Overview

This paper examines the total cost of ownership (TCO) of servers using the Qualcomm Centriq 2400 system-on-chip (SoC) running the Armv8 instruction set architecture (ISA). TIRIAS Research compares an estimated three-year TCO for servers based on the Qualcomm Centriq 2452 SoC against a mainstream x86-based server using Intel Xeon Gold 5120 processors. The performance basis for this comparison is the Redis in-memory database.

Other papers in this series compare the three-year TCO of HHVM and estimated SPECrate2017_int_base. Methodology and assumptions for the series are detailed in the companion report *Qualcomm Centriq 2400 Server TCO: Methodology & Assumptions*.

Purpose

The target audiences for this TCO comparison are social media, Software-as-a-Service (SaaS) and Platform-as-a-Service (PaaS) providers. While databases are as old as computing, in-memory databases have evolved over the past two decades as enterprise datacenter architecture evolved into cloud architecture and as system memory cost per bit shrank and bit density grew.

Application

In-memory databases lower transaction latency by eliminating lengthy disk access latencies. Social media and ecommerce sites like Airbnb, Twitter, Flickr, Weibo, Pinterest, and Snapchat are continually competing to improve service. They deploy open source key value store systems like Redis to implement reliable, distributed, low-latency content caching across their fleets.

Figure 1 shows that a 12-kilowatt (kW) rack full of two single-socket Qualcomm Centriq 2452 motherboards per chassis (36 chassis total) should have an estimated 2.2x performance advantage and a 1.4x performance per dollar advantage over the 29 dual-socket Intel Xeon Gold 5120 chassis that consume the same power.

Figure 1: Full 12kW Rack Performance Comparison

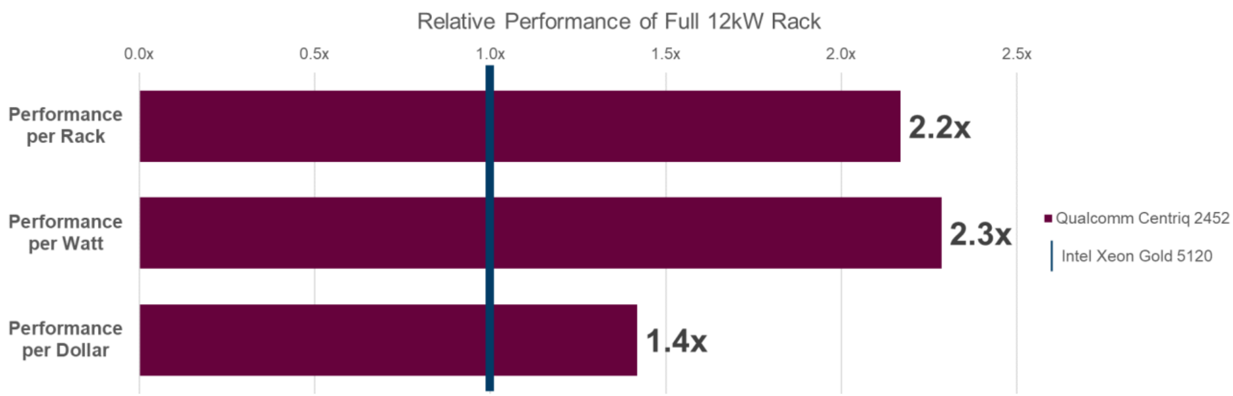

Sources: Qualcomm Datacenter Technologies & TIRIAS Research (See Appendix for notes regarding all Figure & Table sources)

TIRIAS Research observes that servers based on the Qualcomm Centriq 2452 SoC, using two single-socket Qualcomm Centriq 2452 motherboards per chassis, should exceed the Intel Xeon Gold 5120 processor dual-socket motherboard's Redis benchmark suite performance by 2.2x at 0.9x the power consumption (Figure 2) across the real-world workloads represented in Table 1.

