Sun's Midframes: Still Holding the Crown?

Research Note

Gordon Haff

20 November 2001

Last March, Sun Microsystems rolled out a bevy of Sun Fire "Midframe" servers built around the UltraSPARC III processor. What was most attention-grabbing about these systems was the degree to which Sun brought "rocket science" high-end features into the midrange. Capabilities like the ability to split the computer into separate domains, online upgrades, and support for mixed-speed CPU's were uncommon enough in the highest-end Unix systems—much less the midrange. Sun made itself the vendor to beat in this space.

On November 6, Sun announced 900 MHz UltraSPARC III processors for the Sun Fire Midframe server line. A new processor isn't usually an earth-shattering event. However, this announcement offers a logical point at which to take a new snapshot of the Midframe server line. How successful has Sun's delivery of these products been? How does new competition stack up in terms of features and performance? What pieces are still missing?

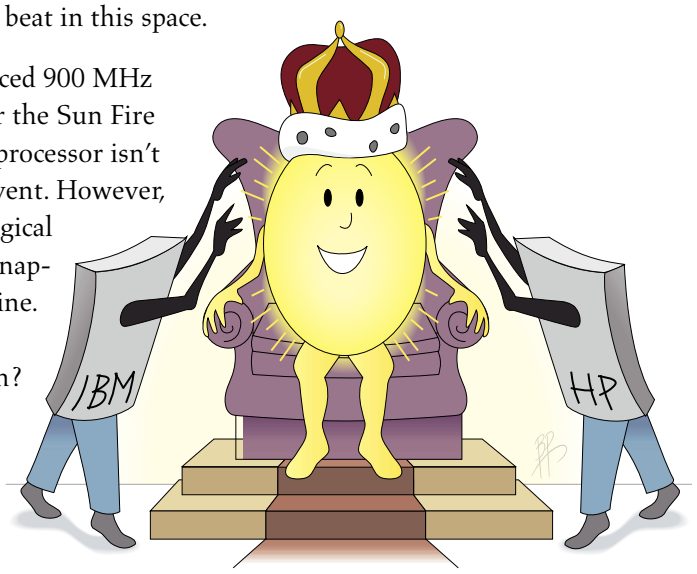

New Processors

With this announcement, Sun's Midframe systems get the same 900 MHz UltraSPARC III processors as the new high-end Sun Fire 15K (a.k.a. Starcat)—rolled out in September—and certain Sun Blade workstation models. Sun achieves speedier performance relative to its current 750 MHz processor by combining a copper interconnect, low-k dielectric, and a 100-nanometer gate transistor. Other processor features remain unchanged, including 8MB of ECC-protected cache, an integrated memory controller, and a 9.6 GB/sec system interprocessor bus. The chip is manufactured using a .15 micron process at Texas Instruments, a 13-year Sun partner in the development of advanced process technologies.

The syncing up of the Midframe processors with Starcat is of more than passing interest because of what Sun refers to as its "Uniboard" design—a concept that provides for common CPU/memory boards across a wide range of UltraSPARC III-based systems. In other words, a CPU/memory board from a low-midrange Sun

Fire 3800 can be removed and placed into a high-end Sun Fire 15K to accommodate changes in workloads or other provisioning needs. The Sun Fire design does not require that all processors in a system be the same speed; one CPU/memory board can have 750 MHz processors and another 900 MHz in a single system.

Sun is currently shipping 900 MHz processors on Midframe systems for revenue; it expects to have the products in general availability[1] by early December. As is typically the case, 750 MHz processors will continue to ship as a lower-priced companion offering to the 900 MHz CPUs, which were announced with higher pricing than the current products.

Benchmark Games

Processor technical features may be intellectually interesting, but they don't really indicate much about application performance at the system level. Generally, the larger the system, the more influence factors other than the processor have on total throughput. So, system-level measurements are needed if we are to make any meaningful performance comparisons among systems.

