


What are we doing to optimize OpenSolaris*?

December 5, 2007

Frank.Wang@intel.com

Presenter : Paul.Guermonprez@intel.com

For Sun Tech Day 2007



Agenda

- Intel enterprise server advantage
- Intel and Sun collaboration
- Key development areas
- Summary/Call to Action



Intel is a trademark of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.



Intel® Quad-Core - A Superior Design

Dual-die vs Monolithic:

Faster to design: 6-9 mos

Lower Cost

- Smaller die size
- Better yield (~20%)
- Lower mfg cost (~12%)

Better supply

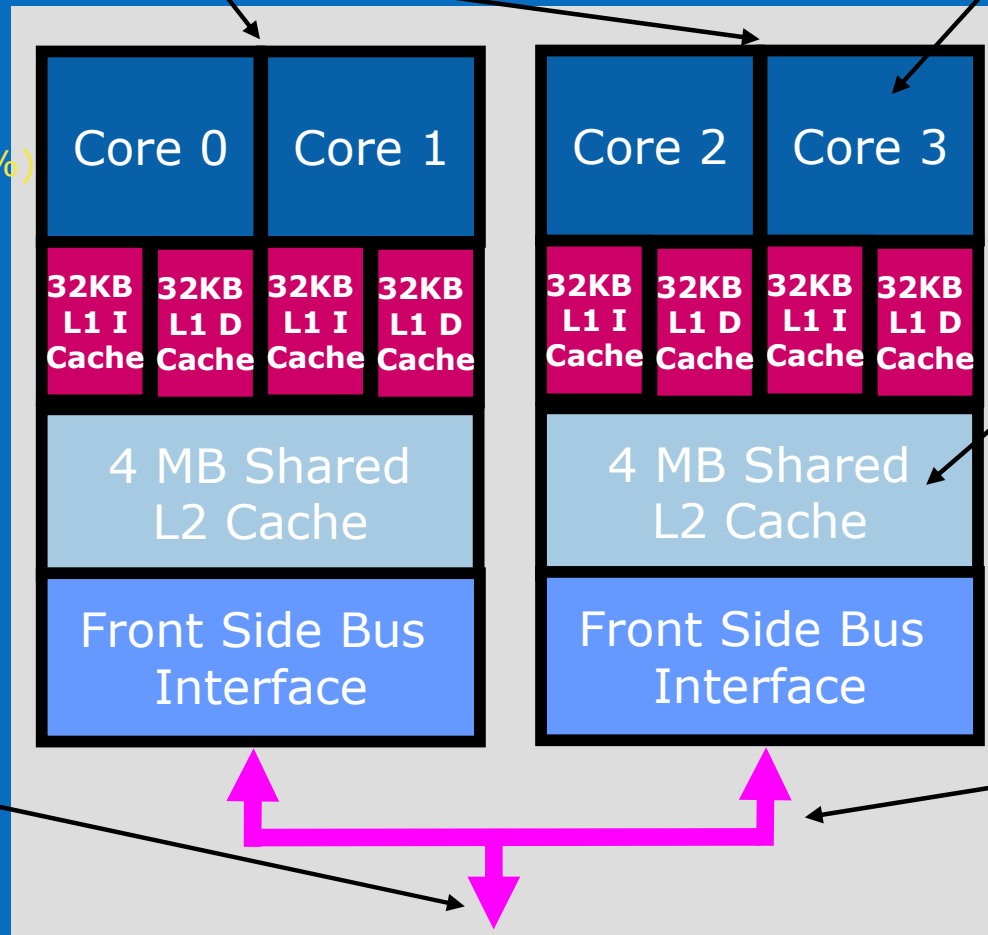
Extends to 45nm

Intel Core™ uArch:

Leading Perf and Perf/W

64-bit

Intel Virtualization Tech.



Large L2 cache:

2X competitors size

Lower latency (vs L3)

Fewer cache misses

More efficient inclusive design

Reduces bus traffic

Socket compatible:

From dual-core through to 45nm quad-core

Front-side Bus: up to 1333MHz

Enables uniform access to shared memory

Leading performance, low cost and extensible



Legal Disclaimers

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit <http://www.intel.com/performance/resources/limits.htm> or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

All dates and products specified are for planning purposes only and are subject to change without notice

Relative performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

SPEC, SPECint2000, SPECfp2000, SPECint2006, SPECfp2006, SPECjbb, SPECWeb are trademarks of the Standard Performance Evaluation Corporation. See <http://www.spec.org> for more information.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor series, not across different processor sequences. See http://www.intel.com/products/processor_number for details.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications. All dates and products specified are for planning purposes only and are subject to change without notice

