

NAS and Data Protection


ZFS Storage Technology

Jason Schaffer, Vice President, Product Engineering

May 17, 2017

Oracle's ZFS Storage Technology

Engineered Systems




Exadata Exalogic SuperCluster Big Data Appliance

NAS and Data Protection



Oracle ZS5 series

SAN



Oracle FS1

Tape and Virtual Tape



Storage Cloud

Deployment Options: Private, Public, Hybrid

Services: IaaS, PaaS, SaaS

Consumption Options: Build, Manage, Subscribe

Storage Software

Storage Management: OISP, OEM, ASM, Storage Analytics, System Management, ACSLS, ELS

Automated Tiering: HSM, Partitions, HSP, HCP, VSM

Data Reduction: ACO, ADO, HCC, RMAN, ZFS Storage Appliance Dedup/Comp

Data Protection: Data Guard, RMAN, OSB, ZS Snap/Rep, MaxRep

Security/Encryption: ASO, Oracle Key Manager, Disk/Tape Encryption

System Software

Data protocols

- OISP
- Fibre channel
- iSCSI
- Infiniband
 - NFS/RDMA
 - IPoIB
 - iSER
 - SRP
- Object API
- NFS V3 and V4
- SMB1, 2 and 3
- HTTP
- WebDAV
- FTP/SFTP/FTPS
- ZFS NDMP V4

Data services


- Hybrid Cloud Pools (ZFS Cloud)
- Hybrid storage pools (HSPs)
- Single, double and triple-parity RAID
- Mirroring and triple mirroring
- End-to-end data integrity
- Local and Remote replication
- Snapshots and clones
- Quota(s)
- Deduplication
- Compression
- Encryption
- Thin provisioning
- Antivirus via ICAP protocol
- Online data migration
- Clustering

Management

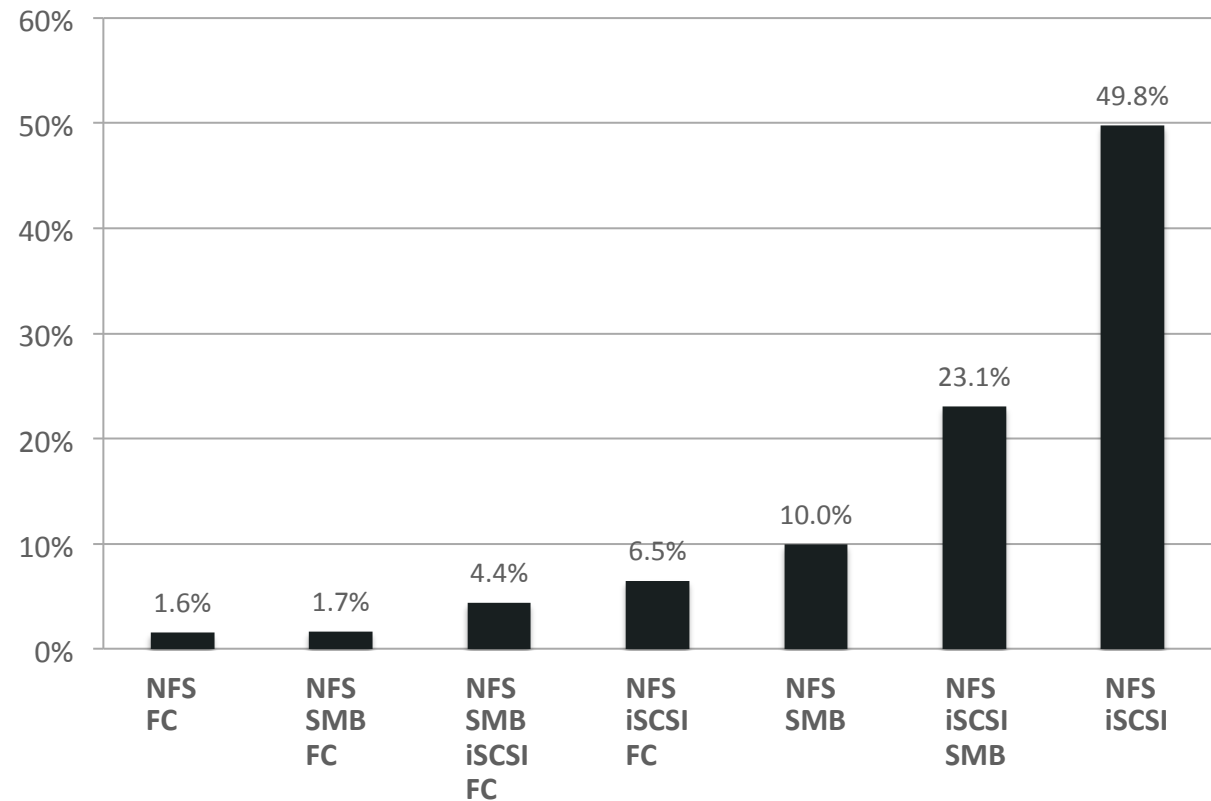
- RESTful API, OpenStack/Cinder
- Browser and CLI interface
- Management dashboard
- Hardware/component view
- Role-based access control
- Phone home
- Event and threshold based alerting
- Dtrace analytics
- Scripting
- Workflow automation
- Advanced networking
- Snap Management Utility
- Source aware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