Software

Redis includes the "redis-benchmark" utility in its source code distribution. Redis-benchmark simulates clients running database commands and logs responses to those commands. The common "Set" (write a key / value pair) and "Get" (use a key to retrieve a value) commands were benchmarked. Qualcomm Datacenter Technologies (QDT) used Redis server version 3.2.8 default server configuration for this analysis. Table 1 lists the parameter settings.

Table 1: Redis-Benchmark Parameter Settings

Fixed Parameters	Value	Default	Description
Key Space	10,000	1	Range for random key generation
Data Size	10 Bytes	2	Payload size of the values stored in or requested from memory
Keep Alive	Persistent	Persistent	Persistent or stateless (reconnect) sessions
Pipeline	1 & 100	1	Request queue depth; '1' means no queue, wait for response
Memory Allocator	Tcmalloc	–	Threaded cache memory allocation for lockless, high concurrency
Client	50	50	Number of simulated simultaneous client sessions
Requests	7.5M & 75M	100,000	Total requests generated by combined number of simulated clients

Sources: Qualcomm Datacenter Technolgies & TIRIAS Research

Key Space, Data Size, and Keep Alive settings are based on cloud customer queries.

QDT varied two parameters, also based on interest from potential cloud customers:

- Requests were run at settings of 7,500,000 and 75,000,000
- Pipeline was run at a depth setting of 1 (no queue) and 100 (100 outstanding requests)

There is another parameter set by database operators: "instances" is a Redis runtime parameter used to set the number of copies of Redis simultaneously running on a single server. Redis is a single-thread database—one instance runs on one hardware thread (a logical core). Multiple instances are needed to fully utilize a modern multi-threaded, multi-core processor. There is no direct communication between Redis instances—Redis instances serve remote queries.

A Redis best practice is to set instances to the number of cores in a processor for best efficiency and utilization. QDT's benchmarking confirmed this practice. TIRIAS Research chose to use QDT's redis-benchmark results for Qualcomm Centriq 2452 instances set to 46 and Intel Xeon Gold 5120 instances set to 56. This asymmetry maximized the efficiency of both products for this TCO comparison, without compromise (more detail in Appendix).

Hardware

QDT benchmarked one single-socket Centriq 2452 motherboard against a single Intel Xeon Scalable Gold 5120 dual-socket motherboard. The Intel Xeon Gold 5120 processor and dual-socket motherboard are representative of solutions often bought by cloud customers to run Redis.

