Sun's Platform Strategy

Ambreesh Khanna
Chief Technologist
x64 Servers and Solaris, US Client Solutions
Sun Microsystems, Inc.

Sun microsystems

The Infrastructure Stack

The Integrated Platform

The Infrastructure Stack

The Integrated Platform

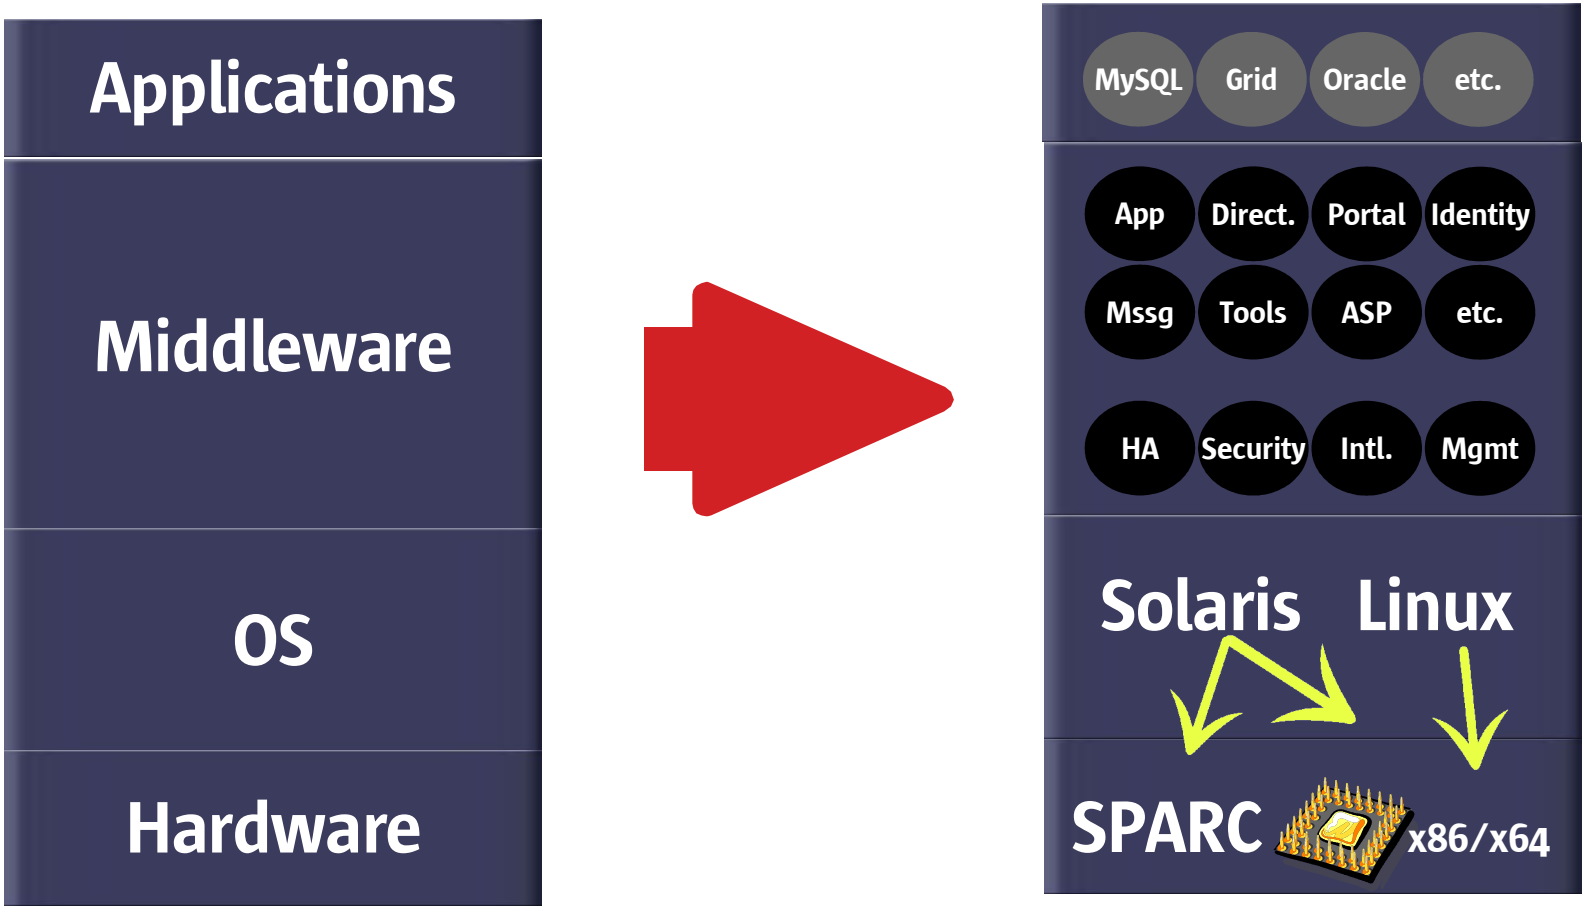

The Infrastructure Stack

The Integrated Platform

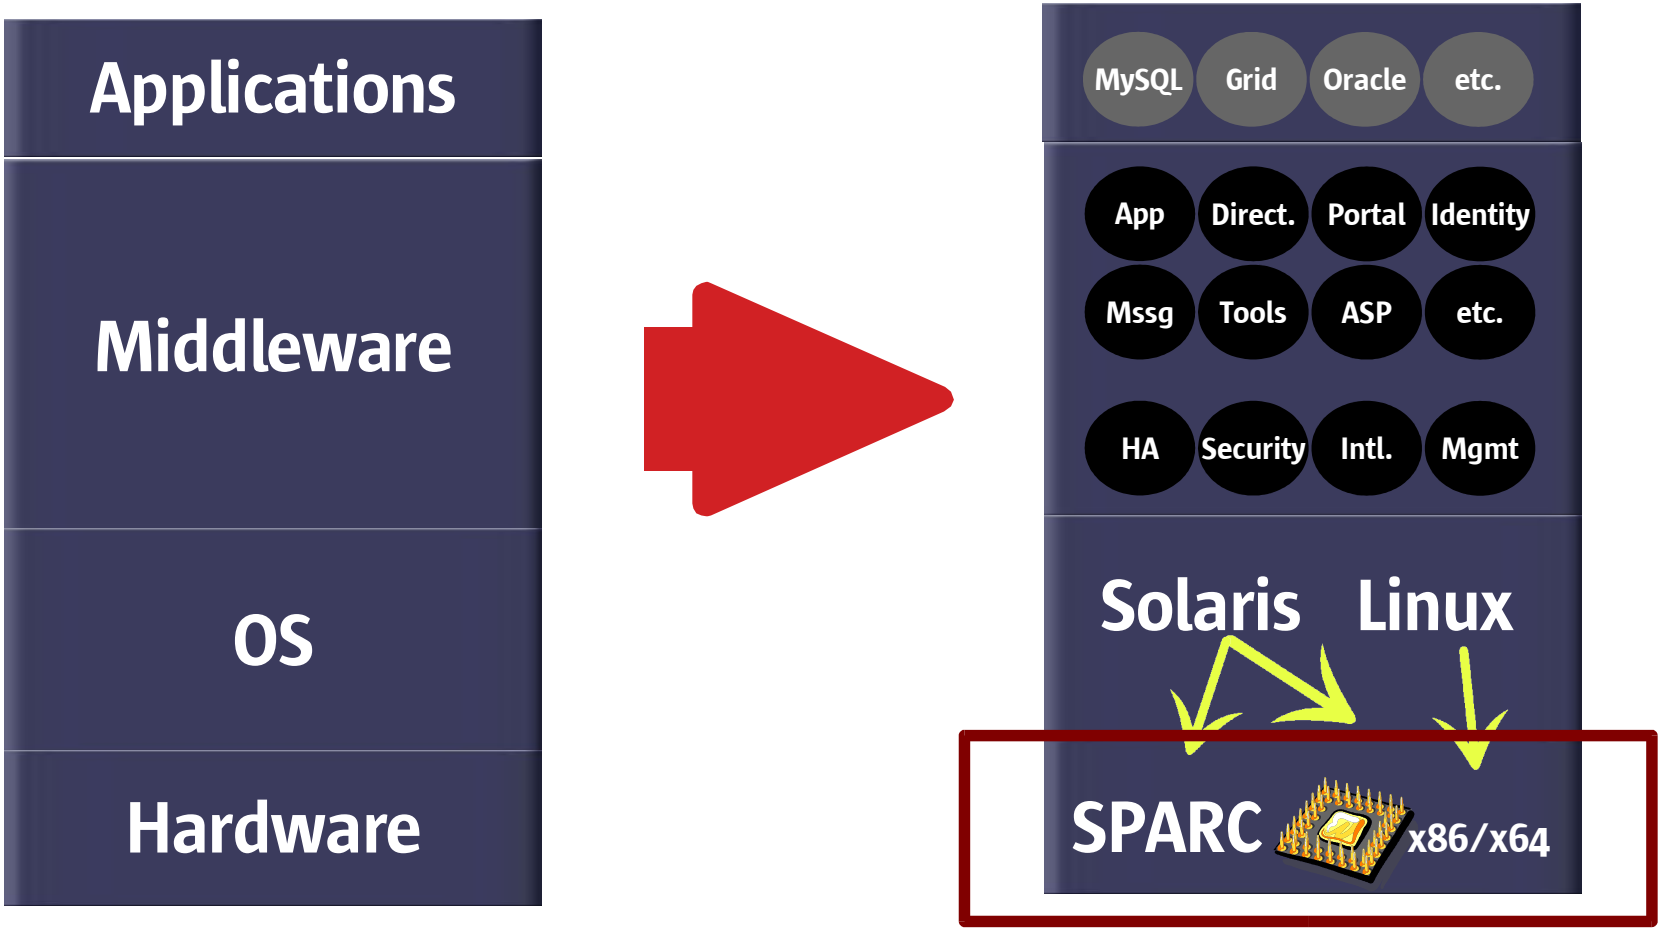

The n-tier Architecture

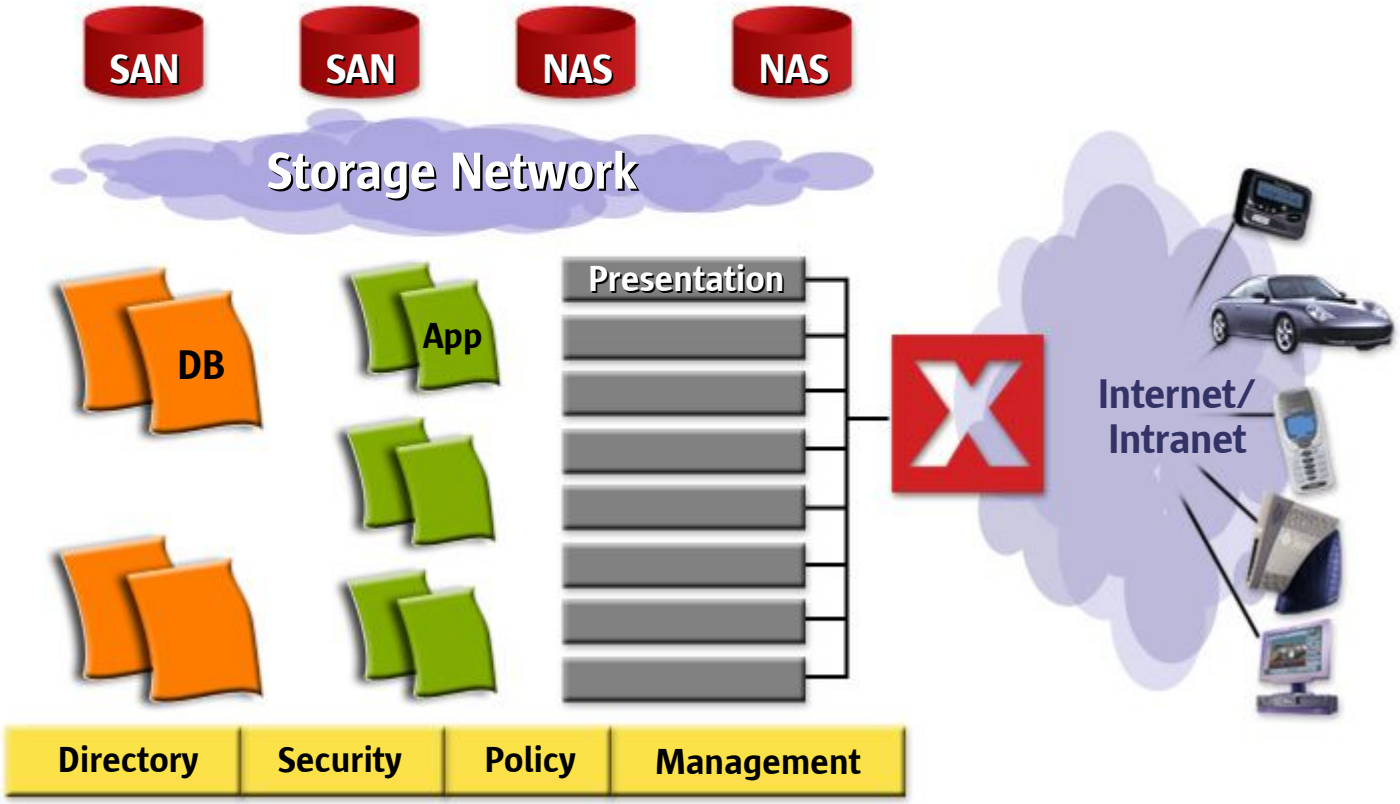

Attributes of Commercial Workloads

	Web Services			Client Server		Data Warehouse
Attribute	**Web** (Web99)	**App Serv** (JBB)	**Data** (TPC-C)	**SAP 2T**	**SAP 3T** (DB)	**DSS** (TPC-H)
Application Category	Web Server	Server Java	OLTP	ERP	ERP	DSS
Instruction-level Parallelism	Low	Low	Low	Medium	Low	High
Thread-level Parallelism	High	High	High	High	High	High
Instruction/Data Working Set	Large	Large	Large	Medium	Large	Large
Data Sharing	Low	Medium	High	Medium	High	Medium

Intel Agrees

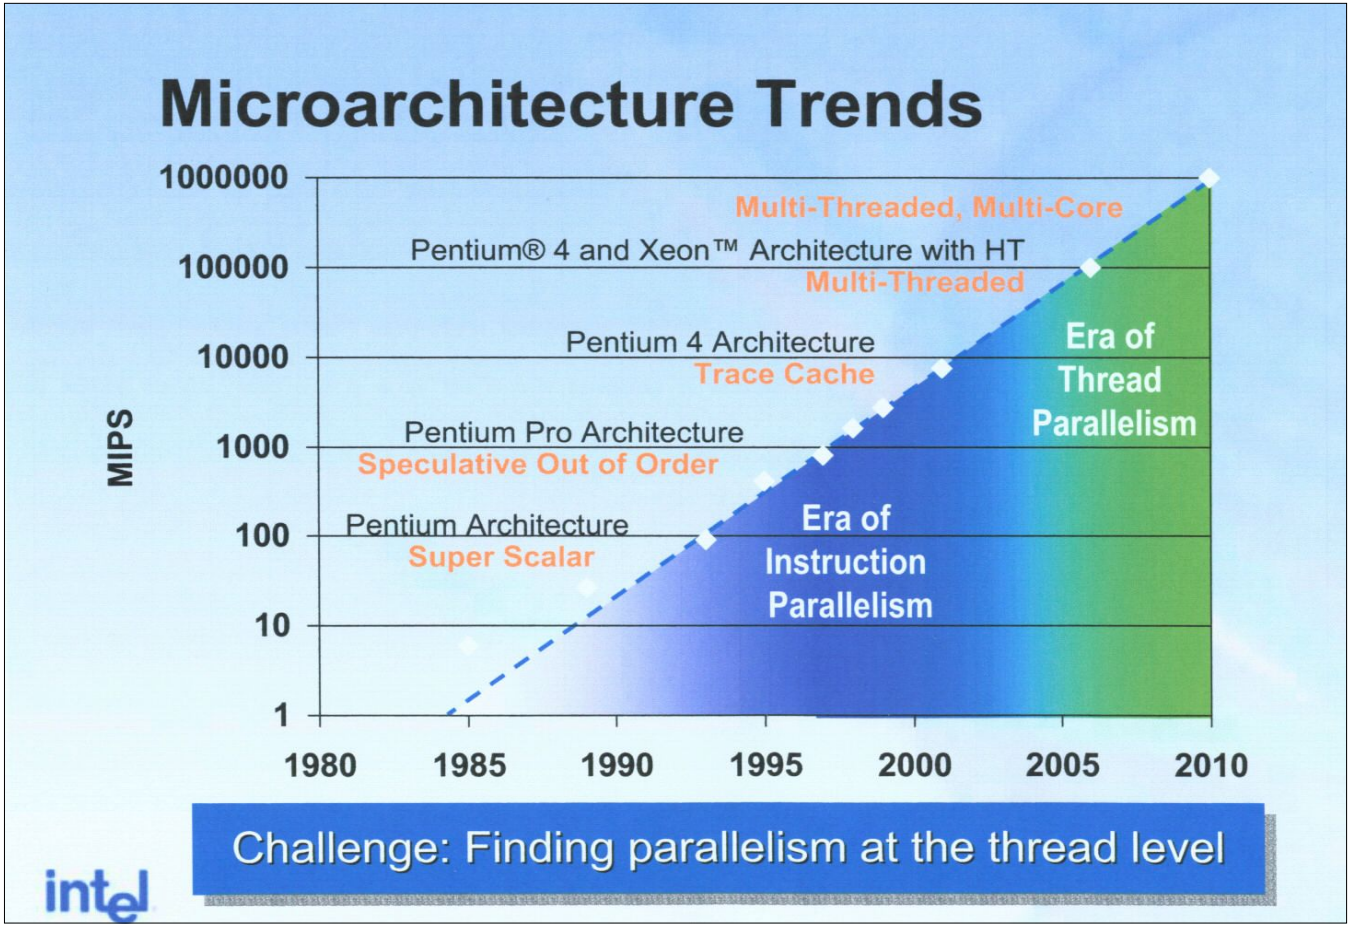

Source: IA32 Processor Architecture Trends and Research, IDF Spring 2003

Memory Bottleneck

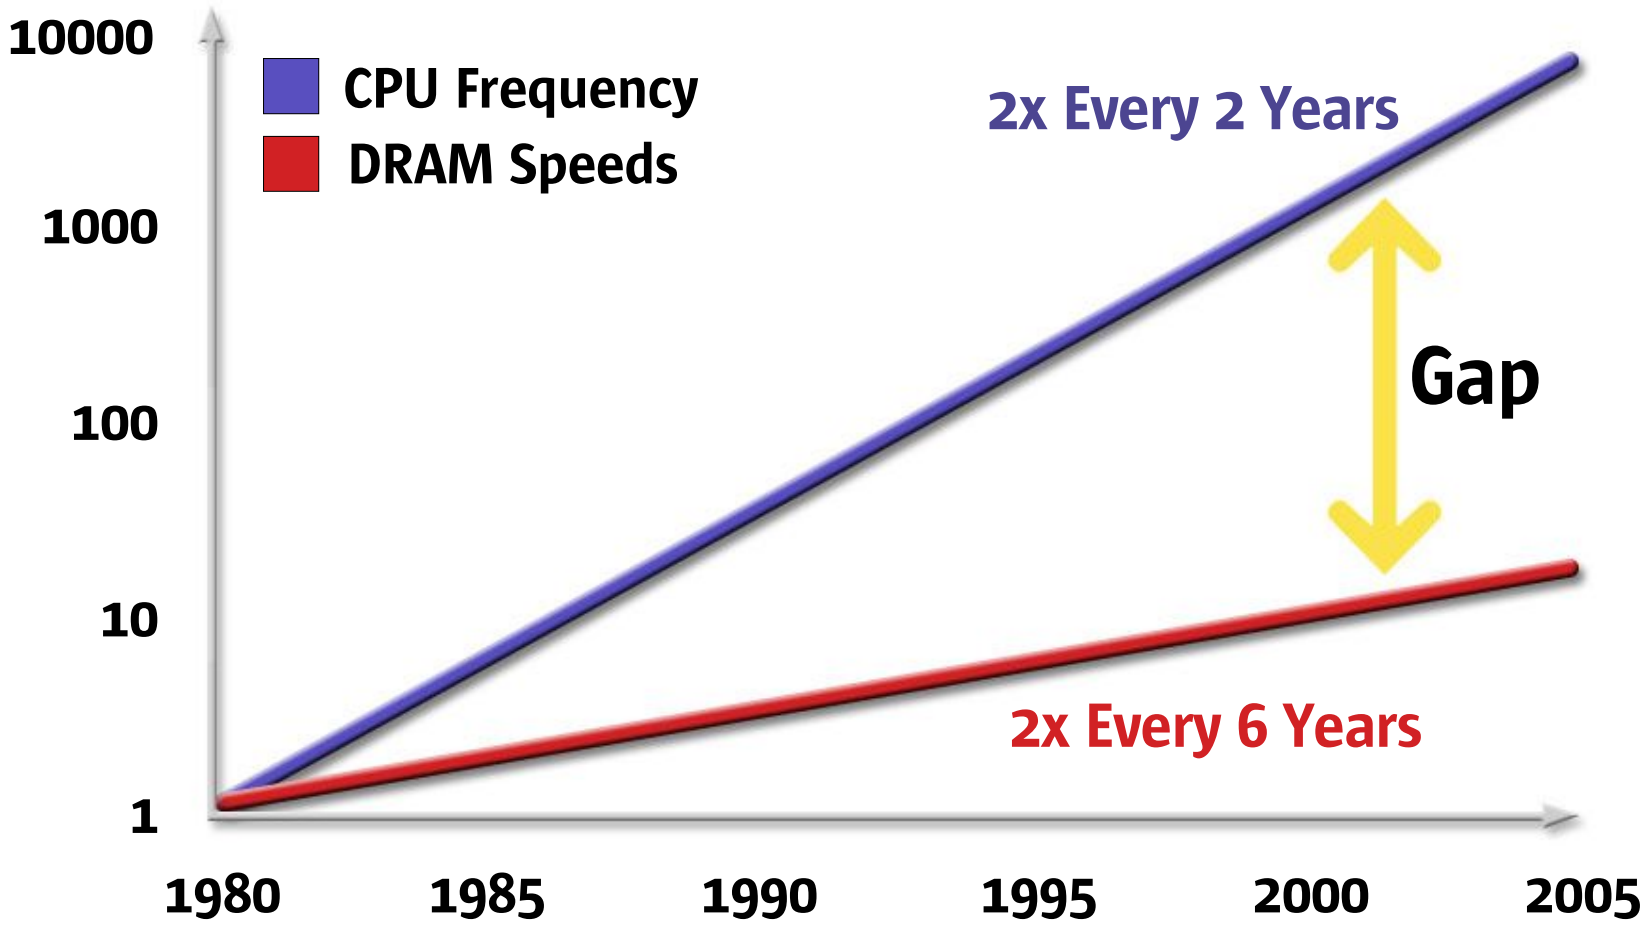

Comparing Modern CPU Design Techniques

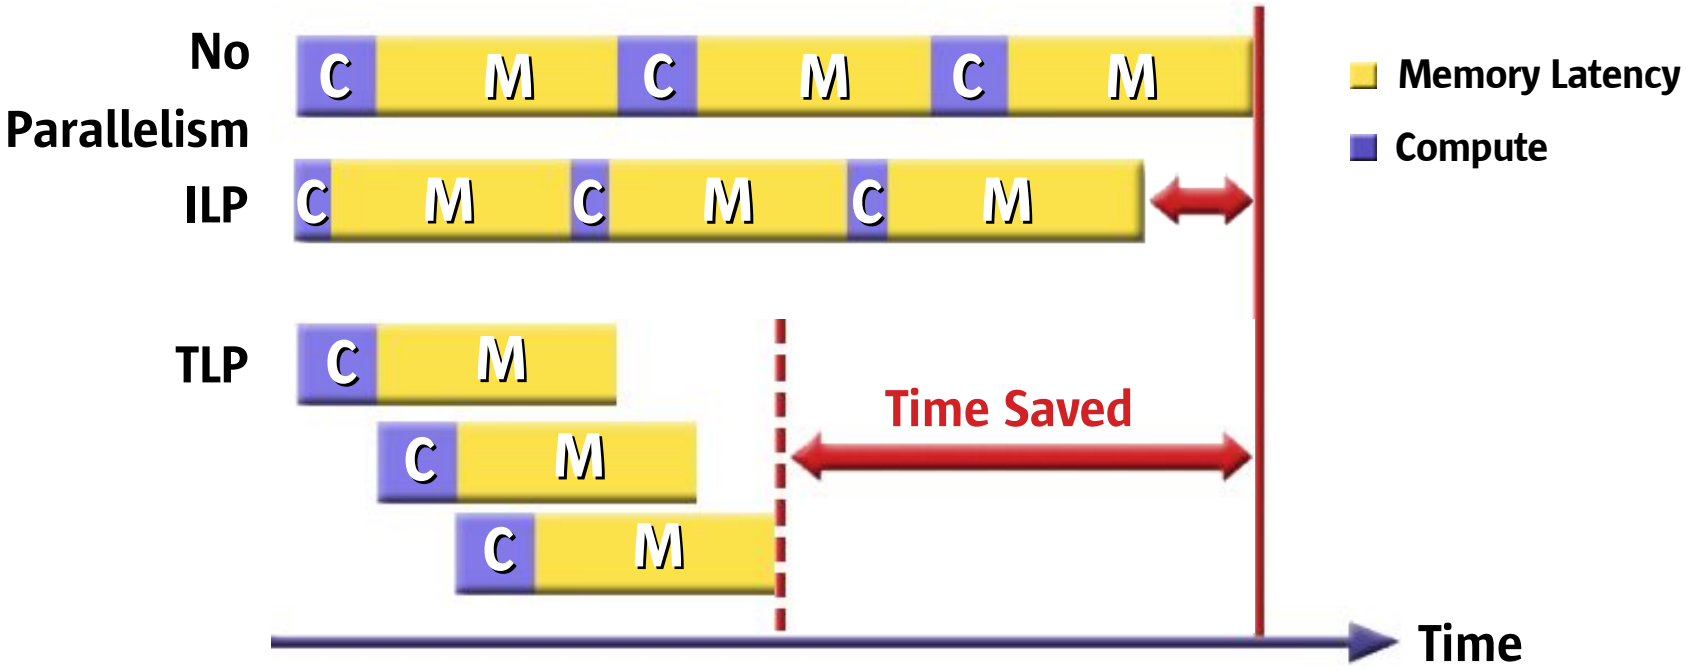

- ILP Offers Limited Headroom
- TLP Provides Greater Performance Efficiency

How Can CMT Deliver?