Software Protocol Use Demographics

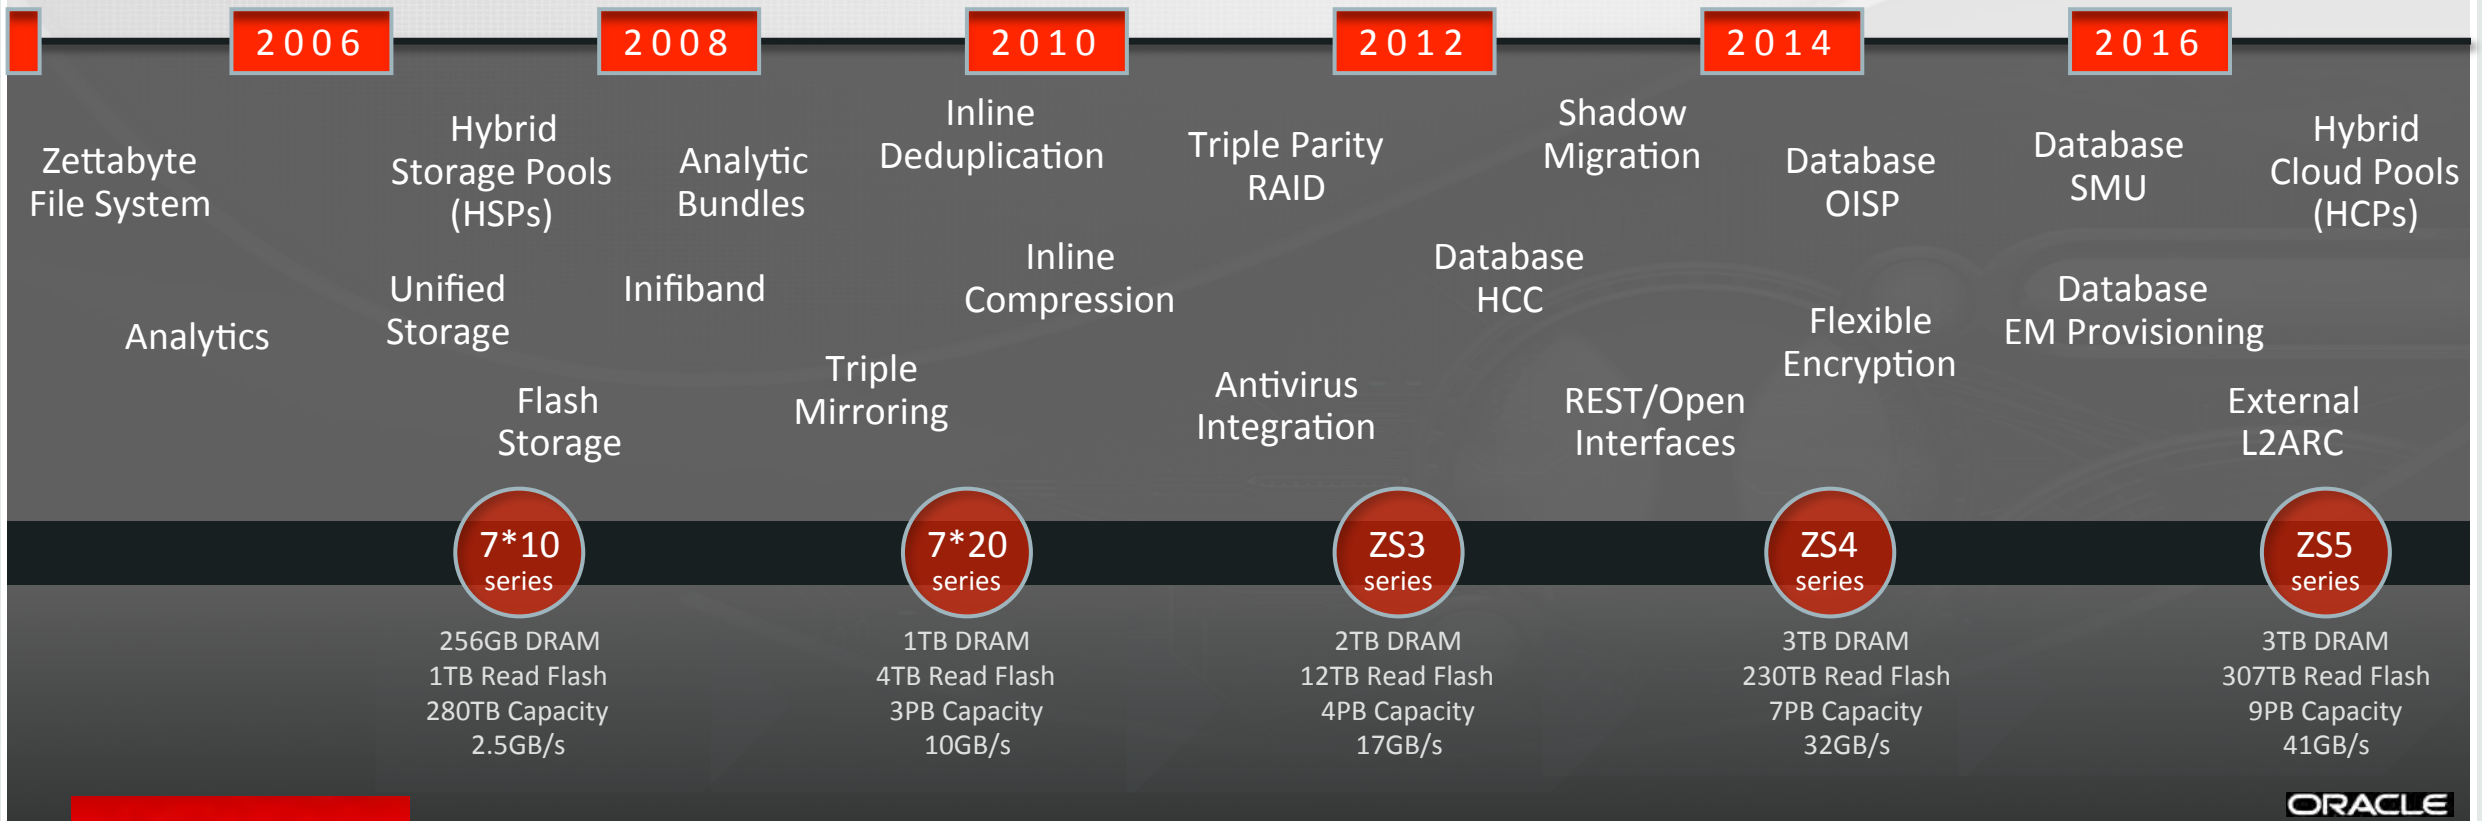
12.3% of systems have only one protocol in use



87.7% of systems have multiple protocols in use simultaneously



Built on a Decade of Storage Innovation



DESIGNED FOR SCALE . . .



Proven Cloud Scale and Efficiency

Oracle's Public Cloud relies on ZFS Storage to deliver database performance with cloud economics.

35,720,000
Snapshots and Clones

28,520,000
Concurrent Users

18,031,000
Hours in Production

1,090,134,000
Transactions per Day

1.01 Exabytes
ZFS Storage

Core cloud-scale design elements

- POD-level performance with a balanced scale and risk profile
- Zero-downtime architecture via system redundancies and sufficient overhead
- Administrative operations per minute (e.g., LUN deletes >800 Ops/min.)
- Rapid (e.g., <35 seconds for 100 clones) snap/clone capability
- Automated I/O prioritization and performance (e.g., db log bias) tuning
- Dynamic and embedded extrinsic storage resource utilization

Engineering ZFS Storage for Cloud Scale

SCALABILITY DEMANDS MULTIPLE ARCHITECTURAL DIMENSIONS

Logical “front-end” Scale

- Multi-protocol
- Hybrid Storage Pools (HSP)
- Dynamic Data Reduction
- Virtualized ZFS machines

Logical “back-end” Scale

- Zettabyte Filesystem
- Decorrelated RAID
- Hybrid Cloud Pool (HCP)
- Efficient Data Mobility

Management Scale

- Analytics (trouble-shooting)
- OISP 2.0 (integration)
- POD/C@C Scale via QoS
- Bring-up/down operations


Scalable System Performance / Hybrid Cloud Pools (HCPs)

Automated use public clouds for data protection

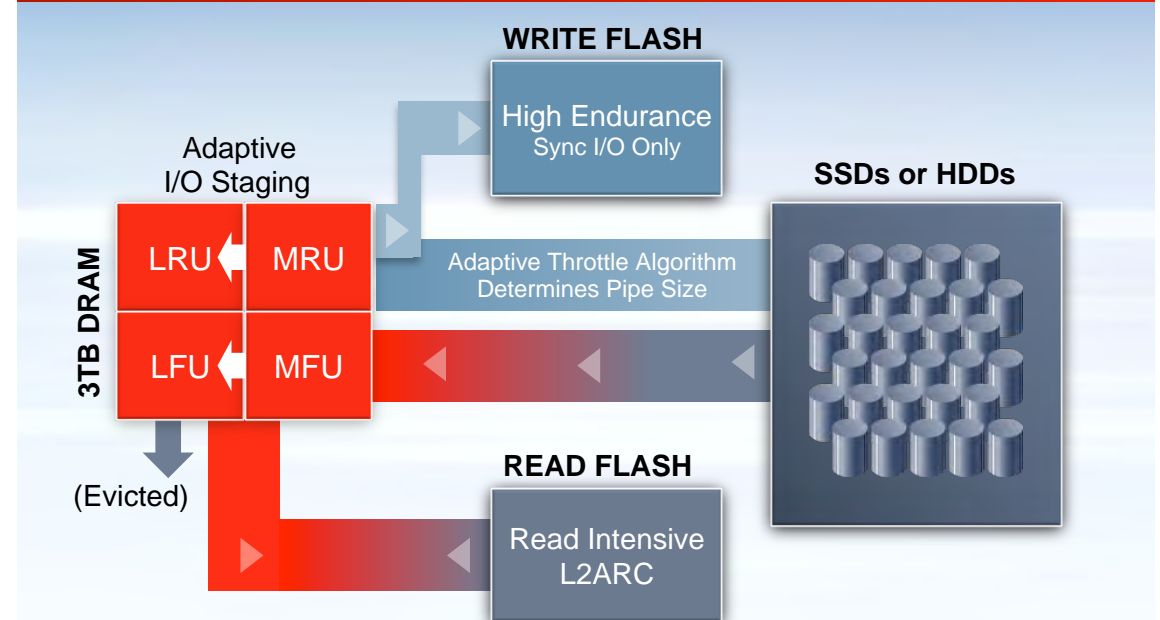
Design Challenge: Enable virtually unlimited storage pool scalability via external cloud integration without impeding primary storage performance

