

# Flexible Laufzeitumgebungen für Software

- Einfach aufgebaut

Detlef Drewanz  
Master Principal Sales Consultant  
Systems Sales Consulting Northern Europe  
17. November, 2015



# 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

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration
- 4 Use Cases

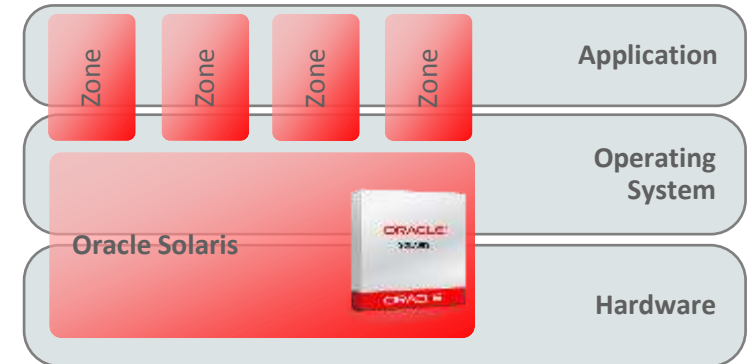
# Runtime Environments for Software (DB and MW)

... so what ?

- Just run my Software
  - Support, Licensing, Performance, ...
- Store Management
- Network Communication
- Enable mobility between physical Systems
- Security Isolation
- What about High Availability ?

# Oracle Solaris Zones (a.k.a. Solaris Containers)

- Available on every system where Oracle Solaris runs
- OS-Virtualization, Isolation and Resource Limitation
- Low Overhead
- Sub-thread Granularity
- Shared Kernel\*
- Instant Restart
- Fast and easy cloning and migration of Zones
- ...



# Evolution of Zones in Solaris

## Solaris 10 (2005)

- First widely adopted OS-level virtualization technology
- Best-of-breed
- Container technology
- Shared kernel

## Solaris 11 (2012)

- Exclusive IP/  
Virtual networking
- Zones on shared storage
- Parallel updates
- Immutable zones
- Mature technology

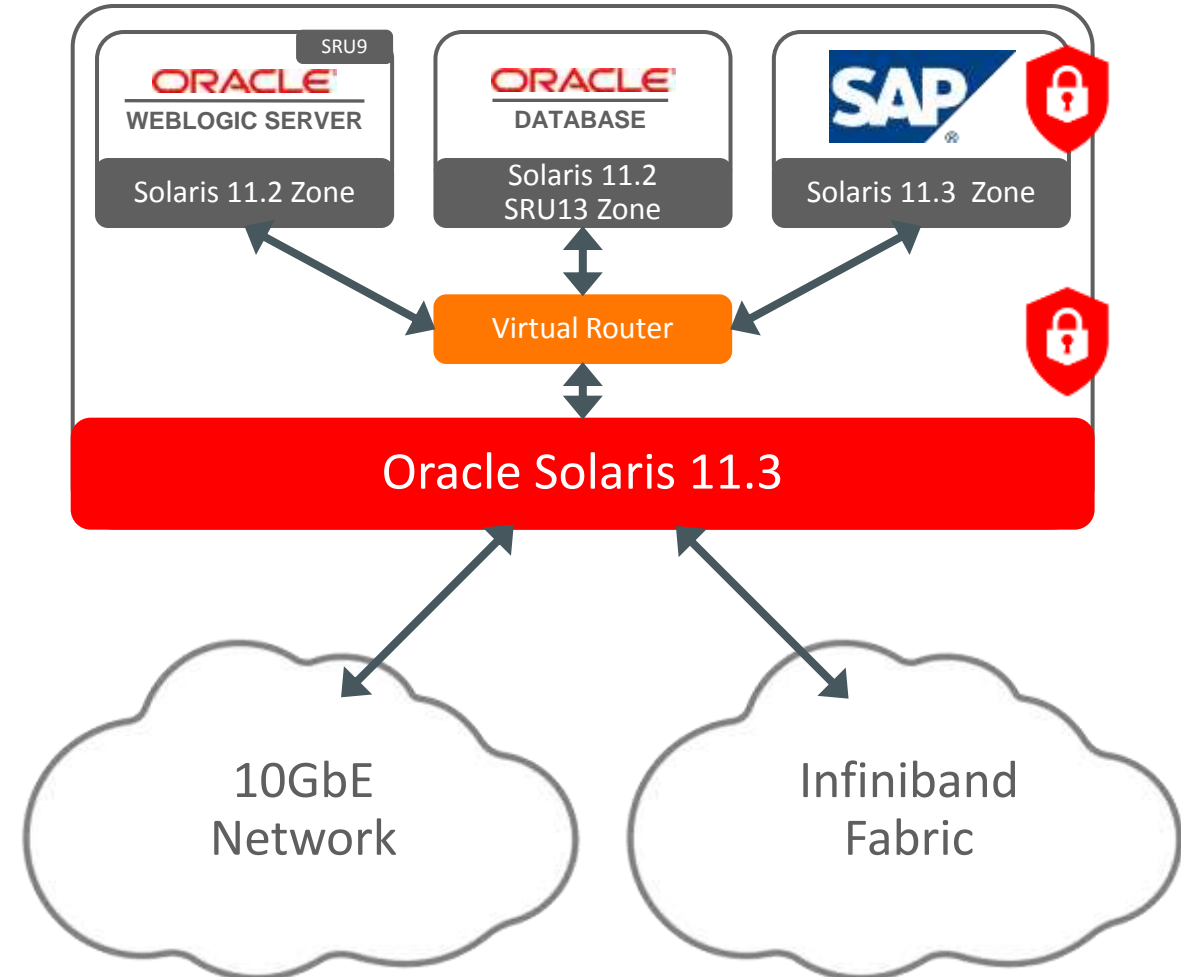
## Solaris 11.2 (2014)

