RMS-350

- Flash Capacity: 4TB, 8TB, or 16TB
- PMR style User NV-RAM: 4GB or 12GB
- User NV-RAM access modes:
  NVMe or byte addressable (dword)
- NVMe PCIe x4 Gen3 interface
- 2.5″ U.2 SSD Drive form factor
- Dual Port 2x2 or Single Port x4 mode
- Hot Plug, Live Insertion, Surprise Remove
- Cooperative Flash Management (CFM)
- DiaLog™ OEM Diagnostic Lifecycle Monitoring

Cooperative Flash Management (CFM)
Dual or Single Port U.2 NVMe SSD

Cooperative Flash Management

By taking a system-driven approach to Flash storage, Radian has pioneered Cooperative Flash Management (CFM) technology as an alternative to traditional Flash Translation Layers (FTLs). Radian's Symphonic™ CFM redistributes functionality and responsibilities between the host system and SSD.

Built on the premise of operating in logical host address space, Symphonic leverages the intelligent segment cleaning and macro scheduling that many modern host software stacks are already performing. This enables the host to control and schedule processes such as garbage collection, but offloads process execution and lower level media management to the SSD.

Based upon Symphonic CFM technology, the RMS-350 SSD is available in configurations for either Zoned Namespaces (ZNS), the original Symphonic API, or an Open-Channel 2* compliant interface.

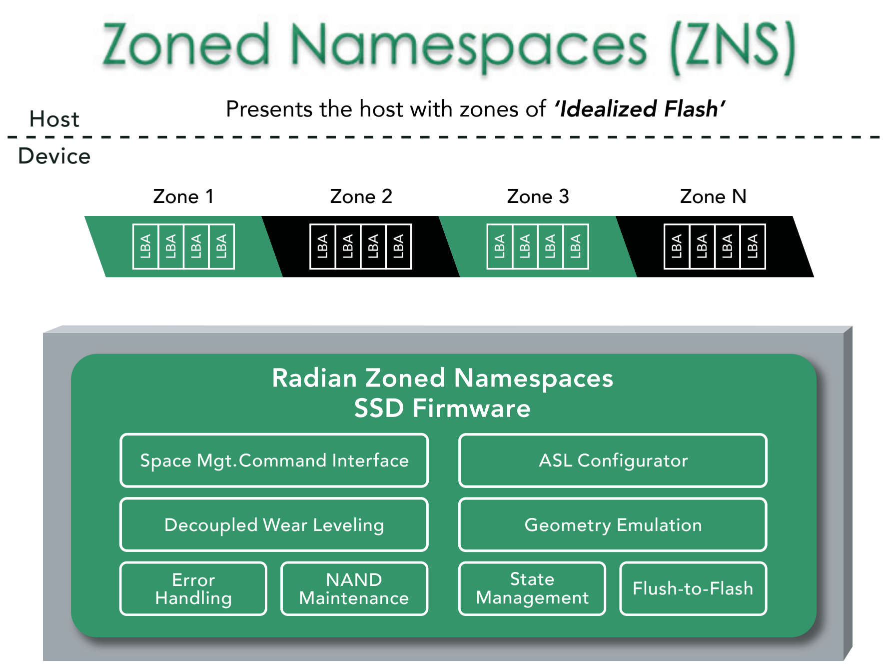

The RMS-350 is a dual or single port, 2.5" U.2 NVMe SSD with up to 16TB of Flash capacity and up to 12GB of PMR style, NVMe byte addressable User NV-RAM that supports Cooperative Flash Management (CFM) with various U.2 modes of Hot Swap functionality. It provides host systems with direct control over both its byte addressable (dword) User NV-RAM, high capacity Flash storage, and transferring data between the two medias.

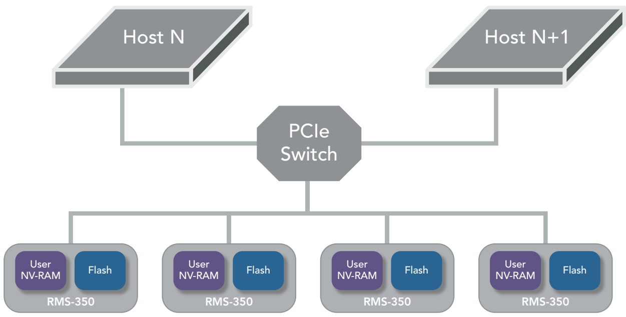

NVMe-over-Fabric (NVMe-oF)

Creating fabrics over NVMe 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. The RMS-350 takes this a step further by locating the byte addressable NV-RAM on the same device as the high capacity Flash storage, providing a NVMe DMA engine, and a specialized Delegated Move function.

