The Science & Technology (S&T) Risk Matrix [2026 revision]

In December 2018, the Deputy Secretary of the Department of Energy (DOE) signed the International Science and Technology Engagement Policy memorandum. This policy memorandum called for the creation of a Science and Technology (S&T) Risk Matrix. The purpose for the creation of the S&T Risk Matrix was to identify and protect critical emerging research and technologies that do not otherwise have control mechanisms, such as the rules associated with classified information, International Traffic in Arms Regulations, Export Administration Regulations, or 10 CFR Part 810 (Part 810)[1]. The S&T Risk Matrix is intended to highlight areas of emerging and potential concern associated with economic and/or international competitiveness, *and not to overlap or supersede existing controls associated with national security or export controls. The S&T Risk Matrix, moreover, is intended to be a living document, evolving based on ongoing dialogues between DOE and its Laboratories*. The S&T Risk Matrix was updated in December 2022, in 2024, and again in 2025 by the National Laboratory Chief Research Officer (NLCRO) committee in coordination with the DOE Laboratory Operations Board. In addition to general refreshes associated with the evolution of research frontiers and assessed risks, the 2022 revision emphasized updates related to Bioscience and Biotechnology; the 2024 revision highlighted Quantum Information Science & Technology; and the 2025 revision prioritized Machine Learning (ML) and Artificial Intelligence (AI). This version of the S&T Risk Matrix is effective April 2026. In the future, NLCRO will continue to coordinate annual updates with the Department of Energy.

As it relates to relevant DOE Orders, the S&T Risk Matrix only applies to interactions with and nationals of specified Countries of Risk. Currently, DOE's Countries of Risk are limited to China, Russia, Iran, Belarus and North Korea, as authorized under 42 U.S.C. § 18912.

To date, six areas of research have been identified as within scope of the S&T Risk Matrix:

- Quantum Information Science & Technology
- High Performance Computing
- Machine Learning/Artificial Intelligence Science & Technology
- Battery Science & Technology
- Bioscience & Biotechnology
- Accelerator Science & Technology

---

[1] Part 810 authorization requirement applies to all persons subject to the jurisdiction of the United States who directly or indirectly engage or participate in the development or production of any special nuclear material outside the United States. Pursuant to § 810.3, Definitions, the term "persons" does not include DOE. As such, DOE is not required to obtain Part 810 authorization for its own exports of nuclear technology and assistance. However, DOE must maintain program oversight (federal funding and direction) of such activities undertaken at National Laboratories to ensure consistency with U.S. national security and nonproliferation objectives. See, "Statement of Advice to Department of Energy (DOE) and National Nuclear Security Administration (NNSA) Laboratories, Plants, and Sites Regarding 10 CFR Part 810 (Part 810) Compliance."

As part of future annual S&T Risk Matrix reviews, the addition of new technology areas to the matrix will be explored in coordination between DOE and NLCRO.

The S&T Risk Matrix uses a Red/Yellow/Green categorization scheme to quantify the risk associated with a given topic and the associated level of controls that are required. When evaluating the status of a specific topic within the S&T Risk Matrix, in addition to examining the Red/Yellow/Green tables below, the reviewer must also review the additional information and broader context for each section provided in Appendix 2 of the S&T Risk Matrix. Further, in addition to reviewing relevant content with respect to the S&T Risk Matrix, users of the matrix, consistent with existing practices and processes of their home organizations, must also consult other documents, including relevant classification and export control guidance, which are beyond the scope of this document, in order to understand and implement the full suite of associated controls and protections required by DOE.

In addition to its function of quantifying risk and defining necessary controls and mitigation, the S&T Risk Matrix can also be used as a resource for education and awareness for Laboratory staff, providing insight into when research might move from Green to Yellow, highlighting the need for additional protections so that staff receive appropriate credit and protection for their innovations.

RED: Red (restricted) emerging technology topics have sensitivities associated with economic and/or international competitiveness that could cause significant harm to critical national interests of the United States (U.S.) if shared with a country of risk without appropriate vetting and approval. These red topics are considered "restricted" for purposes of increased vetting and controls involving interaction with Countries of Risk, and their representatives, as defined in various DOE Orders. Access to restricted technologies by Country of Risk nationals or entities requires enhanced vetting and approvals by DOE at both the local/field and Headquarters levels.

A topic is deemed red when it is developed to a sufficiently high level that it poses such a risk. Initial research, early-stage prototypes etc. are typically not expected to be red. Technology readiness levels (TRL) are often a reasonable guide as to the readiness of a given technology. Note, it is the actual readiness that matters. It is not sufficient for a project to have the aspiration to achieve a given performance level for that work to be deemed to be red.

Additional Protections: Restricted topic areas trigger compliance requirements with elements of other DOE orders and policies (details are provided in Appendix 1). Laboratories will develop access management plans and ensure regular oversight/monitoring of these restricted projects to ensure they remain appropriately protected throughout their lifecycle. Access management plans for restricted S&T topics shall include at a minimum:

- Description of the work that has been identified as restricted
- Responsible principal investigator
- The work in red areas in the subject matrix is not itself designated as Controlled unclassified information (CUI) while it is being carried out, though it will be controlled as outlined in the access management plan. Once manuscripts or other document form publications are being developed, these should be designated "CUI//SP-S&T Matrix"

until such time as they have been approved for release, when this marking will be removed because under the matrix the information is only being controlled for a certain period of time, with the intention of it ultimately being released.

- The process by which the restricted S&T will have intellectual property protection prior to release/publication of the work, and appropriate notifications about the release/publication.
- DOE O 241.1C describes the process by which science and technology information (STI) is appropriately identified, categorized, disseminated, and preserved. Requirements for STI are laid out in the Contractor Requirements Document to inform each Laboratory's processes, and these processes must include appropriate review and approval steps for restricted S&T topic areas. For research that is determined to be a Restricted S&T topic, notification to the funding DOE Program Office(s) should occur prior to publication.

When possible, Laboratories should use existing protective measures, processes, or programs to implement needed controls on these projects.

YELLOW: Emerging technology topics that have the potential to become red (restricted) from an economic and/or international competitiveness standpoint or represent areas in which enhanced vigilance is appropriate.

Additional Protections: Yellow topic areas may require additional controls under certain circumstances. Yellow topics should be thought of as being on a 'watch list.' It is the responsibility of each Laboratory to establish processes for monitoring and controlling technologies in the yellow category. Depending on the technical area, yellow emerging technologies may be defined by a well-defined parameter space; other areas may require the judgement of technical and security subject matter experts (SME) to determine their degree of sensitivity. Laboratories may develop access management plans for projects in yellow technology areas (see elements described under red technology area access management plans). When possible, Laboratories should use existing protective measures or programs to implement needed controls on these projects. Yellow topics may require specific coaching/awareness training of performers and managers on how to engage with, involve and/or share information with individuals from designated Countries of Risk.

GREEN: Emerging technology topics that do not have sensitivities associated with economic and/or international competitiveness. Fundamental scientific studies or technologies at a low TRL are often – but not always – in the green category.

Additional Protections: Green subjects do not require additional controls beyond those already in place and will be handled with existing mechanisms.

Quantum Information Science & Technology

Quantum Information Science and Technology (QIST) is a highly interdisciplinary field that builds on quantum mechanics and information theory to explore the fundamental limits for computation, networking, and sensing. QIST incorporates research in a broad number of areas that include (but are not limited to) materials, computer science, mathematics, laser physics, atomic physics, cryogenics, electrical engineering, systems engineering, and application specific software development. The QIST research, development, and demonstration field has witnessed explosive growth in the last five years and continues to rapidly progress. Novel quantum applications to sensing, computing and simulation, and communications could potentially disrupt many aspects of current information science and technology, although in many instances realization of these approaches is well in the future. The QIST section of the S&T risk matrix provides guidance in the five broad areas of 'Quantum Computing Hardware,' 'Quantum Error Correction, Algorithms and Simulation,' 'Sensing, Clocks, and Metrology,' and 'Communication,' For all topics below, information already protected by either classification or export control (or, in some cases, vendor-specific non-disclosure agreements) is not reflected here, but scope should be evaluated relative to these controls on a case-by-case basis prior to release or dissemination.

