



D2.1: Requirements and use cases refinement

Document Properties

Contract Number	101189612		
Contractual Deadline	M4 (April 30, 2025)		
Dissemination Level	Public		
Nature	Report		
Edited by :	Olivier Déprez, SiPearl		
Authors	Peter Gray, Cloud Sigma Ivy Peng, KTH Blagovest Tushev, CloudSigma Xavier Teruel, BSC Thomas Van Liefde, 2CRSi	Dimitris Theodoropoulos, ICCS Iakovos Mavroidis, Exapsys Olivier Déprez, SiPearl Fabien Chaix, FORTH Stelios Louloudakis, FORTH	
Reviewers	Sven Schenk, Extoll Cagatay Yilmaz, RISE		
Date	30/04/2025		
Keywords	RISC-V, Cloud, Acceleration, Network Object Store, Containers, PCIe, CXL, IaaS, PaaS, OCP, DC-MHS, HPM, DC-SCM, BMC, Memory disaggregation		
Status	Final		
Release	1.0		



History of Changes

Release	Date	Author, Organization	Description of Changes
0.1	03/03/2025	Olivier Déprez, SiPearl	Skeleton document (Table of Contents, assignments)
0.2	17/03/2025	Contributions by SiPearl, 2CRSi, Exapsys, Extoll, FORTH	OCP form factors and system management (round 1)
0.5	24/03/2025	Contributions by KTH, ICCS, RISE, BSC	Use-cases
0.7	31/03/2025	Contributions by CloudSigma	Cloud infrastructure and services
0.75	15/04/2025	Contributions by FORTH	OCP system management (round 2)
0.8	16/04/2025	Olivier Déprez, SiPearl	First version available for review
0.9	25/04/2025	Olivier Déprez, SiPearl	Revision to address review comments (from Extoll and RISE)
0.91	29/04/2025	Olivier Déprez, SiPearl	Finalized contents
1.0	30/04/2025	Olivier Déprez, SiPearl Stelios Louloudakis, FORTH Manolis Marazakis, FORTH	Minor edits, Submitted version

Table of Contents

DOCUMENT PROPERTIES	1
HISTORY OF CHANGES	2
TABLE OF CONTENTS	2
LIST OF FIGURES	4
LIST OF TABLES	6
1 EXECUTIVE SUMMARY	7
2 CONTEXT	8
2.1 PURPOSE OF THE DOCUMENT	8
2.2 REQUIREMENTS ELICITATION	8
3 MARKET CONSIDERATIONS	10
3.1 MARKET SIZE AND GROWTH ESTIMATES	10
3.1.1 <i>Cloud Services Market: IaaS and PaaS</i>	10
3.1.2 <i>Open Compute Project (OCP) Servers: Potential market impact</i>	10
3.1.3 <i>RISC-V as a potential driver force</i>	12
3.2 INDUSTRY ADOPTION	12
3.2.1 <i>The Open Compute Project (OCP) Adoption</i>	12
3.2.2 <i>RISC-V Adoption</i>	13
3.3 COMPETITIVE LANDSCAPE	14
3.3.1 <i>OCP-compliant compute and accelerator modules</i>	14
3.3.2 <i>Server management and security solutions</i>	14
3.3.3 <i>Industry trends and differentiation</i>	15
3.4 STAKEHOLDER DEFINITION	15

4 REFERENCE SERVERS.....	17
4.1.1 Arm-based reference servers.....	17
4.1.2 x86 reference servers	17
4.1.3 RISC-V Platforms.....	17
5 USE CASES	18
5.1 ACCELERATED DATA PROCESSING.....	18
5.2 INFRASTRUCTURE AS A SERVICE (IAAS).....	20
5.2.1 Networked Object Store Use Case.....	20
5.2.2 Regular Storage Use Cases.....	21
5.2.3 Container Use Case.....	23
5.3 PLATFORM AS A SERVICE (PAAS).....	24
5.3.1 Multi-device approach.....	24
5.3.2 HPC domain approach.....	25
5.3.3 ML/AI domain approach.....	26
5.3.4 Cloud domain approach.....	26
5.3.5 Evaluation Approaches.....	27
5.4 REMOTE CXL-BASED DISAGGREGATED MEMORY	29
5.4.1 Efficient Memory pool size management.....	29
5.4.2 Memory Access monitoring.....	30
5.4.3 Data Protection.....	30
5.4.4 Accelerated data processing.....	30
5.4.5 CPM Configuration.....	30
5.4.6 Efficient memory access.....	30
6 KEY RESULTS.....	31
6.1 TECHNICAL READINESS LEVELS (TRL).....	31
7 REQUIREMENTS.....	33
7.1 OCP STANDARDISATION LANDSCAPE	33
7.2 OCP HPM FORM FACTORS.....	34
7.2.1 M-FLW.....	34
7.2.2 M-DNO.....	35
7.2.3 M-SDNO.....	36
7.2.4 HPM Market Considerations.....	37
7.3 DUAL SOCKETS RHEA2-BASED HPM REQUIREMENTS.....	40
7.3.1 General.....	40
7.3.2 Interface with DC-SCM	42
7.3.3 OCP-NIC Support.....	44
7.3.4 High-Speed I/O Connectivity.....	44
7.3.5 Platform Infrastructure Connectivity.....	45
7.3.6 Software.....	48
7.4 SINGLE SOCKET RHEA2-BASED HPM REQUIREMENTS.....	51
7.4.1 General.....	51
7.4.2 Interface with DC-SCM	52
7.4.3 OCP-NIC Support.....	54
7.4.4 High-Speed I/O Connectivity.....	54
7.4.5 Platform Infrastructure Connectivity.....	54
7.4.6 Software.....	54
7.5 EPAC/EUPILOT-BASED HPM REQUIREMENTS	54
7.5.1 General.....	54
7.5.2 OAM Interface & Connectivity	55

7.5.3 <i>Management and Monitoring</i>	56
7.5.4 <i>Firmware and Security</i>	56
7.6 MANAGEMENT MODULE REQUIREMENTS	57
7.6.1 <i>Hardware Requirements</i>	58
7.6.2 <i>Firmware/Software Requirements</i>	59
7.6.3 <i>Security and Access Control</i>	59
7.6.4 <i>Root of Trust</i>	60
7.7 INTEGRATED SERVERS REQUIREMENTS.....	60
7.7.1 <i>Decision path</i>	61
7.8 SDNO CHASSIS.....	68
7.8.1 <i>Dual class-A Configuration</i>	68
7.8.2 <i>Single class-C Configuration</i>	71
7.8.3 <i>Upper layers</i>	73
7.8.4 <i>SDNO chassis Requirements</i>	75
7.9 FLW CHASSIS	80
7.9.1 <i>Acceleration HPM</i>	82
7.9.2 <i>FLW chassis Requirements</i>	83
8 CONCLUSION AND NEXT STEPS.....	85
9 APPENDIX.....	86
9.1 ACRONYMS AND ABBREVIATIONS.....	86
9.2 REFERENCES.....	87
9.3 POSITION TOWARDS A PROCESSOR SOCKET STANDARD.....	88

List of Figures

Figure 1 - OCP Infrastructure spending	11
Figure 2 - resource bottleneck or limitations.....	11
Figure 3 - OCP product split by 2028.....	13
Figure 4 - ecosystem of the HIGHER system and software stack.....	18
Figure 5 - CXL-base disaggregated memory	29
Figure 6 - OCP DC-MHS paradigm.....	33
Figure 7 - DC-SCM and HPM in DC-MHS chassis.....	34
Figure 8 - M-FLW chassis.....	34
Figure 9 - M-DNO chassis.....	35
Figure 10 - M-DNO chassis modularity.....	35
Figure 11 - M-SDNO chassis	36
Figure 12 - M-DNO Common chassis interval (CCI).....	36
Figure 13 - M-HPM demonstrated at OCP'24	37
Figure 14 - nvidia MGX HPM	38



Figure 15 - MGX & DC-MHS similarities and differences.....	38
Figure 16 - Rhea2-based HPM - SDNO Class C.....	40
Figure 17 - Coverage of SBMR by DC-SCI.....	42
Figure 18 - 12V remote power distribution architecture	46
Figure 19 - Battery backed voltage interface	47
Figure 20 - Rhea2 firmware/software stack - high-level view.....	48
Figure 21 - Rhea2 Hypervisor stack - high-level view.....	49
Figure 22 - Offloading kernels for acceleration in multi-devices context.....	49
Figure 23 - Rhea2-based HPM - SDNO Class A.....	51
Figure 24 - RHEA2-BASED HPM form factor decision path.....	61
Figure 25 - RHEA2-BASED HPM - SDNO Class A555, shadow mode.....	62
Figure 26 - shadow mode pros and cons.....	63
Figure 27 - M-FLW vs. SDNO Class C	64
Figure 28 - single SDNO chassis supporting either x2 class A335 or x1 class C335	65
Figure 29 - 2 chassis: SDNO dual class A335 / M-FLW	66
Figure 30 - HIGHER CHASSIS final set.....	67
Figure 31 - SDNO chassis with Host HPM and Accelerator HPM.....	68
Figure 32 - SDNO chassis with x2 Accelerator HPMs	69
Figure 33 - SDNO chassis with x2 Host HPMs	70
Figure 34 - SDNO chassis with x2 Accelerator HPMs combined with another chassis.....	71
Figure 35 - SDNO chassis with Class C Host HPM.....	72
Figure 36 - SDNO server - PCIe slot layer (informative)	73
Figure 37 - SDNO server - PSU and Storage Layer (informative)	74
Figure 38 - SDNO server - Rear View (informative).....	74
Figure 39 - SDNO server - Power Supplies	77
Figure 40 - SDNO server - Power Distribution Board.....	78
Figure 41 - NVMe storage cage and E1.S drives (Informative).....	80
Figure 42 - M-FLW chassis with 3rd Party M-FLW Host HPM	81
Figure 43 - FLW chassis – Second U with Accelerator HPMs.....	83



List of Tables

Table 1 - Classification of Requirements	8
Table 2 - MSI Reference Servers	17
Table 3 - SuperMicro Reference Servers	17
Table 4 - Asrockrack Reference Servers	17
Table 5 - Network Object Storage reference performances	21
Table 6 - Regular Storage Reference Performances	22
Table 7 - Key results from the HIGHER project and their expected technical readiness levels	32
Table 8 - Excerpt of minimum requirement for ARM System Ready servers	57
Table 9 - Acronyms and Abbreviations	87
Table 10 - References	87



1 Executive Summary

Building on top of outcomes from the EPI (ARM RHEA2) and EUPilot projects, HIGHER will adopt the Open Compute Project (OCP) Server family of standards to build processor modules for computation and acceleration, alongside a system security/control module, all operating with fully-featured operating systems and runtimes in project-designed OCP server mechanics.

To achieve this we must first define the high-level specifications for the HIGHER platform in support of the selected use-cases: Accelerated Data Processing, Infrastructure-as-a-Service, Platform-as-a-service, Remote CXL-based disaggregated memory. This document represents the initial steps in collecting the requirements that will inform the definition of the architecture and specification of the modules at component level. The document will serve as a baseline or point of reference, and guide the technical work over the course of the project.

We start by assessing the market as it stands so we can gain some perspective, in terms of market size, stakeholders, the competitive landscape and ultimately the potential scope for commercial exploitation. Market expectations are an important factor that must not be overlooked, since it will help us confidently position all outputs and key results of the project. The next step is to identify the key performance metrics, as well as the requirements relating to the computation module, acceleration module, and security/control module to ensure support for the selected use-cases: (i) accelerated data processing and analysis for converged Cloud and HPC platforms, (ii) Infrastructure-as-a-Service with standardized management and monitoring, (iii) Platform-as-a-Service facilitating large-scale data processing for ML inference and data analytics, and (iv) memory pool management at the server rack level, with access control safeguards aligned with maturing CXL standards.

Each requirement is assigned a unique identifier that will allow us to track progress over the course of the project. The requirements and performance metrics collected here will inform T2.2 “Specifications and Architecture Design”, and will set targets and KPIs for the evaluation phase of the project.

2 Context

2.1 Purpose of the document

This document represents the first step in defining measurable system requirements that will drive the design and specifications for the HIGHER platforms in support of (i) an OCP-Compliant processor module dedicated to computation and hosting 2 RHEA2 EPI chips, (ii) an OCP-Compliant processor module dedicated to acceleration and hosting EPAC 2.0 EPI chips or EUPilot chips and (iii) an OCP Compliant Management module, hosting a RISC-V processor inside an FPGA for server management, security, and control features.

This report is the first WP2 deliverable covering the work carried out in T2.1: Requirements and use cases refinement that is used to gather inputs from stakeholders relating to the four use-cases as well as input from external stakeholders. This report will inform T2.2: Specifications and Architecture Design, by defining a list of high-level specifications for the HIGHER platforms resulting in D2.2 due for submission in M6. We will also define here a list of target goals to be referenced during the evaluation phase of the project and for setting additional platform KPIs. This will inform T2.3: Verification Framework, resulting in three iterations (D2.3, D2.4 and D2.5) of the Verification Framework & Evaluation, due for release in M18, M24 and M30 respectively.

2.2 Requirements elicitation

The requirements are listed in tables made up of 4 columns:

ID	Description	Attribute	Status
UCR-HW-GEN-UIDE	Description of the requirement		

Column “ID”: Each requirement is identified by a unique identifier:

- 1st field is the source of the requirement: UCR stands for Use Cases and Requirements
- 2nd Field is the requirement type: either UC (Use Case), HW (Hardware), DIST (Software Distribution), SYS (System), ETH (ethernet), IO (I/O Peripherals) or ENV (Environmental)
- 3rd Field corresponds to the category of the requirement according to Table 1 below:
- 4th field is a unique number from 0000 to 99999

Req. Category	Interpretation
GEN	General
R2-1S-HPM	Single Socket Rhea2-based Host Processor Module
R2-2S-HPM	Dual Socket Rhea2-based Host Processor Module
EP-HPM	EUPilot based Host Processor Module
MM	Management Module
SDNO-CHS	SDNO Chassis
FLW-CHS	FLW Chassis
ACC	Acceleration Use Case
IAAS	IaaS Use Case
PAAS	PaaS Use Case
CXL	CXL-memory disaggregation Use Case

TABLE 1 - CLASSIFICATION OF REQUIREMENTS

Column “Description”: The literal description of the requirement

Column “Attribute”: Indicates the importance of the requirement:



- “M-Mandatory”: Those features are mandatory; Not supporting a mandatory requirement will seriously affect the project
- “R-Recommended”: Those features are nice to have, but will be implemented (or not) depending on the ratio benefits versus cost / resources / time to market
- “O-Optional”: refers to a feature that belongs to a set of options, one of which should be implemented

Column “Status”: Indicates the status of the requirement, and will be tracked over the project timeline:

- “Draft”: This requirement is in discussion and is subject to changes
- “In review”: this requirement is clearly defined and under review before validation
- “Validated”: The requirement will be implemented as described
- “Rejected”: The requirement will not be implemented



3 Market Considerations

3.1 Market size and growth estimates

The HIGHER project operates within a rapidly evolving technological landscape, where cloud services, open hardware standards, and emerging processor architectures are driving innovation and market growth. This section provides an analysis of the market size and growth estimates for two key areas relevant to the HIGHER project: 1) the cloud services market, with a focus on Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) and 2) the Open Compute Project (OCP) servers market. We also outline the potential impact of RISC-V as a disruptive force in the processor industry.

3.1.1 Cloud Services Market: IaaS and PaaS

The global cloud services market has experienced exponential growth over the past decade, driven by the increasing adoption of cloud computing across many diverse industries. Within this market, IaaS and PaaS have emerged as important cloud service models, enabling businesses to scale infrastructure and develop applications efficiently.

According to a report published by Gartner Research, Global end-user spending on public cloud services is projected to grow significantly, reaching \$723.4 billion in 2025, up from 595.7 billion in 2024. This growth underscores the expanding role of cloud computing. The market is expected to grow by 21.5% in 2025, fuelled by increasing demand for distributed, hybrid, cloud-native, and multi-cloud environments supported by cross-cloud frameworks. Gartner highlights that 90% of organizations will adopt a hybrid cloud approach by 2027, with data synchronisation across hybrid environments emerging as a significant challenge for implementing generative AI (GenAI). This reflects the growing importance of cloud services in supporting advanced AI workloads and ensuring seamless data management across diverse environments.

Some key factors fuelling this growth include:

- An ongoing shift from on-premises infrastructure to cloud-based solutions.
- Increasing adoption of hybrid and multi-cloud strategies by enterprises.
- A rise of data-intensive applications, such as AI, machine learning, and big data analytics.
- The need for standardised management and monitoring tools, as highlighted in the HIGHER use-cases.

The HIGHER project's focus on converged Cloud and HPC platforms, as well as its emphasis on IaaS and PaaS use-cases, aligns well with these market trends. By addressing the need for accelerated data processing, standardised management, and large-scale data analytics, HIGHER is well-positioned to make a valuable contribution targeting the growing demand for advanced cloud services.

3.1.2 Open Compute Project (OCP) Servers: Potential market impact

The Open Compute Project (OCP) is making a significant market impact on the data centre industry by promoting open-source hardware designs and standards. OCP-compliant servers are gaining traction due to their cost efficiency, scalability, and adaptability to modern workloads.

IDC's growth predictions for OCP infrastructure

International market research and consulting firm, IDC indicates that compute platform investments are driving significant growth across a broad set of solutions, and predicts the market will grow to \$73,513 billion by 2028, up from \$41,625 billion in 2024 (see Figure 1).

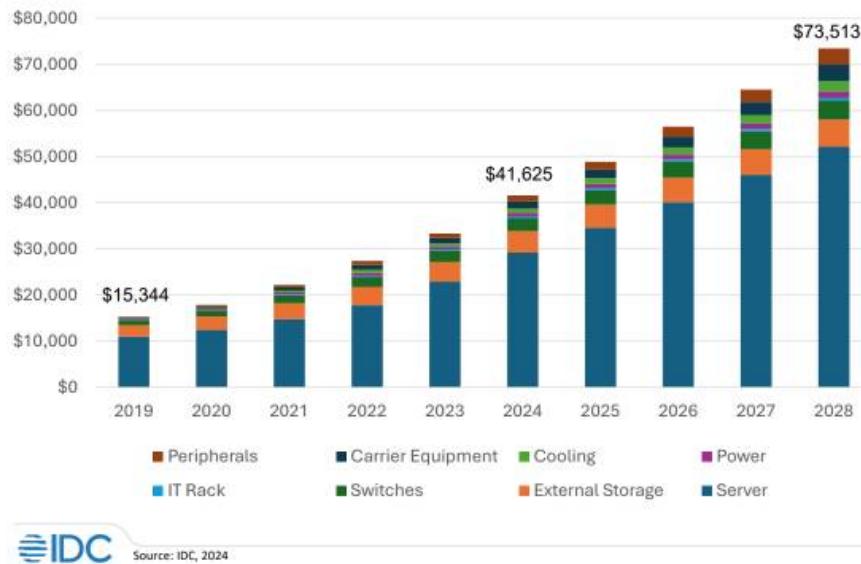


FIGURE 1 - OCP INFRASTRUCTURE SPENDING

IDC defines their predictions further into the following segments of the market:

- Server spending to rise to \$52.16 billion in 2028 (CAGR = 35.7%)
- External storage and switching equipment represent the next largest segment
- Categories showing the strongest spending growth include cooling solutions, with a CAGR of almost 46%, and peripherals (CAGR of almost 51%)

One key factor contributing to the growth of the OCP server market includes increased AI investments. IDC believes that AI investments will drive demand for OCP technology and states that an open-source approach to infrastructure development is needed to unlock the full potential of AI. According to IDC's Enterprise Infrastructure Pulse survey, which was published in April 2024, 62% of organisations stated they plan to increase their use of AI workloads by at least 10%. Additionally, 35% of respondents stated that they are currently using accelerated compute infrastructure within their on-premise environments, while 29% stated they intend to implement it within the next 12 months.

Respondents to the survey were also asked what their greatest resource bottleneck or limitation is with their on-premise compute/server infrastructure (see Figure 2).

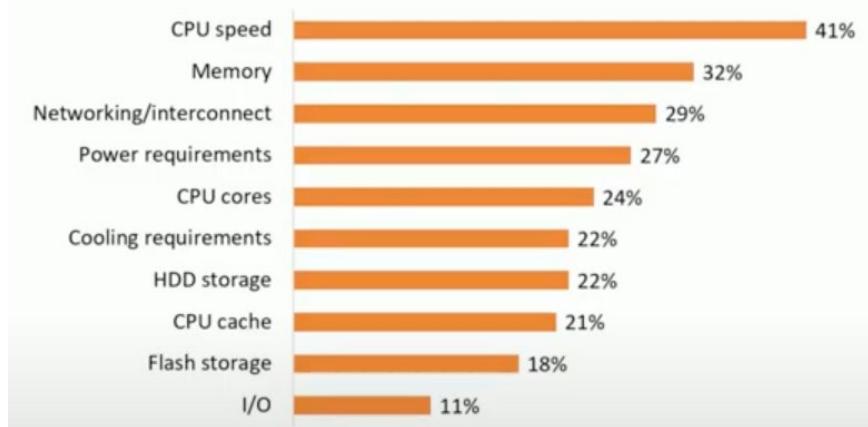


FIGURE 2 - RESOURCE BOTTLENECK OR LIMITATIONS



Over 40% of organisations identified CPU speeds as the biggest limitation for their on-premise compute infrastructure, with memory limitations and restrictions related to networking, power and cooling being other major challenges. OCP members are involved in various collaborative and innovative initiatives that play an important role in enabling AI infrastructure growth. With relevance to the HIGHER project, EUPilot aims to deliver a pre-exascale OCP-based set of HPC and AI accelerators designed and manufactured in Europe. HIGHER's adoption of OCP server standards for its processor modules and system architecture aligns with the industry's shift towards open hardware. By leveraging OCP standards, HIGHER can ensure compatibility, scalability and cost-effectiveness, making it an attractive solution for both hyperscalers and enterprises.

3.1.3 RISC-V as a potential driver force

As the cloud market continues to expand, the underlying hardware architecture that powers cloud infrastructure is becoming a more important factor in shaping its future. Among the emerging technologies in this space, RISC-V has the potential to revolutionise the cloud computing landscape by enabling innovation, reducing costs, and encouraging customisation. Unlike proprietary ISAs such as x86 and ARM, RISC-V is freely available for anyone to use, modify and implement. This openness encourages a collaborative ecosystem where companies, researchers, and developers can innovate without the constraints of licensing fees or proprietary restrictions.

In the context of cloud computing, this innovation can lead to customised hardware solutions with processors optimised to specific workloads, such as machine learning, data analytics, or video processing. This level of customisation can optimise performance and energy efficiency, giving Cloud Service Providers (CSPs) a competitive edge. Cost efficiency is a foundation of cloud computing, and RISC-V has the potential to significantly reduce hardware costs. By eliminating licensing fees and enabling optimised processor designs, CSPs can lower their capital expenditure (CapEx) and operational expenditure (OpEx). This cost reduction can make a CSP more competitive by making cloud services more affordable and accessible.

Market growth predictions

In 2023, the RISC-V market was valued at \$1 billion, with projections indicating a CAGR of 30-40% over the next five years. By 2030, RISC-V is expected to capture 10-15% of the global processor market, driven by its adoption in IoT, AI and data centre applications (ref. RISER D2.1 + <https://www.gartner.com/>)

3.2 Industry adoption

3.2.1 The Open Compute Project (OCP) Adoption

Founded in 2011, The Open Compute Project (OCP) is an organization set up to support data centre product designs and the operation of large-scale computing facilities, as well as to encourage the sharing of industry best practices. Adherence to Open Compute Project (OCP) standards ensures compatibility across multiple OCP servers, chassis, and vendors, enhancing wide adoption and market penetration. At the time of writing, OCP has over 400 members worldwide. Some key contributors and adopters include: Arm, Meta, IBM, Intel, Nokia, Google, Microsoft, Dell, Rackspace, HPE, NVIDIA, and Cisco.

There are clear benefits of OCP adoption. By leveraging open-source designs and standardised components, organisations can significantly reduce hardware costs. OCP's collaborative model eliminates vendor lock-in and promotes competition, driving down prices. OCP-compliant hardware is designed with energy efficiency in mind, reducing power consumption and operational costs. Innovations in cooling systems and power distribution have further enhanced sustainability. OCP's modular designs allow organisations to scale their infrastructure efficiently. The flexibility to customise hardware based on specific workloads has made OCP particularly appealing to cloud service providers

and data centres. Additionally, OCP encourages a culture of collaboration, enabling companies and research institutions to share knowledge and accelerate innovation. The open-source nature of OCP encourages continuous improvement and rapid adoption of new technologies.

As stated in the previous section, IDC Research found that the purchases of OCP-recognized equipment are projected to exceed \$73.5 billion by 2028.

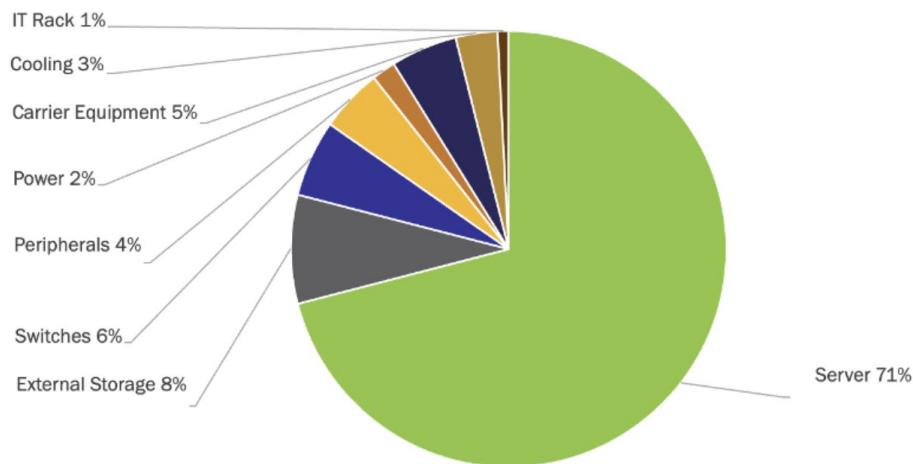


FIGURE 3 - OCP PRODUCT SPLIT BY 2028

As evident from the diagram above, product adoption is underpinned by server sales which IDC noted has a 35.7% CAGR over the period. Cooling and peripherals were the fastest-growing segments at 46% and 51% CAGR respectively.

