ORCD Office of Research Computing and Data

OLYMPIA: A User-Centric Computing Platform for Research

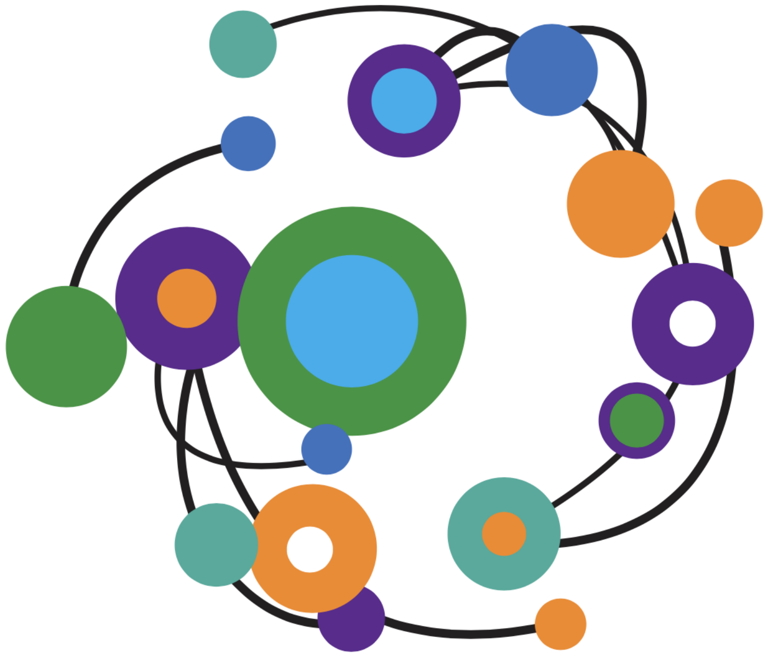

February 2024

Professor Peter Fisher, *Head*
James Cuff, *Executive Director*
Chris Hill, *Director of Scientific Consulting*

Office of Research Computing and Data
7th Floor, NE36, 105 Broadway Cambridge
https://orcd.mit.edu

Contents

February 2024

Executive Summary

The MIT Office of Research Computing and Data (ORCD) introduces Olympia, a model and overarching architecture to solve pressing and unmet high-performance computing needs for the MIT research community in the AI and machine learning (AI/ML) era.

Too often, MIT researchers who need to access and process complex data face multiple roadblocks. Many problems lie in a region too large for a personal computer and too small to warrant competing for a grant for access to large-scale national high-performance computing facilities. Olympia aims to bridge this gap by providing seamless, instant access to foundational computing resources and the ability to scale work with and share resources across departments quickly.

In other words: by deploying easy, frictionless access to virtual computer instances that are sized and pre-provisioned appropriately for modern research, Olympia will significantly reduce the time from ideas to innovation.

Olympia is a bold approach to this challenge, taking inspiration from MIT's highly influential historic computing environments: Athena (1983–91) and Project Mac (1963–76) (**Appendix I**). As a professionally staffed, continuously maintained, and supported core service for MIT located at the MIT-owned Massachusetts Green High-Performance Computing Center (MGHPCC), Olympia will provide an extensive, transformative central computing system that will bolster MIT's research and act as a model for other institutions.

We are seeking funding for this ambitious endeavor. An initial budget estimate has two phases: *Mens et Manus et Machina* (M^3) as the pilot project, followed by the sustaining project *Olympia*.

- **$15M** in seed funds to launch pilot project, Mens et Manus et Machina, over three years
- **$5M** in seed grants for MIT researchers to execute research projects that prove out M^3 and Olympia over the three- years course of M^3
- **$100M** to deliver and sustain Olympia over the coming five years
- **$20M/yr** to sustain Olympia after the end of M^3, in five years

Olympia: A New Approach to a Rising Challenge

If MIT is to continue to attract the most talented scientists who make world-changing discoveries, we must accelerate scientific discovery and make compute more accessible to a broader range of researchers. Today, the Institute includes over 1,080 faculty members and more than 16,000 employees supporting approximately 12,000 students (4,600 undergraduates and 7,200 graduates). The computing needs of this population increase every year, and MIT's computational and in silico research intensity has also scaled exponentially since the 1990s. Computing is an integral part of almost every discipline, and its influence continues to grow.

Currently, each computing system at MIT relies on various identity management systems that lack consistency and limit utility, use different cost models that can limit access, and do not engage in regular hardware renewal. Over the years, MIT principal investigators (PIs) and departments, labs, and centers (DLCs) have invested in additional on- and off-campus data centers. Cloud computing spending has also continued in concert with credits allocated by Amazon, Google, and Microsoft, each in the $1M to $2M range, consumed in an ad-hoc fashion. These decentralized investments have led to computational "haves" and "have nots" on campus throughout the last ten years of rapid advances in computing.

This situation must change if we are to continue to attract the best and most talented researchers to MIT. Furthermore, the rise of AI/ML workloads and sensitive datasets needing secure storage make the current situation even more severe. It requires investment to maintain global leadership in modern research computing.

These challenges are not unique to MIT. But with our existing computing resources, vibrant research landscape, long tradition of innovating "timeshare computing," and the leadership of the ORCD, we are uniquely positioned to tackle it.

Our response, Olympia, is a bold approach that will empower researchers to do their best work in a rapidly evolving computing landscape. Residing in the MGHPCC, the program will appear to the user like a laptop: always on, connected to allocated resources, and enabling interactive code development and execution, while having vastly more processing and storage capability.

Appendix II provides more technical details of Olympia's services, with the following qualities characterizing its distinctive approach:

- We aim to provide Olympia with a **unified software stack**: a tiered system of software libraries ranging from those utilities that do essential operations and communication with devices up through device drivers, system libraries, programming languages, and application software, all available to all users of the shared resource.

- Olympia will aim to provide an **economy of scale** that will make managing, servicing, and upgrading computers easier, ensure the maximum use of computing resources, and allow experiments by allowing the marshaling of significant resources for short periods, encouraging trying out new ideas.

- Olympia will ease **interdisciplinary collaboration** by allowing common software stacks and resources across different research groups, making sharing data and software possible.

- Olympia will allow **better cybersecurity** through centralized management of security resources.

- **Support from ORCD Research Software Engineers** (RSEs) will ease the transition for young researchers who need access to the latest tools, as well as for experienced researchers who will not have to leave the Olympia environment to find the resources they need.

- As a centralized project, Olympia will allow **better data management by providing security and open access,** both now required by our funding agencies.

- Every researcher will have a **resource allocation** and have access to their allotted resources immediately upon connecting with Project Olympia. They may also have access to unallocated resources, depending on their availability. DLCs will make allocations to their researchers on top of the minimum guarantee, which will be something like one or two GPUs, 50 CPUs, and 10 TB of disc space. DLCs will fund allocations through hardware purchases or other means and allocate to their researchers through a governance process unique to each DLC. ORCD will also seamlessly provide cloud access for a fee through Project Olympia.

All of these go together to create **a single unified computer access resource commitment** to all MIT researchers, enabling them to share and communicate seamlessly across MIT. At the same time, Project Olympia will preserve the autonomy of DLCs—whose decentralization has also been a strength—by allowing them to purchase fit-for-purpose hardware and software and determine how their researcher will access their resources.

Project Olympia must offer conventional parallel computing, high-performance computing, and GPU-based resources, in addition to secure storage. ORCD does not have clarity for what proportions of each best serve the MIT community, and we look to the M^3 pilot to inform us, which we describe below.

Using a Pilot to Technically and Scientifically De-Risk Project Olympia

The Olympia platform aims to provide a comprehensive cyberinfrastructure (skilled staff and technologies) developed to support all scientific research at MIT. Olympia requires the coordination of a wide range of resources and expertise that will allow researchers to access the tools and data they need to conduct their work more efficiently and effectively.

The Olympia team proposes a 10–15% pilot project over the next three years to achieve this goal. The pilot has a technical part, **Mens et Manus et Machina (M^3),** and a scientific part consisting of seed funds for research projects to inform on the development of Olympia. **Appendix III** describes the technical aspects of M^3, and here we describe the overall strategy and AI seed fund projects. This approach—technical and scientific[1] de-risking to prove the full Project Olympia—provides a new methodology for building university-based computing infrastructure. We have developed this method in collaboration with Schmidt Futures.

The M^3 pilot needs to answer two big questions for ORCD:

1. What system configuration—computing, network, storage, and system software, with RSE support—will best enable research at MIT?
2. What is the optimal technical implementation of the system configuration defined by the answer to the first question?

