



Unleashing Oracle Performance using Enterprise Flash at Scale

Brad Forcier

Technical Account Manager, Global Accounts
Hitachi Data Systems

Agenda

- ✓ Introduction – The Customer’s Dilemma
- ✓ Customer Asks – The Challenge
- ✓ The Solution – Purpose Built Intelligent NAND Flash device that is leveraged in many different use cases
- ✓ Results – Demonstrate at scale and time
- ✓ Q & A

Introduction – The Customer’s Dilemma

Fortune 500 company in highly competitive e-commerce landscape that requires the innovation of a startup, but has scale of a multi-national corporation. Constantly needs to re-invent itself while providing it’s customers with an “frictionless” online experience.

To stay relevant in this climate, the business needs to deliver

- ✓ Inspire customers beyond intent – multiple customer channels
- ✓ Personalize the user experience across the entire eco-system
- ✓ Continue to deliver a seamless yet secure experience

The Challenge

Performance and Scale

- ✓ Provide linear scale without compromise. Databases are at 70TB+ usable today, and growing rapidly.
- ✓ Mission Critical OLTP workloads required average read response times to be >1 ms.
- ✓ Mission Critical OLTP Oracle Redo log writes sync operations to consistently be serviced under 1ms to avoid session stacking.
- ✓ Handle highly unpredictable dynamic cache unfriendly workloads without performance degradation.

Availability


- ✓ Databases requires 99.99% availability uptime
- ✓ Minimal to zero impact from maintenance activities, hardware failures, and patching.

Agility

- ✓ Don't force me to change my database schema to optimize for your flash technology.
- ✓ Reduce my deployment time from P.O. to Production



Our Solution – Hitachi G1000 AFA Mission Critical OLTP Environments

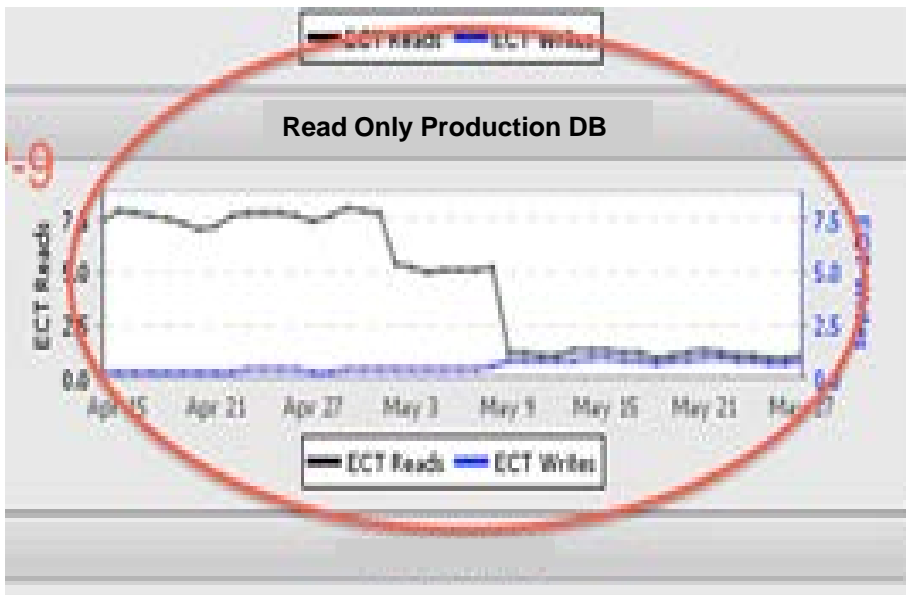


- ✓ HDS G1000 Storage Platform – 4th Generation Controller Based Virtualization Platform
 - ✓ Scales to Petabytes of addressable capacity
 - ✓ 4 Million IOPS @ 100% Random 8K Reads
 - ✓ Superior environmentalals
 - ✓ Hardware ASICs complement Intel x86 architecture to deliver consistent performance @ scale
- ✓ Hitachi Accelerated Flash Storage Intelligent Device
 - ✓ Purpose Built ASIC to offload flash housekeeping for maximum and consistent performance at all times.
 - ✓ Excellent Density – 38.4 TB in 2U
- ✓ Hitachi Dynamic Provisioning and Tiering Software
 - ✓ Storage Pooling Technology that delivers wide striping, thin provisioning and automated policy based tiering to provide maximum flexibility, density and data agility.

