



Bart Sjerps
Principal Systems Engineer
EMC EMEA
bart.sjerps@emc.com
+31-6-27058830
Blog: <http://bartsjerps.wordpress.com>

INCREASING ORACLE PERFORMANCE WITH FLASH

AND RETIRING OBSOLETE PERFORMANCE PRACTICES

AGENDA

- Introduction
- Database & Application Challenges
- The XtremIO Solution
- Integrated Copy Data Management
- Self-Service Options for Databases
- Q&A

DATABASES AND APPLICATIONS TODAY

THE GROWTH CONTINUES

100+

Number of DBs
in typical large
organization

64%

Organizations seeing
increase in number
of Oracle DBs



800

Hours to
provision fully
loaded database

22%

Organizations pursuing
Database as a Service
Initiatives

Source: 2015 IOUG Survey on Database Manageability; Oracle surveys


© 2016 EMC Corporation - All Rights reserved | INFRASTRUCTURE AS A SUPERPOWER | EXECUTIVE FORUM | Bart Sjerps

EMC²


ORACLE CUSTOMER CHALLENGES

TOO MUCH TIME & BUDGET SPENT MAINTAINING ORACLE

Which Tasks Are Consuming Most DBA Time?




Which Areas Would You Like To Invest More?



(Source: **IOUG** - DBA Survey)
independent oracle users group

Time Spent on Maintenance




Top 5 Database Activities

(Source: 2014 IT Resource Strategies Survey)

EMC²

WHY LOOK AT LICENSING?

- Oracle DB licensing is expensive
 - Midsize server (24 cores):
HW ~ \$ 50,000
SW ~ \$ 483,000 @ 50% discount
5Y maintenance ~ \$531,000
(Enterprise Edition + basic options)
- What if we add RAC? Active DG?
Pluggable DB?
- Large part of the TCO of a database infrastructure stack




If we can save 10% on db licenses...

We easily justified 50% more expensive infrastructure

DATABASES SHOULD NOT HAVE I/O WAIT

- Adding CPU does not speed up I/O bottlenecks
 - Memory does somewhat
- IOPS are relatively (!) cheap
- CPU cycles are expensive
 - Because of licenses
- Consolidation leads to
 - Higher IO requirements
 - I/O bottlenecks
 - Bandwidth issues
- Flash storage can solve these limitations




STORAGE IS NO LONGER THE BOTTLENECK

EMC²

SPINNING DISK VS FLASH STORAGE

SPINNING DISK IS 60+ YEARS OLD AND SLOWLY DYING



- **One seek at a time**
- **Mechanical latency (7ms)**
- **Max 150 random IOPS**
- **Good at sequential IO**
- **No write penalty**


VS.

- **Parallel requests (0.1 ms)**
- **No mechanical latency**
- **Virtually unlimited IOPS**
- **Good at sequential IO**
- **Write penalty compensated by controller (cache, leveling)**

**"When 900 years old you reach
Look as good you will not"
-- Yoda (Return of the Jedi)**

IO-BOUND DATABASE SERVER

4 hours time wasted by high-latency random reads



Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
db file sequential read	1,420,110	14,362	10	92.7	User I/O
free buffer waits	89,072	986	11	6.4	Configurat
DB CPU		158		1.0	
library cache lock	138	51	371	.3	Concurrenc
write complete waits	9	17	1919	.1	Configurat

AFTER MOVING TO ALL FLASH ARRAY

Event	Waits	%Time -outs	Total Wait Time (s)	Avg wait (ms)	Waits /txn	% DB time
direct path read	361,425	0	930	3	7,376.0	16.4
db file parallel read	365,088	0	898	2	7,450.8	15.9
db file scattered read	64,628	0	114	2	1,318.9	2.0
db file sequential read	56,748	0	84	1	1,158.1	1.5
gc cr multi block request	137,524	0	62	0	2,806.6	1.1
direct path read temp	19,942	0	44	2	407.0	.8
read by other session	17,389	0	31	2	354.9	.6

Event	Waits	Time (s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		3,427		60.5	
direct path read	361,425	930	3	16.4	User I/O
db file parallel read	365,088	898	2	15.9	User I/O
db file scattered read	64,628	114	2	2.0	User I/O
db file sequential read	56,748	84	1	1.5	User I/O

- ✓ Avg. latency went down by more than 70%
- ✓ CPU utilization doubled

THE IOPS RACE IS OVER

MILLION IOPS? 2 MILLION IOPS? 5 MILLION IOPS? WHO CARES?



Every flash vendor these days offers:

Huge IOPS numbers
Very low latency

Average small AFA: >> 200K IOPS
Average large DB: << 50K IOPS

Competing in the FLASH business is not about IOPS and latency any more *)
-- Bart Sjerps, 2016

*) Considering generic Oracle workload consolidations here - Large scale DWH/BI is a special case.

