

OCP HPC SubProject New Frontiers Proposal

Introduction	2
Background	3
Scope	4
Problem statement (All)	5
Market Gaps (All)	5
Methodology - Simplicity (Allan)	5
First Principles Thinking	5
Simple Modular Building Blocks	6
Computing Wall Infrastructure - From Data Centers to Family Homes	6
Proof Of Concept Research & Development Program	8
System Integration Specifications, Interoperability & POC Lead	9
Power & Thermal Mechanical & Electrical Design & Test Vehicle	9
System Management Hardware and Software Development	10
Universal Interconnect Development	10
HPCM Module Development	10
Evaluation	11
Individual components (respective companies)	11
Integrated system (Terizza)	11
Success Criteria	11
Data analysis	11
Interpretations	11
Discussion	11
Results	11
Limitations	11
Future work	11
Commercials	11
Summary	11
Contributions to the field	11
Recommendations for further action	11

====

```
{ < This Section is to be removed once all work is complete>
Timeline targets for proposal completion
v0.0 - Jan 15 - Outline is set, and who is going to what is decided
v0.3 - Jan 19 - All stakeholders are informed/Document filling is in progress
v0.5 - Jan 26 - No unknown issues left - Document writeup is well in progress
v0.8 - Feb 2 - Draft is almost complete - Not reviewed yet.
v1.0 - Feb 9 - Review Done
}
```

Introduction

Demands of computing have been changing dramatically for more than a decade and this has resulted in some seismic shifts in approaches to computing architecture. We have seen custom architectures emerge from many companies from early efforts in Facebook, resulting in the creation of OCP, to the many proprietary architectures including Google's TPU, NVidia's DGX, Tesla's Dojo etc.

Factors that have driven the need for computing architecture innovations include :-

- AI is Increasing Power demands including delivery and density
- Changing Cooling requirements, both Air and Liquid
- Managing Processor & Memory Heterogeneity
- Bandwidth memory and latency memory within the same compute domain
- Memory centric transaction model/fabrics
- Increasing need for Domain Specific Architecture, DSA, Composability
- System Management tools to support increasingly complex DSA Composability and security.
- Aligning evolving supply chain with the cost model (lower) and performance model (higher).

Specifically in the Government HPC/SuperComputing sector we have seen custom HyperConverged Heterogeneous Nodes in the Summit and Frontier machines. Each node has a very specific mix of CPUs and GPUs including the interconnect between them. These expensive configurations do not generate much external commercial demand and this has resulted in many of the large vendors exiting this business, like IBM, as it has become commercially unviable.

Additionally the custom machines being built by incumbents, like NVidia, are not necessarily a good fit for government application needs. They are also very proprietary, meaning that they cannot be tailored with a mix of heterogeneous processors from different vendors, as may suit many applications.

This response to the ORNL New Frontiers RFP aims to resolve the above challenges by redefining an Open Standard Computing Architecture that addresses the needs of a DSA computer, with extreme composability through a very modular design, that can support Processors, Accelerators, Memory and Media from any and all vendors. Machines based on this new architecture will also be easily serviced and upgraded allowing for a far longer extended service life with the significant embodied carbon benefits that this will bring.

Background

The 19" Chassis was invented 100 years ago by the telecommunications Industry for Telephone equipment. It was first adopted for computing in the 1950's and eventually evolved into the ubiquitous format we know today for all sorts of electronic equipment, including servers, networking devices, and Storage.

The 19" Racks modularity is defined in Rack Units, U's, and In the recent decades we've seen server designs shrink from typical 4U size down to 1U and now frequently multiple servers in a 1U Chassis. We've also seen accelerators be added to these servers in the form of PCIe cards as heterogeneous computing becomes ubiquitous.

If we delve into the underlying silicon technology of today, we are seeing virtually the entire Server effectively subsumed into a single Chip Package. Motherboards are now motherchips'. Chiplets within these packages represent what used to be separate components within the server, like processors, accelerators, memory and IO. Yet despite the relentless shrinking of the silicon and components, the System level architecture remains unchanged with its 19" rack and Rack Unit chassis building block.