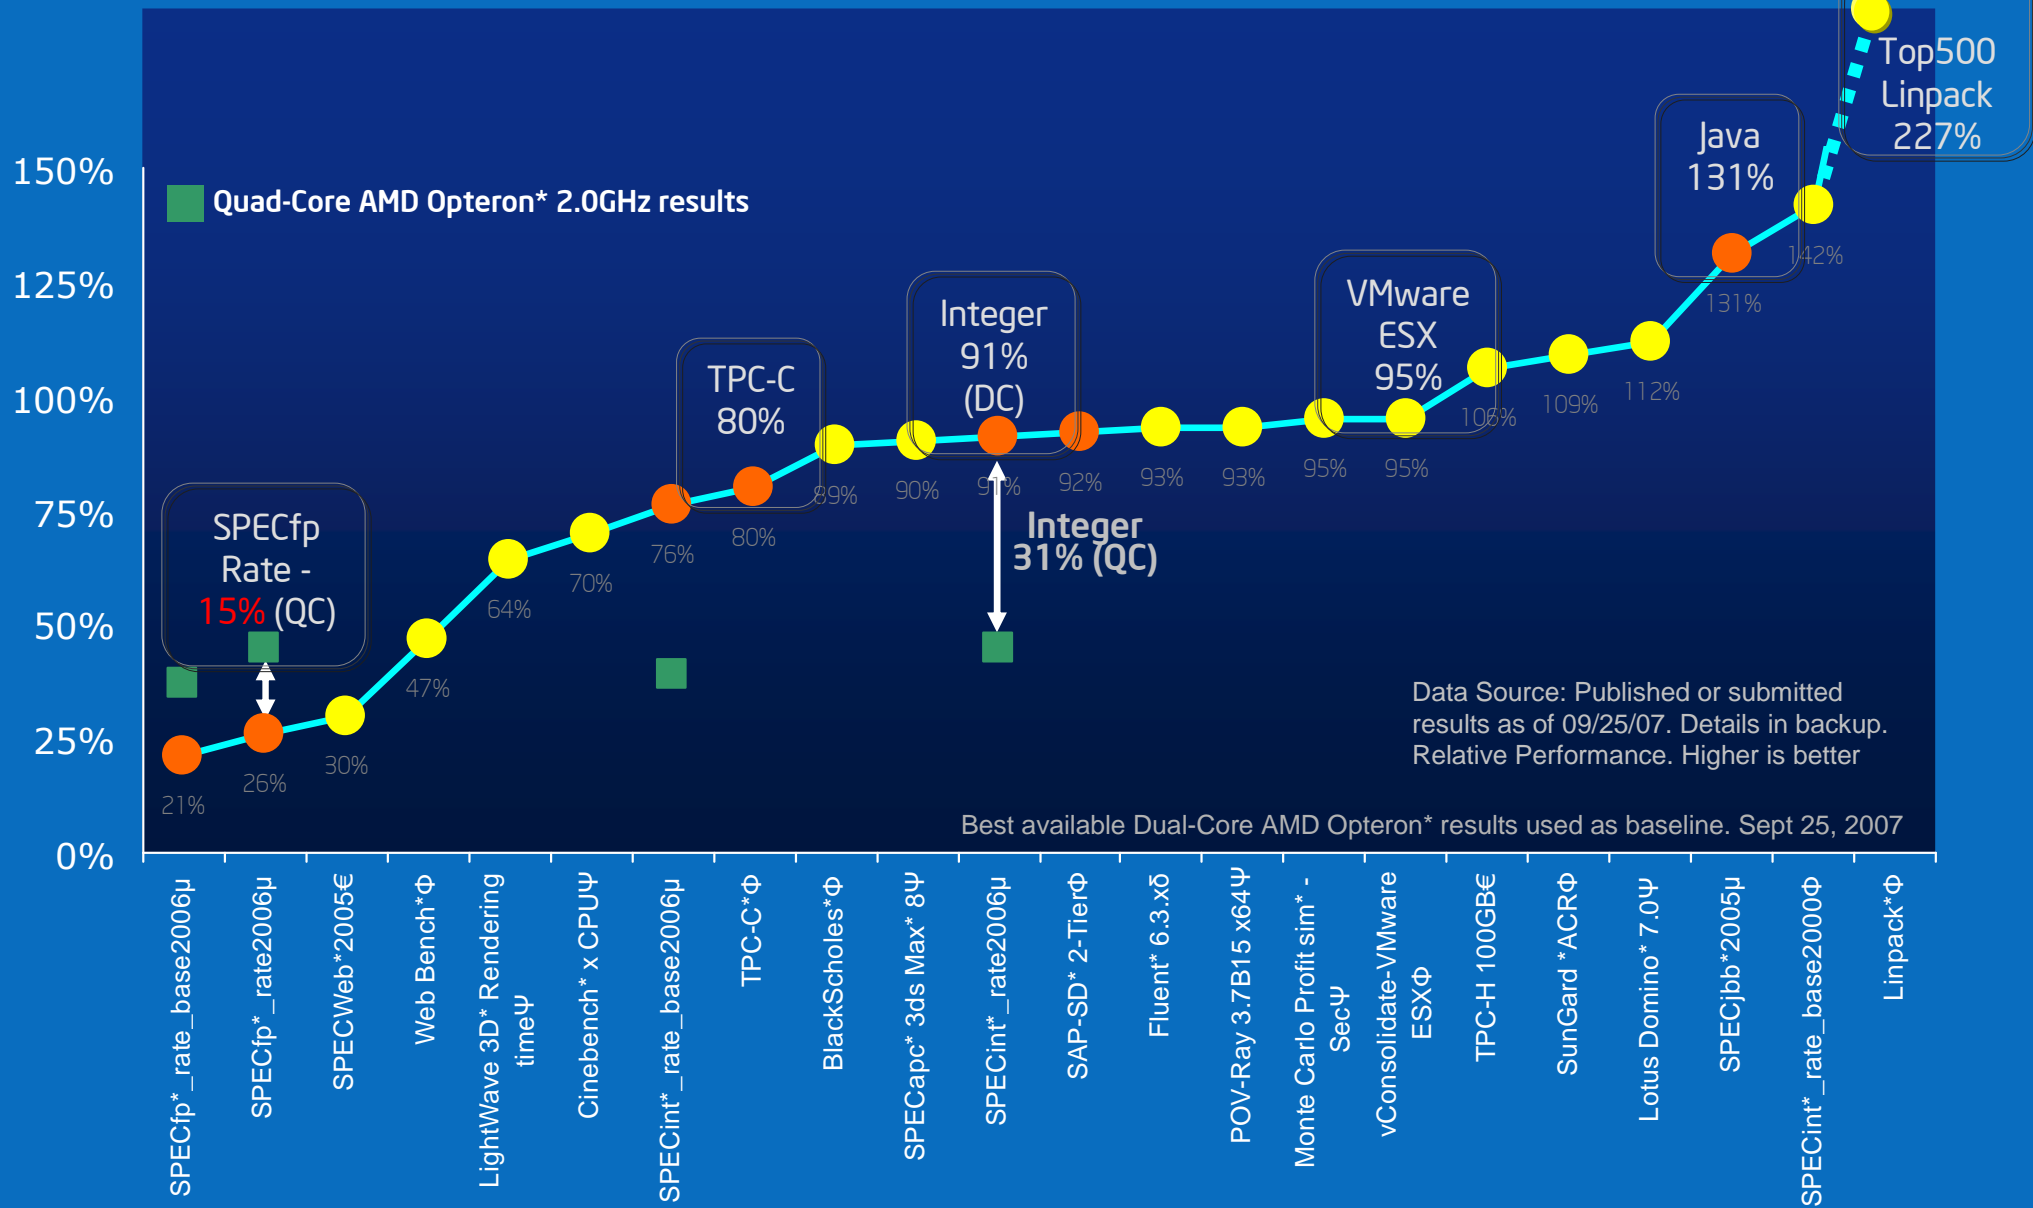
* Other names and brands may be claimed as the property of others.

Copyright © 2007-2007 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon and Intel Core are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.



Quad-Core Intel® Xeon® Processor 5300 Series based Servers

Performance Comparison today on X5365 and X5355



Performance tests and ratings are for informational purposes only and do not represent a warranty. Performance may vary due to system configuration and other factors. © 2007 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Other names and brands may be claimed as the property of their respective owners.

