



Solaris 10
Something for Free and More!!!

Bob Netherton
Sun Microsystems

Solaris 10 Something for Free and More!!!

Bob Netherton

Solaris Adoption, US Client Solutions

Sun Microsystems

SOLARIS 10 Boot Camp

June 6, 2006

- 8:30 - 9:00 Continental Break and Registration**
- 9:00 - 9:10 Welcome & Introductions**
- 9:10 - 10:15 Solaris Update**
Solaris 10 Patching
- 10:15 - 10:30 Break**
- 10:30 - 11:30 Solaris Zones – Use cases**
- 11:30 - 12:00 Lunch**
- 12:00 – 12:45 DTrace performance scenarios**
- 12:45 - 1:15 ZFS Overview and configuration examples**
- 1:15 - 1:45 Migrating a Legacy RC Service**
- 1:45 - 2:00 Q&A**

<http://blogs.sun.com/bobn>

Sun's Value Proposition

Open, Integrated Systems

Innovation Leadership

Industry's Leading OS


Three Key Assumptions

Open is Good


Innovation Pays

**The OS
Matters**

Why the OS Matters



Why the OS Matters



Why the OS Matters



"I think operating system research at the kernel level was done in the the 1960s and 1970s, and there really isn't much point anymore. The "Big Things" aren't on a kernel level."

**Linus Torvalds
Linux Magazine, Jan 2005**

A Multi-Platform OS Strategy



A Multi-Platform OS Strategy



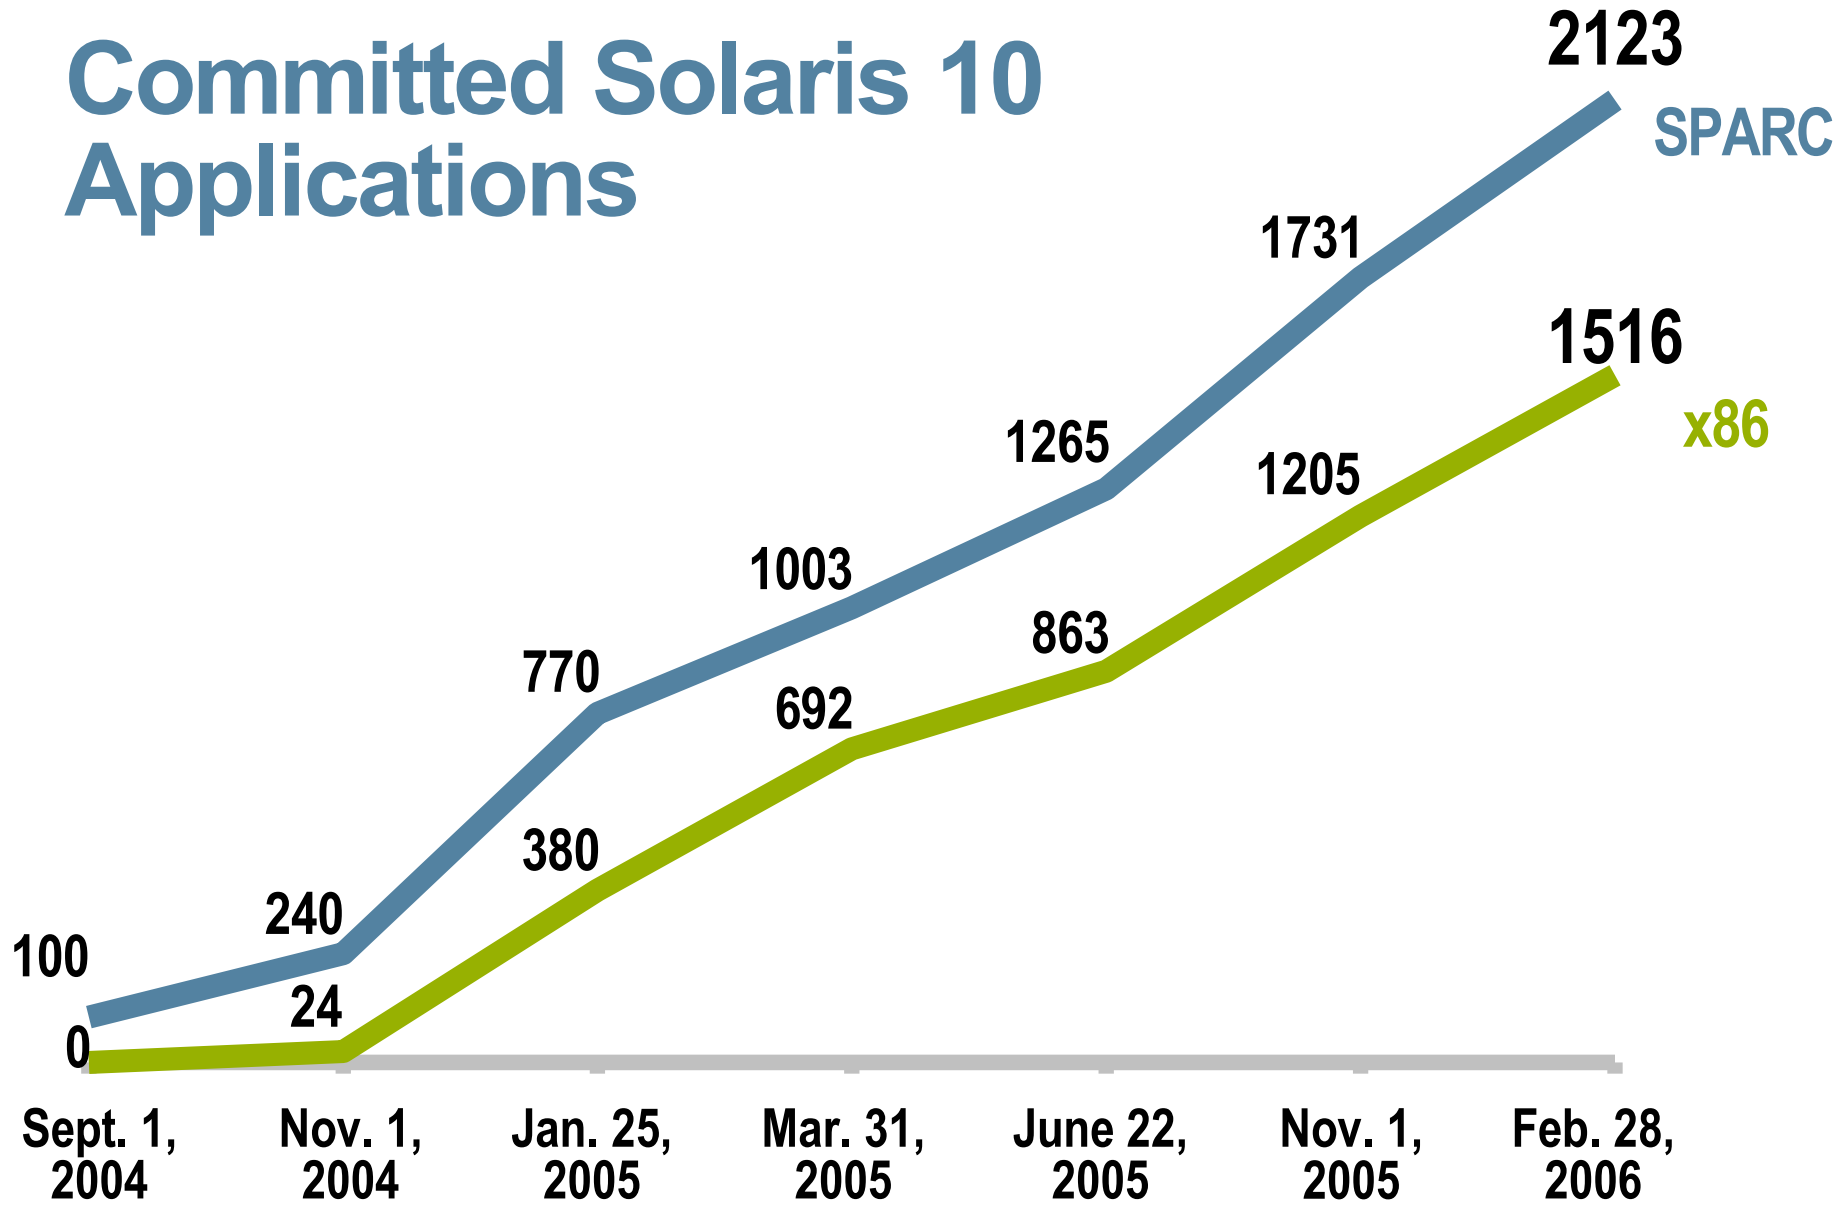
Hundreds of Systems




Thousands of ISVs and Partners



Committed Solaris 10 Applications



Data provided by Sun Microsystems



Solaris 10

4,700,000+ Licenses

(37% SPARC)

