



OSL Storage Cluster and RSIO

A perfect couple

Christian Schmidt
Systems Engineer



Agenda

1. OSL Storage Cluster – an overview
2. What can you do with RSIO Block I/O over Ethernet
3. OSL Storage Cluster and RSIO – an integrated solution
4. Applications in a cluster
5. Backup and Instant Recovery

- *Solution to integrate UNIX servers with modern RAID-based storage infrastructures*
- *adds virtualisation, management and cluster functions to the operating system*
- *storage management, virtualisation, system and application management as well as clustering are understood as one unit*
- *simplification of procedures and administrative tasks in the data center*

Core of the OSL Storage Cluster:

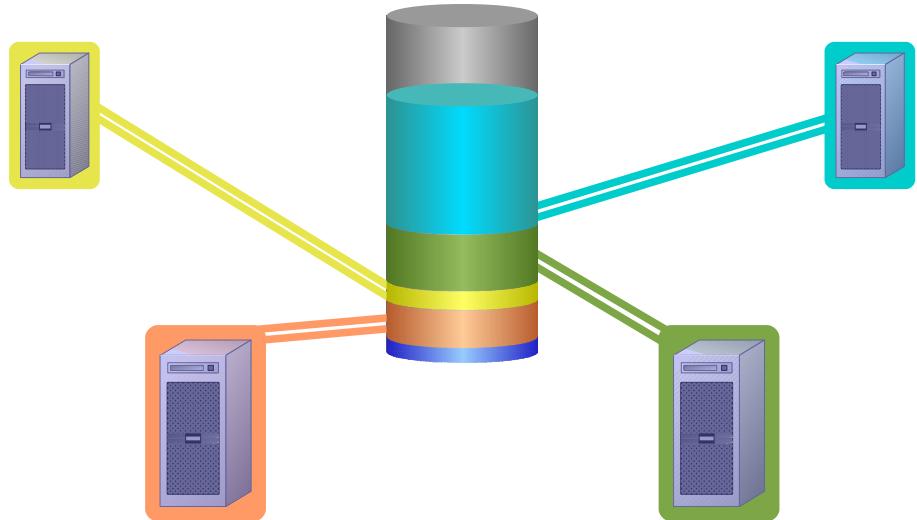
global, host-based storage virtualisation

Storage virtualisation

The OSL solution



- *hardware-independent storage virtualisation with standard components*
- *global storage pool*
- *volumes with unique, arbitrary names*
- *storage administration from any host possible*
- *application aware storage management*



OSL Gesellschaft für offene Systemlösungen mbH

www.osl.eu

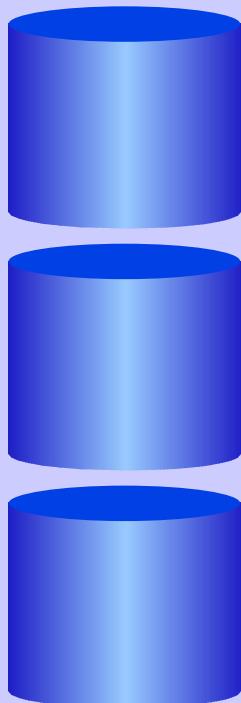
Storage virtualisation

The OSL solution



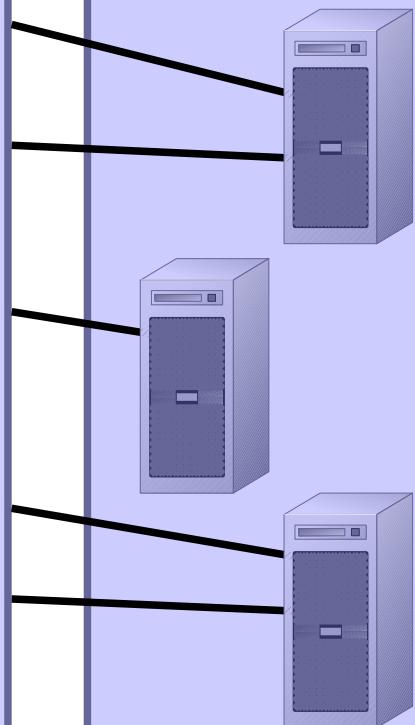
RAID-View

representation of internal resources



Host-View

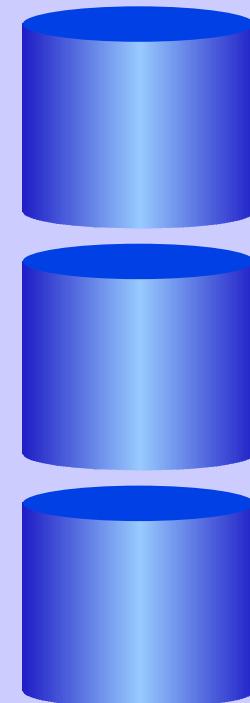
hardware- and host-dependent representation of external resources



Step
1

Virtual Disk-View

cluster-wide, unique, hardware-independent representation of RAID resources

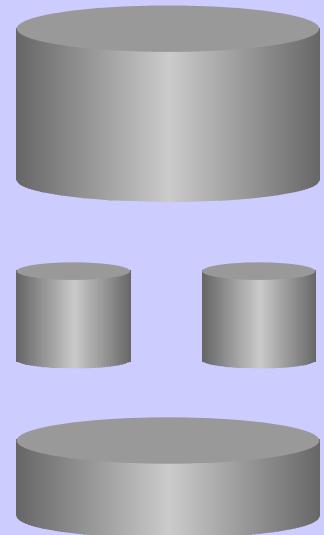


Physical Volumes

Step
2

Free defined Virtual Volumes

cluster wide, uniform, hardware independent deployment of arbitrary virtual volumes



Application Volumes

What it looks like

- *Physical Volumes*

```
[root@rssrv1] pvadmin -lvv
0 rsio_ccf (ok) 4193400 blocks over 1 path(s)
 >[ 1] (ok) /dev/rlofi/1
0 stripe_d1 (ok) 209680380 blocks over 1 path(s)
 >[ 1] (ok) /dev/rdsk/c1d0s3
```

- *Application Volumes*

```
[root@rssrv1] avadmin -lvv
0 clnt1_ora 409600 of 409664 blocks "concat,master" in 2 pieces, 32 block clusters
 [ 1] disk1_s4 [0...307263]
 [ 2] disk1_s4 [161125280...161227679]
0 clnt1_sap 4194304 of 4194368 blocks "simple,master" in 1 pieces, 32 block clusters
 [ 1] disk1_s4 [307264...4501631]
```

- Application aware storage overview

- How much storage uses my application?

```
#> smgr -qa bp1@0
```

```
#> smgr -qa bp1@0
used by bp1@0      :          33544 MB          33 GB          0.032 TB
-----
TOTAL STORAGE POOL SUMMARY
free:          8603311612 bl          4200836 MB          4102 GB          4.006 TB
totl:          10359112860 bl          5058161 MB          4940 GB          4.824 TB
-----
```

- *Application aware storage overview*

- *How much storage uses my application?*
- *Has my application any path failures?*

```
#> pvadmin -lvva bp1@0
```

```
#> pvadmin -lvva bp1@0
0 p01 (ok) 585920024 blocks over 1 path(s)
  >[ 1] (ok) /dev/rdsk/c1t5000402001EC04F4d1s1
    [ 2] (ok) /dev/rdsk/c2t5000402001EA01F4d1s1
0 p04 (ok) 585920024 blocks over 1 path(s)
  >[ 1] (ok) /dev/rdsk/c1t5000402001EC04F4d4s1
    [ 2] (ok) /dev/rdsk/c2t5000402001EA01F4d4s1
0 p02 (ok) 585920024 blocks over 1 path(s)
  >[ 1] (ok) /dev/rdsk/c1t5000402001EC04F4d2s1
    [ 2] (ok) /dev/rdsk/c2t5000402001EA01F4d2s1
INFO (pvadmin) : all physical volumes for application bp1@0 attached
```

Storage virtualisation

The OSL solution – application aware storage management



- Application aware storage overview

- How much storage uses my application?
- Has my application any path failures?
- What is the status of the mirror application?

