

MICROPROCESSOR *report*

Insightful Analysis of Processor Technology

LX2160A Is NXP's BIGGEST MULTICORE

New 16-CPU Processor Has 100GbE and Integrated Ethernet Switch

By Tom R. Halfhill (October 16, 2017)

NXP is chasing high-end networking with its newest QorIQ processor, the LX2160A. Sporting 16 ARM Cortex-A72 cores, 100 Gigabit Ethernet, a 16-port Layer 2 switch, and faster acceleration for cryptography and data compression, it will be the company's largest and fastest multicore embedded processor when it begins production—in mid-2019, by our estimate.

Announced at the recent Linley Processor Conference, the LX2160A is the most ambitious QorIQ design since the 12-core T4240, which began production five years ago. That chip implements the Power Architecture, not ARMv8. Although the LX2160A surpasses the T4240's core count and performance, the older product still offers more threads thanks to its dual-threaded Power e6500 CPUs. (No vendor has released a multi-threaded ARM core.) Nevertheless, the LX2160A has twice as many ARM CPUs as any existing QorIQ and ranks among the largest ARM-based embedded chips announced to date.

The LX2160A is nominally a general-purpose embedded processor but is optimized for networking. It targets enterprise-class virtual customer-premises equipment (vCPE), edge routers that host virtual machines, switch controllers, storage controllers, mobile edge computing, and 5G base stations (building on NXP's strong position in 4G). It also has features for industrial applications and advanced driver-assistance systems (ADAS)—sockets that lesser QorIQs are already winning.

Its powerful CPU complex is particularly suited to network functions virtualization (NFV); the chip can run up to 16 heavyweight virtual network functions (VNFs). NXP strengthened

the offload engines, too: cryptography processing is 67% faster than the T4240's, and the compression and decompression engines are 5x faster.

The LX2160A is also the company's first chip built with FinFETs. All existing ARM-based QorIQ models and Power-based T-series models are 28nm planar devices. For the LX2160A, the company chose the 16nm FinFET-C (16FFC) process at TSMC. Compared with the foundry's original 16nm FinFET+ technology, 16FFC sacrifices a small amount of performance for lower cost. It will enable the LX2160A to undercut the prices of competing chips that beat it to earlier FinFET technologies.

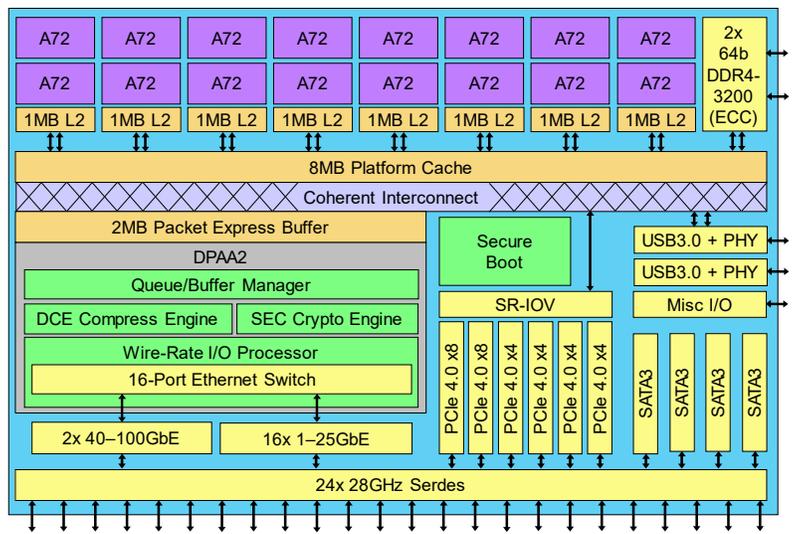


Figure 1. Block diagram of NXP QorIQ LX2160A. In addition to 16 CPUs, major features include 100G Ethernet, a 16-port Ethernet switch, large ECC-protected caches and buffers, PCI Express Gen4, and faster accelerators for cryptography and compression.

