

Solaris Zones on OpenSolaris

Javier Conde

Swiss OpenSolaris User Group
Geneva, Uni Dufour
October 2nd 2008

<http://wikis.sun.com/display/chosug/>

Solaris Zones with OpenSolaris

- Solaris Zones Overview
- How does it work ?
- Demo

Solaris Zones Overview

- Basic concept: isolated execution environment within a Solaris instance
 - Resource, security and fault isolation
 - Lightweight, flexible, efficient
 - One OS to manage

Solaris Zones Overview

- Virtualizes OS layer: file system, devices, network, processes
- Provides:
 - *Privacy*: can't see outside zone
 - *Security*: can't affect activity outside zone
 - *Failure isolation*: application failure in one zone doesn't affect others
- Lightweight, granular, efficient
- Delegated and simplified administration
- Complements resource management

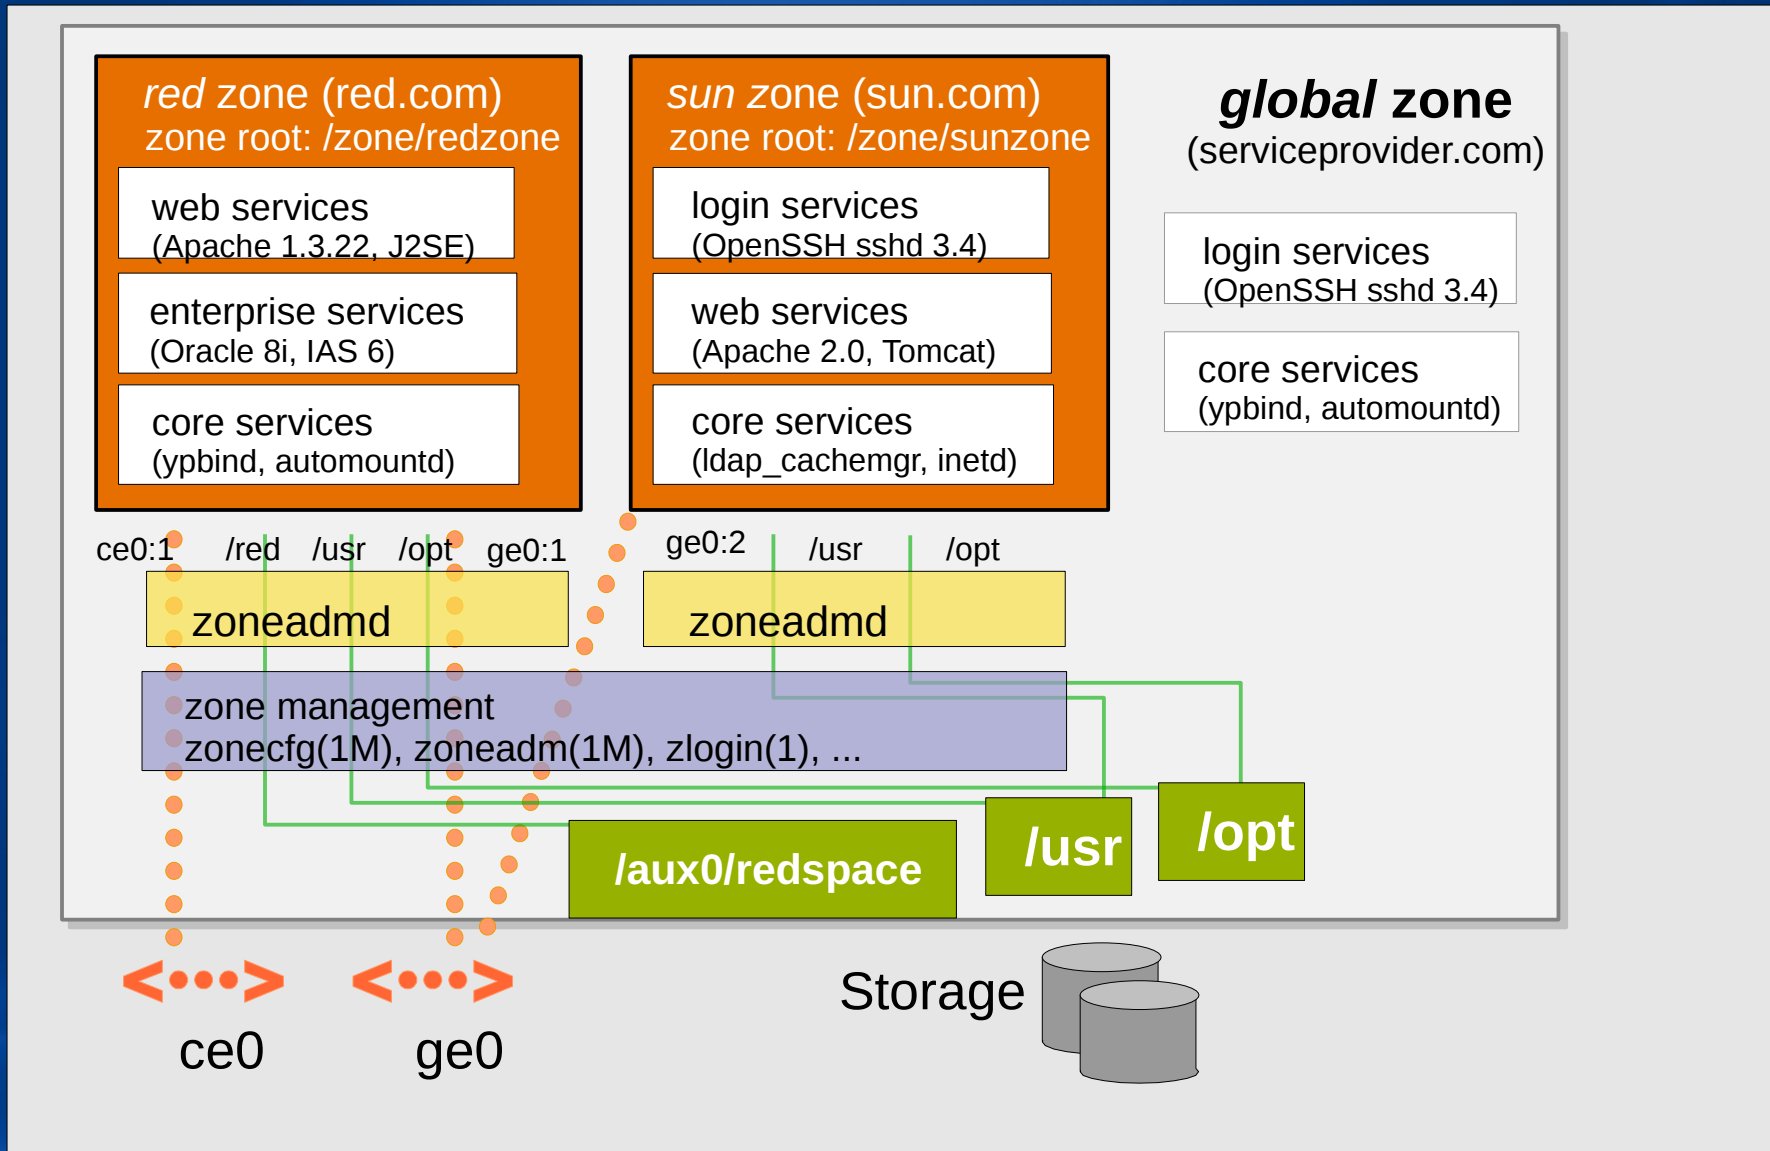
Solaris Zones Security

- Each zone has a security boundary
- Runs with subset of `privileges(5)`
- Compromised zone unable to escalate its privileges
- Important name spaces are isolated
- Processes running in a zone are unable to affect activity in other zones

Resource Management

- Combined with resource management, zones provide a more complete isolated environment
- Zones can be bound to a resource pool to isolate their effect on system resources
- Processors can be partitioned with arbitrary granularity using the fair share scheduler
- Resource limits can be set on a zone as well

Solaris Zones



Solaris Zones with OpenSolaris

- ZFS integration
- Rename and Move
- Clone
- Migration
- Configurable privileges
- DTrace support
- Boot arguments
- IPS

Solaris Zones with OpenSolaris

- Zones direct RM to build a pool on a temporary basis.
 - Pool only setup when the zone boots
 - Pool must be *dedicated* to the zone it is serving
- With specification in zonecfg, pool settings propagate during migrations

```
zonecfg:ex> add dedicated-cpu
