J A N U S L A B S

Autonomous Scientific Intelligence for Everyone

FOUNDING THESIS

Building the Intelligence Layer for Scientific Discovery

DOCUMENT	Founding Thesis — Janus Labs
VERSION	0.1
CLASSIFICATION	For Investor Review
PREPARED BY	Janus Labs

This document contains forward-looking statements regarding technology development, market opportunity, and commercial strategy. Actual outcomes may vary.

CONTENTS

I. EXECUTIVE SUMMARY

Abstract

Quantum computing has entered a commercial phase, yet the software ecosystem surrounding it remains fragmented across hardware architectures, SDKs, calibration methodologies, and pricing models. As a result, a disproportionate share of research effort is spent on infrastructure evaluation rather than on scientific discovery.

Janus Labs is developing an autonomous, AI-assisted intelligence layer that sits between researchers and quantum hardware. The platform interprets research questions, designs candidate experiments, executes them across multiple providers, quantifies statistical confidence, and integrates results into a continuously growing Knowledge Graph.

This document presents the founding thesis of Janus Labs, an analysis of the problem space, the platform architecture, the commercial model, and the deliverables associated with the initial $75,000 funding milestone.

Key Findings

- Hardware progress in quantum computing has outpaced the maturity of its software infrastructure, creating a structural inefficiency across the industry.
- Researchers repeat the same infrastructure work — provider comparison, calibration review, benchmarking — in isolation across thousands of institutions.
- A neutral, hardware-agnostic intelligence layer captures compounding value: every experiment executed improves the platform for every future user.
- An open-core commercial model aligns scientific progress with recurring enterprise revenue.

KEY TAKEAWAY Janus is designed to become foundational scientific infrastructure — beginning with quantum computing and extending, over time, to any discipline built on the iterative scientific method.

II. THE PROBLEM

Quantum Computing Has Entered the Commercial Era — But Its Software Infrastructure Has Not

Over the past decade, quantum computing has transitioned from a purely academic discipline into a rapidly developing commercial industry. Multiple organizations now provide cloud access to quantum hardware, including superconducting, trapped-ion, neutral-atom, and photonic quantum computers. Researchers, universities, and enterprises can execute quantum circuits remotely without owning specialized laboratory equipment.[1,2]

Despite these advances, the surrounding software ecosystem remains fragmented.

Unlike classical cloud computing — which has matured around common development practices, standardized tooling, and decades of operational knowledge — today's quantum computing landscape consists of multiple independent hardware providers, each exposing different programming frameworks, calibration methodologies, performance characteristics, and device architectures.[3]

As a result, organizations seeking to explore quantum computing face significant barriers before meaningful research can even begin.

Fragmentation Across the Quantum Ecosystem

Each quantum computing platform presents unique tradeoffs. Examples include:

- Different physical qubit technologies (superconducting, trapped ions, neutral atoms, photonics)
- Independent software development kits (SDKs)
- Different compiler stacks
- Different calibration procedures
- Different coherence times
- Different gate fidelities
- Different queue management systems
- Different pricing models

No single hardware architecture has emerged as universally optimal for every application. Instead, hardware performance depends heavily on the characteristics of the problem being solved.[4]

Consequently, researchers often spend substantial time evaluating infrastructure before performing meaningful scientific work.

The Cost of Fragmentation

Today, a researcher wishing to investigate a new quantum algorithm typically performs several independent tasks before executing the first experiment:

1. Review published literature.
2. Compare hardware providers.
3. Understand provider-specific software tools.
4. Evaluate calibration reports.
5. Select appropriate qubit layouts.
6. Implement provider-specific code.
7. Execute benchmark experiments.
8. Analyze performance manually.

Many of these steps are repeated independently by thousands of research groups around the world. This duplication represents a significant opportunity for automation.

III. THE MISSING INTELLIGENCE LAYER

Every Computing Revolution Has Required a Software Layer

Throughout the history of computing, transformative hardware has only reached broad adoption after complementary software infrastructure emerged.

The personal computer became widely accessible through graphical operating systems. The Internet became usable through browsers, standardized protocols, and search engines. Cloud computing accelerated through virtualization, orchestration platforms, APIs, and developer tooling.