Quantum Computing Hardware

GREEN	YELLOW
Single Qubits and Qudits Methodologies, designs, control processes, and fabrication processes that achieve qubits, qudits or gate operations with physical error rates greater than 10^{-5} on a single device with at least two connected qubits. The development of novel qubits or new approaches to existing qubit technology. Theoretical and/or experimental understanding of environmental limitations affecting qubit performance. Research developing single isolated physical qubit/qudit devices.	Methodologies, designs, control processes, and fabrication processes that achieve qubits, qudits or gate operations that operate with physical error rates between 10^{-5} and 10^{-8} (where different from those used for higher error rates) on a single device with at least two connected qubits.

Qubit Devices and Systems Theoretical or experimental demonstrations of devices with up to 2,000 physical qubits.	System-level designs for experimental demonstrations of non-commercial quantum computing systems, subsystems, or application-specific quantum processors with up to 20,000 physical qubits, including control electronics, if not otherwise controlled. Analyses of such system-level designs that include specific, realistic qubit models and operational details.
Controls Research quality device operating software for devices with up to 2,000 physical qubits.	Non-commercial device operating software for devices with up to 20,000 physical qubits. Development of novel cryogenic amplifiers for <1K operation.

Quantum Error Correction, Algorithms and Simulation

GREEN	YELLOW
Quantum Error Correction Basic research on error correction codes that is architecture-agnostic and with abstract qubit (or qudit) models. Algorithm-specific error mitigation on NISQ devices for chemistry, computational physics, and machine learning. Error correction subroutine on NISQ systems demonstrating breakeven performance.	Quantum error correction and logical qubit (or qudit) development in laboratory experiments achieving logical error rates within a total logical error of 10^{-8} or better for scientific demonstrations. Laboratory experiments that run portions of error correction codes on systems of physical qubits that are significantly smaller than what would be needed for fault tolerant operation.
Quantum Algorithms Basic research into quantum algorithms that show promise for no to modest degrees of speed-up	Development of quantum algorithms with quantum advantage over the best-known classical counterpart while being applied to research that screens

from heuristics or asymptotic complexity arguments. Quantum algorithm development for chemistry, computational physics, and machine learning. Evaluation of algorithms for small-scale and low-fidelity applications and algorithms for chemistry, materials, particles and other physics, and machine learning. *Quantum Simulation* Exploration of QIST foundations – fundamental questions related to the formulation of quantum mechanics and its relation to information. Digital and analog quantum simulation and optimization.	yellow on the risk matrix in any section. General-purpose quantum simulators (classical computer simulations of quantum computer operation) for systems with more than 33 physical qubits and implementing full noise models.

Sensing, Clocks, and Metrology

GREEN	YELLOW
General quantum sensing for basic research applications. General research into atomic clocks and networked atomic clocks. Use of quantum techniques for enhanced microscopy in a laboratory setting.	Development of entangled photonics capabilities useful for kilometer range quantum imaging. Development of quantum navigations systems capable of 100 m positional uncertainty. Development of quantum clocks with accuracy surpassing ps/day. Development of system components with the goal of enabling fieldable quantum sensors.

Communication

GREEN	YELLOW
Development of fundamental building blocks of quantum networks including single and entangled photon sources, single photon detectors, quantum memory, squeezed light sources, and protocols. Development of error mitigation strategies. Integration of multiple quantum networking devices. Fundamental quantum repeater research. Development and proof-of-principle demonstration of use cases for quantum networks for basic science and sensing applications. Proof-of-principle demonstration of entanglement distribution, including multi-partite entanglement and entanglement across multiple nodes.	Robust and efficient entanglement distribution at practical data rates of up to 1,000 ebits/s. Development of security applications using quantum communication other than quantum key distribution. Experimental fault-tolerant quantum communications via quantum error-corrected quantum repeaters. Development of networking interfaces and protocols for scaling up quantum computer operations.

High Performance Computing

High Performance Computing (HPC) is a critical technology used for both predictive simulation and large-scale data analysis, essential to national security, scientific discovery, and economic competitiveness. Within the U.S., DOE and the National Laboratory complex have a leadership role in HPC, fielding some of the most powerful computing platforms in the world and providing expertise in computational modeling, data analytics, algorithms, and software. The Labs also partner closely with the computing industry, academia, and other government entities to rapidly adopt new commercial innovations at scale and to drive the development of advanced HPC technologies, systems and expertise. DOE Labs also partner with other industries and government entities to accelerate the application of HPC to engineering design, drug and materials discovery, infrastructure security, energy security, manufacturing, and a wide array of other applications. HPC is critical to ensuring the safety and security of the nuclear weapons stockpile and to the basic science and applied energy missions of DOE. The HPC section of the

S&T risk matrix provides guidance in the areas of system research and development (R&D); hardware components and systems; and software, including system software, applications, and libraries for platforms including tightly integrated systems for both modeling and simulation and for AI, as well as systems with components at the edge, including processing at experimental and observational sites and the inclusion of quantum resources for hybridized computing tasks.

Importantly, when HPC capabilities are used to contribute to another area of S&T, guidance from those areas, not the HPC section, is applicable in determining what topics, if any, are restricted. For all topics below, information already protected by either classification or export control (or, in some cases vendor-specific non-disclosure agreements, which are particularly prevalent in HPC) is not reflected here, but scope should be evaluated relative to these controls on a case-by-case basis prior to release or dissemination.

More specifically, DOE has historically engaged in development/procurement activities with vendors in realization of next-generation systems that are proprietary. Though individual hardware and software components under development by HPC vendors are typically explicitly protected by nondisclosure agreements (NDA), care should be exercised to identify information that could be discerned from publicly available sources about the details of these components before their release. This could include production software components that are being specifically tailored to run on NDA-protected hardware platforms that might contain Application Programming Interface (API) usage specific to that new hardware, or even explicit workarounds for testing purposes (e.g., “hard-coded” enumerations of new hardware components in a “testing” routine). This could also include performance predictions for software to be run on the new components that appeal to simple scaling arguments but implicitly include scalings that specify architectural features.

This document describes the underlying HPC technology landscape and provides an initial assessment of criticality to economic and international competitiveness. It is intended to be a living document, as this analysis will need to be evaluated regularly to include emerging technologies as they become viable for HPC deployment and to remove technologies or adjust performance thresholds as commodity computing technology continues to improve.

The scope of this document is intentionally focused on traditional digital computing, storage, and communication, but touches on technologies with the potential to significantly enhance HPC systems, such as neuromorphic and approximate computing, as well as new models for hardware devices (transistors, etc.) for digital computing, memory, and storage. Areas where quantum computing and AI make direct use of or tie directly to HPC areas of concern should be evaluated keeping in mind the totality of all the identified issues in each respective section of the Risk Matrix. Both AI and quantum computing provide abundant opportunities for overlap with HPC, including in the form of hybrid and edge computing cases. Some other topics may warrant a similar analysis.

Other U.S. agencies, notably the Department of Defense, the National Science Foundation, the National Oceanic and Atmospheric Agency, and the National Aeronautics and Space Administration invest in High Performance Computing, and some coordination with those agencies is appropriate, recognizing that their different missions may lead to different technology access control approaches.

HPC System R&D

GREEN	YELLOW
General R&D not associated with realization of specific systems, e.g.: • System management R&D for resilience, energy efficiency. • Integration of runtime software with out of band system health and monitoring, and performance counter networks.	None currently identified that is not otherwise protected

Vendor-Developed Hardware Components (processors, memory/storage, and interconnects) with National Laboratory Co-Design

GREEN	YELLOW
Component design details that are publicly available. Processor architecture (Instruction Set Architecture, or ISA) and performance characteristics. Memory/storage technology interface and performance. Interconnect performance characteristics and user level programming interface. Lab suggestions for co-design improvements to vendor hardware component designs that are not adopted and are generally applicable.	Lab suggestions for co-design improvements to vendor hardware component designs – prior to adoption/integration into product design.