What are differentiating features of Flash solutions for databases - if it isn't IOPS and latency?

PERFORMANCE: PREDICTABLE & SCALABLE


ARCHITECTED FOR PREDICTABLE & CONSISTENT PERFORMANCE

Latency

Consistent performance even as SAP workloads (IOPS) increase & scale

Benefits


- No tuning or overprovisioning for SAP
- Size for utilization, not SAP peak performance
- Confidently meet SLAs
- Run SAP analytics and suite always, anytime, faster



* non-SAP workload

Actual data, 8KB OLTP workload*

WHAT IF YOU COULD SIMULTANEOUSLY...



Consolidate
everything: dev,
test, production,
analytics



Meet DB
performance
SLAs



Simplify DB
operations;
deploy apps
sooner



Reduce product
& administrative
costs

XTREMIO DATA SERVICES

ALWAYS-ON, INLINE, ZERO PENALTY

ALWAYS ON THIN
PROVISIONING



INLINE DATA
DEDUPLICATION



INLINE DATA
COMPRESSION



INLINE DATA
ENCRYPTION



FLASH OPTIMIZED
DATA PROTECTION



WRITABLE
COPIES




*Up To 6:1 Database Reduction On Average
Just For Production Databases*

EMC²

EMC XTREMIO REDEFINES DATABASES

CONSOLIDATE **EVERYTHING** ON THE SAME INFRASTRUCTURE



Production databases
Development databases
Functional & performance testing
Analytics / OLAP
Oracle & SAP applications

All while maintaining performance & availability SLAs!

EMC²

XTREMIO IS #1 WITH 34% MARKET SHARE

XtremIO has

Of Effective
Capacity



XtremIO has

Of the
Fortune 100

"Magic Quadrant in Solid State Analysis" – **Gartner 2014**

"#1 Market Share 2014" – **IDC**

"Leader in Strategic AFA Technology" – **451 Research**

"DBTA Trend-Setting Product for 2016" – **DBTA Magazine**



EMC²

COMPLEXITY IS NO LONGER NEEDED

SIMPLIFY ARCHITECTURE AND OPERATIONAL MANAGEMENT

HISTORIC COMPLEXITIES

Disk Type Selection

RAID Protection Selection

RAID Group or Storage Pool

Striped or Concatenated Meta
Devices

Auto Tiering

Array Caching

Thick or Thin LUNs

XTREMIO

Create Volume of Required Size

DONE!

EMC²

OTHER OBSOLETE PRACTICES

SIMPLIFY ARCHITECTURE AND OPERATIONAL MANAGEMENT



Under-allocation
(disk "short stroking")



Data Separation



Pre-fetching
(attempting to
predict the future)



Striping



Monster caches
(preloading everything)

MORE OBSOLETE PRACTICES



Classic RAID
(Having only one extra safety rope)



Full data copies
(Avoiding I/O interference)




Storage Layers / Tiering
(choice between performance
or economics)



Direct Attached Storage
(Trading redundancy & sharing
for bandwidth & cost)

FLASH INCREASES LICENSE ROI

INCREASED CPU UTILIZATION REDUCING I/O WAITS



**2X
HIGHER
UTILIZATION
OF DATABASE
LICENSE INVESTMENT**

Source: XtremIO Customer Case Study

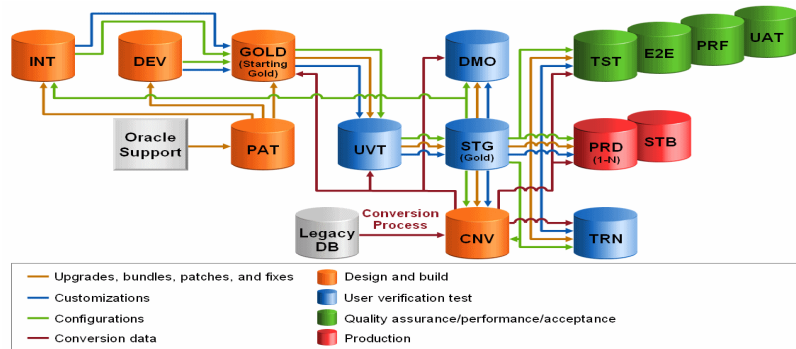
EMC²

WHY DATABASE CLONING?


