



HPC FEATURE OVERVIEW

Fujitsu PRIMEHPC FX700 Smart Network

Why ViON?

Experience, Expertise, Solutions

- Veteran-Owned business with over 40 years' experience serving government and commercial customers
- 100% dedicated focus on delivering optimum solutions to support our customers and their missions
- Industry pioneer in private cloud solutions
- Certified Storage, System and Network Engineers with 24/7/365 Support Center, Cleared and U.S. Based

HPC Solutions from ViON

- Complete HPC portfolio, including Hybrid HPC, ViON Forever Data Cloud, Private Cloud and HPC Solution Services
- HPC solution partners include: Fujitsu, Supermicro, Arista, DDN, Mellanox, Pavilion I/O and more

Achieving peak performance from your HPC environment is critically important to effectively solving the complex challenges targeted by HPC solutions.

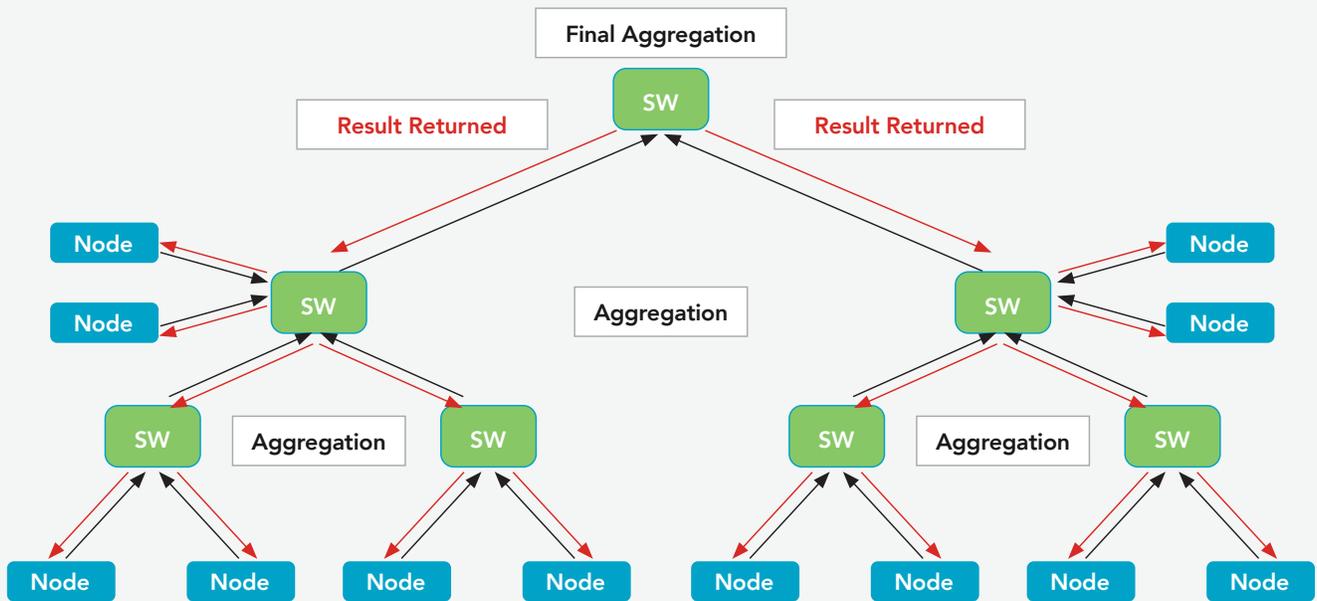
The interconnect fabric in an HPC cluster plays a vital role in the overall performance of the system. Over the years, the network elements in HPC solutions have made tremendous advances, helping to accelerate overall system performance. A few examples include: host network adapter protocol offloads, Remote Direct Memory Access (RDMA), allowing data movement directly to the local memory without CPU involvement, and NVMe Over Fabric, providing virtual local high speed NVMe SSD storage.

HPC and AI applications continue to expand into new areas that call for even higher performance acceleration. Solution design must meet the challenges of faster compute, greater scale efficiency, higher memory bandwidth, accelerated storage performance, faster data staging and de-staging, and the ability to effectively manage and process massive amounts of data. The networking fabric becomes a crucial element to successfully navigating this journey.

Network design increasingly focuses on approaches to accelerating computational performance by offloading data collection and aggregation and accelerating performance between all solution elements and compute nodes. The goal is to maximize compute node resources for their core computational mission. Offloading MPI communications to the network layer increases parallelism and drives faster overall system performance.

Network Accelerated Computing

Another significant advancement in the interconnect fabric is the concept of "In-Network Computing," which was a key element in the previously leading Top Supercomputer, "Summit." This technology is now embedded by Mellanox in its InfiniBand switches and host software package HPC-X. This feature offloads all the MPI Primitives (Aggregation and Reduction) from the local host, moving it into the IB switches (MPI_Bcast, MPI-Reduce, MPI_Allreduce, MPI_Allgather, MPI_alltoall, Barrier communication). The switches perform the vector aggregation operation on data as it moves through the network. The computation includes all the various data types.



In-network Aggregation for MPI Group Primitives (Barrier, Reduce, All-Reduce, Broadcast, etc.)

Parallel computation workload is allocated to multiple processes in a Communication group; each process sends its computation result (vector) up the tree to be aggregated with the results from other processes. Aggregation is done in stages at designated aggregation switches and then the combined partial aggregation is sent to the next ladder up the tree. The final aggregation is completed at the top of the tree. The data has already been reduced in size so there is no need to store all the original data generated by each node to do the full aggregation.

The software creates a virtual network tree on top of the physical network and routes data issued via the MPI interface to an aggregation node where data from various nodes are collected and aggregated then sent up the tree until the final aggregation node is reached. Here the final aggregation step is completed, and the result is sent back down the virtual tree to the MPI communication group. This data streaming technology supports both small and large messages (100MB).

Another benefit of this approach is that it reduces the amount of traffic sent over the network (data is sent only once) and the compute nodes are free to run other computation while it waits for the prior aggregation results.

The FX700 HPC solution design, with high memory bandwidth and industry-leading computational power, leverages this Smart network to deliver a new era of HPC/AI solution scale and performance – capable of handling massive amounts of data and delivering unparalleled multi-processor efficiency in a low energy consumption design.



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