Results

Figure 2: Redis Performance

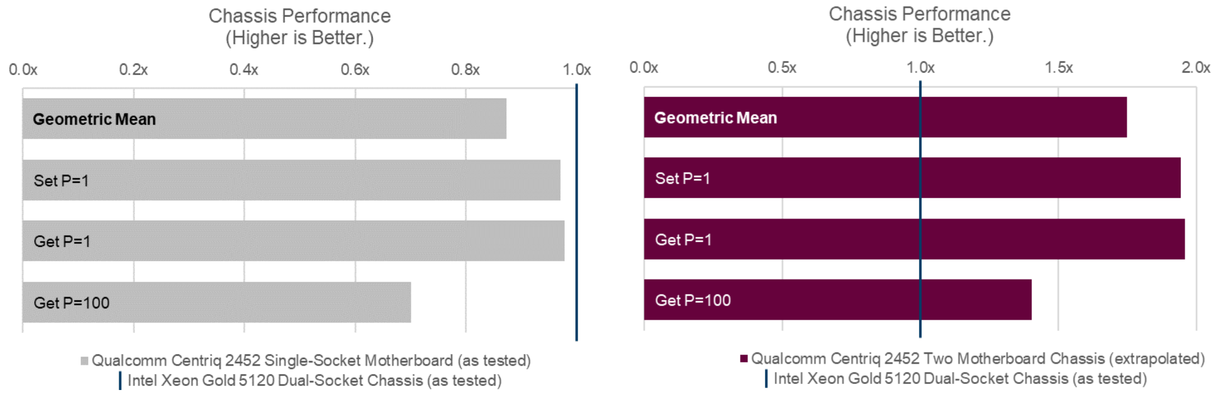

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

Figure 3: Redis Power Consumption

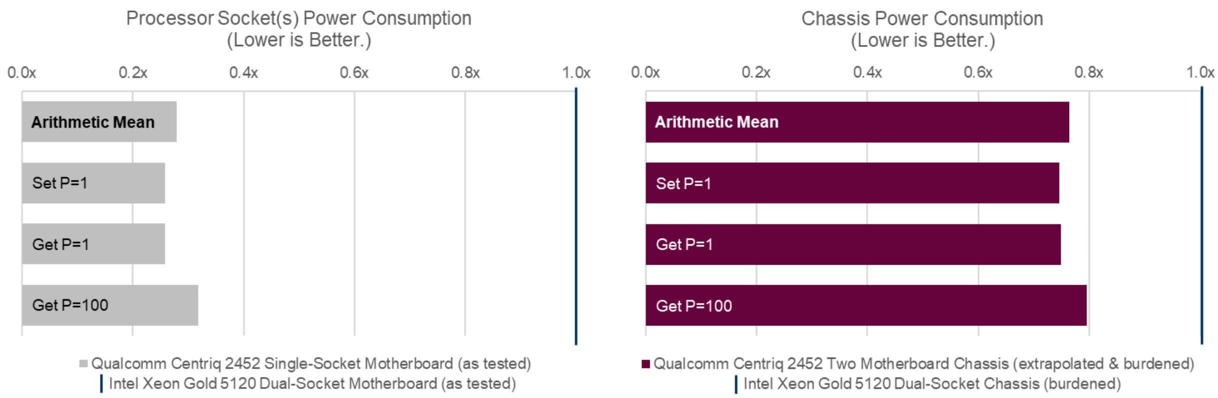

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

Figure 4: Redis Performance per Watt

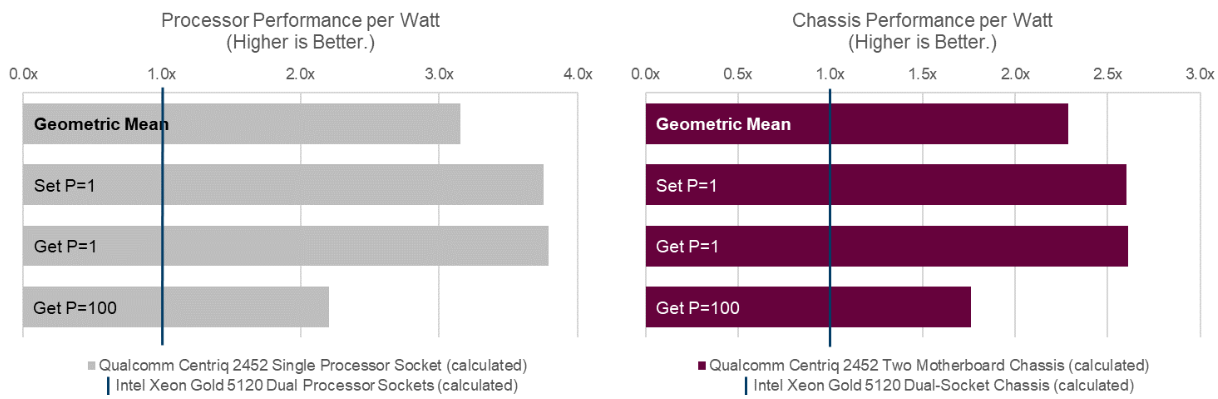

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

See Appendix for a summary of the measured data.

Redis is designed to respond to remote queries. Redis-benchmark client was run on a dedicated server, connected to the chassis under test by in-lab switches. Network connectivity is described in the companion report *Qualcomm Centriq 2400 Server TCO: Methodology & Assumptions*.

