Co-design, deployment and operation of a Modular Data Centre (MDC) with air and direct-liquid cooled supercomputers*

Sadaf R Alam†
University of Bristol, UK
sadaf.alam@bristol.ac.uk

Ross Barnes
University of Bristol, UK
ross.barnes@bristol.ac.uk

Naomi Harris
University of Bristol, UK
naomi.harris@bristol.ac.uk

Dave Moore
University of Bristol, UK
dave.moore@bristol.ac.uk

Simon McIntosh-Smith
University of Bristol, UK
S.McIntosh-Smith@bristol.ac.uk

Emma Akinyemi
HPE, UK
emma.white@hpe.com

Janine Over
HPE, UK
janine.over@hpe.com

Martin Podstata
HPE, Czech Republic
martin.podstata@hpe.com

ABSTRACT

The Bristol Centre of Supercomputing (BriCS) has deployed its first HPE modular data centre (MDC), also known as a Performance Optimised Data Centre (POD), in March 2024. This has been a collaborative, co-design project between HPE and the University of Bristol. The MDC has enabled the rapid commencement of operations for the research community for the direct liquid cooled (DLC) Isambard-AI phase 1 (HPE Cray EX2500) and the air-cooled Isambard 3 (HPE Cray XD224), with NVIDIA Grace-Hopper and Grace-Grace superchips, respectively. The second set of MDCs have been deployed for Isambard-AI phase 2 containing 5,280 NVIDIA Grace-Hopper superchips in HPE Cray EX4000 DLC cabinets, together with the management and storage ecosystems. This manuscript outlines key features of the HPE POD MDCs for sustainability, efficiency, flexibility and observability in the era where data centre cooling and power needs are changing with growing demands for AI and HPC. We leverage the community efforts, specifically, the Energy Efficient High Performance Computing Working Group (EE HPC WG) that aims to sustainably support science through committed community action by encouraging the implementation of energy conservation measures and energy efficient design in HPC [1]. We outline notable advantages of the MDC approach for constraints and requirements that are unique for the Isambard-AI project that led to a co-design approach. We conclude by highlighting the key lessons drawn from this work.

* Co-design, deployment and operation of a modular data centre with air and direct-liquid cooled supercomputers
†Sadaf R Alam, Ross Barnes, Naomi Harris, Dave Moore, Simon McIntosh-Smith, Emma Akinyemi, Janine Over, Martin Podstata. 2025.

CCS CONCEPTS

• Power and Energy • Systems security • Artificial Intelligence • Energy consumption (Green computing)

KEYWORDS

Artificial Intelligence (AI), Supercomputing, Modular Data Centre, Sustainability, AI Sovereignty.

1 Introduction and background

During the panel session on the future of direct liquid cooling power at the International Supercomputing conference in November 2024 (SC24), the panelists highlighted the growing demand for water that could be up to 25% of the total available supply in some local municipalities [2][3]. At the same time, there are headlines regarding the push for nuclear power plants from hyperscalers [4]. In this era, research resources for AI and HPC such as Isambard-AI and Isambard 3, need to consider the responsible and sustainable use of resources alongside providing supercomputing capabilities. Isambard-AI is an AI Research Resource, or AIRR, which was announced prior to the international AI Safety summit on November 1-2, 2023 [5], to address the lack of national AI-capable supercomputing facilities for open research in the UK. Its mission includes providing sovereign AI compute capacities for UK academia, research and government entities. Figure 1 outlines the site layout presented a year ago at CUG24; this has since been successfully deployed on-site [6]. Within the same facility, as shown in Figure 1, which we call Isambard Park 1.0, Isambard-AI phase 2 has been deployed as a combination of two PODs, one Direct Liquid Cooled (DLC) for 12 HPE Cray EX4000 cabinets, which we call the POD DLC Data Hall, and an air-cooled POD DC10 for the ecosystem, including management, networking, storage and site integration.

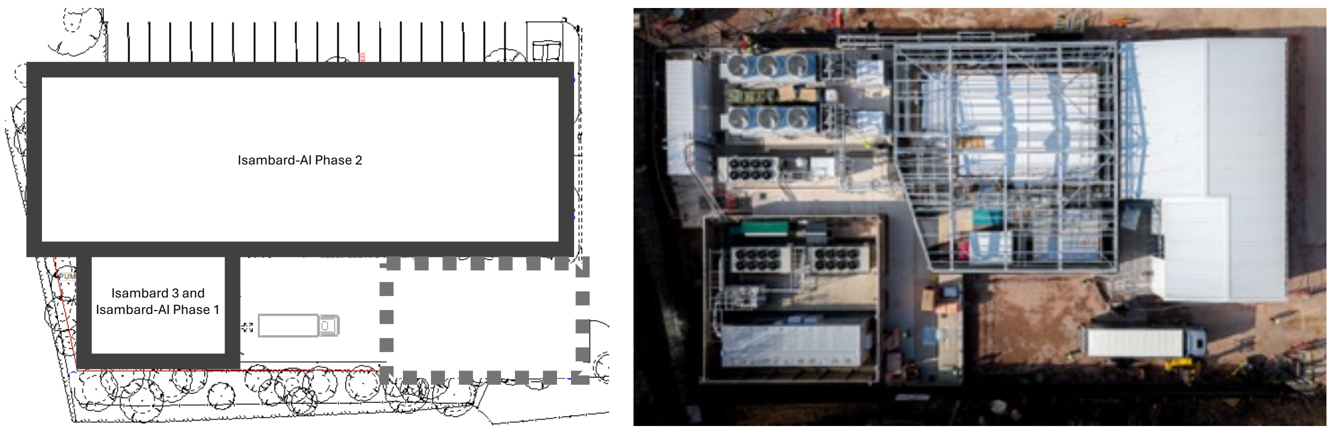

Figure 1: A schematic of Isambard Park 1.0 in the National Composite Centre (NCC) car park, showing two Self Contained Units (SCUs), one for Isambard 3 and Isambard-AI phase 1 (deployed in March 2024) and another for Isambard-AI phase 2 scale out arriving in 2025. The figure on the left has been derived from the Isambard-AI CUG24 paper. The figure on the right represents an ariel view of the status as of late March 2025. Isambard 3 and Isambard-AI phase 1 are in the bottom left corner, sharing a POD DC20, power and cooling ecosystem. The rest of the area shows Isambard-AI PODs, power and cooling distribution units and their supporting Mechanical, Electrical and Plumbing (MEP) infrastructure. Additional physical security measures have been included to fulfil compliance requirements for both projects. The switch house is the rightmost structure in the figure that brings utility power distribution for 5,280 NVIDIA Grace-Hopper GH200 superchips.