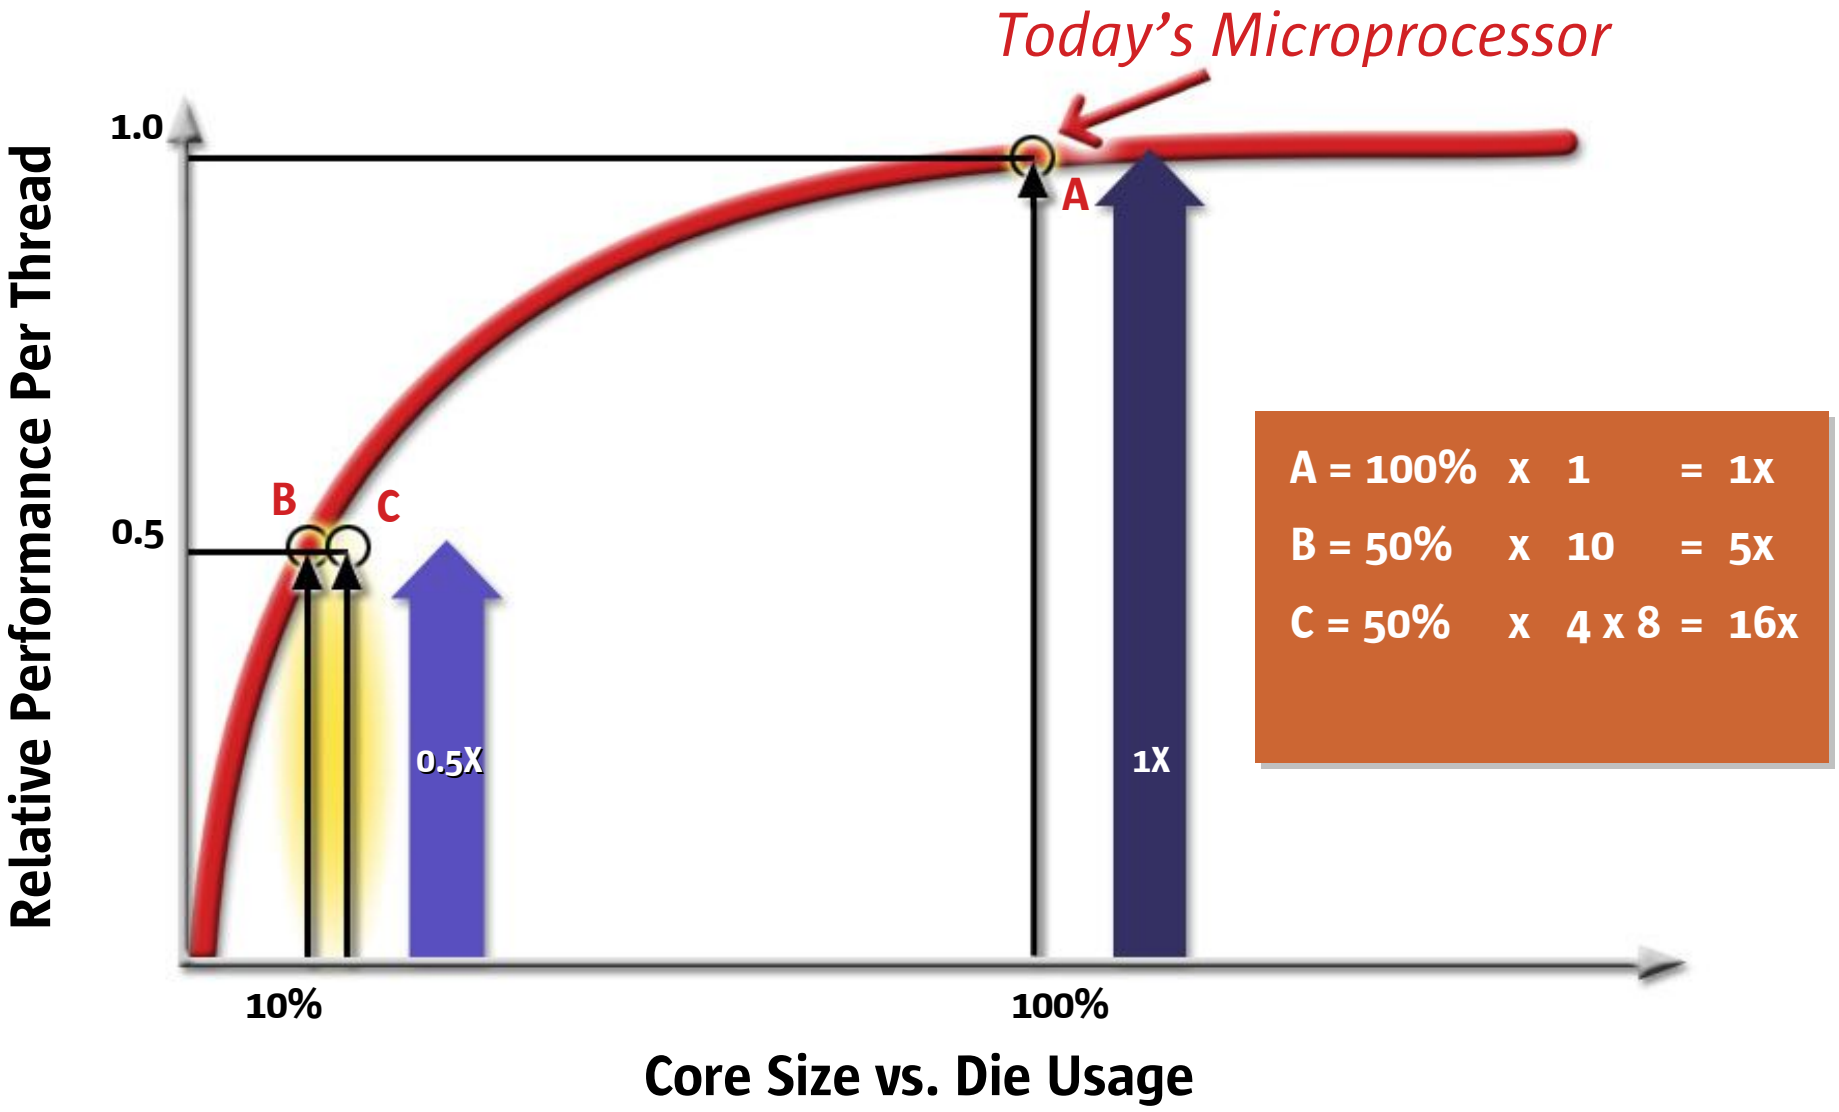

CMT – Multiple Multithreaded Cores

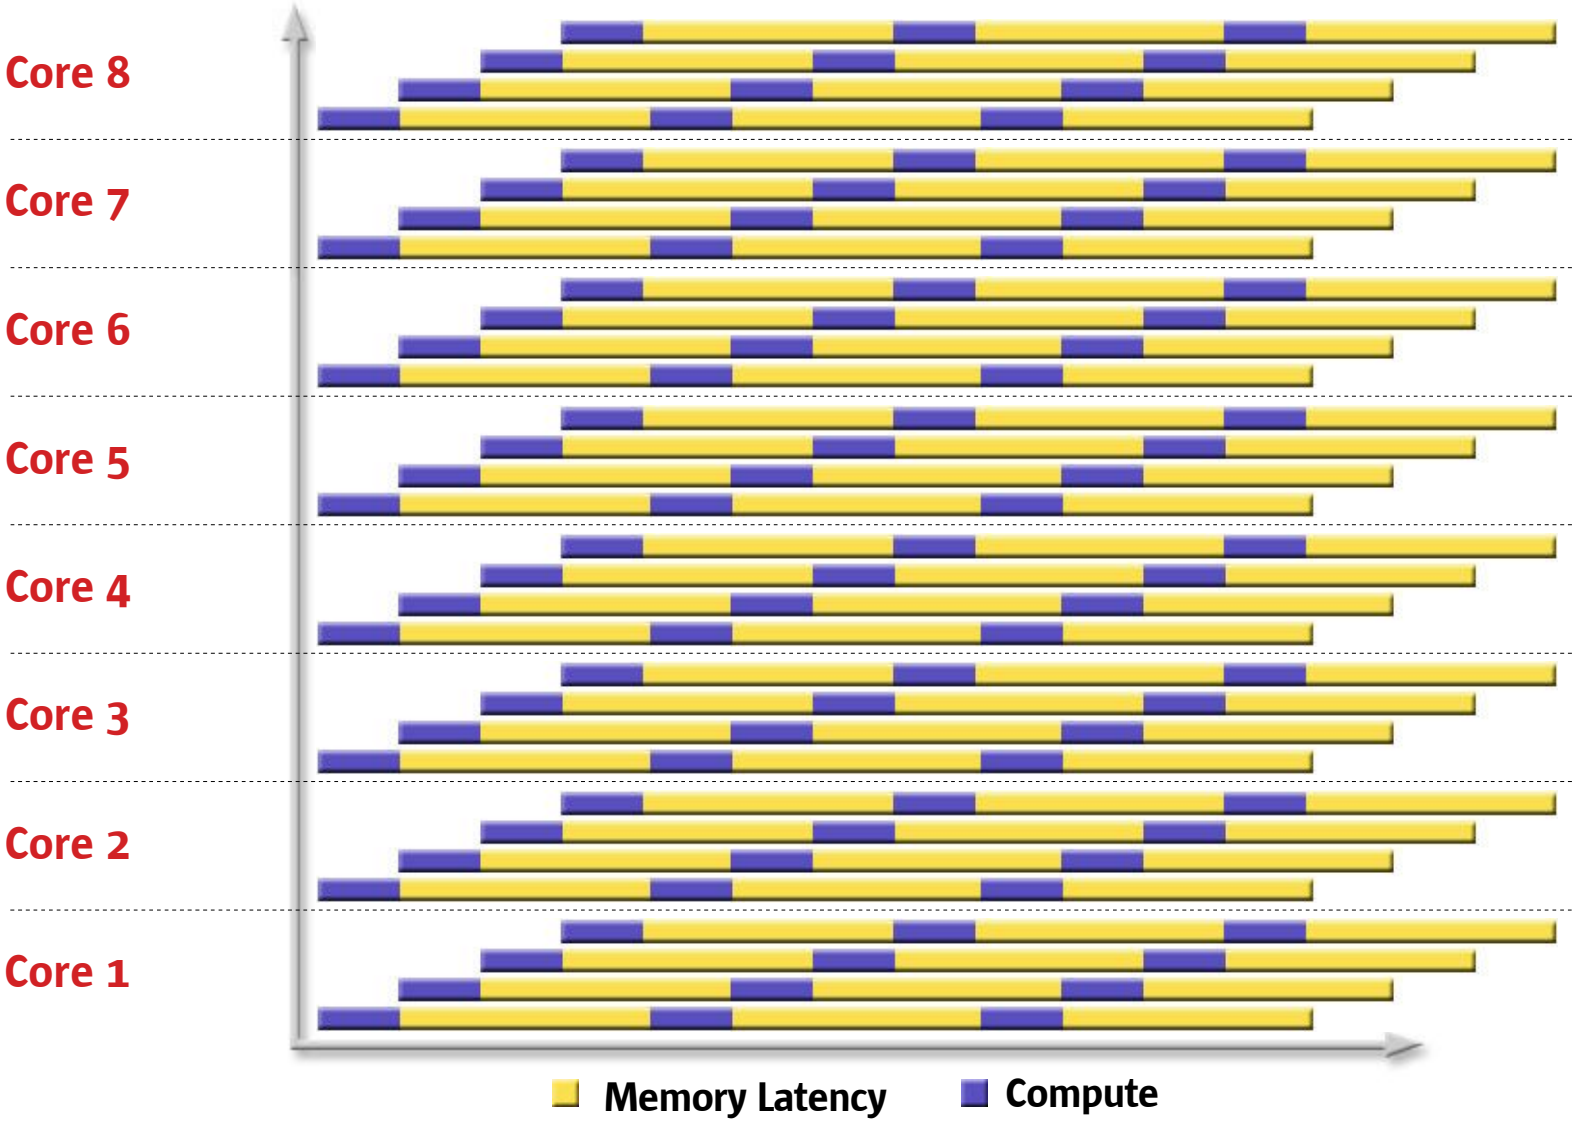

Expanding a 20-Year Relationship

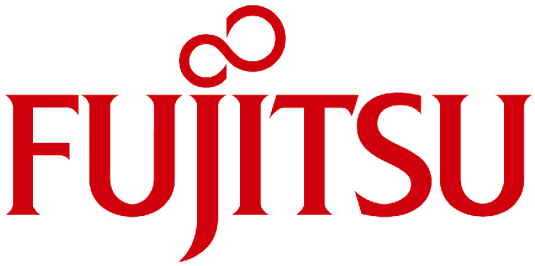

- Will jointly develop and deliver a new generation of SPARC systems debuting in mid-2006
- An expanded distribution of both companies' existing product lines
- Expanding the Solaris community

Joining SPARC Forces for a Bright Future

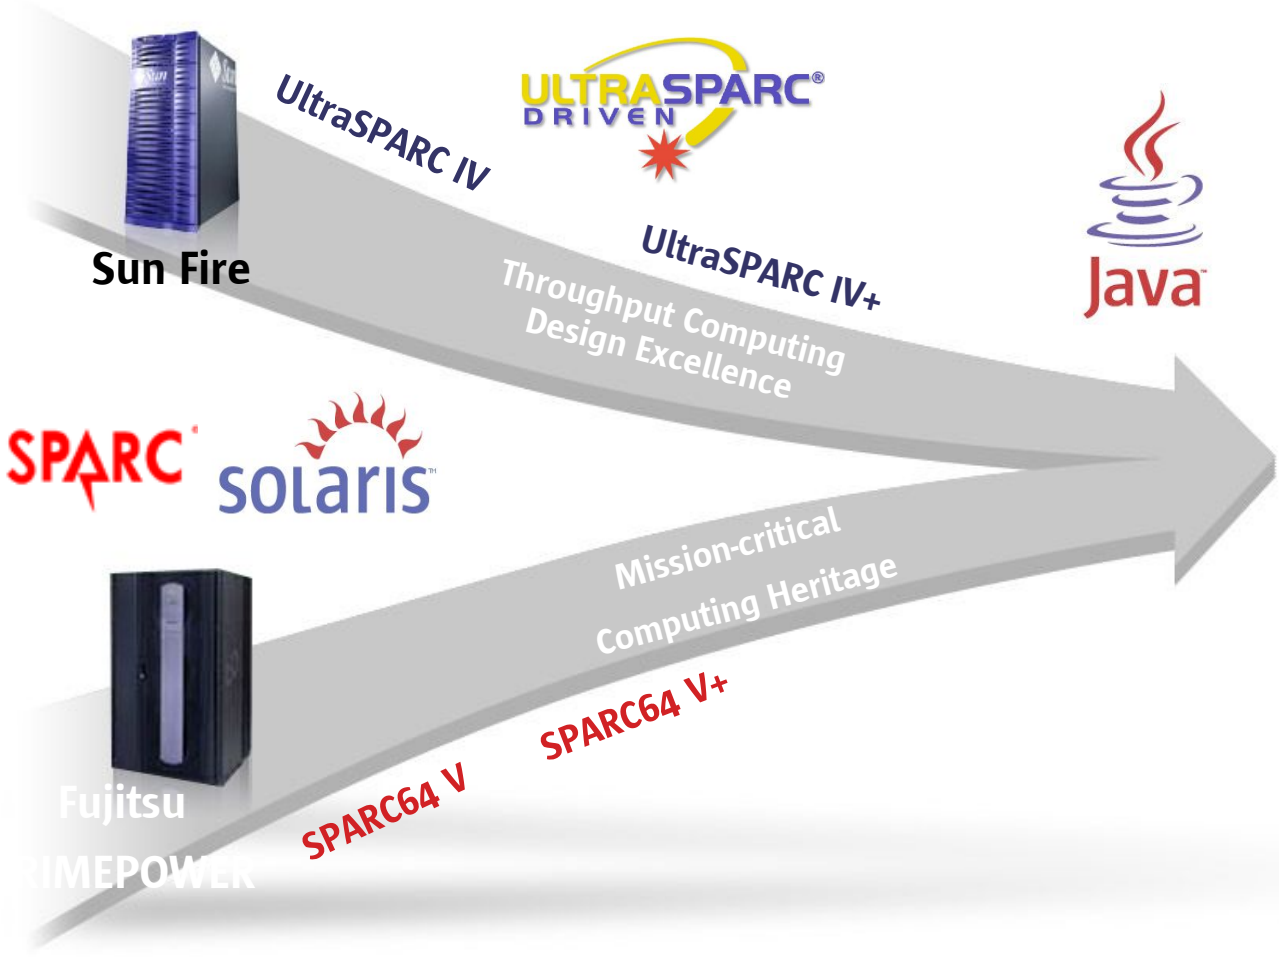

Advanced Product Line (APL)

- Optimized to address all network computing workloads
- Multiple product families (low-end, midrange, high-end)
- Systems based on SPARC V9 architecture: SPARC64 (jointly developed) and Niagara (Sun developed)

Advanced Product Line (APL)

- Complete line of SPARC-based systems to debut in 2006
  - Entry level to mission-critical data center
- Jointly designed based on SPARC V9 processor architecture
  - Capitalize on history of collaborating on open SPARC standards
- Manufactured globally by both Sun and Fujitsu
  - Same specifications, standards
- Capitalize on Sun technology leadership
  - Next-generation CMT, system interconnect, SMP scalability, JES stack
- Capitalize on Solaris leadership
- Leverage Fujitsu heritage of mission-critical computing
  - Mainframe RAS
- Customers to benefit from investment protection, Solaris compatibility, etc.

CMT Acceleration: Advanced Product Line Systems from Sun and Fujitsu

	Network Centric Systems	Data Centric Systems	
		Mid-Range	High End
Processor Design and System Interconnect	Sun (Niagara)	Fujitsu (Olympus)	
Enclosure Design	Sun		Fujitsu
Operating System (Solaris)	Sun		
Manufacturing	All systems in 3 locations: Oregon, Scotland and Japan		