- New type of zone:  
**Kernel Zone**
- Live zone reconfiguration
- CMT awareness in zones
- OpenStack integration

# Solaris Kernel Zones

## OS and Virtualization – Engineered Together

- Next Generation Virtualization
  - Simple administration
  - Leverages Oracle Solaris resource management and network virtualization
  - Seamless P2V and V2P
  - Locked-down root file system for both guest and host
  - Run “any” version, forward and backward compatibility
  - Recognized as a License Boundary



# Kernel Zones Platform Support



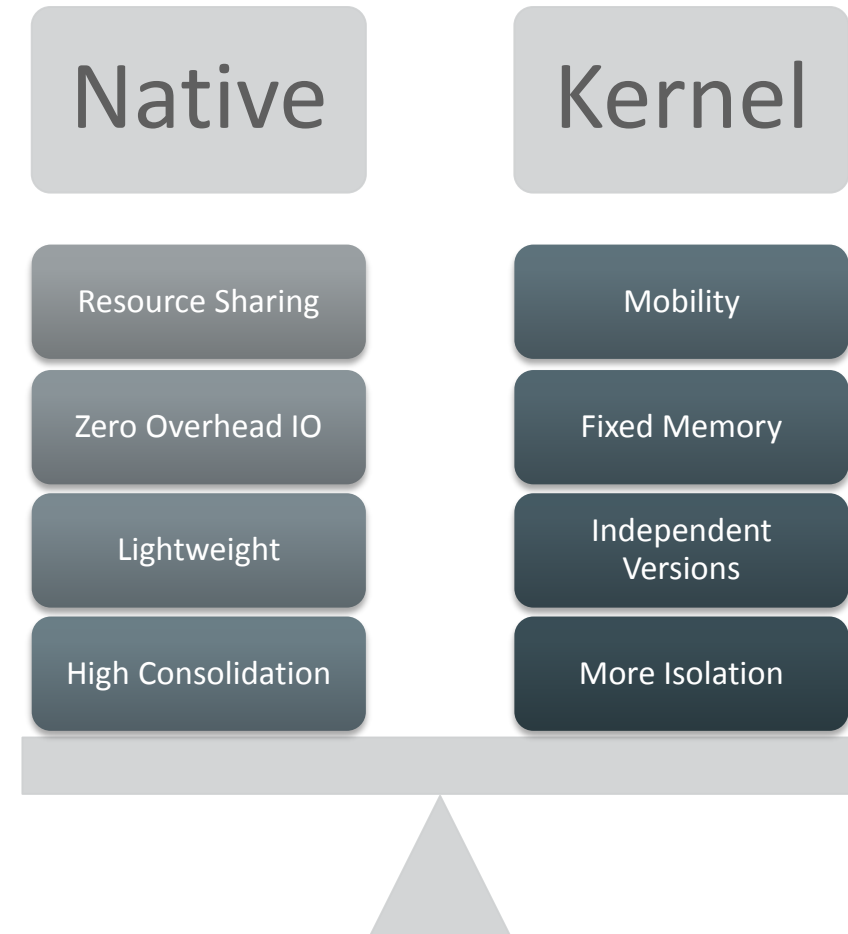
- SPARC: T4 or newer (will need new FW)
- SPARC: M5 or newer (will need new FW)
- SPARC: M10 (Solaris 11.2.4+)
- x86: Intel Nehalem or newer
- x86: AMD Barcelona or newer





# Native or Kernel Zones

- Use application requirements to select
  - **All** applications should go in a zone
  - Most installs will be a combination of Native and Kernel Zones side by side
- Some differences
  - Kernel Independence
  - Memory
  - Live Migration
  - Virtual and Dedicated CPU

