


Network Interconnect for Exascale Systems



Context: will network be the next big bottleneck?

Network interconnect is the backbone of an HPC system, linking the compute nodes together. To reach the exascale with an acceptable energy footprint, supercomputers will include a **huge number of hybrid nodes** (GPUs/CPU), and as a consequence many network interfaces to match GPU throughput and memory bandwidth. Integrating heterogeneous nodes also demands a **smarter interconnect**, with additional features to accelerate connectivity between servers and storage.

About us

RED-SEA brings together the **top European academic centres** and the **key European industrial forces** in the domain of interconnect networks, with a consortium of 12 partners from 6 countries


Project timeframe: 01/04/2021 – 31/03/2024
Project budget: € 7 993 710

We are one of the **SEA projects** working together to develop complementary European technologies for future heterogeneous exascale supercomputing architectures: <https://sea-projects.eu/>

Check the RED-SEA publications and network architecture:



The four pillars of RED-SEA research:



Architecture, co-design and performance


Optimizing the fit with other EuroHPC projects and with the EPI processors

- Analyse network requirements of representative HPC applications, select relevant benchmarks to co-design the RED-SEA network architecture
- Optimize HPC applications and mini-apps to take full advantage of the RED-SEA hardware testbeds and simulation platforms
- Coordinate the various hardware testbeds and simulation platforms used to evaluate the RED-SEA network architecture
- Holistic evaluation of the RED-SEA network design for future exascale systems



High-performance Ethernet

- Develop a **high-performance low latency bridging solution with Ethernet**
- Study RDMA communication over Ethernet using state of the art RoCE semantics
- Build an FPGA prototype for the gateway to offer direct interoperability with Ethernet switch or endpoint, demonstrating TCO and performance benefits
- Develop the necessary IPs for FPGA or ASIC implementation
- Develop the software components: a driver presenting an Ethernet virtual NICs and a virtual switch management software.



Efficient Network Resource management

Congestion management and Quality-of-Service for the challenging traffic patterns produced when mixing HPC with storage workloads on the same interconnect and at scale

- Reducing incast congestion by hardware and software support for collective communications
- Isolation of traffic from different applications through virtual networks and link schedulers
- Optimizing injection throttling mechanisms
- Reducing in-network congestion using adaptive routing
- Network power management



Endpoint functions and reliability

- Scalable end-to-end reliability protocols for BXI
- Protected sharing of clusters using BXI
- Tight integration of network interfaces with RISC-V cores and accelerators, such as those of EPI
- Optimized MPC-MPI and ParaStation MPI libraries
- Advanced programming models for in-network compute