Note: Any performance data or performance analysis for pre-release systems should be covered by NDAs and such information should be embargoed until general availability of the technology.

Note: Includes all post-release information not covered by NDA.

National Laboratory R&D on HPC Components and Systems

GREEN	YELLOW
Conceptual design of HPC, sensors or instruments. This includes any ancillary performance projections, models, or any other analysis of an incomplete or conceptual design. Laboratory technical advances that have resulted in awarded patents.	Detailed design (or documentation/blueprint thereof) that includes all design elements and descriptions such that a reader could recreate or manufacture, but has not yet been manufactured, prototyped, or otherwise demonstrated. (See Note) Lab designs that warrant filing of technical advances because they are potentially valuable for DOE and U.S. Government use but have not progressed to patent filing stage.

HPC System Software

GREEN	YELLOW
General purpose system software (source and executables) developed by a Lab. General purpose open-source software (source and executables) used or modified by a Lab	None currently identified that is not otherwise protected.

HPC Application Libraries and Frameworks

GREEN	YELLOW
Open-source libraries and frameworks.	Domain-specific modules of these libraries/frameworks that are designed specifically to benefit applications that may include, but are not limited to, applications that are already covered by export control, classification guidance, intellectual property processes, other sections of the S&T risk matrix or other controls.

HPC and Hybrid Solutions at the Edge

GREEN	YELLOW
Open source or publicly available hardware or software including libraries and frameworks.	Work at the level of definition and continued development of lab-developed application specific integrated circuits. Tests or solutions that include production hardware/software or developing edge technology. Prototype deployment and development, including the development of standards.

Artificial Intelligence and Machine Learning

Artificial Intelligence (AI) and Machine Learning (ML) are increasingly a key part of virtually all of the ongoing research efforts in the mission programs of DOE. AI and ML are also vibrant areas of research in industry and academia. The scope of ongoing work indicates that significant economic value can be associated with such research outcomes.

This topical section is organized into three sub-areas: Fundamental AI Technology, Data, and Models. Fundamental AI Technology includes algorithms, software, and system libraries. The Data portion covers both raw data as well as methods to curate data. The Models portion addresses generative models, surrogate models, and agentic systems.

For greater clarity, below are some explicit definitions:

- *Algorithm* - A mathematical or procedural recipe for solving a problem. It exists independently of data. For example, gradient descent is an algorithm: it specifies how to iteratively minimize a function. Algorithms are abstract procedures, and their success depends on the context and data to which they are applied.
- *Model* - A model in AI/ML is a structure whose parameters or weights are learned from data through the application of algorithms.
- *Surrogate model* - An approximation of a more complex or expensive-to-run function or system.
  In AI/ML, this can mean: (1) training a model on a limited number of high-fidelity simulations to cheaply approximate the system, or (2) using an approximate function during optimization when the true function or loss is costly to evaluate. In both cases, the surrogate depends directly on the data it is trained on.
- *Training*: The process of applying algorithms to data in order to adjust the parameters of a model. During training, algorithms such as gradient descent or backpropagation iteratively update weights so that the model captures patterns in the dataset. Training is what transforms an abstract model structure into a functioning model. The effectiveness of training depends both on the chosen algorithm and on the quality and characteristics of the data and the structure/architecture of the model.

The use of AI and ML tools and approaches is widespread in DOE research, as well as in some operations. Data science research is often intertwined with AI/ML in their respective uses. AI/ML is also an active area of research in industry and academia. This AI/ML section of the S&T Risk Matrix provides guidance in the areas of foundational algorithm development, the data being used or produced, and the models being built.

Generically, there are two classes of advancement: those that improve AI techniques in general, i.e. new model architectures, new algorithms, optimized implementations, etc., and those that provide new capabilities to an application area, i.e. providing faster surrogate models, speedy approaches for curating data, etc. The economic value of the latter class is predominately the purview of the corresponding application domain since it requires deep domain expertise and knowledge. Consequently, the AI/ML portion of the matrix defers to the application areas in these cases.

Fundamental AI/ML Technology

	GREEN	YELLOW
Algorithms	Basic research in AI/ML algorithms considered to be fundamental research of the same type, level of detail and essential content as basic and applied research in science and engineering, the results of which are ordinarily published and	Fundamental algorithms trained on or specifically tailored to data or domains designated YELLOW in other topics associated with the matrix. Algorithms applicable to data or domains separately designated as

	shared broadly within the scientific community, with no foreseeable adverse impact to any interest in U.S. commercial development of the technology. Research explicitly designated as publicly releasable, i.e. as part of funding requirements or other prior agreements.	YELLOW and which are expected to accelerate training or inference of existing architectures on current hardware or reduce data and/or energy requirements for equivalent performance with significant economic interest. Neural network architectures applicable to data or domains designated as YELLOW expected to reduce the size of models, training or inference time, or data and/or energy requirements for equivalent performance with significant economic interest. Data distillation techniques applicable to data or domains designated as YELLOW expected to reduce resources (time, power, compute, etc.) to build derived models by an order of magnitude.
Software	Implementation of algorithms designated as GREEN.	Implementation of algorithms designated as YELLOW. Implementations of algorithms with known application to specific hardware of commercial interest, i.e.,SWaP- (Size, Weight, and Power) constrained devices, latest generation GPUs, etc., expected to improve performance of training or inference by an “order of magnitude.”
System libraries	Standard development, debugging, and porting to new architectures.	Improvements expected to reduce the cost (computational, acquisition and/or power) of AI capabilities applicable to data or domains designated as YELLOW by an “order of magnitude” over the current baseline.

Data

	GREEN	YELLOW
Data	Data readily available in the current format not designated as YELLOW or RED by the application area. Data created under existing agreements designated as openly available, i.e. those created through government funding and not explicitly restricted.	Data designated as YELLOW by the application area. Any data collection derived from data designated as YELLOW.
Methods to curate data	Standard data curation or collection methods not tailored to or demonstrated on YELLOW or RED data.	Data curation tailored to data specifically designated as YELLOW by the application area in the risk matrix. Data curation or collection methods applicable to data or domains designated as YELLOW expected to result in an "order of magnitude" larger dataset than what is currently available.

Models

	GREEN	YELLOW
Generative models	Models approximating data distributions designated as GREEN by the application area.	Models approximating data distributions designated as YELLOW by the application area. Models whose output is designated YELLOW in an application area. Models enabling application capabilities designated as YELLOW (i.e. distilling models representing data that would be designated YELLOW).

Surrogate models	Models approximating simulations or experiments designated as GREEN by the application area.	Models approximating simulations or experiments designated as YELLOW by the application area.
Agentic systems	Agentic systems tasked with and only in GREEN application domains.	Agentic systems tasked with output targets to explore and operate in YELLOW application areas.
Standards Development	Development of ontologies or standard language for GREEN application domains. Development of interoperability standards limited to specific GREEN domains. Development of standards for machine readability or interpretability limited to specific GREEN domains.	Development of standards specific to YELLOW domains. Development of data standards specifically targeting commercial AI applications resulting in an "order of magnitude" improvement. Development of standards for machine readability or interpretability for YELLOW domains.

Battery Science & Technology

Philosophy: The matrix will examine the status of R&D as compared to the State of the art (SOTA) and determine if new innovations have risen to a threshold where the project is on a trajectory towards economic impact. Such a determination will consider: (i) maturation in TRL and (ii) performance that approaches/beats the SOTA across multiple metrics (e.g., energy density AND cycle life).

Foundational understanding of phenomenon will be considered Green. Performance metrics in lower TRL approaching SOTA metrics will trigger Yellow. Higher TRL activity that shows SOTA-beating metrics in at least one category with comparable performance in other metrics will trigger Red categorization.

DOE has long quantified SOTA numbers using BatPaC simulations coupled with extensive testing of commercial cells. Each year, these numbers are released to the Lab community. The matrix should be updated with SOTA targets once a year, without changing the specific thresholds that trigger a change in color.