QDT measured the network traffic during its Qualcomm Centriq 2452 testing. For the high traffic rates generated while pipeline was set to 100; only about 4% of the client NIC's Ethernet bandwidth was used. Network traffic will increase with payload size (Table 1 shows defaults and tested size). The Intel Xeon Gold 5120 system handled roughly twice the throughput of the Qualcomm Centriq 2452 at peak, and so its NIC is estimated to have experienced less than 10% saturation. TIRIAS Research and QDT believe the network did not impose bandwidth or latency constraints on redis-benchmark results. If any latencies due to network saturation had occurred, those latencies would have been similar for both systems tested.

To show a more accurate TCO comparison for performance per watt consumed, this analysis burdened the two single-socket to one dual-socket processor-based power consumption comparison with estimates for complete system-level power consumption.

TIRIAS Research linearly extrapolated QDT's single-socket measurements to estimate performance and power consumption for two single-socket Qualcomm Centriq 2452 motherboards in a chassis.

TCO Analysis

TIRIAS Research reduces TCO analysis to the smallest set of variables that highlight useful differences between products. For server processor TCO calculations, the number of constants that apply to both configurations in the comparison were simplified. Because processors cannot run workloads without the rest of a functioning chassis, the measured performance, power consumption, and hardware bill of materials costs were extrapolated and burdened to obtain capital expense (Capex) and operating expense (Opex) estimates (Table 2).

These tests do not extend the TCO estimate beyond evaluating simple rack-scale metrics, because identical switches, power distribution, cabling, and rack costs would be used for both configurations.

Table 2 shows a rack-level extrapolation of a three-year TCO based on a 12kW rack power supply. 1kW was subtracted for two (redundant) top of rack (TOR) switches, leaving 11kW available to power servers. The redis-benchmark results for individual server chassis were multiplied by the number of server chassis that can be run within 11kW.

Based on redis-benchmark power measurements, a full rack of 36 Qualcomm Centriq 2452 chassis containing two single-socket motherboards fits within a 12kW rack power budget. Only 29 dual-socket Intel Xeon Gold 5120 chassis fit in the same power budget. The rack full of Qualcomm Centriq 2452 servers should show 2.2x better performance of 29 dual-socket Intel Xeon Gold 5120 chassis at only 1.5x the price and 0.9x the power consumption. Buying fewer servers to meet performance goals may also lower IT software and hardware management costs.

Table 2: TCO Based on Redis

TCO		Two Motherboards† Single-Socket Qualcomm Centriq 2452	One Motherboard†† Dual-Socket Intel Xeon Gold 5120	Difference	Maximum Chassis within 11kW Power Budget 36 Chassis Qualcomm Centriq 2452	29 Chassis Intel Xeon Gold 5120	Difference
Capex	Processor(s)	$2,746	$3,110	0.9x	$98,856	$90,190	1.1x
	Memory	$4,224	$1,787	2.4x	$152,064	$51,825	2.9x
	Motherboard	$935	$820	1.1x	$33,660	$23,780	1.4x
	Storage	$154	$77	2.0x	$5,530	$2,227	2.5x
	NIC	$360	$180	2.0x	$12,960	$5,220	2.5x
	Rest of Server	$290	$290	1.0x	$10,440	$8,410	1.2x
	Infra Power Cost (3yr)	$994	$1,301	0.8x	$35,786	$37,716	0.9x
Opex	Power (3yr)	$630	$825	0.8x	$22,697	$23,921	0.9x
	Total	**$10,333**	**$8,389**	**1.2x**	**$371,993**	**$243,289**	**1.5x**
	Performance	**14,513,419**	**8,307,339**	**1.7x**	**522,483,069**	**240,912,828**	**2.2x**

†extrapolated ††as tested

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

Note that 38 two motherboard Qualcomm Centriq 2452 1U servers fit within an 11kW power budget using redis-benchmark power consumption, but only 36 1U servers fit in a typical 42U rack configuration.

