Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

CAVIUM BRANCHES OUT

New Networking Processors Integrate 2–16 MIPS Cores per Chip

By Tom R. Halfhill {10/5/04-01}

Already a well-regarded vendor of security processors, Cavium Networks is moving in a bold new direction. The company's new Octeon family of networking processors integrates three important functions in a single chip: packet processing, content filtering, and security.

To provide enough muscle for all that heavy lifting, at line rates up to 10Gb/s, Octeon chips will have as many as 16 MIPS-compatible 64-bit processor cores, augmented by numerous coprocessors.

Cavium refers to the Octeon chips as "network services processors"—a new data-plane category that absorbs the functions of separate packet processors, security processors, and content-filtering accelerators. Cavium believes all those Layer 3–7 data-plane functions are becoming so mandatory it's time to integrate them in a single chip. Higher-level integration can dramatically shorten the packet datapaths, eliminate redundant packet processing, simplify board designs, and cut costs. It's a grand strategy, and Octeon chips have ample resources to carry it out.

Target applications are secure network-interface cards, routers, and wireless-LAN switches; server load-balancers, web-service appliances, router blades, and content-filtering switches; and host bus adapters and switches for network storage subsystems. Currently, these applications require multiple chips—often including ASICs and FPGAs—to handle the packet processing, filtering, and security functions consolidated in Octeon. Cavium claims Octeon will typically cut costs to one-fifth of existing solutions and similarly reduce board area and power consumption.

The Octeon family continues Cavium's tradition of designing high-performance processors for vital networking tasks. Last year, Cavium's Nitrox Plus CN1340p won our *Microprocessor Report* Analysts' Choice Award for Best Security Processor of 2002. (See *MPR 2/18/03-09*, "Security By Design.") Cavium also sells a line of chips known as Golden Gate Bridge processors—small I/O bridge chips for SPI-3, SPI-4.2, PCI, and PCI-X. Cavium announced the Octeon family on September 13 and took the stage at **Fall Processor Forum** on October 5 to reveal the first technical details about the new processors, which are scheduled to begin sampling in 1Q05 and enter production in 2H05.

Return of the Marlboro Men

Octeon was largely designed by engineers who created Digital's Alpha 21364, code-named EV7. (See *MPR 3/24/03-02*, "EV7 Stresses Memory Bandwidth.") A hand-picked team of about 35 designers in Marlboro, Massachusetts, worked under Cavium lead architect Richard Kessler—formerly chief architect of the EV7—and Anil Jain, vice president of IC engineering.

Cavium's IC team labored about 18 months on the Octeon project, creating a new full-custom implementation of the MIPS64 processor core, dubbed the cnMIPS64 ("cn" stands for Cavium Networks or content networking). The new core is the first implementation of the MIPS64 Release 2 specification and omits the FPU, which is unnecessary for Cavium's target applications. Cavium says the cnMIPS64 core is three to five times faster than a synthesized MIPS64 core.

Octeon chips will integrate 2, 4, 8, or 16 cnMIPS64 cores. Each CPU has two-way superscalar five-stage pipelines, so a 16-core Octeon chip can execute as many as 32 ALU

instructions per cycle. At Cavium's modest target clock frequency of 600MHz—TSMC will manufacture the first-generation chips in a proven 0.13-micron CMOS process—the peak aggregate performance is an impressive 19.2 billion instructions per second.

Frankly, not many customers need that much processing power today. We expect the dual- and quad-core Octeon chips to be the most popular, with perhaps a few performance-minded customers reaching for the eight-core model. The 16-core Octeon is overkill for all but the most demanding applications, but it gives Cavium bragging rights over the latest networking processors from AMCC, Broadcom, Freescale, IBM Microelectronics, PMC-Sierra, and other competitors. More important, it virtually assures that Cavium's customers won't run out of processing power anytime soon. It's easier for a vendor to sell processors that have plenty of headroom for applications to grow.

