Sun and AMD: The Second Wave

Research Note

Gordon Haff
14 July 2006

With its acquisition of Sun co-founder Andy Bechtolsheim's Kealia in early 2004, Sun set its sights on creating usefully-differentiated AMD Opteron-based systems that, nonetheless, made heavy use of inexpensive, high-volume, "commodity" components.

Such differentiation was in evidence with the first generation of dual-socket "Galaxy" systems. They had nice management, diagnostic, and serviceability features. They came with more network connections than the norm.[1] And, among major vendors, only HP was pushing AMD's Opteron as aggressively into enterprise environments. In short, they were attractive systems, and that was reflected in their rapid sales growth. These servers and their inevitable successors will doubtless remain a major Sun x86 play for a long time to come; this class of systems has come to be the industry's *über*-sweet volume spot and that will be slow to change.

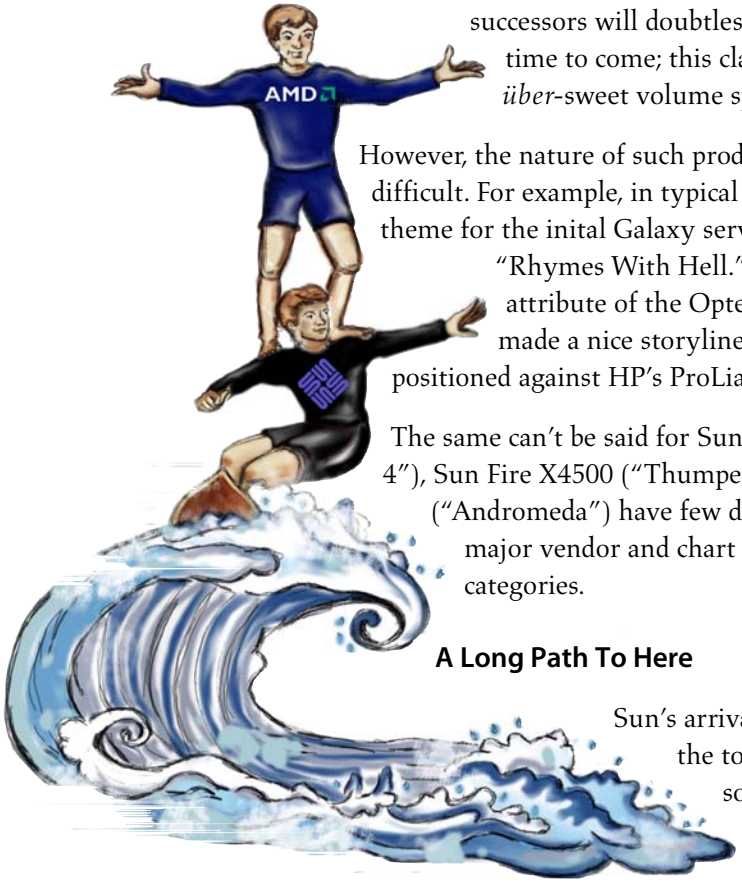

However, the nature of such products makes truly compelling differentiation difficult. For example, in typical amusing Sun over-the-top style, a major ad theme for the inital Galaxy servers was how cool they ran compared to Dell "Rhymes With Hell."[2] Yet power consumption was far more an attribute of the Opteron processor than of the actual servers. It made a nice storyline contra Round Rock, but by necessity it was positioned against HP's ProLiant line largely by ignoring it.

The same can't be said for Sun's latest trio. The Sun Fire X4600 ("Galaxy 4"), Sun Fire X4500 ("Thumper"), and Sun Blade 8000 Modular System ("Andromeda") have few direct analogs to product offered by any other major vendor and chart substantially new paths in their respective categories.

A Long Path To Here

Sun's arrival at this point is, in a way, remarkable given the tortuous, broken path it has taken to get to a solid, compelling x86 product line.

