



# The Solaris Operating System

Larry Wake  
Product Line Manager  
Sun Microsystems, Inc.



# Sun Value Proposition

**Applications**

**Infrastructure  
Services**

**Operating  
System**

**Hardware**

**Support**

- Open, integrated system
- Innovation leadership
- Industry's leading OS



# Three Key Assumptions

**Applications**

**Infrastructure  
Services**

**Operating  
System**

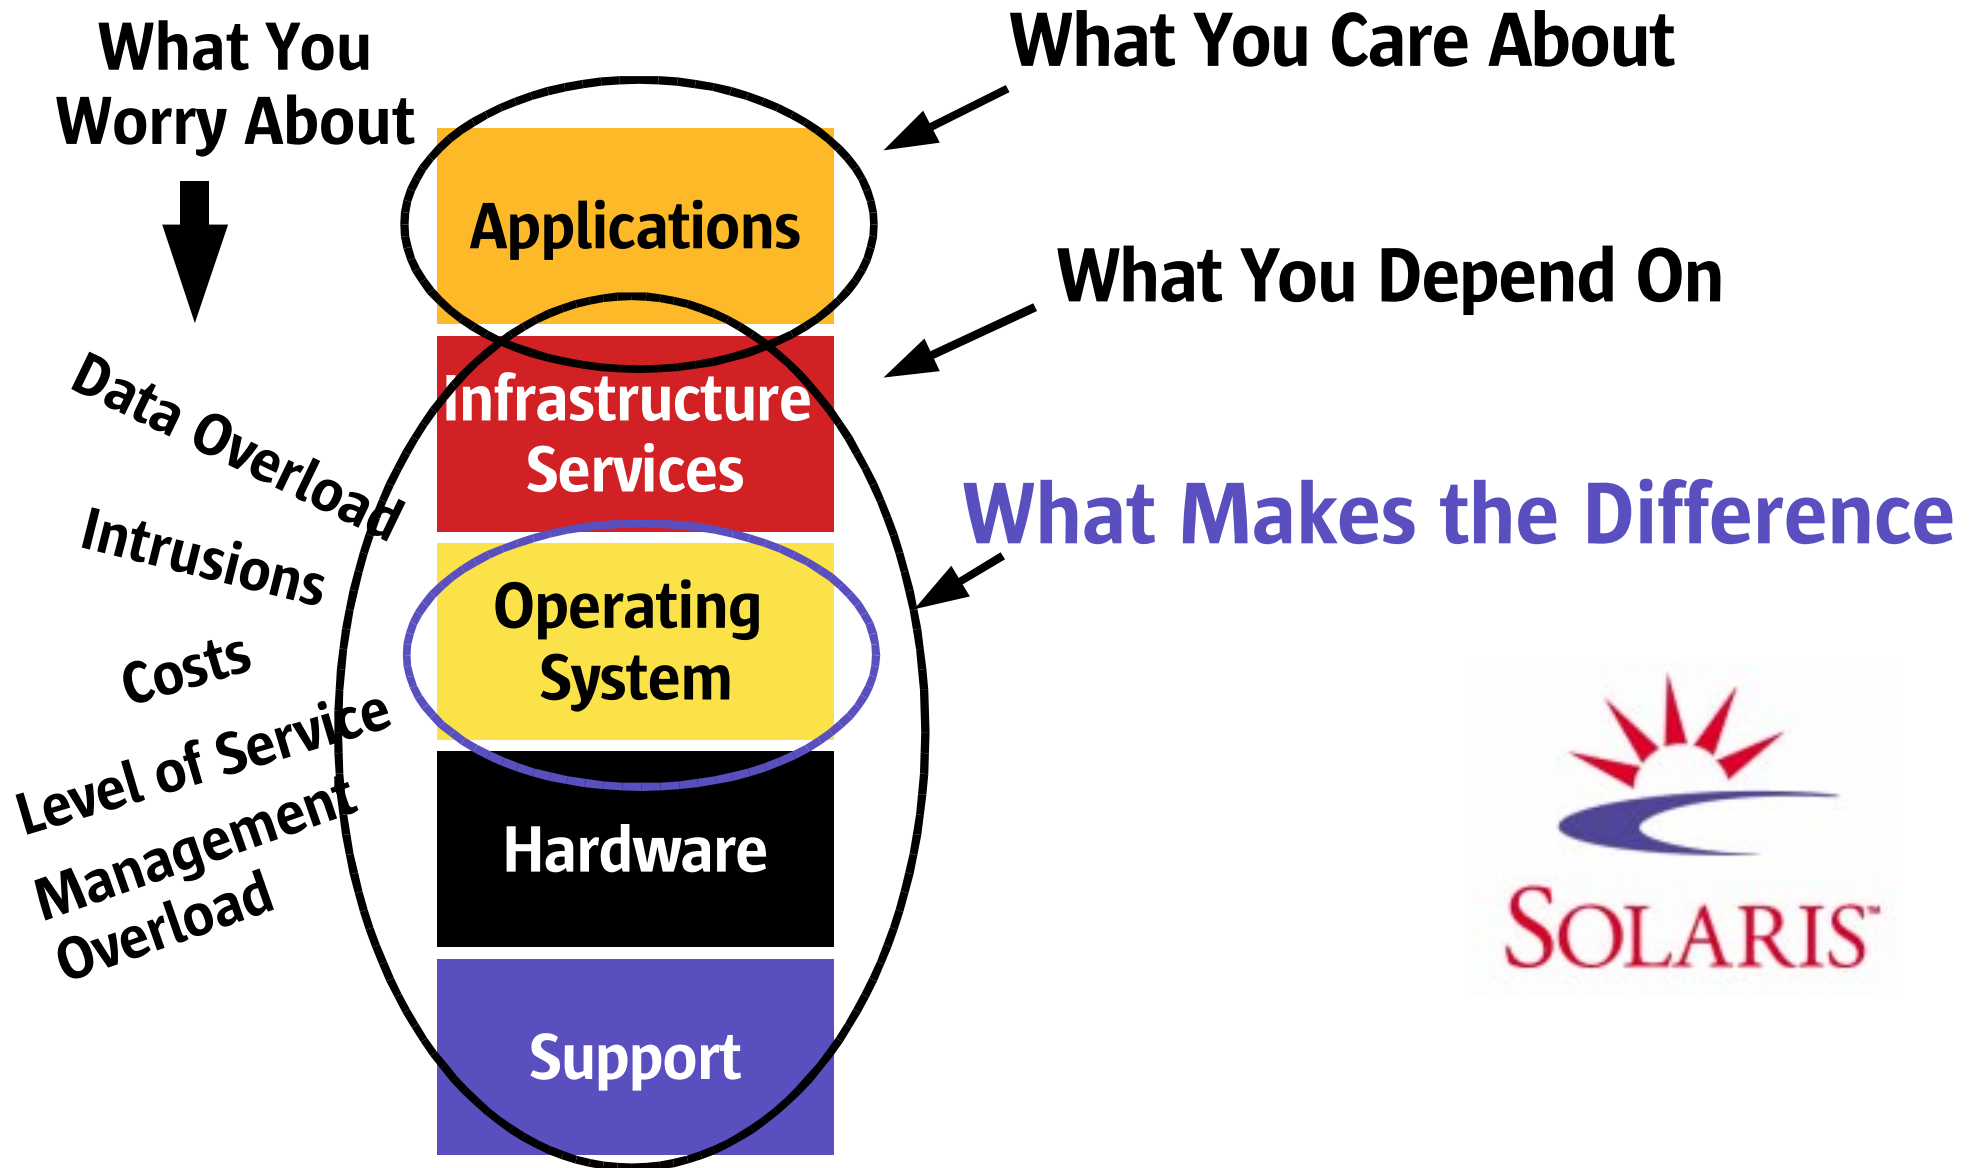