In each case, the underlying hardware represented only part of the technological breakthrough. The software ecosystem ultimately determined how effectively people could utilize that hardware. Quantum computing is approaching a similar stage of development.

Hardware Has Advanced Faster Than the Supporting Infrastructure

Multiple organizations now provide commercial access to quantum hardware using fundamentally different physical architectures:[9,10,11,12]

- Superconducting qubits
- Trapped ions
- Neutral atoms
- Photonic quantum systems

While diversity of hardware is beneficial for innovation, it creates significant complexity for researchers and organizations attempting to determine which system is most appropriate for a given application.

Unlike classical cloud computing — where mature tooling abstracts away much of the underlying infrastructure — quantum researchers frequently interact directly with provider-specific software stacks.[13]

Today's Research Workflow

A researcher attempting to evaluate a quantum algorithm often performs work that extends well beyond the experiment itself:

- Reviewing current literature
- Comparing available quantum providers
- Understanding provider-specific SDKs
- Evaluating hardware calibration reports
- Benchmarking circuits across multiple platforms
- Measuring execution fidelity
- Comparing queue times and availability
- Manually documenting experimental results

Valuable scientific effort is frequently devoted to infrastructure evaluation rather than scientific discovery.

Case Study: Drug Discovery

The 2024 Nobel Prize in Chemistry highlighted the growing role of advanced computation in scientific discovery. David Baker was recognized for computational protein design, while Demis Hassabis and John Jumper were awarded for developing AlphaFold, an artificial intelligence system capable of accurately predicting protein structures. Their work has fundamentally changed structural biology and is accelerating drug discovery.[5,6]

As computational methods continue to advance, quantum computing is widely viewed as a promising future complement to these approaches. Certain molecular simulations — particularly those involving quantum mechanical interactions in complex chemical systems — are computationally expensive for classical computers and are among the long-term applications researchers hope quantum computers will eventually address.[7,8]

Consider a pharmaceutical company developing a new therapeutic. Before meaningful quantum research can begin, the team must answer questions such as:

- Which quantum hardware architecture best fits this workload?

- Which provider currently offers the highest effective fidelity?
- How have recent hardware calibrations affected performance?
- Has this algorithm already been benchmarked elsewhere?
- Which execution environment provides the highest probability of success?

Today, these answers are assembled manually through vendor documentation, published research, internal benchmarking, and repeated experimentation. There is no unified platform that continuously collects this information, benchmarks hardware across providers, and transforms experimental results into actionable recommendations.

KEY TAKEAWAY Janus is designed to become that intelligence layer.

IV. THE SOLUTION

Building the Intelligence Layer for Quantum Computing

Throughout the evolution of computing, the greatest leaps in capability have occurred when software abstracted the complexity of hardware. Operating systems transformed computers from specialized machines into general-purpose tools. Cloud platforms abstracted distributed infrastructure behind simple interfaces. Artificial intelligence has begun abstracting complex reasoning into accessible software systems.

Janus applies this same principle to quantum computing. Rather than developing new quantum hardware, Janus is designed to operate as an independent intelligence layer between researchers and quantum computers. Its purpose is to simplify experimentation, continuously evaluate hardware performance, and organize scientific knowledge into a living system that improves over time.

The Janus Architecture

At its core, Janus automates the scientific workflow. For every research objective, the platform:

- Interprets the research question.
- Designs candidate quantum experiments.
- Simulates expected outcomes locally.
- Executes experiments across available quantum hardware.
- Compares expected and observed behavior.
- Evaluates statistical confidence.
- Publishes reproducible results.
- Updates a continuously evolving knowledge graph.

- Uses accumulated knowledge to design increasingly effective future experiments.

Unlike traditional software, the objective is not simply to automate computation — it is to automate the iterative process of scientific learning.

Hardware-Agnostic by Design

The quantum hardware landscape is evolving rapidly. Different hardware providers continue to improve their systems, introduce new architectures, and optimize for different computational workloads.

Janus is intentionally designed to remain independent of these developments. Rather than competing with hardware manufacturers, the platform continuously evaluates and learns from all available systems. This architecture allows researchers to select the most appropriate hardware for a particular problem while allowing hardware providers to benefit from objective benchmarking and broader ecosystem visibility.

A Continuously Learning Scientific System

Traditional scientific software produces results. Janus produces knowledge.

