Solaris Virtualization: Zones and Containers

Bart Muijzer
Systems Architect, Operating Systems Ambassador
SUN Microsystems Nederland

Server Virtualization Approaches

Hard Partitions

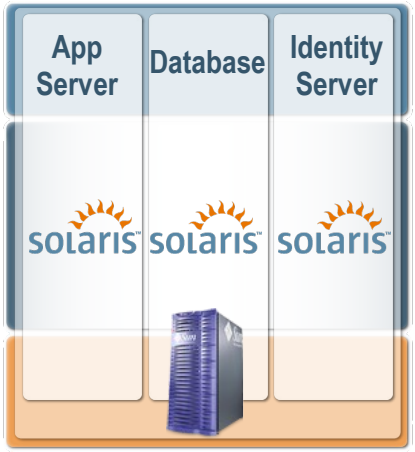

Virtual Machines

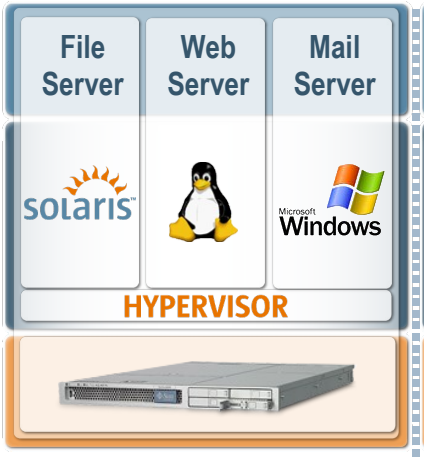

OS Virtualization

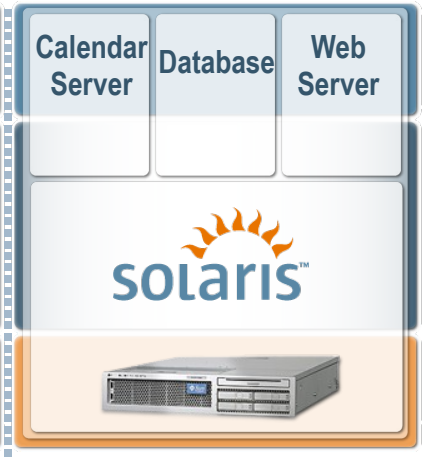

Resource Mgmt.

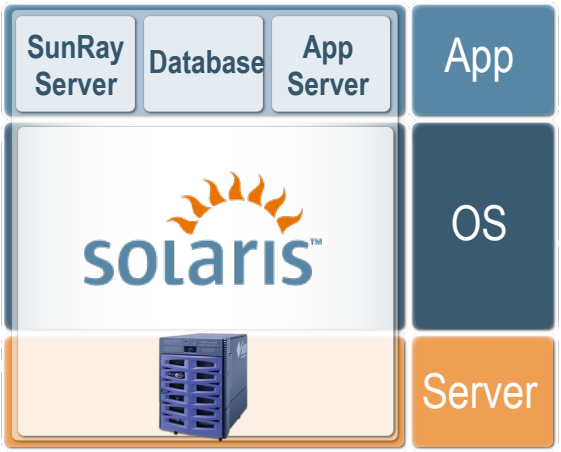

Multiple OSes ← → Single OS

Trend to flexibility →

← Trend to isolation

- > Very High RAS
- > Very Scalable
- > Mature Technology
- > Ability to run different OS versions

- > Ability to live migrate an OS
- > Ability to run different OS versions and types
- > De-couples OS and HW versions

- > Very scalable and low overhead
- > Single OS to manage
- > Cleanly divides system and application administration
- > Fine grained resource management

- > Very scalable and low overhead
- > Single OS to manage
- > Fine grained resource management

Server Virtualization

Solutions from Sun

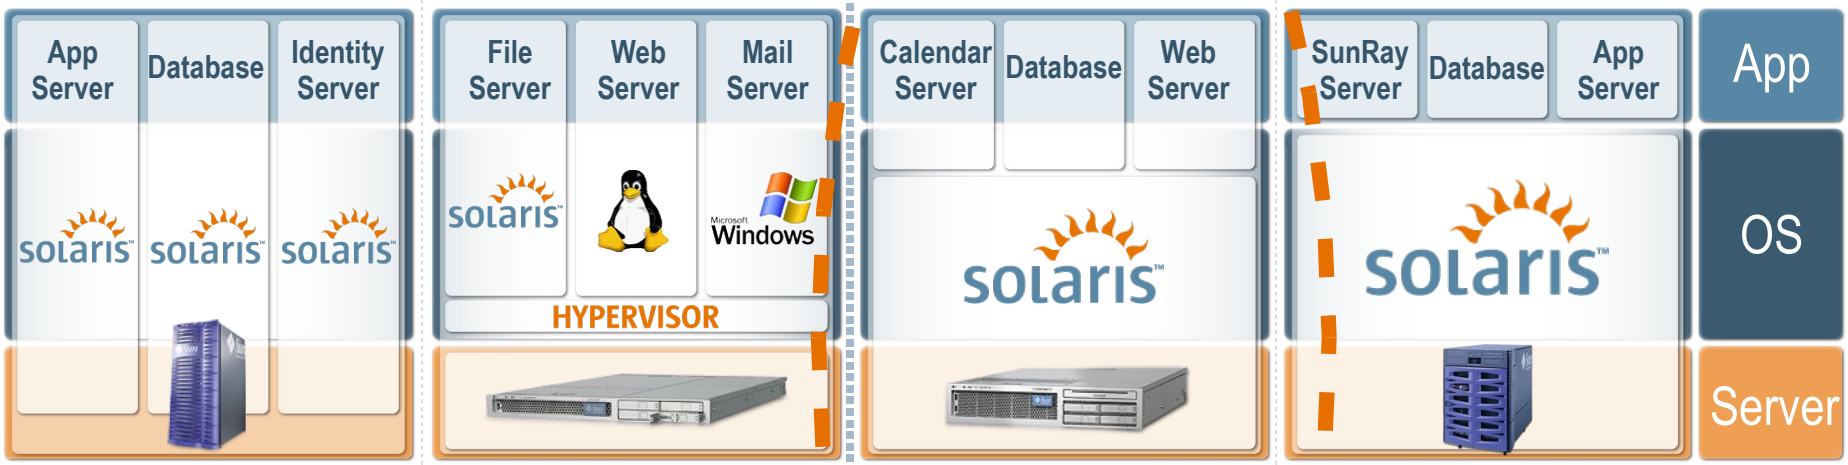

Multiple OSes — Single OS

Trend to flexibility

Trend to isolation

Dynamic System Domains

Logical Domains

Xen

VMware

Microsoft Virtual Server

Solaris Containers
(Zones + SRM)

Solaris Containers
for Linux Applications

Solaris Trusted Extensions

Solaris Resource Manager
(SRM)

OS Virtualization (1)

