

RMS-375

- 16GB NV-RAM
- Lowest latency NVMe device
- Internal ultracapacitor system fully contained in single drive
- PIO mmap and DMA access
- Dual Port 2x2 or Single Port x4 mode
- NVMe PCIe x4 Gen3 interface
- 2.5" U.2 SSD Drive form factor
- Hot Swap, Live Insertion, Surprise Remove
- DiaLog™ OEM Diagnostic Lifecycle Monitoring



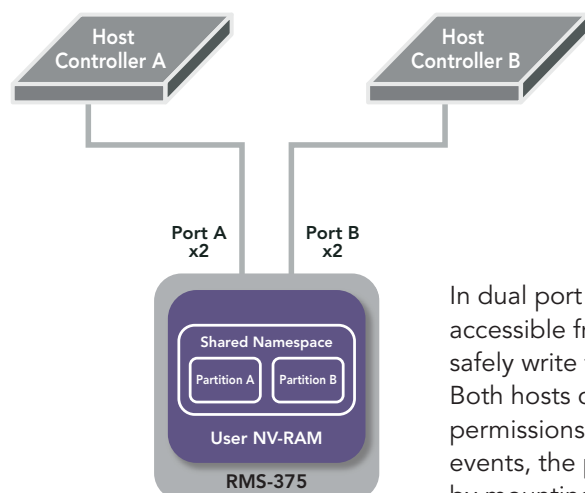
**Dual Port NV-RAM
with unlimited DWPD endurance**

The RMS-375 is a NV-RAM device in a U.2 2.5" form factor with 16GB of the lowest latency memory (DDR4) available on a NVMe PCIe Gen3 x4 device. The memory is fully host controlled and can be accessed via Programmed I/O (mmap) and NVMe DMA transfers. The interface is available in either single port x4 or dual port 2x2 configurations with support for various modes of NVMe Hot-Swap.

Unlimited Endurance (DWPD) + Non-Volatile

NAND Flash and certain storage class memories have a limited number of Drive-Writes-Per-Day (DWPD) that is often inadequate for write intensive applications like caching, logging, buffering and journaling. The RMS-375 utilizes DD4 memory that provides unlimited write endurance and is made non-volatile by an internal power backup system that is fully contained in the same U.2 2.5" form factor SSD. No secondary drive slot is required for auxiliary power or tethering cables to a remote power system.

Dual Port mode for Active/Active Controllers



- Dual Port 2x2 NVMe device
- Shared NVMe namespace accessible from either port
- Namespace can easily be configured into dedicated partitions for each host
- Partner host can access alternate partition on failover
- Avoids complex inter-host coordination

In dual port 2x2 mode, a single shared NVMe NV-RAM namespace is accessible from either port and can be partitioned so that each host can safely write to its own partition without complex inter-host coordination. Both hosts can mount and read from either partition based on sharing permissions. If a host becomes inoperative due to failure or maintenance events, the partner host can access the inoperative host's NV-RAM partition by mounting the partition using standard primitives and permissions.

Hot Swap Functionality

The RMS-375 includes Hot Swap capabilities for High Availability including the following:

- Hot-Plug
- Hot-Swap
- Surprise Hot-Add
- Surprise Hot-Remove

The SSD's fault tolerance capabilities ensure data integrity and protect against shorn writes through these Hot Swap events and unplanned power interruptions. However, each of these features is platform dependent, requiring complementary support and interoperability from the CPU/motherboard, bios, OS, and PCIe fabric.

NV-RAM Applications

With exceptional, consistent performance for small random writes and unlimited write endurance, host controlled NV-RAM can be the ideal solution for applications such as write-ahead logging, caching, journaling, intent logs or any application requiring low latency, persistent storage. The RMS-375 NV-RAM is visible as a standard block device that supports NVMe DMA transfers or can be memory mapped to support Programmed I/O (PIO) operations based on 4-byte

(dword) addressing that are immediately durable.

