



# Zones and Containers

Lightweight virtualization in Solaris

Tomáš Dzik

Solaris Cluster Revenue Product Engineering

Sun Microsystems, Czech

---

# Zones or Containers ?

---

- Container is a zone with resource control
- Resource control could be applied to
  - CPU, memory, NIC
  - See `resource_controls(5)`



# Zones and Containers

---

- Type of OS virtualization
  - Each zone has it's own userspace
  - All zones share the same kernel
- Isolation is provided by kernel
- Same type of virtualization like a Linux vserver
- Not dependent on hardware
- Available by default

# Zones Architecture



# Zones Architecture

---

- Global zone exists by default
  - If no local zones are configured, all processes run in global zone
- Global zone can “see” processes in local zones
- Try:
  - `zoneadm list -iv`
  - `zonename`

# Zones and Containers

---

- Zones are great in combination with
  - ZFS (you can clone zones)
  - Crossbow (network virtualization)
- Zones could be moved from one machine to another
  - Could be even made HA using Sun Cluster

# Crossbow

---

- Allows creating of
  - virtual NICs
  - virtual switches
  - setting bandwidth limits
  - a lot of other things

# Example

---

Creating a zone:

```
pfexec zonecfg -z zone-example-1
```

```
zone-example-1: No such zone configured
```

Use 'create' to begin configuring a new zone.

```
zonecfg:zone-example-1> create
```

```
zonecfg:zone-example-1> set zonepath=/export/zones/zone-example-1
```

```
zonecfg:zone-example-1> set ip-type=exclusive
```

```
zonecfg:zone-example-1> add net
```

```
zonecfg:zone-example-1:net> set physical=zonenic1
```

```
zonecfg:zone-example-1:net> end
```

```
zonecfg:zone-example-1> verify
```

```
zonecfg:zone-example-1> commit
```

```
zonecfg:zone-example-1> exit
```

# Branded zones

---

- lx brand
  - Linux userspace above the Solaris kernel
  - Solaris kernel translate Linux syscalls
- Solaris 8 and 9 brand
  - Available for Sol 10 migration on sparc
- ipkg brand

# Zones Taxonomy

---

- Whole-root zones
  - Copy of the whole filesystem
- Sparse-root zones
  - Inherits some readonly filesystems
  - /lib /platform /sbin /usr
- Shared-ip stack
  - All zones share the same ip stack
- Exclusive-ip stack
  - There are separate data structures in kernel for each zone – allows tuning of tcp/ip stack

# Zones administration

---

## ○ Command line tools

- `zonectfg(1M)`
  - Creates zone configuration
- `zoneadm(1M)`
  - Changes state of zone
- `zlogin(1)`

# Process visibility

---

- `ps -Z`
- `prstat -Z`

# Example

---

Export config:

Create clone zone:

# Find more at ...

---

- <http://www.opensolaris.org/os/community/zones/>
- <http://www.opensolaris.org/os/project/crossbow/>
- System Administration Guide: Solaris Containers-Resource Management and Solaris Zones

# Zones

---

Questions ?



# Zones and Containers

Lightweight virtualization in Solaris

Tomas.Dzik@sun.com

---