SPARC Processor Families

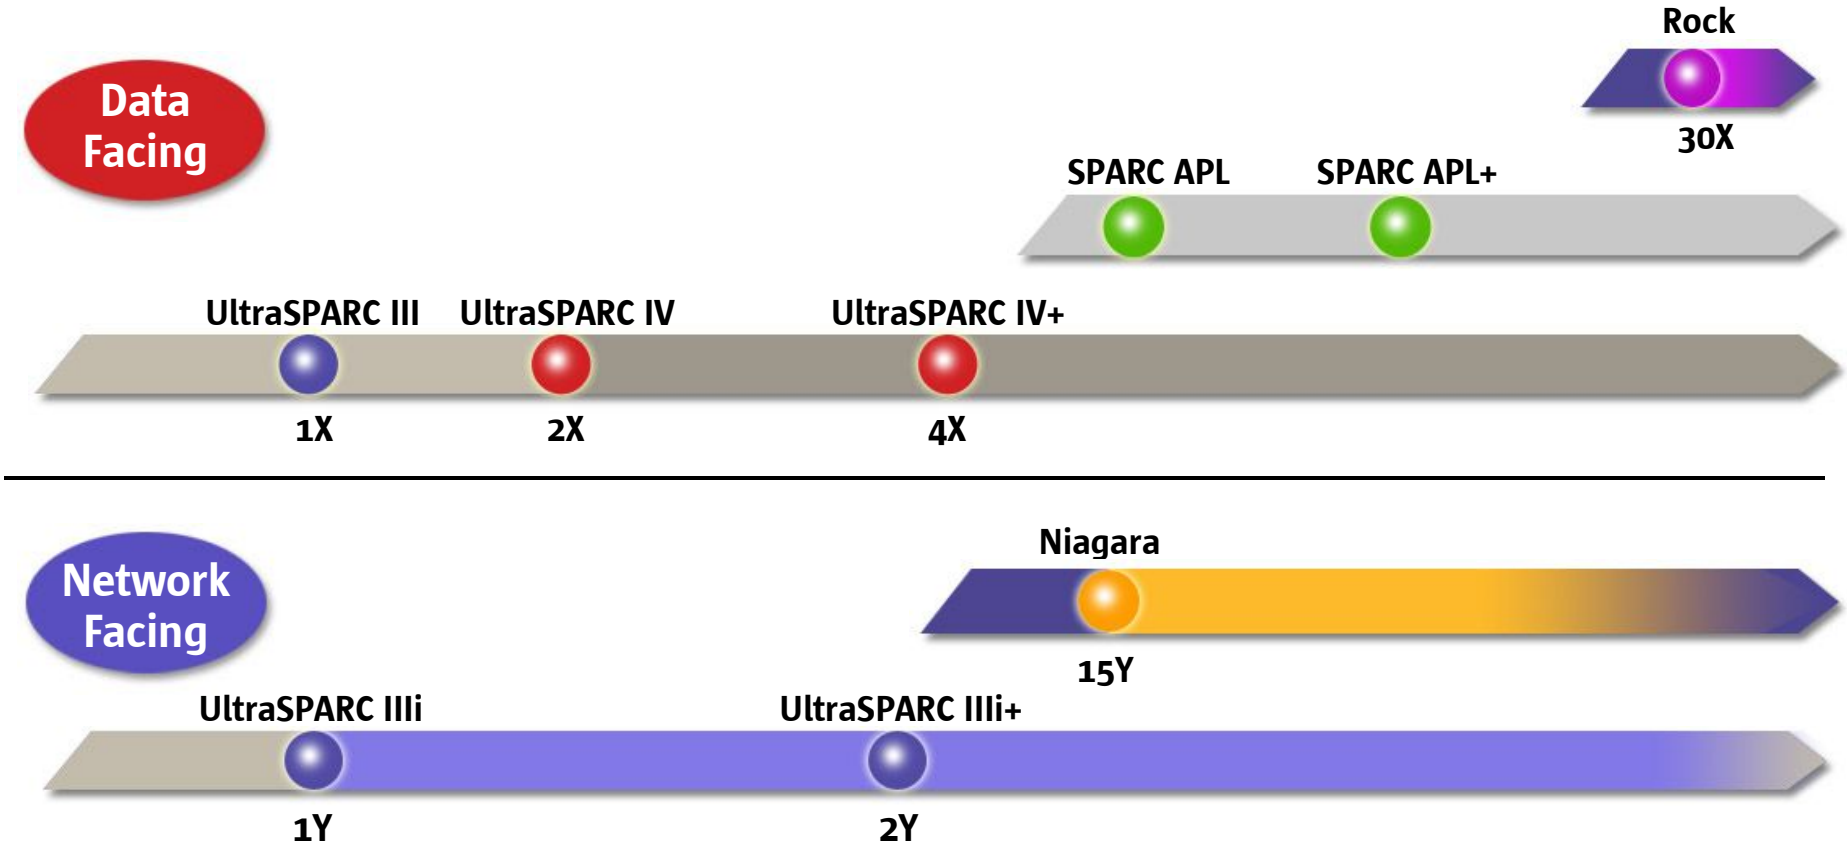

Last Revised: 08.03.04

SPARC: Optimized for Workload Variety

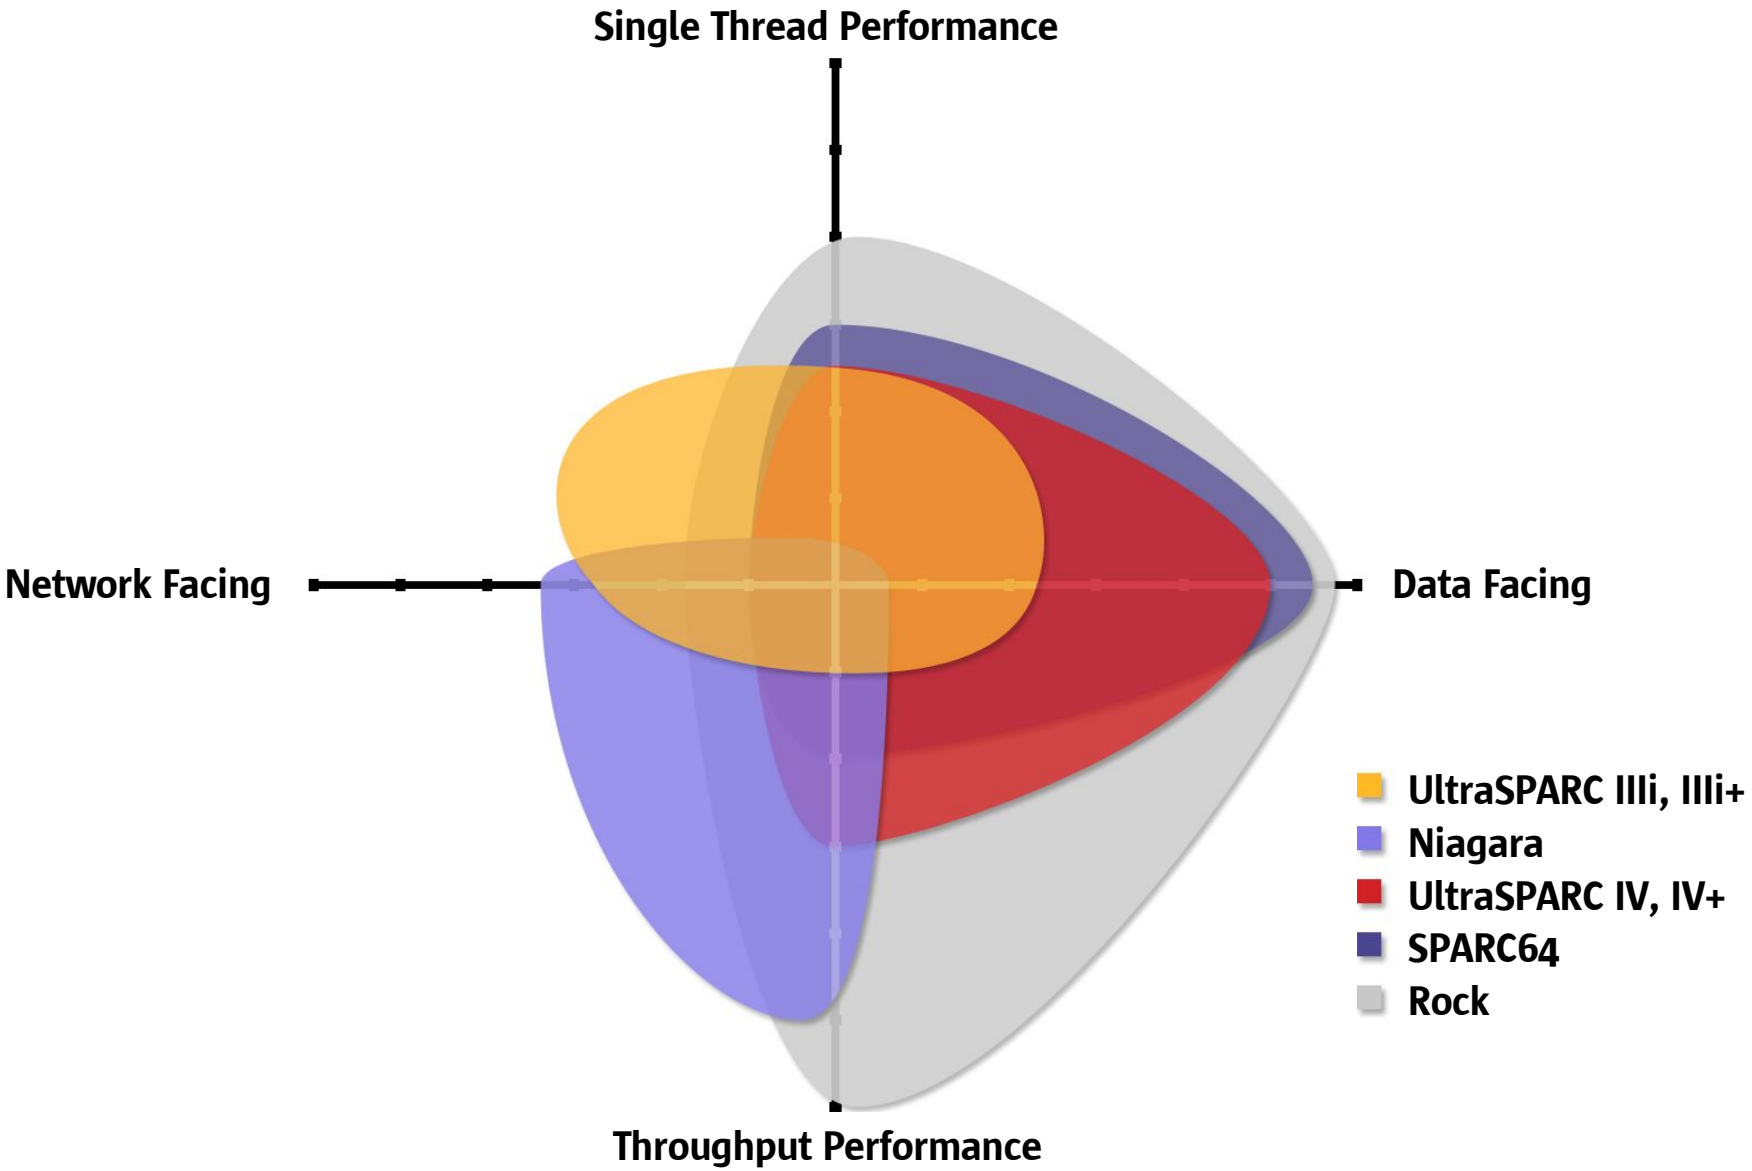

1st UltraSPARC IV

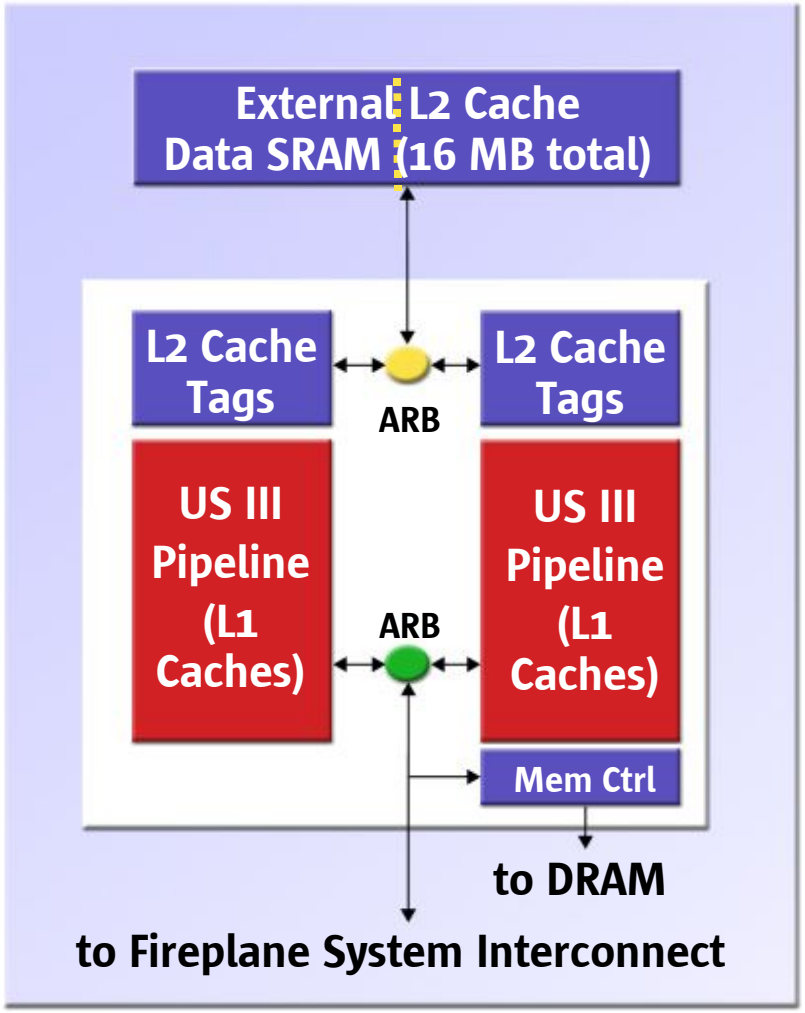

- 2 UltraSPARC III Pipelines
- Level-2 Cache
  - Logically separate
  - Physically shared external SRAM
- Shared system interface
- On-chip Memory Ctrl

UltraSPARC® IV

- Investment protection
  - UltraSPARC III/IV uniboards coexist in same system or domain
  - Leverages UltraSPARC III pipeline
- Up to 2x throughput
  - Improve performance and reliability
  - Follow-on product 4x throughput
- Execute on Sun's CMT Vision
  - Focus on throughput
  - Exploit dual threads & enhanced cache hierarchy

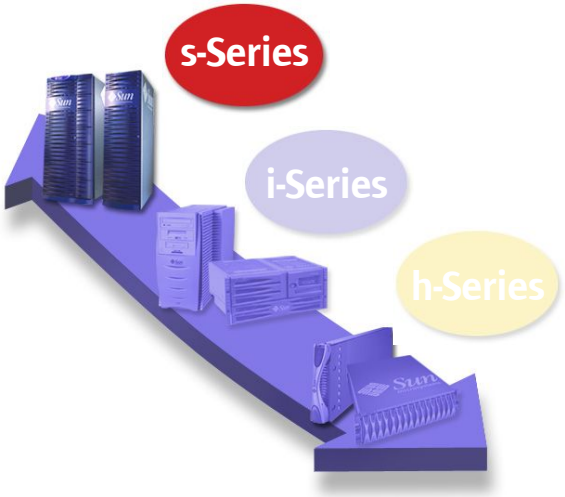

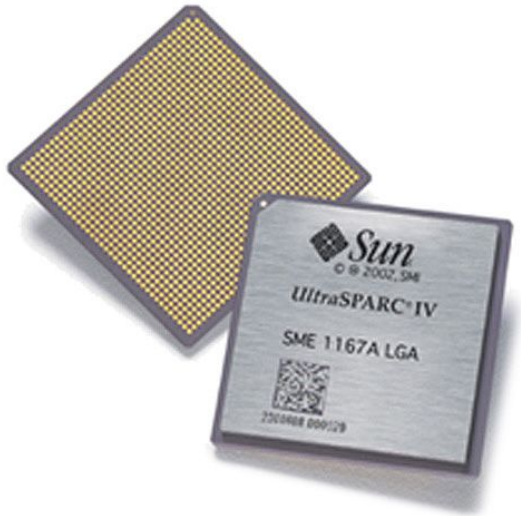

Generation 1 CMT

The Sun Fire Enterprise Server Family (4-12 way)

Solaris™ 10, 9 and 8 Operating System

Sun Fire V490

- Up to 4 UltraSPARC IV1.05 & 1.35GHz processors, 8 concurrent threads
- Up to 32 GB RAM
- 6 PCI slots, 2 FC disks
- 9.6 GB/s Sustained bandwidth
- Rack optimized
- Sun Remote System Control
- Automatic System Recovery
- Solaris Containers

Sun Fire V890

- Up to 8 UltraSPARC IV1.2 GHz & 1.35GHz processors, 16 concurrent threads
- Up to 64 GB RAM
- 9 PCI slots, up to 12 internal disks (1.7 TB)
- 9.6 GB/s Sustained bandwidth
- Tower/deskside, rack mountable
- Sun Remote System Control
- Automatic System Recovery
- Solaris Containers

Sun Fire E2900

- Up to 12 UltraSPARC IV 1.05 GHz, 1.2 GHz & 1.35GHz processors, 24 concurrent threads
- Up to 95 GB RAM
- 6 PCI slots, 2 internal disks
- 9.6 GB/s Sustained bandwidth
- Rugged, compact
- Rack optimized
- Solaris Containers

The Sun Fire Enterprise Server Family (12-72 way)

Solaris™ 10, 9 and 8 Operating System

Sun Fire E4900

- Up to 12 UltraSPARC IV 1.05 GHz, 1.2 GHz & 1.35GHz processors, 24 concurrent threads
- Up to 96 GB RAM
- 16 PCI slots
- 1 or 2 Domains & Solaris Containers
- 9.6 GB/s Sustained bandwidth
- Rack mount or deskside
- SRS Net Connect

Sun Fire E6900

- Up to 24 UltraSPARC IV 1.05 GHz, 1.2 GHz & 1.35GHz processors, 48 concurrent threads
- Up to 192 GB RAM
- 32 PCI slots
- 1 to 4 Domains & Solaris Containers
- 9.6 GB/s Sustained bandwidth
- Datacenter rack
- SRS Net Connect

Sun Fire E20K

- Up to 36 UltraSPARC IV 1.05 GHz, 1.2 GHz & 1.35GHz processors, 72 concurrent threads
- Up to 288 GB RAM
- 36 hot swap PCI+ slots
- 1 to 9 Domains & Solaris Containers
- 86.4 GB/s Peak bandwidth (29 GB sustained)
- Datacenter rack
- SRS Net Connect

Sun Fire E25K

- Up to 72 UltraSPARC IV 1.05 GHz, 1.2 GHz & 1.35GHz processors, 144 concurrent threads
- Up to 576GB RAM
- 72 hot swap PCI+ slots
- 1 to 18 Domains & Solaris Containers
- 172.8 GB/s Peak bandwidth (57 GB sustained)
- Datacenter rack
- SRS NetConnect

POWER PLAY

Long-term strategic alliance with a goal to deliver the most compelling x86 platforms for scaling out compute infrastructure

- Sun will offer a broad portfolio of AMD Opteron™ products
- Sun and AMD are collaborating on the following:
  - Optimize Solaris, Linux, and the Sun Java™ platforms for AMD Opteron processors
  - Scalability above 4-way AMD Opteron processors
  - Coherent HyperTransport technology implementations
  - Creating an industry ecosystem for ISV and IHV support
- Non-Exclusive: Sun will build with Intel™ products that perform

Horizontally Scalable Components

Solaris™ Operating System

Standard Linux Distributions, Windows Certified

Sun Fire™ V20z

- Up to 2 AMD Opteron (242, 244, 248) Processors
- Up to 16 GB memory
- Dual Gigabit Ethernet ports
- Up to 2 Ultra 320 disks (36GB/10K, 73GB/10K)
- 2 PCI/X slots (1 at 133Mhz, 1 at 66MHz)
- Lights Out Management Service Processor, dedicated Ethernet port
- Solaris or Linux operating systems, Windows certified

Sun Fire™ V40z

- Up to 4 AMD Opteron (842, 844, 848) Processors
- Up to 32 GB memory
- Dual Gigabit Ethernet ports
- Dual redundant hot-swap Power supplies and Fans
- Up to 5 Ultra 320 disks (36GB/10K, 73GB/10K)
- 5 PCI-X slots
- Lights Out Management Service Processor with dedicated Ethernet port
- Solaris or Linux operating systems, Windows certified

Sun Java Workstations W1100z/W2100z

- Opteron 1 and 2 series (144,148,150, 246,248,250)
- 1-16 GB, PC3200 ECC (DDR400)
- 5xPCI-X (1 at 133Mhz, 4 at 100MHz)
- AGP8xPro for Graphics
- DVD+CD-RW Combo and DVD Burner Optical media
- UltraSCSI320 (2P) and ATA (1P) Storage
- Tower Form Factor with E-ATX motherboard
- USB 2.0, IEEE 1394, serial, parallel, Audio (AC97)
- Solaris or Linux operating systems, Windows certified

Complete High Performance Opteron Family

- Sun Fire™ V20z
  - Secures top spot on SPECweb99_SSL benchmark in 2 processor (single core) category
  - Best performance among all 2 CPU systems on SPECjAppServer2002 Dual Node benchmark
  - Shines on SPECjbb2000 benchmark – best dual processor result in 64-bit JVM category
  - World record price/performance on SPECjAppServer2002 MultipleNode running Solaris

- Sun Fire™ V40z
  - Best x86 performance on SPECweb99_SSL benchmark(1)
  - Record setting result in 4 CPU (4-thread) category on SPEC OMPM2001 benchmark

- Sun Java™ Workstation W1100z
  - World Record on OCUS Benchmark v4 for PTC Pro/ENGINEER Wildfire 2.0
  - Industry leading results on SPECviewperf 8 benchmark. Wins on most SPECviewperf workloads
  - Best BLAST results on Solaris - outperforming Dell Precision 650 workstation

- Sun Java™ Workstation W2100z
  - World Record SPEC OMPM2001 performance result for all 2 processor (2-thread) systems
  - Fastest run-time and the best Composite Score on the EnSight graphics-oriented benchmark
  - Best BLAST results on Solaris - outperforming Dell Precision 650 workstation

OPTIMIZED SYSTEM PERFORMANCE

(1) x86 category is comprised of Intel 8086, Intel 80186, Intel 80286, Intel 80386, Intel 80486, Pentium, Pentium Pro, Pentium II, Pentium III, Pentium 4, Opteron and Athlon64 processors