- Write-ahead logging and caching for low latency, deterministic response times
- Persistent non-volatile metadata store to support random writes or high frequency overwrites
- Coalescing and buffering writes in NV-RAM, then deterministically scheduling data transfers to large capacity Flash or HDD

Endurance Requirements

These data center logging and caching NV-RAM applications usually involve near continuous write workloads. The following table extrapolates typical logging and caching write data rates into DWPD:

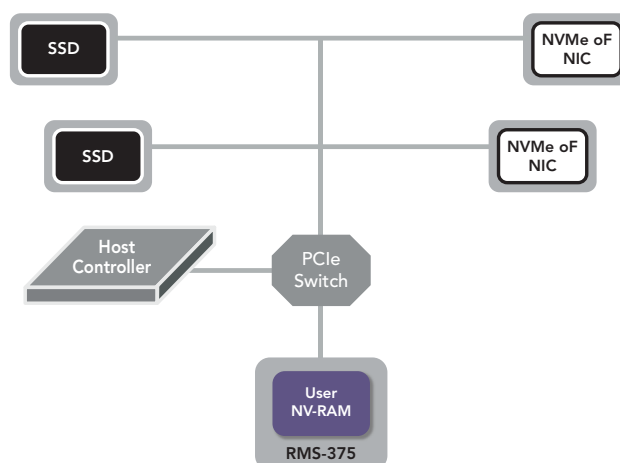
Example Write Data Rates	GB/DAY	DWPD with 8GB device
50,000 IOPS (4K)	17,694	2,212
200,000 IOPS (4K)	70,778	8,847
1 GB/s	86,400	10,800
2 GB/s	172,800	21,600

Data center class Flash SSDs generally provide anywhere from 1 to 10 DWPD. Certain new storage class memories are rated at up to 30 DWPD. While significantly better than Flash SSDs, 30 DWPD is dramatically less than the 2,000 to 20,000 DWPD required to fulfill data center write logging and caching applications.

NVMe-oF Centralized Buffer

As part of an NVMe-oF system, the RMS-375 can act as a centrally accessible repository on the PCIe fabric. DMA based transfers or byte addressable memory windows can be partitioned to buffer and interleave data between different I/O, compute and storage endpoints.

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 NVMe storage endpoints.



Advantages over NV-DIMMs

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.

Performance

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 NV-RAM 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-375 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 recognize as persistent, avoiding the introduction of complex new concepts and semantics.

No Remote Capacitor Packs or Cabling

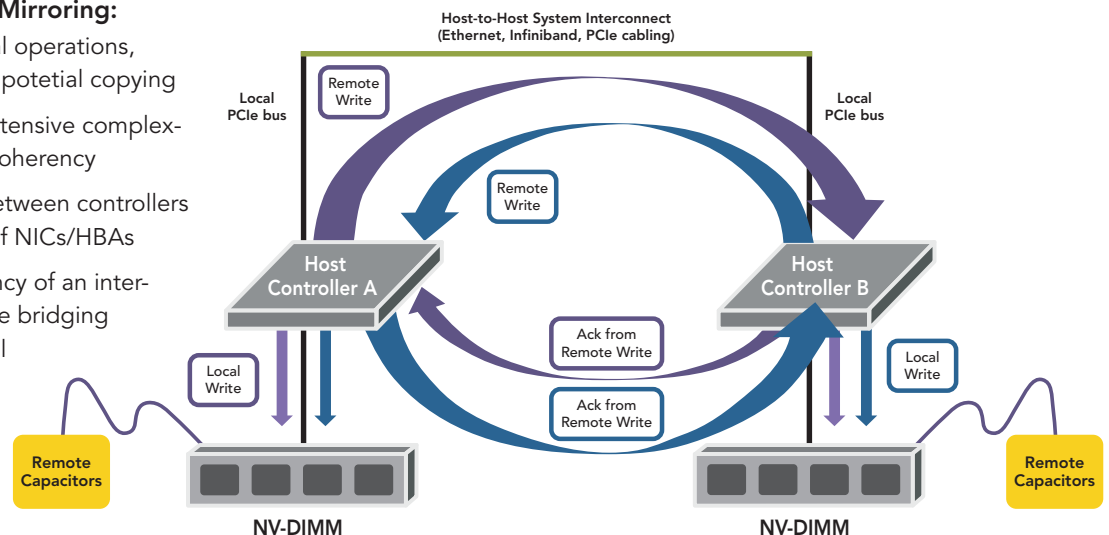
Finding suitable repositories for external capacitor packs is becoming more challenging as platform densities are increasing with each new chassis generation. Unlike a NV-DIMM, no remote capacitor packs or associated cabling are required with the RMS-375's NV-RAM.

