NetApp

Go further, faster®

Multi-Tier Subsystem Management Using SLOs

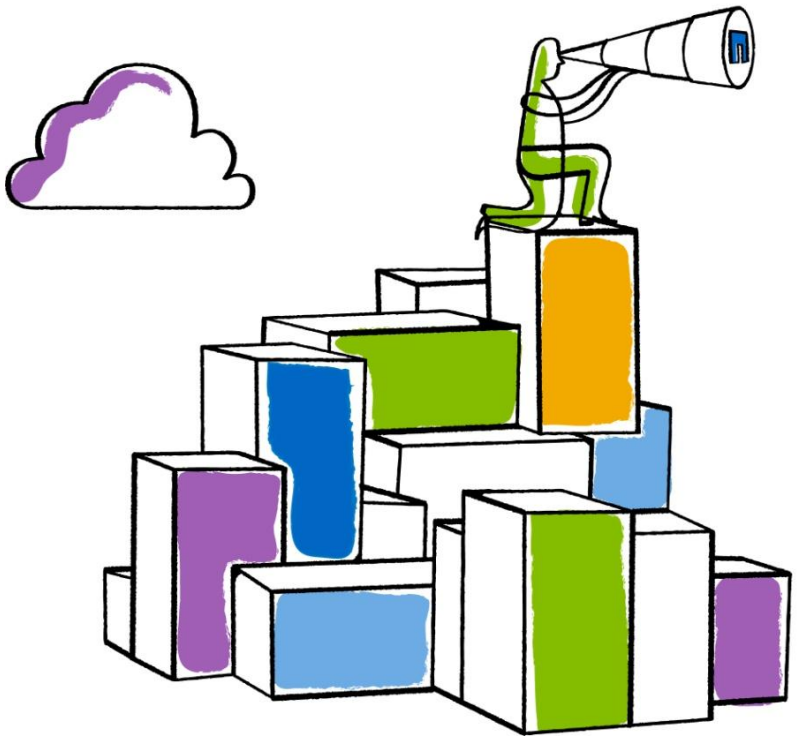

Kaladhar Voruganti

Technical Director, CTO Office

NetApp, Sunnyvale

January 29th, 2013

Talk Outline

- Part 1: Discuss Storage Architecture Trends due to NVM
- Part 2: Impact of the above trends on Storage Management Paradigm (Need for management by SLOs)
- Part 3: How can Standards Help?

Talk Outline

- Part 1: Discuss Storage Architecture Trends due to NVM
- Part 2: Impact of the above trends on Storage Management Paradigm (Need for management by SLOs)
- Part 3: How can Standards Help?

