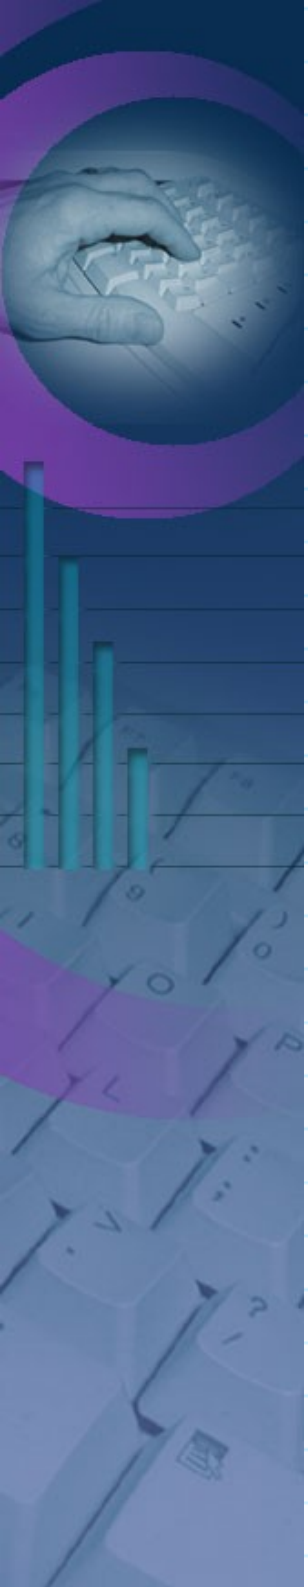

Solaris containers (Zones)