Price and Availability

NXP plans to sample the QorIQ LX2160A in 1Q18. We estimate it will qualify for volume production in 2Q19. The company is withholding the list price for now, but we expect the LX2160A to cost about \$350 in 1,000-unit volumes. Derivatives of the 16-core chip will include the 12-core LX2120A and the 8-core LX2080A, which should begin production at about the same time. For more information, access www.nxp.com/lx2160.

Generous Caches and Faster Offloads

As Figure 1 shows, the LX2160A arranges its 16 Cortex-A72 CPUs in pairs that each share a 1MB L2 cache (see [MPR 2/16/15](#), “Cortex-A72 Takes Big Step Forward”). All the CPUs share an 8MB “platform cache”—NXP’s term for a higher-level cache that stores descriptors, keys, contexts, and other data needed for packet processing. By keeping data-plane traffic out of external memory, the platform cache conserves memory bandwidth and reduces the system cost. In all, the LX2160A has 16MB of ECC-protected L2/L3 cache, which is generous for a processor in this class. The 16-core Cavium Octeon TX, for example, has only 8MB.

The target CPU clock frequency is 2.2GHz, which is conservative for a 16nm processor. Broadcom’s new BCM58808H is the fastest Cortex-A72 chip announced to date—it reaches 3.0GHz in 16nm (see [MPR 8/14/17](#), “Broadcom Hits Smart NICs and Storage”). By limiting the CPU speed, NXP hopes to hold the LX2160A’s power consumption to 30W thermal design power (TDP). That’s about the same TDP as the QorIQ LS2088A, which has

only half as many A72s operating at 2.0GHz. At 30W, the chip doesn’t need a CPU fan.

Acceleration hardware offloads the most tedious tasks from the CPUs. The enhanced SEC security engine can process IPsec streams at 50Gbps, versus 30Gbps for the T4240. Likewise, the new compression and decompression engines can simultaneously operate at 50Gbps, five times faster than existing QorIQs. In addition, the LX2160A implements NXP’s second-generation Data Plane Acceleration Architecture (DPAA2), which accelerates other packet-processing chores and offers developers an abstract programming interface (see [MPR 7/20/15](#), “Freescale Overhauls the Data Plane”).

The LX2160A’s DPAA2 implementation includes the wire-rate I/O processor (WRIOP) but not the advanced I/O processor (AIOP). The latter subsystem contains programmable CPUs and is more capable than the fixed-function WRIOP, but omitting it cuts costs to reach lower market segments. Whereas the WRIOP handles traditional tasks such as parsing, classification, and distribution, the AIOP can also do complex lookups, network address and port translation, header compression, IPsec processing, address resolution, encapsulations, and NetFlow statistics. The LX2160A’s more-numerous ARM cores will handle those tasks when a program calls a DPAA2 function that would otherwise run on the AIOP.

Another omission is NXP’s PME (pattern-matching engine), which accelerates regular-expression (reg-ex) processing. Previously, this feature distinguished higher-end QorIQs from lower-end ones. The PME is nothing to brag about, however—its throughput is only 10Gbps—so the company dropped it from this design. Assigning reg-ex processing to a couple of the 2.2GHz A72 CPUs should deliver similar performance. Most of NXP’s customers don’t need reg-ex, and those that do generally have their own algorithms that don’t map well to a fixed-function engine like the PME.

The dual DDR4-3200 DRAM controllers are 64 bits wide and deliver 51.2GB/s of peak theoretical bandwidth, surpassing the T4240 and all other QorIQ processors. Although the LX2160A provides 72% more DRAM bandwidth per thread than the T4240, it has 25% less bandwidth per thread than the octa-core LS2088A. But the 8MB platform cache and a 2MB packet buffer should relieve pressure on the DRAM interfaces by reducing memory accesses. As usual, those interfaces have ECC protection, and the maximum memory capacity is 256GB.

