



Validating Analytic Write Amplification Models

Thomas W. McCormick
Chief Engineer/Technologist



Santa Clara, CA
August 2016



Background - Embedded Systems

- “Fixed function system”
 - Telecom, automotive, industrial control systems, medical equipment ...
- Commonality: Flash Storage
 - Code & data




[Wikipedia – Creative Commons]



Motivation

- Smaller process, more bits per cells
 - Endurance & ECC
 - Retention
- Challenge: Maintain acceptable service-life for embedded systems





Lifetime & Write Amplification

$$\text{Lifetime} \propto \frac{(\text{Capacity})(\text{Endurance})}{(\text{Write Amplification})}$$

$$\text{Write Amplification} = \frac{\text{Data Written to Flash}}{\text{Data Written by Host}}$$

- Write amplification factor (WAF) is a function of the Flash Translation Layer (FTL)
- FTLs are trade secrets. WAF measurements are largely unpublished [McCormick FMS 2014 & FMS2015]



Analytic WAF Equations

- 2010 Agrawal
- 2012 Luojie
- 2014 Desnoyers
- $WAF = f(OP)$



Over-Provisioning Definitions

$$OP_{\text{Marketing}} = \frac{\text{Physical Flash Size (GB)} - \text{User Data Space (Decimal GB)}}{\text{User Data Space (Decimal GB)}} \cdot 100\%$$

[Smith - FMS2012]

$$OP_{\text{Nominal}} = \frac{\text{Physical Flash Size (GB)} - \text{User Data Space (GB)}}{\text{User Data Space (GB)}} \cdot 100\%$$

$$OP_{\text{True}} = \frac{\text{Total Data Space (GB)} - \text{User Data Space (GB)}}{\text{User Data Space (GB)}} \cdot 100\%$$




Over-Provisioning Values

Configuration	IDEMA Capacities (Decimal GB)
Extended (SLC)	32 GB, 64 GB, 128 GB, ...
Standard (MLC)	30 GB, 60 GB, 120 GB, ...
Enterprise	25 GB, 50 GB, 100 GB, ...

Configuration	OP (Marketing)	OP (Nominal)
Extended (SLC)	0%	7%
Standard (MLC)	7%	15%
Enterprise	28%	37%

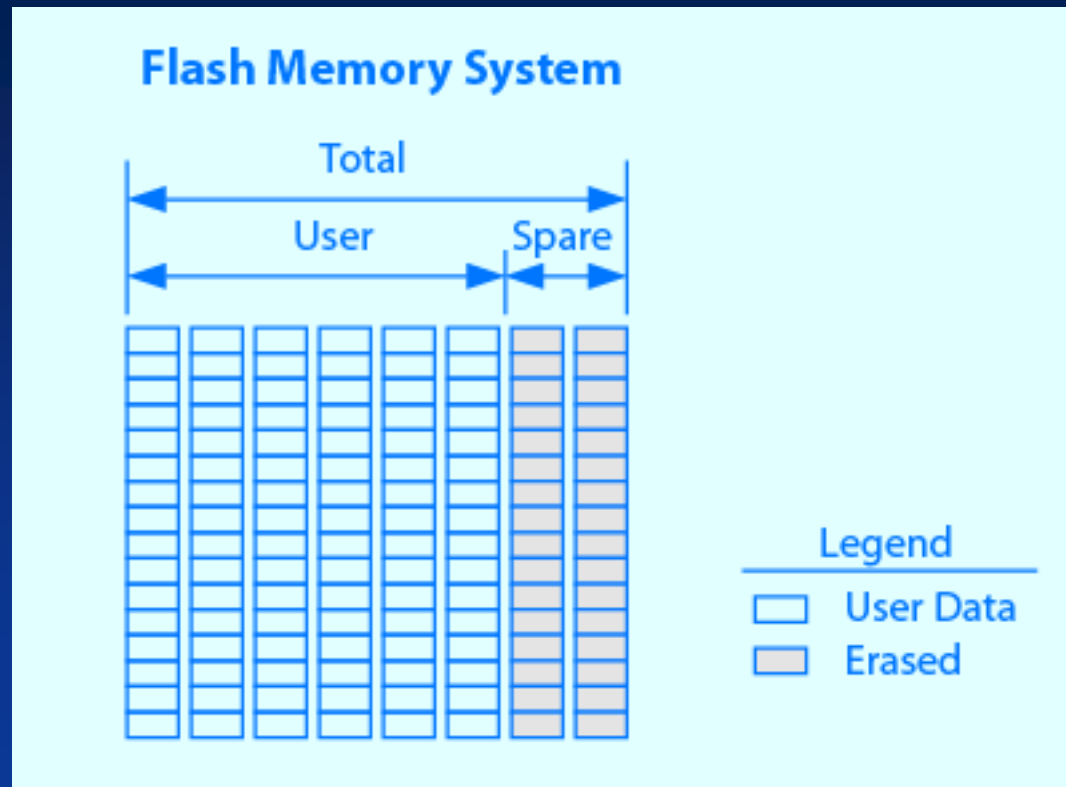


Analytic WAF Model - Development






Analytic WAF Model - Development






Analytic WAF Model - Development

Flash Memory System




Analytic WAF Model - Development



Analytic WAF Model - Development

Flash Memory System



$$WAF = \sum_{i=0}^{\infty} \left[\frac{N_p}{N_{pi}} \right]_i = \left[\frac{N_p}{N_{pi}} \right]_{Avg}$$


