




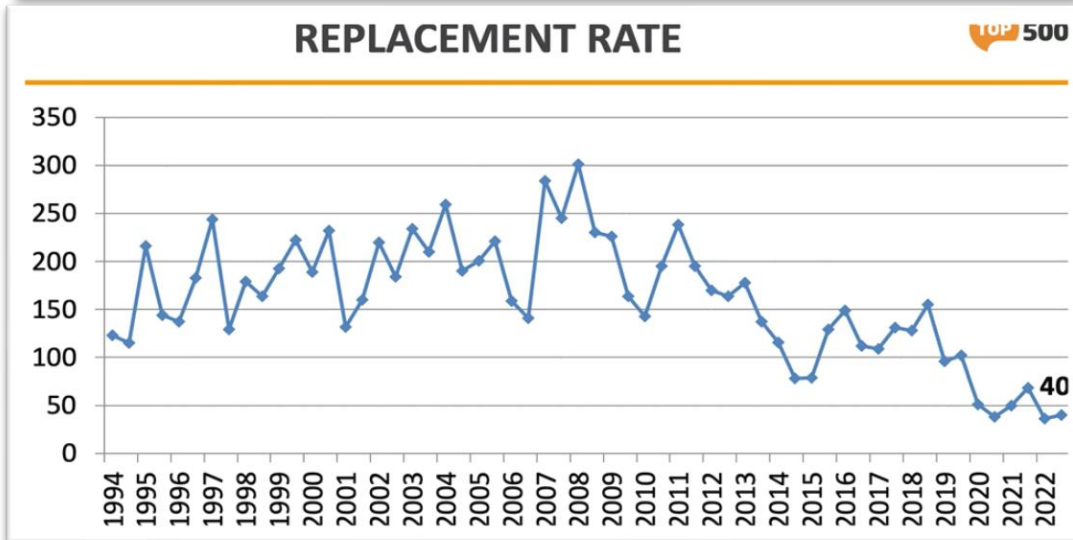

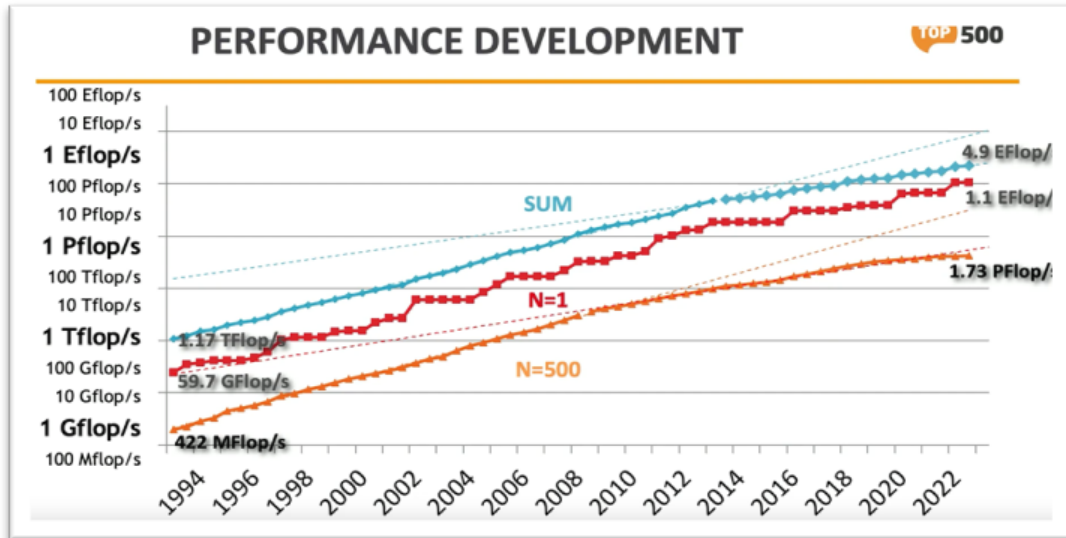
HYPERION RESEARCH

A Strategy for Introducing QC into the Classical HPC Ecosystem

Bob Sorensen
Chief Analyst for Quantum Computing
Hyperion Research Holdings, LLC

Setting the Scene: HPC Today—What a Mess

The Classical HPC Sector Today



But Computational Limitations Are Only Half the Story

- Cost at the highest end of compute continue to climb
 - Recent high-end systems cost upwards of US\$600 million
- Power requirements are increasingly becoming an integral element of TCO
 - 25 megawatts or more for a high-end system
 - *Rule of thumb:* US\$1.5 million/megawatt/year = $1.5 \times 25 \times 5 = \text{US\$187}$ million total utility cost (not counting resulting cooling requirements)
- Spate of new workloads complicate HPC design and operations
 - Modeling/simulation, big data analysis, AI
 - Each with (sometimes spectacularly) varying system hardware and software requirements
- Resulting in growing architectural options
 - CPU-only, GPU-centric, memory-centric, storage-centric
