

ORACLE®

Maximize your Engineered Systems

Jan Poos
Principal Sales Consultant
EMEA Disk BU
@JanStorage

ORACLE
OPEN
WORLD

October 25–29, 2015
San Francisco

ORACLE

Copyright © 2015, Oracle and/or its affiliates. All rights reserved. |

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Program Agenda – Maximize your Engineered Systems

- 1 ➤ ZFS Storage Appliance
- 2 ➤ Database Protection, Provisioning and Offloading
- 3 ➤ Oracle Direct NFS over RDMA
- 4 ➤ Q&A

Oracle's **Storage** Strategy

Deliver the fastest and most efficient storage system for heterogeneous datacenter environments

Engineer advancements that makes Oracle software run fastest and most efficiently

Oracle's Enterprise Storage Portfolio

Engineered for Data Centers. Optimized for Oracle Software.

Engineered Systems




Exadata Exalogic SuperCluster Big Data Appliance

Converged for Ultimate Performance

Storage for Data Transformation

FS1



Engineered for Flash Management & Storage Consolidation

Multi-Purpose Enterprise Flash Storage for SAN

ZS3/ZS4



NAS for High Performance Databases & Files

High Performance Storage Very Low TCO

StorageTek



Tape for Deep Archive

Highest Scalability and Lowest Cost Archive




ZFS STORAGE APPLIANCE - MODELS

Oracle's ZFS Storage Family



ZFS Storage OS

Most powerful storage software suite
Co-Engineered with Oracle software



ZS3-2



ZS4-4

2x Performance
OS8.3 Storage OS Support

- Single or Dual Controllers
- 512GB or 1TB DRAM
- 8 PCIe Slots
- 12TB Read Flash
- 4TB Write Flash
- 3 PB scalability (16 DE's)
- 32 CPU cores

- Single or Dual Controllers
- 3TB DRAM
- 11 PCIe Slots
- 12TB Read Flash
- 10TB Write Flash
- 6.9 PB scalability (36 DE's)
- 120 CPU cores

ZS | Comprehensive Suite of Enterprise Software

Data protocols

- OISP v1.1
- Fibre channel
- iSCSI
- Infiniband
 - NFS/RDMA
 - IPoIB
 - iSER
 - SRP
- NFS V3 and V4
- SMB/CIFS
- HTTP
- WebDAV
- FTP/SFTP/FTPS
- ZFS NDMP V4

Data services

- Encryption
- Hybrid columnar compression
- Hybrid storage pools
- Single, double and triple-parity RAID
- Mirroring and triple mirroring
- End-to-end data integrity
- Local and Remote replication
- Snapshots and clones
- Quota(s)
- In-line deduplication
- Compression
- Thin provisioning
- Antivirus via ICAP protocol
- Online data migration
- Clustering

Management

- Browser and CLI interface
- Management dashboard
- Hardware/component view
- Role-based access control
- Phone home
- Event and threshold based alerting
- Dtrace analytics
- Scripting
- Workflow automation
- Advanced networking
- Snap Management Utility
- Source aware routing
- ZFS Appliance Monitor



WHY IS IT EFFICIENT AND
ECONOMICAL?

ZS | Engineered for Extreme Performance


Most Horsepower Possible

3TB
DRAM

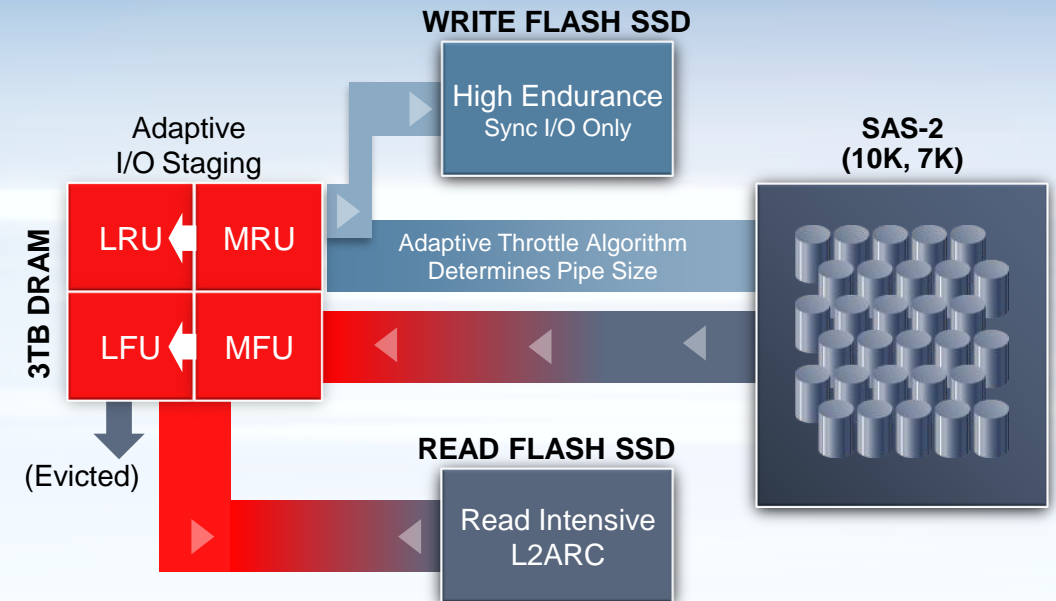
120 Cores
Processing
Power

12.8TB
Read Flash

28TB
Write Flash



Dynamic Storage Tiering (HSP)



- Automated, real-time data migration from DRAM to multi-class flash, to multi-class disk storage
- Software specifically engineered for multi-level flash and disk storage

ZS | Symmetric MultiProcessing


Designed for *Extreme Performance*

20+ years of SMP OS Innovation

- OS uses all 120 cores simultaneously
- Thousands of threads
- Zero CPU bottlenecks
- Enables advanced caching algorithms