100Gbps Networking and Switching

Like some other networking-oriented embedded processors announced lately, the LX2160A steps up to 100GbE and PCI Express Gen4. Faster Ethernet is a particularly big improvement for NXP, whose existing products top out at 10GbE. To implement these interfaces, the LX2160A has 24 high-speed serdes, each operating at up to 28Gbps (effectively

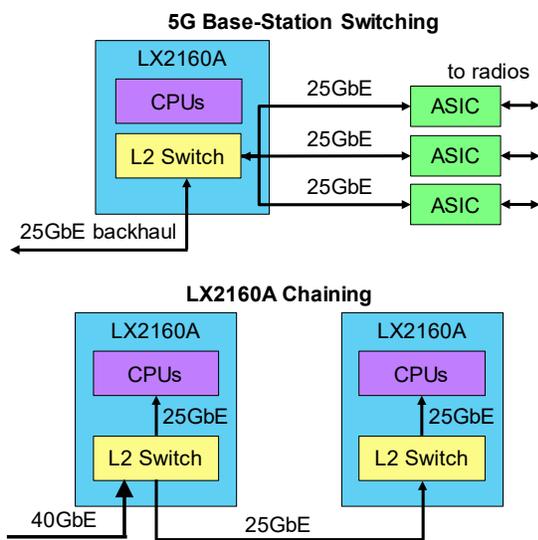


Figure 2. Two common examples of LX2160A switching. In the first case, the processor distributes incoming packets to multiple ASICs in a base station. In the second case, one processor separates the incoming traffic and distributes some of it to a chained processor.

25Gbps with the encoding overhead). Although the T4240 has 32 serdes, they operate at only 10Gbps.

The LX2160A can share its 24 lanes among the Ethernet, PCIe, and Serial ATA (SATA3) controllers in almost any combination. Ethernet is limited to 16 ports, which can include up to two interfaces operating at 100GbE, 50GbE, or 40GbE speeds, plus additional interfaces operating at 25GbE, 10GbE, 2.5GbE, or GbE speeds.

All of the Ethernet controllers connect to a 16-port Layer 2 switch, so any port can redirect packets to any other port. The chip also dedicates an ECC-protected 2MB “packet express buffer” to this switch, reducing DRAM accesses and switching latency. The switch is limited to 130Gbps—for example, it can switch four 25GbE and three 10GbE ports. Although some existing QorIQ models have 16-port switches, they’re limited to 8x10GbE plus 8xGbE, or a total of 88Gbps.

NXP says customers often use the switches built into some of its processors to receive packets through the numerous Ethernet ports, steer them to the CPUs for Internet Protocol (IP) forwarding, and dispatch them through a different port. To NXP, however, true switching means the packets traverse the processor without hitting the CPUs.

One example of true switching is an LX2160A connected to multiple ASICs that require high-speed communication. Since packets can flow through without CPU intervention, the switch latency is only about two microseconds. As Figure 2 shows, this configuration is typical of wireless base stations, and it replaces an external switch chip. It also simplifies the software, because it doesn’t require a separate device driver for an external Ethernet switch.

The internal switch can also chain multiple processors together for higher performance. If the incoming pipe is fat and the packets require more processing than one LX2160A can handle, the switch can divide the traffic using multi-tuple lookups. Redirecting some packets to another processor won’t burn any CPU cycles on the chip doing the switching, and the low switch latency has a negligible effect on the throughput.

Up to 24 PCIe Gen4 Lanes

PCIe Gen4, which is twice as fast as Gen3 (15.75Gbps versus 7.9Gbps), is virtually a requirement for high-performance embedded processors appearing in 2019. The LX2160A has six PCIe controllers; two of them support eight lanes and the others support up to four lanes.

Nominally, that’s a total of 32 lanes. But they share the same 24 serdes as the Ethernet controllers and four SATA3 controllers, so the actual limit is 24 lanes—and then only if the application requires no Ethernet or SATA. Most embedded processors share their serdes among their

high-speed serial interfaces, and some are more restrictive, so these limitations aren’t a competitive handicap.

The LX2160A has the usual low-speed I/O interfaces, including two USB3.0 controllers with integrated PHYs. One new feature is a dual CAN-FD interface. This flexible-data-rate extension to the Controller Area Network (CAN) standard can transfer larger frames at higher bit rates but remains compatible with the popular CAN2.0 standard. Introduced by Bosch, CAN-FD is intended mainly for microcontroller networks in motor vehicles but is also useful for industrial equipment and ADASs.

