CLABBY ANALYTICS

Opinion

And Then There Were Three: POWER, x86, and z UPDATE

Executive Summary
This Opinion is an update of the original, published in November, 2009 — and subsequently updated in March, 2010. The reasons that we are updating this report for a third time are that:

1. Market conditions have changed (our predictions about Itanium were true; our predictions about Oracle were partially on target);
2. There is new sales data that shows the Itanium downturn that we predicted; and,
3. There has been an uptick in migration away from both SPARC- and Itanium-based servers.

Since we last updated this report, Microsoft pulled the plug on Itanium (see our report at: http://www.clabbyanalytics.com/uploads/Microsoft_Withdraws_From_Itanium.pdf); Red Hat discontinued Itanium development; and Oracle will no longer build its business applications and database for Itanium (we correctly predicted that this would happen more than a year in advance of Oracle's withdrawal).

As for Itanium sales, over the past seven quarters we have seen four quarters of revenue declines; a blip upwards when a new Itanium processor finally arrived; then flat sales — and now *a mere 1% improvement* (see Figure 1).

Figure 1 — The Last Seven Quarters

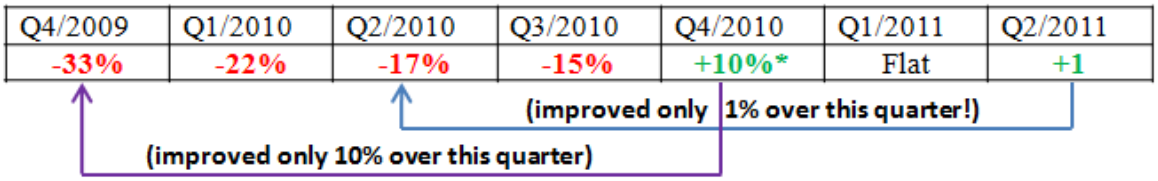

Source: Clabby Analytic — June, 2010

In short: we do not believe that Itanium will ever recover — especially due to direct competition from Intel's own Xeon multi-core servers.

As for Oracle, we correctly predicted that Oracle would cease development on Itanium; but Oracle surprised us by meeting its accretive goals (its financial targets) after acquiring Sun. Our take, however, is that Oracle did this mostly through severe cuts in sales, general, and administrative costs — mostly by laying-off a lot of Sun personnel and/or through attrition. We think that this massive cost-cutting has had a negative effect on the installed Sun base — and has helped triggered migration away from the platform.

In this *Update*, *Clabby Analytics* reinforces the position that we took a year and a half ago when we told our readers that we believe that the server market would consolidate around three architectures: x86 multi-cores, POWER, and z.

Our Position

We believe that there is a major shakeout underway in the midrange/high-end server market place as sales of Sun SPARC/CMT (cellular multi-threading) and Hewlett-Packard (HP) Itanium-based servers decline significantly — and as a new, more powerful versions of Intel's Xeon and IBM's POWER micro-architectures, and IBM's System z have come to market.

> ***We now believe that the midrange/high-end market segment of the computer market place will gravitate toward three microprocessor/server architectures (Xeon, POWER, and z). If Oracle can't right the Sun UltraSPARC/ chip multi-threading (CMT) ship (and we think it won't) — and if Intel's Itanium architecture continues to struggle (and we think it will) — then Sun and HP midrange/high-end customers will have little choice but to move to IBM POWER, IBM z, or Intel Xeon-based servers.***
>
> ***Information technology (IT) executives who don't recognize that this shift is underway — and who continue to invest in Itanium- and UltraSPARC/CMT-based servers — may find that they have wasted their precious information systems hardware budgets on dead-ended server architectures and operating environments.***

A Requiem for Sun

Starting out as a maker of Unix workstations, Sun's greatest success in the computer market came by being in the right place at the right time. In the late 1990s, Sun was perfectly positioned with small Unix-based just as the need for small Web servers took off. (At the time, both IBM and HP had positioned their Unix servers for the midrange/highend Unix market). And, as the need for more powerful Unix servers grew in order to address enterprise resource planning (ERP), customer relationship management (CRM), and supply chain management (SCM) requirements, Sun was able to build on its success by acquiring Cray Business Systems' Starfire server — an extremely well designed SPARCbased server that proved ideal for addressing Sun's mission critical computing needs.