Zettabyte File System


- Highest data integrity and massive scale



Hybrid Storage Pool

Dynamically Serve 90% of I/O from DRAM

Introduced first by Oracle on ZFS Storage Appliances in 2008



Oracle ZFS Storage: An Ideal Integrated Solution

Benefits




Performance

World Record Performance
Extreme Network Bandwidth
ZFS Storage Analytics



Capacity

Hybrid Columnar
Compression
Immensely Scalable
Architecture



Simplification

Enhanced Reliability
Simplified Management
EM Integration



Oracle ZS4-4 RMAN Backup Restore Test Results



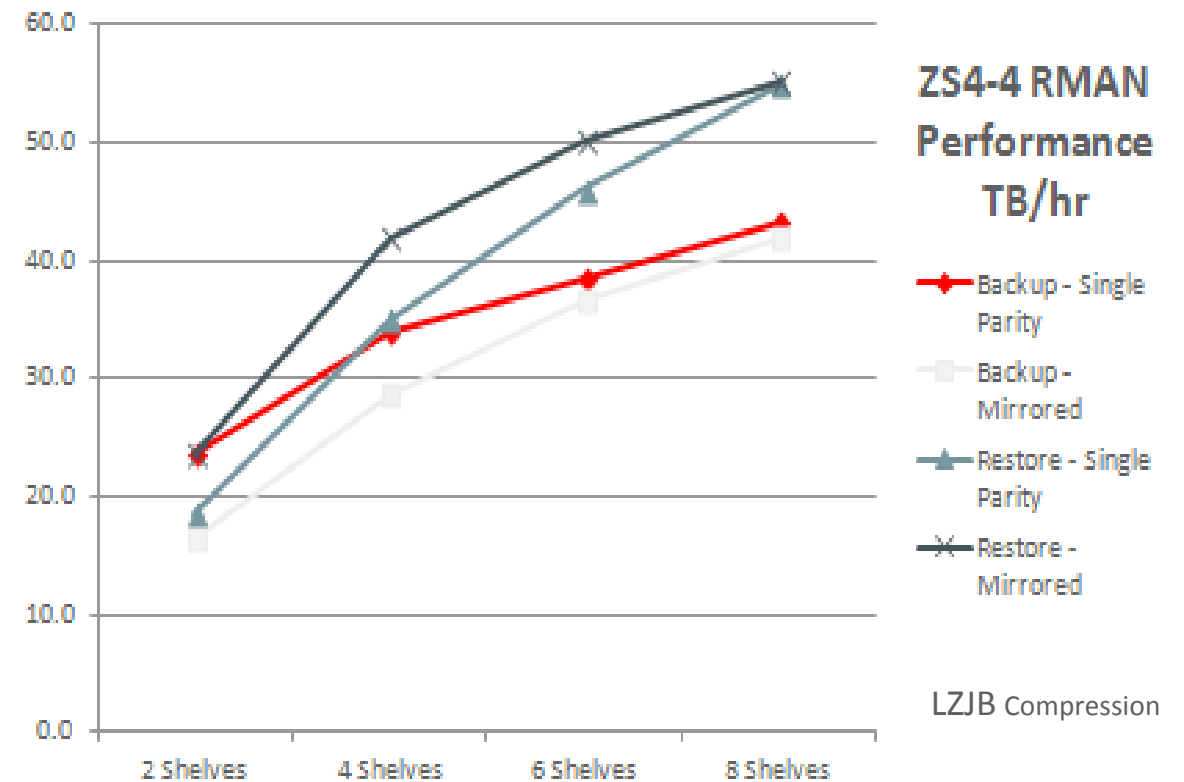
Controllers

- Clustered ZS4-4
- 3 TB DRAM
- InfiniBand Connectivity
- ZFS Storage OS8.4 Software
- No Read Optimized Flash

Enclosures

- 8 High Capacity Disk Shelves
- 7200 RPM drives
- No Write Flash Accelerators

- Peak Sustained Backup: 42 TB/hr
- Peak Sustained Restore: 55 TB/hr



ZFS SA Hardware Compression

Inherited Properties

Mountpoint	<input type="text" value="/export"/>
Read only	<input type="checkbox"/>
Update access time on read	<input checked="" type="checkbox"/>
Non-blocking mandatory locking	<input type="checkbox"/>
Data deduplication (warning)	<input type="checkbox"/>
Data compression	<input type="text" value="LZJB (Fastest)"/>
Checksum	<input type="text" value="Fletcher4 (Standard)"/>
Cache device usage	<input type="text" value="All data and metadata"/>
Synchronous write bias	<input type="text" value="Latency"/>
Database record size	<input type="text" value="128K"/>
Additional replication	<input type="text" value="Normal (Single Copy)"/>
Virus scan	<input type="checkbox"/>
Prevent destruction	<input type="checkbox"/>
Restrict ownership change	<input checked="" type="checkbox"/>

ZFS SA Hardware Compression

BUI VALUE	CLI VALUE	DESCRIPTION
Off	off	No compression is done
LZJB (Fastest)	lzjb	A simple run-length encoding that only works for sufficiently simple inputs, but doesn't consume much CPU.
GZIP-2 (Fast)	gzip-2	A lightweight version of the gzip compression algorithm.
GZIP (Default)	gzip	The standard gzip compression algorithm.
GZIP-9 (Best Compression)	gzip-9	Highest achievable compression using gzip. This consumes a significant amount of CPU and can often yield only marginal gains.

Oracle Direct NFS

- dNFS provides integrated performance tuning when backing up an Oracle database to ZFS Storage
 - Bypasses the operating system
 - Data is cached just once in user space with no 2nd copy in kernel space
 - **New in 12c dNFS over RDMA** provides increased network bandwidth and reduced CPU overhead
- Distribute throughput across multiple network interfaces and stripe buffers over multiple addresses

OISP

- Oracle Intelligent Storage Protocol (OISP) integrated with dNFS in 12c dynamically tunes record size and logbias variables on the Oracle ZFS Storage Appliance to optimize performance

Use IPv4 Protocol

Configure with:

IPv4 AddressMask (192.168.1.2/24)

Use IPv6 Protocol

Interface State up

Interfaces 4/14 available IP MultiPathing Group

aie-test-l71-ibp0
IPv4 static, 0.0.0.0/8, via pffff_ibp0

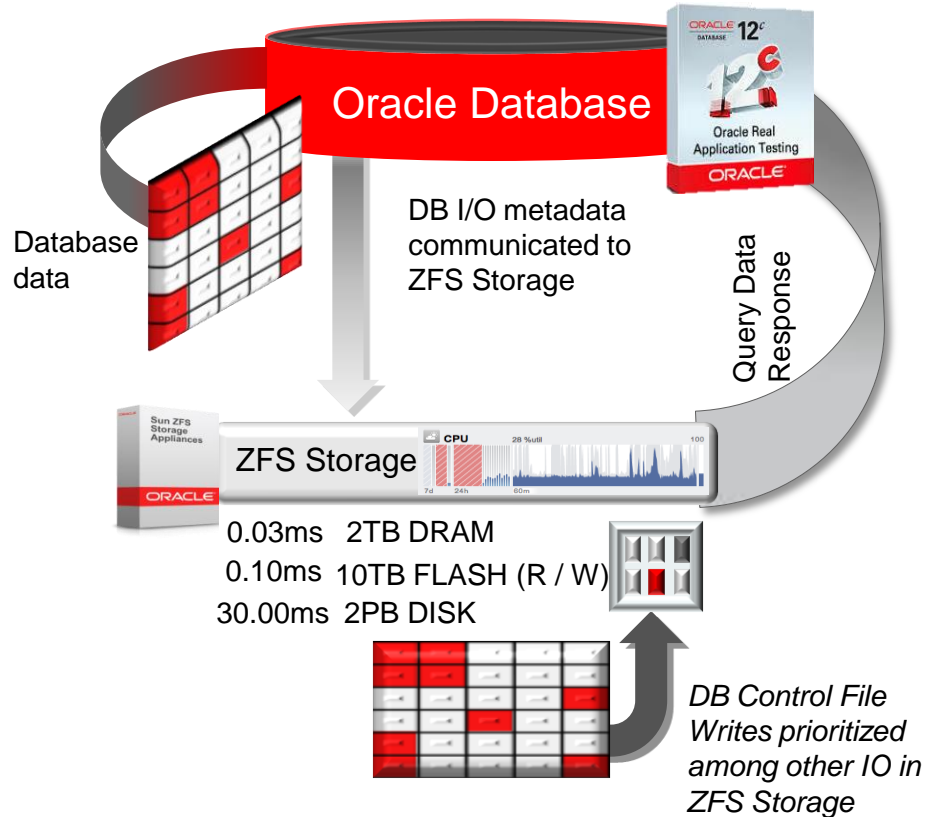
aie-test-l71-ibp3
IPv4 static, 0.0.0.0/8, via pffff_ibp3

Example: Oracle Direct NFS Striping

Number of Addresses	1	2	4
Backup	4209 MB/s	5261 MB/s	5417 MB/s

Oracle Intelligent Storage Protocol (OISP)