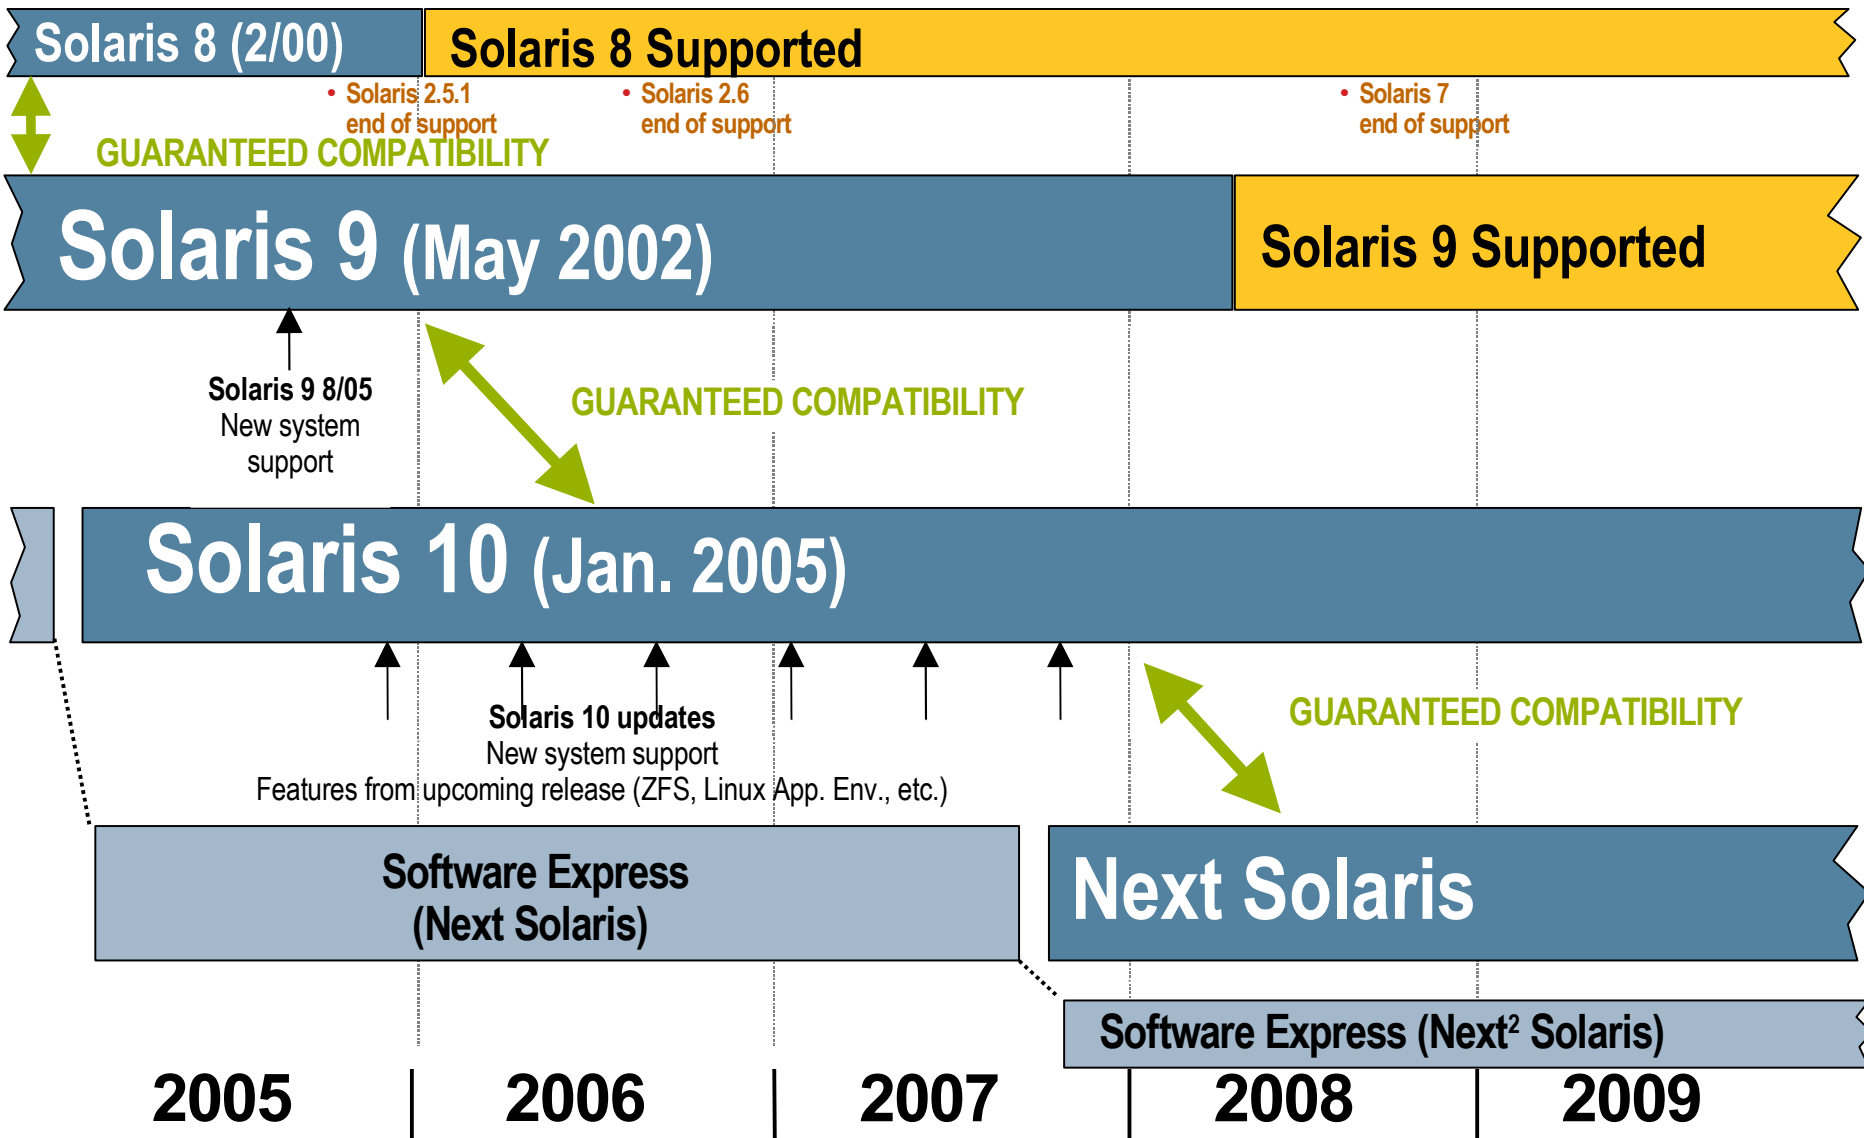
500+ Supported
Systems

400+ New ISVs

2,200+ Applications

40+ OEMs

Solaris Roadmap, 2005 - 2009

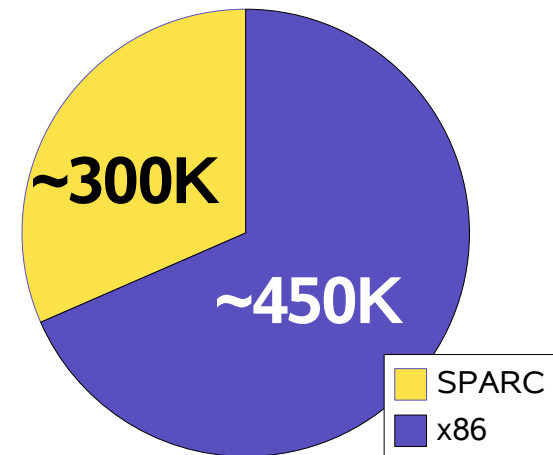


Software Express for Solaris

Timely Access to Leading-Edge Technology

- Monthly builds of next Solaris release
- Early access to new features
- Free access for downloads
- Subscription support for \$99/year

Solaris Express
Registered Licenses, 31-JAN-2004 >750,000



Accelerated delivery of disruptive technology

Software Express for Solaris

Timely Access to Leading-Edge Technology

- Built from the OpenSolaris source repositories
- Other consolidations integrated
- Binary only components (drivers, etc)
- Package based installer
 - > Including Live Upgrade mechanisms
- Internal testing for quality
- Released monthly (generally)
- Current release is Nevada Build 38

OpenSolaris



- OSI approved license (CDDL)
- Source code for the OS and Network (ON) consolidations
- Free compilers
- Active development communities
 - > Desktop
 - > Dtrace
 - > Service Management
 - > Xen

Solaris 10: A Generation Ahead



Extreme Performance



Optimal Utilization



Unparalleled Security



Relentless Availability



Platform Choice

Solaris 10: A Generation Ahead


- **Predictive Self Healing**
 - Solaris Fault Manager
 - Solaris Service Manager



Relentless Availability

The Solaris logo, consisting of a stylized sun with orange rays above the word "SOLARIS" in blue, sans-serif capital letters with a trademark symbol.

