



6 Reasons why to run Oracle on NetApp

HROUG 2007, Rovinj

Helmut Putzenlechner

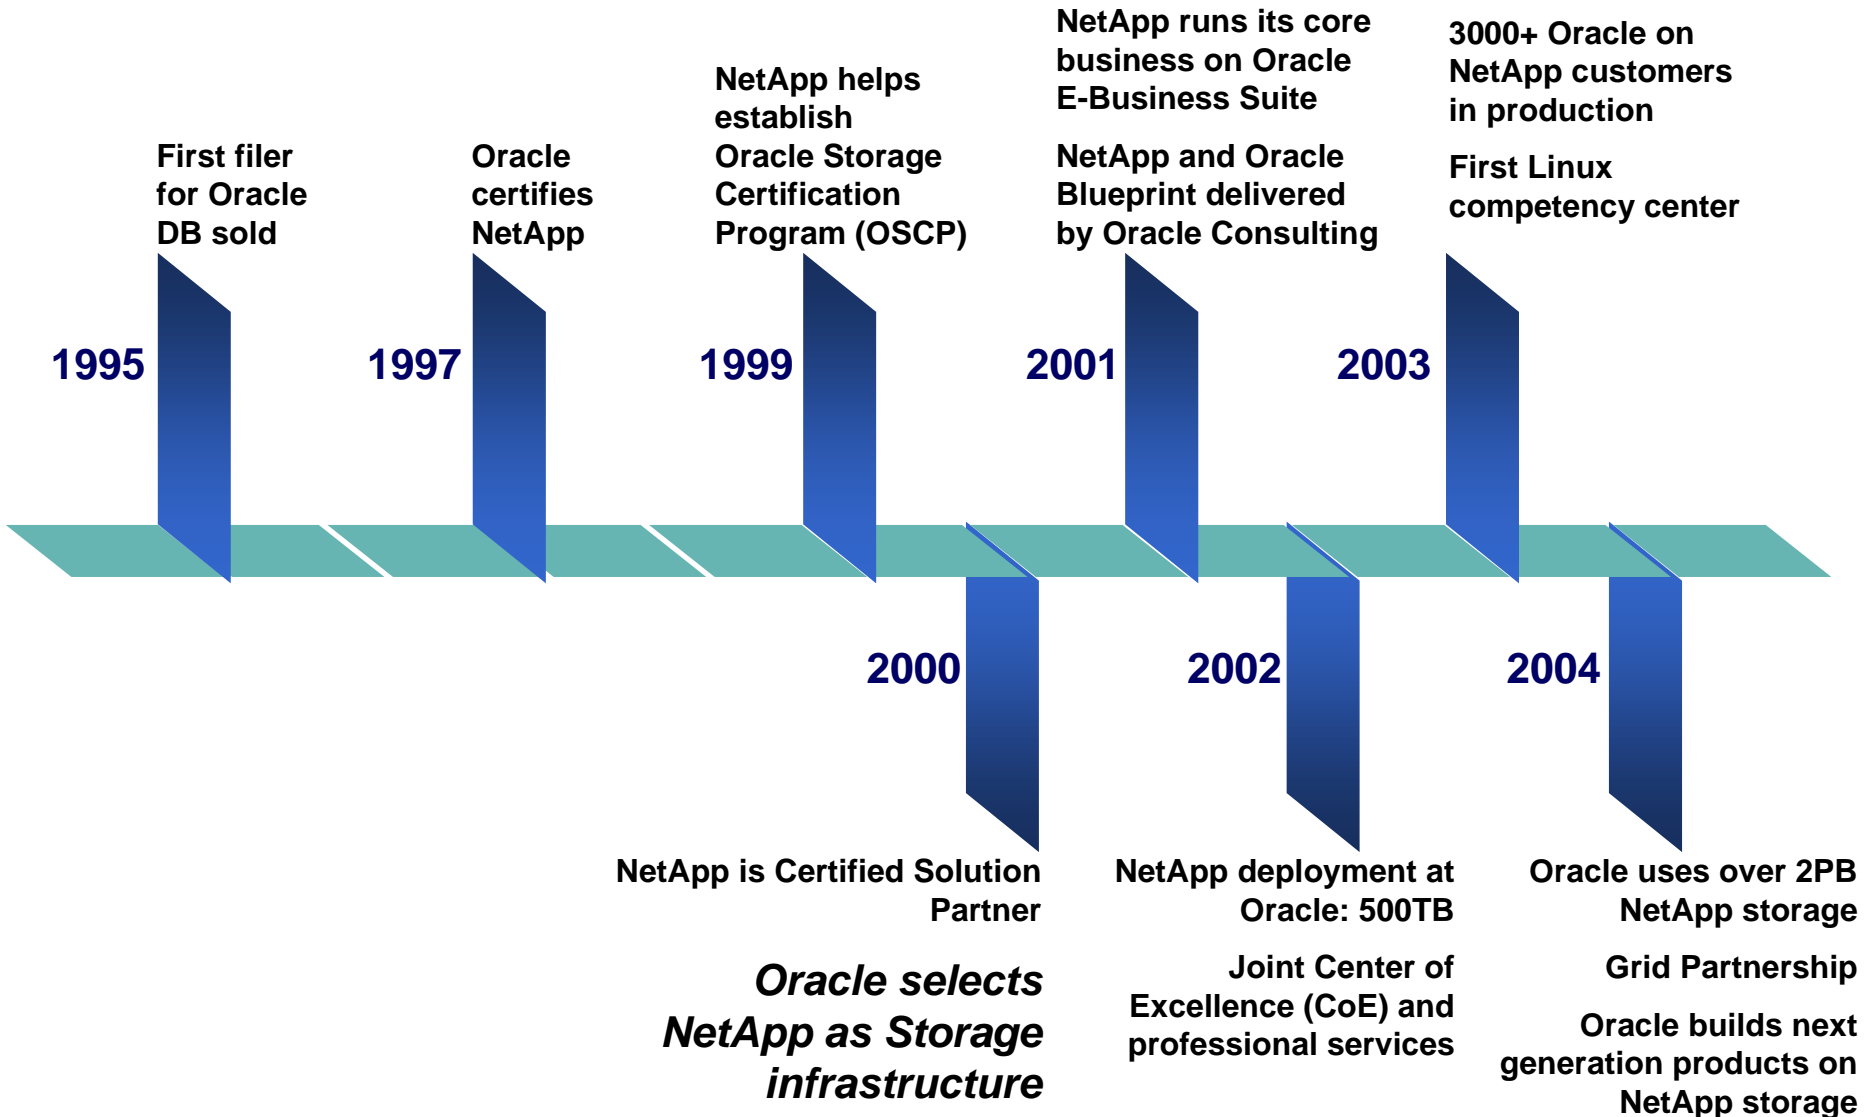
Presales Consultant South-East-Europe

Introduction – Did you know ...?

- ▶ **Shorter Test and Development Cycles**
- ▶ **Easy Restore of partial data loss**
- ▶ **Reduce Restore-Downtime**
- ▶ **Simplify Space Management**
- ▶ **DR-Solution with Transparent Site Failover**
- ▶ **„Boost“ random-write-performance**



Oracle and Network Appliance® Evolution of a Strategic Alliance



Did you know that Oracle “runs” on NetApp? e.g. The Austin Datacenter



- ▶ 15,000 servers
 - ▶ 4,200 TB of NetApp storage
 - ▶ Adding 100 servers and 15TB per week
 - ▶ 400+ On Demand customers
- More than 50% of Oracles total Storage is on NetApp