 - On-prem, cloud, hybrid, continuum, real-time, distributed
- Driving new architectural paradigms
 - Cloud, AI, analog, optical, dataflow, and quantum

QC From the HPC End User perspective

Primary End User Motivations from QC

New algorithms and future classical performance concerns lead

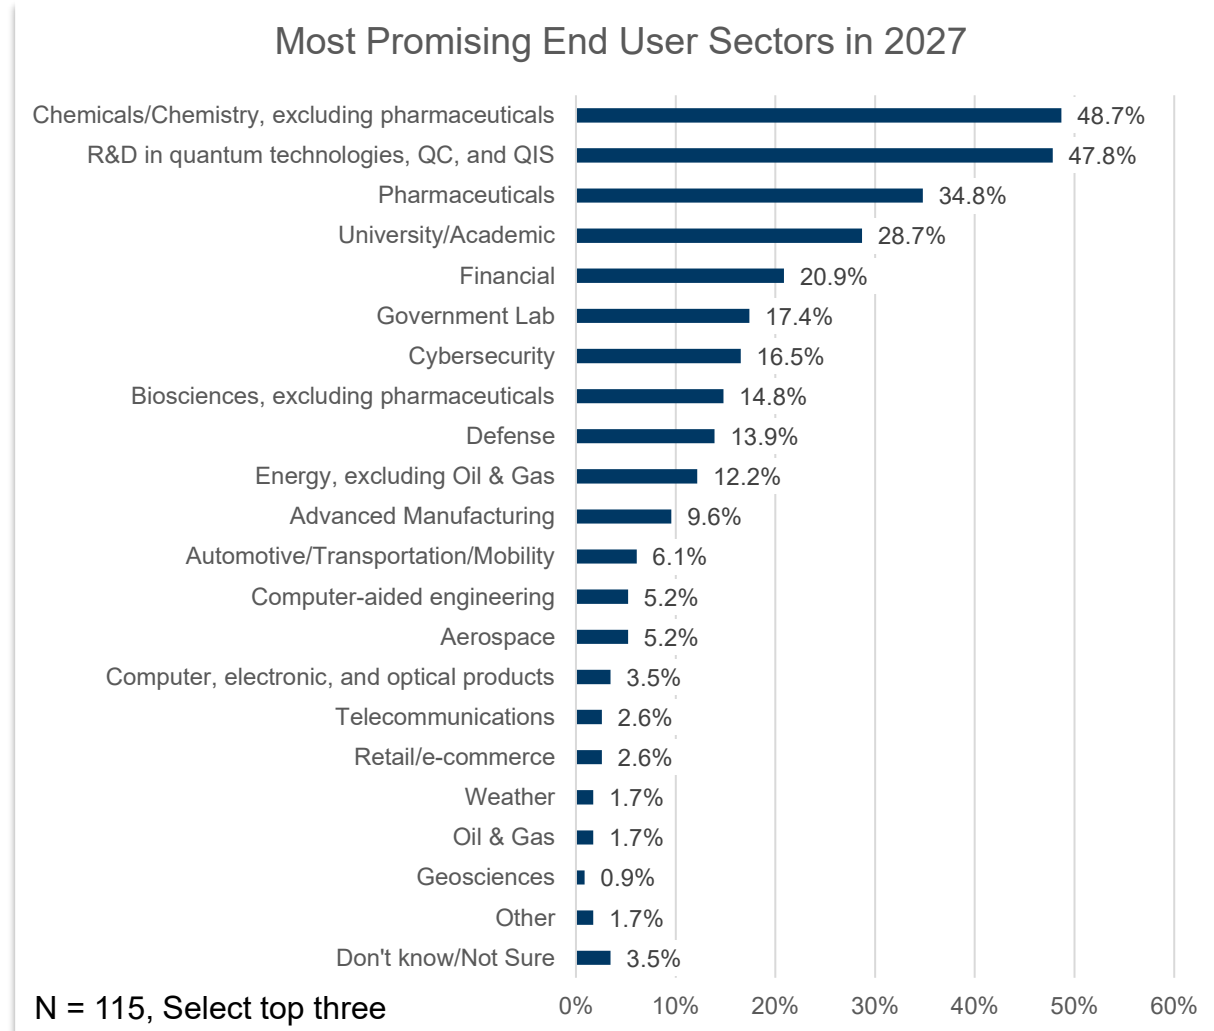
Option	% Selected
Implement new algorithm(s) not possible on classical counterpart systems	56.5%
Address concerns with future performance capabilities of classical computing systems	51.3%
Explore organizationally relevant QC use case potential with no expectations of near-term advantage	47.0%
Develop in-house familiarization with QC skills with no expectations of near-term end use deployment	45.2%
Engage with the QC vendor community for future activities	31.3%
Enable better real-time computational capabilities	27.8%
Realize faster turnaround time on existing classical counterpart systems	27.0%
Reduce overall computing systems costs	23.5%
Reduce overall computational power and cooling requirements	14.8%
Don't know/Not sure	5.2%
Other	2.6%

