

# Chapter1 Solaris Overview

—— *Feature and  
architecture*

**Huimei Lu**

blueboo@bit.edu.cn

# Outline

- **Introduction to Solaris**
- Solaris Kernel Features
- Solaris Kernel Architecture
- Solaris 10 Features
- Performance and Tracing Tools

# Introduction

## □ What is Solaris?

- Sun Microsystems, Inc.
- A complete operating environment, built on a modular, dynamic kernel

## □ SOE- Solaris Operating Environment

- 3 major components:
  - > SunOS- the kernel
  - > Windowing – desktop environment, CDE default, OpenWindows still included. Gnome, KDE
  - > Open Networking Computing

# A brief history

- BSD UNIX
- 1982, Sun UNIX 0.7
- 1983 SunOS 1.0
- 1992 SunOS 4.1.3
- 1992 Solaris 2.0
- 1998 Solaris 7
- 2000 Solaris 8
- 2001 Solaris 9
- 2005 Solaris 10

# Outline

- Introduction to Solaris
- **Solaris Kernel Features**
- Solaris Kernel Architecture
- Solaris 10 Features
- Performance and Tracing Tools

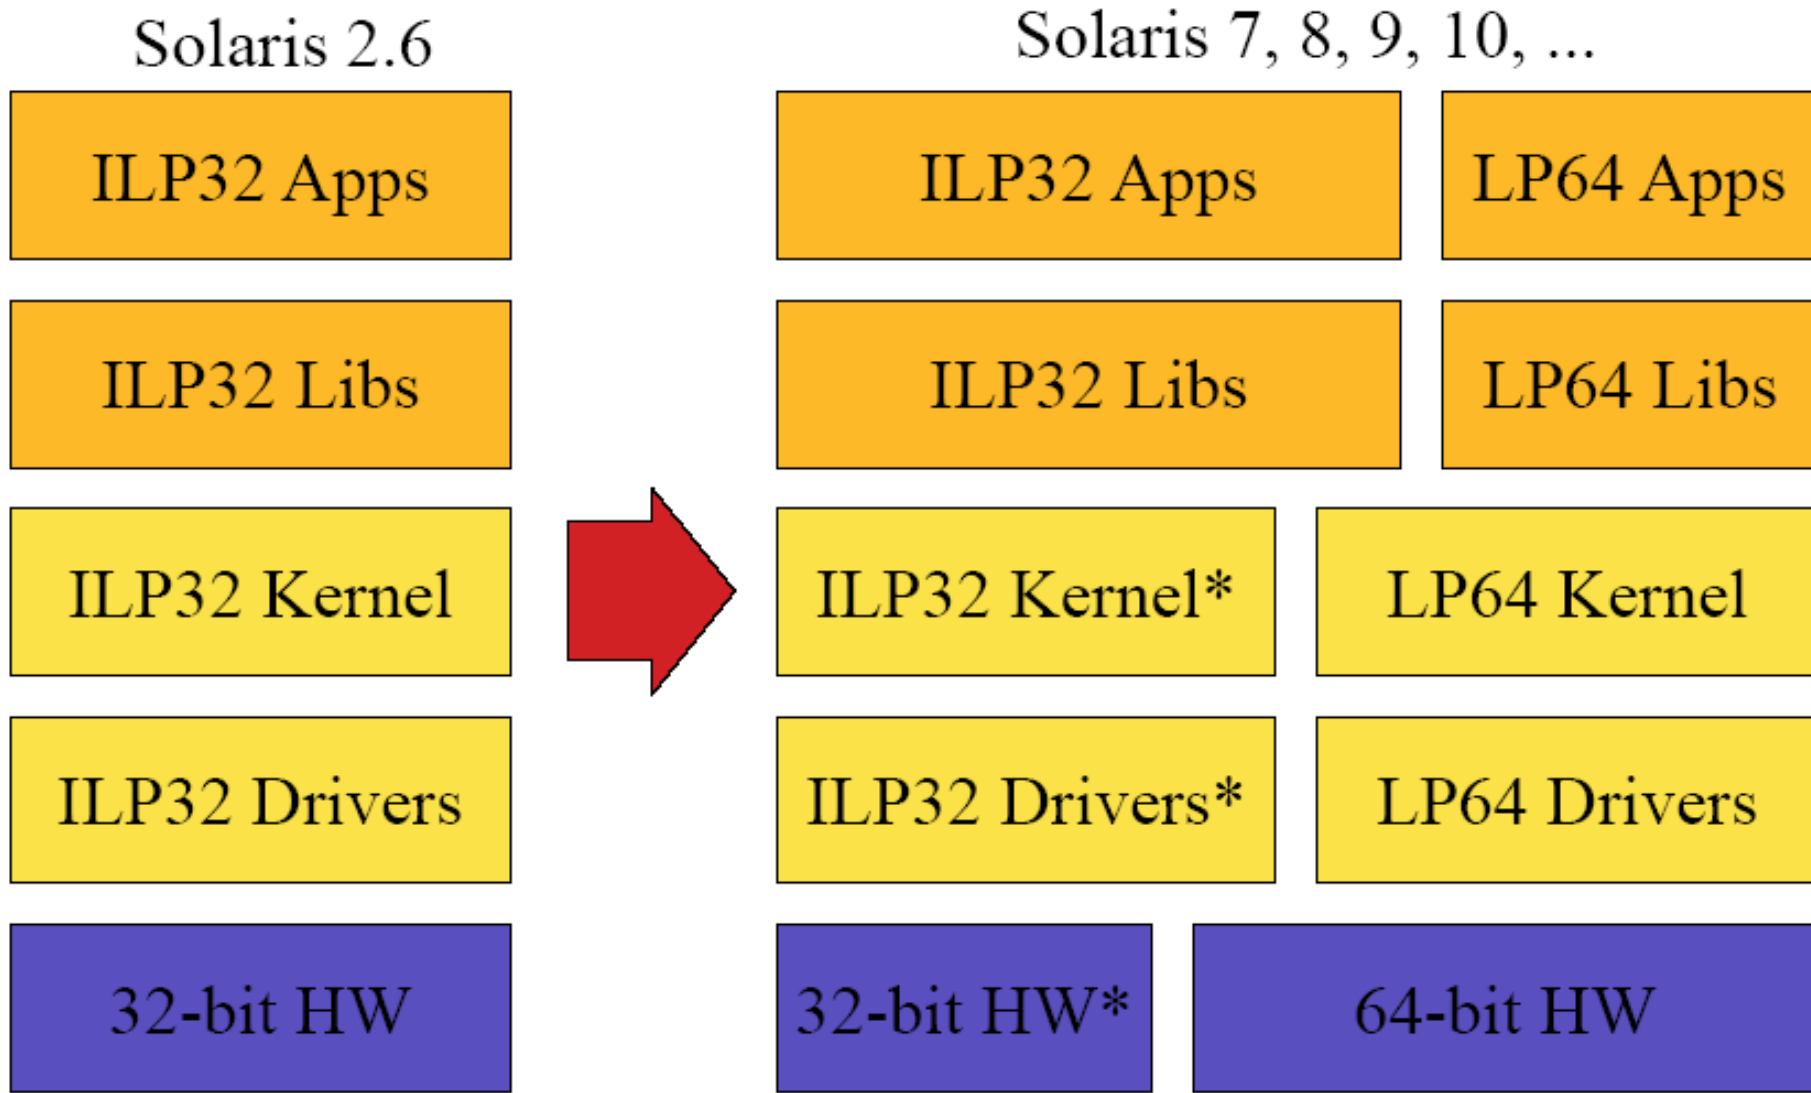
# Solaris Kernel Features

- ❑ Dynamic
- ❑ Multithreaded kernel
- ❑ Preemptive kernel
- ❑ Multithreaded Process Model
- ❑ Multiple Scheduling Classes
  - Including real-time support