Galaxy: Sun's Next-Generation Opteron Systems

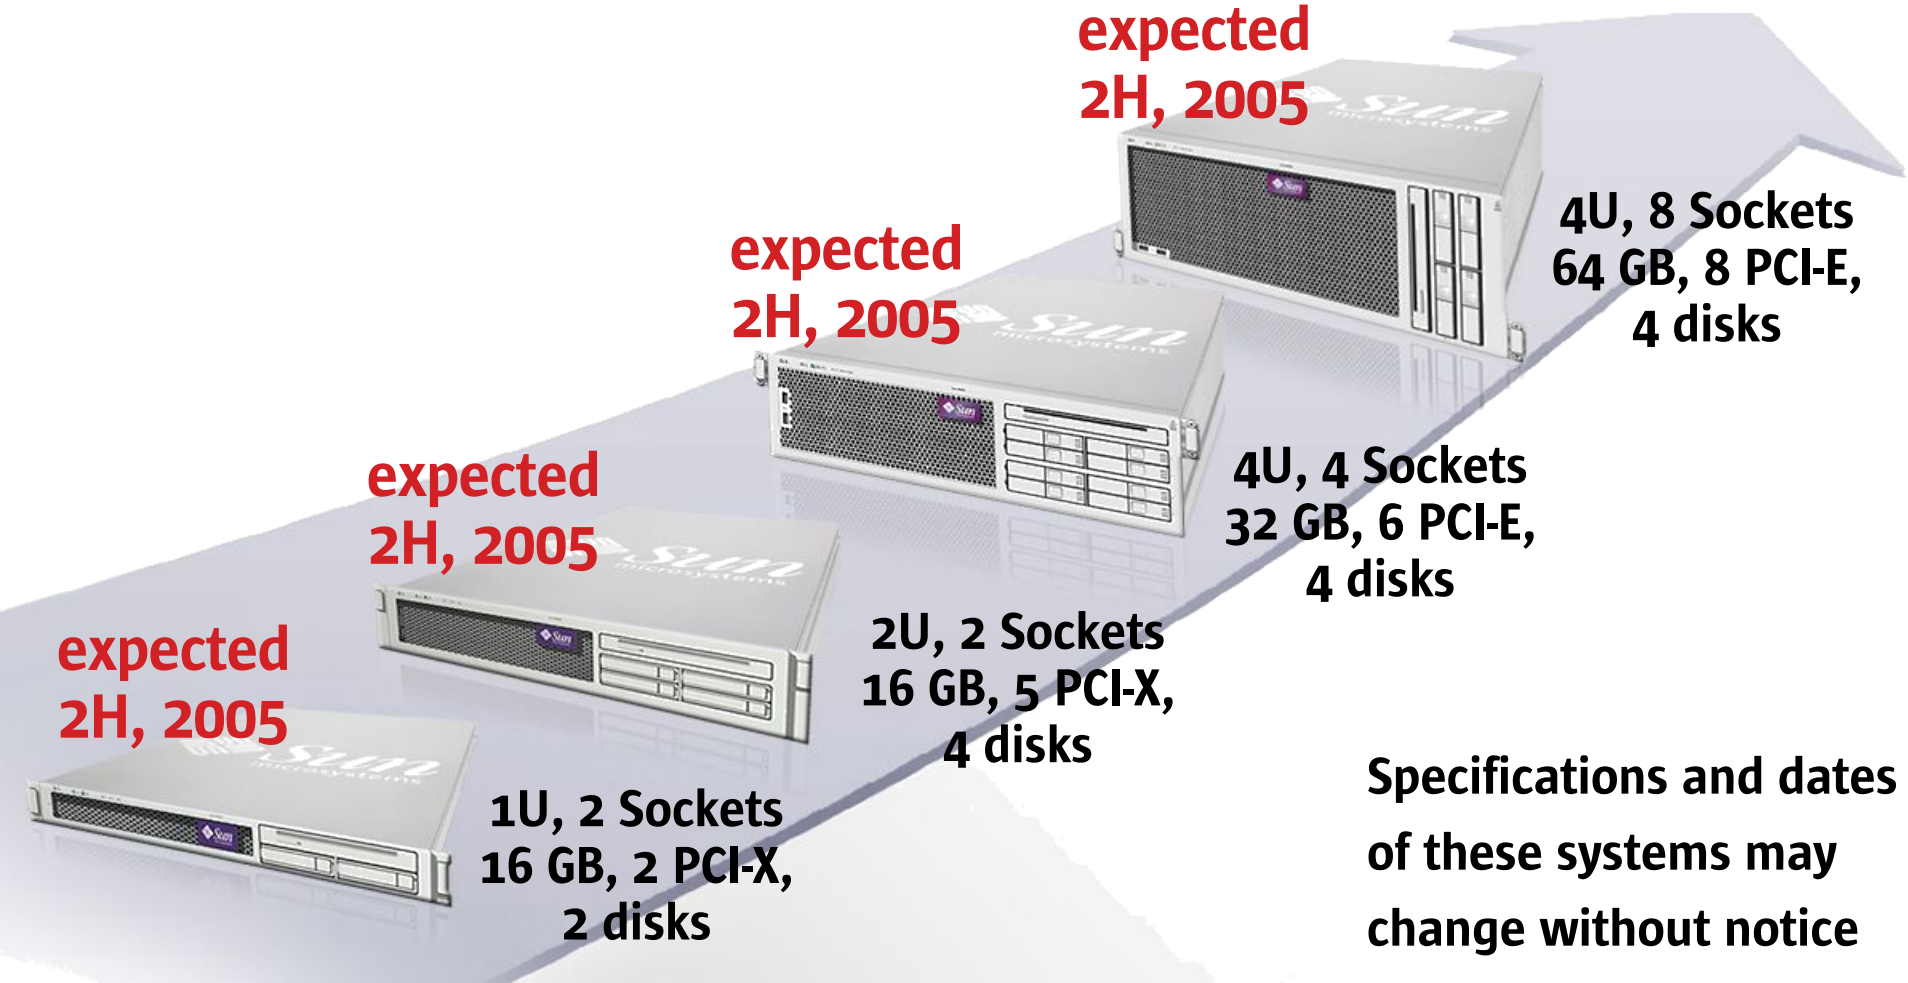

Specifications and dates of these systems may change without notice

The Infrastructure Stack

The Integrated Platform

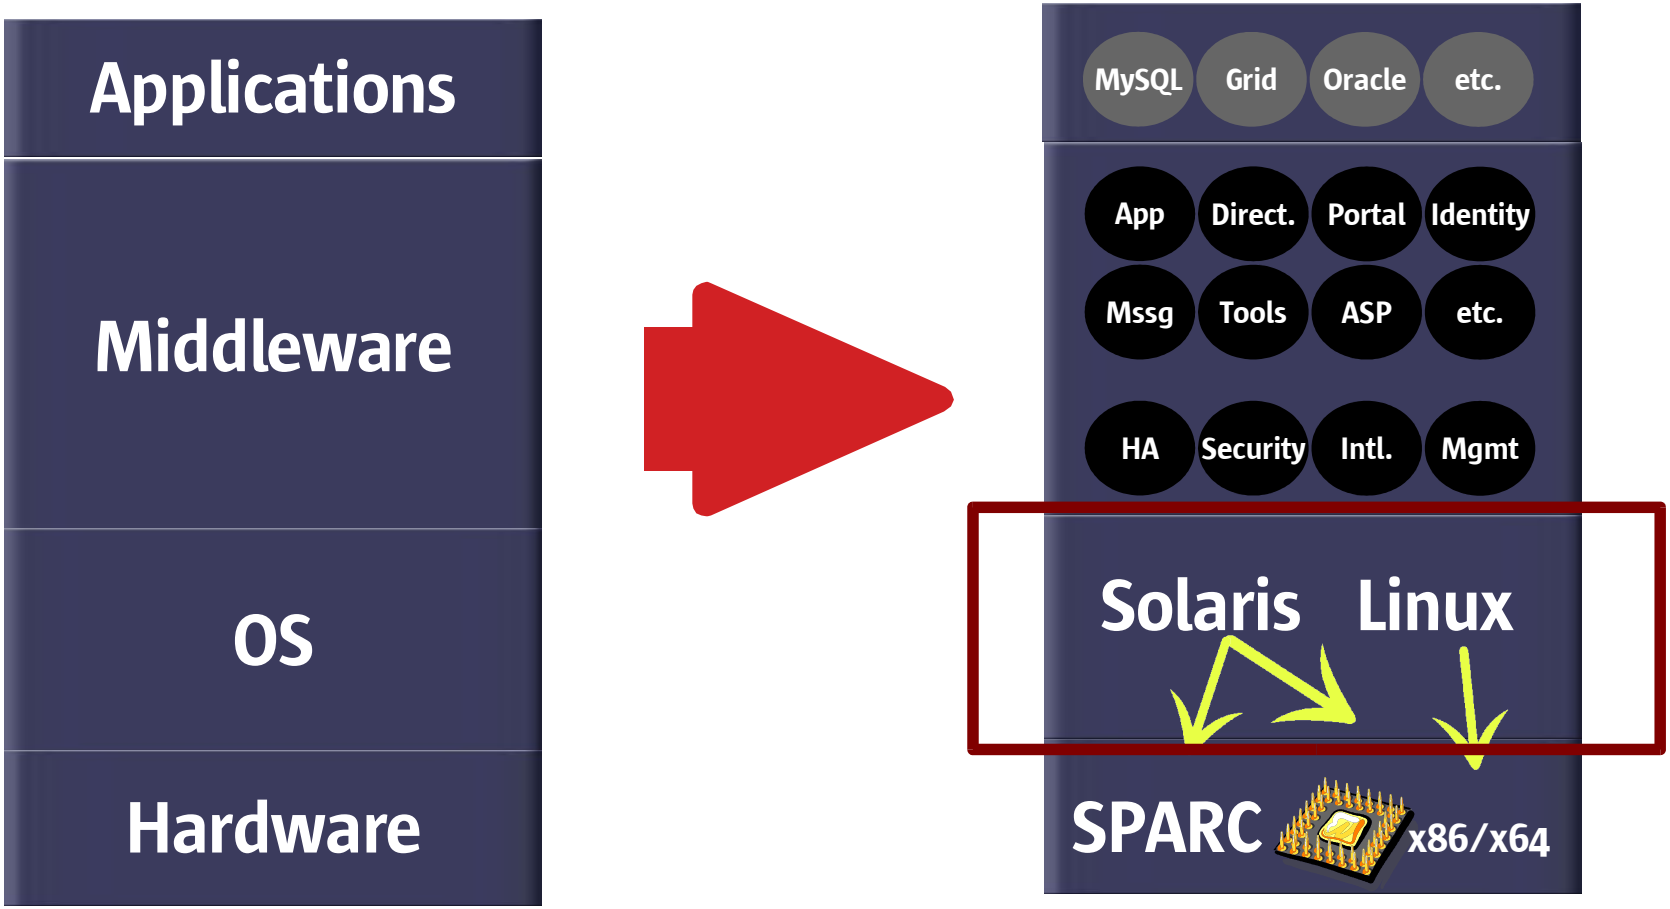

TCO

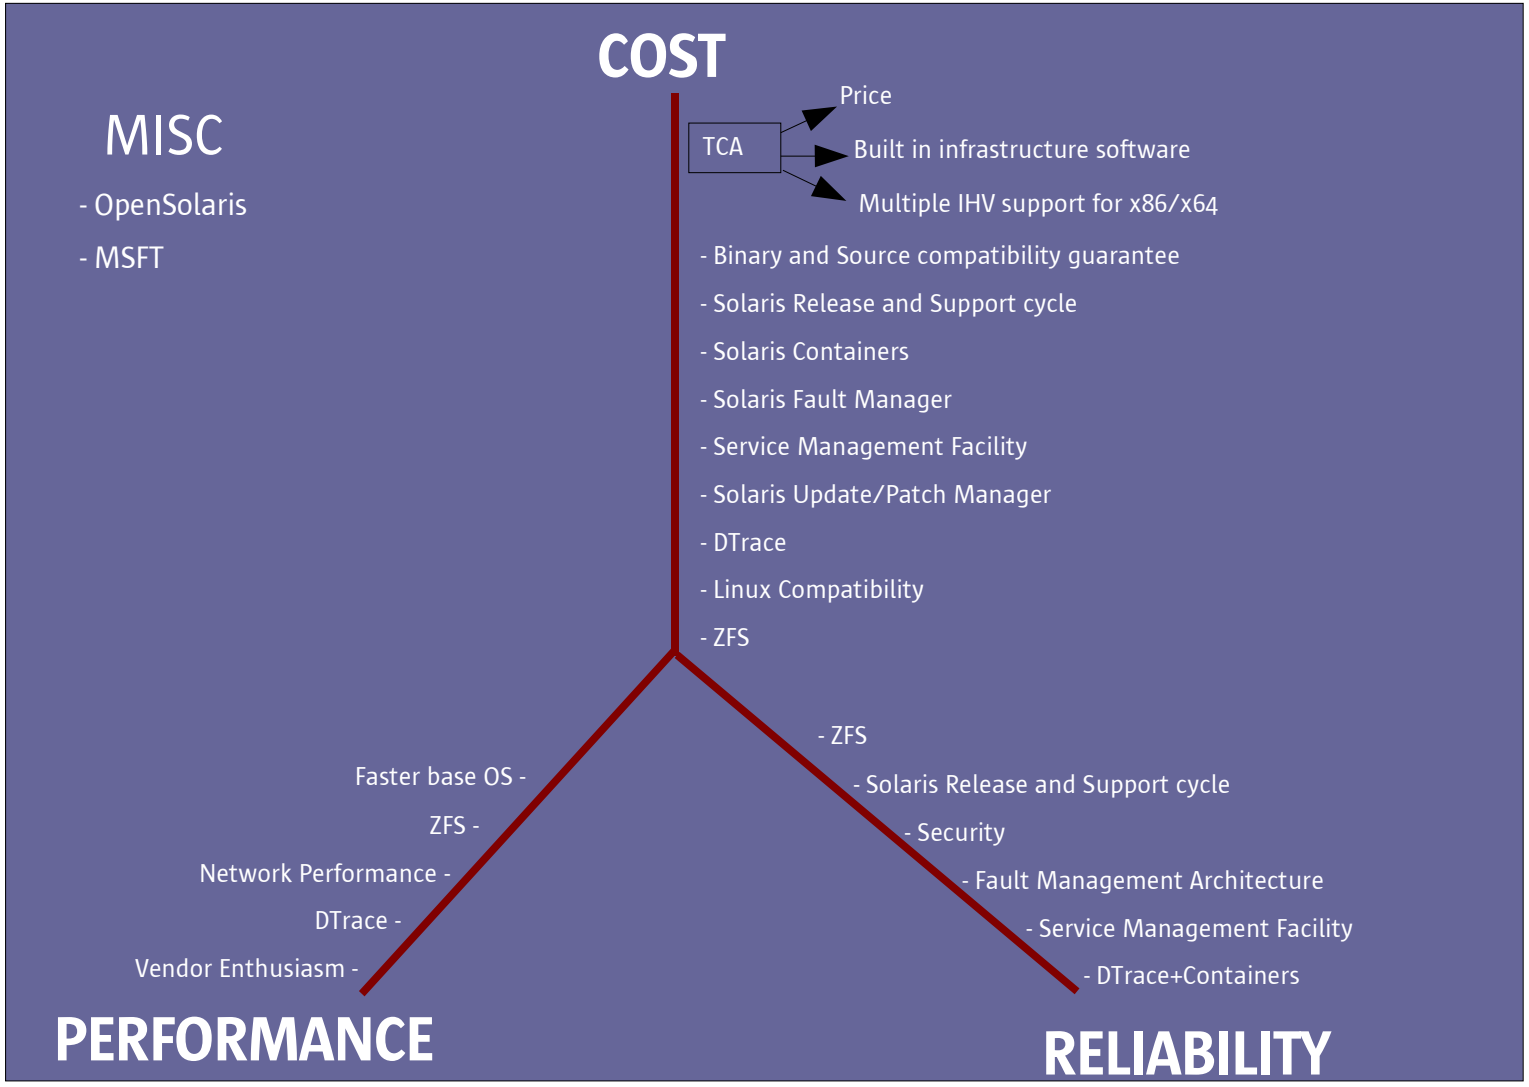
COST
MISC
- OpenSolaris
- MSFT
TCA
Price
Built in infrastructure software
Multiple IHV support for x86/x64
- Binary and Source compatibility guarantee
- Solaris Release and Support cycle
- Solaris Containers
- Solaris Fault Manager
- Service Management Facility
- Solaris Update/Patch Manager
- DTrace
- Linux Compatibility
- ZFS
Faster base OS -
ZFS -
Network Performance -
DTrace -
Vendor Enthusiasm -
PERFORMANCE
- ZFS
- Solaris Release and Support cycle
- Security
- Fault Management Architecture
- Service Management Facility
- DTrace+Containers
RELIABILITY

Solaris 10 Schedule

- Released on 01/31/05
  - SPARC, x86/x64
- Media kits shipping as of 03/04/05

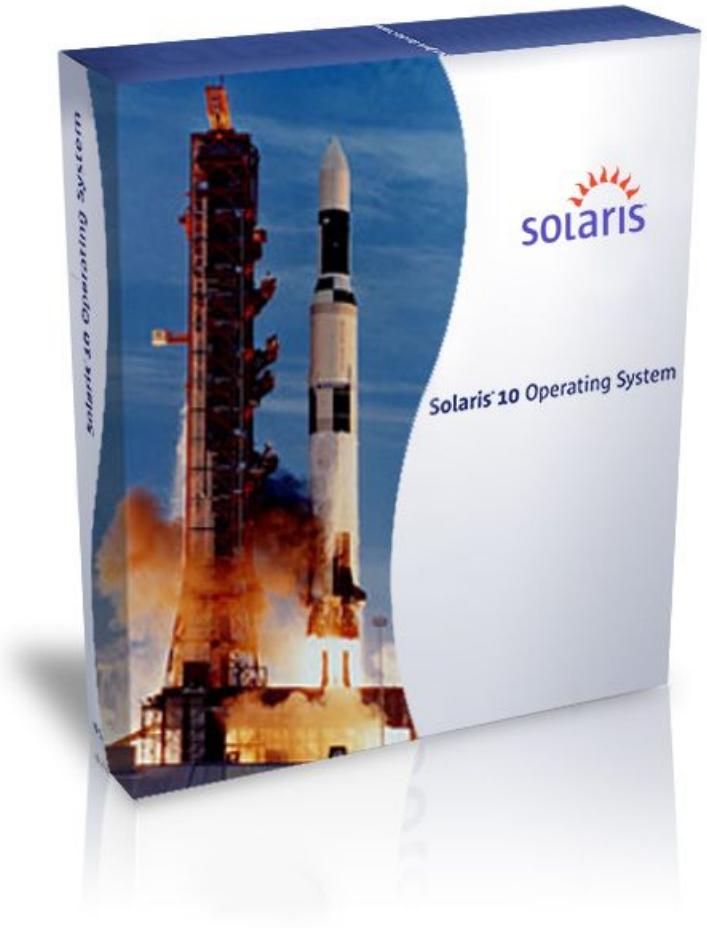

Solaris 10

900,000+ Installs

360+ Platforms

400+ New ISVs

1,100+ x86 Applications

40+ OEMs

http://sun.com/solaris/fcc/lifecycle.html

Solaris Roadmap, July '04 - June '09

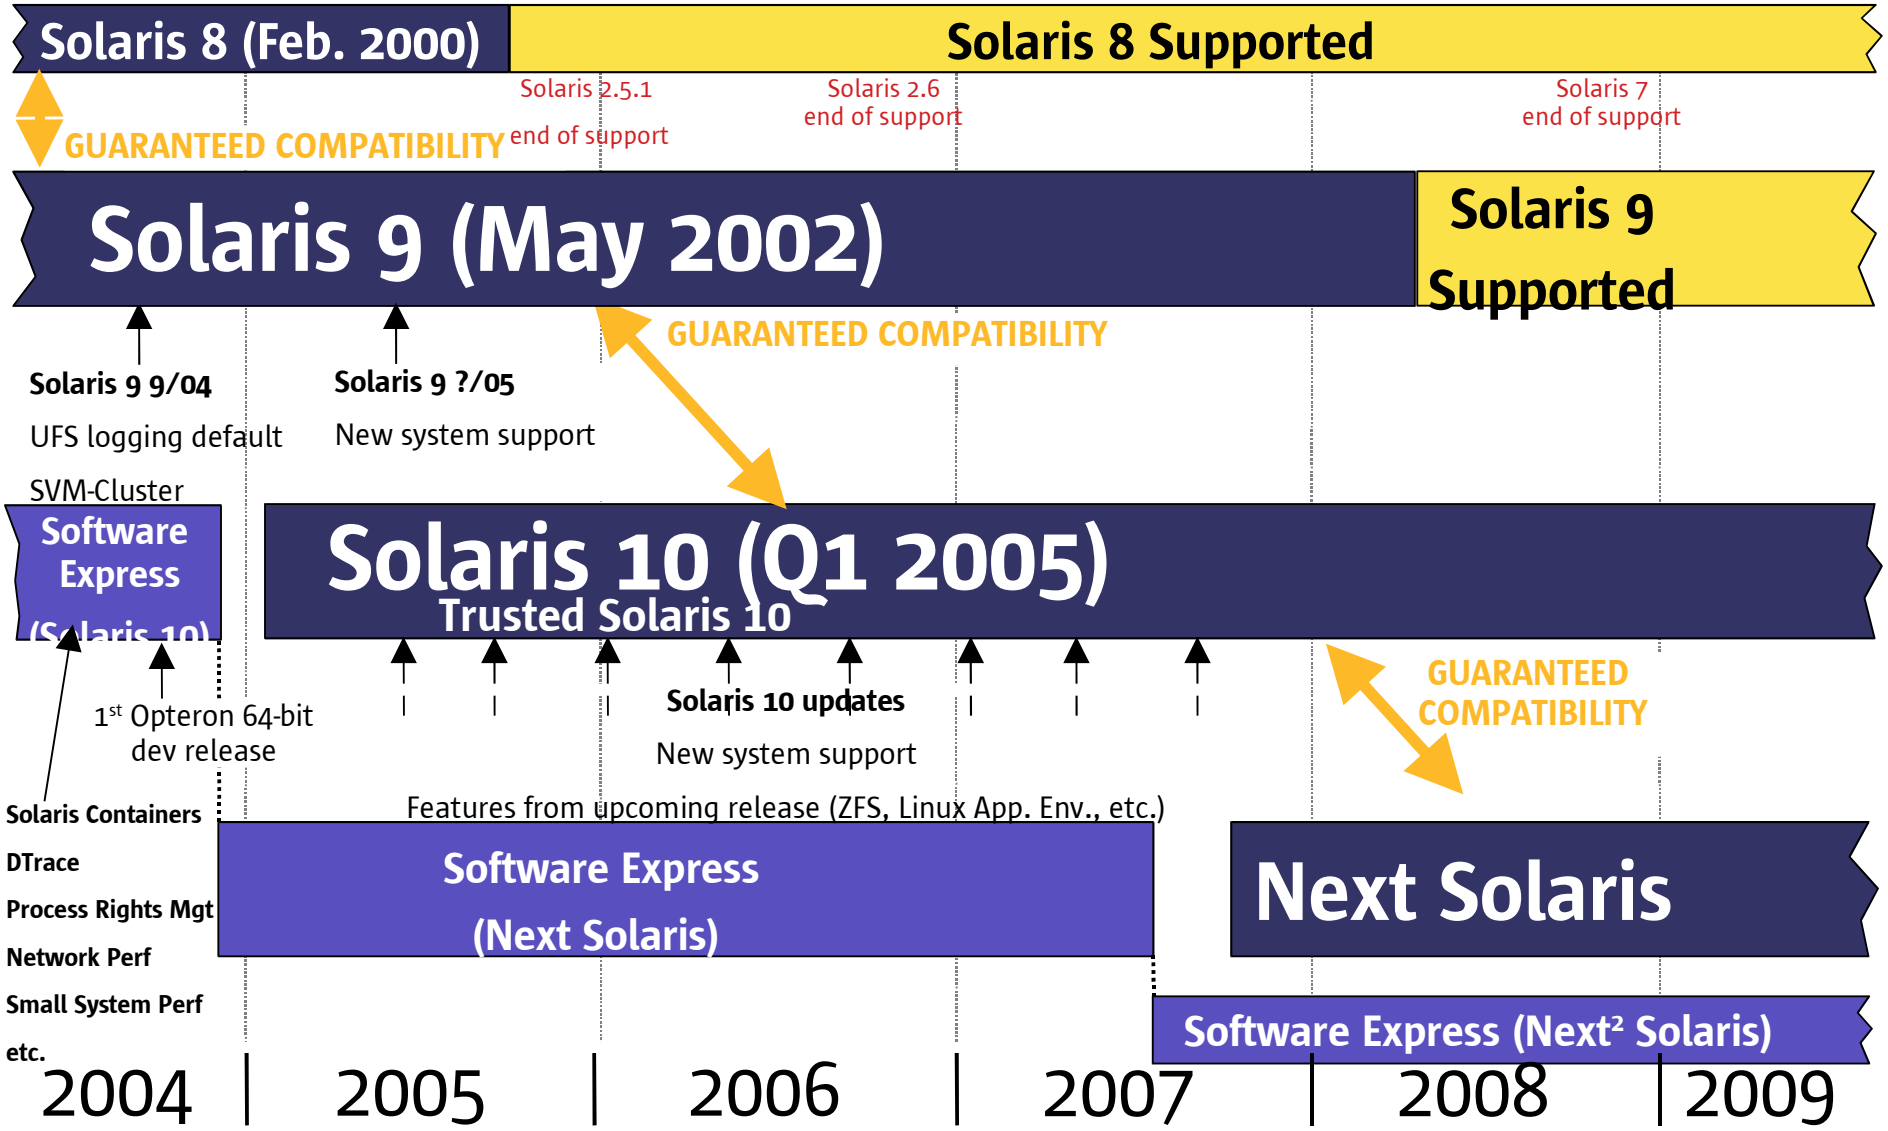

Solaris Investment Protection

Guaranteed Source Compatibility

SPARC to x86/x64

x86/x64 to SPARC

Solaris 10: A Generation Ahead

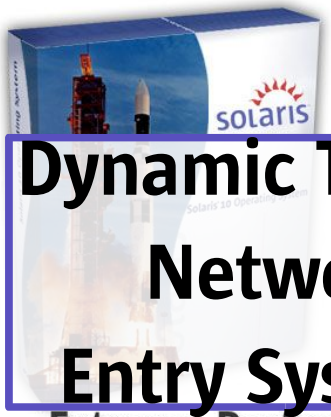

Dynamic Tracing
Network
Entry Systems

Extreme Performance

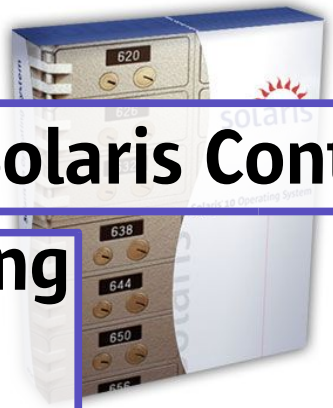

Solaris Containers

Optimal Utilization

Process Rights Mgt.
Crypto Framework
IP Filter

Unparalleled Security

Predictive Self Healing
Solaris ZFS

Relentless Availability

Linux interoperability
Next-gen SPARC
AMD Opteron

Platform Choice

Solaris Containers

Consolidation Made Simple, Safe and Secure

- Breakthrough approach to virtualization
  - Host 1,000s of applications/services on one system, with a single OS instance
  - Hardware independent
- Superior resource utilization
  - Dynamically adjust to business goals
  - Less than 1% system overhead
- Significant increase in uptime and security
  - Each service fault- and intrusion- isolated
  - Instant Restart: containers start in seconds
- Reduced costs

