portrix systems

# Utilizing Oracle Solaris Containers with Oracle Database

Björn Rost

ORACLE Gold Partner

ORACLE 10g Certified Professional

portrix systems

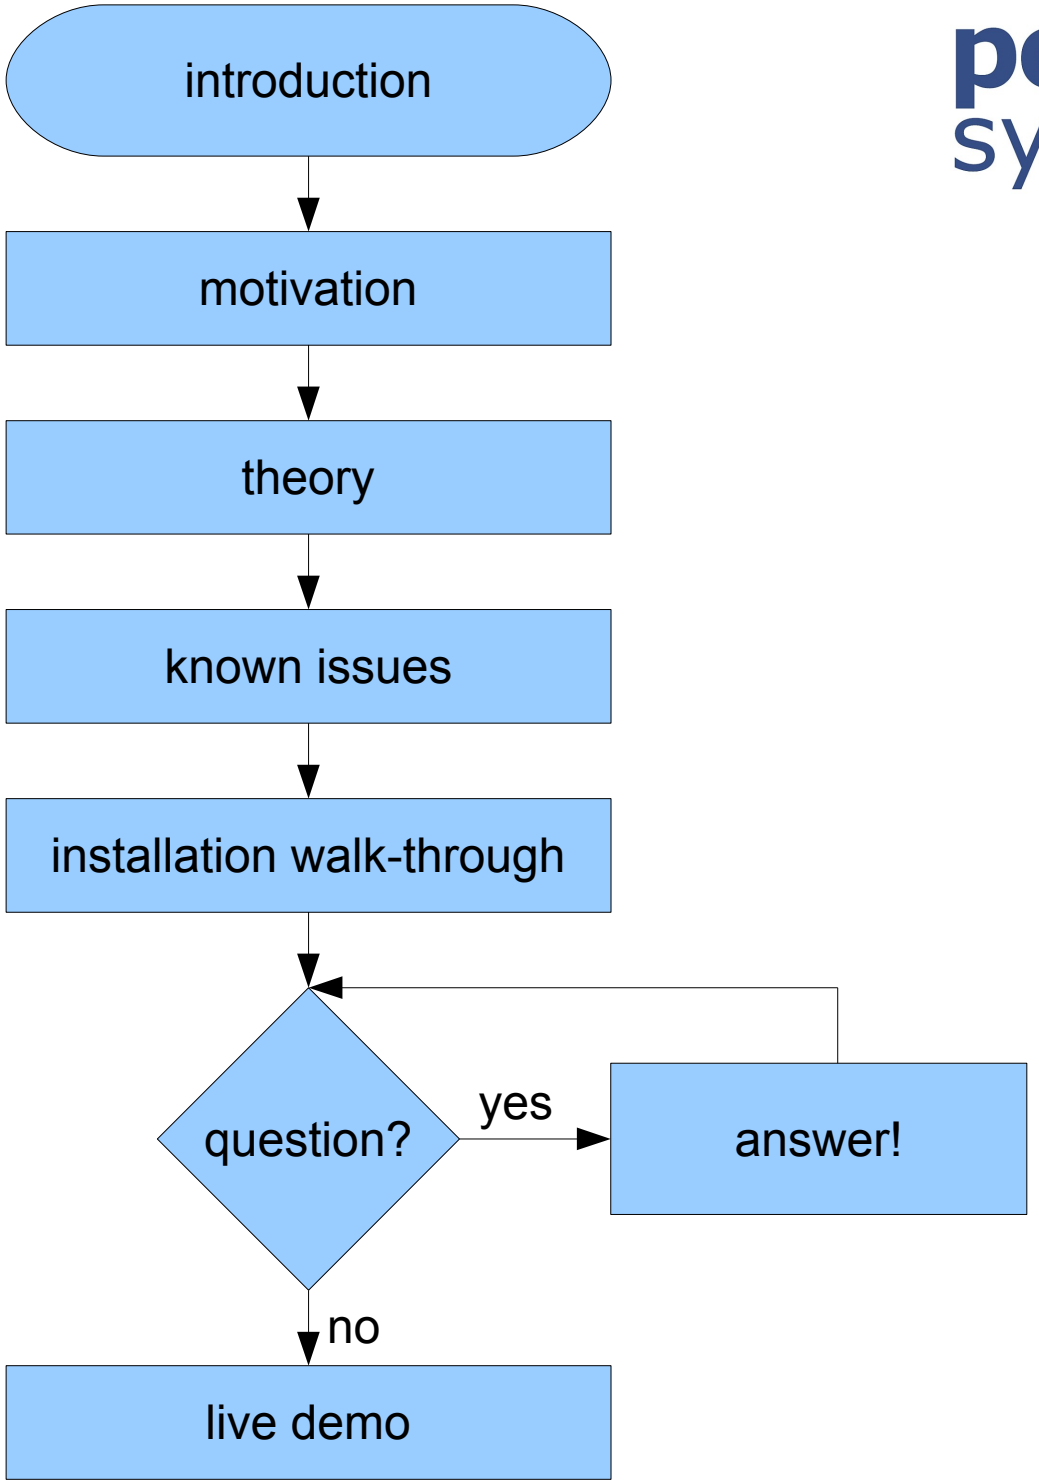

# about portrix

portrix systems

- Software Production company founded 2001
  - mostly J2EE
    - logistics
    - telco
    - media and publishing
- customers demand full lifecycle support
  - hardware resale
  - datacenter operations
  - 3rd party software

Adobe Solution PARTNER

CISCO PARTNER Select Certified

ORACLE Gold Partner

Sun microsystems Partner Advantage ASSOCIATE PARTNER