Of late, Sun has been on the warpath against certain widely used benchmarks—especially those from the Transaction Processing Council (TPC). Sun correctly points out these synthetic benchmarks often do not accurately reflect the complexity of real-world environments. These benchmarks are also notoriously expensive and arduous for vendors to run. Sun's solution has been to tout a mix of application-level benchmarks as more indicative of actual customer environments: a PeopleSoft 8 Financial Online Benchmark on the Sun Fire 6800 (in March 2001 on the original 750 MHz processors), an Oracle Applications Standard Benchmark (ditto), and a two-tier SAP-Standard Sales and Distribution Application Benchmark (on the Sun Fire 15K in September 2001). These measurements can give customers a good handle on the performance they are likely to see on the benchmarked application—or at least as close as they can get without the huge expense and effort of a custom benchmark. They also help system vendors like Sun to highlight key ISV partnerships. The main downside of such benchmarks is that they lack a sizable reference sample and so are difficult to use as a cross-vendor performance comparison.

Among simpler and more widely used benchmarks, Sun's current favorite is SPECjbb2000. This benchmark emulates a 3-tier system with an eye to representing the most common type of server-side Java application. The work of the middle tier— business logic and object manipulation—predominates. SPECjbb2000 aims to model a wholesale company, with warehouses that serve a number of districts. It includes transactions like placing new orders, requesting the status of existing orders, entering customer payments, and checking stock levels. The benchmark exercises the implementations of Java Virtual Machines (JVM), Just-In-Time (JIT) Java compilers, garbage collection, threads, and some aspects of the operating system. It also measures the performance of CPUs, caches, memory hierarchy, and the scalability of shared-memory architectures. Many vendors are releasing SPECjbb2000 results these days because it is a relatively simple-to-run benchmark that nevertheless does a credible job of measuring system-level performance on a common application type.

Benchmarking the Competition

In conjunction with the 900 MHz processors for the Midframe systems, Sun announced a new SPECjbb2000 benchmark for the Sun Fire 6800, the high-end of the Midframe line, of 174,658 operations per second, a result that Sun's press release headline screams, "Leaves IBM in the dust." But several aspects of this benchmark result are worth closer examination.

For starters, one has to forgive a certain hyperbole on the part of the headline writers. The IBM result on the newly announced high-end p690 Unix system (codenamed Regatta) was in fact 169,794

1. Sun defines "general availability" as the time when a system is shipping in volume to all geographies, within stated lead times, and within company guidelines for spares and service.

operations per second. Sun's 3 percent margin of victory over IBM's metric looks to be more of a photo finish than a blowout that raised a lot of dust. Sun achieves its result with a fully-configured 24-processor Sun Fire 6800 set against an IBM 16-processor p690. That's right; Sun required almost 50 percent more processors to achieve a result that was essentially a wash with IBM's.

The Star Fire 6800 also claimed a 31 percent price-performance advantage against the p690 (based on the SPECjbb2000 benchmark). But IBM's published benchmark with which Sun compared itself used the p690 HPC model. The HPC option 4-processor Multi-Chip Modules (MCM) provide twice the amount of Level 2 and Level 3 cache, as well as additional memory bandwidth, relative to the 8-processor MCMs in other p690 models. Because the premium-priced HPC option was designed with technical and high-performance computing in mind, it likely jacks up the system price without necessarily providing a commensurate performance gain in a commercial benchmark. Moreover, this option limits IBM to a maximum of 16 processors rather than 32.

But it's important not to lose sight of the fact that Sun's Star Fire 6800 did, after all, pull off a performance win against the new high-end of IBM's pSeries line—as measured by a widely-used industry standard benchmark. As for the specific IBM configuration that Sun chose for its comparison, Sun had little choice here; the 16-processor p690 HPC is the only SPECjbb2000 benchmark that IBM has released for this class of system. If nothing else, this result clearly points out that Sun's midrange performance is right up there with the largest IBM Unix systems. And Sun has the Star Fire 15K for customers who need more power. Even if Sun's UltraSPARC III processors aren't the speediest on a per-CPU basis, it usually doesn't much matter how many processors it takes to achieve a result as long as the system scales and the price—including software licenses—is right.

