Opinion

The Sunsetting of SPARC

Executive Summary

Clabby Analytics has reviewed two versions of the Oracle SPARC roadmap — and we still find these statement-of-direction roadmaps to be vague (at best). In our opinion, these roadmaps raise more questions than they answer. For instance: "What is the role that Fujitsu will play in future SPARC designs?" (Fujitsu designs and manufactures the SPARC64 chip and co-designs and manufactures SPARC64 servers — but the company is now strategically committed to x86 multi-core architecture). Or: "Is Oracle planning to kill the current M-series chassis when it moves to its M4 processor?" Or: "Are Oracle's performance projections based on comparisons to its own underperforming servers — or are they based on projections of where Intel and IBM will be in five years?"

What our analysis indicates is that Oracle's SPARC64 M-series and its UltraSPARC-based T-series product lines will eventually merge. This will help eliminate confusion related to sorting out the differences between the multi-threading T-series and the single-threaded M-series. (Doing the workload analysis between the two was a real time-sink for Sun field engineers before the Oracle acquisition — and now that many of them are gone, it must be nearly impossible for the remaining ranks to sort this out).

But even if the M-series and T-series merge, will they still be viable? We believe that the performance claims on these roadmaps are being measured against current Oracle SPARC servers (which are far behind their current competitors in terms of performance per core).

And given our SPARC server competitive doubts — combined with the decline in demand for Oracle RISC servers — we also question whether Oracle should continue to invest in the SPARC architecture.

In this Opinion, Clabby Analytics describes why we believe that Oracle will discontinue its SPARC architecture (the company's SPARC64 M-series, and the company's UltraSPARC chip multi-threading T-series) within three years.

We start by taking a look at why various computer architectures have failed over the past several decades. What history shows us is that computer architectures fail when competitors gain a significant technological edge; when buyers move away from proprietary solutions toward standards-based solutions; and/or when it becomes economically unfeasible to continue building a given architecture. *Oracle's SPARC line suffers from all of these maladies*.

We then take a closer look at Oracle's SPARC roadmap.

Finally, we conclude with a view of what we think the server market will look like 3 to 5 years down the road.

Failed Computer Architectures and Historical Precedence

Over the past 40 years almost a dozen major computer architectures have moved into end-of-life (EOL). Some of these architectures were EOL'd due to major technological breakthroughs, such as the invention of CMOS technology (complementary metal oxide semiconductor), which ended bipolar designs; others were EOL'd when microprocessor technology replaced transistors (for instance, the Wang 2200 used hundreds of TTL [transistor-transistor logic] chips to perform calculations — but gave way to microprocessor-based designs). Then came the demise of popular systems from Apollo, Data General, and Digital Equipment Corporation. All of these systems fell out of favor as users moved away from proprietary solutions to standards-based platforms. And more recently, Hewlett-Packard (HP) killed off its own PA RISC line of servers — replacing PA RISC-based servers with the now-struggling Intel Itanium-based servers in order to reduce costs related to microprocessor development and manufacture.

There are several lessons to be learned from the collapse of all of these architectures. These lessons fall into three general categories:

1. *The effect of technological advances*;
2. *The effect of standardization*; and,
3. *The cost of design and development*.

The remainder of this section examines each of these categories more closely.

Technological Advances

To be competitive, microprocessor vendors need to constantly push the envelope in terms of scalability, performance, and reliability. In years gone by, improvements in scalability and performance were achieved by following "Moore's Law" (halving the amount of real estate that a chip would use while roughly doubling the performance). But, at present, individual microprocessors have run into physical limitations (heat dissipation, stability) — maxing out the clock rate at around 5 GHz. So, in order to continue to increase performance, microprocessor makers now use multiple core (multi-core) scaling techniques. Sun was in the forefront of multi-core designs when they introduced the innovative UltraSPARC T1 8-core processor in December of 2005, but it came at a cost: they sacrificed single-threaded (latency-based) performance to gain leading-edge multi-threaded (throughput-based) performance with low power consumption, thus leading to the bifurcation of the SPARC product lines mentioned earlier.

Systems design is also a critical element when it comes to server competitiveness. The technological competitiveness of a given system can be gauged using several metrics such as performance per core, memory proximity, the design of its input/output (I/O) subsystem, the availability of redundant components (for reliability and availability), the availability of specialized virtualization processing facilities, and more.