Cut Your Database and Storage Tuning Time in Half



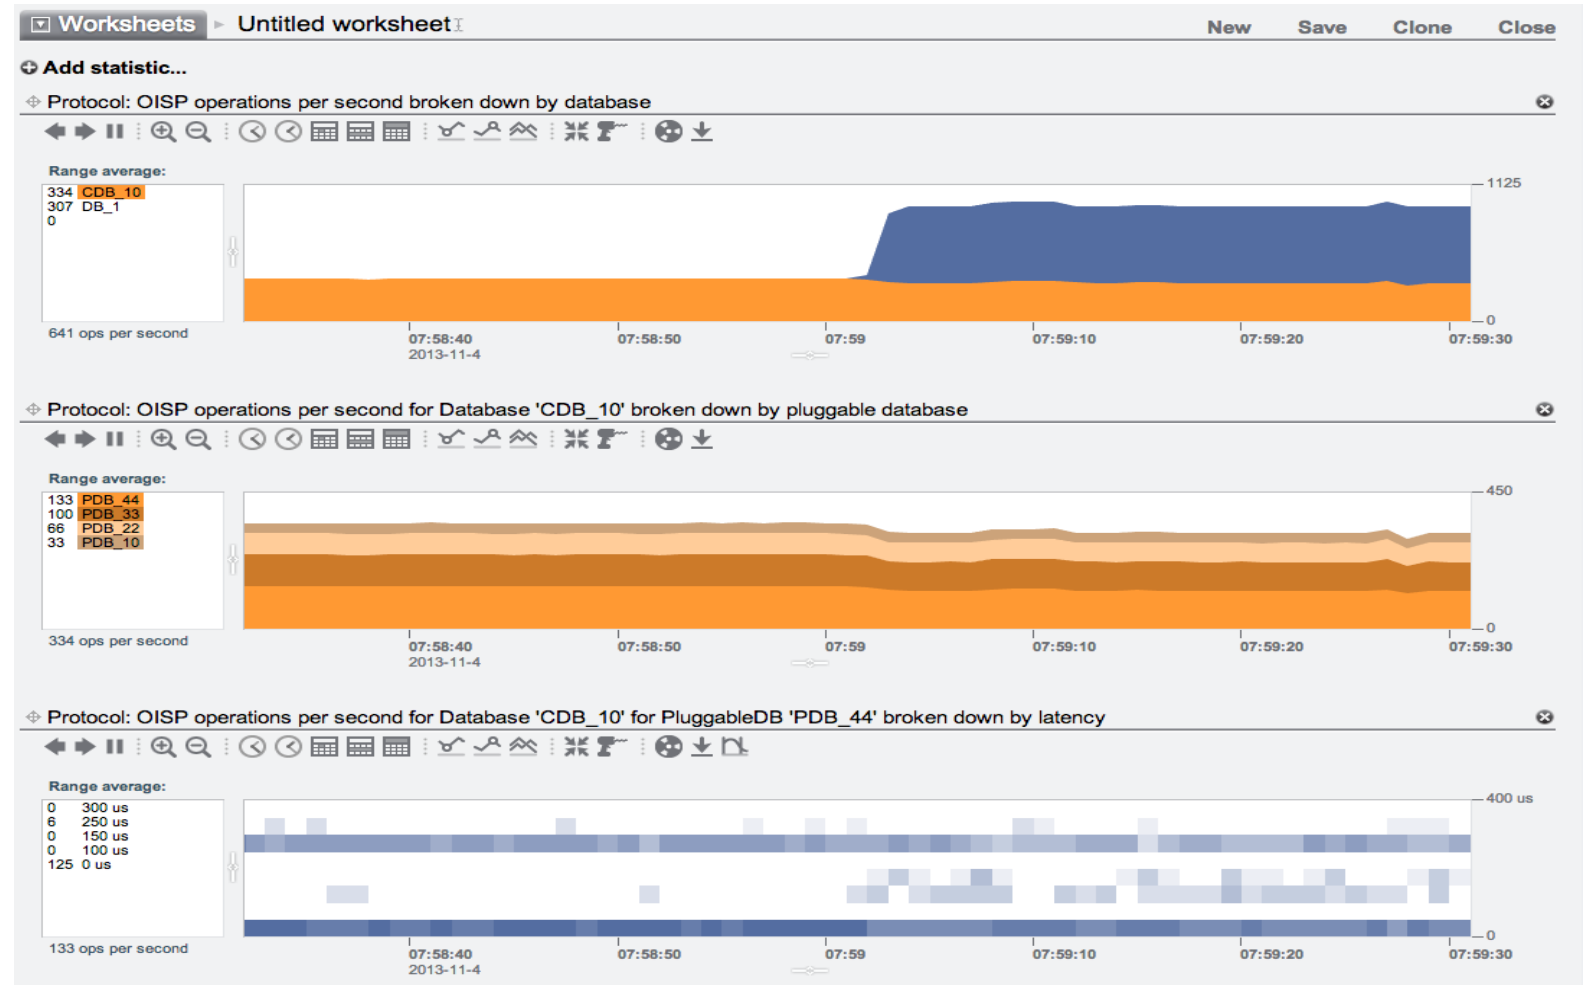
Oracle Intelligent Storage Protocol: Unique language that enables dynamic communication between an Oracle Database and Oracle's ZFS Storage Appliances.

- Available *only* for Oracle Database 12c customers using Oracle Direct NFS (dNFS) with Oracle ZFS Storage Appliances that are running software version OS 8.1 or above

ZFS Storage Appliances dynamically assign system resources to optimize Oracle Database performance and efficiency

Database / PDB Level Analytics

- Enhancements specifically for Oracle Database 12c of our industry leading DTrace-based storage analytics
- Breakdowns of I/O operations by database or PDB name
- Standard DTrace drill-downs of per-database metrics are available (e.g. break downs by latency, operation, client, etc)



RMAN Incremental Backup Strategies

Selecting the Right Strategy for Your Situation

Traditional Incremental Strategy

- Uses **BackupSet** method
- RMAN backupset backup operation to:
 - ✓ Disk - retention from days to weeks
 - ✓ Tape – retention from weeks to years
 - ✓ Cloud – retention from weeks to years
- Smaller backup size than an image copy backup due to RMAN null and unused block compression
- RMAN “`BACKUP BACKUPSET`” to make additional copy of the backup to disk or tape
- **Built-In Backup Deduplication**


Incremental Update Strategy

- Uses **Image Copy** method
- Initial RMAN image copy backup to disk
- Subsequently, backup only the changes (incremental)
- Merge the incremental with the base image
- Block to block copy of database except for temp files
- Allows “`SWITCH TO COPY`” capability
- On-disk retention of 1-7 days
- Synergies with database cloning
- **Built-In Backup Deduplication**

Optimizing for Traditional RMAN Backup Strategy

Preferred Solution Offers Performance and Flexibility

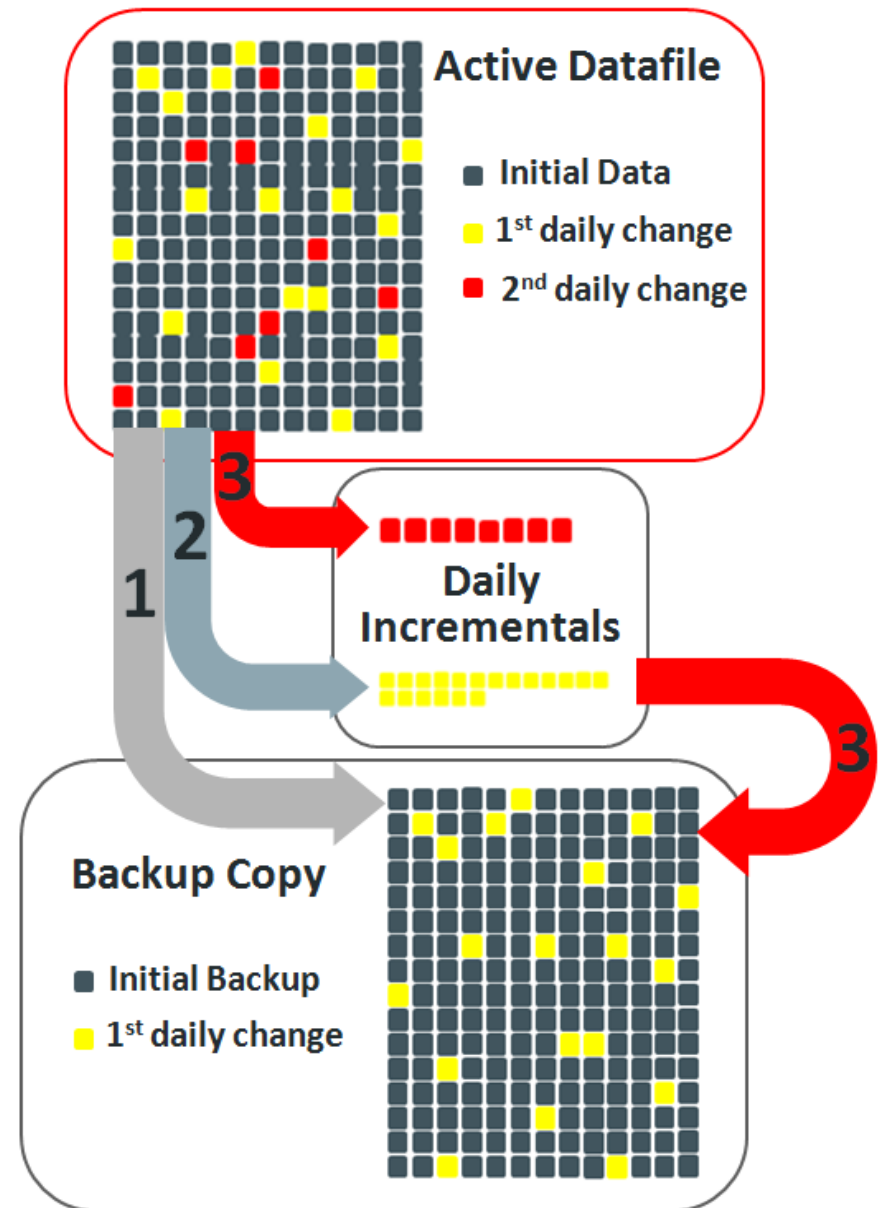
- Single share per storage pool
- 1MB ZFS Record Size
- Synch write bias = throughput
- No read optimized flash
- LZJB share level compression
- Competitive price point
- High performance restore rates satisfy demanding RTOs
- Multiplexing archive logs provides 20mins or less RPOs