The UK's ability to meet the rising power demands of data centres is a very topical subject. It is an area where significant investment is needed to keep up with the rest of the world in the space of AI Data Centers. Thus an approach that is not only cost-effective but also prioritises heat reuse and sustainability of these Data Centres is even more critical.

Isambard 3 and Isambard-AI phase 1 have been fully operational since January 2025 and June 2024, respectively. These systems are available to UK academia, research and government through open calls for proposals [7][8]. Successful projects have been taking full advantage of these MDC-based supercomputing resources for several months at the time of writing of this manuscript. These share an air-cooled Performance Optimised Datacentre (POD) called POD DC20.

In this manuscript, we outline our co-design approach, implementation details and the operational configuration of the MDC that are common across the two deployments and where they vary to share our experiences and lessons learned for future MDC deployments. The use of containers or MDC for IT deployments is not new [9][10]. Neither is their deployment in unusual places, for instance, a car park in our case. However, recently large-scale supercomputing deployments, with direct liquid cooling (DLC), have been gaining momentum. Even in the DLC domain, there are several MDC vendors offering a range of MDC products and solutions to address the growing needs.

MDCs are adaptable and configurable. These can be placed virtually anywhere alongside power and connectivity, allowing customised computing to be closer to users, meeting local regulatory requirements, and supporting remote operations. They can be delivered within a few months to the required specifications. This modular concept can help to easily expand the capacity of the data center. The configuration options for the Isambard site are shown in Table 1. They not only highlight configurability and flexibility, but also the co-design approach taken to meet the requirements and constraints of the Bristol Centre for Supercomputing (BriCS), as an AI and HPC service provider, as well as to meet national and local regulatory and compliance needs. The details of UK, European and international building, electrical and mechanical standards are considered outside the scope of this paper. Suffice to say that it is similar to the classic, bricks-and-mortar data centre processes in order to ensure safe deployment, maintenance and operations of a facility under local weather, fire protection and in general health and safety concerns. Periodic, on-demand and incident driven safety and maintenance checks are included as part of operational support with HPE.

HPE provided the MDC components while the University of Bristol's (UoB) Campus Division designed and supported site-wide integration in collaboration with the National Composite Centre (NCC) since their car park has been repurposed for the Isambard deployments. Table 1 lists the configuration options and separation of concerns between HPE and UoB. These are specific to the Isambard projects. For instance, vendors can offer fully fabricated and self-contained solutions based on the requirements and constraints of the project. Due to the scale and complexity of Isambard-AI, commissioning and integration of the IT equipment was conducted on-site, i.e. the Isambard projects did not use a fully fabricated option.

	HPE Options	HPE provided	Site provided
POD	• For the POD IT module or All in on POD MDC • Outdoor or indoor • Standard products or customised • Scale-out options through spines and corridors	• Outdoor, self-contained PODs for Isambard 3 & AI Phase 1 and AI phase 2 • Standard products or customised • Isambard-AI phase 2 is scaled out to fit 12 EX4000 cabinets plus connectivity with the air-cooled cabinets	• Supporting pre-installation, deployment (in factory integration) and post-deployment support • Liaising with site contractors
Power	• MV and LV Switchgear boards • UPS options • Diesel generators • Power stabilisers • Transformers	• LV power modules with UPS and Batteries to enable graceful shutdown of HPC kit • Diesel generators • Dedicated free cooling units • Fire detection • Downstream LV power distribution	• Infrastructure for utility power connectivity including MV distribution transformers and upstream LV bus bars, etc. • Statutory provider upgrades
Cooling	• Chillers with free air cooling • Dry/hybrid coolers for DLC • Pump skids • Heat recovery infrastructure • On-site pipe works	• Chillers with free air cooling • Dry/hybrid coolers for liquid cooling • Pump skids • On-site pipe works • Heat recovery interlinks	• Supporting pre-installation, deployment and post-deployment support
Site	• Concrete slab • Network and cabling • Secure Cage • Fencing • Access controls • Security Systems	• Networking in between PODs • Storage facility	• Concrete slab • Network and cabling • Secure Cage • Fencing • Access controls to PODs • Water treatment

Table 1: A list of configuration options that were presented by the HPE team for MDC solutions. HPE and UoB teams co-designed a final set of options to meet the requirements and constraints of Isambard 3 and Isambard-AI projects and the overall deployment program.

To support the custom requirements of the Isambard projects and local integration needs, HPE and the UoB had several contractors, most notably Contour Advanced Systems, Oakland Construction Limited and Kendall Kingscott. This co-design project would not have been possible without their commitment and support for this extraordinary endeavour.

Fortunately, having the same team for the deployment of Isambard 3, Isambard-AI phase 1 and then Isambard-AI phase 2 was hugely advantageous: the lessons learned were quickly applied to later phases of the project since the modular building blocks of an MDC remained largely the same, even though the scales changed significantly in terms of power, cooling and space requirements.
Comparing to a classical bricks-and-mortar data centre, the Isambard MDCs have no significant differences. They have a compute floor or area, a distribution zone for power and cooling and a resourcing zone for bringing in power and cooling. In a classical setup, there are typically floors (sometimes raised), which are also feasible in the MDC approach e.g. stacking. This was not considered for Isambard due to the timeline for obtaining building permissions.

The outline of the paper is as follows: the flexibility and versatility of the MDC solution for fulfilling the unique Isambard 3 and Isambard-AI requirements are detailed in Section 2. Section 3 covers sustainability, while Section 4 provides insights into how an MDC solution is highly efficient, not only for deployment but for the usage of resources such as energy and water. Section 5 details observability via DCIM and integration to site-wide monitoring and logging solutions. Section 6 provides details on the operational configuration, and discusses extensibility and scalability options. We conclude in Section 7 with a summary and checklists for deploying an MDC solution from the lessons learned as part of the Isambard-AI and Isambard 3 projects.