ZFS | Hybrid Storage Pool (HSP) Technology

Most Horsepower Possible

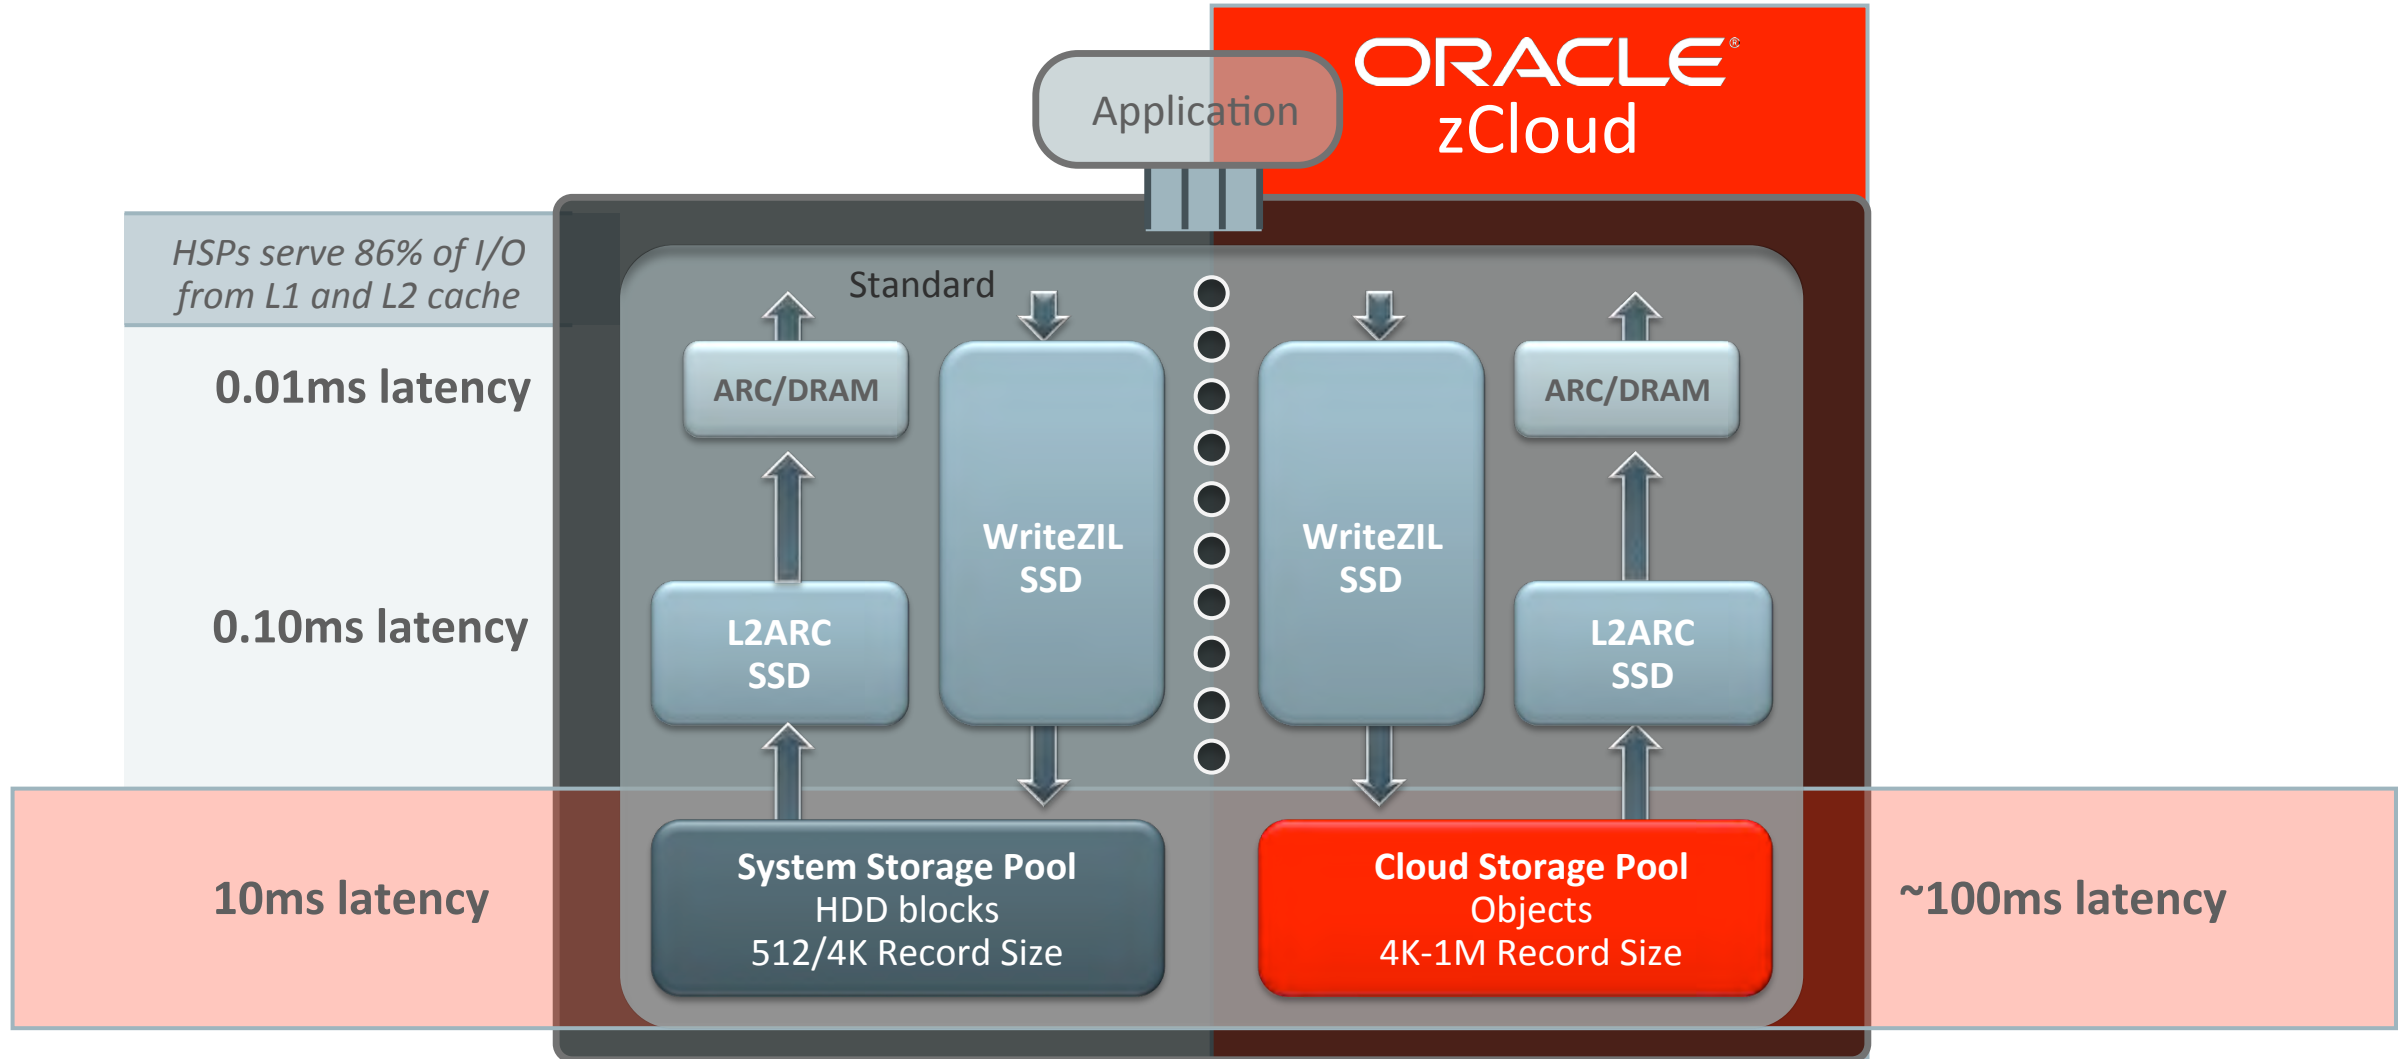


Hybrid Storage Pool

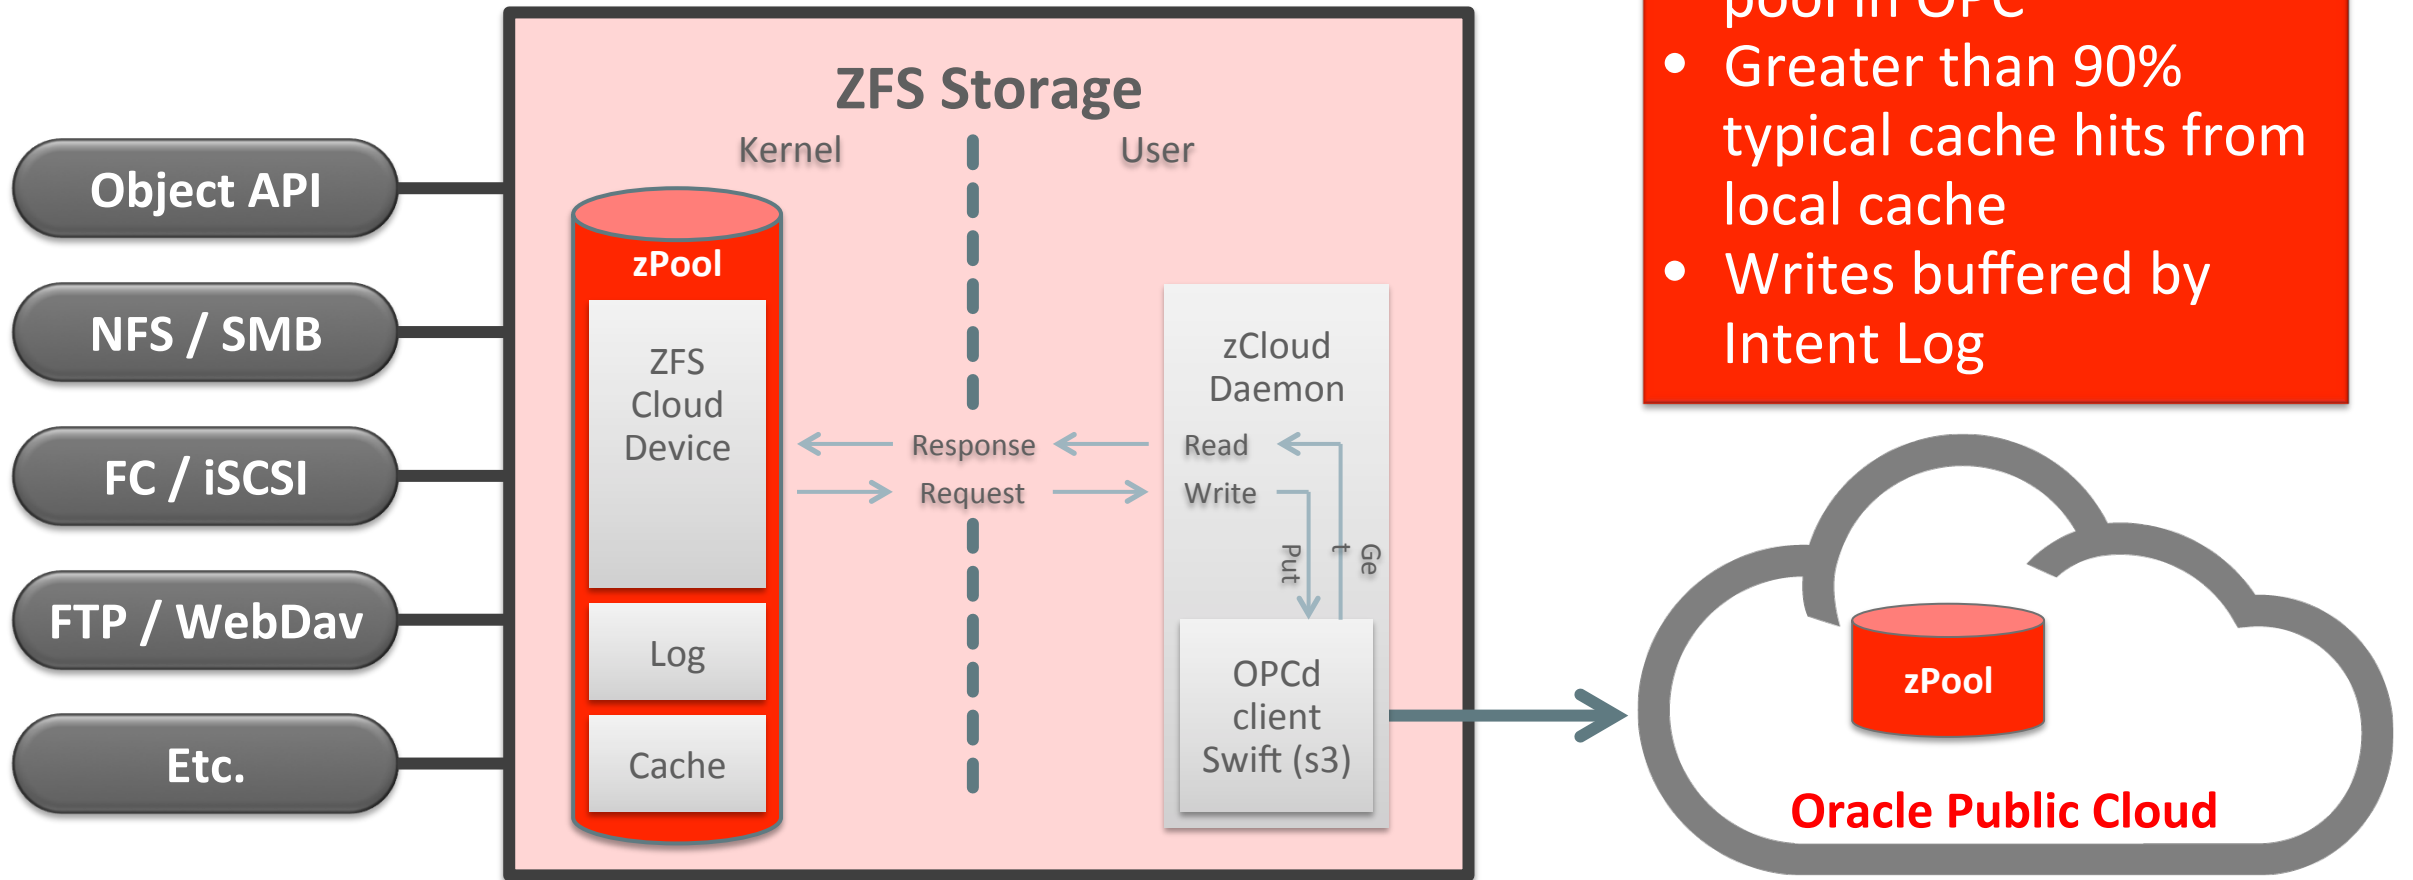


- Automated, real-time data migration from DRAM to multi-class flash, to multi-class disk storage
- Software specifically engineered for multi-level flash and disk storage