Trend 1: NVM Accelerating Multi-Layered Storage Architectures

Traditional Network Storage

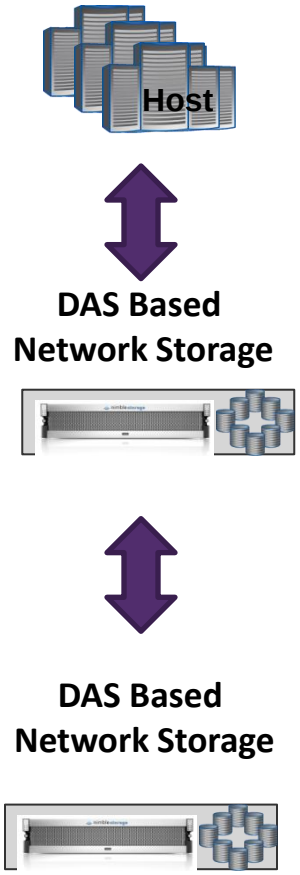

Tries to Optimize For Both Capacity & IOPs

IOPs Value

Capacity Value

Emerging World

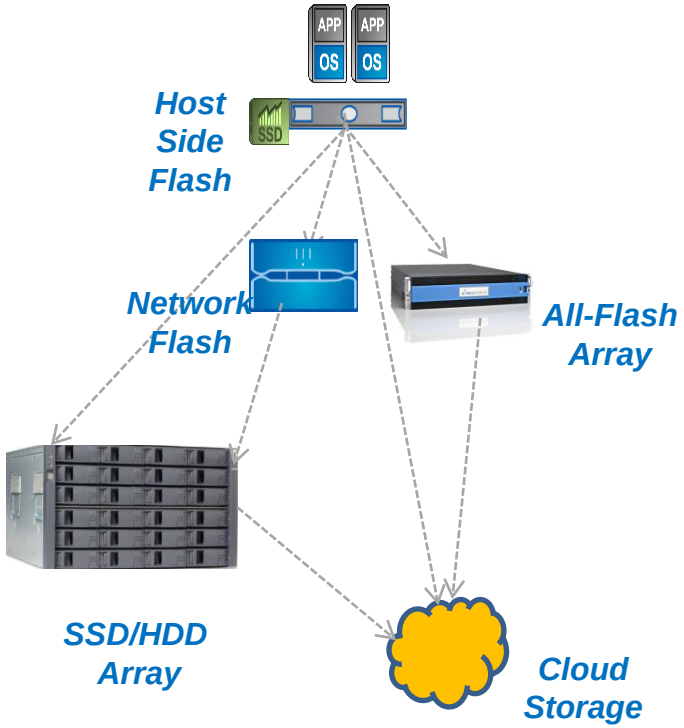

5

NVM Solution Alternatives

Category	Solutions
Server Flash Cache (Dedicated/Clustered)	**NetApp Flash Accel**, Fusion-IO/IO-Turbine, LSI, VFCache, Flashsoft,EMC Lightning
Network Flash Cache (Near-server / in-network)	**NetApp CacheIQ** EMC Thunder,ION, Avere, GridIron
All-Flash Array	Violin WhipTail XtremIO (EMC)
Storage Flash Cache/Tier	**NetApp Flash Cache** **NetApp Flash Pools,** EMC FAST, HDS