Other function blocks in the Octeon chips benefit from the same custom handiwork that went into the cnMIPS64 cores. Cavium's IC team designed its own SPI-4.2 interface, Gigabit Ethernet media-access controller (MAC), and PLL. For less critical on-chip peripherals and function blocks, the designers used automated place-and-route tools instead of hand-packing the circuit layouts.

Meanwhile, Cavium's vice president of software and system engineering, Raghib Hussain, led a team of about 35 programmers in Santa Clara, California, and Hyderabad, India, to write software for the new processors. They wrote custom application stacks that include microcode drivers, device drivers, and dynamic-link libraries (DLL). Cavium says Octeon will come with enough software to get a virtual private network (VPN) up and running 30 minutes after installing a development board. More important, the same software is compatible with every Octeon, so customers can upgrade to chips with additional cores without changing the code.

Software development for Octeon should be relatively easy, by network processor standards. The CPU cores are familiarly MIPS compatible, and there are no proprietary microengines requiring special programming. Moreover, software will be compatible across all members of the Octeon family, no matter how many CPU cores they have. The biggest challenge will be writing code that takes advantage of all the CPUs and their coprocessors. Fortunately, there's a great deal of inherent parallelism in packet processing, unlike in general-purpose computing. Cavium is already shipping early customers GNU software-development tools with customized back ends.

Optimized for Data-Plane Processing

Cavium calls Octeon a network services processor (NSP), not a network processor (NPU), because the chips combine so many different, but complementary, functions. First, Octeon is a packet processor. The dual-core chips will be suitable for low-end packet-processing applications, such as home routers, while the 8- and 16-core chips will be powerful enough for Internet routers. In addition, Octeon has more pattern-matching hardware for Layer 3–7 content filtering than a typical NPU has, so it will excel at those tasks, too. Octeon's third role is to serve as a watchdog, because it has special security hardware to support intrusion detection, antivirus scanning, firewalls, and cryptography.

Customers can use Octeon in different ways. At the two extremes, every CPU core could run data-plane code or a separate instance of an operating system for control-plane duties, because each core has its own memory management unit (MMU) and translation lookaside buffer (TLB). Of course, any combination between those extremes is possible. One or two CPUs could run an operating system, freeing the other CPUs to perform data-plane tasks. Octeon will run new multiprocessor versions of MontaVista Linux and Wind River's VxWorks, scheduled to ship in 1H05. However, Cavium isn't touting Octeon as a control-plane processor. Other vendors' chips that operate at higher clock frequencies and are less burdened with data-plane logic are more suitable for that purpose. Cavium hints that an Octeon-based line of control-plane processors is likely in the future.

There are good arguments for integrating multiple networking functions in a single chip. It eliminates redundancy, because different functions relying on similar operations (such as packet filtering and antivirus scanning, both of which perform bit-field pattern matching) can share the same acceleration hardware. In addition, packets don't have to meander through several different processors to be checked for infections, sorted by content, and routed to their destinations. An Octeon chip can perform all those functions locally and deliver a "blessed" (filtered and secured) packet to the network.

One challenge of integrating all these services in a single chip is ensuring the processor has enough performance to keep from clogging the network pipes. That's why Cavium spent so much time developing a full-custom processor core, custom I/O controllers, and special-purpose coprocessors.

CPU Cores Have Coherent Caches

All chips in the four-member Octeon family share a similar design, except for the number of CPU cores, memory buses, and packet I/O interfaces. Each CPU core has its own 32KB instruction cache, a fully coherent 8KB write-through data cache, 2KB write buffer, MMU, and 32-entry TLB. All cores share a 1MB L2 write-back cache that's eight-way set-associative and supports 128-byte line locking and partitioning. CPU cores access the L2 cache over a coherent bus at their core clock frequency of 600MHz, providing 230Gb/s of bandwidth. CPUs can access the L2 cache on every clock cycle, with a typical latency of 12 cycles between the L1 and L2 caches.