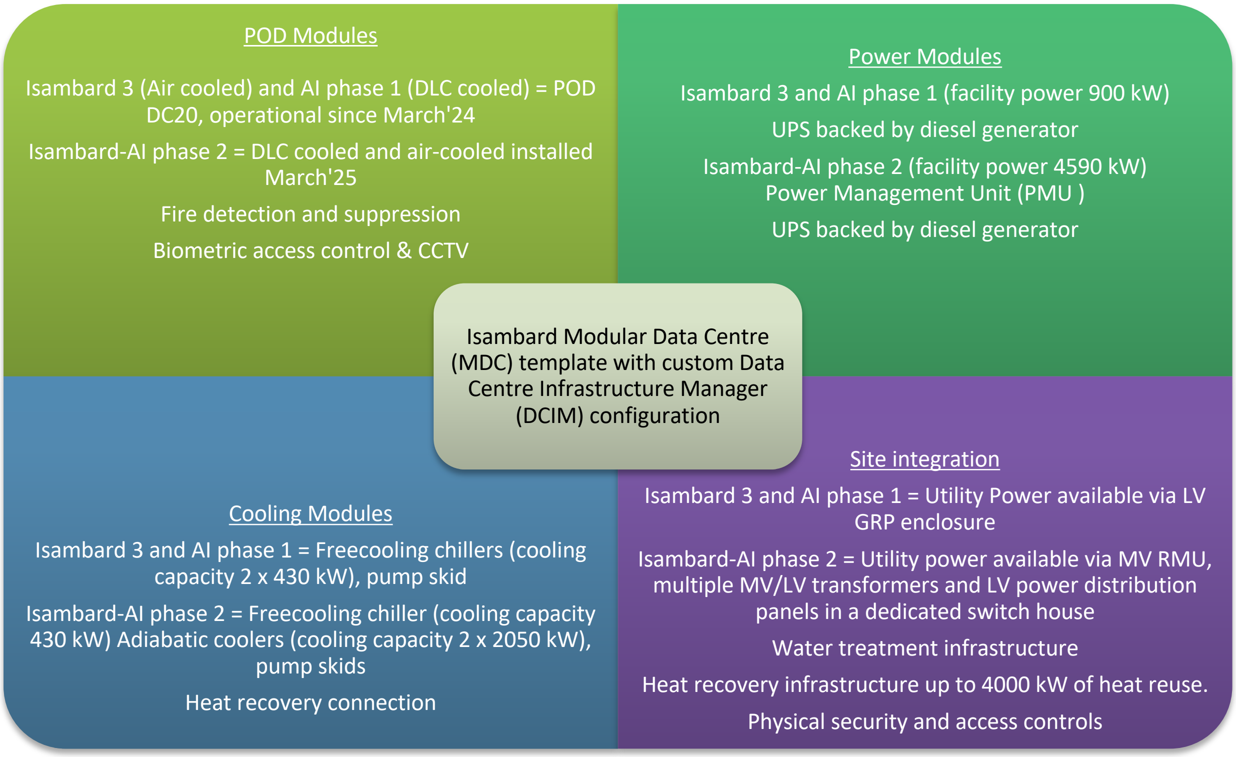

Figure 2: A template from HPE outlining the building blocks for the Isambard 3 and Isambard-AI phase 1 and phase 2 platforms. There are similarities and differences considering the significant increases in the AI compute, thereby power and cooling capacities. Isambard-AI phase 2 also required additional physical security. The 4590 kW facility power specified is related to HPE deliverables only, i.e. these do not include site needs such as the switch-house demands, water treatment power demands, etc.

2 Flexibility and Versatility

The versatility of MDC solutions has been clearly demonstrated by the unique requirements of the two different Isambard projects. From a supercomputing design perspective, these two projects are very distinct. One is largely CPU based and air-cooled, while the other is a dense, direct liquid cooled AI system based on GPUs. The portfolio of POD and MDC ecosystems allowed for a consistent deployment of both solutions, as shown in Figure 2. The MDC solutions are grouped into modules that are co-designed with the requirements and constraints of each project. The specifications of our two MDCs (Isambard-AI phase 1 plus Isambard 3, and Isambard-AI phase 2) are listed alongside power, cooling and site integration modules. The POD modules accommodate not only air-cooled and DLC cooled components, but can serve at different scales, from 168 NVIDIA GH2000 in Isambard-AI phase 1, up to 5,280 GH200 in Isambard-AI phase 2. Likewise, the power modules offer a combination of main and backup power sources, such as UPS. The scale of Isambard-AI phase 2 necessitated a Power Management Unit (PMU) to support 4MW+ operational capacities of AI GPUs. The cooling modules have been installed with necessary free air-cooling capacities and chillers (details on chillers and water usage can be found in the efficiency section).

In fact, the flexibility and versatility extend beyond the POD design to the entire site integration activities. These include co-designing requirements for physical security that the HPE team could accommodate as part of the design, delivery and deployment process. Thus, these MDC technology-based projects provide the flexibility and opportunity for significant parallel working such that the civil works, MDC and IT could proceed simultaneously.
Isambard-AI has been deployed in two phases. The phase 1 was targeted to come online within a few months of finalising the contract, in the existing air-cooled data centre called DC20. Thanks to the flexibility of the MDC, the HPE team co-designed a HPE Cray EX2500 (direct liquid cooled) option with 168 NVIDIA GH200 superchip, or 42 nodes, within the DC20 module. This includes additional management and storage capacities. As a result, the service was installed and came online within four months of signing the contract and within just a couple of months of the installation.

MDCs in general offer a great deal of external, environmental flexibility for deployment. For instance, in the case of Isambard it is in a car park with power and network connectivity. Hence, the solution contained the entire ecosystem including power modules, cooling modules, switch house, water treatment infrastructure, heat recovery, pump skids and connectivity to the on-site utilities. Custom options such as physical security can be incorporated as needed. There are examples where everything including power, cooling and physical security is integrated as part of the solution. This highlights the "modularity" aspect of an MDC approach, where individual modules are co-designed and provisioned according to the requirements and constraints of a project. At the same time however, the work schedule that was dependent on local authorities to bring in 4+ MW of power to the site didn't scale and it required additional interventions to ensure that the project didn't suffer from significant delays.

