

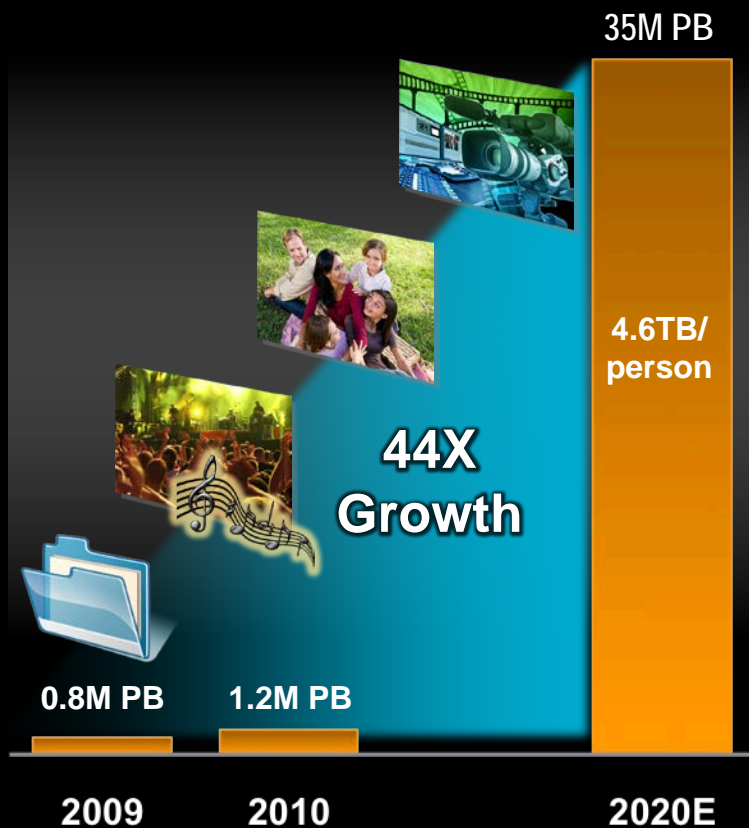
Accelerating Big Data Analytics: The Role of Intelligent Flash Solutions in the Data Deluge

Steve Fingerhut, Vice President of Marketing, Accelerated Solutions Division
LSI Corporation