Figure 5: Compute Density & Power Opex within 12kW Rack

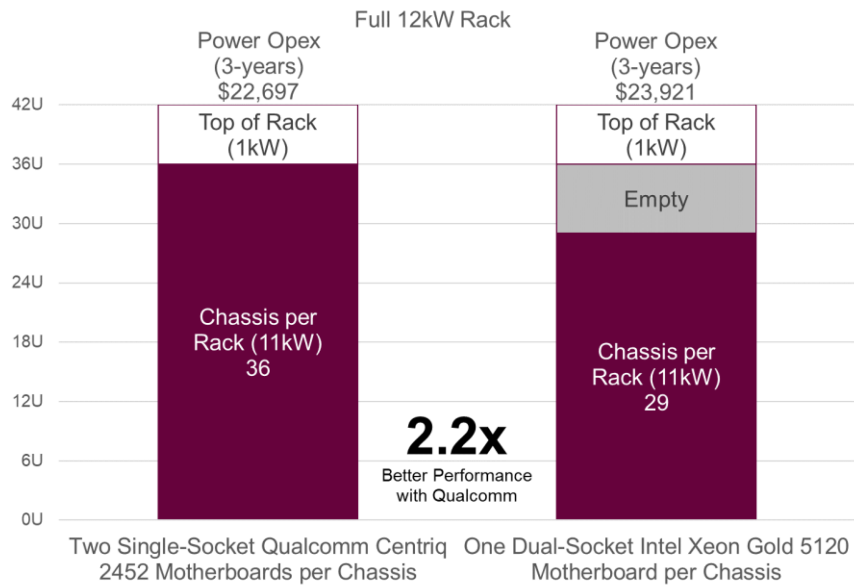

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

Conclusion

Using redis-benchmark results to generate a rack-level three-year TCO comparison of the Qualcomm Centriq 2452 SoC against the Intel Xeon Gold 5120 processor allows a fair comparison of the selected modern cloud workloads. The Qualcomm Centriq 2452 SoC shows 1.4x greater estimated performance per dollar than Intel Xeon Scalable in this comparison, based on its redis-benchmark performance and power consumption. However, actual TCO will vary widely in practical use.

QDT chose to use the same amount of memory on each Qualcomm Centriq 2452 motherboard as was tested on the dual-socket Intel Xeon Scalable motherboard, and QDT used marginally faster memory on its own motherboard. The result is a significant price premium to the Intel Xeon Scalable configurations, with 2.4x the memory cost per two motherboard chassis and 2.9x the total memory cost at rack-level. This choice strongly favors Intel's performance per dollar results. Even with this self-imposed handicap, Qualcomm Centriq 2452 server configurations show 1.4x better estimated performance per dollar than dual-socket Intel Xeon Gold 5120 server configurations. Twice the memory per socket should show measurable advantages for real-world in-memory database workloads.

In a nutshell, 2.2x rack-level performance advantage for Redis in-memory databases can help datacenters achieve higher density, so it will take longer to fill a datacenter to capacity within Capex and Opex constraints. Similarly, a smaller datacenter can be designed to meet a specific level of compute performance.

Companies deploying in-memory database workloads should consider benchmarking their workloads on Qualcomm Centriq 2400 Armv8-based servers.

Appendix

Figure & Table Sources

Unless otherwise noted, all Figures and Tables are based on Qualcomm Datacenter Technologies (QDT) benchmark measurements, Qualcomm Centriq 2400 SoC and motherboard specifications, public competitive processor, motherboard, and chassis specifications, and TIRIAS Research calculations and formatting.

Table A1: Summary of Measured Data

As tested	Qualcomm Centriq 2452 Single-Socket Motherboard		Intel Xeon Gold 5120 Dual-Socket Motherboard	
	Chassis Performance	Processor Socket Power Consumption	Chassis Performance	Processor Sockets Power Consumption
Set P=1	3,478,470	45.2	3,582,254	174.7
Get P=1	3,474,360	43.9	3,548,964	169.9
Get P=100	31,619,580	60.6	45,094,940	190.7
Geometric Mean	**7,256,709**	**n/a**	**8,307,339**	**n/a**
Arithmetic Mean	**n/a**	**49.9**	**n/a**	**178.4**

Source: Qualcomm Datacenter Technologies

Performance Based on Number of Instances

Cloud datacenter managers will default to setting each of the two benchmarked servers to the number of hardware threads in each server, so that each server achieves its highest Redis throughput.