#> avmirror -qa bp1@2

```
#> avmirror -qa bp1@2
2      ora_BP1 ( simple, 1pc, 4096m) image   -      s0123 disconnected
2      sapmnt_BP1 ( simple, 1pc, 500m) image   -      s0123 disconnected
2      usrsap_BP1 ( simple, 1pc, 500m) image   -      s0123 disconnected
2      origlogA_BP1 ( simple, 1pc, 200m) image  -      s0123 disconnected
2      origlogB_BP1 ( simple, 1pc, 200m) image  -      s0123 disconnected
2      mirrlogA_BP1 ( simple, 1pc, 200m) image  -      s0123 disconnected
2      mirrlogB_BP1 ( simple, 1pc, 200m) image  -      s0123 disconnected
2      saparch_BP1 ( simple, 1pc, 6144m) image  -      s0123 disconnected
2      sapdata1_BP1 ( simple, 1pc, 20g) image   -      s0123 disconnected
9 mirror instances (0 av missing), 0 synced, 5 active, 0 need maintenance
```

Storage virtualisation

The OSL solution – application aware storage management



- *The application aware storage management helps the administrator*

- Overview of applications storage usage
- Meet service level agreements
- Easy to use – consistent command line syntax

#> pvadmin -qa <APP>

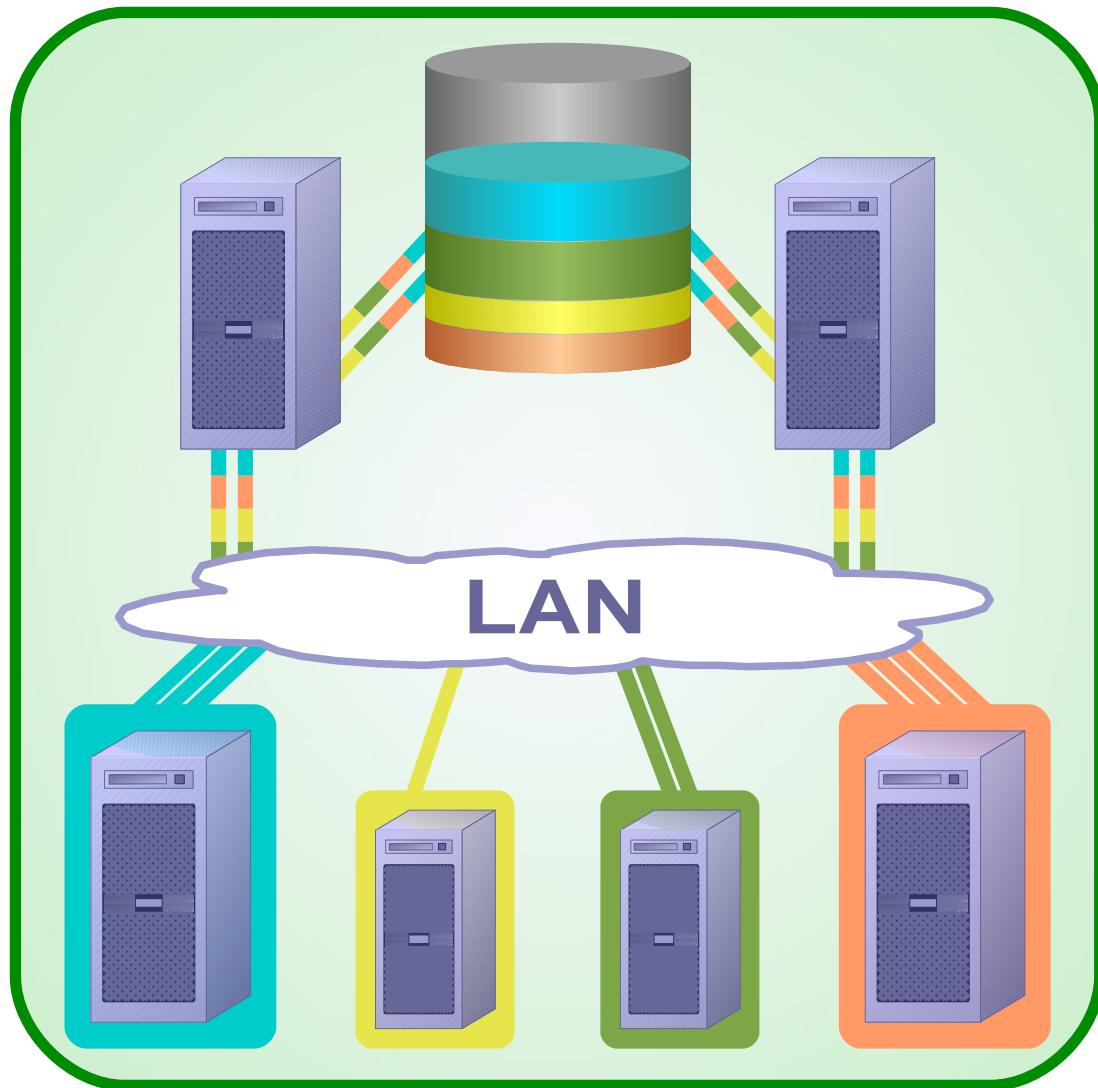
#> avmirror -qa <APP>

#> smgr -qa <APP>

#> smgr -c volume -S 1g -a <APP>

#> pvadmin -lvaa <APP>

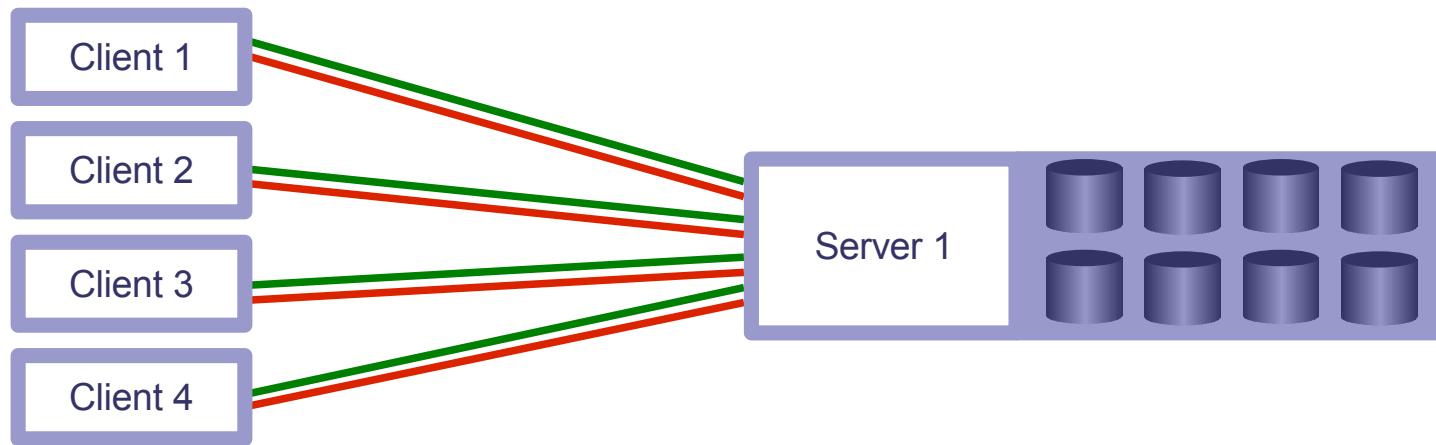
RSIO – Block I/O over Ethernet



OSL Gesellschaft für offene Systemlösungen mbH
www.osl.eu

RSIO – Block I/O over Ethernet

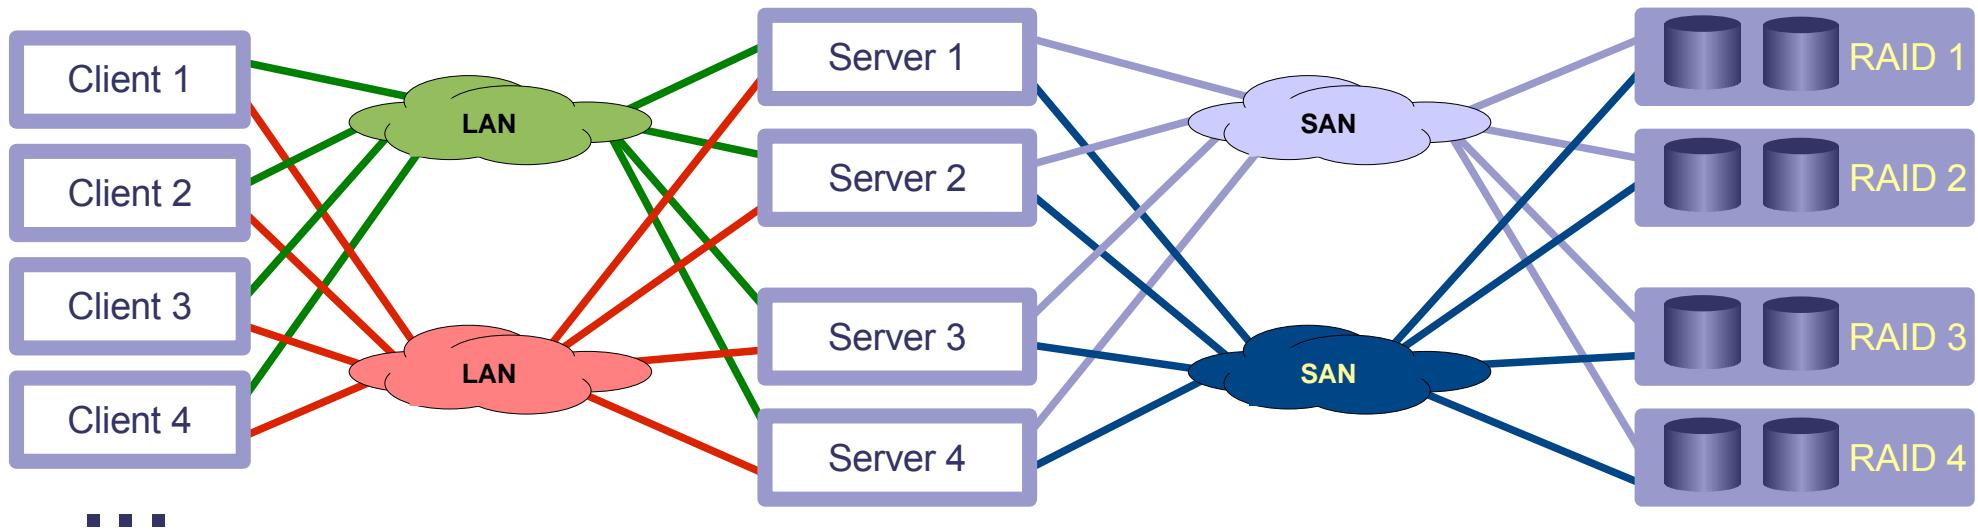
Conquer new platforms with RSIO



- Access to a central storage system → *global pool, global namespace*
- A central point for backups, snapshots, ...
- reasonable storage connect with good performance
- redundant data paths, scalable throughput