Every completed experiment contributes new information regarding hardware performance, experimental reproducibility, algorithm behavior, and statistical confidence. This information is permanently integrated into the Janus Knowledge Graph, where it becomes immediately available to future researchers and future AI systems.

Unlike static software, Janus becomes more valuable through use.

Long-Term Vision

The architecture being developed for quantum computing is intentionally broader than quantum computing itself. Any scientific discipline following an iterative cycle of hypothesis, experimentation, analysis, and refinement can benefit from the same underlying framework. Over time, the Janus architecture may extend into:

- Computational chemistry
- Drug discovery
- Materials science
- Fusion energy
- Climate modeling
- Aerospace engineering
- Robotics

KEY TAKEAWAY Quantum computing represents the first implementation of a larger vision: building the intelligence layer for scientific discovery.

V. HOW JANUS WORKS

Automating the Scientific Method

At its core, Janus is designed to automate one of humanity's oldest and most successful processes: the scientific method. Scientific progress has traditionally followed an iterative cycle. Researchers begin with a question, formulate a hypothesis, design an experiment, collect observations, analyze the results, and use those findings to refine future experiments.

Janus automates this cycle while preserving the researcher as the decision-maker. Rather than replacing scientists, the platform continuously reduces the time spent on repetitive technical work, allowing researchers to focus on interpretation, creativity, and discovery.

The Janus Scientific Loop

STEP	STAGE	DESCRIPTION
1	**Research Question**	Every investigation begins with a scientific objective, such as evaluating hardware performance for a given algorithm or measuring decoherence effects on Bell-state fidelity.
2	**AI Experiment Design**	The Janus AI Engine generates candidate experiments informed by literature, previous experiments, hardware capabilities, and desired statistical confidence.
3	**Local Simulation**	Classical simulations detect implementation errors, validate circuit construction, and filter out low-value candidates before consuming quantum resources.
4	**Multi-Provider Execution**	Approved experiments are submitted to one or more providers — including IBM Quantum, IonQ, Rigetti, IQM, and Quantinuum — for objective cross-platform comparison.
5	**Experimental Analysis**	Observed results are compared against theoretical expectations and simulated predictions, quantifying fidelity, error rates, decoherence, variance, and hardware-specific anomalies.
6	**Confidence Scoring**	Every experiment receives a reproducibility score reflecting statistical significance, cross-provider agreement, and calibration quality.
7	**Publication**	Validated experiments — successful or not — are

STEP	STAGE	DESCRIPTION
		permanently recorded with structured metadata for reproducibility.
8	**Knowledge Graph**	Published experiments become interconnected nodes capturing circuits, algorithms, architectures, error sources, and related literature.
9	**Continuous Learning**	Accumulated knowledge improves the platform's ability to design future experiments, prioritize open questions, and match workloads to hardware.

Figure 1. The nine stages of the Janus autonomous research loop.

Beyond Automation

Janus is not intended to become an automated laboratory that replaces scientists. Its purpose is to become an intelligent research partner. Scientists continue to define important questions, evaluate surprising observations, and determine the direction of research. Janus handles the repetitive tasks of experiment generation, execution, benchmarking, organization, and continuous learning.

VI. THE JANUS INTERFACE

Making Quantum Computing Accessible

One of the greatest barriers to quantum computing is not hardware availability, but usability. Today's quantum ecosystem often requires researchers to navigate multiple software frameworks, provider-specific tools, command-line interfaces, and fragmented documentation before meaningful experimentation can begin.

Janus is designed to abstract this complexity behind a unified, intuitive interface that allows researchers to focus on scientific questions rather than technical infrastructure. Whether the user is a graduate student, enterprise research team, hardware provider, or national laboratory, the underlying workflow remains familiar, consistent, and easy to navigate.

Personalized for Every Organization

Janus is designed to adapt to the needs of its users through customizable dashboards, workflows, permissions, and integrations:

- Personalized research dashboards
- Custom experiment pipelines
- Organization-specific knowledge graphs

- Team collaboration workspaces
- Configurable alerts and reporting
- Enterprise branding and integrations

This flexibility allows Janus to scale from individual researchers to large enterprise and government research teams without fundamentally changing the user experience.