- Implement new algorithms and address concerns with future classical performance selected by majority of respondents
 - N.b. Classical developments could impact QC uptake
- An average of 3.3 options selected per respondent
- Many are still exploring for the sake of exploration
- One in four are looking at real-time compute opportunities
- Reduce overall compute systems cost:
 - 2023 Survey: 9.0%
 - 2025 Survey: 23.5%
- Reduce power/cooling costs
 - 2023 Survey: 17.3%
 - 2025 Survey: 14.8%

N = 115, Select all that apply

QC Market 2027: Top End User Sectors


Chemicals and QC R&D on top, but broad applicability envisioned



- Chemical/Chemistry sector hits number 1
 - Up from #2 last year, #4 year before that
 - Reinforces early emphasis on mod/sim, especially computational chemistry, as major algorithm
- Likewise, pharmaceuticals continues its upward climb
 - 21% last year, 35% this year
- Applicability spans academic, commercial, and government spaces
- Finance drops from 30% to 21%
 - Optimization issues, saturation, or contrived lack of visibility?
- Government labs hold steady, for now
- Although nearly every sector choice deemed important by some, there are clear concentrations in key areas

Greatest Hurdles to QC Adoption by HPC Sector

Led by complexity with IT integration and a lack of in-house QC expertise




- Typical concerns found in many major advanced computing survey going back decades
 - Spanning classical HPC, GPUs, AI, and now QC
- An average of ~3 identified hurdles per respondent
- Only 8% of US respondents selected low business priority
 - Compared with 28% of UK respondents
- No requirements for new compute capability: a perennial low number
 - Indicates significant pent-up demand for new solutions

N=303 Respondents could select all options that apply.

Utility-Class QCs: Expected Time Frame

About half of the respondents indicate a time frame of 2-4 years



- About half of the respondents (52%) expect the availability of utility class QC in the next 2-5 years
- About one in three say five years or more
- About 8% say they are already here or will be in the next year
- Utility class QCs are defined as systems used for mission- or business-critical, production-level compute jobs
- Utility class definition says little about specific QC metrics and benchmarks and more about perceived end use capabilities

Initiating the HPC/QC Adoption Process

Step One: Look Inward

This is not the time to be evaluating modalities, qubits counts, fidelities, etc.

