

MICROPROCESSOR *report*

Insightful Analysis of Processor Technology

THUNDERX RATTLES SERVER MARKET

Cavium Develops 48-Core ARM Processor to Challenge Xeon

By Linley Gwennap (June 9, 2014)

Cavium is nearing tapeout of its first high-end ARM processor, originally called Project Thunder but now branded as ThunderX. The first iteration, the CN88xx family, will have up to 48 ARMv8 cores running as fast as 2.5GHz. Individually, each ThunderX CPU is far less powerful than a Xeon CPU, but in total, they should generate performance similar to that of Intel's mainstream server processors, as Figure 1 shows. This performance would make Cavium the first ARM-compatible vendor to challenge Xeon E5, even if only for scale-out applications. Other ARM-compatible competitors, such as AMD and AppliedMicro, compete mainly against Intel's Atom server processor.

Drawing on its experience with the MIPS-based Octeon processors, Cavium packs these CPUs into an SoC design that includes 100Gbps of Ethernet capability plus PCI Express and SATA interfaces. Like Xeon E5, ThunderX includes a coherent interface to enable dual-socket (2S) configurations. For greater scalability, the chip also has a 200Gbps fabric switch. ThunderX features a cryptography accelerator and other offload engines. Despite this greater level of integration, we expect the 28nm processor will draw considerably less power than Xeon.

Cavium expects general sampling of ThunderX to commence in 4Q14. It will sell the processor in several versions with 24 to 48 cores and different combinations of accelerators and I/Os; these details will emerge closer to production. The company is also developing a lower-cost design that it will sell with 8 to 16 cores; this trimmed-down design will compete against Atom in microservers. The 16-core ThunderX will roll out about one quarter later than the 48-core chip. ThunderX will also sell into embedded applications such as networking and communications equipment.

Long Days of Thunder

When Cavium announced Project Thunder, its plan was to adopt its next-generation Octeon III design for the ARMv8 instruction set, maintaining as much similarity as possible between the two products (see [MPR 8/20/12](#), "Cavium Joins ARM Server Fray"). The company was already shipping 65nm Octeon II processors with up to 32 cores and an integrated SoC design, giving it a solid chassis for the new high-end designs. Using this plan, it expected to sample Octeon III by the end of 2012, and it hoped Thunder would race to market about a year later.

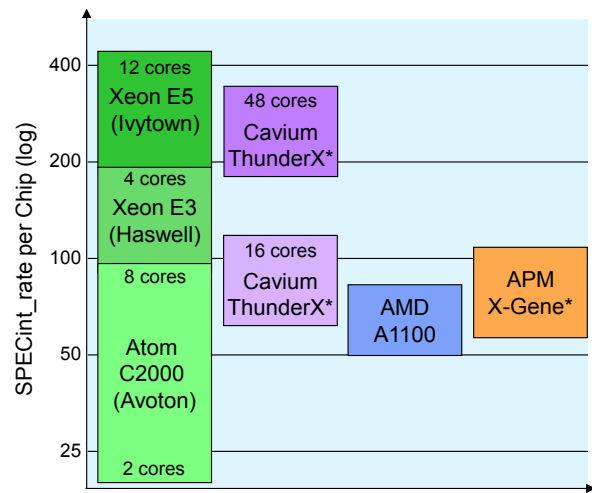


Figure 1. Server-processor performance. Whereas other ARM competitors are targeting Atom in the microserver market, Cavium's 48-core ThunderX processor targets Xeon E5. ARM performance is measured using GCC; Intel GCC performance is generously estimated as 10% less than reported ICC performance. (Source: vendors, except *The Linley Group estimate)

But the Octeon III development effort was a bumpy ride. The two-node jump required big changes in the circuit design. Furthermore, the design includes new features such as out-of-order execution (OOO) and a cache-coherent interconnect as well as a 3x increase in the internal-fabric bandwidth (see [MPR 2/13/12](#), “Cavium Octeon III Sizzles at 100Gbps”). Cavium’s engineers soon realized they couldn’t merely swap out the MIPS decoder for an ARM decoder; the ThunderX CPU needed major reworking to achieve ARMv8 compliance. Other aspects of the SoC needed tweaking to meet server-market requirements. These changes siphoned resources from Octeon III development.

As a result, the Octeon schedule slipped by more than a year, delaying ThunderX as well. (In the interim, Cavium managed to release a low-end version of Octeon III, but this chip lacks OOO as well as 2S support, and it scales to only four cores.) The company is now sampling the original Octeon III, which should make ThunderX bringup an easy cruise, since many of its function blocks come from Octeon III. Thus, we expect ThunderX to enter production in 2H15.