Solaris zones

Resource and namespace isolation

Very scalable

Available on all platforms

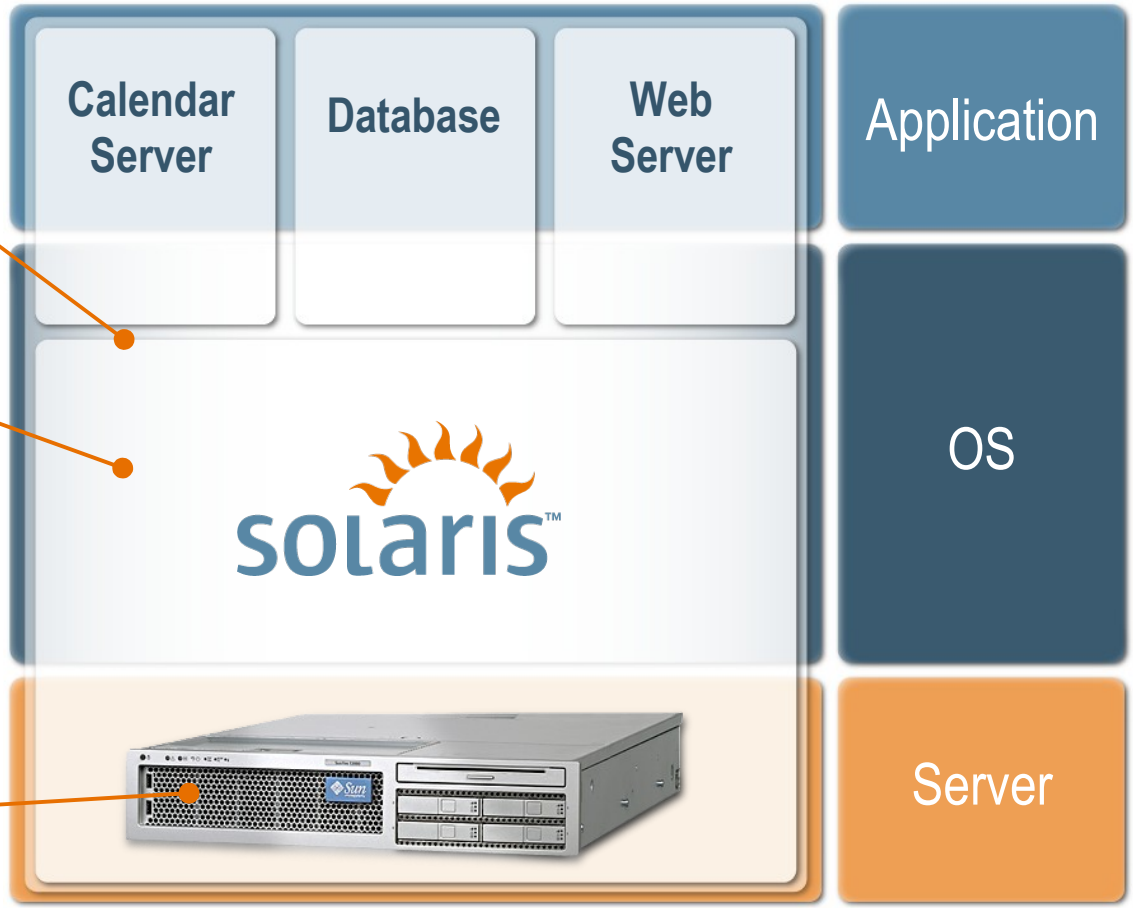

Solaris Zones

- Namespace isolation
- Virtualized OS
- Sharing for better utilization
- Application fault containment

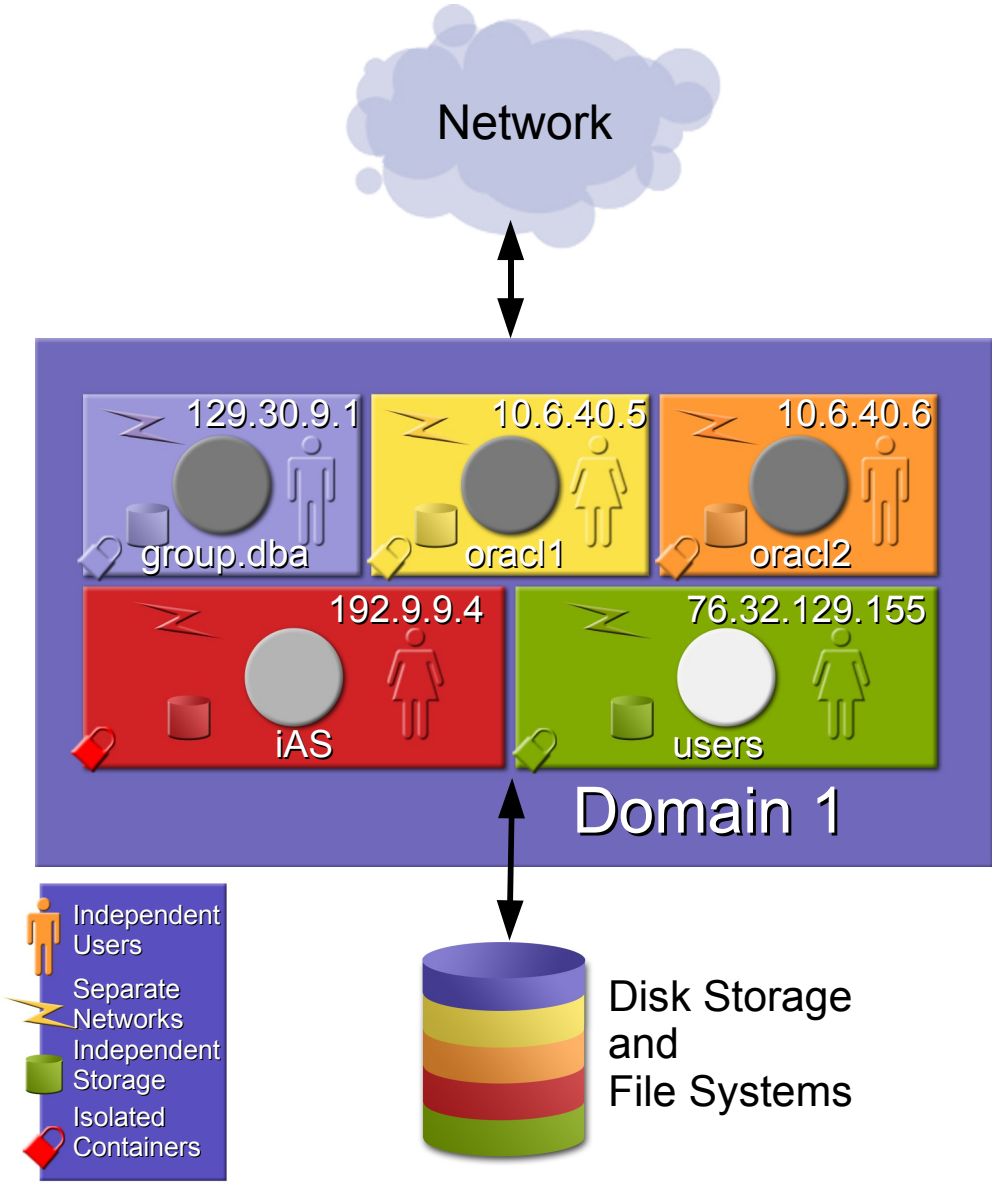

OS Virtualization (2)

