



# High Performance Oracle Database in a Flash

Sumeet Bansal, Principal Solutions Architect

# AGENDA

- **Why use NAND flash?**
- **How to know if NAND will help**
- **A perfect fit with Oracle ASM**
- **Single Instance Oracle implementation**
- **Oracle RAC implementation**
- **High Availability - Oracle Data Guard**
- **Real world examples**



## WHY USE NAND FLASH

- Overcome traditional I/O bottlenecks
- Create simpler architectures
- Save on energy costs
- Drive optimal use of resources

FUSION-io




# THE DATA SUPPLY PROBLEM LEADS TO IDLE CAPACITY


FUSION-io

Processing performance doubles every 18 months, according to Moore's Law

37% of servers are massively underutilized<sup>1</sup>...



...because the performance gap continues to grow



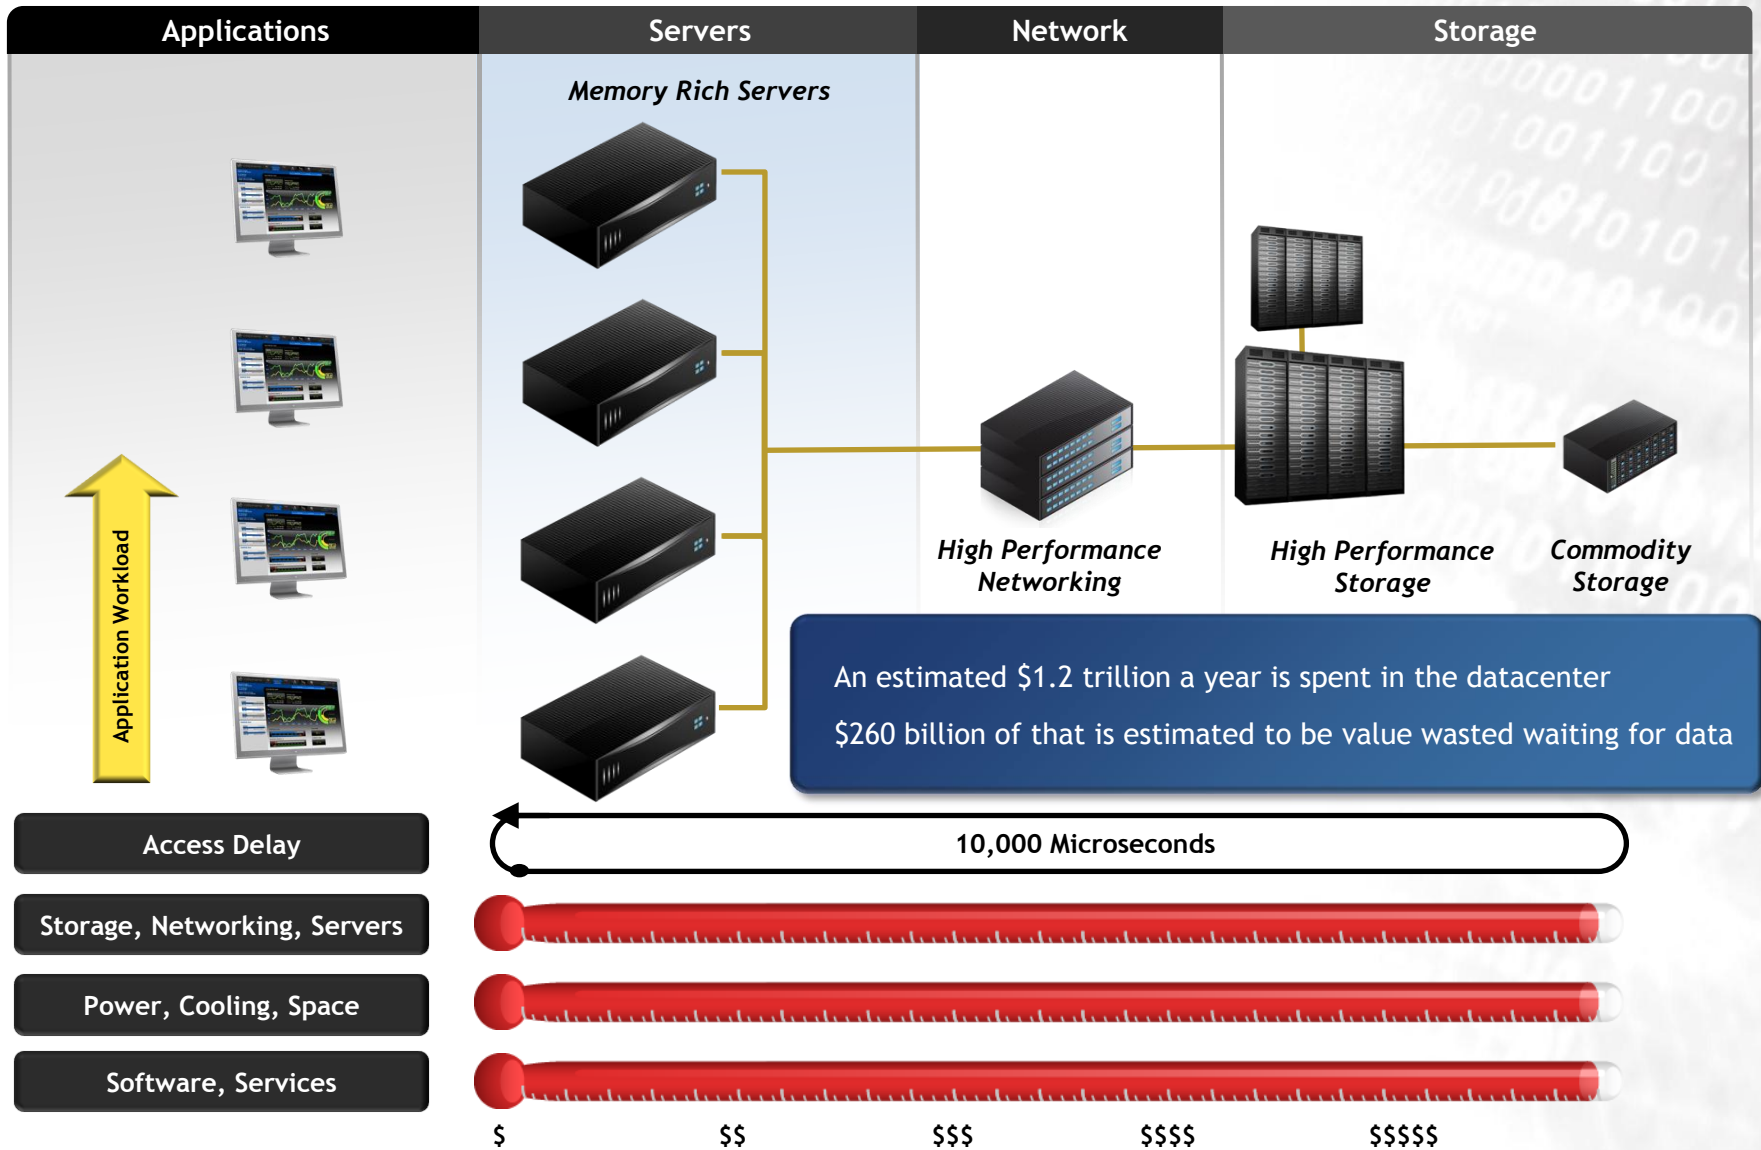
<sup>1</sup> Source: IDC's Server Workloads 2010, July 2010