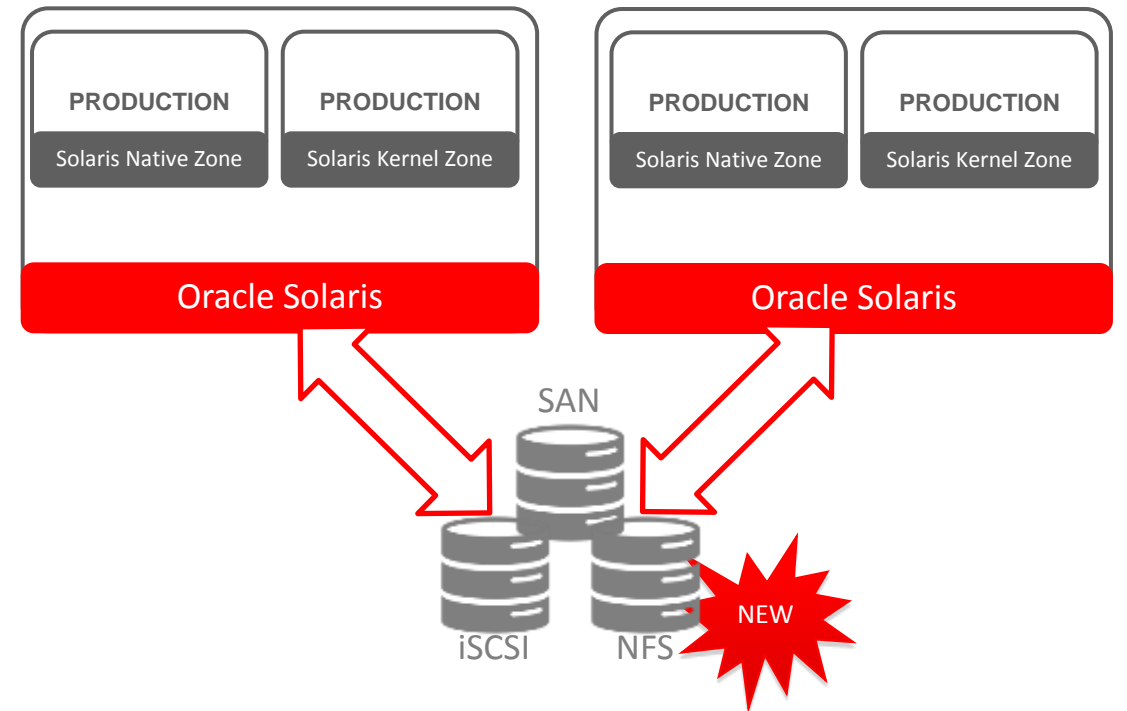


# Program Agenda

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration
- 4 Use Cases

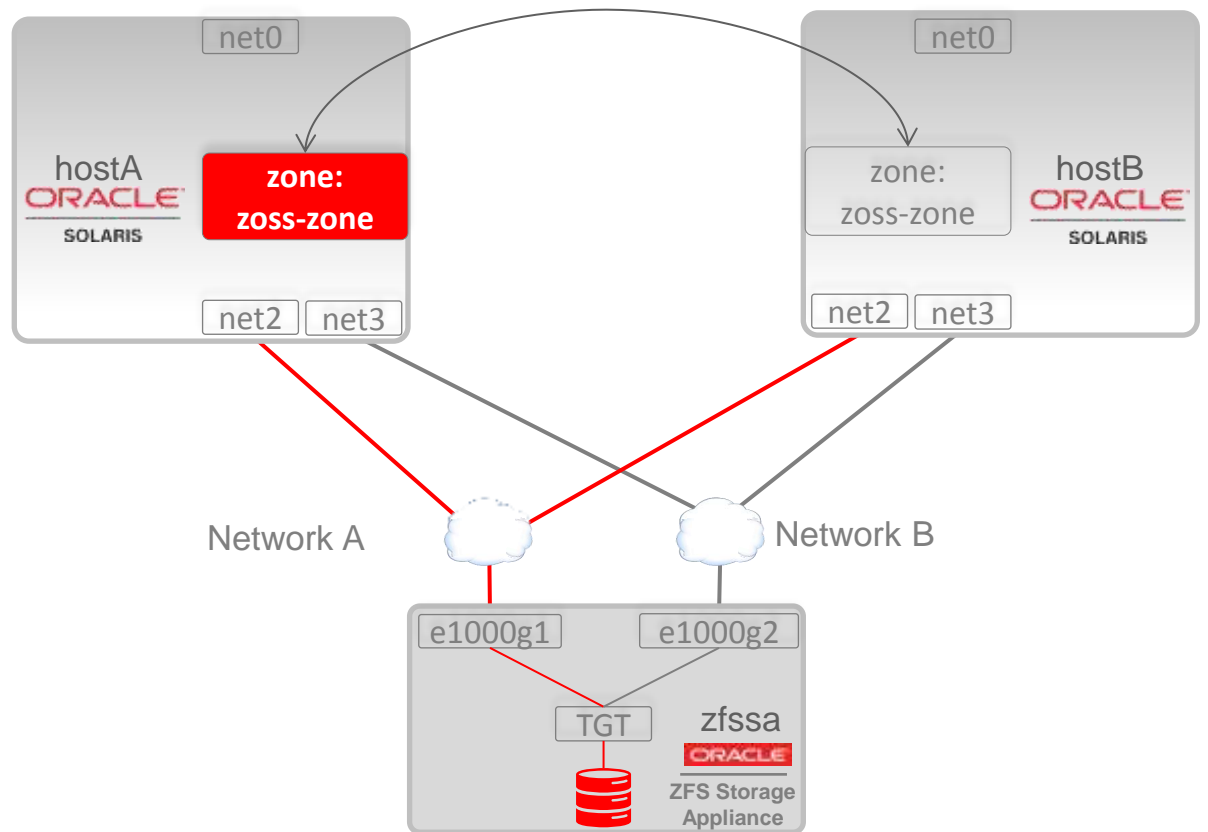
# ZOSS - Zones on Shared Storage

- Dedicated zpools on shared storage
  - Root Pool (rpool)
  - Data
- Managed by Zones framework
- Supported shared storage
  - DAS, iSCSI, NFS, FC, SAS
- Total storage type flexibility
- Avoid configuration errors



# Planned Configuration

- Each Zone has its own zpool
  - <zone>\_rpool
- Consider using dedicated network links for Storage-Traffic
- The Zone is easy to migrate between hostA and hostB



# Program Agenda

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration
- 4 Use Cases

# Short iSCSI Primer

- iSCSI = Block-based Storage over an IP-network
- Initiator
  - The client that initiates SCSI requests to the iSCSI target
- Logical Unit
  - A numbered component in a storage system, identified by Logical Unit-ID
- Target
  - The server, that receives iSCSI requests



# 1. iSCSI ZOSS Setup

## ZFSSA Storage Configuration: Create iSCSI target

The screenshot displays the Oracle ZFS Storage VM configuration interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Configuration' tab is active, and the 'SAN' sub-tab is selected. The 'iSCSI' sub-tab is also highlighted. The main content area shows the 'Storage Area Network (SAN)' configuration, with options for 'Fibre Channel', 'iSCSI', and 'SRP'. The 'iSCSI' option is selected. Below this, there are sections for 'Targets' and 'Initiators'. The 'Targets' section indicates that no iSCSI targets are defined and provides a button to create one. The 'Initiators' section shows a table of target groups:

NAME	TARGETS
default	[ ALL TARGETS ]

An 'Edit iSCSI Target' dialog box is open, showing the following configuration:

- Target IQN: iqn.1986-03.com.sun:02:12c5dc60-1854-62c-aa78d0815dc2
- Alias: zoss-zone
- Initiator authentication mode:  None,  CHAP,  RADIUS
- Target CHAP name: [Empty field]
- Target CHAP secret: [Empty field]
- Network interfaces: e1000g0 (selected), e1000g1, e1000g2

## 2. iSCSI ZOSS Setup

### ZFSSA Storage Configuration: Create and export Logical Unit

The screenshot displays the Oracle ZFS Storage VM management interface. The top navigation bar includes 'Configuration', 'Maintenance', and 'Shares' (highlighted with a red box). The 'Shares' page shows a 'Projects' dropdown set to 'All Projects' and a 'LUNs' section with '0 Total' (highlighted with a red box). A 'Create LUN' dialog box is open, showing the following configuration:

- Project: default
- Name: zoss-zone-rpool-primary
- Volume size: 10 G
- Thin provisioned:
- Volume block size: 8k
- Online:
- Target group: All targets
- Initiator group(s): All initiators
- LU number:  0  Auto-assign

The 'Shares' configuration page for the 'default' project shows a usage of 13.6% of 73.4G. The 'LUNs' section shows '1 Total' LUN. The following table lists the LUN details:

NAME	SIZE	GUID
zoss-zone-rpool-primary	10G	600144F0949056290000529625DD0001



### 3. iSCSI ZOSS Setup for Solaris Zones



The screenshot shows the Oracle ZFS Storage VM web interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' tab is active. Below the navigation, there are tabs for 'Projects', 'default', 'Shares', 'General', 'Protocols', 'Access', 'Snapshots', and 'Replication'. The 'default' project is selected, showing a usage bar for 'zfssa/local/default' at 13.6% of 73.4G. A table of LUNs is displayed with one entry: 'zoss-zone-rpool-primary' with a size of 10G and a GUID of '600144F0949056290000529625DD0001'. The GUID is highlighted with a red box.

