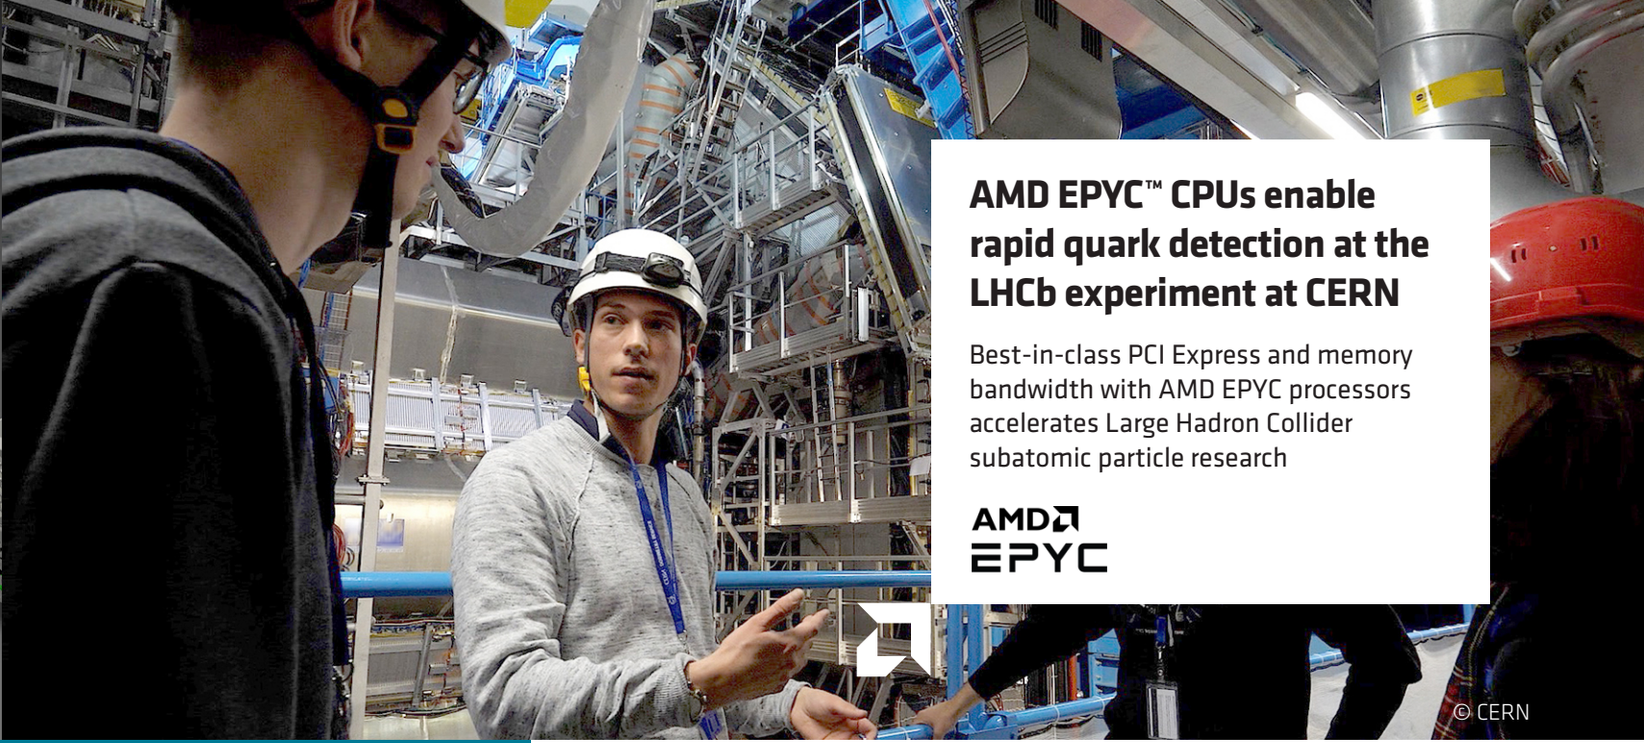

AMD EPYC™ CPUs enable rapid quark detection at the LHCb experiment at CERN

Best-in-class PCI Express and memory bandwidth with AMD EPYC processors accelerates Large Hadron Collider subatomic particle research

PARTNER

CERN

INDUSTRY

Particle physics research

CHALLENGES

Collecting huge throughput of data from the detectors on the Large Hadron Collider

SOLUTION

Deploy AMD EPYC™ 7742 processors in detector data collection servers

RESULTS

More compact systems with cost savings, allowing more money to be spent on science

AMD TECHNOLOGY AT A GLANCE

AMD EPYC 7742 processor with 64 cores

PARTNER

Gigabyte

Science is all about testing theories against experimental data. But when your test apparatus consists of the particle detectors of the Large Hadron Collider (LHC) at CERN, the raw collision data arrives at the enormous rate of 40TB every second. Much of this collision data will not be relevant, so job number one is sifting through the information as it arrives and pulling out the relevant results that have the highest probability of providing critical insight. This is a hugely taxing high-performance computing task, and it requires the most capable server hardware available.

