



Extending Datacenter Capacity To Drive ROI

Ajay Bhatia

Director, Product Marketing Management

Aug 6, 2014

Hyperscale Datacenters & SSD Adoption




2013
Enterprise SSDs
Over \$3B USD*




Performance ↑
Endurance ↑
Most critical levers

*Forward Insights

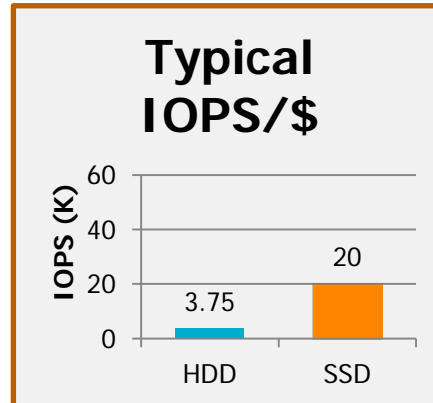
Flash in Datacenters



Flash is clearly faster than HDDs




However \$/GB alone is too high to replace all HDDs



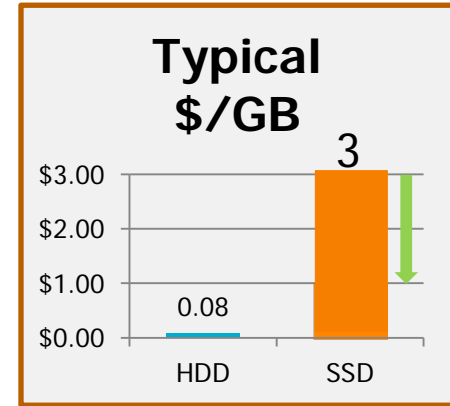
Although the performance tier has over 5x better performance/\$

Flash in Datacenters

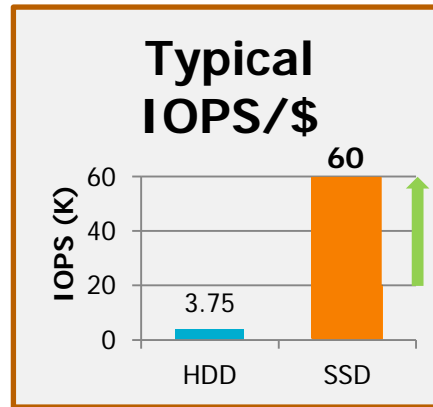


You get the same performance over HDDs

What if you could triple the capacity of the same flash?



The \$/GB drops to a third to replace all HDDs



And the performance tier is 16x higher per \$

Is this really feasible in an SSD?

Most data does not have 100% entropy and data reduction helps

Linux® Web hosting distributions Show ~60% entropy¹


Binaries and Shared Libraries



Database entropy as low as only ~30%²



Entropy by file type³



¹ LSI calculations on Oracle Database solutions

² http://www.thesmarterwaytoaster.com/pdf.php?c=LSI_WP_NyroWD-DuraWrite_040412

³ Measured by WinZip


Is this really feasible in an SSD?

Most data does not have 100% entropy and data reduction helps

ORACLE®



Entropy by file type³



¹ LSI calculations on Oracle Database solutions

² http://www.thesmarterwaytoaster.com/pdf.php?c=LSI_WP_NyroWD-DuraWrite_040412

³ Measured by WinZip



On-the-fly Compression Issues

Compression Engine Location	Issue	Solution / Concern
All Locations	<ul style="list-style-type: none">▪ Need some level of predictable entropy▪ Operating systems today don't understand variable max capacity	<ul style="list-style-type: none">▪ Datacenters now have more knowledge of their data▪ Customize drivers or modify already custom OS
CPU	<ul style="list-style-type: none">▪ Expensive resource and limited max throughput	<ul style="list-style-type: none">▪ Use a device-based solution
HDD	<ul style="list-style-type: none">▪ No variable mapping table; assumes 1:1 map from host	<ul style="list-style-type: none">▪ Major re-architecting required
SSD	<ul style="list-style-type: none">▪ FTL must be re-architected to support non-constant data size	<ul style="list-style-type: none">▪ Complex FTL must be created for data placement, garbage collection, etc.

This is not a simple problem to solve



What SSDs today support variable map sizes?