On its back side, the L2 cache connects to an on-chip SDRAM controller, which addresses up to 16GB of main memory. The dual- and quad-core Octeon CN3420 and CN3430 chips have a single-channel 64-bit DDR1/DDR2

interface (72 bits with ECC). The 8- and 16-core Octeon CN3840 and CN3860 chips have a single-channel 64/128-bit DDR1/DDR2 interface (72/144 bits with ECC). The memory controller supports DDR1 or DDR2 SDRAM at bus frequencies up to 400MHz, effectively 800MHz at double data rates. Therefore, maximum theoretical bandwidth to main memory on an 8- or 16-core Octeon is 12.8GB/s. Maximum bandwidth is 6.4GB/s on the dual- and quad-core Octeon chips. Figure 1 shows a block diagram of an Octeon processor.

A second on-chip memory controller optionally supports up to 1GB of Reduced Latency DRAM (RLDRAM) or Fast Cycle RAM (FCRAM). This is another feature distinguishing lower-end Octeon chips from higher-end members of the family. The dual- and quad-core CN3420 and CN3430 have a single eight-bit RLDRAM/FCRAM bus (9 bits with ECC); the 8- and 16-core CN3840 and CN3860 have a pair of eight-bit buses (18 bits with ECC).

RLDRAM, developed by Infineon and Micron, runs at speeds up to 535MHz and reduces row cycle times to as little as 15ns, less than half the cycle time of standard SDRAM. FCRAM, developed by Fujitsu, uses pipelining, hidden precharging, and simultaneous row-and-column accesses to reduce random cycle times by a similar degree. Neither RLDRAM nor FCRAM is intended to replace large amounts of SDRAM for main memory. Instead, they store smaller amounts of frequently accessed data that the processor might otherwise hold in more costly SRAM or in locked regions of the data cache. Cavium suggests using these low-latency memories to store data for graph walks and string searches—common pattern-matching operations for packet processing, content filtering, and virus detection.

Packet I/O interfaces are the third feature distinguishing lower-end members of the Octeon family from their higher-end brethren. The dual- and quad-core CN3420 and CN3430 have four Reduced Gigabit Media-Independent Interfaces (RGMII), which support four Gigabit Ethernet (GbE) MACs. The 8- and 16-core CN3840 and CN3860 have eight RGMIIs, which support eight GbE MACs or two SPI-4.2 interfaces. All Octeon processors have a 64-bit 133MHz PCI-X host/slave controller, which can act as a data or control interface. Unfortunately, PCI Express won't make it into this Octeon generation, but a $10–$15 bridge chip is a stopgap solution, and PCI-X will continue to dominate networking systems for a while. There's no HyperTransport interface, either, but Cavium says its target customers at the high end are mainly using SPI-4.2. Table 1 lists the features of all four Octeon processors.

Custom Hardware Boosts Performance

Cavium's cavalcade of CPU cores is attracting the most attention, but even the 16-core processor couldn't deliver the performance Cavium promises without help from additional custom logic. That's why Octeon chips have more special-purpose hardware than any other networking processors. These accelerators offload critical tasks from the CPU cores and eliminate the need for extra chips, such as regular-expression (reg-ex) engines. The coprocessors speed up packet I/O, cyclic redundancy checks, memory allocation, DMA transfers, header processing, security scanning, pattern matching, cryptography, data compression, and decompression.

Octeon's packet I/O processor is programmable, has multiple DMA engines per port, supports IPv4 and IPv6, and handles L2–L4 parsing and exception checking. It can strip off packet headers for storage in memory and move packet payloads into memory for content filtering and other processing. The whole datapath is designed to rapidly process the packets as a continuous datastream without making redundant copies, unlike some other network processors. Indeed, packets requiring no special processing can bypass the MIPS CPU cores altogether while other packets get the full treatment.