$$\left[\frac{N_p}{N_{pi}} \right]_{Avg} = \frac{T - U}{T}$$

$$\rho = \frac{T - U}{U}$$

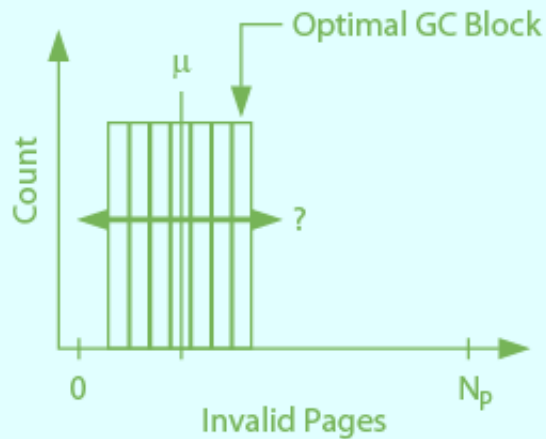
$$WAF = \frac{1 + \rho}{\rho}$$

[Bux 2009]

Analytic WAF Models - Development



Analytic WAF Models - Development



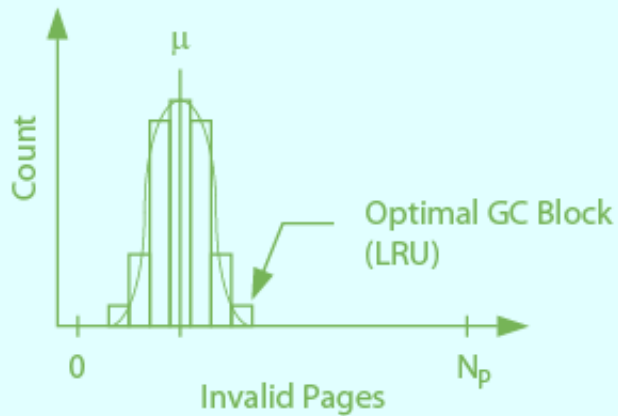
$$WAF = \sum_{i=0}^{\infty} \left[\frac{N_P}{N_{PI}} \right]_i$$

$$\left[\frac{N_P}{N_{PI}} \right]_i = 2 \left[\frac{N_P}{N_{PI}} \right]_{Avg} = 2 \left[\frac{T-U}{T} \right]$$

$$WAF = \frac{1}{2} \left[\frac{1+\rho}{\rho} \right]$$

[Agrawal 2010]

Analytic WAF Models - Development



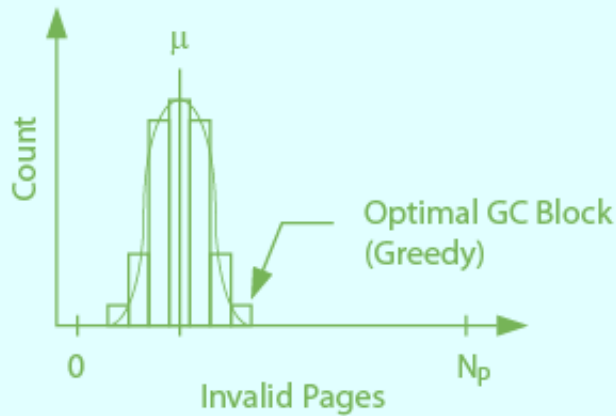
$$\rho = \left[1 - \frac{1}{UN_p} \right]^k$$

$$WAF = \frac{(-1 - \rho)}{(-1 - \rho) - W((-1 - \rho)\exp(-1 - \rho))}$$

[Luojie 2012]

W(): Lambert W function (Google:VBA Lambert)