ZFS Cloud (Hybrid Cloud Pool)



ZFS Cloud (Hybrid Cloud Pool)



- Incremental forever pool in OPC
- Greater than 90% typical cache hits from local cache
- Writes buffered by Intent Log

ZFS Cloud (Hybrid Cloud Pool)

Data Services

NFS	Online	2017-2-21 05:44:05	↺ ⏻
iSCSI	Online	2017-2-21 05:43:58	↺ ⏻
SMB	Disabled	2017-2-21 05:42:45	↺ ⏻
FTP	Disabled	2017-2-21 05:42:46	↺ ⏻
HTTP	Online	2017-2-21 05:44:13	↺ ⏻
NDMP	Online	2017-2-21 05:44:07	↺ ⏻
Remote Replication	Online	2017-2-21 05:43:52	↺ ⏻
Shadow Migration	Online	2017-2-21 05:43:57	↺ ⏻
SFTP	Disabled	2017-2-21 05:42:46	↺ ⏻
SRP	Disabled	2017-2-21 05:42:48	↺ ⏻
TFTP	Disabled	2017-2-21 05:43:59	↺ ⏻
Virus Scan	Disabled	2017-2-21 05:42:46	↺ ⏻
Cloud	Online	2017-2-21 05:43:35	↺ ⏻

Please wait...

Directory Services

SERVICES STORAGE NETWORK SAN CLUSTER USERS PREFERENCES SETTINGS ALERTS

Services Cloud - Accounts Properties Logs

Back to Services 2017-2-21 05:43:35 Online REVERT APPLY

Cloud Service
Provides the ability to manage Cloud services. Users can register Cloud accounts and create virtual data disks on the Cloud.

Cloud Accounts Total: 1 1-1

ID	STATE	LOCATION	USER	DISKS
demo	Online	http://192.168.56.202/object/auth/v1.0/export/zcloud/	zcloud	4

See Also
Help: Cloud
Wikipedia: Cloud Computing

Confirm that all devices are present and minimally functional, and add them to an existing storage pool. ABORT COMMIT

Verify and allocate devices Step 1 of 2

Verify that storage is correctly attached and functioning. If devices are missing or malfunctioning, they will not be available for use and cannot be added without reconfiguring the pool. It is recommended that you fix any problems before configuring storage on the appliance. Mixing device types and speeds is strongly discouraged.

ID	MODEL	TYPE	RPM	DATA	LOG	CACHE	META
demo	swift	Cloud	-	1 (64M)	-	-	-

Add Cloud Account CANCEL ADD

ID

Location

Cloud Administrator User Name

Password

Use web proxy If your system communicates to the web through a proxy, check this box and enter the configuration information below.


Host : port :

Username


Password

Uses: Hybrid Data Management, Migration, and DR

Basic Cloud Gateway: Moves data from on-prem to OPC/object storage. Can be used for data migration as well.



Data Protection using Replication: Create a DR copy in the cloud.



Data Protection/Backup: Local/Remote Targets for NBU, Commvault, and Veeam.



Innovation | Most Secure and Flexible Technology

SECURE



Only Storage System That Delivers . . .

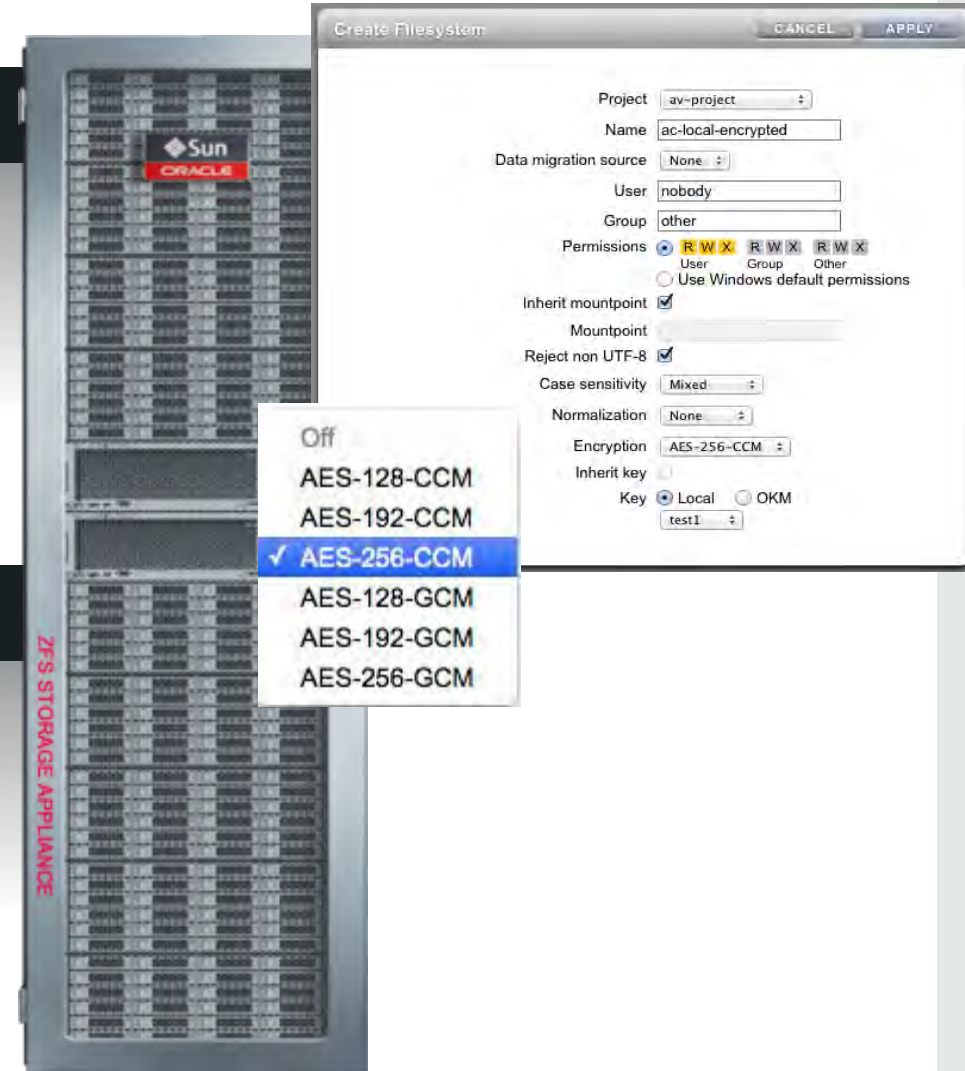
- Strong AES 256-bit encryption keys
- Two tier encryption key architecture
- Integrated local and central key management

FLEXIBLE



Only Storage System That Delivers . . .

- Encryption at project, share and LUN level
- Complete access control options
- Clear text and encrypted data on same system



Configuration Options

Hybrid Cloud Pools (HCP)


- Comprised of Oracle Cloud Capacity
- Functions as a private system resource
- Compatible with all system data services

Deduplication (Meta Data) Flash

- Comprised of Oracle Cloud Capacity
- Functions as a private system resource
- Compatible with all system data services



Security level and control established for shares, LUNs, or pools or profiles



Flash Storage Pools

- 8TB to 3PBs of Flash Storage
- Consistent low latency
- Database indexes, redo logs, active partitions

External Read Flash

- 2TB to 300TBs of Read Flash Storage
- Apply Read flash to any pool
- Ideal for Database warehousing and mining

High-Performance Controllers

- 3TBs of DRAM, 144 processor cores
- Up to 24x10GbE, 16x40GbIB, 16x16Gb FC
- Full suite of enterprise software

Hybrid Storage Pools (HSP)

- Up to 9PBs of Storage Capacity
- Up to 450,000 IOPS at <1ms - 4ms response
- Database data; primary, test/dev, backup

Scalable Storage Integration/O.I.S.P.