Platform Choice



Predictive Self-Healing

Stops problems before they happen

Self-diagnosis in milliseconds

Fine-grained recovery

Zero source code
changes required

Solaris Fault Manager

Solaris Fault Manager

- Automated error handling
 - > Detect, aggregate, diagnose, report, and mitigate faults
 - > Components offlined before failure
 - > Reduces downtime
- Reduces complexity
 - > Simplified error reporting
- Reduces costs
 - > 24 x forever uptime
 - > Increased utilization
 - > Higher server-to-administrator ratio



Relentless Availability


Before Solaris 10

```
[AFT0] errID 0x0000004d.23105c04 Corrected Memory Error on
U1004 is Intermittent
```

```
[AFT0] errID 0x0000004d.23105c04 ECC Data Bit 14 was in
error and corrected
```

```
WARNING:/io-unit@fe0200000/sbi@0,0/dma@0,81000/
esp@0,8000000 (esp0): Connected command timeout
for Target 0.0
```

```
NOTICE: correctable error detected by pci0
(upa mid lf) during DVMA read transaction
AFSR=40f40000.1f800000 AFAR=00000000.a25b4000
...
```

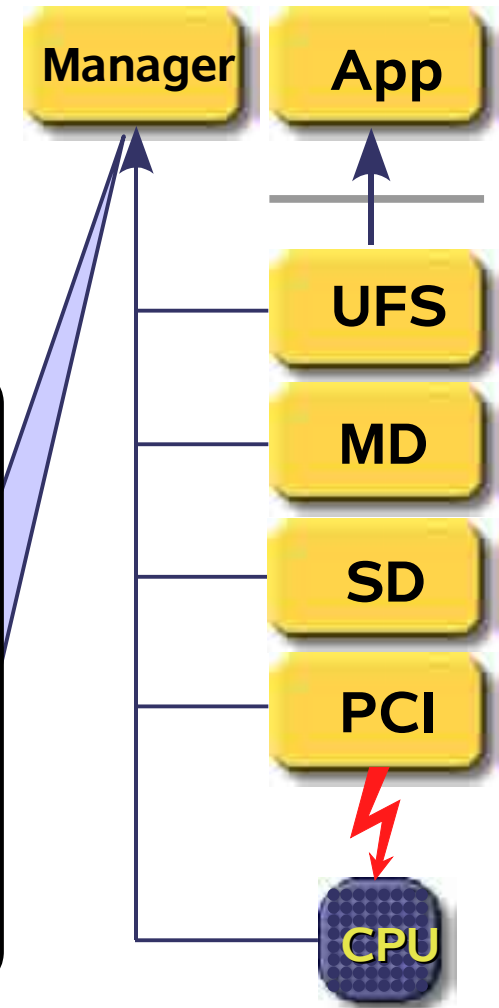


Fault Management Architecture

User-friendly Messages

- Include impact and action statements in a consistent format
- All events are managed and coordinated through a single management service

SUNW-MSG-ID: SF20000-W84N-KP3A-TF; TYPE: Fault, VER: 1, SEVERITY: Minor
AUTO-RESPONSE: Removal of the faulty memory resources has been initiated
IMPACT: Reduction in available memory resources
REQ-ACTION: A service call should be scheduled to inspect/replace the suspect components
DESC: A correctable memory data error occurred which has been diagnosed to be caused by a fault in a memory hardware component.



Fault Management Example

```
/*
 * Inject a kernel instruction uncorrectable ecache write-back error
 */
# mtst kiwdu
/*
 * New fault diagnosis message produced by fmd(1M) upon problem
  diagnosis
 * by the cpumem_diagnosis module
 */
SUNW-MSG-ID: SUN4U-8000-AC, TYPE: Fault, VER: 1, SEVERITY: Major
EVENT-TIME: Tue Apr 13 13:51:07 PDT 2004
PLATFORM: SUNW,Sun-Fire-V210, CSN: -, HOSTNAME: barback
SOURCE: cpumem-diagnosis, REV: 1.0
EVENT-ID: dd022489-9053-6aa5-a562-adf1959a7b0a
DESC: The number of errors associated with this CPU has exceeded
      acceptable levels. Refer to http://sun.com/msg/SUN4U-8000-AC for
      more information.
AUTO-RESPONSE: An attempt will be made to remove the affected CPU
      from service.
IMPACT: Performance of this system may be affected.
REC-ACTION: Schedule a repair procedure to replace the affected CPU.
            Use fmdump -v -u <EVENT_ID> to identify the CPU.
```

Solaris Service Manager

- Defines relationships among:
 - > Applications
 - > Solaris components
- Simplifies administration
 - > Consolidates “application profile”
 - > Helps manage components
- Eases installation
- Increases service reliability
 - > Detects service outages quickly
 - > Recovers services accurately



