



Solaris 10 Zones Overview

Chris Shepard
Technical Specialist
Data Center Practice
Sun Microsystems, Inc.




Introduction to Solaris 10 Zones

- Virtual Operating Environments
- Single Solaris Instance
- Single Global Zone
- Up to 8192 Local Zones
- Lightweight Server Consolidation Vehicle


Traditional Resource Management

- One application per server
- One or more servers per customer
- Every server sized for peak workload
- Low average utilization rates

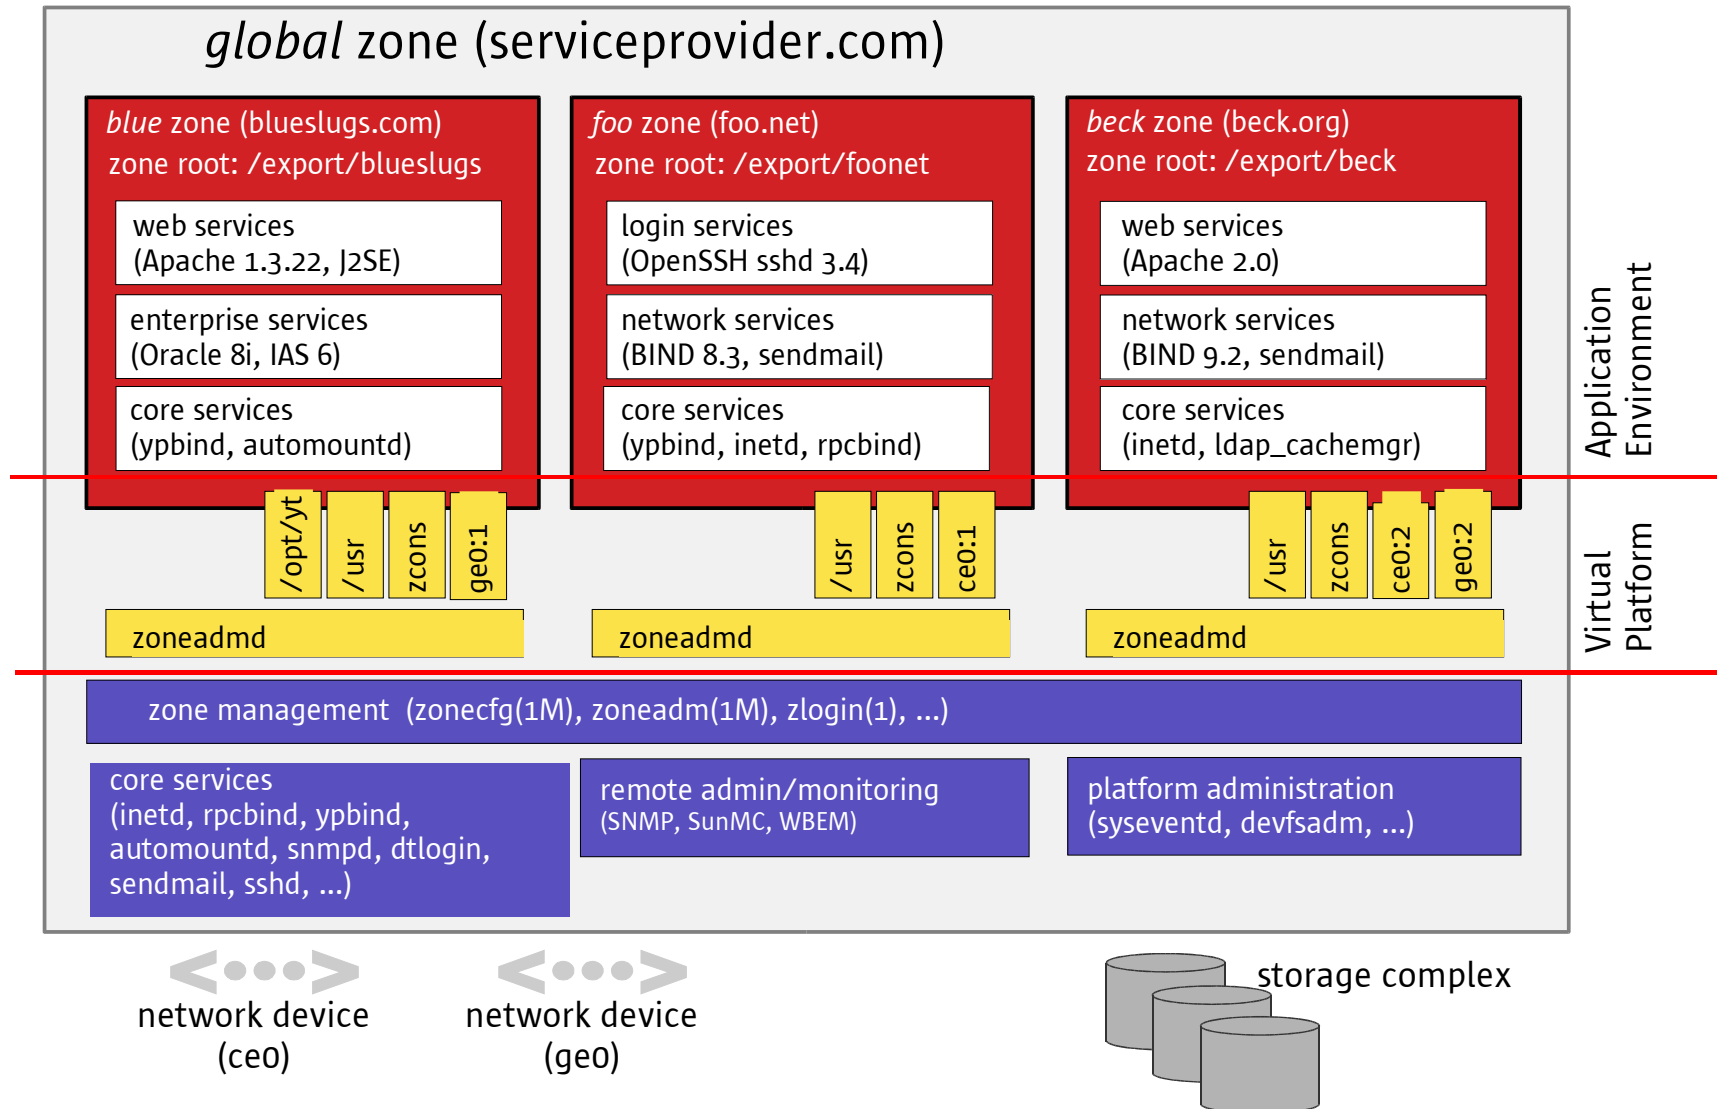


Solaris 10 Zones

- Allow applications to run in secure isolation from other applications.
- Each Zone Has Private
 - Logical IP Address
 - Root File System
 - Process Table
- Default Inheritance of Global Zone Attributes
- Increased Utilization



Zones Block Diagram



Zone Properties

- Lightweight
- Virtualized
- Isolated
- Transparent
- Granular

Zone Properties

- Lightweight
 - Bundled With Solaris 10
 - Private User Space
 - Shared Kernel
 - New System V IPC Resource Controls
 - Dynamic – No Reboot Required
 - Per User or Per Project
 - Was `/etc/system/shmsys_shminfo_shmmax`
 - Now `/etc/project/project.max-shm-memory`