Optimizing for Incrementals Forever

Only for Databases with Small Daily Change Rates

- Mirrored storage profile
- Write flash accelerators
- Multi-share configuration
 - Daily Incrementals
 - synch write bias = throughput
 - large record size (1MB)
 - Backup copy
 - synch write bias = latency
 - align record size to average network I/O size



Exadata Backup ZFS Storage Best Practices

Clear Concise Guidelines and Recommendations

- Performance Sizing Guidelines

- RMAN backup/restore with ZS3-2 or ZS4-4

- Selecting a storage profile

- Mirrored, Single Parity or Double Parity

- How to configure Oracle dNFS, IP multi-pathing and IB network


- oranfstab configurations, IPMP groups and integration with Exadata IB infrastructure

- RMAN tunables and example run blocks

- RMAN backup formats, channels, buffers, multi-section support

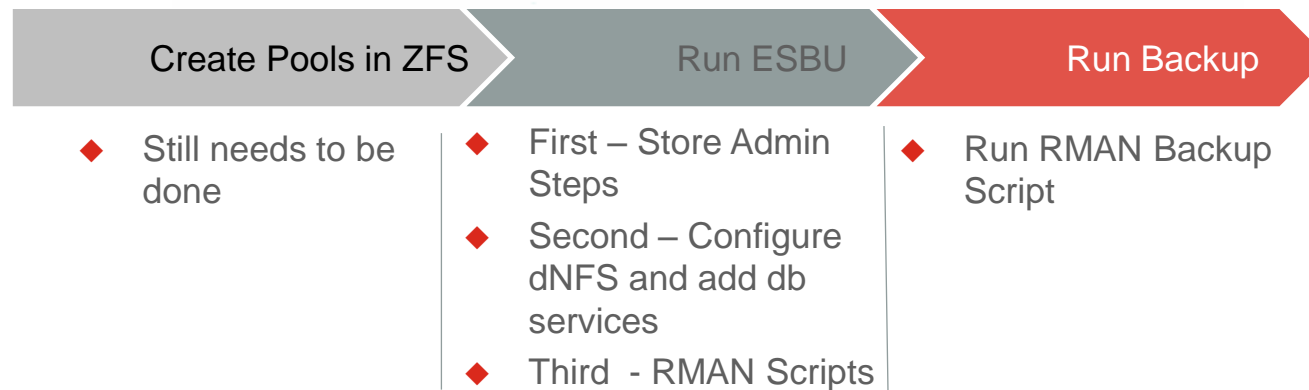

- White paper on Oracle Technology Network

- <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/exadata-backup-zfssa-0715-2620351.pdf>



Oracle ESBU Approach

- Oracle Engineered Systems Backup Utility for Oracle ZFS Storage Appliance
- Almost all done in 3 steps



Oracle ESBU Approach

- Running ESBU from one Exa DB Node

```
[root@ex01dbadm01 esbu]# ./esbu
```

```
This utility configures Engineered System database backup on  
Oracle ZFS Storage Appliance
```

```
The following types of Engineered Systems are supported for database  
backup:
```

0. Exadata
1. SuperCluster
2. Oracle Database Appliance (ODA)

```
Enter the number corresponding to the type of Engineered System database to  
be backed up: 0
```

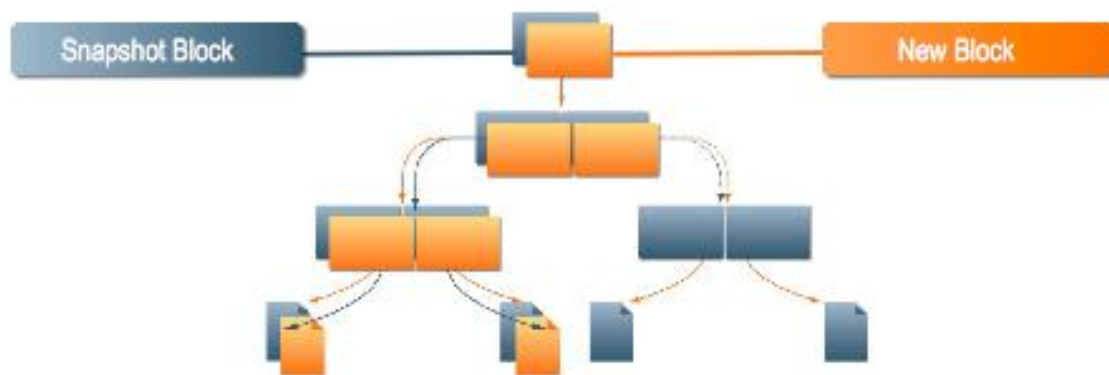
Step 1 Configure Oracle ZFS Storage Appliance

```
In this step a project and shares will be created in the appliance.  
Prior to this step, a storage pool and a network interface must be  
configured first using the BUI of the appliance  
(https://appliance-mgt-ip:215).
```

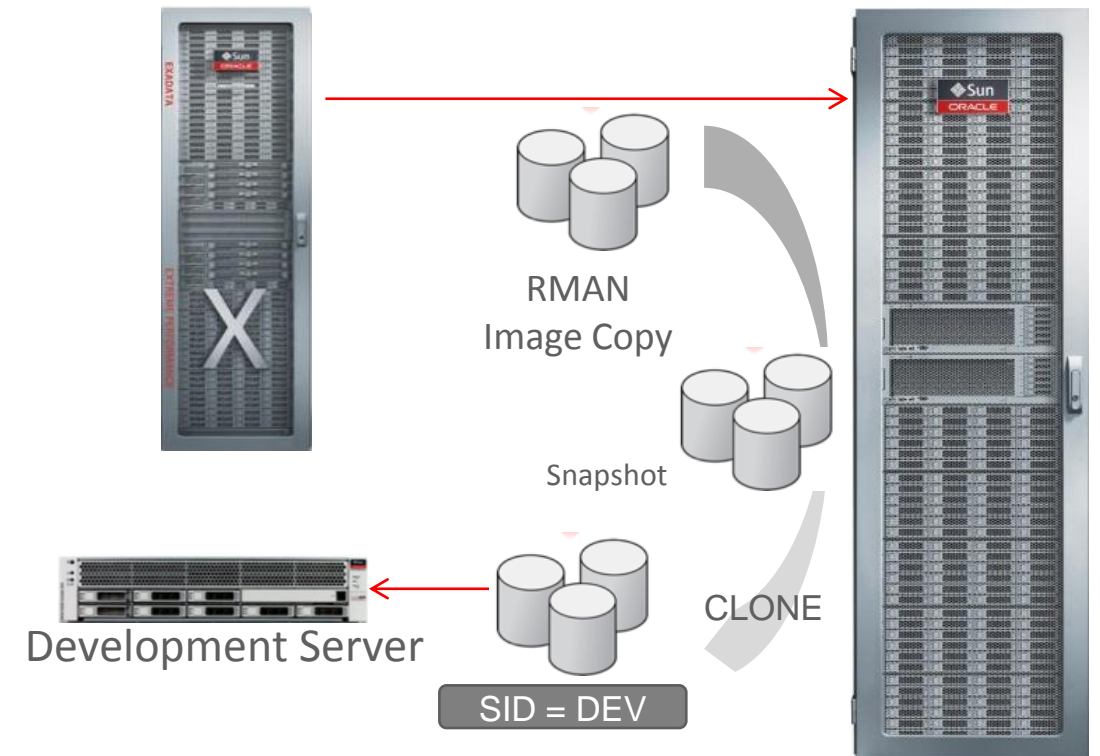
Database Cloning for Dev/Test Provisioning

