

A blurred city street at night with light trails and binary code overlaid. The background shows a city street with light trails from cars and buildings, overlaid with a stream of binary code (0s and 1s) in various colors (red, blue, yellow) moving across the frame.

# The Flash Transformed Data Center & the Unlimited Future of Flash

John Scaramuzzo

Sr. Vice President & General Manager, Enterprise Storage Solutions

Flash Memory Summit – 5-7 August 2014

**SanDisk®**





# Forward-Looking Statements

During our meeting today we will make forward-looking statements.

Any statement that refers to expectations, projections or other characterizations of future events or circumstances is a forward-looking statement, including those relating to market growth, industry trends, future products, product performance and product capabilities. This presentation also contains forward-looking statements attributed to third parties, which reflect their projections as of the date of issuance.

Actual results may differ materially from those expressed in these forward-looking statements due to a number of risks and uncertainties, including the factors detailed under the caption “Risk Factors” and elsewhere in the documents we file from time to time with the SEC, including our annual and quarterly reports.

We undertake no obligation to update these forward-looking statements, which speak only as of the date hereof or as of the date of issuance by a third party, as the case may be.








**100 hours** of  
video are  
uploaded to  
**YouTube every  
minute**

Multiple **private  
clouds** are already over  
**50 Petabytes** in  
size

**Analytics**  
market expected  
to grow to **\$16.9  
Billion** in 2015



**90% of the Data** in the  
world today was **created**  
within the last **2 years**

**Unstructured  
information is 90%  
of Big Data** and is  
human information like  
e-mails, videos, tweets,  
Facebook posts...

BILLIONS OF USERS  
200 BILLION DEVICES



## 3<sup>RD</sup> PLATFORM

MOBILE DEVICES

MOBILE · CLOUD · BIG DATA · SOCIAL

MILLIONS OF APPS



**3<sup>rd</sup> Platform:**  
Big Data and  
Cloud-based Data  
Centers

HUNDREDS OF MILLIONS  
OF USERS



## 2<sup>ND</sup> PLATFORM

PC

LAN / INTERNET · CLIENT / SERVER

TENS OF THOUSANDS  
OF APPS



**2<sup>nd</sup> Platform:**  
Client/Server-based  
Data Centers

MILLIONS OF USERS



## 1<sup>ST</sup> PLATFORM

TERMINALS

MAINFRAME · MINI COMPUTER

THOUSANDS OF APPS



**1<sup>st</sup> Platform:**  
Mainframe-based  
Data Centers

Source: IDC, 2014. Directions "Key Battles — and Strategies — for Dominance on the 3rd Platform" Doc # DR2014\_GS2\_FG. Mar 2014

# Flash Enables the 3rd Platform Transformation

## 3<sup>rd</sup> Platform – Cloud/Internet of Things

### Private & Public Cloud

- Instant resource provisioning is the new expectation

### Service Delivery

- Instant access
- Memory-like speed performance

### In-Memory Computing

- Near real-time results needed by business
- Large data sets

## 2<sup>nd</sup> Platform – Client/Server

### Virtualization

- Performance challenge for shared storage infrastructure

### Infrastructure

- Need reduction in server and storage systems, power, cooling, and floor space

### New Math: \$/GB Replaced by TCA & TCO

- Total Cost of Acquisition (TCA) = drives + enclosures + power supplies + ...



# All Flash Zone



10K HDDs



15K HDDs

Too Slow




Short-stroked 15K HDDs



NEARLINE




Real Time Analytics



Stock Trade Optimization

Currency Exchange

Real Time Data Stream for Video Editing



High Speed Messaging



DRAM



DRAM

Too Expensive



DRAM



DRAM


# Flash Transforming External Storage

**Over 1.5 Million IOPS in a 4u storage array**

▶ **All Flash Aggregated PCIe Performance**

- Scalable: up to 32 nodes
- Dense: 4u configuration
- High throughput: >20GB/sec bandwidth
- Fast Response: <60 microsecond latency!
- Massive Capacity: >25TB!