ORACLE

NetApp - Data ONTAP 7G



...joint Development

→ „FlexVol“ & „FlexClone“



**Reduce
complexity**

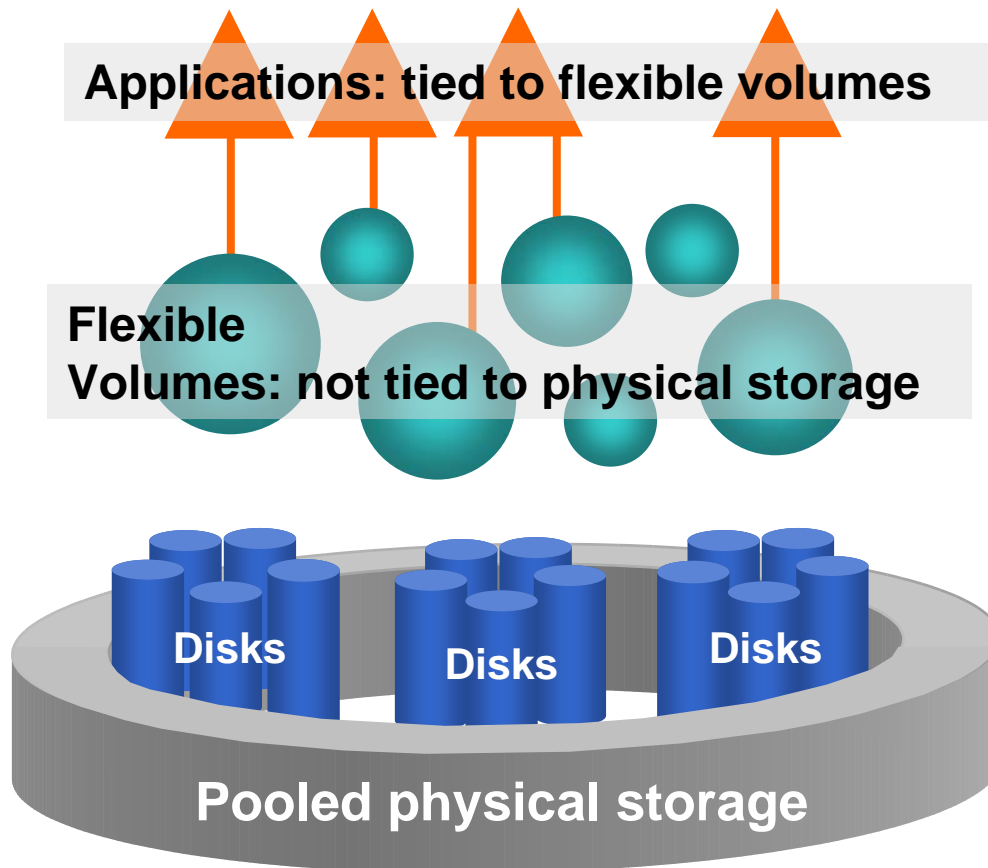
Minimize risk

**Control
change**

- ▶ **“Data Management”**
 - **NOT “Storage” Management**
- ▶ **Application Integration**
- ▶ **Unified Storage Systems**

FlexVol


Manage Data, not Disks!



- ▶ Grow and **shrink** Volumes as you like
 - **Better Utilization**
- ▶ Stripe over big number of spindles (NOT concatenate)
 - **Better Utilization**
- ▶ Allows „Thin Provisioning“
 - **Deploy only as much capacity as really used**


FlexClone: Stop creating more copies - and be fast!

General:



- ▶ Ideally multiple copies of “real production DB” for high quality test results and development

NetApp DataOntap 7G:



- ▶ **FlexClone Technology:**
 - create clones in seconds
 - no data has to be moved

DB Development Test Cycles: FlexClone makes it fast!

Example: Cloning a 100GB LUN

Legacy

- Create new LUN of same RAID type and size
- Create clone
- Change sync priority to high
- Sync LUNs
- or: copy the whole DB!

FlexClone

```
> vol clone create <volname>
```

```
→ all LUNs in volume cloned
```

```
→ step back again and again
```

shorter Testcycles

= higher Application Quality

= faster Business Processes




~ 1 hour (?)




30 seconds

Most efficient Storage Utilization with FlexVol and FlexClone



- ▶ **Writable copies consume x-times space of DB-size**



- ▶ **FlexClones consume space only for the changed data**
- ▶ **FlexVol sizes can grow and shrink fully online**
 - **adjust as you need!**

Restore partial data loss

FlexClones makes it fast and easy!

***Large Databases (DWH etc.):
Partial Data Loss caused by DB-Users,
hard to restore when traditional backup is used***


But with FlexClone, “a single command”:

- `vol clone create <volname>`
- **Mount the instance temporarily from the clone
(concurrently with the production instance)**
- **Export the data which was lost from the temp. instance**
- **Import the data to the production instance**


**SnapManager for Oracle does even this few steps fully
automatically by just a few mouse clicks**

Today's typical Backup/Recovery Strategy

Conventional Backup: 1x/day to Tape Library or Backup-to-Disk



Conventional Restore:



What is required:

Q: What is the solution to avoid that kind of downtimes?


Need to introduce fast-restorable Backups!

- Hardware is protected very well
- Data loss happens often during upgrades/migrations
- Upgrades can not be avoided
- Reason for Restore is „logical“ corruption of data
- Restore scenario must be verifiable proactively


NetApp offers the solution!

Reduce Downtime: Extended Backup Concept with Snapshot™

Snapshots as additional, frequent „online“ Backups:



Restore from SnapShot:

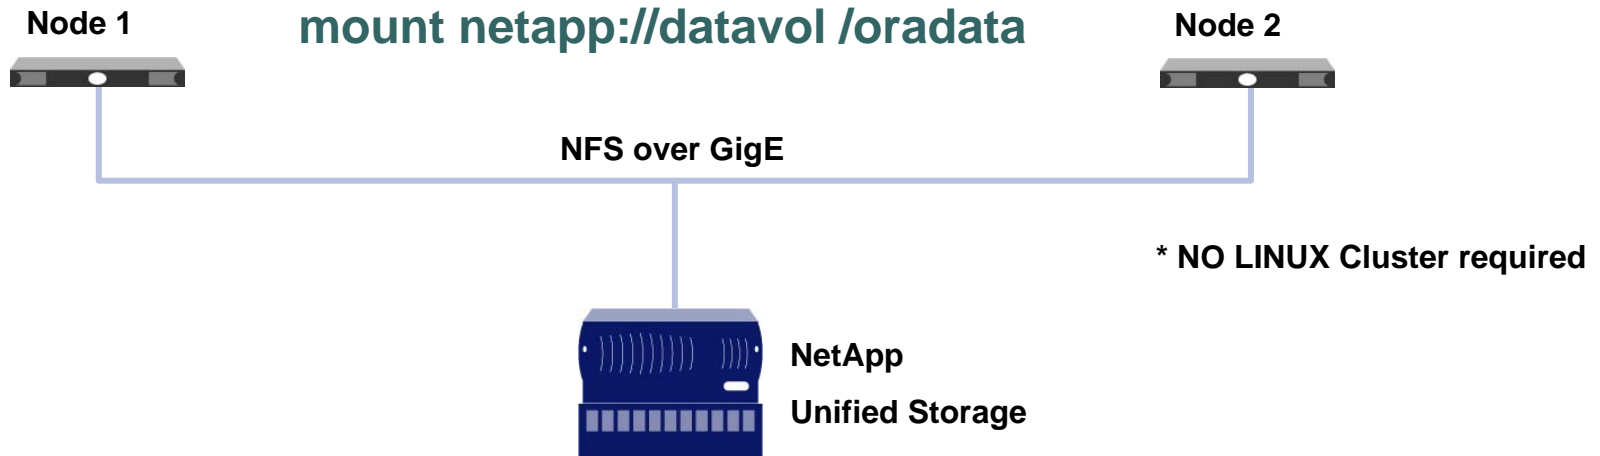


20min vs 6.5 hours Downtime → This is business value!

Simplify Management:

Example: Oracle RAC on Linux with NFS


Oracle RAC using **NFS** as shared Filesystem
(supported by Oracle **only** if running **on Netapp**)*



+ Integrated Framework of ASM, RMAN etc. with
SnapManager for Oracle

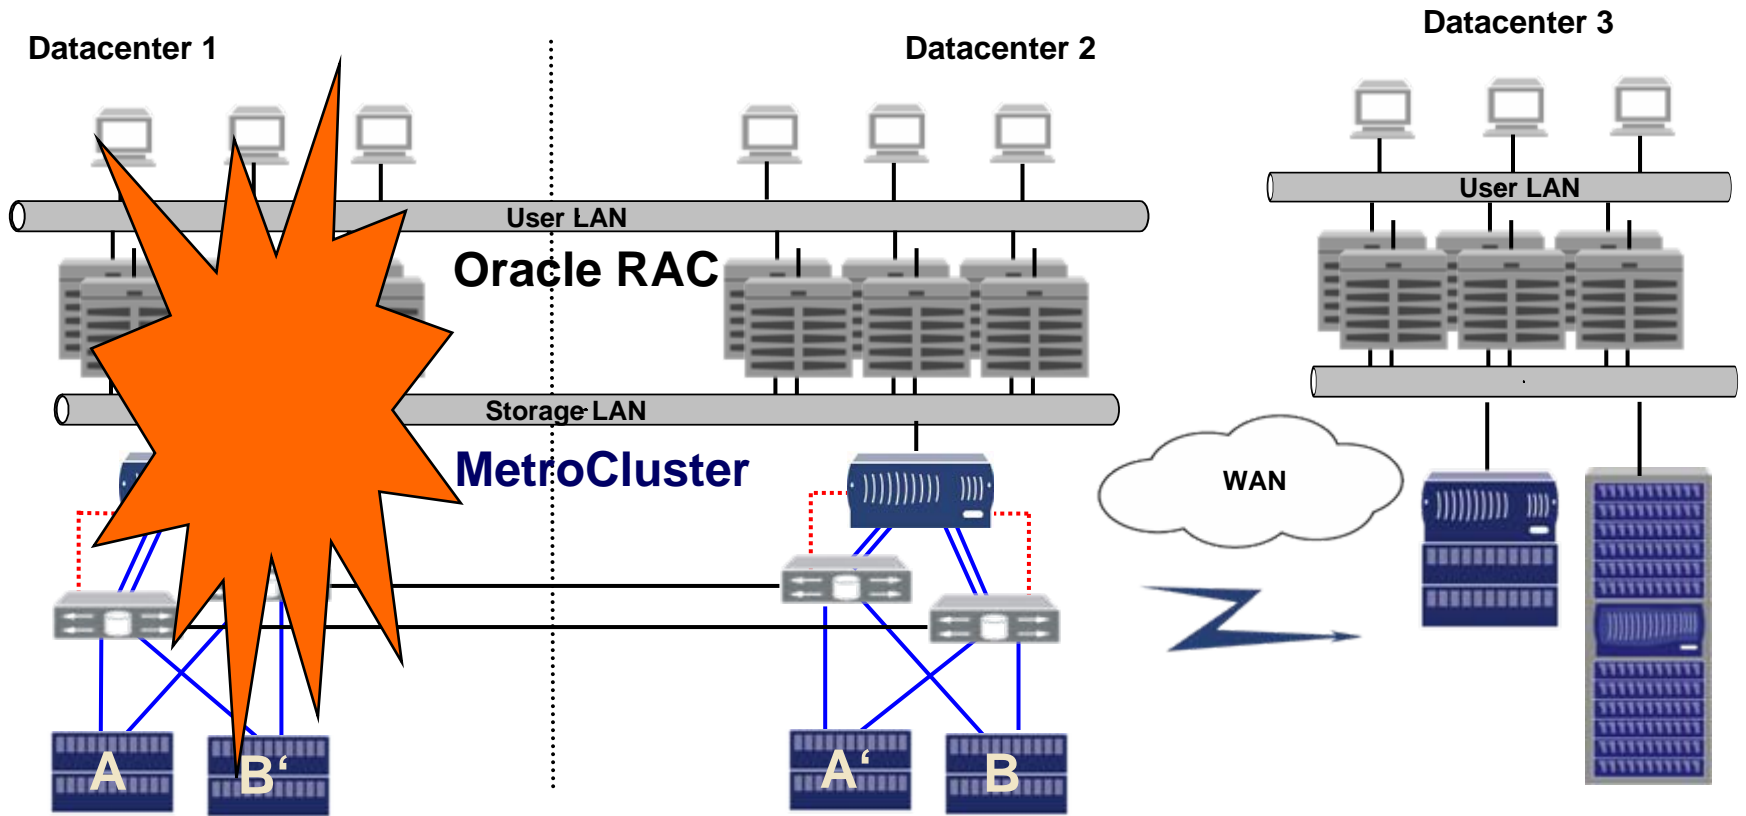
NFS? - What about performance?