- "Serverless backups" & Quick restores
 - Zero production performance impact
 - RTO seconds to minutes
 - Protected & Instant restores
 - Out-of-order restores
 - Instant Restore from remote copies
- Firefighting
 - Creating a quick production copy to solve application problems
 - Without messing with production data
- Creating Test / Dev / Acceptance copies
 - Automated, no tape restores, low people effort
- Creating copies for reporting / staging
 - Datawarehouse queries can bring production performance down
 - Moving reporting workload to copy relieves production
- Application / Database Upgrades
 - Creating application "checkpoints" avoids having to fall back to starting point due to small errors
 - Easy upgrade testing




Sooam is the world leader in dog cloning. Since 2005 they have produced more than 550 cloned puppies.




BUSINESS DATA CLONING - INNOVATIONS



1997: EMC Timefinder (Full copy)



2000's: Snap (incremental)



XTREMIO VIRTUAL COPIES



Efficient In-Memory Copies


- > No space reservation or moving blocks
- > No metadata bloat, shared in-memory metadata
- > Immediate creation with no impact on prod
- > Thin provisioning, dedup, compression

100% Performance


- > Same performance as prod
- > Predictable performance on prod or copies
- > No impact on prod when copies are created
- > Optimal SLA

XVC VS. TRADITIONAL SNAPSHOTS

XtremIO Vs. Vendor "A" DB Volumes, Steady State, IOPS Over Time



XtremIO Vs. Vendor "A" DB Volumes, Steady State, Latency Over Time



XtremIO

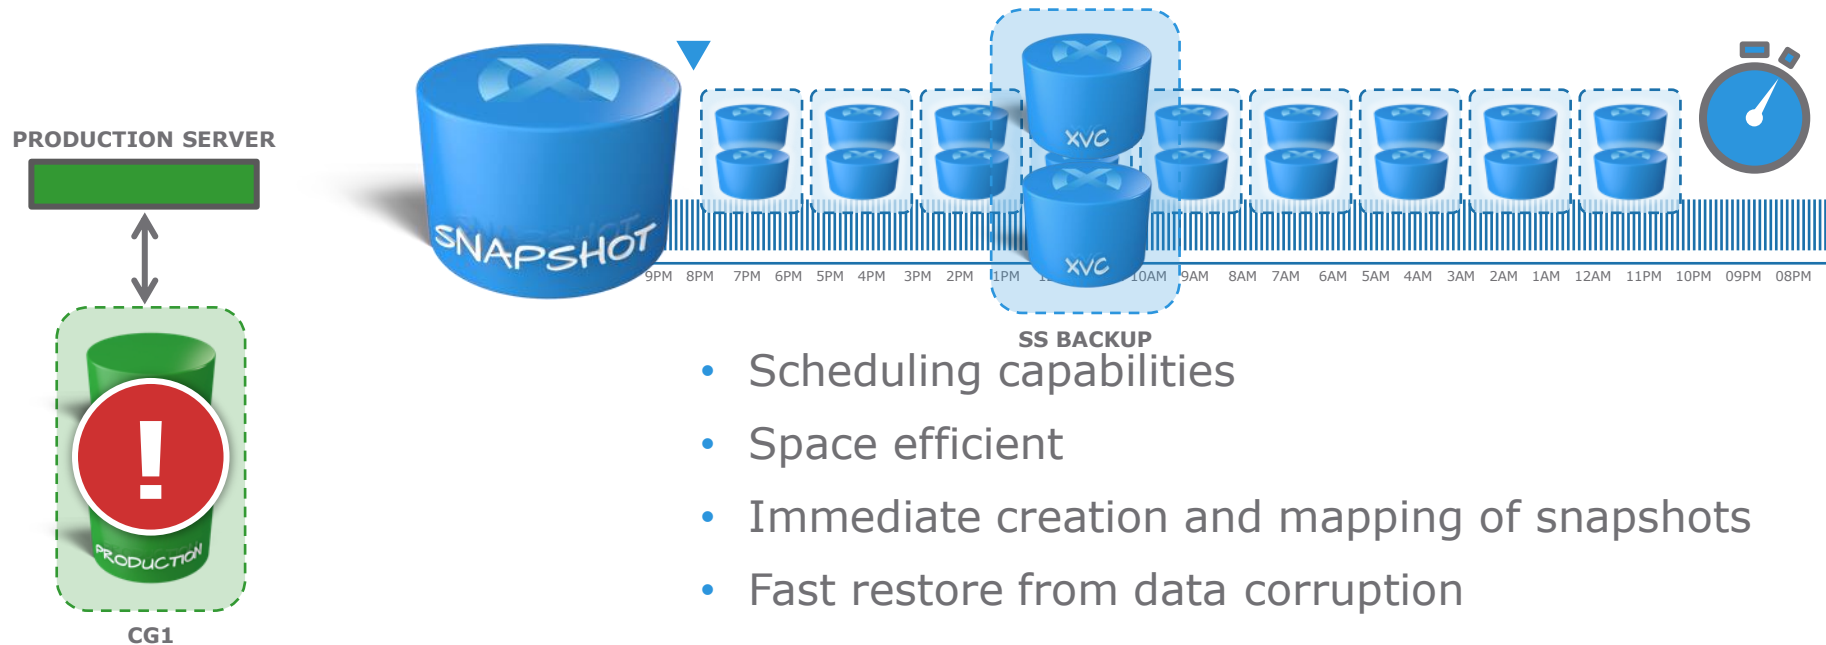
- **No impact** on copy creation
- **Consistent** performance on prod and copy