Demand for Storage Keeps Growing

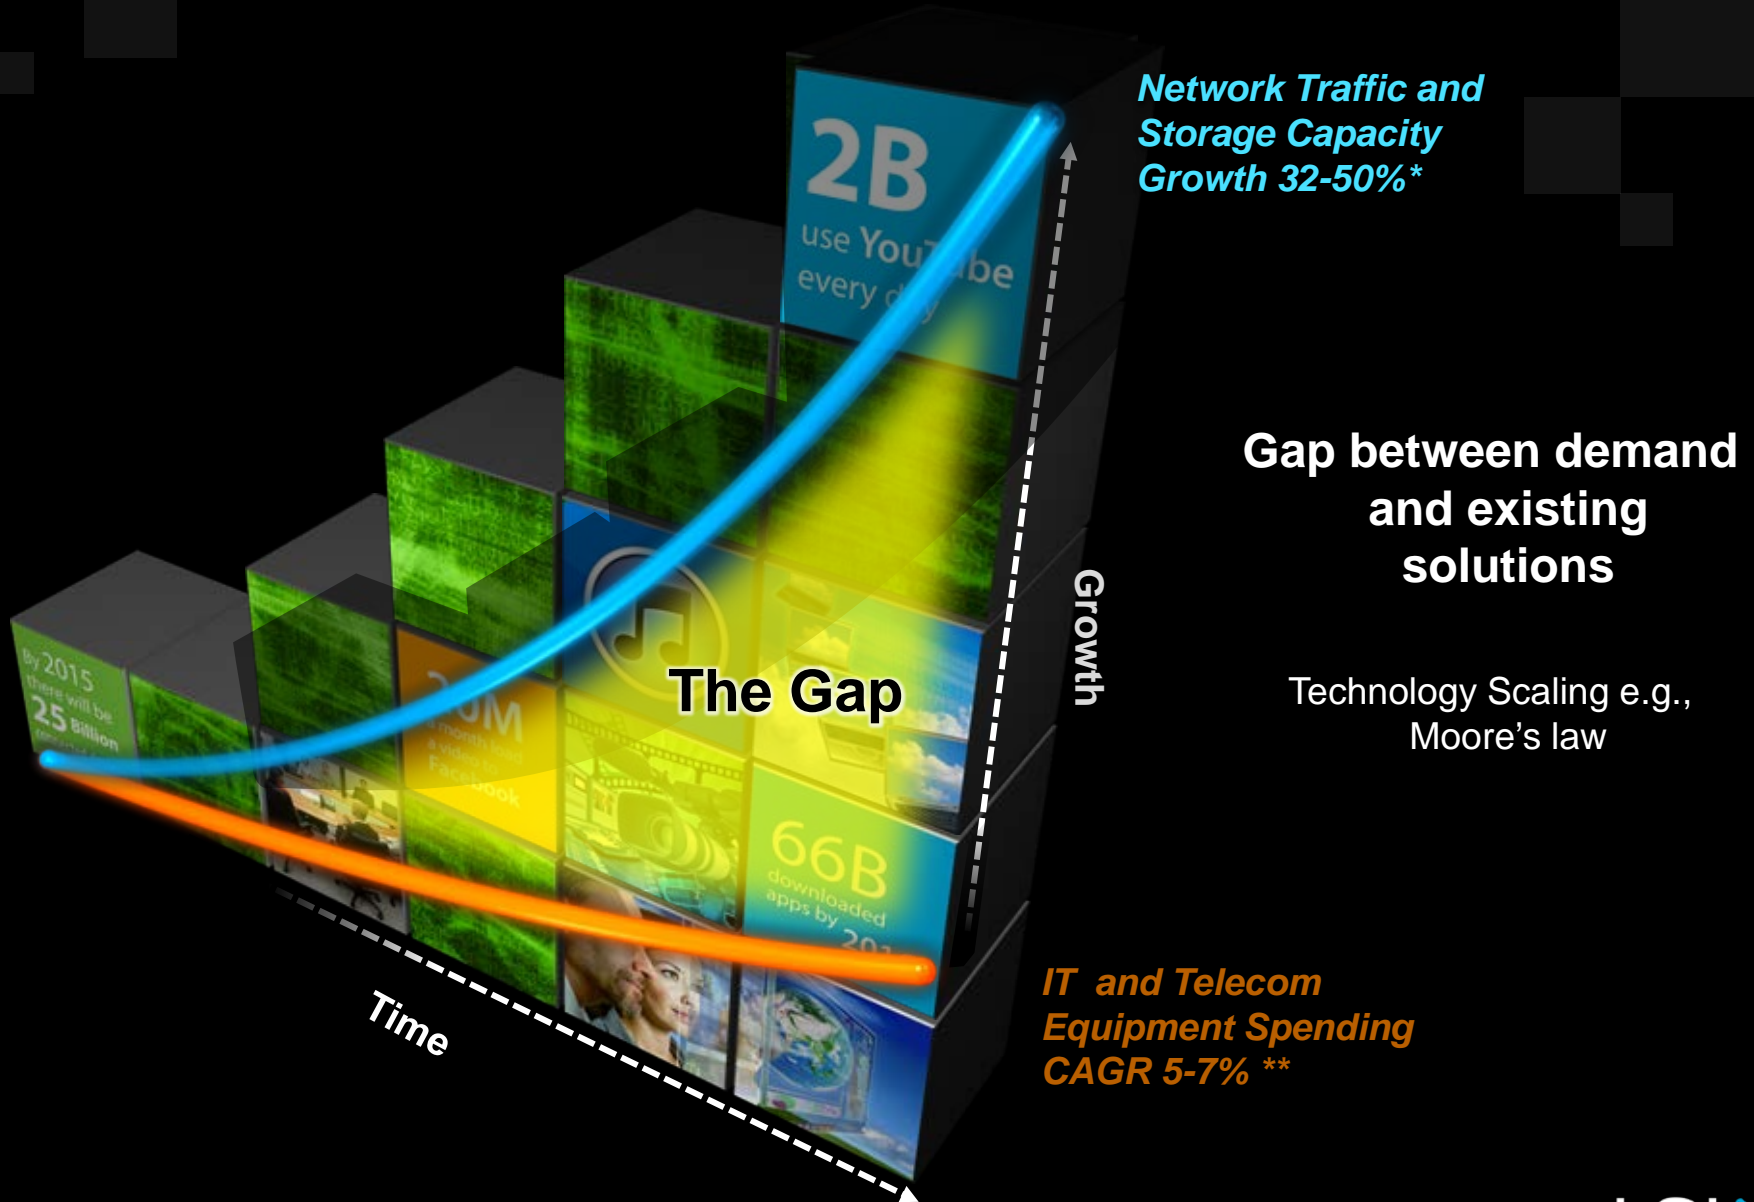
Demand for Storage Exploding



Source: IDC, "The Digital Universe Decade – Are You Ready?" May 2010

- Datacenter managers report that application performance not adequate
- Hard drive capacity gains slower than processor performance gains
- Storage performance is a significant bottleneck
- I/O delays cost money and time
- Traditional approaches are costly half-measures

The Result: The Data Deluge Gap



Closing the Gap with Intelligent Flash Solutions

System Requirements

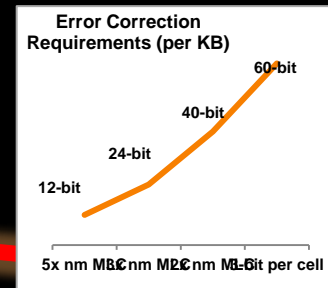



- Increasing Writes
- Higher Reliability
- More Performance

- All flash is different and changes every 18 months!
 - Per bit cost dropping 30% YoY
 - Endurance, Reliability, Performance, Interface




Flash Cell Geometries



- Less Writes
- Lower Reliability
- 30% Cost Reduction Y-Y

Time

Computing Applications: The Next Big Driver of Flash



- Historically Consumer driven, but Compute is next big driver
 - Compute is much more difficult!
- More Flash will be used in Compute in 2014 than all Flash applications last year
- LSI Flash Storage Processors managed 20% of Compute capacity in 2011

LSI Nytro™ Application Acceleration Architecture