The HIGHER project will develop innovative, predominantly open-source hardware and software modules and systems, leveraging technologies advanced in projects such as EPI and EUPilot, as well as standards established by the Open Compute Project (OCP). These developments will enable the deployment of comprehensive, modular open cloud and edge infrastructures. Additionally, HIGHER will assess adoption a socket-based approach for EPI and EUPilot processor chips, aiming to create an EU socket reference. This initiative will facilitate interoperability and seamless processor upgrades, further enhancing system flexibility and scalability.

3.2.2 RISC-V Adoption

RISC-V, an open-source instruction set architecture (ISA), has gained significant momentum since its introduction in 2010 at the University of California, Berkeley. In 2019, it was transferred to RISC-V International, a Swiss non-profit entity. Unlike proprietary ISAs such as ARM and x86, RISC-V offers a modular, extensible, and royalty-free framework, making it an attractive option for organizations seeking customization and cost efficiency. The RISC-V ecosystem has seen rapid growth, with adoption spanning industries such as semiconductors, IoT, artificial intelligence (AI), and edge computing. Major technology companies, startups, and academic institutions are increasingly embracing RISC-V to drive innovation and reduce dependency on proprietary architectures. Some key contributors and adopters (as RISC-V International members) include: SiFive, Andes Technology, NVIDIA, Synopsys, Alibaba's Damo Academy, as well as numerous academic and research institutions worldwide.

The benefits of RISC-V adoption include cost efficiency, a high level of customisation and flexibility, and reduced vendor lock-in. RISC-V eliminates licensing fees, reducing the cost of developing and deploying hardware. This is particularly beneficial for startups and SMEs. The modular design of RISC-V allows organizations to tailor the ISA to specific use cases, optimizing performance and power



efficiency. RISC-V's open-source nature encourages collaboration and innovation, enabling rapid advancements in hardware and software development. By adopting RISC-V, companies can avoid dependency on proprietary ISAs, gaining greater control over their technology stack.

The RISC-V Software Ecosystem (RISE) Project is a newly formed global consortium of industry leaders dedicated to accelerating the adoption of RISC-V by enhancing its software ecosystem. Led by Google's Amber Huffman, the initiative aims to improve the availability of software for high-performance, power-efficient RISC-V cores by coordinating efforts across various companies. Key focus areas include compilers, system libraries, virtualization, Linux integration, and debugging tools. Major tech players such as Google, Intel, NVIDIA, and Qualcomm have already invested in the project, either financially or through engineering contributions, signalling strong industry commitment.

RISE is positioned as a complementary effort to RISC-V International, helping eliminate roadblocks to adoption while encouraging collaboration. The consortium sees potential in expanding RISC-V across industries such as data centres, AI, robotics, and automotive, with mobile and consumer applications driving early development. Given the increasing demand for specialized processing RISC-V is emerging as a competitive open-source ARM and x86 alternative. Tirias Research considers RISE a crucial step toward mainstream adoption, as strengthening the software ecosystem will be essential for RISC-V's long-term success. (ref. <https://riseproject.dev/>)

3.3 Competitive landscape

The development of the HIGHER platforms occurs within a highly competitive landscape where multiple industry players and research initiatives are advancing processor modules, accelerators, and management solutions. This section provides an overview of existing solutions that align with the objectives of the HIGHER platforms, highlighting key differentiators, industry trends and areas where HIGHER can provide a competitive advantage.

3.3.1 OCP-compliant compute and accelerator modules

The Open Compute Project (OCP) has driven the standardisation of high-performance computing modules, capable of supporting an ecosystem of interoperable and energy efficient components. Existing processor modules include:

- Intel Xeon and AMD EPYC OCP Modules – These solutions currently dominate the hyperscale computing market with high core counts, advanced security features and well established software ecosystems. However, their proprietary architectures limit flexibility for customisation.
- NVIDIA GH200 Grace Hopper Superchip – Designed for AI acceleration, this module integrates CPU-GPU technology, competing in the same space as EPAC 2.0 and EUPilot chips, but with a distinct focus on AI workloads.
- SiPearl RHEA Processor – As a key component of the European Processor Initiative (EPI), RHEA aims to deliver high-performance and energy-efficient compute capabilities, positioning itself strategically as an European alternative to non-EU-based solutions.

The HIGHER platform distinguishes itself by integrating RHEA2 and EPAC 2.0/EUPilot chips within an OCP-compliant design, optimising for computational and acceleration use cases while maintaining an open, modular architecture.

3.3.2 Server management and security solutions

RISC-V-based management modules are gaining traction in the data centre landscape as an alternative to proprietary Baseboard Management Controller (BMC) solutions. Existing competitors include:



- OpenBMC – A widely adopted open-source firmware stack for server management, supported by companies like Microsoft, Intel, IBM and Google
- Proprietary BMCs (ASPEED, Nuvoton, AMD) - These dominate traditional server infrastructure, offering robust remote management and security features but with limited customisation and a degree of vendor lock-in.

The HIGHER platform's management module leverages an FPGA-hosted RISC-V processor, offering flexibility for customization, enhanced security features and compliance with open hardware standards.

3.3.3 Industry trends and differentiation

Several industry trends influence the competitive positioning of the HIGHER project, including:

- Energy efficiency and sustainability – With increased emphasis on power efficiency, OCP-compliant platforms must optimise energy consumption. HIGHER designs will integrate efficiency enhancing technologies to meet sustainability goals.
- Open hardware and security – The shift toward open-source architectures is driving demand for RISC-V solutions and transparent security implementations. The HIGHER approach aligns with these trends by promoting open standards.
- AI and HPC acceleration – Competing solutions such as NVIDIA's AI-focused platforms and Google's TPUs dominate AI acceleration. HIGHER's EPAC 2.0-based accelerators provide an alternative optimised for European AI and HPC workloads.

By leveraging OCP compliance, open-source principles, and European-developed processors, the HIGHER platform aims to deliver a competitive, scalable, and secure computing solution tailored for next generation data centre and edge computing environments.

3.4 Stakeholder Definition

The development of the HIGHER platform involves a diverse set of stakeholders who play a crucial role in defining requirements, guiding design choices and ultimately driving adoption. These stakeholders are comprised of industry leaders, academic research institutions, policy makers and end-users. By engaging with these stakeholder groups, the HIGHER project ensures that its development aligns with industry needs, regulatory requirements, and technological advancements, facilitating broad adoption across various levels of the industry. We define the stakeholder groups as follows:

The core technical stakeholders - These stakeholders are directly involved in the technical design, development and integration of the HIGHER platform. Their primary focus is ensuring compliance with OCP standards, while optimising performance, security and energy efficiency. 1) The hardware developers include companies and research teams designing the processor modules, accelerators and management modules based on RHEA2, EPAC 2.0 and EUPilot chips. 2) The software and firmware developers include engineers working on firmware, system software and management tools, ensuring compatibility with industry standards such as OpenBMC and RISC-V ecosystems. 3) System integrators include organisations responsible for assembling and deploying HIGHER-based solutions in cloud and HPC environments, ensuring seamless integration with existing infrastructure.

End user and market stakeholders - End-users are critical stakeholders whose business and technical needs influence the platform's requirements and adoptions. Cloud Service Providers (CSPs) require high-performance scalable and energy-efficient computing solutions, while organisations looking to deploy cost-effective, open-standard computing infrastructure require enhanced security and flexibility. Other end-users include institutions focused on AI and high-performance computing applications that benefit from OCP-compliant accelerator solutions.

Policy and regulatory stakeholders - Regulatory bodies and policy makers play a crucial role in defining compliance requirements and influencing adoption through funding and standardisation efforts. This includes: 1) Entities supporting the European Processor Initiative (EPI), promoting sovereignty in



computing through funding and policy initiatives. 2) The governing body responsible for maintaining OCP standards and ensuring the interoperability of HIGHER with a broader OCP ecosystem. 3) Organisations ensuring that the platform meets security and privacy standards, particularly in sensitive applications such as government and healthcare.

External stakeholders and ecosystem partners - Beyond direct developers and end-users, several external stakeholders influence the success of the HIGHER platform. Semiconductor manufacturers, FPGA providers and cooling solutions vendors contribute essential hardware components. Universities and research labs conducting research on processor architectures, AI acceleration and open-source computing frameworks play a significant role, as do developers and contributors working on firmware, system software and security frameworks that enhance the capabilities of the HIGHER platforms.



4 Reference servers

Here is a list of reference x86, Arm and RISC-V based servers whose performance on the acceleration, IaaS and PaaS use cases will be used as a comparison for HIGHER server. Ideally, such reference servers shall be OCP Compliant.

4.1.1 Arm-based reference servers

- AWS Graviton3/3E/4 (Amazon EC2 C7g/C7gn/X8g bare-metal instances)
- Ampere One from SuperMicro: [ARS-211M-NR](#)
- NVIDIA Grace-Grace from SuperMicro: [ARS-221GL-NR](#)

4.1.2 x86 reference servers

The MSI product line is based on Intel Xeon 6700 or AMD EPYC 9005.

Model	HPM Form Factor	SoC
D3071	DNO Type2	Intel Xeon 6700
D3066	DNO Type 4	Intel Xeon 6700
D5062	M-FLW	X2 Intel Xeon 6700
D4051	DNO Type2	AMD EPYC 9005

TABLE 2 - MSI REFERENCE SERVERS

SuperMicro has the following references:

Model	HPM Form Factor	SoC
SYS-112C-TN	SDNO Class A	Intel Xeon 6700
AS-1116CS-TN	DNO Type 2	AMD EPYC 9005

TABLE 3 - SUPERMICRO REFERENCE SERVERS

Asrockrack also has a product line based on OCP DC-MHS specification :

Model	HPM Form Factor	SoC
GNRAPD12DNO	DNO Type 2	Intel Xeon 6900P
GNRD16DNO	DNO Type 2	Intel Xeon 6700P
GNR2D32FLW	M-FLW	X2 Intel Xeon 6700P

TABLE 4 - ASROCKRACK REFERENCE SERVERS

ASRockrack also proposes dual socket MGX form factor with x2 Intel Xeon: [GNR2D32MGX](#)

4.1.3 RISC-V Platforms

- Banana Pi BPI-F3 (V-ISA 1.0)
- Milk-V Pioneer (V-ISA 0.7.1)
- SiFive P550 (no V-ISA)

5 Use Cases

5.1 Accelerated Data Processing

The use case of accelerated data processing will demonstrate the advantage of the HIGHER platform and its software stack on accelerating use case performance on different cloud servers utilizing ARM and RISC-V host CPUs and acceleration in ARM scalable vector units and RISC-V vector units. It will also demonstrate cost efficiency improvements by enhancing resource utilization and energy improvements from offloading data processing close to the data source on the edge. The use case will use a suite of applications comprised of memory-intensive workloads in CloudSuite, machine learning training and inference in MLPerf, and a molecular docking application for drug discovery.

We will develop a bottom-up approach to evaluate the use case by starting with a set of representative micro kernels extracted from these applications, such as scoring function for molecular docking, graph analytics, and ML inference, and ported to be able to run and accelerate on the available ARM and RISC-V platforms. For this task, we will leverage profiling for identifying important kernels, porting them and validating the compiler toolchain such as GCC and LLVM for auto-vectorization for utilizing vector units. We will validate these implementations on the HIGHER software stack, including OS, Kubernetes, Meta-OS and other relevant library and toolchain support. Figure 4 below illustrates the envisioned use case in the ecosystem of the HIGHER system and software stack.

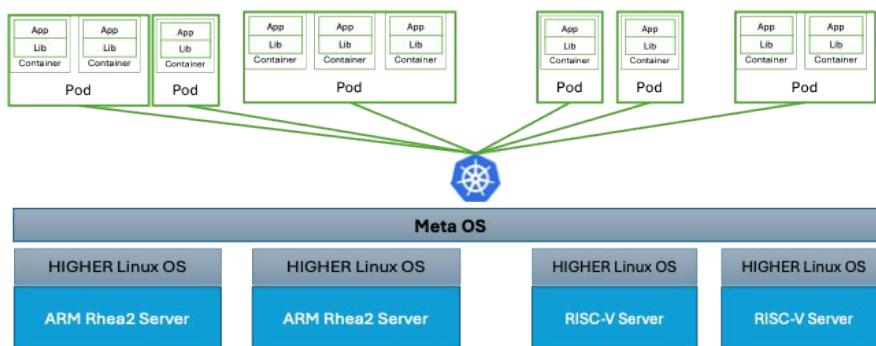


FIGURE 4 - ECOSYSTEM OF THE HIGHER SYSTEM AND SOFTWARE STACK

We will measure metrics in performance, cost, and energy to quantitatively assess the efficiency of the HIGHER stack. Currently we target to assess the execution time on different platforms, relative speed-ups and vectorization ratios. These measurements will be compared with the hardware peak performance to provide a normalized comparison basis across different hardware, e.g., the achieved utilization of the peak floating point computing throughput. For quantifying cost and energy efficiency, depending on the maturity of performance hardware counter support on the ARM and RISC-V platforms, we may also develop analytical models to quantify the metrics.

ColonyOS, as an open-source meta-operating system designed to facilitate seamless execution of computational workloads across diverse platforms, including cloud, edge, and high-performance computing (HPC) systems, will be utilised for data processing and analysis for converged platforms combining cloud and edge processing.

ColonyOS's architecture comprises Colonies servers and Executors. Users submit specification of computational tasks to Colonies servers, which then assign tasks to Executors based on specified conditions. These Executors, functioning as microservices, can operate across various platforms, including Kubernetes clusters, HPC systems, and edge devices.



By integrating ColonyOS into HIGHER platform, we aim to achieve a modular, hyperconverged infrastructure capable of deploying containerized applications that leverage the full spectrum of computing resources, from powerful cloud servers to resource-constrained edge devices.

The requirements identified for this use case are reported below:

ID	Description	Attribute	Status
UCR-UC-ACC-52000	The HIGHER project shall develop a use case that accelerates performance on ARM platform with improved utilization of resources	M	Validated

ID	Description	Attribute	Status
UCR-UC-ACC-52001	The HIGHER project shall develop a use case that accelerates performance on RISC-V platform with improved utilization of resources	M	Validated

ID	Description	Attribute	Status
UCR-UC-ACC-52002	The HIGHER project shall develop a use case that reduces energy cost by offloading processing to data on the edge	M	Validated

Note: This use case involves ColonyOS

ID	Description	Attribute	Status
UCR-UC-ACC-52003	The HIGHER project shall develop a use case that improves resource utilization by mapping to suitable computing and memory resources	M	Validated



5.2 Infrastructure as a Service (IaaS)

The HIGHER project will develop and validate a use case focusing on demonstrating performance and efficiency parity between mainstream x86, ARM and RISC-V-based cloud server platforms.

The use case will focus on benchmarking and optimization to achieve better performance metrics when compared to established x86-based cloud server platforms, particularly in scenarios where I/O operations are critical, such as storage-intensive applications.

CLOUDSIGMA's proprietary cloud computing platform will be used during the evaluation and validation phase to compare the performance results of the HIGHER platform with a fully featured commercial public cloud infrastructure.

For the sake of consistency, the IaaS use cases as defined for [RISER] project will be incorporated as foundational use cases for the [HIGHER] project, as follows:

5.2.1 Networked Object Store Use Case

The HIGHER project will develop and validate a use case focusing on networked key/value object stores for fast data ingest and retrieval.

The starting point for this use case is to establish an evaluation baseline at the cloud node level (with ARM and x86 hosts), using the YCSB data access workload generator and parts of the CloudSuite benchmark collection (esp. data serving and data caching). The end-goal is to demonstrate performance and efficiency parity (or better) with mainstream x86 cloud server platforms, for I/O-intensive workloads - serving at least 1000 object access (read/write) requests per second per core.

Informative options for the implementation of the storage backend include:

- The MinIO object store, often used with Kubernetes, to enable applications to store and retrieve data objects of any size and format.
- The memcached memory object caching system, to enable in-memory caching of variable-length data, as a means to speed up applications by minimizing reads and writes to persistent storage systems.
- Optimized key/value stores, such as Kreon.

ID	Description	Attribute	Status
UCR-UC-IAAS-52100	For the Networked Object Store use case(s), the API to be provided shall be compatible with the APIs of the Kubernetes runtime and AWS's S3 object storage service, thus facilitating deployment in public, private and hybrid (private + public) cloud environments	M	Validated

ID	Description	Attribute	Status
UCR-UC-IAAS-52101	For the Networked Object Store use case(s), the key-value store shall provide a single global namespace and a consistent object storage interface (accessible from multiple cloud providers, including hosted, on-premise, and edge infrastructure)	M	Validated



ID	Description	Attribute	Status
UCR-UC-IAAS-52102	For the Networked Object Store use case(s), the key-value store shall demonstrate performance and efficiency parity (or better) with mainstream x86 cloud server platforms, for I/O-intensive workloads - serving at least 1000 object access (read/write) requests per second per core	M	Validated

CloudSigma offers S3 compatible all HDD MinIO based object storage solutions. Below are test results from one of their production cluster and that is what would be expected per node in an all HDD MinIO cluster with NVMe used for caching:

Test	Throughput	IOPs
PUT	258 MiB/s	4 objs/s
GET	9,664 MiB/s	151 objs/s

TABLE 5 - NETWORK OBJECT STORAGE REFERENCE PERFORMANCES

Node hardware and test configuration:

- Test used - mc support perf
- Object size - 64MiB
- CPU - 2 x Intel Xeon E5-2697 v4 @ 2.30GHz 16 cores
- HDDs - 90 x HGST HUH728080AL5204 8T
- NVMe for caching - 6 x SAMSUNG MZ7KM1T9HAJM 2T
- Memory - 512G DDR4-2666

5.2.2 Regular Storage Use Cases

When deploying new storage clusters, CloudSigma internally performs multiple fio tests (Flexible IO Tester Benchmark - OpenBenchmarking.org) to validate the performance of the underlying storage. The following are the expected base results for the different tests performed on an all NVMe storage and the fio configuration used for the tests:

Test	Configuration used	results
lat-r-4k-1	norandommap randrepeat=0 rw=randread bs=4k iodepth=1 runtime=10s	0.044 ms
lat-w-4k-1	norandommap randrepeat=0 rw=randwrite bs=4k iodepth=1 runtime=1m	0.012 ms
rand-r-4k-64	norandommap randrepeat=0 rw=randread	274077 IOPS



Test	Configuration used	results
	bs=4k iodepth=64 runtime=1m	
rand-rw-4k-64	norandommap randrepeat=0 rw=randrw bs=4k iodepth=64 runtime=1m	257917 IOPS
rand-w-4k-64	norandommap randrepeat=0 rw=randwrite bs=4k iodepth=64 runtime=1m	256186 IOPS
seq-r-1M-64	rw=read bs=1M iodepth=64 runtime=1m	7059 MB/s
seq-w-1M-64	rw=write bs=1M iodepth=64 runtime=1m	3839 MB/s

TABLE 6 - REGULAR STORAGE REFERENCE PERFORMANCES

All the test cases configurations above have the following parameters in common:

- ioengine=libaio
- direct=1
- sync=0
- time_based

Hardware used for the tests:

- Storage: HPe VO007680KYDNA 8T NVMe
- CPU: AMD EPYC 7282 16-Core
- Motherboard/Server: HPE ProLiant DL325 Gen10 Plus

It is expected that the performances above are minimally met by the HIGHER server in its nominal configuration, as defined in chapter 7.8.

ID	Description	Attribute	Status
UCR-UC-IAAS-52200	For the Regular Storage use case, the key-value store shall demonstrate performance and efficiency parity (or better) with mainstream x86 cloud server platforms, for I/O-intensive workloads. The performance enumerated in Table 2 shall be met or overpassed	M	Validated



5.2.3 Container Use Case

The HIGHER project will develop and validate a use case for the HIGHER server in nominal configuration, as defined in chapter 7.8, focusing on containerized execution of user-developed micro-services encapsulated within Linux containers.

ID	Description	Attribute	Status
UCR-UC-IAAS-52300	For the Containers use case(s), HIGHER shall provide a port of the Kubernetes container orchestration engine for automating deployment, scaling, and management of containerized applications	M	Validated

ID	Description	Attribute	Status
UCR-UC-IAAS-52301	For the Containers use case(s), HIGHER server shall enable a combination of Kubernetes-managed computing infrastructure with networked object storage and provide a serverless runtime environment based on Kubernetes-managed containers.	M	Validated

The starting point for this use case is to establish an evaluation baseline at the cloud node level (ARM and x86 hosts), using the K-Bench workload benchmark for Kubernetes and the MLBench distributed machine learning benchmark. The end-goal is to demonstrate performance and efficiency parity (or better) with mainstream x86 cloud server platforms, for I/O-intensive workloads that also includes compute- and memory-intensive phases. Specific performance goals to consider include: (i) At least 50% improvements over both control-plane and data-plane metrics of K-Bench; and (ii) At least 30% improvements over runtime of MLBench for binary-classification models.

ID	Description	Attribute	Status
UCR-UC-IAAS-52302	The HIGHER server shall demonstrate at least 50% performance improvements over both control-plane and data-plane metrics of K-Bench, compared with mainstream x86 cloud server platforms	M	Validated

ID	Description	Attribute	Status
UCR-UC-IAAS-52303	The HIGHER server shall demonstrate at least 30% improvements over runtime of MLBench for binary-classification models, compared with mainstream x86 cloud server platforms	M	Validated

The software implementing this use case will be integrated in an easy-to-deploy software distribution and will be demonstrated in a representative environment of hosted cloud services. HIGHER partner CloudSigma will integrate HIGHER platforms in their infrastructure-as-a-service (IaaS) environment, integrated with standard system management tools for fault and performance monitoring.

5.3 Platform as a Service (PaaS)

The HIGHER project will develop and validate a use case focusing on demonstrating a complete development and deployment environment supporting large-scale data processing for HPC, ML inference and data analytics, including utilization of multiple both ARM and RISC-V host systems and RISC-V accelerators.

The PaaS environment developed on top of the HIGHER platform will support the entire lifecycle of applications, including build/development, test, deployment, run-time management, and updates.

The actual goals of the use case are twofold:

1. a first goal to explore and enhance the efficient offloading of memory-intensive operations, commonly found in cloud computing workloads, to the attached devices. That specific target will continue the work already started in the RISER project in terms of the RISER's Acceleration Use Case.
2. the second goal is to develop several domain specific platforms targeting HPC, ML/AI and cloud workloads.

5.3.1 Multi-device approach

To fully exploit the capabilities of the multi-devices, HIGHER will continue the work already started in the RISER project and adopt the reached status in terms of the hierarchical parallelization strategy.

In the OpenMP language, device programming is typically initiated on the host system through the target directive, enabling to offload blocks of code to the associated accelerators. The directive also allows variable mappings between host and device memory environments. In addition, the target directive can be further extended with a hierarchical distribution of the work using leagues (set of teams) and teams (set of threads) as additional layers of parallelism.

In the RISER project, we proposed a new level of parallelism based on the spread construct. This first level of parallelism allows to distribute work among different devices (instead of only one, as described in the current standard). Based on the actual extend the RISER project will reach at its end, the HIGHER project will:

- Extend the set of results of the existing benchmark to complete the evaluation carried out by the former project.
- Extend the set of benchmarks to further evaluate new aspects of the OpenMP extension.
- Porting the existing benchmarks to a new OpenMP functionality developed under the umbrella of the HIGHER project. New functionalities could involve: better scheduler policies, support of new mapping techniques allowing to explore device-to-device communication, coherent data mapping among different constructs, etc.

The final direction adopted by this project will be determined once the extension of the RISER project, in terms of the OpenMP multi-devices approach, will be clarified. The multiple options considered in that particular aspect of the task respond to the need of avoiding, in all the cases, the double funding risk and clearly identify the boundaries between both projects.

In order to develop that specific use case, the software stack must offer an OpenMP infrastructure able to offload data and code to different devices. This infrastructure relies on the corresponding point-to-point (i.e., host and device) system services:

ID	Description	Attribute	Status
UCR-SW-DIST-53100	The HIGHER project shall provide an implementation of the OpenMP framework (compiler and runtime) on the HIGHER multi-device platform.	M	Validated



ID	Description	Attribute	Status
UCR-SW-SYS-53101	The HIGHER project shall provide a tuple of driver (host) and service (device) able to support the OpenMP offloading on the HIGHER multi-device platform.	M	Validated

The detailed technical requirements identified to guide the implementation and evaluation of this use case are:

ID	Description	Attribute	Status
UCR-UC-PAAS-53102	The HIGHER project shall continue the development of the RISER's multi-devices use case accelerating the performance on the HIGHER multi-device platform with improved functionalities.	R	Validated

5.3.2 HPC domain approach

The main objective of the HPC domain approach is to build a containerized system able to execute HPC workloads. The target platform as a service will enable a linear algebra library capable of exploiting the underlying hardware resources. The use case will rely on the Infrastructure as a Service foundations.

