

SSDs in the Cloud

Dave Wright, CEO
SolidFire, Inc.

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 / flexibility
 - ❑ 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

- Examples:

- EMC FAST, 3par, Compellent

EMC²

- Pros:

 3PAR

- Potentially large capacity available

 compellent

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




SolidFire Solution Architecture

- ❑ Clustered Scale-Out Storage Solution
- ❑ Leverages Industry-Standard Hardware
- ❑ Native 10 GigE iSCSI
- ❑ All Solid State (SSD)
- ❑ **Usable \$/GB similar to traditional SANs**



Combine 100 nodes
Over 5M IOPS

Bridging the Performance Gap




Increasing gap between Server and Storage IO

- Increasing latency, unhappy customers
- Unable to guarantee performance
- Performance degrades with disk / device failure

SolidFire – Guarantee Storage Performance

- ❑ 50,000 IOPS/node. 5,000,000 IOPS/cluster.
- ❑ System-wide Performance Virtualization
 - ❑ Guarantees sustained performance (QoS) to every volume
 - ❑ Enables volume-level performance throttling
 - ❑ Create and bill against performance tiers
 - ❑ Real-time performance adjustments via API
- ❑ SolidFire Helix™ Data Protection – manages device failures without performance impact

New Level of Storage Efficiency



Inefficient use of purchased capacity

- Deduplication and Compression unavailable or turned off to maintain performance
- Capacity Sprawl = large power and cooling cost

SolidFire – Operational Savings

- High capacity utilization - 85%
- Proactive inline Deduplication, Compression and Thin Provisioning
 - Executed across the entire data store, w/o performance impact
 - Deduplication gains (30% - 70%)
 - Compression gains (30% - 70%)
- Drastic reduction in power, cooling, and floor space

Cloud Specific Management



Administration Interface

Proprietary Management

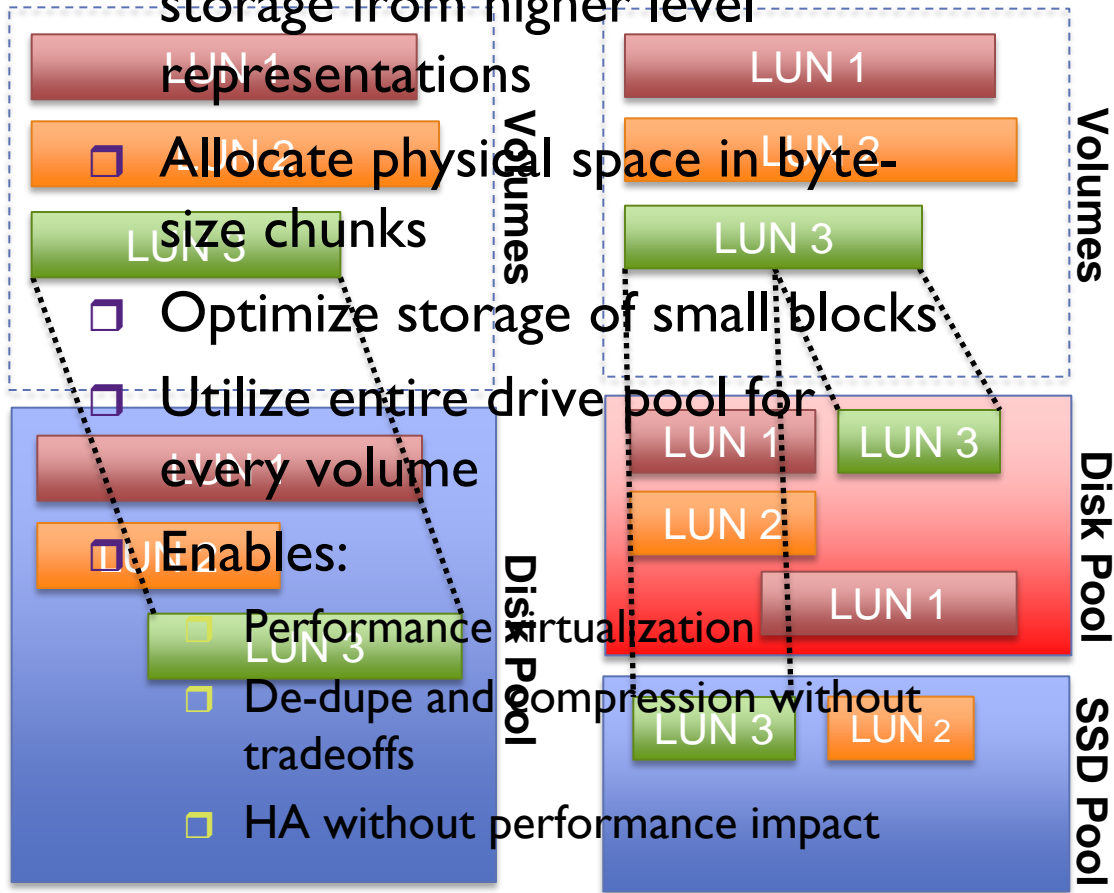
- Complex, Individual GUI interface
- Manual deployment and provisioning
- Lack of “cloud specific” API capability