RSIO – Block I/O over Ethernet

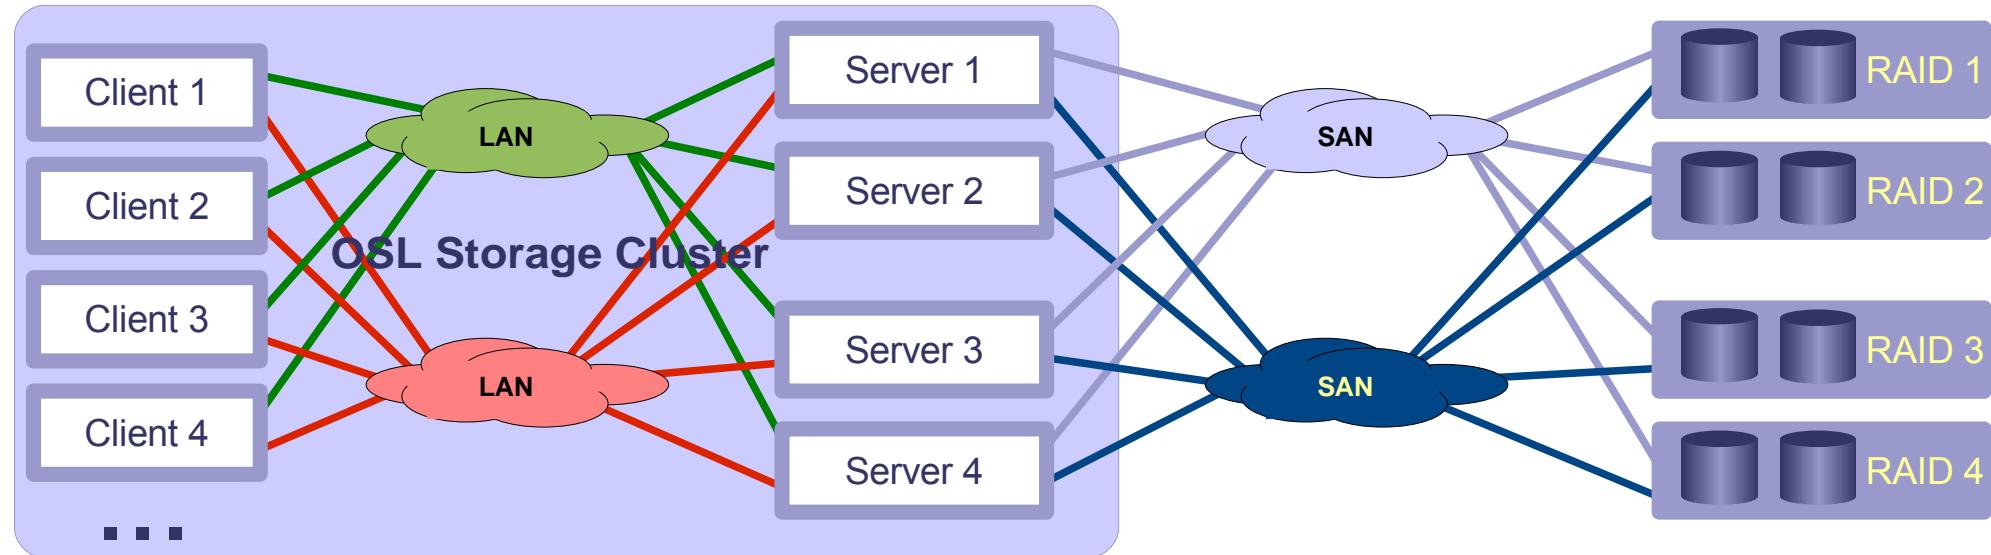
Conquer new platforms with RSIO



- SAN servers act as gateways from LAN to SAN
- improved utilization of SAN, performance rightsizing
- high performance and high availability at low cost for RSIO clients
- available server memory used as disk cache improves performance and system utilization

RSIO – Block I/O over Ethernet

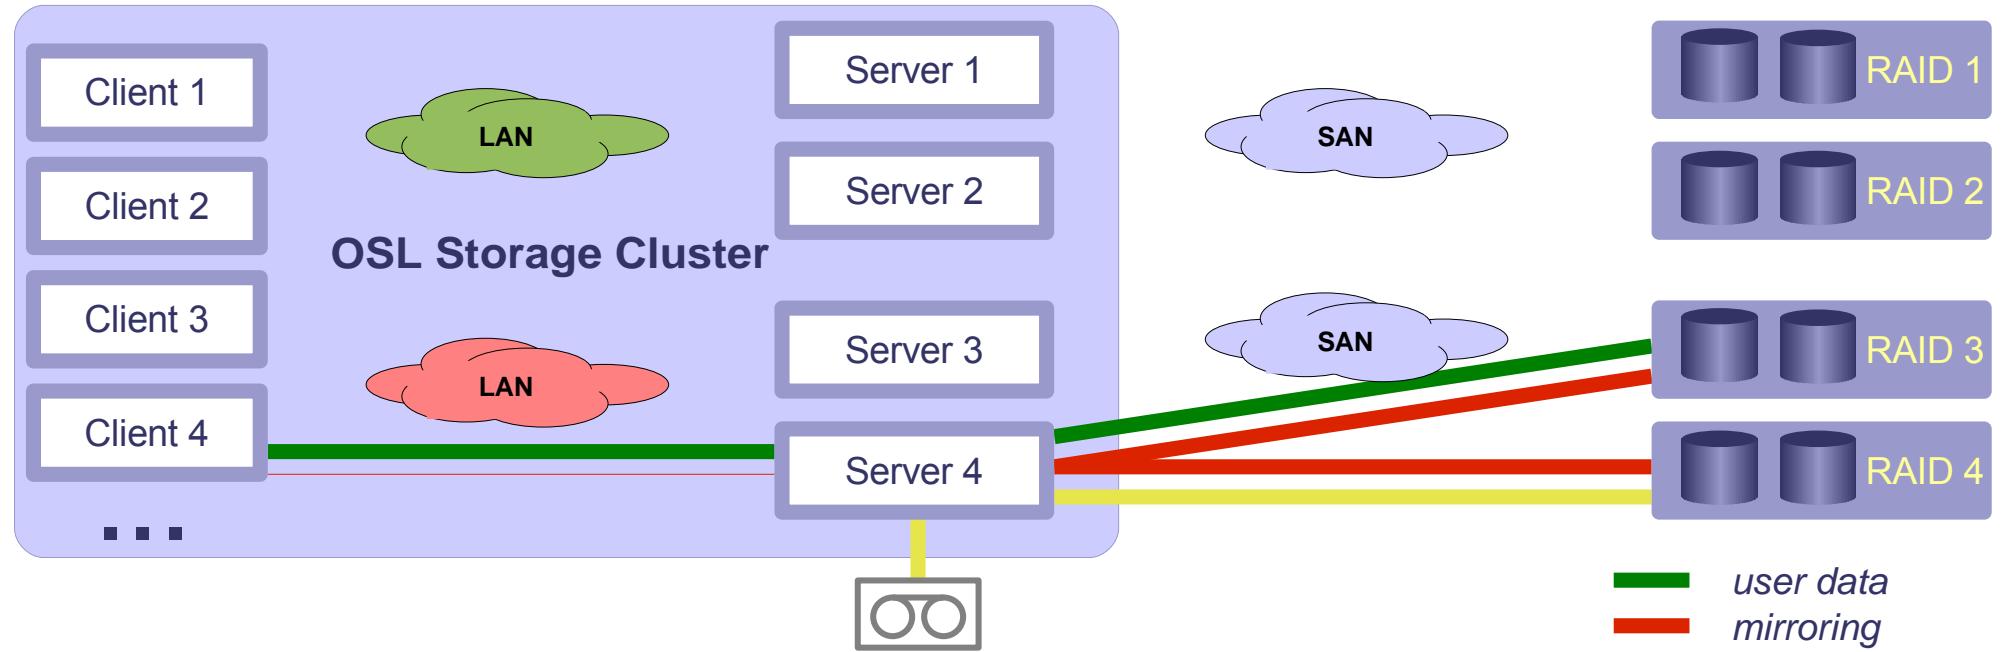
Conquer new platforms with RSIO



- *all functions as in the previous example*
- *additional storage management features:*
 - *storage allocation and management from client side*
 - *application aware storage virtualisation*
 - *host based storage mirroring, cloning and moving on the client side*
- *client and server merge into one unit*
- *run applications everywhere*