With regards to architectural shifts Memory and Storage have historically been treated as separate entities often housed in separate 19" racks in different parts of Computer rooms and Data Centers. However, as data centric applications have become more prevalent we are seeing a need to increasingly treat memory and storage as simply memory that is efficiently accessible by all types of compute. Data Movement is the single largest consumer of power in computing and this has ushered in the need to tightly couple all types of memory and compute in order to minimize that power consumption.

As components have become more tightly packed together, we are seeing exponential growth in power density with Computing Chips consuming close to 1KW today and expecting to move well beyond 1KW in the coming years. This has resulted in major challenges in both delivering power to these components as well as effectively thermally cooling them.

Finally, as interconnect speeds have relentlessly marched higher, we are seeing major challenges in reliably transporting data across distances needed for 19" rack Unit compatibility. This has led to very dense hyperconverged nodes, whose composability is severely limited, as well as the need to more efficiently integrate optical communication at low cost.

When one takes a step back and holistically assesses the changes in computing architecture over the last few decades it becomes manifestly clear that a revolutionary change in our industry standard computing architecture is urgently needed.

In Summary the standard server is going through the following disruptions

- Motherboards to Motherchips (Chiplets)
- Copper based fabrics to Optics (increasingly)
- Air Cooled to Air + Liquid cooled

The challenge is advancing performance, lowering cost, increasing security while improving the velocity of the supply chain and creating a circular electronics economy without loss of availability. We will show our proposal advanced all these 5 dimensions

1. Performance
2. Cost
3. Supply chain
4. Security
5. Serviceability/modularity
6. Availability

In a significant and disruptive way, but with a roadmap that mitigates the risk to our end defined state.

Scope

This is the OCP HPC SubProjects Response to the ORNL New Frontiers RFP. The primary objective of the proposed project is to create a POC for a revolutionary Open Computing Architecture that replaces the 1981 IBM PC Architecture that is still used in Servers today.

A concept for this New Architecture, called HPCM (High Performance Compute Module), has been developed within the OCP HPC SubProject, and is underpinned by 4 primary tenets, being :

1. Sustainability - through Energy Efficiency and Energy Recovery.
2. Cost Reduction - through Open, Flexible, & Modular System Architecture building blocks.
3. Composability - driven by 1 & 2 above yielding Domain Specific Architecture, DSA, solutions.
4. Infrastructure Durability - through modular and replaceable system components.

The project focuses purely on a Modular System Level Hardware infrastructure along with innovations around advanced System Level Management Hardware and Software to accommodate a new era of DSA Composability.

Silicon, Interconnect Protocols, and Application software are not part of the scope of this project. These parameters will be dictated by the components that are used to construct differing HPCM Modules. The HPCM architecture is such that ANY of today's existing Hardware and Software

stacks will have the potential to be easily replicated with this architecture as well as providing an unparalleled opportunity for revolutionary innovation going forward.

A core tenet of this project is its openness under the OCP umbrella and this project will involve the collaboration of multiple OCP participant companies who have contributed to and are cited in this proposal. The Open Specifications resulting from this project will be contributed back to OCP for everyone to access.

Problem statement (All)


We address the following problems

- A. **Workloads:** Enable new workloads (including AI) without compromise, and seamless integration of existing CPU based workloads including higher thermal densities
- B. **Roadmap:** Seamless migration and adoption of new technologies - Optics, water cooling
- C. **Eco-system:** Open reference specification
- D. **Modular system** - Accommodate varying lifecycle of different components (CPU, GPU, HBM, DDR and switches) without whole scale upgrade of the platform.
- E. **TCO:** Align supply chain and cost model

In essence, we take these problems/ constraints and make the system simple to build, deploy and life cycle.

