



HPC FEATURE OVERVIEW

Why Consider ARM-Based Servers For HPC?

ARM offers three specific, highly attractive advantages for HPC system designers:

- CPU Instruction Set Architecture (ISA) that can be optimized for specific workloads – resulting in substantial performance advantages
- Significantly lower energy consumption profile
- Substantially smaller footprint consumption design

Fujitsu's proprietary design added even more capability with the introduction of Scalable Vector Extensions (SVE) to the ARM ISA.

Further, Fujitsu leveraged a System-On-A-Chip (SOC) design using a 52 core ARM CPU with SVE, on-chip high bandwidth 32GB memory and an on-chip network for intra-node and inter-node communication.

The high-bandwidth on-chip memory is a key element in the compute node design as the high memory bandwidth significantly improves the computational capability while the high-speed interconnect network is used to meet the memory capacity needs for the MPI aggregation and collective operations.

As researchers are finding ways to augment traditional HPC applications, the need for high bandwidth memory and large-scale systems is growing at a rapid rate. This simplified ARM System-On-A-Chip design makes it easy to scale horizontally while maintaining a low energy consumption profile.

In summary, the major goal of the ARM Instruction Set Architecture (ISA) is to provide the flexibility to tailor to specific workloads/domains. It is found embedded in low-end devices, mobile phones, some server class machines and, most notably, in the world's fastest supercomputer, Fujitsu's Fugaku. This has become a game changer, as HPC and Supercomputer system designers now have access to a low cost, energy efficient platform that can be tailored to meet their specific workloads. In an environment where applications are getting more complex, the ability to scale without concerns about energy and space consumption is a big advantage over general-purpose systems that carry along excess resources not needed for the workload task but integral to the system design.

A64FX Technologies: Scalable Architecture

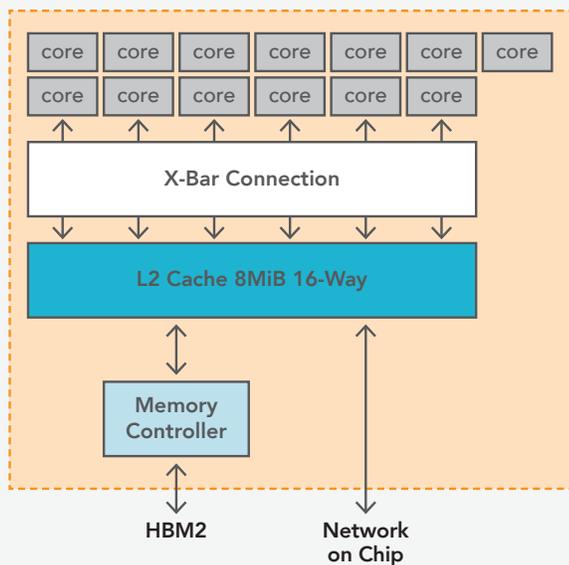
Core Memory Group (CMG)

- 12 compute cores for computing and an assistant core for OS daemon, I/O, etc.
- Shared L2 cache
- Dedicated memory controller

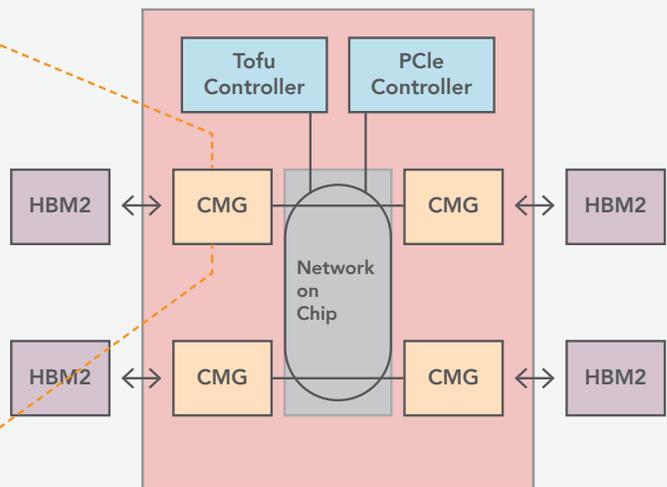
Four CMGs maintain cache coherence with on-chip directory

- Threads binding within a CMG allows linear speed up of core's performance

CMG Configuration



A64FX Chip Configuration



About ViON Corporation

ViON Corporation is a cloud service provider with over 40 years' experience designing and delivering enterprise data center solutions for government agencies and commercial businesses. The company provides a large portfolio of IT as-a-Service, including infrastructure, multi-cloud and artificial intelligence (AI) solutions. Focused on supporting the customer's IT modernization requirements, ViON's Enterprise Cloud is changing cloud management for the market, providing a streamlined platform to audit and control technology in an evolving multi-cloud world. The ViON Marketplace® allows customers to research, compare, procure and manage a full range of everything as-a-Service solutions from leading manufacturers via a single portal. ViON delivers an outstanding customer experience at every step with professional and managed services, backed by highly-trained, cleared resources. A veteran-owned company based in Herndon, Virginia, the company has field offices throughout the U.S. (vion.com).

To learn more, go to vion.com/cloud or email us at info@vion.com