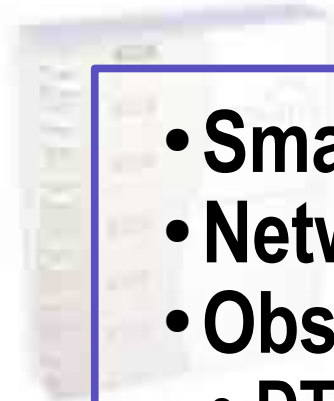
Relentless Availability

Solaris 10: A Generation Ahead



Extreme Performance

- Small system
- Network
- Observability
 - DTrace
 - libMicro



Optimal



Unparalleled Security



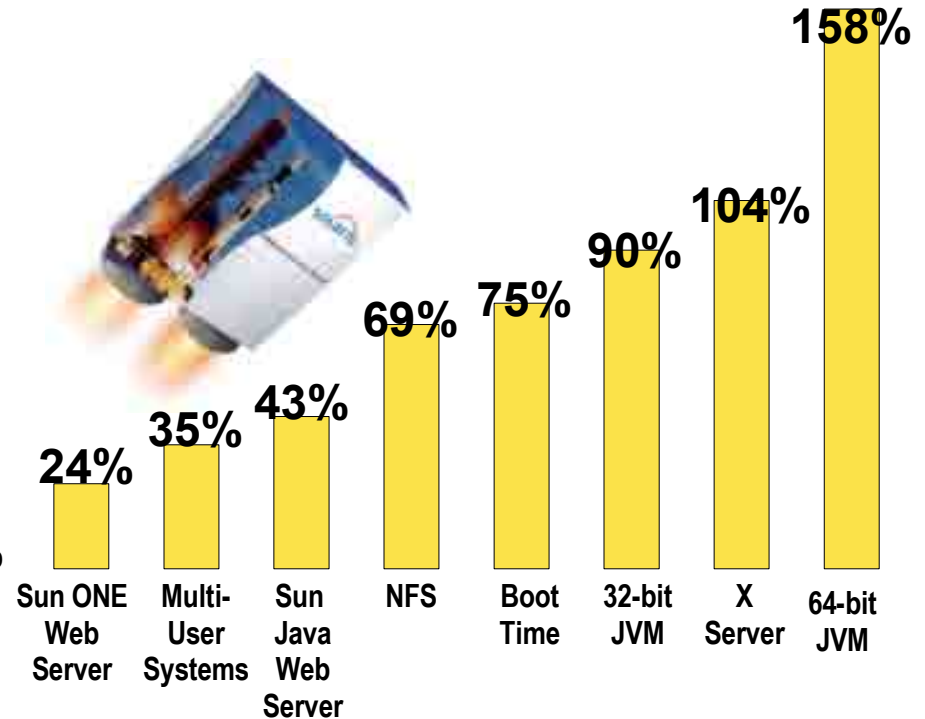

Relentless Availability



Platform Choice

Extreme Performance

Solaris 9 vs. Solaris 8, % Improvement

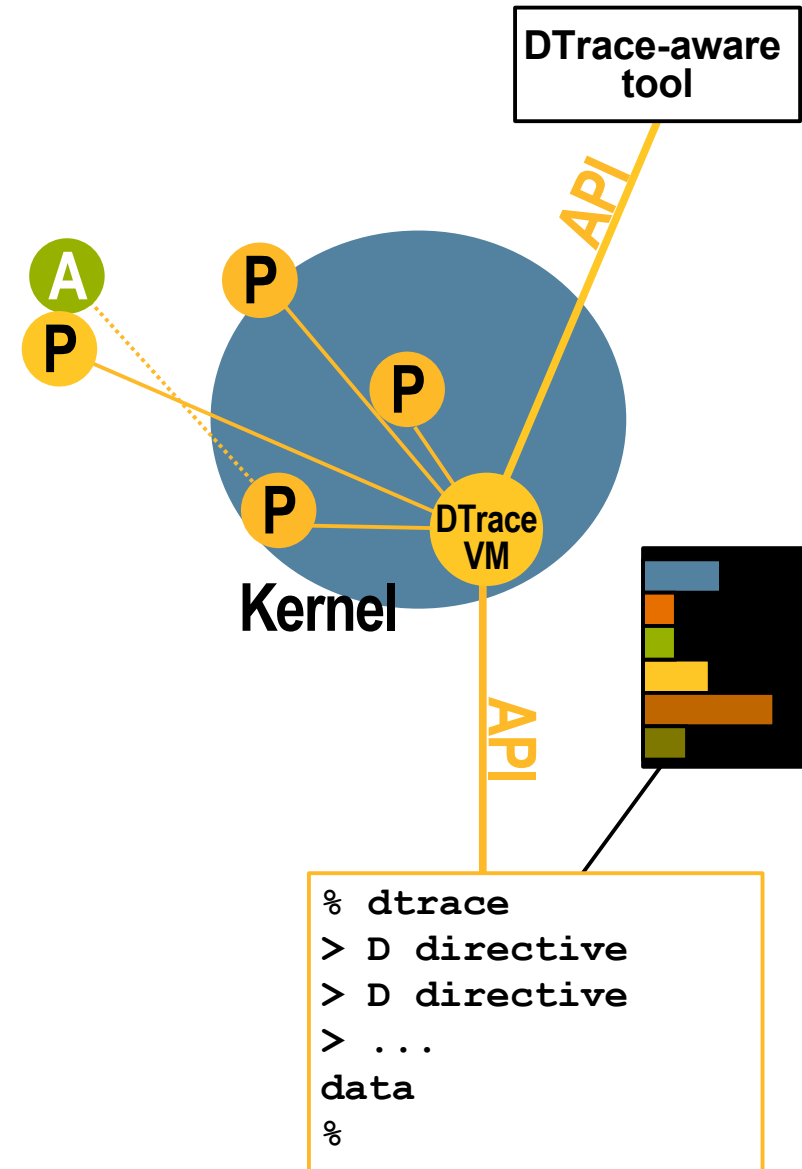


Solaris 10 vs. Solaris 9, % Improvement

Dynamic Tracing

Designed for Production Systems

- Safe; always there
 - > No performance hit
 - > No app or OS changes
- Views system as a whole
 - > Comprehensive
 - > Extensible; scriptable
- Debug, analyze, optimize in real time
 - > Solutions in minutes or hours, not days or weeks
 - > 3-30x speedups seen



Dtrace example

```
#!/usr/sbin/dtrace -s
```


```
syscall::write:entry
```

```
/procname == "esd"/
```

```
{
```

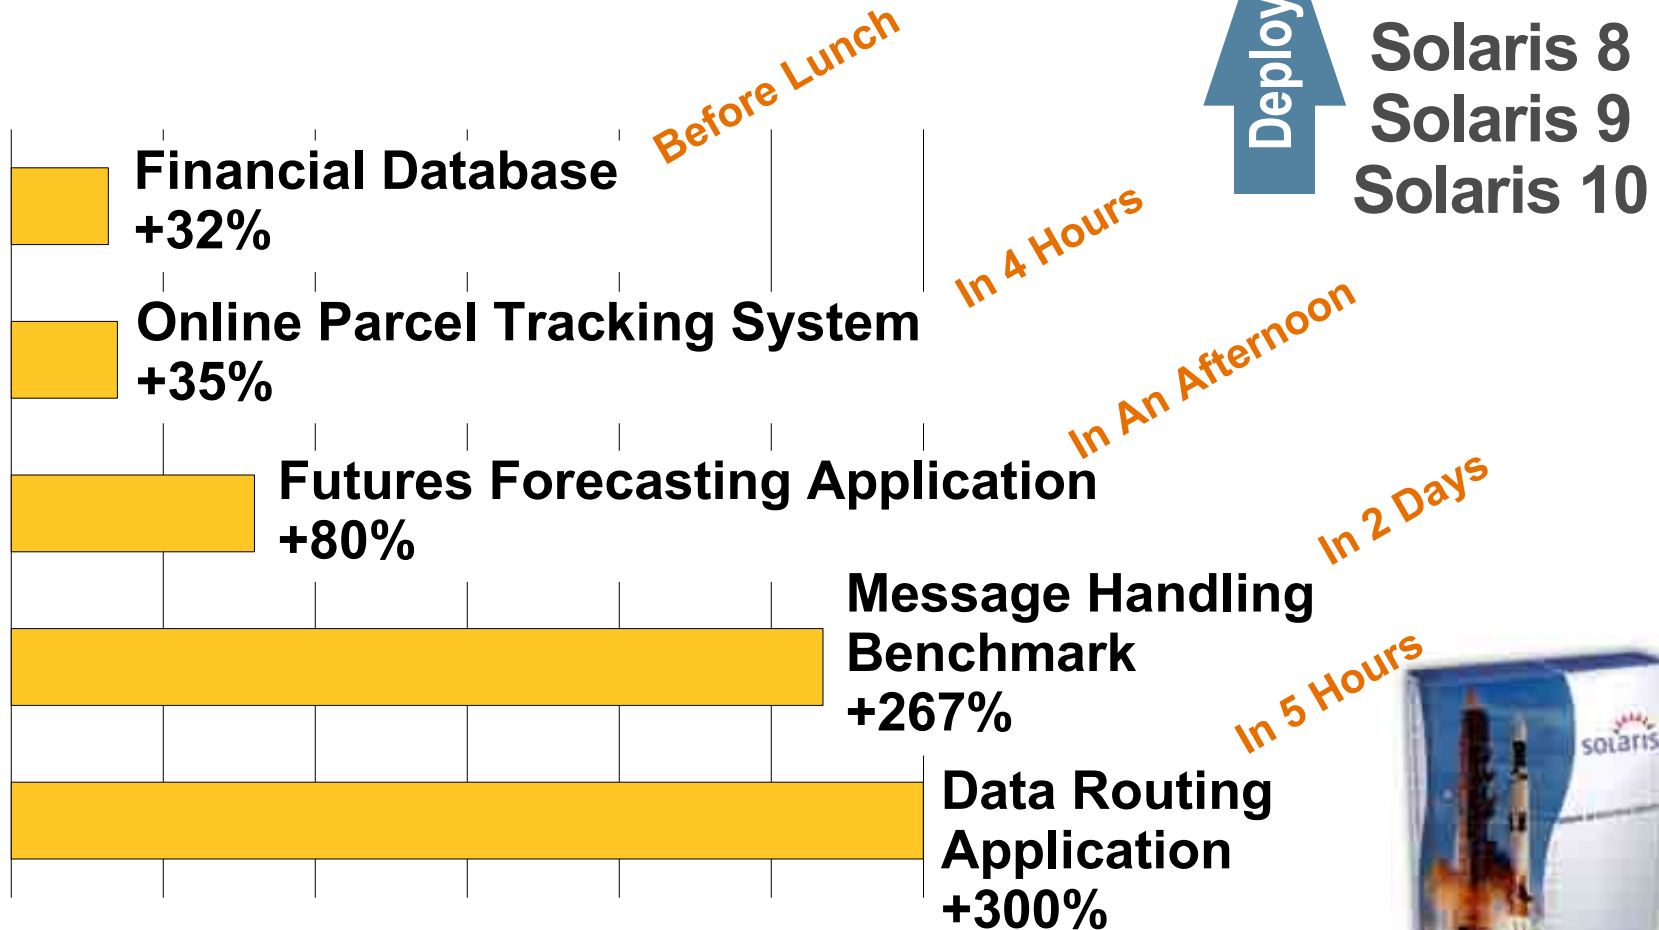
```
    @[execname] =  
    quantize(arg2);
```

```
}
```



Latest DTrace Wins

Simple Tool, Extreme Performance



Solaris 2.5
Solaris 2.6
Solaris 7
Solaris 8
Solaris 9
Solaris 10

Deploy

Maintain



Extreme Performance

Exploring DTrace, cont.