3 Sustainability

The Isambard MDC deployments are fully aligned with the community efforts such as the Energy Efficient High Performance Computing Working Group (EE HPC WG) that aims to sustainably support science through committed community action by encouraging implementation of energy efficient design in HPC. We refer to their documentation "Energy Efficiency Considerations for HPC Procurement Document: 2021" that covered systems to be deployed in about three to five years (2022–2025) in terms of features and capabilities [11]. Recently, a chapter has been added on sustainability [12], which categorises as follows :

- *Basic Requirement*: A requirement that is possible to achieve with 2021 technology.
- *Advanced Requirement*: A requirement that stretches today's technology somewhat.
- *Innovative Requirement*: A requirement that is hard to achieve today but would be nice to aim for in the near future.

The terminology for procurement documents contains guidelines for sites preparing an RFP as well as the vendors and suppliers responding to the RFP. The *Request for Information* can be used in a special document 'Request for Information' or in the 'Request for Proposal'. The *Requirements for Information* do not contain any other conditions that need to be met but the information provided can be used in the evaluation of the proposal if that has been stated in the procurement documents.

The University of Bristol has set a target for achieving NetZero by 2030, which meant the sources of GHG emissions need to be classified per requirements and standards that HPE has provided comprehensively in compliance with the Energy Efficient HPC Working Group's (EE HPC WG) Procurement guidelines on sustainability [13][14][15][16][17][18]. For instance, we received estimates and projections for the GHG emissions based on their sources. Scope 1 emissions are defined as emissions generated onsite from the activities BriCS owns or controls. Scope 2 includes indirect emissions generated from purchased energy. Scope 3 emissions are upstream and downstream emissions associated the supply chains and we needed this information from HPE. Such a level of reporting by vendors is considered as advanced or innovative according to the EE HPC WG guidelines. This is often referred to as Life Cycle Analysis (LCA), which is a method used to evaluate the environmental impact of a product or system considering its entire life cycle. Likewise, the end-of-life reuse and recycle potential for the MDC is at the advanced level, where the majority of the components can be reused or repurposed, resulting in a reduction of carbon relative to traditional bricks-and-mortar building. There are detailed plans for local recycling of IT equipment aligned with the Circular Economy concept. Product Carbon Footprint (PCF) is a methodology for assessing the total GHG emissions, or CO2-equivalent (CO2e), emitted throughout the life cycle of a product using the LCA framework. These are considered standard methods. HPE is among a growing number of companies that are working to improve the ethical responsibility of their supply chains through open communication and transparency. An ESG report covers all three areas: environment, social and governance (ESG).

The projected value for the carbon emissions due to operations (use or scope 2) of the IT equipment represents 90% of the entire footprint. The rest, including the entire POD electrical and mechanical infrastructure, as well as estimated scope 3 emissions of the POD and its mechanical and electrical components (manufacturing, assembly and transportation) and the end of life (a small fraction) of the PCF is estimated to be about 10% of the total. In summary, the Isambard-AI project procurement fulfilled advanced and innovative requirements for sustainability by fulfilling the following:

- *Advanced requirement* embodies water usage associated with production and water usage report.
- *Innovative requirement* for providing an estimated measurement of embodied greenhouse gas emissions of the processes associated with the production according to some standard such as ISO 14067 as well as usage report
- *Innovative requirement* for providing a product Life-Cycle Assessment (LCA) disclosure, including water consumption, pollution, and eco-toxicity. The disclosure should include a statement on the methodology used for the product LCA. There should be assurance that the product LCA is verifiable.
- *Innovative requirement* for providing a supply chain responsibility report specific to the procured product, based on a documented disclosure framework

Furthermore, the 5,280 DLC cooled NVIDIA Grace-Hopper GH200 superchips are utilised as a valuable source of heat through the integrated heat recovery infrastructure, enabling heat reuse for neighboring NCC facilities and in a planned future also to the residential area in the site vicinity, which has a remarkable impact on the sustainability of Isambard-AI projects. Once deployed, this heat recovery system will enhance the sustainability credentials of the Isambard-AI deployments.

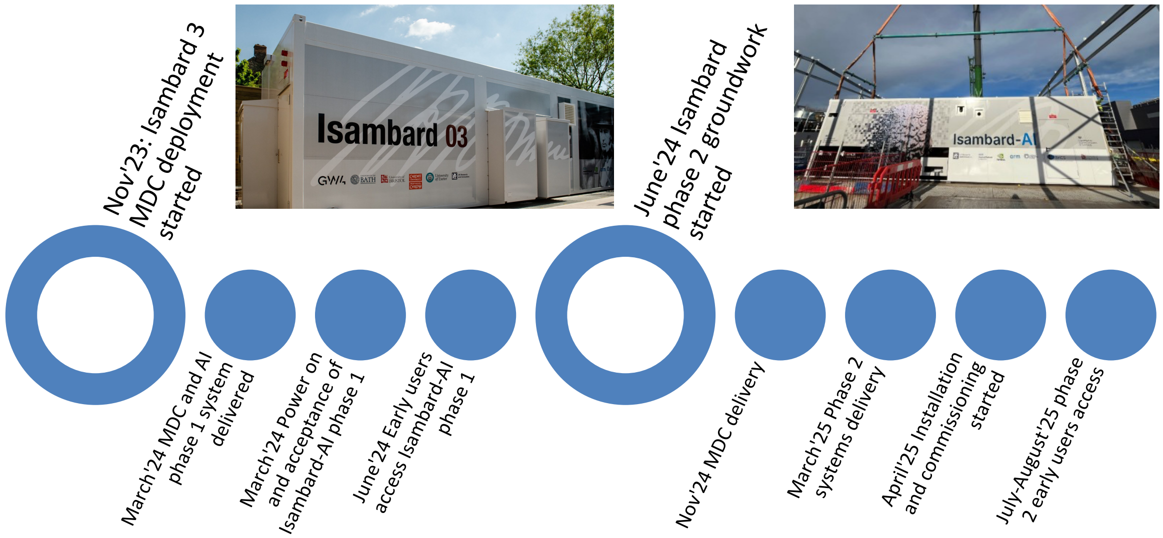

Figure 3: A timeline for Isambard-AI phase 1 and phase 2 deployment demonstrating how quickly AI compute resources became available via the UK AI Research Resource to the user communities. To achieve this aggressive schedule, the data centres exploited MDC technologies, which were constructed in parallel with the build and deployment of the supercomputing ecosystems, as well as integration of power, cooling and physical security needs at the site.