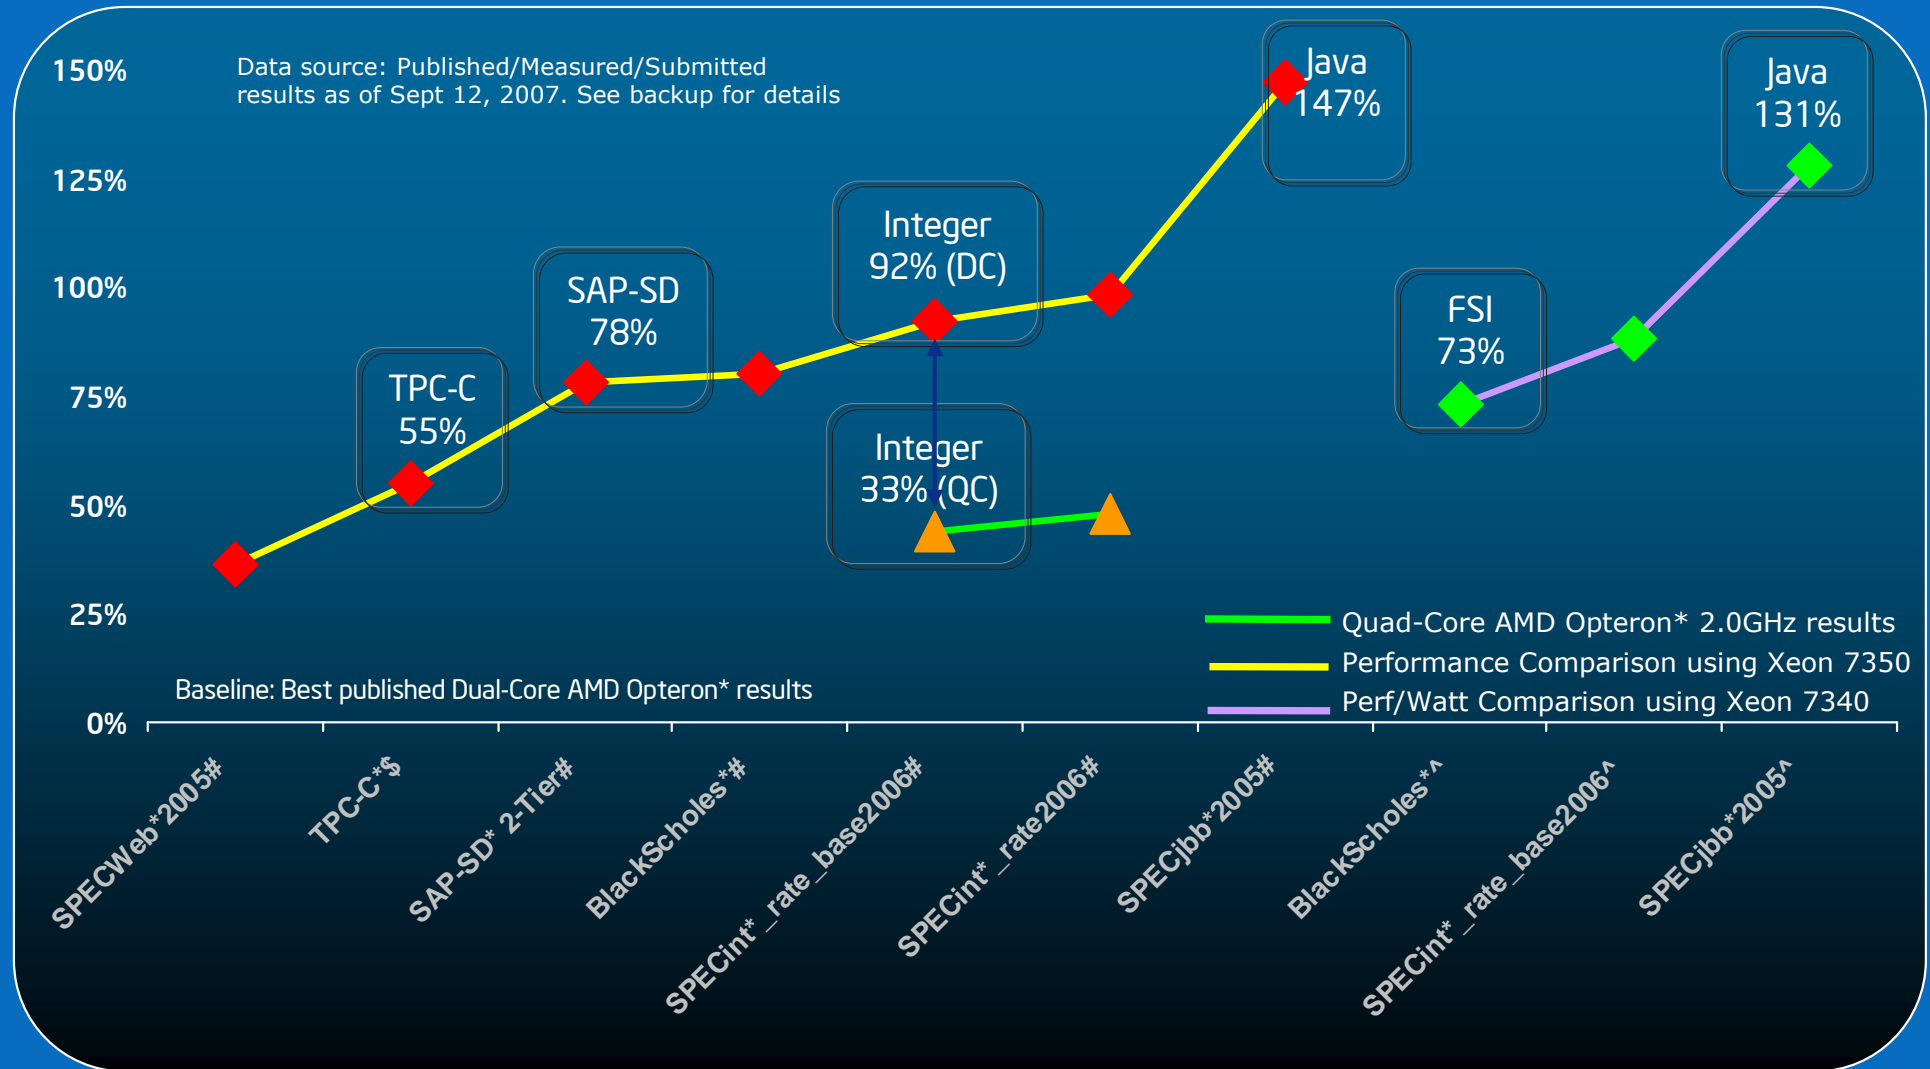


Xeon X5365 - Quad-Core Intel® Xeon® Processor X5365; Xeon X5355 - Quad-Core Intel® Xeon® Processor X5355; μ Dual-Core AMD Opteron* Model 2224SE; € Dual-Core AMD Opteron* Model 2222SE φ Dual-Core AMD Opteron* Model 2220SE (2.80 GHz); ψ Dual-Core AMD Opteron* Model 285 (2.60 GHz) δ Dual-Core AMD Opteron* Model 280 (2.40 GHz);



Quad-Core Intel® Xeon® Processor 7300 Series based Servers

Comparison to AMD Opteron* MP on Performance and Energy-Efficiency (Perf/Watt)



Xeon 7350 – Quad-Core Intel® Xeon® Processor X7350; Xeon 7340 – Quad-Core Intel® Xeon® Processor E7340 ; # Dual-Core AMD Opteron* Model 8222SE (3.0 GHz); \$ Dual-Core AMD Opteron* Model 8222SE(2.80 GHz); ^ Dual-Core AMD Opteron* Model 8220(2.80 GHz, 95 Watt TDP);

Performance tests were conducted using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit <http://www.intel.com/performance/resources/limits.htm> or call (U.S.) 1-800-628-8686 or 1-916-356-3104. Copyright © 2007, Intel Corporation. * Other names and brands may be claimed as the property of others.



Sun-Intel Solaris Collaboration is Significant



**VALENTINE'S DAY IS
ALMOST HERE, BUT WE HAD
TO PRE-ANNOUNCE A
MATCH MADE IN HEAVEN.**

In a landmark move, Intel® has endorsed Solaris™ as the mission-critical, enterprise-class UNIX operating system for the Intel Xeon® processor. Sun will deliver a comprehensive set of Intel Xeon processor based systems. Solaris is open source, free, secure and available with world-class support and over 4,000 applications. We think this makes Solaris the better choice for all your enterprise computing needs, including virtualization, high-performance computing and the Web.