Resource Management

Resource controls only

Very scalable

Available on all platforms

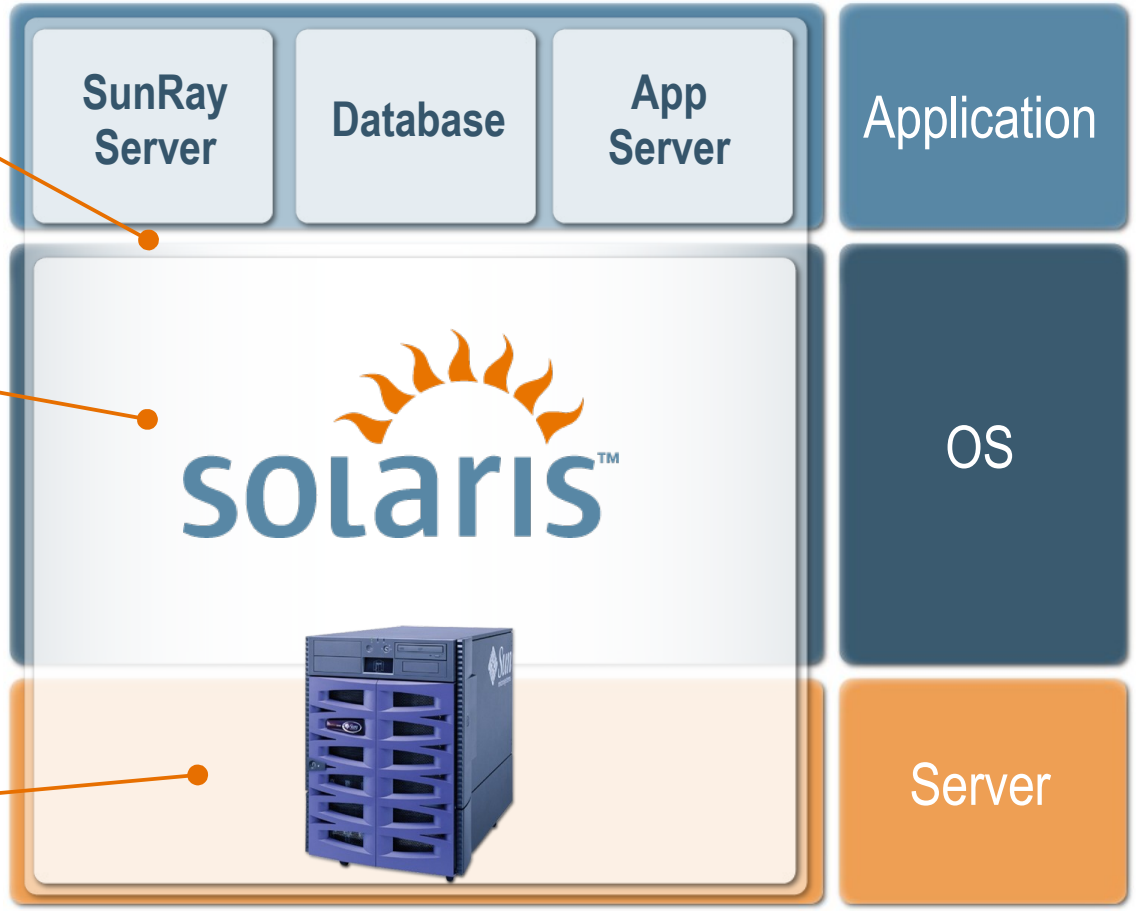

Resource Management

- Resource set
  - Partitions of the hardware resources
  - Can be: CPU
  - Future: memory and SWAP
- Resource pools
  - Logical partitions of different resource sets
  - Multiple pools can link to the same set
  - Dynamic: resources are re(allocated) to meet demand and objectives
- Projects
  - Workload labels linked to a Resource Pool
  - Enables processes running in a project to have specific resource sets
  - Mechanism of “shares” to assign right amount of CPUs to workloads

Resource sets

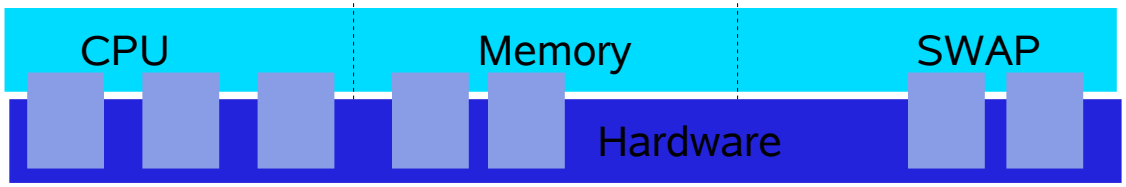

Resource Pools

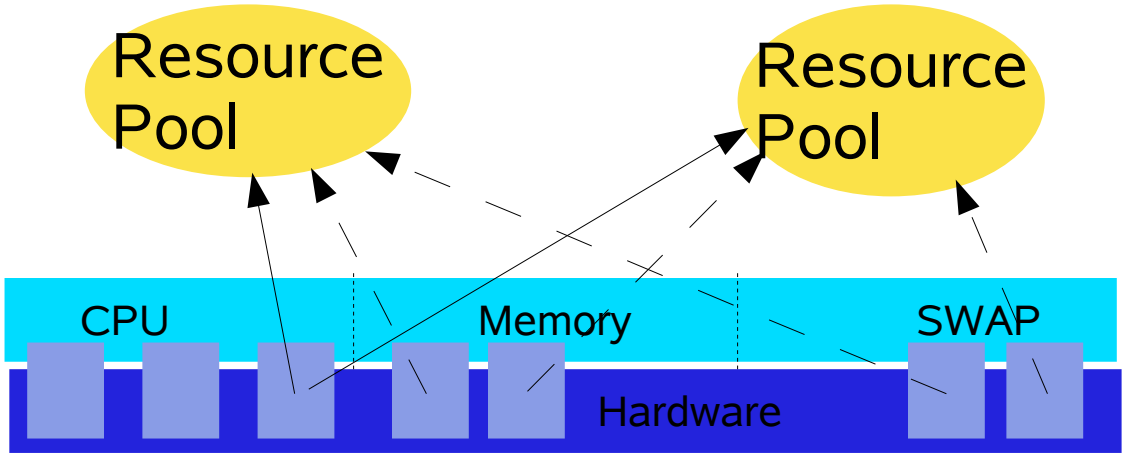

Projects

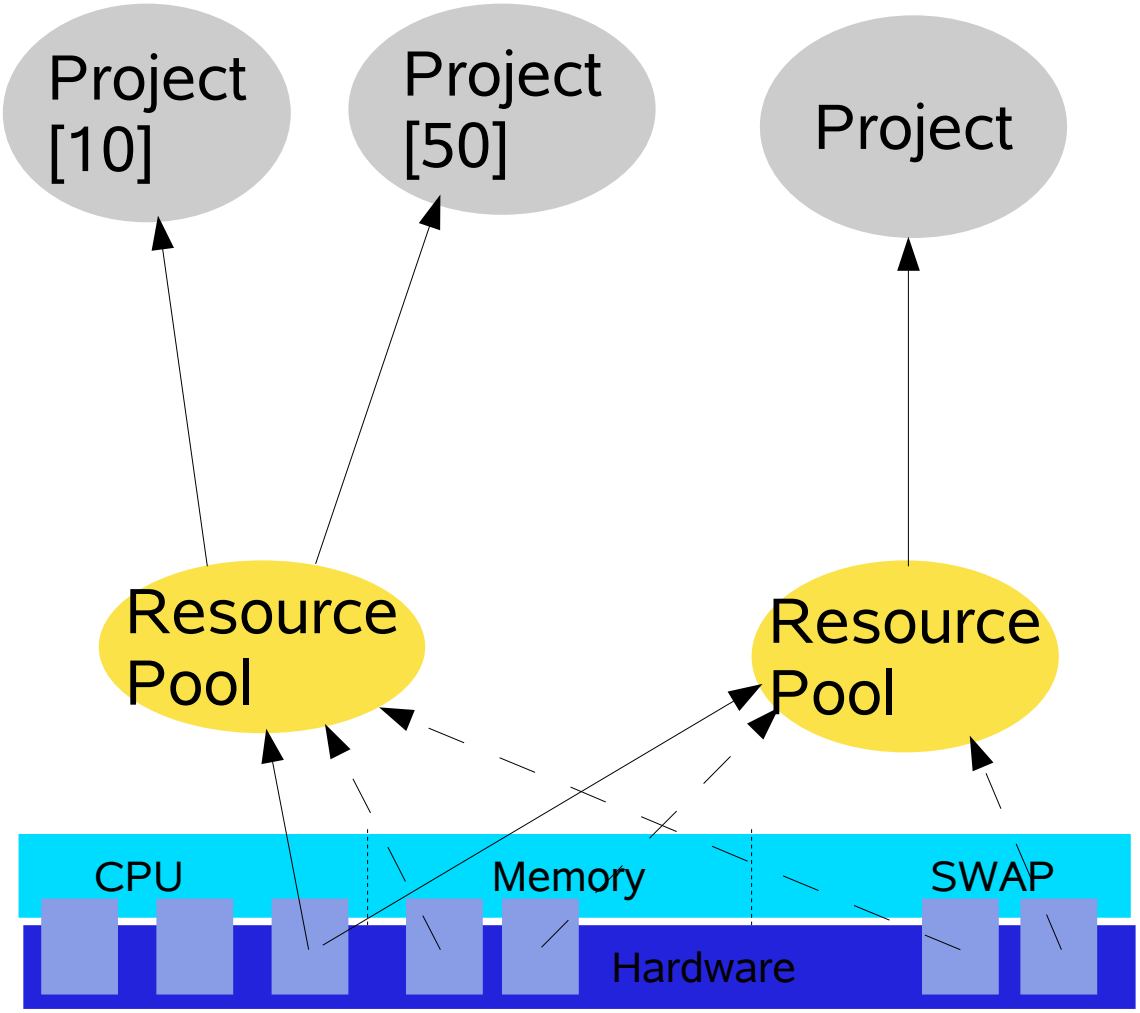
Project
[10]
Project
[50]
Project
Resource
Pool
Resource
Pool
CPU
Memory
SWAP
Hardware

Zones and Resource Mgt

Solaris Zones

+

Solaris Resource Manager

=

Solaris Containers

OS Virtualization (3)