- Aspiring QC end users in the HPC community need to accurately characterize their existing and planned workloads
 - What are the key computational jobs?
 - Defined as jobs that take the most compute time, are most important to end users, are most linked to ROR or ROI, etc.
 - What are the key pain points for those jobs?
 - Not fast enough up (long time to solution, too big for current systems)
 - No viable options for performance increases (algorithmic dead ends)
 - Limited classical vendor roadmaps (not AI or GPU-centric, needs high precision, etc.)
 - Requires too much power (utility cost, data center upgrades, etc.)
 - Requires too much cost
 - Not enough fidelity (too many simplifying assumptions, etc.)
 - Needs major code updates, upgrades (refactoring, GPU or AI conversions, ISV vs open source, etc.)

Step Two: Assess the Classical Remedy

Outline a plan to address the identified pain points classically

- What computation changes are needed to remedy current pain points with classical solutions?
 - Consider range of classical options: cloud, heterogeneous, GPUs, AI, etc.
 - Hardest question: Is it even feasible?
- What is the estimated cost to remedy pain points classically?
 - In terms of hardware, software, data center and most important, personnel
 - Personnel may be the major limiting factor
- What is an estimated time frame for implementation of classical remedies?
- What would be anticipated returns?
 - Technical: Time to solution, power costs, higher fidelity solutions, etc.
 - Non-technical: Faster time to science/solution, lower TCO, increased ROI/ROR
- Steps One and Two have benefits on their own

Step Three: Build an In-House Support Base for QC

Begin to assemble an in-house team to help support exploration of QC potential

- Enlist the aid of HPC end users and HPC center planners to support the exploration of QC opportunities
- Assemble a team of those with an interest in exploring QC opportunities
 - Use outside hiring of QC experts (hard)
 - Tap into in-house base of QC-interested HPC and SME users
 - Rally HPC planners and thought leaders' support for new technology
 - Encourage interested HPC end users to tap into various cloud access models, take on-line QC training/tutorials, attend QC conferences, etc.
 - Begin limited discussion with range of potential QC suppliers
- Lay the groundwork for building an economic QC adoption argument
 - Establish contact with C-suite to communicate potential QC opportunities from an ROI/ROR perspective
 - Leverage classical results of steps 1 and 2 to open door for QC exploration

Step Four: Explore Current QC Algorithmic Opportunities

Explore vendor-agnostic QC algorithm/application opportunities...for now

- Map noted QC opportunities against identified classical HPC pain points
- According to *Quantum Computing Technology Roadmaps and Capability Assessment for Scientific Computing* (arXiv:2509.09882):
 - At least 50% of DOE's Office of Science production workload could benefit from QC
 - Aggregates the computational needs of more than 12,000 NERSC users across the DOE landscape
 - Targeted problems for which exponential improvements in terms of problem complexity, time-to-solution, and/or solution accuracy
 - Report identified that materials science, quantum chemistry, and high-energy physics as the science domains and application areas that stand to benefit most from quantum computers

Step Four: Explore QC Algorithmic Opportunities (cont.)