Scale NV-RAM proportionately with Flash

Flash densities are increasing at an incredible rate, and many applications require NV-RAM capacity in proportion to Flash capacity. The RMS-350 provides high densities of NV-RAM, and an elegant way to scale NV-RAM incrementally in proportion to additional Flash storage.

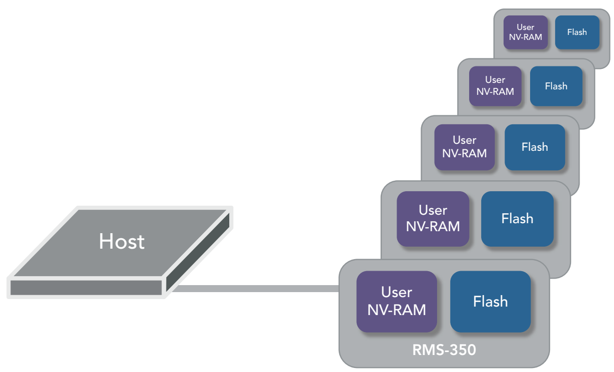

The example below illustrates how the RMS-350 enables a host to write data into its on-board NV-RAM, then issue a single Delegated Move command to transfer the data from NV-RAM to local Flash storage.

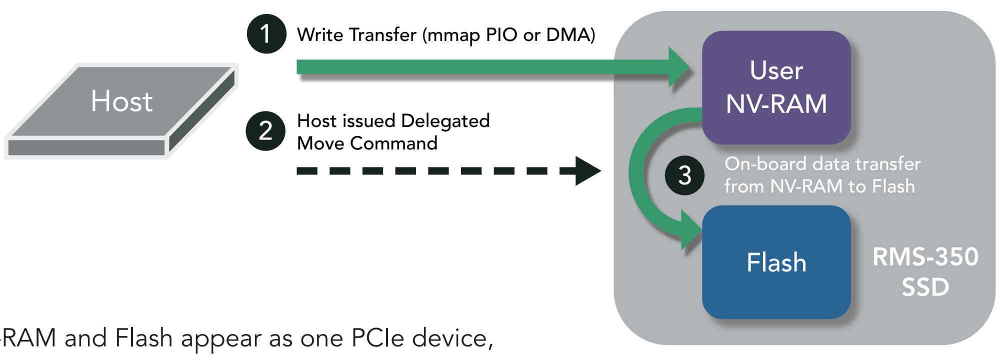

The User NV-RAM and Flash appear as one PCIe device, but as two distinct block devices.
The architecture enables the host to control:

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

The NV-RAM is visible as a standard block device that supports block DMA or dword 4byte addressable via mmap PIO. With exceptional, consistent performance for small random writes and unlimited write endurance, host controlled NV-RAM can minimize latency spikes and updates to Flash, improving wear out.

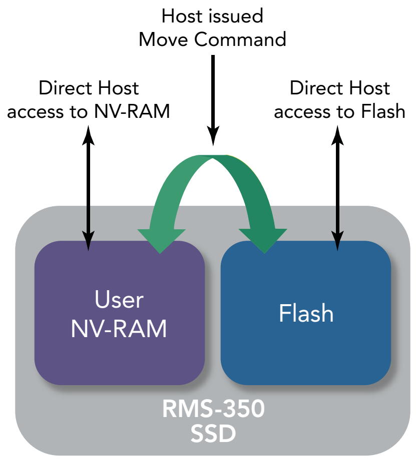

NV-RAM/Flash Hybrid Applications:

- Coalescing and buffering writes in NV-RAM, then deterministically scheduling data transfers to Flash
- Write-ahead logging and caching to NV-RAM for low latency response times
- Persistent NV-RAM store for metadata involving high frequency or random updates
- Ideal for efficient check pointing of NV-RAM 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 NV-RAM 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 Peer-to-Peer, RDMA, or NVMe-over-Fabric implementations

---

Cost Effective System-Level Memory Tiering

While many FTL based SSDs include NV-RAM in addition to Flash, their NV-RAM 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-350 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.