To drive its new server-processor business, the company hired Gopal Hegde as VP and GM in January. Hegde is experienced with ARM servers, having served as chief operating officer at Calxeda, an earlier startup in this market. He also held senior engineering roles on Cisco’s UCS servers and Intel’s Xeon processors. Calxeda hit the wall in December after customers snubbed its 32-bit ARM chips, but now, Hegde gets a second pass at this market using a more powerful 64-bit engine.

Assembling a Championship Roster

To take on the server market’s defending champion, Cavium needs a strong combination of features. Leading Thunder down the court is a small but speedy CPU. Cavium withheld microarchitecture details, but it has said the ThunderX CPU is similar to the Octeon III CPU, a simple dual-issue design with a nine-stage pipeline. Despite this relatively short pipeline, Cavium is targeting an impressive 2.5GHz using a custom circuit design in GlobalFoundries’ 28nm HPM process. The fastest Octeon processors currently shipping run at only 1.6GHz, so this higher speed is no slam dunk.

Although the current Octeon CPUs run strictly in order, Octeon III and ThunderX add limited instruction reordering. The company has not specified details, but we expect its reorder buffer includes about 40 entries—the same size that Cortex-A15 uses but far smaller than in Cortex-A57 (128) and Intel’s Haswell (192). This limitation, combined with only two instruction decoders, will significantly restrict ThunderX’s instructions per cycle (IPC) compared with these more sophisticated designs.

By focusing on performance per watt and per square millimeter, however, Cavium can pack 48 cores into a single chip. In contrast, Xeon E5 tops out at 12 cores, and AMD’s Seattle processor, which uses Cortex-A57, has only 8 (see

[MPR 2/24/14](#), “AMD to Sample First ARM Chip”). This shortfall may leave some Seattle fans pining for the former Thunder.

Following its small-ball philosophy, Cavium won’t try to match Xeon’s enormous caches. Each ThunderX CPU has 78KB of instruction cache and 32KB of data cache, and all share a 16MB L2 cache that operates at the same speed as the CPUs. Xeon E5 chips provide 256KB of cache per CPU plus a roomy 25–30MB shared cache. This configuration works out to 2.5MB per core for Xeon E5 versus just 0.33MB of shared cache per core for ThunderX. The larger caches reduce the number of time-wasting DRAM accesses, in turn reducing system power.

Fortunately, Cavium did not skimp on DRAM bandwidth; it gave ThunderX four 64-bit memory channels (with optional ECC) that can handle a total of 512GB of up to DDR4-2400 DRAM. These channels also support lower-cost DDR3 memory. In memory bandwidth, ThunderX outscores the current Xeon E5, which has four DDR3-1866 channels (see [MPR 9/16/13](#), “Xeon E5 Gets 22nm Refresh”). But Intel is likely to update its offering to the newest DRAM standard before ThunderX reaches production, erasing this advantage.

To connect multiple processors in a cache-coherent configuration, each ThunderX chip provides four ports implementing the Cavium Coherent Processor Interconnect (CCPI), which is similar to Intel’s QPI. Cavium has withheld the bandwidth and other details of its new interconnect. Introduced in Octeon III, the CCPI was originally intended to support up to eight coherent sockets, but the company will initially validate ThunderX only in 1S and 2S configurations. Sticking with two sockets increases bandwidth and reduces latency compared with larger configurations, which comprise less than 10% of all servers.

All Wrapped Up in a ThunderBall

The rest of the chip is where ThunderX shows its advantages. Xeon E5 offers up to 40 lanes of PCI Express Gen3, but for the server to have networking and storage connections, these lanes must connect to external Ethernet and SATA adapters. In contrast, ThunderX integrates these important I/O connections, as Figure 2 shows, eliminating the extra adapter cost. The chip includes several SATA3 ports that operate at up to 6Gbps, making it well suited to Hadoop and other big-data applications. The Ethernet ports use ten 10Gbps serdes that will typically be configured as up to 10x10GbE, although they can support 40G Ethernet as well. If Ethernet is not enough, the chip provides two eight-lane ports of PCIe Gen3.

The integrated Ethernet fabric is another feature that Xeon lacks. ThunderX has an additional 20 serdes and Ethernet MACs that can be configured as 20x10GbE, 5x40GbE, or even 2x100GbE. These ports connect to an internal 200Gbps switch that can analyze incoming traffic and route it back out without CPU intervention. This fabric

can link servers in mesh network, replacing the top-of-rack (ToR) switches that many data centers use.

ThunderX also includes several coprocessors from Octeon III. These coprocessors offload certain functions from the CPUs, such as cryptography, regular-expression (reg-ex) matching, and data decompression. Web servers handling secure information (e.g., HTTPS) can use the crypto engine, eliminating the need for an SSL add-in card. The other engines are available for more specialized workloads, such as deep packet inspection. Cavium plans to offer different versions of the processor that have various combinations of coprocessors enabled, so customers need only pay for the ones they plan to use, providing at least a quantum of solace.

