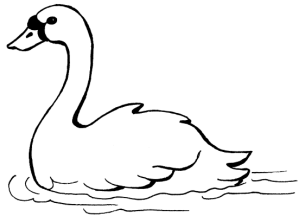

SWAN: Alleviating Garbage Collection Interference through Spatial Separation in All Flash Arrays

Jaeho Kim, Kwanghyun Lim*, Youngdon Jung,
Sungjin Lee, Changwoo Min, Sam H. Noh

VIRGINIA TECH

DGIST

UNIST

*Currently with Cornell Univ.

Image: https://clipartix.com/swan-clipart-image-44906/

All Flash Array (AFA)

- What is AFA?
  - Storage infrastructure that contains only flash memory drives
    - Also called Solid-State Array (SSA)

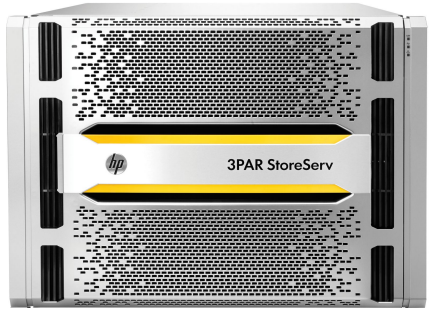

https://images.google.com/
https://www.purestorage.com/resources/glossary/all-flash-array.html

SSDs for Enterprise

Manufacturer	Product Name	Seq. Read Throughput	Seq. Write Throughput	Capacity
Intel	DC P4800X	2.5 GB/s	2.2 GB/s	1.5 TB
	DC D3700	2.1 GB/s	1.5 GB/s	1.6 TB
	DC P3608	5 GB/s	3 GB/s	4 TB
Samsung	PM1633a	3.5 GB/s	3 GB/s	1.6 TB
	SM883	6.3 GB/s	0.9 GB/s	960 GB

Intel: https://www.intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds.html
Samsung: https://www.samsung.com/semiconductor/ssd/enterprise-ssd/

Previous Solutions

Solutions	Write Strategy	How Separate User & GC I/O	Disk Organization
Harmonia [MSST'11]	In-place write	Temporal (Idle time)	RAID-0
HPDA [IPDPS'10]	In		RAID-4
GC-Steering [IPDPS'18]	In-place write	Temporal	RAID-4/5
SOFA [SYSTOR'14]	Log write	Temporal	Log-RAID
SALSA [MASCOTS'18]	Lo		Log-RAID
Purity [SIGMOD'15]	Log write	Temporal	Log-RAID
SWAN (Proposed)	**L**		**2D Array**