The detailed technical requirements identified to guide the implementation and evaluation of this use case are:

ID	Description	Attribute	Status
UCR-UC-PAAS-53200	The HIGHER project shall develop a use case that enables HPC workloads on the HIGHER RISC-V platform. The HPC services will target linear algebra optimized libraries.	M	Validated

ID	Description	Attribute	Status
UCR-UC-PAAS-53201	The HIGHER project shall develop a use case that enables HPC workloads on the HIGHER ARM platform. The HPC services will target linear algebra optimized libraries.	R	Validated

A complementary aspect of the HPC approach will be to leverage a native Message Passing Interface (MPI) implementation. That imposes a software requirement and an additional use case requirement. The detailed software requirement and technical requirement identified to guide the implementation of this use case are:

ID	Description	Attribute	Status
UCR-SW-DIST-53202	The HIGHER project shall provide an implementation of the Message Passing Library (MPI) on the HIGHER platforms.	R	Validated

ID	Description	Attribute	Status
UCR-UC-PAAS-53203	The HIGHER project shall extend the PaaS use case that enables HPC workloads to leverage the	R	Validated



ID	Description	Attribute	Status
	native MPI implementation on the HIGHER platforms.		

5.3.3 ML/AI domain approach

The main objective of the ML/AI domain approach is to build a containerized system able to execute ML/AI workloads. The target platform as a service will enable a data analytic framework capable of exploiting the underlying hardware resources. As in the previous approach, the use case will rely on the Infrastructure as a Service foundations.

The detailed technical requirements identified to guide the implementation and evaluation of this use case are:

ID	Description	Attribute	Status
UCR-UC-PAAS-53300	The HIGHER project shall develop a use case that enables ML/AI workloads on the HIGHER RISC-V platform. The ML/AI services will target data analytic frameworks.	M	Validated

ID	Description	Attribute	Status
UCR-UC-PAAS-53301	The HIGHER project shall develop a use case that enables ML/AI workloads on the HIGHER ARM platform. The ML/AI services will target data analytic frameworks.	R	Validated

5.3.4 Cloud domain approach

The main objective of the cloud domain approach is to build a containerized system able to execute commonly used cloud services. The target platform as a service will enable a web server that can be used as a front-end application. As in the previous approach, the use case will rely on the Infrastructure as a Service foundations.

The detailed technical requirements identified to guide the implementation and evaluation of this use case are:

ID	Description	Attribute	Status
UCR-UC-PAAS-53400	The HIGHER project shall develop a use case that enables cloud fundamental services on the HIGHER RISC-V platform. The cloud fundamental services will target a web server.	M	Validated

ID	Description	Attribute	Status
UCR-UC-PAAS-53401	The HIGHER project shall develop a use case that enables cloud fundamental services on the HIGHER ARM platform. The cloud fundamental services will target a web server.	R	Validated

In addition to have an isolated web server, we can further extend the service by combining it with HPC and ML/AI workloads. The optional technical requirements identified to guide the implementation and evaluation of this use case are:



ID	Description	Attribute	Status
UCR-UC-PAAS-53402	The HIGHER project shall develop a use case that enables cloud fundamental services (front-end) combined with HPC workloads (back-end) on the HIGHER platforms.	R	Validated

ID	Description	Attribute	Status
UCR-UC-PAAS-53403	The HIGHER project shall develop a use case that enables cloud fundamental services (front-end) combined with ML/AI workloads (back-end) on the HIGHER platforms.	R	Validated

5.3.5 Evaluation Approaches

The evaluation of this offload use case will follow two different approaches. For the multi-device approach, we will continue the evaluation carried out in the RISER project on top of the multi-device platform. For the domain specific use cases, we will develop the expected domains on top of the task Infrastructure as a Service (IaaS). Both approaches target similar scenarios (i.e., HPC-, ML/AI- and cloud- domains), although the implementation of these benchmarks could be different due to considering other parallelization approaches (e.g., exploiting the host parallelism instead of offloading kernel to devices).

The multi-device evaluation will leverage the existing kernels provided as results in the RISER project. The evaluation will be adapted to one of the three different directions included in the Multi-device approach section.

The HPC domain evaluation will leverage the provided HPC frameworks from the use case implementation and will port representative benchmarks to exploit them. An example of this evaluation could be to exploit the provided linear algebra library to execute a benchmark leveraging these services (e.g., High-Performance Linpack, HPL).

The ML/AI domain evaluation will leverage the provided ML/AI frameworks from the use case implementation and will port representative benchmarks to exploit them. As an example, we can consider the exploitation of a tensor library to execute models leveraging these services (e.g., Yolo).

The Cloud domain evaluation will leverage the cloud environment from the use case implementation to fully develop a proof-of-concept based on these tools. An example could be the deployment of a web page on top of a web server.

As a final stage of evaluation, the project will also consider to combine the previous domains to offer a multi-layer solution with a web interface (front-end) interacting with the HPC or ML/AI domains (back-end). We classify these final requirements as optional due to the limited amount of efforts of this task.

In all the previous cases we will follow an iterative validation and refinement approach at increasing levels of complexity. The process begins with the development and adaptation of small micro-kernels and will continue with representative benchmarks per domain.

To assess the effectiveness of the performance strategy, several metrics will be used:

- **Execution Time:** A primary metric for performance comparison, enabling computation of speed-up ratios across versions.
- **Percentage of Peak Performance:** Especially relevant for compute-bound kernels, this measures how closely an implementation approaches the theoretical peak of the hardware (e.g., floating-point operations per second).



- **Profile- and Trace- based analysis:** A more detailed performance analysis could be deployed according to the availability of tools. That evaluation will allow to focus the analysis in certain regions of the code and explore additional metrics. The proposed tool to carry out this analysis is Extrae.

In order to carry out the third level of this proposed analysis methodology, the evaluated use cases should provide an additional tool component enabling the generation of traces.

The optional technical requirements identified to guide the performance evaluation of this use case is:

ID	Description	Attribute	Status
U C R - U C - P A A S - 53500	The HIGHER project shall extend the HPC and ML/AI approaches to support the Extrae library on the HIGHER platforms.	R	Validated

We decided to have Extrae available as part of the platform software stacks because it is a very flexible tool enabling different approaches: summary profile tables, time-line exploration and efficiency metrics. Extrae could be enabled with different configurations depending on the available tool chain. For instance, having PAPI supported by the platform, will allow Extrae to show hardware counters.

This structured approach will provide a solid foundation for demonstrating the benefits of the Platform as a Service in representative workloads, while also informing further development in compiler support, framework availability, and performance tuning strategies.

The obtained performance metrics on the HIGHER platforms will be compared with other reference platforms in order to clearly identify their specific behavioural aspects.

5.4 Remote CXL-based disaggregated memory

HIGHER will develop a memory pool that will be available to HPMs over CXL-based links. As shown in the image below, HPMs with VMs can ask the DC-SCM to allocate memory regions in the memory pool. The latter will be controlled by the CXL memory Pool Manager (CPM), and will (i) provide safeguards towards ensuring that memory regions are isolated and protected against illegal accesses, (ii) support dynamic memory allocation for HPMs to follow time-varying workload needs, and (iii) enable optional user-configurable Near-Data Processing (NDP). In terms of security, the envisioned threat model considers cases where, although the DC-SCM will be responsible for the valid memory slice allocation to VMs, adversaries may have physical access to the disaggregated memory pool, being capable to tamper such requests. Therefore, the CPM will act as an additional protection layer that monitors ingress transactions and checks if they are valid and legal.

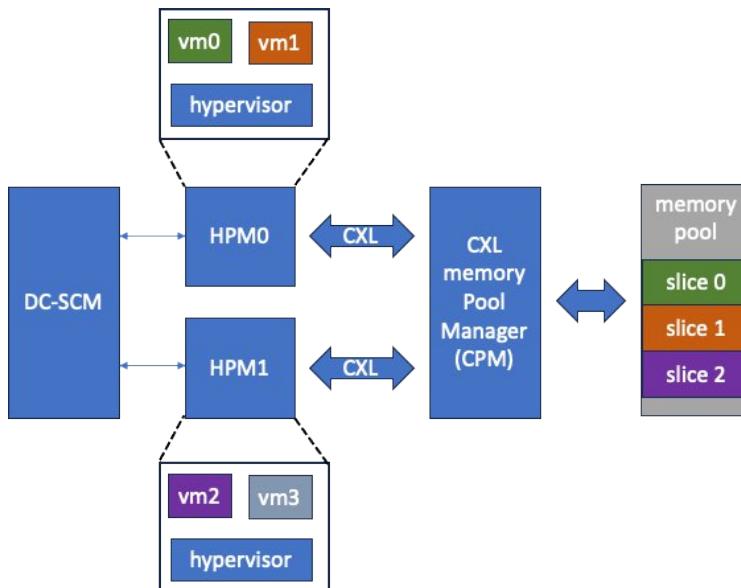


FIGURE 5 - CXL-BASE DISAGGREGATED MEMORY

As such, in order to provide the above functionalities, HIGHER identifies the requirements listed below:

5.4.1 Efficient Memory pool size management

ID	Description	Attribute	Status
UCR-UC-CXL-54100	<u>Memory Pool Granularity</u> : The HIGHER project shall define the minimal memory size to be managed by the CPM (We call this size “slice”)	M	Validated

ID	Description	Attribute	Status
UCR-UC-CXL-54101	<u>Update allocated slices</u> : The HIGHER project shall define a mechanism for (i) allocation of a new slice for HPMs when required, (ii) remove an already allocated slice from an HPM, (iii) move data from a source slice to a destination slice	M	Validated



5.4.2 Memory Access monitoring

The purpose is to Monitor memory access requests and ensure legitimate accesses.

ID	Description	Attribute	Status
UCR-UC-CXL-54200	<u>Apply permissions to slice</u> : The HIGHER project shall support an Access Control List (ACL) for the purpose of update access permissions to memory slices, such as exclusive read-only, write or shared read-only, write	M	Validated

5.4.3 Data Protection

The purpose is to establish hardware security layer for data protection against tampered requests.

ID	Description	Attribute	Status
UCR-UC-CXL-54300	<u>Stored data protection</u> : The HIGHER project shall enable optional inline data encryption before they are stored in memory / data decryption after they are loaded from memory	R	Validated

5.4.4 Accelerated data processing

The purpose is to allow near-data processing at the CPM level.

ID	Description	Attribute	Status
UCR-UC-CXL-54400	<u>Near-Data Processing</u> : The HIGHER project shall enable user-configurable NDP before data get stored in the memory (e.g. compression, multiply-accumulate)	O	Validated

5.4.5 CPM Configuration

ID	Description	Attribute	Status
UCR-UC-CXL-54500	The HIGHER project shall provide an interface for CPM configuration from the DC-SCM with respect to data encryption / decryption, slice permissions, memory allocation, inline NDP	O	Validated

5.4.6 Efficient memory access

The purpose is to enforce low latency and high-throughput memory accesses.

ID	Description	Attribute	Status
UCR-UC-CXL-54600	The CPM should provide low-latency access with high throughput to memory pool slices	M	Validated



6 Key Results

Although we make no exhaustive enumeration of the project's KPIs (these are listed in the Description-of-Action), the fundamental goal remains that all the KPIs, as listed in the DoA, are to be met. The architecture and component-level specification work that is to follow this deliverable will elaborate on technical design decisions and implementation constraints.

6.1 Technical Readiness Levels (TRL)

HIGHER technologies are designed to start from EPI and EUPilot chips targeting Technical Readiness Level (TRL) 6-7 in order to develop 2 compute modules and one management module and achieve at most TRL6 for most of their hardware and software components (project results) by the end of the project. Table 3 below (derived from the Description of Action for easy reference) summarizes the starting and ending TRL of the main project results, and whether there is already existing related work. In this deliverable, we consider each of these key results, and elicit the requirements that need to be elaborated upon for the purpose of defining in detail the architecture and specification the technical work of the HIGHER project.

Project results	Description	TRL
R1: RHEA2 based processor module	PCB of the RHEA2-based processor module which will be a pre-production board so that several similar boards can be built after the project and exploited. (T3.1)	6
R2: EPAC/EUPilot based processor module	PCB of the EPAC/EUPilot-based module which will also include a 1st prototype of a European Socket, while being in a pre-production phase so several similar boards can be built after the project and exploited. (T3.1)	6
R3: HIGHER Management Module	Pre-production deployment including the board firmware supporting Root-of-Trust and proper interfaces to the secure boot modules (see R7 and R8) (T3.3)	6
R4: HIGHER OCP Server	OCP-compliant chassis and overall server efficiently incorporating all HIGHER modules and optimized for low energy consumption (T3.4)	4 to 6
R5: PCIe/CXL and OCP-NIC connectivity	Efficient PCIe/CXL connectivity provided for the processor modules so that a host can be accessed externally and exchange data with other hosts and firmware elements to facilitate device discovery by Linux, and networking driver, for devices following the OCP NIC 3.0 standard. (T3.4, T4.1, T4.2)	2 to 6
R6: European Socket CPU	Reference implementation of a CPU socket which will be compatible with the advanced and next generation European CPUs (T3.2)	2 to 6
R7: Arm Secure Boot and Linux	Provide the low-level boot support software for the ARM processor module, incl. security features and interfaces with management module (T4.2)	4 to 6
R8: RISC-V Secure Boot and Linux	Provide the low-level boot support software for RISC-V processor module, incl. security	2 to 6



Project results	Description	TRL
	features and interfaces with management module (T4.2)	
R9: Meta-OS	Provide the codebase of resource management framework and distributed execution engines for computing continuum application scenarios	4 to 6
R10: 4 real world Cloud Edge Use cases	Provide the codebases of four proof-of-concept use cases: Accelerated data processing and analysis for converged Cloud/HPC platforms, IaaS, PaaS, Rack-level Memory Pooling. (WP5)	6
R11: 3 Open Installations of HIGHER Servers	3 installations of HIGHER systems in 3 different premises open to the wider R&D community for experimentation, further development etc.	4 to 6

TABLE 7 - KEY RESULTS FROM THE HIGHER PROJECT AND THEIR EXPECTED TECHNICAL READINESS LEVELS

7 Requirements

7.1 OCP Standardisation landscape

The data centre - Modular Hardware System (DC-MHS) is an OCP project providing data centre-ready integrated systems for edge, private cloud, and large data centres.

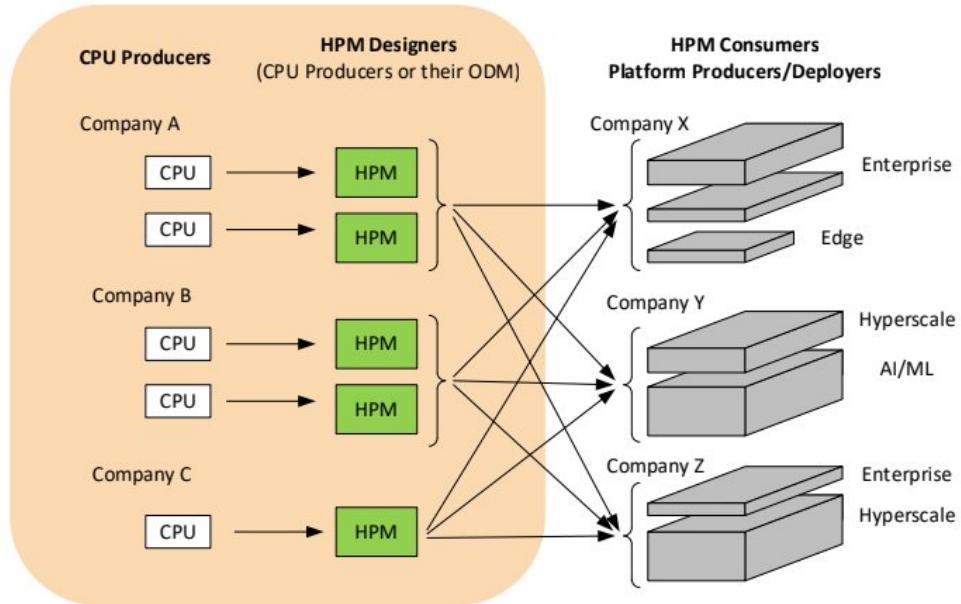


FIGURE 6 - OCP DC-MHS PARADIGM

The constituent of the DC-MHS specification landscape is as follows:

- Modular – Density Optimized (M-DNO) Host Processor Module
- Modular – Full Width (M-FLW) Host Processor Module
- Modular – Common Redundant Power Supply (M-CRPS)
- Modular – Platform Infrastructure Connectivity (M-PIC)
- Modular – eXtended I/O (M-XIO) – PCIe/CXL I/O Connectivity
- Modular – PESTI – (M-PESTI) – Sideband Signal Virtualization
- data centre – Secure Control Module (DC-SCM R2.0)
- Open Compute Network Interface Card R3.0 (OCP NIC 3.0)

Following the DC-MHS which offers specifications for modules compatible across servers, chassis and vendors, HIGHER will develop the following hardware modules:

- Processor Modules - Two OCP-compliant processor modules, which will be based on the OCP Host Processor Module (HPM) or the OCP Universal Baseboard (UBB) standard, one hosting the RHEA2 EPI chip and the other hosting the EPAC2.0 EPI chip and the pin compatible EUPIilot chip.
- Management Module - A data centre-ready Secure Control Module (DC-SCM), hosting a RISC-V processor inside an FPGA for server management, security, and control features.

Those modules will be assembled together in a chassis developed in the scope of the project.

Figure 7 below illustrates server integrating DC-SCM module and HPM module:

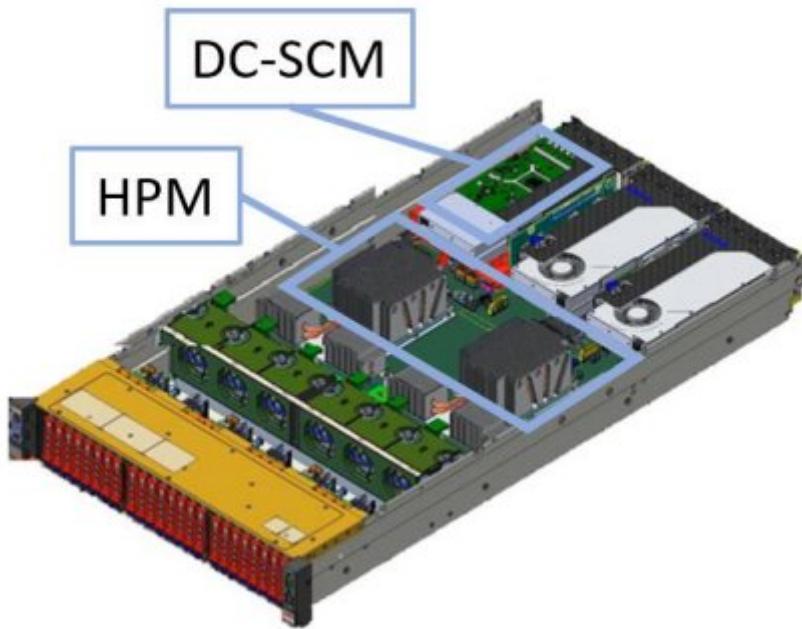


FIGURE 7 - DC-SCM AND HPM IN DC-MHS CHASSIS

7.2 OCP HPM Form Factors

For the HPMs, OCP has defined 3 successive generations of specifications that are still active and maintained by the DC-MHS project.

7.2.1 M-FLW

M-FLW (Modular Full Width) was officially introduced in September 2022, and defines full width fixed form factor HPM optimised for 19" racks. It is dedicated to dual socket/CPU architecture and allows for combination with x1 DC-SCM and x2 OCP-NIC. It is also possible to have configurations with 1x DC-SCM, 1x OCP NIC and 2x E1.S for storage (MSI D5062 for example).

The **Figure 8** below is a high level representation of the construct of a chassis as a combination of a M-FLW HPM with x2 NICs, one DC-SCM, PSUs and Storage.

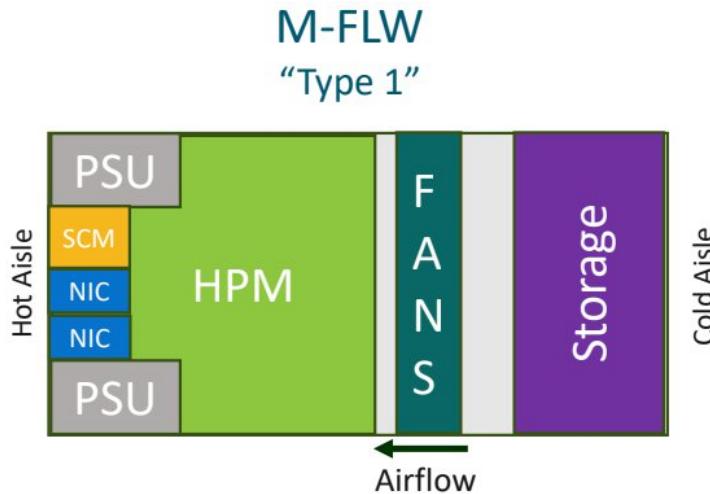


FIGURE 8 - M-FLW CHASSIS



PSUs are not mandatory for power up of the M-FLW motherboard. The M-PIC connectors on the board give the possibility to power up the board with cables, thus having the PSUs deported in the chassis and not plugged into the M-CRPS connectors which are on the board. This option implies the use of a Power Distribution Board on which the PSUs are connected.

7.2.2 M-DNO

M-DNO (Modular DeNsity Optimized) was introduced in September 2022 as well and defines partial width fixed form factors optimised for 19" racks. It is dedicated to densely organised HPM with single (Type1) or dual socket architecture with interoperability in mind. For the dual sockets use case, M-DNO supports both shadow (Type3) or side-by-side (Type4) organisation.

The **Figure 9** below is a high level representation of the construct of a chassis as a combination of a M-DNO Type2 HPM with one NIC, one DC-SCM, PSUs and Storage.

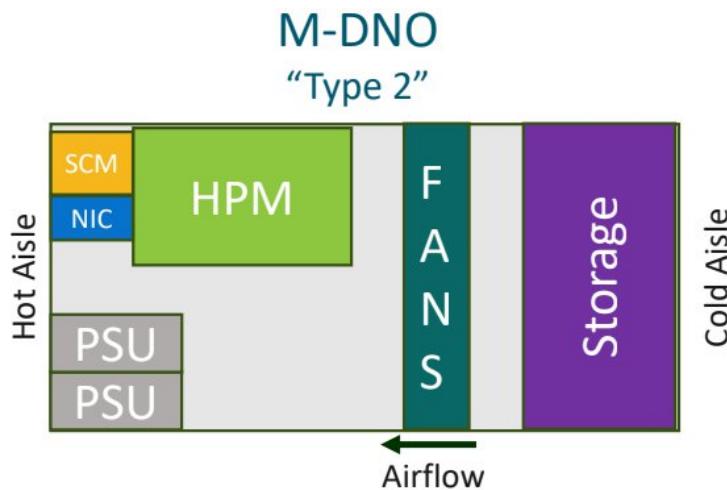


FIGURE 9 - M-DNO CHASSIS

The **Figure 10** below gives an example of interoperability of a chassis that can host without changes DNO HPM of either Type 2 or Type 4.

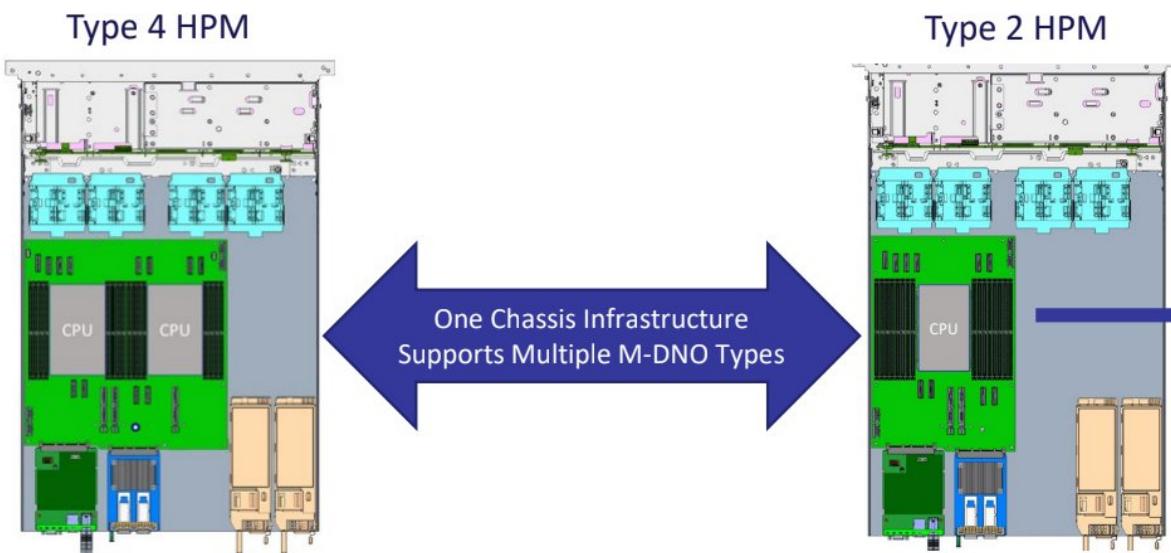


FIGURE 10 - M-DNO CHASSIS MODULARITY

7.2.3 M-SDNO

M-FLW and M-DNO focus on fixed form factors optimised for 19" racks. M-SDNO was introduced in June 2024 for the purpose of enabling ORv3 21" optimized form factors as well as allowing HPM designers with better capabilities to "right size" their boards and create common Half/Full width HPMs for 19" and 21" racks.

The **Figure 11** below is a high level representation of the construct of a chassis as a combination of a M-SDNO Class C335 with one NIC, one DC-SCM, GPU(s) and Storage.