The Star Fire 6800 result also leapfrogs Hewlett-Packard's newest midrange Unix server—the HP Server rp8400. Announced in September with a SPECjbb2000 performance result of 118,547 operations per second (edging out the initial 750 MHz Sun Midframe benchmark), the 16-processor rp8400 result now trails both Sun and IBM by a significant margin.

Beyond Benchmarks: Partitioning for the Midrange

In many respects, Sun's Midframe announcement last March set new expectations for Unix midrange product feature-sets. Suddenly last Spring, midrange systems had to have features previously reserved for the high-end, such as partitioning and high levels of redundancy and availability.[2] However, the competition has not been idle in the months since then. Both HP and IBM have delivered systems with features that address at least some of what was new and unique about the Midframes.

IBM mainframes have partitioning capabilities that remain the gold standard against which Unix-based systems are measured. Mainframe-based LPAR (logical portioning and resource) provides almost complete abstraction of the underlying hardware—enabling system administrators to slice up individual processors and tie them to an arbitrary pool of other system resources independently of considerations such as physical board locations within the system. Of course, mainframe LPAR is completely dynamic—you just don't reboot mainframes—and conservative IBM mainframe customers have a level of confidence in the partitioning facility that only comes through years of use in the most demanding commercial environments. By the standards of the mainframe, Unix partitioning has a ways to go.

The complex details and feature sets associated with partitioning fall into three broad classes of complementary capabilities: physical partitioning, logical partitioning, and workload management. Physical and logical partitioning can both be either dynamic (changeable while the system is running) or static (set at system startup, thus requiring a reboot to alter).

2. Or at least had to provide a roadmap that committed to the timely delivery of high-end features.

Partitioning the Midrange

	HP rp8400	IBM p690	Sun Midframes
Total processors and max CPU speed	16 750 MHz PA-8700	32 (16 with HPC option) 1.3 GHz POWER4	24 (Sun Fire 6800) 900 MHz UltraSPARC III
Operating system	HP-UX 11i	AIX 5L Version 5.1	Solaris 8
Physical (backplane-based) partitions	2 (nPar) available today	None	4 (Sun Fire 6800) 2 (Sun Fire 3800,48x0) Static available today, dynamic planned 1Q02
Logical partitions	vPar planned for rp8400 in 2002 (on N-Class and L-Class today)	16 (LPAR) static in Dec 01, dynamic planned for mid-2002 with AIX 5L Version 5.2	None
Workload management	HP-UX Process Resource Manager (PRM), Workload Manager (WLM)	AIX Workload Manager (WLM)	Solaris Resource Manager (SRM)

All of the major Unix vendors offer one or more workload management products across their Unix-based product lines—HP-UX Workload Manager, AIX Workload Manager, and Solaris Resource Manager from HP, IBM, and Sun respectively. These products, as a group, allow administrators to dedicate specific levels and types of system resources to a specific application or process, thereby helping to ensure consistent performance and to minimize interference between applications. These capabilities provide some of the benefits of partitions in a non-partitioned system—but without nearly the same degree of protection from either hardware or software failures. They also provide a finer granularity of sharing computing resources—important since no Unix partitioning scheme can split a CPU yet.

Sun supports static hardware-based domains on the Midframe servers today, as many as two partitions on the Sun Fire 3800, 4800, and 4810 and four on the Sun Fire 6800. Sun has announced dynamic partitioning capabilities and is currently shipping Midframes that will support dynamic partitioning when Solaris 8 Update 7 ships, an event now targeted for the first calendar quarter of 2002.