With pipeline depth set to 1, *i.e.*, no pipeline (P=1) (Figure A1), each client request must be completed and sent back to the client before an instance can process its next client request. Both systems display virtually identical behavior as throughput converges to a function of the latency of the client server in generating each request, not on any technical limitation of the servers.

Figure A1: Throughput with No Pipeline

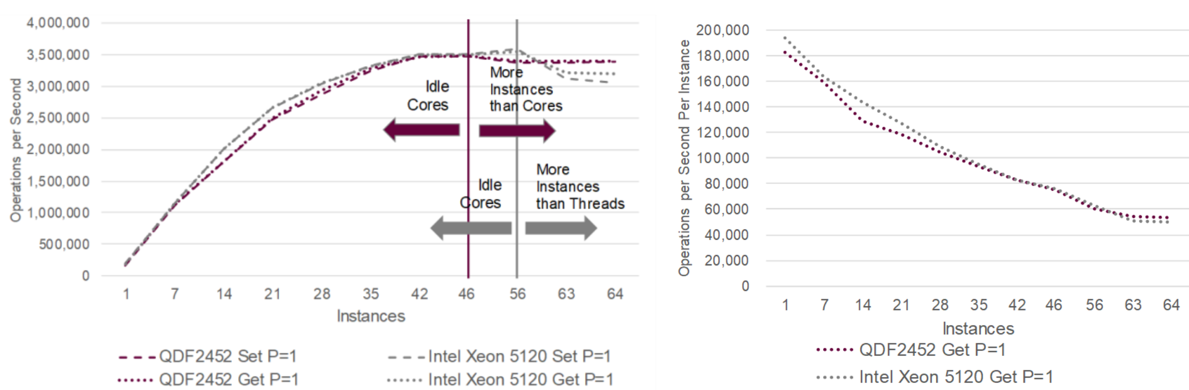

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

With pipeline set to 100 (Figure A2), hardware thread performance and number of hardware threads give Intel Xeon Gold 5120 a raw throughput edge. However, Intel Xeon Gold 5120 performance per core drops substantially after the number of instances exceeds the number of processor cores (half the number of hardware threads), while processor throughput stays relatively flat after that. This effect is due to Intel Hyperthreading. As the number of instances exceeds the number of cores, instances are scheduled on the second thread for each core, resulting in a performance drop per thread on that core.

With no pipelining, Intel Xeon Gold 5120 power consumption running 56 instances (one per hardware thread) draws only 5-7% more power than with 28 instances (number of cores). With pipeline depth set to 100, power consumption is only 11% higher at 56 instances than at 28 instances.

The Intel Xeon Gold 5120 reached peak throughput when the number of instances was set to the number of hardware threads (dual sockets, each with 14 cores and with Hyperthreading turned on equals 56 hardware threads). Power consumption continues to increase slightly when the number of instances exceeds the number of threads. With instances equal to the number of threads, Intel Xeon Gold 5120 loses about 4% efficiency with no pipelining and about 8% with pipeline depth set to 100. These processor efficiency (performance per watt) losses are very small compared to the total power consumption of the entire server.