NAME	SIZE	GUID
zoss-zone-rpool-primary	10G	600144F0949056290000529625DD0001

```
root@hostA:/# zonecfg -z zoss-zone create -t SYSsolaris
```

```
...
```

```
zonecfg:zoss-zone> add rootzpool
```

```
zonecfg:zoss-zone:rootzpool> add storage iscsi://zfssa/luname.naa.600144F0949056290000529625DD0001
```

```
zonecfg:zoss-zone:rootzpool> end
```

```
...
```

### 3. iSCSI ZOSS Setup for Solaris Kernel Zones



The screenshot shows the Oracle ZFS Storage VM web interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' tab is active. Below the navigation, there are tabs for 'Projects', 'default', 'Shares', 'General', 'Protocols', 'Access', 'Snapshots', and 'Replication'. The 'default' project is selected, showing a usage bar at 13.6% of 73.4G. A table lists the LUNs, with one LUN named 'zoss-zone-rpool-primary' having a size of 10G and a GUID of '600144F0949056290000529625DD0001', which is highlighted with a red box.

NAME	SIZE	GUID
zoss-zone-rpool-primary	10G	600144F0949056290000529625DD0001

```
root@hostA:/# zonecfg -z kzone create -t SYSsolaris-kz
```

```
...
```

```
zonecfg:kzone>zonecfg:kzone3> select device 0
```

```
zonecfg:kzone:device> set storage=iscsi://zfssa/luname.naa.600144F0949056290000529625DD0001
```

```
zonecfg:kzone:device> set bootpri=0
```

```
...
```

# 4. ZOSS Setup

## Oracle Solaris Zone Installation

- Installation includes:
  - iSCSI initiator configuration
  - iSCSI target and Logical Unit discovery
  - Creation of logical device links
  - Zone rpool creation
  - Zone installation

```
root@hostA:~# zoneadm -z zoss-zone install
Configured zone storage resource(s) from:
    iscsi://zfssa/luname.naa.600144F0949056290000529625DD0001
Created zone zpool: zoss-zone_rpool
```

# 5: Option to add: Control Logical Unit Discovery

## Create and Assign Initiator Group

```
root@hostA:~# iscsiadm list initiator-node
```

```
Initiator node name: iqn.1986-03.com.sun:01:e00000000000.5295332b
```

The screenshot displays the Oracle ZFS Storage VM management interface. It features a top navigation bar with tabs for Configuration, Maintenance, Shares, Status, and Analytics. Below this, there are sub-tabs for SERVICES, STORAGE, NETWORK, SAN, CLUSTER, USERS, PREFERENCES, and ALERTS. The main content area is divided into several sections:

- Storage Area Network (SAN):** Includes a description and buttons for REVERT and APPLY.
- Targets and Initiators:** A table with two columns: Targets and Initiators. The Initiators column contains two entries: hostA (iqn.1986-03.com.sun:01:e00000000000.5295332b) and hostB (iqn.1986-03.com.sun:01:e00000000000.5295959e).
- Initiator Groups:** A table with two columns: NAME and INITIATORS. The NAME column contains 'default' and 'zoss-zone-h...'. The INITIATORS column for 'zoss-zone-h...' contains the same two initiator names as above.
- Edit iSCSI Initiator Dialog:** A modal window titled 'Edit iSCSI Initiator' with 'CANCEL' and 'OK' buttons. It shows 'Initiator IQN' as 'iqn.1986-03.com.sun:01:e00000000000.529' and 'Alias' as 'hostA'.
- Projects Section:** Shows a project named 'zoss-zone-rp...' with usage information (12.1% of 82.5G) and sharing options. The 'Initiator group' is set to 'LU number' with a value of 'zoss-zone-hosts:0'.

# Program Agenda

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration
- 4 Use Cases

# 1. NFS ZOSS Setup

## ZFSSA Storage Configuration: Create NFS Share

The screenshot shows the Oracle ZFS Storage VM interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares' (highlighted in red), 'Status', and 'Analytics'. The 'Shares' tab is active, showing a table of filesystems. A 'Create Filesystem' dialog box is open, with several fields highlighted in red: 'Name' (kzone3), 'User' (100), 'Group' (10), and 'Mountpoint' (/export/kzone3).

