



OPEN
Compute Project



WHITE PAPER: INTEGRATING QUANTUM PROCESSING UNITS INTO DATA CENTER INFRASTRUCTURE

Author (s) :

Jonathan Burnett, NQCC, Project co-lead
Boniface Yogendran, Project co-lead (OCP)
Allison Schwartz, D-Wave
Burns Healy, Dell Technologies
Andre Saraiva, Diraq
Solyman Ashrafi, Fujitsu
Too Vira, Inflection
Coleman Collins, IonQ
Eric Mansfield, IQM
Günther Leschhorn, IQM
Kyle Thompson, Maybell
James Fletcher, ORCA Computing Ltd
Philip Clarke, Oxford Quantum Circuits
Owen Arnold, Oxford Quantum Circuits
Elisabeth Eude, Pasqal
Aleksander Wennersteen, Pasqal
Bruno Gouvêa Taketani, QBlox
Florian Preis, Quantum Brilliance
Marvin Hirschel, Zero Point Cryogenics
Nilesh Shah, ZPT Corp
Cliff Grossner (OCP)
Steve Helvie (OCP)

Date: April 14, 2026

Executive Summary

“The Analytical Engine does not occupy common ground...it holds a position wholly its own”

- Ada Lovelace (English Mathematician and Writer, 1843).

This quote was a response to how the programmable computer isn't just an incremental improvement but rather an entirely new category and something that has never existed before. The same could be said about quantum. Quantum computing, communication, sensing and simulation aren't just incremental improvements to the classical ways of doing things but represents a significant paradigm shift.

The datacenters needed to deploy quantum require a similar paradigm shift. Current deployments of quantum computers are typically done in bespoke environments such as lab environments, small private datacenters, etc. and not at scale in larger third party datacenters (otherwise known as colocation providers).

In today's environment, the strategic planning cycle of building a datacenter is moving from the traditional time horizon of 2 - 3 years to 5 - 10 years (JLL Global Data Center Outlook, 2026). This puts added pressure on the ability to ensure datacenter fungibility to allow for eventual quantum deployments at scale. Without a proper understanding of the physical requirements of quantum systems, there is a risk of bi-directional stranding between a facility's lack of readiness that limits the overall growth of the quantum sector.

The Open Compute Project (OCP) is actively addressing these integration hurdles through its Future Technologies Initiative (FTI), specifically within the "Data Center Integration of Quantum Information Infrastructure" workstream. This white paper serves as the beginning of a forward-looking Community-led framework, bringing together a global community of academia, vendors, industry researchers, and startups to define the physical and logical requirements for housing quantum systems in production datacenters.

While datacenter operators and quantum providers already realise their importance in their respective sectors, this white paper begins to bridge the deployment gap ensuring next generation facilities are physically equipped to host diverse quantum modalities at scale and quantum providers design their next generation of equipment to match what the operators can practically provide. With these pillars in place, the OCP Quantum Ready Ecosystem can “hold a position wholly its own” as a global standard for mutual readiness.

Table of Contents

1. Purpose	5
2. Motivation	5
3. Introduction to Quantum Computing	7
3.1. Quantum Computing Full Stack	8
3.2. Logical QPU	10
3.2.1. Real-time Host and Interconnect	11
3.2.2. Quantum System Controller	11
3.2.3. Classical Architecture	12
3.2.3.1. AI Acceleration	14
3.2.3.2. Control Plane	15
3.2.3.3. Syndrome Extraction	16
3.2.3.4. The Role of Storage	17
3.2.4. Physical QPU	17
3.2.4.1. Modalities	18
3.2.5. Environmental Control	19
3.3. Deployment	19
4. Current Challenges in QPU deployment	20
4.1. Facility Selection	21
4.1.1. Site Selection	21
4.1.2. Service Boundaries and Demarcations	23
4.2. Datacenter deployment challenges for different quantum modalities	24
4.3. Variations Within Modalities	27
4.4. Risk to the Facilities	27
5. Quantum System Deployment	28
5.1. Deployment of modalities: Neutral Atoms and Trapped-ions	29
5.1.1. Facility and Zone	29
5.1.2. Neutral Atoms	31
5.1.3. Trapped Ions	32
5.2. Deployment of modalities: Photonics	34
5.3. Deployment of modalities: Annealing, Silicon Spin, and Superconducting qubit architectures	36
5.3.1. Facility and Zone	38
5.3.2. Superconducting Qubits	40
5.3.3. Annealing	42

5.3.4. Silicon Spin Based Qubits	44
5.4. Deployment of Modalities: Diamond based architectures	46
5.4.1. Diamond spin qubits	46
5.4.2. Diamond spin qubits with cryogenic electronics	48
6. Future Efforts	50
6.1. Non-Quantum OCP Projects Enabling Scalable and Reliable QPU Deployment	50
6.2. Quantum Specific Efforts at the OCP	51
7. Conclusion	54
8. References	54
9. Creative Commons License	55
10. About Open Compute Foundation	56

1. Purpose

Technologies based on exploiting quantum properties, known as quantum technology, are approaching maturity where deployment, interoperability, and integration with modern digital infrastructure is a serious consideration. Quantum technologies span computing, sensing, timing, and communications. To address these integration challenges and technical translation, the OCP Future Technologies Initiative has commenced a working group on “Data Center Integration of Quantum Information Infrastructure”. The first output of this working group is this white paper which focuses on quantum computing, subsequent outputs will expand towards other quantum technologies.

The purpose of this white paper is to identify the common architectural, operational, and integration considerations for deploying quantum processing units (QPUs) in datacenter infrastructures. It is intended to inform and guide datacenter architects, hardware designers, cloud providers, and IT leaders on the specifications, strategies, and best practices required to integrate QPUs as first-class components of modern datacenters as quantum computing transitions from research into early practical deployment.

Aligned with the Open Compute Project (OCP) Foundation’s tenets of efficiency, impact, openness, scalability, and sustainability, this white paper provides actionable guidance for designing hybrid quantum classical environments that support quantum computing. Quantum computing architectures are expected to be capable of supporting workloads such as optimisation, machine learning, simulation, and cryptography. This paper addresses how QPUs can be integrated with existing compute, storage, and networking resources to ensure scalable, efficient, and manageable deployments.

Quantum computers are based on different physics-implementations which are commonly referred to as modalities. These modalities are each expected to map to different application spaces. However, each modality has different characteristics, leading to significant variability in system design, performance indicators, and integration requirements. As a result, this white paper concentrates on architectures that have started to explore co-location with classical compute within datacenter type environments. Consequently, these modalities are reaching maturity around architectural patterns, deployment models, and interface standards and are now sufficiently stable to inform practical planning. By focusing on near-term, realistically deployable technologies, the document aims to equip practitioners with actionable insights that can guide infrastructure decisions, reduce integration risk, and accelerate the adoption of quantum accelerators within open, interoperable datacenter environments.

2. Motivation

Quantum computing platforms provide computational advantages for well-defined problem classes that are fundamentally inefficient to solve on classical architectures. Similar to other computing technologies, quantum computers are improving year on year. Present systems are categorised as Noisy Intermediate-Scale Quantum (NISQ) devices [1]. A NISQ device

means a functioning quantum computer that is exploring potential application spaces. Vendor's roadmaps then track their anticipated progress toward fault-tolerant operation, at which point the application space expands. Throughout this process, application exploration is expected to empirically find more use cases similar to the empirical nature of AI use-case discovery. In practice, any application of a quantum computer is expected to have a classical computing aspect. Therefore, integration of the quantum processor itself to the surrounding classical infrastructure such as datacenters, networking and high performance computing (HPC) becomes critically important. As such, AI datacenters are the natural integration point for quantum computing, as they already provide the power, cooling, networking, and orchestration environments required to support accelerator-driven workloads.

Quantum computers have been available over the cloud for almost a decade. Here, a number of vendors explore and demonstrate how to operationalise quantum computers towards a computing service. Over the past few years, quantum computers have also been deployed and integrated into a number of datacenter type environments such as supercomputing datacenters, HPC datacenters, and co-location datacenters. Importantly this represents a maturing of the quantum computing landscape from laboratory environments towards environments capable of realising production computational workflow. Significant investments are beginning to flow into the quantum computing ecosystem, signaling a growing appetite to accelerate the state of quantum computing [2].




Figure 1: Recent investment trends in quantum computing. Capital deployment follows a trajectory similar to historical AI investment growth, offset by an estimated ten-year lag in maturity and scale [2].

Many countries consider quantum technology and specifically quantum computing to be a critical technology. Consequently, significant government investment spans the academic research, industrial research, and other funding streams. These act to de-risk private funding streams which have been scaled in recent years. Collectively, these have enabled a rich ecosystem of quantum computing vendors and supply chain around the core building blocks of quantum computing processors to facilitate the transition from lab oriented research work to commercial grade datacenter ready solutions.

Quantum computing system design tends to stem from intellectual property (IP) around the QPU itself, leading to QPU vendors effectively functioning as system integrators. As systems integrators, they own the design choices and risk around component choice, facility design, control stacks, and classical compute integration. This ecosystem momentum build up, coupled with datacenter deployments and the emerging quantum datacenter stack, underscores a clear industry requirement for OCP Community driven good-practice towards standardisation in areas the specialised facility envelopes that handle unique requirements. By integrating QPUs directly into the datacenter fabric, these deployments establish the foundation for a scalable, interoperable supply chain that ensures quantum technology becomes a modular component of the global compute ecosystem.

OCP historically standardised racks, pods, power, cooling, and service demarcation. Given the emerging quantum computing supply chain stack, several opportunities are now ripe for standardisation within OCP:

- Hybrid quantum classical pod layer: form factors, interfaces, mechanical and thermal envelopes
- Classical integration layer: telemetry formats, storage profiles, control planes, and networking models
- Facility layer: zoning, vibration classes, cooling architectures, and environmental envelopes.

Without OCP, each QPU vendor builds bespoke datacenters, hyperscalers create vertically locked stacks, and modularity collapses. With OCP, quantum becomes infrastructure rather than lab equipment, enabling interoperability, supply chain scaling, and true integration into the global compute ecosystem.

3. Introduction to Quantum Computing

The motivation for quantum computing lies in its potential to accelerate complex, computationally intensive tasks that strain even the most advanced HPC and AI infrastructures. This power comes from its ability to exploit quantum phenomena such as superposition and entanglement, enabling a more efficient exploration of complex solution spaces and the execution of algorithms that are intractable on classical systems. By leveraging these principles,

quantum computing is expected to integrate with traditional datacenters as a complementary computational platform, addressing key problem areas such as:

- **Machine learning and AI** – Reducing large datasets to improve decision-making, predictions, and automation, generating high quality synthetic data for model training, and implementing quantum subroutines in hybrid quantum-classical machine learning models to improve model performance and reduce training cost.
- **Simulation of quantum systems and processes** – Modeling molecular processes, material science, and life sciences to drive innovation in pharmaceuticals, energy, and advanced materials.
- **Mathematical optimisation** – Accessing larger-scale or enabling near real-time optimisation for applications such as financial modeling, network planning, energy grid management, logistics planning, and supply chain management.
- **Factorization for security** – Factoring large numbers exponentially faster than classical computers. At fault-tolerant scale, this capability could break widely used public-key encryption schemes, highlighting the urgent need for post-quantum cryptography to secure digital communications.

