Using an Analytics Pipeline to Monitor NAND Health Data in the Field

Behzad Amiri
behzad@purestorage.com

August 2019

PURESTORAGE®

Overview

- Why do we monitor NAND health?
- Our solution: Analytics Pipeline
- Sample of results
  - Raw bit error rate vs. P/E cycles
  - Raw bit error rate vs. data retention

Why monitoring NAND health is important to us?

- Avoid data loss and maintain throughput
  - Rate of uncorrectable reads
  - Rate of retry reads
- Data refresh rate
- Detect worn-out systems
  - When to provide new systems to customers

NAND characterization

- For each new NAND parts:
  - **P/E cycle effect** on raw bit error rate (RBER) and throughput
  - **Retention effect** on RBER and throughput
    - Thermally accelerated retention test
  - **Read Disturb** effect

Nand characterization is not enough

- Characterization is done in a **small scale** => rare NAND issue are not captured
- Some of the NAND features should be monitored **over time**
  Example:
  - Bad blocks, bad planes, bad LUNs
- Characterization tests does not capture the effect of **whole SW stack**:
  - P/E cycles distribution, hot blocks, GC, etc.

Solution: Monitor the health of the nand by looking at the **logs** generated by our systems in the field

Overview of Analytics Pipeline

Analytics Pipeline

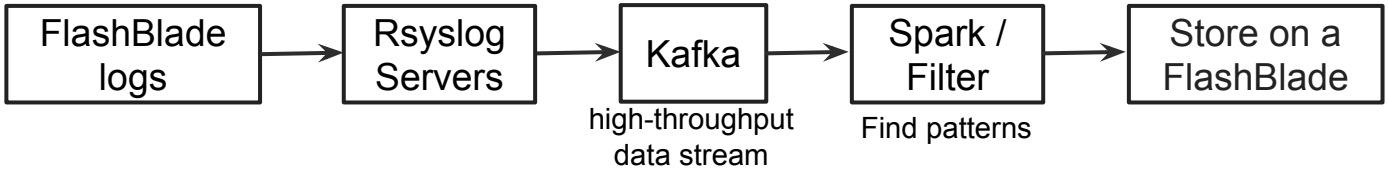

Overview of Analytics Pipeline

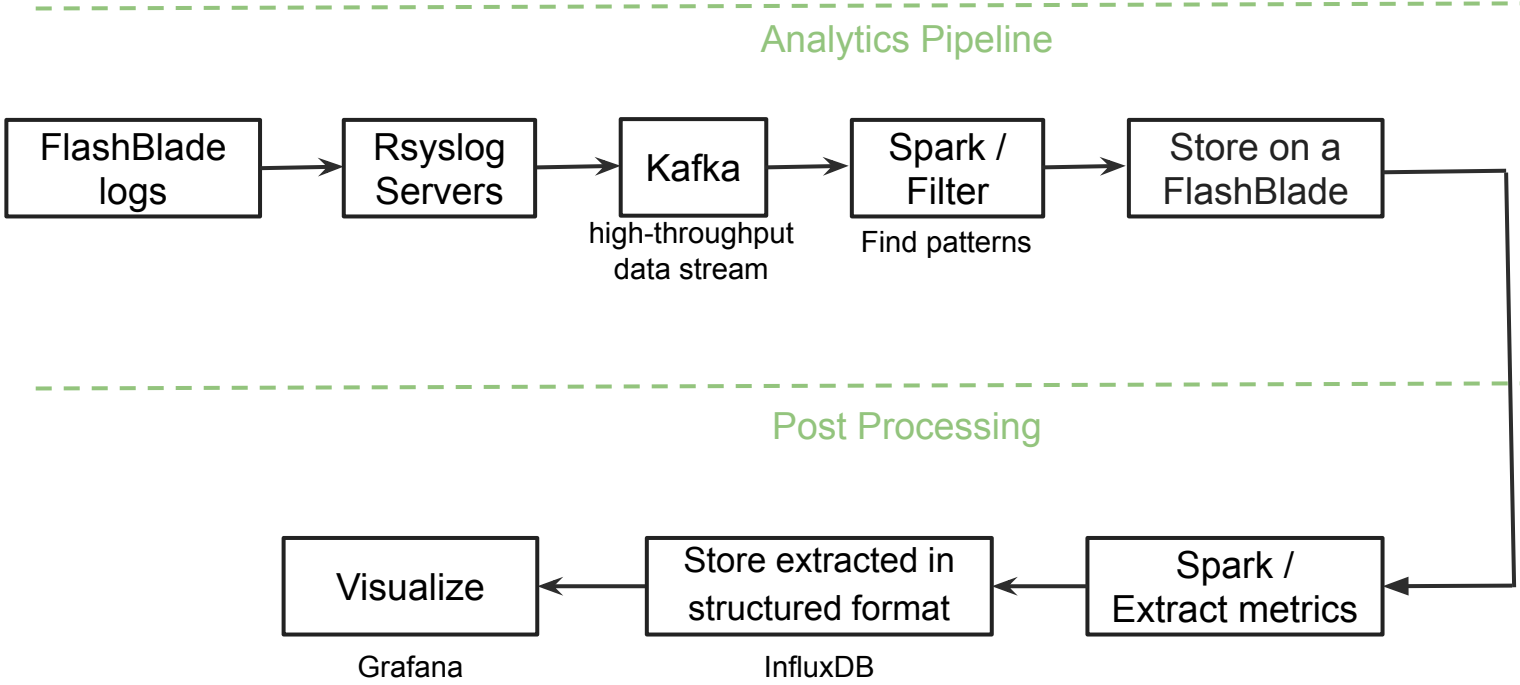

NAND data that we are interested in

- We are currently collecting the following NAND health parameters:
  - Raw bit error rate for different P/E cycle and retention points
  - Growing bad blocks
  - Retry read, soft read statistics
  - LDPC decoding latency
  - 0 -> 1 and 1 -> 0 bit flip statistics

Example of results:

Normalized Raw Bit Error Rate vs. P/E cycles, NAND part 1

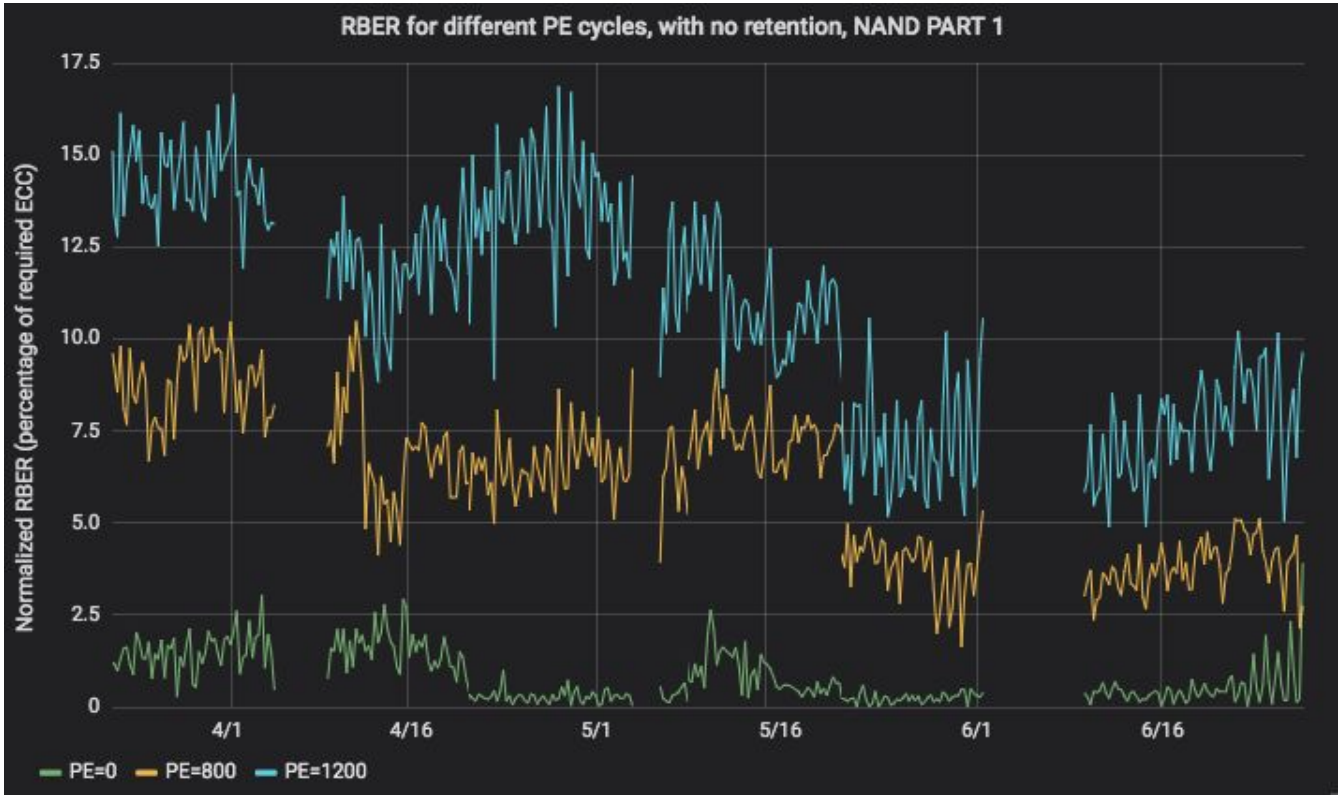

PURESTORAGE®

Example of results:

Normalized Raw Bit Error Rate vs. Retention, NAND part 1

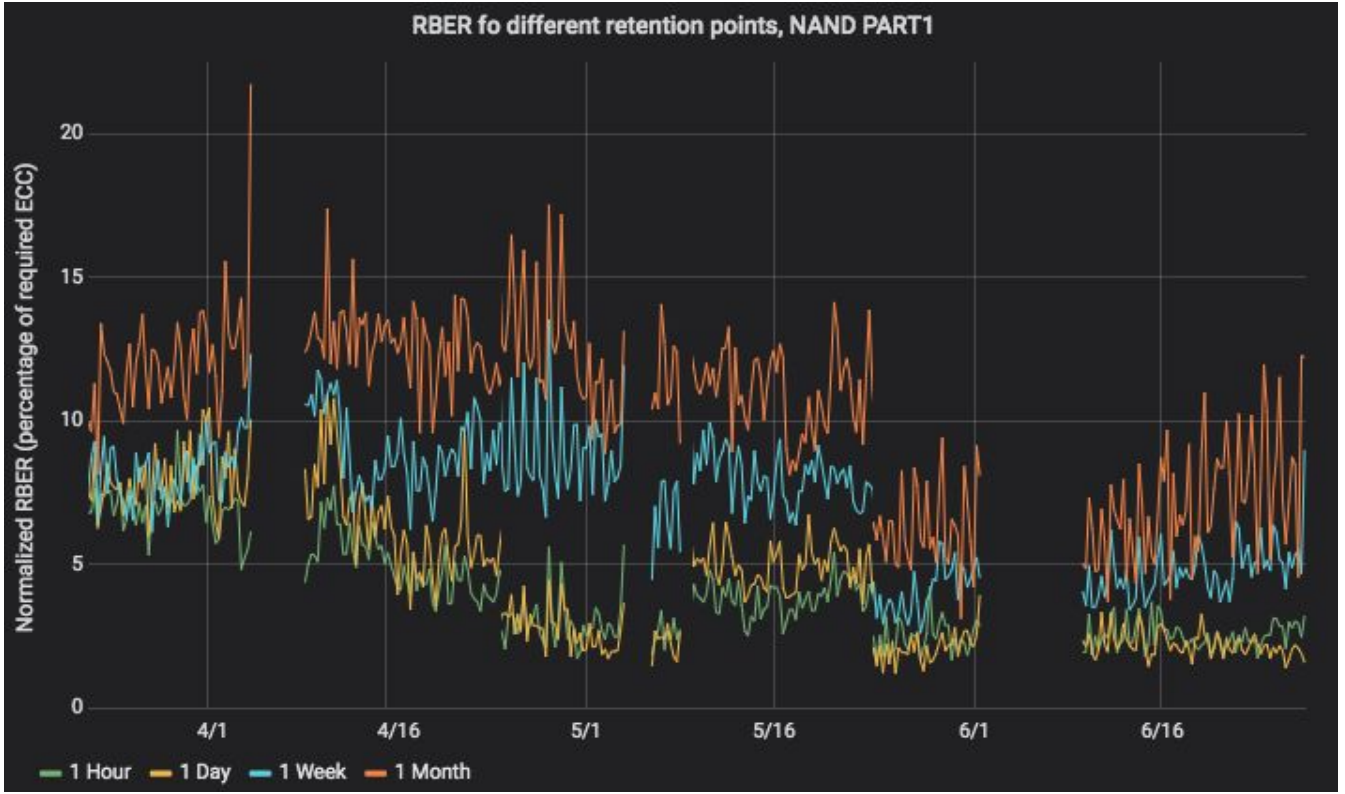

PURESTORAGE®

Example of results:

Normalized Raw Bit Error Rate vs. P/E cycles, NAND part 2

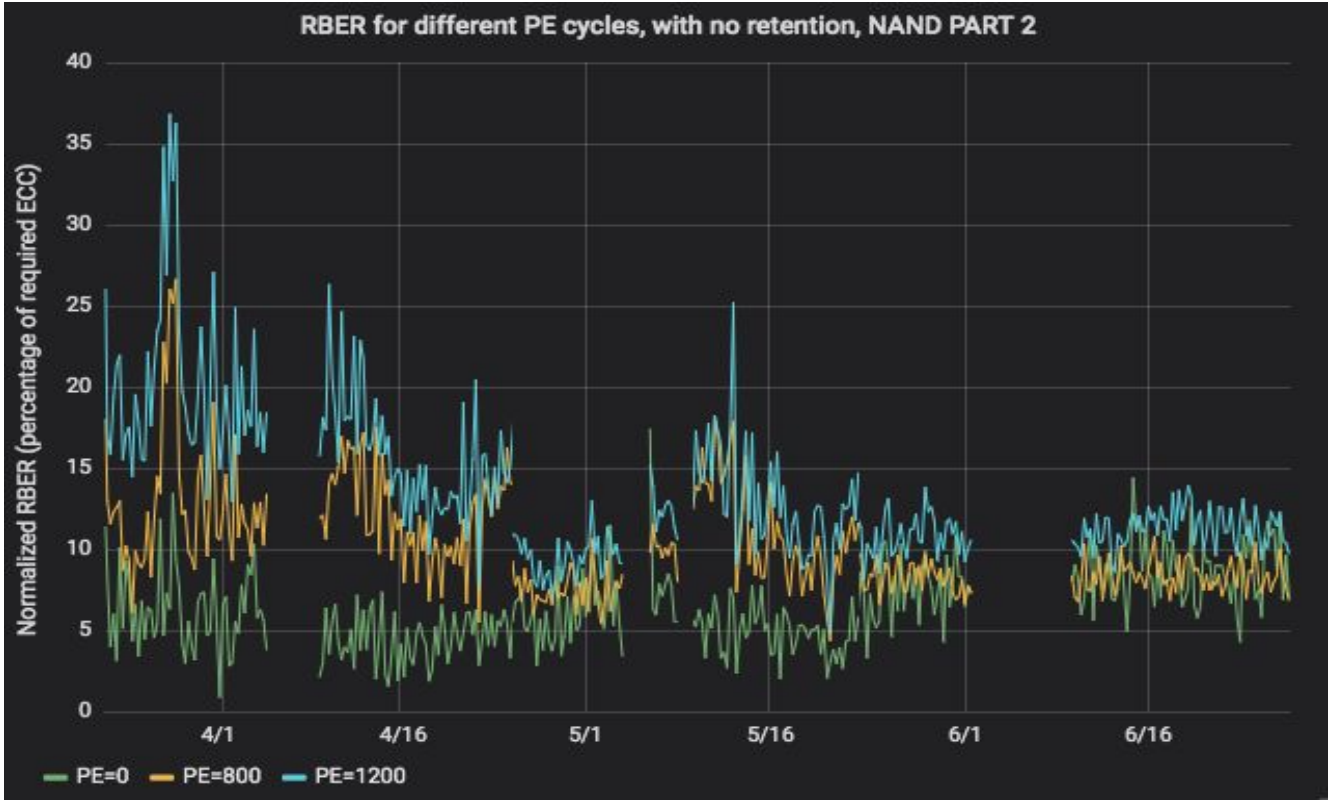

PURESTORAGE®

Example of results:

Normalized Raw Bit Error Rate vs. Retention, NAND part 2

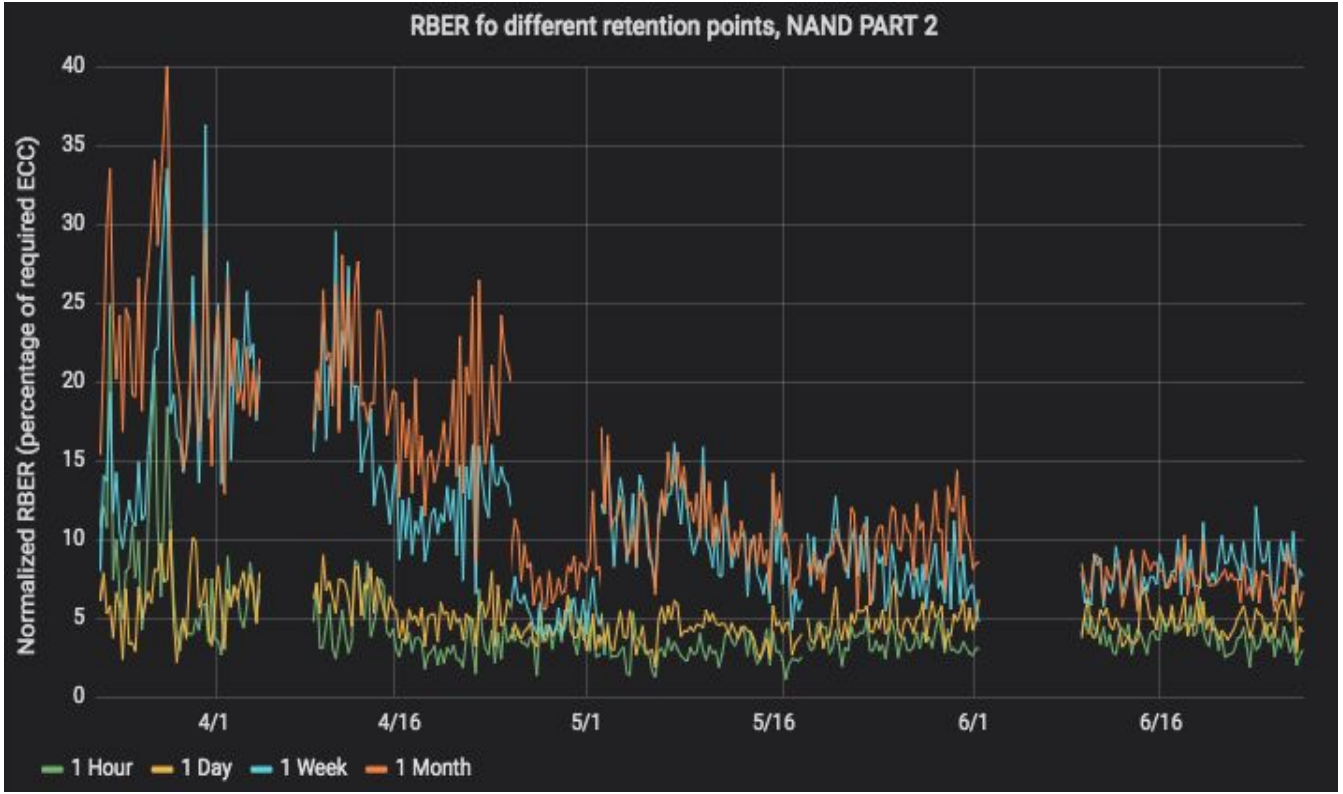

Example of results:

Normalized Raw Bit Error Rate vs. P/E cycles, NAND part 3

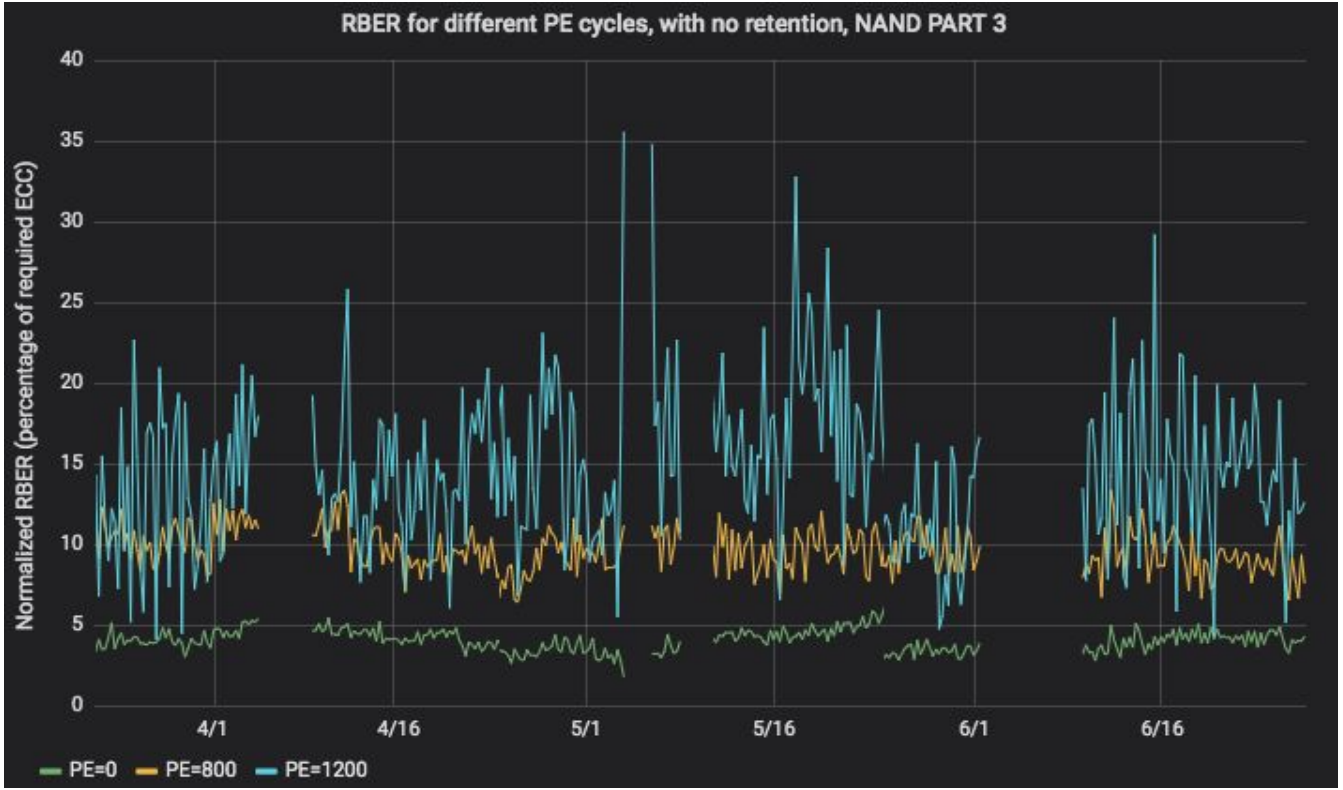

Example of results:

Normalized Raw Bit Error Rate vs. Retention, NAND part 3

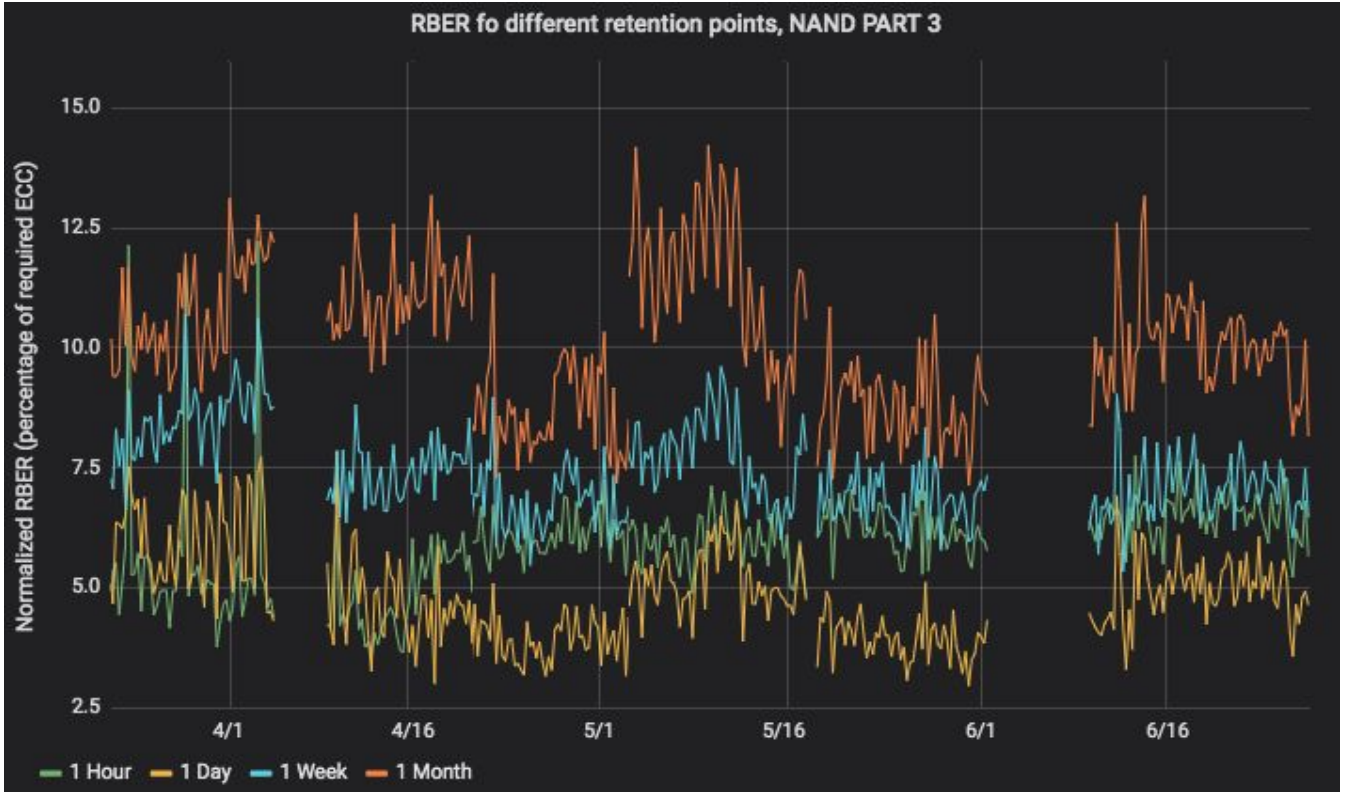

Main Takeaways

- Monitoring NAND health in the field is necessary for us
- An analytics pipeline provides a scalable solution
  - Visualization enables us to find patterns/anomalies
- The analytics pipeline can be used to monitor other HW components of the system

Thank you