Server virtualization

What zones are and how
they are used in ECE/CIS
at the University of Delaware

www.eecis.udel.edu

Ben Miller

Zones - Containers

Server Consolidation

Run multiple workloads on system

Improve utilization of resources

Reduce costs

Run workloads in isolation

Cannot observe others

Security Isolation

Running apps as different user not enough

- privilege escalation bugs

Traditional Resource Management

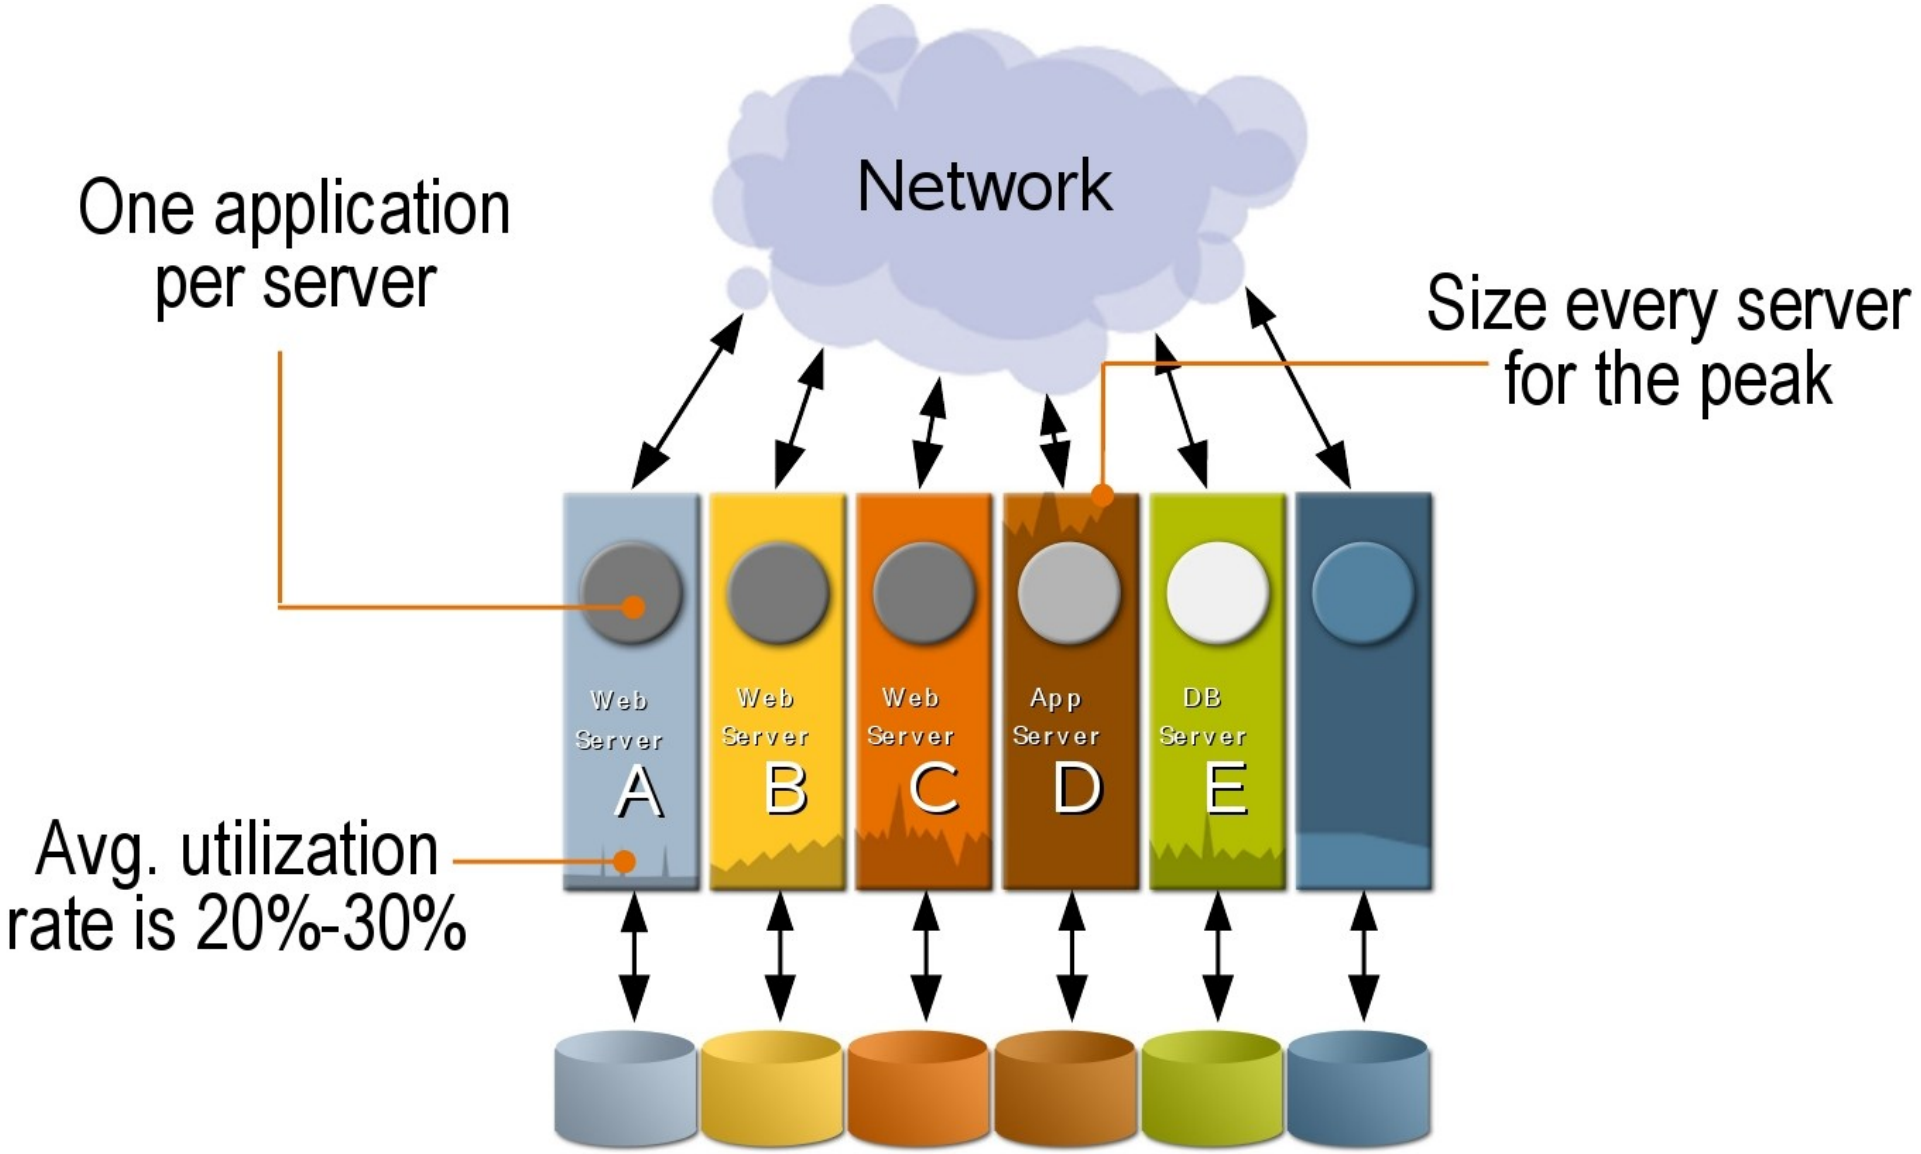

Solaris Zones

Part of Solaris 10 and beyond

