



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Fast Track to Solaris 10 Adoption: 64-bit Performance A Sun Expert Exchange Discussion


In a recent Sun Expert Exchange, Sun engineers answered questions about 64-bit Performance, the powerful new technology upon which the Solaris 10 OS is based. On top of strengthened 32-bit support for Xeon Processor MP, the Solaris 10 OS includes added 64-bit support for AMD Opteron processors and ensured optimal performance for our new UltraSPARC IV processors lineup.

This summary includes highlights of the hour-long Q&A,* organized into the following sections:

- General Information..... Pages 2-4
- Documentation & Training..... Page 5
- Performance Issues..... Pages 6-9
- Installation & Configuration..... Page 10
- Compatibility Issues..... Pages 11-13
- Functionality & Usability Issues..... Page 14

In addition to questions and answers, you'll also find references and links to additional resources provided by Sun. For future reference, you can download a pdf version of this summary.

*Note: The information contained in this transcript, taken directly from a live Sun Expert Exchange event, has been edited for clarity and adherence to trademark guidelines



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Fast Track to Solaris 10 Adoption: 64-bit Performance

General information

1. When we may expect the official release of the Solaris 10 OS?
2. Other than very large virtual address spaces, what other benefits can 64-bit applications provide?
3. What is Sun's intent for Java Desktop System- which is SuSE Linux-based — compared to the Solaris 10 x86? What should customers choose?
4. What are Sun's plans for AMD Opteron technology in the blade market?
5. Why did you make the decision to use D instead of an existing language for DTrace?
6. Does Sun sell the AMD64 system?
7. Are there plans to introduce any distributed file system, multi-writers features in ZFS? QFS is so expensive, but multi host writing, a single file system, would be a killer feature.
8. Are Zones supported on all currently shipping systems, in particular the Illi servers? Are Zones supported on single-processor systems?
9. When is Sun Studio 10 software expected to ship — and ship to maintenance customers?
10. Given that the kernel in the Solaris 10 OS is larger (more lines of code) because of new functions such as DTrace, how can it perform faster?

Q: When we may expect the official release of the Solaris 10 OS?

A: January 31, 2005.

Q: Other than very large virtual address spaces, what other benefits can 64-bit applications provide?

A: Unlike SPARC, where the number of registers is the same between 32- and 64-bit apps, with AMD64 you get eight additional registers. Plus, the calling convention for 64-bit C apps is faster since the first six parameters are passed in registers; 32-bit i386 apps must pass all parameters on the stack.

Q: What is Sun's intent for Java Desktop System- which is SuSE Linux-based — compared to the Solaris 10 x86? What should customers choose?

A: The Sun Java Desktop System is also available on the Solaris OS. With the Solaris 10 OS, customers can take advantage of all of the new release's breakthrough technologies, with the affordable, user-friendly, fully integrated desktop client environment, which lets customers choose the right tool for the job.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: What are Sun's plans for AMD Opteron technology in the blade market?

A: Stay tuned for more about Sun's blade roadmap. This is under consideration, though nothing has been announced.

Q: Why did you make the decision to use D instead of an existing language for DTrace?

A: I'm not a DTrace expert, but DTrace has a major advantage in that it can be used in production environments. To do this, we needed to ensure the safety of the tool in production systems. D helps avoid shooting yourself in the foot by limiting certain things you could do in, say, C, the prime example being looping constructs. In D, you cannot create loops; in C or most other languages, you can create loops - infinite ones in fact. The D language also has a number of other unique and useful features, such as aggregations, strings as first-class citizens, and thread local variables. See <http://www.sun.com/bigadmin/content/dtrace/> for more details.

Q: Does Sun sell the AMD64 system?

A: Yes, today Sun sells a line of AMD64-based systems, including the Sun Fire V20z server, Sun Fire V40z server, Sun Java workstation W1100z and Sun Java workstation W2100z. More systems to come in the next calendar year, so stay tuned. Also, the Solaris 10 OS will run on systems from other AMD system vendors. Watch the Solaris OS hardware compatibility list grow.

Q: Are there plans to introduce any distributed file system, multi-writers features in ZFS? QFS is so expensive, but multi host writing, a single file system, would be a killer feature.