RIPE NCC Member

portrix systems

# project lifecycle

- consulting
- specification
- documentation
- feasibility studies

- J2EE
- php
- database

- Hardware
- SW-Licenses
- Installation
- Benchmarking
- Tuning

- Monitoring
- Support
- Updates

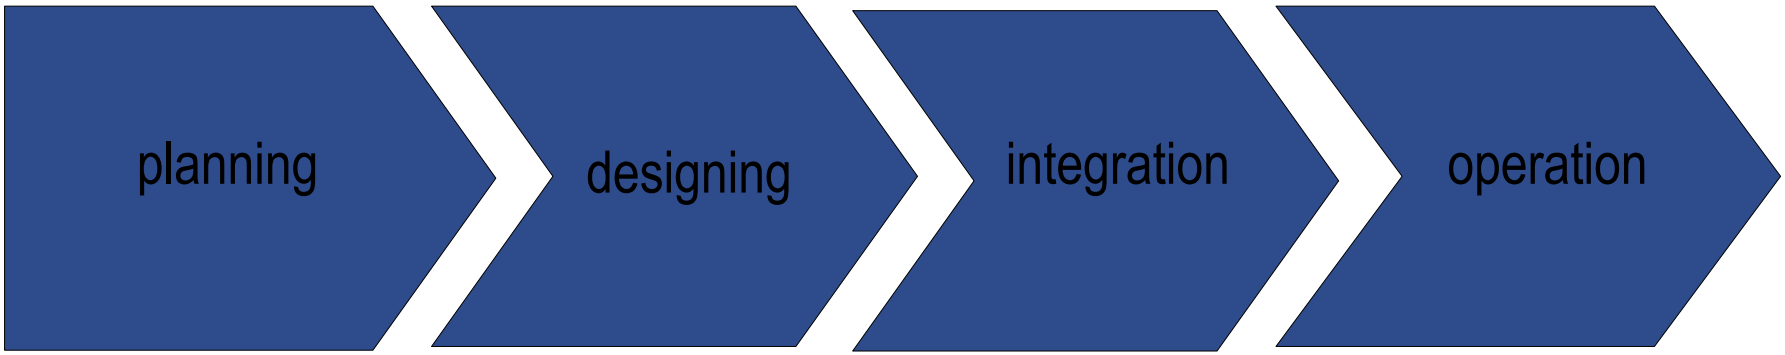


blau
Mobilfunk GmbH


MEEDIA


portrix
systems


ANRUFDIREKT
predictive dialing by call


EDEKA


VIRTUALREEL


e+p


SevenSeas International Publishing Agency


LichtBlick
die Zukunft der Energie


apo-rot
Ihre VersandApotheke


Klett


EatSmarter!