RSIO – Block I/O over Ethernet

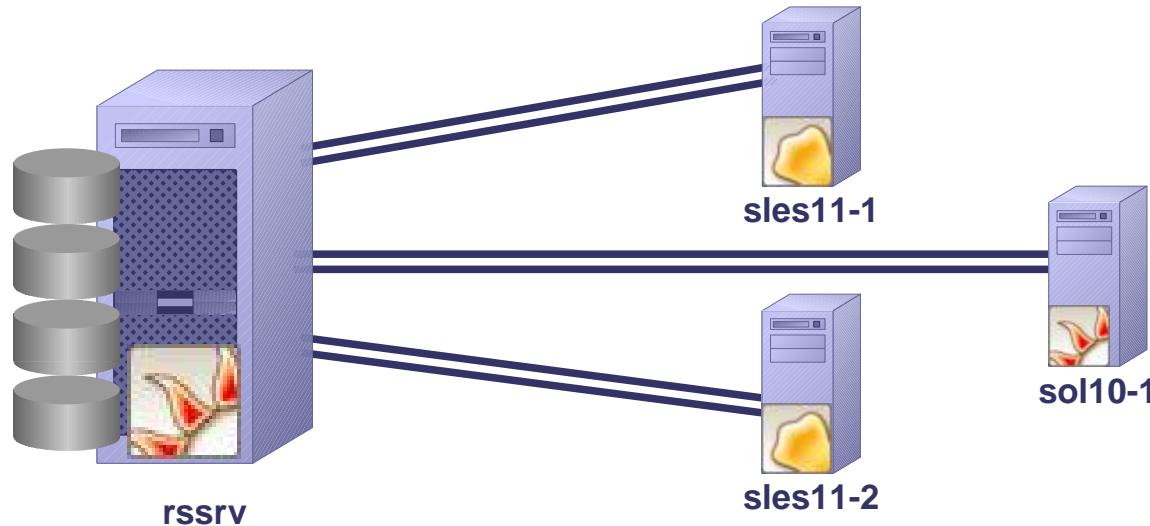
Conquer new platforms with RSIO



- only user data and control commands over LAN
- LAN-less backup:
 - the client has full control over backups, mirrors, clones and moves
 - high speed backup
 - application aware storage actions

RSIO and OSL Storage Cluster

Simple integration of Linux and Solaris nodes in one cluster



- Two steps to integrate a RSIO client in the OSL Storage Cluster

1. Installation of RSIO and configuration

- After installing the RSIO package the configuration can be done with the pseudo graphical tool `rsmgr`
- Available namespaces, servers and interfaces will be discovered and served for configuration

2. Installation of OSL Storage Cluster Software

- After installing the OSL Storage Cluster packages the `rsio` client can immediately start cluster applications

RSIO and OSL Storage Cluster

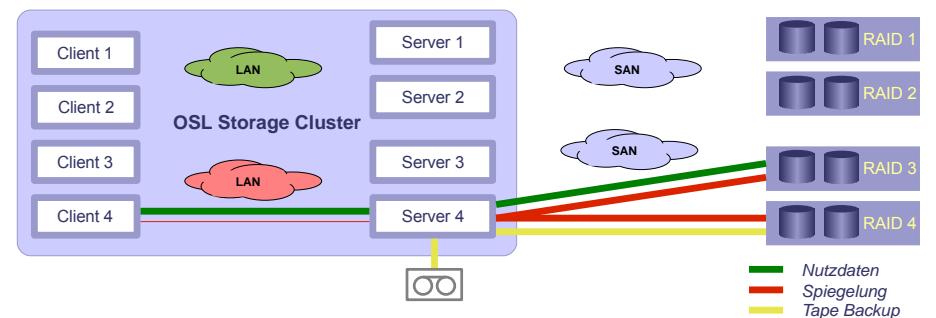
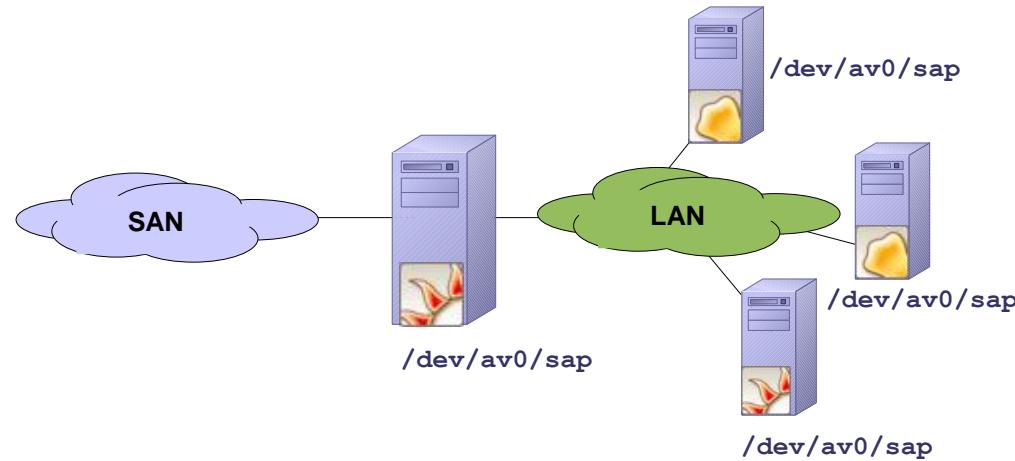
A perfect couple



OSL Storage Cluster adds virtualisation, management and cluster features to the operating system

You can use these features with RSIO on Linux!

- consistent, cross-platform namespace
- global storage pool is available on all RSIO storage clients
- create volumes anywhere (client/server)
- XDM operations controlled by client