Despite these transformative possibilities, most current quantum systems remain highly resource-intensive and inefficient, constrained by limited system size (number of qubits) and limited system performance (qubit error rates). These practical limitations prevent today's hardware from realising the full theoretical potential of quantum computation. Instead, the industry is operating within the NISQ phase—a stepping stone toward scalable, fault-tolerant architectures capable of executing reliable, large-scale quantum workloads.

In this document, the term “qubit” is used as a technology-agnostic abstraction for quantum information carriers. This convention encompasses both discrete-variable systems and continuous-variable systems, including qumodes, unless explicitly stated otherwise. Physical qubits represent the hardware-level quantum elements, whereas logical qubits are encoded, fault-tolerant computational units constructed from ensembles of physical qubits via quantum error-correction (QEC) mechanisms.

3.1. Quantum Computing Full Stack

The quantum computing stack comprises tightly interdependent layers that bridge physics, engineering, and computation—spanning from qubit control to hybrid algorithm execution. At the foundation, the quantum hardware layer implements the basic building blocks of a quantum computer: qubits, qumodes and quantum logic gates. In practice, how these are implemented varies across the various physical modalities.

The quantum hardware layer tends to require an engineered environment consisting of aspects such as vacuum, cryogenics, or low-vibration systems. The quantum hardware layer

sits within this environment and is interfaced with via a signal delivery layer. Above the signal delivery layer, control and readout systems translate digital instructions, often into precise analog signals that manipulate and measure quantum information states. This layer integrates microwave, optical, or electronic control hardware, and plays a critical role in timing synchronisation, noise suppression, and calibration [3].

The error correction and middleware layer provides logical qubit abstraction, mitigating the effects of noise through real-time feedback, error decoding, and qubit redundancy schemes such as surface codes or bosonic encodings. These systems rely on low-latency classical co-processors tightly coupled with the quantum hardware [3].




Figure 2: Conceptual full-stack architecture of a quantum computer, illustrating the tightly coupled layers from the physical QPU hardware to the classical control systems and external application interfaces [3].

On top of this foundation, the software stack encompasses quantum compilers, runtime systems, and programming frameworks that map high-level algorithms onto the underlying quantum architecture. Middleware orchestrates hybrid workflows—coordinating classical pre- and post-processing with quantum execution—while resource managers optimise qubit allocation, routing, and circuit scheduling [3].

Finally, the application and integration layer bridges quantum systems with datacenter infrastructure. This includes APIs, orchestration tools, and interconnect protocols that seek to allow QPUs to operate as integrated accelerators within HPC and AI environments. Achieving this requires open, interoperable standards that align innovation across all layers, enabling scalable deployment of quantum technologies within future heterogeneous compute ecosystems [3].

3.2. Logical QPU

In the previous section, we describe quantum computers within the NISQ regime as possessing high error-rates. A key challenge to enable lower error rates is the realisation of logical qubits. Here, the quantum information is encoded across numerous physical qubits together with ancilla qubits, which facilitate error detection and subsequent correction. Conceptually, these error-correction schemes follow classical error-correction within electronics design. The reference architecture below for a “logical QPU” stack provides a pattern for understanding the facility requirements for QPU integration. This layer goes from physical QPU to middleware that controls the quantum information states. Under the logical QPU, there are 4 main components. They are: real-time host, real-time interconnect, quantum system controller and physical QPU. Not all quantum computing hardware expected to deploy in the next 18-24 months will require all of these components.




Figure 3: HPC and logical QPU schematics. This figure is a reference architecture, inspired by NVIDIA’s NVQLink, that is intended to communicate the components of a logical QPU stack. Not all components will be required for every modality or deployment [4].

3.2.1. Real-time Host and Interconnect

The real-time host acts as the classical “brains” that keeps a QPU running correctly. It receives three main kinds of information: error signals coming from qubits (called syndromes), health and calibration data that show how the quantum information states are drifting or changing over time, and the logical instructions that describe what quantum program the system is supposed to run. Using these inputs, the real-time host continuously decides how to correct qubit errors, updates the control settings so the hardware stays well-calibrated, and organises the sequence of logical operations that the quantum processor should perform next. It then sends these decisions—such as error-correction commands, updated pulse parameters, and execution schedules as digital signals to the quantum system controller through an interconnect, which converts them into precise physical signals for the quantum hardware that implements the quantum information states. This allows the QPU to operate reliably even in the presence of noise.

3.2.2. Quantum System Controller

Quantum computing relies on the precise manipulation of quantum information using microwaves, lasers, optical fields, or other physical control mechanisms specific to each hardware modality. It receives high-level instructions from the classical computing system as

well as feedback from the QPU. Using these inputs, it generates precisely timed control pulses that manipulate quantum information. It may also incorporate data from environmental sensors—such as temperature, magnetic fields, and vibrations—to fine-tune control signals and optimise hardware performance.




Figure 4: Example of a quantum controller system by Qblox. Note this form factor is modular with a chassis and plug-in cards. The chassis is rack mountable. Here the connectors on the front correspond to SMA-ports for microwave signals.

3.2.3. Classical Architecture

QPU deployment must explicitly account for classical infrastructure integration, as quantum workloads are inherently hybrid. Classical compute and storage systems are required for control, orchestration, error handling, data preprocessing, and post-processing, making compute and storage scalability critical design considerations. As quantum capacity increases, the supporting classical infrastructure must scale proportionally to avoid becoming a performance, reliability, or operational bottleneck. The growth of the classical infrastructure is determined by the following:

- **QEC Overhead (FTQC Scaling):** For leading fault-tolerant quantum computing (FTQC) architectures—particularly surface-code families—the number of physical qubits per logical qubit scales approximately with the square of the code distance, while the required code distance increases as lower logical error rates are targeted. This produces superlinear system growth, where increases in logical capability drive disproportionately larger increases in physical qubits, control complexity, and classical processing demand. As a result, doubling “useful” quantum capacity can require more than doubling classical infrastructure capacity.

— **Growth characteristic:** Non-linear

- **Planning implication:** Capacity must be designed in discrete tiers (e.g., pilot → pre-FT → FT), with step-function expansion in:
 - Hardware-accelerated decoder pools (e.g., FPGA/GPU/accelerator resources)
 - Control-plane networking capacity and latency margins
 - Write-intensive storage systems for syndrome data and calibration history
- **Decode Compute Scaling:** Decoding pipelines typically scale near-linearly with the number of stabiliser measurements per cycle under nominal operating conditions. However, scaling behavior changes when systems encounter tighter latency budgets, more complex noise models, multi-QPU synchronisation, or higher concurrency. At this point, systems exhibit non-linear utilisation cliffs, where latency violations trigger retries, expanded logging, checkpointing, and orchestration overhead—creating compounding load amplification effects.
 - **Growth characteristic:** Non-linear
- **Logging, Observability, and Storage Growth:** Even when raw measurement volumes are modest, operational requirements drive accumulation of syndrome histories and decoder states, calibration and diagnostic traces, simulation and modeling artifacts, experiment provenance and audit records (particularly in regulated or sovereign environments). As systems scale, organisations should expand retention horizons to improve yield optimisation, debugging efficiency, uptime, and model quality. Consequently, observability and storage requirements grow with system scale, independent of raw data rates.
 - **Growth characteristic:** Linear (with scale-driven retention expansion)




Figure 5: A standard quantum job workflow within a HPC environment provided by Dell, The workflow largely relies on the classical system, which handles three out of the four actions.

The classical architecture dedicated to supporting the quantum deployment falls into three broad categories. Here, we do not include any standalone HPC environment which the QPU is augmenting.

- Hardware required to run the quantum processor
- Dedicated AI infrastructure for adding value via customised ML models
- Low-latency FPGA / GPU compute for doing error correction

3.2.3.1. AI Acceleration

ML models integrated into the quantum workflow are applied across both pre-processing and post-processing stages:

- Pre-processing (Classical Control and Optimisation Layer):
 - Resource allocation: workload-to-QPU mapping, qubit assignment, and scheduling
 - Transpilation: circuit optimisation, hardware-aware compilation, and gate mapping
 - Queue optimisation: latency management, throughput maximisation, and job prioritisation
- Post-processing (Classical Inference and Mitigation Layer):
 - Shot post-selection: statistical filtering and noise-aware result conditioning
 - Error mitigation: classical correction models for noise suppression and fidelity enhancement

Custom ML models specialised for the QPU are deployed to perform these tasks and they can be performed asynchronously relative to the quantum job. Since, unlike QEC, these tasks are not required during QPU execution, the latency requirements can be loosened.

The classical infrastructure to support these workloads most closely aligns with conventional HPC architectures and operational models. As a primarily AI-oriented compute layer, it will conform to standard rack-level specifications in terms of power density, networking, cooling, and mechanical integration, aligning with existing datacenter infrastructure profiles

Quantum - classical system design is still a topic for exploration and research, here is an example deployment for AI Acceleration setup by Dell:

- Based on GB200 NVL4, implemented in the IR7000 rack series
- Support for 144 Blackwell GPUs, this NVL72 system is the current highest end solution for dedicated AI compute
- Liquid cooled
- Powered up to 264kW per rack
- Up to 36 nodes in a 50OU IR7000

3.2.3.2. Control Plane

The nature and extent of the control plane requirements are strongly dependent on modality and architecture. The scaling characteristics of these requirements directly dictate host-node design, capacity planning, and integration constraints. These control-plane requirements can be viewed at two levels: internal control of the quantum system itself, and external integration into the surrounding compute infrastructure.

Concerning the integration of quantum resources into the broader control plane of the heterogeneous compute infrastructure, one needs to enable resource management capabilities including discovery, allocation, scheduling, and observability that are compatible with established practices. This determines how quantum systems participate in hybrid infrastructures alongside purely-classical resources.

In most modalities, the path to fault tolerance is to drive down individual qubit-errors beyond logical thresholds, using as few physical qubits per logical qubit as possible. The classical requirements scale with the number of syndromes and therefore with this ratio of physical qubit to logical qubits.

3.2.3.3. Syndrome Extraction

Syndrome extraction is the process of detecting error information encoded in a quantum state without disturbing the logical state of the system. This information is used by the classical control system to determine and apply the appropriate corrective action, if any is required. This error correction decision needs to be made at the pulse-level instruction in real-time for each qubit.




Figure 6: Overhead shot of a Dell XE9780, which can be deployed for large scale QEC.

Low latency is essential for maintaining effective error correction. Classical accelerators, including GPUs, are being designed to receive syndrome data directly and compute corrective actions within the quantum system's coherence window. These accelerators are expected to be directly accessible to the pulse-processing units and to incorporate high-performance PCIe switching to minimise communication latency and maximise throughput.

Example deployment for quantum error correction (assuming part of a stack with FPGAs, supporting tens of thousands of physical qubits):

- 4U, 2xCPU w/up to 64 cores per processor
- 32 RDIMM slots
- 4TB memory maximum
- 8x NVIDIA H200 connected with 4th generation NVLink (900GB/s)
- ~11-12kW per server

3.2.3.4. The Role of Storage

As mentioned earlier, the number of logical qubits needs to scale significantly for practical application. This means physical qubits need to scale to the millions. This growth exerts demand on storage capacity. It is estimated that future fault-tolerant quantum computers could generate approximately 100 terabytes per second of error correction data, equivalent to Netflix's entire global streaming bandwidth [5].