Dynamic and automatic database-storage tuning

Design Challenge: Enable dynamic system (i.e., database-storage) I/O sizing and memory use for optimum performance with backup (RMAN) and production (dNFS) database use concurrently

OISP Technology: Database Storage Protocol




Oracle Database OISP

- Extensions built on top of NFSv4 protocol
- Each I/O is tagged with data context
 - I/O reason
 - I/O priority
 - File type (redo log, data file, control file, etc.)
 - Database block size for file
 - Database and/or pluggable database id
 - Cache hints
 - Prefetch hints

Oracle ZFS Dynamic Actions


- Dynamically set Record Size, log bias
- Pluggable Database Analytics
- Database OEM provisioning
- Database Snap Management
- Analytics AWR Feed
- I/O FSS & DB Priority
- I/O Caching/Prefetch



OISP 1.0 | Dynamic Database-Storage Tuning

Eliminates up to 67% of time spent tuning Database and ZFS Storage for optimal performance

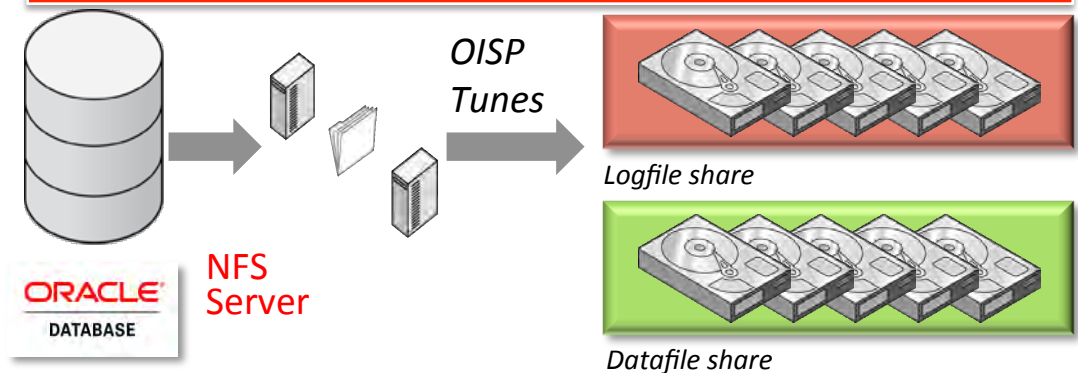
Database Storage Without OISP



10+ shares, each with manually tuned Record Size and LogBias settings

mnt/dbname/redo (Admin sets RS, LB)
mnt/dbname/control (Admin sets RS, LB)
mnt/dbname/archivelog (Admin sets RS, LB)
mnt/dbname/datafile (Admin sets RS, LB)
mnt/dbname/tempfile (Admin sets RS, LB)
mnt/dbname/chgtrack (Admin sets RS, LB)
mnt/dbname/undo (Admin sets RS, LB)
...

Database Storage with OISP Enabled

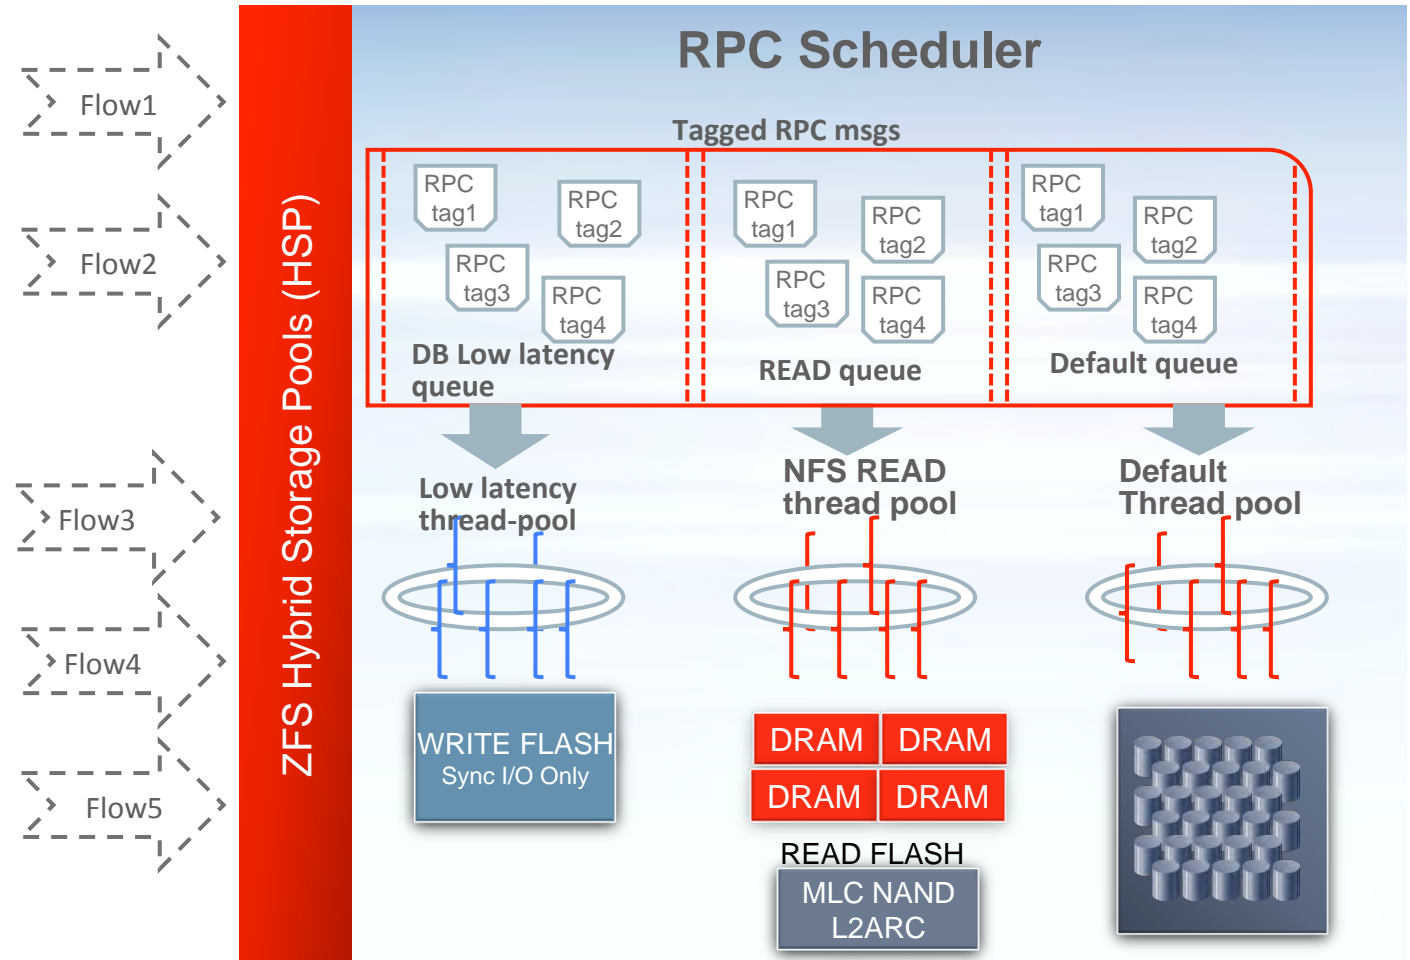


2 shares, dynamically and automatically tuned for optimal performance

mnt/dbname/logfile (RS, LB automatic)
 (redo
 (control
 (chgtrack
mnt/dbname/datafile (RS, LB automatic)
 (datafile
 (tempfile
 (undo
 (archivelog


