



Consolidate Database & Analytics Workloads on a Scale-Out All Flash-Array

Andy Fenselau

August 5, 2014

XtremIO

EMC²

#1 All-Flash Array in the Market

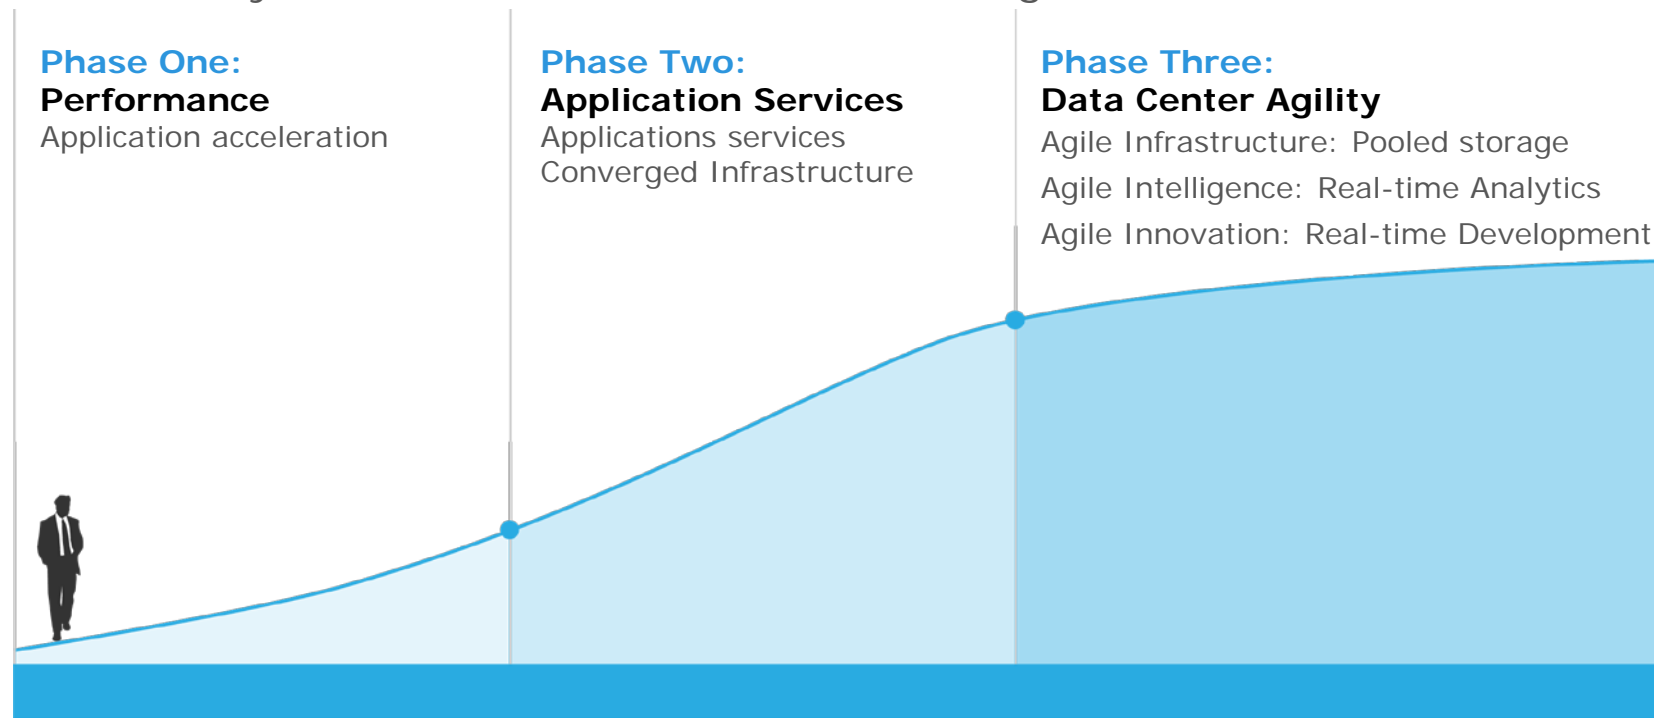
XTREMIO

>\$100,000,000 <6 months >70PB sold

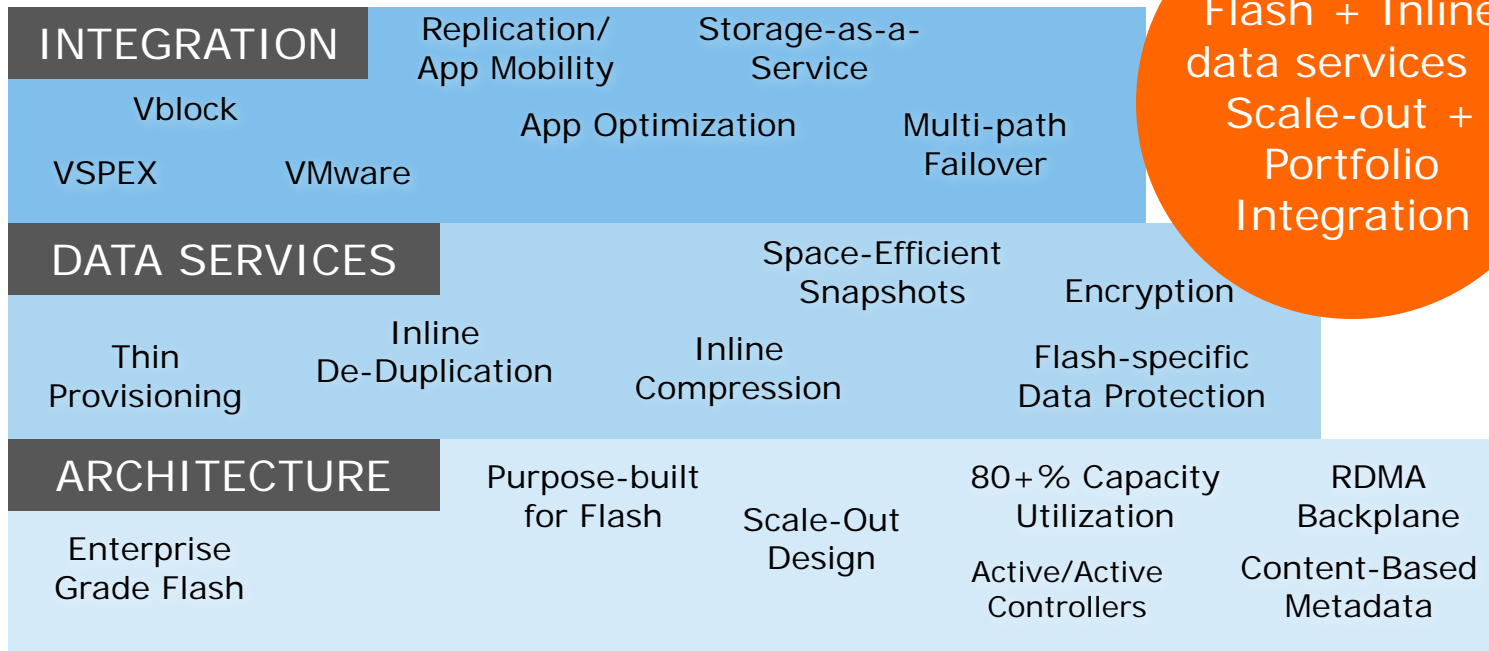
EMC²

Redefine Your Data Center With XtremIO

The Journey to The Private Cloud Part 2: Storage At Last!



Architecture Matters



Flash + Inline
data services +
Scale-out +
Portfolio
Integration



What Customers Are Saying

“When you’re doing a 100,000 IOPS why wouldn’t I put everything on this? You get a boat-load of business value for OLTP. I’m having a hard time saying ‘why not?’ ... I am looking for StarTrek technology. If we are going to get to that level, we can’t be twittling around in the details. We need things that just work. Between virtualization and the XtremIO we have boxes we just plug in. We don’t worry about that moving disk drives and creating LUNS. I think that is what’s going to allow us to get to the next level...We have to find a way to move everything over to it.”