<sup>2</sup> Source: Taming the Power Hungry Data Center, Fusion-io White Paper



# THE DATA SUPPLY PROBLEM CAUSES DATACENTER BLOAT

FUSION-io




Source: Based on 3<sup>rd</sup> party forecasts and company estimates



# THE SOLUTION STARTS WITH MINIATURIZING THE SAN


FUSION-io

The performance of a SAN...



10,000 microsecond access time

...in the palm of your hand



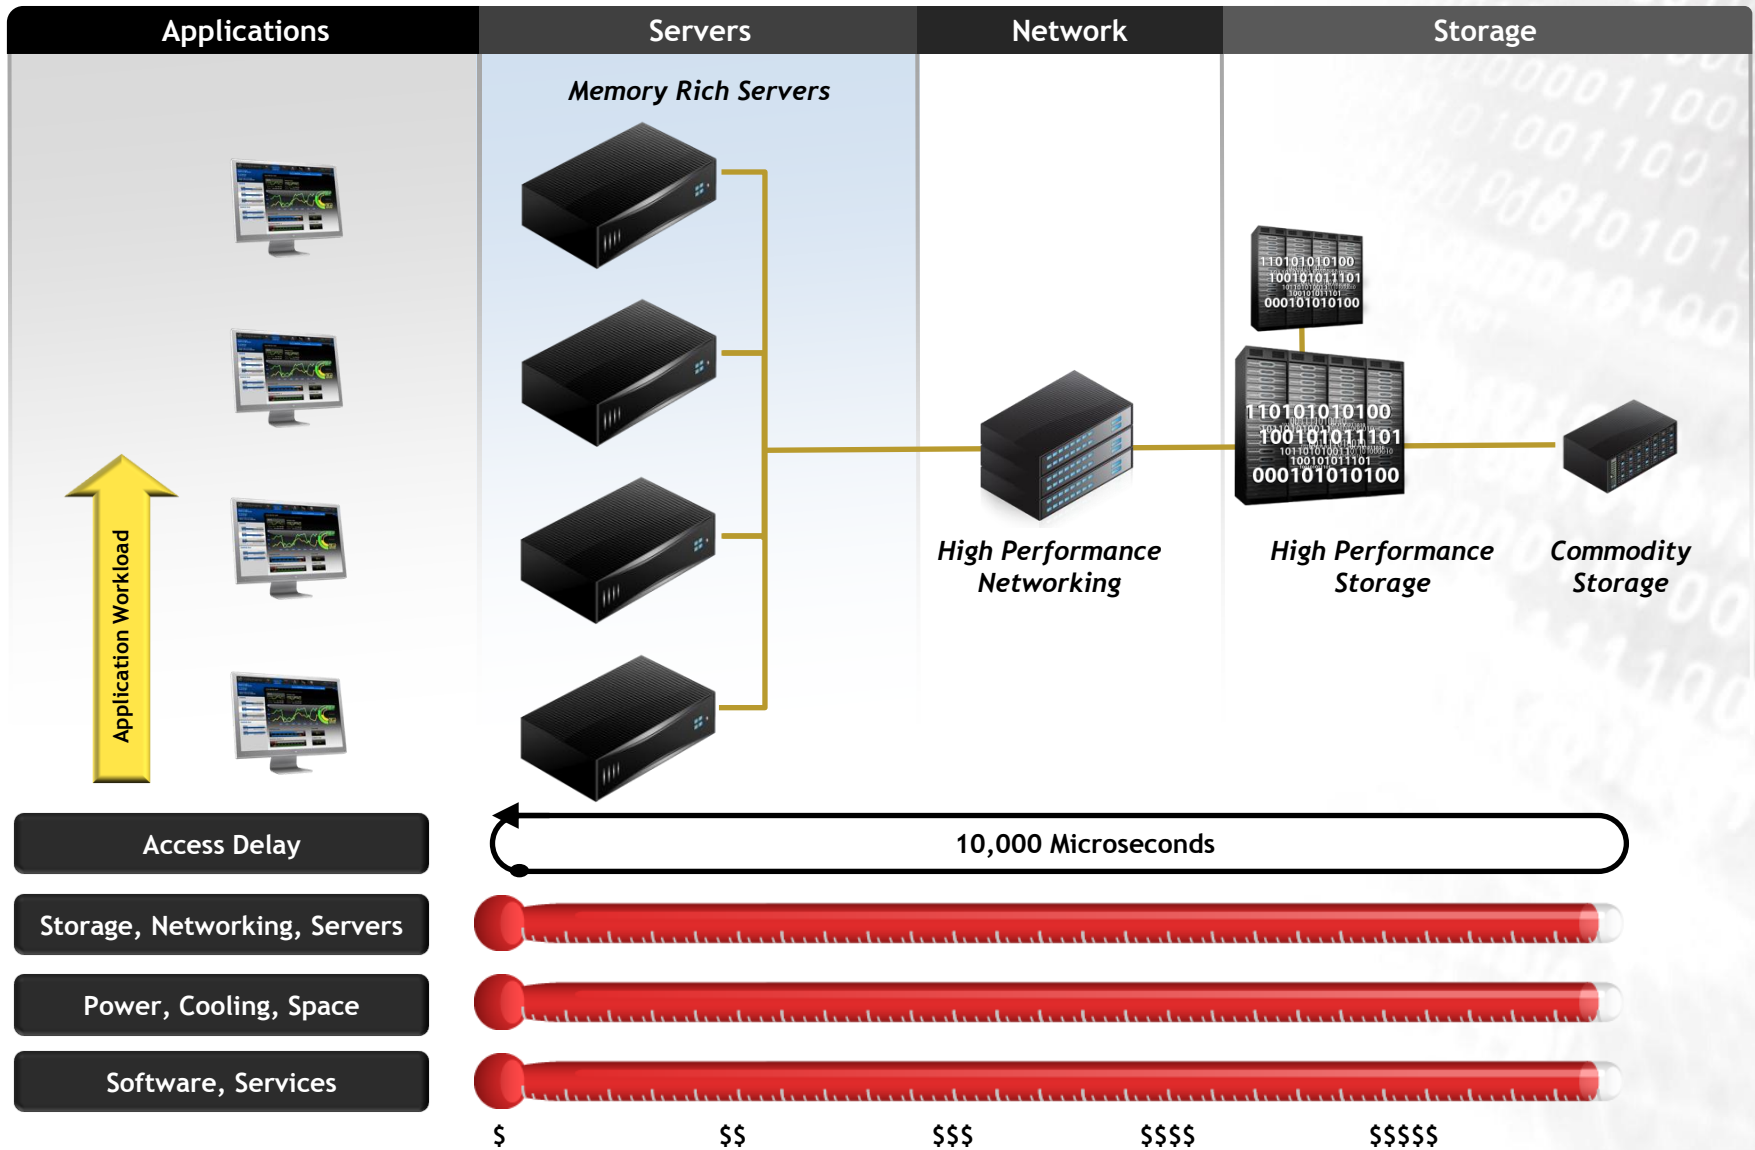
< 30 microsecond access time

Source: The Storage Performance Council and Fusion-io Estimates



# DECENTRALIZATION PUTS THE DATA CLOSER

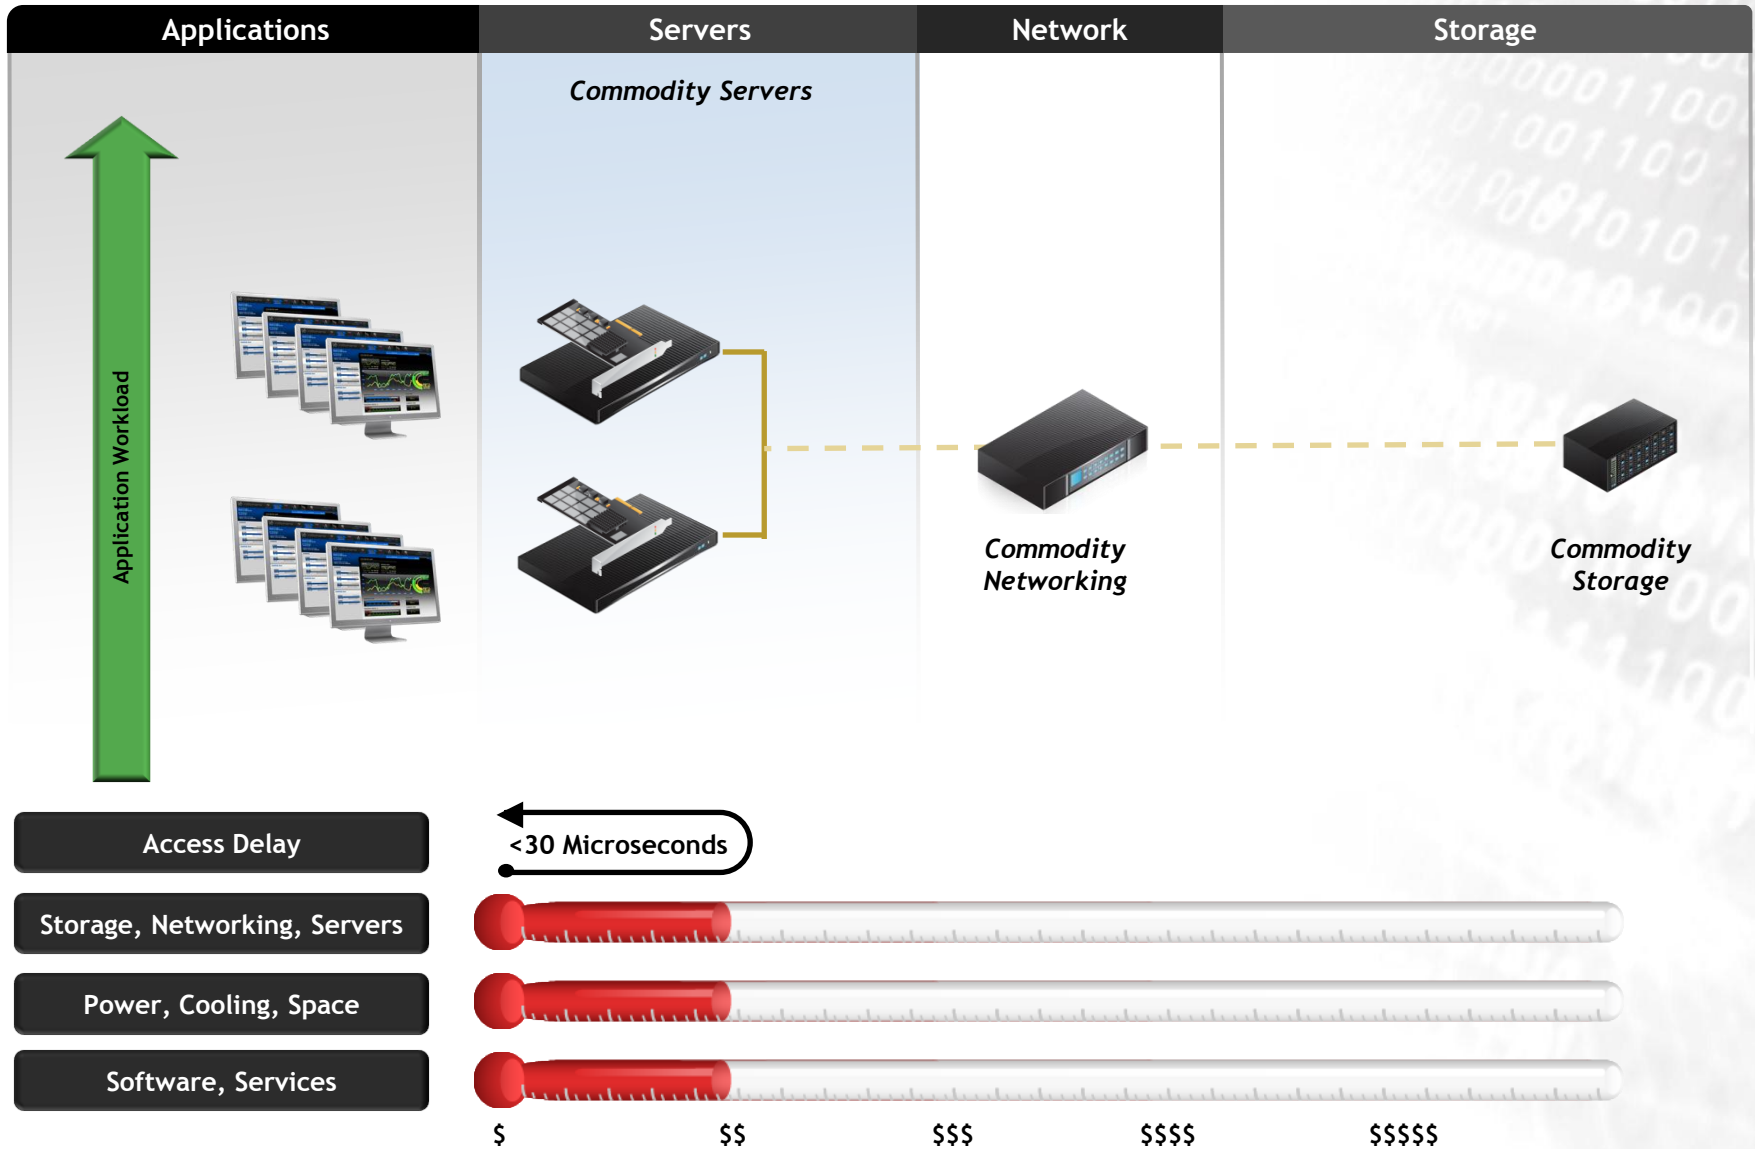