- DTrace source is available at:
 - > **<http://www.opensolaris.org/>**
- BigAdmin has a page and discussion forum dedicated to DTrace:
 - > **<http://sun.com/bigadmin/content/dtrace>**
- Solaris Dynamic Tracing Guide
 - > In the Software Developer Collection
- Brendan Gregg's sysadmin page
 - > <http://users.tpg.com.au/adsl4yb/dtrace.html>

IP Stack Streamlining

- Goal: saturate 10 Gb Ethernet with reduced overhead
- Classification engine for IP packets
- Specific, efficient methods to handle different packet types
- Solaris Network Cache and Accelerator functionality merged into IP stack



Extreme Performance

Scale Up, Scale Out

- Traditional focus: scalability
- Not always optimal for latency
- Goal: Industry-leading Solaris performance on all systems, including 1P/2P




solaris™

Extreme Performance

Throughput Computing Ready

- Ready for *multi-threaded* applications
- Ready for *multiple* applications
- MT-hot kernel, Java VM, system libraries
- Fast thread library
- 64 bit memory addressing
- Memory Placement Optimization
- Development & tuning tools
- Containers



Extreme Performance

Solaris 10: A Generation Ahead



Optimal Utilization

- Solaris Containers
- Solaris ZFS



Unparalleled Security



Platform Choice



SOLARIS™



Relentless Availability

Solaris Containers



Limitless partitioning – One license

Thousand's of applications on
one system

Fault & Security Isolation


Instant restart

Accelerate Consolidation

Multi-core aware

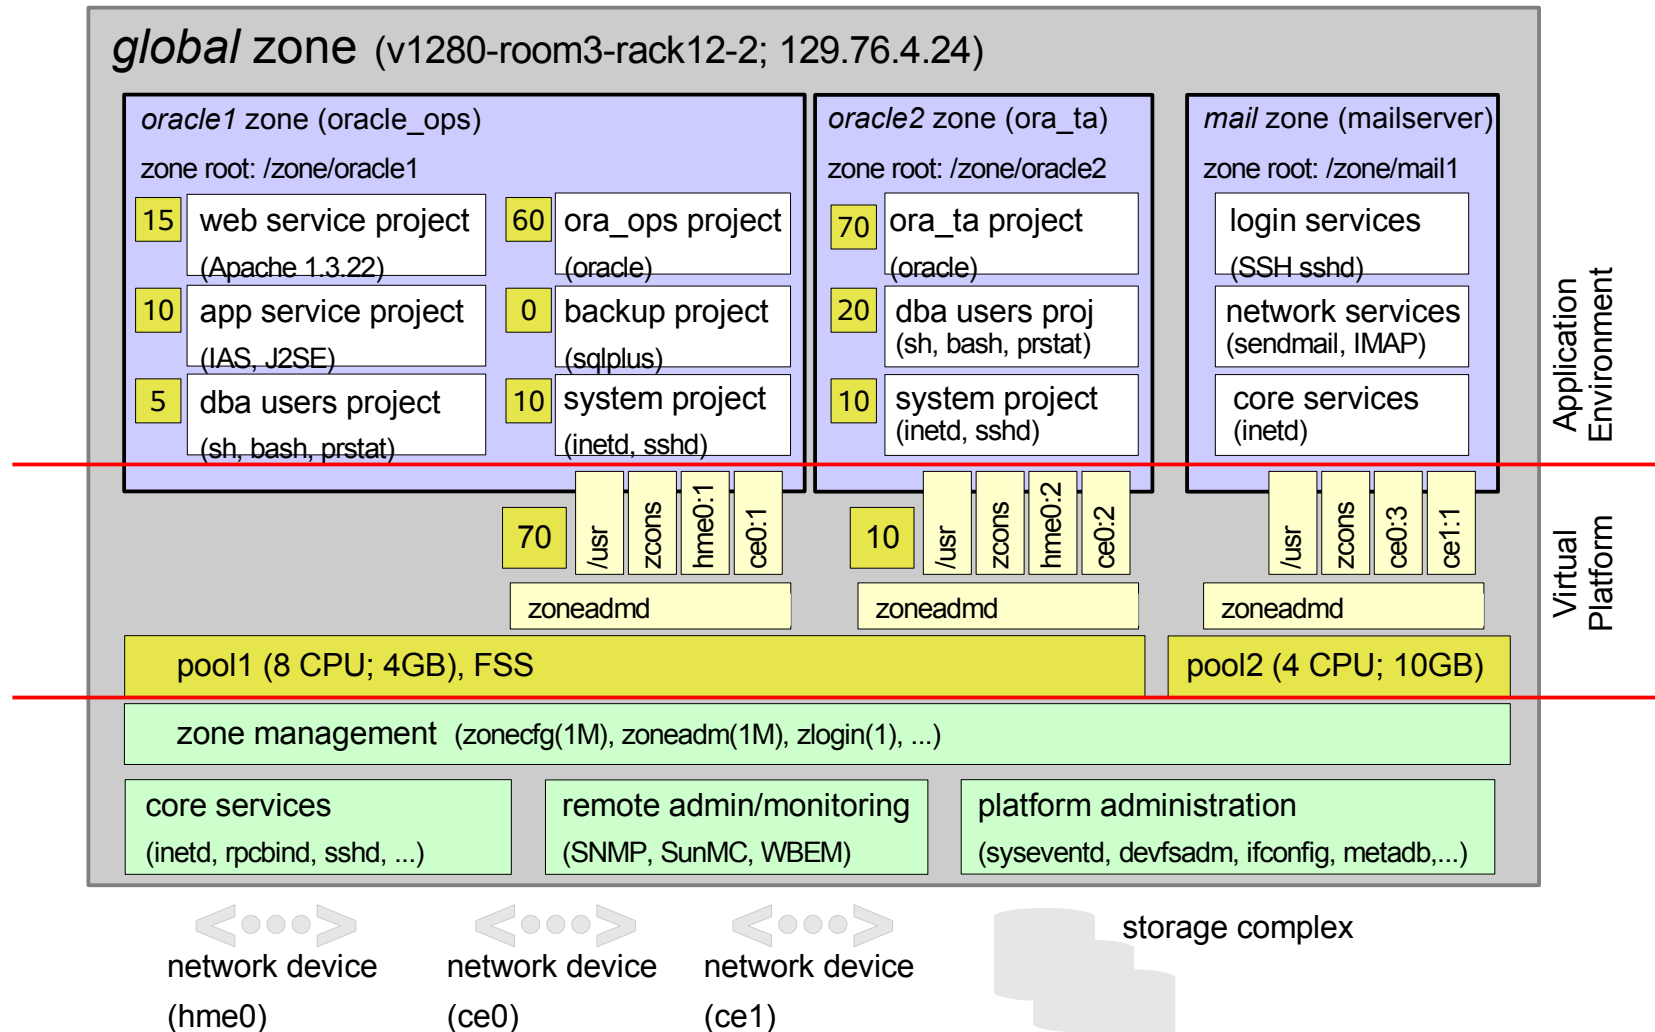
Solaris Containers

- Run multiple applications on one system
 - > Improve utilization by up to 4x
- Isolate applications from:
 - > Faults
 - > Intrusion
 - > Resource contention
- No performance impact