Oracle Snap Management Utility: Clone from RMAN Backup or DG Standby DB

- ZFS Snapshot
 - Read-only point-in-time copy of file system
 - Only changes tracked
- ZFS Restore
 - Restore to any 'point in time' snapshot
- ZFS Clone
 - Writeable copy of a snapshot ideal for storing many private copies of mostly-shared data



- Database Snap Backup
 - Application aware snapshots based backup of db files/shares



Install SMU

Simple Steps

- ◆ Install oracle-smu rpm file on SMU Server
- ◆ Start, Stop and Status simple:
 - Linux:
`/opt/oracle/smu/bin/smu start|stop|status`
 - Solaris:
`svcadm enable svc:/application/management/smu`
`svcadm disable svc:/application/management/smu`
`svcs application/management/smu`

<http://192.168.56.203:8443/smu>

ORACLE® Snap Management Utility



Login

User Name

Password

SMU Login Main Page

ORACLE Snap Management Utility

About Help Accessibility admin Logout

Workgroup

- Applications
- mydb
- Accounts
- Activity Logs
- Administration

Users

Notification General Settings

Actions View + - Detach

User	Type	Full Name	Directory Server	Directory
admin	Local	Administrator		
pportugal	Local	Paulo Portugal		

Tasks

ID	Application	Task	Description	Start Time	End Time	Status
31	mydb	Clone Backup	clone mydb:mydb_snap_bkp_offline to myclone2	Sep 4, 2015 12:35:50 BRT	Sep 4, 2015 12:37:49 BRT	✓
30	mydb	Clone Backup	clone mydb:mydb_snap_bkp_offline to myclone01	Sep 4, 2015 12:29:50 BRT	Sep 4, 2015 12:32:00 BRT	✓
28	mydb	Backup	backup mydb as mydb_snap_bkp_offline	Sep 4, 2015 12:18:33 BRT	Sep 4, 2015 12:19:27 BRT	✓
27	MYDB	Delete Backup	delete MYDB:snap_mydb_offline	Sep 4, 2015 12:14:59 BRT	Sep 4, 2015 12:15:08 BRT	✓
25	MYDB	Backup	backup MYDB as snap_mydb_offline	Sep 4, 2015 12:10:21 BRT	Sep 4, 2015 12:11:15 BRT	✓



SMU Create a SnapClone

Create the clone DB from a snapbackup

The screenshot displays the Oracle Snap Management Utility (SMU) interface. The main window shows a table of Snap Backups for the application 'mydb'. A dialog box titled 'Create Primary Clone' is open, showing the 'Specify Database Account' step. The dialog box contains the following fields:

Name	Type	Mode	Creation Time	Clones
mydb_snap_bkp_offline	Offline	Manual	Sep 4, 2015 12:19:27 BRT 2	

Create Primary Clone

Steps

1. Specify Database Account
2. Primary Clone Summary

Specify Database Account

- * Host Account: host_databases
- ? * Application Name: myclone3
Type: Oracle Database
- ? * Database Name: myclone3
- ? * SID/SID Prefix: myclone3
- * Listener Port: 1521
Cluster:
- ? * Database Home: /u01/app/oracle/product/12.1.0.2/
User: sys
- * Password: [masked]
- * Confirm Password: [masked]

Buttons: < Back, Next >, Finish, Cancel

Total: 1

SMU Create a SnapClone

New clone created. Almost no space added.

The screenshot shows the Oracle ZFS Storage VM web interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares' (highlighted), 'Status', and 'Analytics'. Below this, there are sub-sections for 'SHARES', 'PROJECTS', 'ENCRYPTION', and 'SCHEMA'. The main content area is titled 'Projects' and 'All Projects'. On the left, there is a 'Usage' section showing 10.7% of 74.5G used, with 7.84G of referenced data and 109M of snapshot data. The main table displays 'Filesystems' and 'LUNs' with 4 total entries. The table has columns for NAME, SIZE, MOUNTPOINT, and ENCRYPTED.

