

# Micron XPERT Enhancements for SSDs

Doug Rollins, Senior Strategic Applications Engineer

Micron Technology, Inc.  
Technical Marketing Brief

## Introduction

**eXtended Performance and Enhanced Reliability Technology (XPERT)** is a suite of Micron-designed storage architecture enhancements that greatly improve SSD performance and reliability. XPERT extends drive life and ensures data integrity.

With XPERT, Micron aligns storage media design, SSD firmware development, and hardware integration to create a comprehensive architecture that enables enterprise-class SSDs to meet the unrelenting demands of 24/7/365 data centers.

The XPERT feature set gives us the flexibility to design SSDs to meet exact application requirements. Only the XPERT features that are appropriate to a particular application are designed in, so our SSDs can precisely match data center usage models. From booting general purpose servers and storing mission-critical databases, to ensuring long-term, trouble-free operation of appliances and managing petabytes of cloud data—XPERT-enabled SSDs meet the stringent usage requirements demanded by the massive growth of digital data.

## XPERT Feature Set – Overview

The XPERT feature set includes:

- **Adaptive Read Management/Optimized Read (ARM/OR):** ARM/OR is a dual-faceted data management technique with both proactive and inline protection. In proactive mode, ARM/OR ensures that data stored on the SSD is immediately available by sampling the stored data and dynamically tuning the NAND device in the background. This tuning, based on Micron-specified thresholds, is completely transparent with no effect on drive or system performance. ARM/OR inline mode provides additional protection by making fast and precise foreground tuning when the host reads data from the XPERT-enabled SSD.
- **Redundant Array of Independent NAND (RAIN):** By calculating and storing additional protection that is parity, Micron's RAIN provides data protection well beyond common error correction code (ECC) implementations. To ensure RAIN does not impact performance, Micron SSD controllers have built-in data acceleration engines that streamline RAIN.
- **DataSAFE:** DataSAFE protects user data as it moves from the host to the SSD interface, through the data path inside the SSD, into the storage media, and back. In addition, DataSAFE embeds the host logical block address (LBA) with the host data before storing both in the drive. Storing this additional information (also known as metadata) helps to ensure that XPERT-enabled SSDs return the exact data requested.
- **Reduce Command Access Latency (ReCAL):** ReCAL enables lower maximum command/access latency by managing internal operations at a more granular level. This ensures smooth, stable performance for enterprise operations.
- **Media Customizations:** XPERT storage media can be optimized through specialized testing and production methods. As a global leader in storage solutions, Micron leverages our engineering expertise in all aspects of design to ensure we choose the best media for our enterprise-class SSDs. Micron customizes NAND product lines and backend test flows as required for read-centric enterprise applications to the most extreme, I/O-intensive enterprise environments. Working in lockstep with our global design and qualification teams, Micron's SSD development and test engineers ensure that our XPERT-enabled SSDs offer state-of-the-art performance and reliability.



## XPRT Feature Set – In Depth

### ARM/OR

The optimal methods of reading data on an SSD are not static. Specific characteristics of the READ command can—and should—be dynamically tuned. Such tuning has a direct impact on data reliability. Proper tuning improves READ command performance in terms of immediate data access and long-term data reliability—which are key requirements of enterprise applications.

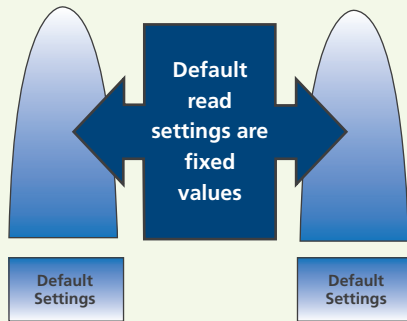


Figure 1a: Default Read Settings

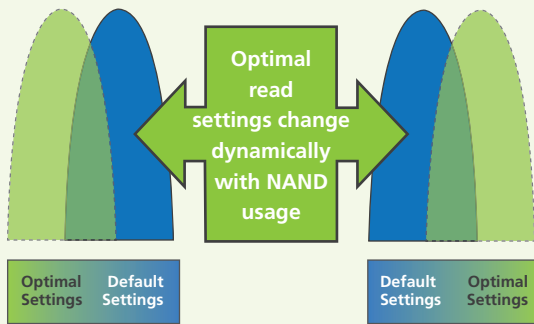


Figure 1b: Optimal Read Settings Change with NAND Use

Figure 1a shows the default settings used to read data from the media on the SSD. These are factory presets and are optimal for new NAND devices. Figure 1b shows how optimal read settings can change as the drive is used. The amount of data written to and read from the drive, for instance, can impact the optimal settings. ARM/OR dynamically tunes these settings to ensure best performance and data integrity for the SSD.