That treatment may include virus scanning, firewall intrusion detection, various types of content filtering, decryption, and data decompression. For instance, 16 regular-expression engines accelerate the pattern- and signature-matching code that scans for telltale signs of viruses, worms, firewall attacks, and other malware. The same engines are useful for many content-filtering tasks, such as blocking spam

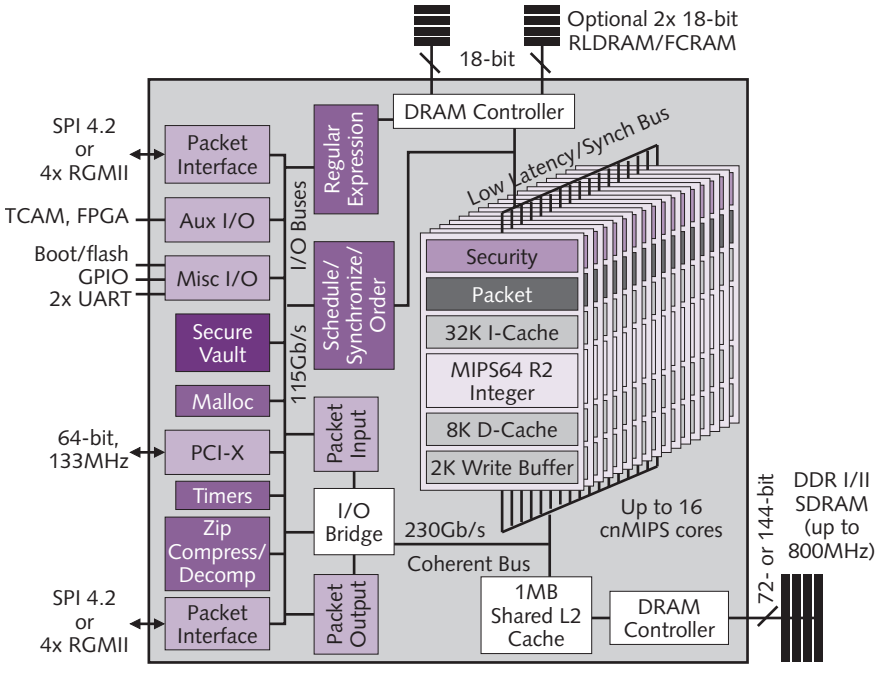

Figure 1. Octeon chips share a similar microarchitecture, differing mainly by the number of CPU cores (2, 4, 8, or 16), packet I/O interfaces, and memory buses. Note the on-chip L2 cache, memory controllers, I/O controllers, and acceleration hardware.