The principal approach to determine the color will be through intellectual property (IP) or knowhow generation. This will leverage the existing processes in the Labs for identification and filing of IP and will bring together lab SME's and business development executives to make the determination of the nature of the IP (foundational vs. approaching economic value) based on the SOTA and the trajectory of the R&D. The table below provides guidance to SMEs on the thresholds that should trigger a change in color. *Ultimate determination of color will be up to the SMEs based on the philosophy outlined here; the table below is a guide that can be contravened as necessary.*

The thresholds below are determined to ensure that projects have time to adapt to changes in requirements necessitated by a change in color designation. Moving from Green to Yellow is triggered by promising data vs. SOTA but at low TRLs. This will require the project to make plans for ensuring information protection and to determine pathways for personnel from Countries of Risk to move to other projects, if needed. Changing from Yellow to Red occurs when the TRL is high AND the metrics are approaching (but not yet meeting/exceeding) the SOTA. This ensures that all economically valued IP will be conducted under Red designation.

Multi-lab projects will require SMEs from all the Labs to coordinate the colors to ensure consistency. We propose that these are done in the summer, in coordination with yearly planning for the next fiscal year. Projects can, in theory, have Green, Yellow, and Red activities simultaneously, based on the focus of the work and the TRL of the activity.

Batteries for Commercial Transport

GREEN	YELLOW
Foundational understanding of phenomenon.	Materials/process technologies tested in either half- or full cells of scale <50 mAh that demonstrate comparable performance and/or cost to SOTA for the specified application. Such cells should also demonstrate <80 percent capacity fade to 300 cycles.

High Energy Batteries for Military/Space Applications

GREEN	YELLOW
Foundational understanding of phenomenon.	Materials/process technologies tested in either half or full cells of scale <50 mAh that demonstrate comparable performance to SOTA for the specified application.

Energy Storage for Grid Applications

GREEN	YELLOW
Foundational understanding of phenomenon.	Materials/process technologies tested in either half or full cells of scale <50 mAh that demonstrate comparable projected cell-level energy density to SOTA with <80 percent capacity fade over 1000 cycles

Low-cost Energy Storage for Grid Applications

GREEN	YELLOW
Foundational understanding of phenomenon.	Materials tested in either half or full cells that demonstrate similar projected pack/rack costs at SOTA. Validated power electronics topologies that have pathway to lower cost and improved performance compared to baseline.

Battery Recycling/Supply Chain/Processing

GREEN	YELLOW
Foundational understanding of material processing phenomenon for battery materials	Material and scale up methods capable of producing < 1 kg of material for battery materials (cathode, anode or electrolyte component) that meet SOTA performance requirements in a cell < 50 mAh

Battery Safety

GREEN	YELLOW
Pre-conceptual research into thermal runaway mitigation, propagation, and short circuit prevention approaches	Technological approaches that reduce the risks and/or consequence of thermal runaway in battery storage systems in which, under a defined abuse condition (nail penetration, overvoltage, over temperature, etc.), cell to cell propagation is prevented. Demonstration on full cells with capacity > 2 Ah with similar electrochemical performance to SOTA.

Battery Life Prediction and Data Protection

GREEN	YELLOW
Implementation of methods to predict life using limited data from duty cycles or test data which mirrors the final use case	Implementation of accelerated test methods and accelerated test methods to predict life or performance

Bioscience & Biotechnology

Biology research is an exceptionally broad and rapidly expanding area, supported by many sponsors, including but not limited to DOE. The Bioscience & Biotechnology section of the S&T Risk Matrix provides guidance in the areas of synthetic biology, omics and automation technologies, data and advanced computational biology, biomanufacturing and biomaterials, biomedical research and technologies, bioscience and biotechnology applications of artificial intelligence, and agricultural and environmental technologies. Information already protected by either classification guidance or export controls (or, in some cases vendor-specific non-disclosure agreements) is not reflected in the risk matrix, but scope should be evaluated relative to these controls on a case-by-case basis prior to release or dissemination. Similarly, reviewers should be cognizant of best practices developed by other segments of the bioscience and biotechnology research community and apply them appropriately.

Synthetic Biology

Keywords: CRISPR, gene editing, genetic modification, protein engineering, gene drive, gain of function, biological parts, transformation, DNA assembly, library, cloning, metabolic engineering, genome-scale modeling.

GREEN	YELLOW
Manipulation and optimization of microbes, algae, fungi, viruses and plants using widely or commercially available synthetic biology tools, technologies and methods for the incremental performance improvement of bioproduction processes.	Engineering and optimization of commercially viable microbial, algal, fungal, viral, microbiomes and plant chassis using proprietary synthetic biology tools, technologies, and methods for pilot-scale bioproduction with significant intermediate-term commercial value.
Use of publicly available methods, reagents, data, genetic information, or tools to manipulate microbes and/or plants for basic research regulating gene expression.	Early (pre-commercial release) access to gene editing reagents from commercial vendors (e.g. vectors, nucleic acid modifying enzymes, genetic code expansion, or gene drives) or early (pre-commercial) access to equipment and/or upgrades.
Internationally distributed microfluidic or engineering systems to improve biomanufacturing technologies and products.	Optimization or validation of reagents or equipment that are NOT available in sensitive countries and potentially suitable for reverse engineering. This may include technology that is widely distributed in the U.S.
Research involving booting or rescue technology of synthetic microbes or viruses at laboratory scale.	Demonstration of booting or rescue technology of synthetic microbes or viruses at laboratory scale.
State-of-the-art models and analysis techniques for predicting genetic mutations for a desired, non-harmful phenotype.	Demonstration at laboratory scale of disruptive synthetic biology methods, tools or technologies that have the potential for significant economic or national security impact (e.g., improved genome editing technologies, DNA cryptography, or viral genome

	modification to increase payload size and transduction efficiency).
Publicly accessible databases, such as registries of standard biological parts, GenBank, and DIVA, which are informatic repositories of microbial strain DNA and protein sequences and designs thereof.	
Publicly available techniques that enable the same genetic construct to be deployed/assessed in parallel across biological phylogeny - currently across some groups of selected non-pathogenic bacteria and fungi for discovery or proof of concept. These capabilities are used in support of basic research, e.g., in establishing the molecules biosynthesized by secondary metabolite clusters for use as biofuels or bioproducts.	

Omics & Automation Technologies

Keywords: DNA sequencing, genomics, transcriptomics, proteomics, lipidomics, glycomics, metabolomics, metagenomics, metatranscriptomics, metaproteomics, laboratory automation.

GREEN	YELLOW
Fundamental research involving omics measurements of biological systems. These can include: • Research on the structure, function, or chemical properties of molecular species that drive prototype technological advancements in analytical measurements or computational pipelines • Omics sample preparation methods that require human intervention or evaluation or	Demonstration at prototype scale of advances in nucleic acid sequencing and data processing technologies that significantly advance the state of the art. Development of automated, multi-omics methods for battlefield, clandestine, or defense applications

those that require multiple hours to accomplish • Incremental advancements in omics instrumentation that still require multiple measurement modes for confident identification (e.g., mass followed by fragmentation with additional mass measurements) and combinations of instrument systems (e.g., vacuum pumps and chromatography systems) • Small sample size or single cell omics measurement platforms that measure multiple biomolecules across different platforms (e.g., proteins, metabolites) • Research using omics-based measurements and technologies for a mechanistic understanding of biological and environmental systems	
Development of fully automated instruments that take cells, tissues or other biological matrices as input and provide data results without human interaction	Development of controls for automated laboratory instrumentation, databases, or analytical tools that could be used for nefarious purposes.

Data & Advanced Computational Biology

Keywords: predictive models, medical countermeasures (MCM), digital biosecurity, signatures, detection probes

GREEN	YELLOW
Research, development and applications of computational and modeling tools including AI, for fundamental science.	Development of proprietary or controlled datasets used for MCM development and predictive mechanistic MCM modeling.