Sun pitches hardware-based domains that partition at the board level as offering greater robustness and fault-tolerance than more granular software-based approaches. But this pitch seems more a case of making a virtue out of necessity than of a real belief that static domains on hardware boundaries offer any significant benefits to customers other than more flexible software-based dynamic domains. After all, Sun simultaneously insists that dynamic partitions also have great robustness and tolerance of faults. Furthermore, the dynamic reconfiguration associated with dynamic system domains is integral to the complete delivery of many Midframe customer-level benefits such as the ability to make hardware changes without system downtime.

Don't get us wrong. Static partitioning is one way to carve up a system. It is available today and customers whose needs are met by a specific partitioned configuration may find that it's all they really need. And buyers will be able to start using dynamic partitions on their systems as soon as the software is ready.

The p690 is IBM's first introduction of the LPAR concept to the Unix-based pSeries. Revenue shipments are planned for December. The initial release of LPAR for the p690 will support up to 16 logical partitions that do not need to conform to the physical boundaries of the system boards—they are, however, limited to processor-chip boundaries. Changes to the resources allocated to partitions are currently static, thus require administrative intervention. In part to address these limitations, IBM explicitly intends to make the AIX Workload Manager (WLM) an integral part of their LPAR strategy.

IBM makes much to-do of their mainframe heritage and their creation of the whole LPAR concept on those mainframes. The company even cleverly brought over the LPAR name to the pSeries—just as a reminder. However, IBM is the new kid on the block with Unix partitioning, and it will take field experience to prove that they have indeed translated their mainframe experience and software to the p690.

With the rp8400, HP has followed a similar strategy to Sun and IBM in that they are migrating downward features from their flagship high-end system—Superdome. Currently announced partitioning capabilities on the rp8400 (nPars) are relatively limited compared to both Superdome and competitive systems. The rp8400 can currently be divided into two independent hardware partitions connected through the system's crossbar fabric. The partitions must be split on a physical processor/memory "cell" boundary. HP plans dynamic partitions (vPars) for a future release of HP-UX in 2002.

Sun's Midframes Hold Their Lead

The debuts of new Sun systems this year have been tarnished by a slow economy—especially in many of the market segments that were making Sun a shining star just a year before. Simultaneously, Sun worsened its own lot with early missteps with UltraSPARC III. The company has also had to contend with another consequence of its recent successes—being target *numero uno* for aggressive competition. Both HP and IBM have taken aim at what they perceive as Sun's exposed flank with new midrange systems combined with much of the same performance, availability, and partitioning patter that Sun helped to pioneer.

Benchmark claims among competitive vendors are always a game of leapfrog and often difficult-to-compare "facts." The controversy and fog surrounding HP, IBM, and Sun's SPECjbb2000 benchmarks is nothing new. But the bottom line is that right now the 900 MHz processors give Sun competitive midrange performance advantage, at least in those application areas where SPECjbb2000 is a relevant benchmark. And if doubters point out that it takes Sun more processors than its competition to achieve a given performance level—well, that's neither particularly relevant nor new for the buyers of Sun systems.

No Unix systems vendor is shipping dynamic partitions yet on their midrange systems, although they all promise that it is on their near-term roadmaps. If Sun can stick to its committed delivery schedule for dynamic partitioning, Q102, any competitive feature gap will be small to nil. And when Sun does ship dynamic partitioning, customers will have a choice of deploying a deeper set of dynamic partitioning and dynamic reconfiguration capabilities on a far broader range of systems—the Sun Fire 3800 through the Sun Fire 15K—than any competitor.

The 900 MHz processors give Sun's Midframe servers a nice bump in performance that may not crush the competition but does position Sun solidly against HP's and IBM's midrange systems. This performance combines with product features that are as good as or better than the competition's. Sun's challenge is to maintain its leadership position in the face of a focused competitive threat—a feat that will require a continued emphasis on timely and smooth product delivery in addition to technological innovation.

Through subscription research, advisory services, speaking engagements, strategic planning, product selection assistance, and custom research, Illuminata helps enterprises and service providers establish successful infrastructure in five key areas: Server Technologies, Information Logistics, Application Strategies, Enterprise Management, and Pervasive Automation.