Analytic WAF Models - Development

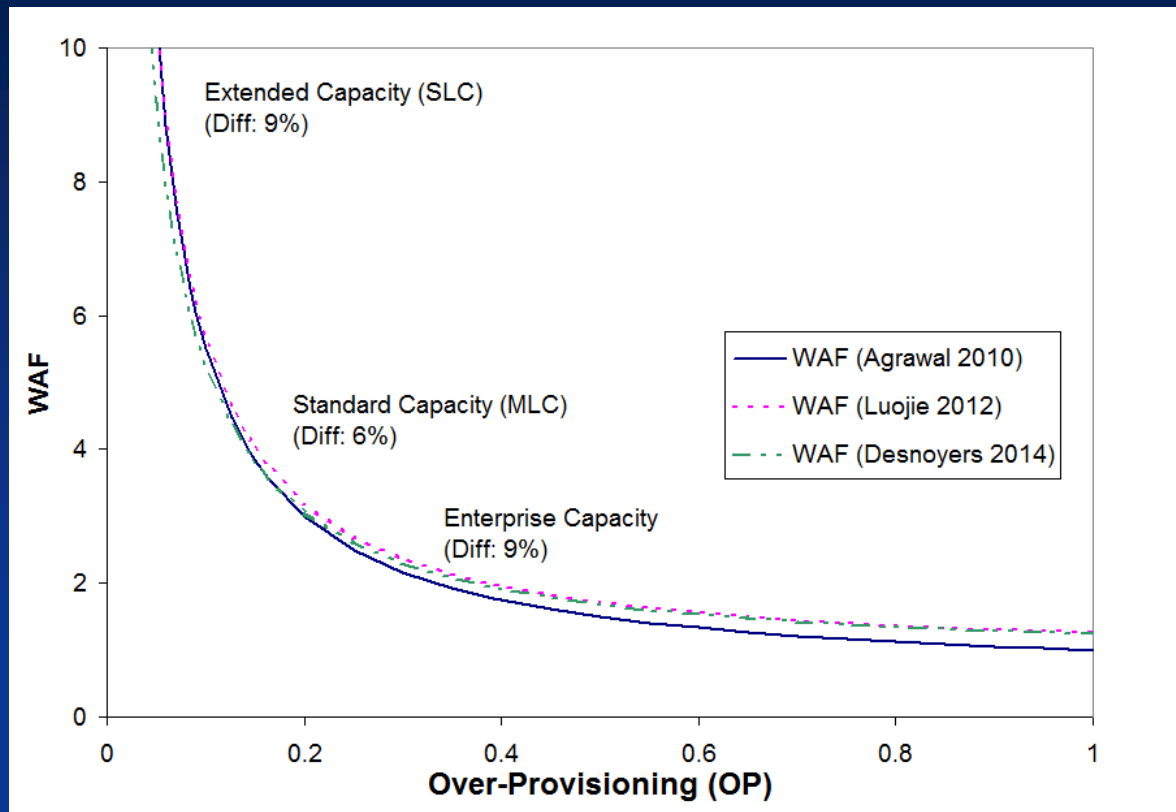


$$X_0 = \frac{1}{2} - \frac{N_p}{\alpha} W\left[-\left(1 + \frac{1}{2N_p}\right)\alpha \exp\left[-\left(1 + \frac{1}{2N_p}\right)\alpha\right]\right]$$

$$WAF = \frac{N_p}{N_p - (X_0 - 1)}$$

[Desnoyers 2014]