VII. WHO USES JANUS

Building Infrastructure for an Entire Quantum Ecosystem

The value of Janus increases with every participant that joins the platform. Unlike traditional software built for a single customer type, Janus is designed as shared scientific infrastructure that serves researchers, enterprises, hardware manufacturers, developers, universities, and government laboratories simultaneously.

Each participant contributes new experiments, data, and insights while benefiting from the continuously growing body of scientific knowledge generated by the platform. This creates a powerful network effect in which the platform becomes increasingly valuable as adoption grows.

Enterprise Organizations	Pharmaceuticals, financial services, energy, aerospace, advanced manufacturing, and defense teams evaluating optimization, molecular simulation, cryptography, logistics, and machine learning.
Universities & Research Institutions	Academic researchers seeking standardized workflows, AI-assisted experiment design, shared experiment libraries, and reproducible benchmarking.
Quantum Hardware Companies	Independent benchmarking, calibration monitoring, longitudinal performance analysis, and cross-provider comparisons for hardware manufacturers.
Government & National Laboratories	Reproducibility, transparency, and standardized benchmarking for organizations evaluating multiple quantum technologies simultaneously.
Developers	Standardized APIs and templates that lower the barrier to integrating quantum computing into existing software systems.

Table 1. Primary user segments and the value each derives from the Janus platform.

Why This Creates a Powerful Platform

Most software companies serve one type of customer. Janus connects multiple participants within the same ecosystem: researchers generate experiments, universities validate discoveries, hardware providers improve performance, enterprises apply scientific knowledge to commercial problems, developers build new applications, and governments establish standards and fund foundational research.

Every participant contributes new knowledge that increases the value of the platform for every other participant. Unlike traditional software, where value is created primarily through new product releases, Janus becomes more valuable through the continued growth of its scientific community.

VIII. COMMERCIAL MODEL & LONG-TERM OPPORTUNITY

Building Sustainable Scientific Infrastructure

Janus is founded on a simple principle: scientific knowledge should remain open. Historically, some of humanity's greatest technological advances have emerged from open scientific collaboration. Open protocols helped establish the Internet. Open-source software accelerated modern computing. Public scientific research continues to advance medicine, engineering, and physics.[14,15,16]

Janus adopts this philosophy by ensuring that the scientific knowledge generated by the platform remains broadly accessible, while enterprise organizations pay for advanced capabilities that improve productivity, collaboration, and operational efficiency.

Open Science Layer

The core scientific infrastructure of Janus is intended to remain publicly accessible:

- Public Experiment Database
- Knowledge Graph
- Scientific Publications
- Community Research
- Benchmark Reports
- Basic Dashboard
- Open APIs
- Educational Resources

Enterprise Platform

Janus Enterprise provides a secure commercial platform designed for research organizations operating at scale:

- AI Research Assistant
- Live Hardware Analytics
- Intelligent Provider Recommendations
- Enterprise APIs
- Private Knowledge Graphs
- Team Collaboration
- Workflow Automation
- Dedicated Technical Support
- Organization-specific Dashboards
- Private Experiment Management

Revenue Model

Channel	Description
Enterprise SaaS	Annual subscriptions to enterprise capabilities including collaboration tools, AI-assisted research, benchmarking, workflow automation, and private research environments.
Enterprise API	Programmatic access enabling laboratory software and internal research platforms to request hardware recommendations, execute experiments, and retrieve benchmarking data.
Professional Services	Hardware benchmarking, workflow optimization, integration, scientific consulting, custom dashboards, and AI workflow implementation.
Scientific Partnerships	Independent benchmark reports, hardware validation, standards development, collaborative research, technical publications, and longitudinal performance studies.

Table 2. Complementary revenue channels generating recurring and project-based income.

Why the Business Model Works

Open scientific resources attract researchers. Researchers contribute experiments. Experiments expand the Knowledge Graph. The Knowledge Graph improves AI recommendations. Improved recommendations increase enterprise value. Enterprise revenue funds continued development of the open scientific platform. This creates a self-reinforcing cycle in which commercial success directly supports scientific advancement.