Technical de-risking will answer the second question via the M^3 pilot. M^3 will start in the first year with a 5% implementation of Project Olympia, and answers to those two questions will guide the development of M^3 over the subsequent years, following the well-established practice of technical de-risking.

ORCD will answer the first question by choosing four to five excellent AI seed projects from across MIT. The performance and manner of execution of these projects will answer our first question: What system configuration will best enable research at MIT? In addition to scientific excellence, we will choose projects that test the capabilities of M^3 and inform us on how to develop the whole Project Olympia to meet the research demands of the community.

Scientific de-risking constitutes a unique approach to this problem, and we are not aware of anyone else in the high-performance and academic hyper-performance computing business undertaking a scientific de-risking.

ORCD has listed critical activities that will provide clear direction for the pilot efforts. These activities focus on advancing science in vital areas, such as climate change, health care, and energy. They will also confirm the Olympia hypothesis that there is a core set of common computational cyberinfrastructure (skilled staff and technologies) that, if implemented and maintained thoughtfully, would accelerate all types of scientific endeavors.

[1] The research projects used for scientific de-risking will come from across MIT—from management to political science to chemistry to environmental engineering to architecture. The “science” in scientific de-risking refers to the scientific analysis of how a broad set of research topics fare in the pilot to inform the development of the full system.

In conjunction with Olympia open-source technical and process blueprints, these projects would provide end-to-end case studies of the ongoing impact of our proposed work. These studies will allow the Olympia team to refine the platform and ensure that it meets the researchers' needs.

Technical and Scientific Risks and De-Risking

Olympia presents a unique opportunity for MIT and a model for the delivery of research. Computing infrastructure de-risking using science means choosing research projects specifically for their scientific excellence and ability to stress or probe different novel parts of the new deployment—in addition to technical de-risking, which is relatively straightforward.

An example was in the 1950s when numerical analysis on electronic computers first began; the scientists who were using these computers to make numerical predictions faced the question of "How do you know it's right?" or "How do you know this job will run in a finite amount of time?" To answer these questions, they spent quite a bit of time mathematically proving that their algorithms did what they were supposed to do and ran a series of test jobs to time out how long execution would take.

In the 2020s, AI and machine learning present a similar opportunity to numerical analysis in the 1950s: providing dramatically different new tools of unknown power and utility to solve research problems. AI and ML provide universal approximation functions—the ability to model almost any function with ample parameter space. These models do many new things. For example, they can produce human-like text and drawings, prove mathematical theorems by writing code, and rapidly explore massive parameter spaces of models. Just like the numerical model in the 1950s, we are asked, "How do you know this is right?" and "How long will this take to run?" and "What problems can this solve?" The de-risking presented by the M^3 pilot will allow us to begin asking these questions and see how the answers change with the scale of the machine.

The AI seed fund in this proposal will encompass four or five crucial problems across the research field group, size, and computing demands to stress M^3 and see where it fails. It will inform the next step in developing Project Olympia from a scientific standpoint, as the pilot informs other technical aspects.

In addition to the AI seed funds, ORCD will also offer 5–10% of our resources as allocations—called AI Starters—awarded on a competitive basis to our younger researchers: undergraduates, graduates, and postdocs. An allocation of computing resources for a relatively short time—days to a week—will allow them to try their ideas quickly. We would assess requests on a rolling basis and give further resources to successful experiments.

We plan to use both the AI seed funds and AI Starters to explore a new opportunity presented by the widespread availability of AI tools: exploring new workflow apologies. The left panel of Figure 1 shows a traditional HPC workflow from High Energy Physics: previous data and theory inform a simulation program that makes predictions for the results of new experiments. The incorporation of new data into the simulation requires extensive physics interpretation and coding. The results of a new experiment further inform the model for subsequent experiments in coming decades. The comparison of prediction and experiment may use AI as a powerful tool but does not require AI.

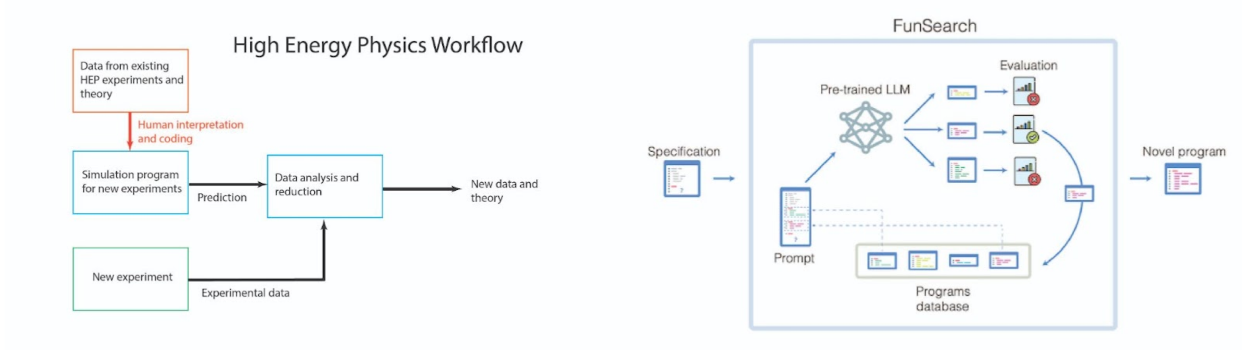

Figure 1—Left: Typical workflow for high energy physics. Right: Workflow for DeepMind FunSearch.[2]

The Google DeepMind FunSearch project uses a workflow shown on the right of Figure 1 and has three striking differences from the previous example:

- The AI produced executable code as output that solves the problem at hand.
- Benchmarking of the code against a program library takes place in the evaluation step, integrating previous coding examples into the next iteration of the code
- The benchmark results in a prompt to the AI element, leading to improved code.

The topology of the workflow aims to carry out the work of improving the code based on a vast body of experience embodied in the program library. In this way, the topology of the workflow merges the code library with the AI output. The DeepMind FunSearch project serves as one example of a new workflow topology.

FunSearch also serves as an example of another important feature of AI-driven research: the way AI merges datasets in specific ways to produce new datasets. Abstraction, encoders, and tokenization allow AI to merge datasets without human intervention, making iterative processes like FunSearch possible. We hope the AI seed funds and AI Starters will encourage this kind of experimentation.

The value proposition for Project Olympia is that ORCD can deliver the computing and storage resources needed to support MIT's research community by creating a large, shared computing and storage facility and providing research software engineering support. ORCD must accomplish this goal financially and environmentally sustainably through careful monitoring and allocating all resources.

The following enumerates our value proposition, the risks, and how M^3 informs those risks from a scientific and technical perspective:

1. ORCD proposes a $100M shared, scalable cluster that gives users access to various computing and storage resources.

[2] https://deepmind.google/discover/blog/funsearch-making-new-discoveries-in-mathematical-sciences-using-large-language-models/, https://www.nytimes.com/2023/07/02/science/ai-mathematics-machine-learning.html, https://arxiv.org/pdf/2302.12433.pdf

a. **Risk 1:** Project Olympia will not scale how we think it will, even though we believe we have a conservative design. **How the** M³ **Pilot informs us: Technical**—The M³ proposal will double the size of our current cluster over three years. Over that time, continued renewal will effect a second doubling. Project Olympia will double again in three years unless we find problems in M³. M³ will show us early enough to adjust our architecture to counter. **Scientific**—We will learn which projects become tractable by doubling the system size and if scientific progress is non-linear in system size. **Metrics: Technical**—Overall usage fraction as a function of cluster size, number of users as a function of cluster size, peak and average latency to all storage elements. Peak and average latency between MGHPCC and campus. **Scientific**—Size of different models as a function of system size.

b. **Risk 2:** Seamless operation of Project Olympia relies on a unified software stack across MIT, low latency networking within the cluster, and the ability to handle the required data volume. Can Olympia adequately meet these requirements? **How the** M³ **Pilot informs us: Technical**—as M³ grows the cluster and the researcher base grows, we will learn if we can maintain an adequate software stack that meets user needs. Similarly, the ORCD staff will see the weak points of the network and storage systems early enough to adjust for the whole project. **Scientific**—Does the software stack grow as the system scales? We will see what must be added to the stack to ensure operation as the system doubles. **Metrics:** The number of elements in the software stack as a function of cluster size and the number of users and requests for additional packages. Number of non-stack installs.

