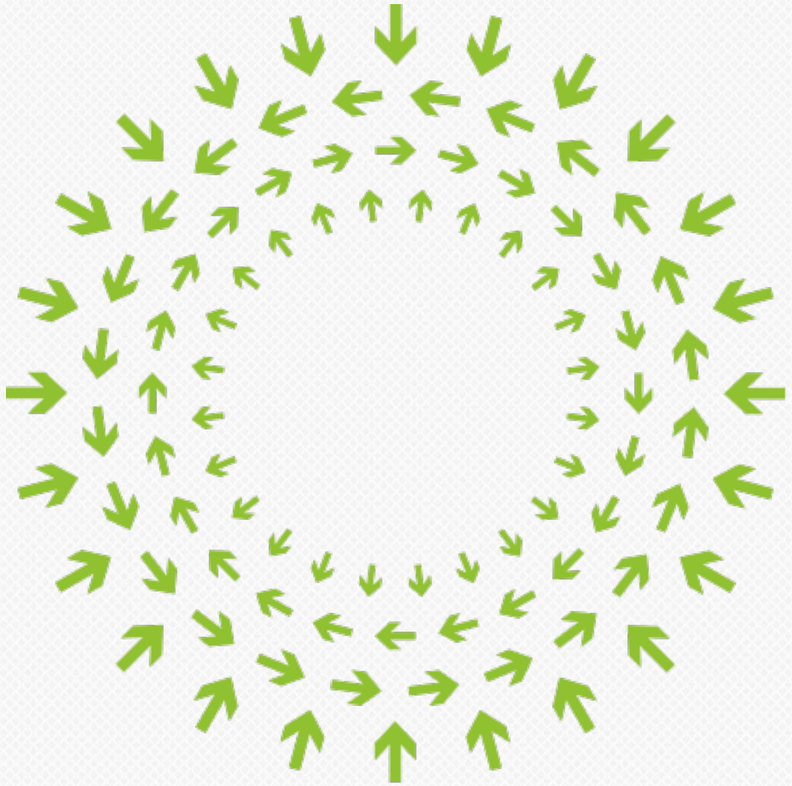

OPEN Compute Summit

March 10–11, 2015
San Jose

OPEN Compute Summit

A New World of Big Data Flash

Open Disaggregation

Ravi Swaminathan
SanDisk Corporation
VP & GM, SanDisk Systems & Software Solutions

Challenges & Considerations

Traditional Approaches Fail

- Over design SW for HDD failure
- Forced massive redundancy
- Poor performance
- New steps not addressing

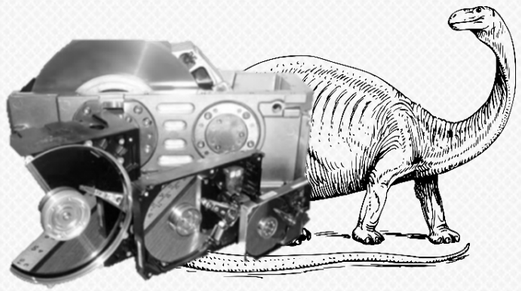

Performance with Capacity

- Forced trade-offs and high complexity
- Billions of users, objects & immediate gratification required
- Your cold data is my hot data

Costs & Complexity at Scale

- Look to DISAGGREGATION of compute and storage
- Simplification is paramount
- At massive scale costs and complexity increase exponentially

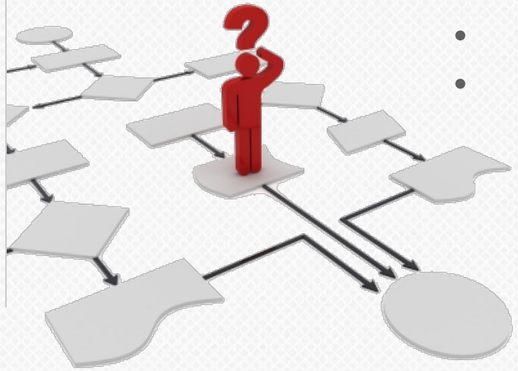

Big Data NDO Scale-out

- File systems break at massive scale
- Object Storage optimal efficiency
- Multi-protocol access to a common storage substrate: Block, Object and File interfaces

Flash excels in 4 areas to transform the data center

Huge Opportunity with Flash for Big Data

CONTENT REPOSITORIES

BIG DATA ANALYTICS

MEDIA SERVICES

Key Workloads that Require Both Performance & Capacity

SanDisk InfiniFlash™

Removing Barriers to Innovation @Scale

Architected for Massive Scale and QoS

- Scale-out, High Density, All-Flash appliance
- Designed for maximum efficiency
- New Flash form factor not SSD based
- Compelling TCO with space and energy savings

Leaving Traditional Approaches Behind

- 512 TBs of Flash in 3U
- Up to 1M IOPS, <1ms Latency, 7GB/s Throughput
- 5x the density, 50x the performance, 5x lower power

Breaking Price Barriers

- $1/GB for InfiniFlash hardware with API's
- $2/GB for InfiniFlash with Scale-out or Scale up SW

- Innovative Big Data Flash Platform
- Breakthrough Economics @ Scale

SanDisk®

Systems and Software Solutions
for Infinite Possibilities

For questions http://www.sandisk.com/bigdataflash