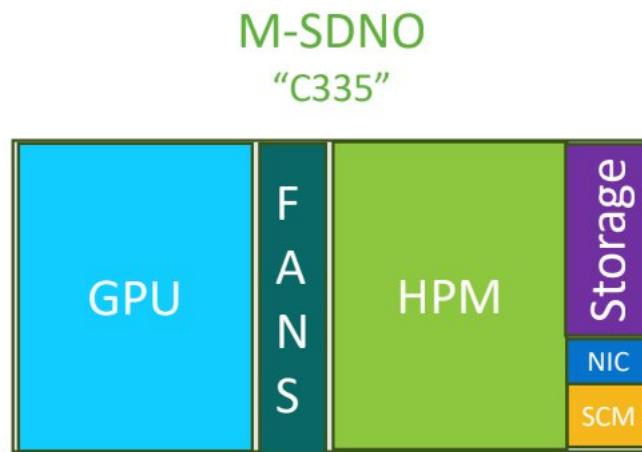


FIGURE 11 - M-SDNO CHASSIS

M-SDNO defines HPM Classes A, B, C, D. All classes have consistent width and feature set and allow for variable length with the concept of Common Chassis Interval (CCI). The **Figure 12** below depicts the various possibilities of sizing offered by SDNO.

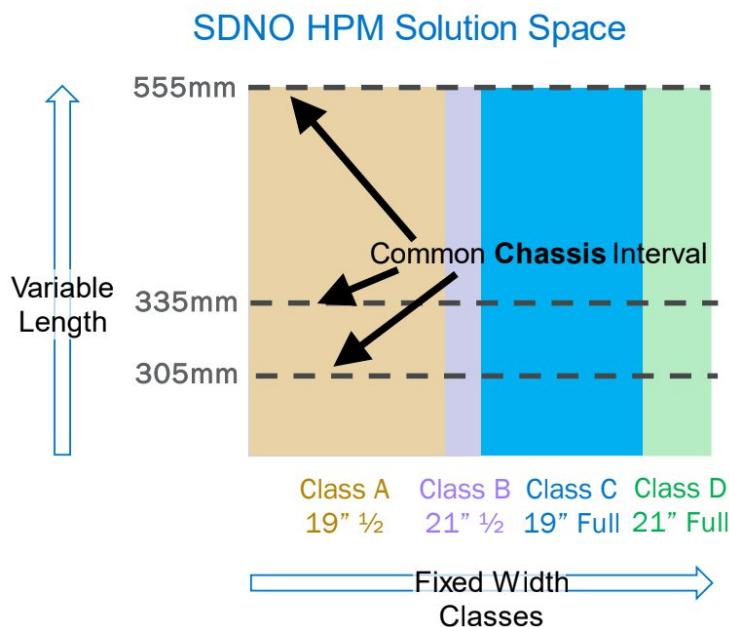


FIGURE 12 - M-DNO COMMON CHASSIS INTERVAL (CCI)

7.2.4 HPM Market Considerations

Right after the publication of the M-FLW specification at the end of 2022 some dual socket reference platforms based on intel Xeon were showcased by Wistron and QCT at OCP Summit 2023. The year after, MSI showcased a set of DC-MHS boards with dual sockets (M-FLW) and single socket (M-DNO Type 2 and Type 4) which turned into constituent of DC-MHS servers actively sold by MSI.

From Q2 2025, Chenbro announces availability of chassis ready for hosting single or dual sockets motherboards compliant with M-FLW.

Since the publication of M-SDNO, there appears to be a lot of traction on the half-width form factors (class A and class B). Ampere, Pegatron MiTAC and Supermicro already have such type of implementation in their upcoming portfolio.

At OCP 2024, there were still numerous M-FLW implementation presented by Jabil, Asus, Intel, Dell, QCT and Wistron. On the side of SDNO Class A, Hyve, Inventec and Supermicro were also presenting solutions. Interestingly, VVDN was presenting a long length SDNO Class B with 2 sockets in shadow mode.

The **Figure 13** below shows an inventory of the HPM that were presented at OCP 2024.

M-HPM In the Expo Hall

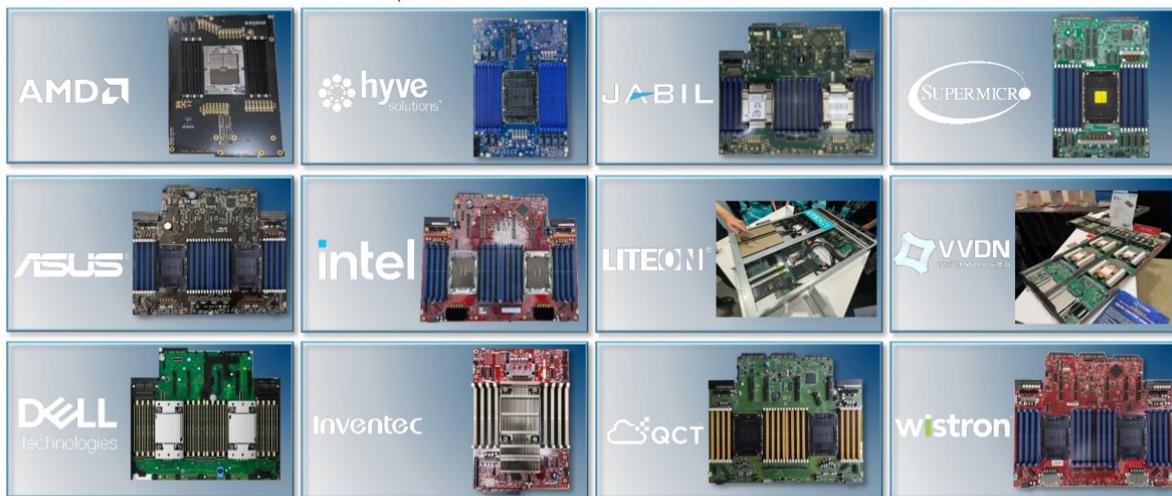


FIGURE 13 - M-HPM DEMONSTRATED AT OCP'24

Concurrently, Nvidia are promoting their server modularity framework labelled MGX which also offers de-facto standardisation for both dual socket and single socket motherboard, as shown in the **Figure 14** below.

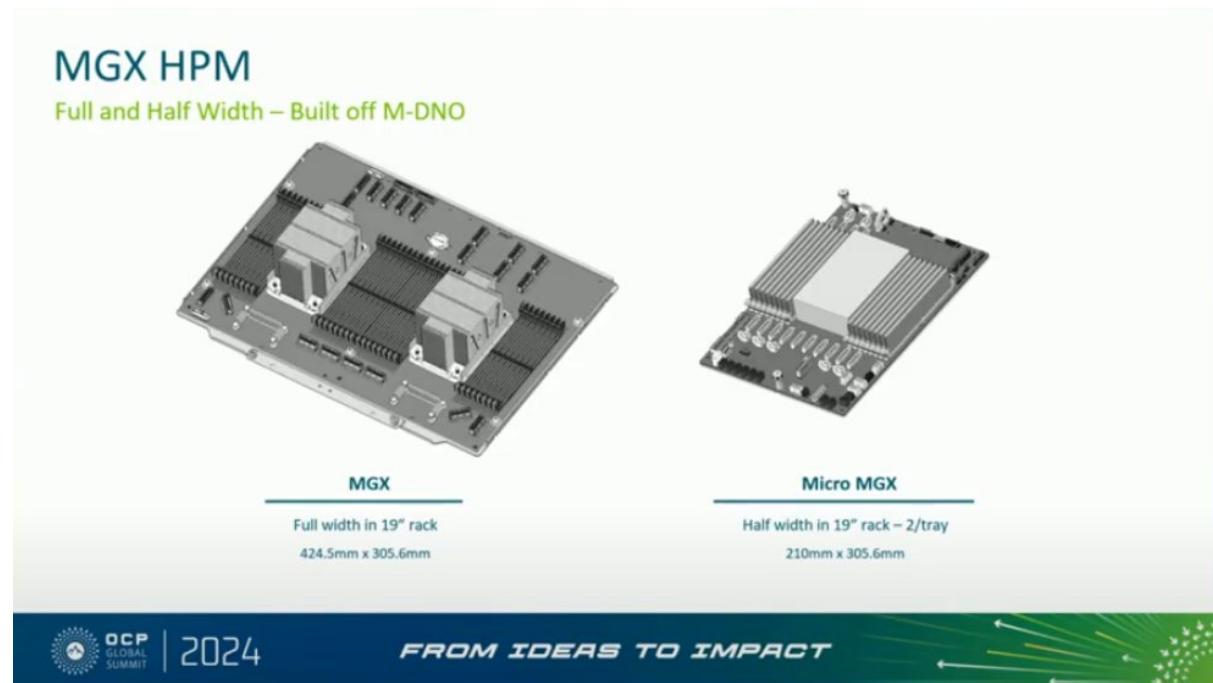


FIGURE 14 - NVIDIA MGX HPM

nVidia is promoting MGX within OCP, advocating that Micro-MGX and M-DNO Type2 (and therefore SDNO class A305) are very similar, as shown in the **Figure 15** below:

MGX & DC-MHS: HPM Similarities & Differences

Half Width – Built off M-DNO Type 2

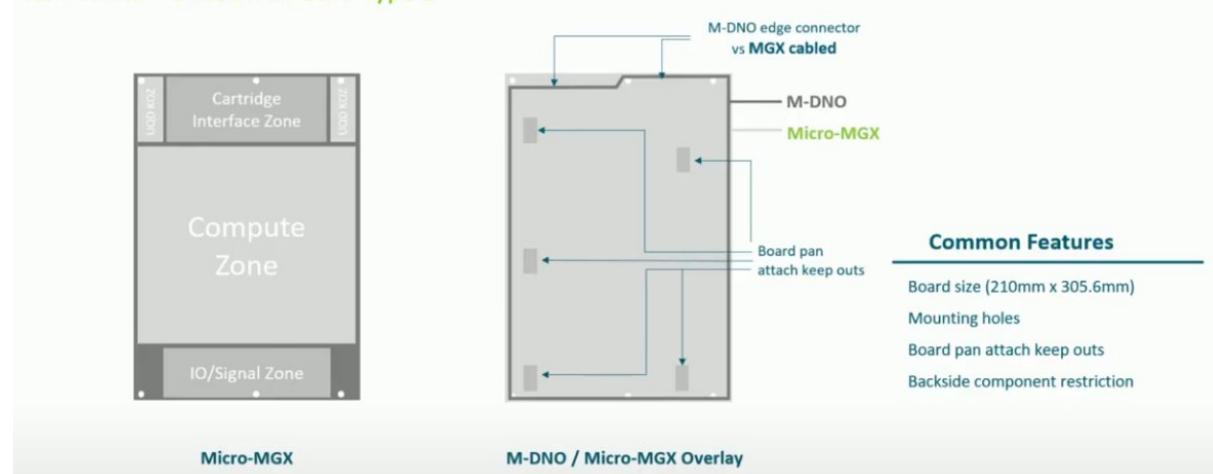


FIGURE 15 - MGX & DC-MHS SIMILARITIES AND DIFFERENCES

Also, MGX (424mm x 306mm) and SDNO Class C (426mm x 305mm) are very similar in size which is a likely reason for SDNO Class C to generate less traction than expected on the roadmap of motherboard suppliers and server manufacturers. Quite a number of players, such as ASRock, Compal, QCT and Chenbro are currently focusing on offering chassis solutions for MGX.

As a summary, the current trend is as follows:

- The dual socket solutions are driven by HPC and intensive dual sockets use cases, and most of the players are currently sticking to M-FLW as a pretty much established and conservative



standard from which M-SDNO Class C/D does not constitute yet a significant breakthrough. However, actors such as Scaleway and GigaComputing confirmed they are looking at SDNO Class C/D (CloudFest 2025 feedback from SiPearl)

- The compact single socket solutions are driven by AI and cloud intensive use cases where the memory bandwidth is more crucial than the interconnect bandwidth, and in this context, the compact DNO form factors and more recently the half-width SDNO Class A/B are getting traction and rising on the roadmaps of OEMs (2CRSi, Scaleway, OVH) as well as ODMs (MSI, AsRock) as confirmed after attendance to CloudFest 2025 (SiPearl & 2CRSi)
- The promotion of nVidia MGX is intercepting the DC-MHS roadmap, and we see server manufacturers engaging in MGX support as a potentially dominant solution, delaying for the time being the adoption of the full width form factors of SDNO (class C and D)

7.3 Dual sockets Rhea2-based HPM Requirements

The dual socket Rhea2-based HPM integrates 2 sockets dedicated to hosting 2 RheaR2 SoCs.

The dual socket Rhea2-based HPM ensures High Speed connectivity between both RheaR2 with x4 PCIe6 x8 CML-SMP links.

The dual socket Rhea2-based HPM also allows for high speed connectivity with up to x4 Accelerators or GPUs and x2 Backplane Storage devices, with a total of x6 PCIe6 x16 links.

The dual socket Rhea2-based HPM allows for connectivity with x1 DC-SCM and x2 OCP-NIC, one for each of the SoCs.

Finally, the dual socket Rhea2-based HPM makes provision with x2 M.2 connectors which will typically be used with M.2 SSD as Boot Storage devices for the Hypervisor or the OS.

The dual socket Rhea2-based HPM complies with SDNO Class C335.

The **Figure 16** below gives a high level overview of the dual socket Rhea2 HPM components and interfaces.

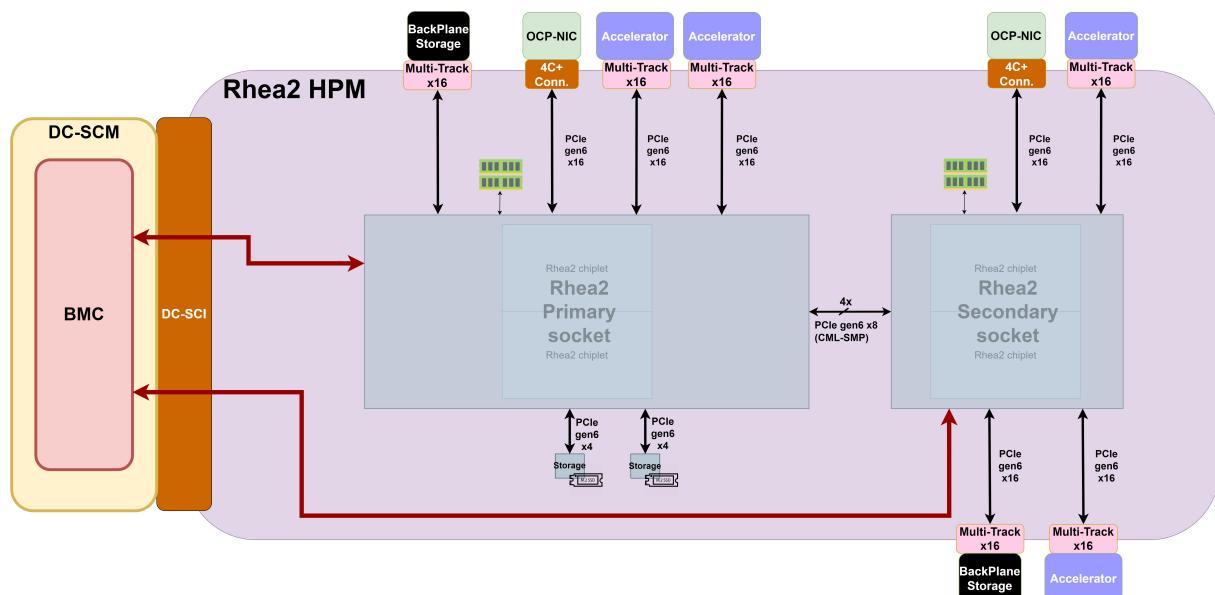


FIGURE 16 - RHEA2-BASED HPM - SDNO CLASS C

The following chapters formulate the requirements applicable to the dual socket Rhea2-based HPM and its associated software.

7.3.1 General

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73100	The [Rhea2-HPM] shall have a form factor compliant with M-SDNO Class C335, as defined in [M-SDNO]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73101	The [Rhea2-2S-HPM] shall comply with the M-SDNO Base Specification	M	Validated



ID	Description	Attribute	Status
	Compliance table, as defined in [M-SDNO] chapter 4		

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73102	The [Rhea2-HPM] shall have x2 sockets dedicated to hosting RheaR2 Host processors	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73103	On the [Rhea2-HPM] the x2 RheaR2 Host processors shall be interconnected with 4x PCIe Gen6 x8 CML-SMP links	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73104	The [Rhea2-HPM] shall support 12x RDIMM connectors for each RheaR2 socket	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73105	The [Rhea2-HPM] shall support one QSPI flash for firmware boot, dedicated to the Primary RheaR2 socket	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73106	The [Rhea2-HPM] shall support one QSPI flash for firmware boot, dedicated to the Secondary RheaR2 socket	M	Validated

7.3.2 Interface with DC-SCI

The RheaR2 SoCs will both communicate with the BMC hosted on the DC-SCI, in compliance with the Arm Server Base Management Requirements, supporting the Level M4 as defined in [SBMR].

The **Figure 17** below highlights the interfaces involved, that are expected to be passing through the DC-SCI.

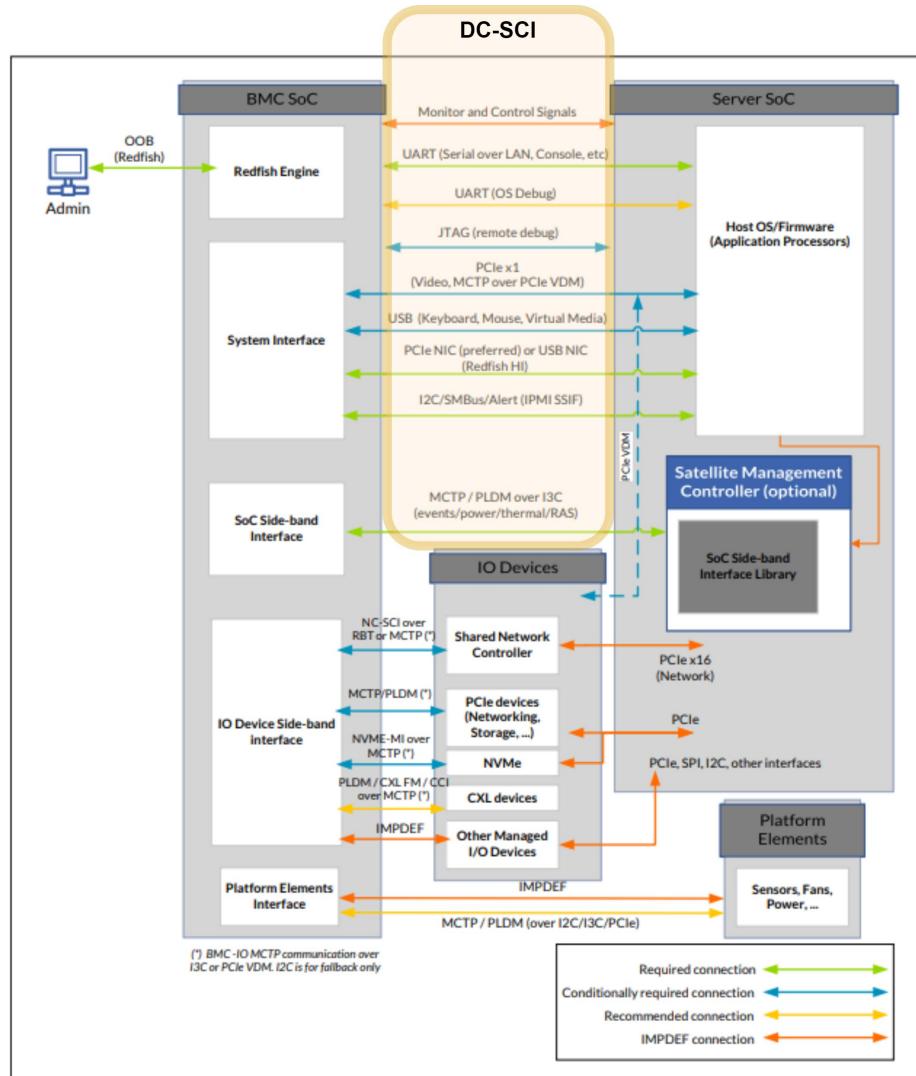


FIGURE 17 - COVERAGE OF SBMR BY DC-SCI¹

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73200	The [Rhea2-HPM] shall support one DC-SCI R2.X connector	M	Validated

¹ source: [Arm Server Base Manageability Requirements](#)



ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73202	On the [Rhea2-HPM], the DC-SCI interface shall be usable for each of the RheaR2 SoC to access the BIOS flash memory of the DC-SCM, for firmware boot purpose	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73203	On the [Rhea2-HPM], the DC-SCI interface shall be usable for each of the RheaR2 SoC to access the Trusted Platform Module of the DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73204	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the Monitor and Control Signals (GPIOs) between both RheaR2 SoCs and the BMC	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73205	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of a UART connection between the Primary RheaR2 SoCs and the BMC, for the purpose of serial-over-LAN (SoL), as specified in chapter 2.1.1.2 of [SBMR]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73206	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of a UART connection between the Primary RheaR2 and the BMC, for the purpose of remote OS debugging through the BMC, as specified in chapter 2.4.1.2 of [SBMR]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73207	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of graphic video redirection over PCIe, as specified in chapter 2.1.1.3 of [SBMR]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73208	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of remote Keyboard-Video-Mouse (KVM) over USB, as specified in chapter	M	Validated



ID	Description	Attribute	Status
	2.1.1.4 of [SBMR]		

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73209	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the RedFish side-band interface over a PCIe connection, as specified in chapter 2.2.1.1 of [SBMR]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73210	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the IPMI side-band interface over an SSIF interface (SMBUS), as specified in chapter 2.3.1.1 of [SBMR]	M	Validated

Note: This is involving I2C with SMBAlert support

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73211	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the side-band interface over I3C, as specified in chapter 2.5.1.2 of [SBMR]	M	Validated

7.3.3 OCP-NIC Support

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73301	The [Rhea2-HPM] shall support one OCP-NIC 4C+ connector dedicated to the <u>Primary RheaR2</u>	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73302	The [Rhea2-HPM] shall support one OCP-NIC 4C+ connector dedicated to the <u>Secondary RheaR2</u>	M	Validated

7.3.4 High-Speed I/O Connectivity

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73400	The [Rhea2-HPM] shall support M-XIO or MCIO connectors dedicated to the <u>Primary RheaR2</u> , exposing 48 lanes in total	M	Validated

Note: x2 16 lanes dedicated to accelerators and x1 16 lanes for Storage



ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73401	The [Rhea2-HPM] shall support M-XIO or MCIO connectors dedicated to the <u>Secondary</u> RheaR2, exposing 48 lanes in total	M	Validated

Note: x2 16 lanes dedicated to accelerators and x1 16 lanes for Storage

7.3.5 Platform Infrastructure Connectivity

General

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73500	The [Rhea2-HPM] shall follow the Connector requirements and placement as specified in chapter 10.8 of [M-SDNO]	M	Validated

Power distribution

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73510	The [Rhea2-HPM] shall support PDB with 12V output and HPM power distribution architecture requirements, as defined in chapter 10.1.5 of [M-PIC]	M	Validated

Note: It means the [Rhea2-HPM] shall support Power Connector(s) (PICPWR)

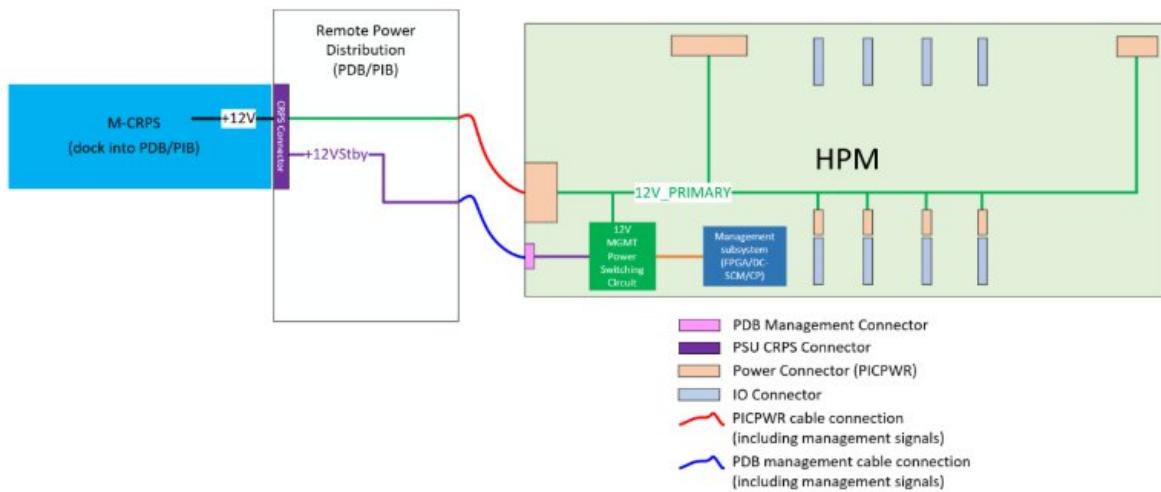
ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73511	The [Rhea2-HPM] shall implement PDB Management connector as defined in chapter 10.2.14 of [M-PIC]	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73512	The [Rhea2-HPM] shall implement PDB Management connector Type 1 as defined in chapter 10.2.14.1 of [M-PIC]	O	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S- HPM-73513	The [Rhea2-HPM] shall implement PDB Management connector Type 2 as defined in chapter 10.2.14.2 of [M-PIC]	O	Validated

Note: The choice between PDB Management connector type 1 and type 2 will be made during execution of task 2.2 and indicated in D2.2.

Note: Use of PDB Management connector type 2 is as depicted in **Figure 18**. For type 1, the PDB management Connector is aggregated with the PICPWR.