For example, a single logical qubit encoded across 1,000 physical qubits, sampled at 1,000 cycles per second with a modest 4 bytes per sample, produces approximately 4 MB per second of metadata: equivalent to roughly 345 GB per day per logical qubit.

This figure accounts only for raw QEC syndrome data and excludes additional sources such as calibration histories, drift models, waveform captures, and AI-assisted prediction or control logs. As a result, storage growth scales with qubit-hours of operation, rather than with the number of user jobs submitted. In operational terms, quantum systems resemble high-rate flight data recorders: even when no application output is large, the continuous need to observe, correct, and verify system state drives sustained, write-intensive storage demand on the classical side of the stack.

3.2.4. Physical QPU

A physical QPU is the hardware subsystem where quantum information is prepared, manipulated, and measured. The way quantum information is prepared is dictated first by the

computational paradigm—for example, whether the system implements gate-based operations, measurement-based logic, or analog/Hamiltonian evolution. Each paradigm, in turn, shapes the choice of modality, the physical mechanism used to realise quantum information. That modality then determines the QPU's environmental envelope and control requirements: cryogenic temperatures, ultra-high vacuum, optical or microwave delivery paths, electromagnetic shielding, timing precision, and vibration sensitivity. In a datacenter setting, the physical QPU therefore behaves as a specialised co-processor whose infrastructure footprint is downstream of its paradigm and modality, informing everything from facility layout to control-system integration.

The key characteristics of a physical QPU include:

- **Qubit fidelity:** Achieving long coherence times and low error rates.
- **Scalability:** The ability to build systems with a large number of interconnected, high-quality qubits.
- **Controllability:** The precision and speed with which the qubits can be manipulated (read and written).
- **Operating environment:** The specialised physical conditions required for the system to function (e.g., extreme cold, high vacuum, powerful lasers).

3.2.4.1. Modalities

Each modality represents a distinct approach to constructing and manipulating quantum information states, shaped by different engineering principles and material platforms. The diversity stems from the varied approaches researchers are taking to meet the demanding requirements for quantum computation [6].

1. **Superconducting circuits** require cryogenic environments and complex microwave control electronics. These systems are typically operated at approximately 10–20 millikelvin in dilution refrigerators and control and readout are performed using room-temperature microwave electronics connected through extensive coaxial cabling to the cryogenic stages. Scaling is influenced by wiring density, thermal loading, and electromagnetic shielding requirements.
2. **Photonic systems** use quantum states of light such as single photons either to implement boson sampling, directly as qubits, or as elements of complex entangled states to implement measurement-based quantum computing. Photonic systems require high quality photon sources, low-loss waveguides (including fiber optics as well as photonic integrated circuits) and electro-optical switches. Photonics systems also require fast, accurate photon detectors which can require cryogenic environments.
3. **Trapped ions** systems use highly stabilised lasers to cool and suspend charged atoms in a vacuum chamber, requiring high precision laser systems and radiofrequency control.

Cryogenics are not required for operation but are often used to achieve optimal vacuum pressure and reduce noise caused by power dissipation.

4. **Neutral Atoms (Rydberg Atoms)** use uncharged atoms, controlled by arrays of highly stabilised lasers, to store and manipulate quantum information.
5. **Semiconductor quantum dot systems** embed qubits in silicon or other semiconductor materials, require cryogenics, typically at millikelvin temperatures in dilution refrigerators. Control is via microwave control electronics in additional magnetic fields.
6. **Topological qubits** are typically implemented via superconducting and semiconducting materials, therefore requiring cryogenics and microwave control electronics with magnetic fields. As this modality has no evidence of near-term deployment in datacenters, it is not covered further in this paper.
7. **Nitrogen vacancy centers** in diamond use defects in the diamond lattice to encode qubits, often allowing operation closer to room temperature. Qubits are typically manipulated using precision microwave and optical control.

This white paper covers the deployment considerations for all major quantum computing modalities except electrons-on-helium and topological qubits; as to best of our knowledge, there is currently no evidence of their deployment outside controlled laboratory environments. However, in both cases the physics is similar to the superconducting qubit and semiconductor quantum dot systems, therefore those are likely to provide a good starting indication on the deployment considerations.

3.2.5. Environmental Control

Quantum computers and their components are generally analog systems, consequently their behavior has high dependency on environmental aspects such as temperature, humidity, vibration, pressure, and electromagnetic fields. It is therefore common for a systems integrator to perform a site-inspection to determine the environmental background at the facility level. The facility then interfaces with systems such as cryogenics or vibration stabilisation. These systems create and maintain a stable and low-noise operational environment that preserves the fragile quantum states essential for computation. This includes regulating temperature, suppressing vibrations, minimising electromagnetic interference, and ensuring clean, low-noise power delivery.

Some of these components can introduce health and safety implications such as laser-system safety or background magnetic fields. The systems integrator should be able to assist with risk-assessment and safety protocol in these circumstances.

3.3. Deployment

Physical deployment refers to how a quantum processing unit is integrated into a datacenter environment, including its spatial placement, connection to classical computing resources, and access to environmental and control infrastructure. There are three primary deployment strategies for QPUs [7]:

- **Standalone Deployment:** The QPU is housed independently and accessed remotely, with no direct integration into classical HPC infrastructure. This model simplifies initial setup but introduces higher latency for hybrid quantum-classical workloads.
- **Co-Located Deployment:** The QPU is positioned near classical computing resources, enabling local connectivity and lower-latency interactions. This approach balances integration complexity with performance, allowing more efficient hybrid workflows.
- **On-Node Integration:** The QPU is embedded directly within a classical computing node, similar to a GPU or other accelerator. This offers the tightest coupling and lowest latency, supporting real-time quantum-classical interactions, but requires the most sophisticated infrastructure and system design.

The vendor defines the deployment strategy, which affects facility requirements, rack layout, cooling and power provisioning, network architecture, and workflow orchestration, making it a critical consideration for datacenter-grade QPU integration.

4. Current Challenges in QPU deployment

Integrating QPUs into operational environments is more complex than adding other accelerators to existing compute infrastructures. The lack of standardisation is the primary contributing factor, affecting hardware interfaces, software stacks, and environmental controls.

While early deployments have taken place within HPC environments, these efforts represent a narrow slice of what future quantum classical integration will demand. In HPC systems, the “tenant” (the systems-integrator or workload owner) and the “landlord” (the infrastructure operator) are often the same entity, allowing requirements and operational practices to evolve organically around the QPU. In contrast, general purpose datacenters and AI factories, where tenants and landlords are separate, have not yet seen QPU deployments. In these environments, defining clear service boundaries is essential, as it establishes the responsibilities and interfaces between the infrastructure operator and the end user, enabling reliable, scalable, and manageable quantum classical operation.

In order to simplify QPU deployment in the datacenter, we first need to define the service boundary. This clarifies what the facility is expected to provide and be responsible for, as well as what the user can expect when deploying and operating workloads on the QPU. HPC deployment cannot be used as a model for general purpose datacenters, due to the different

relationships between tenants and operators. However, these early deployments serve as a valuable reference for identifying the challenges that must be addressed in order to establish clear and effective service boundaries. This section outlines the challenges that vendors encountered during HPC deployments, providing insights that can inform the design and operation of quantum-classical systems in multi-tenant environments.

4.1. Facility Selection

Facility selection for QPU deployment starts with geographical assessment, as the external environment fundamentally determines the stability and controllability achievable within the facility. Factors such as regional seismic activity, vibration levels, electromagnetic noise, and climate variability establish the baseline conditions that the quantum system must operate within. Sources of persistent ground vibration—such as nearby rail lines, subways, or heavy truck routes—can further compound these challenges, as the low-frequency mechanical waves they generate propagate through soil and building foundations, subtly disturbing sensitive equipment. A geographically stable, low-noise, and climatically consistent site that is isolated from such transport-induced vibrations minimises environmental disturbances before any internal isolation or shielding measures are implemented [8].

While QPU deployment considers many of the same criteria as classical or accelerator-based systems—such as access to reliable power, cooling capacity, network connectivity, and physical accessibility—it adds multiple additional layers of environmental constraints to consider. Many quantum systems are far more sensitive to local disturbances, meaning that facility siting extends beyond infrastructure optimisation to include detailed environmental characterisation and mitigation.

Each quantum modality reacts differently to geographical properties: some are highly susceptible to magnetic field gradients, others to ground-borne vibrations or ambient temperature instability. As a result, the suitability of a location depends not only on general infrastructure readiness but also on how its physical characteristics interact with the specific quantum technology.

Geographical selection is typically led or co-led by QPU providers, who possess the most detailed understanding of their systems' environmental sensitivities, shielding and calibration tolerances. They perform comprehensive site surveys—measuring vibration spectra, magnetic fields, acoustic profiles, and thermal stability—to validate compatibility before deployment. This collaboration ensures that the chosen site not only meets traditional datacenter requirements but also aligns with the quantum system's operational envelope, forming the foundation for stable, high-fidelity performance.

4.1.1. Site Selection

Site selection involves identifying the optimal location within a facility to deploy the QPU, along with its associated environment control and quantum hardware control systems. This step

is critical to ensure that the quantum system operates within a stable, controlled environment that maximises performance and reliability.

Site selection is typically guided by QPU vendors and environmental OEMs such as manufacturers of cryogenic, vacuum or laser systems. QPU vendors provide detailed knowledge of the system's shielding requirements and environmental sensitivities, while environmental control vendors specify operational requirements such as cryogenic support, vibration isolation, and temperature regulation. The collaboration between these stakeholders ensures that the chosen site meets the combined needs of the hardware and the supporting infrastructure.

The site selection process can be organised into the following categories [8]:

- **Access/delivery site:** The location must allow safe delivery and installation of the QPU and supporting equipment. Considerations include loading docks, elevators, hallways, door widths, and temporary staging areas, ensuring that any specialised hardware such as heavy cryostats (Figure 7) and delicate components can be moved without risk of damage.
- **Isolation from mechanical/electrical noise:** Depending on the modality, sites may need to be physically separated from vibration sources such as nearby traffic, subway lines, or building mechanical equipment. Electrical noise from high-power equipment or large transformers should also be avoided for some modalities to maintain qubit coherence and gate fidelity.
- **Physical room and structure:** The selected area must meet structural requirements, including sufficient floor load capacity.
- **Electromagnetic and lighting concerns:** For some modalities, the site will need to minimise exposure to electromagnetic interference (EMI) and stray magnetic fields. This includes maintaining distance from radio transmitters, base stations, and fluorescent lighting that could impact system performance. Room temperature magnetic shielding around the cryostat sometimes can be recommended.
- **Thermal and humidity stability:** Temperature and humidity fluctuations must be controlled within the site to prevent degradation of quantum computer performance. Sites with stable airflow, minimal temperature swings, and consistent humidity levels are preferred.
- **Power redundancy:** The site must have access to reliable, redundant power infrastructure. Many quantum systems require extended time to safely restart following a power outage due to cryogenic or calibration processes. Backup power, uninterruptible power supplies (UPS), and failover circuits ensure that the QPU remains operational. Note, however equipment for quantum computers follow form-factors common for laboratory environments, consequently they routinely lack features such as dual

redundant power inputs.