A: Initially ZFS will be host based; distributed file system capabilities are being evaluated for the future.

Q: Are Zones supported on all currently shipping systems, in particular the Illi servers? Are Zones supported on single-processor systems?

A: Yes, Solaris Containers (a.k.a., Zones) is supported on all shipping systems, including single-processor systems.

Q: When is Sun Studio 10 software expected to ship — and ship to maintenance customers?

A: Sun Studio 10 software will ship concurrently with the Solaris 10 OS on Jan 31st. It will be available to maintenance/supported customers at the same time.




Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: Given that the kernel in the Solaris 10 OS is larger (more lines of code) because of new functions such as DTrace, how can it perform faster?

A: There are more lines of code due to added functionality, but much existing functionality (lines of code) have been re-tuned or implemented for optimal performance. For example, we ship multiple versions of libc tuned for AMD64, EM64T, or legacy x86 systems. The most optimal is configured based on the capabilities of the underlying platform at boot time.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members




Fast Track to Solaris 10 Adoption: 64-bit Performance Documentation & Training

Q: When will training tracks from Sun be available for the Solaris 10 OS?

A: Free, Web-based training is available today.

Visit: <http://training.sun.com/US/catalog/solaris10.html> for more information on this and other Solaris 10 OS training offerings.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Fast Track to Solaris 10 Adoption: 64-bit Performance Performance Issues

1. I can recall that one of Sun's goals for the Solaris 10 OS was no more entries in /etc/system. Have you achieved this?
2. Will any of the new trussing capabilities allow a process to be trussed twice? For example, right now a Sun Cluster managed process cannot be trussed while it is under cluster control.
3. Have the SunPro Business Unit Compilers for the Solaris 10 OS been released yet? What kinds of improvements over GCC can be expected in both the 32 and 64-bit spaces?
4. Does the Solaris 10 OS make any performance or scalability improvements? For example, the current OS has a limit of 1024 file descriptors. Will the Solaris 10 OS be able to break through that limitation?
5. For an app like BLAST, which step gives you the best increase in performance? a) move 32-bit app from 32-bit AMD to 64-bit AMD b) convert the app itself from 32-bit to 64-bit?
6. Intel recently announced the cancellation of its 4 GHz Pentium. Do you see 4 GHz as the clock-rate ceiling for single CPUs? If so, where does the AMD Opteron processor go after hitting that ceiling? If not, why is it different?
7. Had Sun run any performance comparisons between the Solaris 10 OS/AMD Opteron processor and Wintel machines running ProEngineer or Autocad? What were the results?
8. Will an application that can fit into the 8MB of cache on Ultra SPARC III 1.2 GHz CPU run faster on AMD64 Opteron where the cache is small?
9. What types of applications (CAD, Imaging, etc.) take real advantage of 64 bits?
10. Since AMD64 is faster than SPARC, why not drop SPARC and put the savings into Solaris OS/services?
11. Are there cases of a 32-bit app running slower on a 64-bit CPU?
12. Are high-end servers coming with AMD64 and the Solaris 10 OS?
13. Can you explain in a little more detail about the performance gains that AMD64 will have over SPARC64?
14. Which has better performance in the Java environment, SPARC 64 bit or AMD64?
15. Has any benchmarking been done for the Solaris 10 OS on AMD64 and Intel platforms?

Q: I can recall that one of Sun's goals for the Solaris 10 OS was no more entries in /etc/system. Have you achieved this?

A: Some mechanism for patching kernel variables will always remain. However, we have purposely eliminated some of the more common reasons to mess with /etc/system, e.g., the shared memory and other IPC tuneables. These are now modeled as resource limits that can be tweaked on the fly (see `rctladm(1M)`), or even managed on the network via the `project(4)` database.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: Will any of the new trussing capabilities allow a process to be trussed twice? For example, right now a Sun Cluster managed process cannot be trussed while it is under cluster control.

A: You still won't be able to truss a process twice; however, the issue with Sun Cluster will be addressed when Sun Cluster migrates to become a service management framework (SMF) enabled service.

Q: Have the SunPro Business Unit Compilers for the Solaris 10 OS been released yet? What kinds of improvements over GCC can be expected in both the 32 and 64-bit spaces?