To increase wafer yields and reach lower market segments, NXP plans to offer two scaled-down versions of the LX2160A based on the same die. As Table 1 shows, the LX2120A will have 12 CPUs and the LX2080A will have 8. All three models share nearly all features, but their L2 caches will vary because of how the company is chopping these products. The 12-core chip disables two CPU pairs along with their L2 caches, leaving only 6MB of L2 cache. The octa-core chip instead removes one CPU from each of the eight pairs, leaving all 8MB of L2 cache. Thus, the low-end chip actually has more cache than the midrange one.

Competitors Need Faster Ethernet

Broadcom and Mellanox will beat the LX2160A to market with 100GbE. The former plans to qualify its BCM58808H chip for production this quarter, and we expect the latter’s BlueField to debut in 2Q18. These processors are more

| | LX2160A | LX2120A | LX2080A |
|-------------------------|---|----------------|---------------|
| CPUs | 16x Cortex-A72 | 12x Cortex-A72 | 8x Cortex-A72 |
| CPU Speed (max) | 2.2GHz | | |
| L2 Cache | 8MB L2 cache | 6MB L2 cache | 8MB L2 cache |
| L3 Cache | 8MB platform cache | | |
| Other Memory | 2MB packet express buffer | | |
| Accelerators | Crypto, compression, packet processing | | |
| IPSec Perf | 50Gbps IPSec | | |
| Compress Perf | 50Gbps compress + 50Gbps decompress (simultaneous) | | |
| DRAM Interface | 2x 64-bit DDR4-3200 with ECC | | |
| DRAM B/W | 51.2GB/s | | |
| Ethernet MACs | 2x up to 100GbE + 14x up to 25GbE | | |
| Ethernet Switch | 16-port 130Gbps Layer 2 switch | | |
| PCI Express | 6x PCIe Gen4 x24 lanes | | |
| Serial RapidIO | None | | |
| Serial ATA | 4x SATA3 | | |
| USB | 2x USB3.0 + PHY | | |
| Serdes Lanes | 24x 28Gbps serdes | | |
| Misc I/O | CAN-FD, SD/eMMC, DUART, I ² C, SPI, GPIO, JTAG | | |
| Other Features | Secure boot, TrustZone | | |
| IC Process | TSMC 16FFC | | |
| Package | 40mm FCPBGA-1,517 | | |
| Power | 30W TDP | 25W TDP* | 20W TDP* |
| List Price (1ku) | \$350* | \$300* | \$250* |
| Availability | Samples 1Q18 (est), production 2Q19* | | |

Table 1. NXP’s QorIQ LX2160A and derivatives. The main differences are their core counts and L2 caches. The company isn’t disabling any serdes or I/O controllers, so the only other differences are TDP and price. These chips are pin compatible with each other but not with any existing QorIQs. (Source: NXP, except *The Linley Group estimate)

specialized, however, primarily targeting storage systems and smart NICs (see [MPR 8/21/17](#), “Mellanox Accelerates BlueField SoC”).

For more-general networking and embedded applications, existing competitors include the Cavium Octeon TX family and Intel Xeon D family. Octeon TX is newer and uses a custom ARMv8-compatible CPU (see [MPR 5/16/16](#), “Octeon Expands Market With ARM”). Xeon D uses the older but still powerful Broadwell x86 CPU (see [MPR 12/21/15](#), “Intel Xeon D Targets Embedded”).

As Table 2 shows, the Octeon TX CN8360 matches the LX2160A’s core count and clock frequency. The NXP chip will lead in single-thread performance because Cortex-A72 is twice as powerful as Cavium’s ThunderX core, which is similar to Cortex-A53. Also, the LX2160A has more than twice the cache and 33% more DRAM bandwidth, so it will waste less time accessing external memory. And for networking, it’s definitely better appointed. The CN8360 is limited to 40GbE, lacks the newer 25GbE standard, has no internal Ethernet switching, and supports fewer of the low-speed Ethernet ports.