Vendor A

- **IOPS** drop by 50% to **35K**
- Spikes to **20ms latency**

ENSURE DATA PROTECTION

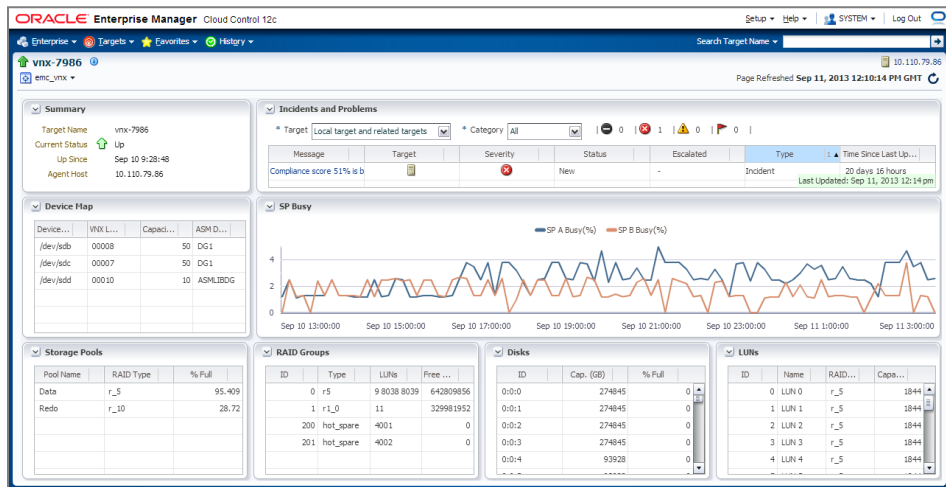
PROTECT FROM LOGICAL DATA CORRUPTION BY CREATING SNAPSHOTS OVER SHORT INTERVALS



- Scheduling capabilities
- Space efficient
- Immediate creation and mapping of snapshots
- Fast restore from data corruption
- Fine Recover Point Objective (RPO)