NV-DIMMs in Active/Active Configurations

The majority of data center storage arrays are deployed in a 'dual head', active/active configuration for high availability. Utilizing NV-DIMMs in these system configurations requires creating a synchronous mirror which usually involves considerable complexity, overhead, and inefficiencies.

NV-DIMM Synchronous Mirroring:

- Overhead from additional operations, acknowledgements, and potential copying
- Additional, potentially extensive complexities in system software coherency
- Often requires cabling between controllers and the additional cost of NICs/HBAs
- Additional potential latency of an interconnect protocol, and the bridging transition to that protocol
- Difficult to service



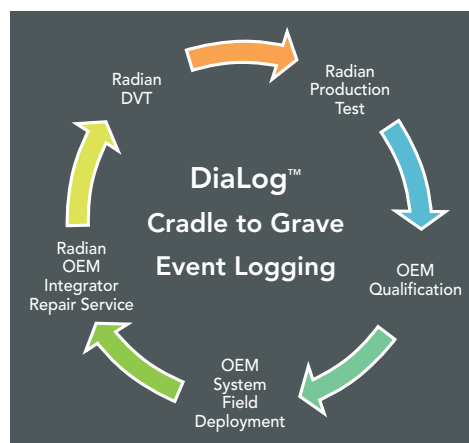
The simplicity of the dual-port RMS-375 overcomes each of these challenges

DiaLog™

The RMS-375 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 I²C 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

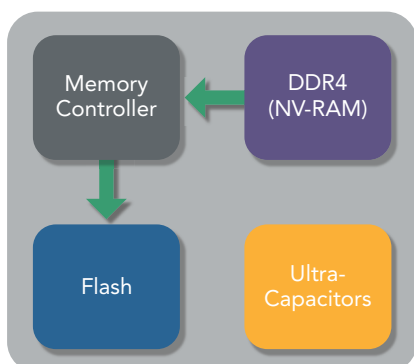
DuraLife™

The RMS-375 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-375 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-375 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, protecting against 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	
U.2 2.5" NVMe form factor	Length: 100mm · Height: 15.0mm · Width: 69.75mm
Dual Port and Single Port Modes	Dual Port 2x2 lane configuration or Single Port x4 lane configuration
PCIe x4 Gen3	Compliant with PCI-Sig PCIe 3.0 Base specification
NVMe Command Set	NVM Express specification 1.0
NV-RAM Capacity Configuration	16GB User NV-RAM based upon DDR4 @ 2,400 MHz with internal backup power provided by DuraLife™ and fault tolerant Flush-to-Flash™ system <i>Future 32GB configuration under consideration</i>
DRAM ECC	64-bit data/8-bit ECC code detects double bit errors and corrects single bit errors
NVMe DMA Engines	Supports NVMe command set, submission/completion queues, and MSI-X vector interrupts
Programmed I/O (PIO)	Support for direct host access 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	Supported on a per packet basis up to 256B packets
BIST and Health Monitoring	DiaLog™ provides OEM hosts the ability to monitor environmental status, component health, 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)
Power Requirements (+12V rail)	Typical Maximum: 13.5W @ 55°C Recharge Cycle: Up to 17W @ 55°C with concurrent r/w operations
Ultracapacitor Recharge Time	38 seconds
Operating Temperature Storage Temperature	0° to 55°C @ 100 LFM 40°C to 85°C
Weight	.291 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

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

Radian Memory Systems, makes no warranty of any kind with regard to the material in this document, and assumes no responsibility for any errors which may appear in this document. Radian Memory Systems reserves the right to make changes without notice to this, or any of its products, in order to improve reliability, performance, or design. All registered trademarks, logos and names are the property of their respective owners. © 2017 Radian Memory Systems. All rights reserved.