Available on sparc and x86/x64

Applications run with no changes

Virtual machine

No significant impact on performance

Little overhead

Zones representation

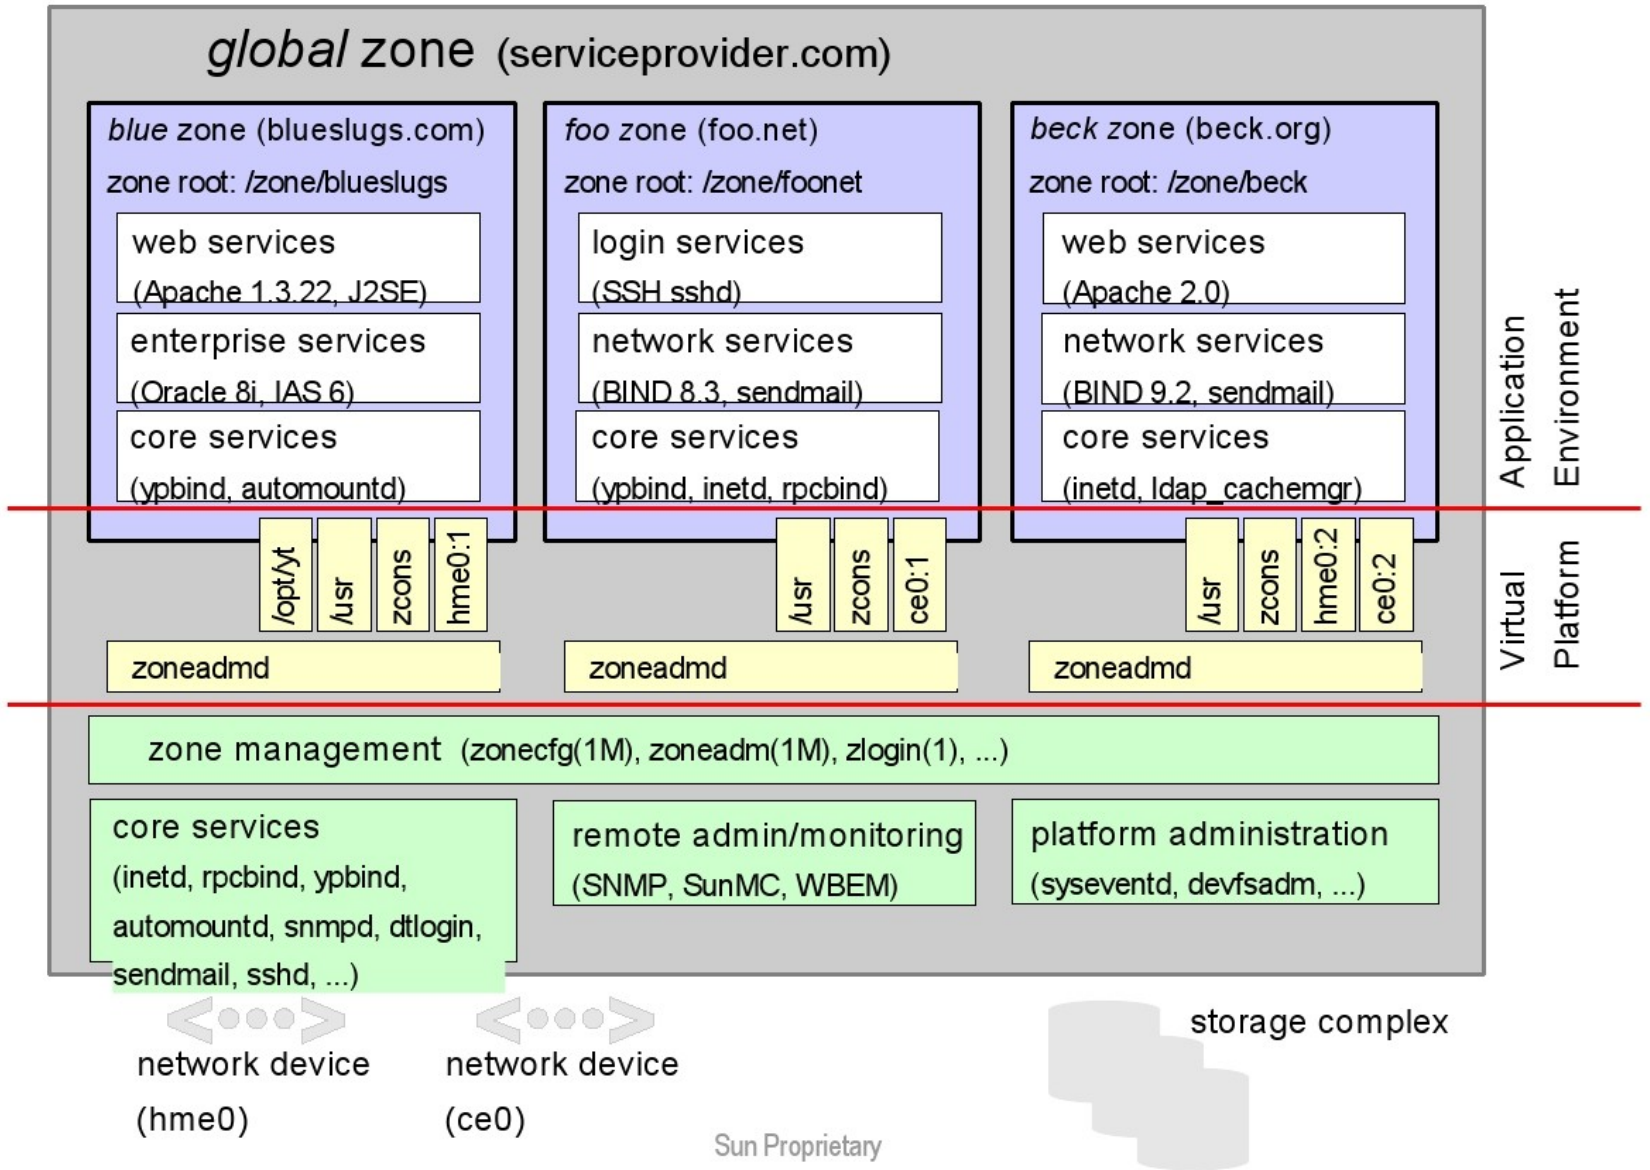

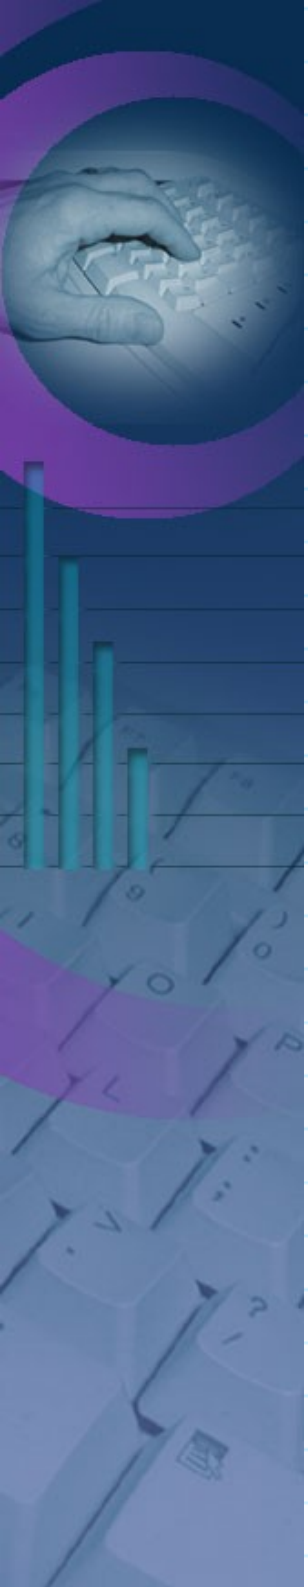

Must solve consolidation problems

Name space isolation
- own hostname, IP, name services

Security isolation

resource allocation

management

Must support commercial applications

A virtual machine

SW in zone should work without change

Admins should not need special scripts