SELF-SERVICE WITH ORACLE ENTERPRISE MANAGER

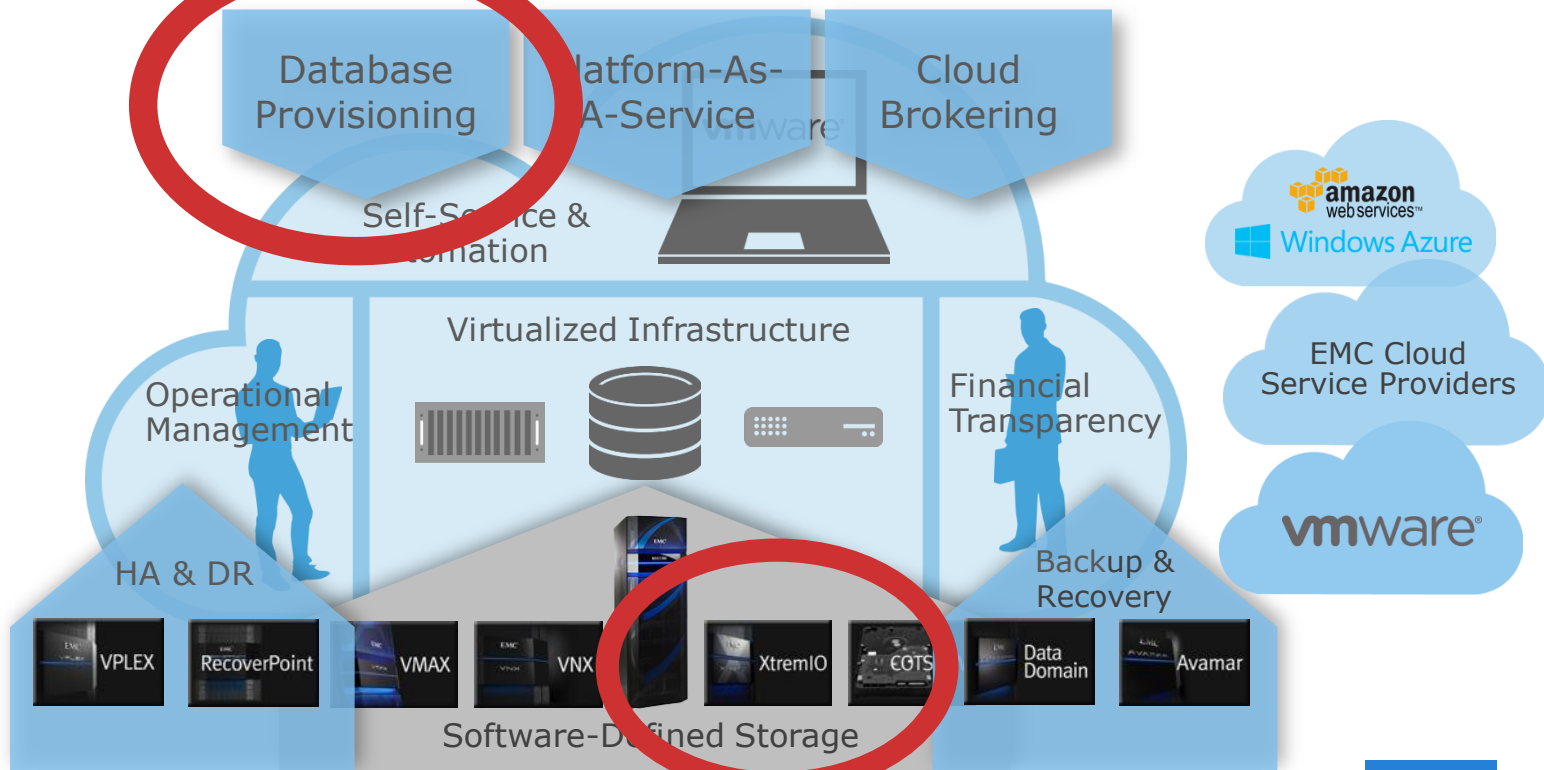
FREE XTREMIO PLUG-IN



- Manage & monitor the Oracle environment, including storage
- Automate tasks and assist on database administration
- Monitor XtremIO utilization, configuration, and performance
 - Throughput and IOPS
 - Response Time
- Set thresholds and notifications for key XtremIO health indicators