> ***But since 2001, when the Internet bubble burst, Sun has been unable to return to its former glory. Sun's inability to get products out-the-door in a timely fashion (as illustrated by the tardiness of its Niagra processors) — and the delays in the release of its CMT architectures — started to erode customer confidence in Sun server hardware schedules and deliverables. Discontent with Sun's inability to deliver up-revs of its micro-architectures and systems led numerous Sun customers to migrate to other platforms.***

The Current Situation — Oracle

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.

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.

For a complete, in-depth analysis of the Oracle/Sun situation, please see our report at: http://www.clabbyanalytics.com/uploads/SPARCfinal.pdf.

The Failed Itanium Experiment

Explicitly parallel instruction computing (EPIC), the computing basis for the Itanium processor, was the brainchild of HP back in 1989. At that time, HP engineers believed that, by preprocessing instructions (using the EPIC approach), they could build a computer architecture that was superior to the reduced instruction set computing (RISC) products offered by IBM, Sun, and even HP itself — and by so doing, they would be able to create a distinct competitive advantage for HP. In fact, HP was so confident in this new approach to computing that it even embarked on a schedule to kill its own RISC processor (PA-RISC) in favor of Itanium.

In 1994, HP partnered with Intel in an attempt to position Itanium as "the industry standard chip" for 64-bit computing. Unfortunately for Intel, computer buyers didn't take to Itanium, partially because Itanium releases were consistently late to market — often by months or even years — while constantly dropping features. As examples, consider:

- Itanium was conceived in 1989 — but didn't make it to market until 2001 — after substantial delays due to structural problems, processor count challenges, and compiler issues (amongst several other issues). Intel and HP had targeted 1998 for the first release of Itanium — so its 2001 release date was almost three years late;
- This lateness trend continues even today. The next build of Itanium, codenamed Tukwila, was originally due in the middle of 2009, then moved to the end of 2009, and finally made it out-the-door in March, 2010).
- From a "dropped features" perspective, one of the most important features dropped in Itanium design was the chip's ability to handle 32-bit computing. During the course of its development, 32-bit emulation mode was removed from the Itanium feature set, making it impossible for IT buyers to run their existing 32-bit applications on the Itanium 64-bit processor. And to make matters worse, Intel was slow to make its own 32-bit Xeon-class processors capable of running 64-bit applications. So, as an IT buyer, if you wanted 64-bit computing, you had to buy Itanium (in effect, Intel precluded a clean migration path to 64-bit computing). AMD, seizing this opportunity, built a highly successful 32-bit/64-bit architecture, ultimately forcing Intel to build 32-/64-bit hybrid Xeon class machines.

> ***By enabling Xeon architecture to run 64-bit applications, Intel has now made it possible for IT buyers as well as IT vendors to deploy 64-bit solutions on an architecture that is considered an industry standard: Xeon. And Intel has, in effect, created a competitor for Itanium within its own product line. This internal competition works strongly against Itanium from an economics perspective as Xeons are high-volume/low-cost processors whereas Itanium is a high-cost/low-volume processor. Low sales volumes mean that Intel needs to charge heavily for Itanium microprocessors in order to recoup its investment; whereas high volume Xeon processors generate lots of money and can thus be sold at lower price points. Accordingly, if Xeon continues to encroach upon the Itanium camp — and if Itanium sales volumes stay static or decline, this question needs to be asked: "how will Intel pay for the continued development of Itanium?" Perhaps this is why the Itanium roadmap gets a little "vague" after two years…***