About the only thing missing from this design is the baseboard-management controller (BMC). This function is usually implemented as a standalone microcontroller that configures and monitors the server motherboard, although Calxeda integrated the BMC in some of its products. Some data-center operators don't even use the BMC. For these operators, a complete server board simply requires adding memory and PHY chips to the processor.

While developing ThunderX, Cavium also kept a finger on the pulse of the embedded market. Many embedded customers are shifting their software development to ARMv8. These customers can use the chip's coprocessors for data-plane processing, security equipment, storage servers, media servers, and cloud RAN, among other applications.

ThunderCats Battle Evil

Although ThunderX offers integration advantages over Xeon, these processors will compete mainly on performance, power, and cost. After much roaring about its leadership position, however, Cavium declined to publicly release any specific data. Given the limitations of its microarchitecture, we estimate the ThunderX CPU will have 20–30% lower IPC than Cortex-A57, albeit at lower power. AMD expects its Seattle processor, packing eight 2.0GHz A57 cores, to exceed 80 SPECint_rate2006 (using GCC). Thus, ThunderX could achieve about 70 SPECint when using eight cores at 2.5GHz.

Linearly extrapolating this score to 48 cores would yield 420 SPECint, or about the same score as Intel's top-of-the-line Xeon E5 with 12 Ivy Bridge cores running at 2.7GHz. Given ThunderX's modest L2 cache, we project its 48-core score at 350, or about 20% less than the linear extrapolation. But whether the chip can claw its way to 300 or 400 SPECint, it will certainly overlap a large portion of the Xeon E5 range. ThunderX performance will be even better for applications—such as security, storage, and networking—that can take advantage of the chip's offload engines.

Price and Availability

Cavium expects its CN88xx (ThunderX) processor will reach general sampling in 4Q14. The company has not announced pricing. For more information on ThunderX, access www.cavium.com/ThunderX_ARM_Processors.html.

For workloads that require high single-thread performance or a large cache, however, the ARM processor will fall well behind Xeon. At its top speed, a Xeon CPU delivers about four times the SPECint performance of an Atom CPU, and ThunderX will be hard pressed to keep up with even Atom in core-to-core combat. Thus, the Cavium processor will be best suited to workloads that can be easily divided among many cores. Conversely, Intel will remain king of the jungle in scale-up applications that require brawny CPUs. We also expect it to win in HPC and other tasks that require strong floating-point performance, as Cavium's microarchitecture is not optimized for FP throughput.

Regarding power, Cavium would say only that its processor will be "well under 100W TDP" even with 48 cores running at their top speed of 2.5GHz. The simple CPU design is considerably smaller than Cortex-A57 and uses less power. We estimate ThunderX could garner an 80W TDP rating at its top speed and closer to 50W at 2.0GHz. By comparison, Xeon E5 has a 130W rating at top speed and 60W in its lowest-power version (1.7GHz), despite using Intel's more advanced 22nm FinFET technology.

Xeon's TDP measurements exclude the south bridge and the Ethernet adapters needed to match the functions ThunderX provides. Intel's C604 south bridge offers six SATA ports (two Gen3 plus four at the slower Gen2 rate)

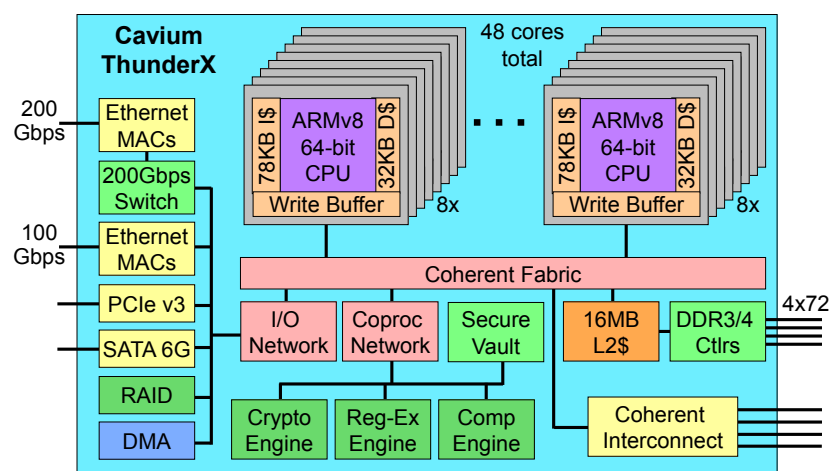


Figure 2. Block diagram of Cavium ThunderX processor. The device includes 48 CPU cores arranged in clusters of eight, reducing the number of ports on the internal fabric. The SoC includes a 200Gbps Ethernet switch and several accelerators, as well as DRAM controllers and other I/O.

