



**ORACLE<sup>®</sup>**

## **Sun ZFS Storage Appliances from Oracle Get More Efficiency from Your Storage**

Jason Weiss

Advanced Principal Support Engineer



# Agenda

- ZFS appliance background/use cases ~20 min
- Demonstration ~20 min
- RMAN Backup recovery architecture drill down ~20 min

# Oracle's Storage Vision

Engineered for Oracle Software / Best of Breed for Mixed Environments



## Seamless Integration

- Oracle applications auto-provision and tune Oracle Storage
- Single system interface for all statistics, alerts, faults, analytics, provisioning
- Dynamic and deterministic QoS driven by Oracle applications

## Highest Performance

- Oracle Database dynamically exchanges I/O requirements with Oracle Storage
- Database query pre-fetch and off-load with optimized data caching
- End-to-end data transport with zero buffer copies over InfiniBand fabrics

## Most Efficient

- Data compression without rehydration from database to disk to tape
- Seamless automated backup and recovery across the entire Oracle stack
- Automated, policy-based tiering across flash, disk and tape

## Most Secure

- Encryption keys auto-generated from application, database, and storage
- Oracle unified security key management across apps, database, disk and tape

ORACLE

# Oracle's **Storage** Strategy

## Run Oracle Software Fastest and Most Efficiently

- Oracle 11g: Up to 50x compression and 5x faster
- Solaris 11: Up to 2x more I/O efficient
- Enterprise Management: 5x lower TCO

## Leverage Oracle Technology

- Enterprise Manager
- Database
- Virtualization
- Operating System
- Servers
- Storage

## Develop Best-in-Class Storage Products

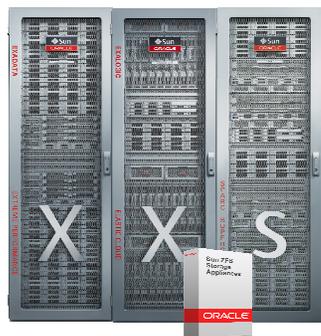
- NAS: Best Performance and Price/Performance
- SAN: Best Multi-Workload & Scalable Performance
- Tape: Best Performance, Price/Capacity

ORACLE

# Oracle's Complete Storage Portfolio

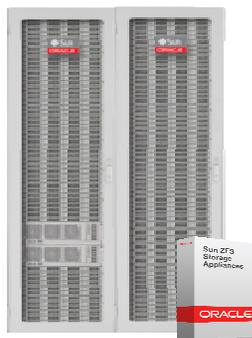
Engineered for Oracle Software / Best of Breed for Mixed Environments

## Engineered Systems



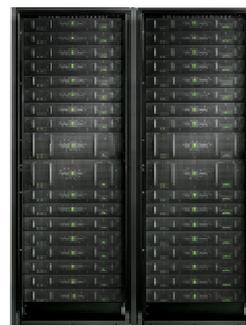
Exadata Exalogic SPARC SuperCluster

## NAS Storage



ZFS Storage Appliances

## SAN Storage



Pillar Axiom 600

## Tape and Virtual Tape



## Flash Storage



SSD

F20 PCIe

F5100 Flash Array

## Storage Software

**Storage Management:** OEM, ASM, DTrace Analytics, ACSLS, ELS  
**Automated Tiering:** Partitions, SAM QFS, Hybrid Storage Pools, VSM  
**Data Reduction:** 11g ACO, HCC, RMAN, ZFSSA Dedup/Comp  
**Data Protection:** Data Guard, RMAN, OSB, ZFSSA Snap/Rep, MaxRep  
**Security/Encryption:** ASO, Oracle Key Manager, Disk/Tape Encryption

ORACLE

# Oracle NAS comes up big—twice



With a lineage that goes back to Sun and StorageTek, Oracle's NAS boxes are meeting, and maybe exceeding, expectations. BY RICH CASTAGNA



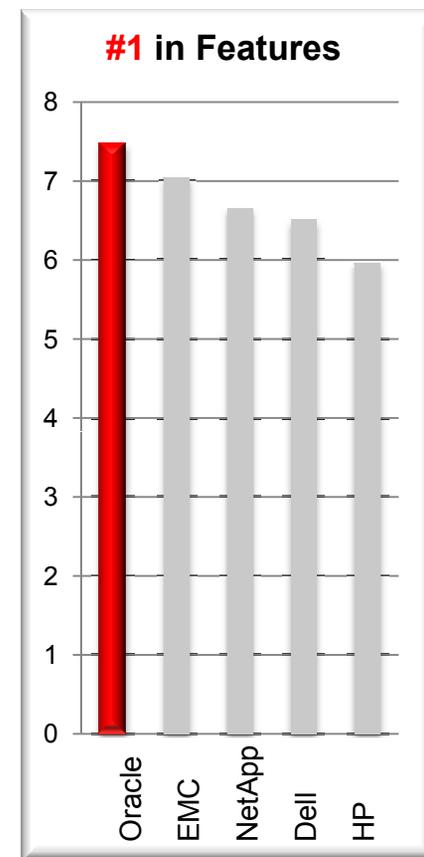
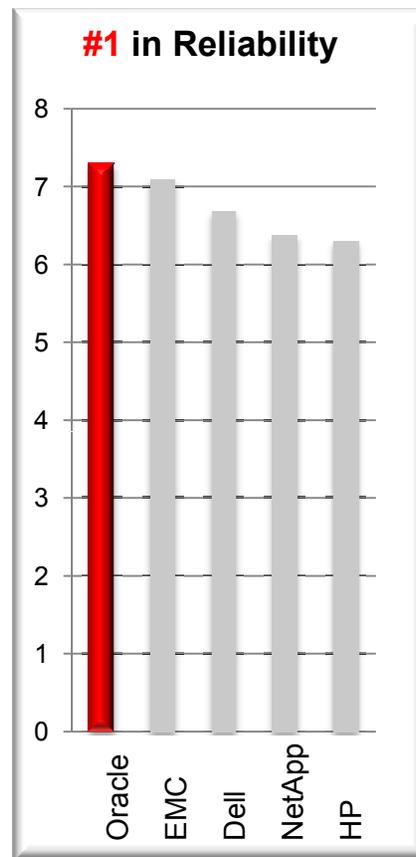
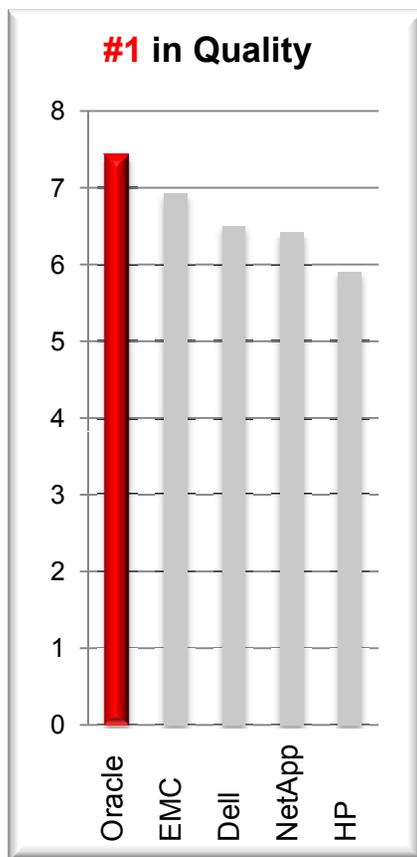
**FILE STORAGE IS** front-page news in the data storage world these days, with talk of "big data" and the relentless growth of end-user information forcing us to focus our attention—and storage infrastructure efforts—on storage systems for file data. The good news is that storage vendors appear to be up to the task, with users generally expressing high satisfaction with their products on the sixth *Storage* magazine Quality Awards for network-attached storage (NAS).

Oracle Corp. storage users are apparently the most satisfied among the legions of NAS users. Their high marks across the spectrum of service and reliability propelled Oracle to the highest ratings among both enterprise and midrange NAS systems.

Over our six NAS Quality Awards surveys, we've seen six vendors garner top honors in the enterprise group and three winners among midrange NAS systems since NetApp Inc.'s earlier dominance when it won the first three Quality Awards. In this year's survey, average ratings for both groups set records with new highs in every category except two.