### RAIN

RAIN is a parity protection mechanism that operates in real time. Using well-proven parity techniques, RAIN embeds protected data with the user data. The details of each RAIN implementation are design-specific, but they include the following core elements:

- **Data-to-Parity Ratio:** Expressed as X data + Y parity (or X:Y), this ratio is optimized for intended drive workload, performance, media type, and several other factors. It is also referred to as the stripe size.
- **Parity Storage Location:** The parity may be stored in a fixed, relative, or rotating location.
- **Protection Level:** RAIN can protect user data from catastrophic media failures.
- **Hardware Acceleration:** RAIN can be managed in firmware or accelerated in hardware.

In Figure 2, the ratio is 7:1—seven elements of user data and one element of parity data. However, RAIN is not limited to 7:1; the ratio can be designed specifically to balance data protection, drive design, intended workload, and cost.



Figure 2: RAIN Implementation; 7:1 Ratio



## DataSAFE

DataSAFE provides data path protection by storing additional information (metadata) along with user data to help ensure that XPERT-enabled SSDs return the exact data requested.

Data is data that is written; and the logical block address (LBA) for that data (or the location associated with it) is called metadata, which literally means data about data. If the host data was a building, the LBA would be that building's street address.

The term data path refers to the logical course or path that the data and metadata follow throughout the SSD as data is either written to or read from the underlying media.

When the host writes data to the SSD, two key elements combine: the actual data to be written and the LBA from which it came.

All data written to the SSD has an LBA associated with it. Micron's DataSAFE logically embeds the host LBA with the associated data before writing the data to the NAND device, as shown in Figure 3.

In addition to host LBA checking and embedding, DataSAFE provides other data path protection methods like memory protection ECC (MPECC). Figure 4 shows how MPECC protects the host data in a Micron enterprise SATA SSD by adding ECC coverage to the data as it enters the DRAM

on the SSD. This additional MPECC follows the host data through the SSD and prevents bit fumbling.