2. ORCD supports Project Olympia with a small staff supported by the MIT General Institute Budget (GIB).
   a. **Risk 1:** The staff the GIB finances can support cannot adequately operate and maintain the system to meet researcher needs. **How the** M³ **Pilot informs us: Technical/Scientific**—M³ and usual renewal will instigate a factor of three scale over the first three years, and the ORCD staff will have to install, maintain, and operate the cluster as it expands. The ability of the RSEs to support the PIs will determine what kinds and scale of projects we can help as the cluster grows. **Metrics: Technical/Scientific**—Number and time to resolution of Service Now tickets as a function cluster size will inform the adequacy of the staff size, size, and diversity of projects successfully undertaken.
   b. **Risk 2:** The ORCD RSE staff will help researchers use ORCD systems in an ever-expanding and changing research world and may be unable to keep up. **How the** M³ **Pilot informs us: Technical/Scientific**—see Risk 2.1 above. **Metrics:** See Metrics 2.1 above, as well as user feedback and usage records.[3]

3. A sustainable funding model provides Project Olympia's needed renewal and growth. Our funding model requires renewal at the 20%/yr level, with half coming from annual renewal by PIs and half from other sources: philanthropy, ORCD proposals, support from MIT schools, or DLCs.
   a. **Risk 1:** PI equipment renewal is needed to meet the need of $10M/y (10%/y of Project Olympia), which does not take place. PIs may choose to develop co-lo systems or need

[3] "Usage records" includes job statistics and analysis of the interactive development sessions, i.e. number and duration of run/debug/edit/compite-build cycles.

to bring in more funds to support Project Olympia. **How the** M³ **Pilot informs us**—**Technical hardware** provided by M³ will double our hardware base in the next few years, creating enough incentive for PIs to locate their hardware in Project Olympia. **Scientific**—Will PIs see Olympia as the place for carrying out their work rather than building up their clusters? **Metrics: Technical/Scientific**—purchase records, hosting arrangements.

b. **Risk 2:** Other sources (philanthropy, grants, support from the schools and the college) do not provide the needed funds to support 10% annual renewal. **How the** M³ **Pilot informs us: Technical/Scientific**—See Risk 1: Metrics-Technical**/Scientific** fundraising, grant success, and school support.

4. Researchers without their hardware allocation (i.e., those who rely exclusively on shared resources) can access sufficient resources to carry out their research.

   Risk: The availability of unallocated resources scales worse than inversely with the number of users, preventing researchers from amassing what they need. **How the** M³ **pilot informs us: Technical**—Delivering significant unallocated computing resources to researchers will be a part of M³, and we will learn how to do so at a lower scale at the start. We can then see how access scales through M³ and the annual renewal cycle. **Scientific**—We will see if PIs can garner ample resources for their projects without an allocation and how this ability changes with a factor of two scaling. **Metrics: Technical**—Turnaround time on large jobs, user reporting. **Scientific**—Changes in the size and nature of projects that rely on unallocated resources.

5. Established researchers who have allocated hardware can access that hardware to carry out their funded research projects.

   Risk: Technical challenges in interfacing with the shared elements of Project Olympia prevent efficient use of allocated resources by their owners. **How the** M³ **pilot informs us: Technical**—Delivering access to PI-owned resources constitutes a part of M³ at a small scale, enabling us to understand how PI access changes with system scale. **Scientific**—ize and nature of projects undertaken with allocated resources and how they scale **Metrics: Technical**—Monitoring and user reporting. **Scientific**—The scale and kinds of projects PIs will understand in M³ as the cluster grows.

6. ORCD must maintain computing and storage infrastructure and control costs well enough to allow needed resource allocation and planning,

   Risk: ORCD still needs to understand the cost of operating Project Olympia and adequate monitoring; thus, our planning and budgeting are estimates rather than measurement-based facts. **How the** M³ **pilot informs us: Technical** development, the needed cost understanding, and planning are a part of the M³ pilot. **Scientific**—Insufficient funds for computing storage and infrastructure will limit the size and scale of projects PI can undertake. **Metrics: Technical**—Quarterly roll-up of infrastructure and power costs, alignment with the budget, and alignment of funding with needs. **Scientific**—The scale and kinds of projects PIs will understand in M³ as the cluster grows.

7. MGHPCC and the High Performance Research Computing Facility (HPRCF)[4] can meet ORCD's hosting needs for the foreseeable future. The recent buildout of 66 x 50kW racks at MGHPCC should be sufficient to house Project Olympia and meet existing hosting commitments. We have a scheme that will free space in HPRCF for co-lo hosting.
   a. **Risk 1:** We need more space in MGHPCC to realize Project Olympia. Running out of space could happen if ORCD has too many co-lo agreements for MGHPCC before resources for Project Olympia show up. ORCD needs support from the Administration to insist the co-lo systems go into HPRCF, which we still need to get. ORCD is in the process of freeing up space in HPRCF to garner the required Administration support. **How the** M³ **Pilot informs us: Technical**—The hardware purchased with M³ support and the AI seed funds will anchor shared space in MGHPCC. **Metrics: Technical—**Space and power monitoring.
   b. **Risk 2:** We may run out of co-lo space in HPRCF, forcing us to host co-lo systems in MGHPCC. We could run out of space in HPRCF if our current mitigation plans for HPRCF fall through. **How the** M³ **Pilot informs us: Technical**—see Risk 1. **Metrics: Technical**—see Risk 1.

8. ORCD, deans in the schools and the Schwarzman College of Computing, and the administration can govern the use of Project Olympia in a way that equitably meets the needs of the MIT research community.

   Risk: Aside from some vague notions, we still need an outline for the governance of the shared system, primarily because we need more to govern. **How the** M³ **Pilot informs us: Technical**—With the appearance of resources from M³ and renewal, building to about 20% of Project Olympia in the coming years, ORCD and the administration will need to develop and adjust the required governance in a way that scales to the whole Project Olympia. **Scientific**—The ability of Olympia governance to match the requirements of scientifically excellent projects to Olympia resources. **Metrics: Technical/Scientific**—Does the research community view the governance process as fair and just, as evidenced by requests to the governance system? Also, usage and user experience.

9. ORCD RSEs will work with PIs and their groups to deliver efficient and robust workflows for the PIs' projects. The RSE staff will grow to five full-time employees in the coming year, supported by student workers, consultants, and partner consultants in the research groups.

   Risk: A team of this size will need help to keep up with user demand for support. **How the** M³ **Pilot informs us: Technical**—M³ contains funds for additional RSE support, and ORCD will adjust RSE resources to meet user demand. Finding the right level of support will tell us how to scale from RSE support from M³ to the whole project in Olympia. Fundraising may need to focus on RSE support as well as hardware. **Science** The demands of PI projects enabled by seed funds will stress how well the research groups' domain experts (postdocs, graduate students) can interface with the ORCD RSE staff. **Metrics: Technical/Scientific**—Do researchers continue to come to ORCD for help, or do they find other means? Also, user experience, usage, and ticket responses.

10. State-of-the-art technology stacks for AI/ML are rapidly evolving in every aspect, from hardware instruction sets to algorithms to applications and software tools.

---

[4] Located at the Bates Research and Engineering Laboratory in Middleton, MA—25 miles from MIT—the HPRCF is a secondary data center that mostly hosts computers for the MIT Laboratory for Nuclear Science.

Risk: Can deployed platforms adapt quickly enough to keep up with the expected level of continuous disruptive technology change? **How the M³ Pilot informs us: Technical**—Over the M³ period, we will refresh hardware platforms and continually evolve a validated standard software stack. We will also create accompanying scalable open-training, use-case, and playbook materials for operations and research use. **Scientific**—It is unlikely that current hardware, tools, algorithms, and application areas will remain anywhere near static over the next few years. A successful platform will need to provide rapid access to cutting-edge developments in all areas to keep scientific research impacts at the forefront. **Metrics:** Regular refresh cycles for hardware technologies, platform deployment agility to change and continuous testing capacity to support changing technology, and elapsed time to research results for newly introduced technologies.

Decision-making for Project Olympia requires collection and analysis of data defined by the metric we have given for each risk. The different risks have different time scales.