FUSION-io





# DECENTRALIZATION ELIMINATES WAIT AND BLOAT

FUSION-io







# HOW TO KNOW IF NAND WILL HELP

FUSION-io

- Look at ADDM Advisory

The screenshot displays the Oracle Enterprise Manager 11g Database Control interface. A confirmation message states "ADDM has been run successfully". The main section is titled "Automatic Database Diagnostic Monitor (ADDM)" and shows a "Database Activity" graph for "Apr 7, 2011". The graph plots "Active Sessions" (0 to 31.8) against time (4:00 AM to 12:00 PM). A legend indicates "Wait" (orange), "User I/O" (blue), and "CPU" (green). Below the graph, a "TIP" explains icon usage. The "ADDM Performance Analysis" section shows "Task Name: ADDM:1273301896\_1\_653" and "Period Duration: 11.4 minutes". A table lists findings with their impact percentages and occurrences:

Impact (%)	Finding	Occurrences (24 hrs ending with analysis period)
94.5	I/O Throughput	3 of 27
63.7	Top Segments by "User I/O" and "Cluster"	5 of 27
58.2	Top SQL Statements	6 of 27
10.4	Undersized instance memory	4 of 27

At the bottom, there is a section for "Informational Findings" and navigation links: "Database | Setup | Preferences | Help | Logout".



