Is The World Ready for ExaScale Flash?

Gary Lyng
Senior Director, Product Management & Marketing
SanDisk Corporation

Or Are We Already There?

Perspective: Market Rapidly Changing

connected, consolidation, convergence

Data creation & *In Motion*

In 2020: 30B interconnected things, effecting every IT vendor*

Data creation & At Rest

Google amazon intuit YAHOO! ebay Microsoft LinkedIn rackspace facebook awsmarketplace

154 Exabyte's by 2016*

50% of IT Spend & 100% of IT growth driven by* ..Mobile, Social, Cloud, Big Data /Analytics

*IDC 3rd Platform & Predictions, Gartner Nexus Predictions

Traditional Technologies Simply Fail!

HDD challenges

- Must Design for Failure & Inefficiencies
- Forces massive redundancy in operations
- Creates more copies of data than needed
- Slows response times

HDD direction not addressing the problem

- Larger disks, longer the rebuild time
- Cloud drives, spin-down – poor performance
- High redundancy still required

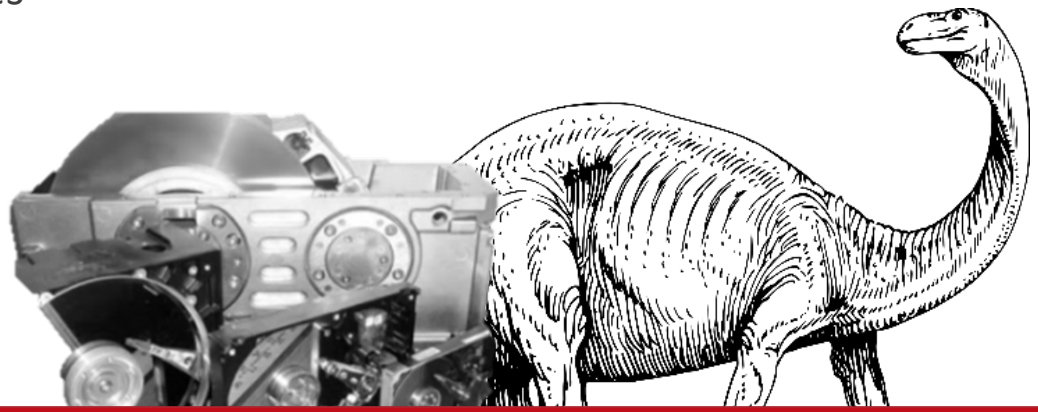

It's Time to Shift to a New Platform that Accelerates Innovation

Petabyte Capacity Flash is Inevitable

Challenges

- Affordability & Cost of Ops
- Reliability
- Complexity creep
- Efficiency
- Quality of Service
- Works at Scale

Capacity & Performance

HDD Arrays

Augment with Basic Caching

The Hybrid

Performance Specific use cases

Market Drivers–All Flash Arrays

- $ per IOPS and latency
- Comparable $ per GB to HDD
- Managed services and QoS
- $ per rack density
- $ per WATT
- $ Power utilization efficiency

ExaScale Requires..

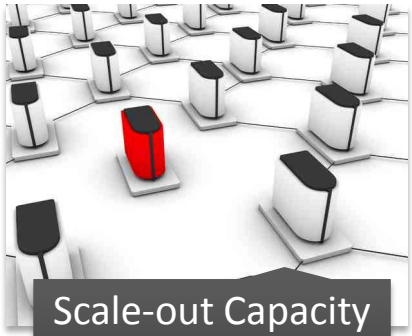

Scale-out Capacity
+ High Density

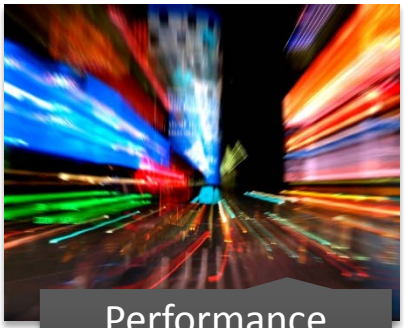

Performance
On-Demand

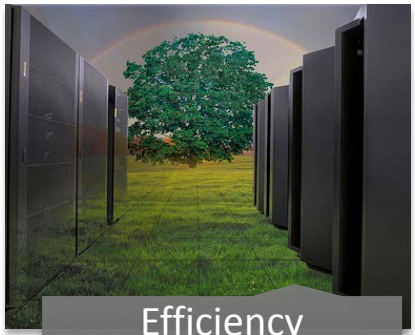

Efficiency
Low operating Cost
(low power, low cooling)

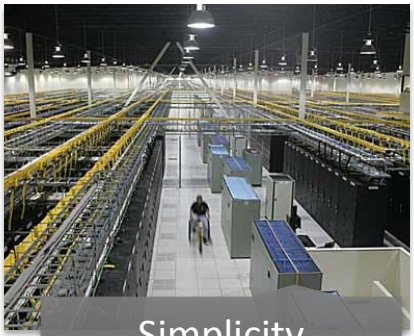

Simplicity
(Self-serve simple administration)

.. with Quality of Service

We Can Build it,
So What's it good for?

Big Data Flash for ExaScale Workloads

Content Repository

- Large containers for long periods with on-demand rapid access
- Mixed Media container, Active-archiving, Backup, Locality of data

Big Data Analytics

- Hadoop, NoSQL
- Massive Read Intensive operations
- Time-to-Value & Time-to-Insight

Media Streaming

- High read intensive access from billions of edge devices
- Hi-def video driving even greater demand for Capacity + Performance

Building Solutions for Massive Data

Software-Defined Storage & Open Source

- Multi-protocol access to a common storage substrate
  - Block, Object and File interfaces
- Configurable Data Protection level
  - Local synchronous replication
  - Remove asynchronous replication
  - Erasure Coding
  - Natively understands physical failure domains
- No single point of failure
- Supports heterogeneous HW nodes

- By 2017 30% of all storage will be under openstack & 70% of all storage deployed will be on standard x86 hardware*
- openstack can be the API for the Data Center operating system
  - Vibrant community of contributors, supporters, consultants, add-ons, …
  - Wide-spread adoption, and at very large scale (10,000+ nodes)
  - Expanding beyond service providers into enterprises

* Gartner,

Perspective for Consideration

1) Cloud & Hyperscale Applications driving unprecedented data growth in both Structured and Unstructured data, demands a new approach to storage -- ***Everything as a service with a service level objective***

2) Cost-optimize Capacity, Reliability and Performance
   - Favor scale-out over scale-up techniques
   - Simple, regular infrastructure leverages volume purchasing,

3) Must consider Real-World TCO not just $/GB - **The Math is Compelling**

4) Open-Source Software often has the highest rate of innovation, ride-it & contribute.

5) It is Your ***Right & Your Applications Right*** to have unprecedented **QoS @Scale**

Flash Memory SUMMIT

Visit SanDisk Booth #204

Thank You & Enjoy the Show!

Questions?:
gary.lyng@sandisk.com

#BigDataFlash

linkedin.com/groups/Big-Data-Flash-6727304

Check out the SanDisk Enterprise IT blog @ http://itblog.sandisk.com/

SanDisk