMMs. This can involve introducing new semantics for how to distinguish volatile memory from persistent memory in terms of re-initialization to a clean state, or protecting against faults involving stray pointers or a kernel panic.

The answer for operating systems and software applications addressing these challenges is often to create a RAM disk, with the NV-DIMM accessed as a block device. However, this introduces the overhead of the kernel block layer and an operating system data path optimized for storage devices that include DMA engines.

Because NV-DIMMs cannot include DMA engines, the additional software overhead incurred from solving the persistent storage challenge can have a tangible impact on system performance.

Alternatively, User NVRAM on a PCIe device can be accessed via a high performance DMA engine, in addition to memory mapped Programmed I/O. The DMA engine on the RMS-325 supports the NVMe command set, involving an optimized data transfer queuing system that provides high performance while consuming minimal host CPU resources. The high performance is complimented by the simplicity of an interface that operating systems understand to be persistent, avoiding the introduction of complex new concepts and semantics.

Scale NVRAM proportionately with Flash

Flash densities are increasing at an incredible rate, and many applications require NVRAM capacity in proportion to Flash capacity. The RMS-325 provides high densities of NVRAM, and an elegant way to scale NVRAM incrementally in proportion to additional Flash storage.

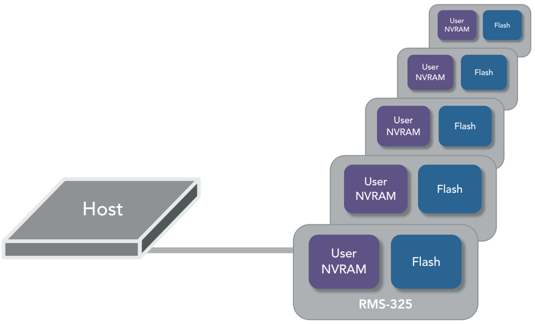

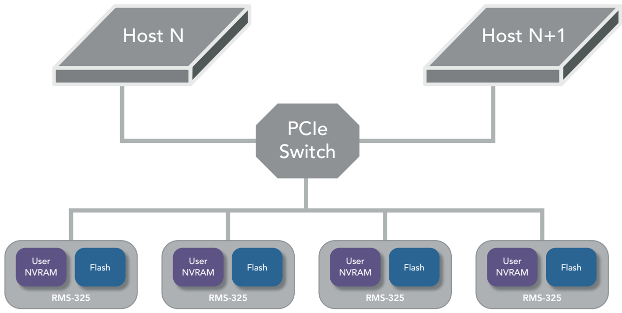

RDMA, NVMe-over-Fabric, PCIe Peer-Direct

Creating fabrics over PCIe and utilizing existing network software drivers can each be simplified and made more performant by having byte addressable memory accessible on the PCIe endpoints. The RMS-325 takes this a step further by locating the byte addressable NVRAM on the same device as the high capacity Flash storage, providing a NVMe DMA engine and the Delegated Move function.

DiaLog™

The RMS-325 includes DiaLog (Diagnostic Logging), a host accessible, embedded diagnostic facility that includes various monitoring functions related to predictive/preventive maintenance, reliability, and continuous process and product improvement.

- Measure & Detect
- Diagnose & Predict
- Record & Notify

On-board health monitoring of components and events are tracked utilizing an I^2C network and communicated to the host. Status is communicated to the host on an on-going basis and when conditions exceed predefined ranges. These capabilities target both Radian's own internal use for Design Verification and Production Test processes, and use by OEM customers throughout their Qualification, Production Test, Field Deployment, Repair, and End-of-Life phases.

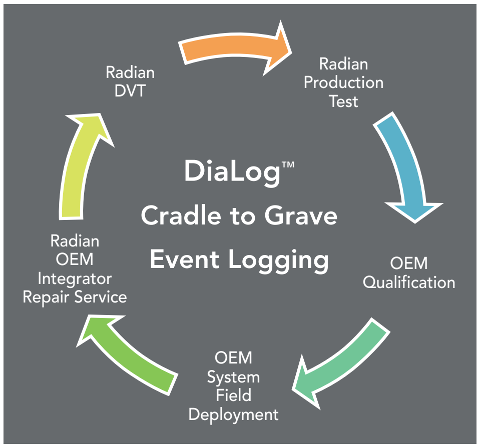

DiaLog provides continuous heart beat monitoring and life cycle management

DuraLife™

The RMS-325 auxiliary power is provided by on-board ultracapacitors, overprovisioned to further ensure data protection in the event of a power failure and to minimize replacement maintenance. Radian's DuraLife power management system addresses ultracapacitor degradation by combining several techniques with mechanisms to significantly extend the useable life of ultracapacitors.