Sun's x86 efforts were long inconstant and confused—dating back to the late *Eighties* with

1 See our *Sun's New Galaxy*.

2 See our *Sun Finds a New Evil Empire in Round Rock*.

System name	*Sun Fire X4600*	*Sun Blade 8000 Modular System*	*Sun Fire X4500*
Codename	Galaxy 4	Andromeda	Thumper
Description	8-socket rackmount server	4-socket blade server	24TB hybrid data server
Max Processors	8 AMD Opteron 800 Series, single- or dual-core (up to 3GHz)	4 AMD Opteron 800 Series, dual-core (up to 2.6GHz)	2 AMD Opteron Model 285 dual-core
Max Memory	128 GB (w/4 GB DIMMs)	64 GB (w/4 GB DIMMs)	16 GB (w/2 GB DIMMS)
Max Storage	Up to 4 SAS (2.5″) 10K rpm	Up to 2 SAS/SATA (2.5″)	24 TB (w/48 SATA II 3.5″ 7200 rpm drives)
Built-in I/O	4 Gigabit Ethernet	6 PCIe interfaces to ExpressModules	4 Gigabit Ethernet
I/O Expansion	6 PCIe (4 8x, 2 4x) 2 PCI-X	2 PCIe ExpressModules per blade; Up to four PCIe Network ExpressModules per chassis	2 PCI-X
Height	4U	Up to 10 blades in 19U chassis	4U

a short-lived workstation codenamed "Road Runner" (the Sun 386i).[3] Five years later, Sun rolled out Solaris for x86, a feature-for-feature port of Solaris for SPARC that the company hoped would prove popular on a wide range of third-party desktops and servers.[4] It didn't. Sun was, however, ultimately successful at giving copies away. This checkered history seemingly came to an end in January, 2002 when Sun gave notice that it was "deferring the productization" of the next Solaris for x86 release.[5] Yet, that wasn't the end. Just a few months later, apparently under pressure from some key customers, and in the context of a new-found x86-Linux push, Solaris on x86 was resurrected.[6]

One would have to have been very generous to call any of these dabbles and digressions a "strategy"—and certainly not a well-formed one—on the part of Sun. Rather they arose from tactical and defensive considerations. The x86 hardware itself was strictly along for the ride; Sun's hardware engineering focus and investments were clearly on SPARC-land, both the processors and the systems that use them.

[3] Sun's one and only attempt to leverage the economies of the volume Intel marketplace using the first 32-bit Intel processor, the 80386.

[4] The rollout was by SunSoft—then a quasi-independent subsidiary that Sun created to develop and market Unix software in general and Solaris in particular (hence, the impetus for moving beyond a SPARC hardware foundation).

[5] See our *Solaris For Intel: Near the End of the Road*.

[6] See our *Solaris for Intel: It Lives!*.

Much has changed over the past couple of years. Sun has resurrected Solaris as a true strategic imperative, no longer presenting it primarily as a largely invisible foundation for middleware.[7] OpenSolaris has gained a far broader community and industry attention than the naysayers predicted. (Sun also continues to sell Linux broadly, if not exactly enthusiastically, on its hardware. Even Microsoft and Windows are tolerated when they're useful to sell hardware.) Add the first round of Galaxies[8] and it is increasingly evident that x86 servers are far more than another sideshow this go-round.

New Horses for New Courses

The second wave of products goes further. Whereas the first Galaxy systems created incrementally superior products on a familiar template, Sun is now taking different and unique approaches. This must be more comfortable territory for Sun, a company that has never viewed itself much as a market-follower. As CEO Jonathan Schwartz put it prior to last Tuesday's launch: "These are a bit of John and Lisa and Andy says"[9] rather than products crafted in response to market surveys.

The three new products may spring from the same technological roots. They're all Opteron-based and share a variety of componentry and management interfaces. But this is essentially a low-level technical view. From a larger fit and function perspective, these products are largely independent of one another; they're part of an integrated whole only in so far as all of Sun's products and offerings are. They're new takes on market segments that, at least today, are largely disjoint.