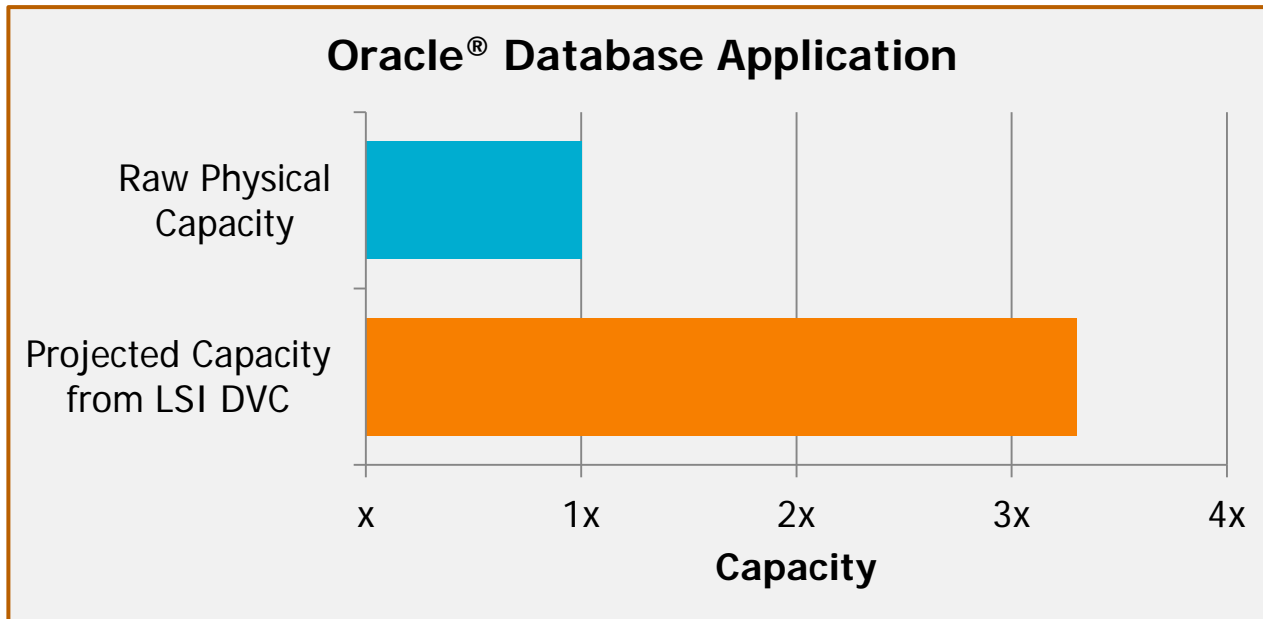
- LSI[®] SandForce[®] Flash Controllers employ DuraWrite[™] data reduction
 - Manages variable-size data placement on the flash

DuraWrite Feature	Benefit
Incorporates a novel Flash Translation Layer (FTL)	<ul style="list-style-type: none">▪ Handles variable-sized data for writing, reading, and garbage collection▪ Manages data spanning flash page boundaries to maximize space utilization
Supports variable data length	<ul style="list-style-type: none">▪ Incoming data size not restricted to be a power-of-two▪ Each LBAs can be a unique size
Provides for multiple flash page sizes to support variable code rate	<ul style="list-style-type: none">▪ Working in conjunction with SHIELD[™] technology



LSI DuraWrite™ Virtual Capacity (DVC)

- Increases the available storage capacity for typical data
- The lower the entropy, the higher the capacity increase
- Continue to get SSD endurance or performance benefits



DVC greatly lowers the \$/GB of available capacity to the user

Source: Projection based on LSI internal testing of Oracle database entropy. Results published in LSI white paper, "[Using DuraWrite™ Technology to Accelerates Flash Performance and Extend Write Lifetimes on the Nytro™ WarpDrive™ Application Acceleration Card](#)", 2012

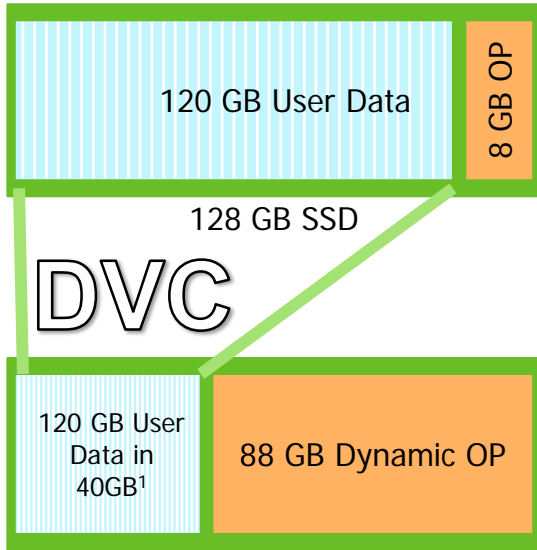


DVC Explained



Speed & Endurance

Enterprise Data Drive




Greater Dynamic Over Provisioning increases performance and endurance



Extra user data space

30% entropy database




Database actually sees 360GB of high performance SSD



Minimize costs

60% entropy boot drive



Able to use a lower cost SSD solution

¹Assumes a 33% entropy. Actual size reduction dependent upon data entropy.

²Assumes a 60% entropy. Actual size reduction dependent upon data entropy.



DVC Use Cases

Use Case	Native User Data Storage	Dynamic Caching
What	<ul style="list-style-type: none"> ▪ Additional flash space appears as a user accessible SSD 	<ul style="list-style-type: none"> ▪ Additional flash space managed by caching application
Pros	<ul style="list-style-type: none"> ▪ User capacity increases for same flash cost 	<ul style="list-style-type: none"> ▪ Increases cache size for same flash cost ▪ No OS changes
Cons	<ul style="list-style-type: none"> ▪ Requires OS driver changes <ul style="list-style-type: none"> – Easier with Linux or custom OS 	<ul style="list-style-type: none"> ▪ Requires integrated caching application
Support	<ul style="list-style-type: none"> ▪ LSI Nytro™ product line ▪ Select SandForce Driven® program SSDs 	<ul style="list-style-type: none"> ▪ LSI Nytro product line ▪ Select SandForce Driven program SSDs
Market	<ul style="list-style-type: none"> ▪ Enterprise with custom OS or Linux ▪ Client with Linux 	<ul style="list-style-type: none"> ▪ Enterprise focus



Summary

- Flash use in datacenters is expanding
- Compression can greatly reduce capacity & cost concerns
- Major complexities around on-the-fly compression
- LSI DuraWrite Virtual Capacity addresses those concerns
- Available in LSI Nytro product line & select SandForce Driven SSDs
- All-flash performance at hard disk price may not be a dream