Solaris Containers

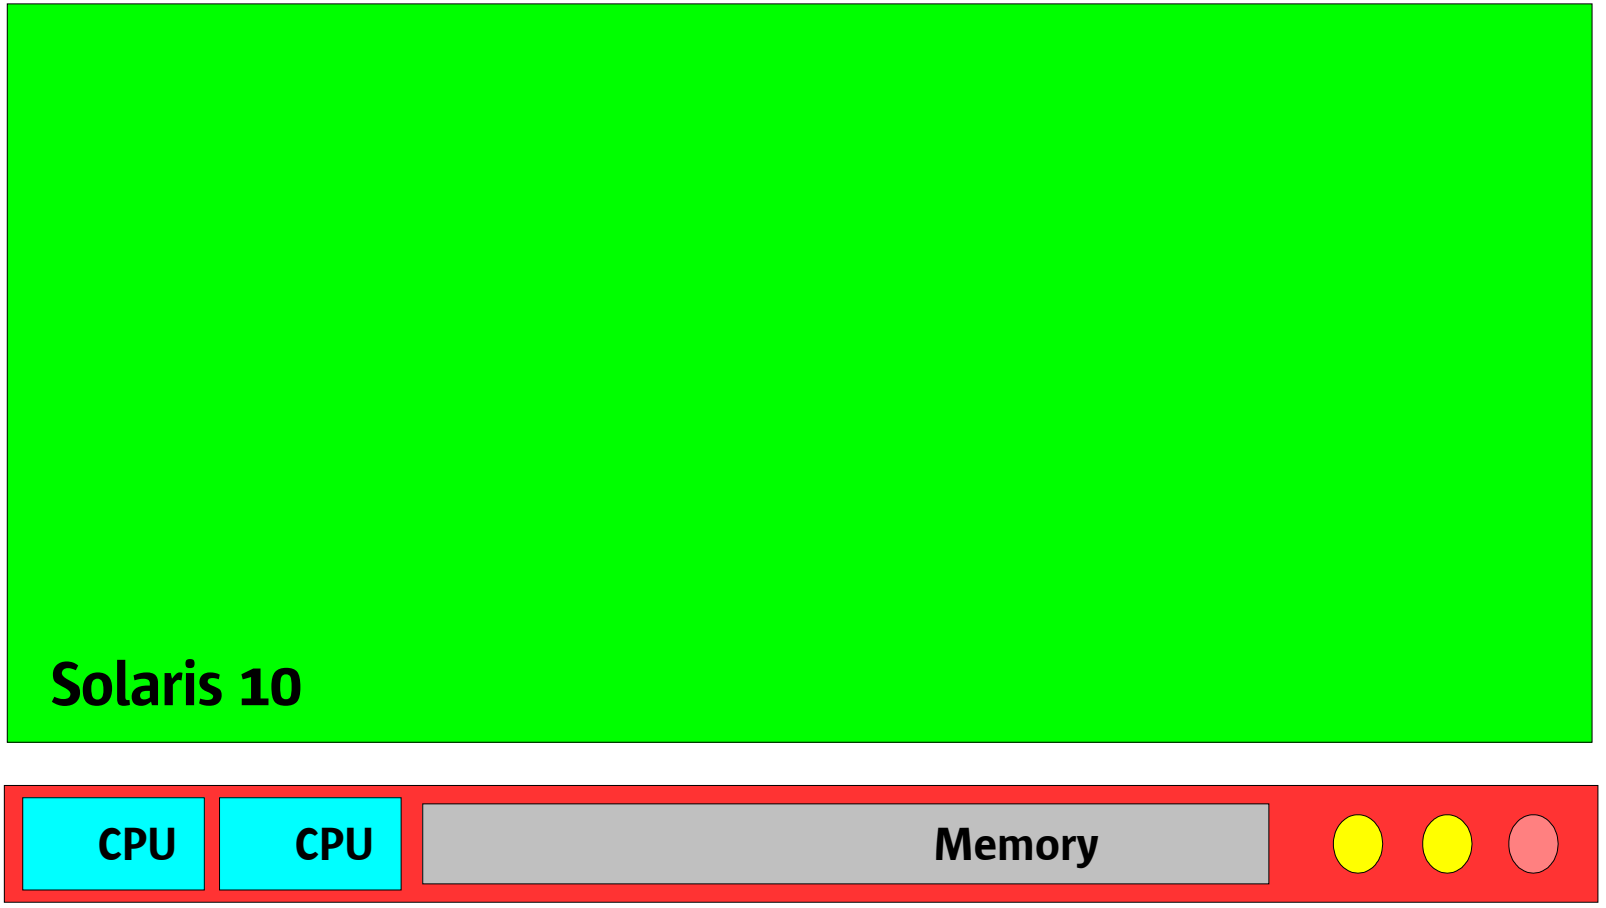

Solaris Containers

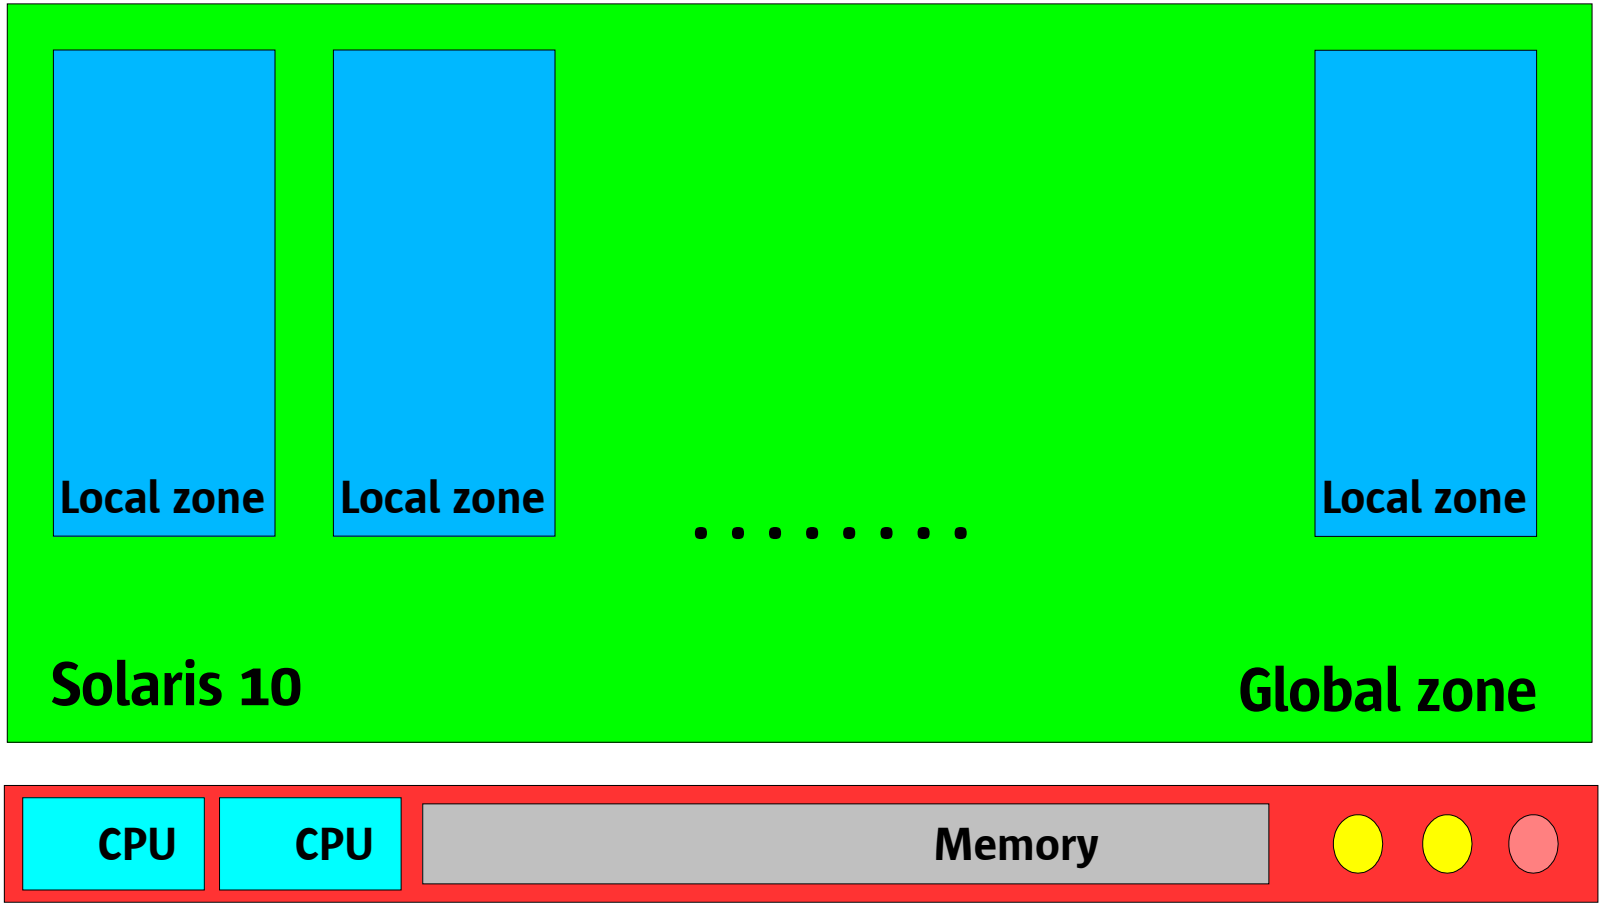

Solaris Containers

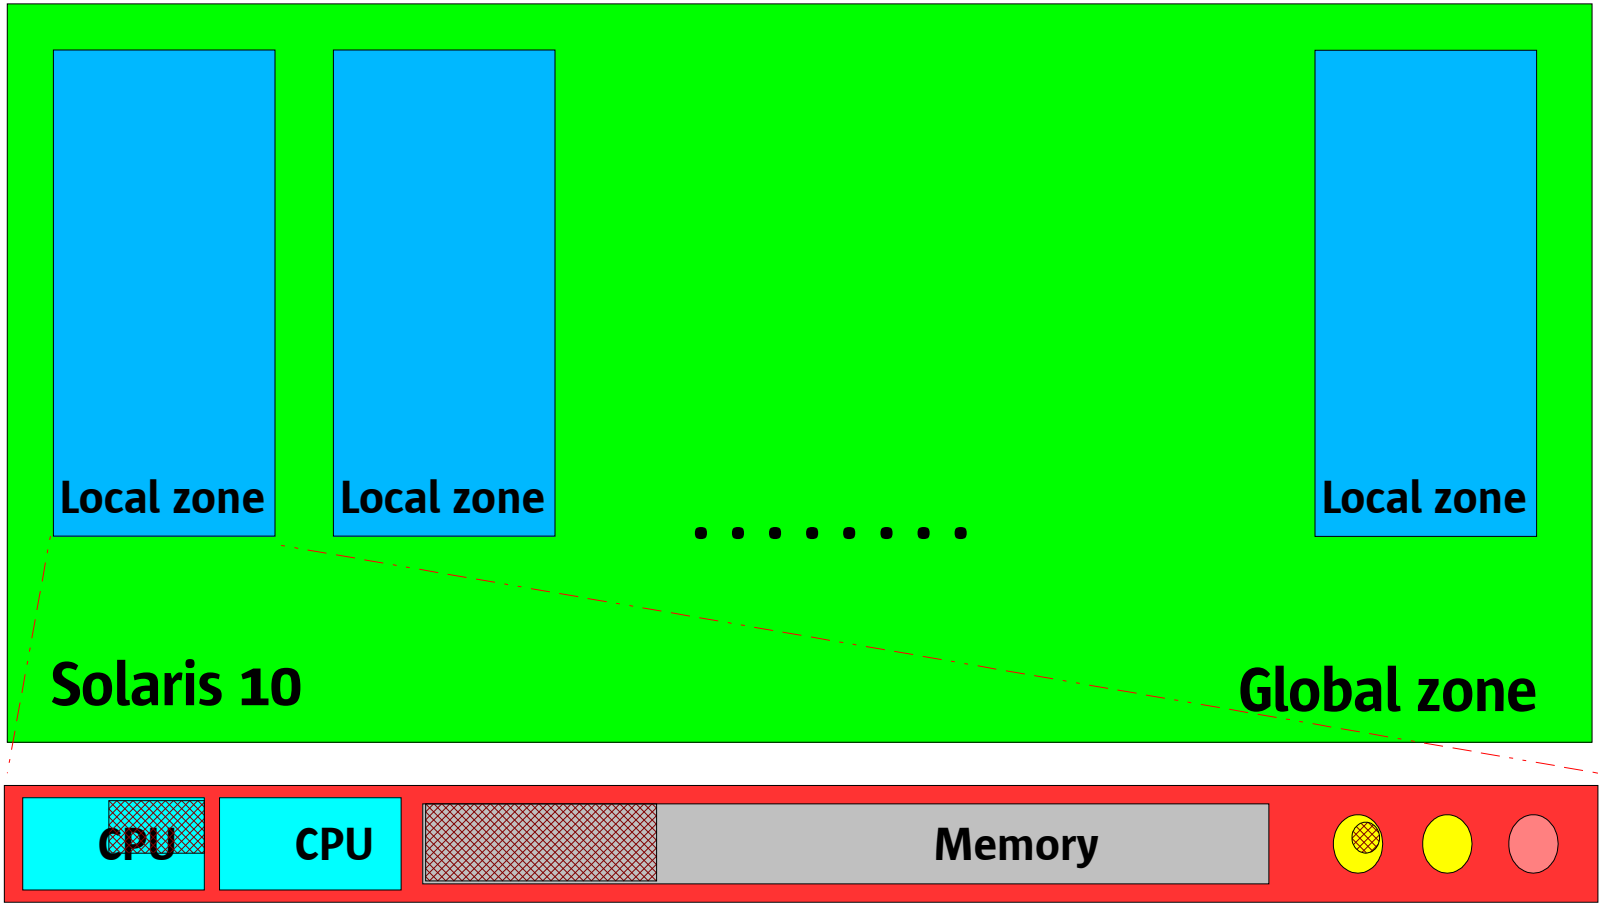

Containers

- Some examples

```
root@vitalstatistix:/# zoneadm list -cv
  ID NAME             STATUS       PATH
   0 global           running      /
   - twilight         installed    /Zones/twilight
   - espn             installed    /Zones/espn
root@vitalstatistix:/# zoneadm -z twilight boot
root@vitalstatistix:/# zoneadm list -cv
  ID NAME             STATUS       PATH
   0 global           running      /
   1 twilight         running      /Zones/twilight
   - espn             installed    /Zones/espn
root@vitalstatistix:/# zlogin twilight
[Connected to zone 'twilight' pts/6]
Last login: Thu Mar 17 21:53:11 on pts/8
Sun Microsystems Inc.   SunOS 5.10      Generic January 2005
#
```

Containers

- Some examples

```
root@vitalstatistix:/# zonename
global
root@vitalstatistix:/# pgrep -lf loop
root@vitalstatistix:/# zlogin twilight
[Connected to zone 'twilight' pts/7]
Last login: Mon Mar 21 19:01:34 on pts/7
Sun Microsystems Inc.   SunOS 5.10      Generic January 2005
# /usr/bin/nohup /twilight_loop&
2792
# Sending output to nohup.out
# pgrep -lf loop
 2792 /twilight_loop
# ^D
[Connection to zone 'twilight' pts/7 closed]
```

Containers

- Some examples

```
root@vitalstatistix:/# zonename
global
root@vitalstatistix:/# pgrep -lf loop
 2792 /twilight_loop
root@vitalstatistix:/# zlogin espn
[Connected to zone 'espn' pts/7]
Last login: Mon Mar 21 19:01:01 on pts/7
Sun Microsystems Inc.   SunOS 5.10      Generic January 2005
# pgrep -lf loop
# /usr/bin/nohup /espn_loop&
2803
# Sending output to nohup.out
# pgrep -lf loop
 2803 /espn_loop
# ^D
[Connection to zone 'espn' pts/7 closed]
```

Containers

- Some examples

```
root@vitalstatistix:/# zonename
global
root@vitalstatistix:/# pgrep -lf loop
 2803 /espn_loop
 2792 /twilight_loop
root@vitalstatistix:/# pkill loop
root@vitalstatistix:/# pgrep -lf loop
root@vitalstatistix:/#
```

Containers

- Some examples – prstat -Z

```
   PID USERNAME  SIZE   RSS STATE  PRI NICE      TIME  CPU PROCESS/NLWP
   954 ambreesh   87M  127M sleep   59    0   0:09:46 1.0% Xorg/1
  1249 ambreesh  361M  308M run     49    0   0:19:50 0.4% soffice.bin/5
  1160 ambreesh   42M   16M sleep   59    0   0:00:05 0.2% gnome-terminal/2
  1156 ambreesh   35M 9720K sleep   59    0   0:00:34 0.1% battstat-applet/1
  2716 root     4816K 4380K cpu0    49    0   0:00:00 0.1% prstat/1
  1113 ambreesh   37M   14M sleep   59    0   0:00:21 0.0% metacity/1
  1122 ambreesh   45M   20M sleep   59    0   0:00:02 0.0% nautilus/7
  1152 ambreesh   35M 9600K sleep   59    0   0:00:03 0.0% mixer_applet2/1
  2259 root     2488K 1652K sleep   49    0   0:00:00 0.0% bash/1
  1120 ambreesh   39M   14M sleep   59    0   0:00:03 0.0% gnome-panel/1
  1150 ambreesh   37M   12M sleep   59    0   0:00:07 0.0% wnck-applet/1
ZONEID    NPROC  SIZE   RSS MEMORY      TIME  CPU ZONE
     0       91 1495M  903M    89%   0:36:12 2.0% global
     2       28   89M   53M   5.2%   0:00:02 0.0% espn
     1       28   92M   55M   5.4%   0:00:03 0.0% twilight

Total: 147 processes, 474 lwps, load averages: 0.09, 0.18, 0.16
```

Solaris 10 Containers: Server Consolidation on Sun Fire V890

- Consolidated 3 diverse workloads using Solaris Containers:
  - Multiple users access J2EE application server that emulate order/inventory and B2B supply chain environment
  - Web serving workload where secure HTTP requests over a wide range of file sizes using SSL
  - A RDBMS-based DSS batch workload using complex SQL queries on a large database
- Solaris Zones used to create private environments to isolate applications for each workload
- Resource Pools used to manage CPU consumption to ensure predictable application service levels even at HIGH system utilization
- Dynamically adjust CPU resources to meet changes in workload demand