OSL Gesellschaft für offene Systemlösungen mbH
www.osl.eu

RSIO and OSL Storage Cluster

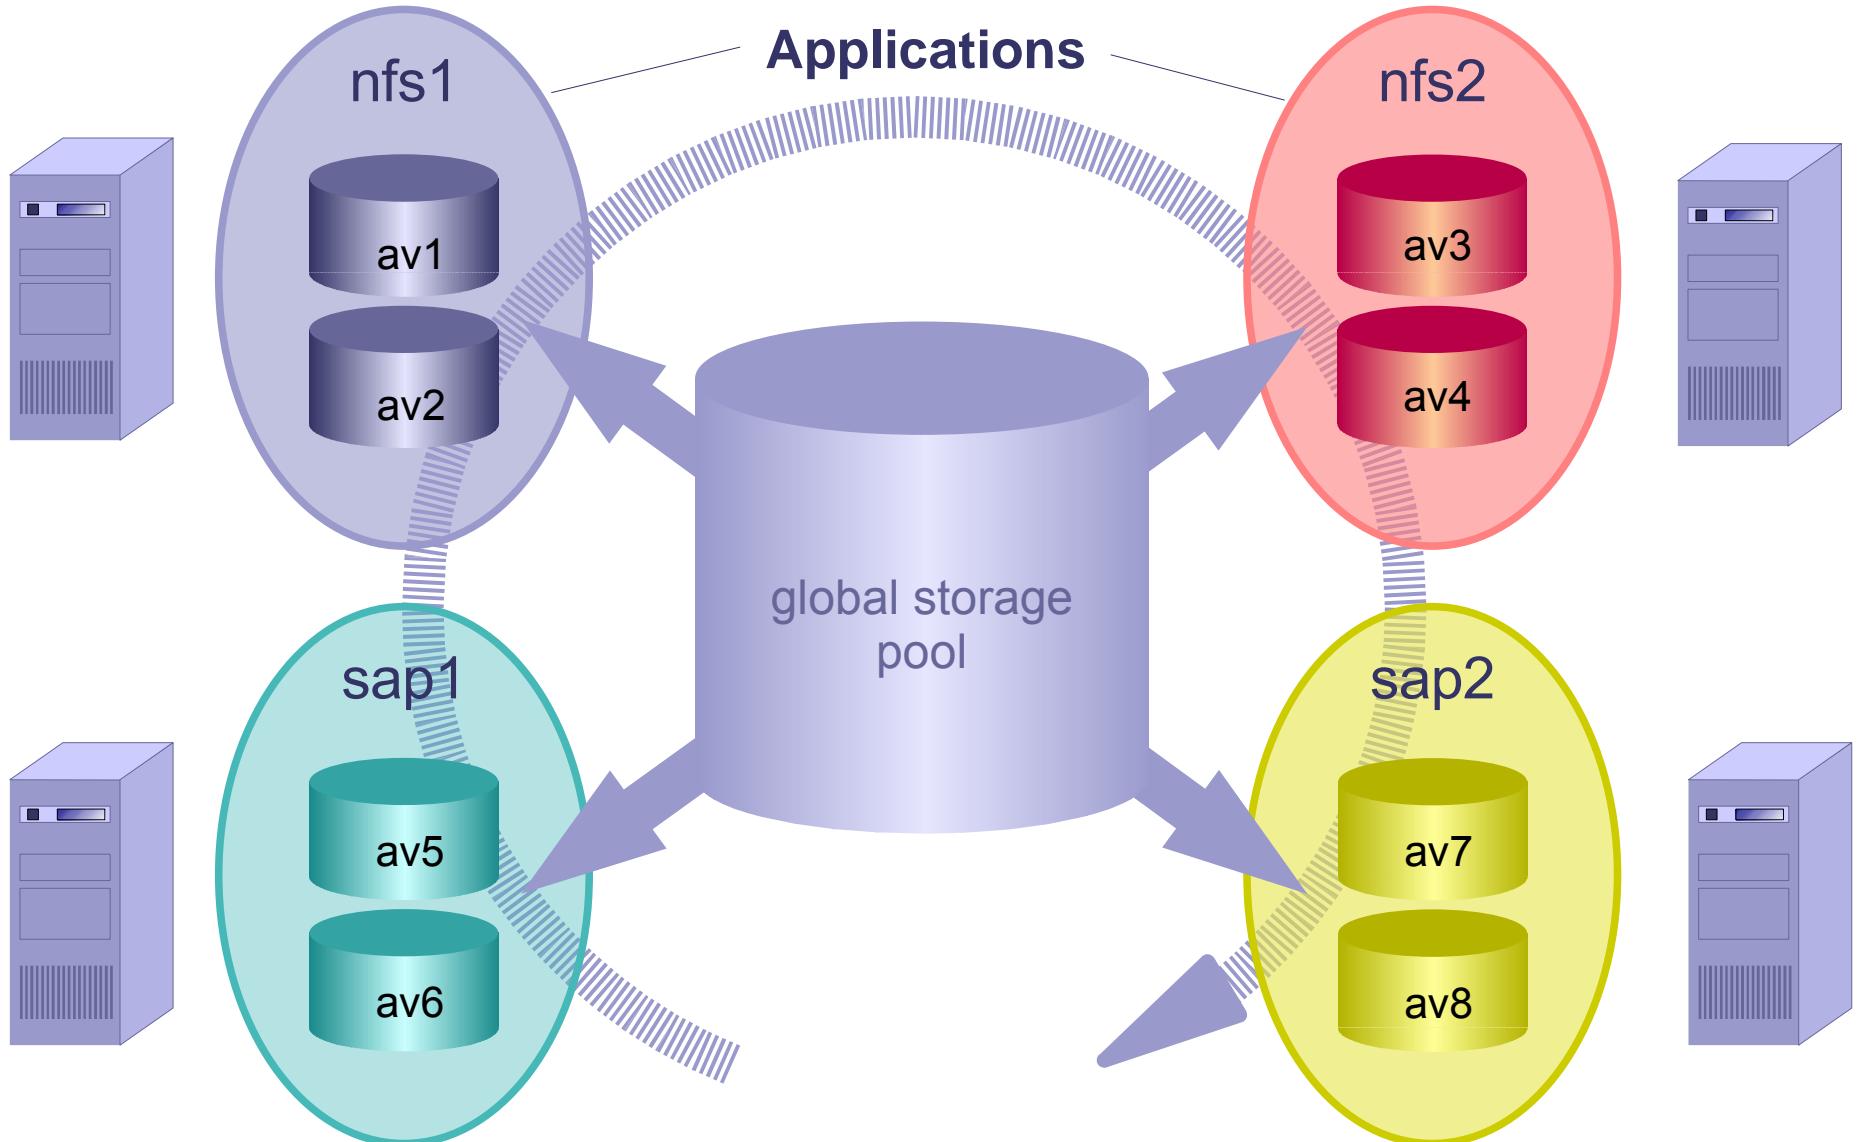
A perfect couple



- *Manage the storage from the RSIO server and any RSIO client*
 - *create and delete application volumes*
 - *mirror, clone and move application volumes*
 - *new created application volumes are automatically attached on all cluster nodes*
 - *The Disk Access Manager (damgr) locks concurrent access*

OSL Storage Cluster

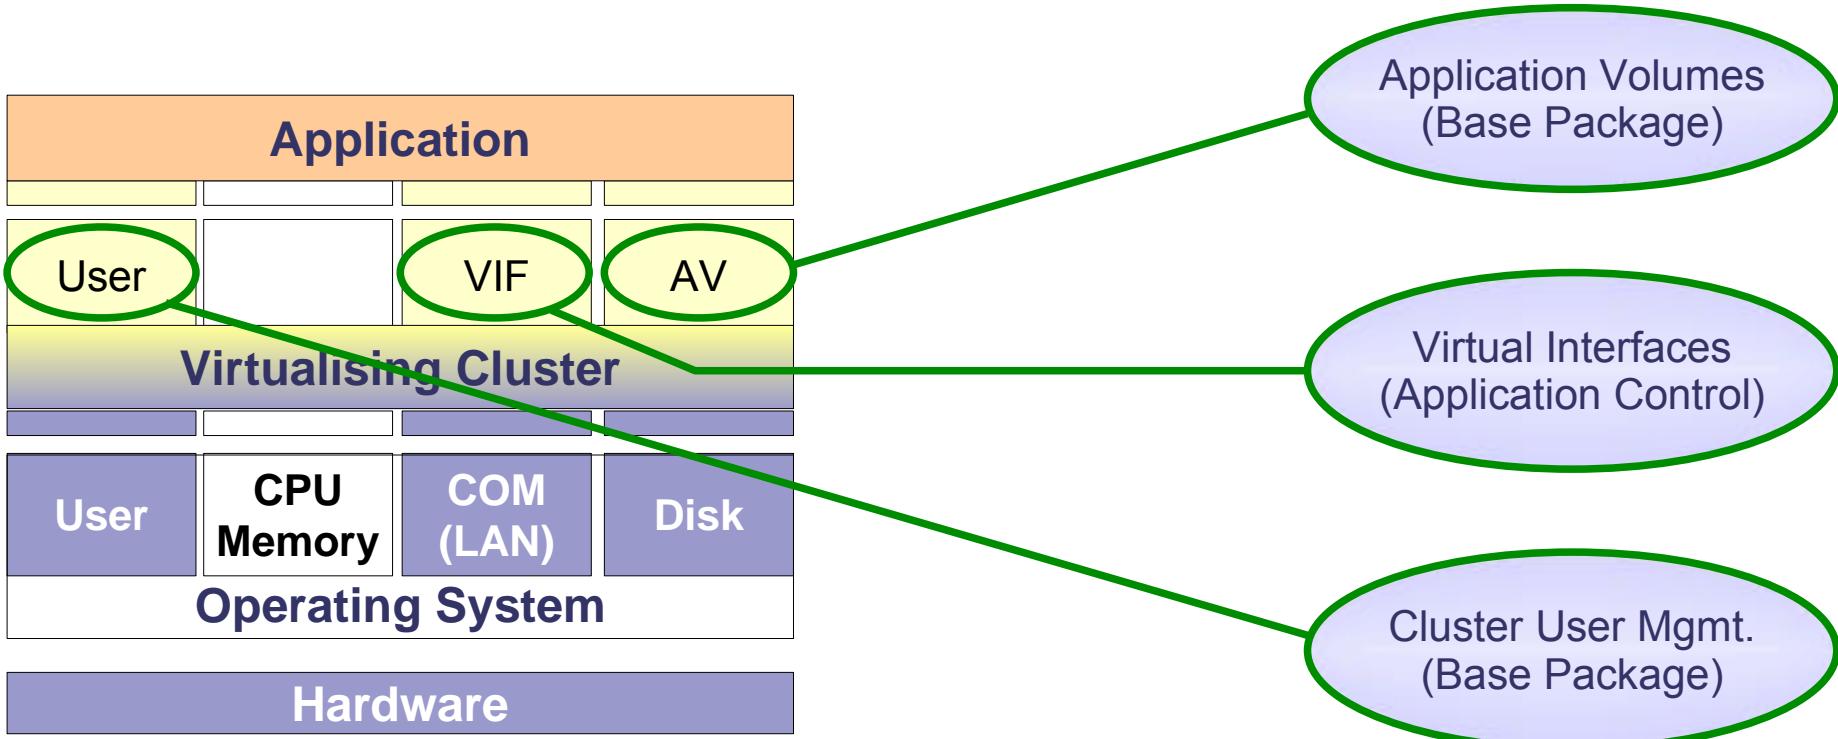
Storage virtualisation – H/A cluster



OSL Gesellschaft für offene Systemlösungen mbH
www.osl.eu

Virtual Runtime Environments

Without zones and virtual machines



Cluster capable, virtual runtime environments consist of:

- Anonym and virtualised resources of the operating system (CPU, RAM, VFS)
- virtualisation and cluster features of the OSL Storage Cluster