Of these considerations, *probably the most important measurement of a computer's performance from a business perspective is the amount of workload processing power it can deliver on a per-core basis.* The more work a microprocessor can do per core, the fewer number of cores will need to be purchased — and, fewer software licenses will need to be purchased. Very often, this can also lead to lower software maintenance costs since

support is often based on the number of software licenses required. To overcome this situation, Oracle might have to make some software cost price "adjustments".

> ***The main reason that the amount of processing per core is so important is this: software licensing is usually priced on a per core basis. So if your system performs poorly on a per core basis, then you need to purchase more cores and more licenses to execute your workload. Oracle's SPARC64 design is notorious for lagging the market in core performance.***

Standardization

Information technology (IT) buyers strive to standardize their computing environments in order to simplify the management of those environments, to overcome vendor lock-in, and to drive down costs. In the late 1980s, through the 90s, and into the early 2000's, Sun Microsystems played this "standardization card" expertly, taking business away from proprietary systems vendors with standards-based UNIX-based workstations and servers.

In today's computing world, however, IT buyers are now looking to standardize their hardware environments (in our opinion, on three architectures: x86 multi-cores, Power Systems, and z). Further, IT buyers are looking to standardize on fewer operating environments — and, accordingly, there is a heavy lean toward the Linux and Windows operating environments. Some IT buyers, however, still require the advanced functionality, scalability and reliability delivered on Unix platforms. These buyers have been gravitating toward IBM's AIX (Unix) operating environment (which has consistently gained market share over the past five years). Hewlett-Packard's HP UX (Unix) operating environment has been down to flat in terms of revenue, but some buyers are still buying HP UX-based Itanium servers (we wonder what the impact of Oracle withdrawing support for HP Itanium-based servers will be…). But there has been clear migration away from Oracle's Sun Solaris (Unix) environment.

What is important to note about Oracle's current server market position is this:

1. Mass migration away from SPARC servers started when Oracle began licensing its software using a per-core multiplier for multi-core systems in the mid-2000s, putting SPARC at a significant disadvantage relative to POWER due to SPARC's lagging per-core performance, and accelerated when Oracle announced its takeover bid of Sun;
2. The Unix market has been steadily declining for the past several years (but has recently started to stabilize in the $20 billion per year range). Sun's share of the Unix market had been steadily declining for three years — and there is little reason to believe that Oracle's acquisition of Sun will reverse that trend (the company would need some very compelling hardware improvements and would need to reestablish itself with customers who have moved off of SPARC technology — and both would be difficult to achieve);
3. Linux represents a logical growth market for SPARC. But the last supported distribution of Linux on SPARC, Ubuntu, was dropped a few years ago — and Oracle Enterprise Linux doesn't run on SPARC. Further, the market has shown little interest in running Linux on RISC architecture (90%+ of Linux runs on x86 architecture). The reason for this appears to be that Linux on SPARC doesn't give

IT buyers any advantages (such as virtualization, performance, or power efficiency advantages) that they can't already get on x86. And, if these buyers were to purchase a Power System, they can take advantage of extensions such as better virtualization, increased performance, and reduced power consumption.

In short, Oracle is not well positioned to ride the next standardization wave with the SPARC architecture.

Sales Volume Versus the Cost of Design and Development

Microprocessor engineering expertise is expensive. But even more expensive are costs related to building of microprocessor fabrication plants (fabrication plants can cost several billion dollars each). And these costs can mount up over time (for example, by some estimates Hewlett-Packard and Intel have spent between $10 and $20 billion designing and manufacturing the Itanium chip set).

To offset these costs, chip makers need to:

1. Sell their microprocessors in volume; and/or
2. Partner with fabrication plant owners to offset plant costs; and/or
3. Charge high prices for their specialty chips (or recoup the cost in systems/support revenue — which Oracle appears to be doing from a support perspective).

From a volume perspective, Intel sells millions of its microprocessors annually — and is able to satisfy demand while keeping microprocessor costs low due to its very high volume of sales. Likewise, IBM sells multiple hundreds of thousands of POWER microprocessors annually, with POWER chips being found not only in computers, but in cars, games, and other devices. We suspect that the volume of SPARC64 and UltraSPARC chips found in other devices pales in comparison to the number of processors sold by Intel and IBM.

As for partnering with fabrication plant owners to offset costs, Sun and Texas Instruments called it quits after about 15 years when TI announced that they were not doing anything past a 65nm node. So Sun went to Taiwan Semiconductor Manufacturing Company (TSMC) for UltraSPARC and Fujitsu announced that they will do the same with their own chips in the future. This guarantees that SPARC going forward will be confined to whatever lowest-common-denominator technology TSMC will offer to all of its broad consumer-oriented client base — and should put Intel and IBM, who invest in their own manufacturing plants, at an even greater technological advantage for enterprise microprocessors in the future.