HPC application area	High level description	Classical computational methods, tools and libraries	Impact of early fault-tolerant QPUs
Materials science	Simulating and predicting material properties (e.g., conductivity, magnetism, phase transitions).	DFT, TDDFT, Tight-binding methods, GW approximation, Molecular dynamics, DMFT	Probability of eFTQC impact: SURE
Quantum chemistry and Molecular dynamics	Simulations of molecular systems to compute energies, reaction pathways, molecular properties and study atomistic time evolution.	DFT, CCSD(T) (Coupled Cluster), Perturbation theory (MP2), DMRG, Hartree-Fock, Force field, Gromacs, NWChem, Molecular dynamics	Probability of eFTQC impact: HIGH
Nuclear fusion	Simulations of ionized matter to model plasma behavior in fusion devices, space, and astrophysical environments, including fluid, kinetic, and microphysical processes.	PIC, TDDFT, MHD solvers, GENE code	(for microphysics modeling only)
Nuclear and Particle Physics	Studies fundamental particles and their interactions, often at high energies, such as those in particle accelerators	Lattice QCD, Monte Carlo	
Computational fluid dynamics	Simulation of fluid flows (air, water, combustion) in engineering (e.g. aircraft design) and physics.	Navier-Stokes solvers, Lattice Boltzmann methods, OpenFOAM, ANSYS Fluent	
Astrophysics	Simulations of stars, galaxies, black holes, universe evolution	N-body simulations, CASTRO,	
Biosciences	Model biological processes at cellular/molecular scales (including protein structure prediction) and High-Throughput Screening (HTS).	BLAST, DOCK, Rosetta, alphafold, Abyss	Probability of eFTQC impact: MEDIUM
Structural modeling	Simulate how structures deform and respond under forces (stress, strain, vibration, failure).	Finite Element Methods	

- HPC applications with significant positive impact on HPC workloads according to:
- *Seizing Quantum's Edge, Why and How HPC Should Prepare for eFTQC* by Alice & Bob and Hyperion Research
<https://alice-bob.com/hpc-report/>
- Additional identified applications with low near-term impact included climate modeling, geoscience, finance, operations research, and AI
- N.b. these are for early FTQC (near-term)
- All these applications, however, are seen as promising based on expectations of future fully fault tolerant QCs