- **Risk 1, 2, 3, 6, 7, and 8**—all having to do with scaling, budgets, support, and hosting—will manifest themselves over six or more months, with straightforward collection of the metrics data.
- **Risk 4 and 5**—related to user success—will resolve in the first few months of operation and require interaction with our users community. The Committee on Research Computing and Data—the CRCD—has the task of collecting this data and ORCD leadership will also meet with key researchers. Collection of what projects PIs undertake and what problem they face requires both an artisanal approach—individual conversations with PIs—and an industrial approach of analysis of ServiceNow tickets.
- **Risk 9**—the ability of ORCD consultants to meet user needs—needs continuous monitoring and adjustment starting from the first week. This risk represents our greatest unknown.
- **Risk 10**—demands for a new tool—also presents a day 1 concern. We estimate that 80% of our researchers will remain in evolving but well-developed workflows using known tool sets. The remaining 20% must forge ahead into new areas and we will have to keep up. The data collection will come form ServiceNow tickets and interactions through our RSEs.

Using science to de-risk the pilot, Olympia and M³ present a unique opportunity for MIT and a model for delivering research computing. Computing infrastructure de-risking using research means choosing projects for their excellence and explicitly for their ability to stress or probe different novel parts of the new deployment.

The AI seed fund part of this proposal will encompass four or five crucial problems across the research field group, size, and computing demands to stress M³ and see where it fails. Analysis of the overall use of M³ will inform us, but the seed projects will act as markers for specific aspects of the pilot.

Olympia and ORCD: The Big Picture

Many top research establishments across the nation need help with supporting advanced computation. MIT, with its long and esteemed history of high-quality engineering and computer design, is uniquely positioned to lead the charge to address this problem. A research-focused program like Olympia must embrace experimentation and a fearless approach to varied possible solutions.

Project Olympia forms a part of the overall ORCD effort to improve the delivery of research computing at MIT. Project Olympia and M^3 provide the centerpiece, but are by no means the only important aspect of ORCD, and many more anodyne activities occur.

ORCD works hard to rationalize data storage and computer usage across the enterprise and develop a hosting scheme that puts computers where they are best able to perform. We plan to have a large shared facility in MGHPCC, a state-of-the-art green computer center where our assets have central management. Some PIs will always want to operate their machines, and we plan to locate these in the HPRCF, where they will be somewhat closer to MIT, and PIs will be allowed easy direct access. These machines will run individually, and PIs will be responsible for software and hardware installation, maintenance, and renewal.

A second significant aspect of ORCD is providing RSEs across the enterprise. The RSEs will be experts in the operation of orchid systems and work with domain experts (postdocs or graduate students) in the research groups to help them get started using and support them using orchid resources, efficiently and meaningfully.

ORCD operates in close partnership with MIT Information Systems and Technology, and we are working with them to improve our data storage cybersecurity and data availability. We also partner with the MIT libraries in data storage and availability.

To make this a reality, MIT must take the following actions:

- **Engage MIT-centric leadership and establish evergreen funding.** Athena and Project MAC needed a base set of funding available. Once industry support ran out, so did each project. We need an enduring framework for modern, highly functional, and professionally supported research computing for MIT researchers.

- **Strengthen partnerships with experts and industry.** MIT is a world leader in computing and research, working with thousands of industry partners in varying capacities. We are confident that the only challenge will be finding the right partners whose goals and expertise align with the project. Best-of-breed vendors, partners, and solutions architects will be selected to help ORCD build up a world-class, exemplary resource that brings together the highest computational capacity, availability, and reliability standards.

- **Develop and leverage measurable use cases.** We live in an AI "summer," one may even say a "heat wave," where research and general audiences alike are eager to engage with new AI tools. The rise of accelerated low-precision computing at scale for deep learning and large language models has disrupted the industry like never before. The shortcomings around highly accessible programs for generative AI have demonstrated a need for improved systems and paved the way for a fertile test ground. We propose a system where research faculty, staff, and students have unencumbered, frictionless access to these accelerated computing devices and the required high-performance networking and parallel storage with all the needed software and expert support to build the next world-changing technology.

- **Launch pilot projects to deliver a high-quality, sustainable platform.** Systems that accommodate thousands of full-time research scientists at scale to provide interactive and just-in-time computing require a blend of interactive and batch-processing approaches for modern research computing. We propose a “crawl,” “walk,” and “run” approach that begins with initial pilot systems dedicated to proving the technology, ensuring adoption, and building with the community in close concert. We must continue to invest carefully and consistently to leverage all avenues of funding but funnel the resources into one tightly orchestrated overarching system—what will eventually become Olympia.

Summary

MIT faces a significant challenge to supporting modern research computing at the new and unprecedented scales needed, primarily due to modern accelerated high-performance computing advances and ever-more power-hungry computing devices, especially general-purpose graphics processors for AI research. We seek funding to launch the pilot phase of Olympia—M^3—and sustain the tool for its first five years as we work to implement this platform fully and meet the rapidly evolving needs of generative AI research and more.

— END —

Appendix I

Athena and Project MAC

We start by looking back at MIT's history to the original mission of Project Athena in 1983, initially labeled as an experimental undertaking by Digital Equipment Corporation and International Business Machines. Each company initially provided $50M in hardware, maintenance, and expertise, with MIT providing faculty, students, and technical staff and $20M in development support. The project eventually cost approximately $100M over eight years (equivalent to roughly $300M today).

"Time-sharing on computers" wasn't a new concept in 1983; MIT's time-sharing Project MAC[5] predates Project Athena by about 20 years. Project Athena initially sought to determine if a network of "high-performance" computer workstations could serve undergraduates and enhance their work. Early skeptics of the proposal were concerned: could one campus system ever serve the diverse needs of everyone on campus—from aeronautical engineers to students of Spanish? Could the assorted products be integrated into a single system so the differences were invisible to users? The questions and naysayers were endless, from concerns about administrative overreach to simple underlying logistical issues of where to locate the various "Athena terminals." Discussions around Project MAC were similarly full of questioning and uncertainty.

Despite these concerns, by 1987, many considered Project Athena a massive and unquestionable success. Arguably, it helped catalyze the deskside workstation era on a broad commercial scale. Project Athena officially ended in 1991, but more than 30 years later, the Athena computing environment is still part of everyday life at MIT. Many technologies evolved from these investments, including the X Window System, Kerberos, network file systems, and software systems still in use today—an impact that could not have been predicted when Project Athena began.

[5] Project on Mathematics and Computation.

Appendix II

The Olympia Accelerant: A Partial List of Services

Olympia brings a core set of ready-to-use cyberinfrastructure services to all these projects that can accelerate progress. These services are broad in scope but standard across many research areas. A partial list of services follows.

- Datasets and curation
  - Fast, efficient storage.
  - Expertise in leveraging large datasets to train and evaluate models.
  - Expertise in techniques for curating and cleaning datasets.
  - Expertise in maintaining tools for managing and sharing datasets.
- Performance engineering training and skills:
  - Expertise in engineering models for performance and scaling.
  - Tools and techniques for measuring and improving model performance.
- Tool expertise:
  - Available expertise in various AI tools, including deep learning and natural language processing.
  - Available expertise in tools for data science, including data wrangling, cleaning, and visualization.
- Robust integrated, optimized platform:
  - A robust platform for training and deploying AI models.
  - Optimized for performance and scalability.
- Automation in software and the lab:
  - Automate tasks in the software development process, such as code generation, testing, and deployment.
  - Automate tasks in the lab, such as data collection, preprocessing, and model training.
- Training:
  - Teaching researchers about training AI models.
  - Supporting tools and workflows for training AI models.
- Inference:
  - Develop techniques for deploying and using AI models in production.
  - Develop tools for deploying and using AI models.
- Synthetic data generation, curation, ingestion:
  - Generate synthetic data to train AI models.
  - Curate and ingest real-world and external synthetic data to train AI models.
- Generative AI models:
  - Expertise and open-source research-ready examples of generative AI models, such as GANs and VAEs.
  - Associated open-source training on the limits and biases of generative AI models.

- Clustering models:
    - Expertise and open-source research-ready examples of clustering models, such as k-means and DBSCAN.
    - Associated open-source training.
- Feature detection models:
    - Expertise and open-source research-ready examples of feature detection models, such as SURF and SIFT.
    - Associated training.
- Language models:
    - Expertise and open-source research-ready examples around language models, such as BERT, GPT-3, and LLAMA.
- Parallelism and distributed processing:
    - Expertise and research-ready examples of techniques for training and deploying AI models in parallel.
    - Expertise in research-ready tools for training and deploying AI models in distributed environments.
- Equation discovery:
    - Expertise in and research-ready examples of techniques for discovering equations from data.
- Explainability:
    - Expertise in and research-ready examples of techniques for explaining the predictions of AI models.
- Reproducible digital research:
    - Expertise in and research-ready examples of practices for open-source reproducible digital research.
- CI/CD and software sharing:
    - Use CI/CD to automate the software development process.
    - Training and support for sharing research software with others using modern practices in version control systems and package managers.
- Ensemble workflows:
    - Expertise in ensemble workflows for automating distributed data generations, data processing, training, and combining the predictions of multiple AI models.
- PDEs:
    - Expertise in and well-documented research-ready examples of methods for solving partial differential equations (PDEs) using AI.
- Bioinformatics:
    - Expertise in and well-documented research-ready examples of methods for using AI in bioinformatics, such as protein structure prediction and drug discovery.