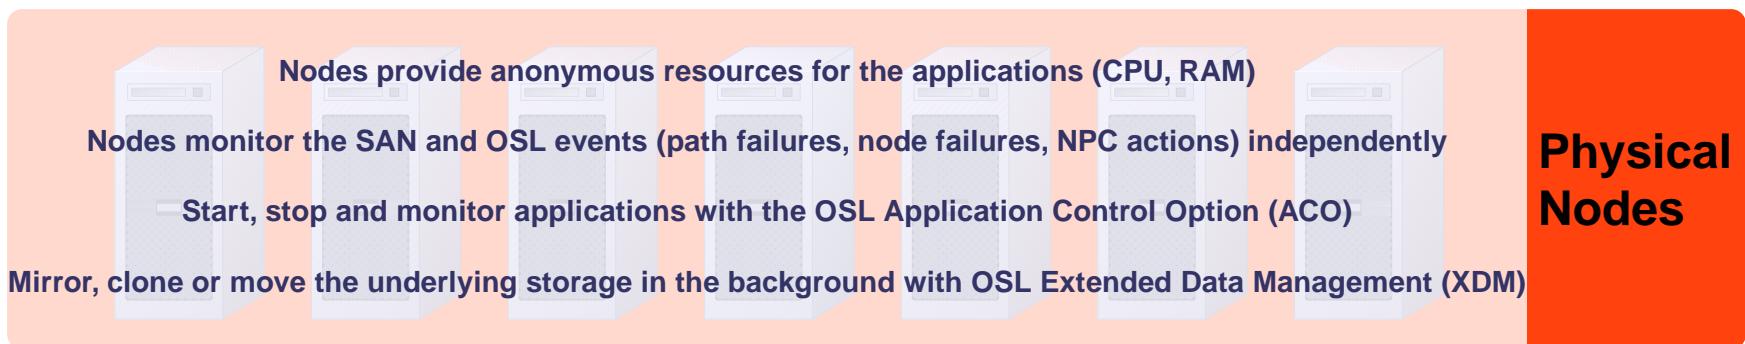
Using zones or virtual machines to run your application encapsulated is supported

OSL Storage Cluster

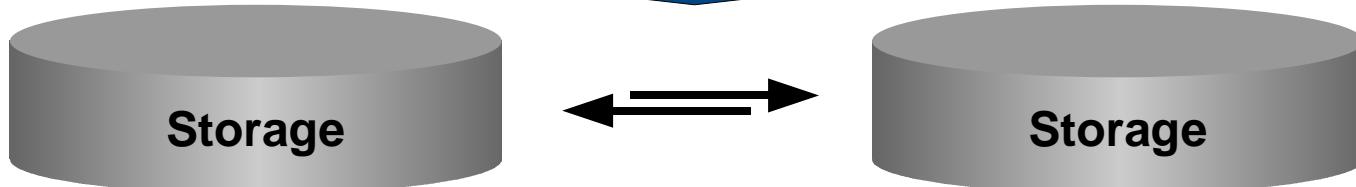
Nodes and Applications



Start, Stop, Monitor



Mirror, Clone, Move



OSL Gesellschaft für offene Systemlösungen mbH
www.osl.eu

RSIO and OSL Storage Cluster

A perfect couple



- *The storage connectivity does not matter – customer cares for applications only*
 - OSL Storage Cluster ACO is also available for RSIO Clients (even on LINUX)
- *What about Windows?*



Headlines:

- Open Virtualization Alliance launched
- Kernel Log: native KVM tool, new kernels galore
- Virtualisation API libvirt hits 0.9.0
- Novell supports KVM

RSIO and OSL Storage Cluster

A perfect couple

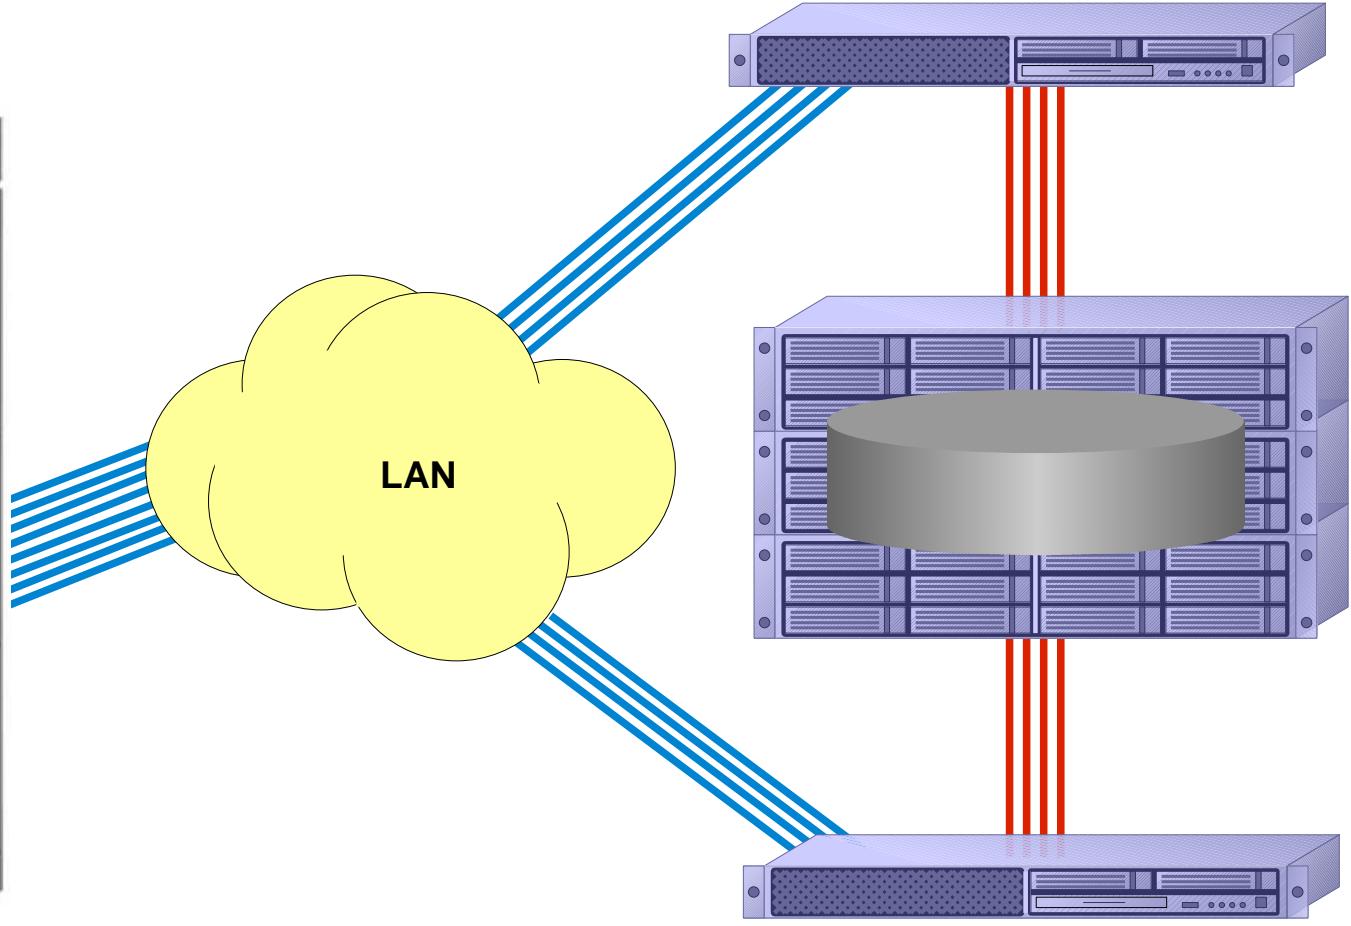


Short demonstration

- *KVM has a future*
 - *preferred virtualisation solution for Red Hat and Novell*
- *Simple administration*
 - *libvirt provides interfaces, even for other hypervisors*
- *OSL Storage Cluster integration*
 - *VM is an application*
 - *live migration between cluster nodes*
 - *benefit from the global namespace*

RSIO and OSL Storage Cluster

A perfect couple



OSL Gesellschaft für offene Systemlösungen mbH

www.osl.eu

RSIO and OSL Storage Cluster

A perfect couple

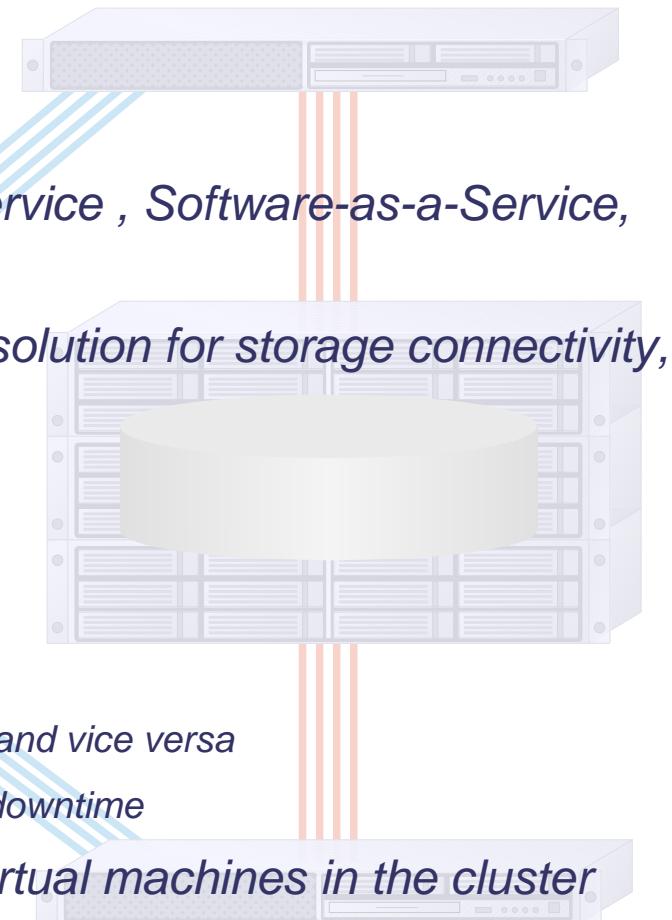


- Sample system: Fujitsu Primergy CX1000

- 38 node cluster with up to 456 processor cores
- useful for VDI, server VMs, Infrastructure-as-a-Service , Software-as-a-Service, Cloud Infrastructures
- OSL RSIO and OSL Storage Cluster is a perfect solution for storage connectivity, node monitoring and application high availability