A: "SunPro Business Unit," or, officially, "Sun Studio 10 Software," will be released concurrently with the Solaris 10 OS, at the end of January 2005. Sun Studio 10 Software will have considerably better FP performance than gcc/g77 (as has Studio 9; it will just continue to improve). We will also have features that gcc doesn't have: OpenMP2.0, automatic parallelization/vectorization, linker scoping, performance libraries (libm and libsunperf), much better debugging support, performance analyzer and a full-featured IDE, etc. Sun Studio 9 Software already improves over gcc for the 32-bit space; Studio 10 will go further and bring all that technology to 64 bits.

Q: Does the Solaris 10 OS make any performance or scalability improvements? For example, the current OS has a limit of 1024 file descriptors. Will the Solaris 10 OS be able to break through that limitation?

A: There are many, many performance improvements in networking, virtual memory, system calls, and filesystems. Some of that work was enabled by applying DTrace to the OS itself. The fd hard limit is 65535 in the Solaris 10 OS. Other scalability improvements include further MPO optimizations (NUMA), more physical memory on both SPARC and x86 machines, large pages, and more address space for 32-bit apps on 64-bit x86 machines.

Q: For an app like BLAST, which step gives you the best increase in performance? a) move 32-bit app from 32-bit AMD to 64-bit AMD b) convert the app itself from 32-bit to 64-bit?

A: We've already tested based on a), and have seen performance increase. For b) we predict improvement based on ABI and calling conventions.

Q: Intel recently announced the cancellation of its 4 GHz Pentium. Do you see 4 GHz as the clock-rate ceiling for single CPUs? If so, where does the AMD Opteron processor go after hitting that ceiling? If not, why is it different?

A: The short answer is, there is no "physically constrained" ceiling at 4GHz. The AMD Opteron processor is currently at 2.4GHz, with future products expected to ship into the 3.xGHz and perhaps 4.xGHz range. However, focusing solely on MHz is the wrong approach. There are many ways to improve performance.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



For example, micro-architectural enhancements to improve per-core performance; multi-core solutions to improve multi-threaded, multi-tasking, and more throughput oriented applications; system architecture improvements (like the current AMD Opteron processor's integrated memory controller and HyperTransport interconnect) to focus on system bottlenecks; power optimizations to improve cost and density. That all being said, of course, MHz cannot be ignored.

Q: Had Sun run any performance comparisons between the Solaris 10 OS/AMD Opteron processor and Wintel machines running ProEngineer or Autocad? What were the results?

A: Our ProEngineer results on the Solaris 10 OS/Opteron is more than 15 percent faster than on the Solaris 8 OS on 32-bit. Expect 64-bit to be faster than 32-bit. Comparison to Wintel is ongoing.

Q: Will an application that can fit into the 8MB of cache on Ultra SPARC III 1.2 GHz CPU run faster on AMD64 Opteron where the cache is small?

A: Yes, but your mileage may vary. AMD has a fast memory subsystem and high clock rate. Depends on what the app is doing.

Q: What types of applications (CAD, Imaging, etc.) take real advantage of 64 bits?

A: Both memory-intensive apps (database), because of increased memory capacity, and computationally intensive ones (media, CAD, imaging, simulation), because of relaxed register pressure and improved large integer multiply performance.

Q: Since AMD64 is faster than SPARC, why not drop SPARC and put the savings into Solaris OS/services?

A: Performance depends a great deal on workload type. SPARC has many performance advantages in scalability and multithreaded workloads. The new multi-core SPARC processors will demonstrate order-of-magnitude performance gains.

Q: Are there cases of a 32-bit app running slower on a 64-bit CPU?

A: No, none that we have seen.

Q: Are high-end servers coming with AMD64 and the Solaris 10 OS?

A: The SPARC product line will continue to range from 2P entry servers to the high end 100+P servers. The AMD Opteron product line includes workstations and servers up to 4P today, with larger systems planned for the future.

Q: Can you explain in a little more detail about the performance gains that AMD64 will have over SPARC64?

A: Our preliminary results show that AMD64 performance is 30-40 percent faster than 32-bit; our 32-bit Opteron processor is faster than SPARC.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: Which has better performance in the Java environment, SPARC 64 bit or AMD64?