Clabby Analytics is not the only organization to have come to the conclusion that Xeon overlaps tremendously with Itanium. The Bright Side of the News organization came to the same conclusion after attending an Intel Nehalem EX press briefing (their article can be found at the following URL):

http://www.brightsideofnews.com/news/2009/5/27/intel-nehalem-ex-xeon-spells-the-voice-of-doom-for-itanium.aspx

Failing to meet delivery schedules not only impacts customer plans, but it also disrupts vendor system build cycles — a circumstance that may cause IT vendors to start looking elsewhere for other microprocessor solutions. And, accordingly, an interesting vendor loyalty situation is developing in the Itanium camp. If the small handful of vendors can obtain Xeon-class servers at a far lower price point than Itanium-based servers — and if IT buyers are showing a clear preference for Xeon — then why should these vendors continue to build Itanium-based servers? One example of this phenomenon is SGI (formerly Silicon Graphics) — a vendor that has, to date, invested heavily in Itanium architecture from a systems design perspective. With Itanium arriving late (again), and with low-cost Xeon processors readily available, SGI could choose to start using Xeon-class processors in its high-end servers. If this happens, and if other Itanium vendors chose to follow suit — HP would essentially be the only vendor continuing to push Itanium solutions.

The Future of Itanium at HP — What Might Happen

From a systems perspective, HP is a two platform company: Itanium and Xeon. HP sells Itanium-based servers to address mission critical computing requirements, while Xeon servers are positioned for general business use. But, Xeon's move into multi-core processing combined with its expanded reliability, availability, and security (RAS) features and capabilities have now enabled Xeon to encroach on HP's Itanium midrange. Sometimes internal competition can be good (it can encourage innovative thinking and new product designs) — but *Clabby Analytics* does not see this particular internal conflict as good for HP in the long run. Here's why:

- There is considerable overlap in the types of applications that can be run on HP/UX on Itanium versus Linux on Xeon. If the same applications are available on two different platforms — and if the Xeon platform costs significantly less than the Itanium platform — which platform are customers likely to choose? (Answer: Linux on Xeon); and,
- HP could counter argue that IT buyers are paying for better RAS if those buyers purchase Itanium-based solutions. But that would imply that HP's Xeon-based solutions offer less RAS than their Itanium solutions — hardly an enviable claim — especially when considering that Dell and IBM compete with HP Xeon offering on the basis of RAS…

> ***Xeon's advance into the midrange/high-end of the computer market has created a Catch 22 for Hewlett-Packard. To claim that Itanium is superior to its Xeon offerings demeans its Xeon offerings; to claim that Xeon is superior to Itanium will weaken its Itanium sales.***

Adding to HP's internal conflict woes are developing migration trends away from Unix toward Linux. Unix market revenues have been flat over the past several years due to competition from the Linux operating environment as it moves upscale in the enterprise. Over the past three years, HP has been losing market share in Unix to IBM. After looking closely at this trend, we think that this loss of Unix market share is a sign that HP/UX

customers are opting for other operating environments rather than adopting HP/UX on Itanium. HP customers whose PA-RISC machines are reaching their end-of-life appear to be gravitating toward IBM AIX (Unix) on POWER (a RISC platform alternative) — or to Linux (but not necessarily Linux on Itanium). And, now that the Xeon E7500 multi-core chip set has been released, we expect HP customers to move toward Linux on x/86 much more aggressively.

> ***Finally, it is important to note that over the past few years there have been over 5000 competitive displacements from Sun (Oracle) and Hewlett-Packard to IBM's AIX/POWER-based servers. And the number is accelerating as these architectures continue to fail.***

For a more in depth discussion on Itanium, please see our most recent Itanium critique at: http://www.clabbyanalytics.com/uploads/Itanium_Fire_CIO_Update.pdf

The Big 3: x86, POWER, and z

As we stated at the outset, we believe that server market is converging around three architectures: x86, POWER, and z. This section takes a closer look at each of these architectures.