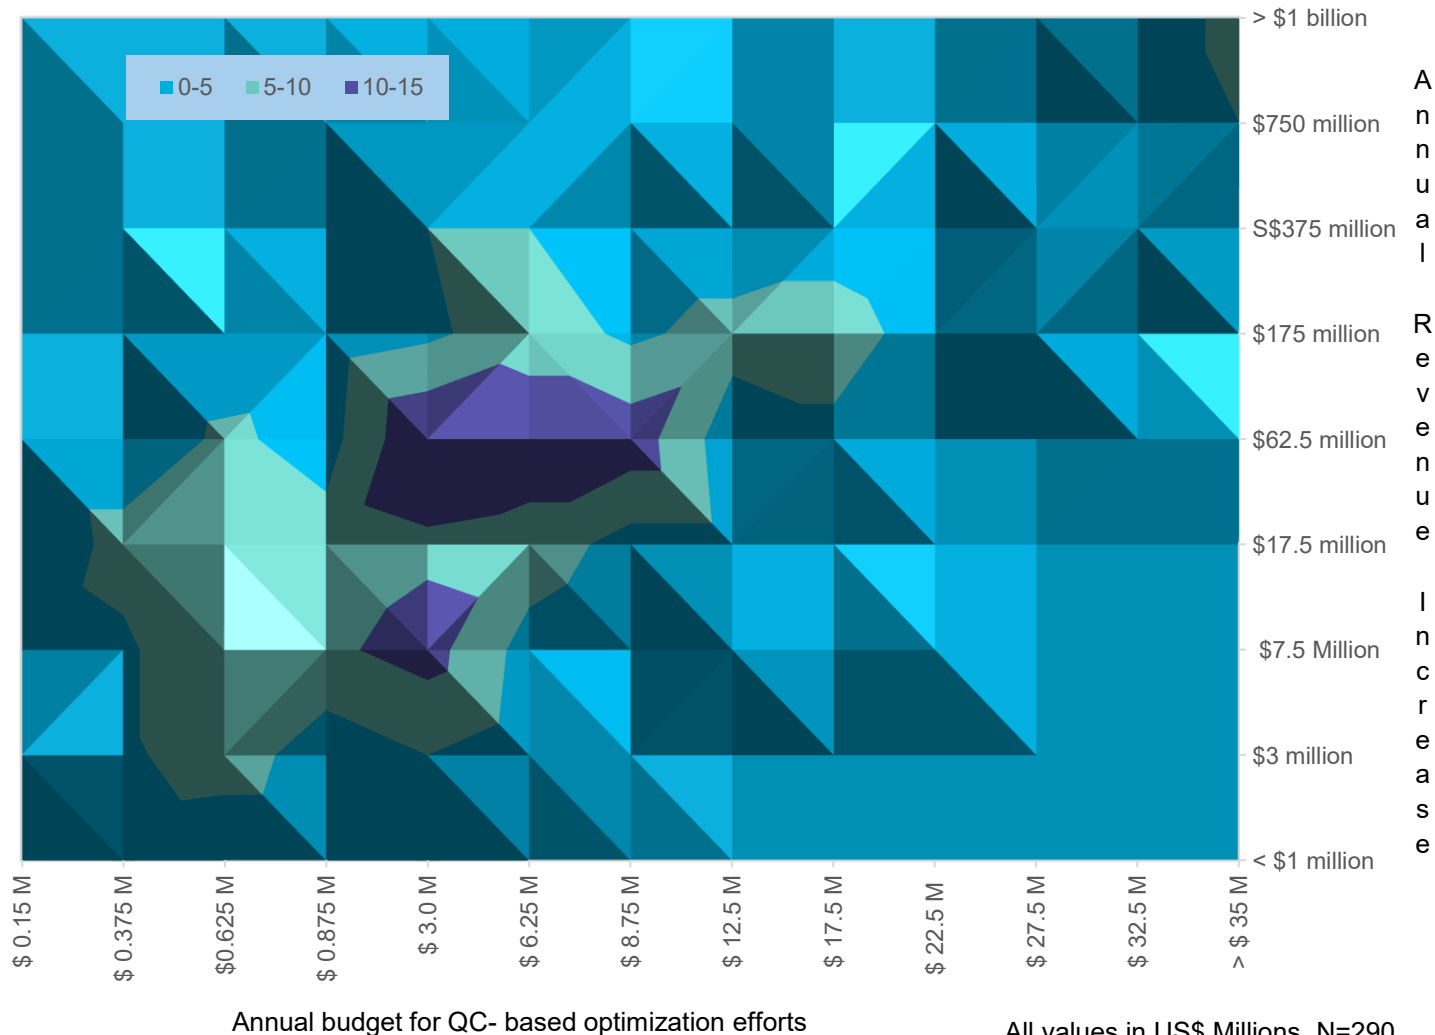
Step Five: Assess QC Value to Identified Classical HPC Jobs

- Map the current pain points and opportunities (identified in Step 1) to determine best-fit opportunities for QC to address classical HPC workloads (identified in Step 4)
 - Assess estimated performance improvements realized by shifting to QC implementations
 - Technical: Time to solution, power costs, higher fidelity solutions, etc.
 - Non-technical: Faster time to science, increased innovation, lower TCO, increased ROI/ROR
 - Factor in considerations for new:
 - HPC center operations, budgets, etc.
 - Hiring requirements, retraining efforts
 - QC facing classical front-end hardware
 - HPC/QC integration/orchestration software
 - Rewriting of existing applications code
 - Assess implications of not implementing QC solutions
 - *The cost of not doing something can be significant*

Step Six: Manage Expectations

- QC is not here today
 - But that gives time to do a measured effort on steps 1-4 starting now
- There are some significant technical obstacles ahead
 - This is not a black box procurement (yet)
- QC sector is Wild West right now
 - 84 QC hardware suppliers, according to Bob Sutor
 - Not all will survive, and the sector will go through some culling, M&As, etc.
 - QC sector funding levels could be significantly affected
- Not all performance gains will be exponential
 - But factors or orders of magnitude still valuable
 - Power and innovation opportunities matter as well
- Performance gains will be limited to narrow range of application/algorithms
 - That is why the classical pain points determination matters
 - The current base of applications matter and will likely expand
 - Many HPC issues will largely remain
- Financial impacts are TBD, but believed to be substantial

