



White Paper

## Why Data Retention is Becoming More Critical in Automotive Applications

*Understanding and Improving Managed NAND Flash Memory for Higher Data Retention and Extended Product Life*

Christine Lee – Kevin Hsu – Scott Harlin  
KIOXIA America, Inc.

Assisted and self-driving vehicles are fully loaded with electronics that support the infrastructure within. They have become mobile data centers that require an immense amount of computing power to capture, process and analyze data 'near-instantaneously' from a myriad of sensors, recorders, algorithms and external connections. A massive amount of this data is either stored locally or uploaded to cloud storage to be leveraged and transformed into real-time intelligence and value. Data storage has now become a critical part of automotive design, placing a precedence on high data retention and continual data integrity.

The automotive environment creates unique challenges and presents a much different scenario from computing equipment in air-conditioned server rooms under controlled temperatures. Due to the extreme temperatures that can affect the NAND flash storage used within vehicles, there are information technology (IT) considerations that require different design approaches. Understanding how data wears out, how temperature and NAND flash memory characteristics can affect data retention and product life, and what practices can be used to improve them, are the focuses of this paper.

### 'Under the Hood' in Automotive Data Storage

Depending on the source, assisted and self-driving vehicles generate terabytes (TB)<sup>1</sup> of data daily. One prediction<sup>2</sup> forecasts that between 5 TB and 20 TB of data will be consumed per day per vehicle, which is overwhelmingly more data than the average person consumes on a smartphone<sup>3</sup>. Automotive electronic designers use NAND flash memory to address a variety of key capabilities within vehicles (Figure 1).





Figure 1: NAND flash storage delivers a number of capabilities within an automotive environment

These attributes make NAND well-suited to meet a variety of automotive requirements (Figure 2), and with new design capabilities on the horizon, will evolve them even further.



\*ADAS = Advanced Driver Assistance Systems

Figure 2: Sample of NAND flash storage uses within assisted and self-driving vehicles

Of the massive amount of data that is generated within assisted driver and autonomous vehicles, a good portion is processed through artificial intelligence (AI). One critical area is ADAS (Advanced Driver Assistance Systems) where AI use is prevalent. The other area is IVI (In-Vehicle Infotainment) which is the delivery mechanism for some of the captured analytics (Table 1). In both areas, a large amount of data is generated that needs to be stored.

| ADAS                                                                                                                                                                            | IVI                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Captures data from their surroundings and uses a variety of sensors, cameras, radars and LiDARs* to help steer, brake and accelerate through traffic                            | Delivers entertainment and information through a combination of systems, audio and video interfaces, touch screen displays, button panels, voice commands, and others |
| Includes a driver monitoring system that recognizes and authenticates the driver, detects drowsiness and distraction, and provides real-time alerts to reduce on-road accidents | Connects with all smart automotive technologies such as ADAS, V2X connectivity solutions, telematics devices, smartphones, sensors, etc.                              |
| Features in-vehicle sensors that read, compare and physically map data to its environment and responds to obstacles in the driving path                                         | Enhances the driving experience by providing better access to maps and media using the vehicle's wireless connectivity, HMI application, and infotainment devices     |

\*Light Detection and Ranging

Table 1: The three key features of ADAS and IVI

Automotive electronic designers are challenged with delivering data between data centers and infotainment systems and use software code to enable this communication. The software code is essential and must address the ability to:

- Provide smartphone-like experiences (e.g., browse internet, receive/send texts & emails, use apps)
- Obtain road / traffic warnings instantaneously
- Navigate the vehicle
- Update software through over-the-air (OTA) communications
- Update sensor and camera data for deep learning
- Enable vehicle safety through vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications

With performance, reliability and safety at stake in assisted and self-driving vehicles, it has now become crucial that NAND storage is used in a way to maintain a high level of data retention and deliver long-term data integrity.

## Data Retention in NAND Flash Memory

In NAND flash memory, data is written to NAND cells. For each occurrence, one Program/Erase (P/E) cycle is generated and the oxide layer within the cell will degrade ever so slightly, resulting in a finite life expectancy of the cell. As more P/E cycles are generated, the oxide layer weakens and cannot hold the electrical charge required for storing data. Single-level cell (SLC), which stores 1-bit per cell, is less impacted by each P/E cycle versus multi-level cell (MLC) and triple-level cell (TLC), which store more bits per cell (Figure 3).