- Stew Gibson, CIO, USI Holdings

Database & Analytics Workloads

1

Breakthrough
Consistent
Low-Latency
Performance



2

Consolidation:
Dev/Test
BI/Analytics
SAP Landscapes



3

Lowest TCO:
Storage
Servers & Licensing
OPEX



4

No Complex Setup
No Tuning



Case Study: Oracle Consolidation

XtremIO Snapshots' Performance, Flexibility to Oracle 11g

Challenges

- Daily Crunch business processing times struggling to update by opening hours
- Query response times
- Challenge for Development/Test scale, performance
- Monthly reporting time too slow, rigid

Solution

- XtremIO All-Flash Array
- Production on One X-Brick
- Test/Dev on 2 X-Brick cluster

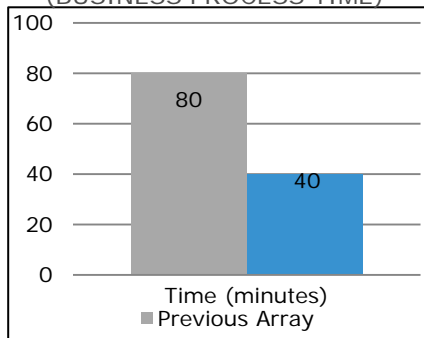
Applications

- Oracle 11G, virtualized Linux, Windows
- High user concurrency
- Large dev/test environment

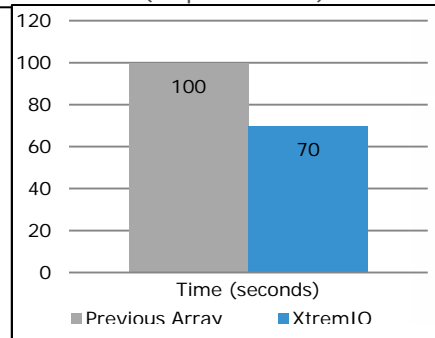
Results

- 60% business process time improvement for daily "Crunch"
- 15-40% user query time improvement
- Monthly batch reporting time cut 50%
- 30% Oracle license savings as CPU wait time reduced, utilization increased
- DBAs "blown away": DB seq read 34ms -> 0.3ms, CPU time improved 10%->96%
- 30% faster dev cycles- Snapshots/Deduplication give "free" performance database copies to all
- Additional consolidation for virtual servers & virtual desktop on same XtremIO
- Planning complete block storage consolidation for new datacenter @ 54% lower cost and 70% smaller footprint


"THE CRUNCH"
(BUSINESS PROCESS TIME)



USER QUERY TIMES
(Improvement)




