

[Sign In/Register](#) [Help](#) [Country](#) [Communities](#) [I am a...](#) [I want to...](#)[Products](#) [Solutions](#) [Downloads](#) [Store](#) [Support](#) [Training](#) [Partners](#)[About](#) [Innovation](#) [SPARC T5 Deep Dive](#)

SPARC T5 Deep Dive

An interview with Oracle's Rick Hetherington

Stay Connected:

Rick Hetherington, Oracle's vice president of hardware development, manages a team of architects and performance analysts who design Oracle's M- and T-series processors. In this interview, Hetherington describes the technical details of the new SPARC T5 processor and expands on the process that is used to design these innovative chips.

Q: What were the design objectives of the SPARC T5 processor?

A: The strategy for the SPARC T5 was to improve on SPARC T4 in every possible dimension. And the team has done this in a bit over one year's time. We wanted to increase processor performance both in terms of single thread as well as throughput. And then increase memory bandwidth, I/O bandwidth, and finally socket scalability.

Q: So what's new in the SPARC T5?

A: You may recall that we designed a brand new core for the SPARC T4 processor called the S3, which was built in 40 nanometer technology. We brought that same S3 core into the SPARC T5 and doubled the number of cores to 16 so that throughput went up by more than a factor of two. And because we were in a new generation of silicon at 28 nanometers, we had the opportunity to increase the clock from 3GHz to 3.6GHz. SPARC T5 integrates PCI express on die much like the SPARC T4 did, but we are now deploying PCI Express Rev 3, which means that we doubled the I/O bandwidth. DDR3 on-die memory controllers are similar to what we have on SPARC T4 but there are now 4 controllers so that will double memory bandwidth.

Q: This is mainly a performance increase story then?

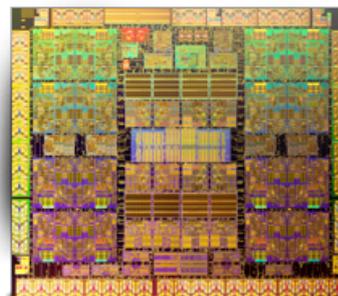
A: It is, and we are delivering what our customers have requested. We've also increased our ability to scale the SPARC T5 to eight sockets. So with 16 cores, SPARC T5 has the capability of gluelessly interconnecting eight sockets for a total of 128 cores. Each core can support up to eight independent threads or strands. From an operating system point of view, there are a total of 1024 CPUs. That means that customers now can enjoy over 100 physical SPARC cores, over 1000 threads or CPUs, and many terabytes of physical memory in a modestly-sized rackable system, which translates into significant savings on power consumption and real estate consumed.

With this large number of cores, memory bandwidth is very important. SPARC T5 was designed to support extreme memory bandwidth: it has nearly 80 gigabytes of usable memory bandwidth. With nearly linear scaling across 8 sockets, SPARC T5 aggregate memory exceeds half a terabyte per second of memory bandwidth. This was measured with a benchmark called Stream. And that's great news for our customers, because this accelerates the performance of virtualized applications that require enormous amounts of memory.

Q: Does the SPARC T5 also support both single-threaded and multi-threaded applications?

A: Absolutely. What we pioneered on SPARC T4 was this notion of a core that had heterogeneous behavior that could supply really good single-thread performance. Thread assignment can be both a manual operation or more automatic when Solaris is left to do that job, so

SPARC T5 Processor

[View larger image](#)

Rick Hetherington on the SPARC T5 Processor

[Watch the video](#)

as customers or the Solaris scheduler assign more of the virtual threads on a core, the S3 cores will dynamically and efficiently transform into a throughput core. The S3 core on SPARC T5 will have exactly the same properties only better with its increased clock rate.

Q: Was there anything in the design process that surprised you, or did things go the way you expected?

A: No, I think things went pretty much as planned. SPARC T5 is our first 28-nanometer product. So we were prepared for some surprises in the technology realm. I focus on architecture and not so much on silicon technology, but from my perspective, development went very smoothly. Some of the potential surprises had to do with the fact that die sizes are getting larger, T5 being the largest die we have ever designed. This would have an impact on yield in the early stages but again we had plenty of good silicon for initial post-silicon check out.

And one of the reasons Oracle's microelectronics has been successful here is that we've learned to refactor our chips, which lets us improve the processor design without taking undue risk. What we try to eliminate is what I would call 'discovery' in post-silicon. You want to be very predictable about what it is you're building because we ask the company to invest a great deal of money in these programs.

My team is responsible for performance modeling and the team does this with great skill and accuracy. These very detailed models not only forecast performance, they assist and guide architects and designers throughout the development process. The resulting silicon matches very well with our forecasts with no surprises.