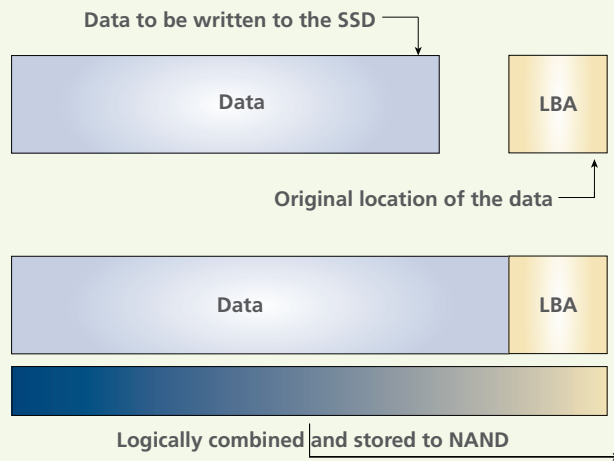


Figure 3: DataSAFE LBA Checking and Embedding

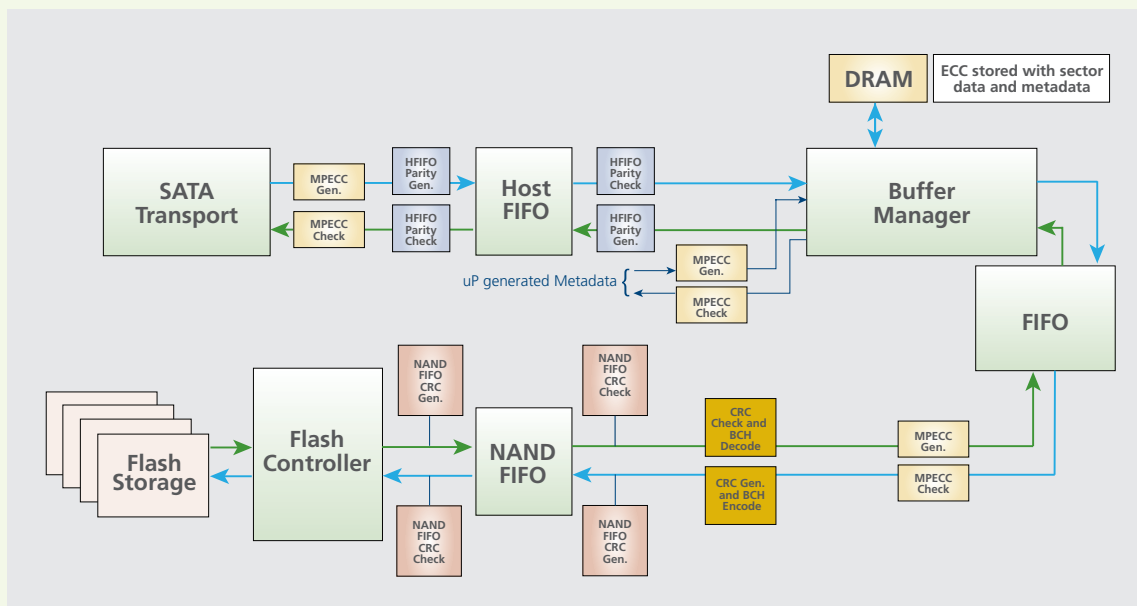


Figure 4: DataSAFE MPECC



## ReCAL

ReCAL uses well-managed background operations to enable faster response times. ReCAL does not interfere with the host I/O, which enables a substantial reduction in maximum WRITE latency. A combination of latency reduction, wear-leveling efficiency, and more efficient internal data structures are shown in Figures 5a and 5b.

XPERT SSDs with ReCAL manage drive wear in much smaller chunks. These small chunks (equal to one RAIN strip)—combined with ReCAL’s ability to interrupt wear leveling to service host I/O requests—reduce the maximum WRITE latency and provide much smoother and consistent performance, as shown in Figure 5a.




Figure 5a: Granular Wear Leveling



Figure 5b: ReCAL L2P – Multiple, Smaller Structures

ReCAL optimizes the logical-to-physical (L2P) internal data structure. As shown in Figure 5b, breaking this data structure into a series of smaller, more-efficient chunks enables the structure to be managed more efficiently and it further reduces maximum latency.

Finally, ReCAL ensures that background housekeeping functions do not interfere with host data storage by managing those housekeeping functions in the smallest unit possible.

## Media Customizations

Like most XPERT features, the level of media customization depends on the SSD design, the intended workload, and the required lifespan. Media optimizations can be simple (e.g., additional factory testing) or complex (e.g., custom designs for specific SSDs).

### Media Customization Example – Customer Media Precycling

When a device contains media that wears, the device will have three distinct failure regions: early-life, long-term use, and end-of-life, as shown in Figure 6a.

Micron’s SSD-optimized NAND is precycled at the factory during the SSD manufacturing process. The precycling step essentially eliminates the early-life failure region, as shown in Figure 6b.

In addition to optimizing the fundamental design of the NAND media, the following can also be tuned to help ensure that the NAND media meets SSD requirements:

- Rated endurance
- Spare area size
- Tight integration to SSD firmware
- Packaging and stacking

Media customizations are not limited to classic storage technologies like NAND. Selecting the correct media is imperative. When selecting XPERT media, consider factors like minimum optimal transfer, level of media maturity, and the type and use of the media (DRAM, NAND, PCM, etc.).

Because different SSDs are designed for different workloads, the design-in process is key to optimizing enterprise SSD designs.