Mixed Workload Performance - Baseline

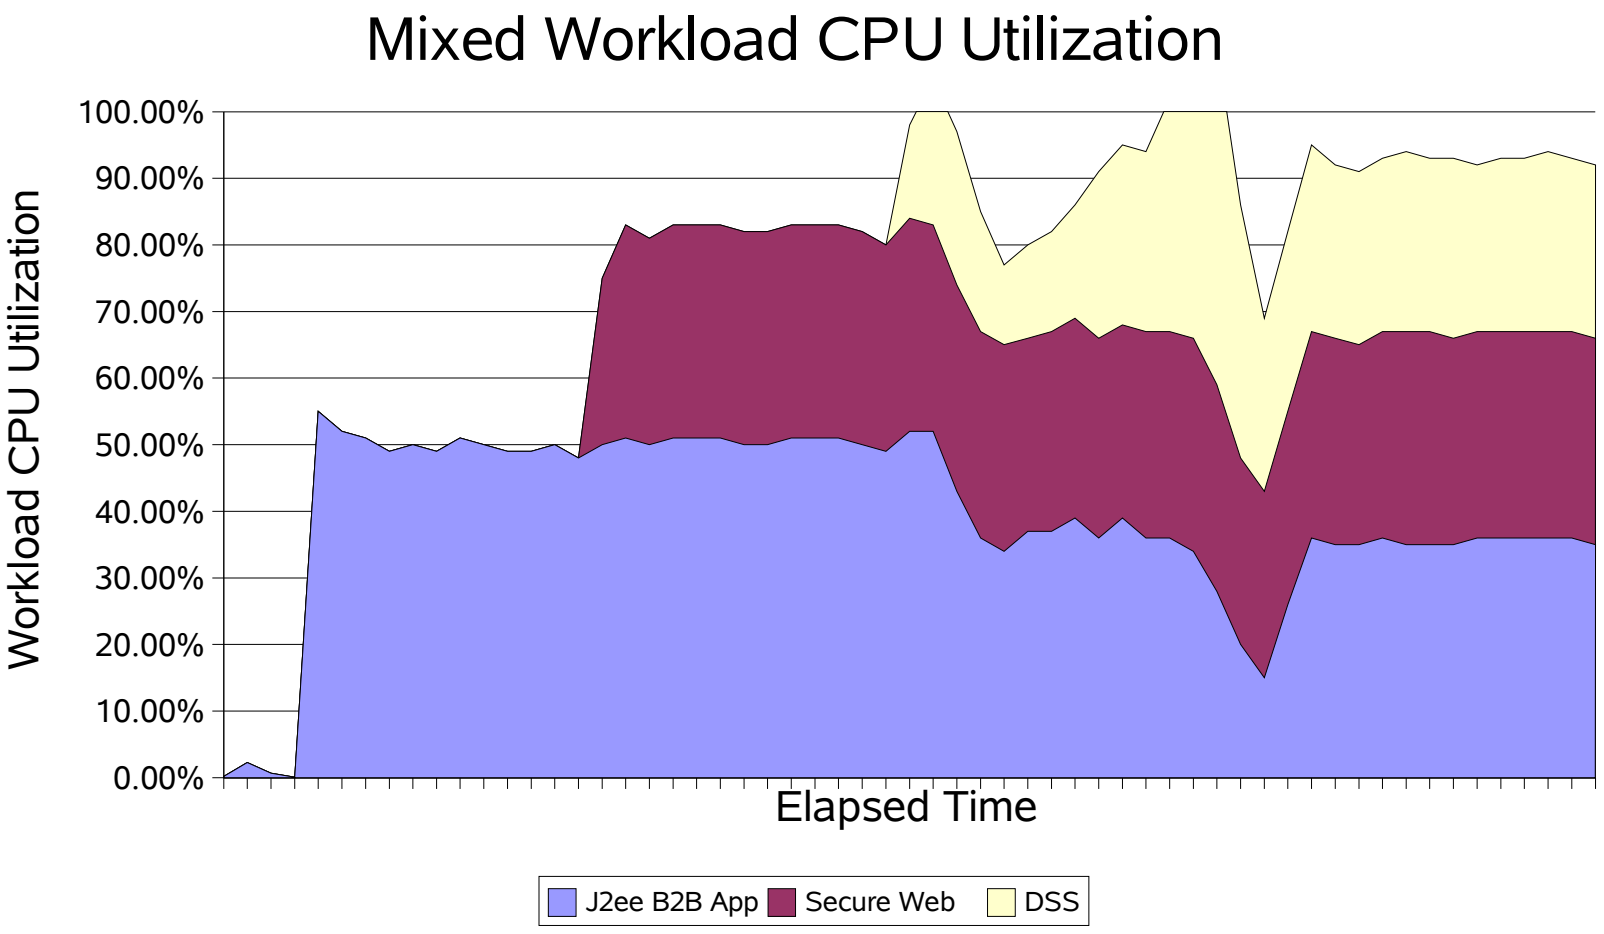

Mixed Workload Performance using Solaris 10 Containers

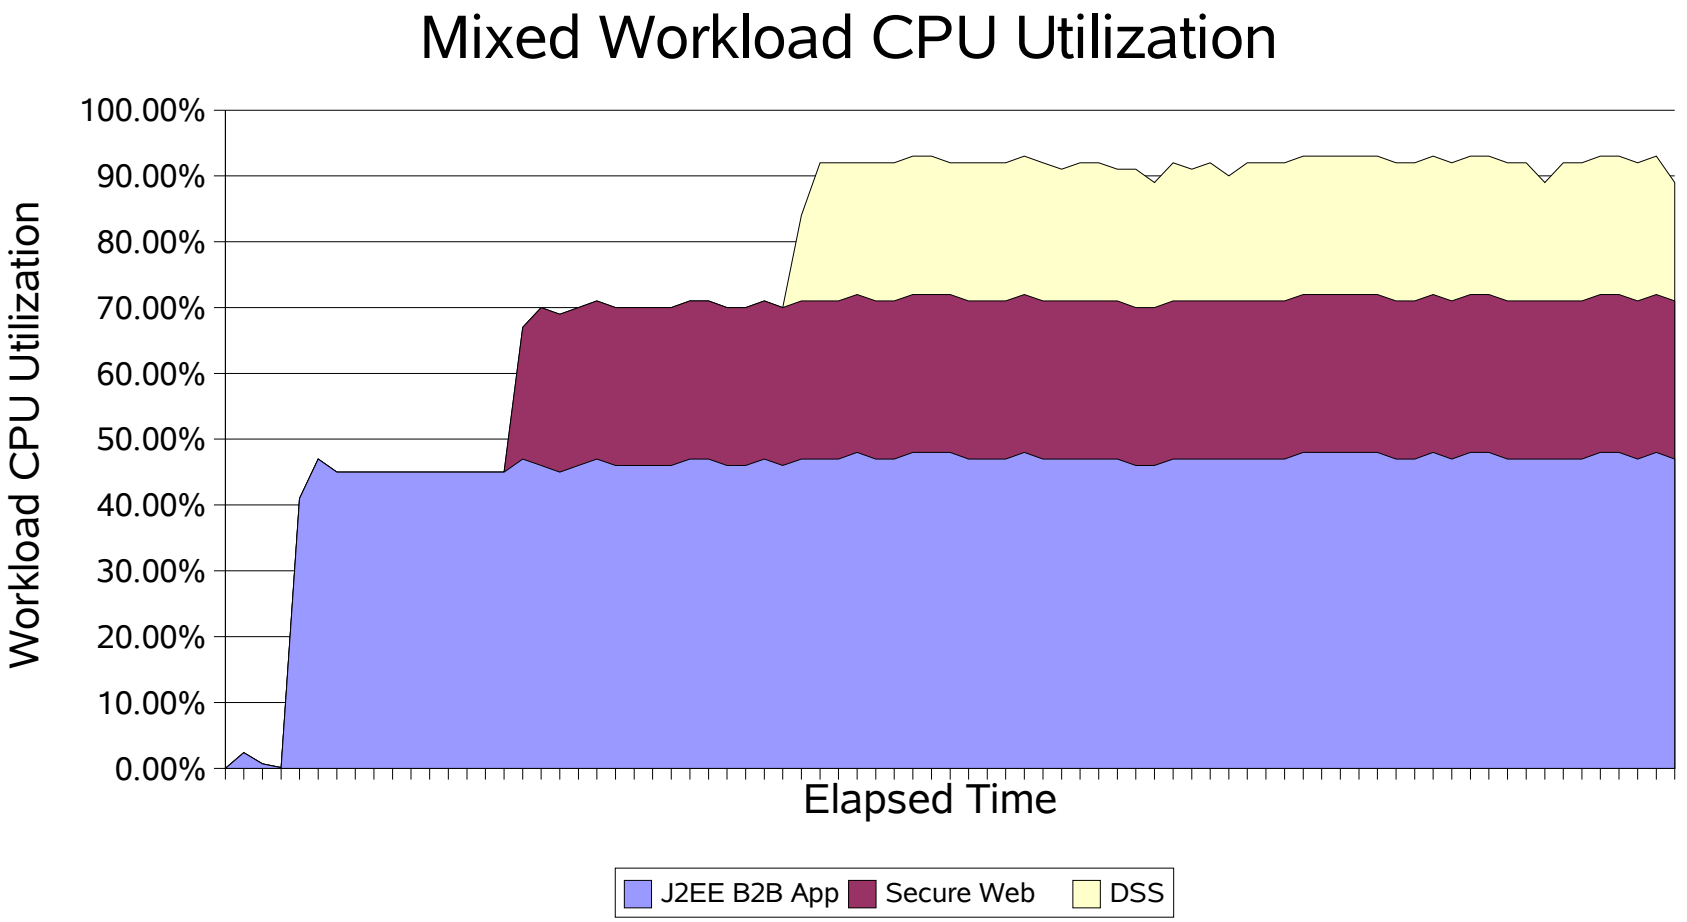

Solaris Containers

vs. Other Virtualization Techniques

Optimal Utilization

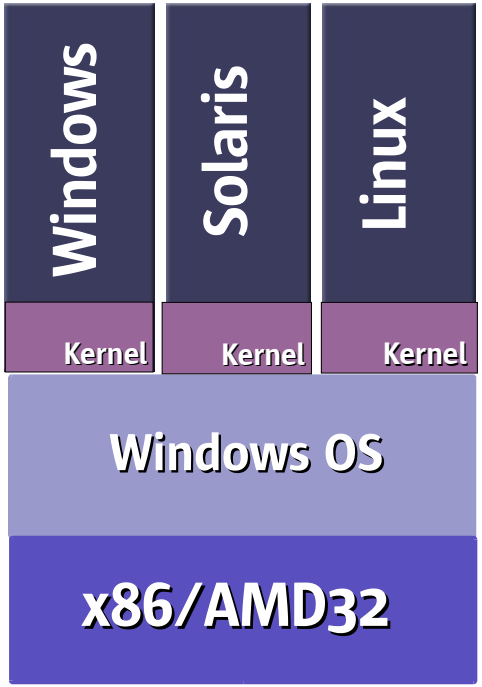

vmware

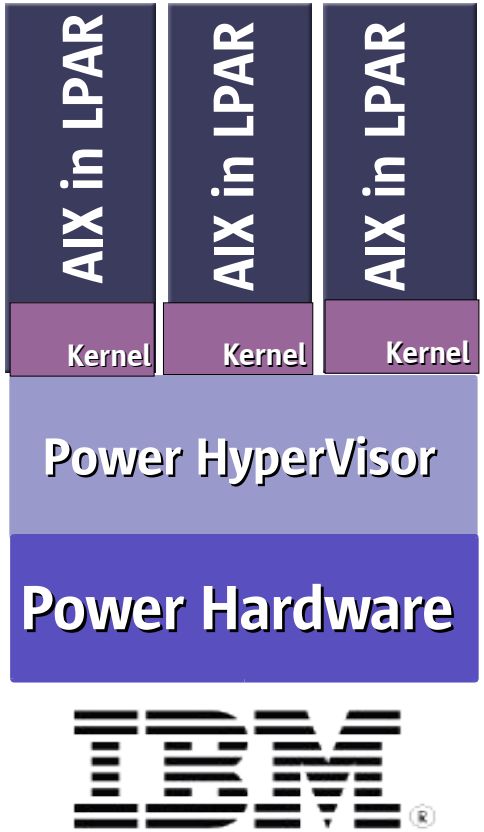

IBM

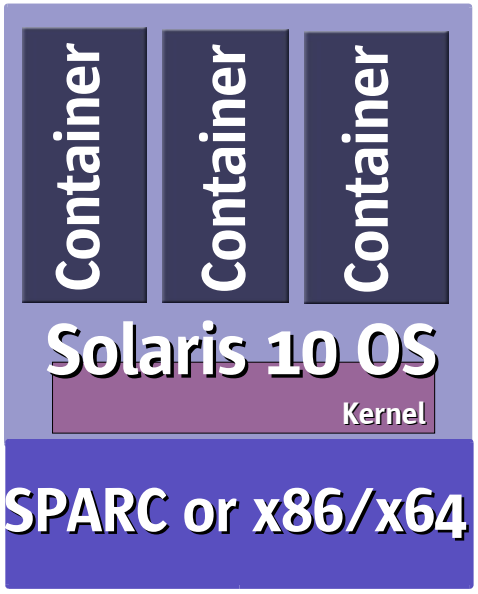

Sun microsystems

Solaris Containers

vs. Other Virtualization Techniques

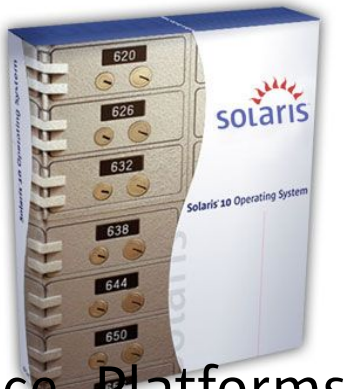

- Contrast these approaches w.r.t
  - Observability, Manageability, Performance, Platforms

Optimal Utilization

20-30%? Overhead

Windows | Solaris | Linux

Kernel | Kernel | Kernel

Windows OS

x86/AMD32

vmware

5-10%? Overhead

AIX in LPAR | AIX in LPAR | AIX in LPAR

Kernel | Kernel | Kernel

Power HyperVisor

Power Hardware

IBM

<1% Overhead

Container | Container | Container

Solaris 10 OS

Kernel

SPARC or x86/x64

Sun microsystems

Dynamic Tracing

Real-time Analysis and Diagnosis

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

--Slashdot post, November '03

- Safe and comprehensive
  - Global view into systems, apps
  - 30,000+ probe points by default
  - Built for use on live production systems
- Reduced costs
  - Solutions in minutes or hours, not days or weeks
  - Optimization: cases of 3-30x customer-app speedups already seen

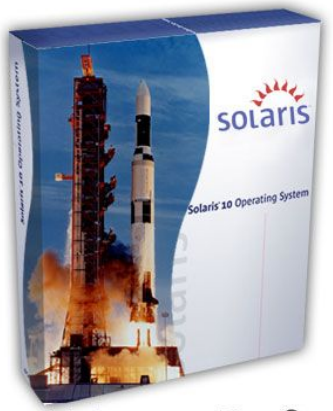

Extreme Performance

DTrace

- Some examples

```
root@vitalstatistix:/# dtrace -l |wc -l
  35066
root@vitalstatistix:/# dtrace -n syscall:::entry'{@[probefunc] = count()}'
dtrace: description 'syscall:::entry' matched 225 probes
^C

  fstat                                             1
  fcntl                                             1
  fstat64                                           1
  putpmsg                                           1
  .......
  ioctl                                           658
  read                                           2162
  pollsys                                        2383
root@vitalstatistix:/#
```

DTrace

- Some examples

```
root@vitalstatistix:/# dtrace -n syscall::ioctl:entry'{@[probefunc,execname] = count()}'
dtrace: description 'syscall::ioctl:entry' matched 1 probe
^C

  ioctl                     gnome-session             2
  ioctl                     gnome-settings-d          2
  .........
  ioctl                     battstat-applet-         22
  ioctl                     gnome-terminal           33
  ioctl                     mixer_applet2            35
  ioctl                     gnome-netstatus-         54
  ioctl                     dtrace                  147
  ioctl                     soffice.bin             284
  ioctl                     acroread                516
root@vitalstatistix:/#
```

DTrace

- Some examples

```
root@vitalstatistix:/# dtrace -n syscall::ioctl:entry'/execname=="acroread"/{@[ustack()] = count()}'
dtrace: description 'syscall::ioctl:entry' matched 1 probe
^C


              libc.so.1`ioctl+0x7
              libX11.so.4`_X11TransBytesReadable+0x11
              libX11.so.4`_XEventsQueued+0x147
              libX11.so.4`XEventsQueued+0x3e
              libXt.so.4`XtAppNextEvent+0x79
              acroread`UnixAppMainLoop+0x13d
              acroread`UnixAppMain+0x2f4
              acroread`main+0x10
              acroread`_start+0x57
              267
root@vitalstatistix:/#
```

DTrace in Action on Wall Street

- **32% in 4 hours** – Midrange SPARC system running a stock market trading, a system and app that the sys admins said was “tuned as fast as it could possibly go”
- **80% in 1 day** – Futures forecasting application running on an x86 platform
- **267% in 2 days** – x86 platform running a message handling benchmark (**beat SUSE by over 40%**)
- **300% in 5 hours** – x86 market order routing engine

DTrace + Containers

- Combination allows unprecedented observability into a multi-tier application
  - easy correlation of events

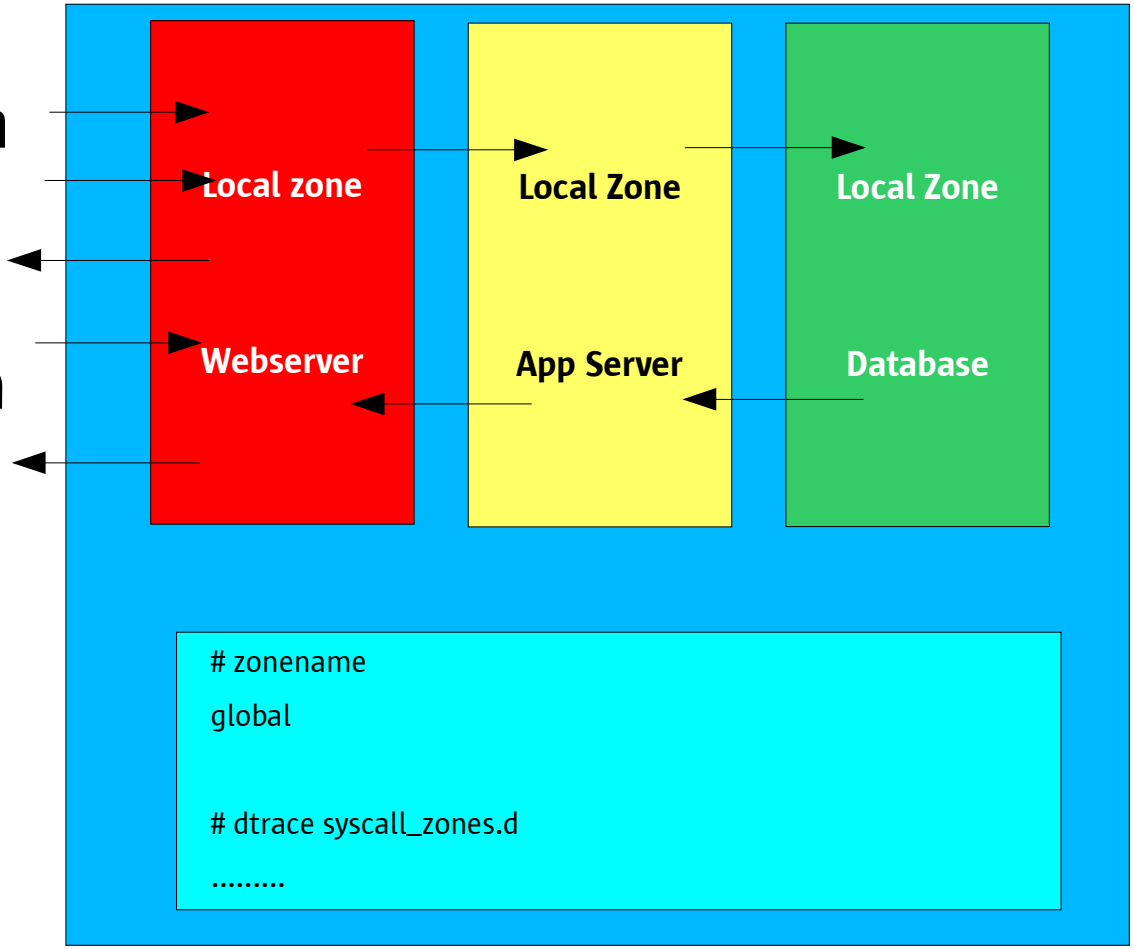

Predictive Self Healing

- Consists of
  - Solaris Fault Manager
  - Service Management Facility

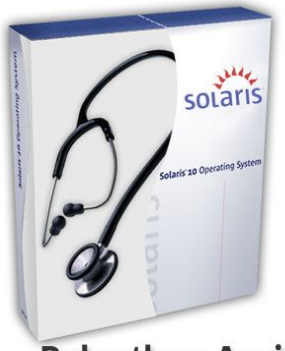

Relentless Availability

Solaris Fault Manager

- Automated error handling
  - Detect faults
  - Aggregate faults
  - Diagnose faults
  - Report faults
  - Mitigate faults

    - takes failing onlines cpus offline, migrates load
    - takes failing online memory offline, migrates load
    - takes failing IO offline, migrates load

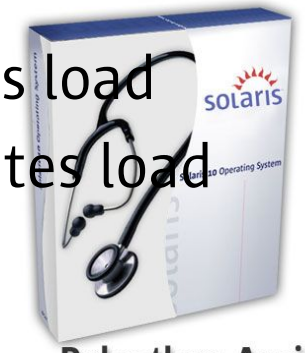

Solaris Fault Manager

sun.com/msg/SF20000-W84N-KP3A-TF

- Customer web-site will provide latest repair procedures for each diagnosis
- Links to information on latest FMA ca`pabilities, updates, and plans

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.

- No passwords – totally free access

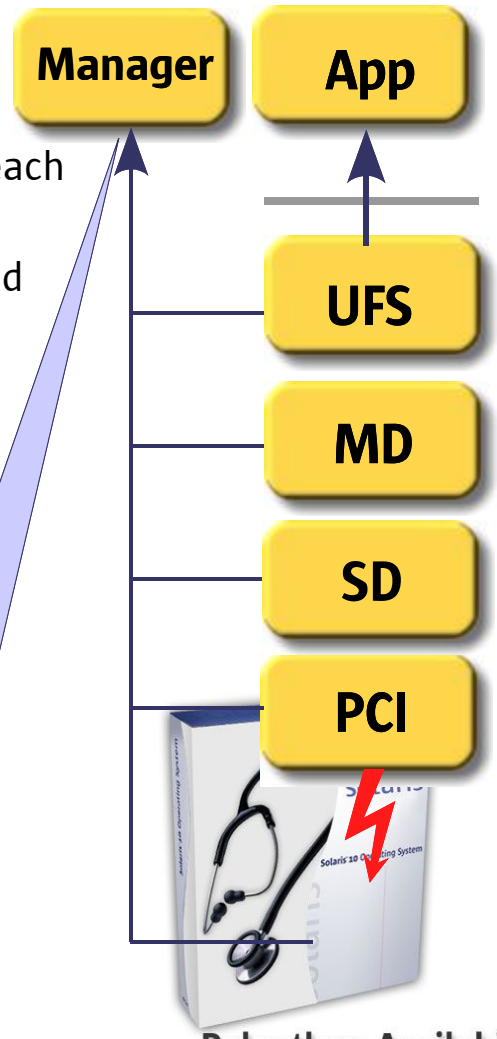

Service Management Facility

- Goal
  - Ease administration of UNIX services
  - Allow “undo” of service configuration changes
  - Automatic restart of failed services
  - Allow admins to get a “system view”
- Elevates UNIX services to managed entities
  - Instead of managing processes, admins manage services
  - Dependencies between services are highlighted
  - Unified management interface

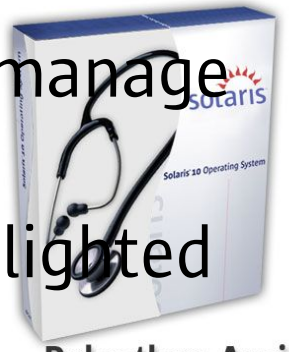

svcs(1) in action

- List active instances, sorted by state, time
- Show dependencies (-d) and dependents (-D)
- Show member processes (-p), additional details (-v)

```
$ svcs
STATE       STIME    FMRI
online      18:18:30 svc:/network/http:apache
online      18:18:29 svc:/network/smtp:sendmail
....
$ svcs -p network/smtp:sendmail
STATE       STIME    FMRI
online      18:18:29 svc:/network/smtp:sendmail
            18:18:29   100180 sendmail
            18:18:29   100181 sendmail
$ svcs -d network/smtp:sendmail
STATE       STIME    FMRI
online      18:17:44 svc:/system/identity:domain
online      18:17:52 svc:/network/service:default
....
```

svcs(1) in action

- List active instances, sorted by state, time
- Show dependencies (-d) and dependents (-D)
- Show member processes (-p), additional details (-v)

```
$ svcs -D network/physical
STATE       STIME    FMRI
disabled     Nov_24   svc:/network/dns/client:default
disabled     Nov_24   svc:/network/dns/server:default
disabled     Nov_24   svc:/network/rarp:default
disabled     Nov_24   svc:/network/rpc/bootparams:default
disabled     Nov_24   svc:/network/slp:default
disabled     Nov_24   svc:/network/shell:kshell
online       Nov_24   svc:/application/print/cleanup:default
online       Nov_24   svc:/system/identity:node
........
online       Nov_24   svc:/network/shell:tcp
online       Nov_24   svc:/network/shell:tcp6only
online       Nov_24   svc:/network/nfs/server:default
$
```

Solaris ZFS
(SOLARIS 10 UPDATE)

- Streamlined system administration
  - Efficient resource allocation via storage pools
  - Automates administrative tasks
  - Extensible: add features such as encryption
- Self-healing data
- Virtually unlimited capacity
  - 2^{128} data blocks: 16 *billion billion* times greater than today
- Breakthrough performance

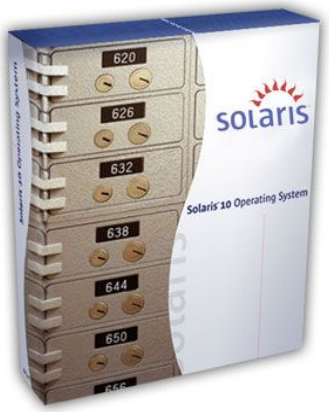

Optimal Utilization

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

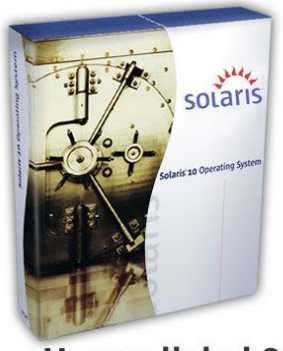

Unparalleled Security

Process Rights Management

- Solaris Process Privileges (ppriv -l)

contract_event	contract_observer	cpc_cpu
dtrace_kernel	dtrace_proc	dtrace_user
file_chown	file_chown_self	file_dac_execute
file_dac_read	file_dac_search	file_dac_write
file_link_any	file_owner	file_setid
ipc_dac_read	ipc_dac_write	ipc_owner
net_icmpaccess	net_privaddr	net_rawaccess
proc_audit	proc_chroot	proc_clock_highres
proc_exec	proc_fork	proc_info
proc_lock_memory	proc_owner	proc_priocntl
proc_session	proc_setid	proc_taskid
proc_zone	sys_acct	sys_admin
sys_audit	sys_config	sys_devices
sys_ipc_config	sys_linkdir	sys_mount
sys_net_config	sys_nfs	sys_res_config

Solaris Linux Application Environment

- 100% binary compatibility with Linux
  - allows Linux apps to run *natively* on Solaris
  - LSB 1.3 compliant
- Works with all Solaris native facilities
  - DTrace, Containers, Networking enhancements, ...
- Available in preview mode soon
- Have successfully tested
  - Opera, Adobe Acrobat reader, Hancom Office, Majesty, BEA Weblogic Server, Oracle Database Server, StarOffice 7, Samba, Apache, and most of the core RedHat Advanced Server 3.0 binaries

Solaris Linux Application Environment

- Continuing to enhance this functionality on all supported platforms
- Using Containers technology to isolate non-native applications
  - Brand Containers as native or non-native

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				■
U.S. $ Price/Socket/Year	**$0**	**$120**	**$240**	**$360**

Subscription List Pricing Comparison:

Solaris 10 vs. Red Hat

	Solaris on x86	Red Hat WS	Red Hat ES	Red Hat AS
Basic	1 CPU: $120/Yr	$179/Yr	$349/Yr	N/A
	2 CPU: $240	$179	$349/Yr	N/A
	4 CPU: $480	N/A	N/A	N/A
Standard	1 CPU: $240	$299	$799	$1,499
	2 CPU: $480	$299	$799	$1,499
	4 CPU: $960	N/A	N/A	$1,499
Premium	1 CPU: $360	N/A	N/A	$2,499
	2 CPU: $720	N/A	N/A	$2,499
	4 CPU: $1,440	N/A	N/A	$2,499

400+ Systems for Solaris... and Growing

Sales Reps are Compensation Neutral.

Performance

- Recent World Records
  - Solaris 10 + SF V40z (852) on SPECjbb2000
    - New high watermark of 116142 JBBops/s
  - Solaris 10 + SF V20z (252) on SPECjbb2000
    - Previous best on SLES9 – 63743 JBBops/s
    - New record on same h/w – 65840 JBBops/s
  - Solaris 10 + SF E6900 (24 1.2GHz USIV) on SPECjAppserver2002 dual node
  - Solaris 10 + SF V20z on SPEC OMPM2001
  - Solaris 10 + SF E6900 on Oracle Apps Batch (HVOP)
  - Solaris 10 + SF E25K on TPC-H
  - Solaris 10 + SF E4900 + Oracle 10g on Oracle Apps Batch

Performance

- 3 primary contributors
  - Faster base OS facilities and tools

Performance

- 3 primary contributors
  - Faster base OS facilities and tools
    - Base OS faster than Solaris 9
    - Networking enhancements
    - libumem
    - Studio 9 and 10
    - ……

Performance

- 3 primary contributors
  - Faster base OS facilities and tools
  - Better observability tools

Performance

- 3 primary contributors
  - Faster base OS facilities and tools
  - Better observability tools
    - DTrace
    - Studio Collectors
    - ....