- New features in Storage Cluster 4.0

- Virtual Node and Virtual Machine Applications
 - start applications on virtual nodes
 - migrate applications from physical nodes to virtual nodes and vice versa
 - move virtual machines from one node to another without downtime
- Toolset to deploy, migrate, backup and recover virtual machines in the cluster

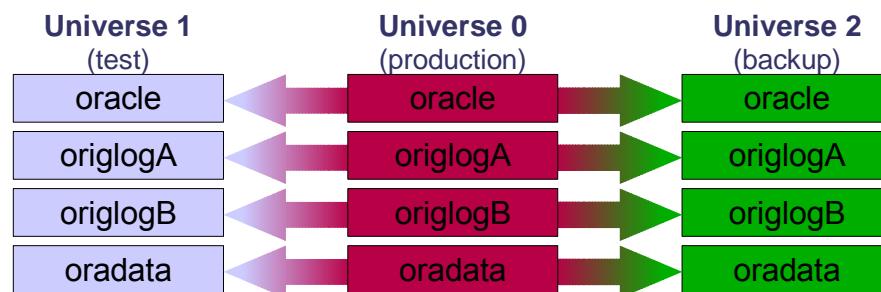


OSL Storage Cluster

XDM – The storage universes concept

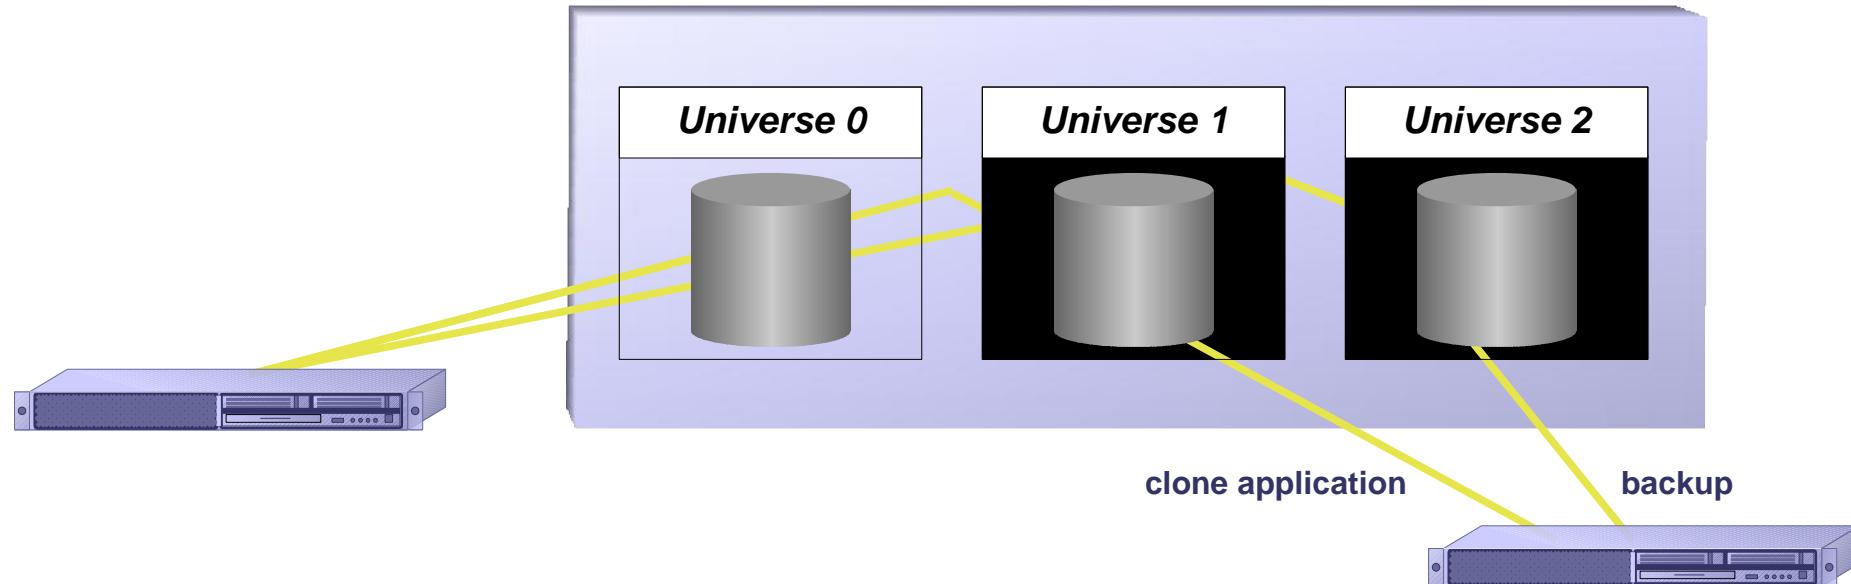


- in typical data centers storage and servers are divided into several sections (production, quality assurance, test)
- there are relationships:
 - test systems are copies from the production
 - backups are full copies of the production
- OSL Storage Universes shows which copy is assigned to which use-case (hardware failure, test, quality assurance)
- copies to another universe can be created at any time – mirror volumes get the same name as the master
- mirror universes can be connected and disconnected at any time (with incremental resynchronisation)
- on production failure, backup copy can be started immediately



OSL Storage Cluster

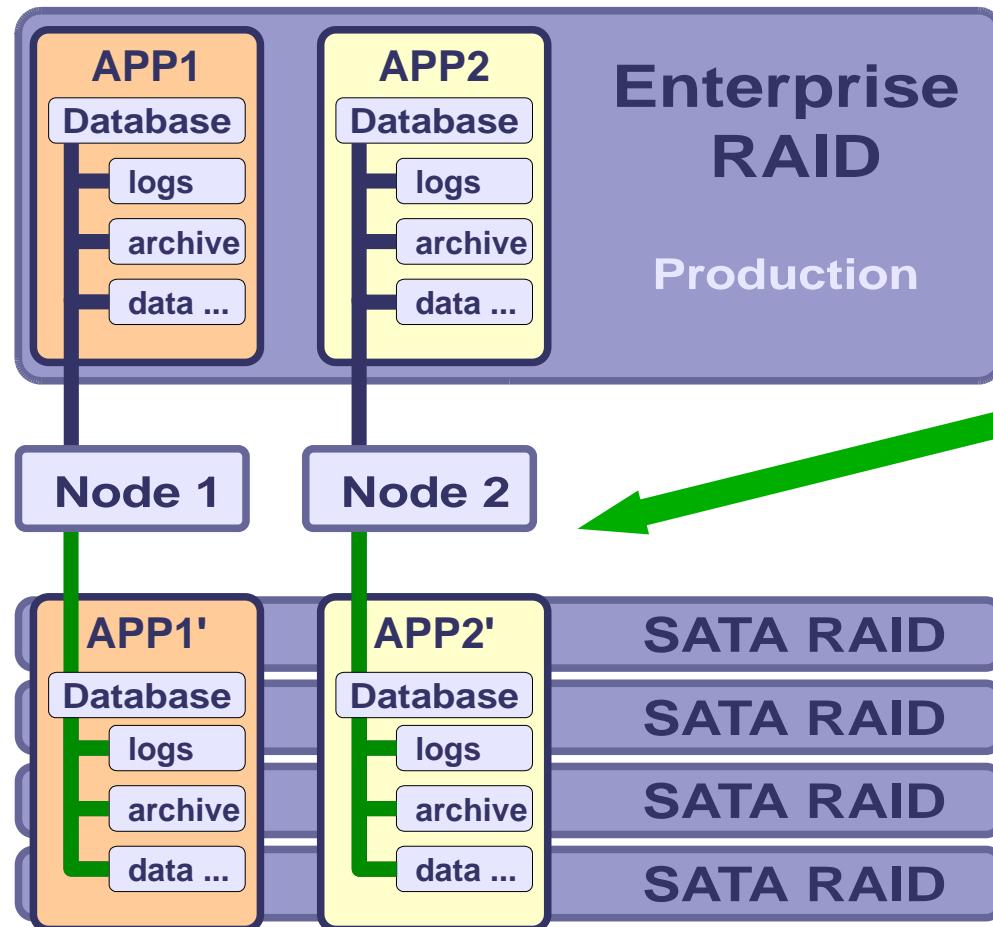
XDM – The storage universes concept



- *no hierarchical device nodes → every device can be mirrored at any time, even when the application is running*
- *image and master have the same name*
- *when the master volume fails I/O continues as long as at least one image is synchronized*
- *disconnected images can be used by every node in the cluster*
- *mirror operations are done when the system is idle*
- *atomic disconnect for a group of volumes*
- *master and image may be of different types and sizes*