With respect to recouping costs by charging higher prices, Oracle has increased SPARC server costs (especially costs related to memory). And support costs are up. And, due to inefficient processing performance per core, software revenue is particularly strong on SPARC for Oracle. Needless to say, customers are not delighted with this scenario…

So, if Oracle: 1) doesn't sell SPARC in volume; 2) has changing or immature relationships with its fabrication partners; and, 3) has had to raise systems/support prices to help fund SPARC development — then, hopefully for Oracle, the increase in SPARC pricing is helping to fund SPARC development — right? Wrong-O. IDC has reported that revenue market share is way down. Gartner has reported that Oracle's server count was down over 40% year-to-year. And an Information Week article points out that Oracle's fiscal third-quarter 2011 earnings showed that new hardware sales came in at about $1 billion, missing

> ***earlier guidance from the Oracle by a whopping $100 million. It also indicates that sales of Oracle/Sun servers declined 14.4% in the fourth quarter of last calendar year, while IBM server sales rose 22% and HP's rose 13% during the same time period.***
>
> ***So, in the end analysis, it appears to us that Oracle is stuck in a position of having to build expensive, low volume chips for a market that is moving away from SPARC systems design and standardizing on other architectures. And this is important, because we believe that at some point in the near future Oracle will have to ask itself the question Sun should have asked itself in 2008 after the glow of its T1 and T2 servers had faded: "Why on earth are we still in the microprocessor design business when SPARC no longer has technological, efficiency, performance, virtualization, or power consumption advantages over its competition?"***

Oracle Roadmap Analysis

First, it is important to note that SPARC is really two binary-compatible product lines: the UltraSPARC-based T-series for multi-threaded, throughput-oriented workloads and the SPARC64-based M-series for single-threaded, latency-based ones. The chip multithreading (CMT) UltraSPARC is designed exclusively by Sun/Oracle, and the SPARC64 chip is designed and manufactured exclusively by Fujitsu.

A copy of Oracle's server/operating environment roadmap can be found, among other places, on page 42 of a Fujitsu presentation given at Oracle OpenWorld in 2010 and is located at: http://www.fujitsu.com/downloads/SPARCE/featurestory/oow2010-keynote.pdf. Figure 1 shows this roadmap.