STORAGE January 2012

## IT Administrators rate their NAS provider . . .



ORACLE

# Oracle IT Runs on Oracle Storage

Over 100 Petabytes of ZFS Storage in Production Across Oracle IT Datacenters

## Results with Oracle's ZFS Storage Appliances

- **3:1** consolidation replacement of existing NAS platforms
- **12x** faster in compute farm workload benchmarks
- **5x** faster in database test suite execution
- **3x** faster with twice the snap/clone data copies in Oracle application environments
- **50%** reduction in tuning and trouble-shooting time

### Global IT

- Global voice, network infrastructure and data center operations
- IT risk and compliance
- Supporting 100K+ internal users and 1.2M+ external users in 145 countries

### Commercial IT Oracle On Demand

- Comprehensive Cloud Products and Services for Business and IT
- Supporting over 5.5M users in more than 4,000 environments
- ~ 3.2 billion peak database transactions per hour; ~ 5.5 million business transactions per day

### Product Development IT

- Product development operations
- Supporting 22K+ developers building over 3K products globally
- 10K hosts executing 90K jobs per day
- 235K compute hours per day

ORACLE

# 3<sup>rd</sup> Generation NAS Storage



**Sun ZFS Storage 7120**  
**Best value**  
Full suite of data services



**Sun ZFS Storage 7320**  
**Best flexibility**  
Single or dual controllers



**Sun ZFS Storage 7420**  
**Best scalability**  
High Availability Environments

## Standard features (all models)

Data protocols: FC, iSCSI, IB, NFS, CIFS, WebDAV, FTP, and more (see data sheet)

Advanced data services\*: Snap, dedup., compression, analytics, and more (see data sheet)

## Clients and applications (all models)

Oracle Solaris • Oracle Linux

Oracle database, middleware, and applications

Oracle VM • VMware • Windows

More than 50 business applications supported

## New benefits

15,000 RPM HDDs now available- advantageous for specific high I/O use cases.

Continued software improvements resulting in better availability and performance.

\* *Replication and Clones are separately licensed features*

ORACLE



### Sun ZFS Storage 7120

Entry standalone

- Up to 60×3TB 7,200 RPM SAS disks
- Or 300GB or 600GB 15,000 RPM SAS disks
- Write flash cache



### Sun ZFS Storage 7320



### Sun ZFS Storage 7420

## Sun ZFS Storage 7120

#### What you get:

- Up to 177 TB HDD capacity
- 4 Core CPU
- 48 GB DRAM cache
- Write flash cache included

#### Value:

Entry level pricing with enterprise class data services and simplified management from 3.3TB up to 177TB to address growing business requirements

#### Uses:

Use for general purpose file sharing, web storage, entry level storage application requirements

Sun ZFS Storage 7120



Sun ZFS Storage 7420



### Sun ZFS Storage 7320 Entry cluster-capable

- Up to 144×3TB 7,200 RPM SAS disks  
or 300GB or 600GB 15,000 RPM SAS disks
- Up to 8×512GB read/16×73GB write flash
- Active-active clustering for H/A

## Sun ZFS Storage 7320

### What you get:

- High availability cluster option
- 2 CPUs (4 cores per CPU) per controller
- Up to 4TB/system of read flash cache
- Up to 1.2 TB of write flash cache
- Up to 144 GB DRAM per controller

### Value:

Entry level modular scalability and high availability up to 432 TB for accelerated business application performance

### Uses:

Use for general purpose file sharing, virtualization, disaster recovery, test, QA, streaming workloads such as video/media

Sun ZFS Storage 7120



Sun ZFS Storage 7320



## Sun ZFS Storage 7420

### What you get:

- Up to 2.59 PB HDD capacity, up to 80 CPU cores/system
- Up to 2TB/system DRAM cache
- Up to 4TB/system read flash cache
- Up to 10.5TB write flash cache
- Up 120Gb/s Ethernet bandwidth connectivity

### Value:

Modular scalability and high availability up to 2.59 PB for accelerated business application performance

### Uses:

Consolidation of virtualization environments requiring multiple data services, heterogeneous file sharing, back up, disaster recovery, and bulk storage

### Sun ZFS Storage 7420 Expandable cluster-capable

Up to 576 × 3TB 7,200 RPM SAS disks  
or 300GB or 600GB 15,000 RPM SAS disks  
Up to 8 × 512GB read/96 × 73GB write flash  
Active-active clustering for H/A

ORACLE

# Industry's Richest Set of Base Data Services

## Data protocols

- Fibre channel
- iSCSI
- Infiniband over IP/RDMA
- iSER
- SRP
- NFS V3 and V4
- CIFS
- HTTP
- WebDAV
- FTP/SFTP/FTPS
- ZFS NDMP V4

## Data services

- Oracle Hybrid Columnar Compression
- Hybrid storage pool
- Single, double and triple parity RAID (RAIDZ, Z2, Z3)
- Mirroring and triple mirroring
- End-to-end data integrity
- Local and Remote replication\*
- Snapshots and clones\*
- Quota(s)
- In-line dedup
- 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
- DFS root support
- Source aware routing

\* Remote Replication and Cloning features are licensed separately. All other features included with purchase of system.

ORACLE

# Sun ZFS Storage Appliances

Engineered for extreme performance.

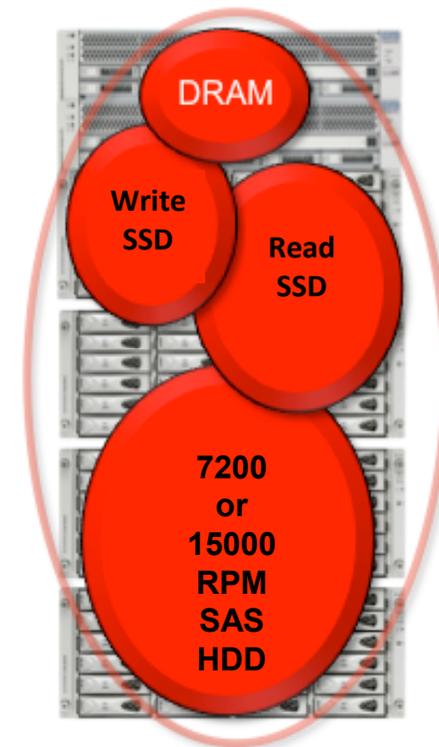
Increase database responsiveness.

Expedite large data queries.

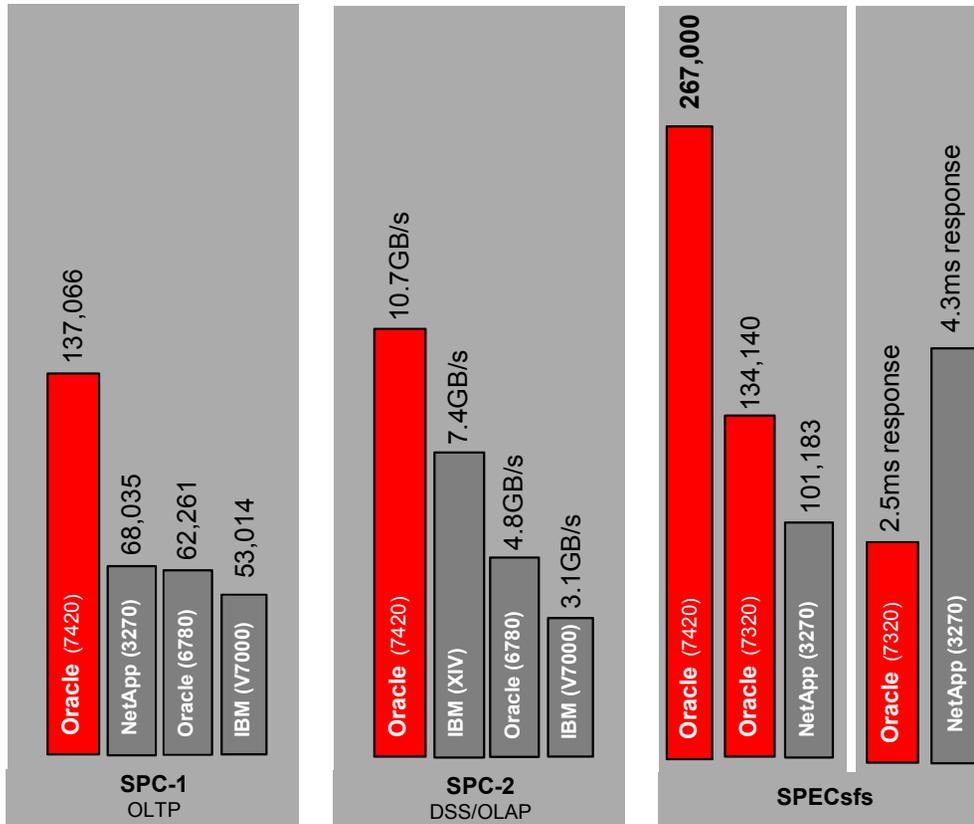
Minimize backup windows.