Figure 3: SLC stores 1-bit per cell and is less impacted by P/E cycles versus MLC and TLC

To maintain data integrity, the level of charge in each NAND cell has to be kept within certain threshold voltages, requiring NAND controller firmware to routinely refresh the NAND cells so it can read just the threshold voltages to its nominal values. However, the electrical charge will naturally degrade with time and the rate of charge loss could be accelerated by extreme temperature variances and the health of the NAND cell.

### Extreme Temperature Variances:

Automotive electronics can experience wide temperature variations depending on their location in a vehicle, and the effects of extreme heat or cold can cause a big challenge to onboard storage and booting of the compute and storage systems. From a data storage perspective, NAND is typically not affected by low surrounding temperatures, such as 40° Celsius (C), however, high temperature can drastically decrease NAND data retention. It is important to note that the represented data retention (Figure 4) requires log scaling (versus linear scaling) to show that it is exponentially reduced<sup>4</sup> (or that it is a logarithmic representation versus a linear one).

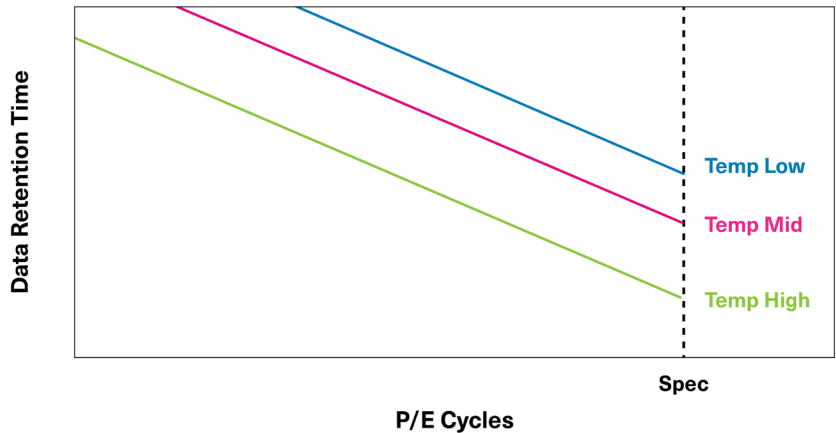


Figure 4: Extreme temperature variances based on data retention time and P/E cycles

**Flash Memory Wear-out**

Before new data can be written to a NAND flash page, the page must be erased first. The erase process requires a relatively large electrical charge to the flash cell. As this erase process repeats, the oxide layer of the NAND chip minutely wears down. Failure is reached when write operations can no longer be completed, and as a block fails. Some of these errors may be corrected, but eventually the error correction code (ECC) routines within the flash controller cannot address them all, resulting in read bit errors or write block failures.

Wear-out refers to the NAND's ability to sustain a finite number of erase cycles before it becomes unreliable and is based on many tests. Wear acceleration in NAND memory can be estimated by calculating its Write Amplification Factor (WAF).

**The Write Amplification Factor**

Write amplification is an undesirable effect associated with NAND-based storage devices, including managed NAND (NAND flash memory and integrated controller in one package) where the actual amount of data written by the controller is greater than the amount written by the host. Ideally, a byte of data written to the NAND storage should equal a byte of data stored in a NAND page. However, in actuality, that data may be written to multiple NAND pages during garbage collection. Garbage collection is a process of re-writing valid data from several NAND blocks to a new NAND block in order to erase the previous blocks so they can be accessible to new data.

The Read-Erase-Modify Write Cycles process can also cause an increase in WAF. In this instance, when one page of a block is edited, the whole block must be erased. The content of the block is stored in cache, where the edited block can be written back to the empty block. Both processes 'amplify' the data written by the host inside of NAND storage and cause an increase in the WAF.

| Write Amplification Factor                                                                      |
|-------------------------------------------------------------------------------------------------|
| $\text{WAF} = \frac{\text{Data written to NAND flash memory}}{\text{Data written by the host}}$ |

The WAF is the ratio of data written to the NAND from the data written by the host. The most ideal WAF is 1, or when the data written to NAND is equal to data written by the host. When the WAF is higher, more P/E cycles are generated that may cause shorter data retention and NAND life.