Figure 1: Oracles's Server/Operating Environment 5 Year Roadmap

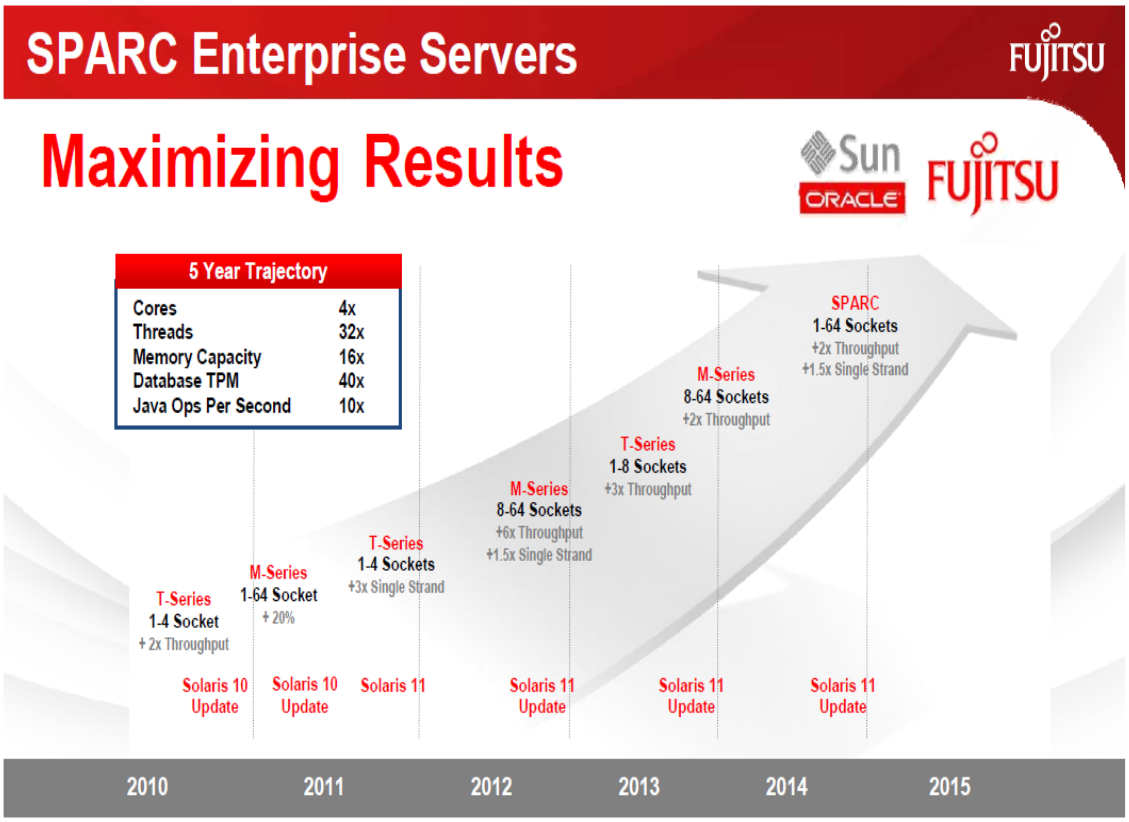

Source: Fujitsu, OpenWorld 2010 Keynote

As we stated at the outset, we believe that this roadmap raises more questions than it answers. From a big-picture perspective, it appears to us that Oracle's ultimate goal is to eliminate its dependence on the SPARC64 architecture and continue M-series and T-series product lines with its UltraSPARC chip series — culminating in a machine labeled SPARC 1-64 Sockets in the 2014 timeframe. Note that this machine promises 2x the throughput and 1.5x single-strand performance — but unfortunately, we don't know which system this 1-64 socket server is being compared with. We suspect that it is being compared to existing Oracle SPARC/UltraSPARC architectures, and this causes us some concern because Oracle SPARC/UltraSPARC architectures already lag the competition in terms of per-core performance. So offering to increase performance versus underperforming architectures does not resonate well with us. Further, it is important to remember that Intel and IBM will not be standing still while Oracle builds its 1-64 core machine. IBM POWER7 already outperforms Oracle in numerous benchmarks — and in some cases by a factor of 3x.

Further, by comparing the OpenWorld roadmap (Figure 1) with an older version (found in a September, 2009 article in "The Register" article, it looks to us like the "M-series" nodes shown on this roadmap in the 2012 and 2014 timeframe are not part of the Fujitsu SPARC hardware engineering collaboration (Fujitsu will continue to manufacture them until their manufacturing moves to TSMC). And this causes us to surmise that the tight engineering collaborative agreement between Fujitsu and Sun/Oracle regarding SPARC may be coming to an end as soon as Fujitsu delivers the M3 processor in early 2011. Potential buyers of M-series nodes should, therefore, be aware that the new M-series nodes may come from different designers and that different manufacturing facilities may be used to produce the resulting chip. *This is particularly noteworthy since Oracle/Sun has not designed a top-to-bottom, volume-to-high-end SPARC product line entirely on its own since the UltraSPARC-III/IV/IV+ series, designed over a decade ago before Sun's troubles began and before many of its SPARC engineers left for greener pastures in the industry.*

> ***The bottom line of our Oracle SPARC roadmap analysis is this: The roadmap is vague. It appears that the Oracle/Fujitsu joint development relationship around SPARC is coming to an end. And it appears that Oracle may be comparing its projected performance against its own, low performance per core machines. This roadmap should not instill a great deal of confidence in the Sun buying community.***

The Computer Marketplace — Our 3-5 Year Outlook

In November, 2009, *Clabby Analytics* published a report entitled "*And Then There Were Three: POWER, x86 and z*" that described why we believed that a major shakeup was underway in the midrange/high-end server marketplace. In that report (which can be found at: http://www.clabbyanalytics.com/uploads/ServerMarketViewMarch2010UPDATE.pdf) we noted that Sun server sales were slumping badly; and, that Hewlett-Packard's Itanium-based server sales were also floundering. Further, we noted that Intel Xeon x86 multi-cores had become capable of challenging both of these server platforms from a performance perspective. We concluded that the days were numbered for SPARC- and Itanium-based servers — leaving x86-, POWER-, and z-based servers remaining to contend for server market share.

Last week's announcement by Oracle that it has decided to discontinue software development on Itanium-based servers should really hurt Itanium. Red Hat, VMware, and Microsoft have also curtailed their Itanium development efforts — but the withdrawal of

Oracle's applications, middleware, and database really hurts because what good is a high-end server environment if it can no longer stay current with business applications, middleware, and database offerings? Good thing Itanium still runs WebSphere… This should be particularly disconcerting to SPARC customers when they realize that Oracle's stated reason for dropping support for Itanium is the low acceptance of Itanium in the marketplace: according to IDC, Itanium-based servers have been out-selling SPARC-based servers since the end of 2009!