With award winning technology and over seven million Solaris licenses downloaded in the last two years, we think Intel made a wise choice. Thanks Intel!



Download Solaris free for HP, IBM, Dell and over 750 other platforms at sun.com/solaris.

- Broad multi-year strategic alliance
- Sun roadmap commitment
 - 1P, 2P, 4P, >4P
 - **Telco, WS, Enterprise**
- Intel endorsement of Solaris* as a mainstream OS for Intel® Xeon® processors
- Joint investment in engineering, design, and marketing alliance for Solaris (and Java*)

Get the best software and hardware for mission-critical applications



Solaris on Xeon Experience

- What worked before (Solaris 10):
 - Over 800 x86 systems supported in Hardware Compatibility List <http://www.sun.com/bigadmin/hcl>
 - 146 Intel-based servers supported by Solaris 10 (vs. 86 AMD servers, 61 SPARC servers)
 - Performance optimizations to win World Record SPECint using Sun Studio 12
 - Majority of existing certified x64 Solaris applications already run on Intel Arch.
 - Fully supported by Sun on Intel servers, workstations, mobile, and desktop from multiple OEMs
- What we're improving:
 - Optimize performance for Intel Architecture processors and platforms (kernel and tools)
 - Support for Intel technologies important for server/workstation
 - VT, DBS, IOAT, TXT
 - Ensure Xeon support is the best for key Solaris technologies
 - Predictive Self Healing (FMA), DTrace, ZFS
 - Provide desktop/notebook silicon drivers
 - Intel graphics, wireless, wired LAN, Management, ICHx Storage



Areas of Development

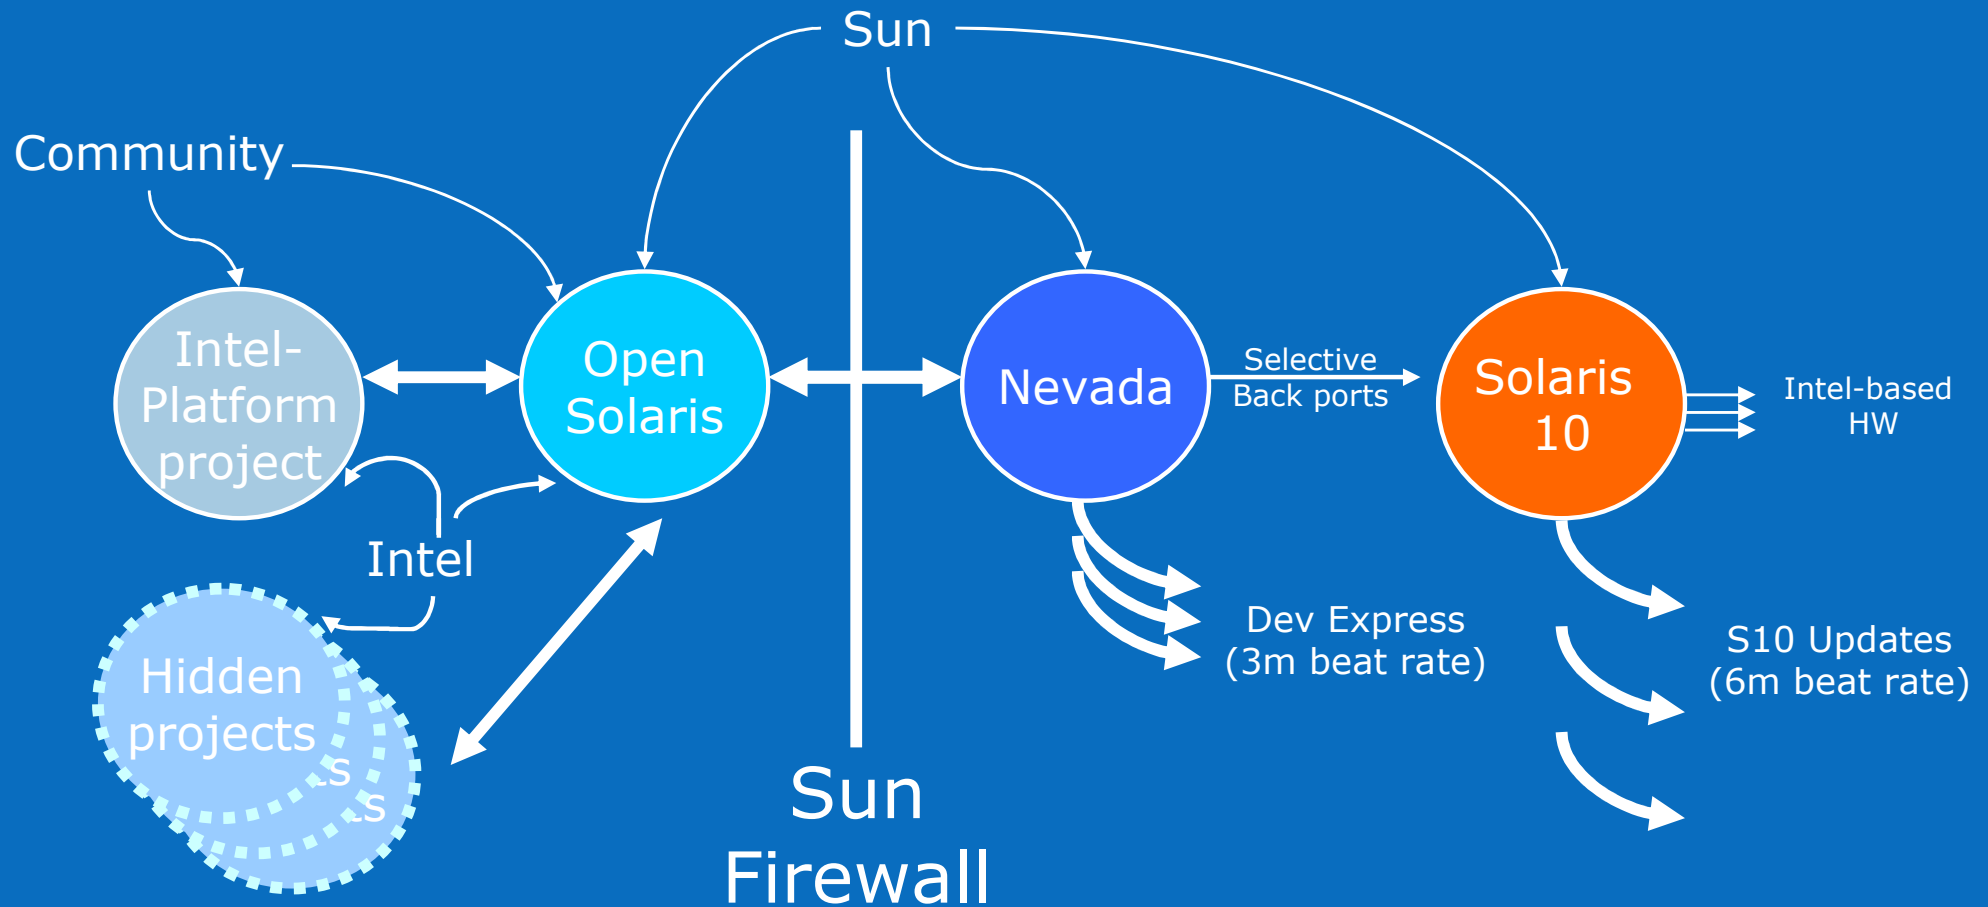
- Performance Enhancement
- CPU Performance Counters
- Compiler Vectorization and Tools
- Power Management
- Driver Support
- I/O Acceleration Technology
- Virtualization
- Predictive Self-Healing