Trend 2: NVM Enabling Load/Store Persistence Programming Model

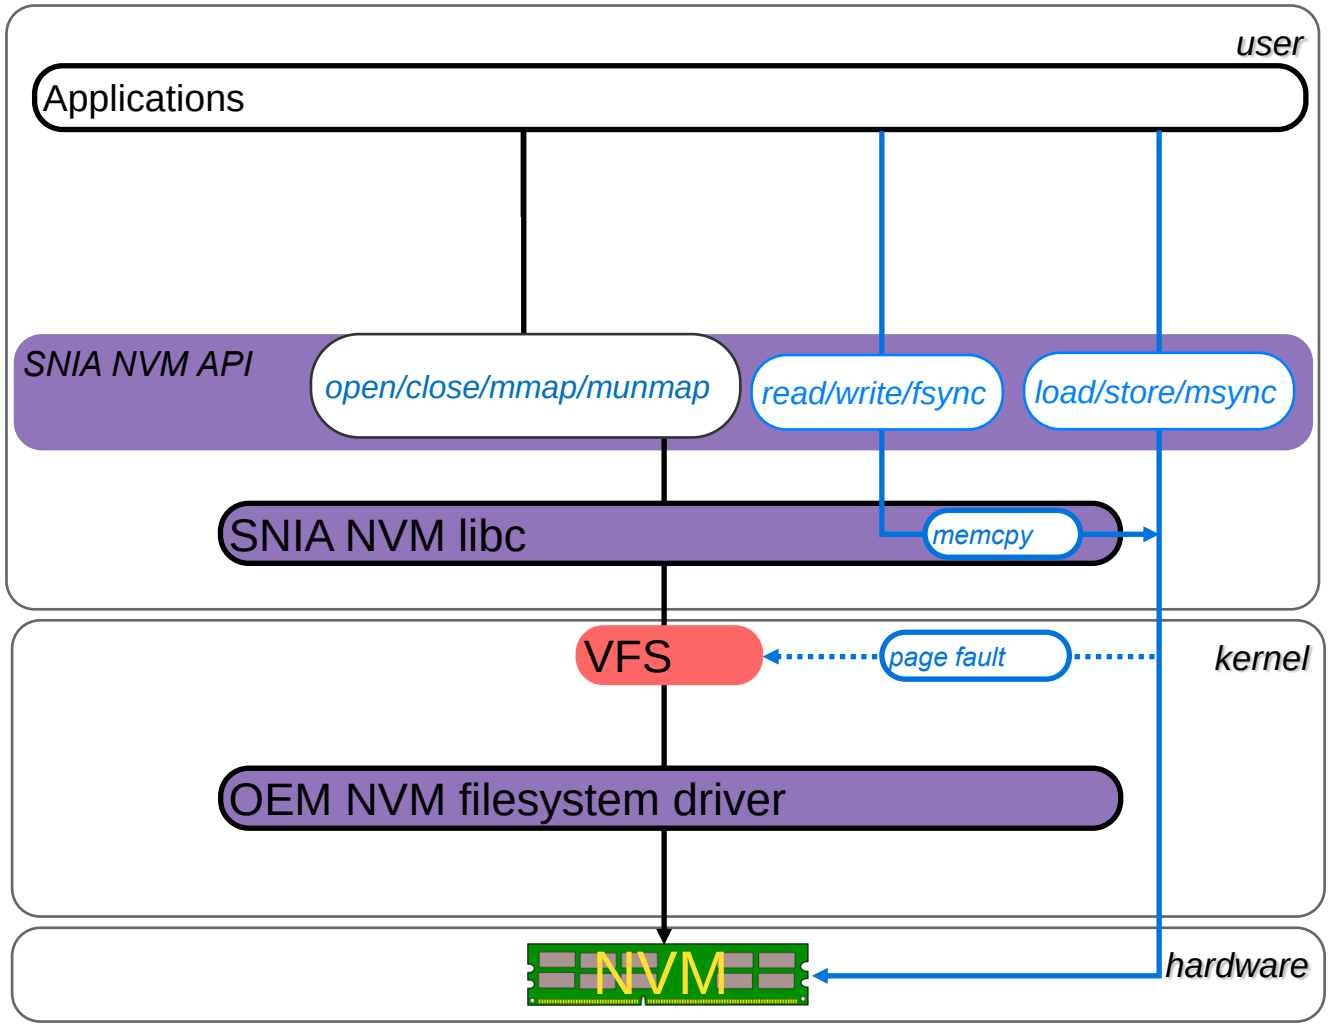

Talk Outline

- Part 1: Discuss Storage Architecture Trends due to NVM
- Part 2: Impact of the above trends on Storage Management Paradigm (Need for management by SLOs)
- Part 3: How can Standards Help?