Kolle Rebbe


LFI
LEICA FOTOGRAFIE INTERNATIONAL


Commercial
Film
Service


JUNG v. MATT

portrix systems

# 7-seas

- global publishing system
  - document management
  - translation
  - timing control

SevenSeas International Publishing Agency

censhare
GLOBAL PUBLISHING

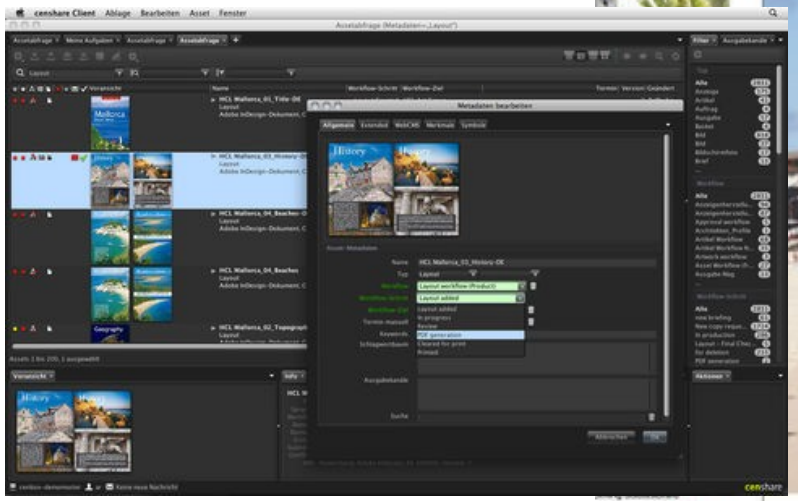

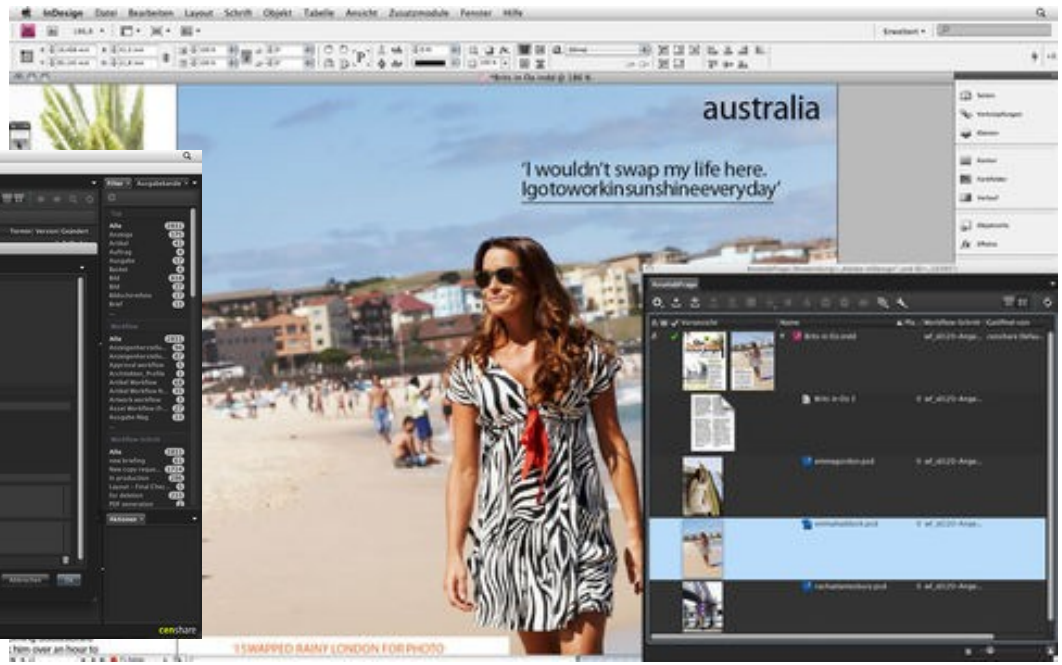

portrix systems

# 7-seas

- client/server based system
  - JAVA application Server
  - Oracle Database
  - Apache Tomcat Webserver

SevenSeas International Publishing Agency

censhare GLOBAL PUBLISHING

Java

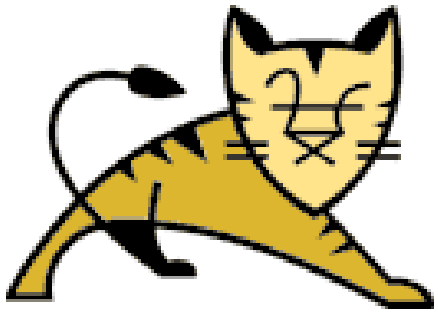

ORACLE DATABASE 11g

# the challenge

portrix systems

- deploy and run this multi-tiered system
  - cost-effectively
    - 3rd party licenses
    - hardware
    - operating costs
  - but stay flexible to scale with demand
  - and have a plan for recovery