SELF-SERVICE DBAAS WITH EMC FEDERATION

ADD-ON MODULES EXPAND CAPABILITIES



EMC²

Pivotal

RSA


VCE

vmware

EMC²


SILLY LITTLE ORACLE BENCHMARK

I/O PROFILING FOR ORACLE – TEST RESULTS



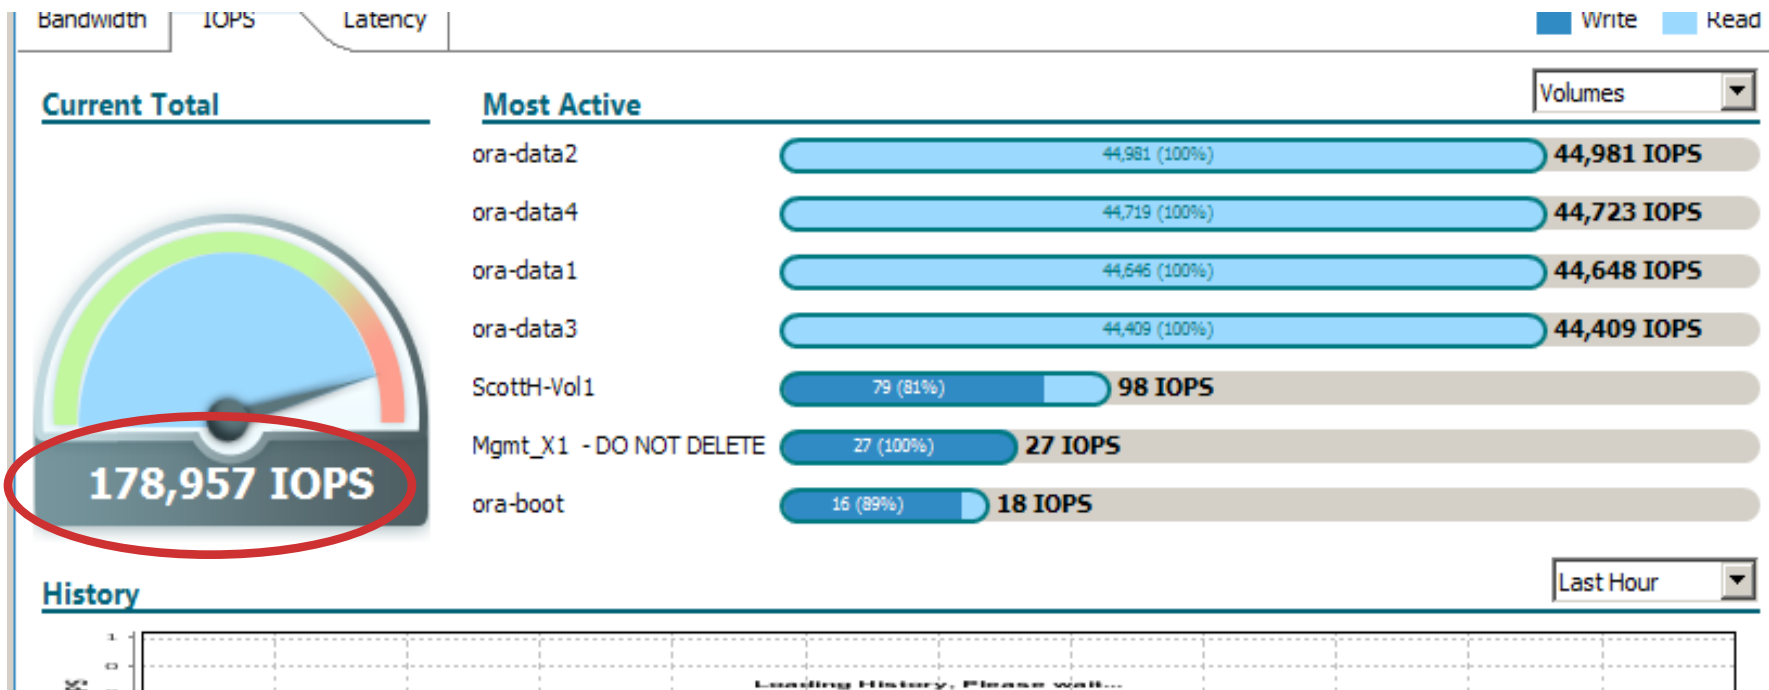
Info & download:
XtremIO.com/slob

RPM pre-packaged version:
[SLOB RPM Wiki on Outrun](#)