Performance

- 3 primary contributors
  - Faster base OS facilities and tools
  - Better observability tools
  - Vendor enthusiasm

Performance

- 3 primary contributors
  - Faster base OS facilities and tools
  - Better observability tools
  - Vendor enthusiasm
    - Benchmark losses treated as **bugs**

Networking enhancements

- Data locality:
  - packets for the same connection are processed on the same CPU whenever possible.
- TCP/IP interaction:
  - Switch from a message passing interface to a function call based interface.
- An IP classifier based approach:
- Active interrupt management and solving the livelock problem

Networking enhancements

- Merge TCP/IP into one module and create a function call based interface.
- Move to a reference based scheme
- Use a serialization mechanism (Vertical perimeters aka squeue) to protect the TCP data structure
- Make IP fully multithreaded to remove dependency on STREAMS protection.
- Use IP classifier to look up connections.

Networking enhancements

- Achieved 45% gain on web like workload on SPARC
- Achieved 43% gain on web like workload on x86 (limited by the VM available)
- Other gains:
  - 10% SSL
  - 10% fileserving
  - 20-40% throughput (ttcp)
- Solaris 10 can fully saturate a 1Gb link with only 8% of 1x2.2Ghz Opteron and
- Solaris 10 can drive a 10Gb link at 7Gbps (limited by PCI-X bandwidth) using 2x2.2Ghz opteron CPUs utilized at less than 50%

New x86 Optimizations in Sun Studio 9

- SSE2 instruction scheduling
- P4, SSE2 instr in assembler
- Handle P4, SSE2 in inlines
- Strength reduction
- Branch prediction
- Induction variable elim
- Invariant hoisting
- Loop interchange
- Loop unswitching
- Alignment of symbol blocks
- Loop unrolling
- Alignment
- Constant propagation
- Vectorization

- -xO4 and -fast gets you all optimizations
- Note: Intel SSE/SSE2 support introduced in Solaris 9 4/04... but you need apps compiled with -xarch=sse2 to enable instruction generation

Sun Studio 9: SPEC Performance

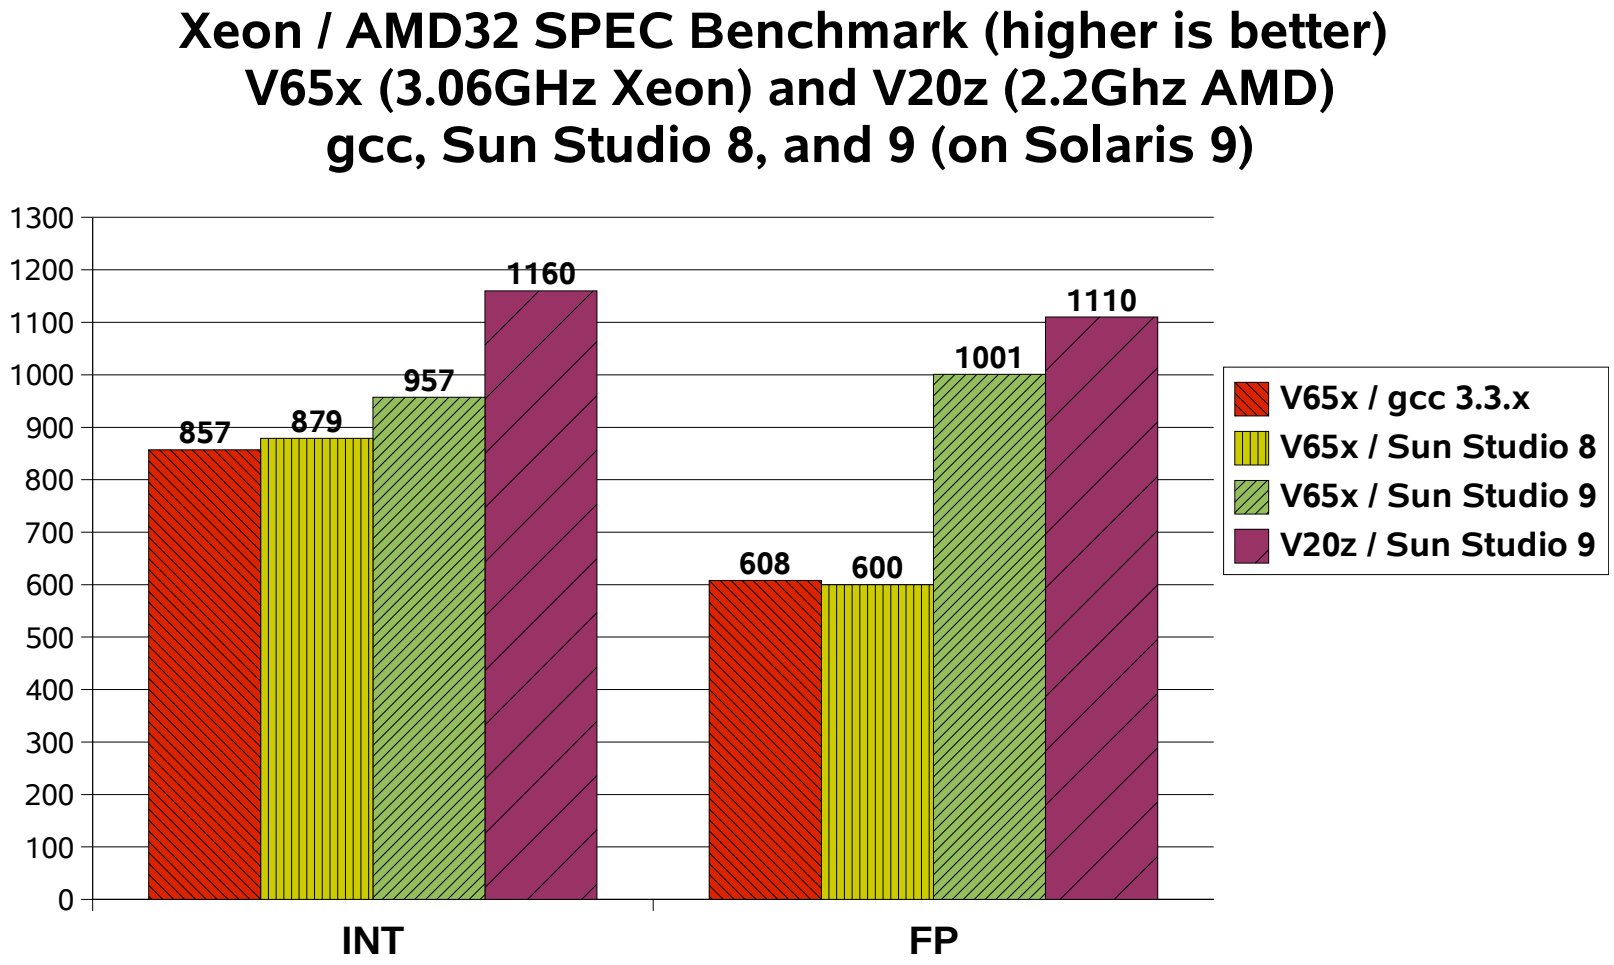

Sun Studio 10 (Vulcan)

Sun Studio 9 for Solaris on AMD64 !

- Goals
  1. 64-bit app support with compelling performance on Sun's V20z, V40z, W2100z
  2. Feature parity w/Sun Studio SPARC to ease ISV adoption
  3. Linux support for seamless development between Linux and Solaris on SPARC, Xeon, and AMD64
     - Better performance on Solaris/AMD64 than Linux/AMD64

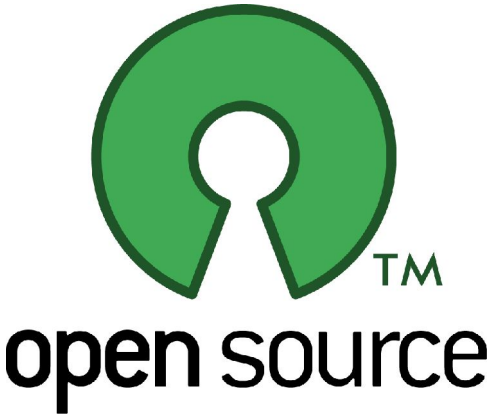

opensolaris™

Solaris Source Code

Over 1600 Patents released

OSI Approved License

Buildable Source
Q2CY2005

Sun's Linux Strategy

- Linux from Sun = Off-the-shelf Linux + Java ES + Sun Service

&

- Latest Sun Java in all distributions
- All Sun x86 hardware certified for Linux
- Run all Sun software on Solaris and Linux
- Service and support worldwide
- Sun continues as a leader in open source community contributions

Sun Delivers on Linux

Commercial-grade HW, OS, Middleware, Tools, Services, Partners: Everything you need — NOW