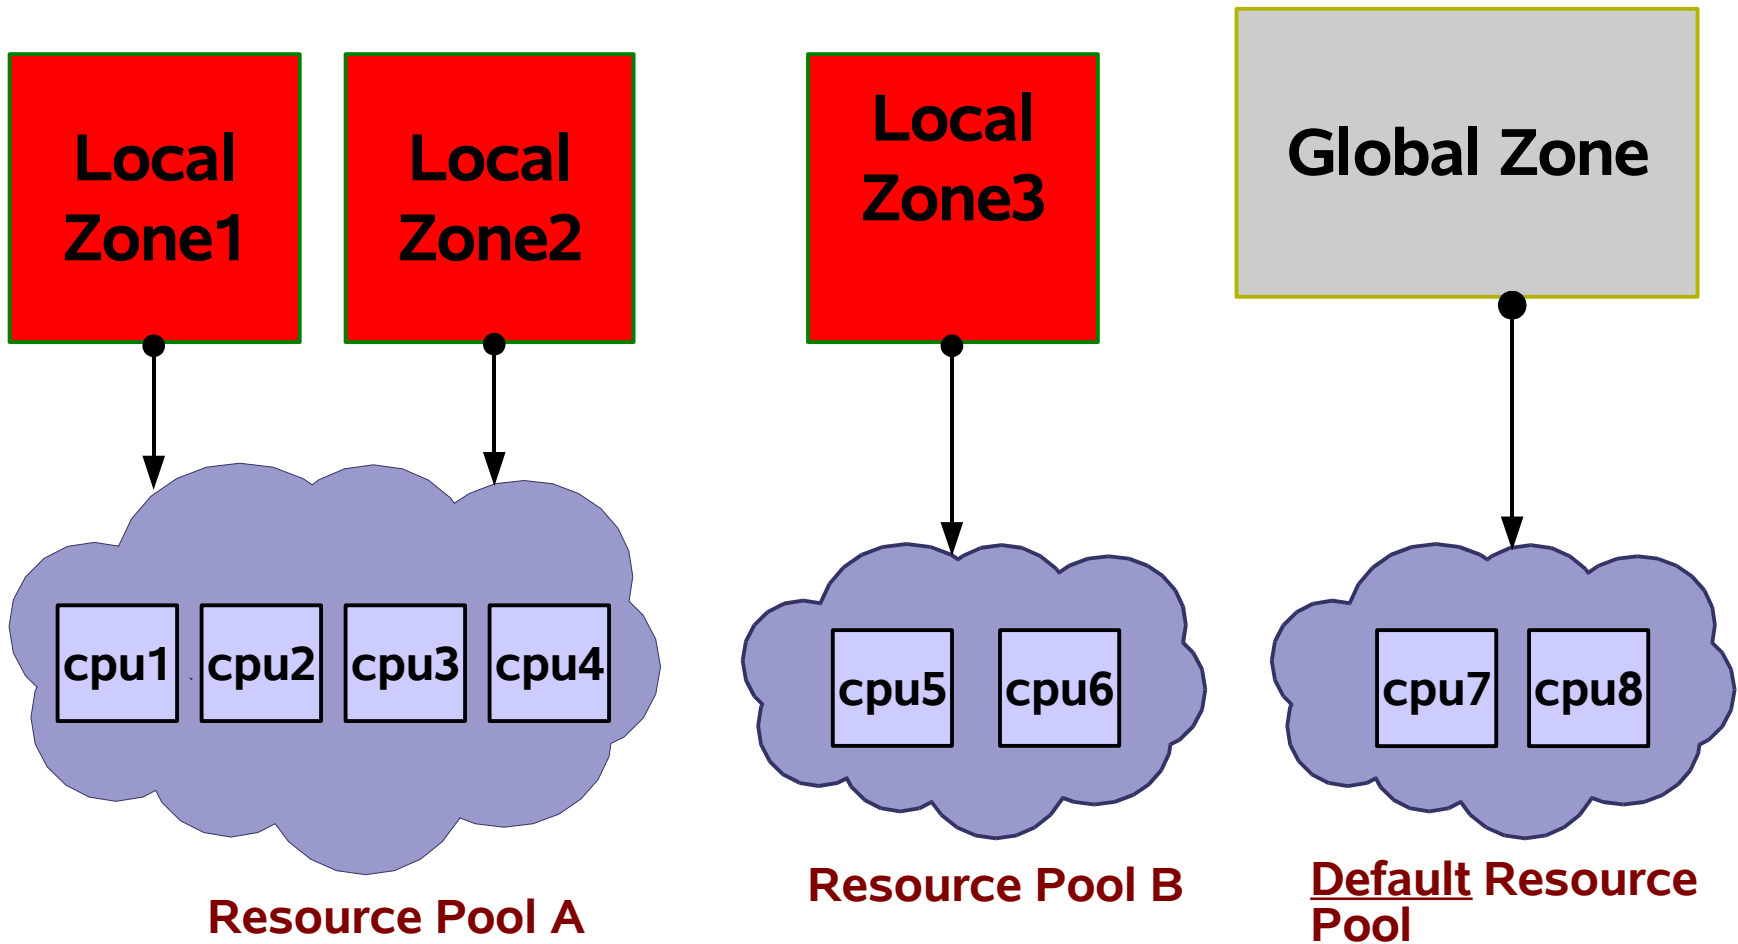


Examples

Multi Application Containers




Resource Mgmt Shallow Dive: Resource Pools



Solaris Containers

- Run multiple applications on one system
 - > Improve utilization by up to 4x
- Isolate applications from:
 - > Faults
 - > Intrusion
 - > Resource contention
- No performance impact



Solaris Containers: Scenarios

- “Shrink wrap” applications
 - > Fault, security, resource isolation
 - > Customize environments
 - > Run mutually hostile applications on same system
- SOA mobility
 - > Predefine environment (node name, addr, users...)
 - > Stop app on server “A,” restart on server “B”
- Safe, rapid, duplicatable prototyping
- Resource allocation—and chargeback

Solaris ZFS

- 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

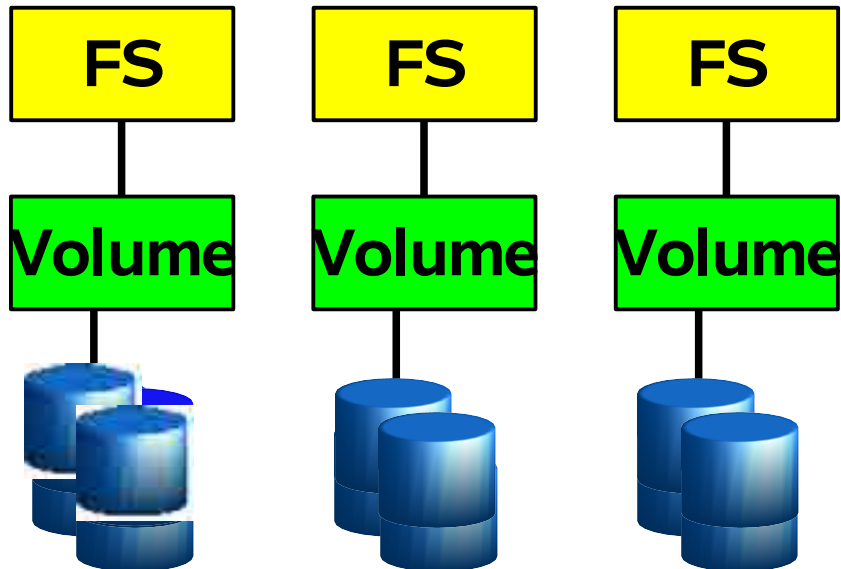


Optimal Utilization

FS/Volume Model vs. ZFS


Traditional Volumes

- Abstraction: virtual disk
- Partition/volume for each FS
- Grow/shrink by hand
- Each FS has limited bandwidth
- Storage is fragmented, stranded




ZFS Pooled Storage

- Abstraction: malloc/free
- No partitions to manage
- Grow/shrink automatically
- All bandwidth always available
- Pool allows space to be shared




Self-Healing Data in ZFS


1. Application issues a read. ZFS mirror tries the first disk. Checksum reveals that the block is corrupt on disk.



2. ZFS tries the second disk. Checksum indicates that the block is good.



3. ZFS returns good data to the application and repairs the damaged block.



Solaris 10: A Generation Ahead

- Administrative
- Application
- Network



Extreme Performance



Optimal Utilization



Unparalleled Security



Relentless Availability



Platform Choice



Proven Security

Over 20 Years of Design, Testing, Refinement and Experience

- Administrative
 - > Secure out of the box
 - > System integrity (BART, Secure Execution)
 - > User rights management
 - > Containers
- Application
 - > Process rights management
 - > Cryptographic framework
- Network
 - > IP filtering



Unparalleled Security

Solaris 10: A Generation Ahead

- Linux interoperability
- New UltraSPARC IV
- New AMD Opteron
- Java Desktop System

Extreme Performance

Unparalleled Security

Relentless Availability



SOLARIS™



Platform Choice

Solaris 10 on x64

- Supports the i386 psABI and amd64 psABI on
 - > AMD 64-bit CPUs, including dual-core support
 - > Intel IA32 CPUs with 64-bit extensions (EM64T)
- Architecture cloned from 64-bit SPARC
 - > Directory organization: ../amd64, ../64
 - > Data model is LP64
 - > All x86 systems are 64-bit development platforms
 - > 64-bit environment supports 32-bit and 64-bit applications
- Leverage hardware features for performance
- Support and use NX (No eXecute) bit for security

AMD64 Specific Features

- **Support for AMD HyperTransport**
 - > Memory bandwidth grows as you add processors
 - > NUMA architecture
 - > Optimizations are identical to Solaris Memory Placement Optimization (MPO) for SPARC introduced in Solaris 9
- **3dNow!, improved prefetching/non-temporal stores: HWCAP**
 - > Instruction set extensions in libraries
 - > Filters or whole libraries
 - > /lib/libc.so.1 is now a mount point(!)
 - > Documented in Linkers and Libraries guide

Linux Interoperability

Coming
Soon

- Solaris Containers for Linux Applications
- Run Linux applications (binaries) without modification
- Dtrace observability
- Solaris performance, scalability and security
- Reduced deployment costs
- Freedom to choose applications



Platform Choice

F/OS Software in Solaris 10

Network Servers & Clients

Apache
Apache2
bind
Mozilla
ncftp
ppp
Samba
sendmail
SER (SIP Proxy Server)
Tomcat
wget
wu-ftpd
xntpd
Zebra

Commands

a2ps
bzip2
footmatic print ppds
ghostscript
ghostscript fonts
Gimp print drivers
GNU patch utility
GNU grep
ImageMagick
IPMItool
Open Printing API
rpm2cpio.pl
System Management Agent

MySQL
patch
texinfo
traceroute
Webmin
gzip
GNU tar
less
texi2html
mkisofs

Libraries

Glib
GTK+
JPEG
Libexpat
Libusb
Libxslt
PNG
Tcl/Tk
TIFF
XML2
XPM
zlib

Compilers & Tools Scripting Languages Security Tools

Binutils
Bison
Flex
Gcc
Gm4
Gmake

Perl
Python

Secure Shell
tcp_wrappers

Shells

bash
tcsh
zsh

Fully supported: Sun provides support in the same way as for Sun owned software.