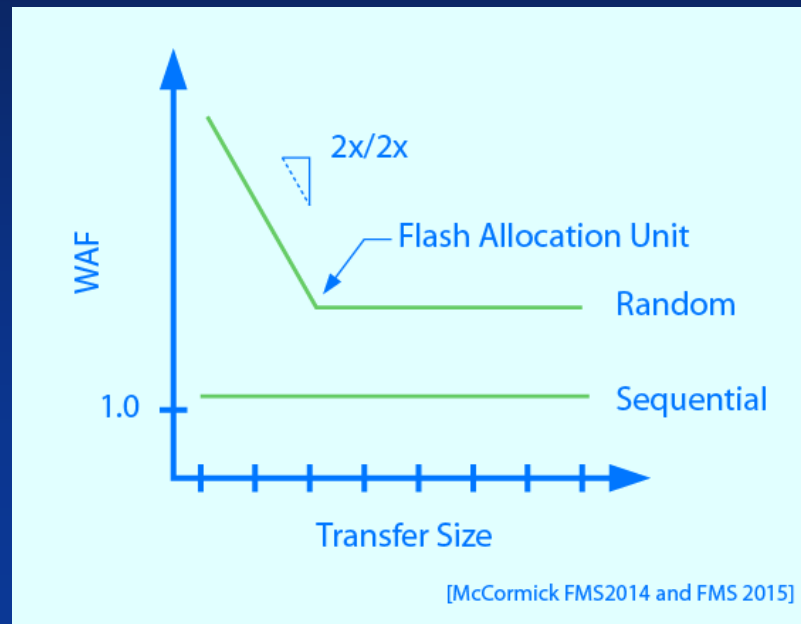
Analytic WAF Models - Predictions





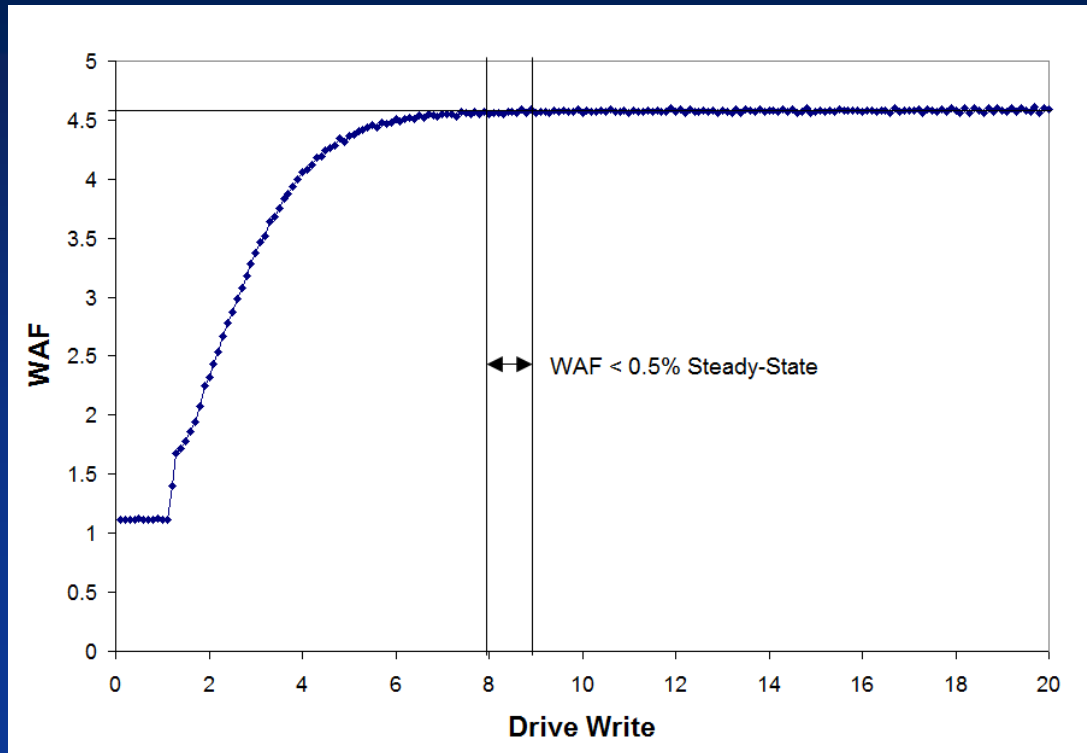
WAF Measurements – Write Size

- WAF is for random writes to Flash Allocation Unit

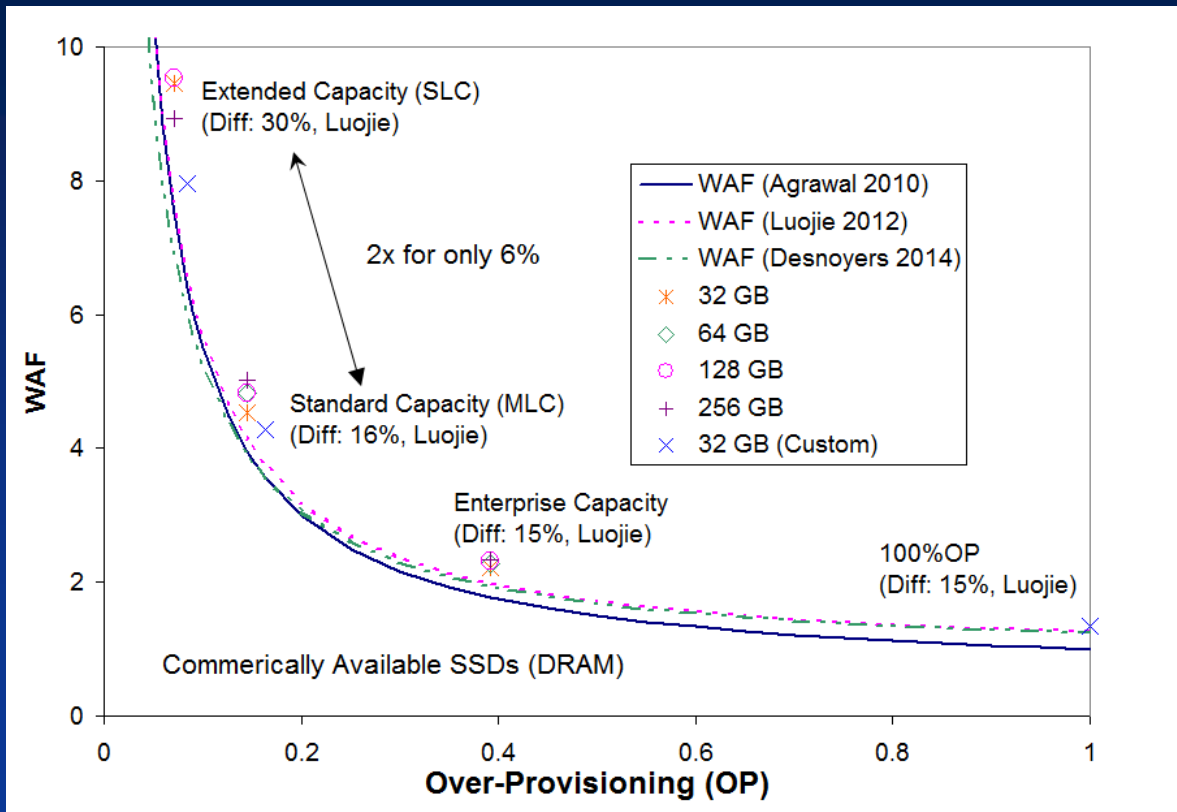





WAF Measurements – Steady-State



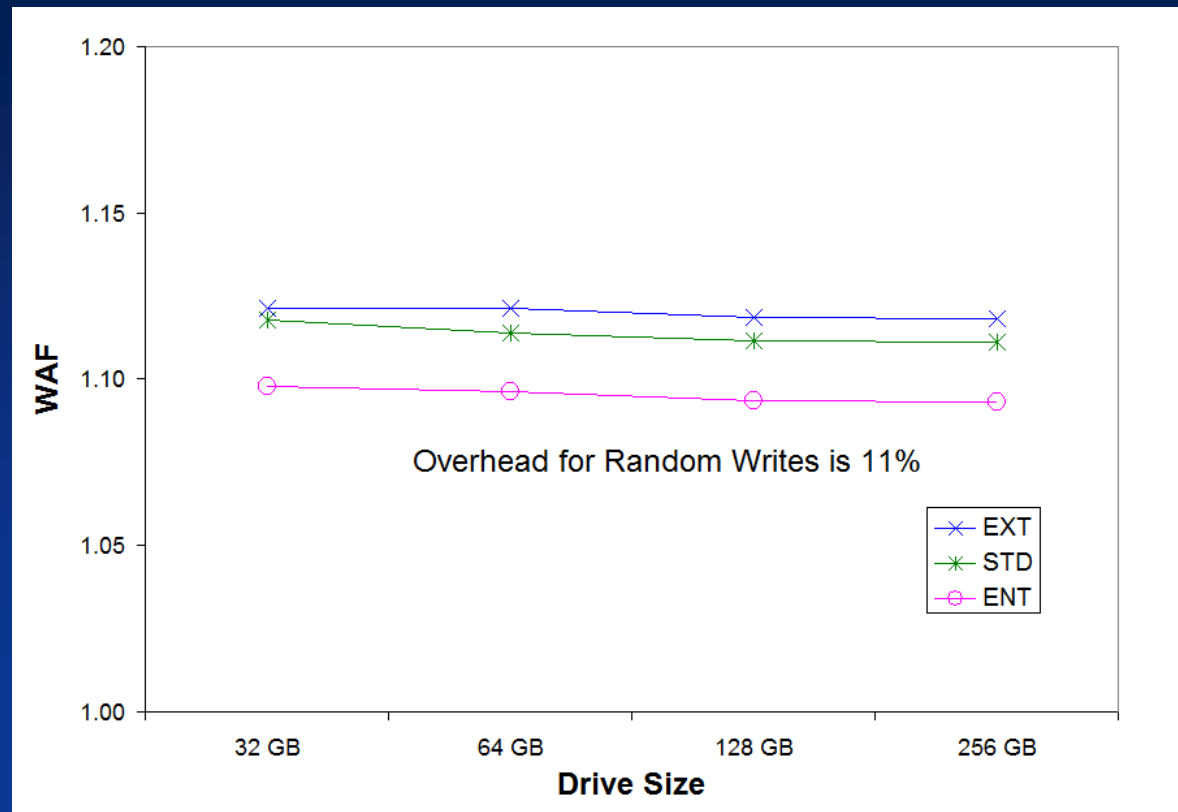
WAF Measurements



