



**OPEN**

Compute Project

Proposed Charter  
For  
High Performance Computing

Draft: December 3, 2014

Version: 1.1

# 1 Revision History

<b>Date</b>	<b>Name</b>	<b>Description</b>
07/15/2014	Devashish Paul, Mohammad Akhter	Base items to serve as input to OCP HPC Kick off for UNH Workshop
08/08/2014	Devashish Paul	Key Items from OCP UNH Engineering Workshop
9/09/2014	Devashish Paul	Updated with discussion items with Aug 2014 HPC group conference call
9/16/2014	Devashish Paul, Thomas Sohmers	Updated content based F2F meeting in SFO with co leads Thomas Sohmers and Devashish Paul
9/17/2014	Thomas Sohmers	Updated sections regarding open silicon devices and silicon photonics
9/18/2014	Thomas Sohmers	Added fabrication and outside involvement pieces
9/22/2014	Thomas Sohmers	Final draft for group review
12/3/2014	Thomas Sohmers	Updated formatting for IC review

## 2 Contents

1	Revision History .....	2
2	Contents .....	3
3	Overview.....	4
4	Scope .....	4
5	Key Values.....	4
6	Relationship to Other OCP Groups .....	4
7	In Scope Technology Categories .....	5
8	Out of Scope Technology Categories.....	5
9	Key Project Focus Areas.....	5
10	Project Phases/Commercialization Strategy.....	5
11	Outside Involvement .....	7



## 3 Overview

The HPC project has been established in the Open Compute Project to service the needs of the High Performance Computing, Supercomputing and Low Latency Analytics needs of the computing industry and to service it with open hardware platforms delivering solutions in the market from system level down to silicon.

## 4 Scope

The project will focus on low latency multi processor systems that can scale to hundreds or thousands of nodes in an energy efficient way to service the needs of the target market. To ensure project successes, it has been divided into multiple phases over a span of a few months to 3+ years, during which the project will deliver open designs/systems to the market. Phase 1 will focus on taking existing industry designs and modifying where needed and opening them up for the general market. Phase 2 will define and deliver new boards and systems based on silicon in the market today. Phase 3 which will run in parallel and will define new low latency interconnect silicon and processor architectures on which Phase 3 system level boards and designs will be built. The project also reserves the right to explore silicon photonics as a Phase 4 roadmap solution.

The HPC project will work within the mechanical and electrical frameworks established in other OCP groups, and where possible will reuse or modify existing OCP approved designs to service the HPC target market and customers. Open silicon solutions developed by the HPC project will be made available to all other OCP projects and the industry at large as a way of driving innovation beyond boards and systems..

## 5 Key Values

- Deliver Open Hardware platform for HPC industry which often must do custom case by case deployments
- Serve as a central point of collaboration for HPC industry in terms of performance and cost optimization of HPC compute and networking platforms
- Group will act as an industry leader to help drive future innovation in HPC market in a collaborative open industry context
- Provide the path for industry standard Open silicon devices.

## 6 Relationship to Other OCP Groups

Where possible use designs from Server/Storage for compute

- Comply with OCP Rack/Hardware/Electrical
- Improve networking designs for latency and scalability, investigate using alternate technologies already in production for near term goals.
- Come up with new OCP Open silicon spec focused on scalability, low latency and energy efficiency

- Open APIs etc
- Innovations from HPC group may flow back to other OCP groups where appropriate

## 7 In Scope Technology Categories

- Low latency top of rack switching
- Combined Switch and micro servers
- Combined compute and switching
- Low latency scalable storage
- Connectivity from HPC Fabrics/clustering technology to out of cluster networking via Ethernet
- OCP mechanicals (19 and 21 inch)
- APIs and software interfaces
- Open Hardware compute and switching
- Development of OCP HPC Interconnect silicon spec (can build on existing mainstream technologies such as PCIe, RapidIO, Infiniband, Ethernet)
- Investigation into Open industry processor specs for HPC Market During this process, the specification will be posted to the wiki.

## 8 Out of Scope Technology Categories

- Items already covered in server, storage, networking and other groups that are not optimized for low end to end latency and multiprocessor compute
- Other items TBD

## 9 Key Project Focus Areas

- High Performance Computing
- Supercomputing
- Data Center Low Latency Analytics
- Mobile Network Edge Computing Analytics
- Low Latency Financial Trading

## 10 Project Phases/Commercialization Strategy

Target time frame indicates when each phase is expected to generate output. All phases start immediately with work from participating contributing companies

1. Phase 1 (6-12 months): Leverage as much as possible, server group compute, look into low latency networking



2. Phase 2 (12-24 months): HPC optimized heterogeneous computing-ARM, GPU, x86, DSP etc

3. Phase 3 (12-24 months): Deliver open interconnect silicon and processor spec to market to drive interconnect innovation optimized for HPC and Supercomputing and other target verticals covered by the scope of the OCP HPC project. In this phase HPC will also work on delivering open industry processor architecture. The work for this is to be defined, During phase 3 new board and system level solutions will be built from emerging OCP HPC compliant silicon. Below are some of the overall themes. The project will build from existing industry solutions available in the PCIe, RapidIO, Infiniband and Ethernet ecosystem and leverage features that are technically and commercially viable to converge on optimal solutions.

- a. HPC needs huge scale of any to any processing nodes
- b. Latency is a primary concern for this market as well as those that need analytics
- c. Energy footprint is an issue
- d. Low hanging fruit is to eliminate latency and power from NIC's and other interconnect devices
- e. Need native protocol termination on processing endpoints, like processors, DSP, GPU, FPGA
- f. Diverse industry initiatives to create proprietary clustering fabrics at many startups and large processor vendors
- g. Desire to open up to remove vendor "lock in" and enhance interoperability
- h. Prefer to start with some industry standard options that scale, have low latency, multi vendor collaboration etc
- i. Take best attributes of PCIe, Infiniband, RapidIO, Ethernet, and other technologies to reach exascale computing
- j. Processor architectures and other low level silicon technologies optimized for the HPC market. This includes specifications for individual components at the intra-IC level and potentially their hardware description language, register transfer level, gate description level, and/or other technical implementation pieces. Higher level specifications and implementations using these base level components will be contributed, and would be in the form of open instruction set architectures, FPGA synthesizable devices, and ready to fabricate ICs.
- i. Due to the relatively high cost of low volume/prototype integrated circuit fabrication, future plans for the group may include organized Multi-Project Wafer (also known as "shuttle run") fabrication, where multiple companies share the costs of a silicon wafer, and each get a fraction of the wafer for their own devices to be fabricated. By negotiating with a foundry as a group, costs can be further reduced, in addition to potential discounts a foundry may give if a design is open sourced and based on their process technology. Having such a program within the group would allow for member companies to reduce time and costs for prototyping, while encouraging open collaboration between members and the greater industry.

4. Phase 4 (24+ months): Silicon photonics has the potential to be a major game changing technology for initially the HPC industry and the greater server and computing industries.

While this is still in early developmental stages, silicon photonics has the potential to radically reduce power consumption, reduce latency, and increase bandwidth of computing systems. By taking this future technology into account now, standards in other parts of the HPC group can plan ahead, and new standards focused on industry use cases can be proposed to influence the development of this emerging technology.

## 11 Outside Involvement

While most server and datacenter systems that the Open Compute Project focuses on commercial systems, most of the largest users and developers of HPC systems come from academic or government backgrounds. As such, the HPC group plans on involving members from these outside communities to both encourage them in using open specifications and standards within their work, and contributing their research and developments back to the Open Compute community.