



# Leveraging SSDs within Cloud Storage Infrastructures

Dave Wright  
Founder & CEO  
SolidFire

# Storage in the Cloud



## Local Storage

- Swap
- Temp Files
- Data Processing
  
- WD
- Seagate
- FusionIO



## Bulk Object Storage

- Media Files
- Content Distribution
- Backup
- Archival
  
- S3
- Atmos
- Nirvanix





## Block Storage

- VM Images
- Application Files
- Databases
  
- EMC
- Netapp
- 3par
- SolidFire

# Block Storage Challenges

- Performance
  - IOPS/GB balance
  - Multi-tenant isolation (seek latency)
- Efficiency
  - Utilization
  - Power/cooling
- Management
  - Automation
  - Reliability



Administration Interface



# Flash Benefits for Block Storage

- Performance
  - Restore balance between IOPS and Capacity
  - Better performance with varied workloads
- Efficiency
  - Remove performance as a limiter on utilization
  - Reduce storage space, power, and cooling
- Management
  - Better reliability?

# What's the downside?

- **Cost**
  - Partially a perception issue
  - Can't be ignored – cloud all about value prop
- **Storage Architectures**
  - Ability to get full performance from flash
  - Write amplification & endurance issues

# How to use Flash today

- Direct-attached
- Host based caching
- Array based caching
- Array based tiering
- All-solid-state arrays

# Direct Attached

- Examples:
  - FusionIO, Virident, 2.5" SSD
- Pros:
  - Best performance
- Cons:
  - Limited capacity
  - No sharing
  - Limited availability / redundancy



# Host based caching

- Examples:

- FB FlashCache, IO Turbine, Adaptec, Marvell



- Pros:

- Larger flash footprint
- Lower latency than array based caching



- Cons:

- Read-only, or affects data integrity/availability
- Huge delta in performance between cached/uncached reads





# Array based caching

- Examples:
  - Netapp FlashCache, EMC FAST Cache
- Pros:
  - Invisible / seamless improvement
  - Provides good boost for hottest data
- Cons:
  - Requires expensive SLC flash due to churn
  - Can cause irregular performance in multi-tenant environment



# Array based tiering

EMC<sup>2</sup>3PARcompellent

- Examples:
  - EMC FAST, 3par, Compellent
- Pros:
  - Potentially large capacity available
  - Ability to automate or manually place data in tiers
- Cons:
  - Requires expensive SLC flash due to churn
  - Controller overhead moving data back and forth
  - Irregular performance when automated

# All-Solid-State arrays

- Examples:
  - SolidFire, Nimbus, Violin, TMS
- Pros:
  - High performance for all data
  - Best \$/IOP
  - Potential for less power/cooling expense
- Cons:
  - Cost/GB?
  - Designed for cloud scale?

NIMBUS DATAviolin  
MEMORYtms

# Overcoming the \$/GB barrier

- Use of lower-cost MLC flash
  - Only 2X 15K RPM drives
- Increase utilization vs. short-stroked disk
- Using data reduction techniques:
  - Compression
  - De-duplication
  - Thin-provisioning

- Flash can help overcome key primary storage issues in the cloud
- Many possible approaches to using it
- If the right techniques are used to reduce \$/GB, all-solid-state solutions are viable