Beyond Quantum Computing

While quantum computing serves as Janus' initial application, the underlying architecture is intentionally domain-independent. The Scientific Engine developed for quantum experimentation can ultimately be adapted to any research discipline following an iterative cycle of hypothesis, experimentation, analysis, and refinement. Potential future domains include:

- Computational Chemistry
- Drug Discovery
- Materials Science
- Fusion Energy
- Aerospace Engineering
- Climate Modeling
- Robotics
- Synthetic Biology
- Advanced Manufacturing

> **KEY TAKEAWAY** Janus aligns commercial incentives with scientific progress. By keeping foundational scientific knowledge open while delivering premium enterprise capabilities, the platform creates a sustainable business model that continuously reinvests in accelerating discovery.

IX. THE FIRST MILESTONE

Building the Foundation of the Janus Platform

The initial $75,000 investment is designed to accomplish a single objective: **demonstrate that autonomous, AI-assisted quantum experimentation can function as a practical software platform.**

Rather than attempting to build a complete enterprise product, this first round focuses on developing the core technology, validating the scientific workflow, and establishing the foundation upon which future capabilities can be built.

Allocation of Capital

ALLOCATION	CATEGORY	PURPOSE
32%	**Founder Execution**	Full-time dedication to product development, platform architecture, scientific research, fundraising, customer discovery, and strategic partnerships.

ALLOCATION	CATEGORY	PURPOSE
20%	**Platform Development**	First production-ready version of Janus Core: experiment engine, AI designer, database, knowledge graph, dashboard, authentication, public API, provider integrations.
13%	**Cloud Infrastructure**	AI inference, quantum simulation, experiment scheduling, database hosting, knowledge graph storage, API services, authentication, monitoring, and backups.
11%	**Quantum Hardware Access**	Cloud quantum credits, cross-provider benchmarking, experiment validation, hardware characterization, and replication studies.
9%	**Scientific Partnerships**	University outreach, technical advisory relationships, workshops, collaborative research, pilot programs, and conference participation.
8%	**Legal & Operations**	Corporate formation, legal agreements, accounting, IP strategy, insurance, regulatory compliance, and financial reporting.
7%	**Strategic Reserve**	Flexibility for unforeseen infrastructure costs, additional quantum hardware usage, specialized expertise, and pilot opportunities.
100%	**Total**	**Initial funding round: $75,000**

Table 3. Use of proceeds — initial $75,000 funding round.

Milestones Achieved

The success of this investment is measured through concrete deliverables rather than abstract progress. By the conclusion of this milestone, Janus expects to deliver:

Product

- Janus Core Platform
- Autonomous Experiment Engine
- AI Experiment Designer
- Public Dashboard
- Knowledge Graph
- Public API
- Experiment Database

Scientific Validation

- 10,000+ autonomous experiments
- Five or more connected quantum providers

- Independent benchmarking framework
- Initial technical white paper
- Reproducible experimental workflows

Commercial Readiness

- Design partnerships with research organizations
- Initial enterprise demonstrations
- Technical advisory board
- Investor-ready product platform

Why $75,000?

This funding round is intentionally modest. Its purpose is not to build the final company, but to answer the most important early questions:

- Can AI autonomously design meaningful quantum experiments?
- Can those experiments be executed across multiple providers?
- Can results be organized into a continuously improving knowledge graph?
- Will researchers find value in an intelligent software layer?
- Can the platform demonstrate a repeatable commercial model?

> **KEY TAKEAWAY** The first $75,000 is not simply funding software development — it funds the validation of a new scientific infrastructure model, establishing the foundation upon which Janus can scale into a platform serving researchers, enterprises, and scientific institutions worldwide.

X. EXECUTION ROADMAP

From Concept to a Functional Scientific Platform

The purpose of the first funding milestone is not to build every feature envisioned for Janus. Instead, it is to validate the core architecture upon which the platform will be built. Over the first six months, development is focused on proving that autonomous, AI-assisted scientific experimentation can operate as a practical and scalable software platform.

Building the Core Platform

The first milestone centers around the development of Janus Core, the foundation upon which all future products and scientific capabilities will be built. Development priorities include:

- A unified research interface

- Autonomous experiment generation
- Cross-provider quantum execution
- Local simulation capabilities
- Experiment storage and versioning
- A continuously evolving Knowledge Graph
- Public APIs for integration
- A real-time research dashboard

Scientific Validation

During this phase, Janus will execute thousands of autonomous experiments across multiple quantum computing providers. The target of **10,000 autonomous experiments** is not simply a usage metric — it represents the beginning of a continuously growing scientific dataset that increases the intelligence of the platform over time.