**Need 5,600 short-stroked  
15K HDDs  
for Equal Performance**






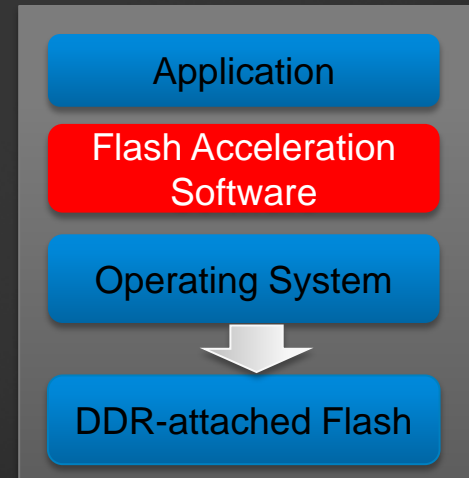
# Flash Transforming In-Memory Compute

Improve \$/transaction by up to 3x

- ▶ Use DDR-attached Flash plus optimized software
  - Increase performance of flash to near-memory like speeds
  - Same or greater capacity at a fraction of the cost
  - Large single server working set size to reduce cost and complexity



**1/2 the Servers,  
Faster Business  
Answers**



# Flash at the NEXUS



2<sup>ND</sup> PLATFORM

PC

LAN / INTERNET · CLIENT / SERVER



3<sup>RD</sup> PLATFORM

MOBILE DEVICES


MOBILE · CLOUD · BIG DATA · SOCIAL





# Density Matters


*Beyond just faster, cooler, and more power efficient SSD*



# Solutions are Already Less Costly than HDD Solutions

*More of the good stuff and less of the bad!*

Example: 50TB minimum of database storage needed



More power and more space needed

30X Performance same cost per GB

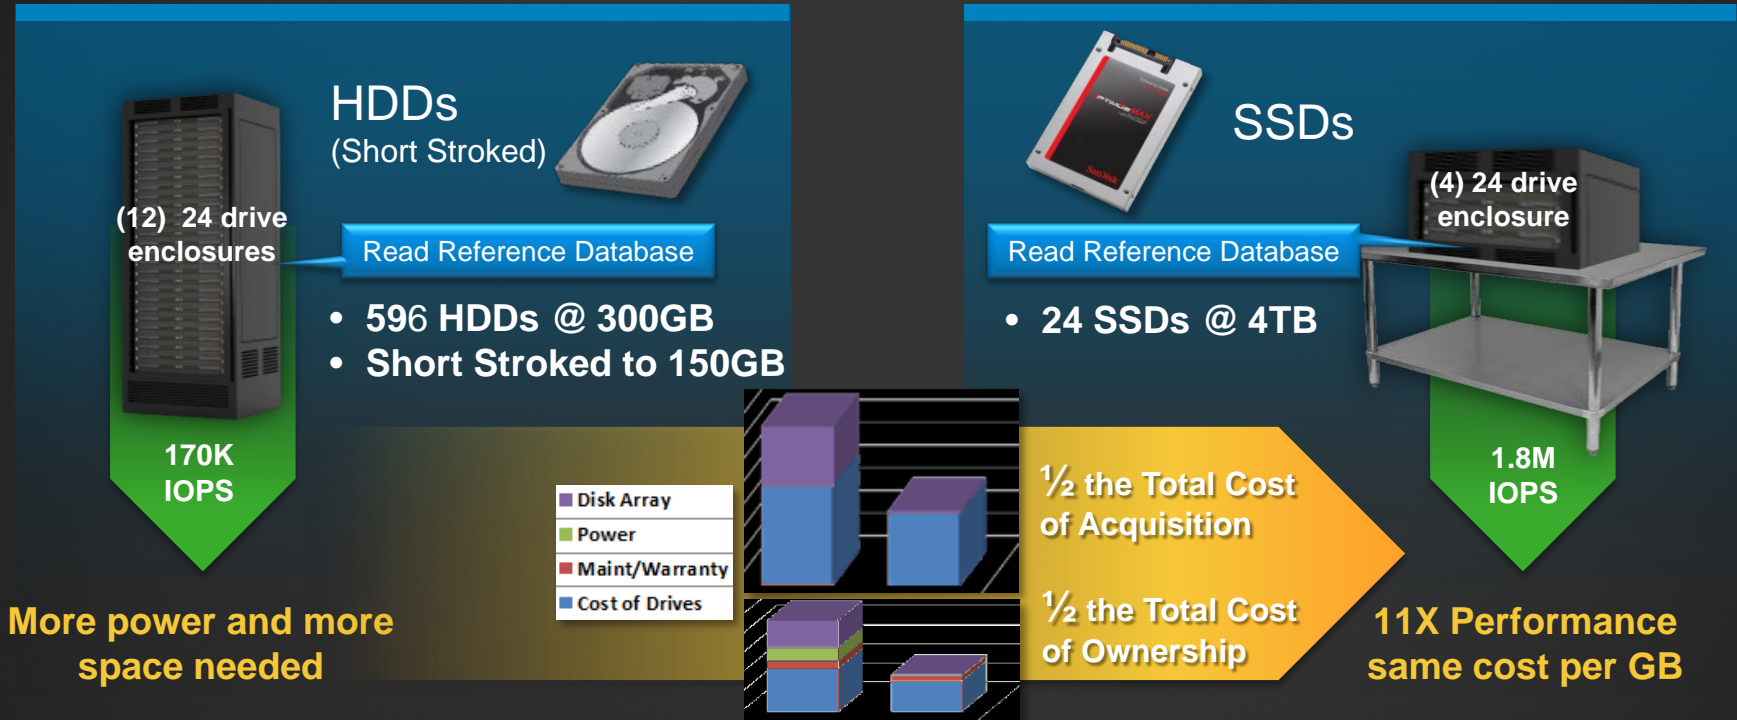
Source: Based on Dell PowerEdge pricing [www.dell.com/us/business/p/poweredge-r720/fs](http://www.dell.com/us/business/p/poweredge-r720/fs) along with HDD studies "Failure Trends in a Large Disk Drive Population" and "Disk failures in the real world": proceedings of 5th USENIX Conference on File and Storage Technologies and incorporated into TCO model developed by SanDisk architect Fritz Kruger—December 2013