# Why virtualize?

portrix systems

- Consolidation saves resources, money
- Flexible deployment and scaling
- Oracle hard partitioning saves licensing
- availability (better to have one SPOF than many)
- poor-man's flashback
- poor-man's compression
- cheap and easy cloning

portrix systems

# Consolidate

- many systems are running at low utilization
  - you might not even know the impact on subsystems at launch
- systems get more and faster cores all the time
  - but not all workflows do, too
- energy costs increasing

# Flexible Deployment

portrix systems

- deploy a system once
- in a virtual environment
- then, move it across servers to scale
- start smart, scale hard
  - initial deployment on single box with multiple VMs
  - scale out by distributing VMs across nodes when needed
- easy cloning of environments
- in case of HW failure restore VM to another box

# Oracle License Partitioning

portrix systems

- Only pay for the CPUs/Cores you actually use
- requires 'hard' partitioning
- grow/scale easily – buy new license and enable cores

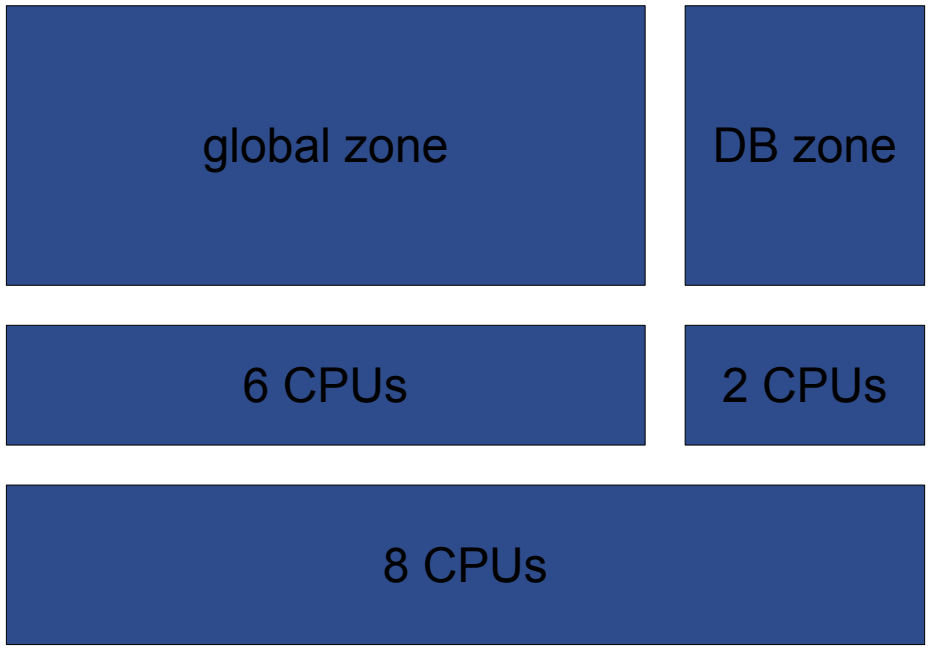


http://www.oracle.com/corporate/pricing/partitioning.pdf

# Solutions

portrix systems

- VMware
  - not even certified/supported according to metalink
- LDom, vPar, nPar, LPAR, DSD
  - vendor- and hardware specific / non-x86
- Oracle VM
  - XEN hypervisor, typically used with OEL as guest OS
  - great for large-scale deployments
- Oracle Solaris Containers
  - native to Solaris, no extra software needed

# Containers

portrix systems

- Virtualization built into kernel – no installation, no costs
  - Resource Management
  - provides namespace, security and fault isolation
- runs on Solaris x86/x64 and SPARC
- no/minumum overhead
  - no hypervisor neccessary
  - no special hardware needed
- mobility: detach/attach and clone
- Oracle supported and hard partitioning possible