**Oracle 10G RAC Performance (3 Nodes)
Order Entry Transactions (OETs) And Filer IOPS**



Note. Test results were obtained in a host limited environment where host CPU utilization was being driven to 100%.

DR with transparent Failover: RAC using a MetroCluster



- **full synchronous, active-active** „Backendmirror“
- Site failover by RAC or HA-Cluster
- HA & DR solution „in One“


Hot Standby ORACLE Failover With SnapMirror

- Asynchronous mirror
- IP based
- e.g. use DataGuard

Performance: Random I/O Comparison

Performance Test on a 400GB LUN done by Veritest

OLTP Workload: 60% Random Read, 40% Random Write



Why is there such a big difference with the ~ same number of disks

Look at the numbers:

CX3-80 LUN delivers 7230 random IOPS

– this is approximately what 32 Disk are able to deliver with random IO (32 x 200)

FAS3070 LUN delivers 14100 random IOPS

– this is more than the disks (28) would ever be able to deliver in random access mode


- ▶ **WAFL: Write Anywhere File Layout**
 - internal Datastructure of NetApp Systems
- ▶ **Implements special Write Cache destage algorithm**

Remark: WAFL is base for NetApp SnapShot Technology


WAFL – “The write performance boost” “Tetris” Cache Optimization

WAFL translates „small random writes“ into „large sequential writes“ !


Host IO Queue



Write Cache



WAFL Filesystem
on Disks



Disk Performance Principle:

Example:

- 1.) Write one 256k block to a disk takes nearly the same time as writing on 8k block
- 2.) Write one 256k block is approx. 30 times faster than 32 x 8k

- ▶ **Shorter Test Cycles: (+ most efficient Utilization)**
 - **FlexClone**
- ▶ **Restore partial data loss:**
 - **FlexClone**
- ▶ **Reduce Downtime:**
 - **SnapShot™**
- ▶ **Simplify Management:**
 - **SnapManager, Oracle on NFS**
- ▶ **Transparent Site Failover:**
 - **MetroCluster**
- ▶ **„Boost“ random-write-performance:**
 - **WAFLs „Tetris“ Cache Algorithm**



ORACLE CERTIFIED
PARTNER

Simplifying Data(*base*) Management

Learn more:

<http://www.netapp.com/products/>