## Best Practices

Given the challenges of extreme temperature variances, flash memory wear-out and the WAF effect that automotive electronic designers must address in vehicle storage, this section provides some best practices that can be applied to managed NAND, helping to ensure that product life meets the design requirements. This includes, but is not limited to, five key best practices as follows:

### Evaluate the WAF:

WAF is a metric for determining the ratio of data that is written to NAND flash memory versus the amount of data that is written by the host. A good practice to perform is collecting workload traces and analyzing them regularly to evaluate WAFs. Due diligence in this area can lead to system optimizations and slower NAND wear.

### Create Larger Chunk Sizes:

Writing small chunks of data can result in a larger WAF because unused dead space is created within the NAND. Write operations become less efficient and NAND wear speeds up. Therefore, another good practice is for the host to perform less frequent small block write operations to reduce the unused dead space it creates. NAND performs well with sequential write operations in big chunks, so assigning the smallest write size to one NAND page size will help to reduce WAF.

If the write size is less than one page, then there will be empty space in the page that becomes somewhat wasted since it does not store any data. To maintain a low WAF, each page should be filled with as much data as possible, and understanding the relationship between pages and blocks, and the operations that each supports, is very important (Figure 5):




Figure 5: Page / Block relationship

For applications, such as data logging, where data typically arrives in small chunks, it is recommended to implement a RAM buffer that can aggregate all of the separate small data chunks into one large block before writing it to NAND. In a managed NAND environment, enabling the write cache function may achieve the same purpose.

### Use pSLC Mode:

One important practice that can enable greater endurance of MLC or TLC based managed NAND is to partition a portion to act as SLC flash memory, or what is known as pseudo SLC (pSLC) NAND. In combination with a refresh policy, pSLC mode can improve automotive data retention even further as it stores only one bit of data (like SLC) within MLC or TLC based managed NAND, and delivers more P/E cycles of cell endurance<sup>5</sup>. The consequence is a reduction of bits that are available for storage. As it relates to MLC, the reduction is ~50%, and in TLC, it's ~66.6%, when bits are partitioned and stored in a pSLC-based partition<sup>6</sup>.

Although NAND density may be reduced by 50% in MLC and 66.6% in TLC, storing a single bit per cell in both can improve write endurance by up to ten times<sup>7</sup>. To implement pSLC mode requires a sophisticated controller with firmware that is optimized for the specific NAND that is used. KIOXIA (formerly Toshiba Memory) embedded Multi-Media Card<sup>®</sup> (e-MMC) and Universal Flash Storage (UFS) solutions support an enhanced data area based on pSLC mode, delivering cost benefits of using MLC or TLC based managed NAND with higher reliability, especially in extreme temperature environments.

### Develop a Refresh Implementation:

A refresh implementation can determine an optimal refresh rate for monitoring and extending data retention in NAND. To achieve this, all of the data that needs to be refreshed for data retention is estimated by P/E cycles and temperature profiles. The refresh process checks the blocks or pages to be refreshed, as well as blocks of data considered to be 'at-risk,' and moves them to a new area where algorithms can extend data retention by many decades. KIOXIA e-MMC and UFS automotive solutions include built-in refresh functionality that can be easily integrated with the host.

## NAND Best Practices

*For achieving higher data retention and extended product life*

- Evaluate the WAF**  
*Collect workload traces and analyze them regularly*
- Create Larger Chunk Sizes**  
*Perform less frequent small block write operations*
- Use pSLC Mode**  
*Store a single bit per cell to improve write endurance*
- Develop a Refresh Implementation**  
*Determine an optimal refresh rate for monitoring/ extending data retention*
- Monitor the Health Register**  
*Monitor the real-time usage of NAND*



## Monitor the Health Register:

The e-MMC and UFS specifications define the features, capabilities and updates for embedded NAND flash memory storage deployable in vehicles. Within both specifications is a Device Health Report that enables users to monitor the real-time usage of the NAND, which in turn, helps to determine the life expectancy of the specification itself. Both specifications provide a report that includes a remaining lifetime indicator ('percentage of life used') and a pre-EOL warning. Therefore, a good practice is to monitor the e-MMC and/or UFS specification health registers on a regular basis to keep track of NAND life expectancy.