# HOW TO KNOW IF NAND WILL HELP

- Look at the I/O Waits (AWR Report)

## Wait Classes

- Ordered by absolute value of 'Diff' column of '% of DB time' descending

Wait Class	% of DB time			# Waits/sec (Elapsed Time)			Total Wait Time (sec)			Avg Wait Time (ms)		
	1st	2nd	Diff	1st	2nd	%Diff	1st	2nd	%Diff	1st	2nd	%Diff
User I/O	86.15	16.97	-69.18	63.12	956.01	1,414.59	16,567.80	639.04	-96.14	750.90	2.48	-99.67
Concurrency	0.21	59.24	59.03	3.74	795.23	21,162.83	40.46	2,230.61	5,413.12	30.96	10.41	-66.38
Other	0.01	3.02	3.01	153.94	367.99	139.05	1.17	113.57	9,606.84	0.02	1.15	5,650.00
Application	0.45	3.38	2.94	0.92	34.73	3,675.00	85.98	127.40	48.17	268.69	13.62	-94.93
Commit	3.97	1.75	-2.23	1.56	65.03	4,068.59	764.41	65.76	-91.40	1,397.47	3.75	-99.73
System I/O	2.01	0.61	-1.40	5.53	62.28	1,026.22	386.78	22.86	-94.09	200.20	1.36	-99.32
Network	0.00	0.08	0.08	288.22	12,258.29	4,153.10	0.09	2.95	3,177.78	0.00	0.00	0.00
Configuration	0.00	0.05	0.05	0.00	0.07	100.00	0.00	1.83	100.00	0.00	96.33	100.00



## A PERFECT FIT WITH ORACLE ASM

FUSION-io

- NAND Flash drives appear as raw devices
- Oracle ASM can consume the NAND raw devices to create Disk Groups and Failure Groups
- Use of ASM provide great use cases for I/O balancing
- ASM provides great levels of redundancy
- Points of failure are minimized



# SINGLE INSTANCE ORACLE IMPLEMENTATION

FUSION-iO

- Can use with ASM or without ASM
- Pure Flash implementation
  - Place entire database on Flash
- Hybrid Flash implementation
  - Place redo logs on Flash
  - Place specific tables/indexes on Flash
  - Use Partitioning to locate a portion of data of a table to Flash



# SINGLE INSTANCE ORACLE IMPLEMENTATION

FUSION-iO

- Oracle Smart Flash Cache
- Dell showed a 45 times (4400%) response time improvement
- <http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/wp-oracle-fusion-io.pdf>
- Must have Oracle 11.2 and OEL/Solaris



# ORACLE RAC IMPLEMENTATION


- Oracle Smart Flash Cache
- Kaminario - k2 Solution

FUSION-iO



# HIGH AVAILABILITY - ORACLE DATA GUARD

FUSION-IO



- 1 Redo is transmitted from Primary to Secondary by Redo Transport Services
- 2 Apply Services update standby Redo Log files
- 3 The Database Writer process updates Primary Database files
- 4 If Standby is not reachable, the Redo data gets archived at the primary, so that it can re-sync once the Standby is reachable again



# DATA GUARD PROTECTION MODES

FUSION-io

## MAXIMUM PERFORMANCE

- ASYNC transport mode is used
- LGWR (Log Writer) never waits for an acknowledgement from the Standby database
- Virtually no performance impact on the primary

## MAXIMUM PROTECTION

- SYNC transport mode is used to provide Zero Data Loss
- This mode does not use the Net\_timeout variable. So, if there is no connectivity between the Primary and the Standby, the Primary will stall (waiting for acknowledgement)





# DATA GUARD PROTECTION MODES

FUSION-io

## MAXIMUM AVAILABILITY

- SYNC transport mode is used to provide Zero Data Loss
- If the Primary database can not connect with the Standby database, the Primary will continue without waiting for an acknowledgement (A variable, Net timeout controls the duration for which the primary will wait before proceeding)
- Once the Standby comes back, Redo is synced and SYNC mode is enforced again



# TYPES OF STANDBY DATABASES

FUSION-io

## PHYSICAL STANDBY

- REDO Apply is used to keep the Standby database updated
- The standby database is an exact, block-by-block replica of the primary database
- A Physical standby server is available for read-only workloads. So a Physical standby database can be used for Reporting purposes. The functionality is called Active Data Guard.
- Backups can now be taken from the Standby database as opposed to the Primary database