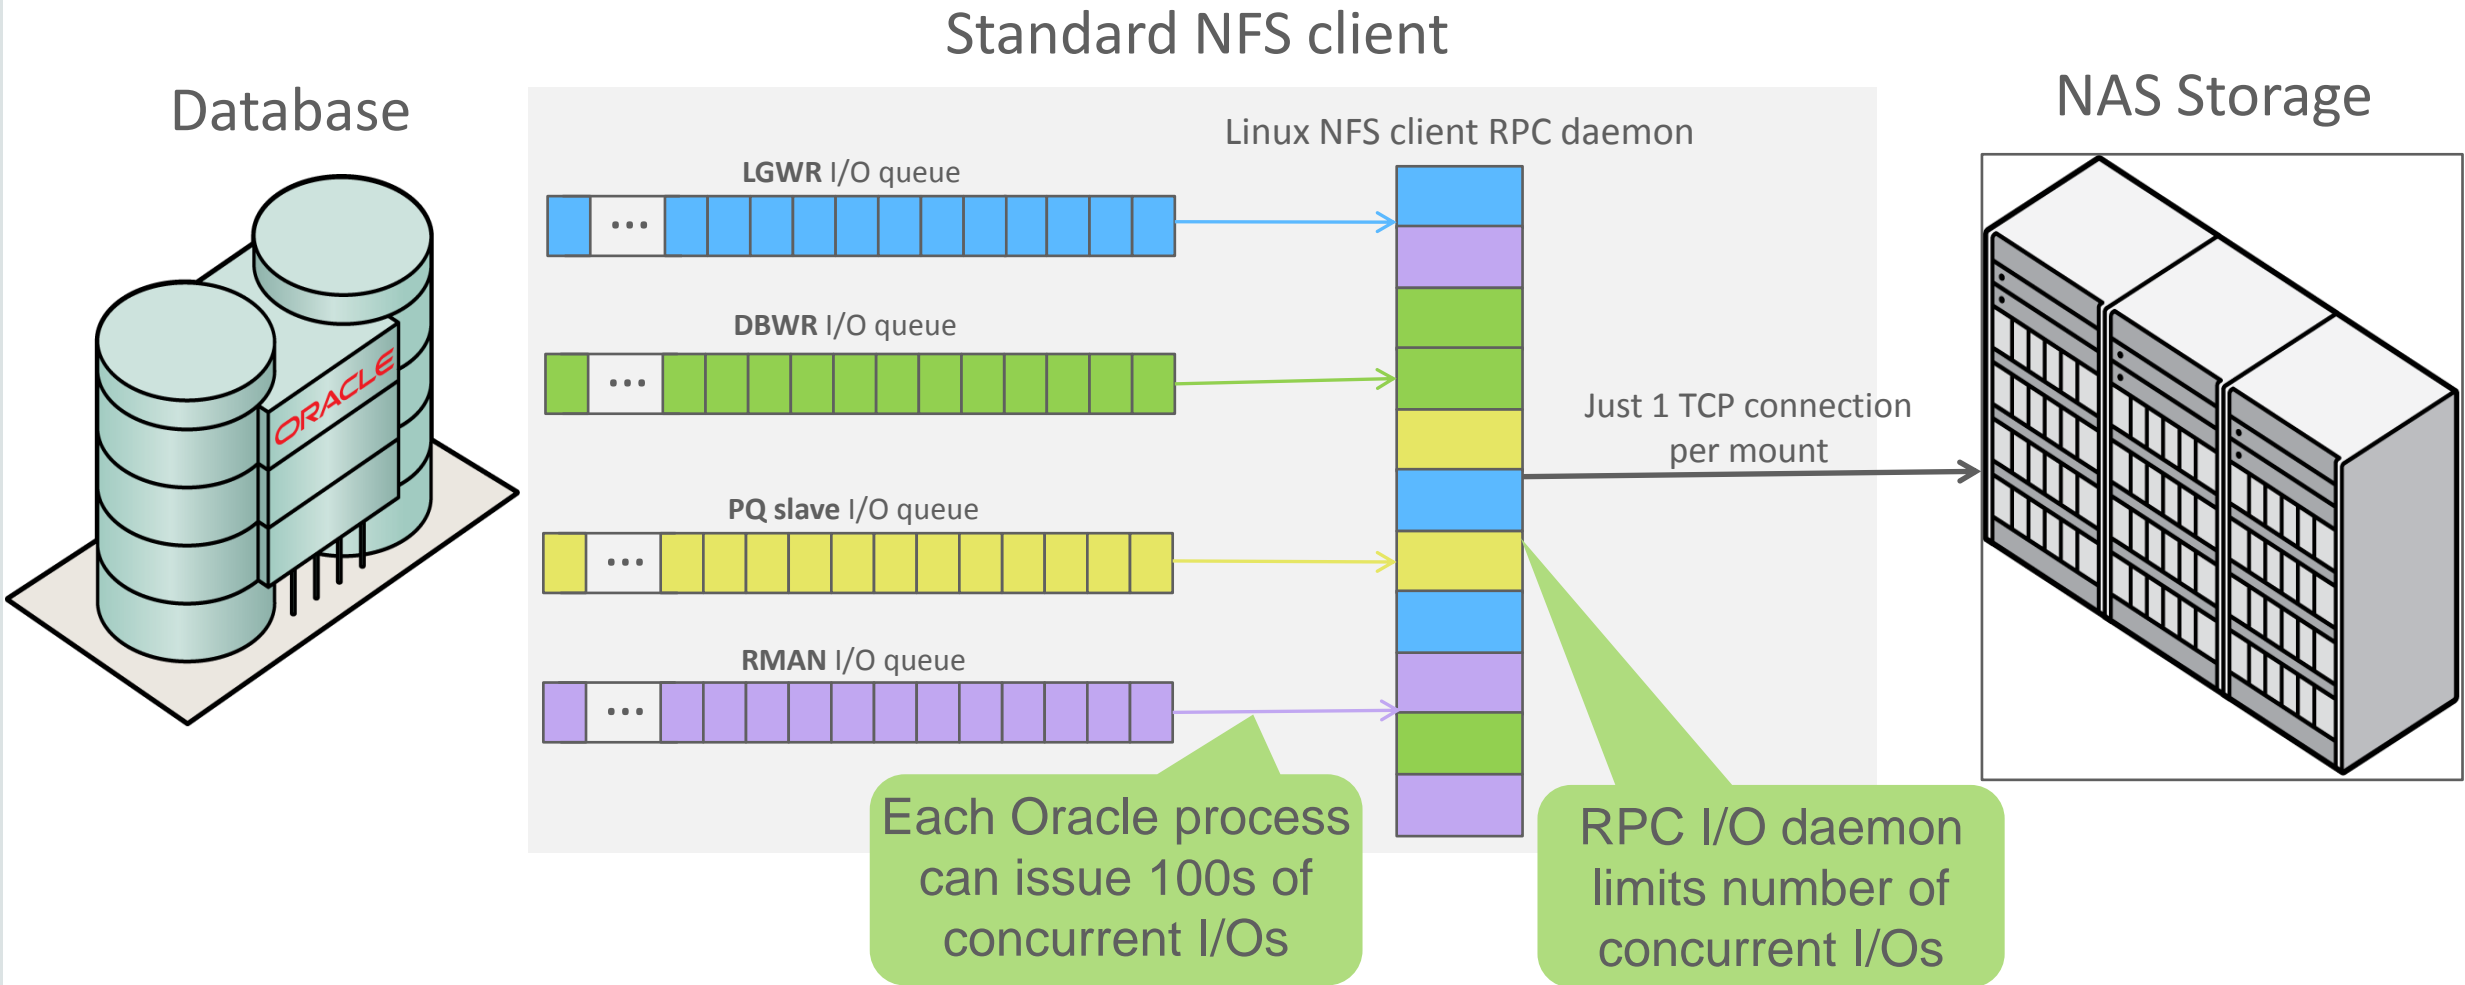
NAME ▲	SIZE	MOUNTPOINT	ENCRYPTED
default /FS1_EXADATA	4.83G	/export/FS1_EXADATA	
default /smu-clone-1441369819088-0	1.03G	/export/smu-clone-14413...	
default /smu-clone-1441370180147-0	1.03G	/export/smu-clone-14413...	
default /smu-clone-1441404732704-0	1.06G	/export/smu-clone-14414...	

Exadata Expansion with the Oracle ZFS Storage Appliance

Storage Offloading to Optimize Space

- Placing tier-2 databases on ZFS Storage
 - Ideal platform to run single instance and RAC databases with less stringent SLAs
- Offload archive logs to ZFS Storage
 - Multiplex archive logs and expire primary copy in the FRA for long-term retention on the ZFS Storage Appliance
- Store RMAN files on ZFS Storage
 - Save space by not placing them in the FRA
- Data refresh and ETL processes
 - Ultra-fast loads, versioning/snapshots, staging and testing copies
 - ZFS SA supports FC, ISCSI, NFS, CIFS etc.