System should look and feel like normal host

Work on single cpu systems and multi cpu systems.

Support several zones on one system.

Address design principles

State model describes life cycle

Configuration engine to describe a zone

Installation support in zone path

Application environment

Virtual platform

Zone States

Configured: configuration done, not installed yet.

Installed: installed according to the configuration.

Ready: zsched process created, initialize network and devices. No user processes yet.

Running: init is created and rest of env can run for apps.

Shutting down: remains in this state until all user processes are destroyed.

Down: remains in this state until virt. Platform completely destroyed. Then go to installed.

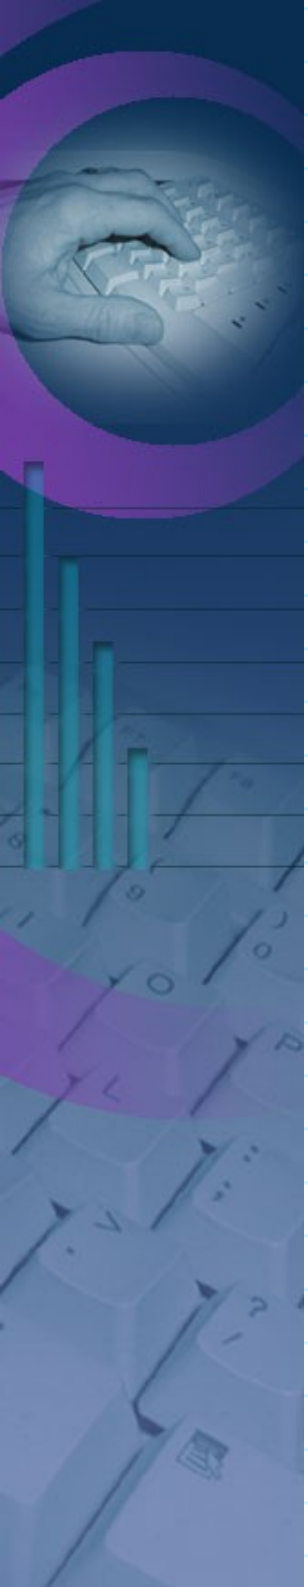

Global zone

The global zone is the default zone

traditional single zone system

Global zone has access to and controls non-global zones

Zone commands

zonecfg - to configure zone info
stored in xml file

zoneadm – used to admin zone
subcommands include install, boot, reboot, halt, shutdown