**Hardware**

**Support**

- Open is good
- Innovation pays
- The OS matters



# Why the OS Matters



# The Solaris Operating System

## UNIX Leadership Today

### Drive Down Cost

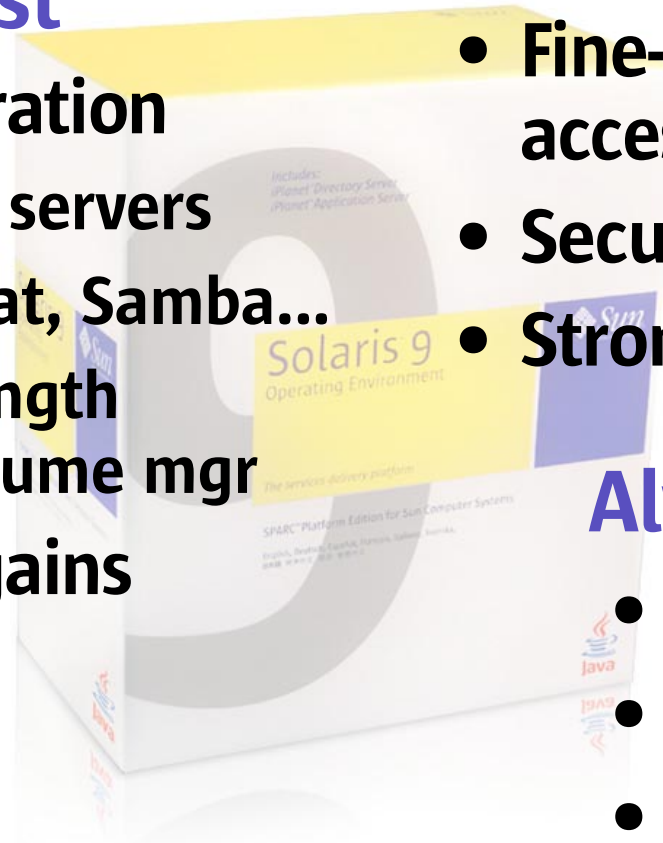
- **Services integration**
  - App, directory servers
  - Apache, Tomcat, Samba...
  - Industrial-strength filesystem/volume mgr
- **Performance gains**
  - Threads
  - Memory
  - Datapath
- **Guaranteed compatibility**