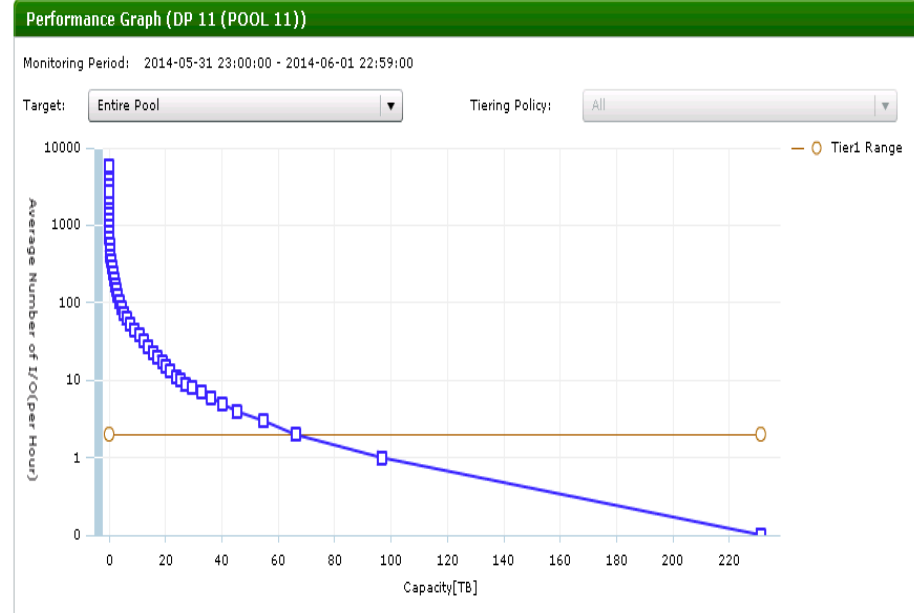


Results for Shared OLTP Workloads

- ✓ **Reduced Random Read response times from 7 ms to > 1ms**
- ✓ **Increased Application I/O Density for OLTP environments and Usable TB per Rack**
- ✓ **Significantly Reduced overall price per \$/GB by leveraging high capacity SAS drives while offloading performance to Flash Tier.**
- ✓ **Maximum Flexibility to allow the customer to adjust utilization for either performance and/or capacity.**



Tier Properties		
Tier 1	Drive	FMD
	Capacity(Used/Total)	50.17 TB / 51.19 TB [98 %]
	Performance Utilization	39 %
	Buffer Space(New page assignment/Tier relocation)	0 % / 2 %
Tier 2	Drive	SAS 10K
	Capacity(Used/Total)	181.13 TB / 268.39 TB [67 %]
	Performance Utilization	1 %
	Buffer Space(New page assignment/Tier relocation)	8 % / 2 %
Tier 3	Drive	-
	Capacity(Used/Total)	- / - [-]
	Performance Utilization	-
	Buffer Space(New page assignment/Tier relocation)	- / -





Results for Mission Critical DBs

Performance and Scale

- ✓ Database Performance Results
 - ✓ Oracle Data Volumes Average Read RT – 320 uSec / Write RT – 190 uSec
 - ✓ Oracle Redo Volumes Average Read RT – 280 uSec / Write RT – 170 uSec
 - ✓ Log Sync Operations Average RT – 680 uSec
- ✓ 7 X Improvement in Performance Cost Density (\$/IOP)
- ✓ Scale to over 500 TB of useable flash within a single system

Availability

- ✓ Near Zero performance impact from the most common maintenance activities and hardware failures.

Agility

- ✓ No changes in application or database design to implement
- ✓ Reduced install time from 5 days to 3 hours

Efficiency

- ✓ Avoided additional compute infrastructure



Thank You

Q & A