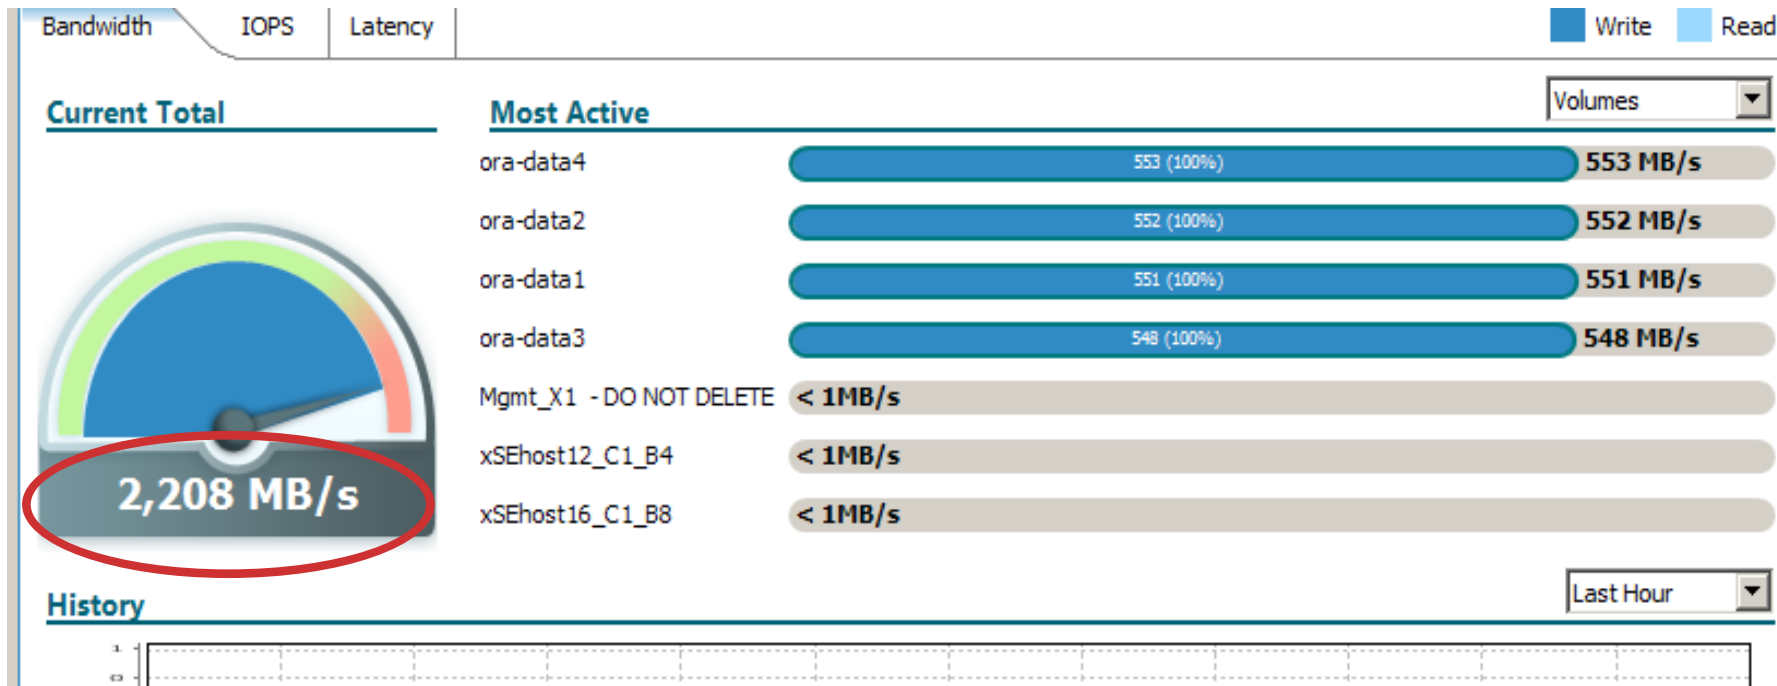
PERFORMANCE EXAMPLE WITH SLOB

LAB TEST: 1 V2.4 X-BRICK, 3 VM'S ORACLE 11.2.0.4.0, VMDKS




PERFORMANCE EXAMPLE WITH SLOB

BANDWIDTH (1 X-BRICK – LINEAR SCALABLE TO 10 X-BRICKS)



PERFORMANCE EXAMPLE WITH SLOB

LATENCY (NEARLY ALL I/O IS 8K RANDOM)



COMPARING COMPRESSION CHOICES

STORAGE VS DATABASE COMPRESSION

XtremIO Compression

Workload



- **All data** is compressed
- **No increase** in database server processor utilization
- No DBA overhead
- No additional cost


Oracle Advanced Compression

Workload



- **Only database** is compressed
- **Increase** in database server processor utilization
- Accelerates queries
- Additional cost: \$11,500 per processor

Does XtremIO and ACO Work Together?



YES!

EMC²

WHAT IF YOU NEED MORE BANDWIDTH?

HUGE DWH/BI WORKLOADS - CLASSIC APPROACH



Network Attached Flash

High Capacity with Enterprise Features

- Latency compromised by networks
- Deliver limited Bandwidth & IO
- ***Forces Compromised Siloed Approach***