By carefully evaluating these categories, operators can define a deployment envelope that optimises QPU performance, ensures compatibility with environmental control systems, and reduces the risk of operational disruptions.

Some parts of the quantum computer consist of components that are rack-mountable, allowing them to be deployed with minimal modifications to existing datacenter infrastructure. Other systems can fit within typical datacenter spaces but require adjustments, such as mounting directly to concrete-slab, reinforced floors, vibration isolation, or additional cooling and power provisioning. In most cases, the power consumption of quantum computers is low such that components are ambiantly cooled rather than following form-factors that support air-containment. The main exception here is in the cryogenic equipment which is typically water cooled. Finally, when comparing within the same modality, depending on vendor IP, some vendors offer systems that integrate more easily, while others require more extensive adaptation.




Figure 7: Oxford Quantum Circuits' quantum cryogenic cooling system deployed within a controlled datacenter facility environment, illustrating integrated dilution refrigerators, structural support frames, and associated classical control and infrastructure required for operational superconducting qubit-based quantum computing systems.

4.1.2. Service Boundaries and Demarcations

In a standard datacenter, the facility provides and manages access control, cooling, networking, power, and physical space. The area between the floor tile and the top of the rack is typically controlled by the user, while any infrastructure or space outside of that boundary remains under the facility's control. This physical delineation establishes the service boundary: everything within the rack, including servers, accelerators, and cabling up to the rack's top, is considered the tenant's responsibility, whereas the facility retains responsibility for the raised floor, power distribution units, cooling systems, and building infrastructure. By clearly defining these physical boundaries, both the operator and the tenant understand their respective responsibilities for deployment, operation, maintenance, and troubleshooting, which is critical for predictable and reliable management of both classical and quantum computing resources.

In addition, some quantum systems may require ongoing operational support that is not common in standard datacenter environments, such as the periodic supply and handling of liquid cryogenics such as liquid nitrogen or filter replacement for some manufacturer deployments. Presently these service requirements either sit with the vendor, or in the case of supercomputer-centers may have been trained to the local facility team. Regardless of who is performing the service action, these do create an ongoing operational requirement for access to the system, for some systems on a more than monthly basis. These service boundaries and demarcations are not always well defined, causing challenges in QPU deployment and creating uncertainty around operational responsibilities, maintenance workflows, and long term system sustainability.

4.2. Datacenter deployment challenges for different quantum modalities

Deployment of QPUs is complicated by the distinct and non-standardised requirements of each modality. Differences in operating temperature and pressure, sensitivity to vibration and magnetic fields, environmental power draw, startup and calibration times, and the use and recovery of consumables such as liquid nitrogen all vary across systems and lack uniform standards. Table 1 summarises these varying requirements, highlighting how they directly impact the deployment and challenges involved with standardisation.

Table 1: Summary of requirements for each quantum modality for datacenter deployment, information compiled based on the vendor participated as listed in the example vendors. Here, sensitivity to magnetic fields, humidity, vibration, or temperature convey the level of care that needs to be taken in terms of extrinsic activity, for example a low score here implies a 3rd-party such as datacenter smart-hands could walk adjacent, power on high thermal load servers in proximity to the modality and expect that operational performance is not affected.

Modality	Example Vendors	Cryogenics	Magnetic field sensitivity	Vibration sensitivity	Humidity sensitivity	Temperature sensitivity	Environmental power draw	Aux requirements
Neutral Atoms	Infleqtion Pasqal	Occasional (cryogenic vacuum) Cryocooler	High	Low to High	High (if free-space optics) Low to med otherwise	Low to medium (free space optics)	3.5kW to 10kW	
NV Centres (cryogenic)	Fujitsu	Always Cryocooler & Dilution fridge	Low	Medium	TBC	TBC	TBC	
NV Centres (room temperature)	Quantum Brilliance	NA	Low	Medium	Low	Low	0.3 to 0.6kW	
Photonics	ORCA	Cryocooler	Low	Low	Low	Low	3kW	
Superconducting (gate based)	Oxford Quantum Circuits IQM	Always Cryocooler & Dilution fridge	Low	Low	Low	Medium (teflon dispersion)	15 to 30kW	Liquid nitrogen transfer every 15 days

Modality	Example Vendors	Cryogenics	Magnetic field sensitivity	Vibration sensitivity	Humidity sensitivity	Temperature sensitivity	Environmental power draw	Aux requirements
Superconducting (annealers)	D-Wave	Always Cryocooler & Dilution fridge	Medium	High	Low	Low	15 to 30kW	Liquid nitrogen transfer every 15 days
Spin-based (semiconductor)	Diraq	Always Cryocooler & Dilution fridge	High	Medium	Low	Medium (teflon dispersion)	15 to 30kW	Liquid nitrogen transfer every 15 days
Trapped Ion	IonQ	Occasional (cryogenic vacuum) Cryocooler	Low	Medium	Medium (free space optics)	Medium (free space optics)	15 to 30kW	

4.3. Variations Within Modalities

Quantum modalities can support multiple operating regimes with materially different infrastructure and deployment requirements. As a result, a single modality may correspond to more than one distinct deployment class.

As Table 1 shows, cryogenic cryocoolers are a common component across quantum computing modalities. Cryocoolers consist of a closed-circuit system that contains a cryogen such as helium. In operation, cryocoolers are capable of run-times exceeding years with no need to top up the cryogen. They typically require water-cooling and have a power draw ranging from 3 to 11 kW, this power difference scales with the required cooling power. The cryocooler is placed within a vacuum chamber, forming an environment that is at room temperature on the exterior, and at cryogenic temperature under vacuum inside.

In quantum computing systems that use cryogenics, this does not describe the entire system. There are other components such as classical computing, networking, microwave-, laser-, opto-electrical- systems that operate at room temperature. The relative scale and complexity of these components as-well-as the cryogenics vary based on the IP of the quantum computing vendor. Consequently, changes across these systems are generally not trivial to integrate, test and deploy.

4.4. Risk to the Facilities

Quantum computing systems introduce non-standard environmental influences that extend beyond the system boundary and interact directly with facility infrastructure. These influences can include magnetic fields, electromagnetic interference (EMI), RF (radio frequency) and microwave emissions, optical radiation, vibration transmission, acoustic noise, thermal gradients, and vacuum- or pressure-related structural loads.

Because these effects propagate into the surrounding environment, risk and safety must be treated as facility-level engineering concerns, not system-level add-ons. Critical risks include interference with adjacent equipment, disruption of building systems, personnel exposure hazards, and interoperability impacts across shared infrastructure. Effective mitigation requires integrated facility design, environmental zoning, shielding strategies, controlled-access areas, grounding and bonding architectures, EMI containment, vibration isolation, and continuous environmental monitoring.

Early and continuous consultation with the quantum system vendor is critical to risk management. Vendors define modality-specific field emissions, environmental tolerances, shielding requirements, safety envelopes, and operational constraints that cannot be inferred from conventional datacenter standards. Facility planning, safety engineering, and compliance frameworks must therefore be co-developed with vendors to ensure that environmental impacts are properly characterised, risks are mitigated at the infrastructure level, and safety controls are correctly implemented throughout the deployment lifecycle.

5. Quantum System Deployment

The quantum stack described in Figure 3 showed the control layer (real time host), the classical compute layer and the environment which houses the QPU. At the facility level, these 3 parts of the system drive the requirements over size, weight, power, and other services. As shown in Table 1, the use of environment aspects such as cryogenics, vacuum systems and vibration stabilisation systems are wide-spread. This section is structured around these paradigms, with deployment requirements and best-practice guidance defined relative to each category.

Typical datacenters are expected to have 3 phase (ϕ) electrical power for dual-redundancy (e.g. phase A, B, and C; each row provides two phases - $\phi A/\phi B$, $\phi B/\phi C$, and $\phi A/\phi C$ in repetition) in case one phase is in need of service. Additionally, this is to support hot-swap power supply functionality, which is also typical in enterprise-grade networking and computing equipment. High service level uptimes are enabled via this dual-redundancy that can each have layers of backups.

Another typical datacenter power infrastructure includes generator(s) and/or dedicated UPS room(s). This is implemented in case of power outage for short and long-term. The operation time for backup power can be back-calculated from each modality's power consumption and backup capacity/capability. Power stability is critical to system performance: unstable or interrupted power can force system recalibration, disrupt environmental control, and degrade operational stability. In some cases, power interruptions may require full system recovery procedures, with certain platforms requiring 48 hours or longer before returning to operational readiness. However, much of the quantum computing supply chain still produces equipment which is suitable for laboratory environments rather than datacenters. Consequently, features such as the dual-redundant power feeds at component level are missing. Due to this, the quantum computers of today tend to not be able to benefit from the full extent of resilient infrastructure within datacenters.

Table 1 highlights that cryogenics are commonly used across all quantum computing modalities. For laboratory systems, the cryogenic cooling comes in the form of a cryocooler, which is responsible for the majority of the power consumption stated in Table 1. These cryocoolers operate around the expansion and compression of a cryogen, usually helium within a closed-cycle system. A typical cryocooler consists of three components; a compressor, hoses, and a cold-head. Of these, the compressor is the component which requires the power and water-cooling to dissipate the heat. Later in this section a table will outline the power and water requirements. When describing cryogenics in the context of quantum computing, it is conventional to describe temperature in kelvin (K) rather than degrees Fahrenheit or degrees Celsius. It can be useful to contextualise cryogenic temperatures by considering the boiling temperature of liquid helium which is 4.15K (-269C or -452F) and the boiling temperature of liquid nitrogen which is 77K (-196C or -321F). A typical cryocooler can obtain temperatures between 2K and 40K.

5.1. Deployment of modalities: Neutral Atoms and Trapped-ions

Modalities such as trapped-ion systems, neutral-atom platforms, and certain diamond spin implementations rely on lasers to reduce motional and thermal energy and to enable qubit initialisation, manipulation, and readout. This section examines the specific deployment considerations for quantum modalities that depend on laser-based subsystems, including requirements for optical infrastructure, environmental stability, and ongoing calibration of laser and control subsystems.

5.1.1. Facility and Zone

Optical components and laser systems can possess high sensitivity to temperature, humidity and particle contamination. For this reason, they are generally deployed with considerable environmental isolation. However, at the facility level, there can be scope for specific additions and protocol in order to enhance performance. At system installation, use of a temporary “clean tent” can be employed to facilitate a low-particle environment. Longer term, the use of sticky mats to reduce dust that is “walked in” can become routine. These systems can be installed and operate in less-clean environments, but experience lower overall performance, higher component failure rates, and increased overall maintenance burden as a result. It is therefore recommended that host facilities maintain as clean an environment as reasonable by implementing controls such as restrictions on food and materials, use of sticky mats, regular High-Efficiency Particulate Air (HEPA) vacuuming and surface cleaning, avoidance of shedding materials and cardboard, appropriate personal protective equipment when working near optics, and adherence to verification of compliance (VOC)-controlled procedures. Periodic particle counting and routine isopropyl alcohol (IPA) wipe-downs can be used to verify and maintain environmental compliance.