	Cavium ThunderX	Intel Xeon E5-2470 v2	AMD A1100	AppliedMicro X-Gene
CPU Type	ThunderX	Ivy Bridge	Cortex-A57	X-Gene
Instr Decode Rate	2 ARMv8	4 x86	3 ARMv8	4 ARMv8
Clock Speed	2.5GHz	2.4GHz	2.0GHz±	2.4GHz
Cores / Threads	48 cores / 48T	10 cores / 20T	8 cores / 8T	8 cores / 8T
SPECint_2006*	350±	320±	80	110±
Coherent Sockets	2S	2S	1S	1S
Power (TDP)	80W±	95W+24W+	25W	35W±
IC Process	28nm HKMG	22nm FinFET	28nm HKMG	40nm G
Production	2H15±	1Q14	4Q14 (est)	2Q14

Table 1. Cavium ThunderX versus current server processors. ThunderX is far more powerful than other ARM server processors and matches up well against a midrange Xeon E5, but it won't ship until late next year. *SPECrate score for GCC (ICC scores derated by 10%); †24W includes C604 south bridge plus four 10GbE MACs. (Source: vendors, except †The Linley Group estimate)

but adds 12W to the platform. The company's 82559 (Niantic) chip provides two 10GbE MACs for 6W. Adding the C604 plus 10x10GbE ports to match ThunderX would increase Intel's processor TDP by 42W, although that many Ethernet ports is overkill for most servers. Taking these factors into account, the Cavium chip should require considerably less board-level power than Xeon E5, as Table 1 shows.

Cavium has not announced pricing for any of the ThunderX versions. Most Xeon E5 products carry a list price of \$885 to \$2,614, although a few (with six or fewer cores) sell for less. These prices are about double what Cavium charges for its Octeon processors, so it should be able to easily undercut Intel on price. Cavium's price advantage becomes greater when considering the cost of the extra south bridge and Ethernet chips that Xeon requires.

By the time Cavium's first server processor leaps into production in 2H15, competitors may have improved their products. In 2015, Intel plans to upgrade Xeon E5 to Haswell, which offers about 10% greater performance at the same power levels. AppliedMicro is already sampling X-Gene 2, a 28nm shrink that reduces power but remains at eight CPUs. Neither of these upcoming products significantly changes the competitive comparison. If ThunderX slips into 2016, however, it will face 14nm Xeon E5 products and AMD's second-generation ARM processors.

Thunderbirds Are Go!

Cavium announced that several third parties will support ThunderX, including Canonical and MontaVista (a Cavium

subsidiary) for Linux; AMI for enterprise systems management and UEFI software; and Gigabyte for reference designs. The company will offer at least two reference designs, an ATX board and a half-SSI board, that can plug into standard chassis. The new processor complies with the ARM SBSA standard (see [MPR 2/24/14](#), "ARM Releases Server-System Spec"), which should ease the porting of server software to the ARM processor.

Sporting 48 cores, ThunderX flies far above the per-socket performance of any ARM server processor yet announced. This achievement draws on Cavium's long experience with large multi-core processors; the company had to rework the ARM interrupt controller and substitute its own internal fabric for ARM's licensed interconnect.

Using its ARMv8 license, Cavium designed a small and power-efficient CPU; a processor with 48 Cortex-A57 cores would have been far too unwieldy.

Other ARM vendors are competing in the microserver market, where price tags are sinking below \$100. With its high performance and multisocket support, ThunderX can rescue ARM from this trap, instead jetting into the stratosphere of four-digit price tags. Cavium plans to deploy a junior version of ThunderX, but it is probably more interested in serving its embedded customer base with this device than wading into the microserver morass.

Compared with Xeon, ThunderX could deliver 50% to 100% more performance per watt and per dollar, particularly when considering the additional chips that Intel needs to complete the server design. Integrating the system fabric switch, a feature that Calxeda also provided, further reduces system (rack-level) cost compared with a Xeon-based design. These advantages should be enough to gain design wins against Intel, particularly in a server market that is hungry for an alternative to the x86 giant.

ThunderX cannot compete with Xeon across the board. Applications that require a muscular CPU, high FP performance, or lots of cache are poorly suited to the new processor, and Windows Server customers need not apply. ThunderX is still several months from sampling, much less production, and so will not be troubling Intel any time soon. But this announcement shows that the ARM camp is making progress in the server market, with Cavium leading the way. ♦

To subscribe to *Microprocessor Report*, access www.linleygroup.com/mpr or phone us at 408-270-3772.