# TYPES OF STANDBY DATABASES

FUSION-iO

## LOGICAL STANDBY

- SQL Apply is used to keep the Standby database updated
- The Logical Standby Process (LSP) reads the Standby Redo Log and creates SQL statements that are then applied to the Standby database
- A logical standby database is available for Reads AND Writes (to tables that are independent of the primary)
- Local Indexes and Materialized views can be created to enhance reporting performance
- SQL Apply is really useful for Rolling upgrades



# AUTOMATIC & TRANSPARENT FAILOVER

FUSION-io

## OBSERVER

- A server that monitors the Primary Server and the Standby Servers

## DATA GUARD BROKER

- A management interface to centrally manage the entire Data Guard configuration
- It consists of:
  - Background processes on each database
  - A set of configuration files
  - A command line interface called DGMGRL

## FAST START FAILOVER

- Brings the automation into the picture
- Avoids the “split-brain” scenario



# AUTOMATIC & TRANSPARENT FAILOVER

FUSION-io

## TRANSPARENT APPLICATION FAILOVER (TAF)

- With TAF, when an existing session detects an error that would result in a disconnect, it will automatically redirect the session to another instance governed by configuration settings
- TAF can failover running Select statements without any break in continuity
- <http://uhesse.wordpress.com/2009/08/19/connect-time-failover-transparent-application-failover-for-data-guard/>



## ADDITIONAL VALUE

FUSION-io

### DATA GUARD BENEFITS EXTEND BEYOND HIGH-AVAILABILITY

- Offload reporting from primary
- Offload backups from primary
- Ensure that the physical corruption on the Primary does not propagate to the Standby database
- Perform maintenance and Upgrades without downtime
- Test on production data without impacting production



## some USE cases

FUSION-io

- A Primary Server with 1 Standby
  - A Physical Standby in Max. Avail. mode (same location)
- A Primary Server with 2 Standbys
  - A Physical Standby in Max. Avail. mode (same location)
  - A Logical Standby in Max. Perf. mode (same location)
- A Primary Server with 2 Standbys
  - A Physical Standby in Max. Avail. mode (same location)
  - A Physical Standby in Max. Perf. mode (remote location)
- A Primary Server with 2 Standbys
  - A Physical Standby in Max. Protection mode (same location)
  - A Physical Standby in Max. Protection mode (same location)



# Fusion-io Oracle Case Studies

FUSION-IO



Premier online marketplace for independent sellers of new and used books, music, and movies (NON-HP,OLTP)

---



Provider of marketing data gathering and analysis solutions (HP,OLAP,OLTP)

---

**PIPKINS**

Supplier of workforce management software for commercial call centers (HP,OLTP)

---



Provides the dataConductor.com hosted data analysis solution for the largest semiconductor companies (NON-HP, OLAP)