# Savings and Performance for Consolidation

- Only Storage Appliance with Hybrid Storage Pools (HSPs)
  - Data is intelligently and automatically migrated between DRAM, Flash and Disk
  - Continuously optimizes storage system performance and efficiency
  - Simplifies management, transparently managed as a single storage pool
  - Now available with 15,000 RPM HDDs; advantageous for high I/O use cases involving random, uncached data



# Building a Reputation for Extreme Performance



**ZFS Storage Demonstrates Industry-leading Performance in all Storage Benchmarks!**

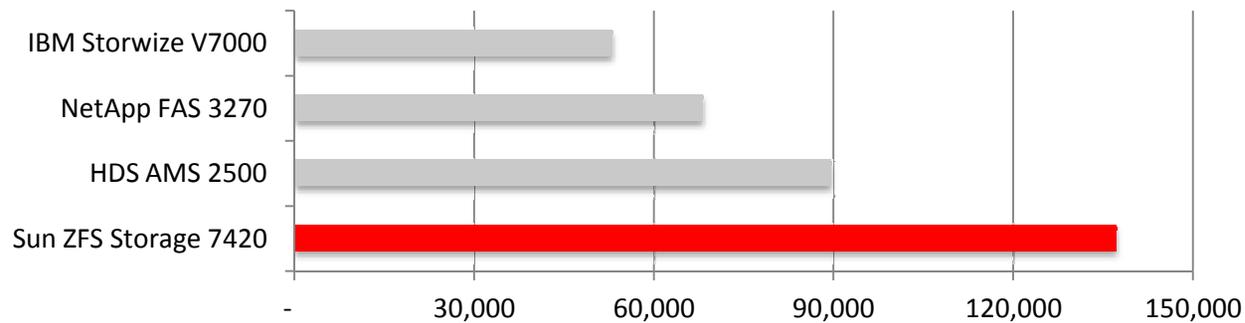


Sources:  
SAN: [storageperformance.org](http://storageperformance.org)  
NAS: [spec.org/sfs2008/](http://spec.org/sfs2008/)

**ORACLE**

## SPC 1 Benchmark: *Storage benchmark that represents a typical database block protocol workload*

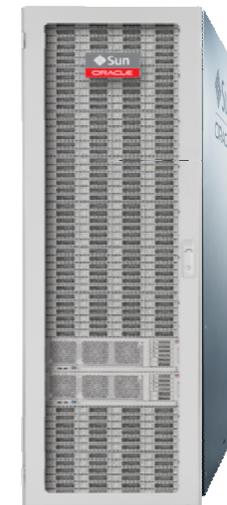
### IOPS



### \$/IOPS



**Sun ZFS Storage 7420:**  
**2X Price / Performance**

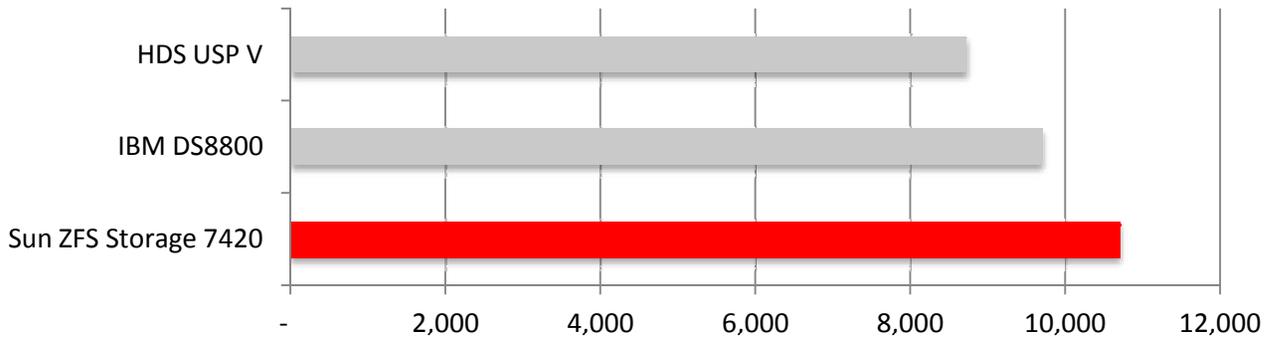


source: storage performance council. [www.storageperformance.org](http://www.storageperformance.org).

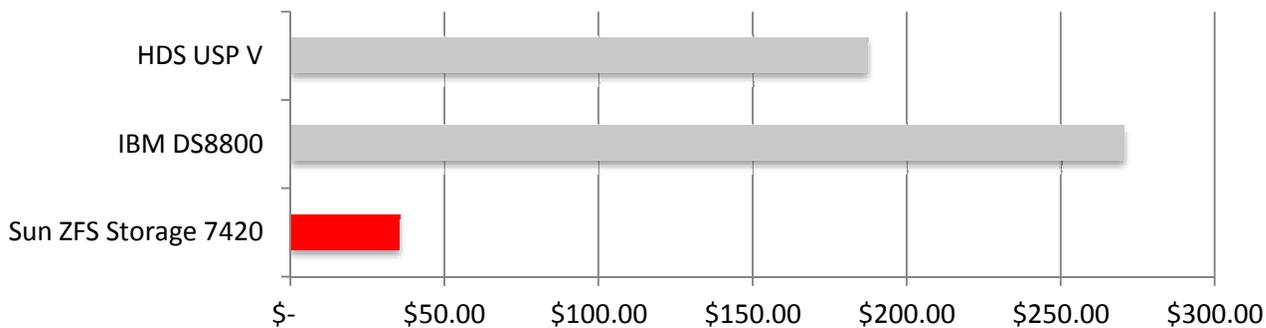
ORACLE

## SPC 2 Benchmark: *Storage benchmark that represents performance of block protocol throughput workloads*

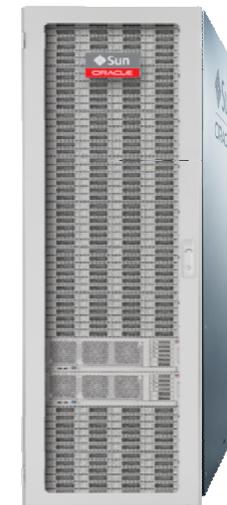
### MBPS



### \$/MBPS



**Sun ZFS Storage 7420:**  
**7X Price / Performance**

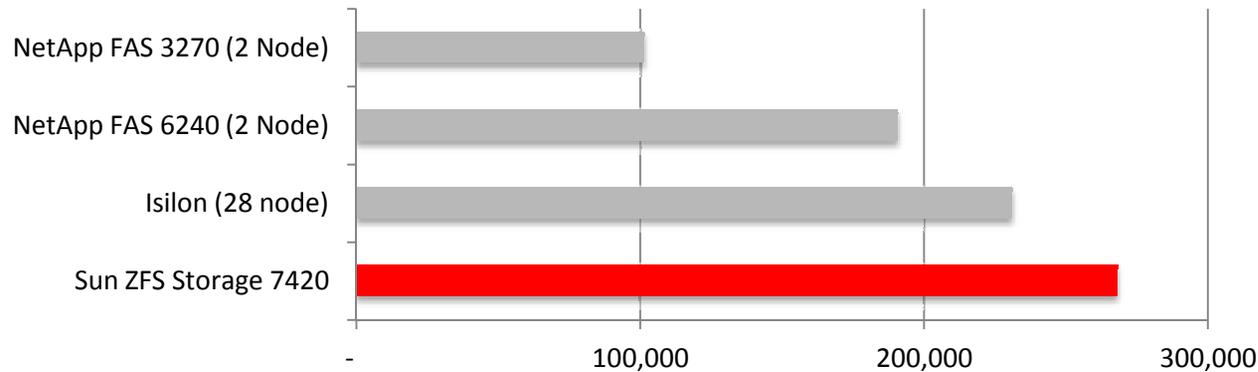


source: storage performance council. [www.storageperformance.org](http://www.storageperformance.org).

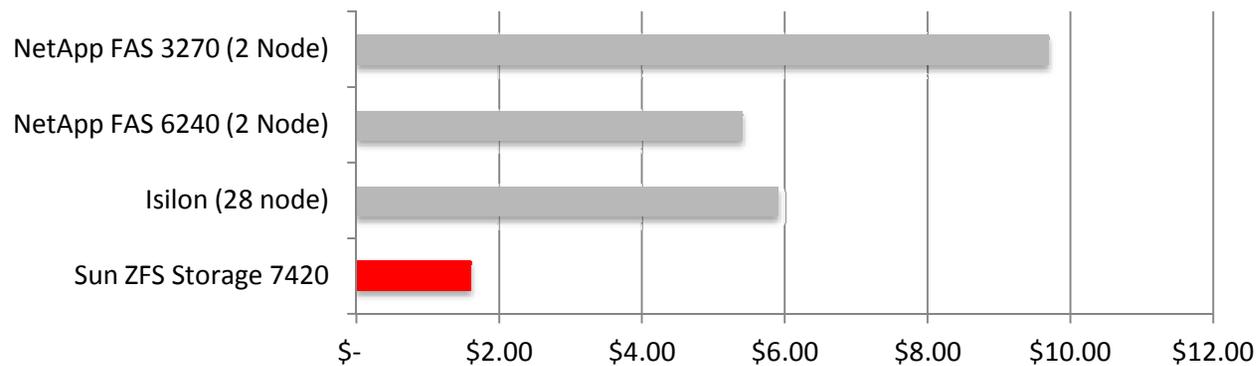
ORACLE

**SPECsfs Benchmark:** Storage benchmark that measures file server throughput and response time, providing a standard method for comparing performance across different vendor platforms