# Oracle Solaris Zones

portrix systems

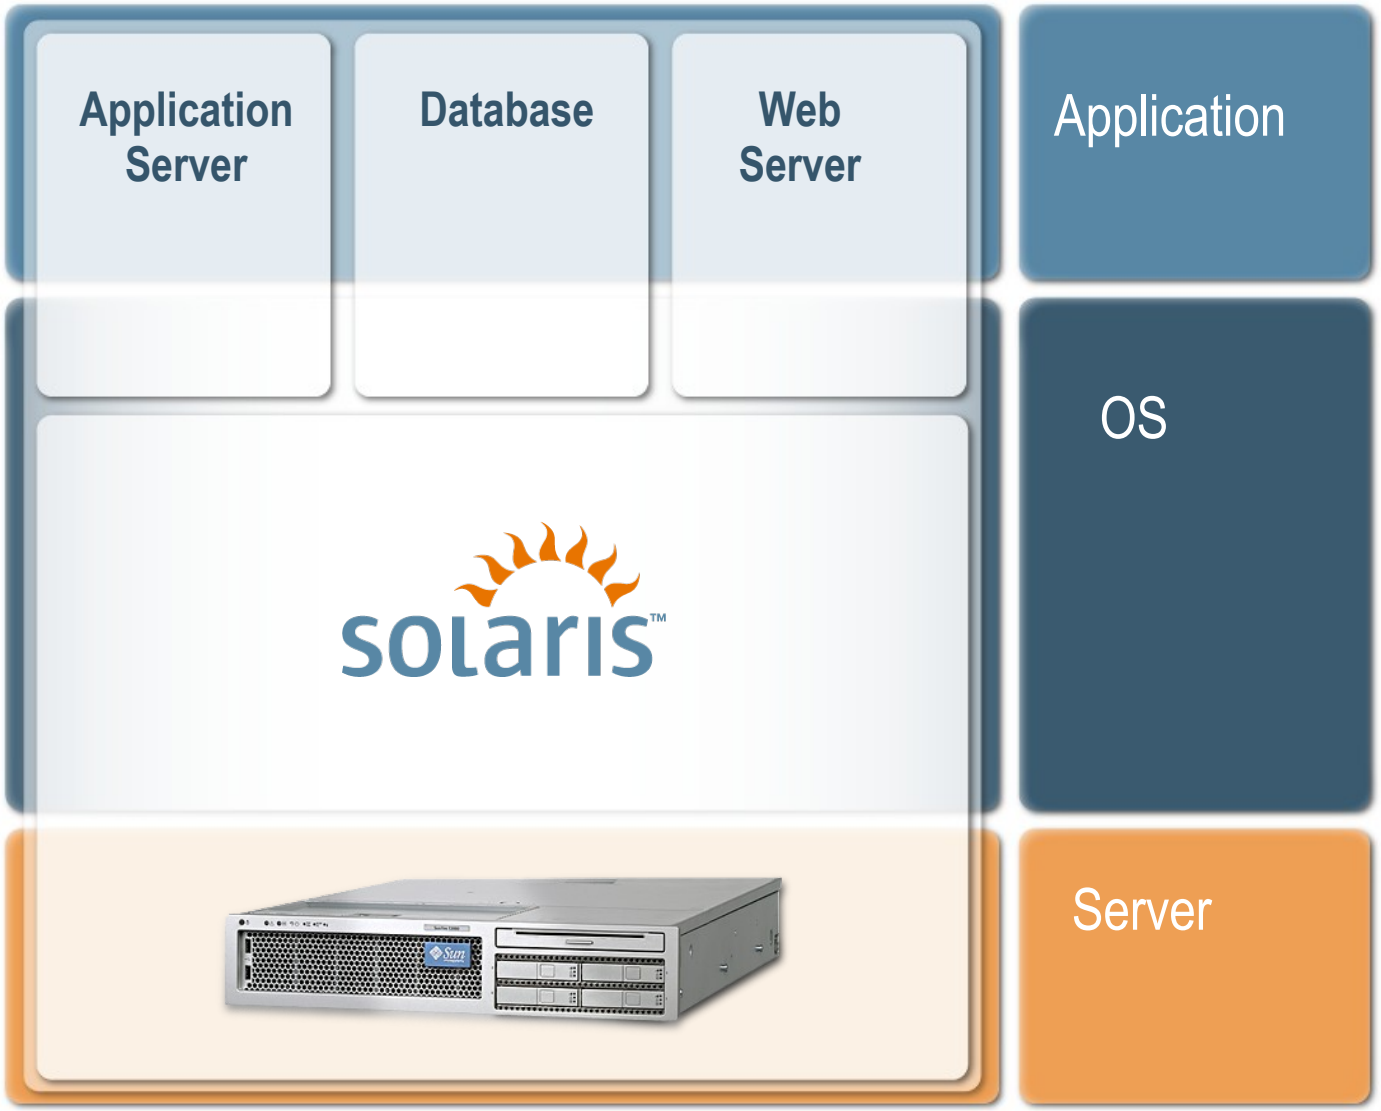

# container config

portrix systems

- global zone – this is the default/main OS instance
- local zone – these are the virtualized environments
  - own IP (and possibly interface)
  - sparse root zone
    - parts of directory structure (/usr, /lib …) inherited from global zone
    - read-only in zone, saves storage space
  - full root zone
    - full directory structure is copied to zone
    - better control over patches in zone

# advanced config

portrix systems

- shared/private IP
- mapping devices / filesystems
- resource capping – container
- migration between servers
- branded zones (linux or solaris 9)

# step-by-step installation

portrix systems

- install latest version of Solaris (or OpenSolaris)
  - zones, resource management and zfs already included
- create zfs and zone(s)
- install oracle into zone
- PLEASE TRY THIS AT HOME!

portrix systems

```
root@hermes:~# zonecfg -z ora11g
ora11g: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:ora11g> create
zonecfg:ora11g> set zonepath=/zp03/zones/ora11g
zonecfg:ora11g> set autoboot=true
zonecfg:ora11g> add capped-cpu
zonecfg:ora11g:capped-cpu> set ncpus=1
zonecfg:ora11g:capped-cpu> end
zonecfg:ora11g> add net
zonecfg:ora11g:net> set address=192.168.42.79
zonecfg:ora11g:net> set physical=e1000g0
zonecfg:ora11g:net> end
zonecfg:ora11g>
```

```
zonecfg:ora11g> info
zonename: ora11g
zonepath: /zp03/zones/ora11g
brand: native
autoboot: true
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
inherit-pkg-dir:
        dir: /lib
inherit-pkg-dir:
        dir: /platform
inherit-pkg-dir:
        dir: /sbin
inherit-pkg-dir:
        dir: /usr
net:
        address: 192.168.42.79
        physical: e1000g0
        defrouter not specified
capped-cpu:
        [ncpus: 1.00]
rctl:
        name: zone.cpu-cap
        value: (priv=privileged,limit=100,action=deny)
zonecfg:ora11g> verify
zonecfg:ora11g> commit
zonecfg:ora11g> exit
```

portrix systems

portrix systems

```
root@hermes:~# zoneadm -z ora11g verify
WARNING: /zp03/zones/ora11g does not exist, so it could
not be verified.
When 'zoneadm install' is run, 'install' will try to
create
/zp03/zones/ora11g, and 'verify' will be tried again,
but the 'verify' may fail if:
the parent directory of /zp03/zones/ora11g is group- or
other-writable
or
/zp03/zones/ora11g overlaps with any other installed
zones.
```

portrix systems

```
root@hermes:~# zoneadm -z ora11g install
A ZFS file system has been created for this zone.
Preparing to install zone <ora11g>.
Creating list of files to copy from the global zone.
Copying <147436> files to the zone.
Initializing zone product registry.
Determining zone package initialization order.
Preparing to initialize <1504> packages on the zone.
Initialized <1504> packages on zone.
Zone <ora11g> is initialized.
The file
</zp03/zones/ora11g/root/var/sadm/system/logs/install_log>
contains a log of the zone installation.
```

portrix systems

```
root@hermes:~# zoneadm -z ora11g boot
root@hermes:~# zlogin -C ora11g
```

now, you get the 'normal' postinstall questions.
language, hostname, root-pw, timezone ...