Baillie Gifford's Traditional Environment





28

DATABASE COPIES



3

ARRAYS/POOLS




1

USE CASE FOR FLASH

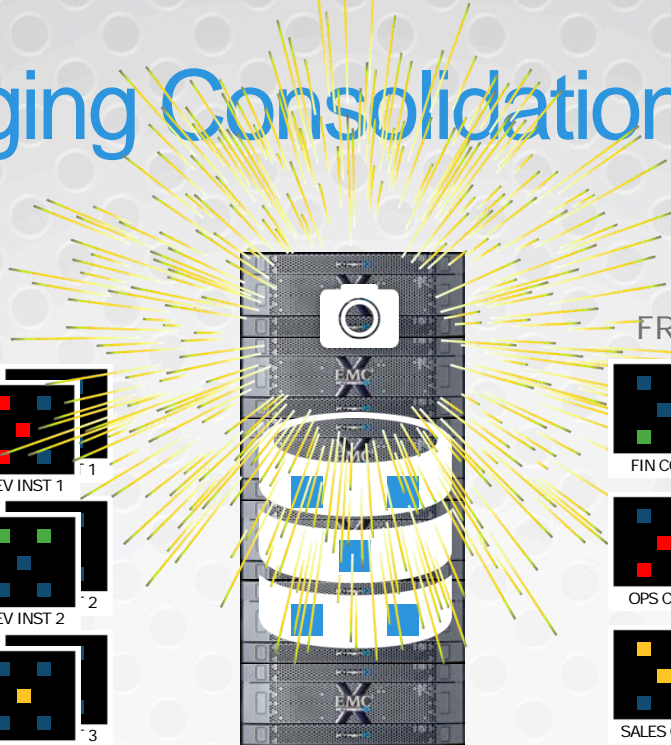
JUST
ONE
APPLICATION

Game-Changing Consolidation, Agility

FREE, FAST DEV/TEST



TEST/DEV COPIES



HIGH PERFORMANCE
PRODUCTION

FREE, FAST ANALYTICS



DATAMART COPIES

SCALE-OUT
IOPS IN
ABUNDANCE

XTREMIO
FLASH FOR
ENTIRE
APPLICATION

MORE
BUSINESS
PRODUCTIVITY

FASTER
APPLICATION
DEVELOPMENT
TIMES

27:

1

Data Reduction

1

XTREMIO CLUSTER

0

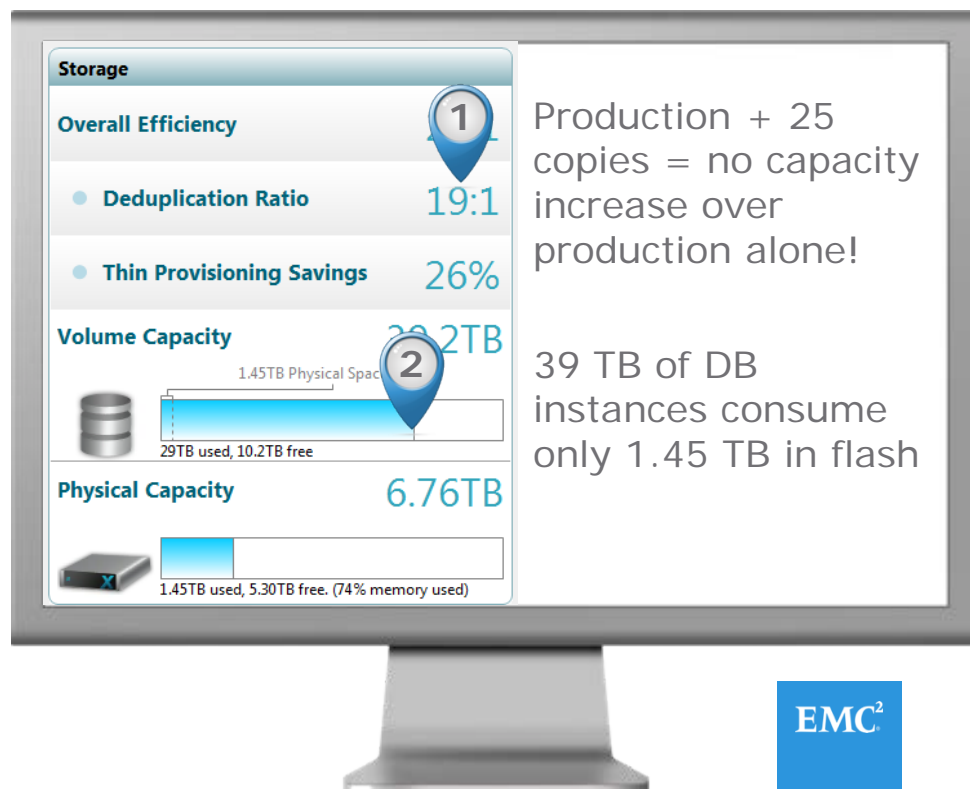
BRUTE FORCE
COPIES

EMC²

Development/Test on XtremIO



- On Demand copies
- Same performance improvements as in production
- Inherent QoS—on other instances even under heavy load
- 30% faster database app development & QA time



Sandy Bryce, Enterprise Platforms Team Lead

- Just to give you an example of the performance gains we've seen, during the testing phase we achieved a 30% reduction in run time for a key business process. That was already amazing—but then, when we put the EMC XtremIO All-Flash Array into production, that number skyrocketed to 60%.”
- “A single copy of the database takes up 1.6TB. Twenty, thirty, forty copies of the database—they still take up only 1.6TB of physical flash, and the performance is on par with production. The inherent data reduction capabilities of XtremIO enable a small Flash footprint to serve a large development environment; it's almost like providing free databases for our developers. This completely solves our dev/test challenge.”
- “From a development point of view, it's flexibility in terms of being able to provision multiple copies without increasing the physical footprint, and reducing the time to do those extra copies from hours or days to minutes. We are also able to give our development platform the same performance that we give production, with no cross-platform impact.”
- We have always been able to add capacity and add more disks, which potentially improves performance; but you always reach a point where, as you add capacity, your performance starts to flat-line or even degrade. With XtremIO, every time you add an X-Brick, the performance scales linearly without impacting things like latency.”