Nytro WarpDrive

Nytro XD
Nytro WarpDrive + XD Caching SW


Nytro MegaRAID

Application



I/O

Operating System

Nytro XD Software



SAS • FC • iSCSI



Intelligent Caching Stores
Stores **Hot Data**

DAS

App Data **Hot Data** Cold Data

LSI Product

THE **Smarter** WAY
TO FASTER™



Storage. Networking. Accelerated.™

The Coming Evolution of Enterprise Flash Solutions

2015 TAM

\$500M

\$1.4B

Additional
Growth
Opportunity

Flash Segment

High capacity, PCIe flash
for primary storage

High Capacity flash with
Seamless Caching /
Acceleration Software

Low capacity flash with
Seamless Caching /
Acceleration Software

Nytro
WarpDrive



Nytro XD



Nytro
MegaRAID



LSI has products and
technologies to address
all segments

Source: IDC 2012

LSI Nytro WarpDrive – Accelerating Analytics



Case Study

The Challenge:

Greenpower, a leading direct marketer of nutritional supplements, needed to improve the performance of SAS analytics beyond what SSDs could deliver.

The Solution:

1600GB LSI Nytro WarpDrive

The Results:

“By adding the Nytro WarpDrive card, we were able to **improve the performance by 50%** for the analytical runs compared to what we were able to get from our existing SSDs.”