zlogin – to log into zone
-C option gives access to zone console

-z or -Z options added to commands like ps, prstat and others for use in global.

Resource controls

Can limit amount of cpu used
- limit # of processors used
- use fair share scheduling to limit % used

Can also limit amount of RAM and swap

Placing limits on network usage also planned for the future.

Non-Global zones

A system can have several non-global zones

Each can be running different set of services

non-global zones are isolated – cannot effect other zones (or even observe them)

All zones in a system share resources.

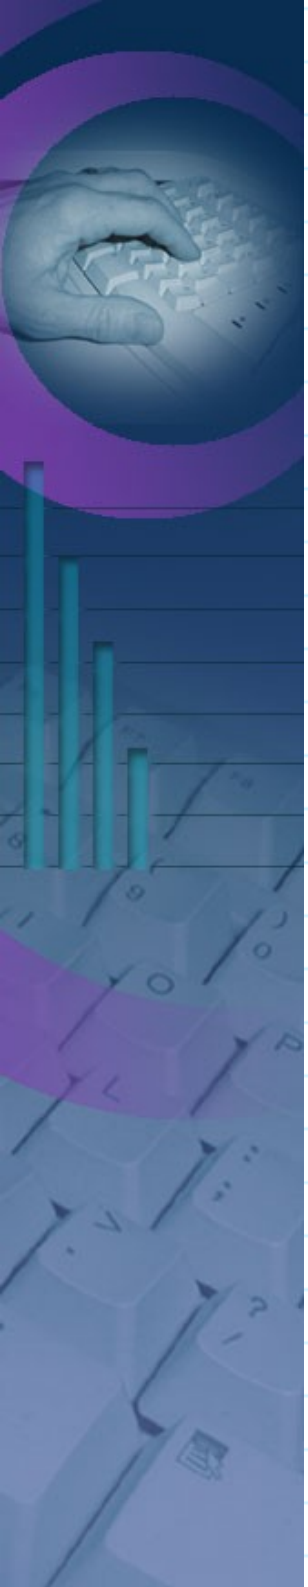

Zone disk usage

Full and sparse zones

zonecfg is given a dir for zone root

read-only lofs mounts used a lot

zones can take very little disk space

By leveraging ZFS more things are possible

Zone Filesystem

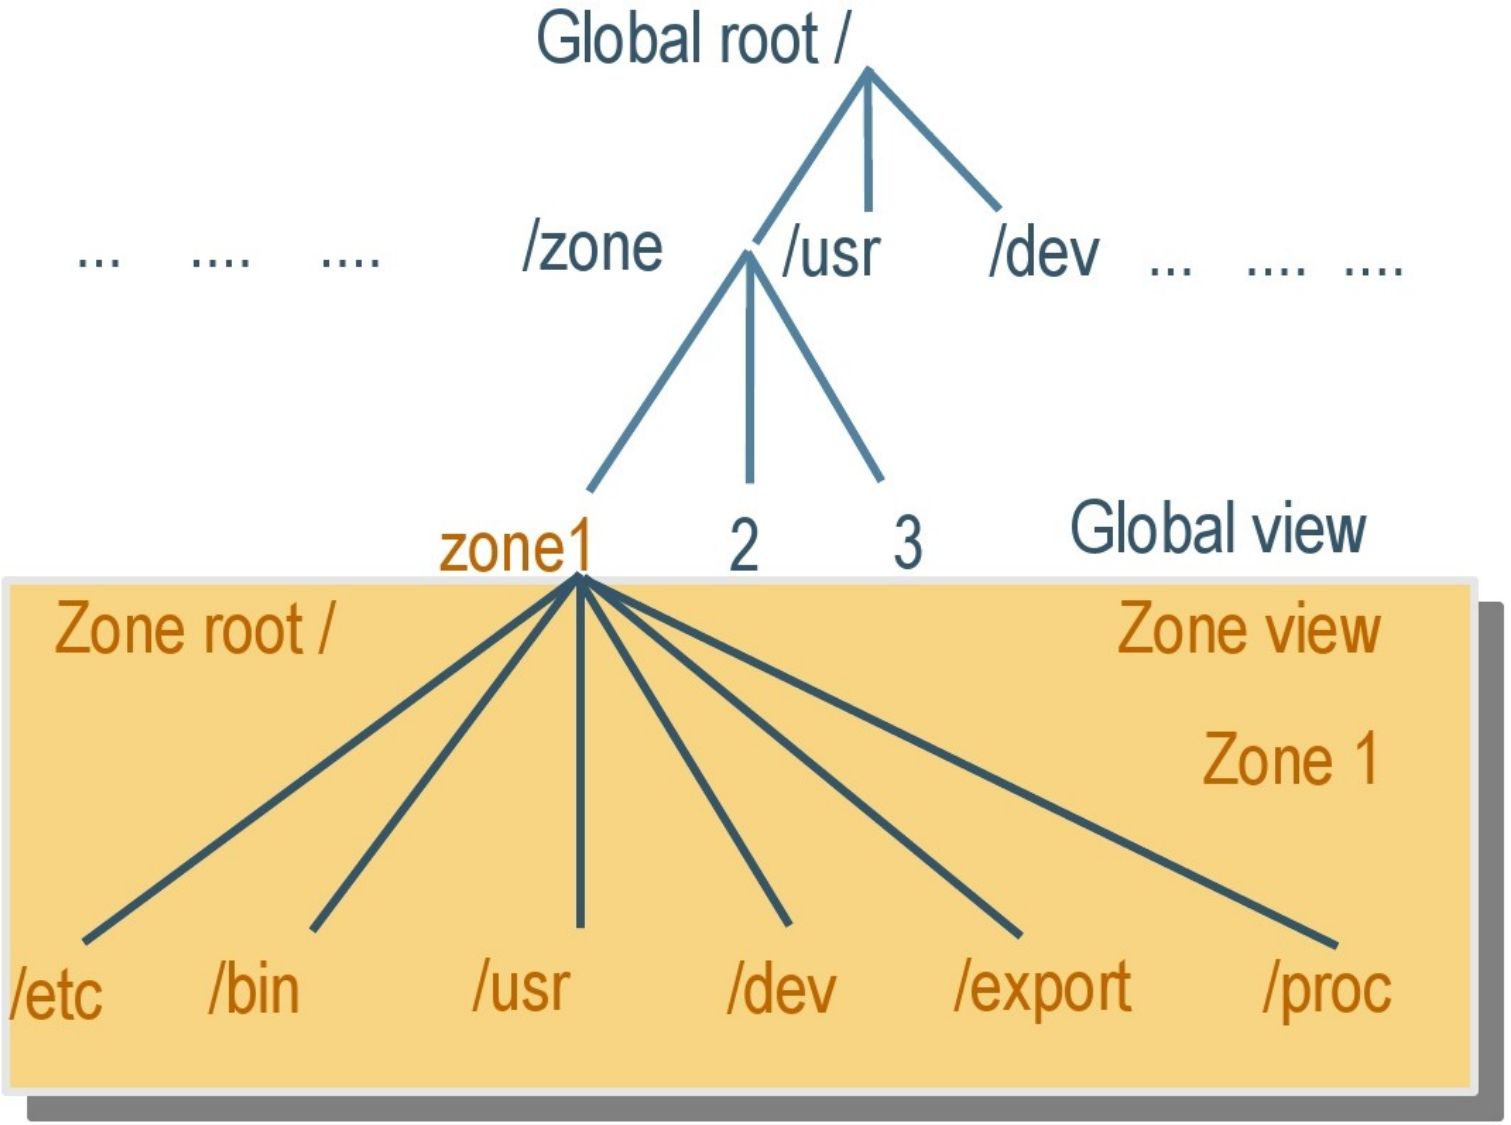