FIGURE 18 - 12V REMOTE POWER DISTRIBUTION ARCHITECTURE ²

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-73513	The [Rhea2-HPM] shall support powering of the E1.S specific board	R	Validated

Note: Refer to 7.8.4 for details on the storage

Boot storage

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73520	The [Rhea2-HPM] shall support 1x M.2 connector accessible from the Primary RheaR2 socket, for the purpose of Boot Storage, as defined in chapter 11.1.2.1 of [M-PIC]	M	Validated

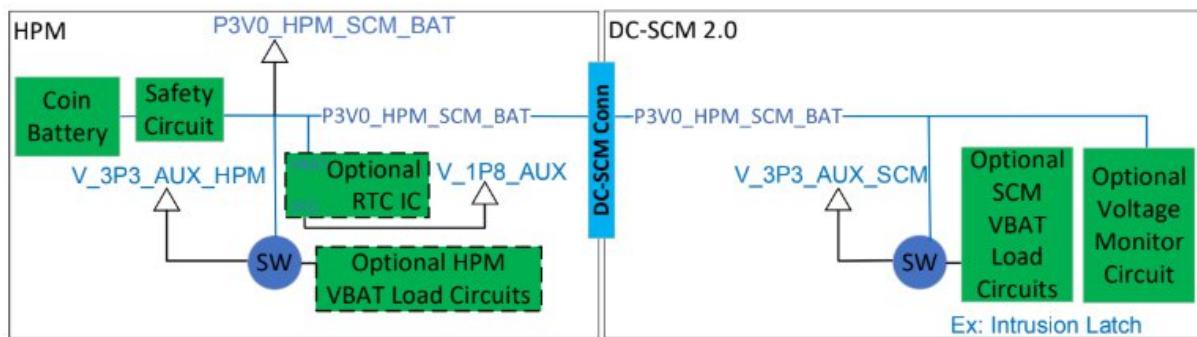
ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73521	The [Rhea2-HPM] shall support a second M.2 connector accessible from the Primary RheaR2 socket, for the purpose of Boot Storage, as defined in chapter 11.1.2.1 of [M-PIC]	R	Validated

Coin cell battery

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73530	The [Rhea2-HPM] shall support a coin cell battery and support the battery backed voltage interface with the DC-SCM, as defined in chapter 10.3 of [M-PIC]	R	Validated

Note: As depicted in [Figure 23](#)

² source: [M-PIC Base Specification](#)

FIGURE 19 - BATTERY BACKED VOLTAGE INTERFACE³

Intrusion switch

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73540	The [Rhea2-HPM] shall support an intrusion switch as defined in chapter 11.2 of [M-PIC]	R	Validated

Internal USB Connector

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73550	The [Rhea2-HPM] shall support an internal Type A USB connector, accessible to the Primary RheaR2 and as defined in chapter 11.3 of [M-PIC]	R	Validated

Control Panel

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73560	The [Rhea2-HPM] shall support the Primary Control Panel (PCP) connector as defined in chapter 11.4.1 of [M-PIC]	M	Validated

Note: Availability of USB connection over primary PCP to be confirmed during task 2.2 activities

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73561	The [Rhea2-HPM] shall support the Secondary Control Panel (PCP) connector as defined in chapter 11.4.1 of [M-PIC]	R	Validated

³ source: [OCP DC-SCM Specification](#)

7.3.6 Software

Leveraging SGA-2 GPP System and Software Specifications

The reference Rhea2 software stack, as developed in the scope of the SGA-2 project will be made available as an input to the [HIGHER] project.

The Rhea2 software stack includes firmware capable of:

- Secure boot
- System Startup, Cold Boot, Power Down
- Chain of trust enforcement
- Power & thermal regulation
- Performance Monitoring
- RAS reporting

The firmware includes a reference BIOS implementation based on EDK2.

The firmware supports the following interfaces, as required by [SBMR] and [BBR]:

- Monitor and Control signals
- Side-band interface
- In-Band interface
- SMBIOS
- PSCI
- TRNG interface
- SCMI
- UEFI
- ACPI

The firmware allows for the boot of established Linux distribution such as RHEL that will be complemented with computational kernel interface such as OpenMP, as shown in the high-level firmware/software stack overview:

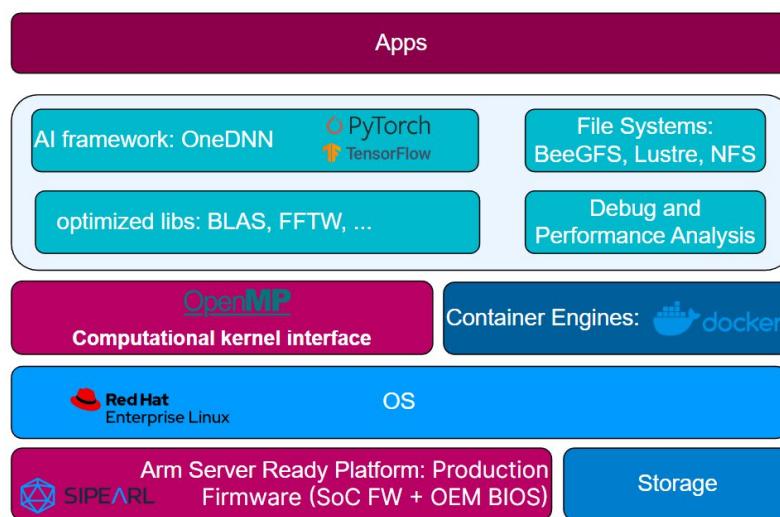


FIGURE 20 - RHEA2 FIRMWARE/SOFTWARE STACK - HIGH-LEVEL VIEW

Alternatively, the firmware allows for the boot of established hypervisors such as KVM or Xen, which will be useful for the CXL-memory disaggregation use case:

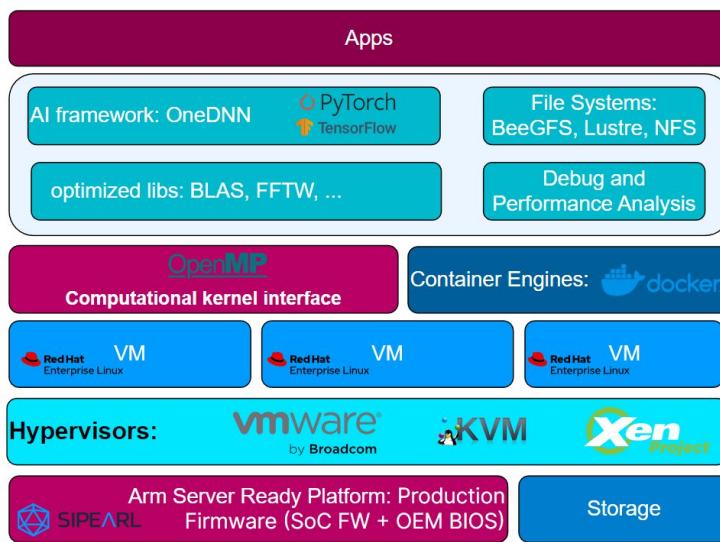


FIGURE 21 - RHEA2 HYPERVISOR STACK - HIGH-LEVEL VIEW

It should be noted that at the time of writing, the availability of KVM or Xen on ArmV9 is not confirmed, hence putting a dependency for the CXL-memory disaggregation use case which relies on availability of an Hypervisor running on ArmV9.

Offloading kernels to OAMs for acceleration purpose

On the specific requirement that the application running on the Rhea2 host is expected to be able to offload kernels to the acceleration units that will be available on the OAMs of the acceleration HPM, some specific acceleration drivers and offloading middleware hooked to the computational kernel interface will be implemented in the scope of the HIGHER project, as depicted in the figure below:

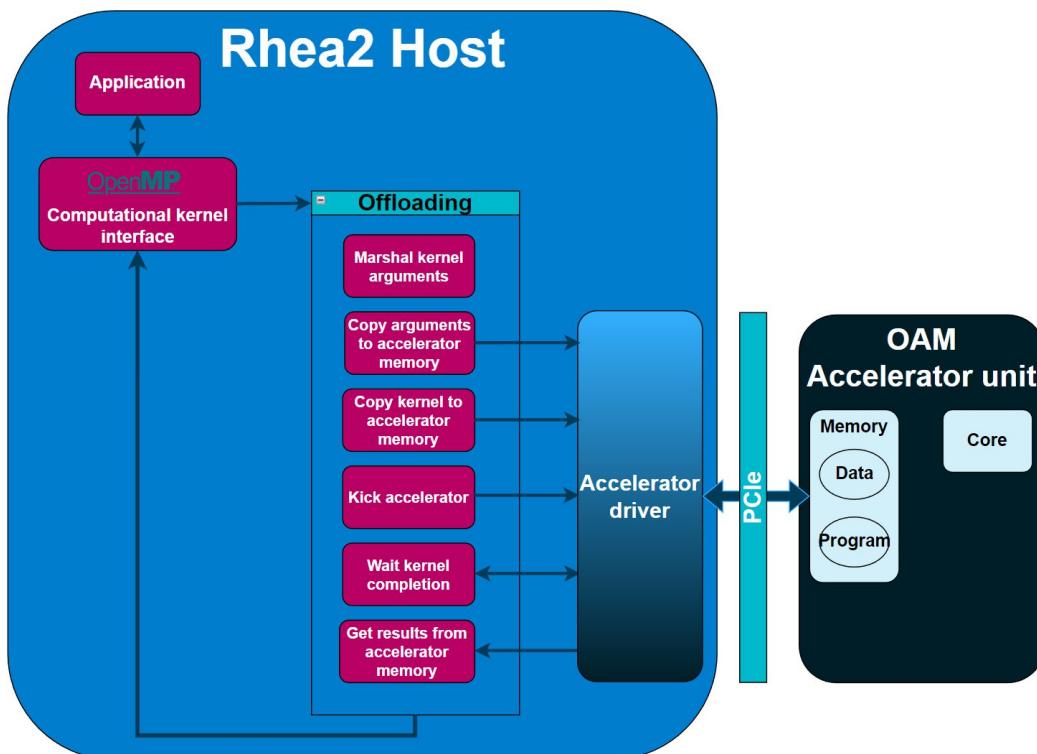


FIGURE 22 - OFFLOADING KERNELS FOR ACCELERATION IN MULTI-DEVICES CONTEXT



Software Requirements

This is the list of software requirements that are specific to the HIGHER project.

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73600	The [Rhea2-HPM] firmware/software stack shall incorporate accelerator drivers and associated offloading middleware to enable interaction between the host-resident software stack and the OAM accelerator's compute and memory resources, following the OpenMP task offload programming model	M	Validated

ID	Description	Attribute	Status
UCR-HW-R2-2S-HPM-73600	The [Rhea2-HPM] software stack shall possibly incorporate an hypervisor ported on ArmV9	R	Validated

Note: This is expected for the CXL-memory disaggregation use case where memory slices are allocated to dedicated VMs.

7.4 Single socket Rhea2-based HPM Requirements

The single socket Rhea2-based HPM integrates one socket dedicated to hosting one RheaR2 SoC.

The single socket Rhea2-based HPM ensures High Speed connectivity with RheaR2 of a secondary HPM with x4 PCIe6 x8 CML-SMP links.

The single socket Rhea2-based HPM allows for high speed connectivity with up to x2 Accelerators or GPUs and x1 Backplane Storage devices, with a total of x3 PCIe6 x16 links.

The single socket Rhea2-based HPM allows for connectivity with x1 DC-SCM and x1 OCP-NIC.

Finally, the single socket Rhea2-based HPM makes provision with x2 M.2 connectors which will typically be used with M.2 SSD as Boot Storage devices for the Hypervisor or the OS.

The single socket Rhea2-based HPM complies with SDNO Class A335.

The **Figure 23** below gives a high level overview of the single socket Rhea2 HPM components and interfaces.

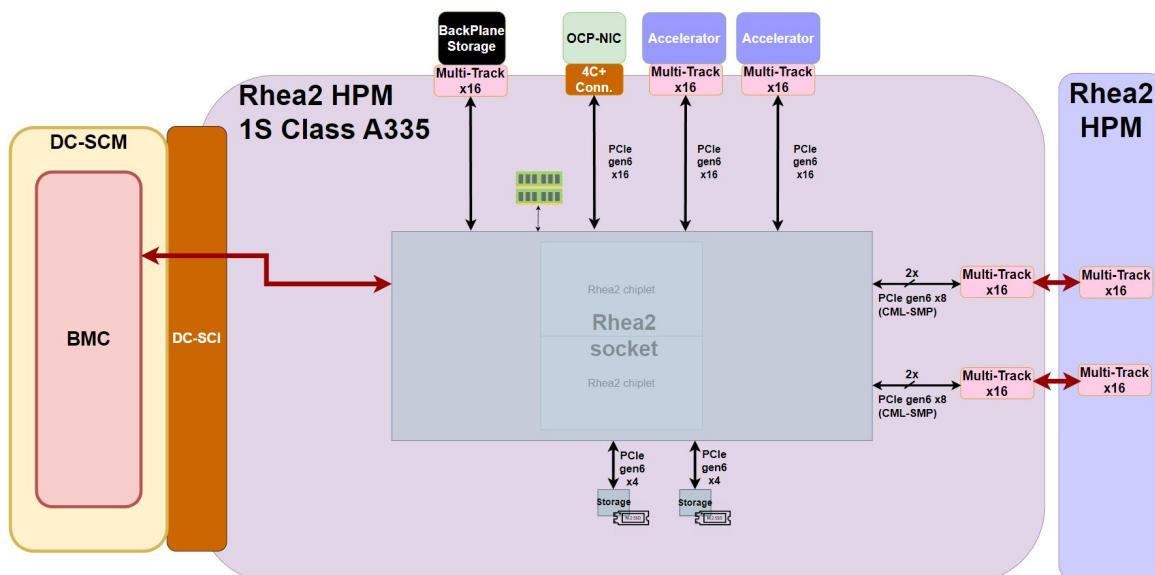


FIGURE 23 - RHEA2-BASED HPM - SDNO CLASS A

The following chapters formulate the requirements applicable to the single socket Rhea2-based HPM.

7.4.1 General

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74100	The [Rhea2-HPM] shall have a form factor compliant with M-SDNO Class A335, as defined in [M-SDNO]	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74101	The [Rhea2-2S-HPM] shall comply with the M-SDNO Base Specification Compliance table, as defined in [M-SDNO] chapter 4	R	Validated



ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74102	The [Rhea2-HPM] shall have x1 socket dedicated to hosting one RheaR2 Host processor	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74102	The [Rhea2-HPM] shall support 12x RDIMM connectors for the RheaR2 socket	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74103	The [Rhea2-HPM] shall support one QSPI flash for firmware boot, accessible from the RheaR2 socket	R	Validated

7.4.2 Interface with DC-SCM

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74200	The [Rhea2-HPM] shall support one DC-SCM R2.X connector	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74201	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the RheaR2 SoC to access the BIOS flash memory of the DC-SCM, for firmware boot purpose	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74202	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the RheaR2 SoC to access the Trusted Platform Module of the DC-SCM	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74203	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the Monitor and Control Signals (GPIOs) between the RheaR2 SoC and the BMC	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74204	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of a UART connection between the RheaR2 SoC and the BMC, for the purpose of serial-over-LAN (SoL), as	R	Validated



ID	Description	Attribute	Status
	specified in chapter 2.1.1.2 of [SBMR]		

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74205	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of a UART connection between the RheaR2 SoC and the BMC, for the purpose of remote OS debugging through the BMC, as specified in chapter 2.4.1.2 of [SBMR]	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74206	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of graphic video redirection over PCIe, as specified in chapter 2.1.1.3 of [SBMR]	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74207	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of remote Keyboard-Video-Mouse (KVM) over USB, as specified in chapter 2.1.1.4 of [SBMR]	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74208	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the RedFish side-band interface over a PCIe connection, as specified in chapter 2.2.1.1 of [SBMR]	R	Validated

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74209	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the IPMI side-band interface over an SSIF interface (SMBUS), as specified in chapter 2.3.1.1 of [SBMR]	R	Validated

Note: This is involving I2C with SMBAlert support

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74210	On the [Rhea2-HPM], the DC-SCI interface shall be usable for the support of the side-band interface over I3C, as specified in chapter 2.5.1.2 of [SBMR]	R	Validated

7.4.3 OCP-NIC Support

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74300	The [Rhea2-HPM] shall support one OCP-NIC 4C+ connector	R	Validated

7.4.4 High-Speed I/O Connectivity

ID	Description	Attribute	Status
UCR-HW-R2-1S-HPM-74400	The [Rhea2-HPM] shall support M-XIO or MCIO connectors exposing 80 lanes in total	R	Validated

Note: x2 16 lanes dedicated to scale-out, x2 16 lanes dedicated to accelerators and x1 16 lanes connector for Storage

7.4.5 Platform Infrastructure Connectivity

Refer to chapter 7.3.5, as requirement for Class C HPM are also application to the class A HPM.

7.4.6 Software

Refer to chapter 7.3.6, as requirement for Class C HPM are also application to the class A HPM.

7.5 EPAC/EUPILOT-based HPM Requirements

The EPAC/EUPILOT-based HPM is designed to serve mainly as a high-performance acceleration platform aligned with the OCP DC-MHS specifications. It supports up to two OCP Accelerator Modules (OAMs), based on the EUPILOT or compatible architecture, providing scalable compute acceleration for AI, HPC, and data-intensive workloads.

This section defines the hardware, firmware, interface, and integration requirements specific to the EPAC/EUPILOT-based HPM module, ensuring compatibility with SDNO Class A335 or equivalent OCP form factors and its correct operation within the HIGHER system architecture.

7.5.1 General

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75100	The [EP-HPM] shall support a form factor compliant with SDNO Class A335.	R	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75101	The [EP-HPM] shall host two OAM modules for acceleration.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75102	The [EP-HPM] shall support thermal design for up to 300W TDP per OAM module.	M	Validated



ID	Description	Attribute	Status
UCR-HW-EP-HPM-75103	The [EP-HPM] shall support PCIe Gen5 connectivity from each OAM to the Host HPM via M-XIO or MCIO interfaces, exposing a total of 16 lanes	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75104	The [EP-HPM] shall provide one DC-SCM R2.X connector to interface with the management module.	M	Validated

7.5.2 OAM Interface & Connectivity

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75200	The [EP-HPM] shall support 2x OAM modules compliant with the OAM r1.0 v1.5 specification.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75201	Each EUPILOT OAM shall connect to the [EP-HPM] via 4-lane PCIe Gen5 interface and multiple 2-lane 32 Gbps OAM-to-OAM interfaces. The HPM can support 16-lane PCIe and 16-lane OAM-to-OAM interfaces, with any unused lanes remaining unconnected.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75203	The [EP-HPM] connectivity shall support direct OAM-to-OAM connectivity between the two OAMs on the HPM.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75204	The [EP-HPM] connectivity shall support OAM-to-OAM connectivity between OAMs on different HPMs.	R	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75205	The [EP-HPM] shall support CXL 2.0 compatibility for memory coherency with the host CPU.	R	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75206	The [EP-HPM] shall support CXL 3.1	O	Validated



7.5.3 Management and Monitoring

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75300	The [EP-HPM] shall expose per-OAM temperature, power, and error telemetry via DC-SCI to the DC-SCM.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75301	The [EP-HPM] shall support dynamic power capping per OAM module based on thermal constraints.	R	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75302	The [EP-HPM] shall support firmware-triggered reset and re-initialization of individual OAMs.	M	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75303	The [EP-HPM] shall support event logging and alerting via the BMC (Redfish/IPMI) on DC-SCM for fault conditions.	R	Validated

7.5.4 Firmware and Security

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75400	The [EP-HPM] firmware shall support secure boot, signed image validation, and rollback protection.	O	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75401	The [EP-HPM] firmware stack shall support field updates using authenticated and encrypted firmware packages.	R	Validated

ID	Description	Attribute	Status
UCR-HW-EP-HPM-75402	The [EP-HPM] shall integrate with the system's Root-of-Trust via the DC-SCM.	M	Validated



7.6 Management Module Requirements

In data centre setups, where availability and operational efficiency are keys, the management module of servers plays a crucial role. It enables administrators to remotely control and monitor a large number of servers in a robust fashion. Management modules rely on a Board Management Controller (BMC), often complemented by an FPGA, to deliver the required services. Those aspects have been formalized for both the ARM and the RISC-V ecosystem under various forms.

For ARM servers, the Arm Server Base Manageability Requirements, also mentioned in [6.3.2 Interface with DC-SCM](#), describes how an ARM-based server may communicate with BMCs, and proposes 5 possible requirement levels of increasing complexity. The Rhea2 HPM follows the M4 set of requirements. In addition, the Arm SystemReady Requirements specification provide a set of minimum requirements for end users to receive the expected experience. The Table below provides an excerpt of the requirements for servers following version 3.0 of that document:

Hardware or peripheral	Functionality	Minimum required
Console in and out	Console for the user to interact with the system to install and boot the OS	At least one input and one output console: Local USB keyboard and video graphics Local UART console BMC remote keyboard and video graphics Remote UART such as IPMI Serial-overLAN (SOL)
OS installation media	Media used for OS installation source	At least two separate media sources from the following groups: Local USB, or BMC remote USB virtual media NVMe drive, or SATA drive, or SAS drive Network (PXE, HTTP, HTTPS, iSCSI, or FCoE, or NVMoF) eMMC or internal SSD
OS boot media	Media used for installing and booting the target OS	At least two separate media destinations from the following groups. For others, at least one media destination from the following groups: Local USB NVMe drive SATA drive SAS drive Network (iSCSI or FCoE or NVMoF)
OS Network support	Network device for OS usage	At least one network device: Integrated Network controller PCIe network card USB network device

TABLE 8 - EXCERPT OF MINIMUM REQUIREMENT FOR ARM SYSTEM READY SERVERS ⁴

On the RISC-V side, the non-ISA RISC-V Server SoC specification⁵ describes what is expected from server-grade RISC-V processors. For manageability purposes, this document makes the following recommendations:

⁴ Source Arm SystemReady Requirements Specification v3.0

⁵ Source: <https://github.com/riscv-non-isa/server-soc/releases/tag/v1.0>



- Use of DMTF Redfish, DMTF Platform Level Data Model (PLDM) and DMTF Management Component Transport Protocol (MTCP) for in-band and out-of-band server management
- Use of DMTF Security Protocol and Data Model (SPDM) for device attestation and encryption of messages between the main processor and the BMC.
- Support of Intelligent Platform Management Interface (IPMI)
- Support open standards such as DC-SCM

In addition, the RISC-V Server SoC Specification includes 3 rules related to manageability:

- Support for an x1 PCIe lane, preferably Gen 5, but at least Gen 3, to establish a connection with the BMC.
- Use of I2C-based IPMI SSIF for in-band communication
- Support for utilizing a UART connection to the BMC, enabling the provision of a host debug console

The Management Module for the HIGHER OCP server must meet the following functional and operational requirements to ensure seamless monitoring, control, and automation of system resources.

7.6.1 Hardware Requirements

ID	Description	Attribute	Status
UCR-HW-MM-76100	Dedicated Management Interface: The module must have a dedicated management Ethernet port, supporting at least 1GbE connectivity.	M	Validated

ID	Description	Attribute	Status
UCR-HW-MM-76101	Remote Management: Independent from the main server operations, allowing remote monitoring and troubleshooting.	M	Validated

ID	Description	Attribute	Status
UCR-HW-MM-76102	Remote Power Control: Support for power on/off, reset, and diagnostics via remote interfaces.	M	Validated

ID	Description	Attribute	Status
UCR-HW-MM-76103	Sensors: Support for hardware monitoring including temperature, voltage, power consumption, and fan speeds.	M	Validated

ID	Description	Attribute	Status
UCR-HW-MM-76104	Expandability: Should support the OCP standard for future upgrades. Compliance: Must meet OCP data centre requirements.	M	Validated



7.6.2 Firmware/Software Requirements

ID	Description	Attribute	Status
UCR-SYS-MM-76200	Compliance: Must be compliant with industry standards such as Redfish, IPMI 2.0, and OpenBMC	M	Validated

ID	Description	Attribute	Status
UCR-SYS-MM-76201	Secure Firmware Updates: Support for authenticated and encrypted firmware updates.	M	Validated

ID	Description	Attribute	Status
UCR-SYS-MM-76202	Configuration Management: Remote configuration capabilities including boot order control, BIOS settings, and power state management.	M	Validated

ID	Description	Attribute	Status
UCR-SYS-MM-76203	Logging and Alerts: Support for event logging, SNMP, and email alerts for hardware failures and security breaches.	M	Validated

ID	Description	Attribute	Status
UCR-SYS-MM-76204	Automated System Recovery: Ability to detect and recover from system failures automatically.	M	Validated

7.6.3 Security and Access Control

ID	Description	Attribute	Status
UCR-ENV-MM-76300	Role-Based Access Control (RBAC): Granular access permissions for different user roles.	M	Validated

ID	Description	Attribute	Status
UCR-ENV-MM-76301	Secure Boot and TPM Support: Ensure the integrity of firmware and software components.	M	Validated

ID	Description	Attribute	Status
UCR-SYS-MM-76302	Encryption and Authentication: Support for TLS 1.2/1.3, SSH, and multi-factor authentication.	M	Validated

7.6.4 Root of Trust

data centres have to provide security guarantees to users despite the sheer complexity and scale of those facilities. Each server consists of many hardware and software components sourced from different vendors, and is maintained mostly in a remote manner by a small team of administrators. In order to hold strong security guarantees, mechanisms are required to ensure that all those components are to be trusted. Software components pose particular challenges since they are by definition mutable. Typically, this mechanism has to assess if a given firmware update is to be trusted or not.

Root of Trusts (RoT) provide such a mechanism. Typically, a Platform Root of Trust placed on the Management Module is used to validate the firmware running on the DC-SCM; and then the firmware running on the motherboard and/or other Root of Trusts, which in turn will ensure that other subsystems are to be trusted.