Join us: <http://opensolaris.org>

Joint work touches most significant areas of OS



Solaris Development Model



non-Sun employees are now OpenSolaris code committers
The next step: Free Nevada from behind the firewall



Intel Core Silicon Enabling

- microcode update – increased serviceability
- iommu – increased DMA capability and security
- extended xAPIC – extend processor addressability for interrupt delivery (up to 4G-1 cores)
- ICHx – LAN, RAID, AHCI, managibility (AMT/ASF)
- CPUID – lead to 4.5% SPECint performance on C2D
- MONITOR/MWAIT – replaced halt leads to 1.2x in certain microbenchmarks

Join us at <http://opensolaris.org/os/project/intel-platform/>



Performance Enhancements

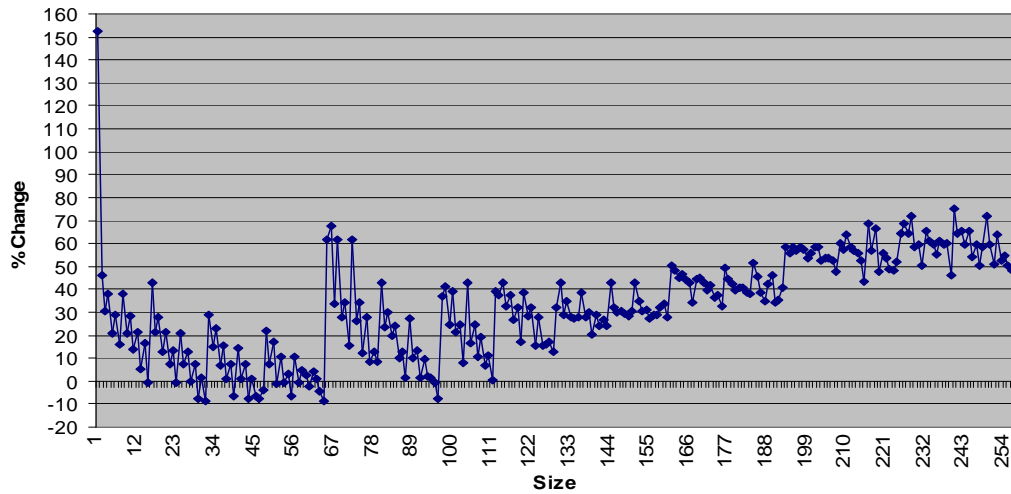
- Goal: Use Intel current and future technologies to improve Solaris performance
 - Libc optimizations
 - memcpy(), memmove(), and memset()
 - Optimized to use SSE2 and/or SSSE3 instructions
 - Significant performance improvements as measured by libMicro
 - Available soon in OpenSolaris
 - Work in progress
 - Libc string primitives (such as, strcpy, strcmp, strlen, etc.)
 - zlib
 - libm
 - Kernel optimizations in progress
 - bzero(), bcopy(), kcopy(), etc.
 - zlib functions for ZFS




Memset Performance Gains

Workload: LibMicro 0.4.0; Baseline: OpenSolaris Nevada Build 70; Processor: Intel® Core™ 2 Duo CPU E6850 @ 3GHz

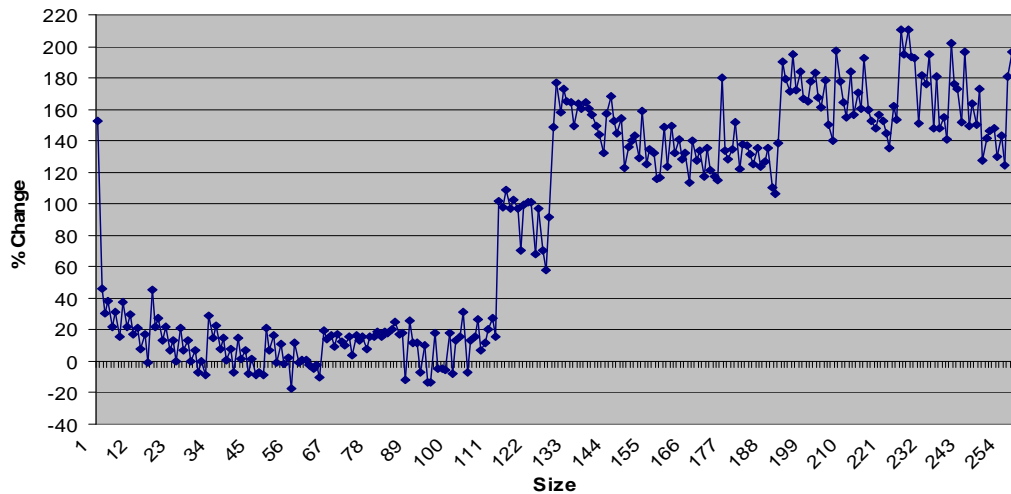
Memset Small (Aligned)



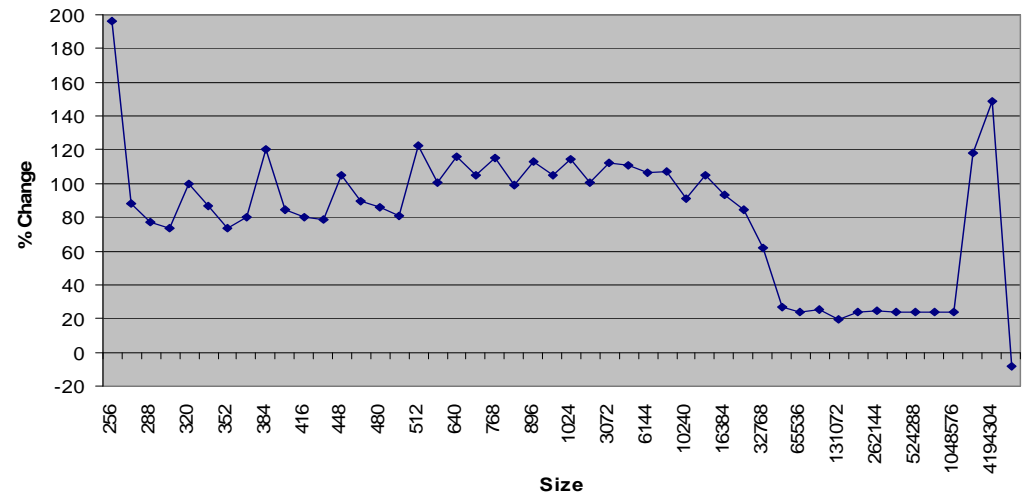
Memset Large (Aligned)



Memset Small (Unaligned)



Memset Large (Unaligned)