Feature	Cavium Octeon CN3420	Cavium Octeon CN3430	Cavium Octeon CN3840	Cavium Octeon CN3860
CPU Architecture	MIPS64-R2	MIPS64-R2	MIPS64-R2	MIPS64-R2
Architecture Width	64 bits	64 bits	64 bits	64 bits
CPU Cores	2	4	8	16
Core Frequency	600MHz	600MHz	600MHz	600MHz
Superscalar Issue	2-way	2-way	2-way	2-way
ALU Pipeline Depth	5 stages	5 stages	5 stages	5 stages
ALU Instr Per Sec	2.4 billion	4.8 billion	9.6 billion	19.2 billion
Instruction Cache	32K, 64-way	32K, 64-way	32K, 64-way	32K, 64-way
Data Cache	8K, 64-way	8K, 64-way	8K, 64-way	8K, 64-way
L2 Cache	1MB, 16-way	1MB, 16-way	1MB, 16-way	1MB, 16-way
TLB	32-entry	32-entry	32-entry	32-entry
FPU	—	—	—	—
DRAM Controller	DDR1/DDR2 800MHz 64b single-channel	DDR1/DDR2 800MHz 64b single-channel	DDR1/DDR2 800MHz 128b single-channel	DDR1/DDR2 800MHz 128b single-channel
Max DRAM Bandwidth	6.4GB/s	6.4GB/s	12.8GB/s	12.8GB/s
Max DRAM Memory	16GB	16GB	16GB	16GB
RLDRAM/FCRAM	1 x 8-bit I/O	1 x 8-bit I/O	2 x 16-bit I/O	2 x 16-bit I/O
Max RLDRAM/FCRAM	1GB	1GB	1GB	1GB
Packet I/O Interfaces	4 x RGMII (4 x GbE MACs)	4 x RGMII (4 x GbE MACs)	8 x RGMII (8 x GbE MACs) (2 x SPI-4.2)	8 x RGMII (8 x GbE MACs) (2 x SPI-4.2)
PCI / PCI-X Controller	PCI-X 64-bit 133MHz Host/slave	PCI-X 64-bit 133MHz Host/slave	PCI-X 64-bit 133MHz Host/slave	PCI-X 64-bit 133MHz Host/slave
PCI Express	—	—	—	—
HyperTransport	—	—	—	—
Other I/O	Flash, UARTs, MDIO, GPIO, TCAM, FPGA	Flash, UARTs, MDIO, GPIO, TCAM, FPGA	Flash, UARTs, MDIO, GPIO, TCAM, FPGA	Flash, UARTs, MDIO, GPIO, TCAM, FPGA
TCP Offload Engine	Yes	Yes	Yes	Yes
ZIP Compress Engine	Yes	Yes	Yes	Yes
Reg-Expression Engine	16	16	16	16
Crypto Engines	DES, 3DES, AES, RSA, DH, MD5, SHA-1, RC4	DES, 3DES, AES, RSA, DH, MD5, SHA-1, RC4	DES, 3DES, AES, RSA, DH, MD5, SHA-1, RC4	DES, 3DES, AES, RSA, DH, MD5, SHA-1, RC4
True RND Generator	Yes	Yes	Yes	Yes
Memory-Alloc Engine	Yes	Yes	Yes	Yes
Fabrication Process	TSMC 0.13µm	TSMC 0.13µm	TSMC 0.13µm	TSMC 0.13µm
Packaging	709 pins	709 pins	1,500 pins	1,500 pins
Power (Worst-Case)	5W	8–10W	14–16W	25W
Price (10K Units)	$125	n/a	n/a	$750
Availability	Samples 1Q05; Production 2H05	Samples 1Q05; Production 2H05	Samples 1Q05; Production 2H05	Samples 1Q05; Production 2H05

Table 1. Octeon processors with 8 or 16 CPU cores have single-channel 64/128-bit main-memory interfaces, two 8-bit RLDRAM/FCRAM interfaces, and two SPI-4.2 interfaces that can operate as eight GbE MACs. Octeon processors with two or four CPU cores have single-channel 64-bit main-memory interfaces, an 8-bit RLDRAM/FCRAM interface, and four GbE MACs, but no SPI-4.2. PCI-X is standard, but first-generation Octeon chips won't have PCI Express. Also standard are several auxiliary I/O interfaces (not shown), such as flash memory, UARTs, management data input/output (MDIO), 2-Wire, FPGA, and general-purpose I/Os (GPIO). n/a = not available

or sorting packets for quality-of-service (QoS) routing. Octeon processors can perform these kinds of analytical tasks at line rates up to 4Gb/s.

Some malware hides in compressed file attachments, which the processor must decompress before scanning and then recompress for further routing. Performing those operations entirely in software wouldn't allow the processor to handle its other tasks at 4Gb/s line rates, so Octeon has a programmable coprocessor to accelerate data decompression and recompression. The ZIP coprocessor supports fixed or dynamic Huffman compression algorithms and other features of the GZIP and PKZIP protocols, at line rates up to 4Gb/s.