To survive, Itanium needs a compelling processor roadmap (the current Itanium roadmap overlaps in several places with Intel's own Xeon roadmap); the ability to control its middleware destiny (Hewlett-Packard, the leading seller of Itanium servers, exited the middleware business several years ago); and it needs database support from market leaders such as Oracle and Microsoft. That support appears to be waning.

For SPARC to survive, Oracle will have to give up selling SPARC as a general-purpose processor. It no longer competes favorably with Intel and IBM in terms of per-core performance. But, this architecture could be redesigned and targeted specifically for parallel data processing jobs (it could become a specialized database appliance architecture). The beginning of this positioning can be seen with the SPARC Supercluster — though this design does not appear to be very well fleshed out at present). Given Oracle's deep database expertise and customer loyalty to Solaris, this appears to us to be the only reason to keep SPARC alive. But justifying the huge cost of hardware development and manufacture, in our opinion, will be difficult for Oracle. The company would be better off building its strategy on the Intel multi-core architecture.

Although the server market will be dominated by x86 multi-cores, Power Systems, and mainframes over the next 3 to 5 years, a new architecture has the potential to make inroads within IT organizations. New server architectures are being designed around ARM microprocessor technology (the same type of microprocessor found in cell phones worldwide). These new designs will be comparatively inexpensive (versus traditional designs), will consume far less power, and will dissipate far less heat. For more information on ARM technology, please see our report at:

(http://www.clabbyanalytics.com/uploads/CalxedaFinal.pdf.

Summary Observations

We started this *Opinion* with a look back at computer architecture failures in the past. And these failures were caused by changes in technology, the effects of standardization, and a breakdown in cost structure (when the cost to build microprocessors and systems environments outweighed the revenue gained by building them).

From our perspective, SPARC is already behind the eight-ball from a technology perspective. It's per-core performance is just plain awful and its virtualization capabilities are far behind what one can get from IBM's PowerVM or from the various software solutions available for x86. And we believe it will take a lot of investment to make SPARC competitive with general-purpose Intel multi-cores, Power Systems, and mainframes.

SPARC is also in trouble from the standardization perspective. First, the market is showing clear preference for three types of hardware: x86-based multi-cores, POWER-based systems, and mainframes. Second, Unix does not appear to be a growth market for Oracle (see sales figures earlier in this report). And third, there has been little demand for Linux on SPARC RISC architecture.

As noted within this report, Oracle is not performing well in hardware sales. IDC and Gartner are both reporting significant declines in revenue and market share. We note that there appear to be pockets of Oracle/Sun loyalists — especially within governments, telcos, and the military — but we believe these pockets of loyalty will dissipate over time if Oracle doesn't make its servers more competitive.

One tactic that Oracle can use with its customers is to sell "packaged" solutions — integrated applications, middleware, operating environments, a database, and hardware (this approach is also known as the "appliance" model). The integration of their solution could be stressed, and customers could be told to focus on the packaged solution (but not the underlying hardware). We suspect, however, that customers might notice that it might take twice or three times as many Oracle servers to do the same amount of work as POWER-based servers — so we don't hold much hope that the appliance model approach will work for Oracle in the long term.

Finally, Oracle's server roadmap raises more questions than it answers. It appears that Oracle will discontinue its SPARC64 collaboration when the SPARC64 M3 becomes available in 2011, and attempt to resolve its bifurcated multi-threaded vs. single-threaded SPARC product lines with its UltraSPARC development in the 2014 – 2015 timeframe. If this is true, then the intervening systems introduced in 2012 will have a very short chassis lifetime of approximately three years, forcing another forklift upgrade in 2014. It appears that Oracle's developmental relationship with Fujitsu is coming to an end — and we are concerned that the continuity in Oracle's SPARC64 line will be disrupted. And it appears as if a chassis change is about to happen shortly — a change that could force users to perform forklift upgrades to accommodate future Oracle servers.

When all is said and done, Oracle will need to look seriously at costs for research and development related to keeping the SPARC architecture alive, versus the revenue gained. The way we see it, Oracle is making its money selling some systems, but also a lot of software licenses and support. Support revenues will decline as systems sales continue to decline. So, at present, investment costs appear to outweigh revenue gained.

Accordingly, we think it's just a matter of time until Oracle recognizes the futility of keeping tne SPARC architecture alive. The simple reason is that they just don't need it — and it's costing them a fortune to continue...

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662