At present, there are two main open RoT implementations:

- OpenTitan is maintained by lowRISC, and is based on a RISC-V ibex core.
- Caliptra is maintained by the CHIPS Alliance project

ID	Description	Attribute	Status
UCR-SYS-RT-76400	The RoT must measure the integrity of the first boot firmwares executed in the DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-SYS-RT-76401	The RoT must provide service to measure the integrity of the HPM hardware and software components	M	Validated

ID	Description	Attribute	Status
UCR-SYS-RT-76402	The RoT must provide secure mechanisms to ensure transfer of ownership	M	Validated

ID	Description	Attribute	Status
UCR-SYS-RT-76403	The RoT must provide anti-rollback protection	M	Validated

ID	Description	Attribute	Status
UCR-SYS-RT-76404	The RoT must support secure A/B updates	M	Validated

7.7 Integrated Servers Requirements

According to the DoA of the HIGHER project, the developed HPMs and DC-SCM modules are expected to be integrated into an OCP-compliant HIGHER chassis, together with off-the-shelf OCP-compliant modules for power distribution, network connectivity, and storage.

The original intent was that the DC-SCM, the Rhea2-based HPMs and the Accelerator HPM based on EUPilot are integrated together in the same chassis. As demonstrated in the following, this original intent will not be fully met as it is practically complex to integrate the dual socket Rhea2-based HPM and the Accelerator HPM together on the same floor. However, this should not be considered as a major

issue or deviation considering that there is no use case that requires this co-integration in the same chassis.

7.7.1 Decision path

The topic of server structure and more especially the combination of HPM form factors that it should support has been intensively discussed in the scope of the task 2.1, ending up with different paths having been explored and either finally abandoned with justifications or still candidate as the solution that will be implemented.

The below flow diagram in **Figure 24** summarizes the various paths that have been explored for the form factor of the Rhea2-based HPM and subsequently in this chapter, rational is brought-up, providing ground to decisions that have been taken at the time of completion of the task 2.1 activities.

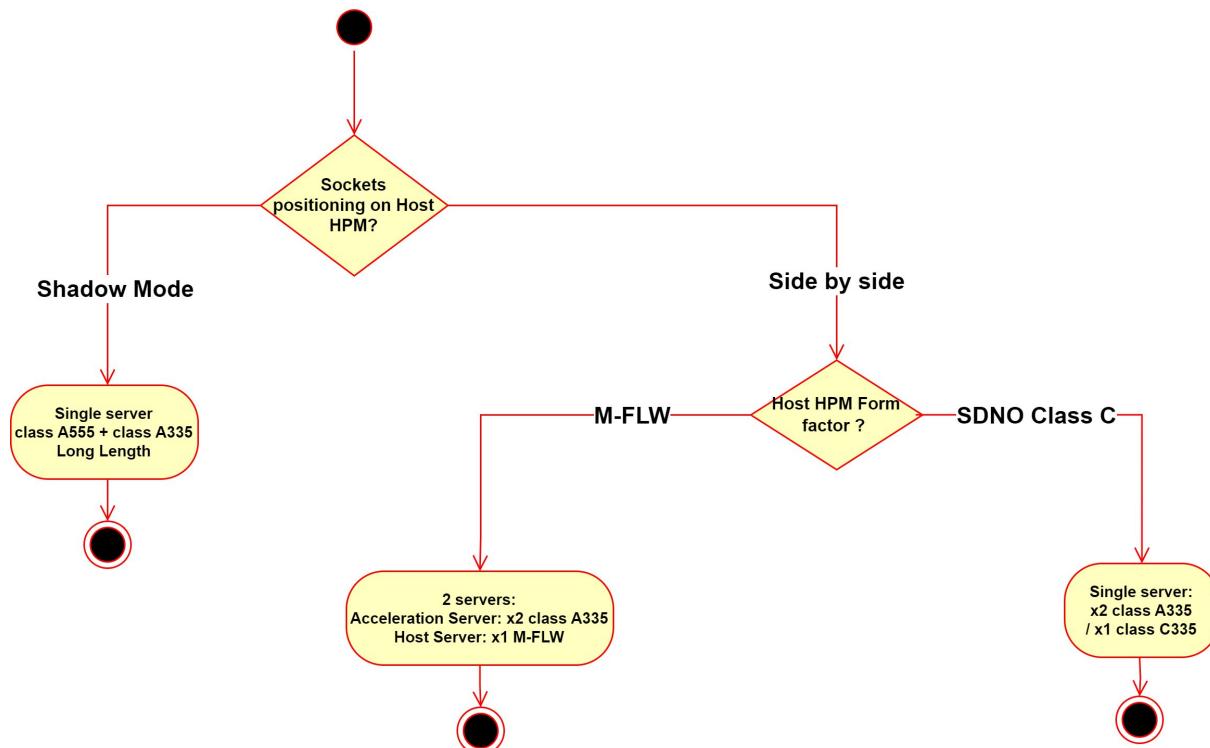


FIGURE 24 - RHEA2-BASED HPM FORM FACTOR DECISION PATH

Dual sockets positioning: side-by-side vs shadow mode

The first level of decision concerns the respective positioning of the Rhea2 sockets on the Host HPM, which can be in shadow mode or side-by-side.

In a shadow mode configuration, a class A form factor has been considered with the maximal authorized length = 555mm, ending up in a chassis combining a Host HPM SDNO class A555 and an Accelerator HPM SDNO class A335, as shown in the **Figure 25** below:

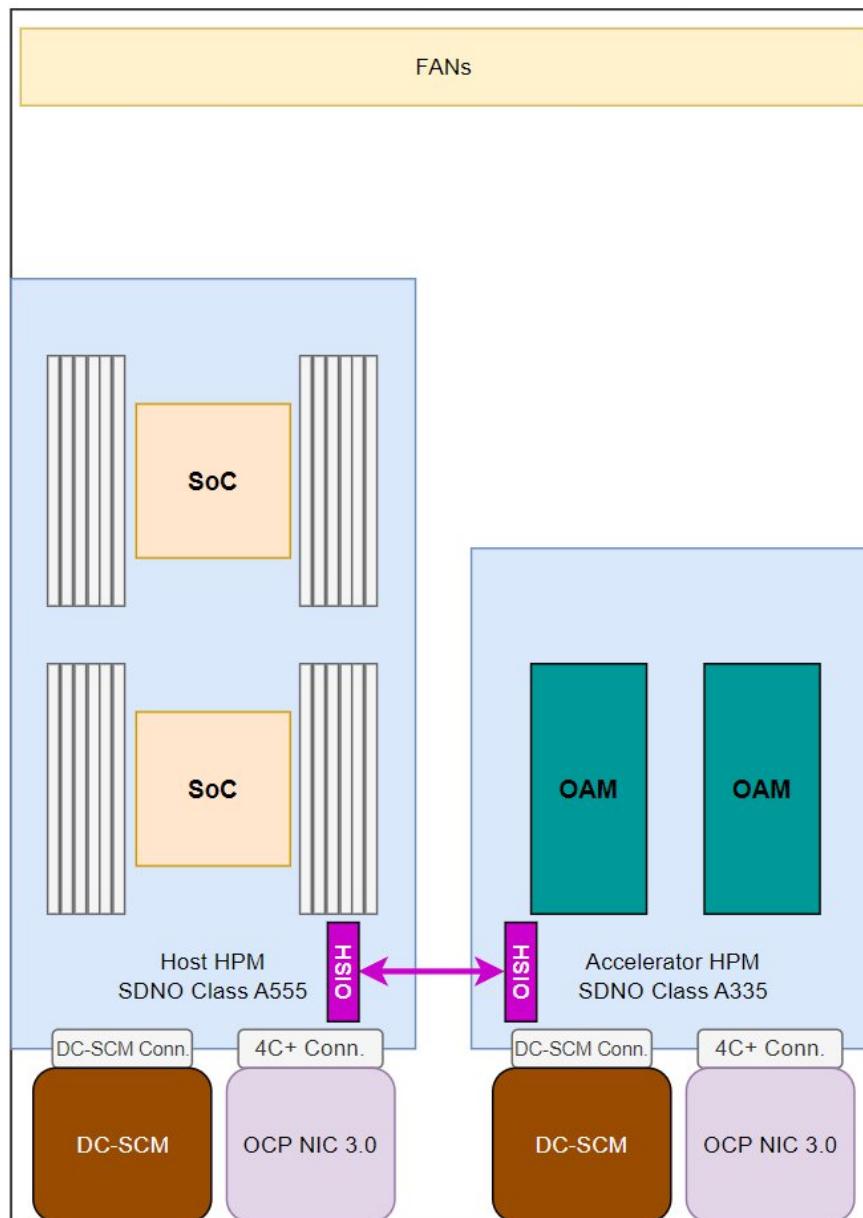


FIGURE 25 - RHEA2-BASED HPM - SDNO CLASS A555, SHADOW MODE

The relevance of the shadow mode configuration has been studied by SiPearl and the conclusions are that it is too much constraints for a prototyping activity, as shown in the **Figure 26**:

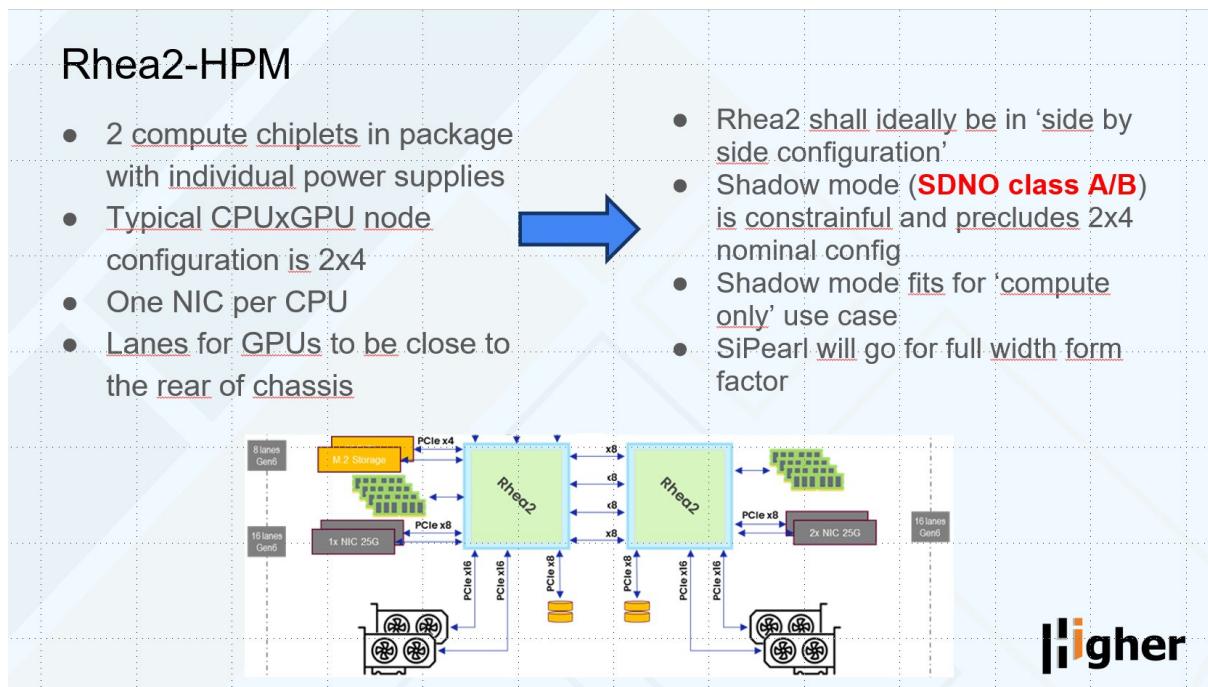


FIGURE 26 - SHADOW MODE PROS AND CONS

Rhea2 has 2 chiplets in the package and comes with the need for 3 'big' power supplies for each chiplet as these latter should be able to be powered differently according to the workloads.

Therefore, the power supply will be dense on top and bottom of the CPUs and will have to come from both sides. In shadow mode, there are constraints on the PCIe lanes for the chip that will be far from the rear of the chassis. This introduces a lot of complexity that cannot be afforded, especially when accelerator cards or GPUs are involved.

Additionally, SiPearl intents to avoid passing the PCIe lanes in the middle of the power supply mentioned previously. Moreover, SiPearl needs one NIC per CPU where class A/B offers only one. Workaround options with slotted NICs, as proposed by 2CRSi, are feasible in theory but this would take up the place reserved for GPUs or accelerator cards and bring undesirable asymmetry between the Rhea2 SoCs for what concerns network access.

As a conclusion, the work group agreed to reject shadow mode configuration and the form factor for the Host Rhea2-based HPM shall be full width (either M-FLW or SDNO Class C) so that the sockets can be arranged side-by-side.

Comparison of M-FLW with M-SDNO Class C

For the Host specific server containing a dual socket HPM, SiPearl has studied the pros and cons of the 2 eligible form factors for the dual socket Rhea2 based HPM, M-FLW and M-SDNO Class C.

The **Figure 27** below is a high-level synthesis of the study, considering a possible combination with GPUs, which is out of the scope of the HIGHER project but a concern for further integration derivative use cases:

	M-FLW	SDNO – C335
Flexibility	-	+++
Power capability	-	++
Server format	+	++
Future	-	+
Today market	++	-
PCB area	+	+++
Rhea Air flow	-	+
GPU Air flow	+	-
Total	5	12

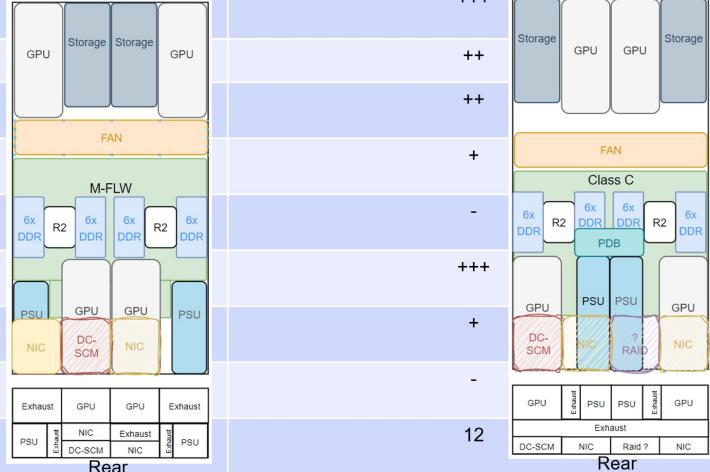


FIGURE 27 - M-FLW VS. SDNO CLASS C

The rational for each evaluated criteria is given here, along with the feedback from 2CRSi:

- Flexibility: On the sizing and dimensioning of the power supply
 - **2CRSi:** sizing and dimensioning of the power supply is the same as for M-FLW, which is up to 3200W, and there is also the possibility to connect a cable to power up the HPM
- Power Capability: With M-FLW, the power of the PSUs is limited, whereas with SDNO Class C, it is possible to use a power supply capable of delivering more current
 - **2CRSi:** The power limit of the PSU is the power limit of a M-CRPS power supply, which is 3200W, be it for M-FLW or M-SDNO Class C
- Server Format: M-FLW is slightly longer and leaves less room in the prototyping front area.
 - **2CRSi:** 348mm for M-FLW vs 335mm for M-SDNO Class C
- Future proof: SDNO Class C is promoted actively by OCP. However, it is recognised that this is not sufficient and the intersection with nVidia MGX is a concern. SDNO Class C also comes with the opportunity for SiPearl to design an SDNO Class C to create a class A rather straight fully.
 - **2CRSi:** M-FLW is mostly adopted by different ODM for dual CPU requirement.
- PCB Area: SDNO Class C leaves more room, especially in those areas that are precluded for the PSUs on the M-FLW
 - **2CRSi:** There is the possibility to add storage in the PSUs' space if power comes from a Power Distribution Board (PDB)
- Rhea2 Air flow: By uplifting the PSUs, SDNO Class C offers a better flow
 - **2CRSi:** There is the possibility to uplift the PSUs with M-FLW form factor

On the side of 2CRSi, it is considered that SDNO class C does not represent a significant breakthrough compared to M-FLW, other than the modularity aspects. Leaning back on the HPM market considerations as described in chapter 7.2.4, 2CRSi is rather supporting focus on M-FLW for the Host Server.

Single chassis

This is the configuration that leverages most of the DC-MHS modularity concept, demonstrating that the same chassis can be used for hosting either (i) x2 SDNO class A or (ii) x1 SDNO class C.

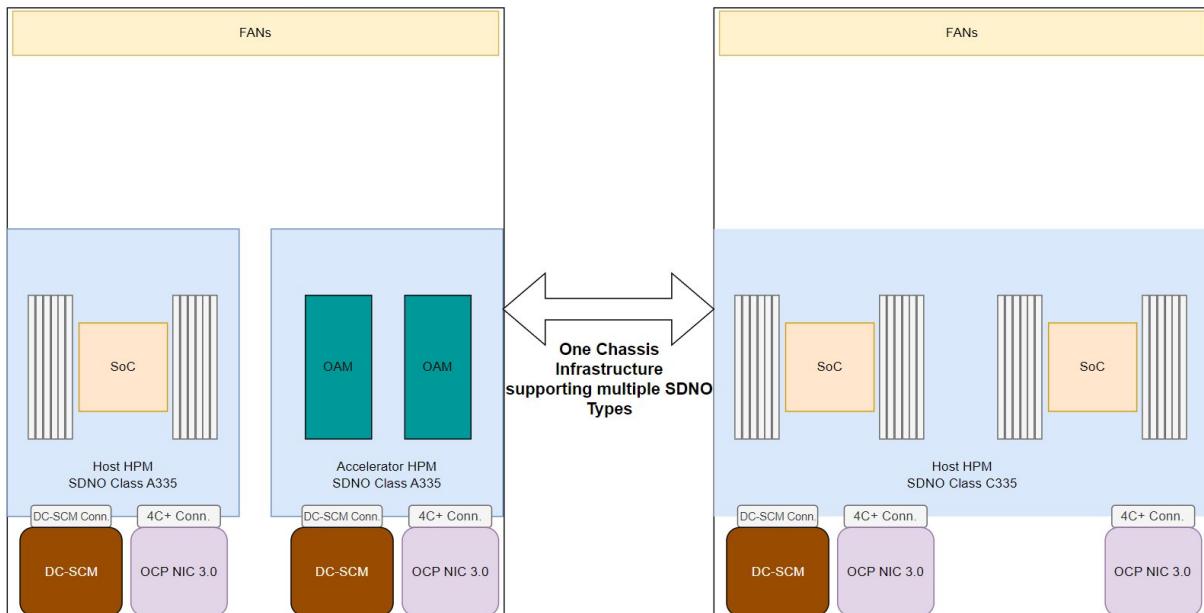


FIGURE 28 - SINGLE SDNO CHASSIS SUPPORTING EITHER X2 CLASS A335 OR X1 CLASS C335

It is important to be noted that by making an SDNO Class C HPM for the dual socket solution, SiPearl has also confirmed the intention to make a class A single socket as a straightforward derivative, which has the advantage of promoting this Rhea2-based SDNO Class A HPM as usable as a Host in the dual class-A configuration.

2 distinct chassis

In this approach, the Rhea2-based HPM is M-FLW and 2 distinct chassis are involved.

- The “Accelerator Server” incorporates 2 half-width HPMs, the first being a third party Host HPM, the second being the Accelerator HPM that will be developed in the scope of the HIGHER project.
- The Host server will include mounting holes, as defined in [M-FLW] for assembly of M-FLW HPM developed in the scope of the project.

In this approach, the Rhea2-based HPM is M-FLW which is not the favoured approach according to SiPearl’s assessment, and it would not allow SiPearl to make a half width HPM derivative, hence for the dual class A server, the Host HPM come from a 3rd party and cannot be provided by SiPearl.

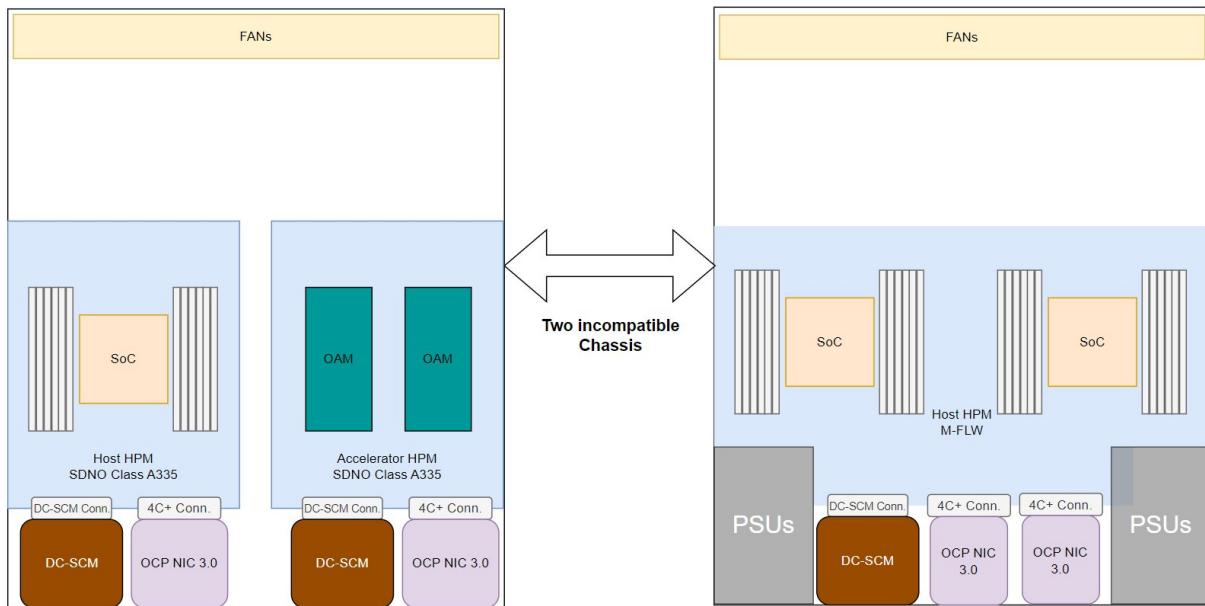


FIGURE 29 - 2 CHASSIS: SDNO DUAL CLASS A335 / M-FLW



Conclusion and decisions taken

Considering the benefit of the modular approach as proposed by DC-MHS with interchangeability of SDNO Class A HPMs with SDNO Class C HPM, the consortium has decided to go for the single “SDNO” chassis option, as described in **Figure 28 - single SDNO chassis supporting either x2 class A335 or x1 class C335**.

Nevertheless, taking into account the market considerations, as expressed in chapter 7.2.4, and in consideration of the original intent of the DoA of the HIGHER project to develop a chassis compatible with M-FLW, it has been agreed that 2CRSi would develop such a chassis that will be equipped with a 3rd party (x86) M-FLW. This M-FLW Host server will be combined with the “SDNO” chassis when equipped with 2 Accelerator HPMs and demonstrate the 2xCPU/4xOAM configuration.

Figure 30 below highlights the structures of the motherboards plan for the SDNO and the FLW chassis:

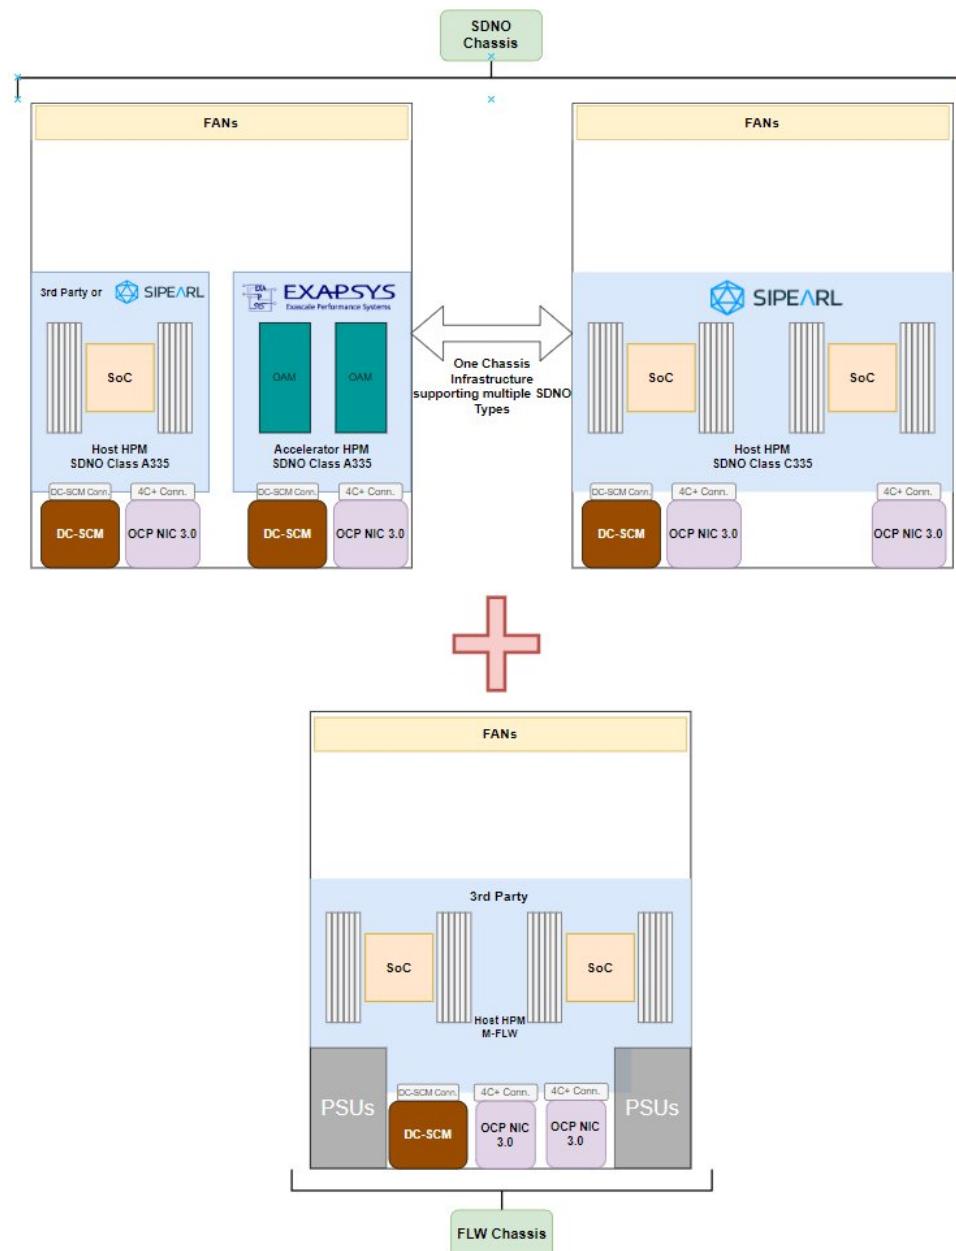


FIGURE 30 - HIGHER CHASSIS FINAL SET

7.8 SDNO Chassis

7.8.1 Dual class-A Configuration

In the dual class-A configuration, the SDNO chassis embeds 2 class A335 HPMs, side by side.