Appendix III

Pilot Project Overview: “Mens et Manus et Machina”

This section outlines the pilot project phase for Olympia. We refer to this pilot project as “Mens et Manus et Machina” (M^3). The goal of the pilot is to demonstrate research impacts and deploy a first-generation architecture that will validate the subsequent Olympia strategy.

The pilot project will be sizable in itself. It will build an exemplary artificial intelligence/machine learning (AI/ML) cyberinfrastructure facility for MIT and identify a sustained operating model. It will immediately produce a needed supercharging of applied AI/ML research at MIT. It will develop a foundational team of people with applied research computing expertise needed to accelerate AI/ML work and beyond at scale.

The design is formulated to be adaptable to future technology and software stack changes. It combines hardware, software, and people to form a whole. It is suited to replication by other organizations.

This appendix outlines the elemental building blocks of the “Mens et Manus et Machina” project, describing what will be built and how the various building-block elements fit together into a whole. The pilot project budget and timeline are summarized in this part.

Pilot project elements

Mens et Manus et Machina is a $15M pilot project to create a first-of-its-kind open platform for at-scale AI/ML in academic research that is fit for purpose and fit for the current era in which powerful generative AI/ML capabilities and practices are evolving at a rapid pace; are starting to positively transform nearly all areas of leading universities’ research and education missions and practices; and are being simultaneously developed and applied, in work that spans all the way from computer hardware systems innovation to applied discipline research.

The pilot project will deploy a complete top-to-bottom architecture and allow MIT’s research community to access more computing power than it currently can, although it is not as much computing power as Olympia will ultimately provide. The project architecture contains all the elements that must be combined to create an effective, accelerated AI/ML research and applications activity. It will provide, demonstrate, and publish a blueprint that can be widely adopted, scaled, and replicated. The proposed design is based on conversations with current AI/ML researchers at MIT, highlighting desired research workflows that current systems do not fully support.

The three key project elements are:

1. **Staff.** Professional staff will support and facilitate efficient use and scaling of researcher workflows and accelerate research progress.
2. **Software.** A comprehensive curated and integrated open software stack must be maintained and evolved as part of the project.

3. **Hardware**. An integrated hardware system and operations activity must be designed to enable iterative research in any discipline. The design will span from development and prototyping to at-scale deployments.

Figure 1 shows a breakdown of these key project elements into major subparts. The figure illustrates how these form a comprehensive system that is fit for research purposes. The subparts form an overall "Scalable and Adaptable AI/ML Platform" that can adjust, scale, refresh, and grow as AI/ML technology and the user landscape evolves.

The implementation is constructed so researchers can iterate between working on experimental ideas at the small scale and scaling up, when promising results and code warrant, with minimal points of friction. The design, software stack, and operations are intended to support both open research and restricted research that is subject to various commonly employed, controlled unclassified data use standards. The overall system design incorporates dialogue with practitioners of applied AI/ML at MIT. The result is a comprehensive virtual space for applied AI/ML research across MIT.

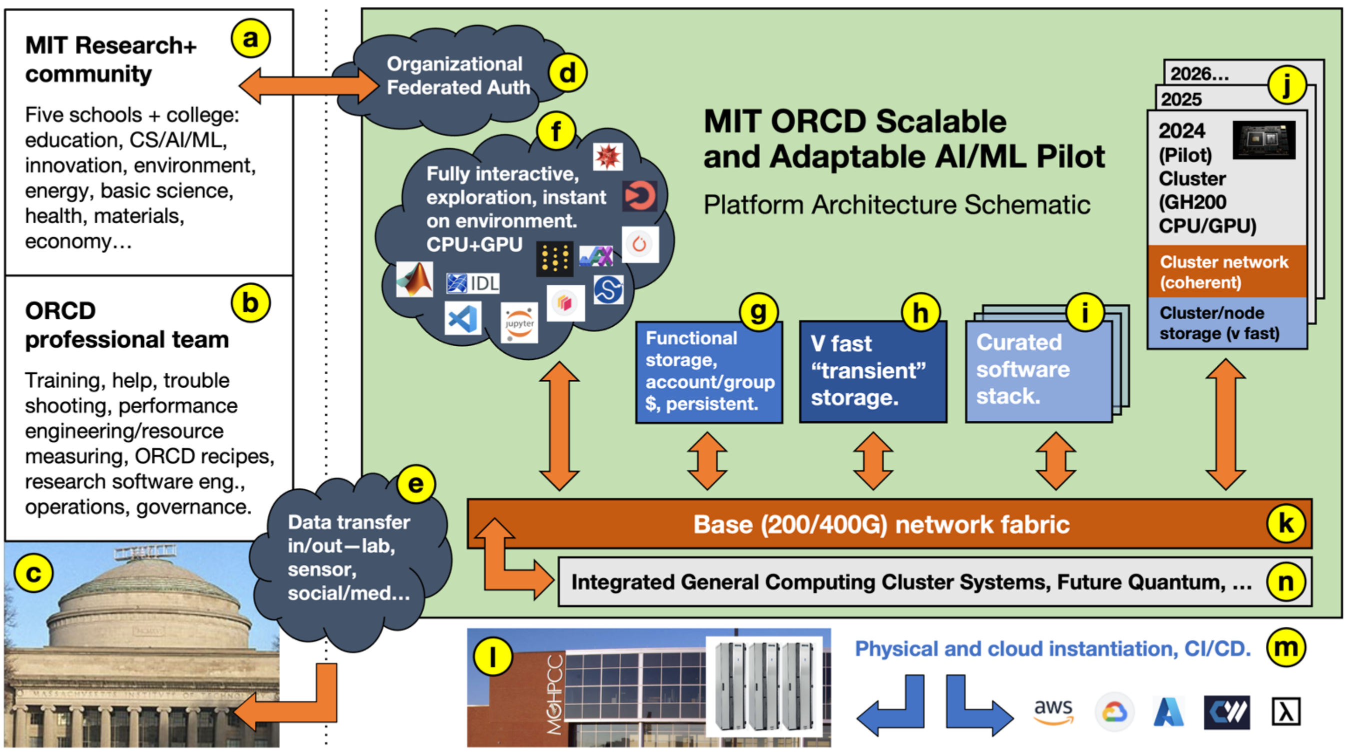

Figure 1: *Schematic of key project pieces and related activities in the plan.*

The project plan diagram in **Figure 1** provides an overview of the subparts of the project. These are labeled (**a**) through (**n**). All map in some way to the key project elements (1) through (3) (**Staff**, **Software Stack**, and **Hardware**). The subparts form a complete whole. The whole is a powerful combined human expertise and technology "meta-instrument" that we are seeking funds to develop and deploy.

The subparts are described on the following pages.

(a) The research community that drives innovation.

This consists of all undergraduate, graduate, postdoctoral, research staff, faculty, and principal investigators from all departments at MIT. We strongly believe that these technologies can and should play a role in every field. A piece of our overall vision is proactive engagement of communities across all of MIT in every department, laboratory, and center. Every department at MIT is finding growing interest in novel applications of AI/ML technology to its field. A goal of this proposal is to create an at-scale resource for grand challenge research that is paired with an "inalienable right" to a base level of AI/ML computation for all researchers at MIT.

The M^3 phase would initially work with key subsets of the MIT community but grow to working with the entire community over time. Projects from the research community would drive the pilot project development, and each project would feed into the creation of a general architecture.

(b) The MIT Office of Research and Computing professional staff team.

The Office of Research Computing and Data (ORCD) is a new Institute-wide initiative and focal point for all applied research computing activity on campus. ORCD provides a center of gravity for bridging between domain research and computational technologies at every level. The ORCD core staff has deep experience in the development of top-to-bottom research computing platforms. They have strong practical knowledge of areas including technical operations of hardware; software stack maintenance; training; and governance, finance, and fundraising. ORCD receives base funding from MIT.

This proposed pilot project would fund additional professional staff members and support a new program of mentored research assistantships/fellowships for graduate students.