	Information regarding the security or vulnerabilities of databases and computational tools that can be used as training/validation sets for computational applications in MCMs.
New methods in HPC architecture and AI/ML methods to manipulate, store, and model data more efficiently.	Development of methods to protect data in cloud computing resources.
Basic research into the use of nucleic acids for data storage.	Testing and evaluation of DNA-based technologies to store, read and display data in a significantly smaller footprint (e.g. 5X or better).
Databases of environmental microbes that may contain sequences not readily available from other sources, when used in support of basic research, e.g., in finding new branches of life and understanding virus/host relationships.	Development of tools used to clean up and identify errors in publicly used biodata sets for activities related to any organisms on Federal Select Agent and Toxins list or the U.S. Department of Commerce export control lists.
	Unpublished and/or proprietary synbio software and workflows (in accordance with U.S. guidelines for providers of synthetic double stranded DNA) for the screening of DNA sequences against lists of organisms and toxins on the select agent and commerce control lists. This includes both software and workflows that pertain to specific signatures, details of (informatic) sensor systems and analysis and annotation methods relating to the identification of biological agents.
Technologies already commercially available or in the public domain.	Development of technologies meant to facilitate genetic testing and personalized medicine including “gene banks.”

	Security of drug products information [e.g., supply chain dependencies, vulnerabilities in the development pipeline, drug product data, regulatory filings, etc.]

Biomanufacturing and Biomaterials

Keywords: process engineering, bioreactor, supply chain, bioproduction, biocatalysis, pilot scale, commercial scale, misuse, dual use, extraction and purification, industrial enzymes, biomimetic, bio-derived material.

GREEN	YELLOW
Foundational knowledge for metabolic engineering and commercial biomanufacturing, including publicly available gene editing and protein expression technologies and standard cultivation conditions for microbes, fungi, yeasts, plants, and animals.	Engineering and optimization of metabolic pathways for commercial bioproduction processes with significant near-term commercial value. Such bioproduction processes are economically feasible with robust biomanufacturing and process controls.
Foundational knowledge for metabolic engineering, biomimetics and biomaterials relevant to National Security, for example, pathways to make precursors for energetic materials or biomolecules that bind rare earth elements (REEs).	Late-stage engineering of biomimetic structures and processes. Optimization of biosynthetic pathways for commercial production of target products that are national security relevant, for example a complete pilot scale system using engineered microbes or plants to mine REEs in a relevant environment.
Predictive tools for scaling biomanufacturing, including modeling bio-manufacturing scenarios to identify areas for technology innovation.	Development of a full process scheme at a commercial scale for valorization of bio-derived materials into energy, security, or industrial applications.
Publicly available tools for biomanufacturing facility modification and development. Techno-economic analysis and life cycle assessments using public data.	Proprietary engineering design and definition of a complete biomanufacturing facility.
Unit operation, commercial parts in bespoke facility design	Advances in biomanufacturing systems using interchangeable equipment designs for scalable and flexible production.

General purpose bioreactor designs and technologies including inline monitoring and real time monitoring, standard downstream processing, and quality assurance optimization.	Modeling and simulation that specifically supports design and optimization of proprietary reactors or separation systems to produce novel bio-derived materials. New strategies for inline monitoring, real-time adjustments, and feedback to processing.
Global regulatory testing standards, advanced process control strategies and raw material characterization	New biocontainment processes with commercial value.
Publicly available methods for the separation of biomass components and harvesting logistics at the molecular and tissue levels including methodology and equipment designs.	Commercial scale activities on extraction and purification of intermediate products derived from biomass with energy and environmental applications.
Enzyme compositions from natural biomass degrading communities (metaproteomes). Design and modification of a enzyme or ligand, binding protein for pre-commercial research. Fundamental research on biomaterials such as nanocellulose and other plant-based polymers. Development of theory and simulation methods that describe and/or predict the behavior or properties of biomaterials.	Protein and process optimization at commercial scale of top performing biomolecules.
New methods for process intensification and quality control techniques for protein production.	Late-stage optimization of production protocols of industrial and biomimetic enzymes that are likely to change reactor design, profits, and markets, etc. at the commercial scale.
General approaches for immobilizing and scaffolding enzymes in cell-free systems.	Subsystem-level production of cytotoxic products in cell-free

	systems that can be scaled for commercial production.
Fundamental research on catalyst materials, chemistry, structure, and modeling including synthesis and characterization techniques.	Pilot or commercial scale production protocols for catalyst materials, catalyst active phases or supports, and process design for biomanufacturing systems with significant, near-term commercial value.

Biomedical Research & Technologies

Keywords: diagnostics, treatments, medicine, bioengineering, cognitive sciences	
GREEN	YELLOW
Fundamental biology research into cellular mechanisms of host function and response to various stimuli such as pathogens, radiation, or chemicals, including chemical weapons agents (CWA).	Applied research on techniques for therapeutics, vaccine design or other medical countermeasures, including personalized medicine, intended for clinical trials and human treatment.
Fundamental research on the use of nanoparticle platforms or other emerging delivery platforms for vaccine and therapeutic delivery including in vivo distribution analysis, immunogenicity, pre-clinical safety/toxicity evaluations and efficacy.	Demonstration in a laboratory environment of potential dual-use delivery systems for therapeutics and vaccines such as: • Nanomaterials • Microneedles • Targeted in vivo delivery • Systems designed to cross the blood-brain barrier
Development of platform technologies, such as: • Monoclonal antibody design and development, natural and synthetic	Applied research on platform technologies, such as: • Monoclonal antibody design and development, natural and synthetic

• Wearable sensors within situ data analytics • Multiplexed, point-of-care, • fieldable devices • CWA countermeasures	• Wearable sensors within situ data analytics • Multiplexed, point-of-care, fieldable devices • CWA countermeasures
Development of high-through put screening tools for R&D applications. This may include tools for: • Biomarker identification, verification, validation • Epigenetics • Small molecule-inhibitor/activator screens • Lab automation	Integration of multi-omics technologies or methods for field-based or *in situ* detection of pathogens or biomarkers.
Fundamental research on materials, biomaterials, or biosynthetic polymers with biomedical applications, including research on implanting biological signals into synthetic polymers.	Advanced development of fabrication methods for nano- and microfabrication of medical devices or therapeutic delivery systems, or implantable bionics.
Development of techniques for 3D printing of biological materials that can replace human components.	Application of techniques for 3D printing of biological materials that can replace human components.
Development of medical implants technology through optimization of material characteristics and function.	Application of brain-computer interfaces and neuroprosthetics to enhance performance/function.
Development of brain-computer interfaces and neuroprosthetics.	

Neuroscience-based human brain mapping studies or methods.	
Implementation of AI systems using HPC to inform clinical decisions and optimize health.	

Agricultural & Environmental Technologies

Keywords: 3-D printing of foods, bioenergy crop, bioremediation, cellular agriculture, field trial, genome-wide associations, genotype, germplasm, marker-assisted plant breeding, microbiome, phenotype, propagules, robotics, seeds, transformation, unmanned aerial vehicle (UAV) hyperspectral detection, vertical agriculture.

GREEN	YELLOW
Research on publicly accessible germplasm or germplasm intended to be made publicly available, i.e., research on component or system validation at Laboratory-scale, to establish proof of concept, and to understand basic biological mechanisms.	Technology demonstration on germplasm at the engineering/pilot scale in a relevant environment, including field testing.
Publicly available in vitro regeneration and propagation techniques. Techniques to manually introduce and propagate recombinant DNA in plants that significantly advance the state of the art.	Demonstration and deployment of high-efficiency, high-throughput, crop transformation technology at pre-commercial scale.
Analysis of samples from laboratory scale systems or of devitalized samples from research field trials and pilot scale processing.	Pilot-scale development of processing technologies for bioenergy crop species.
Plant breeding technologies for specific traits including the use of Genome-Wide Association Studies and the development of marker-assisted breeding techniques.	Collection and mining of large data sets (representing large populations) for traits of commercial interest.