zonecfg:ex:dedicated-cpu> set ncpus=1-8
zonecfg:ex:dedicated-cpu> set importance=50
zonecfg:ex:dedicated-cpu> end
```

Solaris Zones with OpenSolaris

rctl aliases feature aids setting resource controls:

```
zonecfg:endzone> set max-lwps=500  
zonecfg:endzone> set cpu-shares=5  
zonecfg:endzone> set scheduling-class=FSS  
zonecfg:endzone> set max-shm-memory=10M
```

```
$ zonecfg -z global  
zonecfg:global> set cpu-shares=10
```

Solaris Zones with OpenSolaris

```
zonecfg:endzone> set max-lwps=500
zonecfg:endzone> add dedicated-cpu
zonecfg:endzone:dedicated-cpu> set ncpus=1-5
zonecfg:endzone:dedicated-cpu> end

zonecfg:endzone> add capped-memory
zonecfg:endzone:capped-memory> set locked=64M
zonecfg:endzone:capped-memory> set physical=128M
zonecfg:endzone:capped-memory> set swap=256M
zonecfg:endzone:capped-memory> end
```

Solaris Zones with OpenSolaris

- Available today:
 - Linux zones on Solaris
 - x86 only
 - Solaris 8 zones on Solaris 10
 - SPARC only
- Other possibilities
 - Alternate Solaris zones
 - Nexenta/ShilliX/BeleniX
 - Replace Solaris tools in /usr/bin with GNU equivalents

Solaris Zones with OpenSolaris

- Zone Migration Enhancements
 - “update on attach”

```
# cd /export/zones/u6
# tar xf ~/images/s/u4.tar
# rm SUNWdetached.xml
# zoneadm -z u6 attach -u
zoneadm: zone 'u6': The zone was not properly detached.
      Attempting to attach anyway.
Getting the list of files to remove
Removing 1751 files
Remove 560 of 560 packages
Installing 1889 files
Add 609 of 609 packages
Updating editable files
The file </var/sadm/system/logs/update_log> within the zone
contains a log of the zone update.
```

Varies based on the migration delta (e.g. a single patch difference is minimal and fast)

Examples of uses

- Data center workload consolidation
- Software development
 - Test vs. Production
- Hostile or untrusted applications
- Hosting environments
- WAN-facing services
 - Break-in containment

demo