## Summary

Memory and storage requirements are on the rise for various automotive applications, especially relating to assisted and self-driving vehicles that need to capture, process and analyze data in 'near real-time' from a variety of sources. The attributes of NAND flash memory are well-suited for a variety of automotive electronic requirements, and the ability to move to higher densities (2D NAND to 3D NAND) provides data scalability in the same footprint.

To achieve increased data integrity and life expectancy, automotive electronic design engineers should consider several factors that affect the data retention of NAND. These include extreme temperature variances, flash memory wear-out and the WAF effect. As more P/E cycles are generated, NAND reliability worsens and eventually can no longer be used reliably. Automotive designers have several best practices they can use to optimize NAND life expectancy, which include, but are not limited to, evaluating the WAF to monitor NAND wear, creating larger chunk sizes to slow NAND wear, using an enhanced partition based on the pSLC mode, developing a refresh implementation to extend NAND data retention, and monitoring the e-MMC and UFS specification health registers to help maximize NAND life.

Additional KIOXIA automotive product information is available [here](#).

## Notes:

<sup>1</sup> Definition of capacity - KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2<sup>30</sup> bits = 1,073,741,824 bits, 1GB = 2<sup>30</sup> bytes = 1,073,741,824 bytes and 1TB = 2<sup>40</sup> bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

<sup>2</sup> Source: Enterprise AI, "Driverless Cars Generate Massive Amounts of Data. Are We Ready?" Mark Pastor contributor, October 2018, <https://www.enterpriseai.news/2018/10/24/driverless-cars-generate-massive-amounts-of-data-are-we-ready/>.

<sup>3</sup> The average person used 2.9 GB of mobile data per month in 2019 and is expected to use almost 4 GB in 2020. Source: OFCOM UK, "How Much Data Do I Need?" Communications Market Report, July 2019, and published in 3G.co.uk, June 2020, <https://3g.co.uk/guides/how-much-data-do-i-need>.

<sup>4</sup> The equation to calculate data retention as a function of temperature is exponential and available at: [https://en.wikipedia.org/wiki/Arrhenius\\_equation](https://en.wikipedia.org/wiki/Arrhenius_equation).

<sup>5</sup> Cell endurance specifications in NAND flash memory usually assume a minimum data retention duration of 12 months (at the end of drive life).

<sup>6</sup> Source: Hyperstone, "How pseudo-SLC mode can make 3D NAND flash more reliable," published November 2019: [https://www.hyperstone.com/en/How-pseudo-SLC-mode-can-make-3D-NAND-flash-more-reliable-2524.html#:~:text=Pseudo%2DSingle%20Level%20Cell%20\(pSLC,in%20each%20cell%20to%20one](https://www.hyperstone.com/en/How-pseudo-SLC-mode-can-make-3D-NAND-flash-more-reliable-2524.html#:~:text=Pseudo%2DSingle%20Level%20Cell%20(pSLC,in%20each%20cell%20to%20one).

<sup>7</sup> pSLC supports up to 30,000 P/E cycles based on industry specifications, while MLC and TLC support up to 3,000 P/E cycles equaling a 10x improvement.

<sup>8</sup> The embedded Multi-Media Card (e-MMC) v5.1 specification is the current release by the Joint Electron Device Engineering Council (JEDEC®) and published in February 2015.

## TRADEMARKS:

JEDEC is a registered trademark of JEDEC Solid State Technology Association. All other trademarks or registered trademarks are the property of their respective owners.

## DISCLAIMERS:

KIOXIA America, Inc. may make changes to specifications and product descriptions at any time. The information presented in this white paper is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of KIOXIA America, Inc. products as measured by those tests. Any differences in software or hardware configuration may affect actual performance, and KIOXIA America, Inc. does not control the design or implementation of third party benchmarks or websites referenced in this document. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to any changes in product and/or roadmap, component and hardware revision changes, new model and/or product releases, software changes, firmware changes, or the like. KIOXIA America, Inc. assumes no obligation to update or otherwise correct or revise this information.

KIOXIA America, Inc. makes no representations or warranties with respect to the contents herein and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this information.

KIOXIA America, Inc. specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. In no event will KIOXIA America, Inc. be liable to any person for any direct, indirect, special or other consequential damages arising from the use of any information contained herein, even if KIOXIA America, Inc. is expressly advised of the possibility of such damages.

© 2020 KIOXIA America, Inc. All rights reserved.