A: We predict AMD64 will have better JVM performance.

Q: Has any benchmarking been done for the Solaris 10 OS on AMD64 and Intel platforms?

A: We have seen great performance on AMD Opteron processor-based systems with the Solaris 10 OS. 64-bit BLAST (bio-science benchmark) results:

1. The Sun Fire V40z server with a single AMD Opteron 850 processor running 32-bit Solaris Operating System shows up to 42 percent better performance than a Dell 6650 server with a 3.0 GHz Xeon running Linux.
2. The 64-bit Solaris 10 Operating System exhibits up to 31 percent faster performance than the 32-bit Solaris OS with only a recompilation of the application source code.
3. The 64-bit Solaris 10 OS exhibits incredible scalability of 95 percent on 4-way system.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members




Fast Track to Solaris 10 Adoption: 64-bit Performance Installation & Configuration

Q: Is Sun's listed B.C.P. to do fresh installs of the Solaris 10 OS? Is there a penalty for doing an OS upgrade? Any words of wisdom on Solaris 9-to-Solaris 10 OS in-place migration?

A: For general information on moving to the Solaris 10 OS, please visit: <http://www.sun.com/software/solaris/10/move.jsp>. You can also visit the Solaris 10 OS Adoption Kit at the following URL, which includes documentation and resources for testing, support, porting, and application recompile: <http://iforce.sun.com/protected/solaris10/adoptionkit/>

Q: Our SAP/Oracle apps generally have us modifying the /etc/system file to change semaphore and shared memory parameters. Does the Solaris 10 OS have the ability to dynamically change the parameters, or do we still have to tweak /etc/system?

A: In the Solaris 10 OS, the system-wide limits have been remodeled as resource controls. Resource controls can be altered on the fly, controlled by the project database, and thus managed from the network if needed. If you have a Solaris 9 OS, see `rctladm(1M)`.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Fast Track to Solaris 10 Adoption: 64-bit Performance Compatibility Issues

1. Are there features in the recent versions of Sun Studio to support 64-bit porting?
2. If we run the Solaris 10 OS on non-Sun servers, will the drivers for third-party hardware such as NIC and video be available through Sun, or is Sun developing partnerships with other hardware vendors?
3. I read about running Linux binary (x86) on the Solaris 10 OS (SPARC). Is there provision to run SPARC binaries in the Solaris OS x86 environment? This would help migrating older applications to faster AMD-based Sun hardware.
4. Last time I tried the Solaris OS, I had a problem with my LAN card. Will the Solaris 10 OS be more plug-and-play, like what we see in Linux platforms?
5. Has Sun considered making SPARC-based binary and Intel-based binary compatible in the near future?
6. Can I run the Solaris 10 OS on my IBM systems that use the AMD Opteron processor?
7. Do you know when SAP and Oracle will support the Solaris 10 OS?
8. Is the Solaris 10 OS certified for and supported on the F15K?
9. What kernel version is supplied with the Solaris 10 OS, and does it support other software, from IBM for example?
10. Will the Linux compatibility support running 32-bit Linux apps on the 64-bit Solaris 10 OS?
11. I'm assuming that I can run 32-bit applications on the 64-bit Solaris 10 OS. Oracle has 32-bit and 64-bit versions, but many "out of the box" apps don't.
12. What ERP applications (SAP, Oracle, PeopleSoft) have been or are in the process of being qualified with the Solaris 10 OS?
13. When can I use my Sun PciII card with the Solaris 10 OS?

Q: Are there features in the recent versions of Sun Studio to support 64-bit porting?

A: Both C and C++ feature lint-like warnings to indicate problem areas; lint has also been updated for AMD64-specific migration warnings. There is also documentation for migration help, including FAQ items (via the EA site).

Q: If we run the Solaris 10 OS on non-Sun servers, will the drivers for third-party hardware such as NIC and video be available through Sun, or is Sun developing partnerships with other hardware vendors?

A: Typically, we partner with hardware vendors for certification, and we also allow vendors to self-certify using the hardware compatibility test suite (which we provide as a free download) See: sun.com/bigadmin/hcl



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: I read about running Linux binary (x86) on the Solaris 10 OS (SPARC). Is there provision to run SPARC binaries in the Solaris OS x86 environment? This would help migrating older applications to faster AMD-based Sun hardware.