Biggish Iron

"Galaxy 4," the Sun Fire X4600, is the straightforward and familiar one of Sun's new offerings. It's an eight-socket x86—or as Sun likes to call it, x64—design in a 4U-high chassis. Configurable with up to 128 GB of memory, it also comes with four built-in Gigabit Ethernet ports and supports PCI Express (PCIe) for additional I/O.

x86 designs with eight sockets are uncommon today. The only remaining such from a tier-1 vendor is IBM's X3 lineup, a classic NUMA design that connects multiple quad-socket chassis together using scalability ports.[10] However, the real sweet spot of IBM's design approach is even larger scale points; at "just" eight-sockets, an integrated, single-chassis approach can be cheaper and higher-performance. Dell offers no servers with greater than four sockets. For its part, HP employs Intel's Itanium processor for its larger boxes.

In large part, the dearth of larger x86 system designs reflects a market that has moved overwhelmingly to dual-socket servers. The planners for Dell and HP's ProLiant lines care about *volume* and they just don't see a volume opportunity in the eight-socket space.Are they wrong? Not really. At least, not today. But there are signs that the multi-year shift toward smaller systems may be at least moderating. More than one vendor has told us that they're seeing quad-socket server volumes stabilizing and even strengthening and—even better from their perspective—they're going out the door with richer configurations and ASPs.

Server virtualization appears to be the disrupting factor here. Two-socket servers came into vogue largely because only one application ran on a system—and therefore you were typically going to underutilize the server anyway. There was little benefit in buying bigger only to have a more expensive system that was going to be even less utilized. But server virtualization, in whatever exact form it takes, lets multiple applications more easily co-exist on a single system. Given that, all

---

7 See our *Solaris 9: Fortifying the Foundation*.

8 Sun's September, 2005 Galaxy launch also included a high-performance computing (HPC)-oriented server codenamed "Aquarius" (Sun Fire X2100).

9 John Fowler, EVP of Systems Group; Lisa Sieker, his VP of marketing; and Chief Architect Andy Bechtolsheim.

10 See our *IBM's X3 Heads Into the Yukon*.

other things being equal, it's easier to manage fewer physical servers than more, husky servers combined with server virtualization can provide the best of both worlds—fewer physical boxes that are nonetheless effectively utilized. This marriage doubtless explains why Sun gave VMware so much airplay at its announcement (while still playing up the benefits of Solaris 10 containers as an alternative or complementary technology).

Finally, as part of Sun's renewed push into HPC, it sees benefit in being able to offer a "fat node" in the same vein that has garnered business for other vendors such as IBM and SGI for both HPC and more commercially-oriented clustered database roles.[11] Indeed, Sun's featured customer at the announcement was the Tokyo Institute of Technology, which had quickly constructed a supercomputing cluster that is now #7 on the TOP500 supercomputer list using pre-release X4600 servers.[12]

In short, the market for eight-socket x86 servers remains dwarfed by the units of dual-socket boxes that go out vendors' doors. However, there are clear trends and market opportunities that suggest a continuing and even growing opportunity for larger systems, too. Given how well Opteron Hypertransport links lent themselves to Sun's creating such a design "gluelessly" and the fit of such a system with Sun's historical big system design and selling savvy, it's an eminently sensible product direction for them to have taken.

Big Blades

Sun's new "Andromeda" blades, the Sun Blade 8000 Modular System, likewise eschews what has become familiar blade territory.

One notable difference is the level of integration. HP and IBM, the market leaders in blades, have taken their products down an explicit and aggressive integration path. According to their concepts, the blade chassis is a focal point that pulls in network switching and, increasingly, even storage. It's an approach that has found application even in non-datacenter environments such as replicated retail where a blade chassis becomes a sort of "datacenter-in-a-box." One downside to this approach is that it requires customized, chassis-specific switching components from vendors such as Cisco, Brocade, and Nortel—which, in turn,