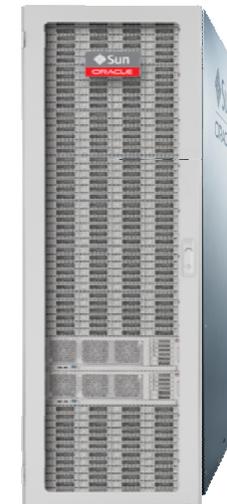
## IOPS



## \$/IOPS

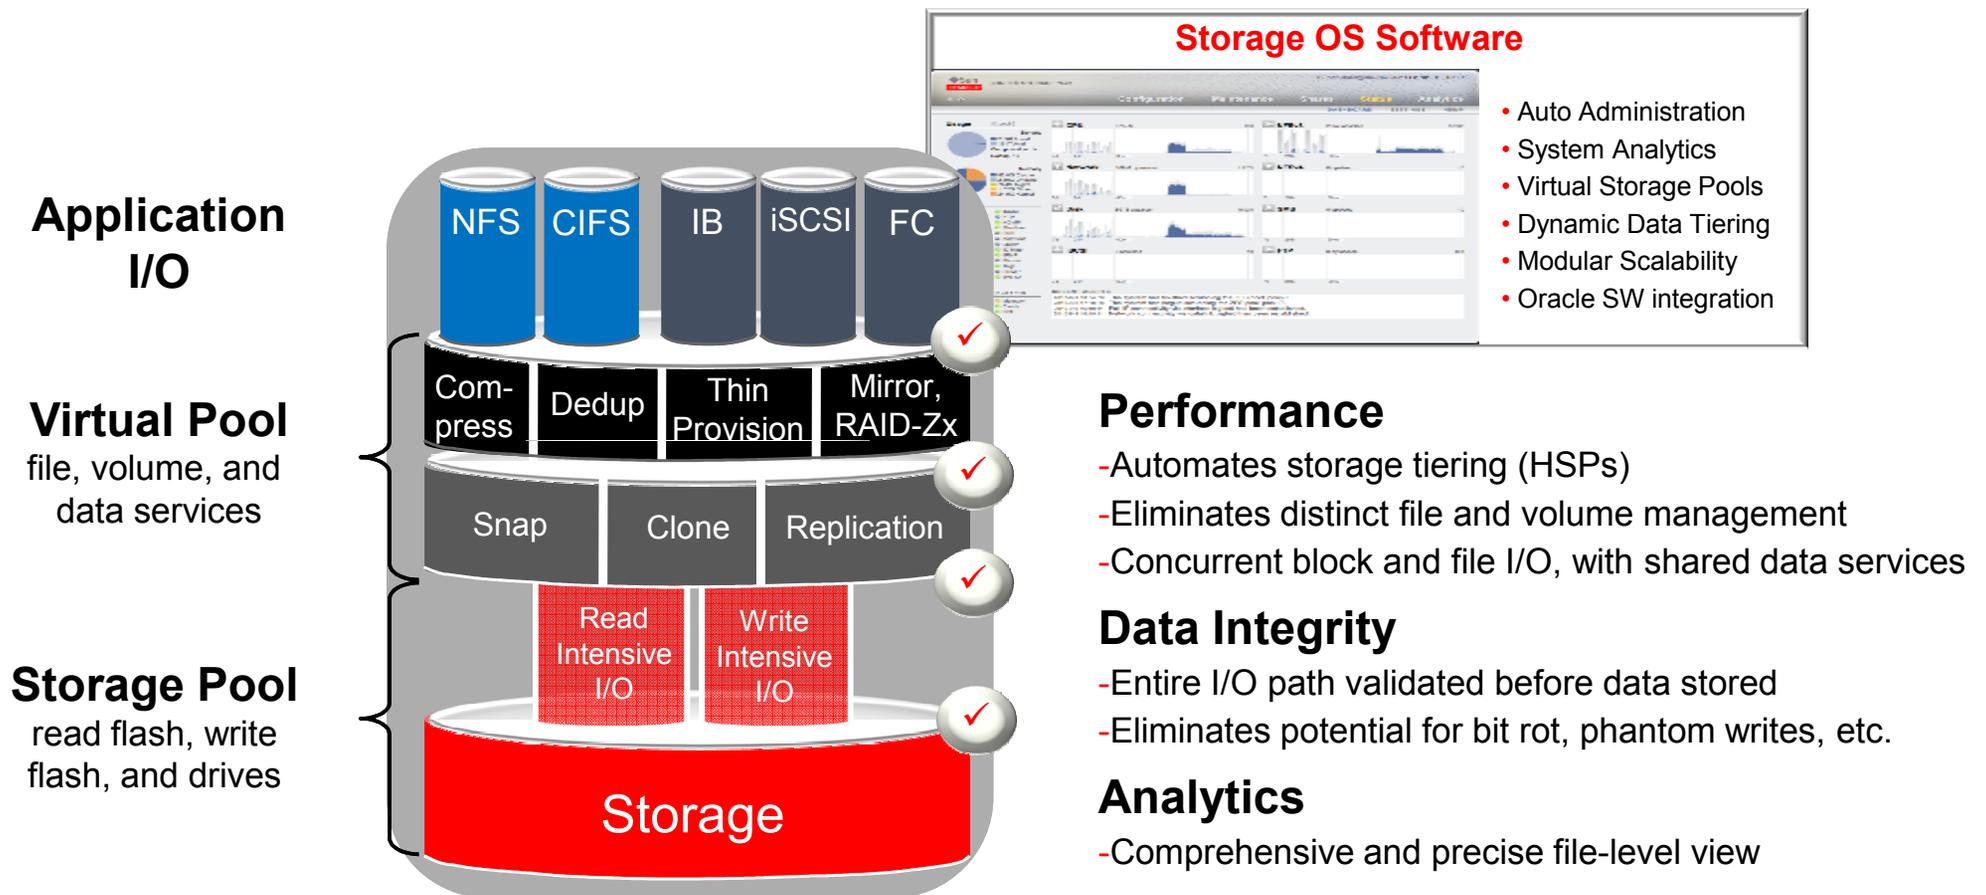


**Sun ZFS Storage 7420:**  
**3X Price / Performance**



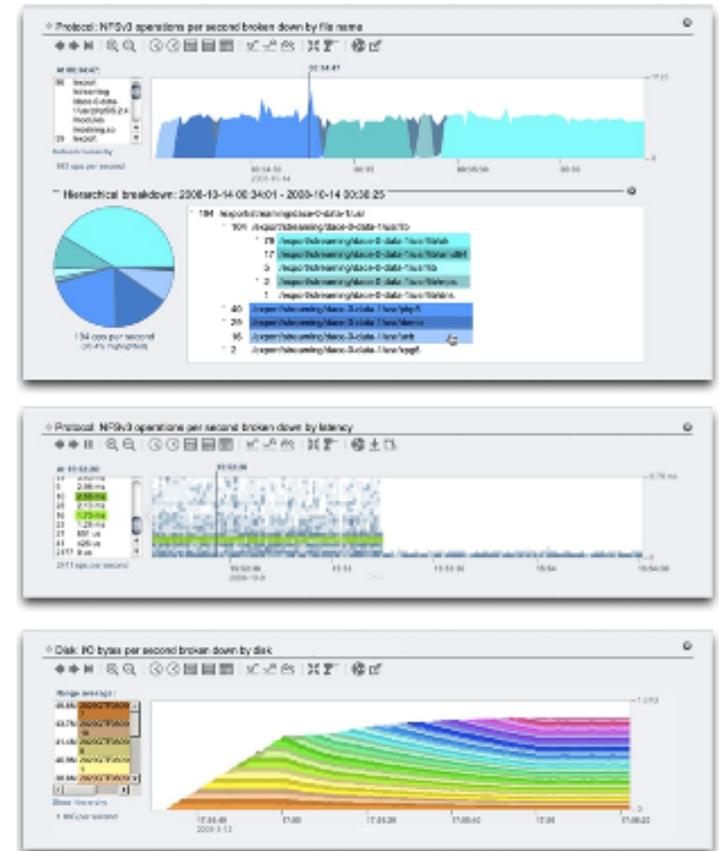
ORACLE

# Industry-leading Storage Efficiency



# Best in Management Simplicity

- **Only storage appliance with DTrace Analytics**
  - Automatic real-time visualization of application and storage related workloads (on right picture of database and table index analytics)
  - Simple yet sophisticated instrumentation provides real-time comprehensive analysis
  - Supports multiple simultaneous application and workload analysis in real-time
  - Analysis can be saved, exported and replayed for further analysis



# Sun ZFS Storage Appliances

## Engineered with Oracle software.

Fastest and most reliable storage for SPARC Supercluster, Exadata, and Exalogic.

Eliminates 3 data management layers between Oracle Solaris and Storage.

Increases 11g Database performance and efficiency by (at least) 3-5x.

Increases RMAN backup performance by 2x.

Over 70 million hours in production environments with zero data loss.