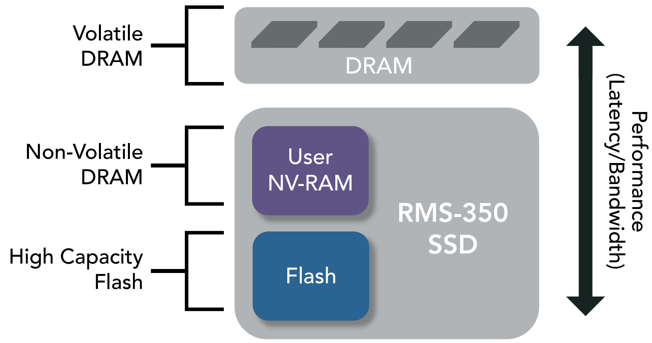

Radian
MEMORY SYSTEMS

NVMe DMA→High Performance NV-RAM

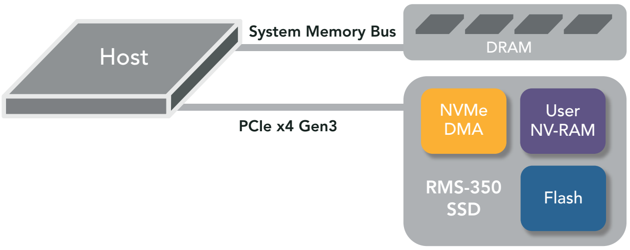

- DMA engine with NVMe command set often outperforms a RAM disk on the memory bus without a DMA engine
- Doesn't require introducing new persistent storage concepts to OS or software applications
- No requirement to cable to a remote power pack

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 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-350 supports the NVMe command set, utilizing 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.

Dual Port

The RMS-350 is available in either a single port x4 or dual port 2x2 configuration.

NV-RAM: In dual port mode, a single shared 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.

Flash: The RMS-350's Flash supports a similar dual port model with two shared namespaces, enabling each host to write to its own Flash namespace without complex inter-host coordination. The partner host can assume control of the inoperative host's namespace in the event of a failure or maintenance event.

Hot Swap Functionality

The RMS-350 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.

DiaLog™

The RMS-350 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.

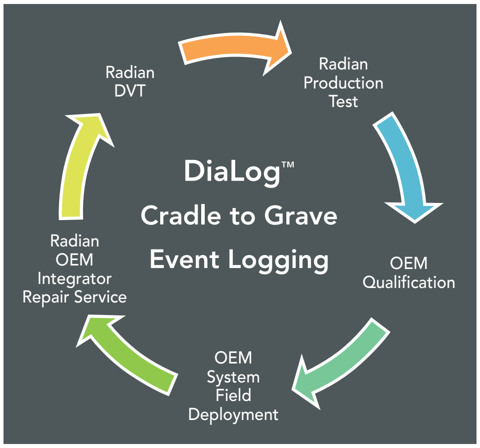

DiaLog provides continuous heart beat monitoring and life cycle management

DuraLife™

The RMS-350 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-350 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-350 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.

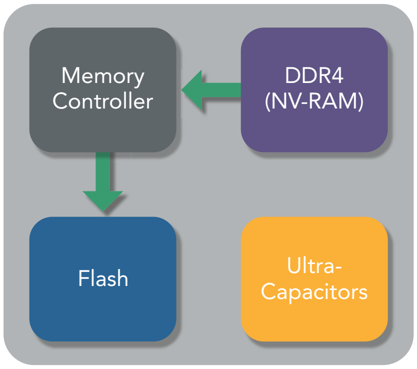

In addition, resources such as ultracapacitor power and NAND capacity are overprovisioned to address events such as repeated system power blackouts and brown-outs, protecting against shorn writes and data integrity errors during the flush process. Extensive monitoring and component checks are performed 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
Flash Capacity Configurations	4TB, 8TB, or 16TB eTLC 3D NAND
NV-RAM Capacity Configuration	4GB or 12GB User NV-RAM based upon DDR4 @ 2,400 MHz with backup power provided by 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 byte addressable (dword) direct host access to NV-RAM by memory mapping (mmap) into host PCI address space via PCIe fixed BAR with configurable window size
Maximum Payload Size	Configurable to 128B or 256B single packet size
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)	16W with 33% Writes/66% Reads at peak bandwidth
Ultracapacitor Recharge Time	32 seconds
Operating Temperature **Storage Temperature**	0° to 55°C @ 100 LFM 40°C to 85°C
Weight	.295 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	Estimated 2.1M hours
Device Drivers	NVMe Linux 3.1 and above

**Open-Channel Specification Revision 2.0, January 29, 2018*

Radian MEMORY SYSTEMS

Radian Memory Systems, Inc.
Tel 818 222 4080 Fax 818 222 4081
sales@radianmemory.com www.radianmemory.com