Host HPM combined with Accelerator HPM

In the nominal use case, the left-hand side HPM is a Host HPM and the right-hand side HPM is the Accelerator HPM. The SDNO chassis shall make provision for one DC-SCM for each HPM and an OCP-NIC dedicated to the Host HPM. The SDNO chassis shall also allow for the SoC of the HPM to be connected to each of the OAMs of the Accelerator HPM by means of PClex4. **Figure 31** below illustrates the floor level of the SDNO chassis in the nominal use case:

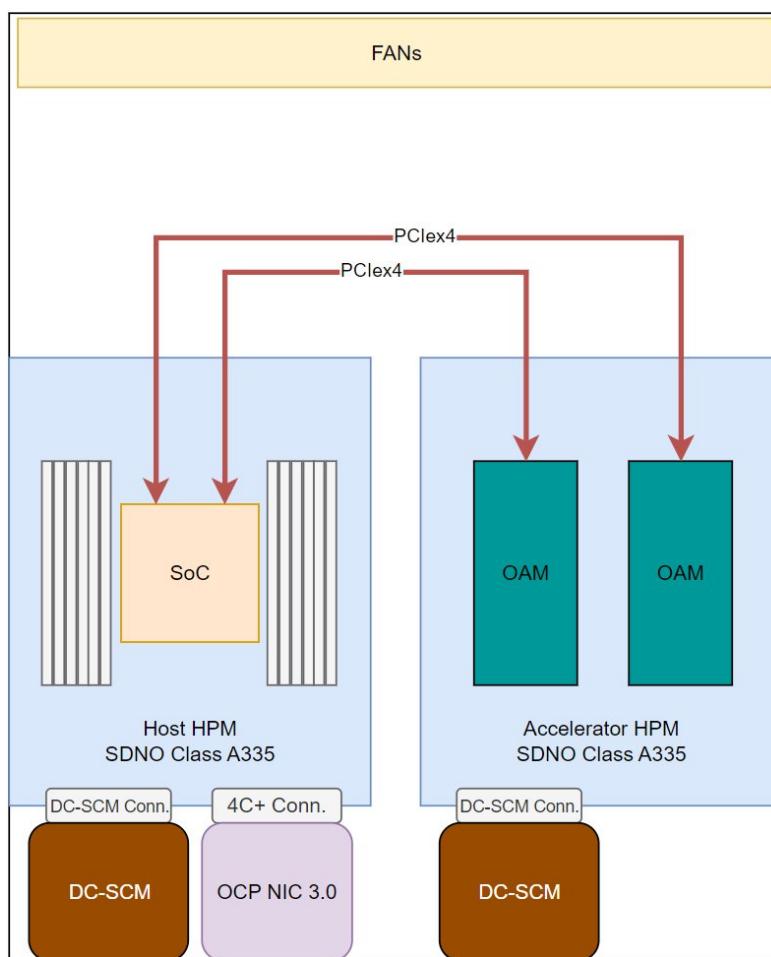


FIGURE 31 - SDNO CHASSIS WITH HOST HPM AND ACCELERATOR HPM

The Host HPM can be any commercially available class A HPM available from the market. Existing references based on x86 Intel Xeon 6700 are available from Supermicro and Pegatron. References are also available on Arm with Ampere showcasing class A HPM based on Ampere One. And finally, as described in chapter 61, SiPearl is also considering a single socket class-A Rhea-2 based HPM, as a derivative of the dual socket class C HPM.

Note: The EUPilot RISC-V chip on the OAM only has 4 lanes and therefore the size of the PCIe connection is limited to x4.

Combining 2 Accelerator HPMs

In an alternative use case, the Host HPM role is played by an Accelerator HPM, where the x4 OAMs shall be connected in an all to all configuration, as shown in the **Figure 32** below:

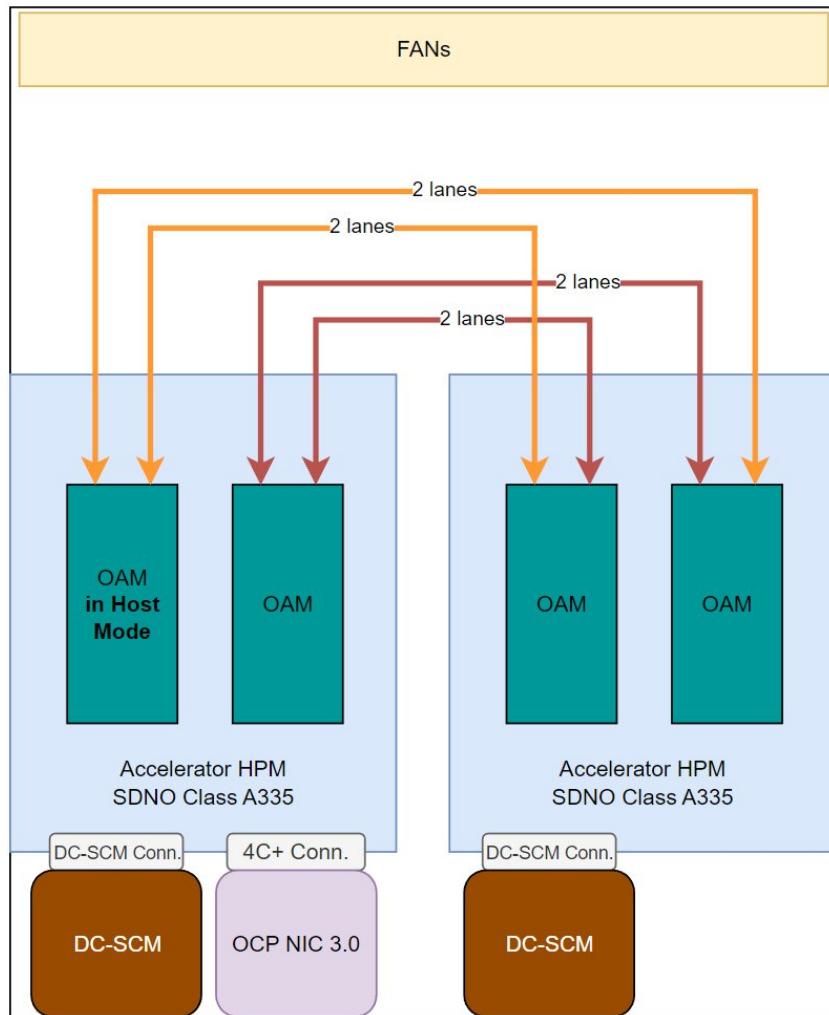


FIGURE 32 - SDNO CHASSIS WITH X2 ACCELERATOR HPMs

Combining 2 Host HPMs

In an alternative use case, the SDNO chassis is equipped with x2 Rhea2-based class A Host HPMs where both Rhea2 shall be interconnected with x4 PCIe gen6 x8, as shown in **Figure 33** below:

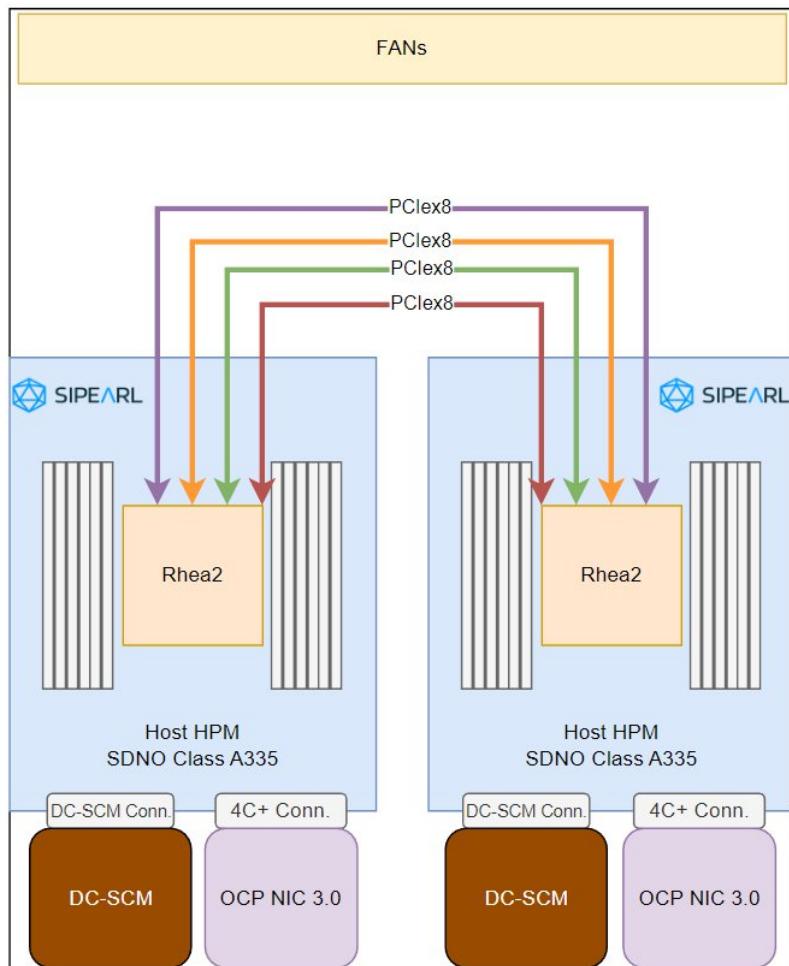


FIGURE 33 - SDNO CHASSIS WITH x2 HOST HPMs

Combining the SDNO Server with a Host Server

In another use case, the SDNO Chassis is equipped with 2 Accelerator HPMs that are controlled by a Host server, be it another SDNO chassis or an M-FLW chassis, as depicted in **Figure 34** below :

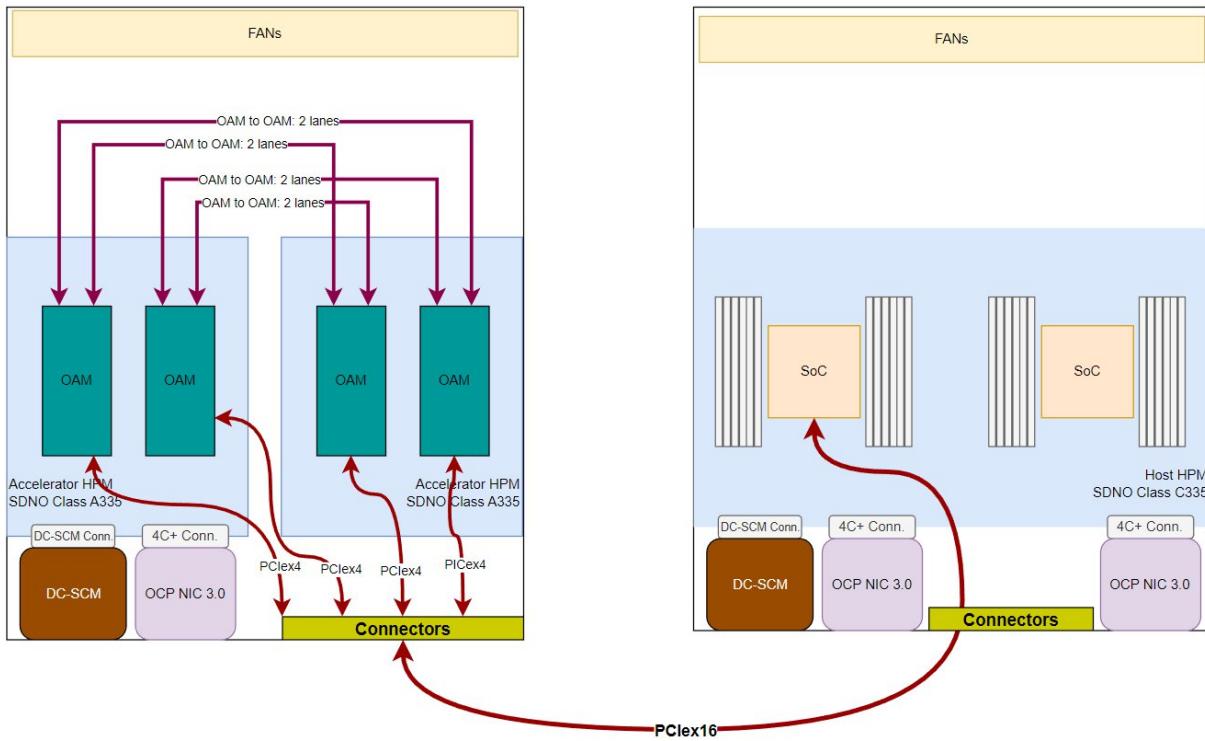


FIGURE 34 - SDNO CHASSIS WITH X2 ACCELERATOR HPMs COMBINED WITH ANOTHER CHASSIS

The configuration above assumes that one of the EUPilot (OAM) plays the role of the master, as described in the Riser project D2.1. The RISC-V chip only has 4 lanes, therefore the size of the PCIe connection for each OAM is limited to x4.

7.8.2 Single class-C Configuration

In the single class-C configuration, the OCP Server embeds 1 class C335 Host HPMs.

The SDNO chassis shall make provision for 1 DC-SCM and 2 OCP-NICs dedicated to the Host HPM. The figure below illustrates the floor level of the SDNO server:

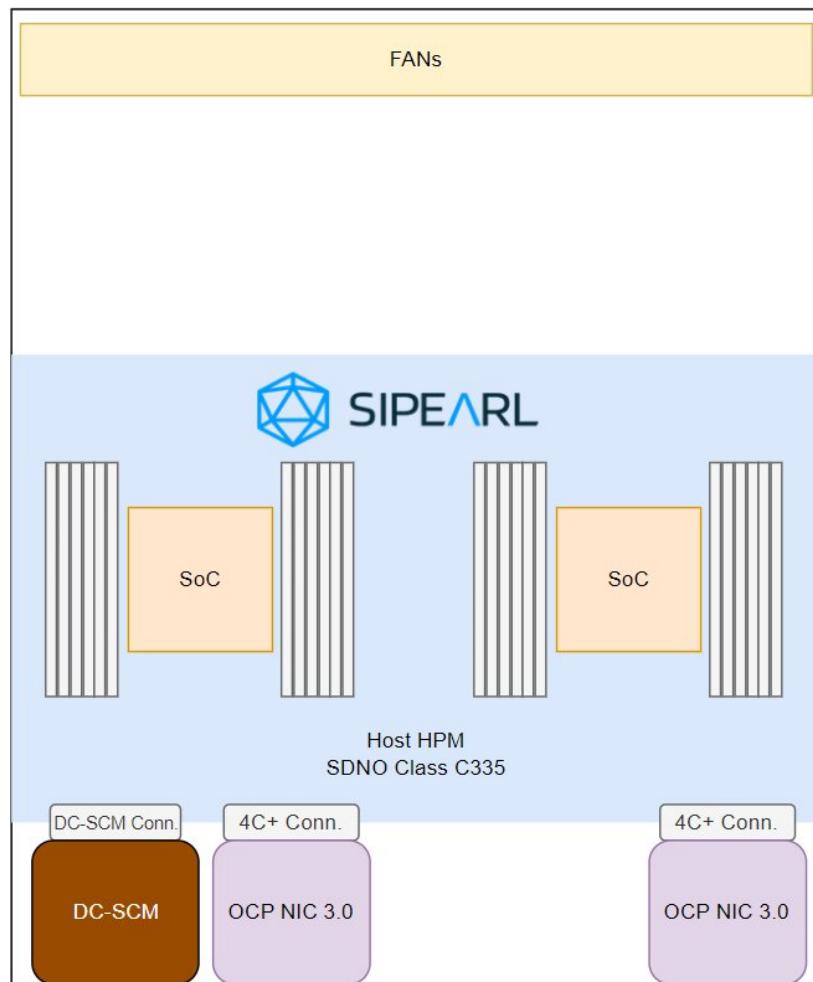


FIGURE 35 - SDNO CHASSIS WITH CLASS C HOST HPM

7.8.3 Upper layers

This chapter describes a possible structure of the chassis on the upper layers. The examples and figures are for the dual class-A HPMs configuration but the same applies to the single class C HPM. It is informative as there are other possibilities that will be explored during the Task 2.2 activities.

PCIe slots layer

The floor layer is 1/2 U and the first U can be complemented by x2 PCIe Single Slot coming on top of the DC-SCM and OCP-NIC, as shown in the top view **Figure 36** below:

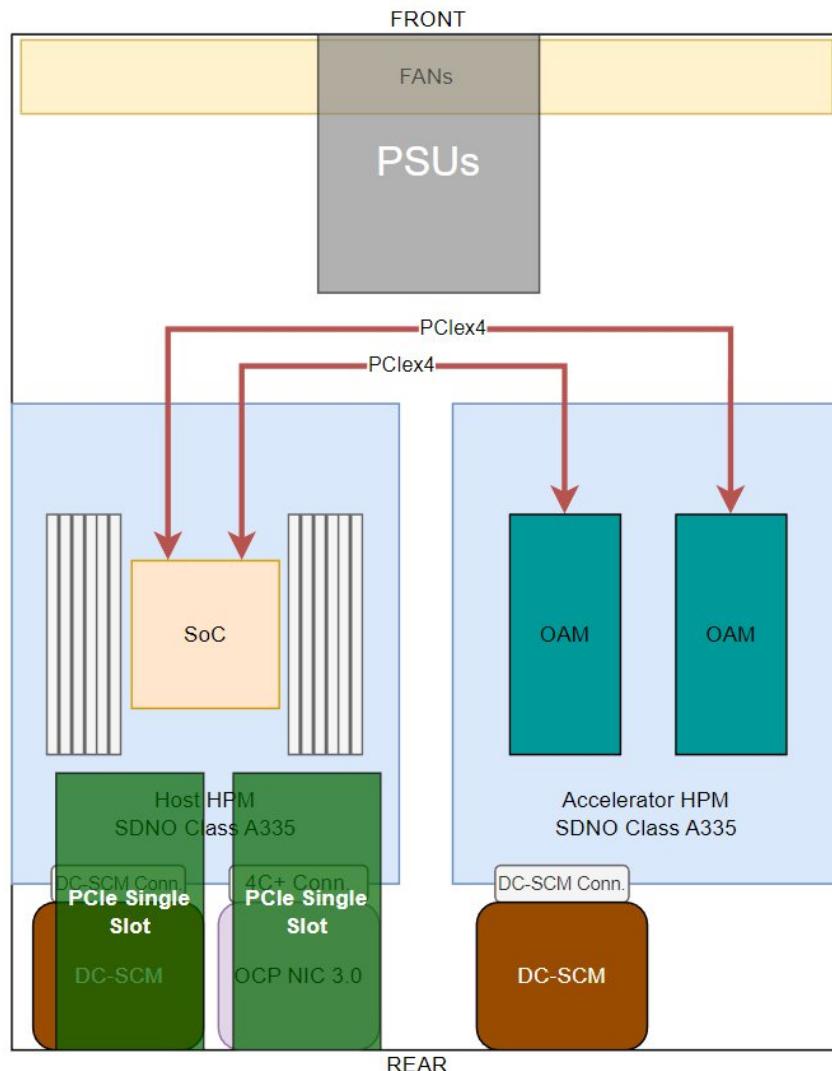


FIGURE 36 - SDNO SERVER - PCIe SLOT LAYER (INFORMATIVE)

Power Supplies and Storage Layer

The server will be completed with a second U dedicated to power supplies and storage as shown in the top view in **Figure 37** below:

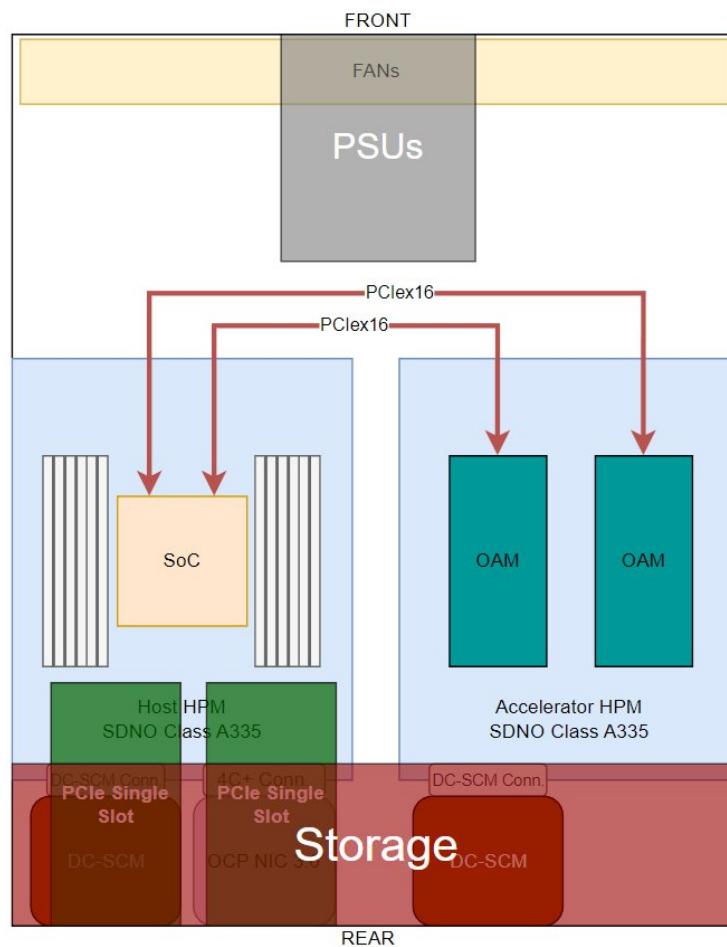


FIGURE 37 - SDNO SERVER - PSU AND STORAGE LAYER (INFORMATIVE)

Figure 38 below is a representation of the rear view of the SDNO server:

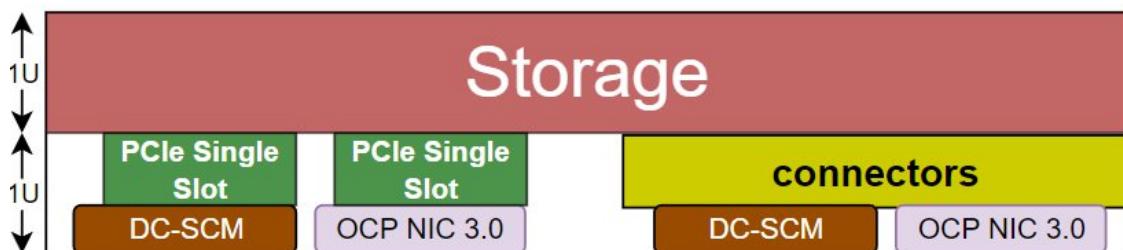


FIGURE 38 - SDNO SERVER - REAR VIEW (INFORMATIVE)



7.8.4 SDNO chassis Requirements

HPM Hosting

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78410	The [SDNO-Chassis] shall be capable of hosting one HPM compliant with M-SDNO Class C335	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78411	The [SDNO-Chassis] shall be capable of hosting two HPMs compliant with M-SDNO Class A335	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78412	The [SDNO-Chassis] shall comply with the M-SDNO Base Specification Compliance table, as defined in [M-SDNO] chapter 4	M	Validated

DC-SCM and OCP-NIC

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78420	The [SDNO-Chassis] shall be capable of hosting a <u>primary</u> DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78421	The [SDNO-Chassis] shall be capable of hosting a <u>secondary</u> DC-SCM	O	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78422	The [SDNO-Chassis] shall allow the SDNO Class C HPM to connect to the <u>primary</u> DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78423	The [SDNO-Chassis] shall allow the <u>primary</u> SDNO Class A HPM to connect to the <u>primary</u> DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78424	The [SDNO-Chassis] shall allow the <u>secondary</u> SDNO Class A HPM to connect to the <u>secondary</u> DC-SCM	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-	The [SDNO-Chassis] shall be capable of	M	Validated



ID	Description	Attribute	Status
CHS-78425	hosting two OCP-NIC		

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78426	The [SDNO-Chassis] shall allow the SDNO Class C HPM to connect to both OCP-NIC	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78427	The [SDNO-Chassis] shall allow each SDNO Class A HPM to connect to one of the OCP-NIC	M	Validated

High Speed I/O Connectivity

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78430	The [SDNO-Chassis] shall allow the 2 SDNO Class A HPMs to be connected with each other via M-XIO (or MCIO) to M-XIO (or MCIO) links, supporting up to 32 lanes	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78431	The [SDNO-Chassis] shall allow the SDNO Class C HPM to be cable-connected to the x2 PCIe slots from dedicated M-XIO or MCIO connectors on the HPM, supporting up to 16 lanes per connection	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78432	The [SDNO-Chassis] shall allow either of the SDNO Class A HPM to be cable-connected to the x2 PCIe slots from dedicated M-XIO or MCIO connectors on the HPM, supporting up to 16 lanes per connection	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78433	The [SDNO-Chassis] shall allow each SDNO Class A HPM to be cable-connected with rear panel connector(s) from dedicated M-XIO or MCIO connector(s), supporting up to 8 lanes per HPM, hence 16 lanes in total	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78434	The [SDNO-Chassis] shall allow SDNO Class C HPM to be cable-connected with rear panel connector(s) from dedicated M-XIO or MCIO connector(s), supporting up to 16 lanes	M	Validated

Power Supply

Each configuration should be able to integrate two M-CRPS of at least 3200W. Redundancy is mandatory in servers and with two PSUs in the server it will be possible to have N+1 redundancy. M-CRPS is the latest specification from OCP. It standardizes the management and monitoring of the PSUs data compared to the CRPS specification.