1. Traditional RAID

2. Log-(based) RAID

3. SWAN

Comparison of RAID Schemes

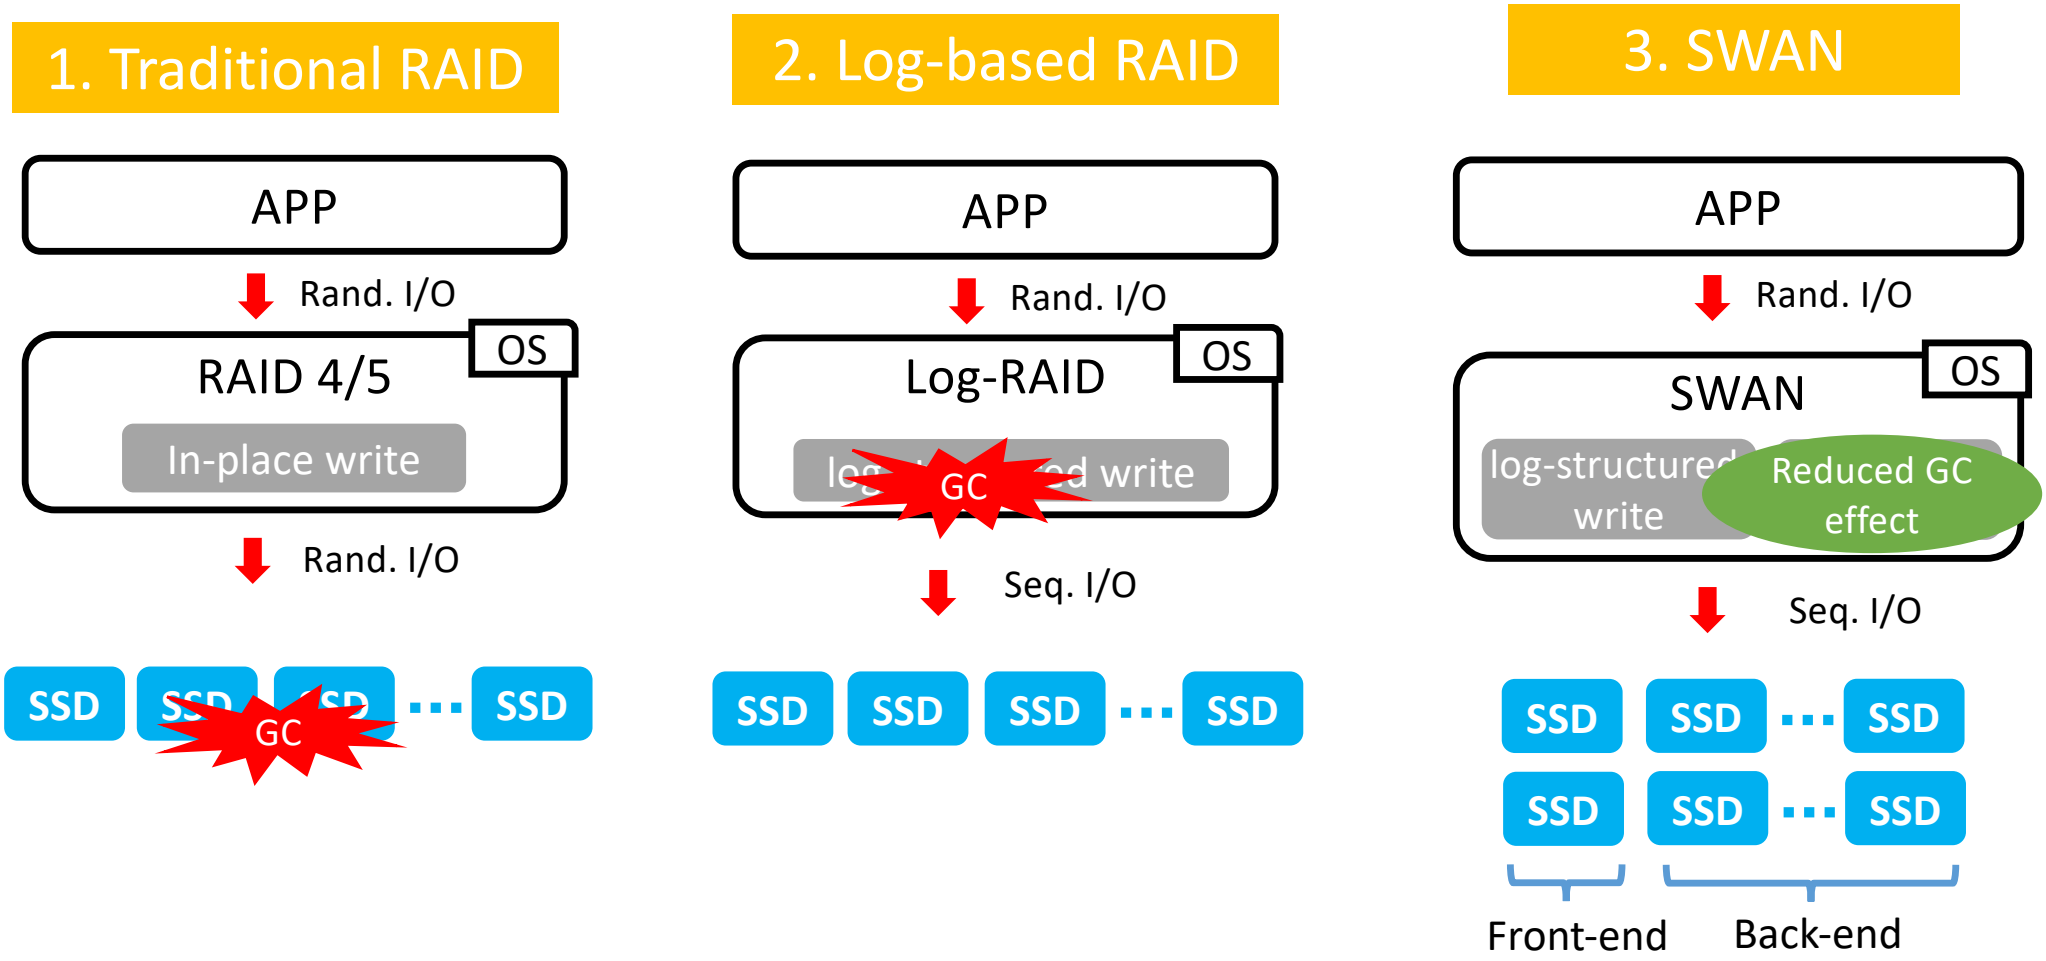

Summary of SWAN

1) Provide full write performance of an array of SSDs up to network bandwidth limit

2) Alleviate GC interference through separation of I/O induced by application and GC of AFA

3) Introduce an efficient way to use SSDs in All Flash Array

9:15 AM, Session Track 1, on July 12th Friday