# ZFS

portrix systems

- ships for free with solaris – no installation needed
- new and innovative FS, Volume Manager
  - snapshots/clones
  - compression
  - quotas
  - raid
  - bootable 

  - share with NFS or iSCSI
- enhances manageability in combination with zones

# Oracle datafiles on ZFS

portrix systems

- might not perform as well as ASM
- but taking snapshots of the whole database is amazing
  - fast and deduplicated DB cloning
  - poor-man's flashback database
  - poor-man's compression
  - backup/recovery

# demo

portrix systems

- cloning
- poor man's flashback

# licenses needed

portrix systems

- Solaris (already included in OPS with SUN hardware)
  - enterprise OS
  - zfs file system and volume manager
  - container virtualization
  - dTrace profiler
- Oracle SE or EE
  - only for sockets or cores actually used

# known issues

- performance management
- limit netservices
- kernel parameters
- ro /usr
- no RAC?

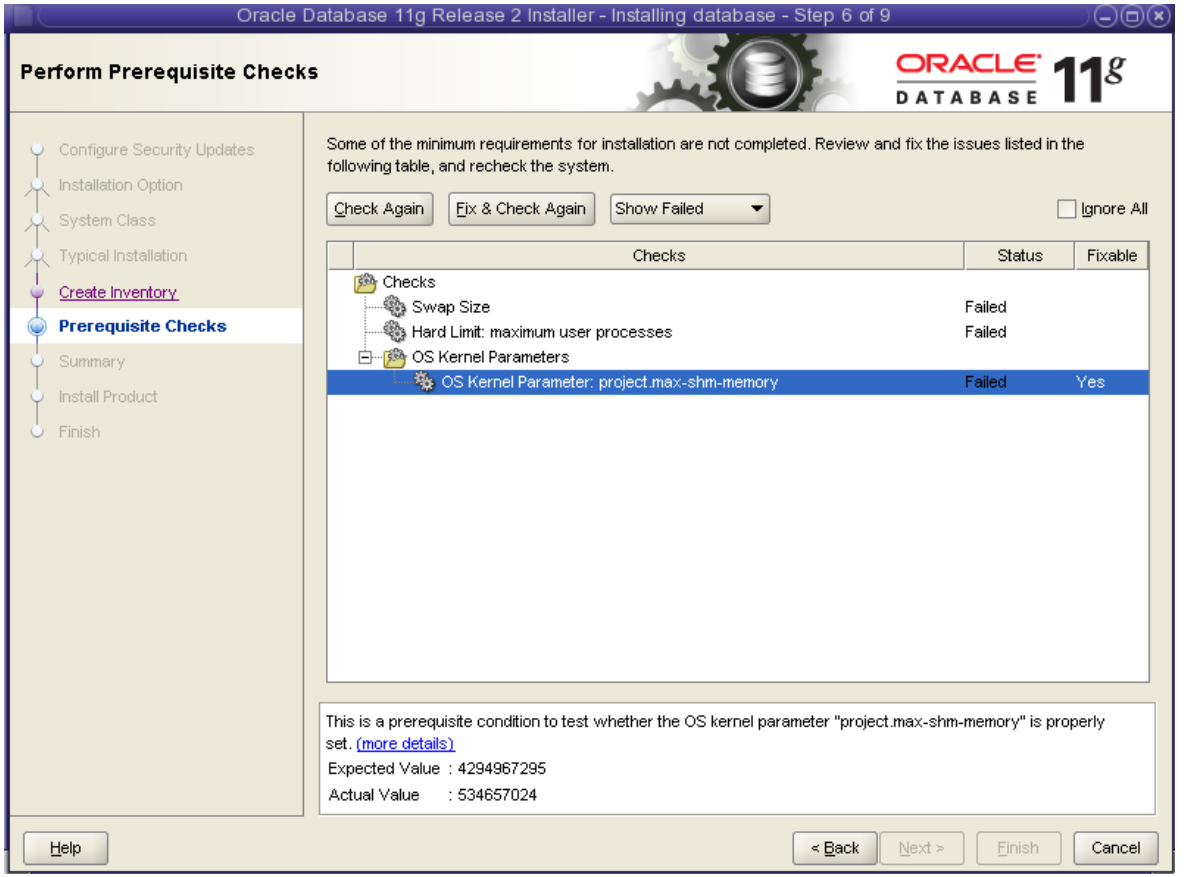

# Solaris 11

portrix systems

- Solaris 10 branded zones
- possibly no more sparse root zones
  - but IPS will change deployments anyway
  - even more flexibility?
- crossbow – network virtualization

portrix
systems

Thank you!

questions?