CERN has been preparing for the upcoming LHC restart in 2021 for a while and has been looking for a hardware platform that is able to cope with the requirement of handling immense data throughput. "You need to save data from every server to every server," explains Niko Neufeld, Project Leader, LHCb Online Computing, CERN. "It's like folding a matrix. You solve a little puzzle, and you have to do this for all the data as it comes in. It is a messy I/O challenge." This led Neufeld to the 2nd Gen AMD EPYC™ processor for its ability to deliver the processing performance, memory speed and bandwidth that his experiment requires.

"With the AMD EPYC CPU, we've been able to show more than a terabit per second of data flowing from the servers sustained over days. Achieving this on a single server rather than requiring a supercomputer—as was the case in the past—is a significant advance."

Niko Neufeld, Project Leader, LHCb Online Computing, CERN

Piecing together a beautiful jigsaw

The experiment that Neufeld is working on is called LHCb. "The B stands for beauty," explains Neufeld. "This is one of the six types of quarks in nature. It's one of the primary targets of the studies we do here." A quark is a type of elementary particle and a fundamental constituent of matter. Quarks combine to form composite particles called hadrons, the most stable of which are protons and neutrons. The LHCb experiment is an investigation into what happened just after the Big Bang that allowed matter to survive and build the universe we know today. Observing the behavior of beauty quarks as they decay helps explain why the universe is chiefly made of matter rather than antimatter. Neufeld adds, "with beauty quarks, the difference is more pronounced than with other quarks."

The beauty quarks are created in the experiment by colliding hadrons in the LHCb detector, but they're so short-lived that the data must be collected quickly to detect their behavior. This raw flow of data is first handled by custom FPGA cards that do the initial interpretation. "Every server is mapped to a geographic piece of the detector," says Neufeld. "You slice up the detector into pieces, and then every server is connected point to point to part of this detector. But then you need to get all the data pieces together in a single location because only then can you do a meaningful calculation on this stuff."

"No single server can do the processing, so you need an all-to-all exchange," continues Neufeld. This activity requires fast processing of the data, high-bandwidth access to lots of memory, and very rapid connections to the I/O devices that attach the servers to each other. For the latter, LHCb uses four Mellanox 200Gbit InfiniBand adapters per server, which need to operate at full capacity via the fastest possible interconnect. "You have data coming from all different directions and that's being handled by different servers that then have to put a jigsaw together."

Ultra-fast I/O and memory with AMD EPYC™ processors

LHCb had prior history with AMD Opteron™ processors, and Neufeld had already been testing an installation of servers based on 1st Gen AMD EPYC processors for the final stage of the workflow, where the data from the collectors is packaged for use by the various research tiers associated with the LHCb experiment. The positive experience from this trial and the improved specification offered by the 2nd Gen AMD EPYC CPU made it an attractive solution for the even more intensive collection stage.

Although the AMD EPYC CPU's high core count paid dividends for data processing, the support for 128 PCI Express 4.0 lanes was the standout feature, enabling the four Mellanox networking cards in each server to run without bottlenecks. "With the AMD EPYC CPU, we've been able to show more than a terabit per second of data flowing from the servers sustained over days. Achieving this on a single server rather than requiring a supercomputer—as was the case in the past—is a significant advance," says Neufeld.

"The total system memory pool is also very important," continues Neufeld. "You have so much data coming in. You cannot stream the data from the FPGA card directly to the network card, so it has to go to main memory and come back. You need a massive amount of memory." LHCb is using 512GB memory on each of its servers. "The Rome platform meets this need because you have many memory channels and you have support for very fast bandwidth. It's a very balanced platform." With up to 8TB of 3,200MHz DDR4, the 2nd Gen AMD EPYC CPU delivers perfectly on the RAM requirements, matching what the PCI Express bus can deliver.

"There is currently no commercially available alternative device that is comparable to the AMD EPYC design, nor platforms offering this many Gen4-capable slots and this kind of massive work power. The 2nd Gen AMD EPYC CPU platform therefore offers an architectural advantage."

Niko Neufeld, Project Leader, LHCb Online Computing, CERN

Room to grow with AMD EPYC™ CPUs

From May through September 2019, CERN qualified its 2nd Gen AMD EPYC CPU platform before selecting it as the solution for the LHCb experiment. "This solution allows us to reduce by one third the number of servers," explains Neufeld. "This not only saves costs, but having fewer servers is also an advantage when you're building a high-speed, low-latency network. With a larger network, you have more collision problems. The more compact you can make the system, the better."

The AMD EPYC CPU will give LHCb the upgrade path it requires to take its experiment even further. "We don't have the computing power of Google or Facebook, but the AMD EPYC CPU enables us to do the processing we need in a relatively small and compact system. This was not possible 10-15 years ago. Now, there is room for growth. With the same EPYC technology, it is possible to double our capacity in the same space. Our plans for subsequent years are to increase the detectors and sensors. It gives you a lot of headroom. EPYC allows us to do more."

"There is currently no commercially available alternative device that is comparable to the AMD EPYC design—nor platforms offering this many Gen4-capable slots and this kind of massive work power. The 2nd Gen AMD EPYC CPU platform therefore offers an architectural advantage," concludes Neufeld.

WANT TO LEARN HOW AMD EPYC™ PROCESSORS MIGHT WORK FOR YOU?

Sign up to receive our data center content

amd.com/epycsignup

© CERN

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit amd.com/epycserver.