As much as possible, we anticipate recruiting the additional professional software engineering staff from MIT's recent graduate student population. We would develop a program around excellence in research software engineering designed to attract the strongest and most capable candidates from across the MIT talent pool and from elsewhere as appropriate. We plan to augment the team with current students (including undergraduates) who we will train in state-of-the-art research computing practices for AI/ML. For this, we envision an approach in the same spirit as the student operations experiential learning system at the MIT Nuclear Reactor Laboratory. This program trains students to take on serious roles and has clear expectations of professionalism. The students work with full-time professional staff. For students who take part, this sort of experiential program can be a career-defining experience.

(c) Laboratories, instruments, national and international collaborations that the MIT faculty and PIs operate and that are connected to large quantities of digital information in every domain.

Data sources are vast. They span everything from electron microscopes in laboratories around campus, to environmental sensors in space, to fusion reactor experiments, to social, financial, and health data and much, much more. These data, from strategic areas such as climate, brain science, materials research, and bioengineering, as well as from projects across MIT, would provide the training data and models that can be incorporated into the proposed M^3 system.

(d) A computational gateway that seamlessly integrates with appropriate institutional cybersecurity and enterprise authentication systems.

The system we are proposing to build would be for open research but would support adequate access controls to ensure compliance with research sponsor privacy and audit requirements. To implement this, the facility would fully integrate with federated authentication services.

This proposal would integrate an open recipe that can easily be replicated and that supports seamless web, desktop application, and terminal access to AI/ML hardware and data resources.

(e) A professionally operated, secure, high-bandwidth ingress and egress network fabric.

This would be a TCP/IP-based network fabric monitored by performance and security tracking nodes. The hardware implementation would leverage MIT-owned dark fiber that connects to peering points in New York. This fiber network provides multiple wide area 100Gb/s links between MIT campus, the MIT data-center this project leverages, and commercial cloud provider peering points. Productivity and automation software for data transfer, such as Globus and rclone, would be configured to support data transfers at scale.

A pilot of a dedicated data transfer node network based on open technologies would be part of the proposed project. It would provide a complete recipe that can be replicated elsewhere.

(f) "On-ramp"/gateway resources that provide rapid and easy access for development workflows and that are integrated with larger-scale resources for research that meets scaling criteria.

The on-ramp resources are envisioned to be open to the entire research community on MIT's campus and to provide a base level of service at no direct charge. They would be the main point of entry to the system for students, researchers, and faculty and would provide web portal, desktop application, and terminal access. Workflows for proactive AI/ML experimentation would be supported and kept up to date by default on these resources. A researcher accessing on-ramp resources would see a private, persistent environment that can be customized to their needs. A default software stack would provide support for common tools such as vscode, pytorch, tensorboard, gradio, and more. Access options would support entirely web-based approaches, as well as simple terminal and graphical windowing approaches. This gateway resource would leverage the ORCD team's extensive prior experience with portals.

The pilot project would support the development of gateway resources scalable to at least 25% of the MIT research community at any one time. It would demonstrate the impacts of a unifying open environment for both research productivity and cost-effective operations. The design would scale so that, subsequent to the pilot, resources could be added to meet needs of 100% of the MIT research community.

(g) A base, long-term, bulk storage infrastructure that has high capacity and performance fit for serving terabyte-sized data collections and beyond.

This storage infrastructure would be based on a fully documented open architecture that can be replicated elsewhere. Account and group storage on this infrastructure would be visible throughout the envisioned platform. The storage architecture would connect to other components of the system over redundant 200 Gb/s ethernet links. This storage layer is planned to leverage the CEPH software-defined storage system. The ORCD team has considerable experience with the use of CEPH as a sufficiently cost-effective but manageable solution for an active bulk storage cloud.

The pilot project would be used to demonstrate a cost-effective, base service supporting all researchers, together with a direct charge system for adding project-specific capacity.

(h) A high-bandwidth and I/O rate storage cache that is globally visible.

Model training workflows in AI/ML often sample terabytes and potentially petabytes of data with a non-deterministic, quasi-random sampling strategy. To support large data for active projects, we need a fast, globally visible storage cache that supports high rates of I/O for model training workflows, which require data that are visible to all resources. This storage will reside on an Infiniband network that would be bridged to the long-term bulk storage infrastructure. This storage layer would capitalize on new generations of flash drive storage that can be used to create dense petabyte storage devices. A farm of these would provide a high I/O rate and bandwidth transient storage fabric for data that are visible to all devices in the system.

The pilot project would demonstrate a base system design leveraging NVMe I/O building blocks that would be designed to support the needs of AI/ML training beyond the MIT community.

(i) A continually evolving and extensible open software stack that forms a default ready-for-research environment.

The stack envisioned in **Figure 1** will contain many layers and will be maintained, versioned, continually tested and evolved by ORCD staff and students, in collaboration with research laboratories across the campus. **Figure 2** (below) illustrates a vision of the scope of such a stack. The stack would be a key component of a vision of a common environment that can serve all the broad interests of the MIT research AI/ML community. It represents an integration of a continually evolving collection of technologies that, when maintained together, can greatly accelerate research in individual projects and catalyze the exchange of ideas, experience, and people between projects. A comprehensive end-to-end integration of this nature does not currently exist at any US university that we know of.

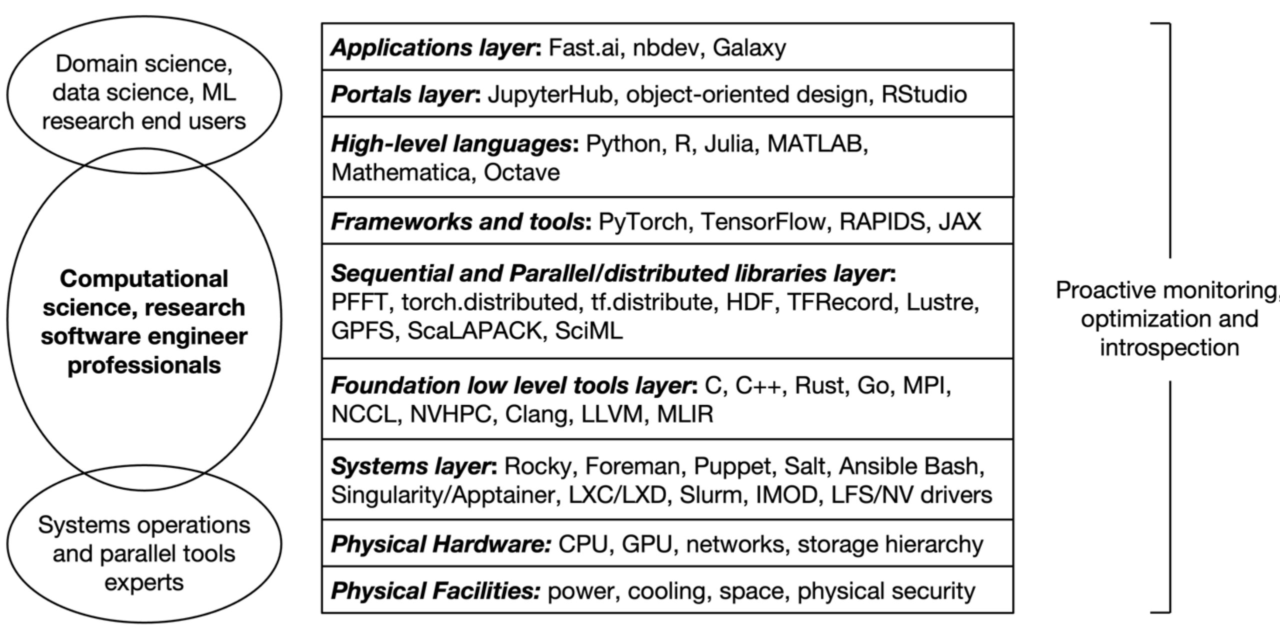

Figure 2: *A schematic of the many levels of a modern AI/ML for research software stack.*

The stack spans everything from a base operating system recipe to high-level environments that are fundamental to productivity. The stack we plan would be fully open and documented to allow others to replicate our platform. It would be designed to align with the overall architecture in **Figure 1** so that an entire clone of the platform can be deployed. It would contain default standard elements that emphasize and encourage adoption of proven productive tools and common practices.