A: The Janus project allows you to run x86 Linux apps unmodified on the Solaris OS running on x86-based platforms, but not on SPARC-based platforms. We currently do not have any way to run existing SPARC apps on x86 platforms.

Q: Last time I tried the Solaris OS, I had a problem with my LAN card. Will the Solaris 10 OS be more plug-and-play, like what we see in Linux platforms?

A: Device support is vastly improved with the Solaris 10 OS; please reference the BigAdmin Services HCL list.

Q: Has Sun considered making SPARC-based binary and Intel-based binary compatible in the near future?

A: Yes, we have looked into this. The most portable approach is to use Java technology, which allows for apps to be run on any platform.

Q: Can I run the Solaris 10 OS on my IBM systems that use the AMD Opteron processor?

A: Yes, the Solaris 10 OS supports a wide range of third-party hardware, including IBM. See sun.com/bigadmin/hc for a list of supported platforms.

Q: Do you know when SAP and Oracle will support the Solaris 10 OS?

A: We expect both Oracle and SAP to support the Solaris 10 OS when it releases.

Q: Is the Solaris 10 OS certified for and supported on the F15K?

A: Yes.

Q: What kernel version is supplied with the Solaris 10 OS, and does it support other software, from IBM for example?

A: The kernel version is SunOS 5.10. That's the Solaris OS kernel. We're working on IBM software commitments.

Q: Will the Linux compatibility support running 32-bit Linux apps on the 64-bit Solaris 10 OS?

A: Yes. We will support 32-bit Linux apps on the 64-bit Solaris OS kernel.



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Q: I'm assuming that I can run 32-bit applications on the 64-bit Solaris 10 OS. Oracle has 32-bit and 64-bit versions, but many "out of the box" apps don't.

A: As with the Solaris OS on SPARC platforms, on AMD64 platforms you'll be able to run your existing 32-bit apps unmodified, side-by-side, with new 64-bit apps.

Q: What ERP applications (SAP, Oracle, PeopleSoft) have been or are in the process of being qualified with the Solaris 10 OS?

A: SAP is qualified, PeopleSoft is not.

Q: When can I use my Sun Pcell card with the Solaris 10 OS?

A: You will be able to view the Solaris 10 OS HCL in January, which will list the cards that are supported. View the HCL in progress via the Software Express program at sun.com/bigdmin/hcl



Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members



Fast Track to Solaris 10 Adoption: 64-bit Performance Functionality & Usability Issues

1. This may be more of a hardware question, but how does the Solaris 10 OS handle the obp issue on Sun's AMD Opteron processor-based servers?
2. Are there any developments on centralized patch management packaged with the Solaris 10 OS?
3. Can I run 32-bit zones in a 64-bit global zone?
4. From an existing SPARC platform, is it that easy for a C/C++ program to run without noticing a 64-bit environment?
5. Replication mechanisms and volume snaps are becoming operating system features rather than add-on products, particularly in the Wintel space. Is Sun moving in that direction with the Solaris 10 OS and SNDR?

Q: This may be more of a hardware question, but how does the Solaris 10 OS handle the obp issue on Sun's AMD Opteron processor-based servers?

A: There are none. Solaris OS for x86 platforms uses the industry standard BIOS.

Q: Are there any developments on centralized patch management packaged with the Solaris 10 OS?

A: You will have access to the Sun Update Network, called CNS (Customer Network System) functionality in the first update release of the Solaris 10 OS.

Q: Can I run 32-bit zones in a 64-bit global zone?

A: Yes, you can.

Q: From an existing SPARC platform, is it that easy for a C/C++ program to run without noticing a 64-bit environment?

A: I'll assume that you are migrating from SPARC 32-bit to AMD64. Yes, there may be problems if the application uses, say, int to hold an address; as for AMD64, pointers are 64-bit.

Q: Replication mechanisms and volume snaps are becoming operating system features rather than add-on products, particularly in the Wintel space. Is Sun moving in that direction with the Solaris 10 OS and SNDR?

A: Yes, we are moving this way. The ZFS file system includes a very fast mechanism for snapshots