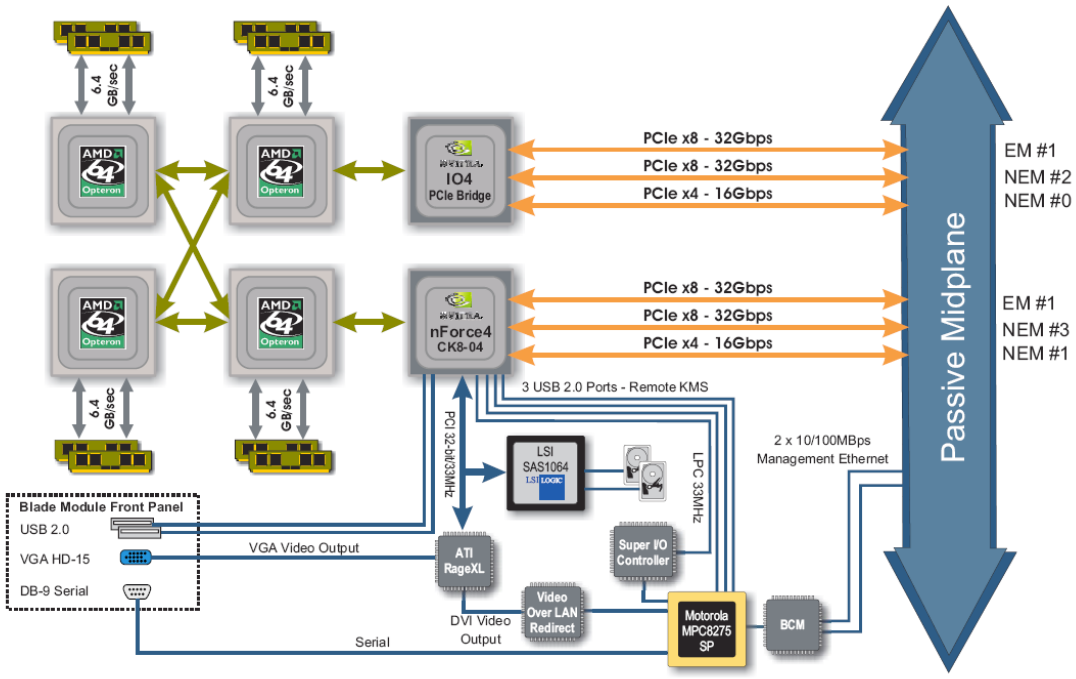

[11] See our *IBM's DB2: Database for Enterprises' Masses* and *High Performance Computing Meets Databases*.

[12] http://www.top500.org/site/690

makes changes in chassis designs that require new switching hardware a step not to be taken lightly.[13]

Sun has taken a fundamentally different philosophical approach. Sun's blades are far more pluggable "servers on their side" than a fundamentally different integrated environment in the vein of HP and IBM. Each of ten four-socket blades that plug into a chassis can incorporate up to 64 GB of memory (using 4 GB DDR1 DIMMs) and can be configured with up to two 2.5-inch SAS or SATA disk drives. Each blade can connect with up to two PCIe ExpressModules, an industry standard for hot-pluggable I/O modules, albeit a relatively new one; I/O then runs from these modules that plug into the back of the chassis to external switches—as with a standard server. Each 19U chassis can also accommodate four 20-port Gigabit Ethernet "pass-through" PCIe Network ExpressModules. Sun's approach to I/O is therefore fundamentally different from more integrated approaches that use custom mezzanine cards on the blades that then connect to integrated switches within the chassis. With Sun's approach, the chassis is on the periphery rather than in the center. In fact, Sun has indicated that new generations and versions—such as two-socket blades—will likely use different chassis optimized for the specific blade or at least family of blades.

Sun has also departed from the norm by initially concentrating on the four-socket space. As with the X4600, this gives Sun product that the competition either doesn't have or, certainly, is far less focused upon. If one views Sun's offering in the vein of a vertically mounted 1U, four-socket server that eliminates a great deal of the cable mess associated with standard rackmounts, it has essentially the same opportunities as standard four-socket servers —fat node HPC, financial services (where Egenera's four-socket blades have had some success), and various types of server consolidation. In the same vein, these are not particularly dense designs; in fact, Sun points to the size of these blades—about 50 percent larger than HP's C-class and 2.5-times IBM's BladeCenter—as a matter of virtue rather than a failing, supposedly because it means more room for memory and I/O. And, as Andy Bechtolsheim puts it, they're still as dense as "mere mortals" can power and cool given that a rack loaded with these blades would still draw in the 10kW range—about the maximum for normal air-cooling measures and, in fact, more than many existing datacenters can easily handle. In other words, there's little reason to make compromises to gain density beyond what can be readily air-cooled.