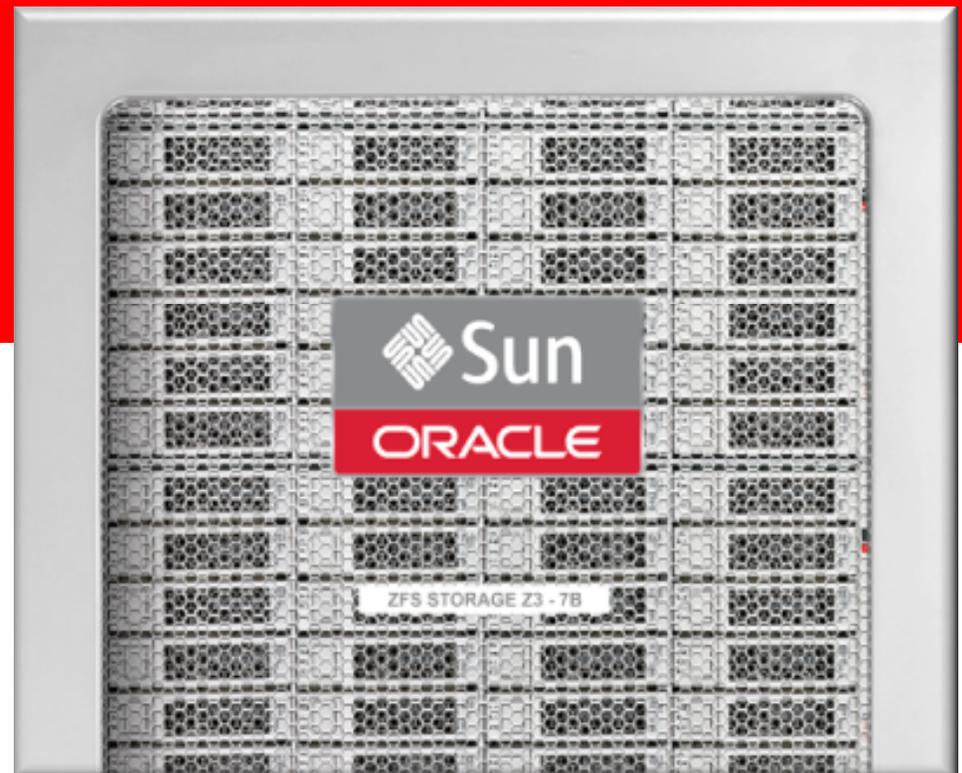
ORACLE

# ZFS Backup Appliance

*powered by ZFS*

NEW

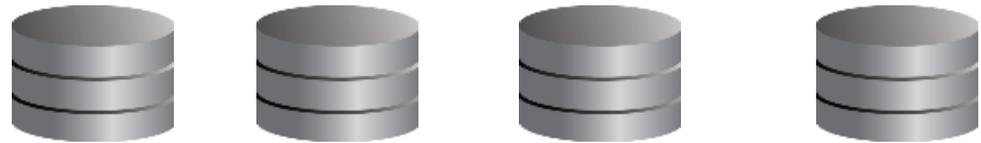
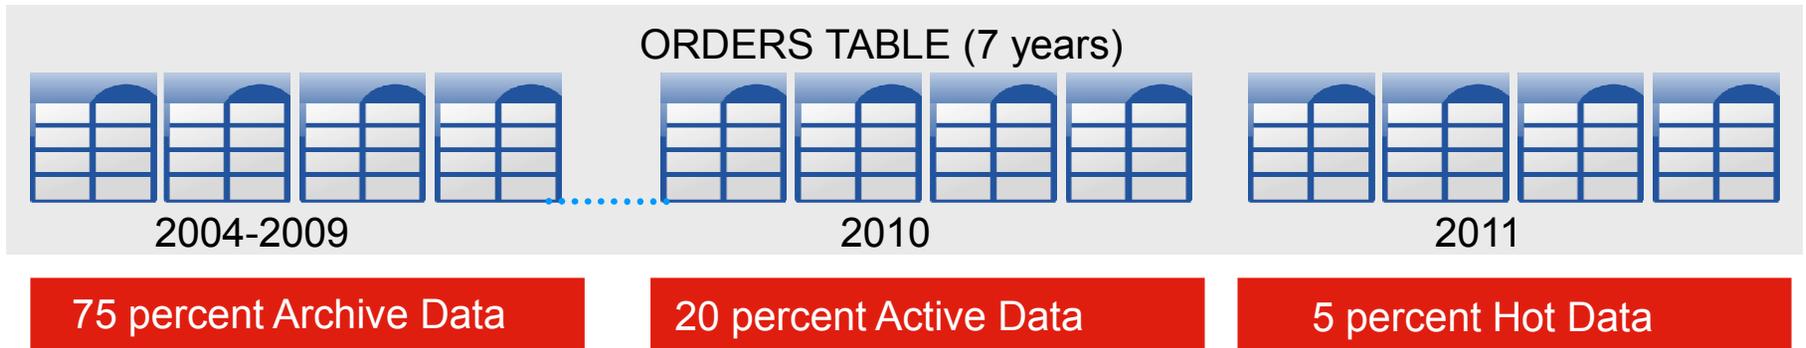
**Up to 20-30TB/h** backup, **~10TB/h** restore  
**2x** faster restores than Data Domain  
**4x** faster NDMP backups than NetApp  
**Scales** to 2.4PBs and 7TB Flash