# Solaris Kernel Features (Cont'd)

- ❑ Tightly Integrated File System & Virtual Memory
- ❑ Virtual File System
- ❑ 64-bit kernel
  - 32-bit and 64-bit application support
- ❑ Resource Management
- ❑ Service Management & Fault Handling
- ❑ Integrated Networking

# The 64-bit Revolution



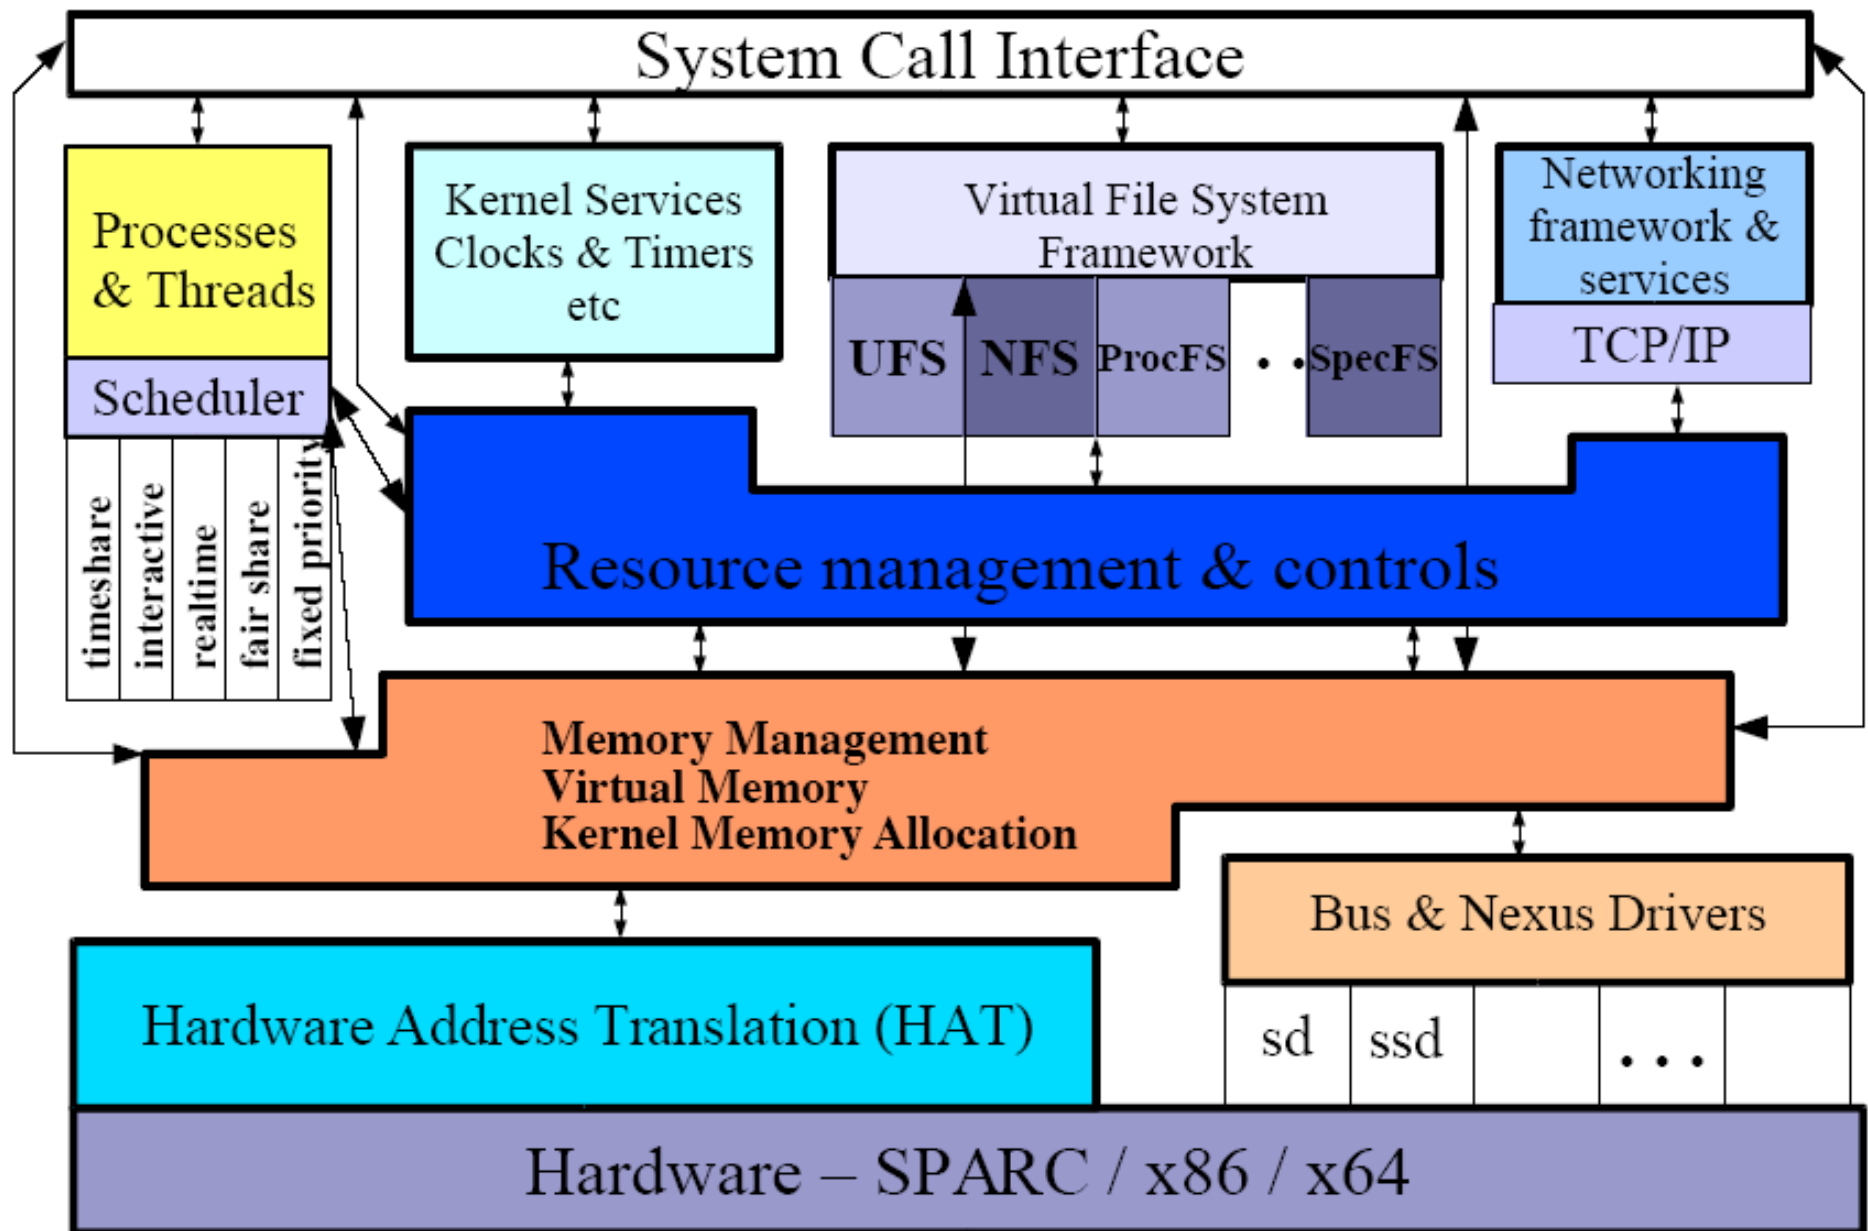
\* Solaris 10: 64-bit kernel only on SPARC