Solaris Containers

Resource and namespace isolation with Resource Controls

Very scalable

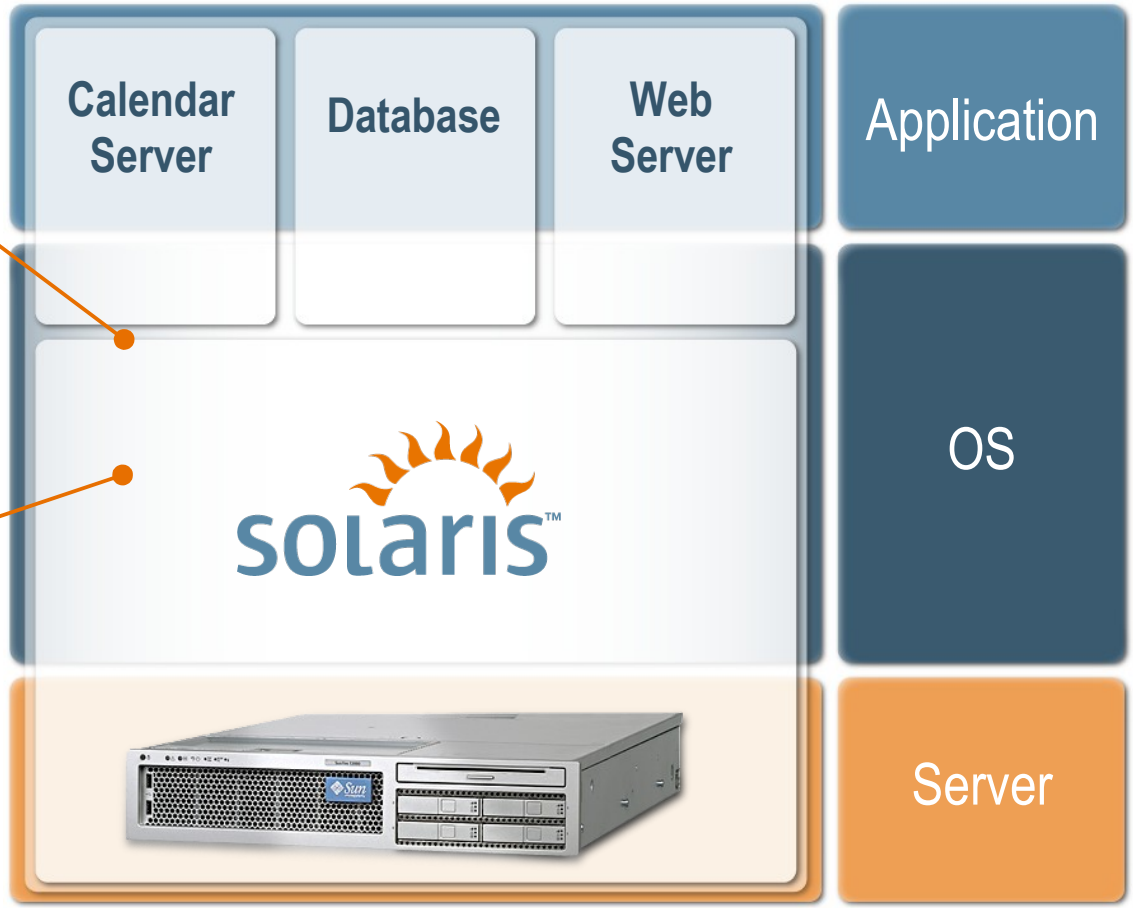

Containers and Resource Management

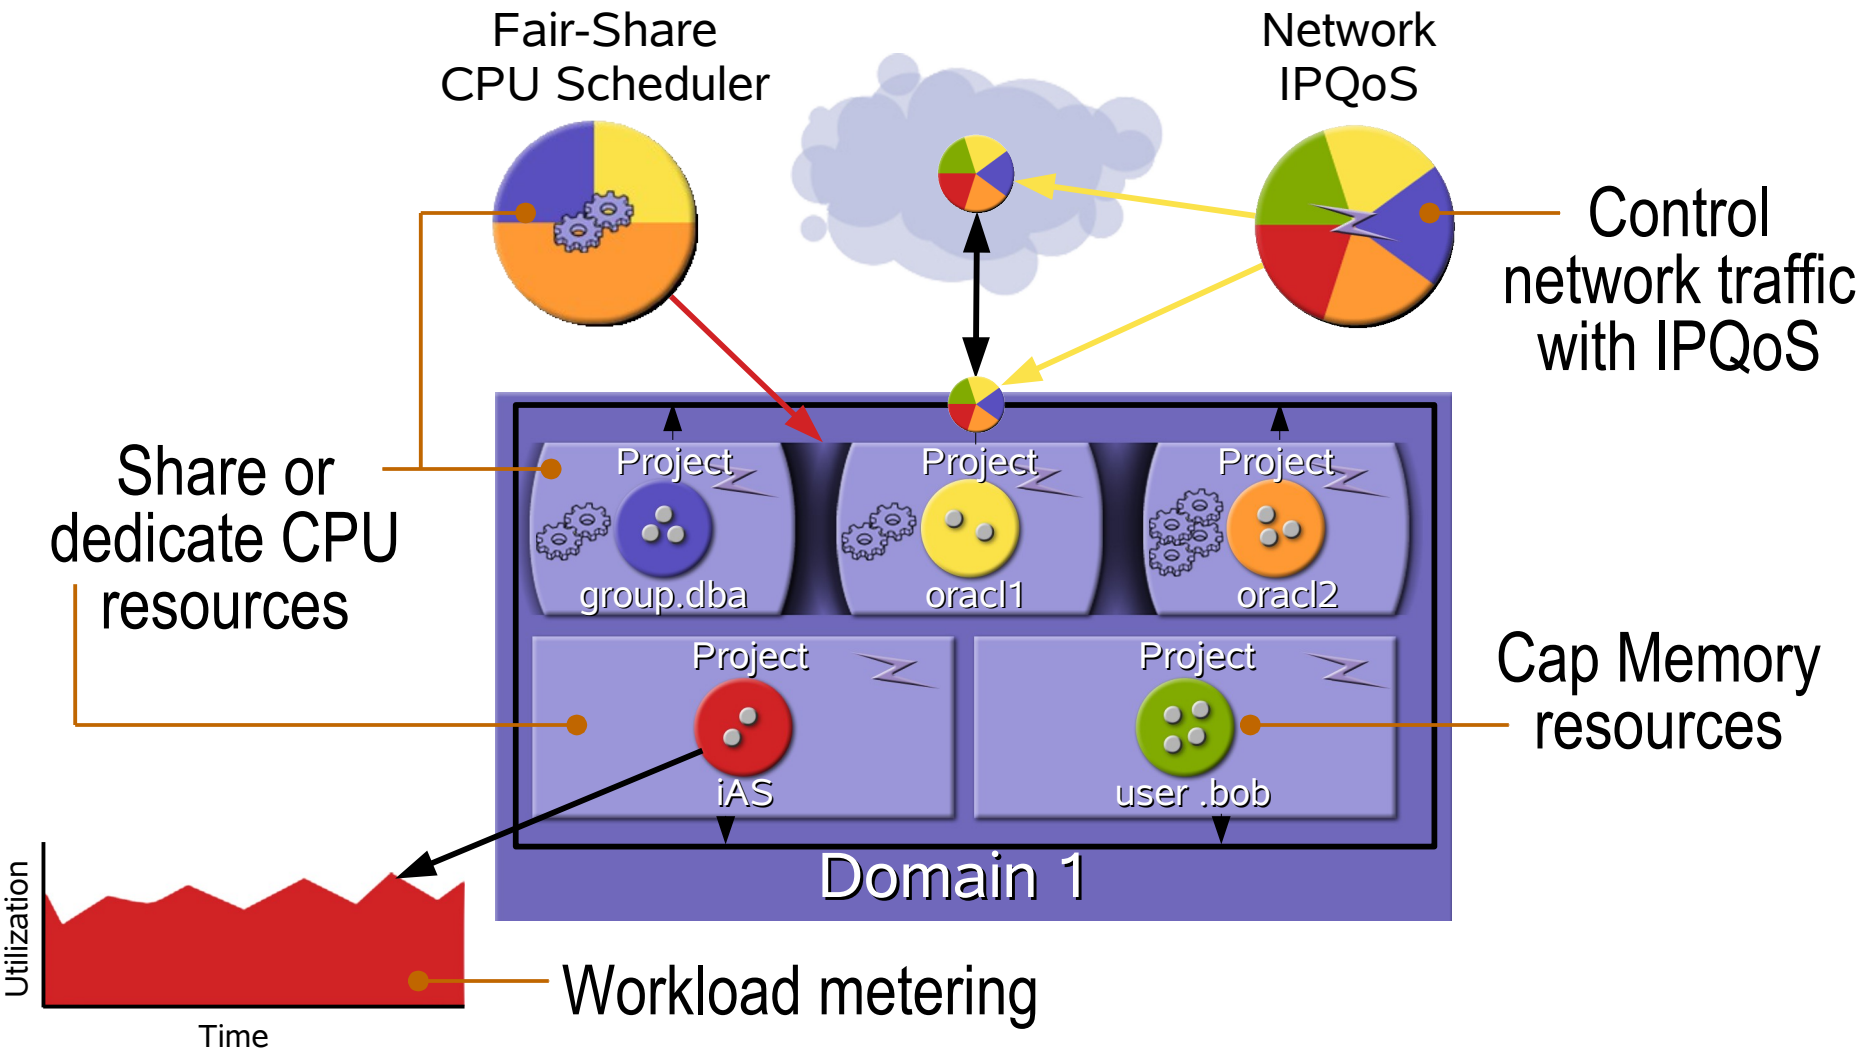

S10 Resource Management

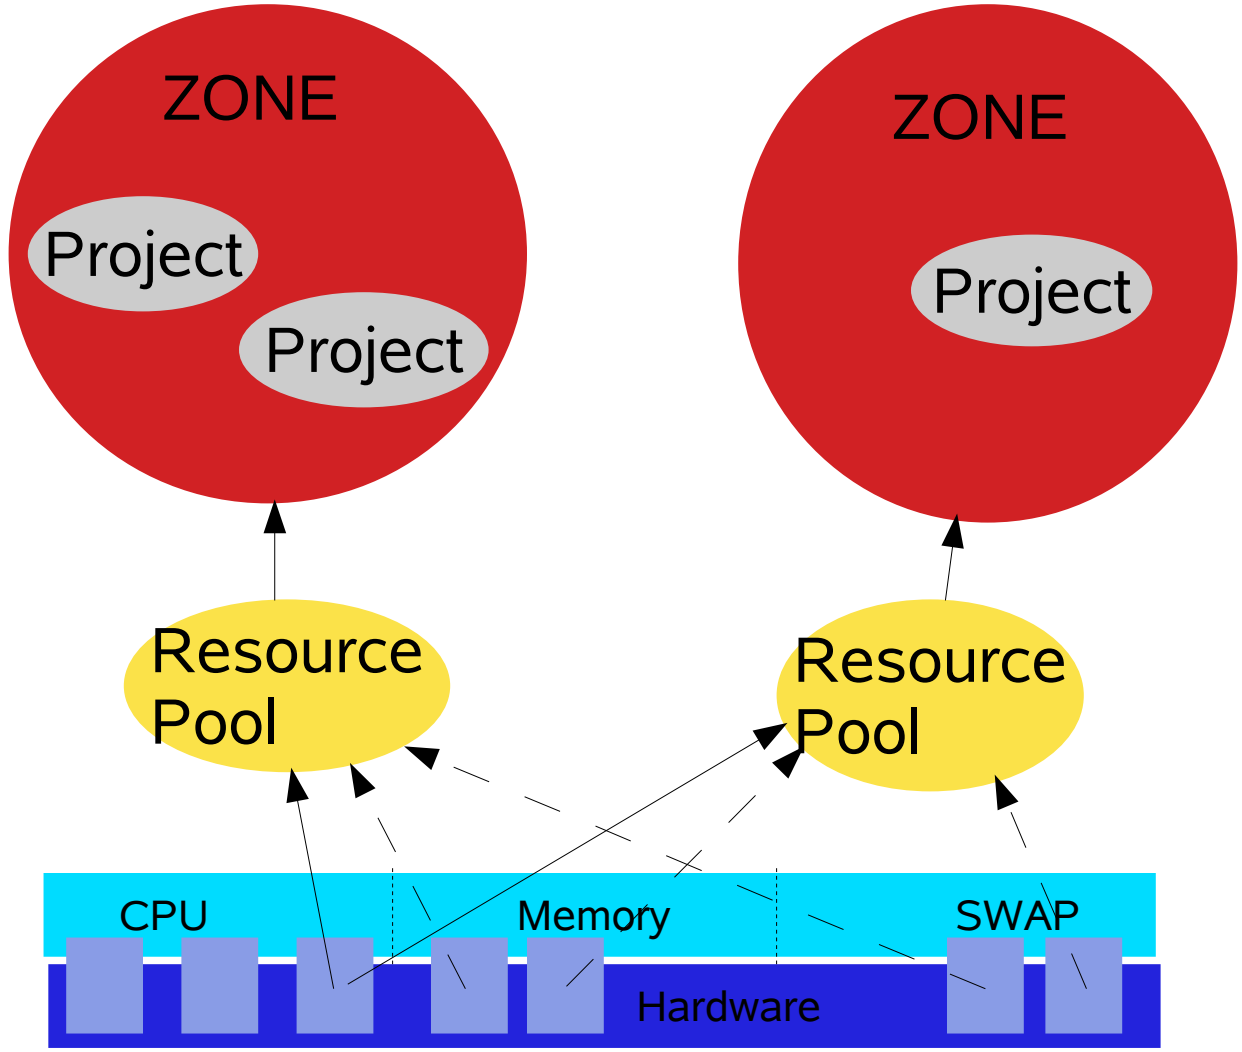

Solaris Container

Resource Management – Fair Share Scheduler

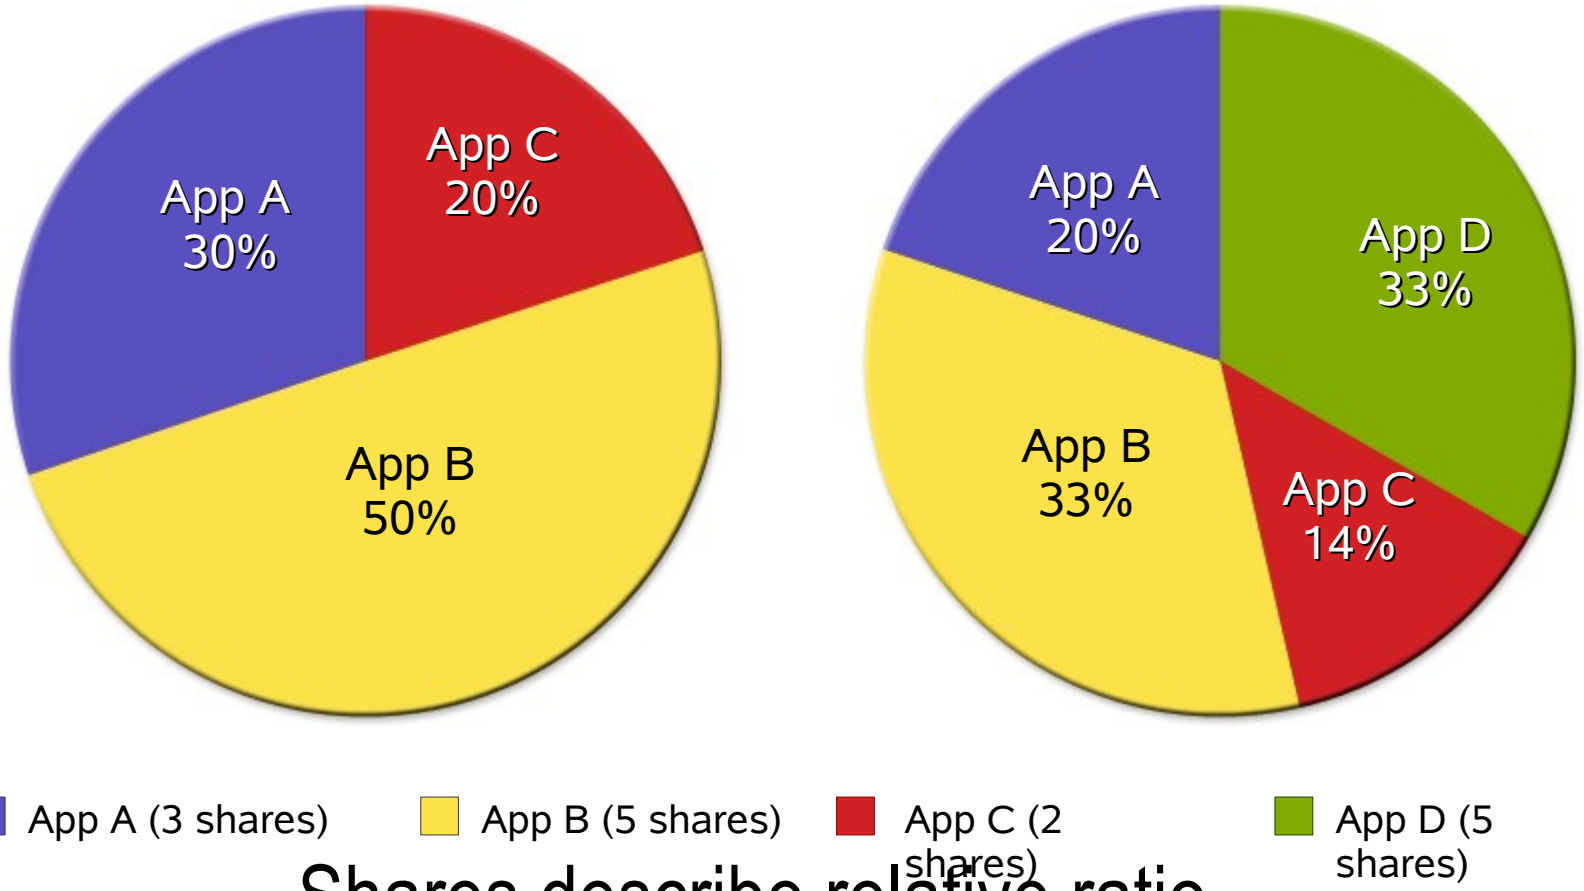

Shares describe relative ratio...

Demo:
Zones with Resource Management

(Graphical, for easy viewing :-))

- Limitless partitioning – one license
- Thousands of applications on one system
- Ultimate consolidation tool
- Container Cloning and Container Migration
- Instant restart

NEW 11/06

Solaris Containers

Solaris Containers

Business Value

- Higher utilization
- Single OS to manage
- Free, no additional license
- Limits required licenses

LOWER TCO

- Resource Controls
- Fault Isolation
- Security Isolation

Predictable Service Levels

Solaris Container

Components

- Full Resource Containment
  - > Provide predictable service levels – Solaris 9
- Security isolation
  - > Prevent unauthorized access – Solaris 10
- Fault isolation
  - > Minimize fault propagation and unplanned downtime – Solaris 10
- Service Management Application
  - > Ease of management – Container Manager