Managed: Sun provides existing patches and escalates new bugs to the developer community

F/OS Software on Solaris 10 CCD

Applications / Accessible	Applications / Networking	Applications / Utilities	Development / Language	Development / Libraries
brltty-3.3.1	cups-1.1.20	afio-2.4.6	bison-1.35	aalib-1.2
emacspeak-18.0	etherreal 0.10.5	amanda-2.4.4	gcc-2.95.3	berkeley-db 1.85
emacspeak-ss-1.9.1	fetchmail 6.2.5	cdrttools-2.01	gcc-3.4.2	berkeley-db 4.2.52NC
freetts-1.1.1	hpjjs 1.6	cupsddk 1.0	libtool 1.5.2	curl-7.10.3
screenbrltty-4.02	lynx-2.8.4	diffutils-2.8.1	m4-1.4 (GNU)	fttk-1.1.3
unwindows-1.1.3	mutt-1.4.2.1	enscript 1.6.1	MySQLpython API0.9.2	flib-0.5
w3-4.0.47	nmap 3.5	expect5.39	php-4.3.2	GD Graphics library 2.0.
yasr-0.6.4	nmh-1.0.4	file-4.10	ruby-1.6.4	guile-1.3.4
	Open LDAP 2.2.17	fileutils-4.1	samp-1.0	imlib-1.9.15
Applications / Editors	Open SLP 1.0.11	findutils-4.1.20	tcIX-8.2.0	libxpt-1.95.7
bluefish 0.12	pine-4.61	Foomatic filters 3.0.2		libmpeg-1.3.1
emacs 21.3	procmail-3.22	Foomatic-ppds 3.0.1	Development / Tools	libpcap-0.8.3
gawk-3.0.6	rsync 2.6.3pre1	gcal-3.01	autoconf 2.59	libsane 1.0.14
joe-3.1	slm-0.9.6.2	gettext-0.10.35	automake 1.8.3	linungif-4.1.0
sed-3.02 (GNU)	snort-2.0.0	gimp-print-4.2.6	binutils-2.15	ncurses-5.2
vim-6.3	tcpdump-3.8.3	gkrellm 2.1.19	cvs 1.11.17	Ogglib-1.0
xemacs-21.4.15		gnuplot 3.7.3	ddd 3.3.8	Perl regex lib 4.5
	Applications / Publishing	ispell-3.2.06	gdb 6.2.1	qt-3.1.1
Desktop / Environment	espgs-7.07.1	lxr-0.9.6.1	global-4.8	readline-4.2
kde-3.1.1a	graphviz 1.10	mpack-1.5	make-3.80 (GNU)	slang-1.4.0
KOffice-1.2.1	groff-1.16.1	mpage-2.5.1		SDL-1.2.5
XFce-3.8.16	xpdf 3.0			Xaw3d-1.5
	System / Daemons	X / Applications	X / Window Managers	
	imap2002d (UW)	asclock-1.0	afterstep-1.8.8	
		etherreal-0.9.11	fwmm2-2.4.3	
		mpg123-0.59r		
		mysql-jdbc-3.0.8		
		netpbm-10.3		
		plotutils-2.4.1		
		pnm2ppa-1.12		
		rpm-4.1		
		sane 1.0.12		
		screen 4.0.2		
		sgrep-1.92a		
		sh-utils-2.0		
		sharutils-4.2.1		
		sudo 1.6.8p5		
		TeX 2.0.2		
		textutils-2.0		
		tnet 1.1.3		
		top-3.5.1		
		uudeview-0.5.20		
		vorbis-1.0		
		wine 20041104		
		xpp-1.1		

System Choice

Optimized Solaris Performance on a Variety of Systems

- AMD Opteron
 - > 64-bit optimizations
 - > Hypertransport
- UltraSPARC
 - > CMT scheduling enhancements
 - > CPU performance counters
- Pentium 4
 - > SSE/SSE2
 - > Hyperthreading




Platform Choice

Java Desktop System

- New highly customizable GNOME desktop
 - > Based on GNOME 2.6 community release
 - > Extensive assistive technologies
- Java 1.5 SDK
- Mozilla 1.7
- Evolution 1.4
- APOC preferences management
- KDE available on the Software Companion
- CDE still available



Platform Choice



Solaris 10:

Relentless Availability

Extreme Performance

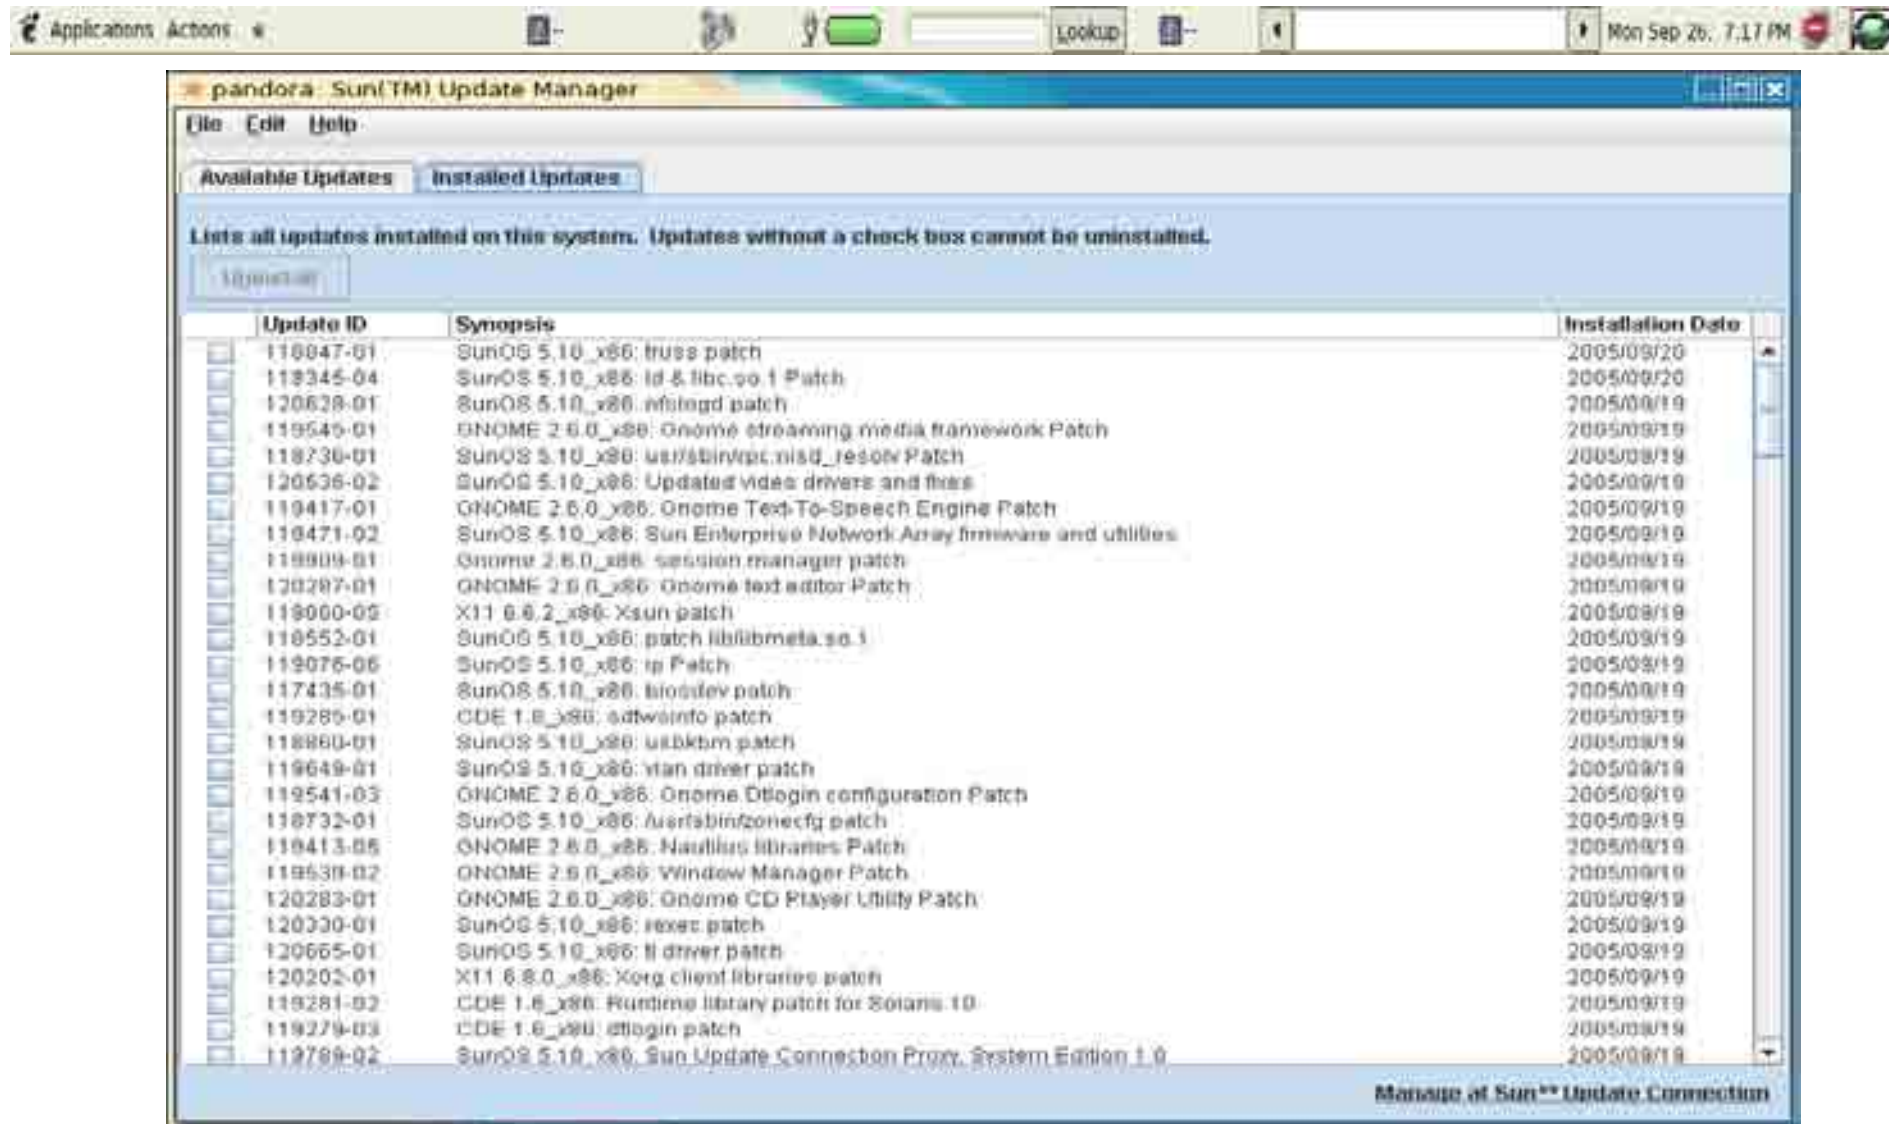
Optimal Utilization

Unparalleled Security

Platform Choice

“The Whole is Greater...”

Solaris Update Manager



The screenshot shows the Solaris Update Manager application window. The title bar reads "pandora: Sun(TM) Update Manager". The menu bar includes "File", "Edit", and "Help". There are two tabs: "Available Updates" and "Installed Updates", with "Installed Updates" being the active tab. Below the tabs, a message states: "Lists all updates installed on this system. Updates without a check box cannot be uninstalled." There is a "Signature" button. The main area contains a table with the following columns: "Update ID", "Synopsis", and "Installation Date".

Update ID	Synopsis	Installation Date
<input type="checkbox"/> 118047-01	SunOS 5.10_x86: truss patch	2005/09/20
<input type="checkbox"/> 118345-04	SunOS 5.10_x86: ld & libc.so.1 Patch	2005/09/20
<input type="checkbox"/> 120628-01	SunOS 5.10_x86: nfslogd patch	2005/09/19
<input type="checkbox"/> 119545-01	GNOME 2.6.0_x86: Gnome streaming media framework Patch	2005/09/19
<input type="checkbox"/> 118730-01	SunOS 5.10_x86: usr/sbin/rpc.nisd_resolv Patch	2005/09/19
<input type="checkbox"/> 120636-02	SunOS 5.10_x86: Updated video drivers and files	2005/09/19
<input type="checkbox"/> 119417-01	GNOME 2.6.0_x86: Gnome Text-To-Speech Engine Patch	2005/09/19
<input type="checkbox"/> 116471-02	SunOS 5.10_x86: Sun Enterprise Network Array firmware and utilities	2005/09/19
<input type="checkbox"/> 118909-01	Gnome 2.6.0_x86: session manager patch	2005/09/19
<input type="checkbox"/> 120287-01	GNOME 2.6.0_x86: Gnome text editor Patch	2005/09/19
<input type="checkbox"/> 118060-02	X11 6.8.2_x86: Xsun patch	2005/09/19
<input type="checkbox"/> 118552-01	SunOS 5.10_x86: patch liblibmeta.so.1	2005/09/19
<input type="checkbox"/> 119076-06	SunOS 5.10_x86: tp Patch	2005/09/19
<input type="checkbox"/> 117435-01	SunOS 5.10_x86: biosdev patch	2005/09/19
<input type="checkbox"/> 119285-01	CDE 1.6_x86: sdwainfo patch	2005/09/19
<input type="checkbox"/> 118880-01	SunOS 5.10_x86: usbkbn patch	2005/09/19
<input type="checkbox"/> 119649-01	SunOS 5.10_x86: vian driver patch	2005/09/19
<input type="checkbox"/> 119541-03	GNOME 2.6.0_x86: Gnome Dlogin configuration Patch	2005/09/19
<input type="checkbox"/> 118732-01	SunOS 5.10_x86: /usr/sbin/zonectl patch	2005/09/19
<input type="checkbox"/> 118413-05	GNOME 2.6.0_x86: Nautilus libraries Patch	2005/09/19
<input type="checkbox"/> 118538-02	GNOME 2.6.0_x86: Window Manager Patch	2005/09/19
<input type="checkbox"/> 120283-01	GNOME 2.6.0_x86: Gnome CD Player Utility Patch	2005/09/19
<input type="checkbox"/> 120330-01	SunOS 5.10_x86: rexec patch	2005/09/19
<input type="checkbox"/> 120665-01	SunOS 5.10_x86: ti driver patch	2005/09/19
<input type="checkbox"/> 120202-01	X11 6.8.0_x86: Xorg client libraries patch	2005/09/19
<input type="checkbox"/> 119281-02	CDE 1.6_x86: Runtime library patch for Solaris 10	2005/09/19
<input type="checkbox"/> 119279-03	CDE 1.6_x86: dlogin patch	2005/09/19
<input type="checkbox"/> 119789-02	SunOS 5.10_x86: Sun Update Connection Proxy System Edition 1.0	2005/09/19

Manage of Sun** Update Connection



solaris™

\$0



open source™

Subscription-based Service Plans for Solaris 10

Subscription Pricing

	Free	Basic	Standard	Premium
Solaris 10 OS security fixes	■	■	■	■
Regular Solaris 10 OS update releases	■	■	■	■
Solaris 10 OS overview Web training course	■	■	■	■
Sun Update Connection Web training course		■	■	■
Real time access to patches/fixes		■	■	■
System Edition of Sun Update Connection		■	■	■
Skills self-assessment		■	■	■
One Web course			■	■
Optional training credits			■	■
5 x 12 telephone support			■	■
7 x 24 telephone support				■
Interoperability services				■
US\$ Price/Physical Processor/Year, 1-8	\$0	\$120	\$240	\$360
9+	\$0	--	\$480	\$600

“... I came away more pleased than I have been in a *long* time after any meeting with a vendor!

“Solaris 10 represents true strategic competitive advantage for Sun and true breakthrough functionality for [our firm]. The combination of D-Trace, Containers, and Fault Tolerance are the holy grails of commercial-grade operating systems and Solaris 10 seems to have all three (and fully integrated to boot)!”

**Executive VP for Technology,
Fortune 500 Brokerage Firm
email to Scott McNealy**

**This is cool, so how
do I get started ???**

What can you do ?

- Download Solaris 10 – it's free!!!
 - > Have a spare system, set up a dual boot environment
 - > Use Live Upgrade to preserve customizations
 - > <http://www.sun.com/software/solaris/get.jsp>
- If you want to see where this may be going
 - > Download Software Express for Solaris
 - > It's free too!
 - > <http://www.sun.com/softwareexpress>
 - > See the new Solaris boot environment
 - > Lots of new Dtrace probes
 - > ZFS soon!
 - > One step to convert to OpenSolaris community edition

Sun Value Proposition

Open is Good

Innovation Pays

**The OS
Matters**

SOLARIS™

Solaris 10 Something for Free – and more!

Bob Netherton

bob.netherton@sun.com

<http://blogs.sun.com/bobn>

The logo for the Suncoast Oracle Users' Group. It features the word "SUNCOAST" in blue, "ORACLE" in white on a red rectangular background, and "USERS' GROUP" in blue. The text is set against a white background with a large, glowing yellow circular arc behind it.

SUNCOAST
ORACLE
USERS' GROUP

Monthly
4th Thursday
6pm – 8pm

IBM Center
Rocky Point

[WWW.SOUUG.NET](http://www.souug.net)