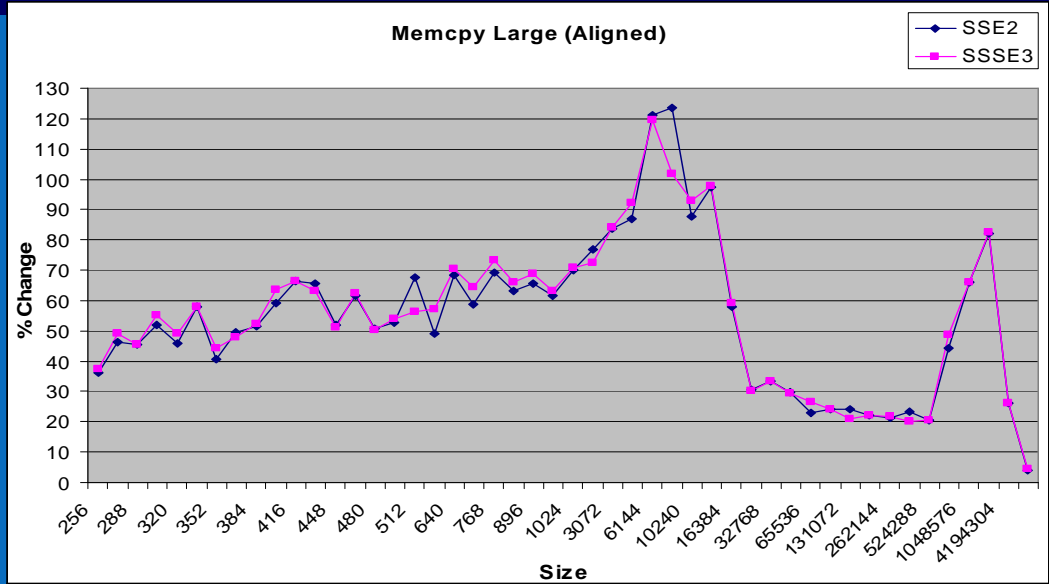
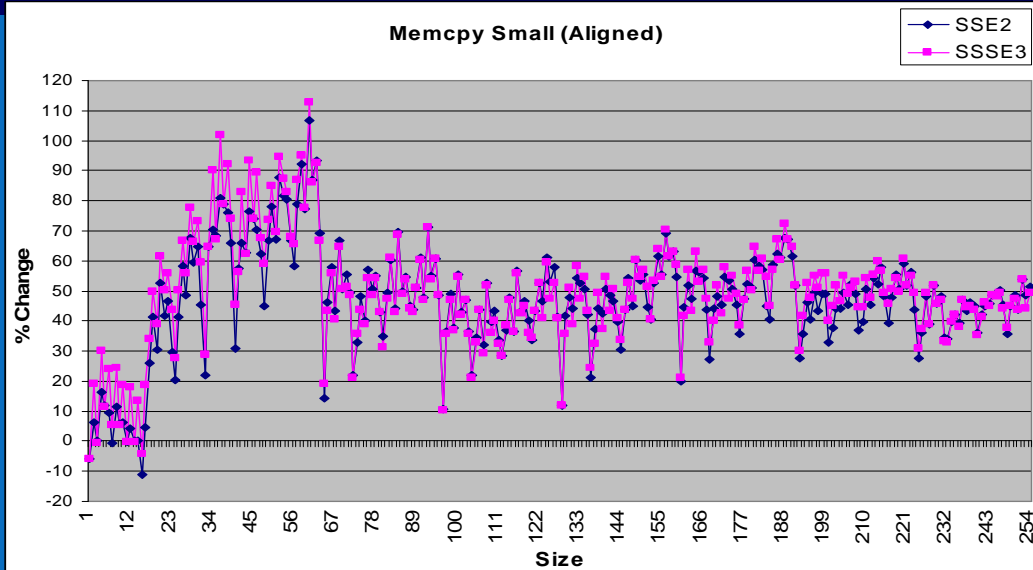


Significant gains on frequent small sizes; Using SSE2 provides significant speed up for larger sizes



Memcpy Improvements

Worload: LibMicro 0.4.0; Baseline: OpenSolaris Nevada Build 70; Processor: Intel® Core™ 2 Duo CPU E6850 @ 3GHz



Power Management

- P-states

- Performance states. Different P-states are at different frequency and voltage. You actually save energy.

- C-states

- C0 - you execute code, no other mode executes code
- C1 - HALT instruction, no instructions get executed
- C2 - like C1, no code executed. Clock stopped
- C3 - FSB shut down. No snooping and caches can be shut off

- T-state

- Emergency brake

- Parts of ACPI:

- Static Tables that the BIOS creates
- Captures the platform power capabilities (how many P/C-states, power, switch latency, etc.)

Stay as long as you can in C-states



Power Management Development Areas

- Establish baseline for current builds
 - Use **SPECPower** and **wall power meter** to identify opportunities
- Solaris Power profiling (like Powertop)
 - **Characterize C-states' behaviors** to identify barriers (for each core)
- Lots of kernel improvement areas
 - **Tickless kernel**
 - **Power-friendly scheduling**
 - **P-State improvement**
 - **C-state support**
 - **HPET timer**
 - **Interrupt binding**
 - **S-states support in client/server (driver compatibility challenge!)**

Join us at <http://opensolaris.org/os/project/tesla/>



IO Acceleration Technologies

Supported Features		Intel® I/O Acceleration Technology		
		Intel® GbE Controller (Gilgal, Ophir)	Intel® 82575 Gigabit Controller, Intel® 82598 10GbE Controller	Next gen Gigabit Controller, Next gen 10GbE Controller
Intel® QuickData Technology		✓	✓	✓
LAN stateless offloads	Header/data split	✓	✓	✓
	Receive Side Scaling	✓	✓	✓
	TX/RX checksum offload	✓	✓	✓
	TCP segmentation	✓	✓	✓
	Header-splitting / replication		✓	✓
	Receive Side Coalescing		✓ (Intel® 82598 10GbE Controller)	✓ (Next gen 10GbE Controller)
Message Signaled Interrupts		MSI	MSI-X	MSI-X
Direct Cache Access			✓	✓
Low Latency Interrupt			✓	✓

IOAT v1 is already supported and IOAT v2 in progress



Today's Virtualization Usage Models

	Static Server Consolidation	Mainframe Migration	Multi-OS Workstation	High Availability/ Disaster Recovery	Dynamic Load Balancing
End User Value	Reduce CapEx, increase utilization	OS and HW freedom for mission critical applications	Workstation Consolidation without compromise on graphics performance	Maintain high levels of business continuity	Reduce OpEx, streamline resource utilization balancing real-time computing demands with capacity
Time					

Virtualization Options in Solaris

Trend to flexibility

Trend to isolation

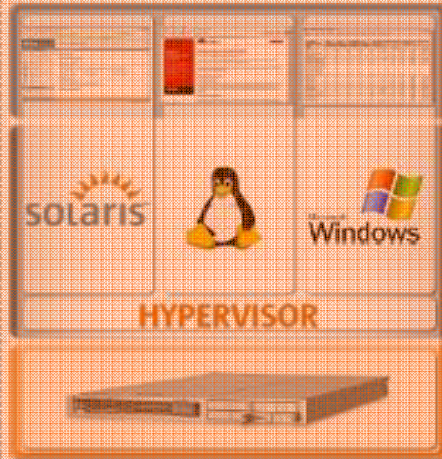
Hard Partitions



Multiple OSes

Dynamic System
Domains

Virtual Machines



Xen
Logical Domains
VMware

OS Virtualization



Single OS

Solaris Containers
(Zones + SRM)
BrandZ
Trusted Extensions

Resource Mgmt.