ORACLE

# HCC: Reduce Storage Costs

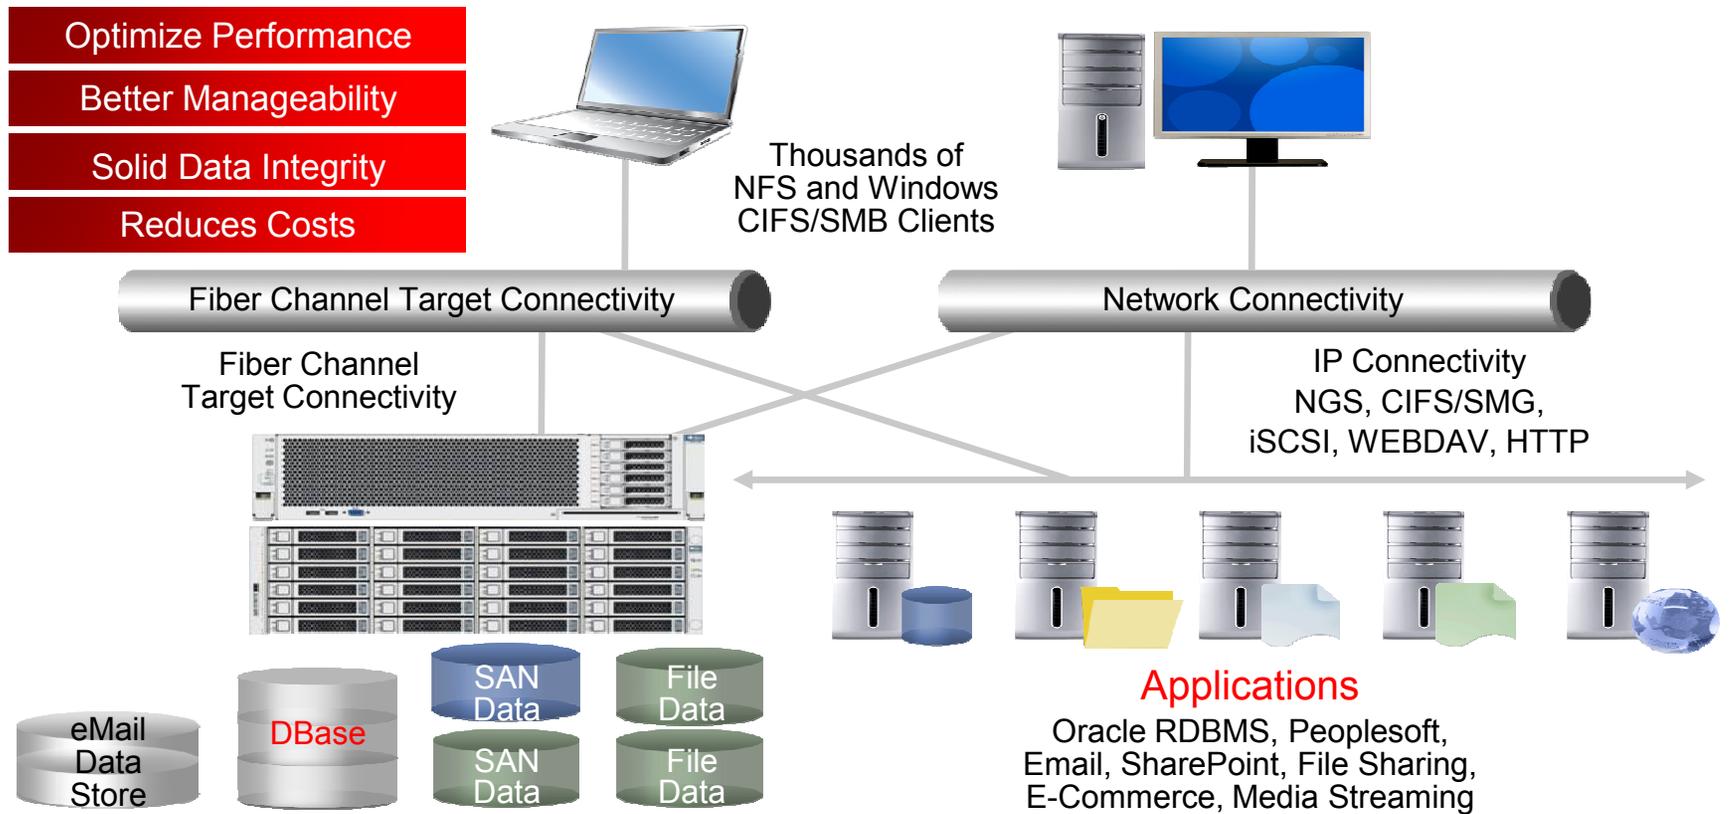
## Database Partitioning and HCC



Existing Storage

Sun ZFS Storage Appliance or Axiom Storage

# Storage Consolidation



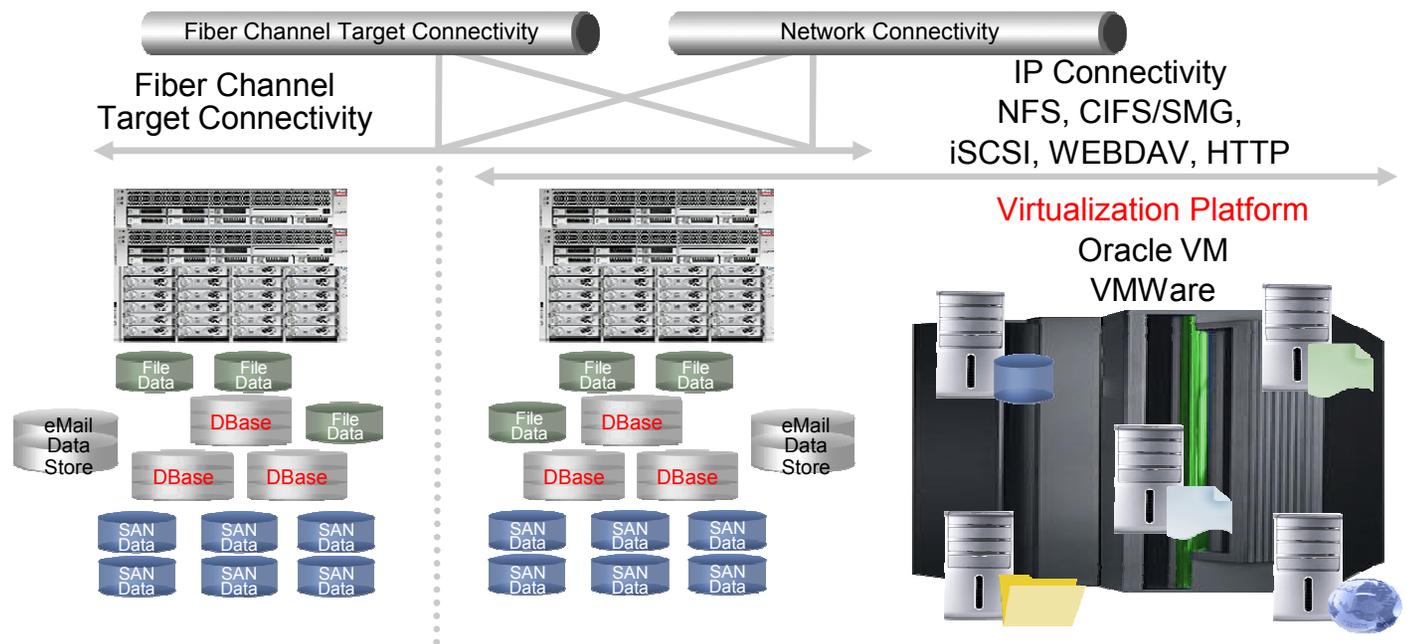
# Virtualization Environments

Optimize Performance

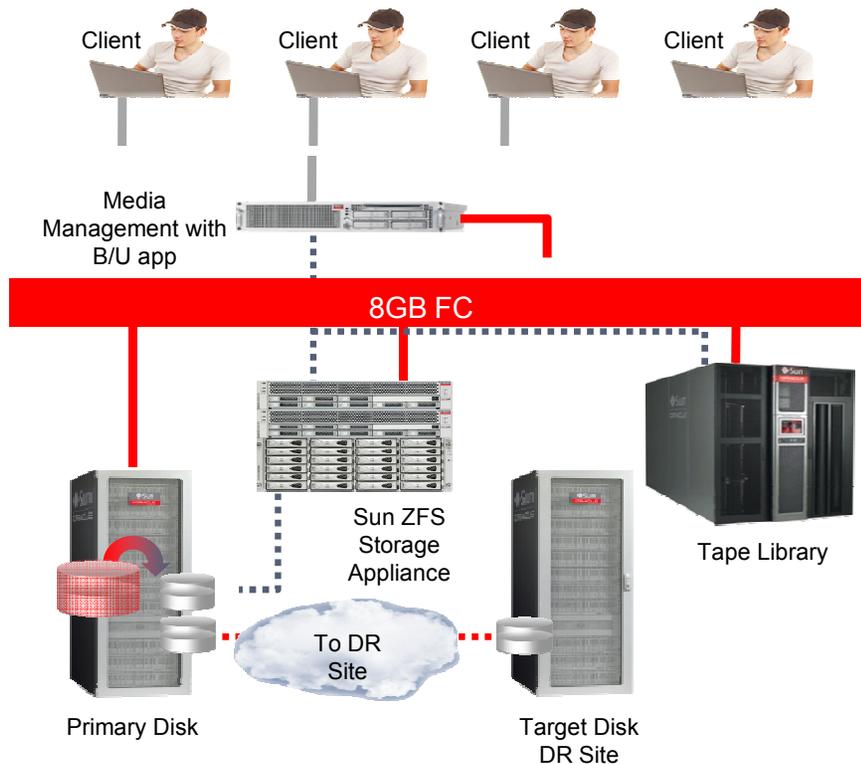
Better Manageability

Solid Data Integrity

Reduces Costs



# Backup Storage and Business Continuity



## Business Drivers

- Reduce risk and increase productivity
- Long-term retention (data mining)
- Shorten back-up windows

## Core Product Requirements

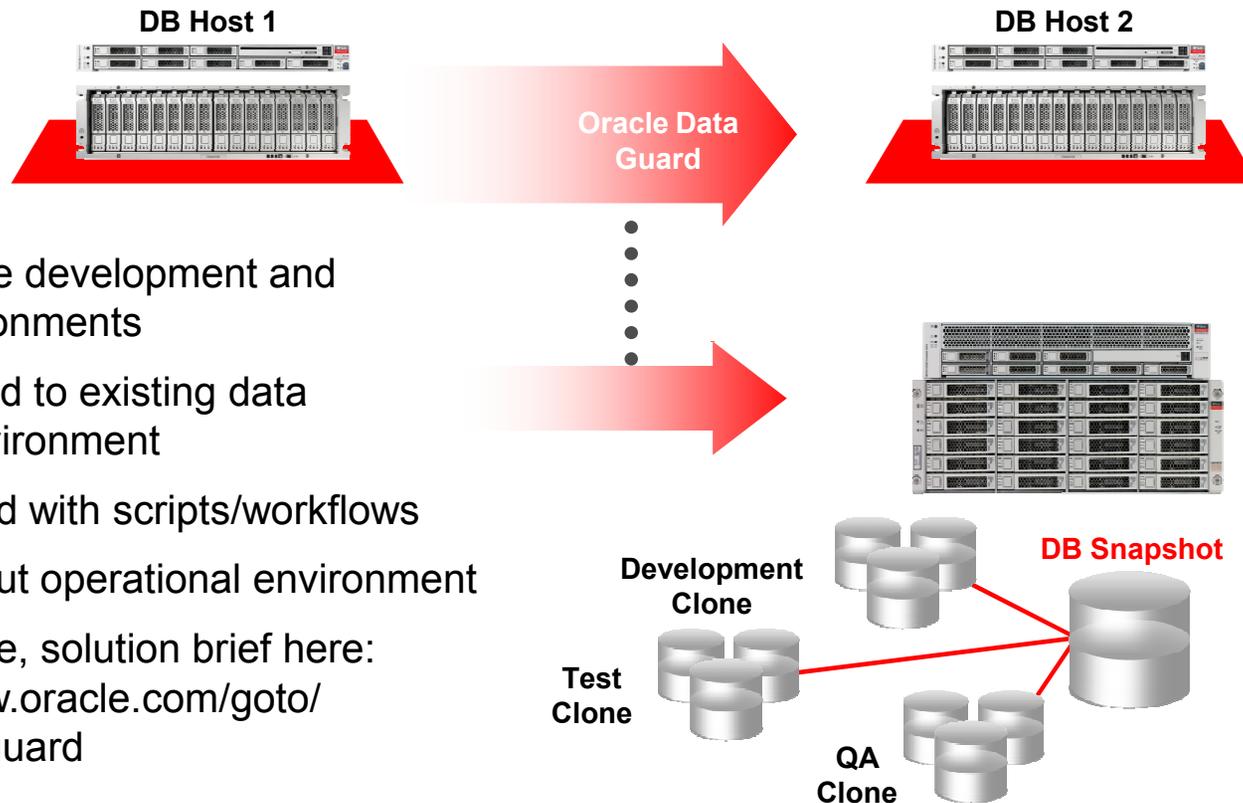
- Highly available architecture
- Scalable, multi-tiered solution
- Data movement software (app. Int.)
- Services (migration, scripting, etc.)