# Solutions are Already Less Costly than HDD Solutions

*More of the good stuff and less of the bad!*

Example: 50TB minimum of 150K IOPS needed



Source: Based on Dell PowerEdge pricing [www.dell.com/us/business/p/poweredge-r720/fs](http://www.dell.com/us/business/p/poweredge-r720/fs) along with HDD studies "Failure Trends in a Large Disk Drive Population" and "Disk failures in the real world": proceedings of 5th USENIX Conference on File and Storage Technologies and incorporated into TCO model developed by SanDisk architect Fritz Kruger—December 2013

# Response Time Matters


## Active Content Repository System Example



**Read-Only Data**



Density



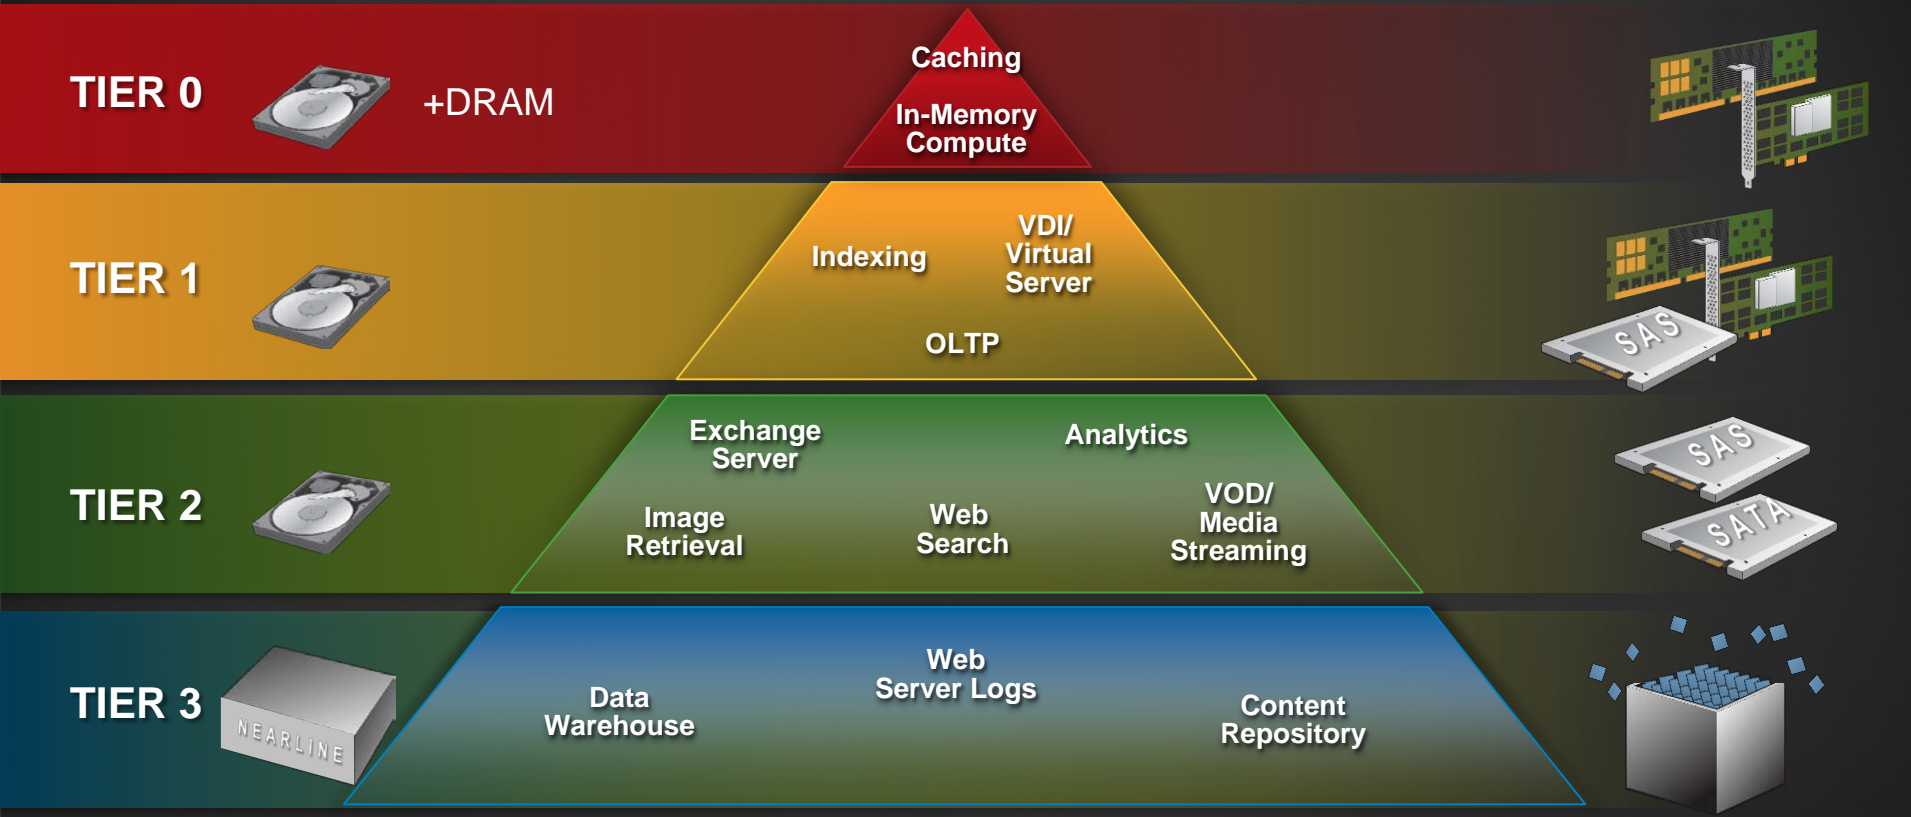
NAND  
Scaling



TCO



# Flash Transforms Applications Across the Tiers of Storage



# Flash Transforms Applications Across the Tiers of the Data Center


TIER 0  
New TIER 0



TIER 1  
New TIER 1



TIER 2  
New TIER 3  
TIER 3



Software  
Optimization



# Software Optimization Matters

**“Hadoop still too slow for real-time analysis applications?”**

-- *SearchBusinessAnalytics*

 **stackoverflow.com**

 **SearchDataCenter.com**



**“Cassandra speed. Why so slow?”**

-- *stackoverflow.com*

**“Google Analytics Code is Slowing Down My Site”**

-- *AnalyticsMarket blog*

**“Removing Memcached because it's too slow”**

-- *blog.serverdensity.com*