Market Gaps (All)

1. Open and non-democratized solution (today). The hyperscalers are going proprietary and closed with none of their sub-systems being made available to other parties. Nvidia is the sole participant among the big 4 semiconductor companies (Broadcom, AMD, Intel, Nvidia) building cloud scale systems, but its product is proprietary (NVlink, CUDA).
2. A cloud-scale system for an emerging new category of AI Datacenter (non-hyperscalers). AI workloads are scale-up and scaleout unlike SaaS which was pure Scaleout and traditional enterprises (3-Tier - SAP, Oracle, SQL based Enterprise apps) are scaleup. A hybrid scale-up/scale-out system with built-in memory coherent/consistent fabrics for Scaleup and scaleout with UltraEthernet that is open with a rich linux-like SW ecosystem is a key gap in the market today.
3. Government applications don't suit the proprietary solutions existing today. Continuity, sovereignty and need to deploy new form factors (drones to robotic systems)
4. Emergence of AI centric data centers. These are green field DC that are being invested by geographies due to availability of power and real estate and sovereignty needs. The visual below shows the growth in various categories of DCs.



Methodology - Simplicity (Allan)

First Principles Thinking

Considering the simple building blocks of all computers we can simply boil it down to the following pieces: -

- Compute, (Including Intelligent Switches)
- Memory
- I/O - Interconnect


The above 3 building blocks also require some fundamental pieces in order to operate, as follows: -

- Power Input
- Heat Removal (Thermal Management)
- System Management

Simple Modular Building Blocks

When considering the above elements, it becomes clear that the Compute and Memory have become heterogenous at the system level. We can separate these components into energy dense components that are primarily Compute with some Memory or relatively energy light components for High Capacity Memory with some or no Compute. Applications require different mixes of these components and so a modular approach with pluggable composability would be

advantageous. This leads us to two composable building blocks as shown below. These modules have been given names, HPCM (High Performance Compute Module) and E3.S, the new industry standard Media Module.



A Data Center or High Performance Computer, HPC, will then be constructed out of many of the 2 above building blocks. The interconnect between these modules will depend on the system level architecture design which can vary between classic Data Center network topologies like Spine-Leaf etc or Memory Centric Architecture, DSA, Topologies, e.g. Frontiers Node CPU/GPU internal interconnect. Flexibility at this level is essential for the easy construction of any DSA that is required for optimal performance to each specific application.

Computing Wall Infrastructure

With such a radical simplification of computing, down to the 2 fundamental building blocks described in the previous section, it is possible to leverage these modules in many infrastructure use cases, including : -


- Data Centers
- HPC Machines
- Edge Computing
 - Offices
 - Schools
 - Family Homes
 - Telecom Base Stations
 - Internet Equipment
 - Automobiles

The design of the infrastructure would simply vary only in the total number of modules that could be supported in the different application domains.

Therefore the infrastructure design would ideally be easily scalable depending on the number of compute modules it needs to support.

The critical parts of the Infrastructure design will be Power delivery and Heat Removal, and ideally Heat Re-use. All infrastructure, in each use case, would be designed into the building that is constructed and this will effectively replace the traditional 19" Rack.

The proposed Infrastructure will take the form of a traditional Wall of any building that will carry the water and power to the Compute Modules. The Compute Modules will be "plugged" into the wall where the Power and Water will be Blind Mate connected as the module is inserted. Interconnection cables between modules will be populated after the Compute and Memory Modules have been installed.



Wall Example : Capable of supporting up to 16x HPCM Compute Modules

Future Data Center Infrastructure could include gantry robots capable of servicing the interconnect and modules as needed for maintenance or upgrade purposes.

Wall of compute Value Proposition (Compared to DGX or Rackmount servers?)

- Cost
- Power/Thermal
- Performance (Interconnect/fabric)
- Serviceability
- Supply chain alignment