Having up to 24 PCIe Gen4 lanes, the LX2160A surpasses the CN8360’s PCIe connectivity in both speed and lane count. Surprisingly, it also beats Octeon TX in cryptography and compression acceleration—Cavium’s traditional

strengths. In addition, Octeon TX is the only family in this comparison that still uses 28nm planar transistors. But the CN8360 is shipping now, whereas the LX2160A is more than a year away.

LX2160A Runs Cooler Than Xeon D

The Intel Xeon D-1548 has only eight cores, but they’re dual threaded, so they match the other chips’ thread count. These Broadwell CPUs are among the industry’s most powerful and can surge to 2.6GHz in turbo mode when a program needs an extra burst of single-thread performance. Although Broadwell is more powerful than CortexA72, the performance gap narrows when using open-source compilers (see [MPR 10/31/16](#), “Adjusting SPEC CPU2006 Scores”). Using GCC, NXP estimates the LX2160A’s performance at 13.8 SPECint2006 and 155–160 SPECint_rate2006, but these estimates are based on simulations, not silicon.

The LX2160A has 29% more cache and 33% more memory bandwidth than the D-1548, negating two of Intel’s traditional strengths. The Xeon D has more PCIe lanes (32), but they implement the slower Gen2 or Gen3 standards, as is natural for a chip launched in 2015. The D-1548 is also at a networking disadvantage even though it’s one of Intel’s most integrated devices. By itself, it provides only two 10GbE interfaces. To get more, system designers can attach an XXV710-AM2 (Fortville-25) Ethernet chip, which adds two ports of up to 25GbE and costs \$133. We estimate it also adds 7W TDP.

High-performance networking systems need packet acceleration, and the D-1548 suffers in that respect as well. NXP estimates the LX2160A will process IPsec streams at 50Gbps, which we figure is 32% faster than the Xeon can. Intel’s DH8900 (Coletto Creek) south-bridge products offer faster QuickAssist crypto and compression acceleration, but they add another 17–20W TDP and \$214–\$308 in list price. A newer hub is the C62x-series (Lewisburg), which has up to four 10GbE interfaces and even faster hardware acceleration, but those chips increase power by 21–29W and list price by \$132–\$329. In return, their improved QuickAssist engines deliver much greater throughput than Coletto Creek’s: as much as 100Gbps for bulk crypto and data compression. But because the D-1548 lacks a DMI3 interface, it must sacrifice some PCIe lanes to attach the new hub.

For storage, the D-1548 offers more SATA3 and USB interfaces than the LX2160A. Its power consumption is 50% greater than NXP’s estimate for the QorIQ chip, though. The 15W difference would accommodate an external SATA or USB controller that would even the score.

In July, Intel quietly launched five Xeon D models in the new Network Series, designated by an “N” suffix (e.g., D-1553N). They have four,

| | NXP QorIQ LX2160A | Cavium Octeon TX CN8360 | Intel Xeon D 1548 |
|---------------------------------------|--|--|-----------------------------|
| CPUs, Threads | 16x Cortex-A72, 16 threads | 16x ThunderX, 16 threads | 8x Broadwell, 16 threads |
| CPU Speed (max) | 2.2GHz | 2.2GHz | 2.0GHz |
| L2 Cache | 8MB L2 | 8MB L2 | 2MB L2 |
| L3 Cache | 8MB platform | None | 12MB L3 |
| Other Memory | 2MB pkt buffer | None | None |
| Accelerators | Crypto, compress, packet processing | Crypto, compress, packet processing | None |
| IPSec Perf | 50Gbps IPSec | 40Gbps IPSec* | 38Gbps IPSec |
| Compress + Decompress Perf | 50Gbps + 50Gbps | 40Gbps + 40Gbps | Not applicable |
| DRAM Interface | 2x 64-bit DDR4-3200 | 2x 64-bit DDR4-2400 | 2x 64-bit DDR4-2400 |
| DRAM B/W | 51.2GB/s | 38.4GB/s | 38.4GB/s |
| Ethernet MACs | 2x 100GbE + 14x 25GbE | 3x 40GbE, 12x 10GbE | 2x 10GbE |
| Ethernet Switch | 16-port 130Gbps | None | None |
| PCI Express | PCIe Gen4 x24 | PCIe Gen3 x22 | PCIe Gen3 x24 |
| Serial ATA | 4x SATA3 | 6x SATA3 | 6x SATA3 |
| USB | 2x USB3 + PHY | 2x USB3 | 4x USB3 |
| Other Features | Secure boot, TrustZone | vNIC subsystem | None |
| IC Process | 16nm FinFET-C | 28nm HKMG | 14nm FinFET |
| Power (TDP) | 30W | 40W* | 45W |
| List Price (1ku) | \$350* | \$300* | \$675 |
| Production | 2Q19* | 3Q17 | 4Q15 |