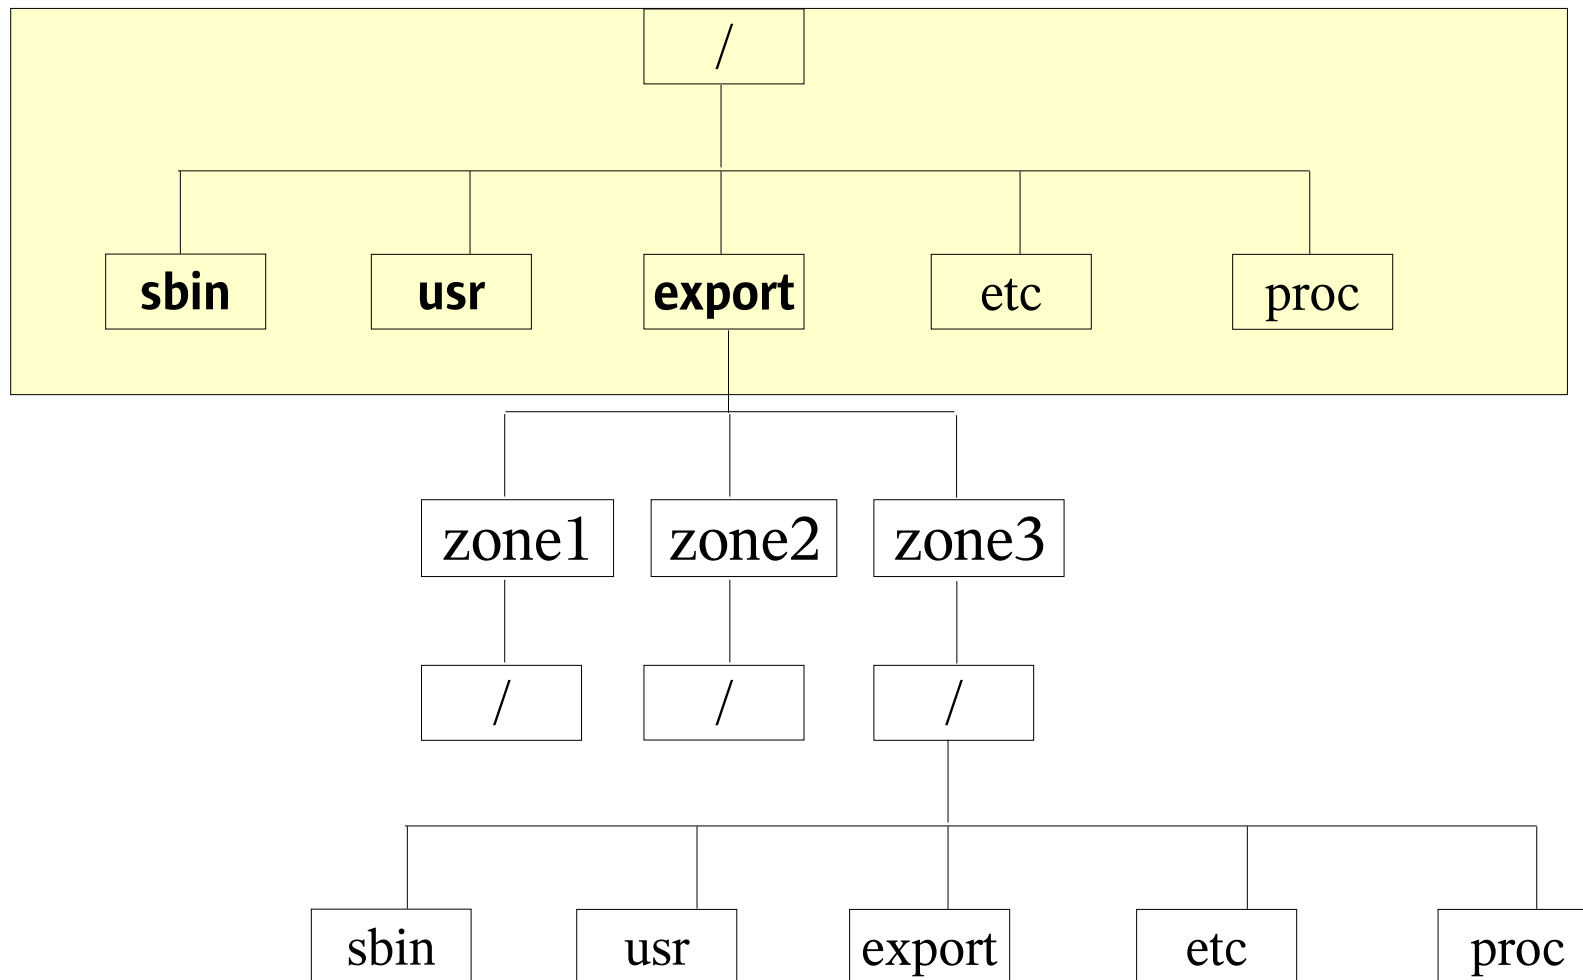
Zone Properties

- Virtualized
 - A zone appears as an instance of Solaris
 - An instance within an instance
 - Zones can be administered separately
 - Separate name services
 - Separate users, patches, packages

Zone Properties

- Isolated
 - Everything in a local Zone is isolated from things in other local zones.
 - A zone contains :
 - processes, files, interfaces, and System V IPC elements.
 - Includes a separate root file system
 - Can include separate users (including root)
- Faults and Security

Isolation: root filesystem



Isolation: Processes

```
# hostname
```

```
demozone1
```

```
# ls /proc
```

```
18332 18359 18389 18427 18437 18444 18459 18491 18546 19777
```

```
18335 18386 18398 18433 18443 18456 18464 18545 18549 20243
```

```
# hostname
```

```
demozone2
```

```
# ls /proc
```

```
19959 19986 20018 20056 20066 20073 20088 20120 20177 20180
```

```
19962 20015 20027 20062 20068 20085 20093 20176 20178 20240
```