- Availability

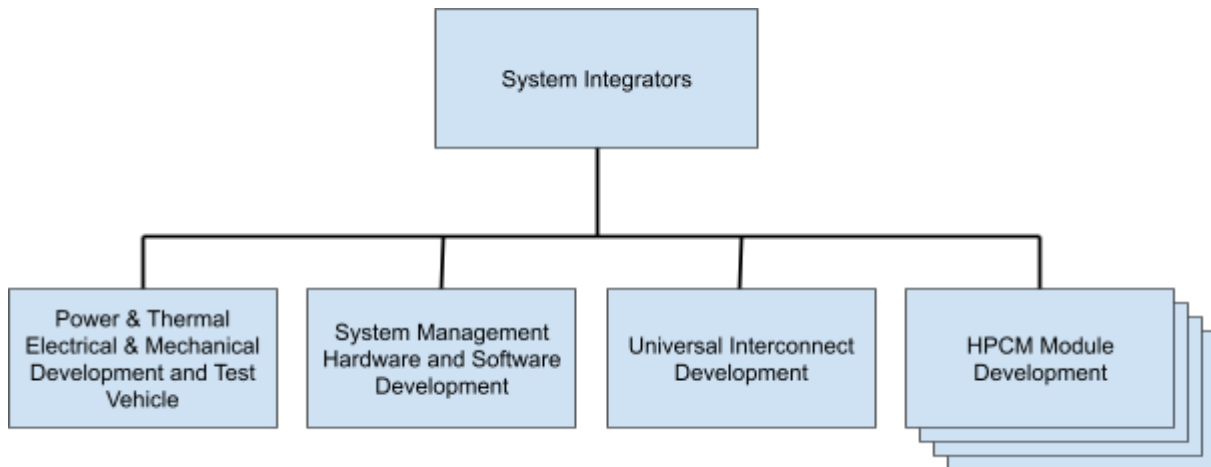
Proof Of Concept Research & Development Program

The OCP HPC SubProject has been advancing a 3D Model Concept of this design and it is planned to use the New Frontiers funding to take this idea from concept to a work POC reality. Several companies have participated in pulling the concept together as well as contributing to this proposal. These companies will partake in the creation of the POC.

We have split the project into 5 separate pieces, being : -

1. System Integrators
2. HPCM Module Development, Multiple
3. Power and Thermal Infrastructure Design and Test Vehicle
4. Universal Interconnect Development (Electrical & Optical, NPO & CPO support)
5. System Management Hardware and Design and Development

The basic Project Structure is as follows : -



The System Integrator participants shall be responsible for pulling all the pieces of the project together into a working POC Wall of Compute. They will be responsible for coordinating all of the specifications for consistency and interoperability and will be responsible for submitting the Open Standard Specifications to OCP for approval and formal release at the end of the Project. They will also create simple working POC Demo's showing interoperability between the HPCM's that have been developed and 3rd party E3.S Modules as appropriate.

The Power and Thermal Test Vehicle participants shall design and develop the core mechanical and electrical pieces of the Wall infrastructure and HPCM Module. These will be the common components that will be used for all HPCM Module development. A Power LoadSlammer Test HPCM Module shall be created as part of this effort. It will be based on one of the real HPCM

Modules that are being developed and developed in coordination with the participant that is responsible for creating the working HPCM.

The System Management Hardware and Software Development participants will initially define a Hardware management infrastructure that will be integrated into each HPCM. This will allow the HPCM to self boot as well as be discoverable by a higher level system & security manager. It will follow the Open LibreBMC efforts using an FPGA at the heart of the controller. System Management software will also be developed but will also leverage industry standard software from OpenBMC and wider community.

The Universal Interconnect Development participants will collaborate on developing a universal interconnect design that has the ability to work with all transceiver IO protocols that are supported in the industry today from PCIe/CXL to Ethernet and proprietary protocols such as NVLink. Additionally it will include support for both electrical and optical interconnects as well as the ability to conductively cool components in active cables.

The HPCM Module Development participants will create multiple HPCM Modules featuring different technologies including A processor, Accelerator and Switch variant. When integrated together into one HPCM System the System Integrators will be able to demonstrate interoperability between the HPCMs and their E3.S modules showcasing a POC Demo application.