Figure 2 also shows the central role of ORCD professional computational science/research software engineering staff. These are a key part of our proposed pilot project. As shown on the left of the figure, these staff bridge all the way between the systems-level hardware and facilities universe and the high-level research environment abstractions and tools. Professional staff with expertise in building, maintaining, applying, documenting, and teaching the software stack shown in **Figure 2** are critical to ensuring new techniques are applied effectively and productively. They are also key to ensuring that both stable and leading-edge environments are kept readily available and tested to ensure performance and correctness.

In the pilot project proposed here, we would create an initial, extensible, open and documented version of a complete stack that spans all layers with a limited selection of tools at each layer. The project would focus on a core of most common tools (primarily Python based in AI/ML) and on allowing for innovation around that core.

(j) Multiple generations of core hardware resources for flagship problems.

These are the computational hardware blocks that are used for scaling up AI/ML workflows and for applying new AI/ML ideas to domain research. They will be used to research new model algorithms and to experiment with inference ideas at large sizes. These resources are the computational heart of the AI/ML platform. The surrounding technologies are in support of the use of the core hardware resources. The technology (and economics) of potential AI/ML hardware is evolving very rapidly at the moment. The future direction of systems is quite unclear, despite the clear need for them. Accordingly, the figure shows a potentially rapid refresh cycle for any hardware or cloud resource that might be employed. For each cycle of hardware resource, a three-to-six-month commissioning period is envisioned. In steady state, this would overlap with two previous generations of systems.

Over the next five years, innovations in memory technologies, in optical silicon interconnects, in internal networking, and more could lead to sizable performance gains, while algorithmic innovations could shift the types of functions that are needed. To allow for this, we take a high-level perspective in **Table 1**. It presents successive overlapping generations of systems in terms of electrical power. In those terms we plan to target a steady-state level for the next five years of around 2 MW of AI/ML computing power. In the table, generational cycling is shown in these terms for the period out to 2032. **Table 1** shows that for each generation of resource, the life cycle would be similar. It would begin with an approximately six-month deployment phase (marked dep). During this phase, research applications would be leveraging the system while a fully tested initial software stack would be developed in parallel. The following approximately 18 months are the peak operation period for research using the system (marked ops). A system would then transition to becoming part of the gateway pool for 18 months before being retired (in phases marked tr, gw, and ret, respectively). This strategy serves to rapidly upgrade as technology evolves but also captures extra utility from resources by migrating them to “on-ramp”/gateway resource operations before eventual retirement.

On **Figure 1**, we sketch a likely [DGX GH200](DGX GH200) supercluster resource as the first such system. Such a system includes an internal tightly coupled interconnect and internal high-speed storage and allows distributed AI/ML workflows to execute efficiently. The actual hardware for the first cycle core resource would be determined by the project team once the pilot was underway.

In the pilot project, we propose to fund the deployment of the first generation of these systems. This first system would be of smaller power capacity but will allow significant work and demonstrate end-to-end system impacts.

		Y1 (2024)		**Y2** (2025)		**Y3** (2026)		**Y4** (2027)		**Y5** (2028)		**Y6** (2029)		**Y7** (2030)		**Y8** (2031)		**Y9** (2032)		
		H1	H2	H1	H2	H1	H2	H1	H2	H1	H2	H1	H2							
sys2024	300KW		dep	ops	ops	ops	tr	gw	gw	gw	ret									
sys2025	500KW				dep	ops	ops	ops	tr	gw	gw	gw	ret							
sys2026	1MW						dep	ops	ops	ops	tr	gw	gw	gw	ret					
sys2027	1MW								dep	ops	ops	ops	tr	gw	gw	gw	ret			
sys2028	1MW										dep	ops	ops	ops	tr	gw	gw	gw	ret	
:	:												..	..	..	..	..	..	..	..

Table 1: *Pipelining of evolving hardware technology system with annual refresh to track rapid development. The dates are based on an assumed start date of 2024. In practice, this outline will be gated by availability of funds.*

(k) Network fabrics that provide both high-bandwidth and low-latency connectivity.

The network would include standard ethernet connectivity to long-term storage and external ingress/egress. It would also include HDR and NDR Infiniband fabrics. This would provide very low-latency connectivity to high I/O rate shared storage for grand challenge training and inference work. The low-latency interconnect is an important piece to support distributed training workflows well.

The proposed pilot would include deploying a small-scale network fabric to support AI/ML workflows and demonstrate system impacts.

(l) An on-premise facility for housing and operating at least 2 MW of power-dense AI/ML computing equipment and 500 KW to 1 MW of additional gateway resources at cost-effective rates.

The pilot project we are proposing would leverage existing building, electrical, mechanical, and cooling infrastructure that MIT owns at MGHPCC. This facility draws most of its power from non–fossil fuel-based energy sources. MGHPCC serves as a purpose-built facility for housing power-dense research computing infrastructure and connects back to MIT over multiple 100Gb/s links and also connects over high-speed fiber links to commercial cloud peering points in New York.

(m) A maintained and tested seamless capability to instantiate all the software and support aspects of this overall platform into commercial cloud environments.

The open software stack shown in **Figure 2** will be designed for automated deployment to physical hardware and to commercial cloud resources. Recent advances in container technologies mean that it is practical to deploy software within simple LXC-based container environments that can streamline maintenance of a sizable stack that is fully portable. Additionally, the recent emergence of comprehensive integration and deployment testing tools allows much of the maintenance burden to be automated.

The proposed pilot will include documented, openly available software that can be deployed on commercial cloud environments as well as on physical hardware.

(n) Integration with non-AI-focused resources used for conventional simulation, data analysis, virtual experimentation, and digital design.

Finally, **Figure 1** shows that the system network fabric also provides capacity to connect seamlessly to general purpose research computing resources. In the MIT case, these are shared resources that the ORCD organization operates. While the generalization of this connectivity to any system lies beyond the scope of the project, the mechanisms envisioned will be fully documented and openly available and would likely extend to use by other organizations.

The proposed pilot will include demonstrations of integration with general purpose research computing resources that MIT ORCD operates.

In total, the subparts create a unified whole system that can serve as an integrated environment for AI/ML research from ideation to large-scale application. The ultimate value of this architecture comes from full, holistic integration under one project with a long-term design for continual refresh.

Surprisingly, at present, no university that we are aware of operates an open, state-of-the-art AI/ML system that is integrated across the entire research organization at the level envisioned here. In part, this reflects the often-decentralized nature of research organizations. Furthermore, possibilities have been limited until recently due to lack of necessary innovations in Linux container technologies that make them easier to maintain. Developing, demonstrating, and openly documenting the software and technical plans of an integrated resource would be transformative for practical AI/ML research computing needs at MIT and beyond.

Budget and Timeline

We proposed a $15M+ initial AI project covering the 2024–2027 timeframe. This will include sizable pilot hardware and dedicated expertise. Cloud resources may bootstrap the activity. Broadly, the funds split between the staff funding ($9M) and the pilot compute needed ($6M) to show impact and demonstrate potential. The field is evolving fast, and the timeline would bring resources together rapidly in the first two years, potentially leveraging cloud rental to kick-start platform deployment. The later years would see a significant addition of hardware assuming that aligned fundraising campaigns are successful.

In the proposed plan, we would follow the **Table 1** model for hardware deployments.

- New hardware would be deployed every year.
- Hardware would be operated for research-intensive work for approximately four years, consisting of several stages:
    - 0–6 mos: Initial operations phase
    - 6–24 mos: Standard operations phase
    - 24–36 mos: Transition to gateway resource
    - 36–48 mos: Gateway resource
    - 48–54 mos: Retirement

Under this approach, the key steps in the seed project by year are summarized below. **Table 1** provides a high-level view of the pilot project budget and a plan for phasing into the Olympia budget.

YEAR 1

Here we lay the core foundation for activities that would grow in subsequent years.

In year 1, we would operate a beta program. An initial version of the software stack we envision would be deployed. A hardware platform with all the components in **Figure 1**, but at a modest scale, would be deployed. We might use cloud services in the first year.

We would solicit beta testers from across the campus via a short request for proposals. Selected projects would work closely with the ORCD team on selected research. Participating groups would meet weekly with ORCD team members who would help harden workflows and examine scaling of training, fine tuning, inference, and other activities. Beta testers' interest in contributing to an open-source portfolio of "Open Research Ware" would be encouraged and incentivized through preferred access to more resources. Beta testers would be sought from every MIT school and the Schwarzman College of Computing, and from strategic areas such as climate, bioengineering, materials, and AI/ML algorithms.

The year 1 phase would advance domain research and also furnish a first batch of "ORCD recipes." These would provide open case studies that would serve as training material for larger groups of researchers in year 2 and beyond.