SolidFire Advantage

- ❑ REST-based API (management and automation)
 - ❑ Designed around the cloud service providers' needs to automate storage management and provisioning
 - ❑ Extremely simple to build user-facing storage controls with your own business rules
 - ❑ Enables multi-tenant provisioning, reporting, and performance isolation
 - ❑ Any SolidFire function can be automated via API

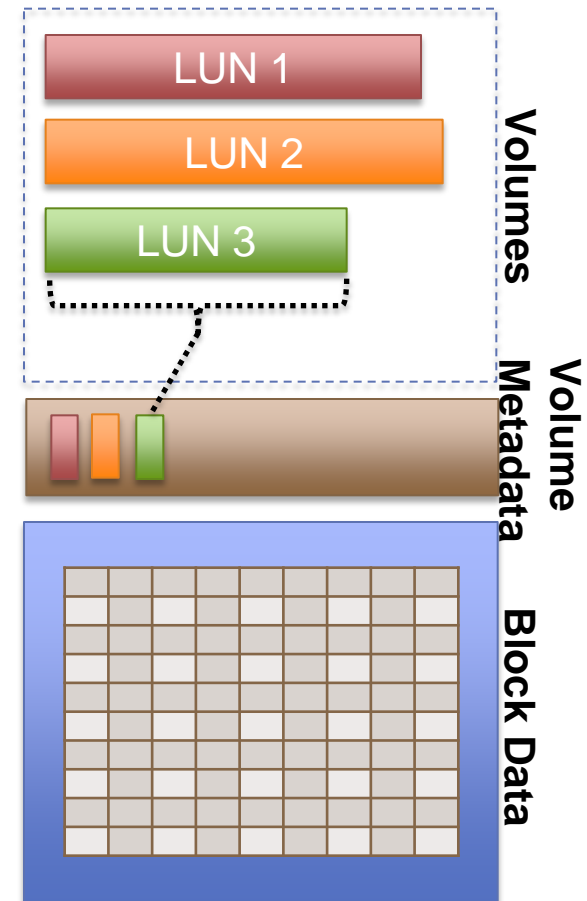
SolidFire Architecture

- ❑ **SAN 1.0** Completely separate data storage from higher level representations
- ❑ **SAN 2.0** Allocate physical space in byte-size chunks



- ❑ Optimize storage of small blocks
 - ❑ Utilize entire drive pool for every volume
- Enables:**
- ❑ Performance virtualization
 - ❑ De-dupe and compression without tradeoffs
 - ❑ HA without performance impact

SolidFire



- ❑ 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
- ❑ Future direction of cloud storage is clear:
 - ❑ Disk for capacity oriented, secondary storage
 - ❑ Flash / SSD for all primary storage