Impact 1: Service Offered by Multiple Layers and Vendors

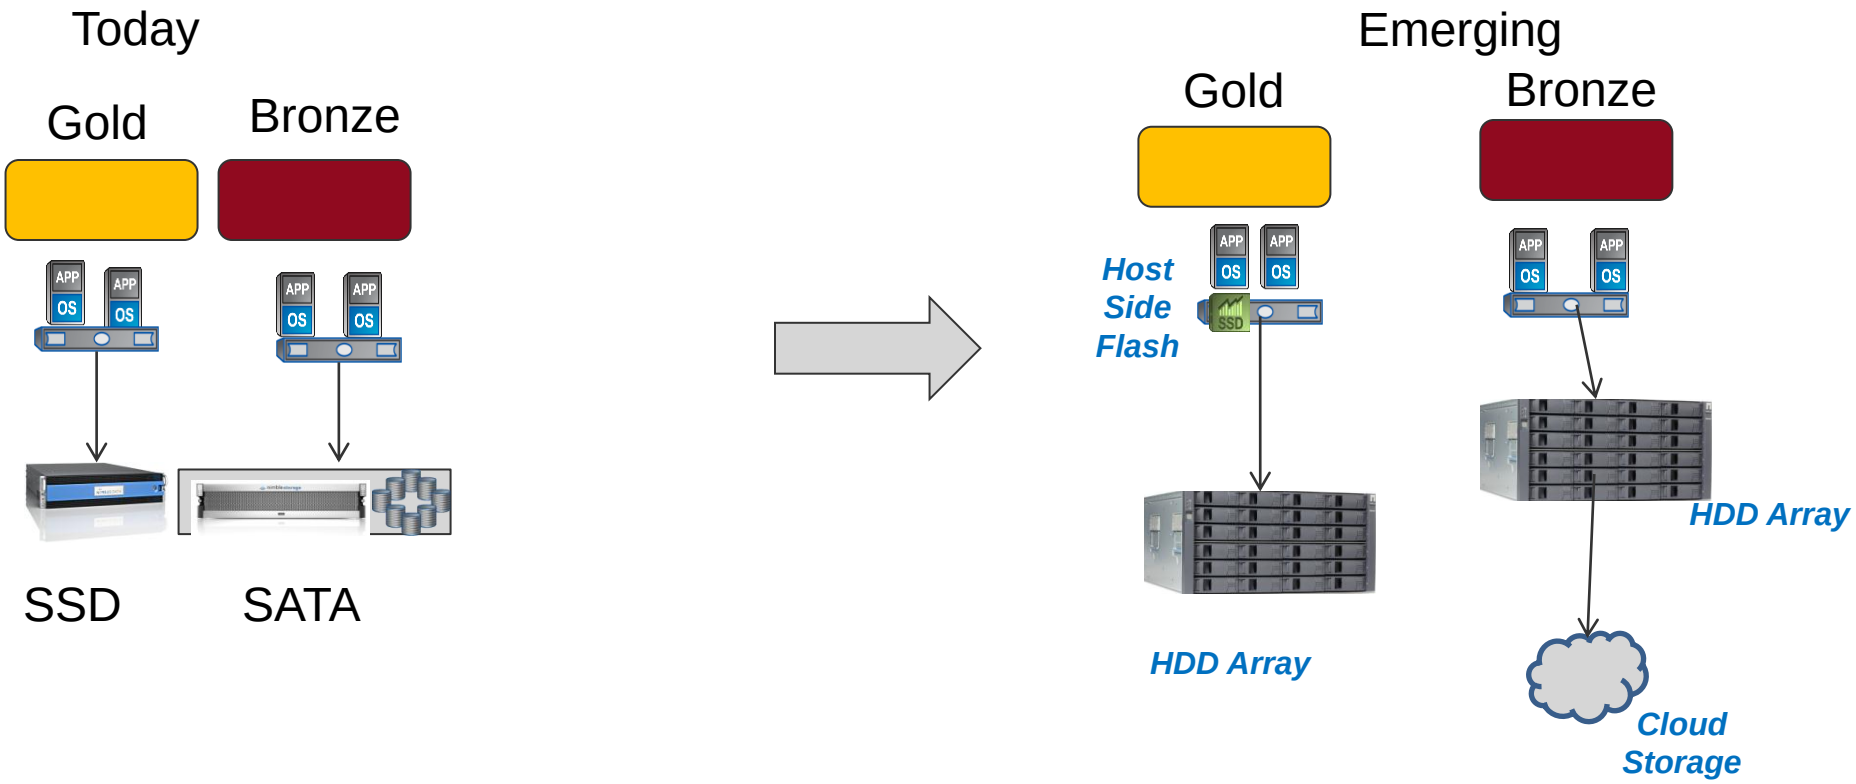

Service Offered By

- Single Layer
- Single Vendor

Service Offered By a combination of

- Multiple Layers
- Multiple Vendors

Impact 2: Service Level To Configuration Binding is Not Fixed

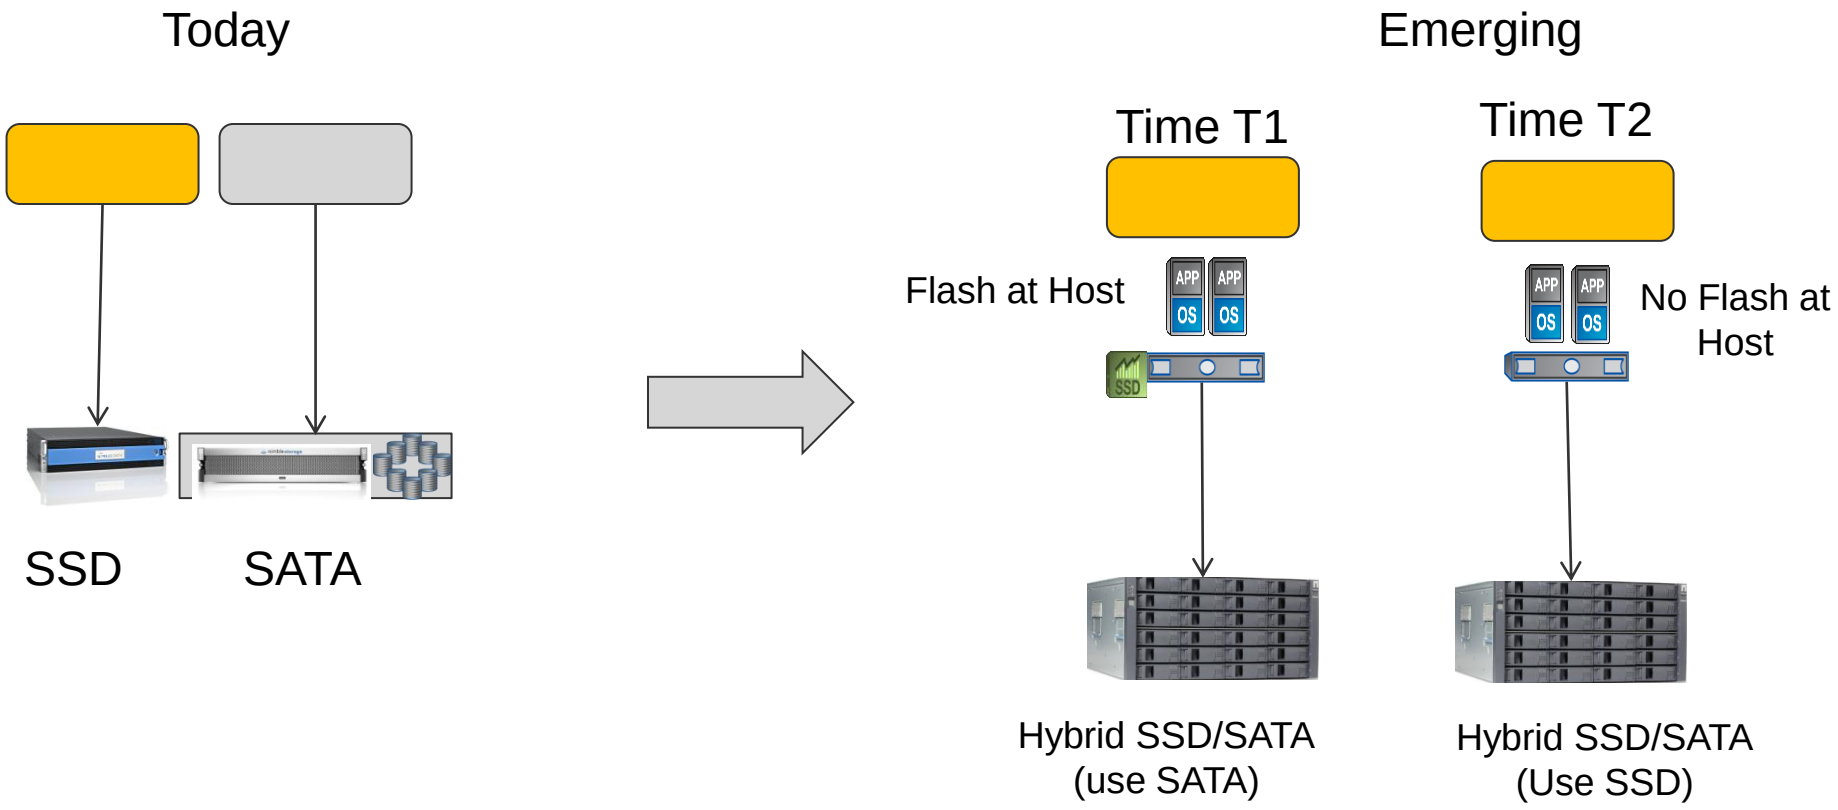

- Fixed Binding between Configuration and Service Level

- Dynamic Binding between Configuration and Service Level for Better resource utilization

Impact 3: In band Policy Binding

Today

Volume
Created using
Gold Service (policy)

Map file/LUN to
Gold Volume

Out of Band
Policy Mapping

Application
reads/writes
to the file/LUN that
is bound a priori to Gold Storage

Application does not specify Policy
In band

Emerging
NVM Load/Store Model

Application
Asks for Gold Policy
NVM region

malloc(size,GoldService);

In Band
Policy Mapping

SNIA NVM API library
Requests
Kernel for Gold NVM region

Software Layer on top of SNIA NVM API Layer
Provides Gold
Level Semantics
To Gold NVM region