---


**20+ case studies at <http://fusionio.com/casestudies>**





6<sup>x</sup>

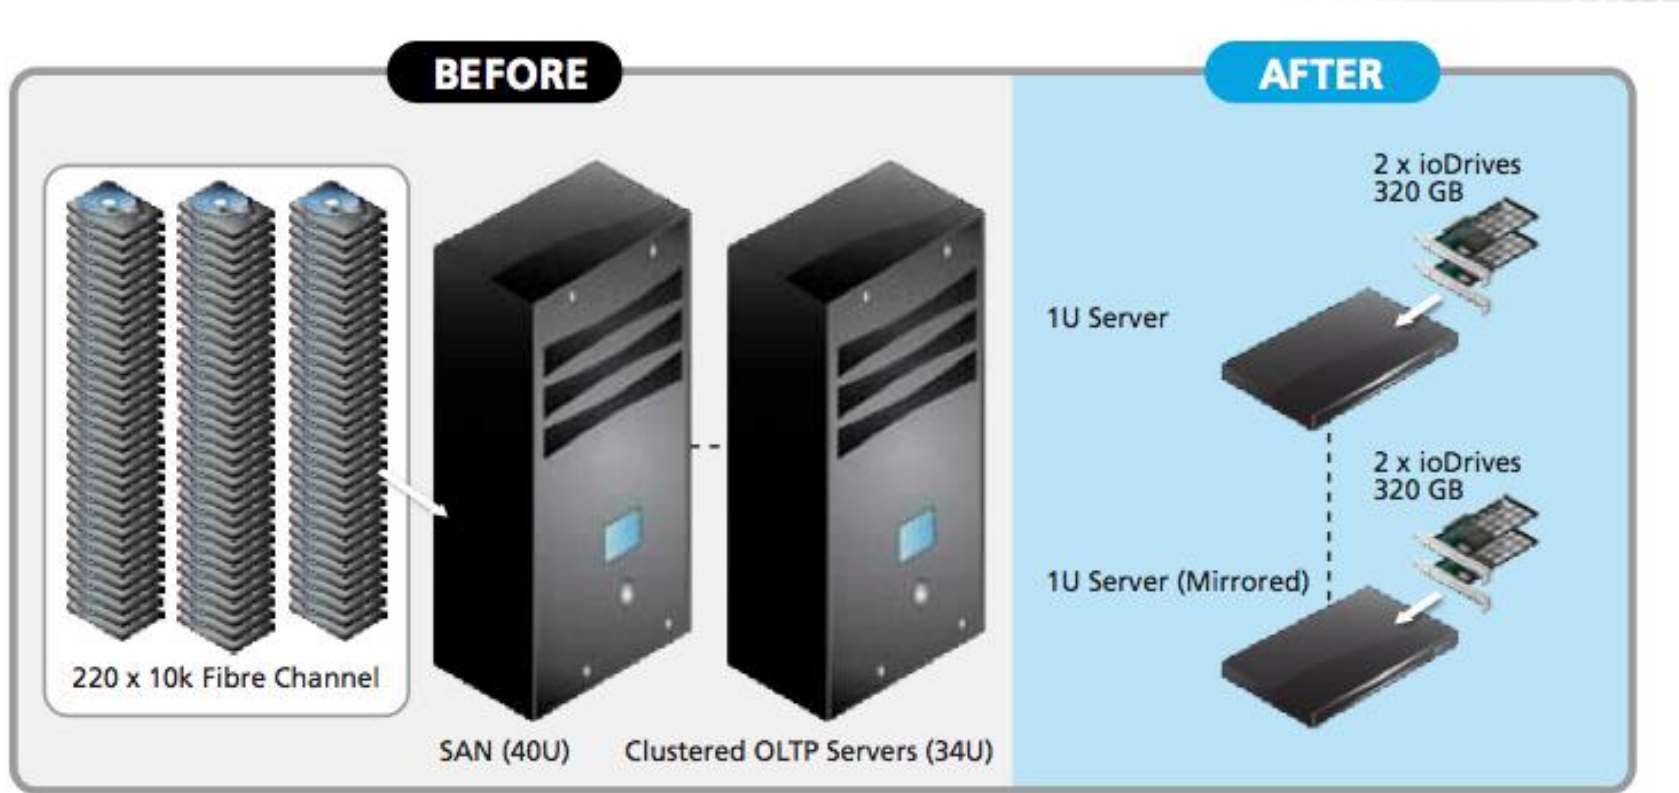
FASTER SEARCHES



“Previously, any searches that had to go back to disk frequently resulted in a lost sale. Now we cache all data in either RAM or Fusion-io. Search speeds have improved over six times, eliminating slow response as a cause of customer abandonment.”

John Shearer,  
Director of IT







# DATALOGIX CASE STUDY

FUSION-iO

# 4<sup>x</sup>

FASTER RESPONSE TIMES

# 40<sup>x</sup>

FASTER DATA WAREHOUSE QUERIES

“Our rapid growth requires massive scalability. The ioDrives’ linear scaling allows us to double our performance capabilities with a single card...”

Dave McCabe,  
Director of Data Center Operations


DataLogix Network  
Prospect. Retarget. Retain.





# DATALOGIX WEB APP SAVINGS

FUSION-io





# PIPKINS CASE STUDY

FUSION-io

# 6<sup>x</sup>

IMPROVEMENT IN DATA LOAD

ELIMINATED DATABASE  
QUEUING  
AND WAIT TIME

**“Our cloud customers are extremely cost-conscious but not at the expense of performance. Fusion-io hit our price point and exceeded our performance needs.”**

Joel Gilbert,  
CIO



**PIPKINS**



5-10<sup>x</sup>

AVERAGE PERFORMANCE

10-30<sup>x</sup>

PERFORMANCE UNDER  
HEAVY LOAD

**“The ioDrives improved application performance by five to ten times nominally, and ten to thirty times on the extreme end. This performance raises the bar on the size of the data sets that our customers can effectively work with and allows us to give them a competitive advantage.”**

Tim Lewis,

VP of Operations





**BEFORE**




30 x 7200K FC Disks




5U Sun V490

**AFTER**



4 x ioDrive Duos  
640GB



4U Dell R910



**Thank you!**

**Questions?**

[sumeet@fusionio.com](mailto:sumeet@fusionio.com)