**Usage** 3.7% of 71.0G

Referenced data 2.64G

Total space 2.64G

**Filesystems** LUNs: 1 Total

SHOW ALL LOCAL REPLICA

NAME	SIZE	MOUNTPOINT	ENCRYPTED
default /kzone2	2.64G	/export/kzone2	

**Create Filesystem** CANCEL APPLY

Project: default

Name: kzone3

Data migration source: None

User: 100

Group: 10

Permissions:  R W X  R W X  R W X  
User Group Other

Use Windows default permissions

Inherit mountpoint:

Mountpoint: /export/kzone3

Reject non UTF-8:

Case sensitivity: Mixed

Normalization: None

Encryption: Off

Inherit key:

Key: Local OKM

### 3. NFS ZOSS Setup for Solaris Kernel Zones

```
root@hostA:/# zonecfg -z kzone3 create -t SYSsolaris-kz
```

```
...
```

```
zonecfg:kzone>zonecfg:kzone3> select device 0
```

```
zonecfg:kzone3> select device 0
```

```
zonecfg:kzone3:device> set storage=nfs://100:10@zfssa/export/kzone3/rpool.img
```

```
zonecfg:kzone3:device> set create-size=16g
```

```
zonecfg:kzone:device> set bootpri=0
```

```
...
```

```
root@hostA:/# zoneadm -z kzone2 install -x storage-create-missing
```

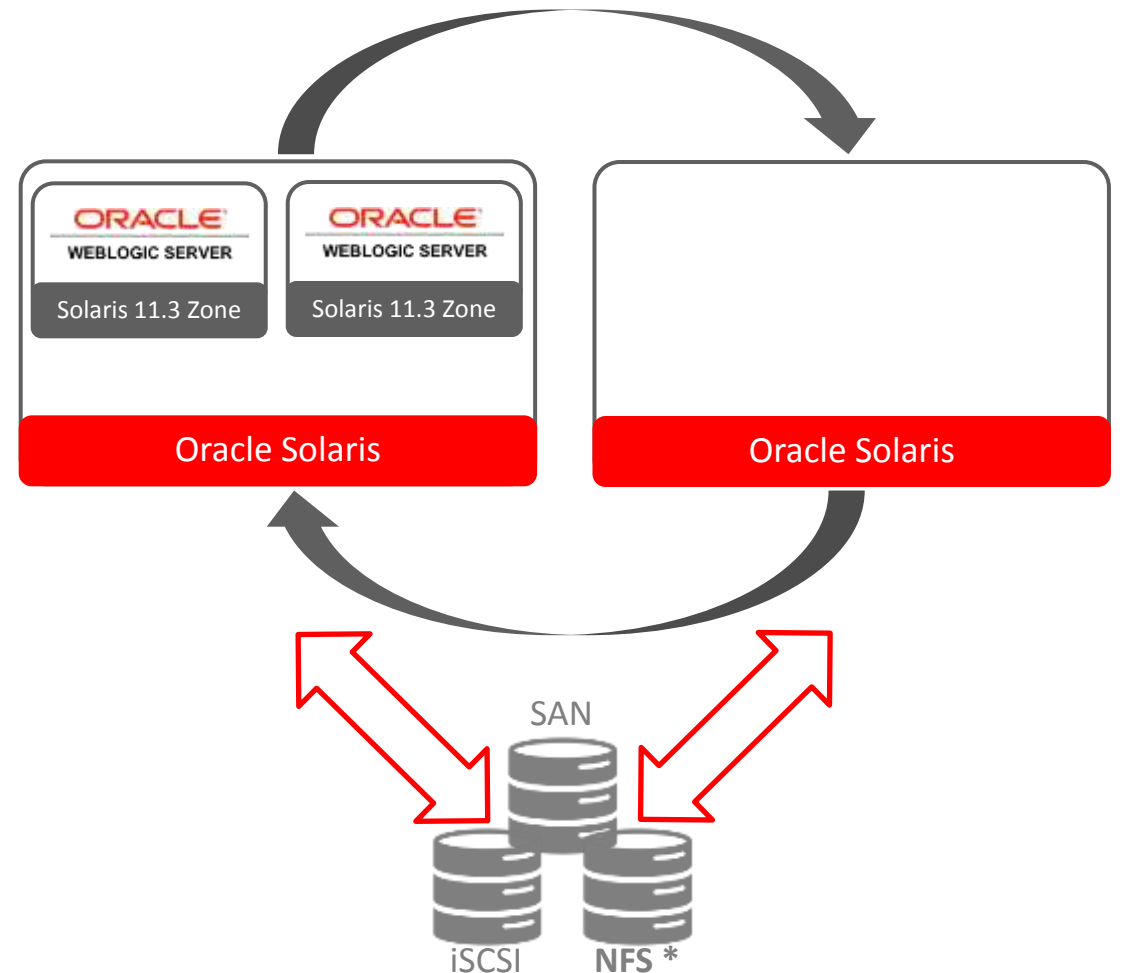
# Program Agenda

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration**
- 4 Use Cases



# Live Migration with Kernel Zones

- Move Kernel Zones without outage
  - Live migrate on SPARC and x86
  - No outage host maintenance
  - Oracle Solaris 11.3 or later required
- Migration Types
  - **Cold:** Resource intensive applications with no local state (Analytical)
  - **Warm:** Resource intensive applications with large local state (Database)
  - **Live:** Applications that require moves without outages (Web services, etc)