Table 2: Facility and zone requirement for Neutral Atom and Trapped Ion based quantum system deployment

Modality	Delivery physical clearance	Staging area	Structural floor support for frame	Electrical details (Peak)	Temperature details	Humidity details	Vibration details	Noise
Trapped-Ion	Hallway width: 1.5m Doorway width: 1.2m to 2m Must be able to move 1.5m x 3.6m object	1.5m around the exterior edge of the system footprint (footprint is appx 2.5m x 4.5m) Height: 3m Setup area: minimum 40m ²	500kg/m ²	27 kW (inclusive of cryo infrastructure)	20°C to 24°C $\Delta T < 2^{\circ}\text{C}$ within a 48h period	30% to 50% RH; range around setpoint of less than $\pm 4\%$ over 24h	12.5 $\mu\text{m/s}$ in the frequency range 1 to 200 Hz	Acoustic: not to exceed 65 dbA Electric Field: FCC Class B
Neutral Atom	Weight of the heaviest component: 750kg (max)	0.45m to 1.5m around the an optical table which can be 1.5m x 3.6m Height: 3m Setup area: 40 m ² to 70 m ²		10 kW	20°C to 21°C 14°C to 18°C water temperature (Pasqal)	30% to 40% RH (Inflection) 40% to 50% RH (Pasqal)	50 $\mu\text{m/s}$ (Inflection)	Acoustic: not to exceed 75 dbA

5.1.2. Neutral Atoms

This section is contributed by Infleqtion and Pasqal. Neutral atom qubits are a laser-cooled, matter-based quantum modality in which information is encoded in the electronic or hyperfine states of individual atoms trapped in optical lattices or arrays of optical tweezers. These qubits can be initialised, manipulated, and measured using precisely controlled laser fields and microwave pulses, enabling high-fidelity single- and multi-qubit operations.

The QPU consists of ultra-high vacuum chambers containing arrays of individually trapped atoms. Laser cooling and optical trapping prepare the atoms in well-defined quantum states and maintain their motional stability. Qubit readout is typically performed via state-dependent fluorescence, while quantum gates are implemented using laser-driven interactions or Rydberg excitations. This approach leverages the natural uniformity and long coherence times of neutral atoms, allowing for high-fidelity control and scalability.

Deployment of neutral-atom systems requires precision optical setups, stable laser sources, vibration-isolated optical tables, and environmental control rather than cryogenic refrigeration. While scaling to very large arrays is limited by optical complexity and laser power, neutral-atom qubits are well-suited for modular quantum computing, analog quantum simulation, and hybrid quantum-classical applications.



Figure 8: Neutral-atom quantum computing systems. Left: Infleqtion system deployed at a research center facility. Right: Pasqal QPU installed at Jülich Supercomputing Center in Germany

- Facility and zone:** The system has sensitivity to background vibrations which makes deployment near to infrastructure such as railroad tracks a challenge. These sensitivities continue into the electromagnetic environment where proximity to equipment like switching equipment or heating, ventilation, and air conditioning (HVAC) can provide challenges. Humidity and temperature stability are critical. In practice a pre-installation site assessment will assess these factors.

- **Environmental control:** System can operate at room temperature but thermal stability is critical as lasers can drift when undergoing large temperature changes. Noise and vibration isolation are critical environmental factors to consider when the system is in operation.
- **Qubit control:** Neutral atom systems use laser-based control for qubit manipulation.
- **Pre-deployment:** This phase takes about 1 to 9 days. The primary effort is a comprehensive environmental site assessment to characterise EMI, ambient and DC magnetic fields, mechanical vibrations, temperature and humidity stability. Measurements are conducted continuously to capture daily and operational variability and to identify transient disturbances. This phase shall include power infrastructure and ensuring that all receptacles work and are properly grounded.
- **Deployment:** This phase spans 6 to 8 weeks and must be carefully scheduled to align with environmental stability, facility readiness, and operational constraints. Activities include the physical assembly of associated control electronics, installation and verification of hardwired electrical and plumbing connections, and coordination with facility interfaces. Following mechanical and electrical integration, the system undergoes initial validation of qubit control paths, and readout integrity. If the vendor makes use of a cryostat, then the installation includes the deployment of this cryostat along with an initial cooldown which will also validate thermal performance.

5.1.3. Trapped Ions

This section is contributed by IonQ. Trapped-ion qubits are a matter-based quantum modality in which information is encoded in the internal electronic states of individual atomic ions confined using electromagnetic fields. These qubits are initialised, cooled, and measured via laser light, and manipulated via either precisely tuned laser pulses or “electronic qubit control” (EQC) microwave antennae built into the ion trap chip, enabling high-fidelity, highly-parallel single- and multi-qubit operations with very high connectivity. This section describes the architecture and deployment requirements for EQC trapped-ion systems, the first of which is expected to begin field deployment and production operation in early 2027. Where they differ greatly, it also notes the requirements for laser-based-gate architectures such as the IonQ Tempo platform.

The QPU consists of a microfabricated, multi-layer silicon ion trap that includes integrated AC and DC electrodes that suspend the ions above the trap and drive qubit transport and gate operations, and laser light delivery for state preparation and measurement (SPAM) via integrated photonics. This architecture exploits the intrinsic uniformity of atomic ions and their exceptionally long inherent coherence times, supporting precise and repeatable gate operations. The trap package is mounted to a cryocooler’s 4 K coldhead in an ultra-high-vacuum chamber. Laser cooling prepares the ions to an acceptably low motional state, and qubit readout is performed via state-dependent fluorescence, while quantum logic

operations are implemented through a combination of common AC magnetic fields and DC fields for local tuning of interactions. [9] and [10] provide additional information on the EQC architecture.

Deployment of trapped ion systems requires high-stability laser systems, ultra-high vacuum infrastructure, low-noise electronics, magnetic field control and isolation, and vibration-isolated optical platforms. Many trapped ion platforms, including IonQ's, utilise cryocoolers in the 4 K range to improve vacuum quality and reduce noise caused by trap power dissipation. This is not an inherent requirement of ions as a qubit modality, but allows for improved performance and reduced environmental sensitivity.

- **Environmental control:** The system's control and laser delivery operate at room temperature. Good HVAC with high thermal and humidity stability and the ability to absorb a heating load of 5 to 8 kW are required — the system has some inbuilt environmental management, but not enough to overcome extremes of set point or rapid changes. The system can tolerate modest noise and vibration, but performs best in environments where these are minimised as much as is reasonable. These details would be assessed within the initial site survey.
- **Qubit control:** Qubits are controlled via electromagnetic fields delivered via chip-embedded electrodes driven by rack-mount control electronics. These racks are considered part of the system and integrated into its footprint and environmental requirements.
- **Pre-Deployment:** This phase takes approximately one week and primarily consists of site readiness evaluation, including physical space assessment, measurements of temperature, humidity, acoustic and electromagnetic noise, vibration, VOC and particulate counts, taken several times over the period.
- **Deployment:** The system is delivered to the site in a series of subassemblies, which are then assembled and commissioned onsite. Install takes approximately 8 to 12 weeks, including verification and validation activities.
- **Maintenance:** The system undergoes periodic scheduled maintenance, primarily to ensure continued performance of the optical and laser subassemblies. This maintenance must be performed by the vendor or by technicians trained and certified by the vendor.
- **Calibration:** System parameters are calibrated over different frequencies and timescales to actively correct for drifts in alignment, observed power, environmental conditions, and similar. This is a mix of frequent “servo” and characterisation “spot check” activities (short to very short duration, multiple times a day) and less-frequent “true” recalibration, which happens much less frequently (order daily to weekly), but take longer to complete. All calibration activities that happen during nominal operation are automated.

5.2. Deployment of modalities: Photonics

This section is contributed by ORCA computing. Photonic quantum processors based on telecom-band fiber optics and time-bin encoding provide a highly practical architecture for deploying within datacenter environments. These systems use single photons generated by laser-driven quantum light sources, which are routed and processed entirely within standard optical fiber infrastructure. Fiber-based optical components and interferometric structures are inherently stable, low-loss, and environmentally robust, supporting high uptime, predictable operation, and straightforward field deployment.

These systems are based on quantum information categorised as qumodes, a generalisation of qubits that are not restricted to binary states (0 or 1), but can occupy higher-order states (2, 3, ... N), with practical limits defined by optical loss within the system. Qumode control in fiber-based photonic platforms is implemented through the generation, routing, and detection of time-bin-encoded single photons. Qumodes are produced using telecom-laser-driven nonlinear optical processes that generate high-quality photon pairs at standard telecom wavelengths (e.g., 1550 nm). Time-bin encoding provides intrinsic robustness against phase noise, environmental fluctuations, and mechanical disturbances, making this architecture well suited for deployment outside laboratory environments.

The modularity of fiber-based photonic architectures enables compact, appliance-style QPUs that integrate within datacenter racks alongside HPC and AI systems. Quantum information is generated using telecom-wavelength photon sources, manipulated through fiber-integrated optical components, and measured using photon-number-resolving detectors supported by closed-cycle, rack-mounted cryogenic subsystems. This architecture leverages the long-standing maturity, reliability, and standardisation of classical telecom infrastructure, positioning photonic QPUs as among the most datacenter-compatible quantum computing platforms currently available.



Figure 9: ORCA system deployed within a datacenter at the UK National Quantum Computing Center.

- **Quantum control:** Quantum logic is implemented by manipulating photons as they travel through fiber-integrated optical switches, phase shifters, and delay lines. These components allow the system to build programmable interferometric circuits without requiring free-space optics or active alignment. Real-time control electronics—typically FPGA-based—coordinate the timing, switching, and synchronisation of optical paths and collect detection events for classical processing.
- **Deployment:** Deployment takes about 2 days. Pre-deployment survey is not required as the system is not sensitive to the environmental factors. The system consists of two standard 19-inch racks treated as an integrated appliance with a combined weight of approximately 800kg. Systems are shipped as fully integrated rack-level units. Deployment requires only rack placement, connection to facility power, and network

integration, consistent with the installation process for conventional datacenter accelerators.

- **Calibration:** Routine calibration covers laser operation, photon generation, interferometer accuracy, and detector performance, and is fully automated without requiring system downtime, thermal cycling, or environmental resets. Time-bin encoding and fiber-integrated components provide inherent long-term stability, minimising calibration frequency and eliminating the need for complex manual optical realignment.
- **Maintenance:** The photon detectors operate inside a closed-cycle cryo unit operating at temperature of ~2-4K. This subsystem requires maintenance on 10,000 and 20,000 hour periods, performed by the vendor.

5.3. Deployment of modalities: Annealing, Silicon Spin, and Superconducting qubit architectures

Table 1 outlined similar requirements across annealing, solid-state spin, and superconducting qubit modalities. In general, this relates to similar usage of control systems and cryogenic platforms - in particular the need to operate below temperatures that a typical cryocooler can obtain. The most common way to reach these lower temperatures is by a dilution refrigerator, which operates at temperatures of 10 mK. These low temperatures enable the quantum phenomena that underpins these modalities to arise, it also enables use of specific components such as low noise amplifiers or detectors for readout.