QC-based Optimization: Steady State Expected ROI



- Chart maps relationship between expected annual budget spent on **QC-based optimization** efforts (x-axis) against expected annual resulting revenues increases (y-axis)
- Most chosen options (purple region) centered on US\$3-6 million expenditures with resulting revenue increases of US\$60-\$65 million
 - Leading to an ROI of approximately between 10-20 to one
- One out of ten respondents foresaw annual QC-based optimization budgets to exceed US\$27 million while expecting an annual ROIs between US\$376 million to over US\$1 billion

Step Seven: Shop for QC Vendors and Systems

- Seriously reach out to QC vendors
 - The optimal vendor selection process entails identifying the QC vendors that can deliver demonstrated performance gains on key workloads that address the pain points and opportunities identified in Step 5
 - Details of qubit modality, qubits counts, gate fidelities, and road maps are best left as 'nice to know'
 - Translating QC hardware specifics into performance impact on end user workloads is complex and often not terribly satisfying
- Consider a range of QC supplier options:
 - Full stack providers, collections of HW/SW providers that contribute to a complete solution, QC integrators, QPU custom providers
 - Consider QC access model: Cloud, on-premises, hybrid
 - This will likely evolve as the sector matures
 - Focus on QC suppliers with vertical-specific expertise
 - Develop and use mini-benchmarks to assess QC vendor offerings
 - Partner with QC suppliers for exploratory projects

QC Partnerships With QC End Users

Most respondent organizations have a range of partnerships with QC end users

Option	% Selected
Explore new QC sector/vertical-specific QC-related opportunities	74.1%
Field test/evaluate new QC hardware	44.8%
Field test/evaluate new QC software	44.8%
Explore key performance gains over classical counterpart	43.1%
Establish sector-specific capabilities	41.4%
Foster public attention	39.7%
Encourage follow-on sales	36.2%
Explore QC/classical integration issues	31.0%
Access QC end user QC expertise	29.3%
Explorer QC sector/vertical-specific performance opportunities on existing classical workloads	27.6%
Access QC end user classical IT expertise	8.6%
Other	5.2%

- 71% of respondent organizations have a partnership with at least one QC end user
- Average respondent selected 4.2 options
- Field testing QC hardware and software both selected 44.8%
- Exploring sector-specific opportunities was overwhelming justification (74.1%)
- Building sector-specific skills (41.4%) and exploring QC performance advantages (43.1%) also key drivers

Step Eight: Procure QC Capabilities

- Take full advantage of your existing HPC procurement infrastructure for guidance on this complementary procurement;
- Treat this as part of the overall HPC-site procurement process, not a stand-alone activity
 - For Cloud:
 - Consider QC vendor cloud, CSP, or orchestration firms (Strangeworks) for QC access
 - Carefully review anticipated QC usage, software stack, and pricing options
 - Short-term vs long term pricing
 - QPU hours
 - QPU availability (time)
 - QPU upgrades
 - QPU variety
 - QC software development kits, etc.
 - For On-Premises:
 - Buy vs. lease
 - Vendor options for upgrades, mid-life kickers, new capabilities roll-outs options
 - On-site QC maintenance requirements
 - Site specific physical requirements
 - Vendor lock in-issues (and vendor long-term prospects)
 - Additional classical hardware requirements

Summary

- Step One: Look Inward
- Step Two: Assess the Classical Remedy
- Step Three: Build an In-House Support Base for QC
- Step Four: Explore Current QC Algorithmic Opportunities
- Step Five: Assess QC Value to Existing Classical HPC Jobs
- Step Six: Manage Expectations
- Step Seven: Shop for QC Vendors and Systems
- Step Eight: Procure QC Capabilities
- Step Nine: Compute, Assess, Repeat

QUESTIONS?

When the great innovation appears, it will almost certainly be in a muddled, incomplete, and confused form. To the discoverer himself, it will be only half understood; to everybody else it will be a mystery.

----Freeman Dyson 'Innovation in Physics', Scientific American, Sept 1958



bsorensen@hyperionres.com