Determining the function of individual genes and gene combinations and/or linking genotypes to phenotypes using germplasm.	
Emerging technologies, such as: • "Cellular Agriculture" for protein production • 3-D printing of foods • Vertical agriculture • Microbiome research to support plant health Biosensors and plant sensing, such as: • UAV hyperspectral measurement of plants in an open research environment Genetically encoded biosensors for natural stimuli (e.g. drought and nutrient stresses or pests)	Of High TRL stage development of emerging technologies for commercial or national security applications. In-field applications of plant or microbial biosensors for natural stimuli at high TRL, with commercial potential.
Bioremediation, bioleaching, biofertilizers, and biopesticides at pre-commercial, low TRL scales.	Deployment of High TRL bioremediation or bioleaching processes that enable a commercial process with toxic products.
Water delivery systems or hydrodynamics for vertical agriculture at pre-commercial scales. [Tentative]	In-field deployment of high TRL hydrodynamic systems with commercial potential. [Tentative]

Bioscience & Biotechnology Applications of Artificial Intelligence

Keywords: Toxin design, pathogen design, novel threat creation, stability enhancement, capability uplift, generative design models, biological design tools, high throughput screening datasets, large language model, in silico pathway optimization, DNA synthesis screening invasion, model leakage, information hazards, adversarial prompts, data poisoning, AI risk evaluation frameworks, proxy testing, misuse.

GREEN	YELLOW
AI-accelerated vaccine/antiviral discovery (e.g., epitope prediction, generative molecule design), or AI-accelerated development of chassis organisms for biomanufacturing and/or medical countermeasures for public health pathogens.	Validation of biological models and specialized datasets linking genotype to phenotype in elite commercial germplasm with significant near-term commercial value.
Development of scientific benchmarks to evaluate AI capabilities.	Development of risk related benchmarks to evaluate AI capabilities.
Utilization of AI for the discovery and analysis of novel chemical reactions or biochemical pathways.	Application or integration of AI-enabled systems for the discovery and analysis of novel chemical reactions or biochemical pathways with significant dual-use applications in an environment representative of an operational or production context.
AI-enabled pathogen-detection, genomic surveillance pipelines and imaging diagnostics.	Integration of AI-enabled pathogen-detection pipelines and imaging diagnostics with significant dual-use applications in an environment representative of an operational or production context.
Development of “auto grading” AI systems for AI capability evaluation.	Development of “auto grading” AI systems for AI risk evaluation.
	Development of safeguards for AI systems focused on misuse in the context of national security.

Development or utilization of AI to automate routine taskings such as scheduling, laboratory information management system management, data clean-up, writing software/coding for biological instruments, or Interfacing with and automating usage of bioinformatics software. Development of methods to integrate AI into wet laboratories to increase efficiency of scientific research.	Development of security methods for AI enabled automated wet labs.
Use of AI for scientific knowledge synthesis such as literature review and integration of information from multiple sources.	Development of AI methods with dual open literature and surveillance applications.

Accelerator Science and Technology

The current breadth of particle accelerator technology began in the 1930's and continues to see dramatic advancement of capabilities through the development of novel concepts and technological improvements. Since their conception, accelerators have been instruments of discovery and scientific advancement, evidenced by the fact that almost one third of the Nobel Prizes granted in physics have been connected to the interaction of accelerated particles, particle detection, or the advancement of accelerator technology. Over the course of 90 years of internationally cooperative and competitive development, machine energies have increased by a factor of 10,000,000. This technology has many important applications in medicine and industry, with more than 40,000 accelerators operating around the world and total accelerator sales of about U.S.$5B [Industrial Accelerators and their Applications," "R. W. Hamm and M. E. Hamm, World Scientific, 2012, https://doi.org/10.1142/7745]. Advanced accelerator technology continues to be developed in the U.S., largely at DOE National Laboratories and at universities, with important anticipated applications in scientific research, industrial and medical processes, and national security. Accelerator technology plays a crucial role in many scientific areas where the U.S. currently has significant technological advantages over other countries. Achieving the right balance of open participation in international development efforts, and discretion in disclosing select aspects of application-related knowledge enables the U.S. to maintain its advantages and continue to realize the commercial and intellectual-property based economic benefits.

The accelerator science & technology section of the S&T risk matrix provides guidance in the areas of superconducting radio frequency technology, laser and plasma wakefield acceleration, superconducting magnets, cryogenic plant design and operation, advanced light source technologies, very high average current beam technologies for electrons and hadrons, and accelerator based nuclear systems and isotope production. When accelerator capabilities are used to contribute to another area of S&T, guidance from those areas is also applicable in determining what topics, if any, are restricted. For all topics below, information already protected by either classification or export control (or, in some cases vendor-specific non-disclosure agreements) is not reflected here, but scope should be evaluated relative to these controls on a case-by-case basis prior to release or dissemination.

Superconducting Radio Frequency Technology

GREEN	YELLOW
Foundational or basic research (corresponding to TRL 1-4) in superconducting radio frequency (SRF) and its applications. including: SRF accelerators for future high energy physics (HEP) and nuclear physics (NP) colliders, neutrino facilities, and rare isotope research. Generic SRF research on niobium surface modification, new materials, coating and processes toward achieving higher Q and/or gradient, unless performed specifically for a project/application. Generic cryomodule engineering and cleanroom assembly techniques. R&D on narrowband SRF cavity resonance control. Research and development on conduction-cooled ($T = 4$ K or greater) SRF cavities and systems. Development of SRF cryomodules using high-Q and/or high-gradient cavities: for CW applications Q_0 up to 10^{11} at intermediate gradients of up to 30 MV/m; for pulsed applications Q_0 up to	Development of SRF cryomodules using high-Q and/or high-gradient cavities: for CW applications $Q_0 > 10^{11}$ at intermediate gradients of > 30 MV/m_(for elliptical cavities, or corresponding peak surface magnetic fields for non-elliptical cavities); for pulsed applications $Q_0 > 5x10^{10}$ at gradients of > 65 MV/m Conduction-cooled (T=4K or greater) SRF cavities and systems at TRL 5 or TRL 6 for near-market/precommercial industrial/medical/security applications.

5×10^{10} at gradients of up to 65 MV/m	

Laser and Plasma Wakefield Acceleration

GREEN	YELLOW
Basic research thrusts, including: Wake accelerating structure excitation and particle injection, beam control. Basic research with physical concepts being developed and used in the research community, potential for long term market in accelerator applications. Plasma target gas valves and capillaries ionized by lasers or >100 ns discharges, systems are being built with commercial components. Small niche markets in the accelerator community; some potential for future licenses. Diagnostics including for radiation sources and positron production. Basic concepts are being developed and used to measure and support area (2), similar status. Applications: High Energy Physics Colliders, Thomson MeV photon sources, and Free Electron Laser type coherent photon sources. Basic research with physical concepts being developed and used in research community, potential for long term markets (>ten y away) in accelerator applications and markets in scientific, imaging & detection (nonproliferation, industrial, medical).	Integrated systems beyond the laboratory proof-of-principle stage of development with application potential in near term markets (<five y away, and at TRL 5-6)

Superconducting Magnets

GREEN	YELLOW
Basic research thrusts, including: Generic research on superconducting materials and superconducting magnets towards higher fields, unless performed for a specific application. Development of diagnostics and the associated signal analysis that provides feedback to magnet design. Development of superconducting materials for high-field magnets, and their performance characterization The basic development of improved superconducting magnet fabrication techniques. Development of high and very high field accelerator magnets and others needed for basic science applications.	Certain systems beyond the laboratory proof-of-principle stage of development, with commercial application potential in near term markets (<five years) and at TRL 5-6), as follows: Advanced superconducting materials with high transition temperature (>10K) and high critical field (>15T), including material science and processing optimization, intended to enhance conductor performance. High field magnet technologies (>12T) with identified commercial applications – for example, high-field nuclear magnetic resonance solenoids. Advanced design tools (non-open source), including modeling techniques, that enable advancement of magnet designs.