Solaris Resource
Manager (SRM)

Intel is supporting Sun on Solaris-based XEN



Without hardware support



- What the VMM Does ...
- Emulates a complete hardware environment for every Virtual Machine
- Allocates platform resources
- Isolates execution in each virtual machine

Virtualization solutions without hardware support work, but there are limitations and require frequent software intervention

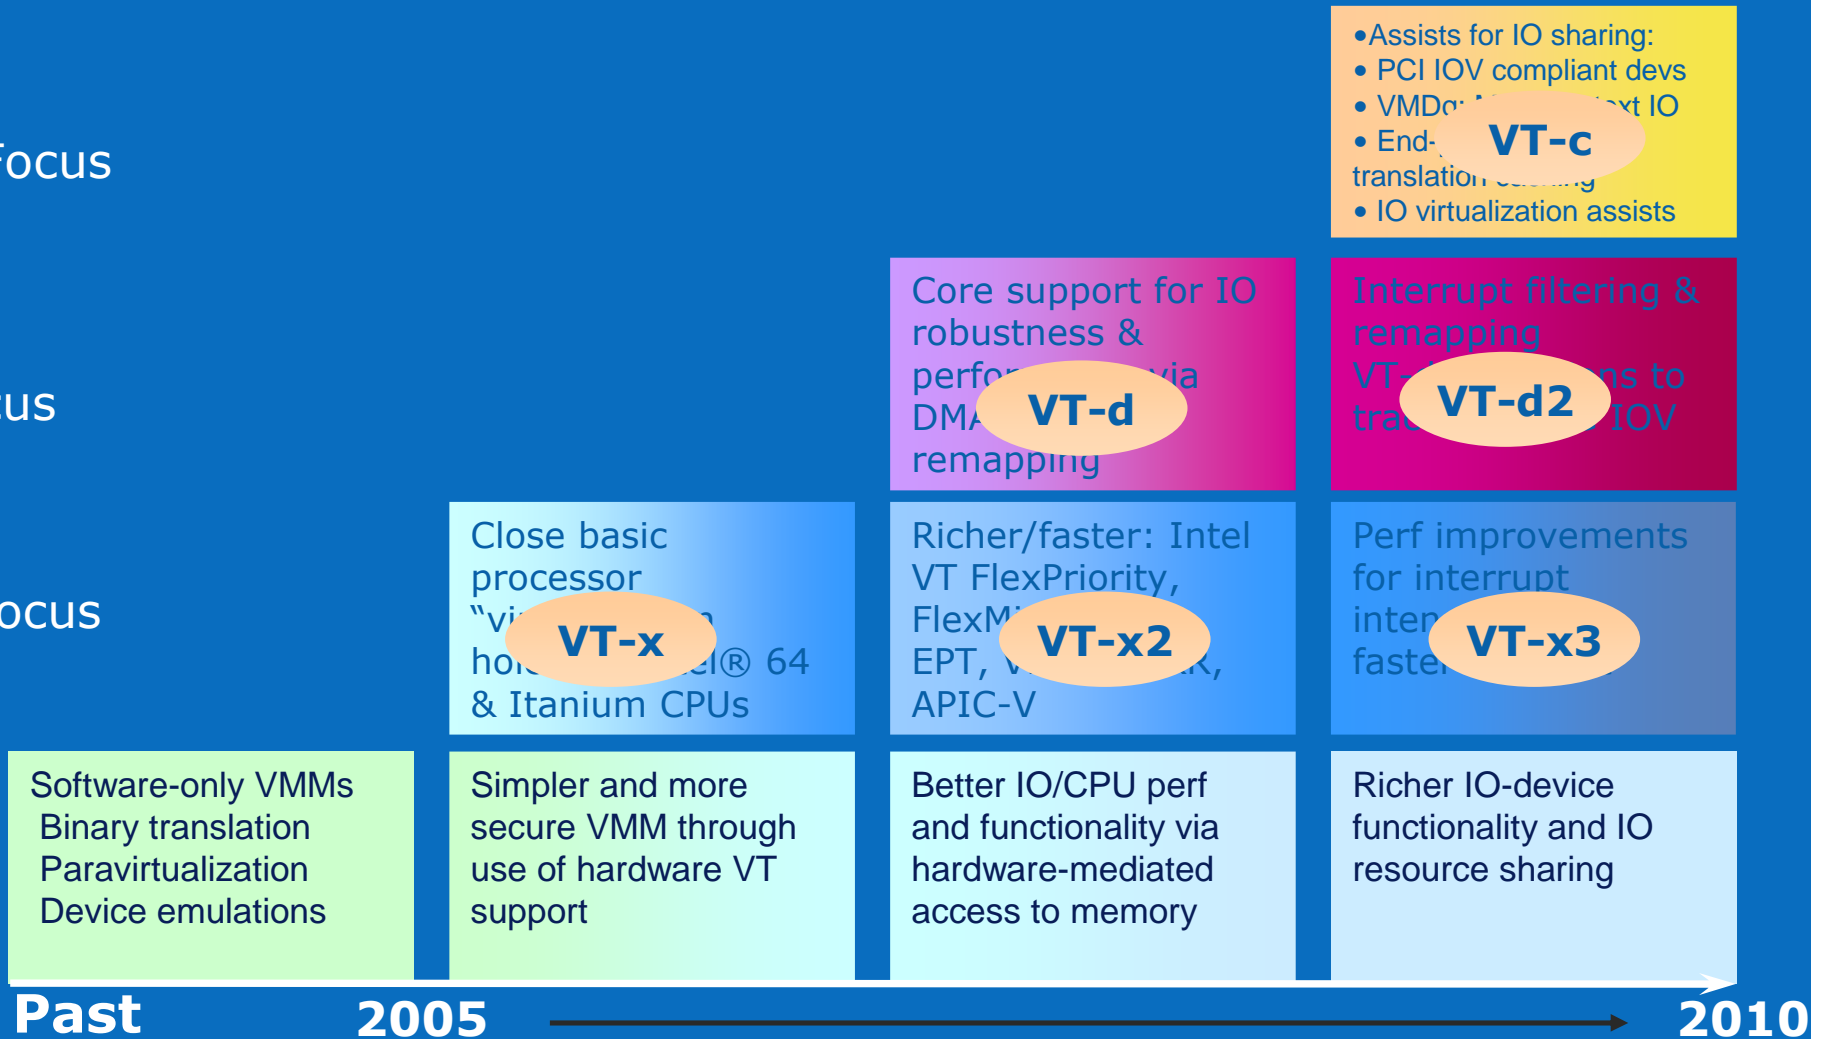
Intel® Virtualization Technology Evolution