System Integration Specifications, Interoperability & POC Lead

The following companies are contributing to this proposal from a Systems Integrator perspective :-

- ABRA Works
- Terizza

ABRA Works and Terizza Proposal Contribution here

Power & Thermal Mechanical & Electrical Design & Test Vehicle

The following companies are contributing to this proposal from a a Power & Thermal Test Vehicle perspective :-

- Boyd Corporation
- Lumenir
- Progranalog
- Power supply Company? (MPS?, Analog?, Renasys? etc)
- Electronic Innovations

We propose that the printed circuit board at the heart of an HPCM module be populated exclusively with chips and low-profile packaged parts; not including big packages or heat sinks or daughter boards or special enclosures. This circuit board may include an organic or a glass substrate. The secret sauce is a hermetic protection layer that covers all the mounted components with a thin conforming layer composed of a well-engineered sequence of layers. This hermetic layer protects against water intrusion across all operating parameters, including high power density, corrosion, turbulence, thermal cycling, mechanical stress due to thermal expansion variations, and pressure variations. The proposed solution builds on recent intensive research by multiple companies into coating methodologies spanning atomic layer deposition (ALD), plasma surface activations (for promoting adhesion), chemical vapor deposition, and electroplating. It combines proven performance of ALD metal oxides and Parylene C, with the strength and durability of Cu/Ni platings. We propose the testing of eight different hermetic solutions in order to select the best performance, combining reliable protection against water intrusion with the lowest possible thermal resistance between the heat generating junctions and the cooling liquid. Combinations of water and ethylene glycol may be employed, and dielectric liquids will be tested for comparison. The proposed solution combines both single-phase and two-phase cooling, with sophisticated AI control. An optimal operating point will be determined, wherein the pressure of the coolant is reduced to promote low-temperature boiling, to support a maximum junction temperature of 90°C.

Discussion on Vacuum Pulling Thermal Water Loop

From HPCM Thermal Meeting on 2/28/2024 we discussed how a thermal loop would operate if we were to successfully pull a vacuum on ~0.25 Atmospheres.

The loop in the wall would consist of the Heat Exchanger at the top that take a single vertical column of 8x HPCM modules and exchanges the heat from this closed loop column with the facility water. The return path from the Heat Exchanger would initially flow into a vacuum chamber which would have a vacuum pump attached pulling the ~0.25Bar vacuum. The water would then pass out of the chamber and into an expansion chamber before flowing into a circulation pump that pumps the water down the cold feed pipe to all 8 HPCM modules and back to the hot return. The cold water feed will pass through a variable valve before entering the

HPCM modules and the valve will control the water flow through the HPCM module in order to maintain an output water temperature of 65C.

Diagram here depicting description in the above paragraph.

System Management Hardware and Software Development

The following companies are contributing to this proposal from a a System Management Hardware and Software Development Perspective : -

- Lattice Semiconductor
- Intel
- LANL

High Level Requirement

Configuration Management

For this section, we need to firmly establish the problem(s) and why they are different from current technological solutions. I find it useful to separate the discussion into the physical and logical Configuration management features. Physical Configuration management features include:

- Ensuring that we have an accurate inventory of every device that is electrically connected to the system, including a secure assertion of the firmware version of each device. At this level, we need a way of disabling any connector that doesn't validate and subsequently raising an alert. - This is where Caliptra, the TEE, and other standards like SPDm come in. Security must be rooted in the hardware, but verifiable in control software.
- Providing an abstraction that allows users to manually compose "Nodes" from all the disaggregated devices.

At the logical Configuration Management level, the system management tooling must provide abstractions that are built on top of the hardware management to ultimately optimize the full system for use by users.

- Providing an abstraction for composability/resource management that allows a single system to be subdivided into a collection of "Nodes" which can then be made available for workloads. It is at this layer where user-centric APIs will be necessary to provide feedback on optimal reconfigurations based on various desired parameters e.g. per-node memory latency vs GPU density.
 - Redfish has a [composability standard](#) which may or may not be a sufficient standard for this.

- Securely managing any certificate or key material required to ensure that communication among devices that are part of one “Node” cannot be intercepted or interfered with by devices associated with another “Node”.
- Securely managing any interconnect technologies that provide intra-node connectivity such that workloads can be effectively isolated from each other.