Q: And how do you know which workloads you want to model on?

A: We have a variety of workloads based on industry standard benchmarks, which we've been able to gather instruction traces on. This is a very complex task and very few processor teams can match our capability in this regard. The nice thing about using benchmarks is they provide comparability across T4 to T5, M5 and M6 as well as our competitors. The database team recently supplied us with a set of micro-benchmarks that represent critical sections of future database releases.

One of the main traces that we have is for Oracle Database 11g. It uses a 256 wide trace to simulate our performance models running TPC-C, which is an OLTP workload. It's generally considered the most well understood OLTP workload available. We use that to forecast performance and also to understand and analyze the behavior of these designs. It takes a tremendous amount of work to develop and acquire the trace, to verify the trace, and then apply that trace to performance models.

Q: What do you mean when you say trace?

A: These are instructions that are gathered from executing an application; a long list of billions and billions of SPARC instructions. We run Oracle applications on an existing SPARC system and gather instructions as they are executed. And then once we have the trace itself, we verify that it is legitimate.

Once we have the verified trace, another team of engineers will apply that trace to a model of the processor in development. This modeling is how we eliminate surprises when we first examine actual silicon.

Q: What do you consider innovative in the SPARC T5?

A: There are a number of areas that we would consider innovative on T5. I would point out the move from a snoopy-based coherence protocol to a directory-based protocol. This has reduced unnecessary coherence chatter and reduced memory latency. The result is our ability to scale in a near linear fashion from one to eight sockets.

T5 has a number of unique power management features, as well. This would include Dynamic Voltage and Frequency Scaling, which allows customers to set a power limit, and then a combination of firmware and hardware will adjust to execute within those set policies.

Q: What kinds of applications will benefit the most from the SPARC T5?

A: I would say anything within the Oracle software stack. Database transaction processing as well as analytics will do quite well. Java middleware performs extremely well. T5 is a well-balanced, general-purpose processor that scales from the simplest task to the most complex and demanding workloads. So it is safe to say that the entire Oracle application stack will run well on T5.

Q: What about the security features in the T5?

A: The team takes security very seriously. Each core on T5 has an encryption engine that will accelerate all of the most common bulk encryption ciphers like AES and DES. SPARC T5 also supports asymmetric key exchange with RSA and ECC and authentication or hash functions like SHA and MD5. We also have hardware random number generation.

With negligible overhead, customers can build an entire 3-tiered data center and never have to communicate between servers in clear text. It will all be encrypted as we go from the edge of the data center to, let's say, the backend of the database. What we are trying to do here is provide security for customers that have avoided using full encryption within their data centers because of performance issues. With SPARC T4 and now T5, there is no real reason not to run a data center that is secure from end to end.

Q: I'd like to know if there was any kind of optimization with Solaris 11.

A: Solaris is specifically designed to take advantage of the features in SPARC T5 like extremely high core and thread counts and on-chip security. SPARC T5 will run Solaris 10 Update 11 as well as Solaris 11. We encourage our customers to stay current with the latest Solaris release.

Q: And the Solaris binary compatibility still applies?

A: Yes, I think that's one of the important things to point out. Customers can be assured that whatever they are running on any previous SPARC system, which includes T1 through T4, will run just fine on T5, only better.

Q. What's the most important thing you want customers to know about the SPARC T5 processor?

A: We want customers to be assured that Oracle continues to invest heavily in SPARC processors and systems. They should also know that the team is delivering exciting new products in a predictable way as we outlined in our public roadmaps. We value our SPARC customers and appreciate the trust they have in our products and we will return that trust with a continuous stream of competitive and innovative products.



ORACLE CLOUD

- Learn About Oracle Cloud Computing
- Get a Free Trial
- Learn About DaaS
- Learn About SaaS
- Learn About PaaS
- Learn About IaaS
- Learn About Private Cloud
- Learn About Managed Cloud

JAVA

- Learn About Java
- Download Java for Consumers
- Download Java for Developers
- Java Resources for Developers
- Java Cloud Service
- Java Magazine*

CUSTOMERS AND EVENTS

- Explore and Read Customer Stories
- All Oracle Events
- Oracle OpenWorld
- JavaOne

E-MAIL SUBSCRIPTIONS

- Subscribe to Oracle Communications
- Subscription Center

COMMUNITIES

- Blogs
- Discussion Forums
- Wikis
- Oracle ACEs
- User Groups
- Social Media Channels

SER

- Log I
- Train
- Beco
- Find
- Purcl

- (
- Sale
- Glob
- Orac
- Partr