In some respects, Sun's new blades are an approach born of necessity. An HP or IBM integrated approach wouldn't have worked for Sun as it would have meant gaining support from a panoply of hardware partners to design switches and other custom components for Sun—when the top two vendors control more than 80 percent of the market. At the same time, Sun's alternative is a rational one that makes a more incrementalist advance on conventional rackmount servers—gaining some density and physical simplification while minimizing broader changes in management or datacenter operations.[14] It also means fewer tough design decisions around I/O and memory capacity—and therefore makes these blades potentially better for resource-intensive database or HPC workloads. At the same time, starting off with a four-socket blade will minimize the direct competition with the blade leaders and will put it more in competition with four-socket rackmounts where it will more easily be able to demonstrate a direct, competitive edge.

[13] Thereby forcing some tough decisions that can involve tradeoffs between compatibility and growth capability. HP opted for the new design path with its new c-Class blades (see our *HP's Integration Point*) while IBM decided to maintain backward compatibility (see our *A Resharpened BladeCenter*) with its latest announcements.

[14] Although, not surprisingly, Sun's depiction of competitors' blades as islands apart from their rackmount servers is an exaggeration. In the case of HP, for example, HP SIM can largely integrate the management of blades with conventional rackmount servers and, in the case of C-Class, Virtual I/O abstracts much of the integrated switching from network and storage administrators.

Lotta Disks

Probably the most intriguing of the new products is the Sun Fire X4500 ("Thumper"). Sun calls it a "hybrid data server." Whatever you call it, it's the ultimate direct-attached storage server—combining up to 24 TB of disk (using 48 SATA II disk drives) running off six controllers with two Operon processors in just 4U of rack space.

The X4500 ignores the traditional NAS playbooks. There's no hardware RAID; by default the X4500 will protect data using the software RAID built into its new 128-bit ZFS file system which also integrates volume management and is designed to guarantee data integrity across very large datasets. (ZFS verifies each data block against an independent checksum—stored in the pointer of the block's parent rather than in the block itself—after the data has arrived in the host's memory.)

In addition, while the disks are hot-pluggable, they aren't accessible from the front of the system as is typical with most disk arrays. The inside-the-box vertical arrangement that the X4500 favors instead is what gains it such unprecedented storage density. The fact is that the X4500 wasn't really designed as a general-purpose storage device in the usual sense; at Kealia, the development project that eventually became the X4500 was originally intended to be a type of media server.

The X4500 has direct and immediate applicability in a few areas—such as HPC workloads that manipulate huge datasets. However, for the most part, it's better thought of more as a *tabula rasa* awaiting its integration with software applications from Sun and from partners. One possibility cited at Sun's Network Computing announcement event where these products were launched was a database appliance running a massively parallel database (an MPP version of PostgreSQL) developed by Greenplum. One could also imagine solutions developed within Sun for archiving and other purposes—perhaps based on its "Honeycomb"[15] developments or other technologies. The greatest value and most interesting uses for the X4500 will be realized as Sun and others build such solutions atop the basic hardware platform.

Conclusion

It's not easy to be different in the x86 product space—and it's even harder to be different in sensible ways that customers will actually value. Given that reality, simultaneously announcing three products that *each* stakes out a new (and interesting) chunk of x86 server territory is an impressive achievement indeed. From a single-chassis eight-socket server to a four-socket blade that's unabashedly solely about heavy-duty datacenter infrastructure to the ultimate manifestation of direct-attached storage, these systems all take mostly standard volume components as their starting point. But only a starting point.

No doubt, much of Sun's volumes and revenues from x86 servers will continue to come from their two-socket boxes. Indeed, much of its volumes and revenues will come from those boxes running Linux or perhaps even Windows. But, however necessary a bread-and-butter component of its business, dual-socket rackmounts aren't really a space in which Sun can really flex its innovation muscles. As Jonathan Schwartz is fond of saying, "Computing is commoditizing; computers aren't." These new systems are anything but commodities.

[15] http://news.com.com/Sun+hopes+for+better+storage+with+Honeycomb/2100-1015_3-5553913.html. (Sun previewed both Honeycomb and Thumper back in May but made no mention of Honeycomb in this announcement.)