## Sun ZFS Storage Appliances

- Shorten back-up windows with integration compression
- Scripts and workflows provided for targeted solutions



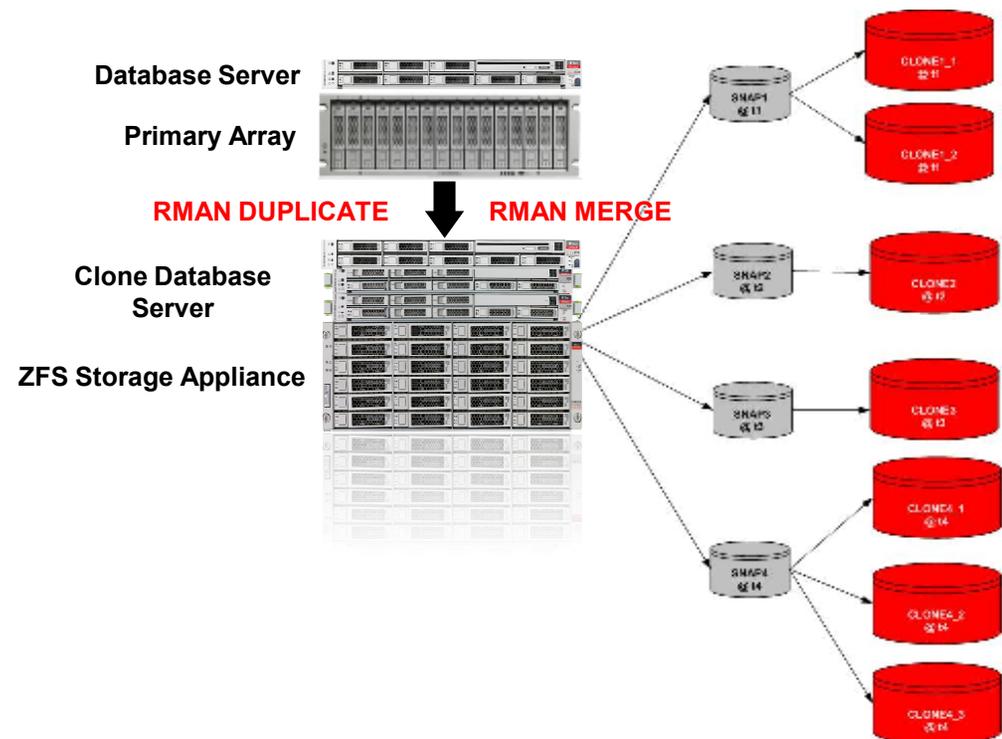
# Adding Value to Data Guard DR Copies with Snaps/Clones



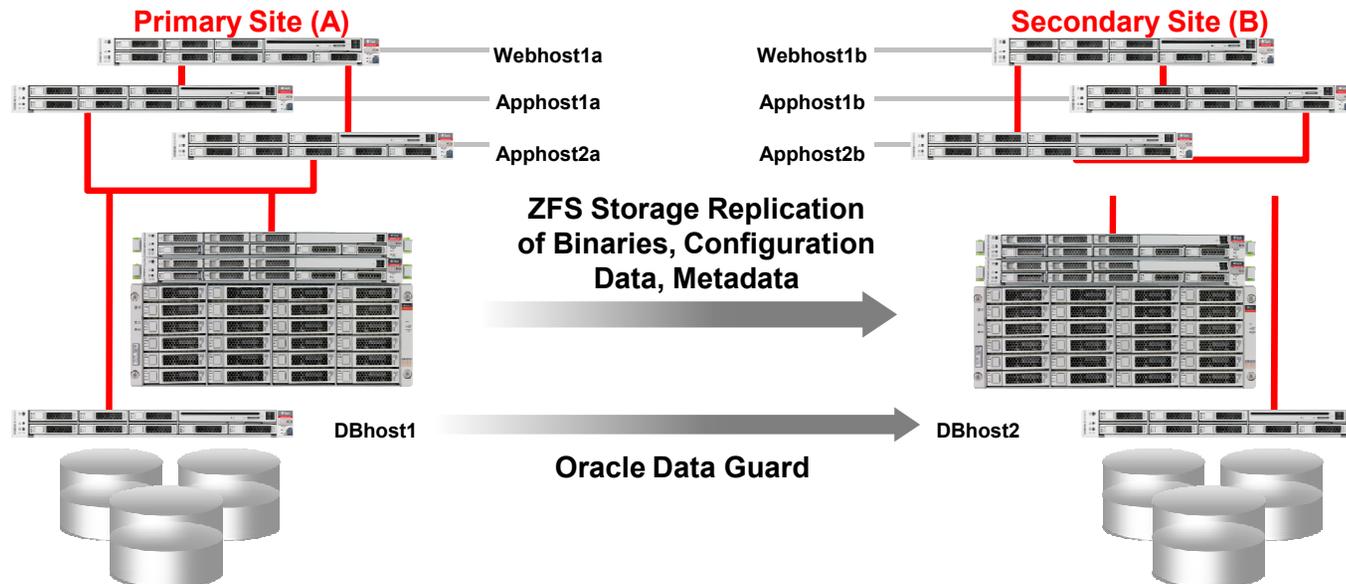
- Accelerate development and test environments
- Simply add to existing data guard environment
- Automated with scripts/workflows
- Rounds out operational environment
- MAA guide, solution brief here: <http://www.oracle.com/goto/zfs-data-guard>

# Accelerate Development and Test environments

- Accelerate deployment of patches, upgrades, and app' s
  - Effectively and efficiently give each Developer their own database
  - ID changes in code that affect storage system performance
- Nearly no procedural changes
- Add value to Data Guard DR copies of database with included snapshot and clone features
- Tested, validated, and supported by Oracle



# Add Value to DR for Fusion Middleware

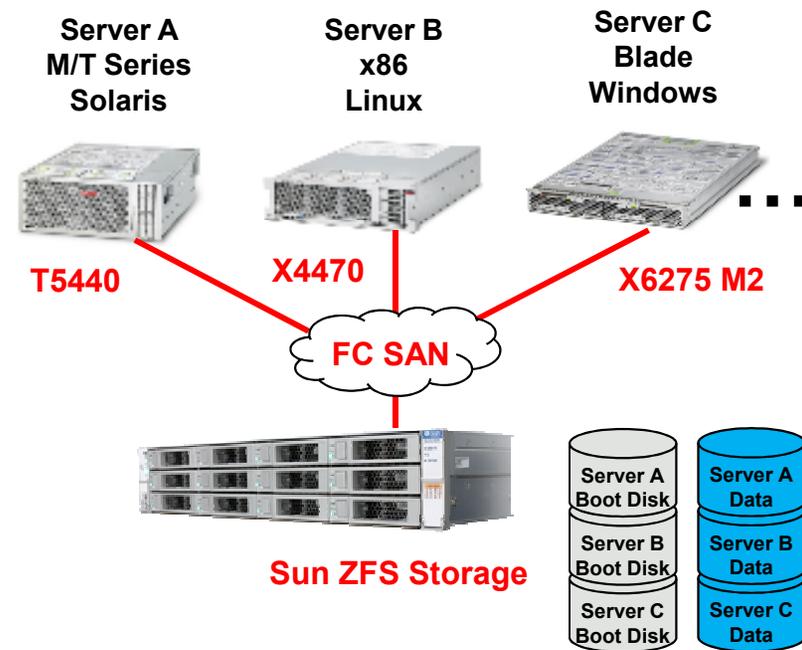


- Protect the whole business process
- Completes data guard environment by replicating data stored outside of database
- Failover/failback of both ZFS storage and data guard coordinated and automated
- MAA Guide, solution brief here: <http://www.oracle.com/goto/zfs-middleware>

# Fiber Channel SAN Boot

SAN Boot support with Oracle servers in virtualization and consolidation environments

- Fully supported boot solution with Oracle servers
- Detailed system analysis with DTrace Analytics and management dashboard
- 10X easier to manage with unlimited snapshots and clones, centralized boot image upgrades
- 25X greater efficiency with thin provisioning for max capacity
- More reliability and lower operating cost with less disk and moving parts, less power and cooling

