TSMC optimizes workloads with AMD EPYC™ CPUs

Compelling performance and validated compatibility for TSMC's global microchip manufacturing business

AMD EPYC

As the leading foundry of semiconductor wafers in the world, TSMC has a huge need for data center infrastructure. The company, headquartered in Hsinchu, Taiwan, not only fabricates the wafers but is also engaged heavily in research and development. Its over 50,000 employees worldwide require strong IT services to communicate and collaborate seamlessly. So when the AMD EPYC™ processor became available, it had clear potential to provide productivity and cost of ownership benefits across all their teams.

"We first introduced AMD EPYC processors into the general workload," explains Simon Wang, Director of Infrastructure and Communication Services Division at TSMC, "and they are being deployed with our research and development team." TSMC was looking for a server solution optimized to implement a hyper-converged infrastructure (HCI) design where all three components – compute, storage and networking – could be provided by the same underlying hardware. Although TSMC is the company that manufactures AMD's technology-leading 7nm products, which includes the groundbreaking AMD EPYC, Ryzen™ and Ryzen™ Threadripper™ processors, as well as the AMD Radeon™, Radeon™ Pro, and AMD Instinct™ GPUs, it didn't mean that AMD EPYC processors were automatically going to be the preferred selection.

"Performance-wise, we had no doubts. We tested with key applications and performance was excellent."

*Simon Wang,
Director of Infrastructure and Communication Services Division at TSMC*

Wide platform compatibility

The 7nm technology used by AMD's latest products is one of the most recent in a long line of revolutionary innovations from TSMC, and the technology itself marked the first time that the world's most-advanced semiconductor technology was made available broadly to all semiconductor companies and market segments. Keeping the constant flow of innovation and delivery requires very efficient IT provisioning. TSMC's infrastructure team is large, with members across several sites in Taiwan, Europe, Japan, mainland China, and the United States. Switching to a new server platform seamlessly requires significant planning and testing effort.

TSMC has to deliver and balance server core density within their data center environment. "Our existing data centers are limited by space and power," says Wang. "We designed and implemented the virtualization of the computing, storage and network, and all of that in our standard template of service architecture. I need to find products that can be directly introduced into this standard template, and then deploy them quickly across our data centers, without redesigning the architecture. We have two VM [ed. virtual machine] configurations for our workloads depending on their complexity." But running lot of VMs on one server is not optimal. "If I had my server supporting 20 machines, then when the server goes down, 20 machines will take the hit. That is why we have two types of VM server configurations to reduce the risk.

In addition to virtualization, we also have other software. For instance, there is path alternation protocol, PAP, that we use for controlling the communication between production machines and the virtual servers. When a server goes down, PAP automatically connects the machines to another VM server in the same cluster."

AMD + TSMC CASE STUDY

AMD

"Compatibility was the key factor before we started the large-scale deployment [of EPYC]," explains Wang. "Performance-wise, we had no doubts. We tested [EPYC] with key applications and performance was excellent. For manufacturing and the general workloads for support functions, there has been no issue with compatibility." The company's testing showed full compatibility with multiple Linux distributions, Windows system, and storage virtualization. "We leverage a lot of virtualization technology, and we already tested that. We're doing a pilot launch with some storage virtualization products, and they also offer certified AMD EPYC™ products."

Reduced total cost of ownership with AMD

For the general workload rollout, TSMC was partnering with HPE. "We had thorough discussions with HPE," explains Wang, "and the AMD product was readily available to use directly on our existing architecture, reducing the cycle time. Otherwise, we need to do more testing, and that will significantly increase the time to deployment." TSMC opted for the HPE DL325 G10 platform running the 2nd Generation AMD EPYC 7702P CPU with 64 cores, a base 2GHz clock frequency and up to 3.35GHz boost clock.

"Two more data centers will be implemented in 2021," explains Wang. "One in Tainan, the other in Hsinchu." For general workloads, memory density is a critical aspect. The results were impressive with the AMD EPYC processors TSMC chose. Fewer resources were required for each virtual instance in the HCI implementation, while achieving better performance. There were cost savings from having fewer physical servers due to the per-processor core density, reducing management and power costs as well as the space required in the data center.

The reduced number of servers will provide increased room to expand in the data center as the company grows further. Switching from dual to single socket also meant hardware, operating and software cost savings as well as reduced power consumption. All AMD EPYC processors offer 128 PCI Express lanes per CPU, so there is plenty of bandwidth available for speed optimized NVMe storage, with each HPE server equipped with ten hard drive slots.

"For our general workloads, memory density is a clear advantage of AMD EPYC processors."

Simon Wang, Director of Infrastructure and Communication Services Division at TSMC

Expanding deployment to manufacturing

TSMC now is expanding its AMD EPYC processor rollout across the company. "We're now introducing EPYC servers into our manufacturing teams," says Wang. The new HF high-frequency AMD EPYC CPU range is a focus of interest, specifically the 7F72 with 24 cores and a base 3.2GHz clock frequency. "That's the CPU we're considering for R&D, because of the high clock rate. This team doesn't necessarily need more sockets or cores. For R&D, if we use two sockets, that might create unexpected effects, because it means one CPU needs to communicate with another CPU, which will create overheads. So, for R&D, we choose a one-socket CPU and the high clock-rate will be an important advantage."

A one-socket AMD EPYC processor delivers 128 PCIe4 lanes, so there is no sacrifice of IO opting for a single-CPU server, plus one-socket EPYC CPU-powered servers can support a very large memory footprint.

"For automation with the machinery inside our fab, each machine needs to have one x86 server to control the operation speed and provision of water, electricity and gas, or power consumption," adds Wang. "These machines are very costly. They might cost billions of dollars, but the servers that control them are much cheaper. I need to make sure that we have high availability in case one rack is down, then we can use another rack to support the machine. With a standard building block, I can generate about 1,000 virtual machines, which can control 1,000 fab tools in our cleanroom." This will mean a huge cost saving without sacrificing failover redundancy or reliability.

"Because of TSMC's leadership in the world," concludes Wang, "our IT department needs to continuously leverage the latest applicable technology to support the company driving the world forward."

WANT TO LEARN HOW AMD EPYC™ PROCESSORS MIGHT WORK FOR YOU?
Sign up to receive our data center content
amd.com/epycsignup