Cryogenic Plant Design & Operation

GREEN	YELLOW
Basic research thrusts, including: Studies of materials properties at low temperatures, design optimization of cryo-cycles and cryo-equipment.	None currently identified that is not otherwise protected

Advanced Light Source Technologies

GREEN	YELLOW
Basic research, long-term R&D, and proof-of principle studies on fundamental light source technologies and compact accelerator technologies. 4th generation storage ring design tools Fast, high power pulse electronics (<10 ns rise and fall time) and associated kicker magnet concepts. General designs and beam physics concepts for storage rings. General designs and beam physics for linac-based (including energy-recovery) light source designs (including bright-beam photoinjector designs, advanced beam manipulations). High accuracy and bandwidth electron beam position monitors and electronics.	Integrated designs of future generation storage rings, as well as linac based light sources. Application/project specific TRL 5 or TRL 6 design demonstrations of original advanced beam physics concepts that lead to advanced performance and capabilities for storage rings, free electron lasers (FEL), and ultrahigh brightness storage rings. New ultrabright electron photocathodes to optimize the performance of FEL facilities at the TRL 5 or TRL 6. Developing next generation superconducting undulators, including conduction cooling technology, at TRL 5 or TRL6. High performance X-ray optics for X-ray free electron laser oscillators at the TRL 5 or TRL 6 level of demonstration.

Very High Average Current Beam Technologies: Hadrons

GREEN	YELLOW
Basic research thrusts including: CW and pulsed H- and proton sources for injection into accelerators. Development of adaptive and machine learning techniques toward autonomous control of	Systems designs of adaptive and machine learning techniques beyond the proof-of-principle stage providing autonomous control of accelerators. Engineered development beyond the proof-of-principle state of technologies required for proton

accelerators to improve their overall reliability.	beam power >five megawatts (MW) and <10 MW.

Very High Current Beam Technologies: Electrons

GREEN	YELLOW
Basic research and development on very high current electron beam enabling technologies. This encompasses areas such as electron source technology, SRF linear accelerators (LINAC), energy recovery LINACs (ERLs), storage-ring concepts, phase-space manipulation (including electron cooling) techniques and engineering. This includes exploring disruptive ideas, conceptual studies and designs. The deliverables of these R&Ds are at Technology Readiness Level up to 4	Applied research and development on very high current electron beam technologies for national security and advanced industrial applications at Technology Readiness Level from 5 and 6. Examples of industrial applications include advanced metallurgy, additive manufacturing, and high throughput lithography. Enabling technologies include electron sources, energy recovery LINACs (ERLs), beamlines for lossless high-current-beam transport.

Accelerator-based Nuclear Systems

GREEN	YELLOW
None currently identified that is not otherwise protected	None currently identified that is not otherwise protected

Accelerator-based Isotope Production and other Industrial and Medical Accelerators

GREEN	YELLOW
Foundational accelerator physics and radiochemistry research of potential relevance to isotope production but for which target and separations optimization has not yet begun.	Technology development of production and processing of isotopes with application potential in near term markets (<five yr away). Novel or compact accelerator technologies at TRL 5 or TRL 6 with the potential of commercial

Foundational or basic research of compact accelerators. Foundational and basic research both correspond to TRL 1-4.	advantages for industrial or medical applications or with the potential of dual-use national security applications.

Appendix 1: Associated DOE Orders and Policies Involving the S&T Risk Matrix

Order 142.3C – Unclassified Foreign National Access Program

A determination of access approval is required before each access request is granted and must ensure that any identified risk to the Government associated with the access granted has been appropriately evaluated and mitigated, including a review against the S&T Risk Matrix.

Country of Risk foreign nationals' access to certain restricted technology or information as identified in the S&T Risk Matrix must undergo an enhanced review process (see below).

Requests for access to conduct research under a User Agreement at an Office of Science, Critical Minerals and Energy Innovation, or Nuclear Energy User Facility at a non- NNSA DOE Laboratory are exempt from the review requirements related to the S&T Risk Matrix discussed below.

1. A request for access is reviewed by the site, in coordination with the head of the cognizant DOE Field Element, to determine if the access request is in an area identified as restricted in the current S&T Risk Matrix before submitting the proposed access request through the standard access request review and approval process promulgated by this Order. When an access request is from a Country of Risk foreign national and in an area identified as restricted in the current S&T Risk Matrix, the cognizant DOE Field Element must agree to proceed with the enhanced review process. Absent this agreement the proposed access request is not pursued any further.
2. When the Head of the cognizant DOE Field Element agrees to proceed with the enhanced review process, a justification and clear description of why the access request benefits the U.S. must be prepared. The access request must then be submitted through the appropriate Program Secretarial Office (PSO) and Cognizant Secretarial Officer (CSO), with final approval/disapproval being provided by the cognizant Under Secretary or their designee.
3. Completion of specialized enhanced vetting conducted by the DOE Office of Intelligence and Counterintelligence (IN) is required prior to final approval of the access request. A copy of the request, along with the required additional information, must be submitted to the cognizant local counterintelligence office to support the review. Indices checks will be conducted as part of the specialized enhanced vetting process; therefore, it is recommended the request be submitted 45 days prior to the start date of the access request.
4. DOE may consider broad approvals for specific categories of these types of access requests, such as those supported under government-to-government agreements and in line with National Security Council policy guidance, to ensure existing priorities are not unduly impeded.

Order 241.1C – Scientific and Technical Information Management

The Contractor reviews S&T Information generated under the contract to determine appropriate release and handling and apply any necessary statutory or program-driven announcement and/or availability restrictions, including those related to nonproliferation, national security, export control, intellectual property, protected Personally Identifiable

Information and privacy. In addition, the Contractor must apply to the STI product any restrictive markings required, include any required legal disclaimers, and, for STI products resulting from DOE-funded work, identify the sponsor as follows: U.S. Department of Energy, [name of DOE program office], [name of DOE subprogram]. For research that is determined to be a Restricted S&T topic, this should include a notification to the funding DOE Program Office prior to publication.

Order 481.1E – Strategic Partnership Projects

The S&T Risk Matrix must be reviewed for each proposed engagement with a foreign entity from a Country of Risk to determine if the engagement is in an area identified as restricted. Project exemption requests for foreign-sponsored work with entities from Countries of Risk in areas identified as restricted in the current S&T Risk Matrix must be submitted to the cognizant Under Secretary or his/her designee. Exemptions must be approved by the cognizant Under Secretary or his/her designee prior to initiating a review of foreign-sponsored work under DOE P 485.1A.

Order 483.1B – DOE Cooperative Research and Development Agreements

Review the current S&T Risk Matrix for each proposed foreign Cooperative Research and Development Agreements (CRADA) project with a foreign entity from a Country of Risk, to determine if that project is in an area identified as restricted. Proceeding with such a CRADA requires an exemption request through the field element, CSO, and PSO for cognizant Under Secretarial approval.

Policy 485.1A – Foreign Engagements with DOE National Laboratories

Restricts DOE National Laboratories from conducting foreign engagements with Countries of Risk in the scientific and technology areas identified as restricted in the current S&T Risk Matrix unless an exemption is granted by the Department.

Order 550.1 – Official Travel

Requires a review of the S&T Risk Matrix, for all proposed official foreign travel to a Country of Risk to determine if the travel involves areas identified in the S&T Risk Matrix as restricted.

DOE Order 471.7 – Controlled Unclassified Information

While this Order does not specifically invoke the S&T Risk Matrix, it provides the requirements for how restricted information identified within the S&T Risk Matrix, which is designated as Controlled Unclassified Information Basic, is properly safeguarded.

Appendix 2: Broader Context for the Research Areas Described in the S&T Risk Matrix

Quantum Information Science & Technology

State of the Art in Quantum Information Sciences and Definitions

Quantum technology uses physical realizations of quantum bits ("qubits"), idealized, two-level quantum systems, to sense, process, and transmit information. Qubits have two important characteristics that distinguish them from their classical analogs: they can form superpositions of the 0/1 values, and different qubits can be entangled with each other in ways that lead to non-classical, long-range correlations. There are many physical systems that can be made accurate enough to serve as qubits, but the most common ones are made from superconductors, semiconductors, trapped ions, trapped atoms, and photonic devices, or combinations of these.

The requirements for practical quantum computation were described by David DiVincenzo and coworkers in several pioneering articles (see D.P. DiVincenzo, Progress of Physics 48, 771 (2000) for a summary), and state that a working quantum computer must satisfy:

1. A scalable physical system with well-characterized qubits,
2. The ability to initialize the state of the qubits to simple fiducial states,
3. A universal set of quantum gates,
4. Sufficiently long decoherence times for gate operation and state evolution,
5. The ability to measure qubits.