**“Turn to in-memory processing when performance matters”**

-- *SearchDataCenter*

 **serverfault**

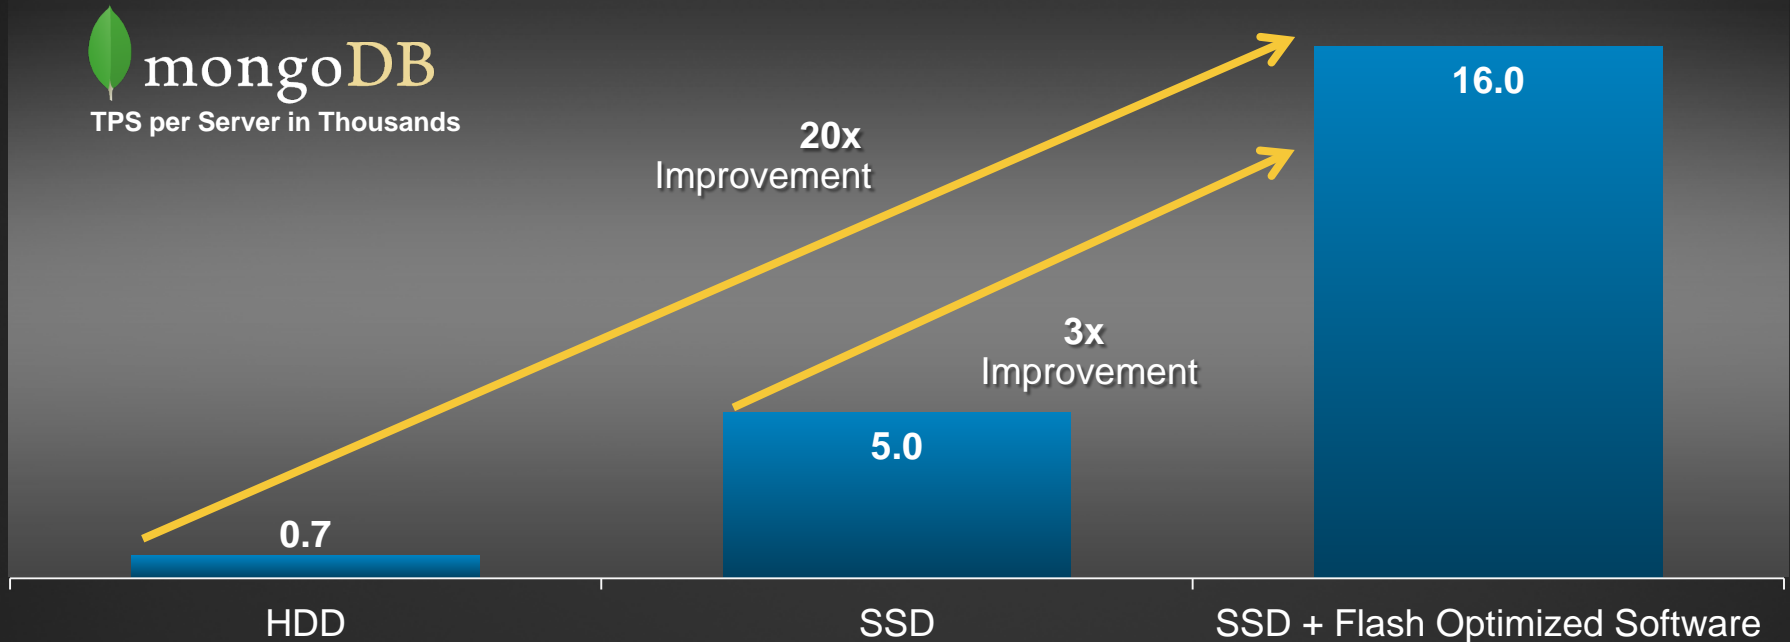
**“what affects Redis speed”**

-- *ServerFault.com*



# Software Optimization Matters

## Impact of Flash-Optimized Software

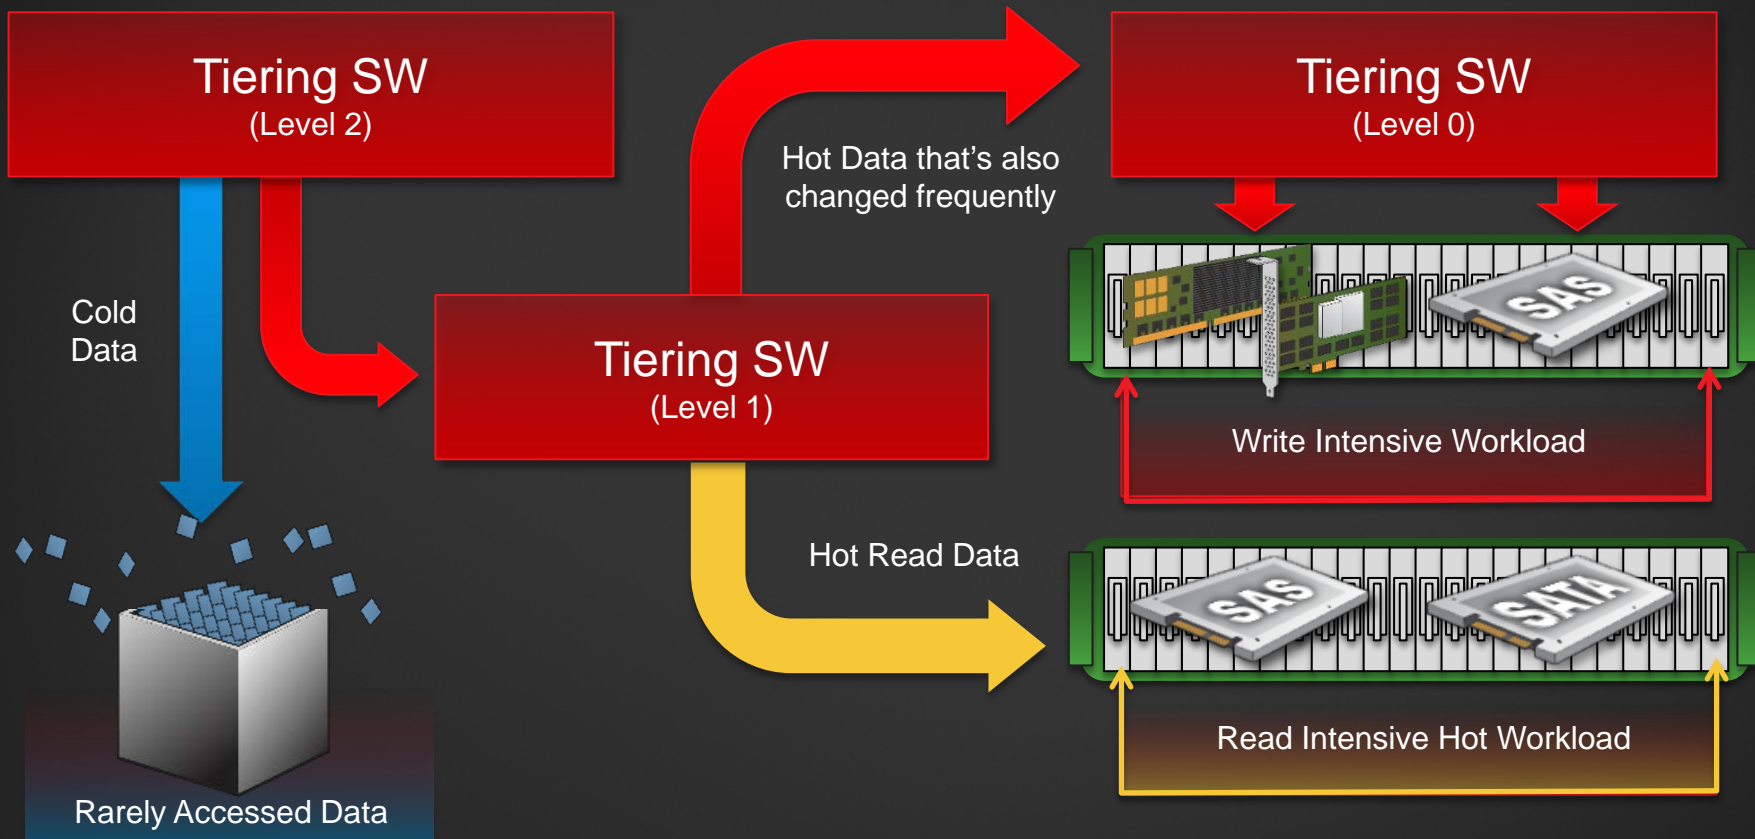


Critical to 3<sup>rd</sup> Platform: NoSQL, Hadoop, SQL, In-Memory Compute, Object Storage

128GB 1K data; 24 core Westmere, 96GB DRAM, 8 x 128GB SSD @ 1K data; YCSB measurements performed by SanDisk using ZetaScale Software



# Software is Required to Optimize the Flash Transformed Data Center



# Flash is Transforming our World

Enterprise Data Center



Internet of Things







# Thank you!

© 2014 SanDisk Corporation. All rights reserved. SanDisk is a trademark of SanDisk Corporation, registered in the United States and other countries. ULLtraDIMM is a trademark of SanDisk Enterprise IP LLC. Lightning is a U.S. registered trademark of SanDisk Enterprise IP, LLC. Cloudspeed, CloudSpeed Ascend, CloudSpeed Eco, CloudSpeed Extreme, CloudSpeed Ultra and Cloudspeed logo are trademarks of SanDisk Enterprise IP LLC. Optimus, Optimus Ascend, Optimus Eco, Optimus Extreme, Optimus Max and Optimus Ultra are trademarks of SanDisk Enterprise IP LLC. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).

**SanDisk®**