The server should be compatible with both types of power supply, CRPS and M-CRPS. For the dimension of the PSUs, the server should accept 73.5mm for the width and 185mm for the length. Both PSUs will be connected to a Power Distribution Board (PDB).

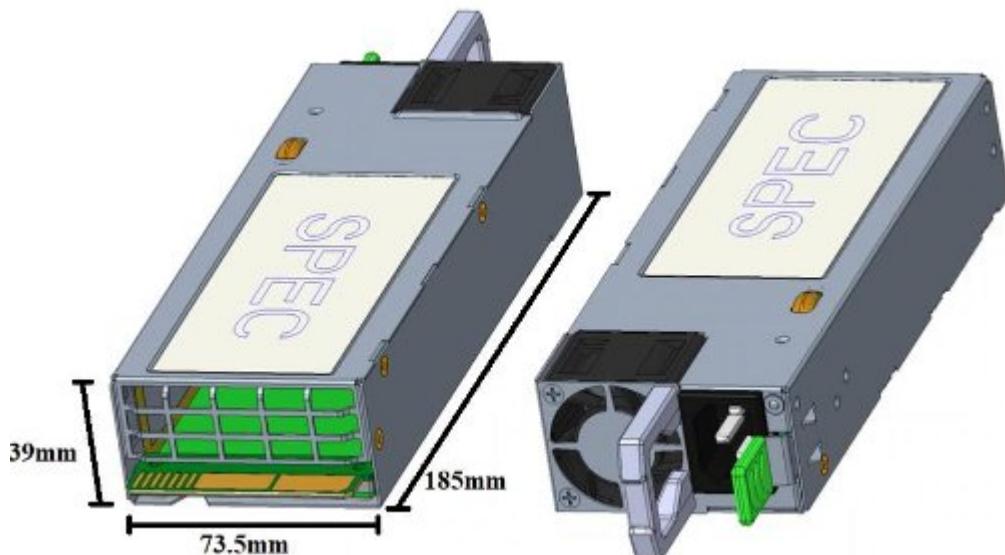


FIGURE 39 - SDNO SERVER - POWER SUPPLIES

All the required connectors (M-PIC connectors, management connectors and fan connectors) should be available on the PDB with enough power budget to power up every component of the server. The management connector will be connected to the HPM with cables (one per HPM SDNO Class A, and one per HPM FLW and SDNO Class C).

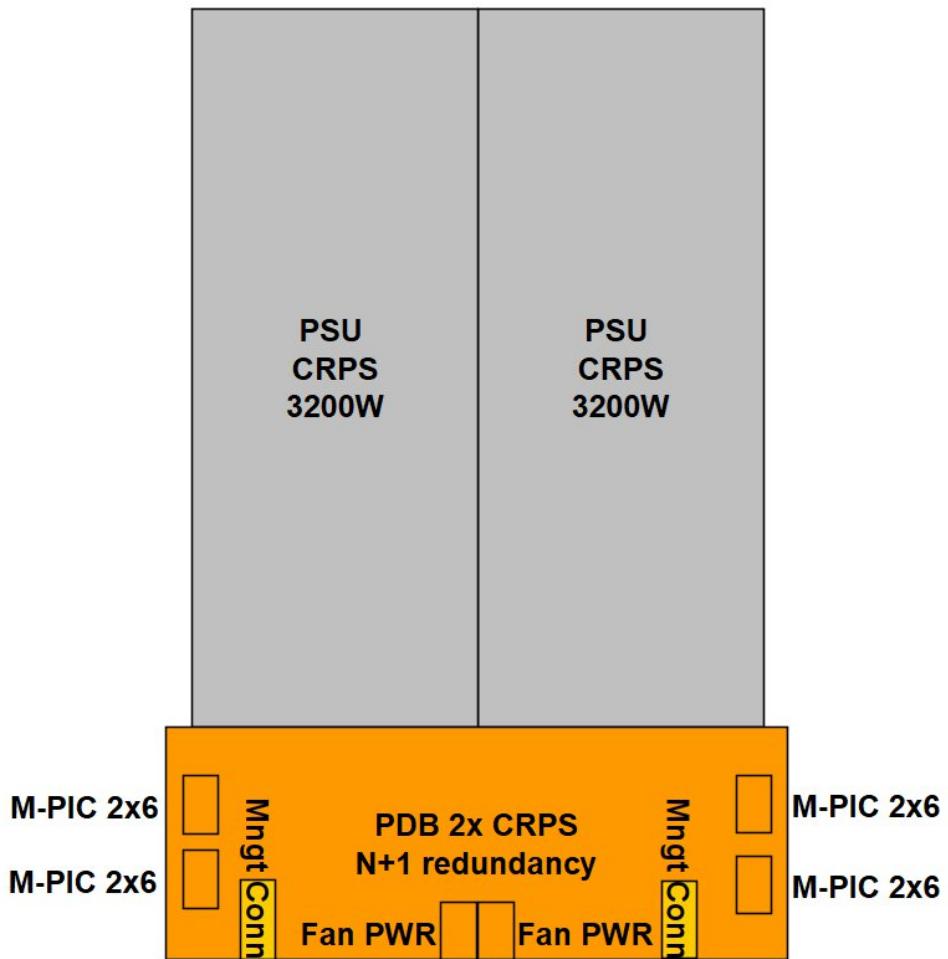


FIGURE 40 - SDNO SERVER - POWER DISTRIBUTION BOARD

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78450	The [SDNO-Chassis] shall support x2 M-CRPS PSUs of at least 3200W each	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78451	The [SDNO-Chassis] shall support either x2 CRPS PSUs of at least 3200W each	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78451	The [SDNO-Chassis] shall support PSUs with 73.5mm for the width and 185mm for the length	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78440	The [SDNO-Chassis] shall support 12V Power supply for the SDNO Class C HPM by use of one PDB, according to	M	Validated



ID	Description	Attribute	Status
	the HPM power distribution architecture requirements as defined in chapter 10.1.5 of [M-MPIC]		

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78441	The [SDNO-Chassis] shall support 12V Power supply for each of the SDNO Class A HPMs by use of a single PDB, according to the HPM power distribution architecture requirements as defined in chapter 10.1.5 of [M-PIC]	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78451	The [SDNO-Chassis] shall ensure connection of both PSUs to the PDB	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78451	The PDB of the [SDNO-Chassis] shall support x2 PDB management connectors Type2, according to chapter 10.2.14.2 of [M-PIC]. The management connector will be connected to the HPM(s) with cables (one per HPM SDNO Class A, and one per HPM FLW and SDNO Class C)	M	Validated

Storage

Storage will be possible with different form factors:

- M.2 drives
- E1.S drives
- 2.5" U.2 drives

For the M.2 drives, if used, they should be located directly on the HPM and should accept 2280/22110 drives. Up to 4x PCIe lanes are required per M.2 NVMe drive. Refer to chapters 7.3.5 and 7.4.5 for related requirements application to the Rhea2-based HPMs.

Other type of storage will be located on the second U of the server. The chassis should be able to integrate different cages for NVMe 2.5" drives. The PDB or HPM should have a dedicated power connector for the storage.

NVMe 2.5" drives can be directly integrated in existing cages such as the [N-49NVMS](#) from Vision3.

E1.S drives will be connected to a specific board. This board can host up to 4x E1.S drives. This board will have 2x MCIO 8x connectors on it. A specific cable will be designed to connect the HPM to the E1.S Host board. It should be possible to connect the HPM to the E1.S Host board with MCIO 8x to MCIO 8x cable or M-XIO 16x to 2x MCIO 8x. Power could come either from the HPM or the PDB.

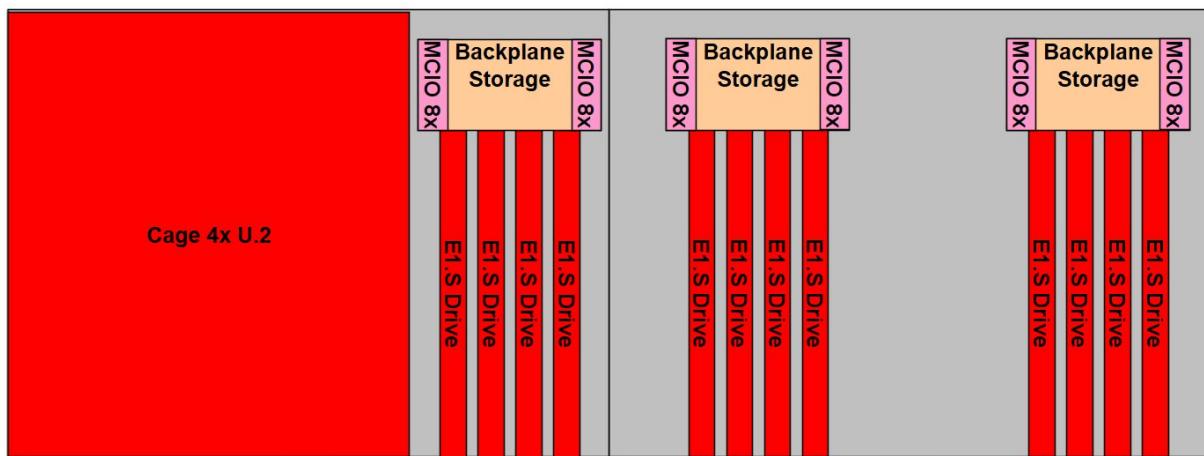


FIGURE 41 - NVME STORAGE CAGE AND E1.S DRIVES (INFORMATIVE)

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78460	The [SDNO-Chassis] shall support integration of M.2 drives located directly on HPM when supported by the HPM	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78461	The [SDNO-Chassis] shall support up to x2 specific boards integrating E.1S drives. Each board shall integrate up to x4 E1.S drives and support 2x MCIO 8x connectors	M	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78461	The [SDNO-Chassis] shall support powering of the E1.S specific board from the PDB	O	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78461	The [SDNO-Chassis] shall support powering of the E1.S specific board from one of the HPM	O	Validated

ID	Description	Attribute	Status
UCR-HW-SDNO-CHS-78462	The [SDNO-Chassis] shall support x1 cage containing x4 2,5" NVMe SSDs.	R	Validated



7.9 FLW Chassis

The FLW chassis embeds 1 M-FLW Host HPM.

The FLW chassis shall make provision for 1 DC-SCM and either 2 OCP-NICs dedicated to the Host HPM or 1 OCP NIC and 2 E1.S drives. The chassis will be 19" wide and of 2U height. The chassis will be air cooled.

The figure below illustrates the floor level of the FLW server:

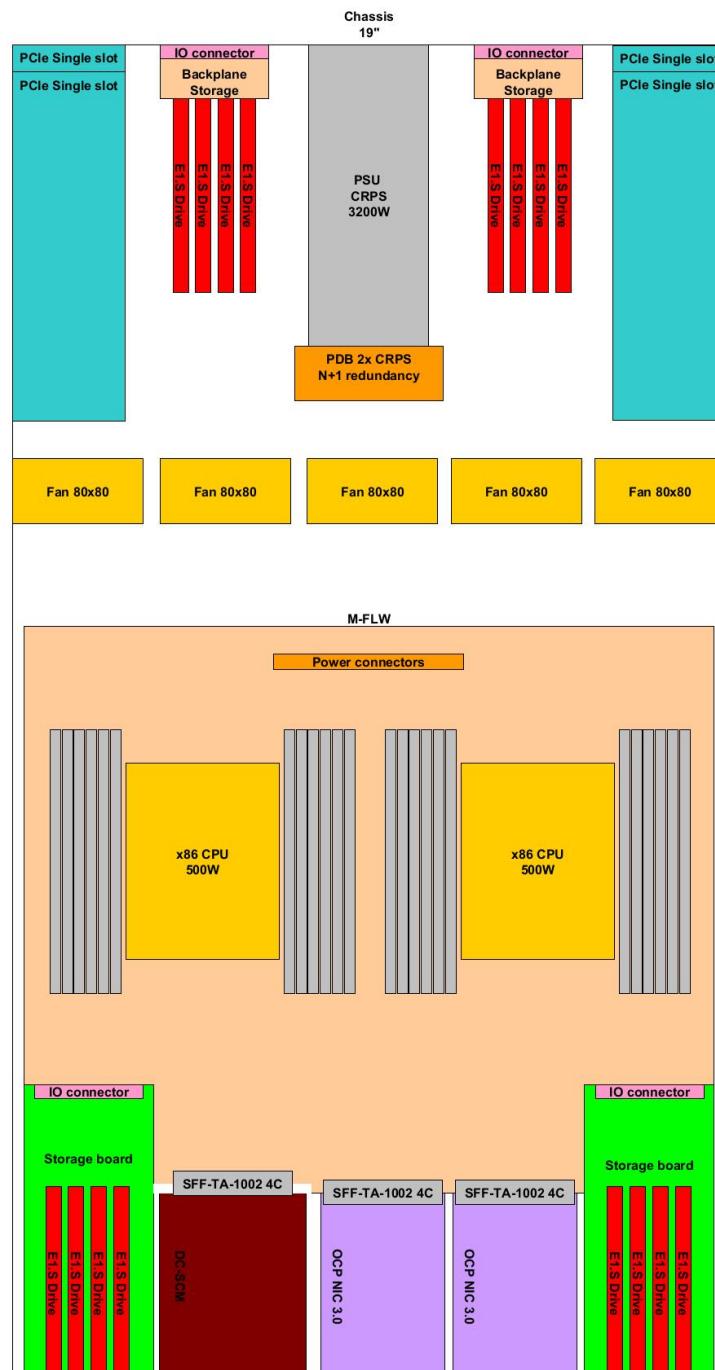


FIGURE 42 - M-FLW CHASSIS WITH 3RD PARTY M-FLW HOST HPM



In order to be as modular as possible, it should be possible to have PSUs connected directly to the FLW HPM or deported and connected to a PDB. If PSUs are deported, a storage board will fill the PSU's space. Fans with 80x80mm dimension will be in the server. Picking fans of 80x80 instead of 40x40mm improves cooling and power efficiency. It should be possible to integrate up to 4 PCIe 16x boards (HHFL or HHHL form factor) at the front of the server, at the floor level. Storage can also be added at the front of the server if there is enough PCIe lanes available from the HPM.

PSUs at the front will be on top of each other. This configuration is chosen to ease a possible transition to a 21" OCP ORv3 form factor server in the future if required.

7.9.1 Acceleration HPM

On the second U of the server, we will have the accelerator modules. The chassis should be able to integrate two M-SDNO Class A HPM with OAM accelerators. If necessary, a Host Interface Board (HIB) will be implemented in the chassis in order to connect the Host (CPU) on the HPM to several PCIe devices. A PCIe switch will be on the HIB.

More options for storage and PCIe CEM cards (network cards or RAID cards for example) will be available on the second U of the chassis.

Cables will be required for the connection between host and accelerator. Cables will also be used if there is connection between Host and HIB and HIB to accelerators.

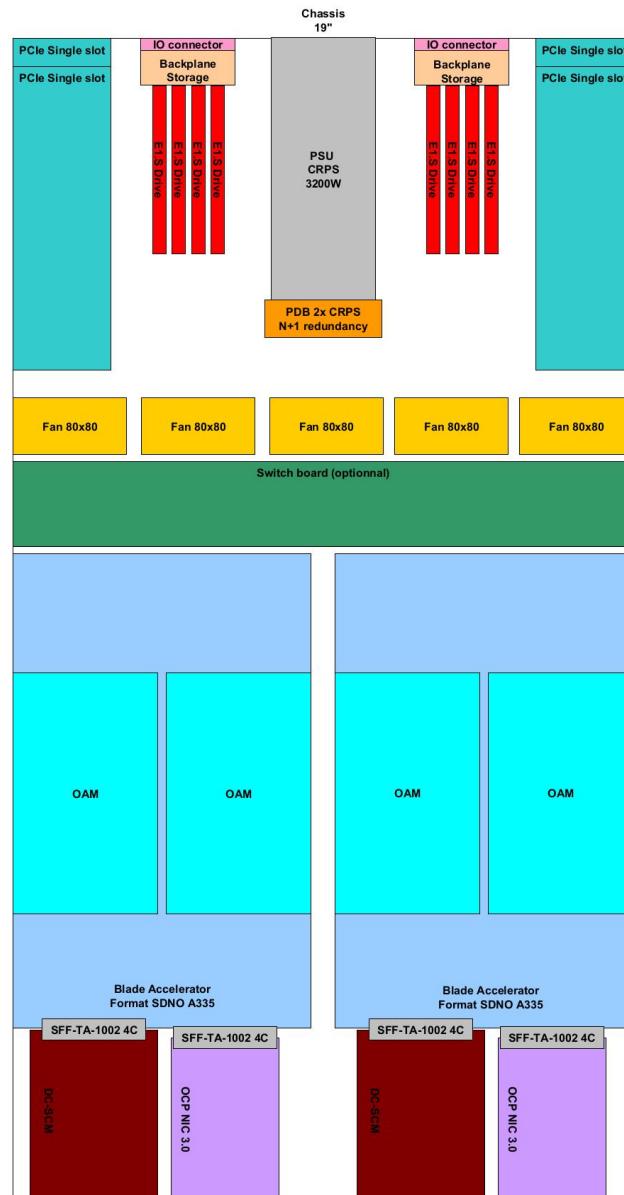


FIGURE 43 - FLW CHASSIS – SECOND U WITH ACCELERATOR HPMs

7.9.2 FLW chassis Requirements

HPM Hosting

ID	Description	Attribute	Status
UCR-HW-FLW-CHS-79210	The [FLW-Chassis] shall be capable of hosting one HPM compliant with [M-FLW]	M	Validated

ID	Description	Attribute	Status
UCR-HW-FLW-CHS-79211	The [FLW-Chassis] shall comply with the W Base Specification Compliance	M	Validated



	table, as defined in [M-FLW] chapter 4	
--	--	--

DC-SCM and OCP-NIC

Same requirements as applicable to SDNO Class C in chapter [7.8.4 SDNO chassis Requirements](#).

High Speed I/O Connectivity

Same requirements as applicable to SDNO Class C in chapter [7.8.4 SDNO chassis Requirements](#).

Power Supply

Same requirements as applicable to SDNO Class C in chapter [7.8.4 SDNO chassis Requirements](#).

Storage

Same requirements as applicable to SDNO Class C in chapter [7.8.4 SDNO chassis Requirements](#).



8 Conclusion and Next Steps

This report is the first WP2 deliverable covering the work carried out in T2.1: Requirements and use cases refinement that is used to gather inputs from stakeholders relating to the four use-cases as well as input from external stakeholders. The requirements and performance metrics collected here will inform T2.2 “Specifications and Architecture Design”, as part of follow-up deliverable D2.2 (due: M6 / June’2025) and will set targets and KPIs for the evaluation phase of the project. The targets defined are to be referenced during the evaluation phase of the project and for setting additional platform-level KPIs. Moreover, T2.3: Verification Framework, will focus on these targets, generating three successive releases of the project’s verification Framework and associated evaluation results (due for release as deliverables D2.3, D.24, and D2.5 / due in M18, M24 and M30, respectively).

9 Appendix

9.1 Acronyms and Abbreviations

Term	Definition
ACL	Access Control List
AI	Artificial Intelligence
AWS	Amazon Web Services
BMC	Baseboard Management Controller
CAGR	Compound Annual Growth Rate
CCI	Common Chassis Interval
CHI	Coherent Hub Interface
CML	Coherent Mesh Link
CPM	CXL-memory Pool Manager
CRPS	Common Redundant Power Supply
CSP	Cloud Service Providers
CXL	Compute Express Link
DC-SCI	Datacenter-ready Secure Control Interface
DC-SCM	Datacenter-ready Secure Control Module
DC-MHS	data centre Modular Hardware System
DNO	DeNsity Optimized
DoA	Description of Actions
EPAC	European Processor ACcelerator
EPI	European Processor Initiative
FLW	Full Width
GCC	GNU Compiler Collection
GPP	General Purpose Processor
GPU	Graphical Processing Unit
HPC	High Performance Computing
HPE	Hewlett Packard Enterprise
HPM	Host Processor Module
ISA	Instruction Set Architecture
KVM	Keyboard-Video-Mouse
LLVM	Low-Level Virtual Machine
LTPI	LVDS Tunneling Protocol & Interface
MCTP	Management Component Transport Protocol
ML	Machine Learning
MPI	Message Passing Interface
NDP	Near-Data Processing
NIC	Network Interface Card
NVMe	Non-Volatile Memory Express
OAM	OCP Accelerator Module
OCP	Open Compute Project
ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
OpenMP	Open Multi-Processing
ORv3	OpenRack v3
PCB	Printed Circuit Board



Term	Definition
PCIe	Peripheral Component Interconnect express
PDB	Power Distribution Board
PESTI	Peripheral Sideband Tunneling Interface
PIC	Platform Infrastructure Connectivity
PSU	Power Supply Unit
RBAC	Role-Based Access Control
RISC	Reduced Instruction Set Computer
RISE	RISC-V Software Ecosystem
SDNO	Scalable DeNsity Optimized
SGA	Specific Grant Agreement
SME	Small and Medium-sized Enterprise
SMP	Symmetrical Multi-Processing
SNMP	Simple Network Management Protocol
TDP	Total Dissipation Power
TPM	Trusted Platform Module
TRL	Technical Readiness Level
UART	Universal Asynchronous Receiver Transmitter
UBB	Universal BaseBoard
USB	Universal Serial Bus
XIO	eXtended I/O

TABLE 9 - ACRONYMS AND ABBREVIATIONS

9.2 References

Reference	Description	Version
[M-FLW]	M-FLW Base Specification	1.2RC3
[M-DNO]	M-DNO Base Specification	1.1RC2
[M-SDNO]	M-SDNO Base Specification	1.0RC2
[M-XIO]	M-XIO Base Specification	1.04RC1
[DC-MHS]	DC-MHS Specifications	NA
[OCP-NIC]	OCP NIC 3.0 Specification	1.5.0
[OAI-OAM]	OAI-OAM Base Specification r2.0	1.0
[DC-SCM]	OCP DC-SCM Specification	Rev 2.1 Version 1.1
[M-PIC]	M-PIC Base Specification	1.11
[ORV3]	Open Rack V3 Base Specification	1.0
[OCP-LTPI]	OCP LTPI Specification	1.0
[SBMR]	Arm Server Base Manageability Requirements	2.1
[SystemReady]	Arm SystemReady Requirements Specification	3.0

TABLE 10 - REFERENCES



9.3 Position towards a processor socket standard

The definition of a processor socket standard has historically shown a great impact in the microelectronic industry, because it allowed multiple motherboards to use a variety of processors. Hence, motherboard manufacturers could reduce design costs and compute owners could upgrade their machines at a fraction of the cost. However, in the last decade, the data centre market has been transformed significantly by the advent of the OCP project first, and the advent of a mature chiplet ecosystem. These changes have significantly reduced the impact of processor sockets by providing other —more relevant— practical levels where modularity can be implemented. Nevertheless, we see a potential benefit to strive for a common processor socket for RISC-V processors, to foster the emergence of RISC-V-based solutions in data centres.

The OCP project aims to improve reusability, compatibility and composability in the data centre in many aspects. With respect to processor modularity, the OCP project proposes to add another scale of reusable component: the motherboard or Host Processor Module (HPM) in OCP parlance. An overview of the current options is presented in Section [6.2 OCP HPM Form Factors](#). In this context, motherboards for a variety of processor vendors are produced, with a strong specification on the external aspects (e.g. motherboard form factor, power supply, connectivity), whilst leaving internal aspects (e.g. processor socket, memory technology, voltage rails and cooling) to the motherboard vendors. It is worth noting that with the current market conditions, those motherboards are much cheaper than the processor and the external memory that they host. In addition, there is significant value in making different motherboards compatible, because servers also require many other elements that can then be sourced independently and match the specific needs of each data centre. This includes management modules (c.f. Section [6.6 Management Module Requirements](#)), Network Interfaces (with OCP NIC), Power Supply Units, chassis and racks for instance.

In terms of modularity, the advent of chiplet-based designs also has had a big impact in the last years, and is now redefining how processors are designed and marketed. Initially, prominent processor vendors leveraged chiplets to reuse pre-existing functions and more efficiently address the needs of their clients. Then, the chiplets ecosystem has undergone a rapid standardization process, of which the most prominent example is the UCIe Die-to-Die (D2D) standard. There are now many actors, who can provide EDA tools, design and manufacture chiplets and finally integrate chiplets together. Incidentally, the OCP project delivered a white paper that details how chiplets may be leveraged⁶. We can therefore expect that in the coming years, chiplets implementing groups of processing cores will be integrated with third-party chiplets and substrates both to improve profitability in large volume products and to create better-suited solutions for smaller markets.

The impact of the OCP and chiplet approaches also holds for RISC-V processors and lessens the need for a standard processor socket. However, that architecture experiences a very high growth rate, where many systems are being developed concurrently or shortly after related aspects are being standardized. In this context, the proposition of a common processor socket for server-grade RISC-V processors may inform the design of fore-coming chips, and reduce the amount of Non-Recurring Engineering (NRE) costs required by each actor. Therefore, in the HIGHER project, we will be in contact with sibling projects such as EUPILOT, and seek to gather common practices and insights from them, with the aim to approach the definition of a server-grade RISC-V processor socket.

⁶ Source: [Business Analysis of Chiplet-based systems and technology](#)