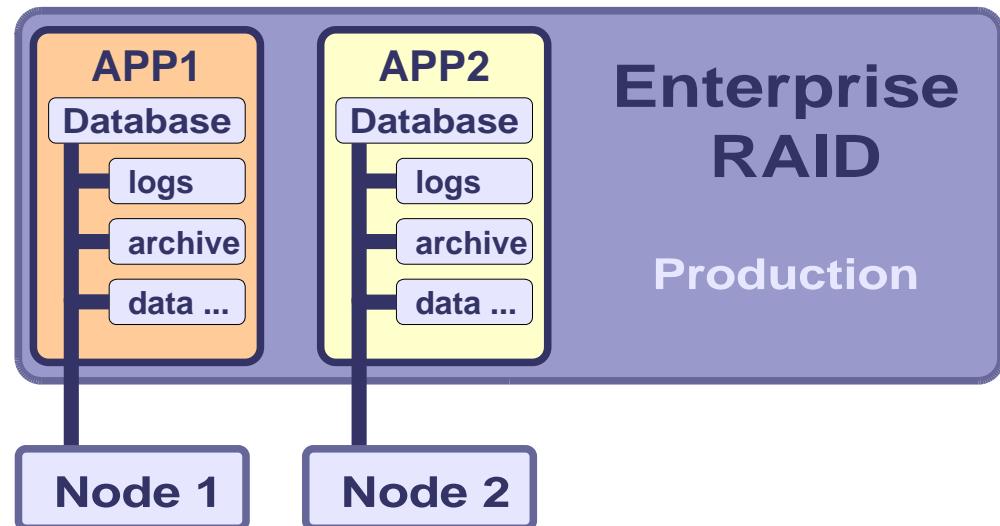
RSIO and OSL Storage Cluster

A perfect couple

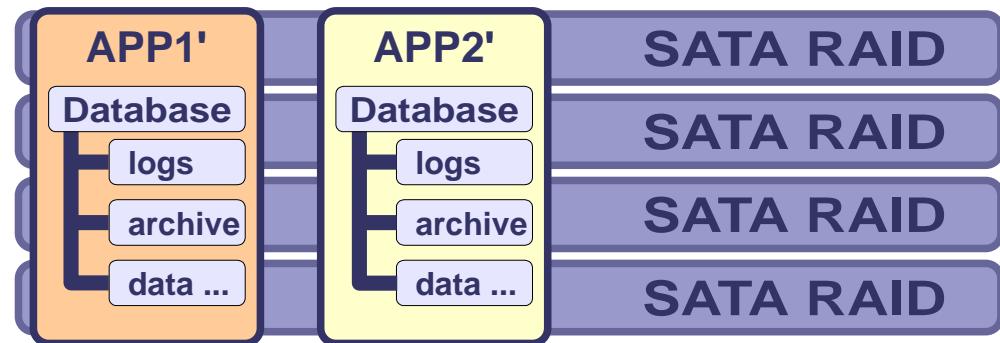


RSIO and OSL Storage Cluster

A perfect couple



***High-speed backup to tape,
without effecting the
productive application***

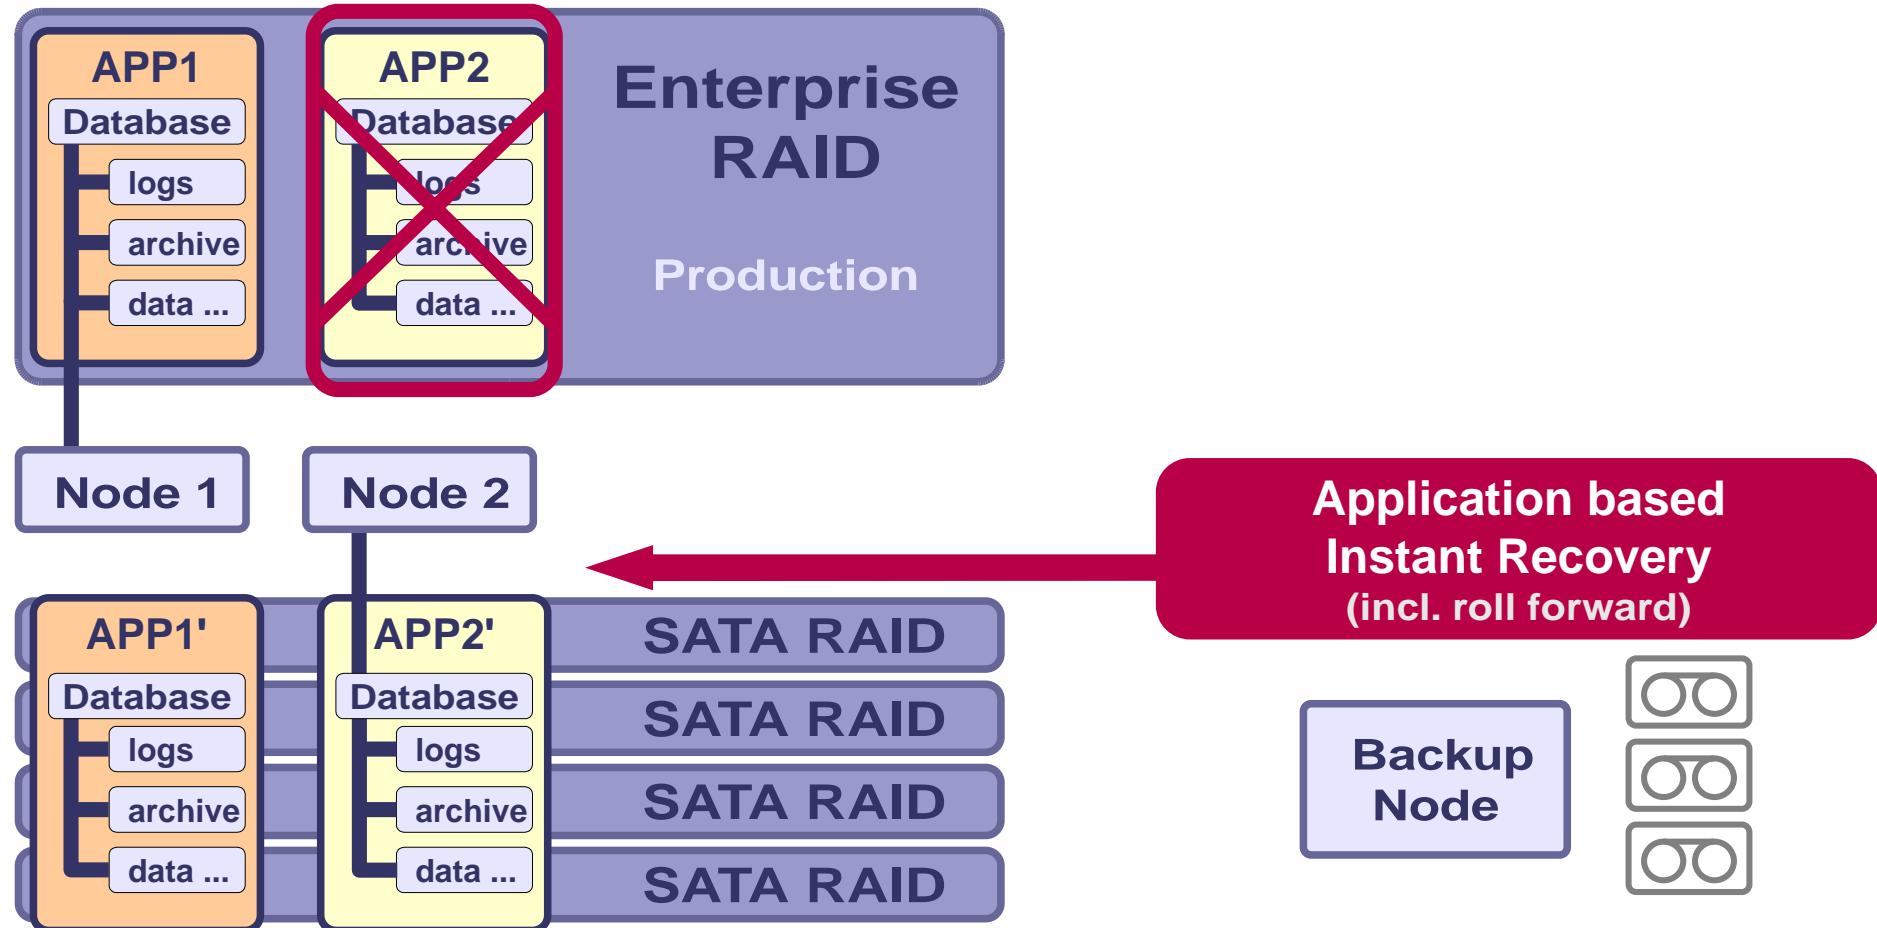


OSL Gesellschaft für offene Systemlösungen mbH

www.osl.eu

RSIO and OSL Storage Cluster

A perfect couple



Instant Recovery with OSL Storage Cluster

- ***What is "Instant Recovery"?***

- *an always restartable backup – a restore is not necessary*
- *it is possible to start a backup application from another storage-universe, without modifying the application*
- *restart the production application, after setting the source to a backup universe*