This includes techniques such as dynamic voltage margining which works in concert with the RMS-325 DiaLog™ monitoring system that provides information on applicable ultracapacitor variables such as temperature, and the number, frequency and duration of charge cycles. DuraLife typically more than doubles the raw ultracapacitor life expectancy

Flush-to-Flash

Upon system power failure, the RMS-325 switches to an auxiliary power mode provided by on-board ultracapacitors and data stored in volatile DRAM is transferred to persistent NAND memory by the Flush-to-Flash firmware. The overall Flush-to-Flash system and underlying NAND array are based on a fault tolerant architecture, utilizing transactional semantics and ACID design principles.

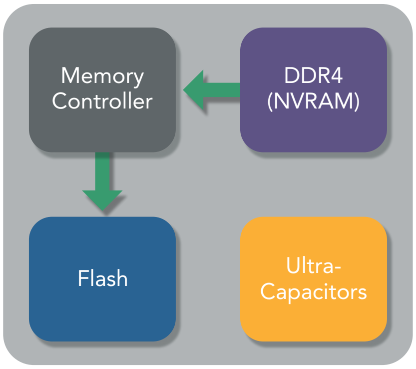

In addition, resources such as ultracapacitor power and NAND capacity are overprovisioned to address events such as repeated system power blackouts and brown-outs, involving shorn writes and data integrity errors during the flush process. Extensive monitoring and component checks are performed on an on-going basis during normal operations to discover predictive anomalies in advance of failures.

Specifications

PCIe Low Profile/Short Length Edge Card form factor	Length: 6.6″ · Height: 2.7″ · Width: 0.8″
PCIe x8 Gen3 Host Interface	Compliant with PCI-Sig PCIe 3.0 Base Specification
NVMe Command Set	NVM Express specification 1.0
Flash Capacity Configurations	MLC 3D NAND 4TB or 8TB TLC 3D NAND 3TB, 6TB or 12TB
NVRAM Capacity Configuration	4GB or 12GB User NVRAM based upon DDR4 @ 2,400 MHz with backup power provided by on-board ultracapacitors with DuraLife™ and fault tolerant Flush-to-Flash™ system
DRAM ECC	64-bit data/8-bit ECC code detects double bit errors and corrects single bit errors
NAND ECC	Hardware LDPC engine
NVMe DMA Engines	Supports NVMe command set, submission/completion queues, and MSI-X vector interrupts
Programmed I/O (PIO)	Support for direct host access to NVRAM by mapping memory (mmap) into host PCI address space with configurable window size
Maximum Payload Size	Configurable to 128B or 256B single packet size
Atomicity	NVRAM: Supported on a per packet basis up to 256b packets Flash: Supported on a per logical block basis, 4KiB minimum
BIST and Health Monitoring	DiaLog™ provides OEM hosts the ability to monitor environmental status, component health, create event notifications and log statistics for continuous product life cycle management
Field Upgradeable Firmware Updates	Mechanism for upgrading firmware in the field via host control (no drive removal necessary)
LEDs	Four LEDs for progress/error codes and initialization with two under host ioctl control during normal operations
Power Requirements (+12V rail)	Typical Maximum: Up to 25W @ 55°C
Ultracapacitor Recharge Time	32 seconds
Ultracapacitor Replacement	DuraLife Power Management System intended to prolong ultracapacitor life expectancy to over ten years
Operating Temperature **Storage Temperature**	0° to 55°C @ 100 LFM 40°C to 85°C
Weight	.465 lbs.
Shock	Operating: 5 G Non-Operating: 10 G
Vibration	Operating: 0.5 G Non-Operating: 1.0 G
ESD	1,500 volts, human body model
MTBF	1M hours
Device Drivers	NVMe Linux 3.1 and above

**Open-Channel Solid State Drives NVMe Specification Revision 1.2, April 2016*

Patent Information: www.radianmemory.com/patents

Radian Memory Systems, Inc.
5010 North Parkway Calabasas, Unit 205, Calabasas, CA 91302
Tel 818 222 4080 Fax 818 222 4081
sales@radianmemory.com www.radianmemory.com