

Reliability Considerations in Enterprise SSD

Michael Laflin

Marketing Director

Enpirion, Inc


mlaflin@enpirion.com

First, Some Definitions

- Reliability
 - “the probability that a piece of equipment operating under specified conditions shall perform satisfactorily for a given period of time” -EPSMA
- Failure Rate (λ)
 - “Number of units failing per time”
 - If constant and random then Reliability is: $R(t) = e^{-\lambda t}$
- MTBF
 - Mean Time Between Failure
 - $MTBF = 1/\lambda$
- FIT Rate
 - Failures per billion hours of operation
 - FITs add linearly: FIT of system = sum of component FITs

The Bath Tub Curve

- Reliability analysis is concerned with steady state region



What Does MTBF Really Mean?

- Reliability: $R(t) = e^{-\lambda t} = e^{-t/MTBF}$
- When $t = MTBF$, then $R(MTBF) = e^{-1} \cong 0.37$
 - For a large number of units, only 37% of their operating times will exceed the MTBF
 - For a single unit, the probability that it will work up to its MTBF is only 0.37
- MTBF is NOT the expected service life

Enterprise SSD Target MTBF

- Historical Target MTBF is 2,000,000 Hrs
- 2,000,000 hrs = **500** FITs
- Embedded is 4,000,000 hrs or 250 FITs
- This is the maximum composite failure rate for ALL components in the SSD.

A Few More Definitions:

- **Soft Failure**
 - Recoverable failure (bits/bytes/blocks, even chips)
 - Data can be recovered using ECC or other redundancy
 - Applies to NAND and DRAM; focus of controller IP

- **Hard Failure**
 - Complete data loss

- **Redundancy**
 - The addition of capacity to enable failure recovery
 - Can be additional bits such as ECC
 - Can be reserve capacity
 - Can be RAID , remote backup, ...

A Few Words on Redundancy

- With enough redundancy reliability can be ~infinite
- Redundancy **ALWAYS** adds cost
- Redundancy **ALWAYS** adds latency and/or requires more processing horsepower
- The amount of redundancy required is proportional to the expected failure rates

Components Subject to Failure

- NAND, DRAM
- Flash Controller, RAID Controller, PCIe Switch, Terminations, Buffers, etc
- Passives
- Power Sub-System
 - DC-DC Converters
 - Linear Regulators
 - Inductors, passives
 - Power Fail (switches, Bulk Caps, Battery,...)



Power Subsystem


"Power is a fundamental need but it can also be the biggest threat to the reliability and operation of any system".

-BitMicro DataSentinal Whitepaper



Power Fail Architecture

- Prevents data loss in buffer memory in the event of an upstream power loss
- Use energy storage device, (battery, SuperCap, or capacitor array) to provide power for orderly shut-down



Power Fail Comments

- MTBF of caps or battery can be very poor.
 - Significant vulnerability.
- It is difficult to assess health of energy storage elements.
 - Must design to worst case.
- Cannot mitigate against down stream supply failure.
 - Catastrophic data loss

Power Converter Reliability Prediction

- Assess using standard methodologies
 - MIL Handbook 217-F
 - Telcordia SR 332
- FIT rate of the converter is the sum of the FIT rates of the components

DCDC Converter Component FITs

ITEM	FIT RATE	COMMENT
Controller	1-4	
MOSFETs	1-4	
Tantalum Capacitors	5-10	Temp, Bias, Ripple Dependent
Aluminum Electrolytic	20-300	Temp, Bias, Ripple Dependent
Ceramic Capacitors	0.1	Typical value for X5R, X7R
Chip Resistors	0.1	
Inductor	38	

Example DCDC Converter MTBF

ITEM	FIT RATE	#	TOTAL FITs	MTBF
Controller	2	1	2	
Tantalum Capacitors	7.5	2	15	
Ceramic Capacitors	0.1	6	0.6	
Chip Resistors	0.1	4	0.4	
Inductor	38	1	38	
TOTAL			56	2,038 yrs

- FIT rate for just one single converter is ~56
- Recall that target for entire drive is 500

Supply Related FITs for SATA Drive

Supply Rail	FIT RATE	MTBF
Controller Core	56	
Controller IO	56	
Controller Memory	56	
FLASH Core	56	
FLASH IO	56	
DDR VDDQ	56	
DDR VTT	56	
TOTAL	392	291 yrs

Typical FIT Rate by Drive Type

Drive Type	Number of Rails	Supply Related FIT Rate
mSATA	3	168
SATA	7	392
SAS	9	504
PCIe	12	672

- Failure rates exceed target or leave no headroom

Supply FIT Mitigation

- Limited traditional options:
 - Consolidate supply rails where possible
 - Try to eliminate Electrolytic and Tant Caps
- Leverage new converter technologies
 - PowerSoC or other module approach

PowerSoC FIT/MTBF

ITEM	FIT RATE	#	TOTAL FITs	MTBF
PowerSoC	3.2	1	3.2	
Tantalum Capacitors	7.5	0	0	
Ceramic Capacitors	0.1	6	0.6	
Chip Resistors	0.1	4	0.4	
Inductor	38	0	0	
TOTAL			4.2	27,200yrs

- Contrast to discrete DCDC with 56 FITs

FIT Rate Comparison

Drive Type	Discrete DCDC	PowerSoC
mSATA	168	12.6
SATA	392	29.4
SAS	504	37.8
PCIe	672	50.4

- PowerSoC provides > 10x improvement
- Increases FIT headroom for other components

- Power conversion is a serious risk to Enterprise class Solid State Storage reliability.
- Power Fail circuits do not protect against down stream supply failures and represent significant reliability risks of their own.
- Discrete DCDC converters consume most or all of the target 500 FITs allowed for 2,000,000 hrs MTBF.
- PowerSoC offers a > 10x improvement in power supply related reliability for Enterprise SSD
- Enpirion is the inventor of, and leader, in PowerSoC
 - Visit us at: WWW.ENPIRION.COM