Example 1: Solaris Containers

Single Application Consolidation

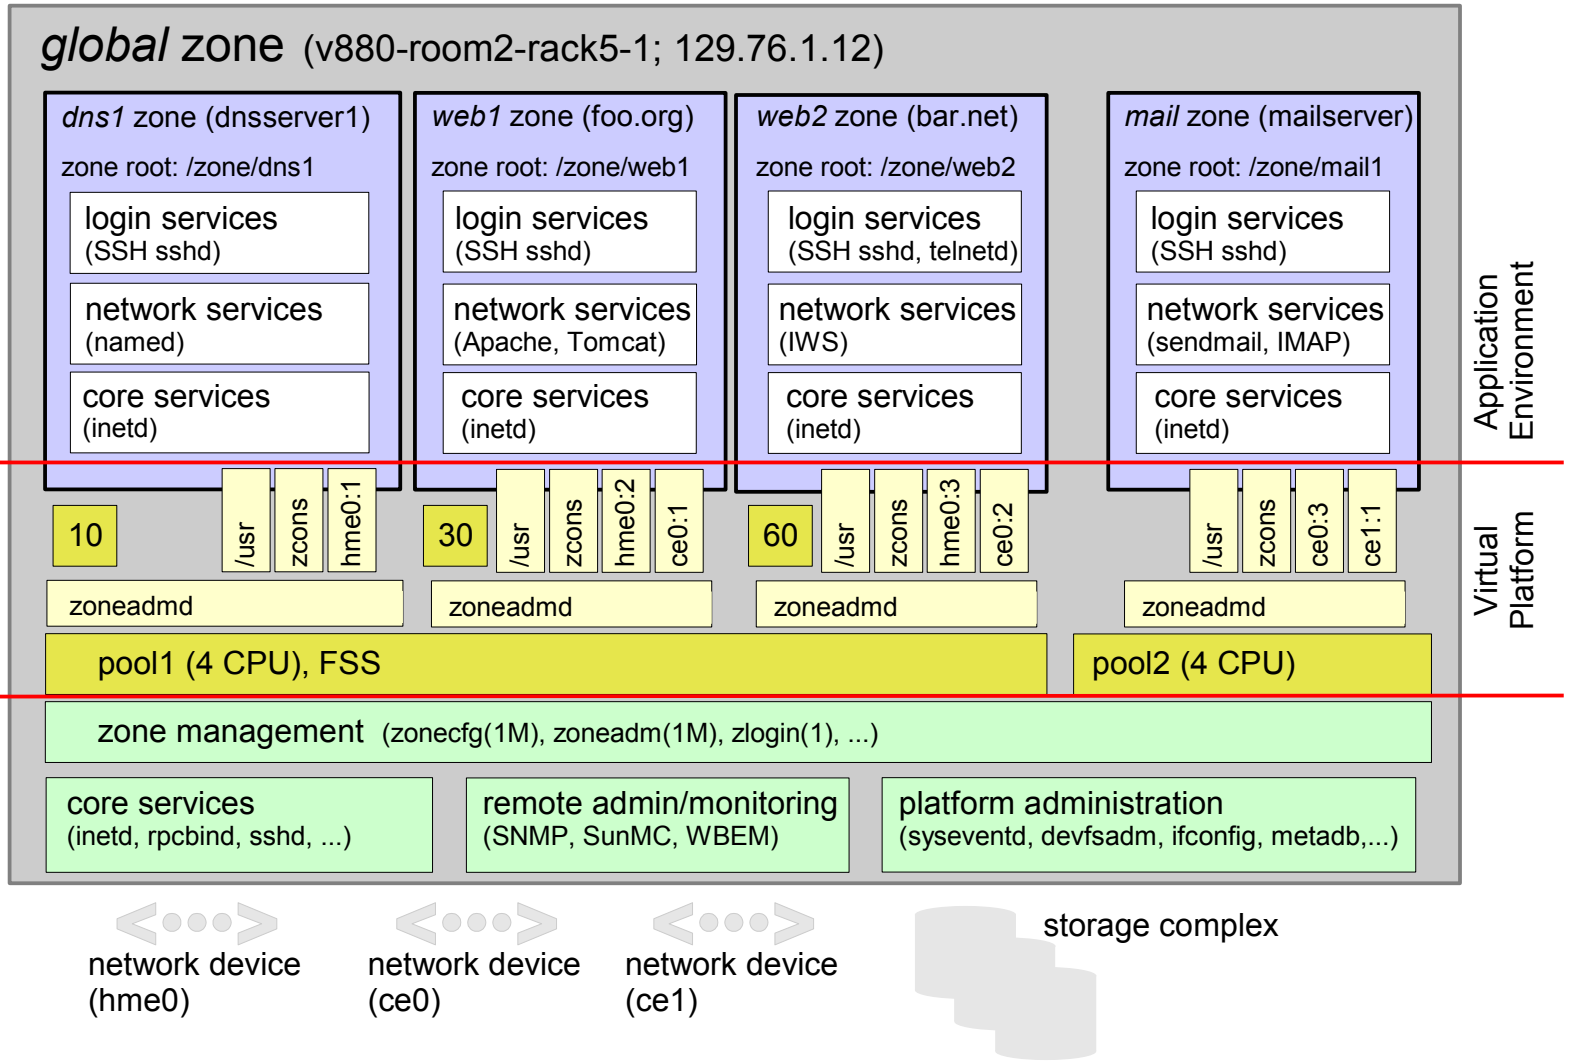

Example 2: Solaris Containers

Test & Development Systems

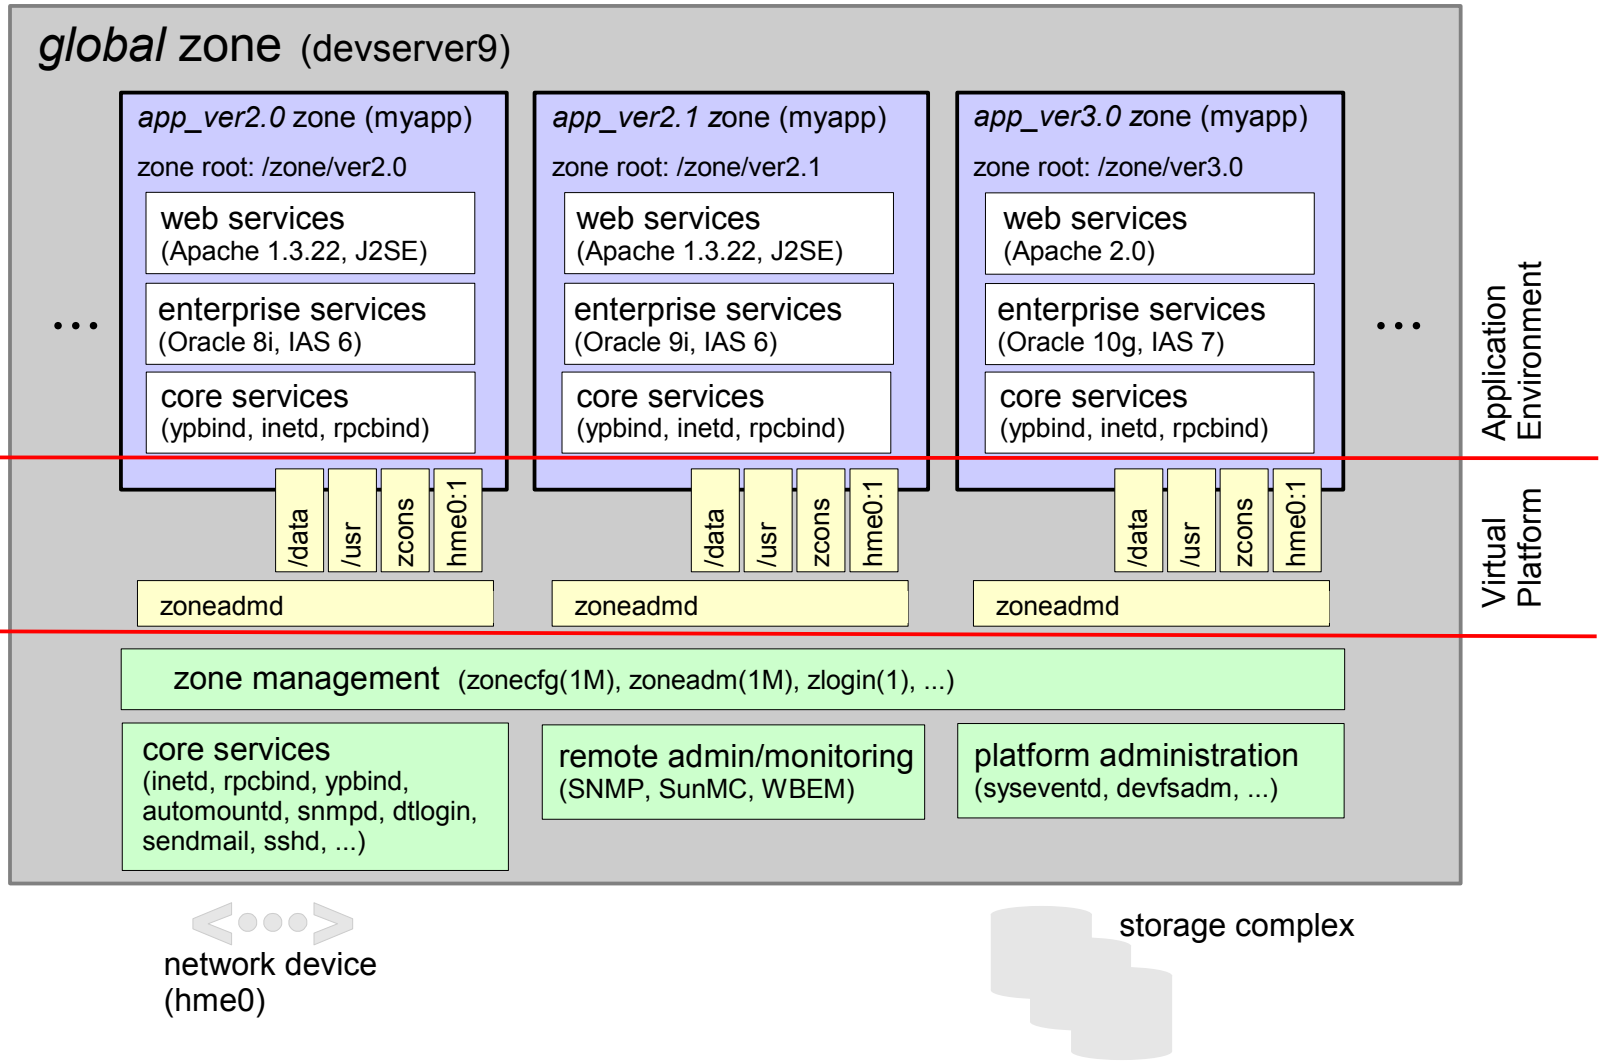

Example 3: Solaris Container

Multi Application Containers

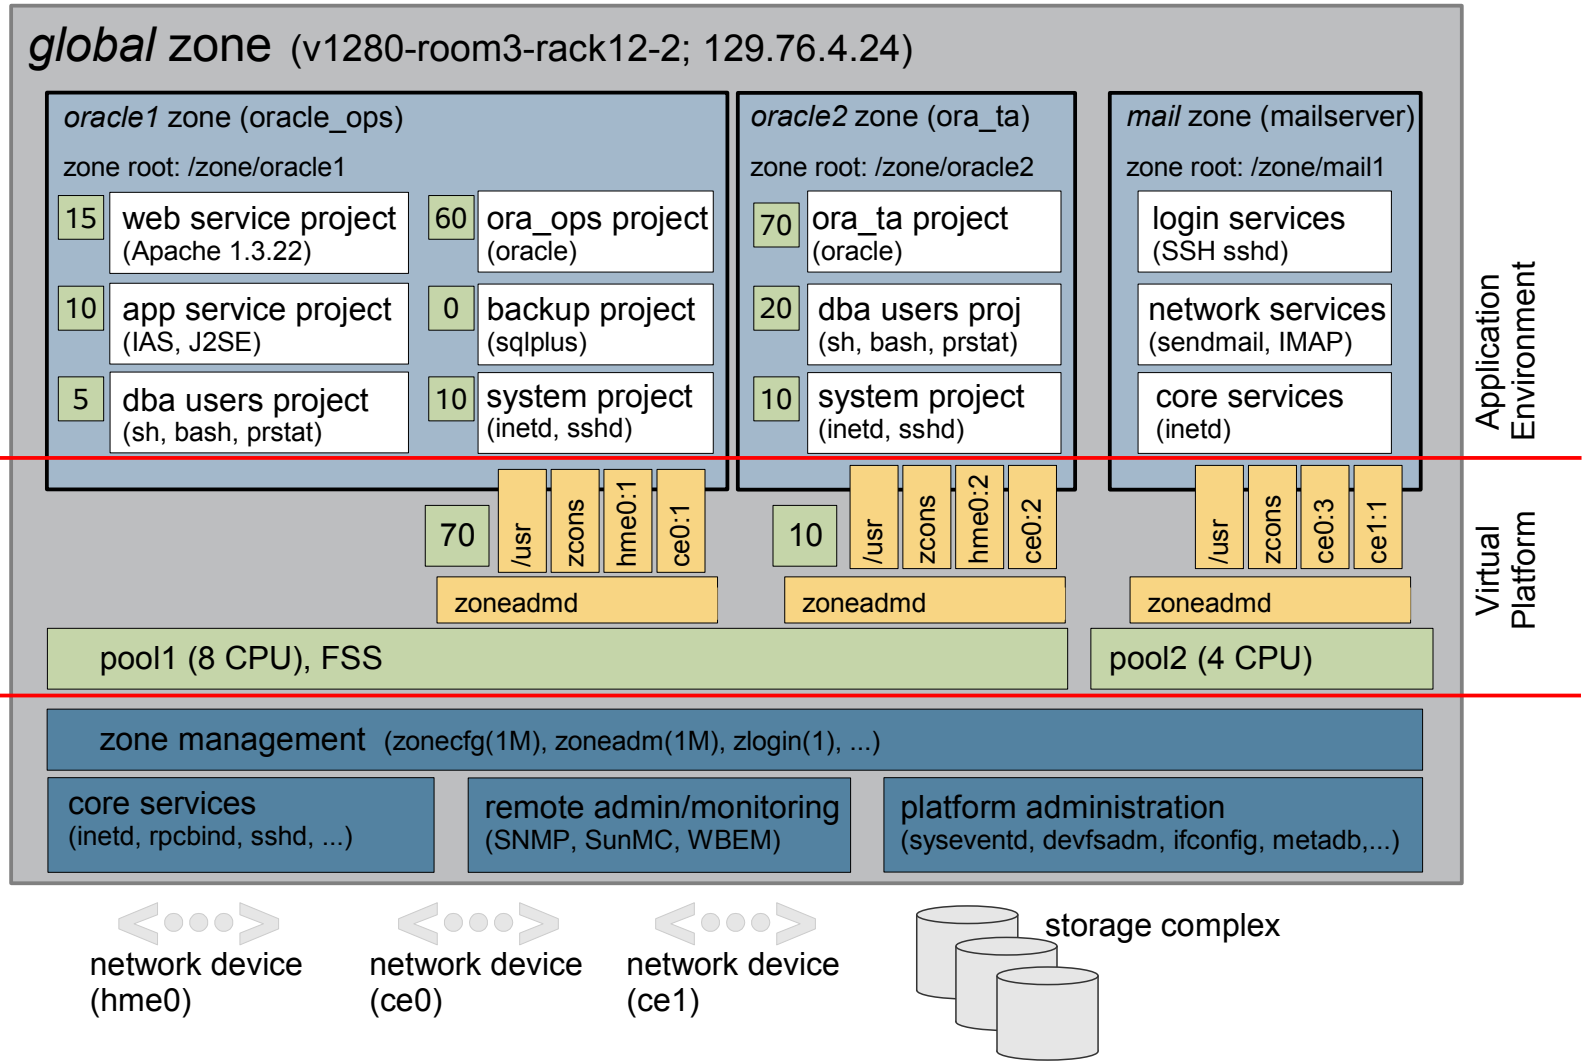

Sun™ MC – Solaris Container Manager

Manage systems that run the Solaris 8, 9, and 10 OS

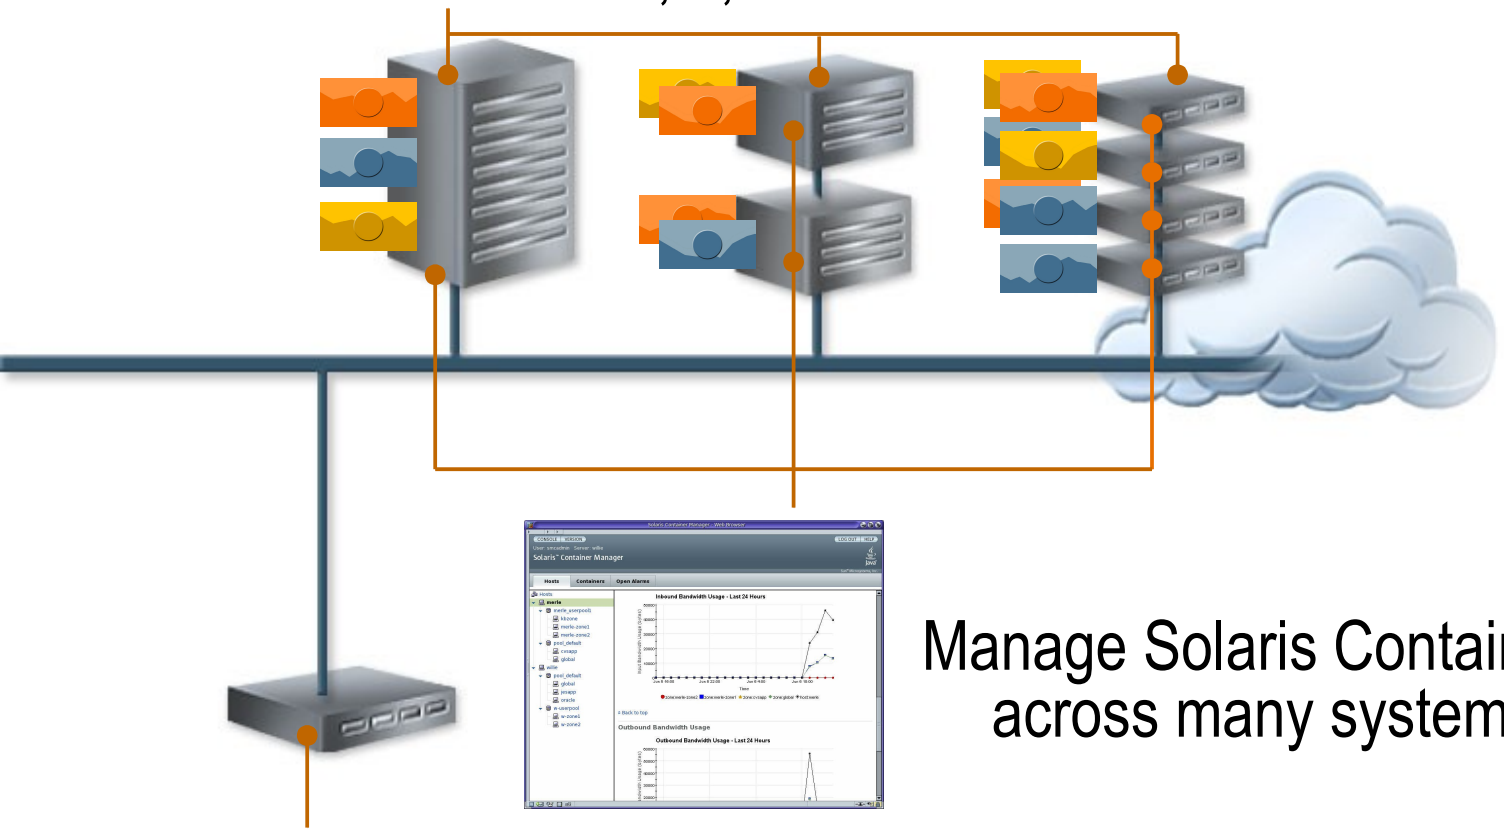

Manage Solaris Containers across many systems

Uses Sun Management Center 3.5 Update 1b

Solaris Container Enhancements

New Management and Security Features

Container Cloning

Container Migration

Per-Container Security Privileges

Simplified Resource Pool Management

Zone Demos

- Creating and booting zones
- Migrating zones
- Cloning zones
- ZFS integration with zones
- Resource Management

Solaris Virtualization: Zones and Containers

Bart Muijzer
bart.muyzer@sun.com

Linux Compatibility

Solaris Application Environment for Linux

- Extension of Solaris 10 Container technology
- Run unmodified Red Hat applications
- Designed to support different “brands”
- Leverages all Solaris security / administrative advantages
- Combines the best of virtualization, resource management and OS flexibility
- Community on opensolaris.org
- Targeted for Solaris 10 update 3

Platform Choice

Run Linux Apps unchanged

- Linux faster than Solaris x86 considered a ....
- Step 1: Linux Application Environment
  - Project Janus
  - Implementation of Linux SysCall interface on Solaris
- Linux Standards Base compliant
- On S10 x86 only
- 100% compatible with RHEL 3
- Needs a valid Linux RTU
  - ... and a Linux CD
- Step 2: branding