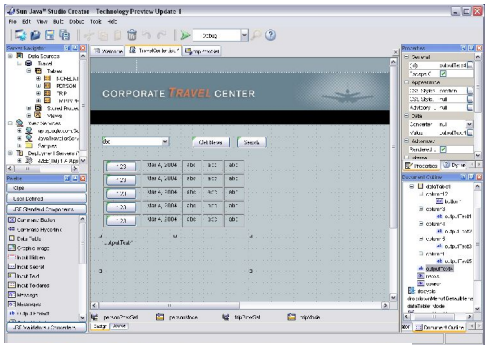

End-to-end Commercial Linux Portfolio

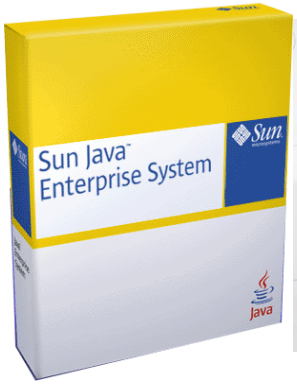

The Infrastructure Stack

The Integrated Platform

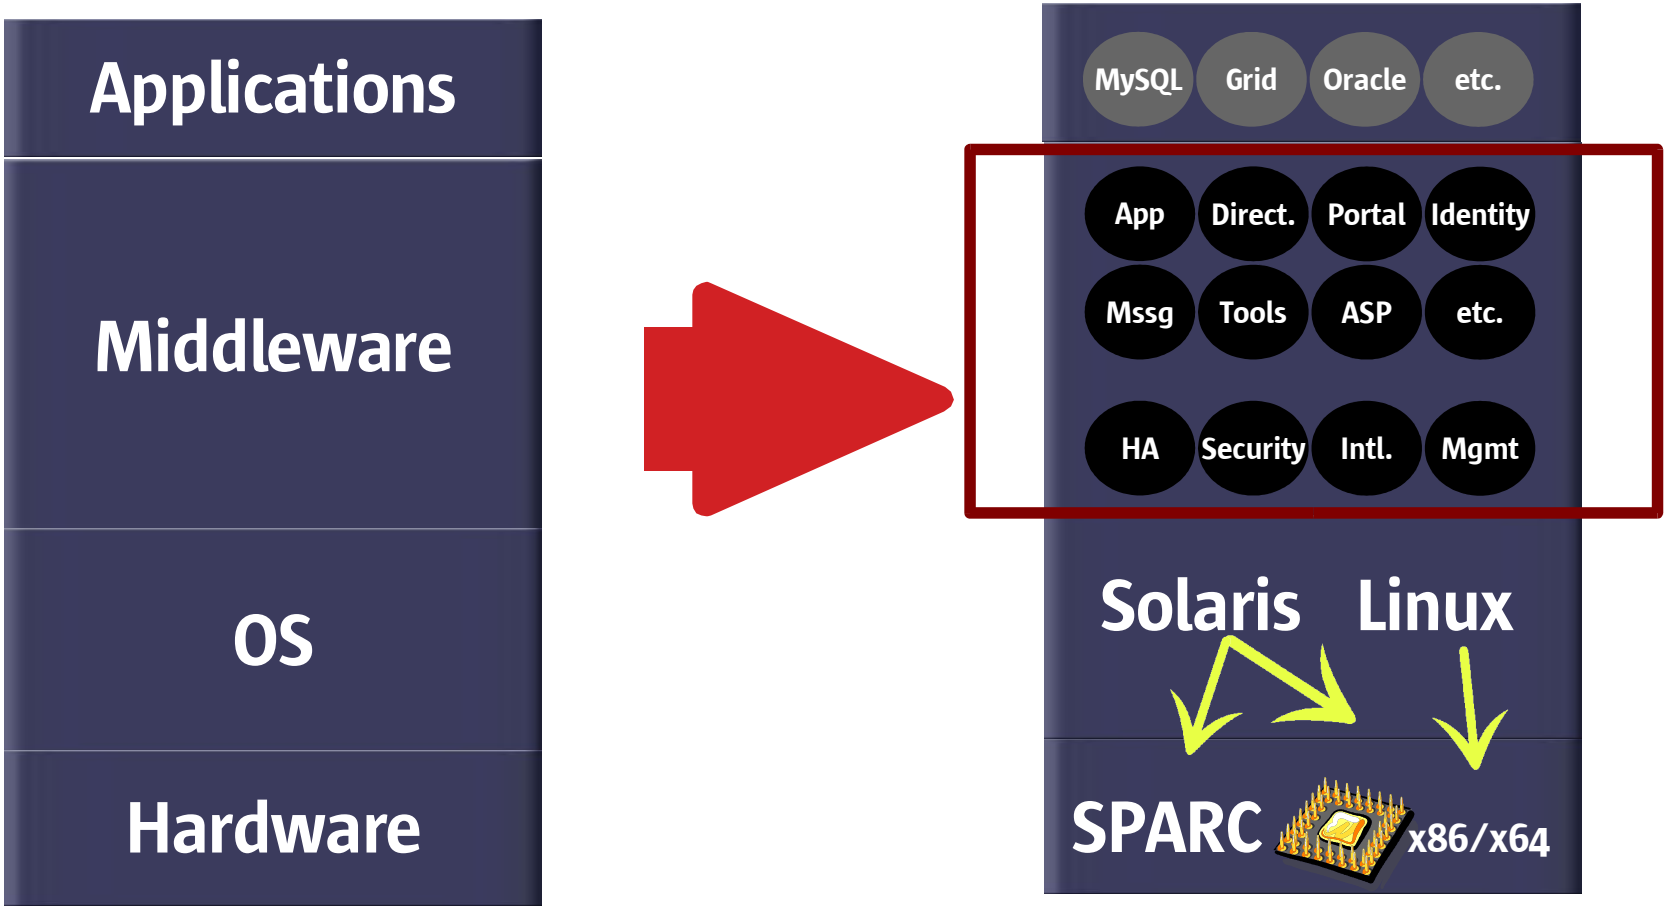

Problem 1: Integration Nightmare

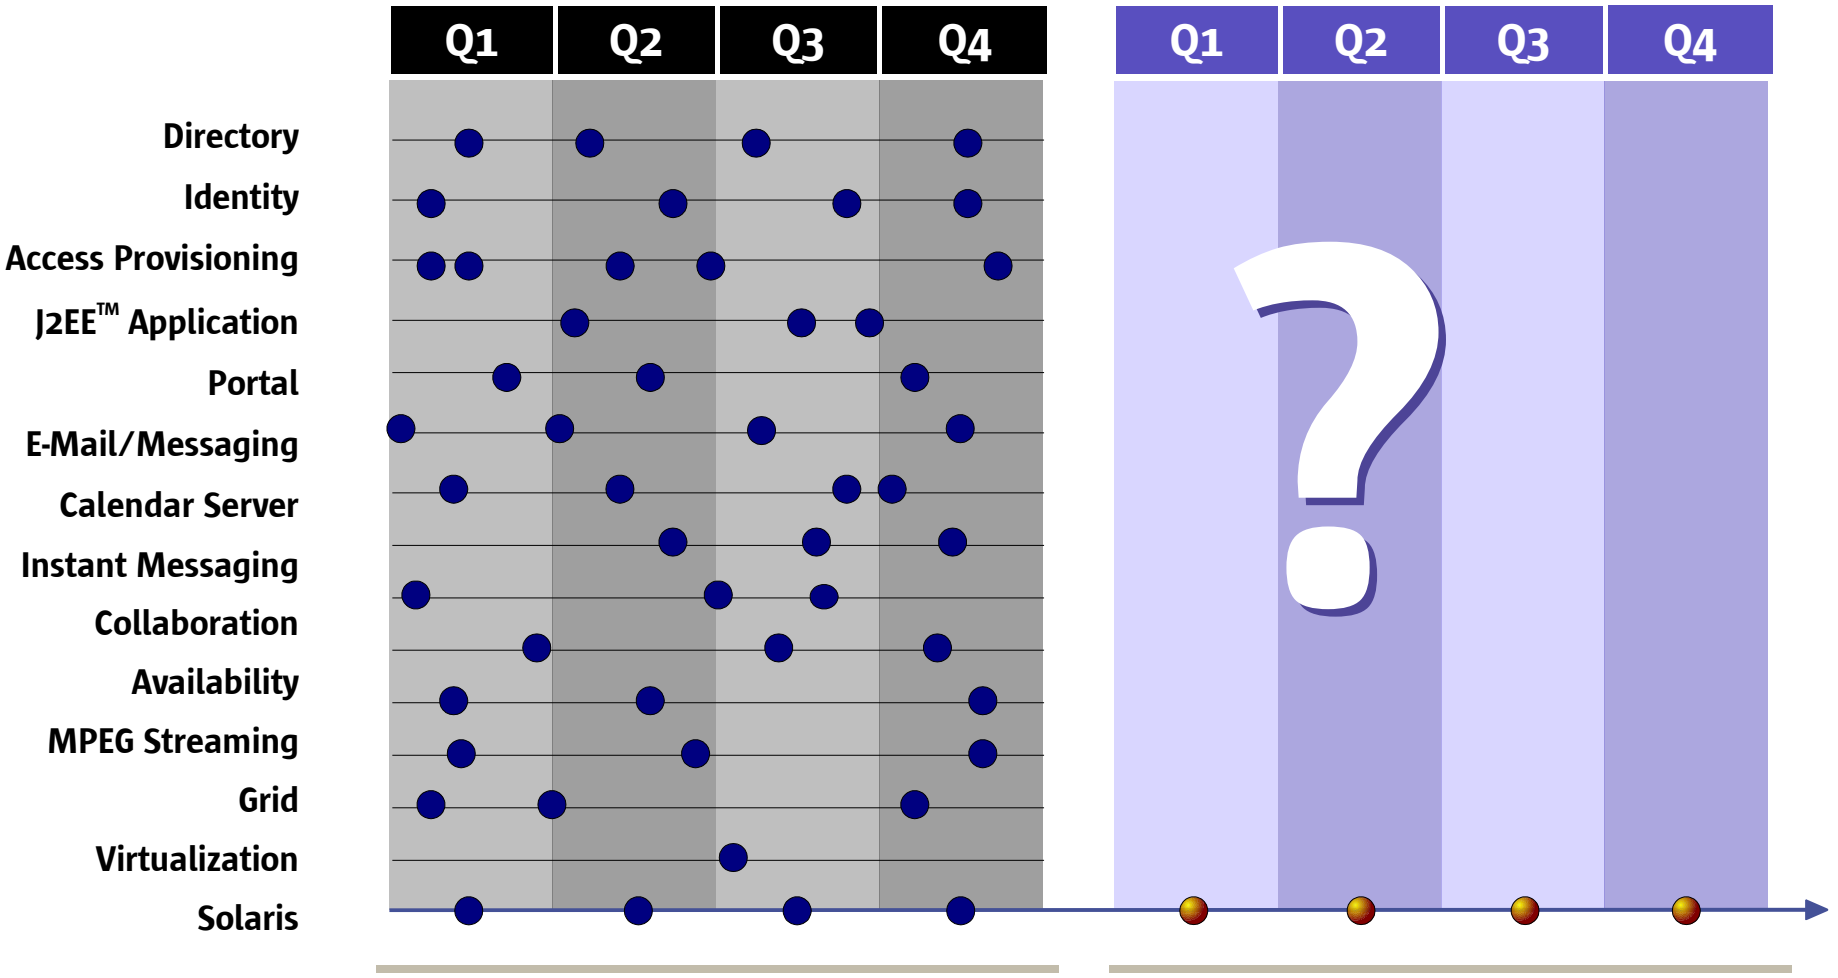

Problem 2: Price, Confusion

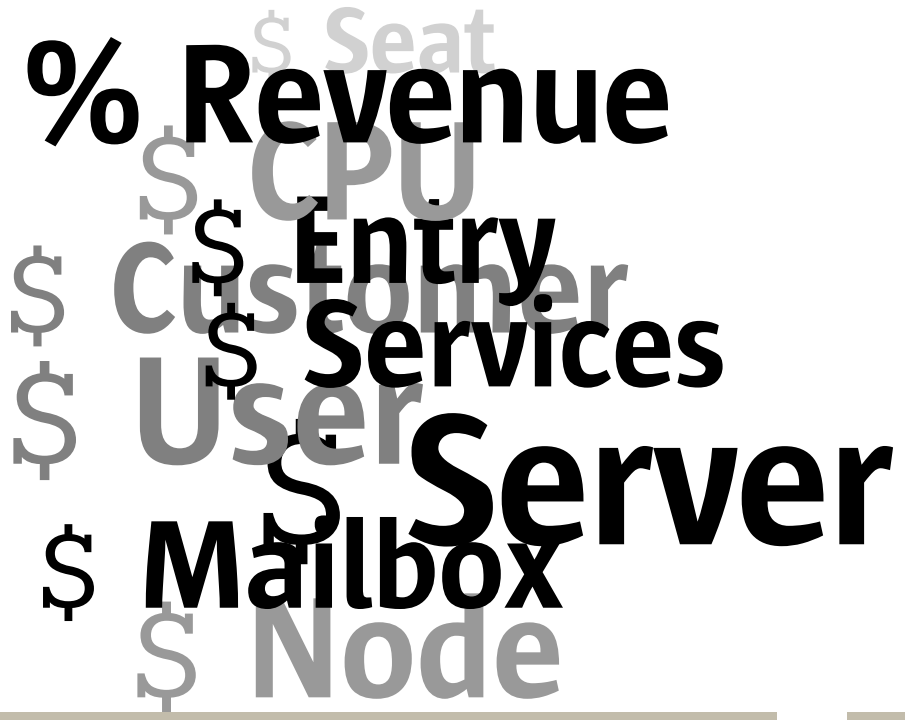

One unit
of measure

Problem 3: Licensing Nightmare

Solution 1: Java ES Delivery

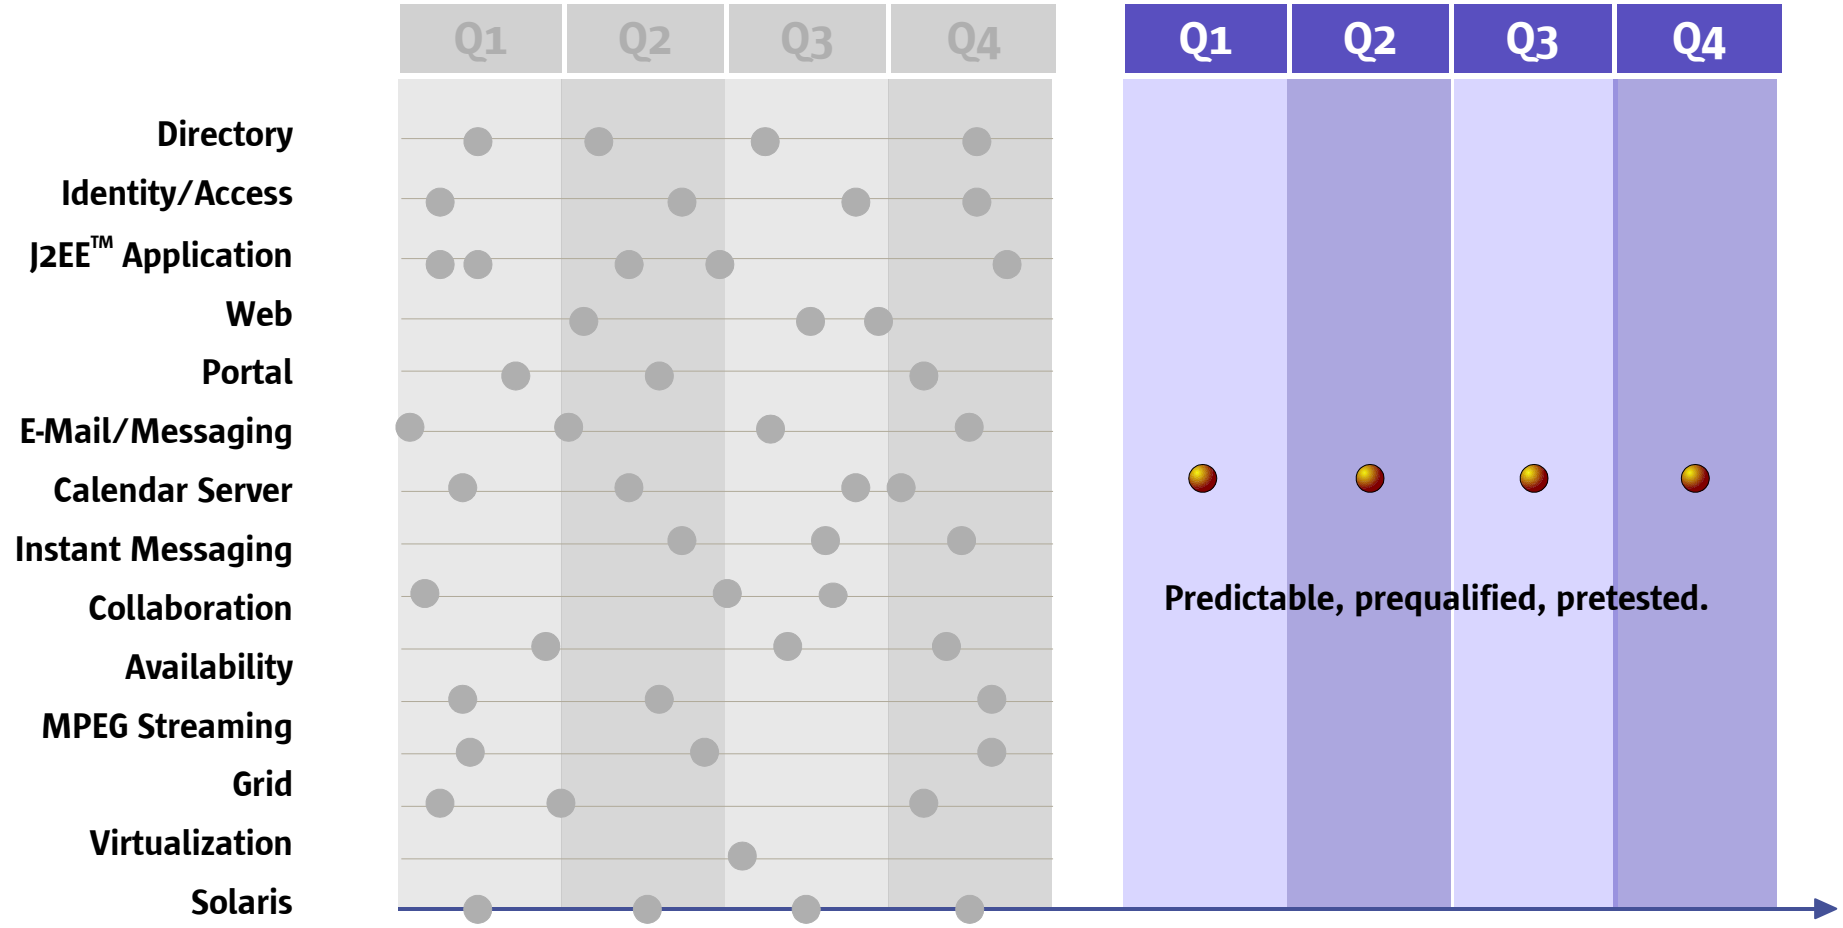

Solution 2: Java ES Pricing

$/Employee/Year

$140

Includes migration services, training, and support

Solution 2: Java ES Pricing

$/Employee/Year

$140

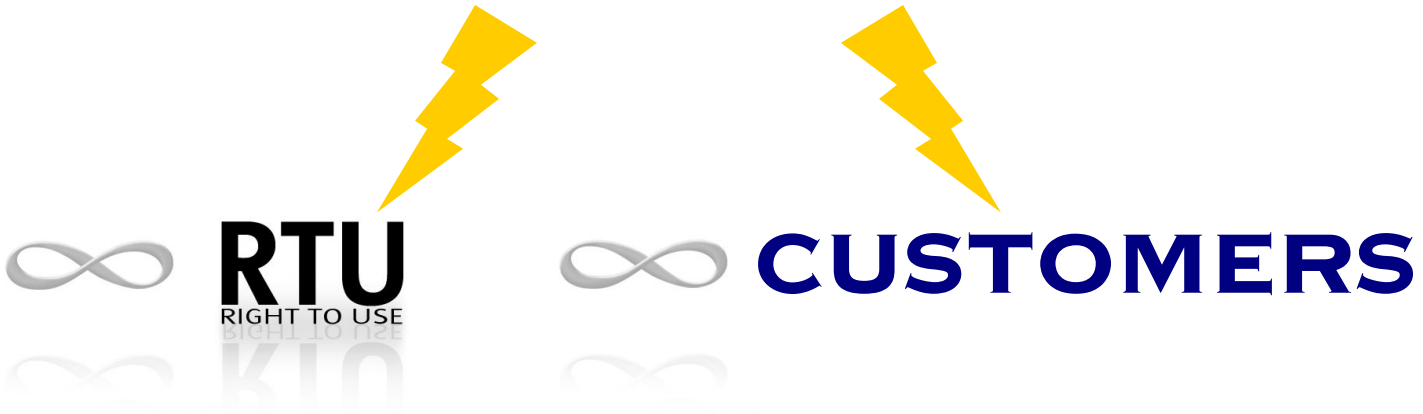

YOU CAN DEPLOY **ALL** INTERNAL & **EXTERNAL** SERVICES

WITH NO ADDITIONAL CHARGE

Value Starts At Acquisition

	Java Enterprise System	IBM	Microsoft
5 Year Total, Including Support, Maintenance, Consulting and Education Services			
1:30 employee-to-external users ratio			
5,000 Employees	$3,500,000	$9,264,226	$9,380,066
10,000 Employees	$7,000,000	$13,866,605	$14,649,553
25,000 Employees	$17,500,000	$26,749,667	$28,483,744
10,000 Employees			
1:10 employee-to-external users ratio	$5,000,000	$8,731,079	$9,011,796
1:30 employee-to-external users ratio	$5,000,000	$13,866,605	$14,649,553
1:60 employee-to-external users ratio	$5,000,000	$19,070,315	$19,555,036

Sun Java Enterprise System

- Directory
- App Server
- Message Queue
- Web Server
- Portal
- Messaging
- Calendar
- Corporate IM
- Portal Mobi
- Cluster

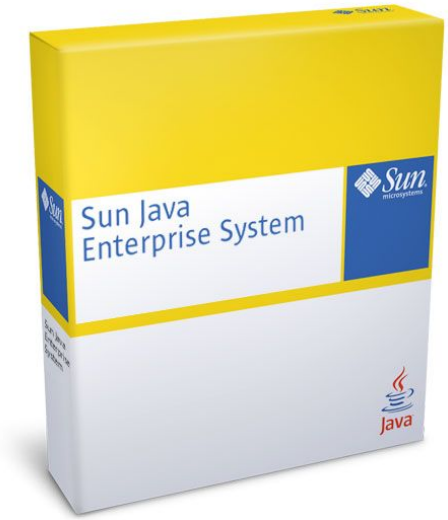

RTU
RIGHT TO USE

$50-140
Empl/Year

Highlights

Readily deployable infrastructure software

Predictable: Pre-integrated, Pre-tested

Multiplatform support (Solaris on SPARC/x86/x64, Linux)

250+ ISV applications ready

Simple $140 employee/year subscription model, Infinite Right-to-Use

New: $50 empl/yr "Suites"

Java Enterprise System Expands to Suites

Java Web Services

Sun microsystems

NEW

- Applications Platform suite
- Availability suite
- Communications suite
- ID Management suite
- Web Infrastructure suite

RTU RIGHT TO USE

$50 per Employee-Yr

$140 per Employee-Yr

The Infrastructure Stack

The Integrated Platform

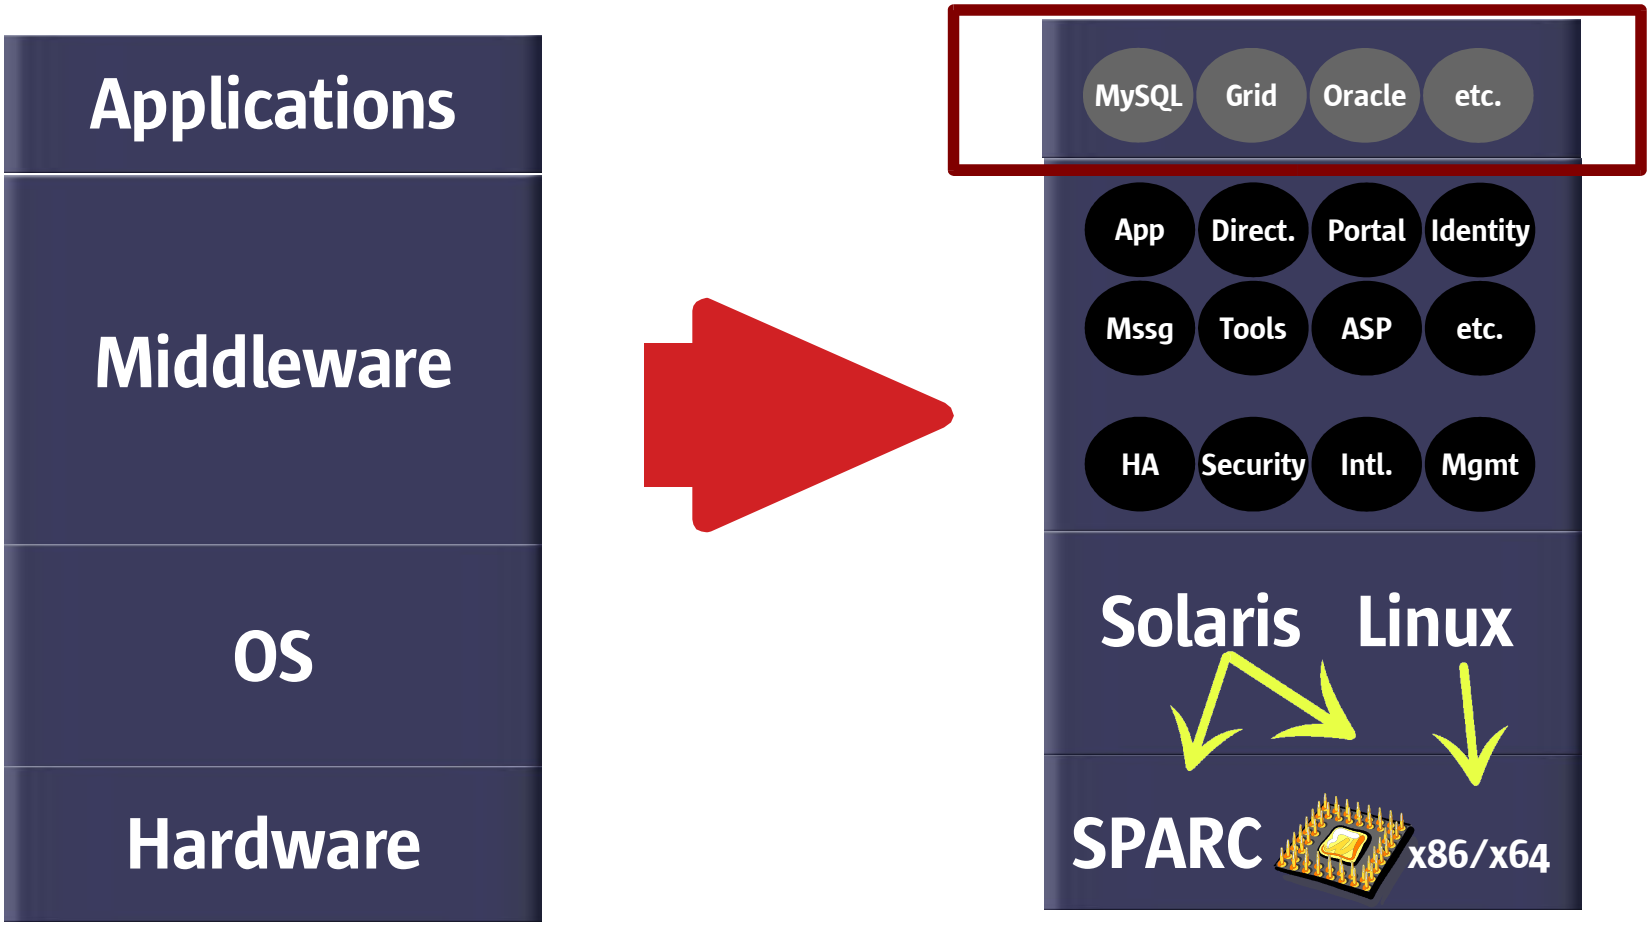

Applications

- 3 major classes
  - Freeware
  - Homegrown
  - ISV
- Goal
  - Ensure availability of all 3 classes of applications on all supported systems

Applications – Freeware

- Ensure availability of freeware on all supported systems
  - http://sun.com/solaris/freeware
    contains list of freeware software included with Solaris 10
    - Included with Solaris 10 and supported
      - eg., Apache, Samba, Tomcat, Perl, Secure Shell, MySQL, Gcc
      - 60+
    - Co-packaged with Solaris 10 via the Companion CD
      - 130+

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

- Binutils
- Bison
- Flex
- Gcc
- Gm4
- Gmake

Scripting Languages

- Perl
- Python

Security Tools

- 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 / Accessibility
brltty-3.3.1
emacspeak-18.0
emacspeak-ss-1.9.1
freetts-1.1.1
screenbrltty-4.02
unwindows-1.1.3
w3-4.0.47
yasr-0.6.4

Applications / Editors
bluefish 0.12
emacs 21.3
gawk-3.0.6
joe-3.1
sed-3.02 (GNU)
vim-6.3
xemacs-21.4.15

Desktop / Environment
kde-3.1.1a
KOffice-1.2.1
XFce-3.8.16

Applications / Networking
cups-1.1.20
ethereal 0.10.5
fetchmail 6.2.5
hpijs 1.6
lynx-2.8.4
mutt-1.4.2.1
nmap 3.5
nmh-1.0.4
Open LDAP 2.2.17
Open SLP 1.0.11
pine-4.61
procmail-3.22
rsync 2.6.3pre1
slm-0.9.6.2
snort-2.0.0
tcpdump-3.8.3

Applications / Publishing
espgs-7.07.1
graphviz 1.10
groff-1.16.1
xpdf 3.0

System / Daemons
imap2002d (UW)
proftpd 1.2.10rc1
squid 2.5.STABLE7

Applications / Utilities
afio-2.4.6
amanda-2.4.4
cdrtools-2.01
cupsddk 1.0
diffutils-2.8.1
enscript-1.6.1
expect 5.39
file-4.10
fileutils-4.1
findutils-4.1.20
Foomatic filters 3.0.2
Foomatic-ppds 3.0.1
gcal-3.01
gettext-0.10.35
gimp-print-4.2.6
gkrellm 2.1.19
gnuplot 3.7.3
ispell-3.2.06
lxrun 0.9.6.1
mpack-1.5
mpage-2.5.1
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
TeTex 2.0.2
textutils-2.0
tnef 1.1.3
top-3.5.1
uudeview-0.5.20
vorbis-1.0
wine 20041104
xpp-1.1

X / Applications
asclock-1.0
ethereal-0.9.11
gimp-1.2.1
rxvt-2.7.10
xcpustate-2.5
xdelta 1.1.3
xmcd 3.2.1
xmms 1.2.10

Development / Languages
bison-1.35
gcc-2.95.3
gcc-3.4.2
libtool 1.5.2
m4-1.4 (GNU)
MySQL python API 0.9.2
php-4.3.2
ruby-1.6.4
samp-1.0
tclX-8.2.0

Development / Tools
autoconf 2.59
automake 1.8.3
binutils-2.15
cvs 1.11.17
ddd 3.3.8
gdb 6.2.1
global-4.8
make-3.80 (GNU)

Development / Libraries
aalib-1.2
berkley-db 1.85
berkley-db 4.2.52NC
curl-7.10.3
fltk-1.1.3
fnlib-0.5
GD Graphics library 2.0.15
guile-1.3.4
imlib-1.9.15
libexpt-1.95.7
libmpeg-1.3.1
libpcap-0.8.3
libsane 1.0.14
linungif-4.1.0
ncurses-5.2
Ogglib-1.0
Perl regex lib 4.5
qt-3.1.1
readline-4.2
slang-1.4.0
SDL-1.2.5
Xaw3d-1.5

X / Window Managers
afterstep-1.8.8
fvwm2-2.4.3
WindowMaker-0.80.2

Applications – Homegrown

- Ensure availability of Development Tools on all supported systems
  - C/C++/FORTRAN
    - Sun Studio 9/10
    - Solaris, Linux
    - http://sun.com/software/products/studio

Applications – Homegrown

- Ensure availability of Development Tools on all supported systems
  - Java
    - Sun Java Studio Enterprise
      - Java, J2EE, Web Services Creation
      - Solaris, Linux *(work in progress)*, Windows
      - http://sun.com/software/products/jsenterprise
    - Sun Java Studio Creator
      - Rapid Development and Deployment w/JSF
      - Solaris, Linux, Windows
      - http://sun.com/software/products/jscreator

Applications – Homegrown

- Ensure availability of Development Tools on all supported systems
  - Java
    - Netbeans
      - Solaris, Linux, Windows, Mac OS X
      - http://www.netbeans.org

Applications – Homegrown

	Solaris Version						Linux
Product Release	2.5.1	2.6	7	8	9	10	
WorkShop 3.0	Y	Y	Y	N	N	N	N
WorkShop 5.0	Y	Y	Y	Y	N	N	N
Forte Developer 6	N	Y	Y	Y	N	N	N
Forte Developer 6 update 1	N	Y	Y	Y	N	N	N
Forte Developer 6 update 2	N	Y	Y	Y	Y	N	N
Sun ONE Studio 7	N	N	‡	‡	‡	N	N
Sun Studio 8	N	N	‡	Y	Y	Y	N
Sun Studio 9/10	**N**	**N**	**N**	**Y**	**Y**	**Y**	**Y**

Applications – ISV
Hundreds of ISVs and Partners

It's All About Choice

The Infrastructure Stack

The Integrated Platform

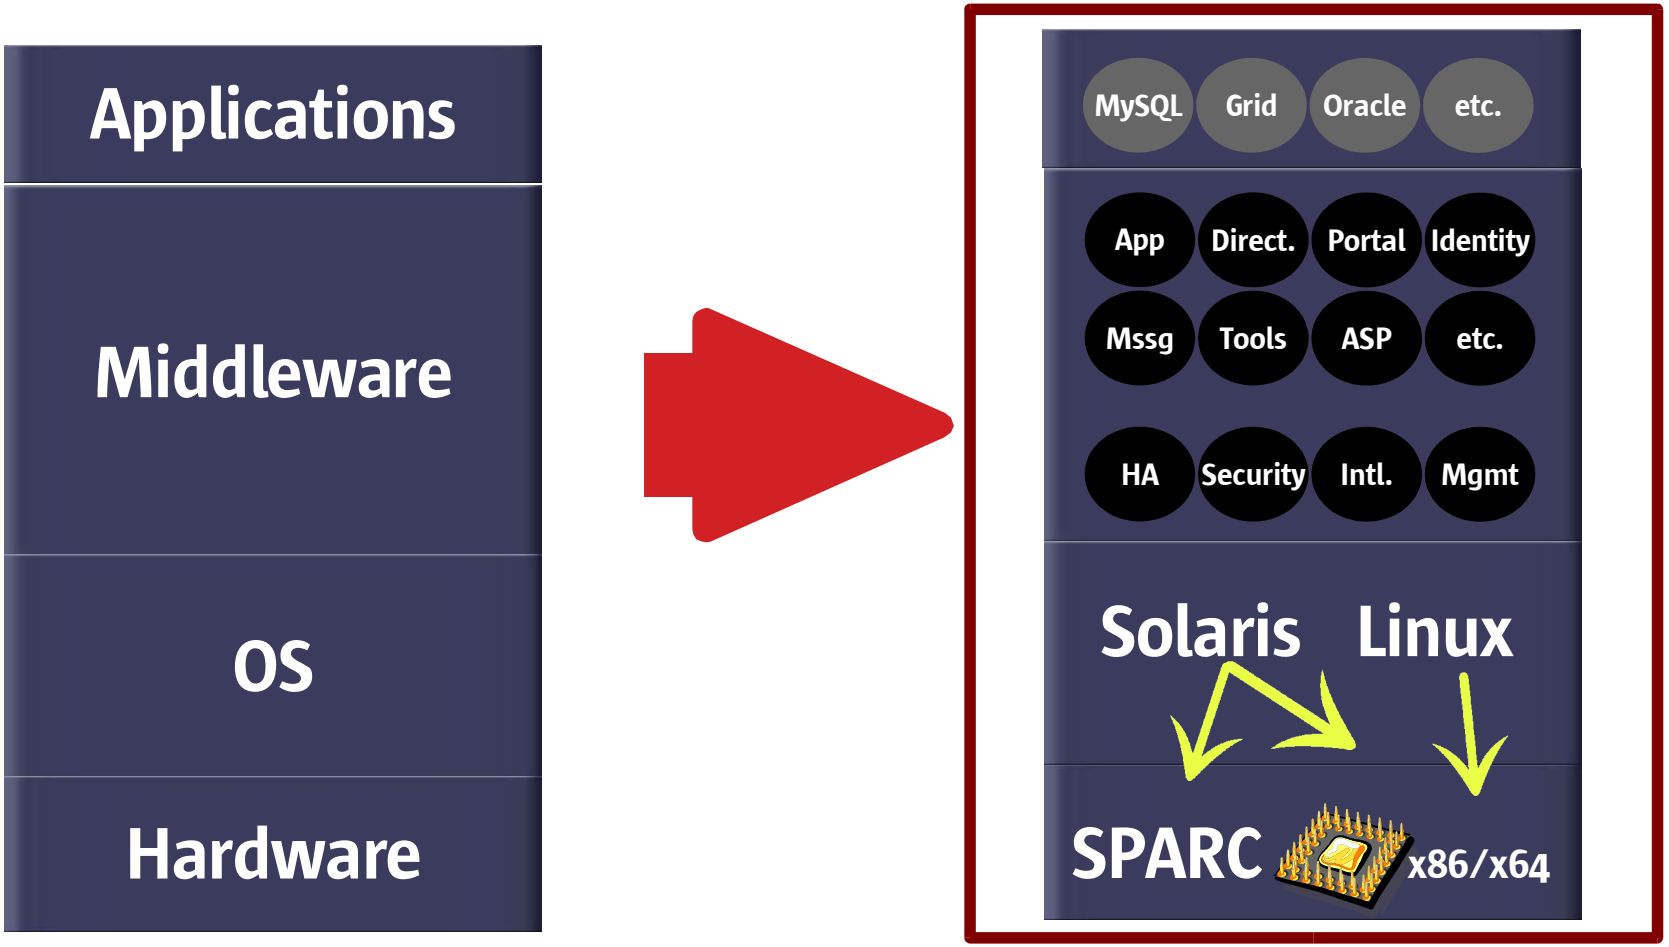

Offering Customers a Choice

Java
Sun Microsystems

Applications

bea™ cadence® Adobe® real SYNOPSYS® VERITAS™
documentum MatrixOne Changing the Way the World Brings Products to Market™ ORACLE® ca Computer Associates SYBASE®

Infrastructure

Java Enterprise System/Java Desktop System

Operating System


solaris™ solaris™

Architecture

intel inside XEON
AMD 64 Opteron
ULTRASPARC® DRIVEN

Sun's Platform Strategy

ambreesh.khanna@sun.com

Sun microsystems