- Guido Meijers, System Engineer at Greenpower



Before

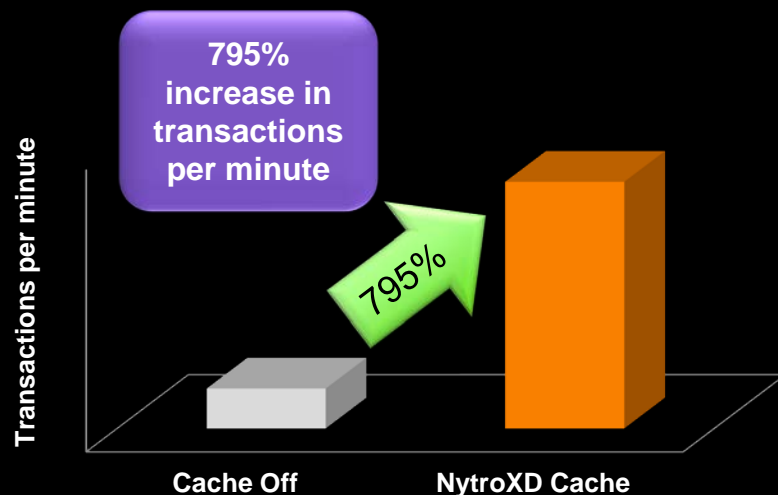
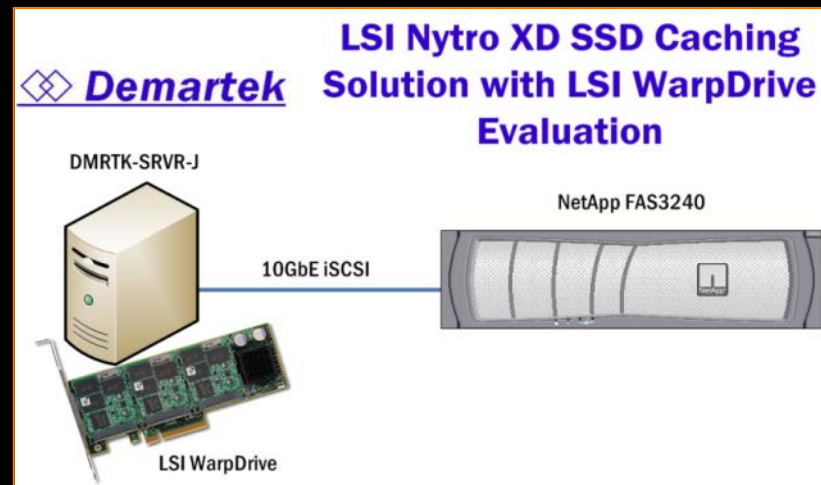


After




LSI Nytro XD - Demartek's Evaluation Highlights

- Decision support system analytics query test
- LSI Nytro XD Solution increased database performance by 795%.
- No change in application or back-end storage system was required.
 - Snapshot /restore of array done during testing to verify server side cache didn't impact SAN storage features




Demartek suggests Nytro XD as an excellent choice for server-side caching solution

LSI Nytro MegaRAID – 75 GB TeraSort Run



18 Disks connected via Advanced RAID controller (presenting RAID 0 VDs): 19 Mins.



With Nytro MegaRAID Caching enabled: 12 Mins. = **37% run time reduction!**

Hadoop job_201208081235_0004 on [MegaraidCI2](#)

```

Job Name: TeraSort
Job File: hdfs://MegaraidCI2:9000/mnt/hdfs1/tmp/mnt/hdfs2/tmp/mnt/hdfs3/tmp=
./mnt/hdfs4/tmp/mnt/hdfs5/tmp/mnt/hdfs6/tmp/mnt/hdfs7/tmp/mapred/stagi=
ng/root/staging/job_201208081235_0004/job.xml
Submit=20 Host: MegaraidCI2
Submit Host Address:=20 172.11.11.114
Job-ACLs: All users are allowed
Job = Setup:=20 Successful
Status: Succeeded
Started at: Wed Aug = 08=20 12:44:48 CDT 2012
Finished at: Wed Aug 08 13:04:09 CDT=20 2012
Finished in: 19mins, 20sec
Job Cleanup:=20 Successful
    
```

Hadoop job_201208081311_0001 on [MegaraidCI2](#)

```

User: root
Job Name: TeraSort
Job File:
hdfs://MegaraidCI2:9000/mnt/hdfs1/tmp/mnt/hdfs2/tmp/mnt/hdfs3/tmp/mnt/hdfs4/tmp/mnt/hdfs5/tmp/mnt/hdfs6/
Submit Host: MegaraidCI2
Submit Host Address: 172.11.11.114
Job-ACLs: All users are allowed
Job Setup: Successful
Status: Succeeded
Started at: Wed Aug 08 13:12:44 CDT 2012
Finished at: Wed Aug 08 13:24:55 CDT 2012
Finished in: 12mins, 11sec
Job Cleanup: Successful
    
```

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	Failed/Killed Task Attempts
map	100.00%	320	0	0	320	0	0 / 0



Storage. Networking. Accelerated.™