Universal Interconnect Development

The following companies are contributing to this proposal from a Universal Interconnect Development Perspective : -

- Amphenol
- Samtec
- ACES
- Avicena
- Lightelligence
- Kandou

The universal interconnector needed by HPCM architecture provides two types of links: HPCM base module to E3.S module connector (Referred as E3 connector) and HPCM to HPCM cable (Referred as HPCM cable). The connector needs to be full SMT to save board space for electrical routine. As HPCM can adopt different protocols, so should the connector and cable. It will support different data rates such as 32G NRZ, 64G PAM4, 112G PAM4. HPCM architecture integrates the tight pitch of media module on a small board estimate. Power delivery and heat dissipation will be implemented on connector interface as well.

Amphenol provide the below design concept for E3 connector and HPCM cable.

- E3 connector is compatible with current EDSFF device following SNIA SFF-TA-1002 and SNIA SFF-TA-1020. The connector is converted to full SMT mounting features. Both 4C and 4C+ EDSFF connector design are illustrated.

Outline dimensions

60.12
57.02


9.10

10.64
8.64

[illegible]

Figure Error! No text of specified style in document.1 4C Vertical Full SMT Connector

Outline dimensions



Foot Print

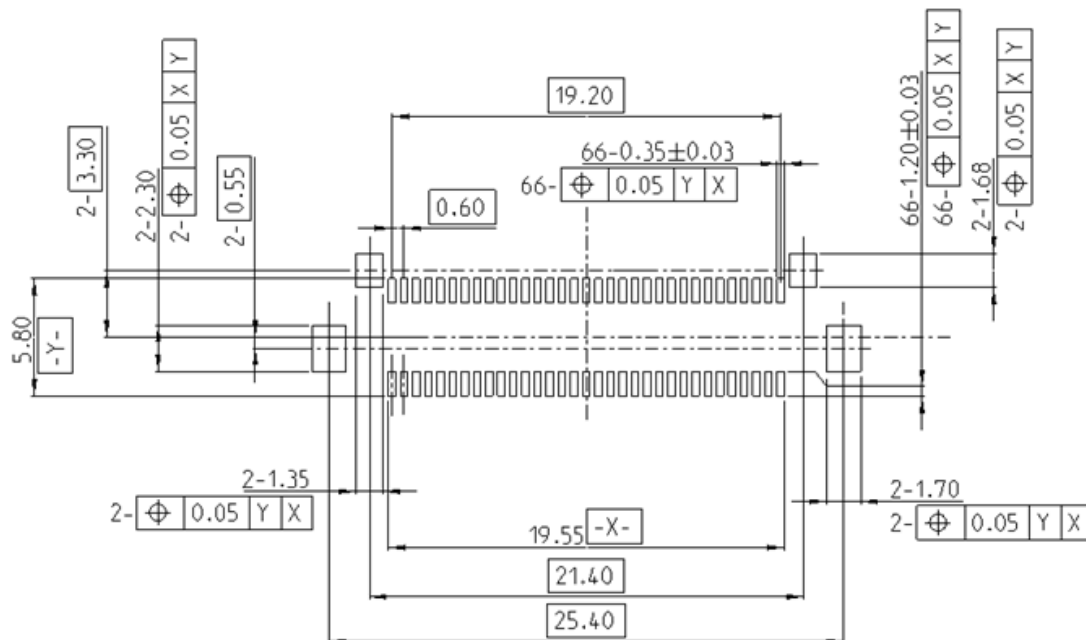


Figure Error! No text of specified style in document.3 66 pin HPCM cable connector

The active cable need dissipate the heat generated by the transmission chip. Such heat can be transferred through the metal gasket feature design as shown in the below picture.




Figure Error! No text of specified style in document.4 HPCM Cable With Metal Gasket Feature

HPCM Module Development

The following companies are contributing to this proposal from an HPCM Module Development Perspective : -

- ABRAWorks
- Progranalog
- Lumenir
- Enfabrica

Evaluation

Individual components (respective companies)

Integrated system (Terizza)

Success Criteria

Data analysis

Interpretations

Discussion

Results

Limitations

Future work

Commercials

Summary

Contributions to the field

Recommendations for further action