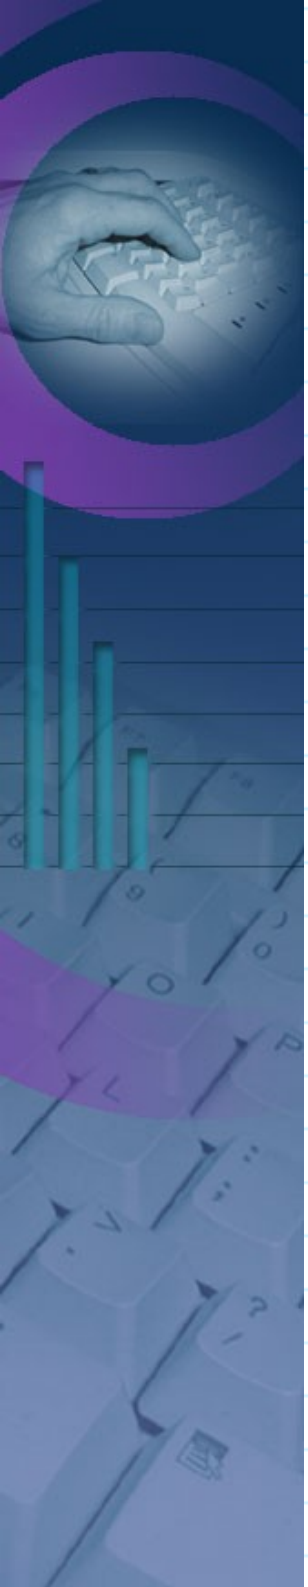

Security considered

As part of the design

devices are limited

only root in global zone can access a zone root filesystem.

A zones hacking contest was held.

Operating systems supported

Solaris 10 and beyond

BrandZ – Currently lx (Linux 2.4.21) is supported. (32bit)
- experimental Linux 2.6 kernel (lacking needed features)

Also Solaris 8 and 9 zones

First use in ECE/CIS

Shortly after Solaris 10 3/05 released

3 old servers upgraded to one
- advstudies, ntp, cgi
- old servers wouldn't run Solaris 10
- used an Ultra 5 as first system
- eventually merged onto new system with several more zones.

More ECE/CIS Examples

Current web hardware is running
26 zones
- eecis, ece, cis web servers
- several research groups
- other web servers for classes, etc
- development zone
- used to host LUG web server
- College of Engineering web servers

Load balance zones
- use zones to create aliases users can use to log into a group.

More ECE/CIS Examples

Samba zones (most on X4540)
- ECE/CIS research
- ECE/CIS academic (classes)
- ECE office
- CIS office

Name services (NIS, DNS, LDAP)
- redundant servers
- separate LDAP zones for samba

More ECE/CIS Examples

Application server/compute servers
- divide up for research/acad
- offer Solaris and Linux* (w/ xen now)
- use another zone to load balance...
- amd64 and sparc hardware used

Linux brandz zones for Matlab, Opnet
- Limited to 2.4 kernel
- 2.6 kernel support lacking
- transitioned to xVM (xen) for 2.6

More ECE/CIS Examples

An X4600
- zone for general research use
- zone where researcher has root
- Solaris and Linux (moved to xen)

Resource controls used
- Fair-Share scheduling (FSS)
- limit memory and swap
- limit processors