### Security Everywhere

- **Fine-grained access control**
- **Secure remote access**
- **Strong authentication**

### Always Available

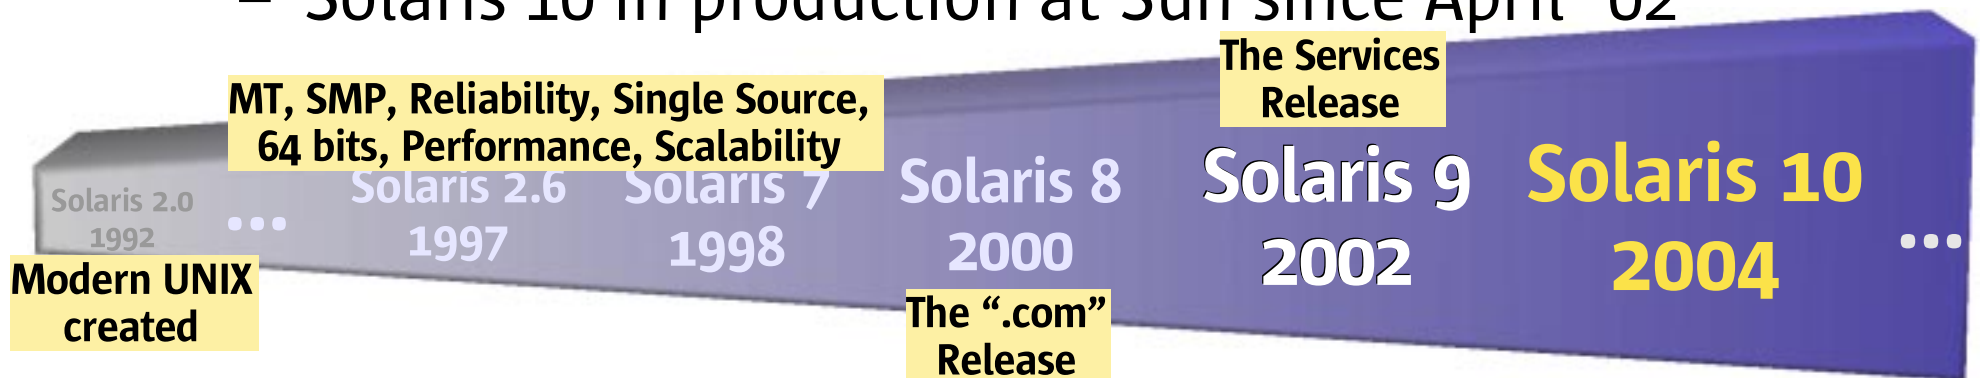
- **Patch manager**
- **Sun Fire RAS**
- **Solaris Provisioning**



# Solaris Development Model

Compatibility, Stability, Performance – by Design

- Product support lifecycle designed for customers' needs
  - Releases extend, rather than replace
  - Guaranteed compatibility
- Extensive testing
  - Continuous, measurable quality gains
  - Solaris 10 in production at Sun since April '02



# Solaris 10: A Generation Ahead

## Extreme

### Performance

*Dynamic Tracing  
Network  
Entry systems*

### Unparalleled Security

*Process Rights Management  
Crypto Infrastructure*

### Optimal Utilization

*N1 Grid Containers*

### Relentless Availability

*Predictive Self Healing  
Next-gen file services*

### Platform Choice

*New UltraSPARC IV,  
New AMD Opteron*



# SOLARIS™

# Solaris 10 Performance

- Observability
  - DTrace, libMicro
- Platform optimization
- Network
- Latency
- Scalability
- Throughput



# Observability: Today

- Who cares?
  - Developers
  - System administrators
  - Resource stakeholders
- How do we observe problem systems?
  - Forced crash dumps
  - Replicated environment
  - Instrumented system / application
- Results:
  - Unhappy users
  - Overburdened administrators
  - Slow or unavailable services

# Dynamic Tracing

## Real-time Analysis and Diagnosis

- Comprehensive; safe
  - Global view into systems, apps
  - Always there
  - No app/OS change
  - No performance hit
- Extensible; easy to leverage
  - Scriptable for re-use
  - Provider model allows for new probes
- Debug, analyze, optimize
  - Solutions in minutes or hours, not days or weeks
  - Cases of 3-30x speedups already seen



“...it's like they *saw inside my head* and gave me The One True Tool.”