The Xeon Class

Nehalem is the codename for the class of Xeon multi-core, multithreading processors that succeeded Xeon Intel Core microarchitecture. The first Nehalem processor was designed to provide advanced desktop support, and was released in November 2008 as Intel's Core i7 processor. Xeon 55xx-class servers followed Core i7, and servers that use these processors are available today from several systems makers.

Nehalem class processors are different from predecessor Intel Core processors in the areas of multi-core support, memory management, interconnect, and hyper-threading capabilities. With Nehalem, Intel has introduced an integrated memory controller that can support up to three memory channels of DDR3 SDRAM or four FB-DIMM channels. Intel has also introduced "QuickPath" on Nehalem, a new point-to-point processor interconnect. And threading capabilities have been improved in Nehalem-class processors.

From *Clabby Analytics* perspective, the first generation Nehalem processors are interesting because they showed that Intel is now serious about playing in the server market high-end with an architecture that can scale well and that can perform a substantial amount of thread processing. But Nehalem-class servers have become much more interesting in their second revision (the newly released E7500 multi-core server design), because not only does this processor offer more cores — it also has four QuickPath interfaces, and is able to support a lot more standard memory (up to 16 standard DDR 3 DIMMS).

IBM's POWER 7

Intel's new Nehalem-class Xeon servers (previous section) will, over time, challenge IBM's POWER architecture and even mainframes for share in the server market midrange and high-end. But, in order to move up the server ladder, Xeon will have a lot of catching up to do — especially in the areas of processing power, virtualization facilities, and systems software. In the mean time, IBM's soon-to-be-released POWER 7 architecture will

set the competitive bar for Unix/Linux servers in midrange and high-end commercial, scientific, and super computer markets.

The POWER 7 design ranges from 2-8 cores per processor — with each core providing 32 gigaflops An 8-core unit would thus provide 256 gigaflops of computing capability and run twice as fast as its predecessor, the highly successful POWER 6 architecture. IBM will build POWER 7 into advanced systems designs that range from water-cooled and air-cooled towers, through distributed supercomputer designs, all the way to blades.

From a competitive perspective, if the premise contained in this *Opinion* is correct (that SPARC and Itanium are "troubled" architectures), POWER 7 should find itself in a class-by-itself, virtually unchallenged in the Unix market — and difficult to challenge in the Linux market based on sheer computing capacity. In addition to its impressive serial, parallel, and data processing capabilities, IBM had also focused on improving interconnect speeds, graphics handling, and virtualization capabilities — further distancing POWER 7 from Itanium and SPARC.

POWER 7 Virtualization Differentiation
POWER 7's virtualization capabilities serve as an excellent barometer to illustrate how much work lies ahead for Xeon as it attempts to move upscale in the midrange/high-end server market place. Some of the advanced virtualization/provisioning features found in POWER 7 architecture and systems designs include:

- Fine-grained dynamic sharing of processors, memory and I/O;
- Resources may be dedicated;
- Shared, dedicated processors;
- Extreme scalability and robustness (especially when compared with x86 virtualization offerings);
- Integrated firmware hypervisor;
- Virtual I/O servers layer;
- Hardware enforced isolation;
- LPARs and WPARs;
- DLPAR and Processor folding;
- Capacity-on-demand functionality; and,
- Partition Mobility.

It should also be noted that IBM Systems Director, as well as IBM's advanced Tivoli management products, offer rich systems/applications management facilities — as well as additional provisioning and process management facilities. This combination of further distances IBM POWER Systems from its competitors.

System z Mainframes
Mainframes have been around for almost fifty years — and despite a recent declines in sales (due to the coincidence of an off-cycle [no new mainframe delivered] year and the economic decline), mainframes are not about to fade into the distance. Unlike Itanium,

mainframes are core to servers in almost every major bank throughout the world, as well as in leading retail and financial institutions. Although pundits have forecast the end of the mainframe for over two decades, the demise of the mainframe is just plain not going to happen in the foreseeable future.

