


# Solaris ZFS & Solaris Zones: The next step in SO

 USE  IMPROVE  EVANGELIZE

**Ezequiel Singer**  
**Sun Campus Ambassador**

開  
放  
的  
열린  
مفتوح  
libre  
मुक्त  
ಮುಕ್ತ  
livre  
libero  
ముక్త  
开放的  
açık  
open  
nyílt  
•••••  
ΠΙΠΠ  
オープン  
livre  
ανοικτό  
offen  
otevřený  
öppen  
ОТКРЫТЫЙ  
வெளிப்படை



## Solaris 10

Dynamic Tracing (DTrace)


Solaris Containers

Predictive Self-Healing

Secure Execution

188 Open Source Apps

Compatibility Guarantee



## Solaris ZFS

---

Transactional System

---

Self-healing data

---

Platter speed performance

---

16 *billion billion*  
times larger than today  
(128 bit)

---

For the “65<sup>th</sup> bit”


---



# Volumes versus ZFS


## Traditional Volumes

- Abstraction: virtual disk
- Partition/volume for each FS
- Grow/shrink by hand
- Each FS has limited bandwidth
- Storage is fragmented, stranded



## ZFS Pooled Storage

- Abstraction: malloc/free
- No partitions to manage
- Grow/shrink automatically
- All bandwidth always available
- Pool allows space to be shared



# ZFS Administration

- Create a storage pool named “home”  
**# zpool create home mirror disk1 disk2**
- Create filesystems “ann”, “bob”, “sue”  
**# zfs create home/ann /export/home/ann**  
**# zfs create home/bob /export/home/bob**  
**# zfs create home/sue /export/home/sue**
- Add more space to the “home” pool  
**# zpool add home mirror disk3 disk4**



# ZFS Administration – Cool Features

- Turn on compression for Ann's data  
**# zfs compression=on home/ann**
- Limit Bob to a quota of 10G  
**# zfs quota=10g home/bob**
- Guarantee Sue a reservation of 20G  
**# zfs reservation=20g home/sue**
- Take a snapshot of Ann's filesystem  
**# zfs create home/ann@tuesday**



## ZFS

- Create a new pool of disks
- Create new filesystems on pools
- Reservations / Quotas
- NFS Sharing
- Snapshots
- Data Compression



## ZFS... more info


- [http://es.wikipedia.org/wiki/ZFS\\_\(sistema\\_de\\_archivos\)](http://es.wikipedia.org/wiki/ZFS_(sistema_de_archivos))
- <http://www.opensolaris.org/os/community/zfs/>
- [http://www.sun.com/software/solaris/zfs\\_learning\\_center.jsp](http://www.sun.com/software/solaris/zfs_learning_center.jsp)



## ZFS

ZFS

Volume Manager + File System



**Let's fly...**



## Solaris Containers

---

Limitless partitioning

---

Thousands of applications on  
one system

---

Ultimate consolidation tool

---

Instant restart

---





# What are Zones? Containers?

Isolated instances of Solaris, > 8000 of them, own/shared filesystem, own root access, own network identity....

Zones

+

Group sets of processors into “pools”, fair-share scheduling of processors....

Resource Management

=

A lightweight OS virtualization facility....

Containers

# Solaris Containers (Zones)

- Isolated application environments within a single Solaris instance
  - Acts like a separate operating system
- Name space, security, and failure isolation
- Dedicated hardware resources not required
- Applications do not change
- Solaris *containers* are zones coupled with resource management

# What are Zones good for?

- Let students have their own UNIX environment to play with (including root access...)
- Create developer “sandboxes” for build and installation
- Test environments for new applications
- Team sharing of compute resources
- Consolidate many servers into one safely

# Solaris Zones: Some details...

- Default *global* zone
- *Sparse-root* or *Whole-root* filesystem
- Zones require 85 MB of disk (sparse-root) and 40 MB of additional memory, < 1% performance hit (see [www.sap.com/benchmarks](http://www.sap.com/benchmarks))
- Theoretically > 8000 zones allowed\*
- No global reboot necessary
- Each zones can have their own IP address
- Each zone has their own root access
- Zones can have local copies of directories
- Zones can provide HTTP, ssh, NIS, sendmail services