4 Efficiency

We consider the efficiency for designing and installing the data centre alongside the operational efficiency in terms of PUE and resource usage. Previous generations of Isambard HPC platforms were hosted by the UK's Met Office. This was no longer feasible for Isambard 3 and subsequently for Isambard-AI. Hence, a data centre solution was needed that was available quickly (in matter of months) and cost-efficient compared to private, commercial data centres options, which are often unsuitable for dense HPC. The MDC option fulfilled these requirements. The University of Bristol quickly explored the availability of power and space as the key requirements. Cooling and the other usual data centre requirements were part of the proposed, efficient MDC ecosystem deployment.

One of the main reasons or differentiators for selecting an MDC option was the aggressive schedule of Isambard-AI. From the signing of the contract, the phase 1 system was deployed within 4 months in the existing Isambard 3 MDC (DC20). At the same time as the operational services started for phase 1 in June 2024, the groundwork for phase 2 began. Throughout the process, the timelines have been adjusted to accommodate the evolving needs of the project, as discussed in the flexibility and versatility section. This rapid deployment timeline is shown in Figure 3. As reported in the CUG24 paper, the Isambard-AI project was awarded to deliver AI capacity to UK communities quickly. The UK Government announced the award of the Isambard-AI project on September 13th, 2023. The University then ran a rapid procurement process for a 5MW, AI-optimised solution, including an MDC solution, and the contract was awarded to HPE in mid-October. In November 2023, the groundwork started at the NCC car park. Within 4 months, a Phase 1 system of 168 identical Grace-Hopper systems was delivered as an EX2500 within the Isambard 3 MDC. Hence the first Isambard-AI users were able to run early experiments on the Phase 1 system in the May/June 2024 timeframe, with Phase 1 put into production in the summer of 2024. At the time of writing this paper, phase 2 MDC commissioning is complete and power on is in progress. This is expected to take longer than phase 1 because the scales for power, cooling and supercomputing ecosystems are considerably different this time. We expect to start early user access later this summer 2025, about a year after starting the groundworks.

To meet the energy, water and operational efficiency needs in general, the MDC solution for phase 2 was specific to the DLC technology i.e. HPE Cray supercomputers reaching up to 400 kW/rack. HPE's patented cooling systems, which combine air and liquid cooling in one system, is expected to improve energy efficiency and reduce cooling waste by typically reducing the carbon footprint by 30-40%. The MDC solution for Isambard-AI is specifically designed to achieve a low PUE and to reduce water usage over the lifetime of the project. The PUE (considering the whole year ambient temperature profile of Bristol and supercomputing usage conditions) is expected to be 1.11. In the cooler months, the PUE is expected to be even lower.

Water usage projections are based on 100% IT load. The usage is negligible when the ambient temperature is no higher than 19°C or 66.2°F. When the ambient temperature is over 19°C the cooling system's wet mode is activated. This could require up to 3.67M litres of water across the ~11% of the hours in a year that are expected to be above 19°C in Bristol. The required flow rates range from about 2.2m^3 at 19-20°C, up to 7.7m^3 at 35°C (2,200 litres up to 7,700 litres per hour once in wet mode). An Olympic-sized swimming pool is taken to have a volume of 2,500,000 litres, so Isambard-AI would use enough water to fill about 1.5 such swimming pools over a whole year. This is considerably less than many large-scale supercomputing facilities, where a year of Isambard projected water usage can be consumed within a few days or weeks. Hyperscalars are for instance were reporting billions of gallons of water usage a few years ago. In short, nearly 90% of the time throughout a year of operation, the Isambard-AI system doesn't consume any water for cooling. Over the course of a whole year, we'll use about as much water as it would take to fill 1.5 Olympic sized swimming pools, or about 20,400 standard bath tubs (180 litres), or 56 baths a day averaged over the year.

5 Observability

One of our key requirements for the MDC, being a new technology for BriCS, is to have full data centre observability as one would expect from an HPC data centre plus a holistic approach for data centre and IT telemetry in operations. The MDC solution comes with an integrated DCIM (Data Centre Infrastructure Manager) solution from HPE partner Cadence [19]. This provides a digital twin of each MDC, including facility and IT metrics within the MDC such as a searchable archive of IT assets, network port capacity, environmental conditions in hot and cold aisles, power related metrics, UPS parameters, water leakage sensors, humidity and flow values, infrastructure weight and telemetry from individual cabinets inside the MDC. HPE/Cadence DCIM provided setup and configuration in collaboration with their MDC supplier (Contour) and DCIM (Cadence). Additional configuration options are possible. The same system incorporates the IT and cooling subsystems like PDUs, temperature of hot and cold isles, etc. Operational support teams at the University of Bristol and the suppliers have access to MDC related information for on-demand and periodic maintenance. In addition, the university Campus Division could incorporate these within the existing Building Management System (BMS) frameworks that it uses across a wide range of academic and research facilities. The BriCS team has been working in parallel on a site-wide telemetry solution that incorporates DCIM and the system level metrics to facilitate operation of the services. These will help in correlating different sources of data for setting alerts, for diagnosing failures and regression, reporting on sustainability metrics, and overall ensuring high availability and quality of services. Figure 4 shows the setup that BriCS has deployed to integrate DCIM metrics for monitoring, logging and alerts for operational support of Isambard-AI and Isambard 3 services via a central data collection repository.