The mainframe does, however, have some weak points in the areas of price and perception. Because the mainframe has been around for almost fifty years, some IT executives consider mainframes "old technology". And many IT executives find mainframes to be expensive. From an "old technology" perspective, IT executives should know that:

- IBM System z mainframes are considered "the gold standard" in the computing industry for *virtualization and provisioning* — the model server that all vendors want to beat, allowing customers to consolidate several hundred servers onto a single mainframe;
  Offer the highest level of security of any commercial server in the IT marketplace (EAL Level 5), plus advanced encryption/public key facilities.
- Networking — instead of having to rely on external switches, routers, and cabling — processors are networked using a switched backplane within the System z. This greatly simplifies System z deployment while reducing networking hardware acquisition and cabling costs and associated network latency issues.
- *MTBF and high-availability* — over the last five decades, IBM has gotten very proficient in the design of highly reliable, highly-available systems design (mainframes can provide 99.999 per cent or greater high availability).
- *Manageability* — fewer people are needed to manage mainframes as contrasted with equivalently configured distributed systems environments;
- *Near-linear scalability* — IBM's System z design along with Parallel Sysplex enables near linear processor scalability (as additional processors are added they perform at nearly 100% of their capacity). Other system designs (such as some SMP designs) can see processor scalability drop by as much as 50% as new processors are added to a system.
- *Floor space* — IBM's System z packs a lot of processing power into a relatively small footprint (as compared to the floor space dozens of networked SMP or PC servers might occupy if equivalently configured). And, finally,
- *Energy efficiency* — When comparing the System z to an equivalently powerful distributed computing environment, the System z uses far less power and cooling to deliver the same amount of computing power.

From a cost perspective, IT executives should know that IBM has recently undertaken several pricing actions to help reduce mainframe costs including the creation of lower cost business class mainframes; the introduction of coprocessors and specialty engines (lower cost processors that operate in mainframes); and Solution Edition pricing (see http://www-03.ibm.com/systems/z/solutions/editions/ for further details on this pricing action).

Clabby Analytics believes that IBM is undertaking these price reduction measures to close the gap between mainframes and midrange/high-end Unix/Linux servers to attract new workloads to the platform.

Summary Observations

Information technology executives and strategic planners need to pay very close attention to what is happening in the commercial server market — both from a microprocessor architecture and system design perspective. Over the next few years, *Clabby Analytics* believes that three microprocessors will come to dominate the midrange/high-end server market: the mainframe "z" microprocessor; the reduced instruction set POWER architecture; and Intel 8-core and beyond Xeon-class chip complex. Sun's RISC processors will fade into the woodwork when Oracle realizes that it is not a hardware company; and Intel's Itanium-class architecture will implode due to internal competition from Intel 8-way and beyond Xeon-based processors.

Failure to grasp these changes could result in the purchase of dead-ended servers and operating environments — and the loss of valuable IT investment funds.

Some readers may also wonder why *Clabby Analytics* believes that IBM POWER Systems make the cut — while Itanium and SPARC did not. From our perspective, it will take several years (if not a decade) to enrich post-Nehalem Intel Xeon processors to a point where Xeon can compete head to head with IBM's POWER architecture and mainframes. POWER and mainframe architectures are way ahead of Intel Xeon in virtualization, memory management, RAS, and other mission critical features. Still, even though Xeon servers are far behind mainframes and POWER Systems in these areas, they will meet the needs of a large segment of the server market (a segment some people call the "good-enough" computing segment) — and they will definitely create new dynamics in the midrange/high-end server market.

Some readers of this document may believe that Clabby Analytics has been unnecessarily harsh on HP and Sun. We understand that many of our readers have devoted their careers to deploying SPARC and Itanium solutions from these companies — and the arguments put forward in this opinion may be hard to take. For those who wish to respond directly to Joe Clabby, President of Clabby Analytics, regarding the content and arguments put forward in this article, please send an email jclabby1@AOL.com.

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