Standard NFS Architecture

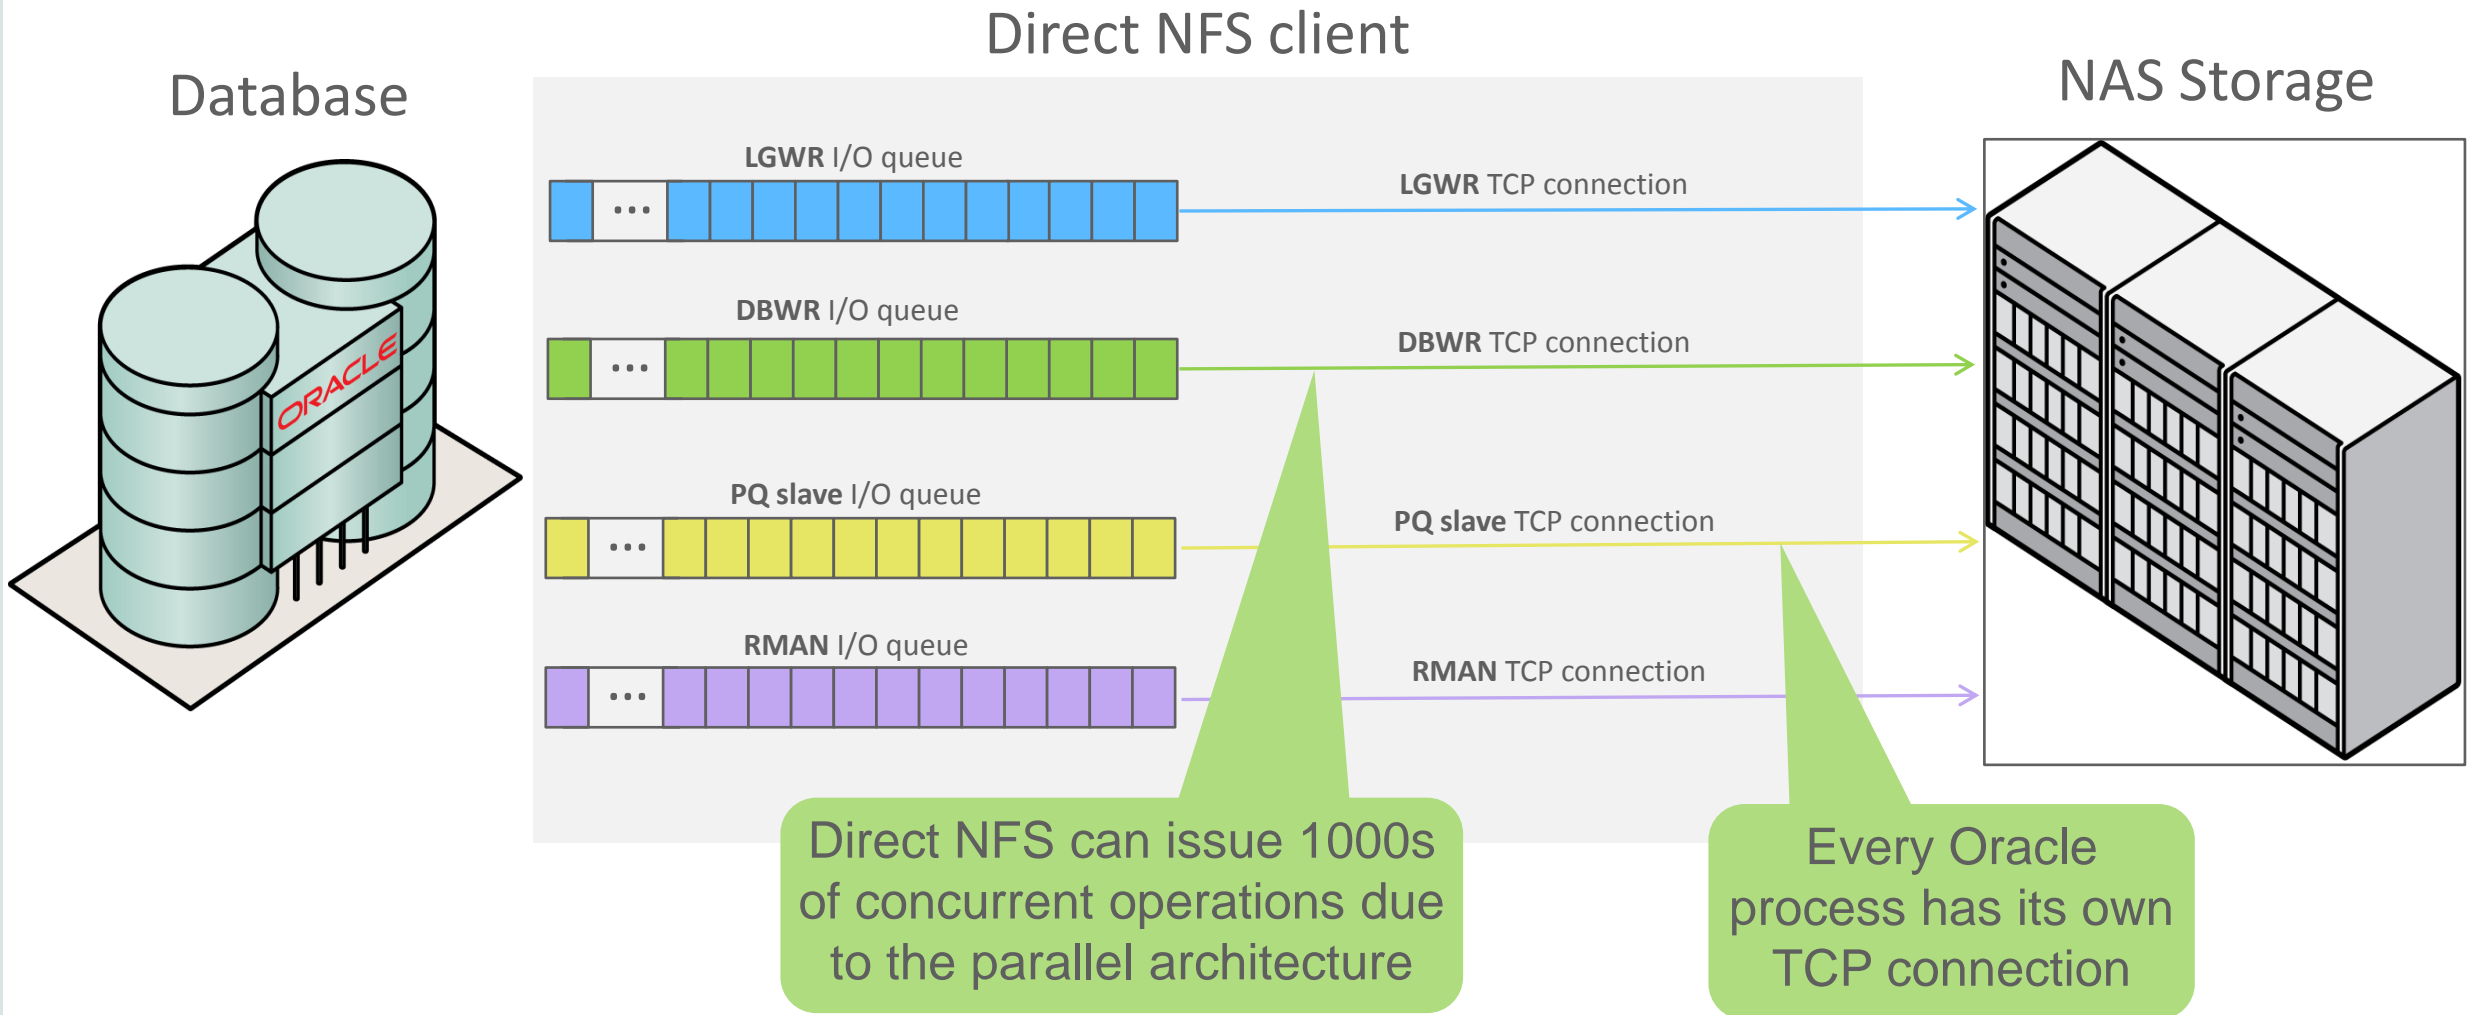


Oracle Direct NFS Eliminates Database NAS Challenges

- An advanced NFS client inside Oracle Database
- Dramatic performance improvements versus host OS NFS clients
- Improves high availability of Database NAS implementations
- Vastly reduces NAS CPU utilization impact
- Simplifies NFS client management uniformly across different host OS platforms


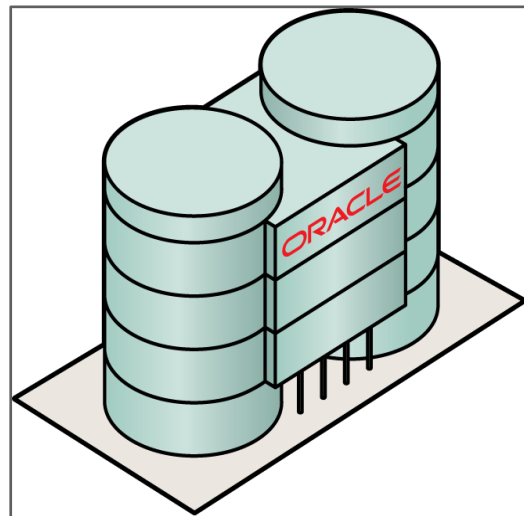
SAN (FC) Performance with NAS Cost and Manageability

Direct NFS Architecture




Direct NFS Scalability and High Availability

- Direct NFS improves **HA** and optimizes server **scalability**
 - Supports up to 4 parallel network paths to storage
 - Automatically load balances across all network paths
 - No extra configuration necessary
 - Failover messages in the event of a path failure