Measuring Success

The success of the first milestone is evaluated through demonstrable technical progress rather than financial metrics alone. By the conclusion of this phase, Janus expects to have demonstrated:

- A functioning autonomous experiment platform
- Cross-provider quantum execution
- AI-assisted experiment generation
- A publicly accessible experiment database
- A continuously expanding Knowledge Graph
- A reproducible benchmarking framework
- The first scientific publication
- Thousands of completed autonomous experiments

Looking Beyond the First Milestone

The first six months represent only the beginning of the Janus platform. Each capability developed during this phase becomes a permanent building block for future releases. Subsequent phases focus on expanding provider support, improving AI reasoning, increasing automation, and extending the Janus architecture beyond quantum computing into additional scientific domains.

> **KEY TAKEAWAY** The first six months are dedicated to validating the core architecture of Janus. By delivering a functional autonomous research platform, publishing scientific results, and building a continuously expanding knowledge base,

this milestone establishes the technical and scientific foundation upon which the broader Janus vision can be realized.

XI. REFERENCES

1. National Academies of Sciences, Engineering, and Medicine. Quantum Computing: Progress and Prospects. Washington, DC: The National Academies Press, 2019. https://doi.org/10.17226/25196
2. Gibney, E. "The quantum computing era is here — sort of." Nature 574, 22–24 (2019). https://doi.org/10.1038/d41586-019-02935-4
3. Preskill, J. "Quantum Computing in the NISQ era and beyond." Quantum 2, 79 (2018). https://doi.org/10.22331/q-2018-08-06-79
4. Blume-Kohout, R. & Young, K. C. "A volumetric framework for quantum computer benchmarks." Quantum 4, 362 (2020). https://doi.org/10.22331/q-2020-11-15-362
5. Royal Swedish Academy of Sciences. "The Nobel Prize in Chemistry 2024." Press Release, 9 October 2024. https://www.nobelprize.org/prizes/chemistry/2024/press-release/
6. Jumper, J., Evans, R., Pritzel, A. et al. "Highly accurate protein structure prediction with AlphaFold." Nature 596, 583–589 (2021). https://doi.org/10.1038/s41586-021-03819-2
7. Cao, Y., Romero, J., Olson, J. P. et al. "Quantum Chemistry in the Age of Quantum Computing." Chemical Reviews 119(19), 10856–10915 (2019). https://doi.org/10.1021/acs.chemrev.8b00803
8. McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. "Quantum computational chemistry." Reviews of Modern Physics 92, 015003 (2020). https://doi.org/10.1103/RevModPhys.92.015003
9. Arute, F., Arya, K., Babbush, R. et al. "Quantum supremacy using a programmable superconducting processor." Nature 574, 505–510 (2019). https://doi.org/10.1038/s41586-019-1666-5
10. Wright, K., Beck, K. M., Debnath, S. et al. "Benchmarking an 11-qubit quantum computer." Nature Communications 10, 5464 (2019). https://doi.org/10.1038/s41467-019-13534-2
11. Henriet, L., Beguin, L., Signoles, A. et al. "Quantum computing with neutral atoms." Quantum 4, 327 (2020). https://doi.org/10.22331/q-2020-09-21-327
12. Madsen, L. S., Laudenbach, F., Askarani, M. F. et al. "Quantum computational advantage with a programmable photonic processor." Nature 606, 75–81 (2022). https://doi.org/10.1038/s41586-022-04725-x
13. LaRose, R. "Overview and Comparison of Gate Level Quantum Software Platforms." Quantum 3, 130 (2019). https://doi.org/10.22331/q-2019-03-25-130
14. Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H. F. & Secret, A. "The World-Wide Web." Communications of the ACM 37(8), 76–82 (1994). https://doi.org/10.1145/179606.179671
15. Raymond, E. S. The Cathedral and the Bazaar. O'Reilly Media, 1999.

16. OECD. Making Open Science a Reality. OECD Science, Technology and Industry Policy Papers, No. 25 (2015). https://doi.org/10.1787/5jrs2f963zs1-en