Figure 10: Examples of dilution refrigerator systems from Zero Point Cryogenics (left; Model L, and middle; Model I) and from Maybell (right; Fridge). The image on the left shows a system without user-wiring, here the upper two stages show integration of cryo-coolers. Within the middle image the cabinet shows the gas-handling system which contains electronics, pumps and storage vessels to enable the dilution fridge to operate. The right image shows a form-factor more aligned with datacenter cabinet layouts.

A dilution refrigerator enables continuous cooling of solid-state devices to sub-kelvin and millikelvin temperatures. It operates using helium isotopes within a closed-circuit. When at low temperatures, a unique phase behavior of a mixture of helium-3 and helium-4 takes place. Below approximately 0.87 K, a mixture of these two isotopes spontaneously separates into two liquid phases: a concentrated phase that is nearly pure helium-3 and a dilute phase consisting mostly of helium-4 with a small fraction (≈ 6 to 7% at zero temperature) of dissolved helium-3.

Individually, the two isotopes have different low-temperature limits when cooled by evaporation. Pure helium-4 can be pumped to reduce its vapor pressure and reach temperatures of roughly 1 K. Pure helium-3, owing to its lower boiling point and higher vapor pressure at low temperature, can be evaporatively cooled to approximately 0.3 K. These limits arise from the thermodynamic properties of each isotope when used alone.

When helium-3 and helium-4 are mixed at sufficiently low temperature, however, their behavior changes fundamentally. Helium-4 becomes superfluid below 2.17 K and has essentially zero vapor pressure at millikelvin temperatures, while helium-3 remains a normal Fermi liquid and retains significant vapor pressure to much lower temperatures. In a dilution refrigerator, helium-3 atoms are continuously transferred from the concentrated phase into the dilute phase across the phase boundary in the mixing chamber. This transfer is driven by a difference in chemical potential rather than by a conventional phase transition.

When helium-3 atoms enter the dilute phase, they absorb energy because the enthalpy of helium-3 in the dilute phase is higher than in the concentrated phase. This enthalpy difference

produces cooling. By continuously circulating helium-3 through the system and removing it from the dilute phase via evaporation at higher temperature stages, the cycle can be sustained indefinitely. The cooling power scales approximately with helium-3 flow rate and the square of temperature, enabling dilution refrigerators to reach temperatures in the single-digit millikelvin range, asymptotically approaching absolute zero.

The companies that produce dilution refrigerators can be referred to as cryogenic original equipment manufacturers (cryogenic-OEMs). Quantum computing modalities that use dilution refrigerators are then a collaborative design between the systems integrator who integrates products from the cryogenic-OEM. Figure 10 shows an example of the dilution refrigerator as a standalone product from a cryogenic-OEM.

5.3.1. Facility and Zone

Facility and zone considerations for cryogenically-cooled quantum modalities are largely similar, with shared requirements for thermal stability, vibration isolation, electromagnetic shielding, and careful integration of supporting infrastructure. These are collected into Table 1 and Table 2 to enable quick referencing and comparison.

Below is a table collecting facility requirements for the following modalities; superconducting qubits, spin qubits, and annealers. For this table, the delivery physicals generally relate to the dilution fridge. Likewise, the electrical and water requirements generally relate to the cryocooler requirements. The humidity requirements generally map to the control layer and its related signal delivery components. The HVAC details relate to the ambient cooling of these control and signal delivery components. Finally, the sensitivity to vibrations relates to magnetic fields at the QPU layer. Within the system integrator's product design, they also design integrated shielding and isolation mechanisms designed to attenuate environmental electromagnetic, acoustic, and vibrational noise. While these features can significantly relax certain siting constraints, they do not eliminate the need for rigorous facility assessment and environmental characterisation during deployment planning.

Table 3: Facility and zone requirement for dilution refrigerator based quantum system deployment, information compiled based on the vendor participated as listed in the brackets under modality column.

Modality	Delivery physical clearance	Storage room	Staging area	Structural Floor support for frame	Electrical details	Temperature details	Humidity details	Vibration details	Coolant details	HVAC details
Superconducting (OQC, IQM)	Height 2.3m Width 0.9m Weight of the heaviest component 750kG (max)	2.4m x 2.4m x 3m	6.1m x 6.1m x 3.6m	1000 kg/m2	30kW (peak) 208V, 60Hz, 60A (standard) 240/400V, 50Hz, 32A (international)	20°C to 25°C ΔT < 1 °C within 12 hours	25 - 60% RH	400 μm/s	Max water pressure 88 psi.	5 kW (17,000 BTU/h) normal mode 12.5 kW (43,000 BTU/h) aux. Mode
Annealing (D-Wave)				20°C to 25°C ΔT < 1 °C within 15 mins		50 μm/s		Min water temperature 15°C at 9.4L/min. Max water temperature 25°C at 20.5L/min.		
Solid-State Spin (Dirac)				500 kg/m2				20°C to 21°C ΔT < 1 °C within 15 mins	Max. water pressure 88 psi Water temperature at 10°C to 28°C at 1 to 2L/min/kW	

5.3.2. Superconducting Qubits

This part of the section is contributed by IQM and Oxford Quantum Circuits (OQC). A superconducting quantum system is the widely deployed quantum modality. These systems use chips containing lithographically fabricated circuits—typically aluminum or niobium—that act as qubits when cooled to millikelvin temperatures. At these ultra-low temperatures, superconductivity eliminates resistive losses, allowing qubits to form highly coherent quantum states suitable for computation. Qubits are manipulated using precisely controlled microwave pulses delivered via engineered control lines.

The QPU resides within a dilution refrigerator that maintains an operating temperature of 10–20 mK. The dilution refrigerator consists of multiple temperature stages, each progressively colder toward the QPU. Thermal isolation is achieved through high-vacuum insulation, while cooling is provided by a combination of pulse-tube cryocoolers and dilution refrigeration stages. In addition to thermal management, the system provides electromagnetic shielding, vibration isolation, and supports specialised wiring and RF componentry, which are often customised for the specific QPU architecture.




Figure 11: Photograph of an OQC system with removed dilution refrigerator cover (left) and a close-up of components inside an IQM host cryogenic system (right).

While a Dilution refrigerator may contain the QPU, it still requires compressors (for the pulse tubes) and equipment to operate the dilution stage (gas handling system). Additionally control electronics that are used to generate RF pulses to control and read the state of the Qubits will need to be placed near to the QPU.

- **Environmental Control:** Superconducting gate-based QPUs require an operational temperature of 10 to 20 mK to maintain qubit coherence, provided by a dilution refrigerator. Continuous and stable power is essential to maintain this temperature, as any rise above specified temperature renders the system unstable and necessitates a full warm-up followed by cooldown. Beyond the base temperature requirement, the cryostat stages must provide sufficient cooling power to compensate for all heat loads, including ohmic heating from the QPU and auxiliary electronics as well as thermal conduction along wiring infrastructure. Managing thermal conduction is particularly challenging for large bundles of coaxial RF lines, which scale with the number of physical qubits.
- **Qubit Control:** Superconducting systems use microwave pulses for qubit manipulation and readout.
- **Pre-Deployment and Deployment:** This phase takes about 7 to 9 days. The primary effort is a comprehensive environmental site assessment to characterise EMI, ambient and DC magnetic fields, mechanical vibrations, temperature stability, and humidity. Measurements are conducted continuously to capture daily and operational variability and to identify transient disturbances.
- **Calibration:** Automated calibration takes about 20 minutes and it typically runs a couple of times a day through workflow. Full optimisations can last up to 100 minutes.
- **Consumables:** Liquid Nitrogen can be used as a consumable as part of a cold trap (to mitigate against ingress of air-based impurities), with an estimated consumption of 6L per day. While gas handling systems which do not use liquid nitrogen traps are available, this may not be part of an integrator's standard product offering.
- **Maintenance:** Maintenance activities are driven primarily by the cooling and support infrastructure, with intrusive procedures defined by the cryogenic or cooling-system manufacturer. Preventive maintenance is typically required every six months, consisting of a one-day service cycle for cryogenic systems. This maintenance includes flushing the liquid nitrogen circuit to remove accumulated ice and contaminants. Additional age-dependent activities may include UPS battery servicing and replacement of cryogenic pump tip seals. System thermal cycling (warm-up and cool-down periods) must be explicitly accounted for in operational scheduling and maintenance planning.




Figure 12: Photograph of an IQM system without covers (left) and a view inside the electronics rack (right). The microwave pulses for qubit manipulation and readout are generated by the quantum control system (QCS). A separate microwave generator (MWG) produces the continuous wave pump signals for the Travelling Wave Parametric Amplifier (TWPA, see Figure 11) used for multiplexed readout of superconducting qubits.

5.3.3. Annealing

Annealing quantum computers were the first commercially available quantum computing systems, providing practical access to quantum-assisted optimisation and probabilistic sampling applications. In quantum annealing, the system is initialised in a low-energy ground state and the parameters of the target problem are introduced gradually through controlled evolution. This adiabatic process increases the likelihood that the system converges to a low-energy state of the problem Hamiltonian, corresponding to an optimal or near-optimal solution.

In optimisation workloads, annealing quantum computers address problems characterised by large combinatorial search spaces, such as scheduling, routing, and resource allocation (e.g., logistics planning, transportation routing, and operational scheduling). In probabilistic sampling workloads, the system samples from multiple low-energy states to characterise the energy landscape of a model. This capability is particularly valuable for machine learning and probabilistic modeling, where sampled distributions provide information about model states for a given parameter set and support iterative model refinement and training.

- **Environmental control:** The QPU operates at temperatures below 20 mK, maintained by a closed-loop cryogenic dilution refrigerator system. This system is protected from environmental interference through integrated RF shielding and magnetic shielding subsystems to preserve qubit coherence and system stability.

- **Qubit control:** Each qubit and coupler on the QPU has several controls that are manipulated by individual on-QPU digital-to-analog converters (DACs). Along with the DACs, a small number of analog lines provide the time-dependent control required by annealing.
- **Pre-deployment:** Similar to 5.3.2
- **Deployment:** This phase spans 5 to 8 weeks and must be carefully scheduled to align with environmental stability, facility readiness, and operational constraints. Activities include the physical assembly of the cryostat and associated control electronics, installation and verification of hardwired electrical and plumbing connections, and coordination with facility interfaces. Following mechanical and electrical integration, the system undergoes initial cooldown and baseline calibration to validate thermal performance, qubit control paths, and readout integrity. Calibration procedures vary by qubit modality.
- **Calibration:** Calibration is performed after the QPU has reached base temperature and persists for the lifetime of the QPU entire cooldown. During calibration of a QPU, a small number of qubits and couplers may not meet the specifications to function as desired, and, thus, they are removed from the programmable structure that users can access.
- **Maintenance:** The annealing quantum computer is typically maintained on a quarterly basis. During maintenance events, software, firmware, and the cryogenic system are updated and serviced.
- **Consumables:** The annealing quantum computer's cryogenics systems require certain gases and cryogenics as consumables. Nitrogen and helium gases are typically needed, with expected usage of approximately one T-size cylinder per gas per year. Liquid nitrogen is also required for cold traps, with an estimated consumption of around 6 L per day.