# Solaris Containers ([www.sap.com/benchmarks](http://www.sap.com/benchmarks))

<u>SAPS</u>	<u>Operating System - Release</u>	<u>RDBMS Release</u>	<u>ERP Release</u>	<u>Central Server</u>	<u>Central Server Memory (MB)</u>	<u>Certification Number</u>
<a href="#">Desc</a> <a href="#">Asc</a>			<a href="#">Desc</a> <a href="#">Asc</a>			
1170	Solaris 10	<u>MaxDB 7.5</u>	2004	<u>Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1 1.2 GHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache</u>	32768	<a href="#">2006032</a> (ran concurrently with 2006029, 2006030, 2006031)
1180	Solaris 10	<u>MaxDB 7.5</u>	2004	<u>Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1 1.2 GHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache</u>	32768	<a href="#">2006031</a> (ran concurrently with 2006029, 2006030, 2006032)
1180	Solaris 10	<u>MaxDB 7.5</u>	2004	<u>Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1 1.2 GHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache</u>	32768	<a href="#">2006030</a> (ran concurrently with 2006029, 2006031, 2006032)
1170	Solaris 10	<u>MaxDB 7.5</u>	2004	<u>Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1 1.2 GHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache</u>	32768	<a href="#">2006029</a> (ran concurrently with 2006030, 2006031, 2006032)
4780	Solaris 10	<u>MaxDB 7.5</u>	2004	<u>Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1 1.2 GHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache</u>	32768	<a href="#">2005047</a>

**4700/4780 = 0,983263598 menos de 1,7% de overhead!!**

El unico otro particionamiento auditado en SAP es la LPAR de un IBM p670 pero nunca auditaron la opcion sin particion...



## Solaris Zones: CLI

% poolcfg – Creates processor sets and resource pools

% pooladm – Instantiates resource configuration

% priocntl – Change scheduling for processes

% zonecfg – Create and configure zones

% zoneadm – Install and boot zones

% zlogin – Log into zones





# Solaris Zones/Containers Demo

Global Zone 1 Processor Share 192.168.1.128	Web Zone 3 Processor Shares 192.168.1.150	Software Zone 2 Processor Shares 192.168.1.151
<b>/usr, /platform, /sbin, /lib</b>	<b>(r/o)</b>	<b>(r/o)</b>
<b>/etc, /var</b>	<b>/etc, /var (r/w)</b>	<b>/etc, /var (r/w)</b>
<b>/usr/local</b>		<b>/usr/local (r/w)</b>
<b>/cdrom</b>		<b>/cdrom</b>

Ultra 20, AMD Single-Core Opteron, Solaris 10 x86

(resource management becomes much more fun on multi-processor machines)



## Fast Track Your Career with Free Training from Sun

- Entrenamiento Web disponible:
  - > Java
  - > Solaris 10
  - > Java Enterprise System
  - > XML, Perl y más
- Materiales para Entrenamiento con Instructor:
  - > Java
  - > Solaris 10
- Registrarse en: <http://learningconnection.sun.com>

# Solaris ZFS & Solaris Zones: The next step in SO



**Ezequiel Singer**  
[ezequiel.singer@sun.com](mailto:ezequiel.singer@sun.com)  
[blogs.sun.com/argentina\\_ambassador](https://blogs.sun.com/argentina_ambassador)

**Sun Campus Ambassador**

開  
放  
的  
열린  
مفتوح  
libre  
मुक्त  
ಮುಕ್ತ  
livre  
libero  
ముక్త  
开放的  
açık  
open  
nyílt  
•••••  
πικρ  
オープン  
livre  
ανοικτό  
offen  
otevřený  
öppen  
ОТКРЫТЫЙ  
வெளிப்படை