Vector 3:
IO Device Focus

Vector 2:
Chipset Focus

Vector 1:
Processor Focus

VMM
Software
Evolution



VMM software evolution over time, with hardware support

We are adding vt-d, vt-d2, vt-x, and vt-x2 into Solaris xVM



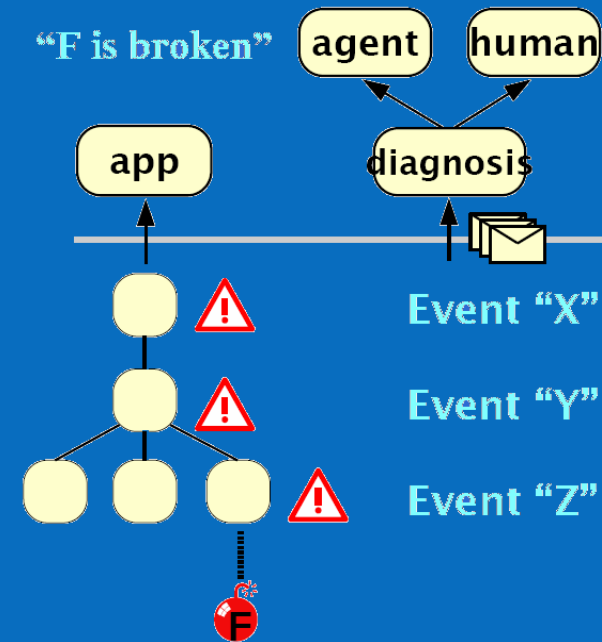
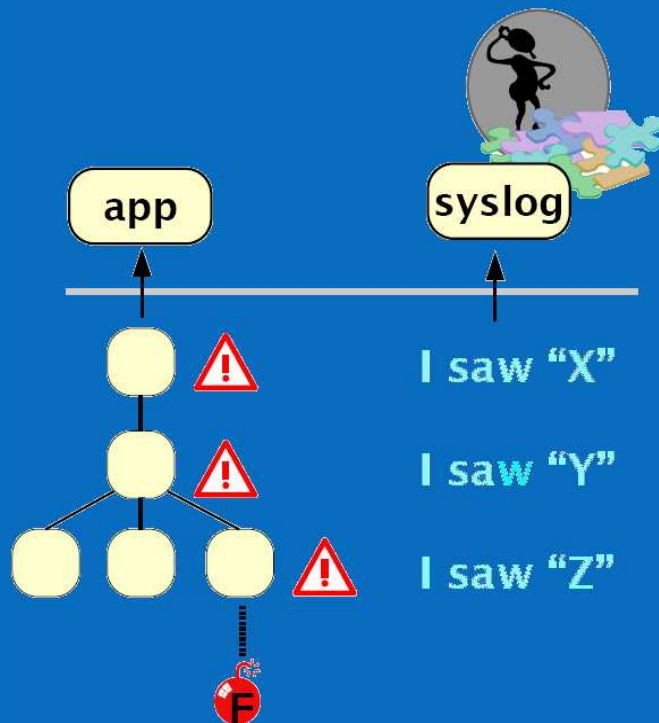
Xen Enabling in Solaris

- Xen project being led by Sun CTO Tim Marsland, staffed and running for past 1 year
- Have complete matrix working, target to be complete by April
 - Solaris as Dom0 (32 and 64bit)
 - Solaris/Linux as DomU (32 and 64 bit)
 - Solaris/Vista/Linux as HVM guest (VT-x)
- Although all combinations engineered, not all tested regularly
- Development opportunities
 - Solaris/XEN testing facilities
 - Solaris/XEN performance measurements and optimization
 - Solaris and Linux equivalence (such as) in Xen security
 - Partnering on new technologies – VT-x2, VT-d, etc

Join us at <http://opensolaris.org/os/community/xen/>



Fault Management Architecture



- Error – an incorrect signal, datum, result
- Observation that is a symptom of a fault
- Old systems only know how to report errors
- Diagnosis left to humans

- Fault – a defect that may produce errors
- The output of the diagnosis of errors
- Something we can associate with an impact and a corrective action
- Diagnosis software automates the steps

Fault Management Architecture

- **Current status**
 - Available in upcoming update to Solaris 10
 - Error Handling and Diagnosis support for Intel Machine Check Architecture
 - Cache, TLB, internal errors, CPU offline
 - PCI/PCI-X/PCIE and Memory Error Handling for the Intel platforms based on the Intel 5000 series chipsets
- **Development opportunities**
 - Error Injection tools and infrastructure
 - Develop for next generation chips and chipsets
 - Error handlers
 - Error injector and test harness
 - Fault diagnosis rules based on supplied error telemetry
 - Topology: DIMM, rank, chip, cores, caches, bridges, buses...
 - Fault Response
 - CPU offline, page retire, online spare...
 - Event and article repository information

Join us at <http://opensolaris.org/os/community/fm>



Developer Tools

- Sun Studio 12 Compiler (released June 2007) with selected Xeon-specific optimizations
 - INT optimizations (world record on SPECint)
 - FP optimizations are available in patch and will release soon
- Sun Studio Performance Analyzer will contain latest Intel Architecture performance counters
- Joint announcement on Solaris support for open source version of Intel's Thread Building Block 2.0 (multi-core programming paradigm for C++ developers)
- Transitive® QuickTransit®-
 - Run Solaris/SPARC binaries on Solaris/Xeon
- Solaris/SPARC to Solaris/Xeon is an easier transition than Solaris to Linux
 - Not just from source recompile, also includes SMF, ZFS, stability...



Sun Studio Compiler Optimization Flags

- **Aggressive – For large projects**
- `-fast -xtarget=woodcrest -m64 -xvector=simd,lib -xipo -xprofile=collect/use -Wu,-sched_first_pass=1`
- `-xtarget=woodcrest` expands to `"-xarch=ssse3 -xchip=core2 -xcache=32/64/8:4096/64/16"`
- SSE3 code generation, core2 architecture optimization and cache configuration selections, O5 level optimization, and inter-procedural optimization
- Enable instruction scheduler for FP calculation on IA
- Profile guided optimization
- **Medium – For most applications**
- `-fast -xtarget=woodcrest -m64 -Wu,-sched_first_pass=1`
- All aggressive optimization but no IPO and profile guidance
- **Low – For extra precise floating point calculations**
- `-O -xtarget=woodcrest -m64`
- O3 (medium) optimization level
- Quickest compilation



Desktop/Mobile Driver Support - Wireless Driver

- www.opensourcewireless.org
- Target both OpenSolaris & BSD communities
- Focus on 4965 specific documentation
- High-level architecture and uCode initialization
- Downloadable uCode and dual licensed header files
- Phase 1 – expected completion Nov.
 - 802.11 A/G
 - infrastructure mode
 - power/temperature calibration (FCC regulatory)
 - WEP
- Phase 2
 - 802.11n
 - IBSS(ad-hoc) mode
 - WPA

4965 wireless driver is working and improvement on the way



Device Drivers Support - Others

- Graphics
 - All Intel graphic silicon is supported
- ASF/AMT
- Matrix Storage – IMSM
- NICs
- Others
 - Audio codec, USB, etc.

Intel platform laptop/desktop is supported.



Summary/Call to Action

- Intel platform and Solaris bring the best technology to end user
- Intel and Sun teams at full strength through the community
- Result is significant in various kernel areas
 - Performance, drivers, FMA, virtualization, etc.
- Call to action
 - Run OpenSolaris/Solaris on latest Intel server platforms
 - Joint development with us at OpenSolaris projects



Leap ahead !