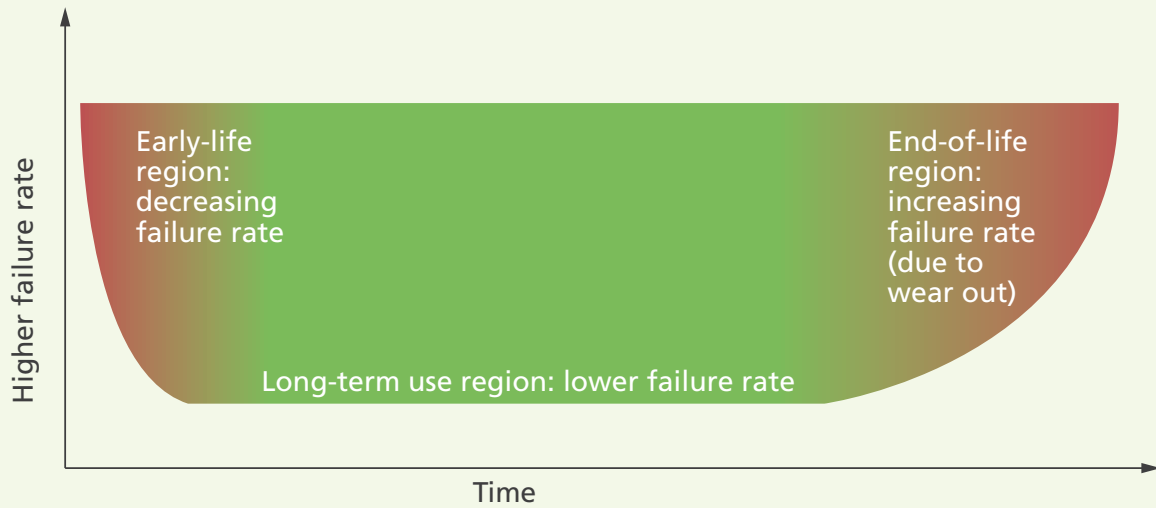


Figure 6a: Device Failure Regions

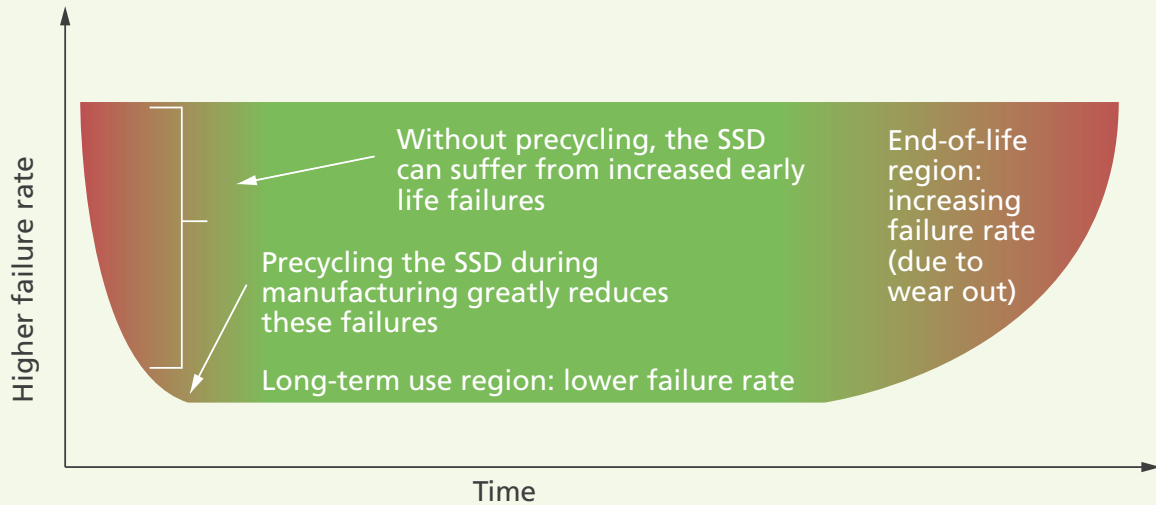


Figure 6b: Device Failure Regions with Precycling

## Conclusion

The XPERT suite of technologies extends performance and enhances the reliability of Micron’s enterprise-class SSDs. The specific XPERT features and storage media

used by an SSD can be customized to provide the performance, security, and flexibility required by various enterprise storage applications.

[micron.com](http://micron.com)

Products are warranted only to meet Micron’s production data sheet specifications. Products and specifications are subject to change without notice.  
©2013 Micron Technology, Inc. Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners. All rights reserved. 02/13