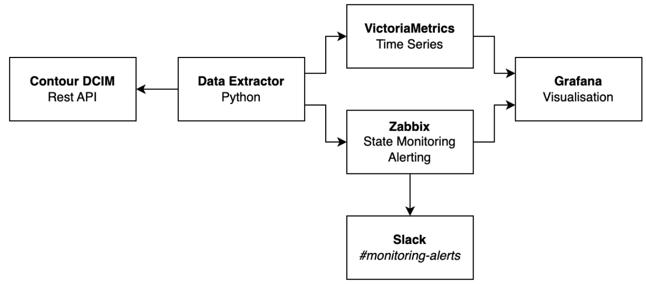

Figure 4: Workflow for capturing MDS DCIM values using the REST API that is then fed into Zabbix and VictoriaMetrics to create Grafana dashboards for the operators. Alerts are generated by Zabbix and sent to infrastructure operators via dedicated Slack channels.

The relevant DCIM states are fed into Zabbix, which is an open-source software to support monitoring of a variety of infrastructure elements by collecting relevant metrics [20]. Alerts and triggers for system operators are available via dedicated channels for the infrastructure operators who can then inspect the data sources including DCIM. The BriCS team has implemented a role-based access control (RBAC) that is compliant with zero trust security architecture [21]. Essentially, the system level information, logs and outputs are available to individuals who have appropriate roles and credentials. Access to DCIM software is governed by the RBAC model for operators from UoB and suppliers. For a time series output, VictoriaMetrics has been deployed, which is an open-source solution for gathering and storing near-real time metrics from the DCIM [22]. These outputs are available via different Grafana dashboards that have been configured by the operators like any other bricks-and-mortar supercomputing facility as shown in Figure 6.

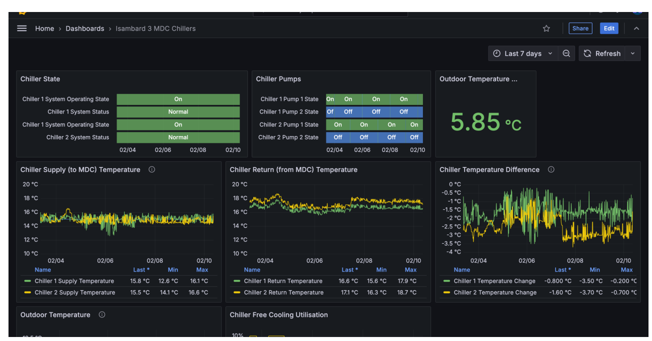

Figure 5: Grafana dashboard showing DC20 chiller metrics, pump states and outdoor temperature using data gathered from the DCIM and stored in Zabbix and VictoriaMetrics.

Anecdotally, since the MDCs have been designed to support a remote operational model, the control systems in some instances may have finer grained metrics than classic bricks-and-mortar systems that may have been deployed decades earlier. This is shown as a screen capture of Zabbix in Figure 7. Here details on chillers and pumps are available where alerts can be setup both for the Campus Division as well as the BriCS supercomputing operations team. Such capabilities are highly useful because MDC as a technology for large-scale computing system deployment is a new concept for both teams who are in the process of setting up operational processes for scheduled and unscheduled maintenance, incidents and outages of the MDCs and the supercomputing ecosystems.

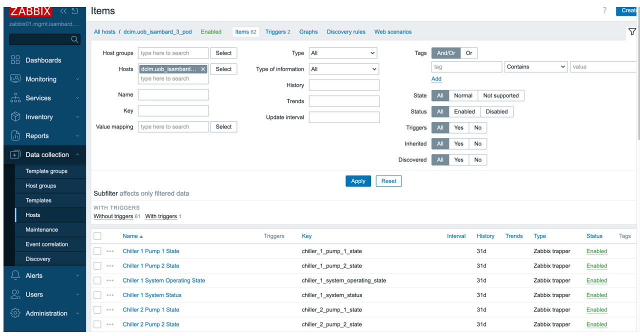

Figure 6: Zabbix dashboard demonstrating fine-grained MDC component level states that are available to Isambard-AI and Isambard 3 supercomputing service operators within a site-wide monitoring, logging and alerting ecosystems to troubleshoot incidents and outages.

6 Operational Configuration, Scalability and Extensibility

In this section we share practical details about the following:

- Water delivery to the site
- Management of water chemistry
- Cooling solution deployment
- Wet mode description
- MDCs commissioning and connectivity

In addition, we will comment on the scaling of this solution beyond Isambard Park 1.

The water is connected via a standard water connection via the NCC (adjacent building on the same campus). There is a reserve tank to ensure service continuity. The Isambard has a raw water tank, that can be filled overnight (where there is limited use), this water is then passed through the RO (Reverse osmosis) plant into a treated water tank for use within Isambard. The local mains (drinking) water utility connection of the MDC site is provided from an existing branch of the neighbouring NCC facility. The capacity of that water supply is therefore shared in between the NCC and Isambard AI POD MDC demands.

The DLC cooling system is engineered to control minimal flow temperature of cooling water (for both chilled water or CW cooling as well as DLC cooling) above the dew point all the time, so the risk of condensation on the surface of cooling components is excluded. For the CW system the chiller plant setpoint is 15°C and the DLC cooling system is utilizing 3 port valves to mix the minimal flow temperature to the CDU as 18°C or more.

The mechanical infrastructure has been provided with no redundant components, this is typical for GPU intensive HPC DC facilities. However, to minimise the probability of unplanned downtime, the secondary circulation pumps are arranged in a redundant configuration 1+1 and Air-cooling POD infrastructure will be energised by a Diesel generator (from the PMU). Additionally, the cooling infrastructure involved in the graceful shutdown of HPC infrastructure is further supported by the UPS, namely pumps, overhead cooling units in POD DC10 and Compute Room Ai Handlers (CRAHs) in the DLC datahall as well as the Electrical Power Management System (EPMS) and Building Management System (BMS) control systems.