Table 2. Comparison of three embedded processors for networking. All these products implement ECC on their DRAM interfaces. The Xeon can surge to 2.6GHz in turbo mode but requires external chips for additional Ethernet ports and acceleration. (Source: vendors, except *The Linley Group estimate)

six, or eight Broadwell CPUs operating at base clock speeds of 1.6–2.3GHz. Unlike their brethren, some have four 10GbE interfaces instead of two, and all of them add QuickAssist acceleration for cryptography and compression. These accelerators appear to be similar to Coletto Creek's. But Intel markets the N-series Xeons as server chips and declines to certify them for the extended availability and temperatures that embedded customers may require. In general, Xeon D is best suited to customers that need x86 compatibility and can live with the lower integration and higher power consumption.

Although the LX2160A looks strong compared with existing competitors, we expect a different landscape by the time it reaches production in 2019. Cavium will likely upgrade the Octeon TX family to the newer ThunderX2 CPU, PCIe Gen4, and 100GbE. Intel will probably upgrade Xeon D to the refreshed Skylake or an even newer core—perhaps the 14nm+ Coffee Lake or even the 10nm Cannonlake.

Winning High-End Sockets

NXP is stepping up its game in several ways with the LX2160A. By doubling the number of Cortex-A72 CPUs that today's QorIQ chips offer and raising their clock speed, the new device delivers more than twice the throughput. Although it still fails to match the thread count of the five-year-old T4240, it will deliver more single-thread and aggregate performance—twice the throughput on SPECint2006, by NXP's estimate. Moving to 16nm FinFET technology enables these gains while consuming at least 33% less power than the T4240 and about the same power as the octa-core LS2088A.

Equally important are the new 28Gbps serdes, which allow Ethernet connectivity up to 100GbE and PCIe up to Gen4. Moreover, the Ethernet MACs connect to a 16-port Layer 2 switch with a dedicated 2MB buffer that can sustain packet traffic at up to 130Gbps without CPU intervention. This high-speed I/O and packet switching positions the LX2160A for next-generation networking equipment and cellular infrastructure. It will hit the market at about the same time that 5G base stations begin commercial deployments in 2019.

To keep up with the faster packet I/O, NXP has upgraded the LX2160A's accelerators to perform bulk crypto and decompression at 50Gbps. These are further examples of a balanced design. The only notable missing feature is the PME reg-ex engine, which we suspect few customers are using in current QorIQs.

Before the LX2160A sees production, however, competitors have a chance to catch up. We expect Broadcom, Cavium, and Mellanox to be shipping processors with similar features in 2018 and 2019, and Intel's Xeon D is due for upgrades as well. Any slippage in NXP's schedule—always a risk when fabricating a large new design in an unfamiliar IC process—would dampen the chip's debut. Qualcomm's delayed acquisition of NXP may add distraction.

If everything goes smoothly, the LX2160A is poised to compete effectively with its likely counterparts. It extends the upper range of the QorIQ family in multiple dimensions, giving customers a reason to stay with NXP instead of changing vendors in search of higher performance. It also improves performance without inflating power consumption—always a welcome achievement. ♦

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