Application Specifies Policy
In band

Impact 4: Data Structure Level Policies

Today

Volume
Created using
Gold Service (policy)

Map file/LUN to
Gold Volume

Application
reads/writes
to the file/LUN that
is bound a priori to Gold Storage

File, Object, LUN, Volume
Level Policy Management

Emerging
NVM Load/Store Model

Application
Asks for Gold Policy
For data structure

malloc(size,GoldService);
new(KVStore, GoldService);

Software Layer on top of SNIA NVM API Layer
Provides Gold
Level Semantics
To Gold Data Structure
In Gold NVM Region

Data Structure Level
Policy Management

Talk Outline

- Part 1: Discuss Storage Architecture Trends due to NVM
- Part 2: Impact of the above trends on Storage Management Paradigm (Need for management by SLOs)
- Part 3: How can Standards Help?

Following Standards Efforts are Required

- Service Definition using SLOs needs to be standardized
  - For coordination between different layers
  - For policy specification in Load/Store access model
- Inter-Layer Control Protocol needs to be standardized
- Inter-Layer I/O Protocol needs to be standardized

SLO Standards for Service Definition

SLO Dimension	Old way of describing Service Attributes using Vendor technology attributes	SLO Way of describing Service Attributes
Performance	SATA, SSD etc	Latency Ceiling Latency Floor IOPs or Bandwidth Ceiling IOPs or Bandwidth Floor
Operational Recovery	SnapVault, SnapMirror,NDMP	RPO Retention Time Backup Mode
Disaster Recovery	SnapMirror Async, MetroCluster	RPO Distance Replication Mode
Protection (Availability)	RAID-DP, HA	Media Failure Protection Shelf Failure Protection Head Failure Protection
Space Guarantee	Dedup, Compression, Fractional Reserve etc	Thin Provisioning Thick Provisioning

- Services Described using Quantitative, Vendor Technology Independent SLO Attributes
-
- Service Definition can encapsulate multiple layers managed by different Vendors
- We need to standardize both SLO dimensions and SLO attributes for Services

Control Protocol Standards Needed

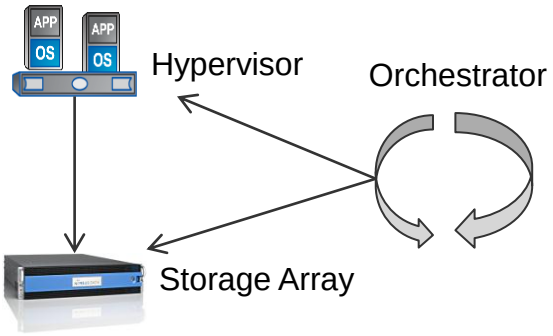

Coordinate Control Actions Between Layers

Migration	Workload Throttling	Caching	Encryption
Which Layer Should Migrate data During SLO violation	Which layer should throttle Workload during SLO violation	Coordinate Cache Sizing across layers	Which layer should encrypt data

- Control Protocol works either directly between different vendor layers or between an Orchestrator and the different vendor layers
- Is SMI-S the right place to work on this standard, or where do we standardize this?
- Need a Common Cost Model for Coordination amongst layers

I/O Protocol Coordination Standards Needed for the following

- Sharing of QoS Class (Gold, Silver, Bronze etc) Information across Layers (tag each I/O)
- Share Caching hints across layers
  - E.g. NFS 4.2 fadvise(), SCSI Disable Page Out
- Share Encryption/Dedup Hashes across layers to get the best of storage efficiency and encryption
- Sharing storage efficiency info (e.g. a priori knowledge that there is a clone relationship and so prepopulate the sharing map at the host)

Conclusion

- NVMs are accelerating the bifurcation of Capacity and IOPs optimized storage architectures
- Need to take a holistic end to end data management view to manage both IOPs and Capacity layers
- SLO Based Management Notions hide the multi-layer, multi-vendor issues from storage subscribers
- Standards have a key role to play for service definition, control and I/O interactions between layers

Acknowledgements

- Steve Byan
- David Dale
- Fred Knight
- Madalin Mihailescu
- Bala Ramachandran
- Doug Santry