Figure 13: D-Wave Advantage2 quantum annealing system deployed in a facility. The system uses a closed-loop cryogenic dilution refrigerator, housed in a magnetic shielding enclosure, to maintain the QPU at its operating temperature below 20 mK.

5.3.4. Silicon Spin Based Qubits

This section is contributed by Diraq. Diraq is a silicon-based semiconductor spin qubit company, implementing electron spin qubits in CMOS-fabricated quantum dots. The qubits are encoded in the spin states of single electrons confined within semiconductor quantum dots, which are typically fabricated in silicon leveraging well-established CMOS manufacturing processes.

Qubit operations are performed using voltage and microwave pulses applied to gate electrodes, allowing for precise control of individual spins and their interactions. These operations are currently performed at millikelvin temperatures, but spins in silicon can achieve high coherence at 1 K — offering the possibility of cryocooler-only systems in the future.



Figure 14: Artist's rendering of Diraq's spin qubit system, encompassing a cryostat, a rack of room temperature control electronics and integrated classical processing.

The QPU comprises a quantum logic unit containing the qubits, and some integrated classical electronics required to control and read out the qubits, all at low temperature. Additional room-temperature electronics are required to decode quantum computations and enable hybrid algorithms. These resources must be connected to low-temperature stages within the cryostat by means of high-speed, low-latency digital links. Effective integration requires short physical paths between the QPU and HPC units, adequate shielding from EMI, vibration and airflow management, and uninhibited access for installation and occasional service.

- **Environmental control:** Spin-based QPUs require an operational temperature typically in the range of millikelvin to a few kelvin, depending on the qubit implementation. Maintaining such low temperatures is critical to preserve spin coherence and minimise thermal excitations that can disrupt quantum operations. If the temperature rises above the operational threshold, qubit performance degrades, and the system may require a full warm-up followed by a controlled cooldown to restore stability. The duration of such thermal cycling can vary from several hours to a day.
- **Qubit control:** Qubits are controlled using a combination of microwave and voltage signals delivered to gate electrodes with high temporal and spectral precision. The control electronics are typically housed in racks adjacent to the cryostat. Careful cabling design is required to preserve signal integrity and minimise thermal conduction to the QPU. The overall control stack is designed to support scalable, low-latency operation, enabling precise qubit rotations, entangling gates, and high-fidelity readout essential for quantum computation.

- **Pre-Deployment and Deployment:** Similar to 5.3.2
- **Calibration:** Installation/bring-up time is 5 to 6 weeks for present day systems, a small-scale prototype, and up to 12 weeks at utility scale.
- **Maintenance:** Non-disruptive, in-person service is required on a weekly basis where liquid nitrogen is involved (for larger-scale deployments). Semi-annual preventative maintenance (may not require full reboot). Disruptive service might be required on a half-yearly basis for system-health evaluation, and on a timescale of 2 to 3 years for compressor and pump servicing.
- **Consumables:** Spin-based quantum systems require specific gases and cryogens to maintain operational performance. Nitrogen and helium gases are typically required, with expected usage of approximately one T-size cylinder per gas per year. Liquid nitrogen is also used for system cooldown and magnetic shielding, with an estimated consumption of ~6 L per day. Additionally, the system requires cooling water maintained at 10 to 28 °C (50 to 82 °F), with a flow rate of 1 to 2 L/min per kW of thermal load.

5.4. Deployment of Modalities: Diamond based architectures

Carbon atoms within diamond provide another route to producing qubits. Here, the diamond material can be functionalised in different ways such as introducing nitrogen to enable nitrogen-vacancy (NV centres) or introducing other atomic species. The result is a system where the electronic spin and nuclear spin states can be initialised, manipulated and read-out as qubits. The sections that follow show two approaches based on this diamond architecture.

5.4.1. Diamond spin qubits

This section is contributed by Quantum Brilliance, which builds NV-based QPUs designed for room-temperature operation. NV centers in diamond are a solid-state quantum computing modality based on localised spin systems embedded in a crystalline lattice. An NV center is a point defect formed by a substitutional nitrogen atom adjacent to a vacant lattice site in the diamond crystal. In its negatively charged configuration (NV⁻), the defect hosts an electronic spin that can be optically initialised and read out and coherently controlled using microwave fields.

Optical excitation is used to polarise and measure the electronic spin state, while microwave control enables precise manipulation of spin transitions. The electronic spin serves as the primary interface between the quantum device and the classical control and measurement stack, enabling initialisation, control, and readout under ambient operating conditions.

In addition to the electronic spin, NV centers provide access to coupled nuclear spin systems, including the nitrogen nuclear spin and nearby C nuclear spins in the diamond lattice. These nuclear spins typically exhibit significantly longer coherence times than the electronic

spin and are used as computational qubits or long-lived quantum memory. Quantum operations are implemented by mediating interactions between nuclear spins through controlled coupling to the electronic spin, which functions as a control and readout bus.

This layered spin architecture, optically addressable electronic spins coupled to long-lived nuclear spins enables quantum information processing in diamond at room temperature and distinguishes NV-center systems from quantum modalities that require cryogenic environments to maintain qubit coherence.




Figure 15: NV-based quantum computer systems. Left: Cluster of three NV-based quantum computers with two qubits each mounted in a standard HPC rack at Oak Ridge National Laboratory (ORNL). Image courtesy of ORNL. Right: CAD diagram of an NV-based quantum computer. The lower-left compartment shows the optical and static magnetic system housing the diamond chip. The upper-right compartment contains an x86-architecture PC, including a GPU. The control electronics are located below and are not visible.

- **Facility and zone:** Electromagnetic stability is the primary environmental consideration. These modalities rely on precise microwave and RF control, as well as stable bias magnetic fields, to ensure reliable qubit initialisation, manipulation, and readout. Deployment zones should therefore minimise exposure to strong, time-varying electromagnetic interference and implement predictable grounding, shielding, and cable-management practices consistent with high-frequency RF systems. Optical components used for qubit initialisation and readout are enclosed and light-tight. Systems are designed to comply with applicable laser and optical radiation safety

standards without necessitating facility-level laser control zones or additional access restrictions.

- **Environmental control:** When environmental baselines (EMI/EMC, mechanical vibration, thermal stability, and air quality) are within specified limits, the system does not require auxiliary shielding for qubit protection.
- **Qubit Control:** The control system integrates RF, microwave, and FPGA subsystems into a single enclosure, operates at ambient temperature, supports direct adjacency deployment to the QPU, and imposes no additional facility integration constraints.
- **Pre-Deployment:** Most aspects of the pre deployment survey follow standard procedures for accelerator installation. The primary distinction for room temperature NV center systems is the need to evaluate environmental stability, with attention to ambient temperature fluctuations, airflow variability, acoustic conditions, mechanical vibration sources, and the local electromagnetic environment, especially near high power electrical equipment. This phase takes about 4 days.
- **Deployment:** Deployment of room-temperature NV-center quantum processing units typically involves a small, specialised on-site team with direct access to the datacenter to perform system assembly, integration, and initial calibration. Following installation in the rack, activities focus on optical alignment, control system bring-up, and iterative calibration of qubit control and readout. Under normal conditions, the on-site deployment phase spans approximately 3 to 4 weeks.
- **Maintenance:** In architectures using free-space optical components, oil-immersion lenses in the optical readout path constitute the primary maintenance driver, requiring routine service intervals every ~3 to 6 months with typical downtimes of ~5 days, and full maintenance cycles conducted annually or biennially with typical downtimes of ~3 to 4 weeks.
- **Calibration:** Fully automated calibration that the system can run for many months without interruption. The system does not require to be shutdown for the calibration.

5.4.2. Diamond spin qubits with cryogenic electronics

This section is contributed by Fujitsu. Diamond spin qubits are a solid-state quantum modality that encode information in the spin states of electrons associated with atomic defects in diamond, most commonly NV centers. These qubits can be initialised, manipulated, and read out using optical and microwave techniques, allowing coherent control of quantum states. In addition, neighbouring nuclear spins can also be used as qubits. Unlike many other solid-state systems, diamond spin qubits can operate at room temperature, though performance can improve under cryogenic conditions.

The QPU is built around engineered diamond crystals containing these defect centers, with each qubit fixed at a precise lattice location. Optical excitation and spin-dependent fluorescence enable high-fidelity readout, while microwave pulses allow single- and multi-qubit operations. This approach leverages the robustness of diamond as a host material, providing long coherence times and resilience to environmental noise.

Deployment of diamond spin systems typically requires modest cryogenic cooling or precision optical setups rather than full dilution refrigerators. Infrastructure needs include magnetic shielding, stable optical paths, and low-vibration environments. While scaling to very large qubit counts remains challenging, diamond spin qubits are well-suited for quantum sensing, small-scale accelerators, and hybrid computational applications.

- **Facility and zone:** This part is similar to section 5.3.1
- **Environmental control:** Diamond spin qubits require laser stabilisation and cryogenic operation at 4 to 10 K. Vibration isolation is also an important component of environmental control for this modality.
- **Pre-deployment:** Since diamond spin qubits incorporate laser components, the pre-deployment survey must evaluate air conditioning systems to maintain laser and optical stability. It should also assess ambient light levels and confirm the presence of appropriate shielding or light-controlled conditions to prevent optical interference. This is done in parallel to cryogenic related site assessment.
- **Deployment:** In addition to installing the cryogenic system, an optical table is set up, followed by the wiring and integration of all control electronics.
- **Maintenance:** Since these systems haven't been deployed, there is no information on maintenance.
- **Calibration:** Short calibrations typically take a few minutes, while longer calibrations can last tens of minutes to several hours.
- **Consumables:** Diamond spin systems do not involve the consumption of gases. Nevertheless, certain cryostat specifications might necessitate a small supply of helium or nitrogen gas during maintenance

6. Future Efforts

As QPUs transition from laboratory systems to production datacenter assets, there is a growing need to standardise both hardware interfaces and the commissioning of workflows. Unlike traditional IT infrastructure, QPU deployments currently lack common physical integration standards, interface specifications, and operational certification processes, creating operational friction for facility operators. At the same time, both datacenter architectures and QPU technology roadmaps are evolving in parallel, often independently. This co-evolution creates a

coordination gap where facility design cycles, infrastructure investment horizons, and quantum hardware development timelines are misaligned.

The OCP intends to lead a coordinated, multi-stakeholder effort to close this gap by driving cross-industry standardisation across power delivery, cooling, mechanical integration, control interfaces, safety systems, and operational acceptance testing, enabling scalable, repeatable, and production-grade QPU deployment models. These efforts will focus on key characteristics of physical QPU identified in Section 3.3.3.

6.1. Non-Quantum OCP Projects Enabling Scalable and Reliable QPU Deployment

While QPU standardisation is a critical focus area, the success of quantum deployment in production datacenters is equally dependent on the alignment of existing non-quantum infrastructure initiatives within the OCP. These foundational domains already define how modern datacenters are designed, built, and operated. They must converge into a coherent integration model that can support quantum–classical hybrid environments at scale.