The year 1 software stack would be based on updated variants of existing ORCD work and experience both in teaching using web portals and notebooks and working with business organizations like Code Ocean and Google collabs and academic groups like Open On Demand (Ohio) and OpenHPC (MGHPCC and others). This technology is well understood by the current ORCD team and could be brought up on hardware or cloud resources rapidly. In year 1, we would hire a number of graduate students as staff to work on systematizing this stack.

During year 1, we would begin to pursue funding strategies, using the beta system as a concrete example of our vision. Funding directions developed would include leveraging institutional indirect costs; direct charge support from business units within MIT; philanthropic gifts for base cyberinfrastructure and for accelerating progress in discipline areas of interest in partnership with departments and groups across campus; and reconfiguring MIT processes around internal new hire and promotion funds.

Budget target: *$1M staff, $1M compute*

YEAR 2

In year 2, we propose to scale up the beta test program and system. We would add a more sizable hardware component capable of useful impacts to a larger group. We would extend the beta program toward being able to support “on-ramp” accounts for all of campus, with the expectation that several thousand “on-ramp” accounts could be active at any one time by the end of year 2. We would ramp up the professional staff/graduate student trainee team.

We would enhance the software stack to incorporate recent tools and to target the year 2 system. A likely candidate for this stage hardware could be an ARM-based DGX GH200 platform, which would add extra dimensions to software stack representation. In addition to instruction set changes on host CPUs, this hardware may support stochastic rounding and/or novel bit representations for numeric values. Both these would provide opportunities for accelerating research activities with an updated software stack and proactive tooling assistance for research groups. The ORCD support team would also introduce more extensive energy and performance measures into the stack.

During year 2, we plan to officially publish our first-generation architecture and workflows recipe and host a summer boot camp for peers interested in adopting and contributing to the architecture.

Budget target: *$2M staff, $5M compute*

YEARS 3–5 AND BEYOND

In years 3–5, we would focus on ensuring that the cyberinfrastructure developed becomes a true “inalienable right” for all members of the MIT research community. This would entail:

- **Ensuring** that there are truly seamless links between academic activities and research activities
- **Developing** professional boot camp activities to build a broad community for open AI/ML research across academic, government, and commercial partners
- **Working** to grow a dedicated base of funds to support future directions over the long term

Subject to external fundraising, we would plan to deploy up to $20M of new hardware each year during this period. Updated software and hardware recipes would be openly published each year. By the end of year 5, we aim to be able to transition all staff to a sustainably funded model that can extend beyond the pilot project phase.

Annual budget target: $2M staff ($20M hardware, outside of pilot)

Appendix IV

MIT Focused Discipline-Specific Science

We provide a list of projects and PIs below to prove existing research projects across MIT that could give scientific de-risking for the M³ pilot. Other projects certainly exist.

- Accelerating Climate and Environment Science with Olympia
  Faculty and PIs: R Ferrari, P O'Gorman, S Wang (CoC dual), A Bodner (CoC dual), S Beery (CoC dual), P Donti (CoC dual), S Ravela, A Fiore

- Accelerating Materials Science with Olympia
  Faculty and PIs: R Gomez-Bombarelli, T Smidt (CoC dual), C Coley (CoC dual), E Jossou (CoC dual)

- Accelerating Synthetic and Real Biology with Olympia
  Faculty and PIs: M Bathe, J Davis, S Vos, C Drennan, E Fraenkel

- Accelerating Health Care with Olympia
  Faculty and PIs: M Ghassemi (CoC dual), D Katabi, R Barzilay, S You (CoC dual)

- Accelerating Fundamental Physics:
  Faculty and PIs: W Detmold, P Shannahan, J Thaler, T Slayter, P Fisher

- Accelerating Brain Science
  Faculty and PIs: J Tenenbaum, J McDermott, N Seethapathi (CoC dual), L Lewis (CoC dual), S Ghosh

- Accelerating Social Understanding:
  Faculty and PIs: M Raghavan (CoC dual), IS Kim

Project Budget and Timeline

Olympia is an integration project of the current best-of-class approaches for containerization, virtual desktop, authentication, data storage (both near line and high-speed), sophisticated software stacks, and most importantly, the expert humans to be able to help construct, document, and guide the MIT community through the system. To achieve the goals set out, we plan a timeline with a pilot phase, Mens et Manus et Machina—M^3—and a steady-state phase, Olympia, following the multiyear timeline laid out below.

			Year 1	**Year 2**	**Year 3**	**Year 4**	**Year 5**	**Year 6**	...
MMM: Pilot Phase	Budget	Staff	$1M	$2M	$2M	$2M	$2M		...
		Computing Resources	$1M	$5M	–	–	–		...
	Activities		Initial 200KW deployment. Target 30 groups across campus. Five from each school plus college. Establish platform CI/CD.	Grow/refresh to 750KW deployment. Expand coverage to 20% of research groups. ORCD platform workshops and training.	Grow/refresh to 2MW deployment. Expand coverage to 40% of research groups. ORCD platform workshops and training.	Refresh at 2MW deployment. Expand coverage to 100% of research groups. ORCD platform workshops and training.	Ramp down MMM pilot funding. Transition to long-term funding profile.		...
Olympia: Long-Term Sustaining	Budget	Staff	ORCD staff	ORCD staff	ORCD staff	ORCD staff	ORCD staff + $1M	ORCD staff + $3M	...
		Computing Resources	–	–	$20M	$20M	$20M	$20M	...
	Activities		Build initial MMM team.	Grow MMM team.	Pilot supplementary annual contracts for research software engineering services.	Develop long-term base funding strategy. Test fee structures for supplemental support.	Refresh 500KW of 2MW. Grow base funding for MMM staff.	Refresh 500KW of 2MW. MMM staff fully base funded.	Continue with ongoing evolution. Funding supported from base and PI funds.

***Table 2**: Pilot project budget and goal for transferring to a long-term sustained budget.*

Olympia in Action

The research community at MIT is broad, deep, and varied, with a joint mission of producing the most innovative work in their fields and collaborating whenever possible to reach their goals. How will Olympia lead to successful outcomes for each community? We take the following as a subset of hypothetical community members to illustrate how this world-class platform will enhance and accelerate their time to research and provide substantial economies of scale.

The Newly Hired Professor

As a newly hired professor, day one consists of rapidly realizing that you are now running a small startup inside a large, complicated company. Startup (and retention) packages typically involve a request for computational capacity and capability and, eventually, the designation of funds to buy hardware. The challenge starts with the arrival of funds. Where should we plug in the hardware, how should we manage it, and is there enough power in the lab? What about software? Are there any staff members available to help? The list continues: in the meantime, the time to science slows down. Olympia will eliminate this process, providing instant, scalable, always-on computing and storage for the new professor, their staff, and students. Data sharing will be seamless, and negotiation allocations will determine the computing capacity —the request for data storage or a pre-negotiated amount of computing capability, as opposed to funds. Systems that are available immediately and able to scale will provide an attractive resource for the top faculty members we want to recruit.

The Tenured Faculty

As a senior faculty member with a strong research program, onboarding new postdoctoral students and staff follows with successful grant proposals. Olympia will ease the incorporation of new capabilities straightforwardly enabled by grant funds. Based on demand and the project's needs, the systems can scale from an initial allocation, growing or shrinking. The team has one environment to “log in” and a sustainable, reproducible software stack between the ebbs and flows of researchers coming in and heading off into their new careers. Long-term staff will gain familiarity with the environment, documentation will be stable and informative, there will be a strong feeling of belonging to the group, and it will be easy to share data among collaborators.

The Student Researcher

When you join a research group as a Ph.D. candidate, master’s student, or undergraduate participating in the Undergraduate Research Opportunities Program, your MIT Kerberos account connects to the principal investigator’s (PI) Olympia allocation, and you can start working on your first day. The documentation you look up is rich, well-vetted by previous users, and easy to digest. You can get up and running immediately. When you have a question about Olympia, you don’t contact an individual group server or service but a large, supported, interactive, and dynamic computer. You can set your work hours—the system is always active and available to try out your latest idea. Sizable standard computing images are big enough to try new things quickly.

The Administration

Each department will control a pool of resources to allocate between PIs, research staff, and scientists. Requests to the deans and funding agencies will pool these resources into Olympia allocations to PIs, with a base pool of significant capacity and capability. Requests for more computing are simply for additional capacity added to the Olympia system. With central Warranty management, nothing expires—the compute will be available on time and at the right time for your research community.