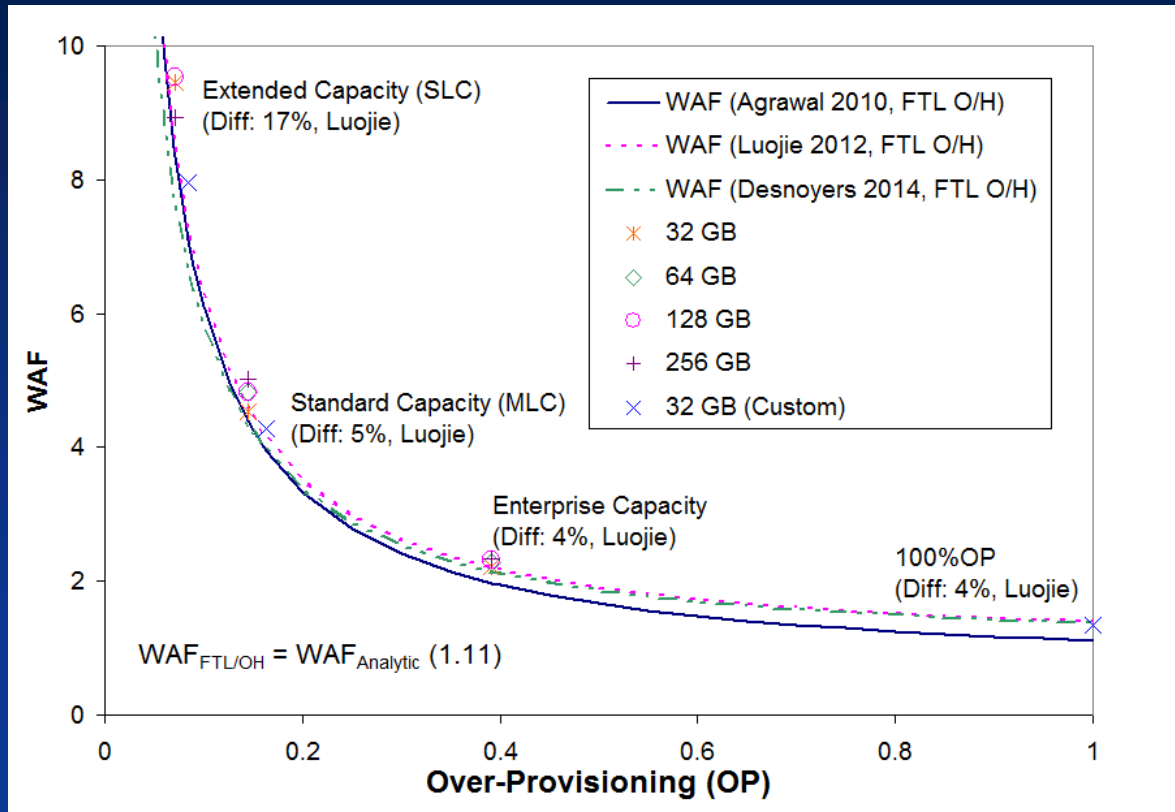
WAF Measurements – FTL O/H



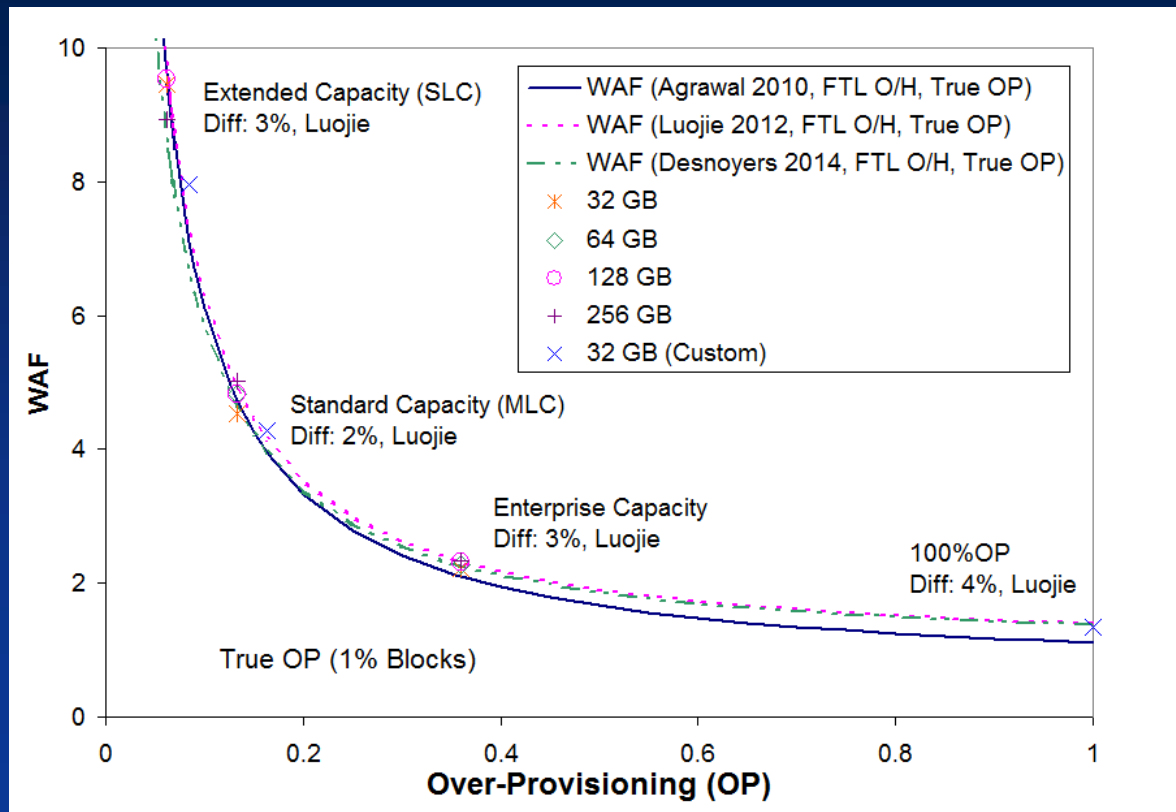
WAF Measurements – FTL O/H



WAF Measurements – FTL O/H



WAF Measurements – True OP





Case Study – 100% OP vs pSLC

$$\frac{\text{Lifetime}_{100\%OP}}{\text{Lifetime}_{Standard}} = \frac{\left(\frac{\text{Endurance} \cdot \text{Physical Drive Size}}{\text{WAF}} \right)_{100\%OP}}{\left(\frac{\text{Endurance} \cdot \text{Physical Drive Size}}{\text{WAF}} \right)_{Standard}}$$

$$\frac{\text{Lifetime}_{100\%OP}}{\text{Lifetime}_{Standard}} = \frac{\left(\frac{3000 \cdot 512 \text{ MB}}{1.24} \right)_{100\%OP}}{\left(\frac{3000 \cdot 512 \text{ MB}}{7.64} \right)_{Standard}}$$