The following OCP project areas form the non-quantum backbone of quantum infrastructure readiness:

- **Facility:** Facility architecture defines the operating environment of physical QPUs—controlling vibration, EMI, spatial zoning, access control, and environmental stability. These parameters directly influence system stability, uptime, and deployability.
- **Hardware management:** Hardware management provides the operational control, observability, and fault-handling frameworks required to operationalise QPUs as managed infrastructure rather than bespoke laboratory systems. This includes standardised interfaces for resource discovery, configuration and firmware updates, telemetry, control, asynchronous events, and hardware fault management—ensuring predictable system behavior under failure conditions and enabling integration into rack- and datacenter-level management planes.
- **Networking:** Classical networking fabrics define the control-plane and data-plane scalability of hybrid quantum systems. Deterministic latency, synchronisation, and bandwidth provisioning directly affect orchestration efficiency and closed-loop quantum control.
- **Cooling environment:** Environmental and thermal frameworks define the thermal operating envelope of QPUs, including cryogenic integration, heat rejection paths, and thermal stability—factors that directly impact qubit coherence, noise floors, and operational reliability.
- **Rack and power:** Rack form factors, grounding, and power distribution define electrical stability, scalability, and deployment density, directly shaping system reliability and

long-term scaling trajectories for physical QPU installations.

- **Sustainability:** Sustainability frameworks define long-term operational viability, lifecycle energy efficiency, and resource optimisation—critical constraints for scaling quantum systems beyond pilot deployments.
- **Storage:** Storage architectures define data-plane scalability, workflow persistence, and classical post-processing capacity, directly shaping the throughput and scalability of quantum workloads.
- **Time appliance project:** Precision timing and synchronisation provide the temporal reference required for coordinated quantum–classical operation. Accurate clock distribution and low-jitter synchronisation enable deterministic control pulses, coordinated classical controllers, and timely quantum error-correction feedback across distributed systems.

Collectively, these non-quantum efforts define the operating environment, scalability envelope, and performance stability of physical QPU deployments. They directly influence qubit fidelity, coherence stability, system reliability, and scale-out feasibility—which are core characteristics of any production-grade quantum platform.

The OCP will bring expertise from these non-quantum domains to help QPU vendors develop systems that are not only functional but also optimised for real-world deployment, ensuring that each system can reliably operate, scale, and deliver predictable qubit performance within production datacenter environments.

6.2. Quantum Specific Efforts at the OCP

Many of today's quantum systems, across multiple hardware modalities, were originally designed for laboratory environments as self-contained platforms and later adapted for HPC facilities. Replicating this model directly in datacenters would lead to duplicated infrastructure, fragmented maintenance practices, and increased facility power demand, creating operational and economic inefficiencies. A more integrated approach is required, one that embeds QPU systems into shared datacenter infrastructure through standardised interfaces, coordinated environmental controls, and optimised use of facility resources to enable scalable, reliable, and maintainable operations.

A fundamental requirement for QPU reliability, regardless of modality, is a clearly defined operating zone with controlled environmental conditions and uninterrupted power. Different quantum technologies have different sensitivities such as temperature stability, vibration, electromagnetic interference, humidity, or acoustic noise, but all require environmental consistency. As vendors deploy systems outside laboratory settings, they are learning how real-world datacenter conditions affect performance and are strengthening shielding, enclosures, and isolation techniques accordingly. At the same time, systems must be designed to limit their own emissions such as vibration, thermal output, or electromagnetic noise to

prevent impact on surrounding infrastructure. Both inward protection and outward containment must be engineered deliberately.

An example of infrastructure level integration is around how best to scale cryogenic platforms. Here, it is possible to replace the cryocoolers with a consolidated liquid helium layer. At scale, this has the potential to offer advantages in terms of cryogenic performance per energy consumption. However, the trade-off is in another closed cycle system with liquid helium that could lead to reduced redundancy. An example of this architecture is under development by Maybell Quantum Industries, illustrated in Figure 16. In this approach, centralised refrigeration supports multiple modular quantum nodes. Rather than each system operating as a fully independent cooling plant, shared infrastructure provides core services and electronics are organised in adjacent racks. While this example reflects a cryogenic implementation, the architectural principle of shared plant resources supporting modular quantum nodes is modality agnostic and illustrates how efficiency and scalability can be improved through coordinated design. Note, at this time it should be expected that the design here sits within the demarcation of the system integrator rather than the facility.



Figure 16: Conceptual render of Maybell ColdCloud distributed cryogenic architecture illustrating centralised 4 K plant infrastructure supplying modular millikelvin nodes.

Another critical layer is the control system. Regardless of hardware type, a quantum controller must interface between the classical host environment, facility infrastructure, and the

physical QPU. Its integration affects performance, scalability, and operational reliability. The time required to measure quantum states, process results classically, and apply corrective actions influences overall system accuracy. As deployments scale, minimising delay, ensuring signal integrity, and coordinating with facility systems become increasingly important. Controller integration must therefore be designed specifically for datacenter environments rather than treated as a laboratory extension.

Operational maturity also requires clearly defined service expectations. Uptime targets, mean time between failures, maintenance intervals, calibration processes, and consumable management must transition from research norms to production standards. Across modalities, positive trends have emerged in longer uptime, improved reliability, and more predictable maintenance cycles as systems are deployed outside laboratory settings. Vendors are incorporating field experience into system hardening and serviceability design. Where consumables are required, resilient supply chains and standardised replacement procedures must be incorporated into deployment planning.

At a larger scale, physical layout and connectivity become performance relevant. Classical networking between controllers, processing systems, and storage must scale with node count. Physical placement, cabling pathways, and rack adjacency can affect signal quality, maintainability, and operational efficiency. Quantum datacenters therefore combine elements of compute clusters and specialised infrastructure environments. Infrastructure design decisions influence not only logistics but also system stability and performance.

The constraints described such as environmental sensitivity, cooling efficiency, input and output density, vibration management, maintenance coordination, and latency are not isolated component issues. They are systems challenges. Achieving utility scale quantum computing requires applying structured reliability and systems engineering principles to quantum infrastructure from the outset.

The OCP will coordinate quantum specific efforts to define common interfaces, integration guidelines, reliability metrics, and monitoring standards across modalities. The objective is to move from laboratory style deployments to standardised datacenter integrated architectures in which environmental control, power, cooling, control systems, and operational processes are engineered as a unified system.

7. Conclusion

The success of production-grade quantum systems depends on more than device performance. Environmental control, power delivery, cooling, networking, storage, and operational processes directly govern qubit fidelity, coherence stability, and overall scalability—the core characteristics that determine whether a quantum system can operate reliably at scale.

The OCP will play a critical role in bridging quantum system vendors with expertise from existing OCP projects. By coordinating standards, integration guidelines, and best practices across facility infrastructure, classical compute, and quantum-specific systems, the OCP will simplify deployment, improve operational reliability, and reduce complexity for datacenter operators. This cross-domain collaboration ensures that quantum systems are not only technically capable but also practical to operate, maintain, and scale within modern datacenter environments.

In short, scaling quantum computing requires a systems approach: one where engineered infrastructure, standardised processes, and collaborative expertise come together to transform experimental devices into production-ready, reliable, and maintainable datacenter assets.

8. References

- [1] J. Preskill, “Quantum Computing in the NISQ Era and Beyond,” *Quantum*, vol. 2, p. 79, 2018. doi:10.22331/q-2018-08-06-79.
- [2] A. Batson and D. Thorpe, “The Future of Quantum Real Estate,” JLL Insights, Jones Lang LaSalle IP, Inc., Jul. 30, 2025.
- [3] H. Soller, M. Gschwendtner, and V. Kermans, “Making Fault-Tolerant Quantum Computers a Reality,” *McKinsey Tech: Forward*, Dec. 8, 2025.
- [4] S. A. Caldwell et al., “Platform Architecture for Tight Coupling of High-Performance Computing with Quantum Processors,” *arXiv preprint arXiv:2510.25213*, 2025.
- [5] J. Gransden, “GPUs, ASICs or FPGAs? Here’s How They Measure Up for Quantum Error Correction,” *Riverlane Blog*, Aug. 28, 2025.
- [6] Microsoft Quantum, “Types of Qubits,” *Microsoft Quantum Insights – Education Concepts*, 2025.
- [7] K. Rallis et al., “Interfacing Quantum Computing Systems with High-Performance Computing Systems: An Overview,” *arXiv preprint arXiv:2509.06205*, 2025.
- [8] E. Mansfield et al., “First Practical Experiences Integrating Quantum Computers with HPC Resources: A Case Study With a 20-Qubit Superconducting Quantum Computer,” *arXiv preprint arXiv:2509.12949*, 2025.
- [9] C. M. Löschnauer et al., “Scalable, High-Fidelity All-Electronic Control of Trapped-Ion Qubits,” *PRX Quantum*, vol. 6, no. 4, 2025. doi:10.1103/h4wk-v31j.
- [10] A. C. Hughes et al., “Trapped-Ion Two-Qubit Gates with >99.99% Fidelity Without Ground-State Cooling,” *arXiv preprint arXiv:2510.17286*, 2025.

9. Creative Commons License

OCP encourages participants to share their proposals, specifications and designs with the community. This is to promote openness and encourage continuous and open feedback. It is important to remember that by providing feedback for any such documents, whether in written or verbal form, that the contributor or the contributor's organisation grants OCP and its members irrevocable right to use this feedback for any purpose without any further obligation.

It is acknowledged that any such documentation and any ancillary materials that are provided to OCP in connection with this document, including without limitation any white papers, articles, photographs, studies, diagrams, contact information (together, "Materials") are made available under the Creative Commons Attribution-ShareAlike 4.0 International License found here: <https://creativecommons.org/licenses/by-sa/4.0/>, or any later version, and without limiting the foregoing, OCP may make the Materials available under such terms.

As a contributor to this document, all members represent that they have the authority to grant the rights and licenses herein. They further represent and warrant that the Materials do not and will not violate the copyrights or misappropriate the trade secret rights of any third party, including without limitation rights in intellectual property. The contributor(s) also represent that, to the extent the Materials include materials protected by copyright or trade secret rights that are owned or created by any third-party, they have obtained permission for its use consistent with the foregoing. They will provide OCP evidence of such permission upon OCP's request. This document and any "Materials" are published on the respective project's wiki page and are open to the public in accordance with OCP's Bylaws and IP Policy. This can be found at <http://www.opencompute.org/participate/legal-documents/>. If you have any questions, please contact OCP.

10. About Open Compute Foundation

The Open Compute Project (OCP) is a collaborative Community of hyperscale datacenter operators, telecom, colocation providers and enterprise IT users, working with the product and solution vendor ecosystem to develop open innovations deployable from the cloud to the edge. The OCP Foundation is responsible for fostering and serving the OCP Community to meet the market and shape the future, taking hyperscale-led innovations to everyone. Meeting the market is accomplished through addressing challenging market obstacles with open specifications, designs and emerging market programs that showcase OCP-recognised IT equipment and datacenter facility best practices. Shaping the future includes investing in strategic initiatives and programs that prepare the IT ecosystem for major technology changes, such as AI & ML, optics, advanced cooling techniques, composable memory and silicon. OCP Community-developed open innovations strive to benefit all, optimised through the lens of impact, efficiency, scale and sustainability. Learn more at www.opencompute.org