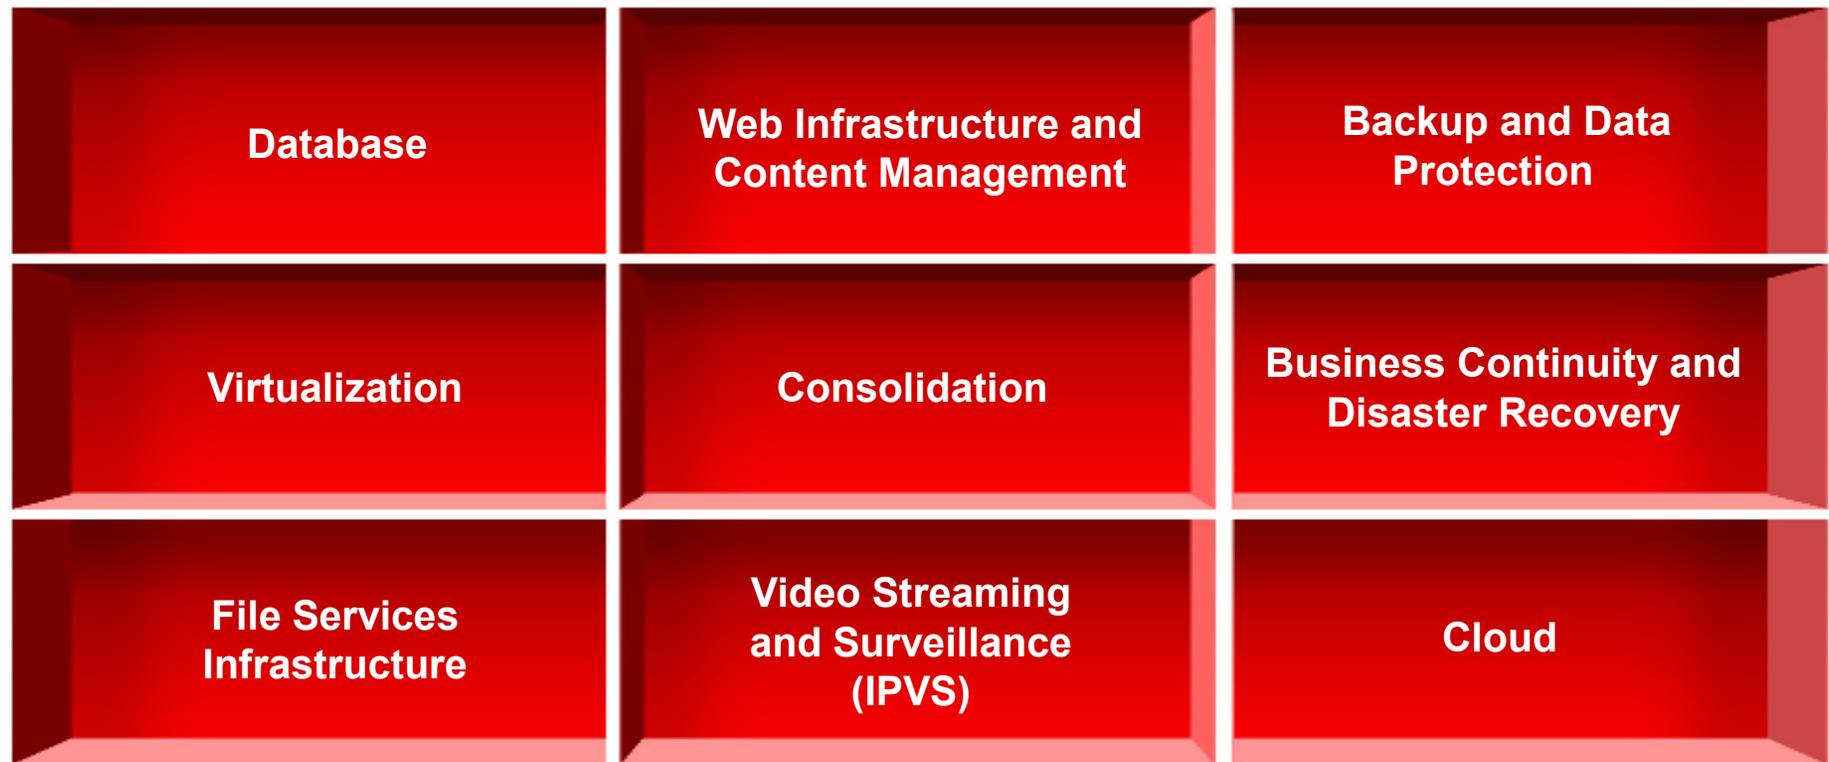


# Oracle GIT Migration from Existing NAS to 7000

## Advantage with Sun ZFS Storage 7000

- 2 to 1 ratio in capacity and performance
- Nearly 50% reduction in power and cooling costs
- Immediate operational efficiency with provisioning, snapshots and clones
- Response to customer issues improved 75-91% when using DTrace Analytics
- On-Demand customer migrations to 7000 were immediate with Shadow Migration feature with nearly zero downtime.

# Supporting a Wide Range of Applications



ORACLE

# ZFS Appliance Monitor

Available  
now on  
App Store

- iPhone application that monitors multiple ZFS Storage Appliances
- Brings the power of DTrace Analytics at your fingertips
- Easily view the health/status of any ZFS Appliance, anywhere
- Ideal for Execs/Storage Admins with need for instant access to information
- Graphical and intuitive display of:
  - System and component status (disks, CPU, RAM, fans, etc)
  - Service status (NFS, iSCSI, SMB, etc)
  - Resource & I/O status
  - DTrace Analytics (real time statistics via saved worksheets)
  - Logs for alert, system, fault, audit, phone home
  - Active systems problems (can be marked as repaired via app)
  - High level group status including storage use, problems, hardware status
  - Storage utilization
  - Remote activation of system and component locator LED



ORACLE

# Sun ZFS Demonstration

- Basic administration and login GUI/CLI
- Project/share/cifs protocol mount
- Data copy with analytics

# Sun ZFS for Exadata Backup

- Architecture
- Tunings
- Performance



# Backup Architecture – Measure twice cut once

- Determine business and technical requirements
  - RPO/RTO
  - Backup windows
  - Tape integration (offsite requirement?)
  - Capacity analysis
    - Retention, Schedule, Full vs incr vs merg
    - Image vs backupset
  - Current infrastructure
    - Physical
      - networking
    - Staffing



## ZFS Configuration - Hardware

- Single head or cluster
- Appropriate CPU/RAM/networking (IB or 10GbE?)
- PCIe SAS cards
- Required disks rotational speed/capacity
- Appropriate qty of disk shelves
- Log and cache SSD devices



## ZFS Configuration - Software

- Create pools and data protection level
  - Active/active vs active/passive
- Projects and share organization
- Share tuning – e.g. logbias, compression, record size etc
- IB network configuration
  - IPMP, active/active, active/passive, multiple IP's

# Exadata/Host Readiness

- dcli is your friend
- fstab mount options
- Kernel/OS tunings
- dNFS `make -f $ORACLE_HOME/rdbms/lib/ins_rdbms.mk dnfs_on`
  - *Oracle instance running with ODM: Oracle Direct NFS ODM Library Version 3.0*
  - `oranfstab` – share path load balancing
    - `gv$dnfs_servers`: Shows a table of servers accessed using Direct NFS.
    - `gv$dnfs_files`: Shows a table of files currently open using Direct NFS.
    - `gv$dnfs_channels`: Shows a table of open network paths (or channels) to servers for which Direct NFS is providing files.
    - `gv$dnfs_stats`: Shows a table of performance statistics for Direct NFS.

# RMAN Approach

- Increase:
  - backup\_disk\_bufcnt"=64 &
  - backup\_disk\_bufsz"=1048576
- Organize channels to write to both zfs cluster heads
- Use 2-4 channels per tray
- Experiment with increasing the channel count
  - Find the knee in the curve
- Use as many db nodes as possible
- Compression? Basic, ACO: low/med/high; zfs comp?

# References

- MOS note 1354980.1
- MOS note 1495104.1
- cnt1727488
- maa-wp-dbm-zfs-backup-1593252
- exadata-7000-367640
- <http://www.oracle.com/technetwork/articles/servers-storage-admin/perf-hybrid-columnar-compression-1689701.html>
- <https://www.virtualbox.org/wiki/Downloads>
- <https://www.virtualbox.org/wiki/Downloads>
- [http://www.oracle.com/webapps/dialogue/ns/dlgwelcome.jsp?p\\_ext=Y&p\\_dlg\\_id=10521841&src=7299332&Act=45](http://www.oracle.com/webapps/dialogue/ns/dlgwelcome.jsp?p_ext=Y&p_dlg_id=10521841&src=7299332&Act=45)

# Hardware and Software

ORACLE®

# Engineered to Work Together

ORACLE®