# Kernel Zones Cold Migration

```
root@hostA # zonecfg -z kzone3 export | ssh root@hostB zonecfg -z kzone3 -f -
```

```
root@hostA # zoneadm -z kzone3 shutdown
```

```
root@hostA # zoneadm -z kzone3 detach
```

```
root@hostB # zoneadm -z kzone3 attach
```

```
root@hostB # zoneadm -z kzone3 boot
```

# Kernel Zones Warm Migration

```
root@hostA # root@hostA # zonecfg -z kzone3 export | ssh root@hostB zonecfg -z kzone3 -f -
```

```
root@hostA # zonecfg -z kzone3 info suspend  
suspend:
```

```
storage: iscsi://zfssa/luname.naa.600144F0949056290000529626DD0001
```

```
root@hostA # zoneadm -z kzone3 suspend
```

```
root@hostA # zoneadm -z kzone3 detach
```

```
root@hostB # zoneadm -z kzone3 attach
```

```
root@hostB # zoneadm -z kzone3 resume
```

# Kernel Zones Live Migration

```
# svcadm enable -s svc:/system/rad:remote svc:/network/kz-migr:stream svc:/network/ntp:default
```

```
root@hostA # zonecfg -z kzone3 export | ssh root@hostB zonecfg -z kzone3 -f -
```

```
root@hostA # zoneadm -z kzone3 migrate hostB
```

Password:

```
zoneadm: zone 'kzone3': Importing zone configuration.
```

```
zoneadm: zone 'kzone3': Attaching zone.
```

```
zoneadm: zone 'kzone3': Booting zone in 'migrating-in' mode.
```

```
zoneadm: zone 'kzone3': Checking migration compatibility.
```

```
zoneadm: zone 'kzone3': Starting migration.
```

```
zoneadm: zone 'kzone3': Suspending zone on source host.
```

```
zoneadm: zone 'kzone3': Waiting for migration to complete.
```

```
zoneadm: zone 'kzone3': Halting and detaching zone on source host.
```

```
zoneadm: zone 'kzone3': Migration successful.
```


# Program Agenda

- 1 Solaris Zones and Kernel Zones
- 2 Zones on Shared Storage (iSCSI and NFS)
- 3 Kernel Zones and Live Migration
- 4 Use Cases

# Oracle Solaris Cluster Virtualization


HA for **your** business, operation and migration model

- **Choice** of VM or application centric model
- **Choice** of technology: OVM domains or zones
- **Built-in** asset optimization with load balancing, affinity and dependency management at application or VM level
- Oracle Solaris 11 and 10 co-existence



**Application Failover**  
With zone clusters or OVM server for SPARC domain clusters

**Workload Failover**  
With failover zones or failover guest domains

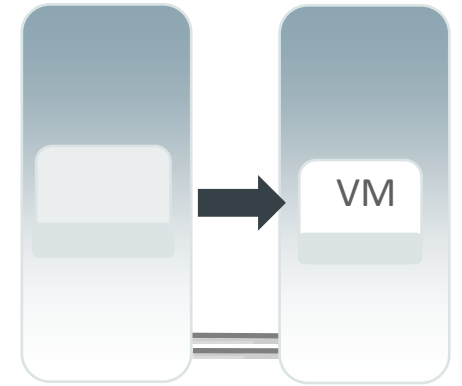


# Efficient Virtualization, Reduced Downtime


## HA for workloads with Zones and Kernel Zones

- Managed zone switchover with cold / warm / live migration (kernel zone)
- Automatic zone restart or zone failover upon node failure
  - **2x** faster recovery
- No modification of workload
- Support for Zones over shared storage with FC and iSCSI
- Dependencies and load management at zone level