Zone Properties

- Transparency
 - Applications run as is
 - Standard Solaris interfaces are provided
 - Restrictions limited to privileged operations
 - No global reboots, access to /dev/kmem, physical devices

Zone Properties

- Granular
 - Partitions at sub-CPU level
 - Only file system space is dedicated
 - HW resources multiplexed across zones
 - Resource controls on cpu, memory, bandwidth can be used

Solaris Containers

Features

- Projects and Extended Accounting
- Scheduling Classes
 - TS, IA, RT, FSS, FX
- Resources Pools
 - Processor Sets, Memory Capping Daemon
 - Dynamic Resource Pools
- IPQoS

Creating a zone

```
global# zonecfg -z zone1
```

```
zone1: No such zone configured
```

Use 'create' to begin configuring a new zone.

```
zonecfg:zone1> create
```

Setting's for the zone

```
zonecfg:zone1> set zonepath=/zoneroots/zone1
```

```
zonecfg:zone1> set autoboot=true
```

```
zonecfg:zone1> add net
```

```
zonecfg:zone1:net> set address=192.9.200.67
```

```
zonecfg:zone1:net> set physical=hme0
```

```
zonecfg:zone1:net> end
```

```
zonecfg:zone1> ^D
```

```
#zoneadm list -c
```

Installing the zone

global# zoneadm -z zone1 install

Constructing zone at /zoneroot/zone1/root

Creating dev directories

Creating dev links

Copying packages and creating contents file

Copying files and directories

Setting up /etc/motd

Setting up /etc/inittab

Setting up /etc/vfstab

Setting up /var/yp/aliases

Configuring files

boot the zone

```
global# zoneadm -z zone1 boot
```

– Took about 30 seconds for first boot on Ultra10.

- global# zlogin -C zone1
- [Connected to zone 'mydesktop' console]
- <Run through sysid tools as usual to do initial customization>

Example: zones and fs

```
#zonecfg -z name
zonecfg:zone1> add fs
zonecfg:my-zone:fs> set dir=/opt/local
zonecfg:my-zone:fs> set special=/local
zonecfg:my-zone:fs> set type=lofs
zonecfg:my-zone:fs>end
zonecfg:zone1> verify
zonecfg:zone1> commit
zonecfg:zone1> ^D
```

this will mount the /local directory from the global to a mount point of /opt/local in the zone

Example: RM+zones

```
#zonecfg -z name
zonecfg:zone1> add rctl
zonecfg:zone1:rctl> set name=zone.cpu-shares
zonecfg:zone1:rctl> add value \ (priv=privileged,limit=10,action=none)
zonecfg:zone1:rctl> end
zonecfg:zone1> verify
zonecfg:zone1> commit
zonecfg:zone1> ^D
```

```
#prctl -n zone.cpu-shares -r -v 25 -i zone zonename
```

Example: zones and fs

```
#zonecfg -z name  
zonecfg:my-zone> add device  
zonecfg:my-zone:fs> set match=/dev/dsk/c0t0d0s0  
zonecfg:zone1> verify  
zonecfg:zone1> commit  
zonecfg:zone1> ^D
```

this will give that zone full permission to the device

Solaris Container Manager

1.1

- Sun Management Center 3.5 U1b
- Current GA Schedule June 2005
- GUI for Resource Mgmt, Zone Mgmt.

New Zone – Web Browser

N1™ Grid Console - Container Manager

New Zone

Steps

Help

Step 1: Provide zone creation parameters.

→ 1. Provide zone creation parameters.

2. Provide additional zone attributes.

3. Provide IPGoS attributes.

4. Select a Resource Pool.

5. Review the selections made for the zone.

* Indicates required field

* Zone Name :
Cannot exceed 32 characters

* Zone Host Name :
Cannot exceed 32 characters

* Zone Path :

IP Address :

Network Interface :

* CPU Shares :
Type an integer value which corresponds to share of CPU resources which will be used by this zone.

* Maximum CPU Shares :
Type an integer value which corresponds to the Maximum CPU resource which will be used by this zone.

Auto Reboot : Enabled

BigAdmin Portal

- sun.com/bigadmin/features/articles/zones_partition.html

Consolidation White Paper

- “Consolidating Applications with Solaris Containers”
- sun.com/datacenter/consolidation

Summary: Zones

N1 Grid Containers provides a significant improvement in the area of **server consolidation**. It allows for hosting different applications for various departments or customers in a **secure isolated container**. The amount of consolidation is much greater because of **low overhead** and that **no dedicated hardware** is required.



Solaris 10 Zones Overview

christopher.shepard@sun.com