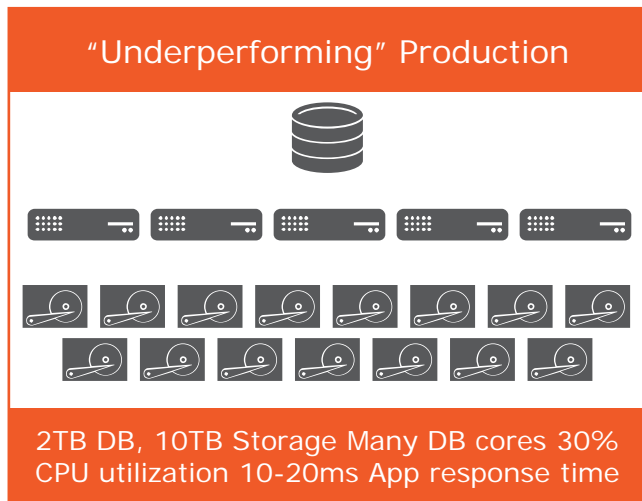
Xpect More Information

www.xtremio.com

- Reference Architectures & Solution Guides
- All Flash Array Testing Guide & Best Practices
- Architecture & Product white papers, videos
- Customer case studies

EMC²®

Why Don't Databases Perform?



- Variability in database IOs
 - Sustained random/sequential I/O
 - High concurrency → Random
- Growth in data and randomness
 - Disk/RAID types, caching, tiering does not work in all cases
 - Keeping up with data hotspots

Chasing Hotspots

Table 6 Aggregate Layout

Controller	Aggregate Name	Option /RG Size	# of Disks/ Usable Size	Purpose	Volume name	Use
O_A_1	aggr0	RAID-DP, 3	no's/350GB	DOT and root volume	ORA_HOM_A_1_A_1	Database Binary
O_A_1	AGGR_OR_A_2_A_1	RAID-DP, 42 no's/13TB		Data files, Redo logs, control files, Database and Cluster Ware Binary	ORA_HOM_A_1_A_1	Database Binary
O_A_2	aggr0	RAID-DP, 3	no's/350GB	DOT and root volume	ORA_HOM_A_2_A_1	Database Binary
O_A_2	AGGR_OR_A_2_A_1	RAID-DP, 42 no's/13TB		Data files, Redo logs, control files, Database and Cluster Ware Binary	ORA_HOM_A_2_A_1	Database Binary
O_B_1	aggr0	RAID-DP, 3	no's/350GB	DOT and root volume	ORA_HOM_B_1_A_1	Database Binary
O_B_1	AGGR_OR_A_2_A_1	RAID-DP, 42 no's/13TB		Data files, Redo logs, control files, Database and Cluster Ware Binary	ORA_HOM_B_1_A_1	Database Binary
O_B_2	aggr0	RAID-DP, 3	no's/350GB	DOT and root volume	ORA_HOM_B_2_A_1	Database Binary
O_B_2	AGGR_OR_A_2_A_1	RAID-DP, 42 no's/13TB		Data files, Redo logs, control files, Database and Cluster Ware Binary	ORA_HOM_B_2_A_1	Database Binary
3170_B_2	AGGR_OR_A_2_B_1	RAID-DP, 42 no's/13TB		Data files, Redo logs, control files, Archive Log	ORA_HOM_B_2_A_1	Database Binary

1. Different RAID groups

7. Short-stroking

6. Multiple Indexes/Schemas

4. Different volumes

3. Different drive types


2. Different controllers

5. Different sizes

Overprovisioning is "THE NORM"

IO-Bound Database Server - Wasted CPU Cycles

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
tree 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

Fewer Cores Need To Be Licensed