**Planned Maintenance Downtime:**  
Workload migration



**Unplanned Downtime:**  
Immediate workload restart or failover

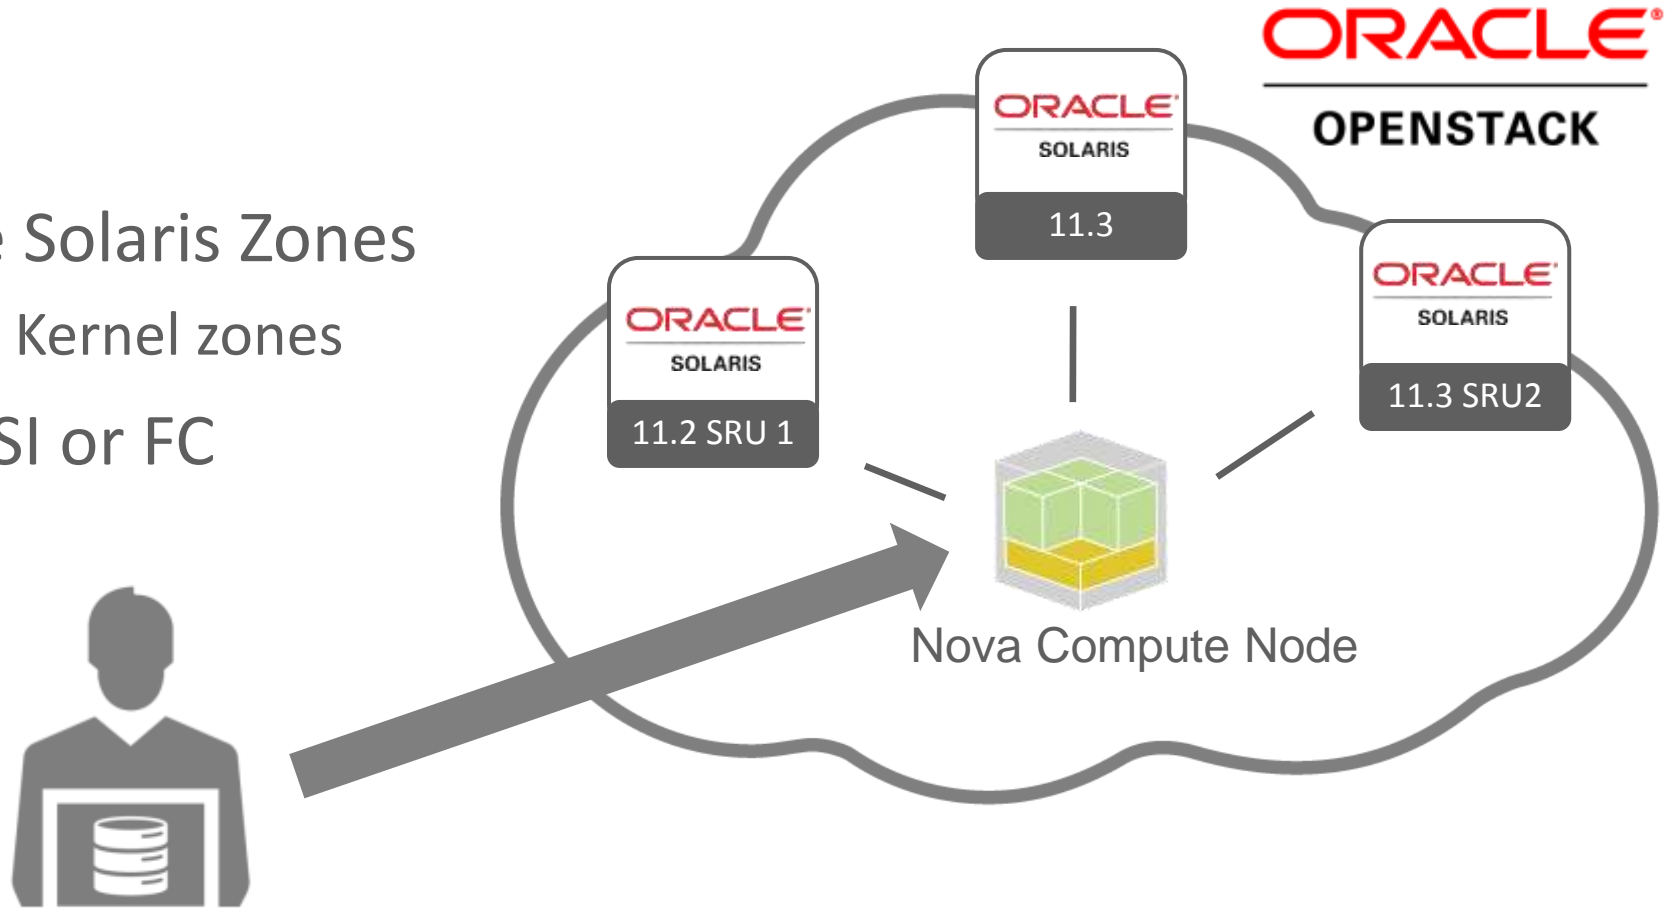


**Maximize consolidation with Zones, Minimize risk with Oracle Solaris Cluster**

# Self Service Cloud with OpenStack Nova Compute

High density virtual environments – ideal for multi-tenant cloud

- Private Cloud
  - Managed by OpenStack
- Integrated with Oracle Solaris Zones
  - Native non-global zone, Kernel zones
- Shared Storage on iSCSI or FC





# Solaris Zones Use Cases

- Consolidation
  - Optimize workload and licences
  - Consolidate small, test, development, learning systems
- Isolation
  - Multiple instances of the same service on a system
  - Development/Test/Quality Ensurance/Production on one system
- Create new architectures
  - Flexible service operation (instant restart)
  - Encapsulate applications
  - Operate legacy applications on new hardware



Q & A

Detlef.Drewanz@oracle.com

# Integrated Cloud

## Applications & Platform Services

ORACLE®