To speed up data encryption and decryption for secure network connections, Octeon has additional coprocessors for the most popular cryptography standards and algorithms: DES, Triple DES (3DES), AES, Message Digest 5 (MD5), the Secure Hash Algorithm (SHA-1), Diffie-Hellman, RSA, and RSA's Rivest Cipher 4 (RC4) symmetric stream cipher, which is part of the Secure Sockets Layer (SSL) in virtually all web browsers. To support those algorithms, ciphers, and protocols, there's also a true random-number generator.

Depending on how many other tasks it must perform, Octeon can encrypt or decrypt packets fast enough to support line rates from 500Mb/s to 10Gb/s. Of course, simpler packet processing makes the faster line rates possible. A TCP

offload engine (TOE) helps to synchronize the packets and manage the packet buffers, providing full TCP packet termination at 10Gb/s.

Higher Integration Was Inevitable

MPR is always skeptical when a vendor claims to have created a new category of processors, and with good reason—it rarely happens. In this case, Cavium has some justification for its marketing pitch. No other processor integrates as many networking functions as Octeon does, and the few processors that aspire to Octeon's level of feature integration don't back up their aspirations with as much integrated logic. Octeon is truly an NSP, because in one chip, it combines the services becoming common in networking systems.

Cavium is arriving at its destination from the opposite direction of most other companies. While other chip vendors are adding security features to their network and communications processors, Cavium is, in effect, adding packet processing and other network services to its security processors. For instance, Freescale recently introduced its new PowerQuicc II Pro family, which integrates a security engine formerly available only as a separate chip. In addition, Freescale has added security to a new PowerQuicc III. (See *MPR 5/10/04-02*, "Freescale Secures PowerQuicc.") But on-chip security still isn't universal across the PowerQuicc product line, and Octeon chips have greater processing resources.

For data-plane tasks, Octeon generally outclasses the new networking processors announced at Fall Processor Forum by Broadcom, Faraday, Freescale, and PMC-Sierra. It's not just that Octeon chips have up to 16 CPU cores. They also have more special-purpose hardware, security, and I/O. Surprisingly, those extra features don't seem to impose a significant power-consumption penalty. Cavium claims Octeon's power consumption will range from 5W to 25W, depending on the number of CPU cores; this compares favorably with the estimated power consumption of Octeon's less integrated rivals. Octeon will almost certainly cost more—production prices range from $125 to $750 in 10,000-unit quantities—but its higher integration will lower the overall system chip count and therefore reduce the total bill of materials.

Price & Availability

Production pricing for the four-part Octeon family ranges from $125 for the two-core CN3420 to $750 for the 16-core CN3860 in 10,000-unit quantities. Sample chips and evaluation boards will be available in 1Q05; production quantities are scheduled to be available in 2H05. A development kit with a simulator, GNU software-development tools, and reference applications is available now. For more information, see *www.cavium.com/octeon.html*.

Additional competition for Octeon will come from EZchip's NP-2 family of network processors. (See *MPR 4/5/04-02*, "EZchip Adds New NPU Line.") Like Cavium, EZchip is targeting line rates up to 10Gb/s with highly integrated chips that can process packets down to Layer 7. Some of the NP-2 chips announced so far have two 10Gb/s traffic-management coprocessors, 10 GbE MACs, one 10Gb/s Ethernet MAC, and two SPI-4.2 interfaces. Those chips are scheduled to sample this quarter. EZchip says future members of the NP-2 family, coming in 2005, will add TOE and security engines—a more direct challenge to Octeon. However, EZchip's processors have a proprietary CPU architecture, which is less familiar to programmers than Octeon's MIPS architecture. And we will be surprised if EZchip can match the capabilities of Cavium's integrated security engines, because security is Cavium's strong suit.

Of course, Cavium hasn't invented the concept of higher integration. Growing transistor budgets and the 40-year trend in larger-scale IC design made it inevitable that someone would unite a handful of separate networking chips into one tight package. Cavium deserves credit for expanding the boundaries of what a single networking chip can do—and for designing an unusually powerful implementation. Octeon is surely a harbinger of things to come.