$$\frac{\text{Lifetime}_{100\%OP}}{\text{Lifetime}_{Standard}} = 6.14x$$

$$\frac{\text{Lifetime}_{pSLC}}{\text{Lifetime}_{Standard}} = \frac{\left(\frac{\text{Endurance} \cdot \text{Physical Drive Size}}{\text{WAF}} \right)_{pSLC}}{\left(\frac{\text{Endurance} \cdot \text{Physical Drive Size}}{\text{WAF}} \right)_{Standard}}$$

$$\frac{\text{Lifetime}_{pSLC}}{\text{Lifetime}_{Standard}} = \frac{\left(\frac{20000 \cdot 256 \text{ MB}}{7.64} \right)_{pSLC}}{\left(\frac{3000 \cdot 512 \text{ MB}}{7.64} \right)_{Standard}}$$

$$\frac{\text{Lifetime}_{pSLC}}{\text{Lifetime}_{Standard}} = 3.33x$$



WAF Measurements - Workload

- Sequential
 - WAF ~1.0
- Random
 - WAF = f(OP)
- Enterprise & Client (JEDEC)
- Embedded?



Enterprise Workload

- JEDEC 219A Workload:

- 512 bytes (0.5K) 4%
- 1024 bytes (1K) 1%
- 1536 bytes (1.5K) 1%
- 2048 bytes (2K) 1%
- 2560 bytes (2.5K) 1%
- 3072 bytes (3K) 1%
- 3584 bytes (3.5K) 1%
- 4096 bytes (4K) 67%
- 8192 bytes (8K) 10%
- 16,384 bytes (16K) 7%
- 32,768 bytes (32K) 3%
- 65,536 bytes (64K) 3%

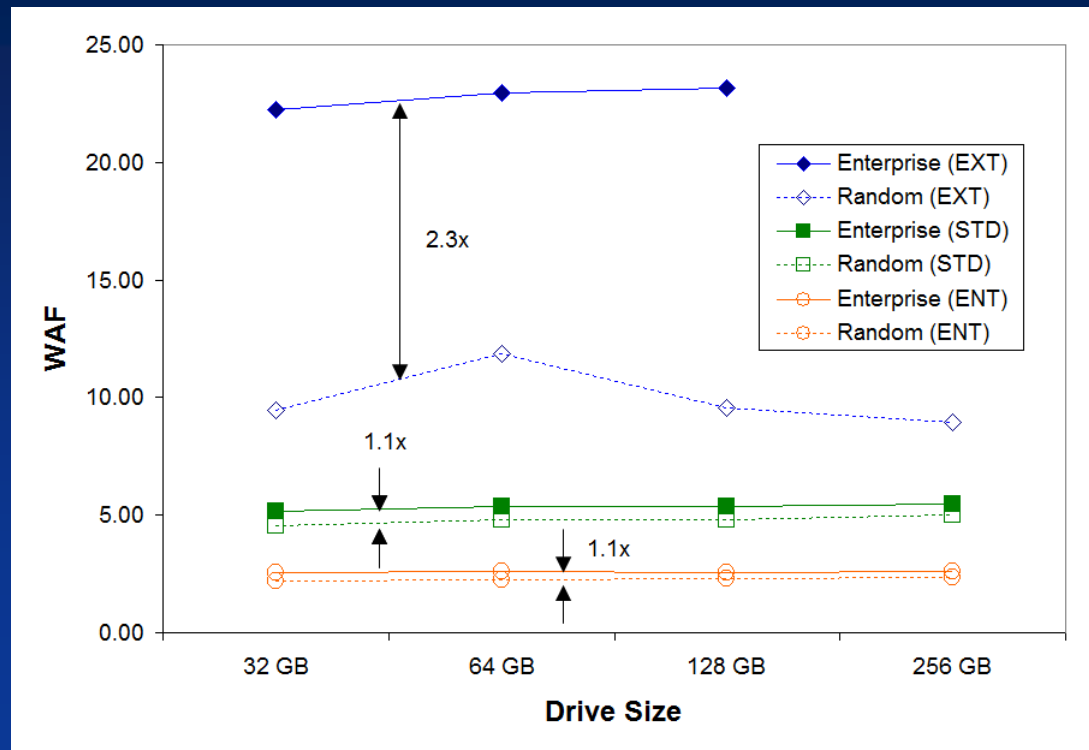
- JEDEC 219A Workload Range:

- A: 5% LBA Range, 50%
- B: 15% LBA Range, 30%
- C: 20% LBA Range, 20%

$$WAF_{\text{Enterprise}} = f(WAF_{\text{Random}})$$

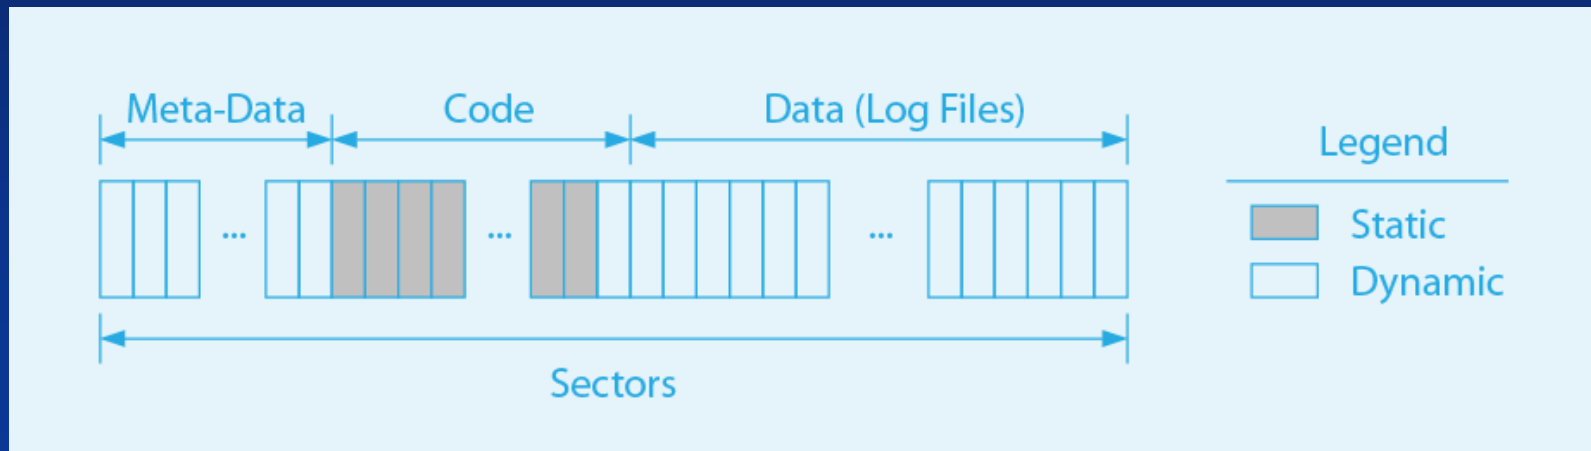


WAF Measurements: Enterprise and Random



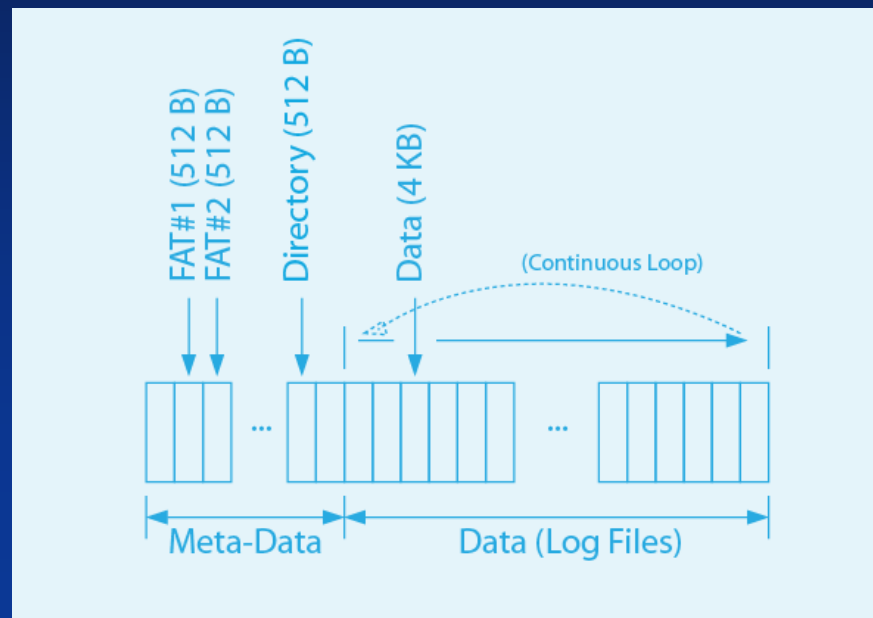
WAF Measurements - Workload

- Embedded Workload
 - File System: Code and Logs



WAF Measurements - Workload

- Embedded Workload: Sequential & “Repeated”





Summary & Conclusions

1. WAF_{Random} equations align with measured values
 - Commercially Available SSD (DRAM) -> “ideal”
 - $WAF_{\text{Random}} = f(OP)$
 - Consider FTL O/H and True OP (<5%)
2. Avoid Extended (SLC)
 - 50% Lifetime Reduction for 6% Capacity Increase
3. 100% OP ~2x Lifetime Improvement over pSLC
4. $WAF_{\text{Enterprise}} = f(WAF_{\text{Random}}, OP)$
5. Need for Embedded Workload definition



More Embedded Sessions

- Flash Memory System Embedded Events:
 - Beer, Pizza, and Chat with the Experts
 - Tues 7:00 – 8:30 PM
 - Embedded Applications, Part 1 (203-B)
 - Wed 3:50 – 4:55 PM
 - Embedded Applications, Part 2 (204-B)
 - Wed 5:10 – 6:15 PM



Questions?

Tom McCormick - Chief Engineer/Technologist
Swissbit

tom.mccormick@swissbit.com