–Slashdot post, November '03

# Extreme System Performance

## Focus on Latency Reduction

- Portable microbenchmarks (libMicro) used to compare, tune OS performance

- Result: faster syscalls (25%+):

3-5x  
faster

~~dup, fcntl, flock, getsockname,~~

7x faster

~~getpeername, gettimeofday, lseek,~~

7x  
faster

select, semop, setcontext, setsockopt,

sigaction, siglongjmp, signal,

sigprocmask, socket, time, times

- ...faster library functions (400%+):

31x faster!

~~strftime, mktime, localtime, getenv, SPARC str\*~~



# Extreme System Performance

...But Even More Scalability, Too

- Per-process TSBs, etc.
  - Handle 100k+ busy processes
- Memory management
  - fsflush, fsync, pager, use of idle CPUs, etc.
  - ISM attach/detach
- Event Ports
  - Unifies poll, async I/O, timers, messaging
  - Handling 20,000 connections is now cheap
- libumem: scalable fast memory allocator
- MPO enhancement: liblgrp
  - 8% database performance gain seen
- More OS self-tuning



# Leading the Way for Throughput Computing

- MT-hot kernel, Java VM, system libraries
- Fast thread library
- 64 bit memory addressing
- Memory Placement Optimization
- N1 Grid Containers
- Development & tuning tools



*Ready for multi-threaded applications*  
*Ready for multiple applications*

# New Network Technologies

- NFSv4
- SCTP
- Infiniband
- BIND 9
- Zebra (OSPF, RIPv2, BGP)
- SIP
- AAA (Diameter client libraries)

# N1 Grid Containers

Consolidation Made Simple, Safe, Secure

- Run multiple applications on one system
- “Shrink wraps” applications, isolating from:
  - Faults
  - Intrusion
  - Resource contention
- No performance hit



# Unparalleled Security

Over 20 Years of Design, Testing, Refinement and Experience

- Administrative
  - Secure out of the box
  - User rights management
  - Containers
- Application
  - Process rights management
  - Cryptographic framework
- Network
  - IP filtering



# Solaris Security

## Process Rights Management

- UID 0 (root user) has no “inalienable rights”
  - But ships with traditional root privileges enabled for compatibility
- Granularity: 40+ privileges associated with users, processes
  - SetUID becomes deprecated practice
- Privilege inheritance is selectable
- Full compatibility with previous environments, applications



# Solaris Security

## Solaris Cryptographic Framework

- Cryptographic services for applications
- Central administration of crypto service providers, security policies
- Plug-in algorithm support
  - AES, DES/3DES, RC4, MD5, SHA-1, DSA, RSA, Diffie-Hellman
- Hardware acceleration support



# Solaris Dynamic File System

## Simple, Reliable, Infinitely Scalable

- Streamlined system administration
  - Efficient resource allocation via storage pools
  - Automates administrative tasks
  - Extensible: add features such as encryption
- Self-healing data
- Virtually unlimited capacity
  - 16 *billion billion* times greater than today
- Breakthrough performance



# Predictive Self Healing

## Solaris Fault Manager

- Automated error handling
  - Comprehensive architecture to detect, aggregate, diagnose, report and mitigate faults
- Reduced downtime
  - Components proactively offlined before failure
- Reduced complexity
  - Simplified error reporting
- Reduced costs
  - 24 x forever uptime, increased utilization
  - Higher server-to-administrator ratio



# Predictive Self Healing

## Solaris Services Manager

- Define dependencies / relationships among applications, Solaris components
- Simplified administration
  - Consolidates a system's “application profile”
  - Tools for component management
- Easy service/application installation
- Increased service reliability
  - Fast, automatic service outage detection
  - Fast and accurate service recovery



# Software Express for Solaris: 7/04

## Optimal Utilization

*N1 Grid Containers*

## Extreme Performance

*Dynamic Tracing*

*Network*

*Entry systems*

*NFSv4*

## Unparalleled Security

*Process Rights Management*

*Crypto Infrastructure*

*IP Filter*

## Relentless Availability

*Solaris Fault Manager*

*Solaris Services Mgr*

*Next-gen filesystem*



**SOLARIS™**

## Platform Choice

*New UltraSPARC IV,*

*New AMD Opteron*

*Linux Compatibility*

## Three Key Messages

**"PS: All you guys at Sun,  
all I can say is WOW!!!!  
Solaris10 is un-[...]-  
believable!!!!!! I feel  
like buying a round of  
drinks for the whole  
staff.... WOW!"**

*—Zones discussion board,  
BigAdmin, February 28, 2004*



# The Solaris Operating System

[larry.wake@sun.com](mailto:larry.wake@sun.com)