OISP2.0 | Database I/O Prioritization

- Prioritize high priority I/O handling over medium and low priority I/O
 - Weighted prioritization of storage system I/O managed by the database
 - High priority examples: RAC pings, LGWR writes, CKPT writes
 - Low priority example: RMAN backups, background DBWR writes



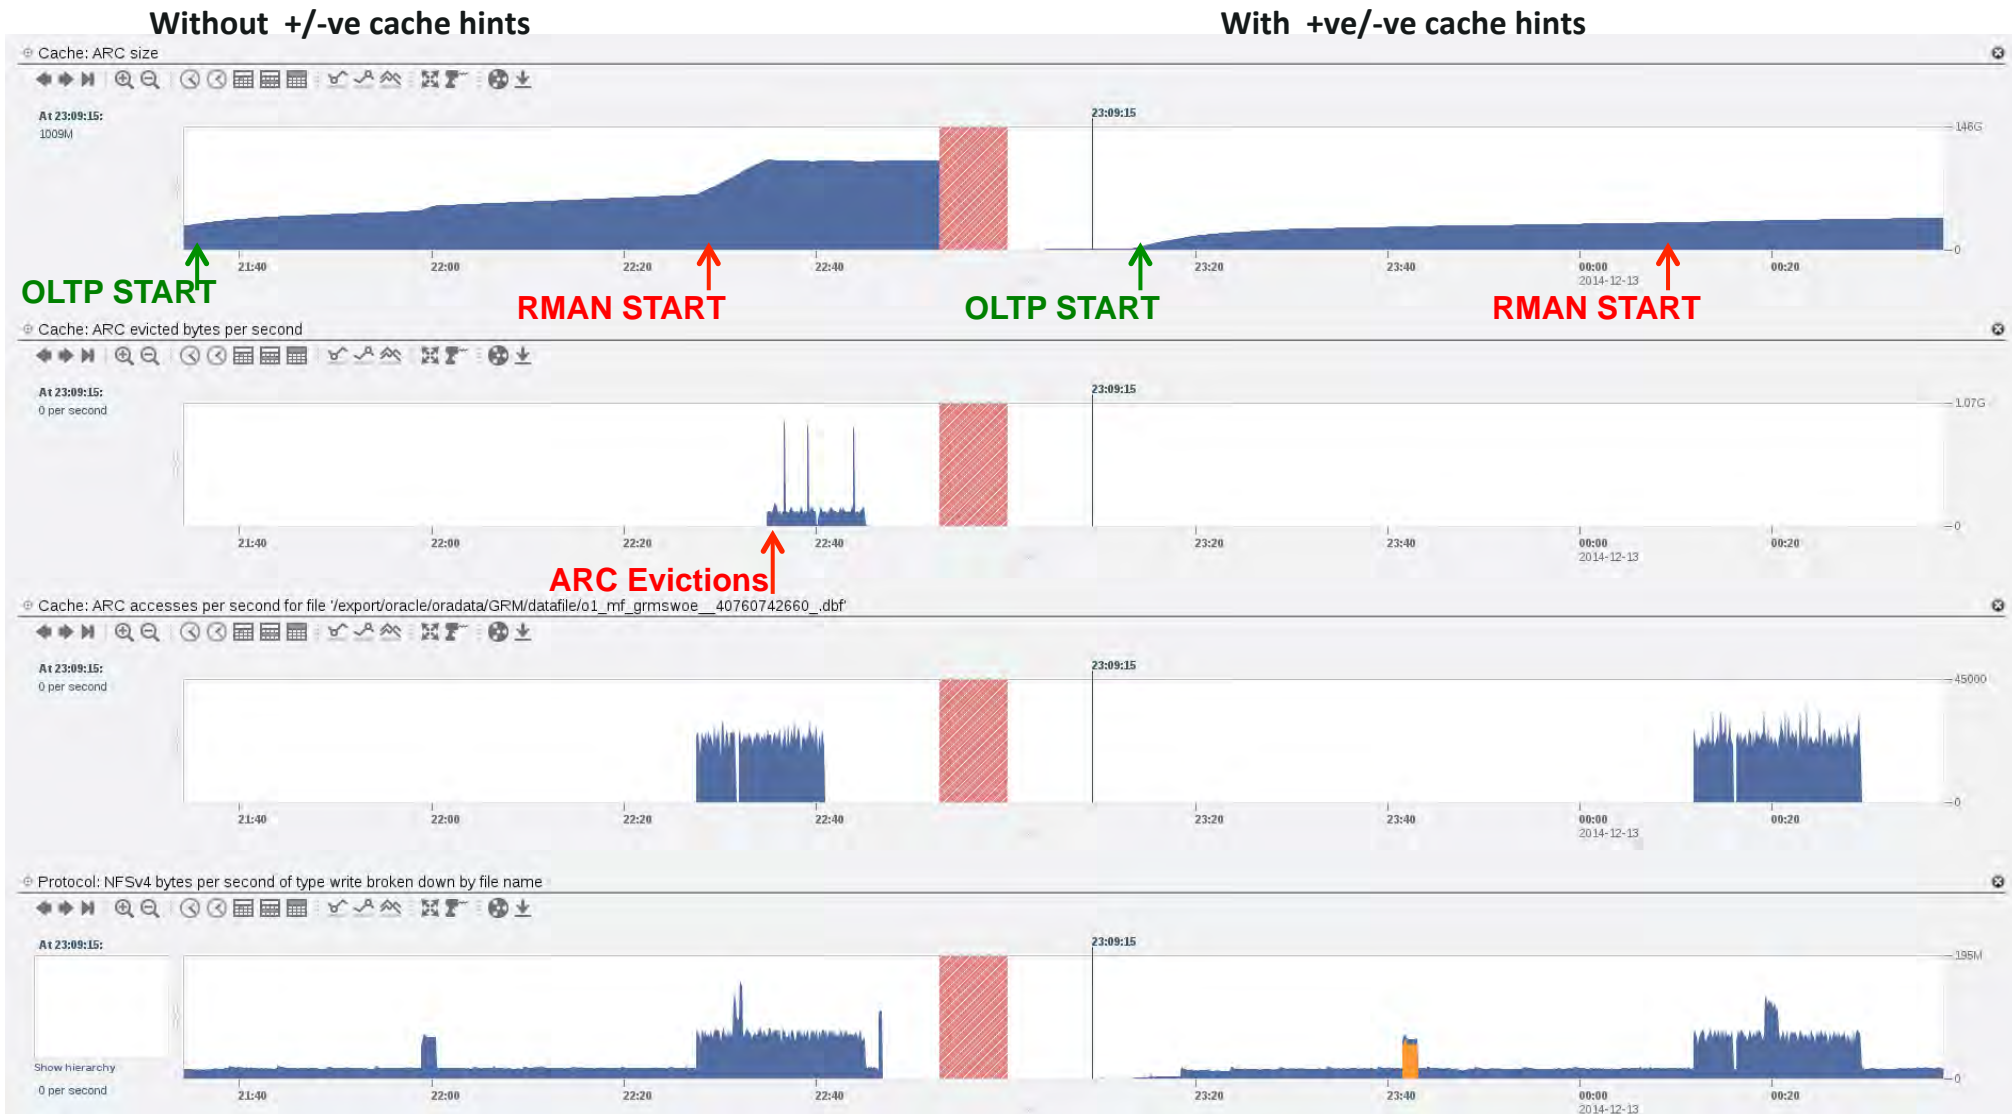
OISP 2.0 | Database Performance Results

Multiple (OLAP/OLTP) Database OLTP Workloads Running Concurrently



OISP2.0 Dynamic Tuning Extensions

- Simultaneous RMAN + OLTP workloads
- Negative cache hints for RMAN and Archivelog I/O
- ARC growth rate steady for OLTP,
- No ARC growth for RMAN, Archivelog I/O
- No ARC evictions!



