Ten Ways to Improve Flash Storage System Performance

Camberley Bates, Evaluator Group

@camberleyB

Panel Moderator

Abstract

- Quite frequently, the wonderful flash storage systems, despite having NVMe, 3D flash, and other new technologies, doesn't run fast enough. The customer needs the performance for their current environment and the next breakthrough now, if not even sooner. This session will look at what are the issues facing systems performance and what approaches might blast through the bottlenecks. We will look at current and futures that address the demands in the Data Center, in the Cloud and in Research applications.

Flash Memory Summit

The Panel

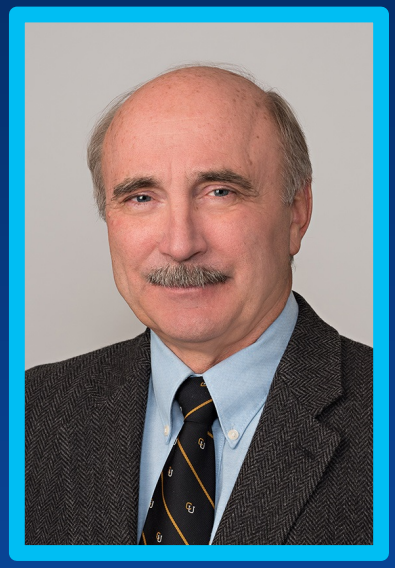

Randy Kerns,
Chief Strategist
Evaluator Group, Inc.

Dan Cobb,
Dell EMC Fellow

Andy Walls
IBM Fellow

Siamak Nazari
HPE Fellow

RANDY KERNS

EVALUATOR GROUP
SR STRATEGIST

Evaluator Group: Limiting Factors Observed

- Customer Environments
  - Network Infrastructure
  - Application
  - Drivers
  - Synchronous Replication

DAN COBB

Dell EMC Fellow

No app is an island & all apps have a lifecycle

Take a systematic approach to app performance & utilization – and include your developers