The Water Treatment plant including raw water tank, sand and active carbon filtration, water softener, precision filter, reverse osmosis and the Treated water storage tank is provided to support DLC cooling demand of the two Adiabatic hybrid coolers. The targeted treated water quality has been determined by the Hybrid coolers vendor. The provision of treated water to the hybrid coolers is essential to secure DLC operation in case if ambient air temperature exceed 18.4 °C. For the operation of Hybrid coolers, the use of town mains (drinking) water or rainwater is generally recommended. Use of harvested rainwater can be limited by the effects of bacteriology, collection surfaces and their properties. A bacteriological analysis of the rainwater Water specification and treatment proposal should be carried out in advance. Furthermore, seasonal variations in water quality and biological load are possible. The use of recycled water (e.g. grey water) has not been recommended because of the quality, biological load, solid residues, and the unknown and varying mineral content. Water treatment has been therefore made according to the local water quality and the demanded treated water specification. Per recommendation, as in most cases, a water treatment system is installed to comply with these limits. The size and method of this system have been agreed based on the local conditions, the local water quality, the achievable concentration factor and as well the investment and operating costs.

Since the Isambard AI system has been designed and built as a "Self-Contained Unit (SCU)" including IT HW kit, Datahall space and necessary power and cooling infrastructure. The facility part of the SCU is as a Modular, Performance Optimised Data Center ("POD") solution consisting of following key POD modules:

- One IT module (POD DLC Datahall), hosting the DLC IT HW
- One IT module (POD DC10-300), dedicated space for the Mission critical air-cooled IT HW (Management nodes, Storage, Networking switches…)
- One Power module PMU400 with UPS and Diesel generator to support mission critical infrastructure
- Two Hybrid coolers supporting DLC heat load
- One Chiller to supporting air conditioning of POD MDC

- Three Pump skid modules
- One Storage module for spare parts and tools

The items above lists factory integrated and tested POD MDC modules which were craned into position and assembled onsite, this included sealing (the Data Hall, DC10, Storage Containers) from water and dust within just a couple of days. The process of craning and sealing is weather dependent, for instance, strong winds, rain or snowfall may restrict the process and influence the POD shells/sections sealing timescale and quality. Careful planning of deployment activities (including contingency planning) has been therefore essential, especially in British autumn/winter weather conditions. The two sets of MDCs for Isambard AI phase 1 and Isambard 3 and Isambard AI phase 2 have network connectivity (data centre backbone) with redundancy. The option to connect via SlingShot across the two MDCs is available should it be needed in the future (cabling is in place).

Considering the above, and since we already demonstrated expansion within Isambard Park 1, we believe such an MDC based solution is designed in a modular manner, to allow scale. The MDC facility concept is enabling scaling in the form of future expansion or partial/full decommissioning very well, especially when comparing to traditional brick and mortar approach. The POD MDC is scalable in matter of physical space, electrical power, cooling as well as physical security. The local zoning requirements, shape and dimensions of available plot of land, access roads to it, adjacent neighbours, availability of utilities, mainly suitable electrical power supply, data connectivity, water and sewage availability are the parameters determining the limits of scalability of modular DC facility. However, the infrastructure, power, water and data must be scaled to meet the demand. In particular, obtaining large amounts of power to site on time, long lead items which are very much dependent on supply chain and can be/are influenced even by global geopolitical situations or turbulent resources cost evolvement, large number of suppliers to co-ordinate parallel work.

We believe the issues are similar at scale. This depends on whether the users / owners are different then this results in independent cooling systems. Space can be another issue, especially like Isambard machine room, power and cooling will be all at one level. Planning where the systems may change throughout time. Security needs for different users. We learned that the local zoning requirements, shape and dimensions of available plot of land, access roads to it, adjacent neighbours, availability of utilities, mainly suitable electrical power supply, data connectivity, water and savage availability are the parameters determining the limits of scalability of modular DC facility.

Our longer-term considerations for future MDC deployment and ongoing operations. We will continue to evaluate and reflect on Isambard Park 1 operations, which started for Isambard-AI phase 1 and Isambard 3 already, and in progress for Isambard Isambard-AI phase 2. The MDC approach is enabling phasing or even "grow as you need" if you will at AI compute capacity as well as at the Datacenter facility space/power/cooling capacities. The reduced CAPEX should be therefore expected compared to typical brick and mortar DC facility, designed and built at day one for expected day ultimate IT HW demand. Notwithstanding IT hardware and its space/power and cooling demands, which are constantly evolving as IT HW infrastructure evolves. The MDC approach therefore better follows the evolution cycle of IT HW. MDC OPEX is to be reduced due to more efficient cooling (closed coupled cooling and DLC) as well as shorter, thus more efficient energy distribution paths. The MDC modules standardisation and factory integration and testing pays also important role into the reduction of on-site efforts, time to deploy as well as operation failure probabilities.

Any service provider considering MDC as an option should note that the key is to be able to define the business demand and correlate it with the IT hardware demand, for instance, AI and HPC. It is supposed to be simpler for initial (day one) stage, while it may be more challenging for day ultimate or any mid to long term plans. The MDC concept enables to follow the specific IT hardware demand at time, when that IT hardware is being deployed. The brick-and-mortar approach typically doesn't have such level of flexibility, capability, and co-design options in a time- and cost efficient manner as we demonstrated in the paper.

7 Summary and Lessons Learned

Several teams across the University of Bristol, BriCS, HPE and our subcontractors have undertaken an innovative, co-design approach for meeting ambitious timelines and requirements for the Isambard-AI and Isambard 3 projects using MDC technologies. Thanks to the flexibility, adaptability and efficiency of this approach, we have been able to manage the objectives of the projects. Operationally, the HPE solution for Isambard-AI using the MDC and supercomputing systems enable BriCS to accommodate sustainability and operational resilience objectives.

In short, the planning and operational needs are no different from a traditional HPC data centre for periodic maintenance and upkeep. We would like to acknowledge that the Isambard projects have greatly benefitted from the EE HPC WG community efforts. Furthermore, we leveraged the rapid rate of change in the data centre requirements and subsequently in the MDC design space that are emerging primarily due to the growth in AI across the board for industry, academia and research.