OISP 2.0 Database (AWR) Analytics Extensions

New Functionality: Storage analytics that incorporates Oracle Database AWR data.

Database I/O Statistics by File Type

- Data File
- Temp File
- Archive Log
- Data Pump/Dump
- Log Files
- Backups
- Control File
- Flashback Log

Database I/O Statistics by Function

- DBWR
- Direct Writes
- LGWR
- Archive
- RMAN
- Recovery
- Streams/AQ
- XDB

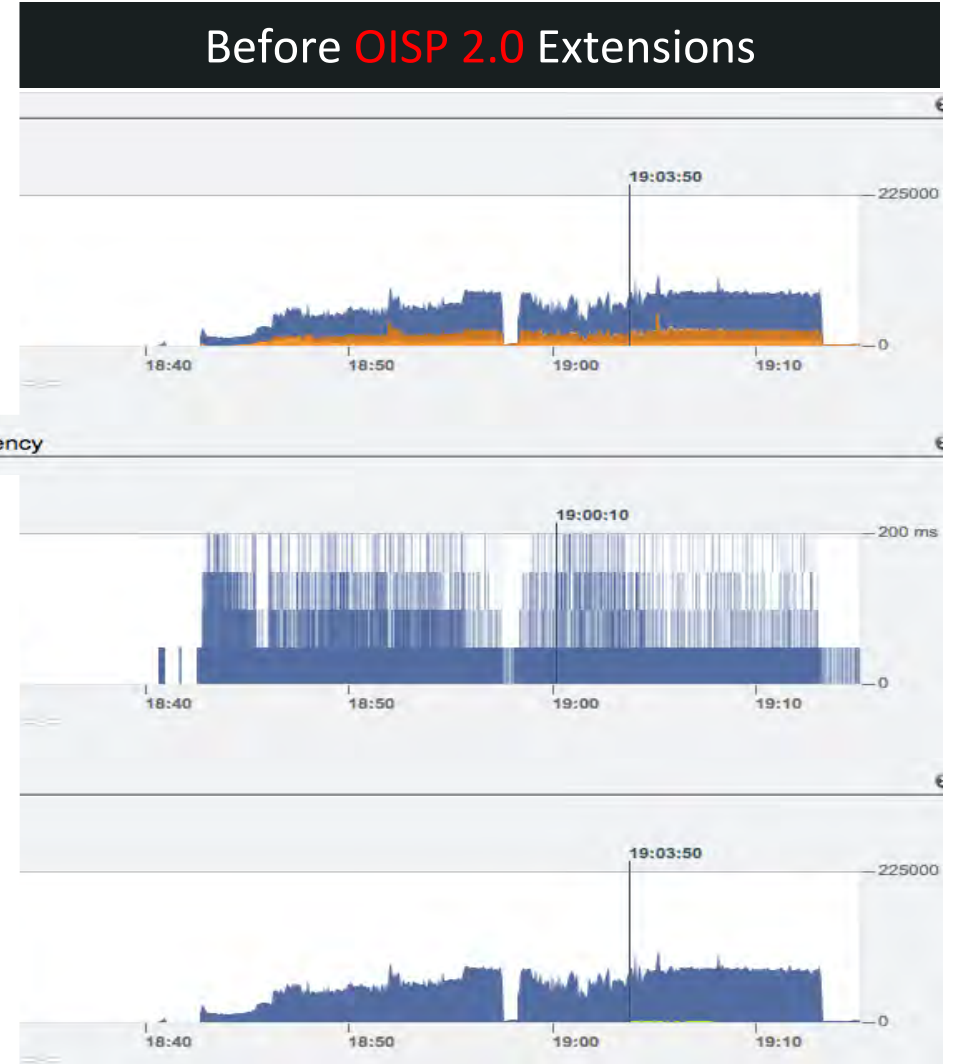
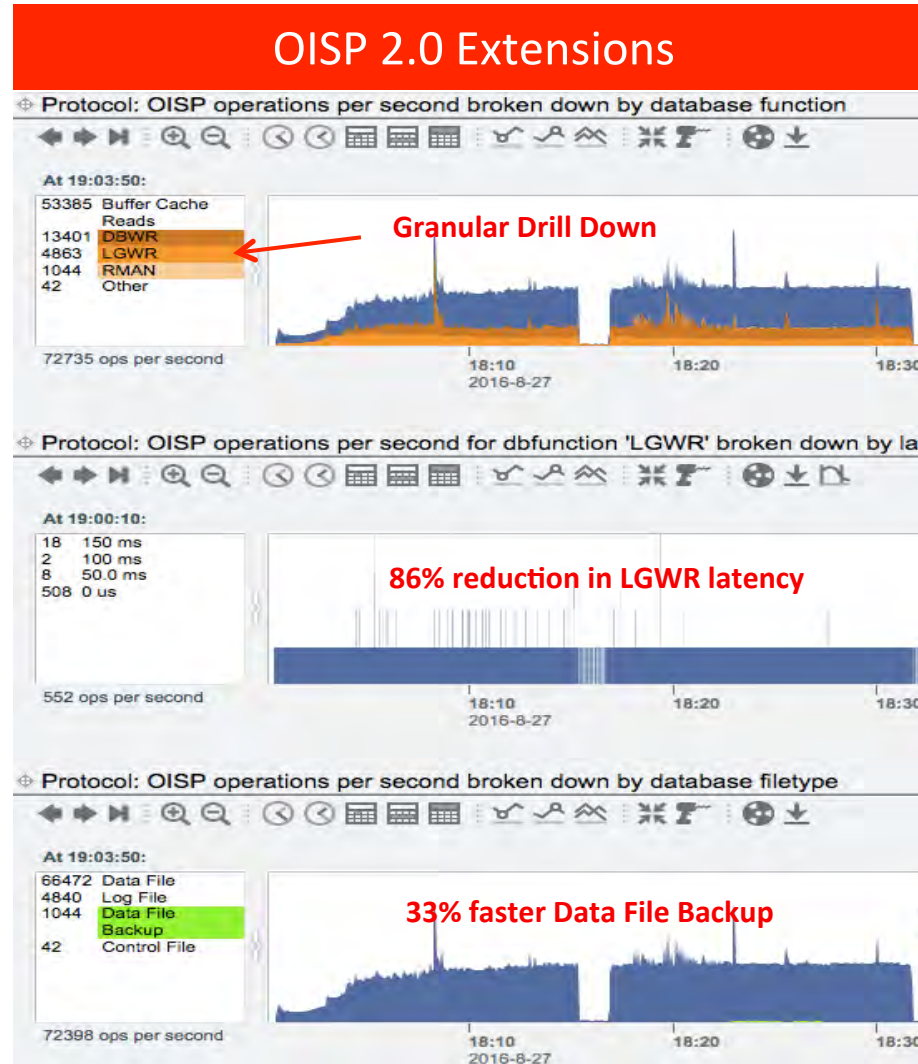


OISP 2.0 Database I/O Prioritization

- **12c Database Analytics**

- DBWR
- LGWR
- RMAN

- Bounded latency for higher priority DB I/O
- Negative cache hints for RMAN and Archive log I/O
- No ARC growth for RMAN, Archive log I/O
- Reduced ARC evictions!



Oracle's **NAS** and **Data Protection** System

- **Best Performance. Best Efficiency. Best Security.**
- Trusted by a majority of the F100 and Oracle IT
- Foundation storage for Oracle Cloud
- Engineered for Oracle Software