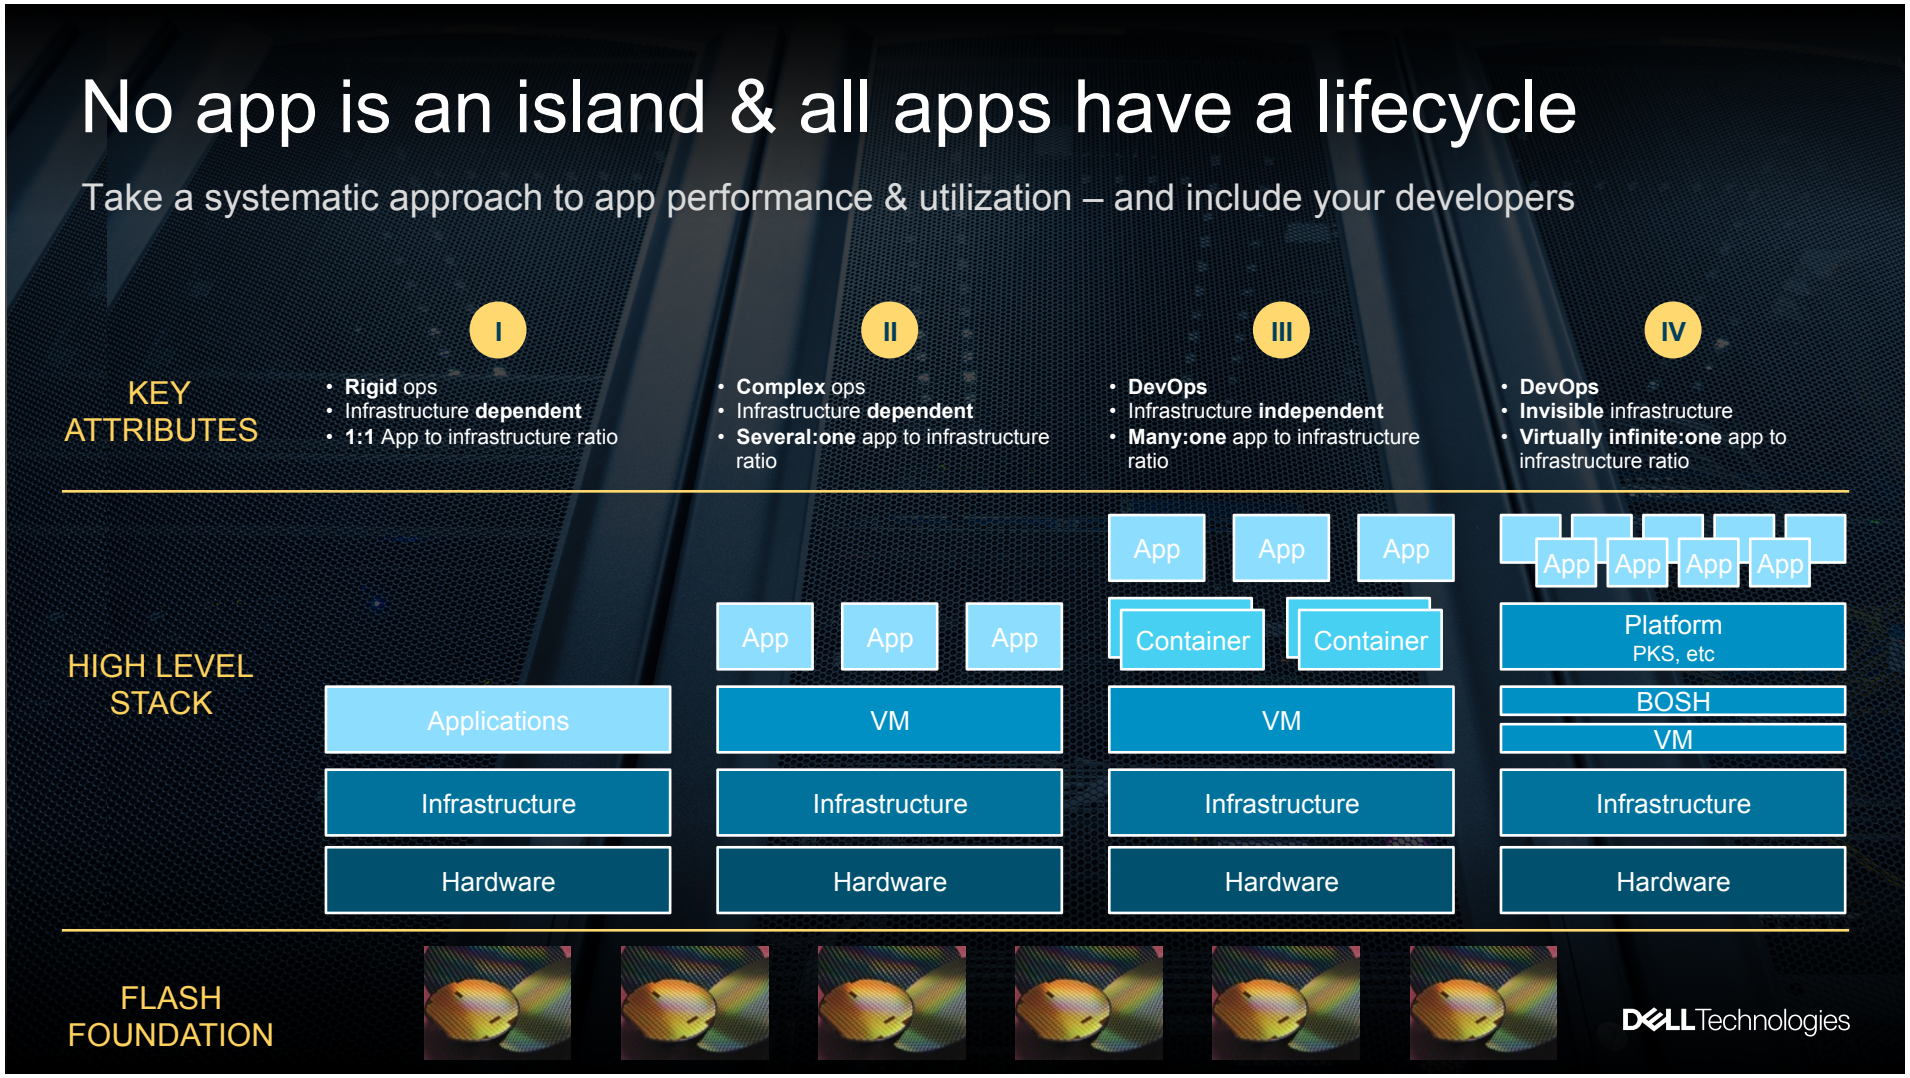

DELL Technologies

The future is not what it used to be...

Solve for today's need AND tomorrow's opportunities

	BEFORE...	...and NOW
Users	People	+ Machines, AI, bots, real-time
User Experience	Websites	+ AR/VR
Application	Java, .net	+ Microservices, 12 factor, API-driven
Data Management	Relational databases	+ In-memory dbs, memory-centric architectures
Orchestration	Single cloud	+ Cross cloud/mega-cloud
OS/Virtualization	VMs, Microsoft, Linux	+ Uni-kernels, functions, containers, serverless
Network	Discrete: FC | Ethernet | IB	+ Open networking, integrated networking, NVMoF
Storage	Block, File	+ Object, key value, memory-centric architectures
Compute	X86	+ GPU, FPGA, TPU, IPU, diverse accelerators, ARM
Media	SSD | HDD | DRAM	+ Non-volatile memories, NVMe, SCM

Flash Storage Performance

Andrew Walls
IBM Fellow, CTO Flash Storage

Flash Performance Characteristics

On the surface everything looks great!

- Flash Read Latency is over an order of magnitude faster than Spinners.
- Write latency is similar, but an AFA can do thousands in parallel.
- NVMe semantics to simplify and speed up memory access

Challenges abound

- Complex storage functions like Data Reduction add metadata that must also be fetched
- Average Flash Latency increases as we go from SLC to MLC to TLC to QLC
- Software data path introduces queuing penalties and latency inconsistencies

Techniques for keeping performance High!

- Hardware Data Path as much as possible with Software management
  - Consistency and fastest possible
- Good Caching and Tiering to hold hot data and metadata
  - DRAM – Optane/LL NAND – Optane/Persistent Memory DIMM
- Efficient Garbage Collection – Do not fill up the array excessively
  - Overprovisioning can be key to consistent performance
  - Keep Write Amplification down
- Metadata copy or using Dedupe to accelerate Snapshots and replication
- NVMe in the server for reducing latency and CPU utilization in the host

SIAMAK NAZARI

HPE Fellow

Flash Memory Summit

Q&A