The key challenges and shortcomings that we have identified include:

- Site work uncertainties for power, cooling and in general utilities connectivity since the campus or estate or facility divisions are new to this concept i.e. a modular data centre (MDC) facility. Some permissions from local authorities took longer than anticipated and required additional resources, time and cost, to ensure the overall program timelines are not adversely impacted. As this

becomes mainstream, we anticipate these can be managed effectively.
- Just-in-time findings as both the MDC and IT solutions were developed and evolved in an agile, co-designed manner. These include requirements for insurance and liability as an example, which necessitated redesigning different aspects of both physical and cybersecurity that were not typically considered when an MDC is deployed within a secure facility (Isambard Park 1 has been built in a car park). These changes were managed with additional resources and expertise.
- Like any agile, co-designed project another issue we experienced was managing scope for a timely delivery. While we could address changing requirements, it often led to managing priorities and expectations as well as the project timelines, within a budget envelope, that includes such contingencies. An overall program governance was and tightly coupled and highly motivated teams at Bristol and HPE were invaluable for the delivery.

The above may have been exasperated for the Isambard projects due rather ambitious, expedited timelines.

We would like to reflect on our experiences by attempting to provide checklists for future deployments based on our lessons learned. These checklists are divided into two parts. The first part covers a set of constraints and requirements that are largely covered in the EE HPC WG procurement and sustainability documents with a few exceptions. For Isambard-AI, these are:

- Articulating priorities clearly from the outset, which are the ambitious schedule, accommodating a targeted number of AI GPUs, and sustainability criteria for meeting the university of Bristol NetZero targets
- Defining a rapid but agile deployment timeline for a research supercomputing system for AI starting from scratch (no data centre facility was available)
- Identifying a physical location (outdoors or indoors) with sufficient power capacities
- Following EE HPC WG procurement guideline documents
- Outlining pragmatic acceptance criteria for the MDC and IT, separately and as a whole
- Identifying project and program management requirements for the above
- Preparing, in parallel, a plan of a handover and operational plan post-acceptance.

The second checklist covers execution of an ambitious project plan such as Isambard-AI that relies on external factors such as the site integration and hence the co-design approach that has been outlined in the manuscript:

- Incorporating co-designing flexibility and agility in the original solution that could be gradually adopted for the initial design choices
- Setting aside planning and budget contingencies to manage uncertainties and unknowns, for instance, complex supply chain issues for MDC and IT
- Managing different site-specific, emerging needs and timelines in the program via frequent checkpoints
- Establishing solid relationships with all contractors and suppliers
- Introducing a leadership structure between university of Bristol, BriCS and HPE for overseeing a pragmatic project management across MDC and IT deployments
- Communicating to various stakeholders including funders, executives at participating institutions, regularly and periodically
- Overcommunicating across different teams throughout deployment using formal meetings and asynchronous methods, such as Slack
- Collaborating to define handover and operational plans for an MDC, which could be a new concept for some AI and HPC service providers.

This collaborative experience and expertise between HPE and the University of Bristol can be an exemplar for other sites who may be interested in one or more features that we discussed in relation to the deployment of MDCs, namely flexibility, sustainability, efficiency and observability.

ACKNOWLEDGMENTS

Isambard-AI is operated by the University of Bristol and is funded by the UK Government's Department for Science, Innovation and Technology (DSIT) via UK Research and Innovation; and the Science and Technology Facilities Council [ST/AIRR/I-A-I/1023]. The project has only been possible thanks to a very talented and extremely hardworking team of people across the University of Bristol, HPE, Contour, NVIDIA, the National Composites Centre, and our building contractors, Oakland.

REFERENCES

[1] Energy Efficiency Considerations for HPC Procurement. https://sites.google.com/lbl.gov/procurement/home
[2] https://www.theregister.com/2025/01/04/how_datacenters_use_water/
[3] P. Li, J. Yang, M. A. Islam, S. Ren. "Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models." 2025. https://arxiv.org/abs/2304.03271
[4] https://diginomica.com/why-hyperscalers-are-going-nuclear-and-why-thats-probably-really-good-idea
[5] "Bristol set to host UK's most powerful supercomputer to turbocharge AI innovation", https://www.gov.uk/government/news/bristol-set-to-host-uks-most-powerful-supercomputer-to-turbocharge-ai-innovation
[6] S. McIntosch-Smith, S. R. Alam, C. W. Woods. Isambard-AI: a leadership class supercomputer optimised specifically for Artificial Intelligence. In proceedings of CUG 2024. https://arxiv.org/abs/2410.11199.
[7] Access to high performance computing facilities – autumn 2024. https://www.ukri.org/opportunity/access-to-high-performance-computing-facilities-autumn-2024/
[8] AIRR Expression of Interest to use UK large scale compute. https://engagementhub.ukri.org/ukri-infrastructure/airr-eoi/
[9] https://datacentremagazine.com/critical-environments/the-drive-towards-global-modular-data-centre-growth-in-2024
[10] https://www.gminsights.com/industry-analysis/modular-data-center-market
[11] Energy Efficiency Considerations for HPC Procurement Document: 2021 https://drive.google.com/file/d/1aB7uv47anaHUcHzw140tJLUSe9xVkQfQ/view?usp=sharing
[12] Sustainability Chapter added to the Energy Efficiency Considerations for HPC Procurement Document: 2021 (version 2.0) https://drive.google.com/file/d/1nJYzxVmIivvG6R1fThophbA_N_fQM97a/view
[13] ISO 14001 https://www.iso.org/iso-14001-environmental-management.html

[14] Sustainability and the U.S. EPA: https://nap.nationalacademies.org/catalog/13152/sustainability-and-the-us-epa
[15] EU Eco-Management and Audit Scheme (EMAS) https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX%3A32023D2463
[16] ISO 14067 Greenhouse gases – Carbon footprint of products – Requirements and guidelines for quantification, https://www.iso.org/standard/71206.html
[17] PAS 2050 Specification for the assessment of the life cycle greenhouse gas emissions of goods and services, https://knowledge.bsigroup.com/products/specification-for-the-assessment-of-the-life-cycle-greenhouse-gas-emissions-of-goods-and-services/standard
[18] WRI GHG Protocol Product Life Cycle Accounting and Reporting Standard, https://ghgprotocol.org/product-standard
[19] https://www.gartner.com/en/information-technology/glossary/data-center-infrastructure-management-dcim
[20] Zabbix:- https://www.zabbix.com/
[21] S. R. Alam, et.al. "Federated Single Sign-On and Zero Trust Co-design for AI and HPC Digital Research Infrastructures". https://arxiv.org/abs/2410.18411
[22] https://github.com/VictoriaMetrics/VictoriaMetrics