Automatic load
balancing across
all paths

Path 3
goes down



Get the most out of Oracle Database 12c

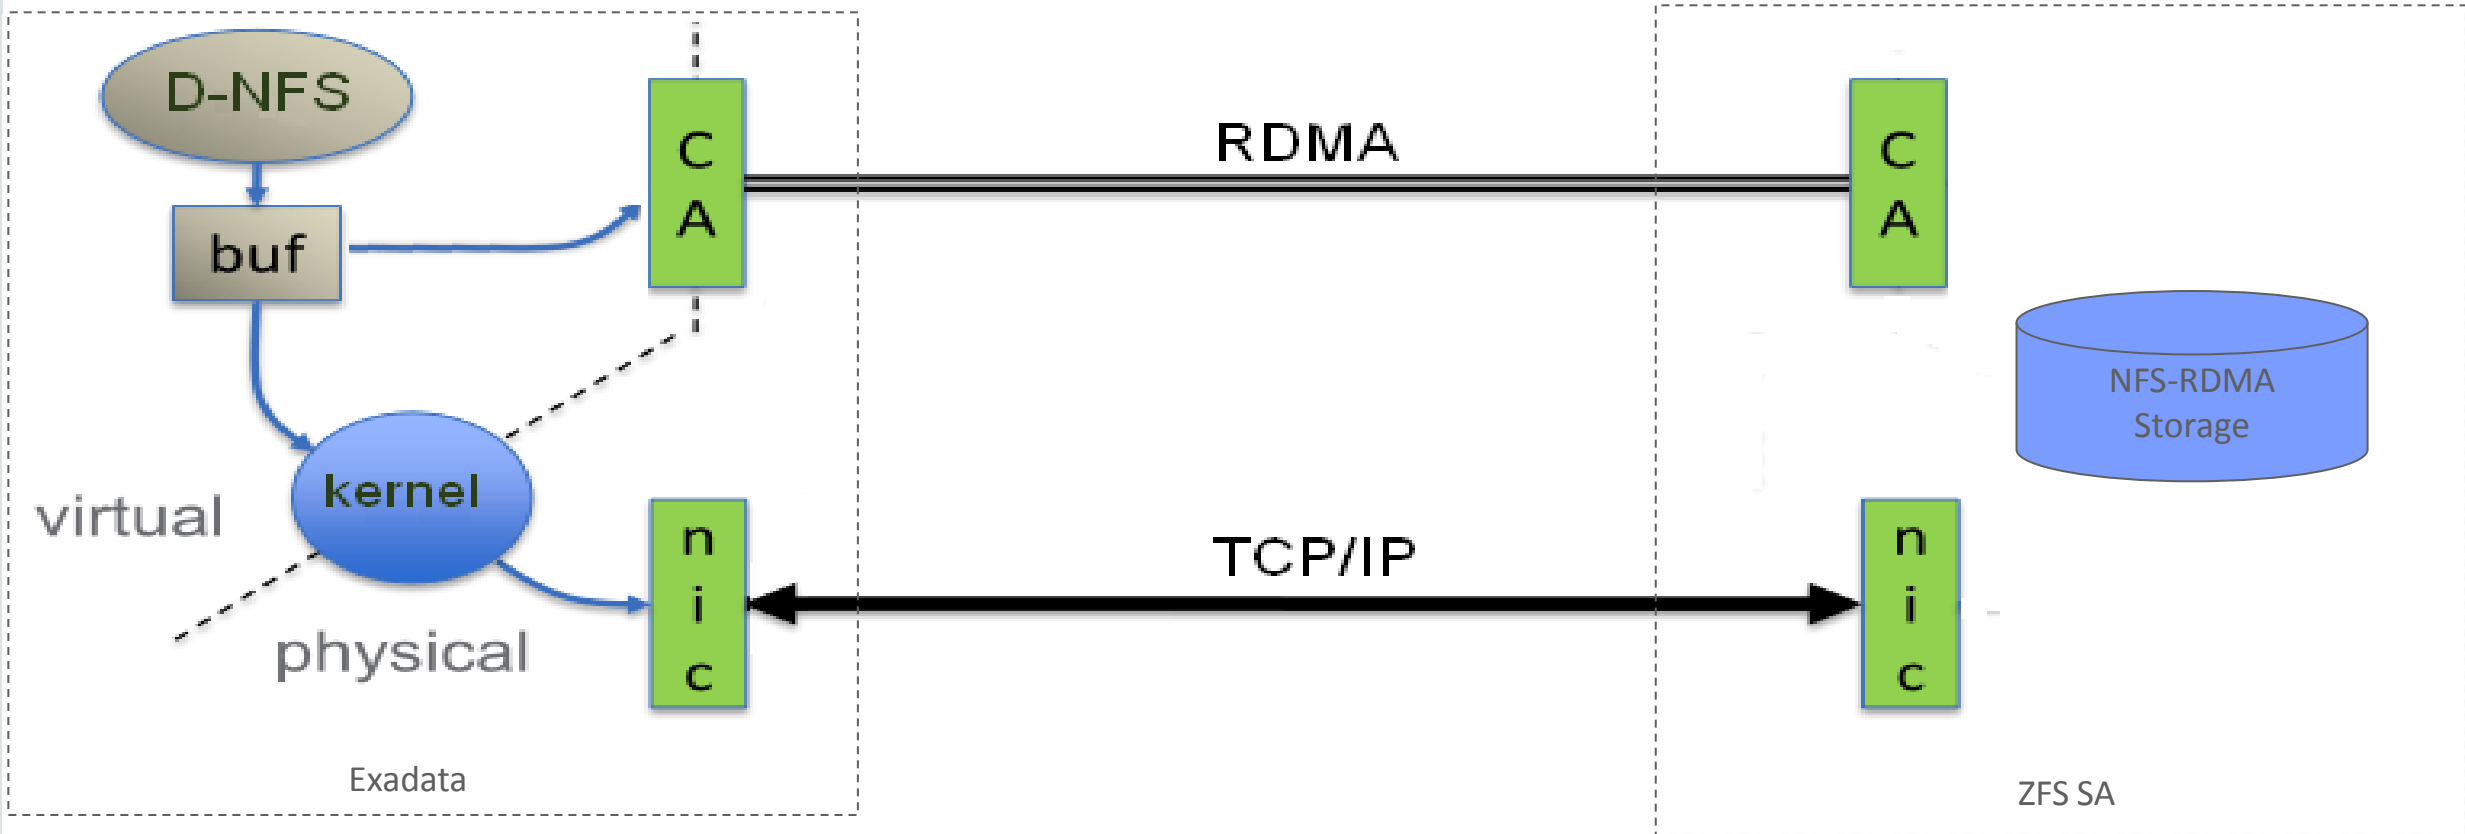
Direct NFS with Remote Direct Memory Access (RDMA)

- New in Oracle Database 12c; Available in Exadata 12.1.2.3
- Retains all Direct NFS benefits
- Implements NFS RDMA protocol over Infiniband
- RDMA zero copy transfers eliminate buffer copies resulting in memory bus savings
- RDMA hardware offloading frees up main system CPUs from networking tasks
- Drastically reduced CPU utilization and improved response times
- Increased bandwidth utilization due to massively parallel IO operations
- Use Case: RMAN backup to ZFS Storage Appliance
 - Direct NFS over RDMA saturates Infiniband QDR (40Gb/s) pipes
 - High performance backup solution for Exadata systems

ORACLE[®] 12^c
DATABASE

ORACLE[®]
ZFS STORAGE
APPLIANCE

DNFS/RDMA: HYBRID ARCHITECTURE



- Based on IO size, DNFS/RDMA decides on best IO path
- Control ops (e.g. MNT) are always sent over tcp

DNFS/RDMA: PERFORMANCE

	D-NFS/IPoIB	D-NFS/RDMA	
Total throughput (mb/sec)	5897	8853	Improvement of over 50%
RMAN IO avg wait (ms)	23.24	9.16	Improvement of over 60%
Avg CPU utilization (%)	51.625	41.125	Improvement of over 20%

Environment: Exadata 12.1.2.1.2, X3-2 full rack with 2 IB channels per DB node / 8 node RAC;
Oracle Sun ZFS 7420 Storage 2013.06.05.4.2.1,1-1.1 64-bit

DNFS/RDMA: VALUE PROPOSITION

- Improve ROI of ZFS Storage Appliance
- Deliver high bandwidth utilization
- Deliver low latency
- Stack bypass
- Copy avoidance
- Reduce CPU utilization
- Reduce memory bottlenecks

Not Discussed today

- ZFS SA and the PCA
- ZFS SA and Sparc Super Cluster
- Database Backup in the Oracle Cloud
- Recovery Appliance (ZDLRA)
- ?

More information

- General information of our storage products:
 - <https://www.oracle.com/storage/index.html>
- OTN White Paper placeholder for ZFS SA
 - <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>
- Latest OTN White Paper for Exabackup:
 - <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/exadata-backup-zfssa-0715-2620351.pdf>
- Oracle ZFS SA Plugin downloads:
 - <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html>

More information

- [Oracle ZFS Storage Appliance: How to Enable Oracle Intelligent Storage Protocol \(OISP\) \(Doc ID 1943618.1\)](#)
- [Oracle ZFS Storage: FAQ: Exadata RMAN Backup with The Oracle ZFS Storage Appliance \(Doc ID 1354980.1\)](#)
- [Snap Management Utility for the Oracle Database - Information and Troubleshooting \(Doc ID 1522925.1\)](#)

A woman with long brown hair and glasses is sitting at a wooden table in a bright, modern office or cafe. She is wearing a brown leather jacket over a blue patterned scarf. She is holding a black smartphone to her ear with her left hand and looking down at a large open book or document on the table with her right hand. In the background, another person is sitting at a table, and there are large windows letting in natural light.

Q & A

ORACLE®