ERP



ERP EMEA



Data



BU Mart



ODS



BI Mart

SOLVING DWH/BI BANDWIDTH PROBLEMS

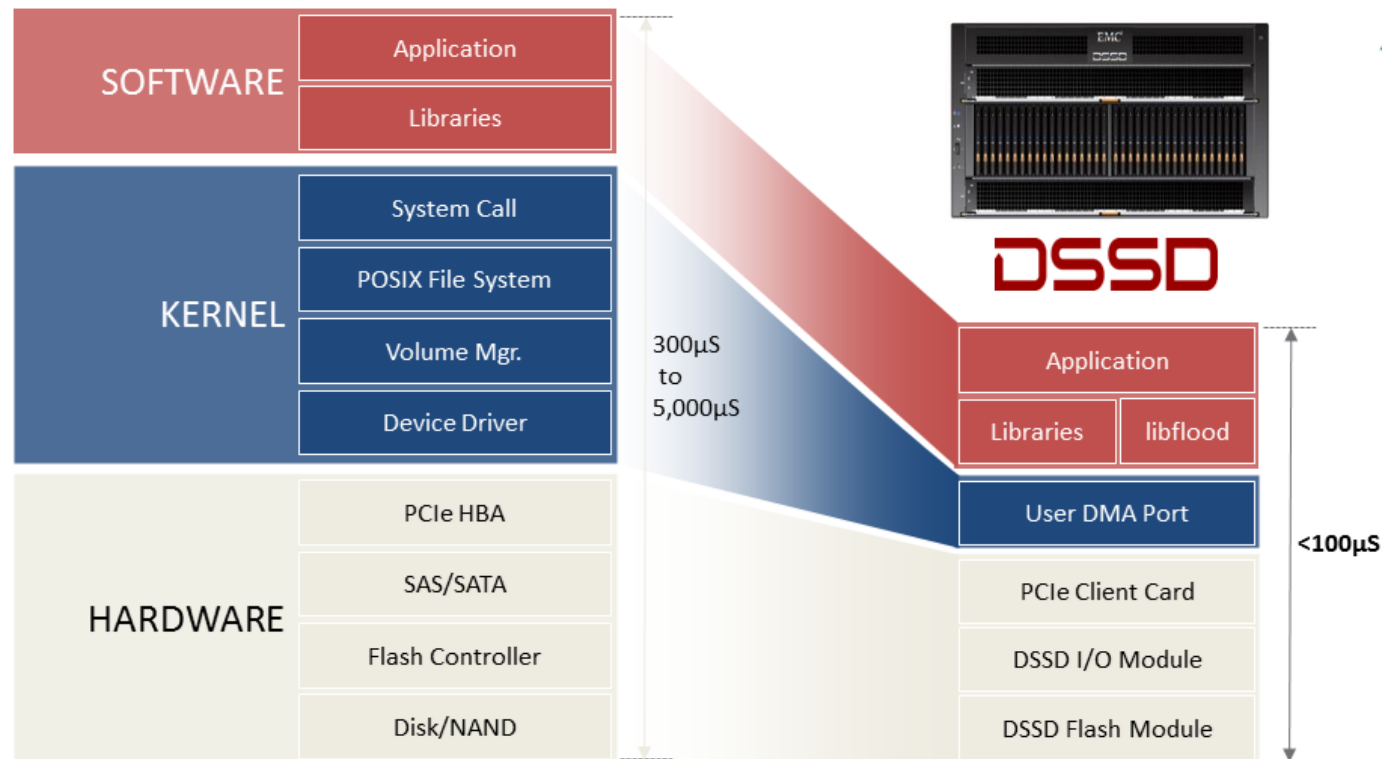
CLASSIC SOLUTIONS

- Materialized views
- Heavy indexing / partitioning
- Multiple copies / subsets
 - data marts
 - Business copies (snapshots)
- Attempt to offload in storage
 - Storage indexes
 - Columnar storage / compression



INNOVATION: NVME I/O PROTOCOL

EMC DSSD: PCIE DIRECT ACCESS TO SSD



- ✓ 10M IOPS
- ✓ 100 GB/s
- ✓ <0.1 ms

Enterprise
Readiness

- ✓ Data protection
- ✓ Redundancy
- ✓ Consistent performance

Tradeoff:
Data services
⇒ Future

EMC/ORACLE SOLUTIONS CENTER

ORACLE

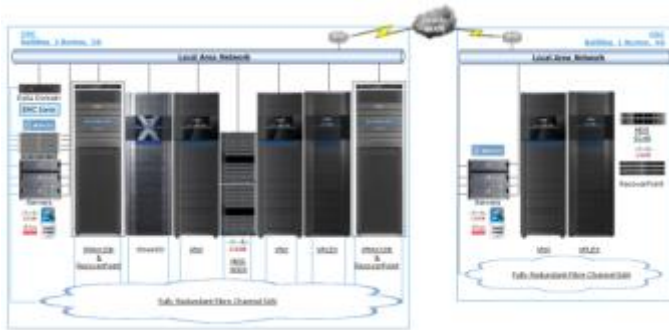


EMC²



Oracle Campus, Reston, VA

- Shared services for Oracle & EMC
 - Over 500 servers
 - Over 1PB EMC storage
 - Fully Virtualized on VMware
- Provides infrastructure for
 - Oracle's Training & demos
 - EMC Demos
 - EMC POCs
- Oracle Integration Demos
 - Storage integration, cloning & replication
 - HA Stretched clusters
 - Management tooling



OSC Infrastructure

[Leverage EMC at Oracle Solution Centers](#)

EMC²

REFERENCES

My Blog "Dirty Cache"

<http://bartsjerps.wordpress.com>

Everything Oracle @ EMC (community):

<http://emc.com/everythingoracle>

XtremIO

<http://xtremio.com/>

Outrun:

<http://outrun.nl>



Dirty Cache

A storage infrastructure perspective on optimizing business applications

HOME INDEX RESOURCES PRESENTATIONS ABOUT
ORACLE PERFORMANCE INNOVATION GENERAL VPLEX FAQ VARIOUS VIRTUALIZATION

← Thank you, Larry Ellison

Stop Idling – Start Saving

OCTOBER 23, 2012 LEAVE A COMMENT



One of my missions is to help customers saving money (Dirty Cache Cash). So considering the average enterprise application environment, I frequently ask them where they spend most of their IT budget on. Is it servers? Networks? Middleware? Applications?

Turns out that if you look at the operating cost of an Oracle database application, a very big portion of the TCO is in database licenses. Note that I focus on Oracle (that's my job) but for other databases the cost ratio might be similar. Or not. But it makes sense to look at Oracle as that is the most common platform for mission-critical applications. So let's look at a database environment and forget about the application for now. Let's say that 50% of the operating cost of a database server is spent on Oracle licensing and maintenance



BLOG
WITH
INTEGRITY
blogwithintegrity.com

» Follow Blog via En

Enter your email address to receive notifications of new posts by email.

Join 37 other followers

EMC²

EMC²®