# Outline

- Introduction to Solaris
- Solaris Kernel Features
- **Solaris Kernel Architecture**
- Solaris 10 Features
- Performance and Tracing Tools

# Solaris kernel Architecture



# Outline

- Introduction to Solaris
- Solaris Kernel Features
- Solaris Kernel Architecture
- **Solaris 10 Features**
- Performance and Tracing Tools

# Solaris 10 (The headline grabbers)

- ❑ Solaris Containers (Zones)
- ❑ Solaris Dynamic Tracing (dtrace)
- ❑ Predictive Self Healing
  - System Management Framework
  - Fault Management Architecture
- ❑ Process Rights Management
- ❑ Premier x86 support
- ❑ Optimized 64-bit Opteron support (x64)
- ❑ Zetabyte File system (ZFS)
- ❑ ... and much, much more!

# Outline

- Introduction to Solaris
- Solaris Kernel Features
- Solaris Kernel Architecture
- Solaris 10 Features
- **Performance and Tracing Tools**

# Performance and tracing Tools

## □ Process stats

- cputrack - per-processor hw counters
- pargs – process arguments
- pflags – process flags
- pcred – process credentials
- pldd – process's library dependencies
- psig – process signal disposition
- pstack – process stack dump
- pmap – process memory map
- pfiles – open files and names
- prstat – process statistics
- ptree – process tree
- ptime – process microstate times
- pwdx – process working directory

## □ Process control

- pgrep – grep for processes
- pkill – kill processes list
- pstop – stop processes
- prun – start processes
- prctl – view/set process resources
- pwait – wait for process
- preap – reap a zombie process

# Performance and tracing Tools (Cont'd)

- Process tracing/debugging
  - abitrace – trace ABI interfaces
  - dtrace – trace the world
  - mdb – debug/control processes
  - truss – trace functions and system calls
- Kernel tracing/debugging
  - dtrace – trace and monitor kernel
  - lockstat – monitor locking statistics
  - lockstat -k – profile kernel
  - mdb – debug live and kernel cores

# Performance and tracing Tools (Cont'd)

## □ System stats

- acctcom – process accounting
- busstat – Bus hardware counters
- cpustat – CPU hardware counters
- iostat – IO & NFS statistics
- kstat – display kernel statistics
- mpstat – processor statistics
- netstat – network statistics
- nfsstat – nfs server stats
- sar – kitchen sink utility
- vmstat – virtual memory stats



# Dtrace – Dynamic Tracing

- ❑ Seamless, *global* view of the system from user-level thread to kernel
- ❑ Not reliant on pre-determined trace points, but *dynamic instrumentation*
- ❑ Data *aggregation* at source minimizes post processing requirements
- ❑ Built for live use on *production* systems

## Dtrace (Cont'd)

- Ease-of-use and instant gratification engenders serious hypothesis testing
- Instrumentation directed by high-level control language (not unlike AWK or C) for easy scripting and command line use
- Comprehensive probe coverage and powerful data
- management allow for concise answers to arbitrary questions

# Modular Debugger - mdb(1)

- ❑ Solaris 8 mdb(1) replaces adb(1) and crash(1M)
- ❑ Allows for examining a live, running system, as well as post-mortem (dump) analysis
- ❑ Solaris 9 mdb(1) adds...
  - Extensive support for debugging of processes
  - /etc/crash and adb removed
  - Symbol information via compressed typed data
  - Documentation
- ❑ MDB Developers Guide
  - mdb implements a rich API set for writing custom dcmds
  - Provides a framework for kernel code developers to integrate with mdb(1)

# References

- Jim Mauro, Richard McDougall. Solaris Internals-Core Kernel Components. Sun Microsystems Press, 2000
- Richard McDougall, Jim Mauro, SOLARIS™ Kernel Performance, Observability & Debugging. Tutorial at UNENIX'05.
- Solaris 10 What's New, Sun Microsystems Press, December 2005

End

- Last.first@Sun.COM