Figure A2: Throughput with Pipeline

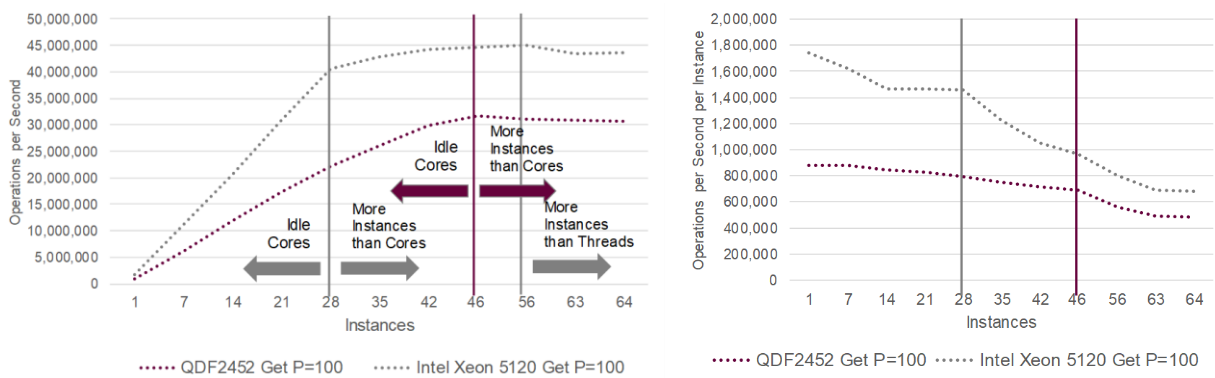

Sources: Qualcomm Datacenter Technologies & TIRIAS Research

Qualcomm Centriq 2452 reached peak throughput, power consumption, and efficiency when the number of instances was set to equal the number of hardware threads (which is 46, the number of cores), *i.e.* the benchmark was running at one instance per core. Because a single-socket Qualcomm Centriq 2400 motherboard has 64% more cores than a dual-socket Intel Xeon Gold 5120 motherboard, it can run many more Redis instances at full core performance. Beyond 46 cores, context switching reduces performance per instance at roughly the same rate as Intel Xeon Gold 5120.

Qualcomm Centriq 2452 reaches efficiently with only about a third of its cores running instances, while Intel Xeon Gold 5120 has a shallower, almost linear efficiency curve, reaching peak efficiency half-way between its core count and its thread count. Also, Qualcomm Centriq 2452 consumes a fraction of the power of the Intel Xeon Gold 5120 (Figure A3).

The result is that datacenter managers experience the best throughput out of Intel Xeon Scalable architectures with the number of instances set to the number of hardware threads, without sacrificing much efficiency. TIRIAS Research speculates that both architectures reach their pipeline-enabled performance limit as a function of hardware thread performance and memory system performance.

Figure A3: Power (Left) & Efficiency (Right)

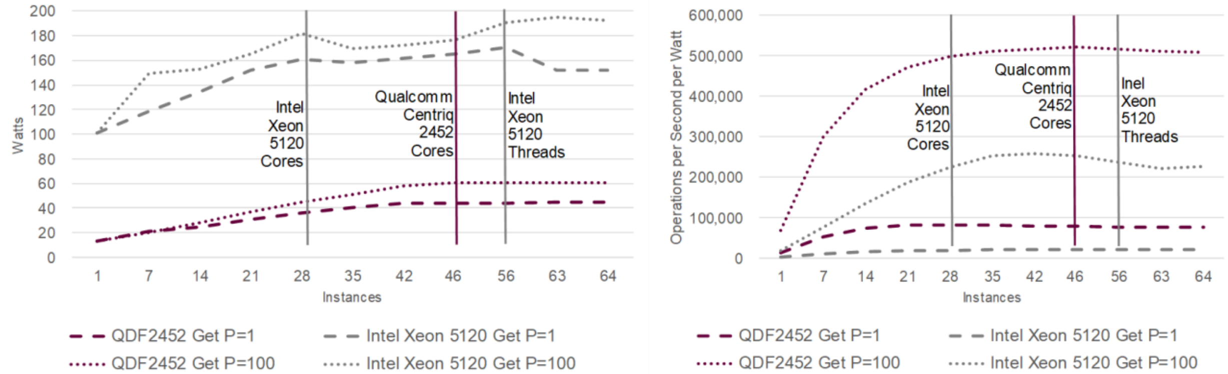

Sources: Qualcomm Datacenter Technologies & TIRIAS Research