Additional requirements for quantum communications involve "flying qubits," i.e. qubits that can transport quantum information across distance:

6. The ability to interconvert between stationary and flying qubits,
7. The ability to transmit flying qubits across distance.

Practical quantum computation involves preparing several qubits in some initial state, applying a variety of one- and two-qubit operations to the collection, and measuring the resulting state. Whereas classical computing devices can be made almost completely error-free for practical purposes, quantum devices are susceptible to many different noise processes that make implementing practical computing platforms from them very challenging. Several physical measures are used to characterize the quality of qubit systems. The decoherence time is the amount of time that a device takes for one value to relax into another value, for example, the time for a "1" state to turn into a "0" state. The decoherence time is shorter for systems that interact strongly with a noisy environment; it can range from nanoseconds to hours in extremely pure systems. The fidelity is the accuracy at which various operations can be applied to the qubits.

Progress toward developing realizations of quantum computing technologies can also be characterized in terms of error rates. The error rate measures the accuracy with which various operations can be applied, and it is routinely used to evaluate the accuracy of storage operations, gate operations, transmission operations, and measurement operations applied to qubits. Varying methods are available for estimating the error rate in experimental settings. Additionally, the capacity of the quantum technology register for storing quantum information, frequently referenced as the number of qubits, is an important metric. Breakthroughs toward large-scale quantum technologies are expected to be indicated by significant increases in qubit and gate fidelities, significant decreases in error rates, and significant increases in register capacities.

Alongside these basic metrics, high-fidelity operation of quantum technology is expected to require quantum error correction (QEC) schemes. Like their classical counterparts, QEC schemes use the redundant encoding of quantum information in multiple physical qubits to represent an individual logical qubit. The encoded logical qubit can then be corrected by monitoring for errors using syndrome measurements, feed-forward measurement, and other methods. While this approach often incurs a substantial overhead in the number of physical qubits and physical gate operations required to perform a calculation, QEC affords the benefit of storing and processing quantum information with sufficiently lower errors. The corresponding logical error rate (LER) and the associated logical gate fidelity are important composite metrics for monitoring progress toward the development of large-scale quantum technologies and their applications.

Presently, the DOE is investing in multiple programmatic areas and in five large National Quantum Information Science Research Centers to vigorously pursue QIST R&D in conjunction with universities and industry. The scale of these efforts, and their multiplier effects, should bring about a faster realization of quantum technologies with the requisite characteristics to realize practical and real-world applications. Additionally, industry has initiated robust R&D programs to scale up and integrate quantum technologies for delivering systems, products, and services.

Application areas where quantum technologies may have dramatic impacts emphasize the efficiency in the resources required to solve a problem, perform a task, or complete a calculation. For example, an original motivation for the development of quantum computers came from Richard Feynman's observation that building computers from devices that obey quantum mechanical laws could execute algorithms that have exponential reductions in the time-to-solution and memory requirements for computing properties of other quantum mechanical systems. Related applications of quantum computing are now termed quantum simulations. Additional algorithms have been developed that promise exponential speed ups for a variety of problems across many different scientific domains. Quantum algorithms with exponential and polynomial speed-ups have also been developed for various algebraic problems, including factoring integers.

The gap between current capabilities and expected requirements for meaningful applications has been deemed the noisy intermediate-scale quantum (NISQ) regime. The NISQ regime is characterized by quantum technologies with small register capacities that lack quantum error correction but may use other error mitigation techniques to ameliorate performance.

A major challenge is the identification of algorithms that can run effectively in the NISQ regime, and there is a broad search in the scientific domains of materials, condensed matter systems, high-energy physics, chemistry, nuclear physics, computational physics, machine learning, and optimization for applications.

Improvements in quantum computing technology, including register capacities, gate fidelities, and error rates, are expected to enable a new regime of quantum computing using early QEC operations. The demonstration of QEC at small scales may still not meet the requirements for logical error rates (LERs) and logical gate fidelities necessary to sustain large-scale applications of quantum technologies. Instead, this early QEC regime continues to emphasize exploratory applications of quantum computing using limited resources.

Computing and Simulation

The paper *Elucidating Reaction Mechanisms on Quantum Computers* [Reiher, Wiebe, Svore, Wecker, & Troyer, Proc. Natl. Acad. Sci. USA 114 (29) 7555-7560 (2017)] outlines some of the resource requirements necessary to compute problems of interest to the chemical or pharmaceutical industries using quantum computing. The authors estimate that both many qubits (over 100 for the case they consider) and qubits capable of operating at high fidelity (error rates below 10^{-13}, likely requiring quantum error correction) are required to have potential to these industries. More recently the same group, [Burg *et al.*, Phys. Rev. Research 3, 033055 (2021)] described resources needed for computing catalysis. Burg *et al.* estimated a need for between 3,700 and 7,400 logical qubits to successfully calculate catalysis energy levels at the mHartree level (that is, with chemical accuracy). A fault-tolerant logical qubit consists of hundreds to thousands of physical qubits. Thus, over a million physical qubits would be required to perform the catalysis calculation. Eventually, we may see a breakthrough that will push us toward this region. In the meantime, NISQ computation coupled with error mitigation is the name of the game for quantum gate computation.

An alternative to the quantum simulation algorithms described so far is analog quantum simulation, where the evolution of one quantum system is used to determine properties of a related Hamiltonian system. One recent application involves a tour deforce 219-qubit simulation of topological states in quantum liquids [Semeghini *et al.*, Science 374, 1242 (2021)]. Analog simulation of this type can provide a more efficient method for simulation of certain physical systems than techniques implemented on classical computers, although the mapping involved is not always realistic, and it is often challenging to quantify the uncertainty or model error in the answer provided by the analog simulation.

A quantum sensor is a device that uses quantum effects to achieve sensitivities higher than can be achieved using purely classical methods. Quantum sensors have already been used to make more accurate atomic clocks, magnetometers, and inertial sensors. As just one example in this area, progress has been made to develop quantum sensing for arbitrary-frequency fields using a quantum mixer [Wang *et al.*, Phys. Rev. X 12, 021061 (2022)], as demonstrated in nitrogen-vacancy centers in diamond. These techniques can provide extreme sensitivity in a variety of environments.

There are also techniques that use quantum superpositions and entanglement to develop communications technologies. These have been used in quantum key distribution to distribute encryption keys in a way where eavesdropping can be detected. Additional techniques have also been developed that use entangled photons to teleport quantum states across distance, or otherwise, couple quantum processors or sensors across a distance. One of the key developments in this area will be in non-linear optics which should enable the development of efficient quantum repeater technology.

As of 2025, quantum technological advances have led to the two-qubit physical gate fidelities of 99.9 percent demonstrated in multiple technologies. Conversely, this fidelity corresponds to a physical error rate of 10^{-3}. Because of the multiplicative error model in quantum technology, executing a calculation with a modest 700 two-qubit operations and a physical error rate of 10^{-3} per operation yields the correct result with a probability of less than 50 percent. Current estimates on the number of two-qubit operations required for large-scale calculations exceed 10^{6} two-qubit gates, placing current technology well above the physical error rates necessary for breakthrough advantages. Additionally, as of 2025, multiple quantum technologies have surpassed the 1,000-qubit threshold. Estimates of register capacity required for large-scale calculations is estimated to exceed 1,000,000 qubits, placing current state of the art three to four orders of magnitude below requirements. Additionally, early demonstrations of LER remain well above the LER requirements of 10^{-12} or lower.

Removal of 'Materials and Fabrication' from the 2025 S&T Risk Matrix

In the previous version the QIST section of the S&T Risk Matrix, fundamental R&D in materials and fabrication was denoted "green" across a wide range of material classes, from superconducting to semiconducting, 2D materials, topological insulators, quantum materials, and isotopes. "Green" research also included theoretical materials studies and simulations, as well as rapid characterization techniques and advanced fabrication approaches. Topics in the "yellow" and "red" categories were those activities that led to advanced devices technologies, largely in computing. There was nothing fundamental about materials or fabrication R&D that was inherently "yellow" or "red," and for that reason it has been removed as a stand-alone section in the current version of the Risk Matrix. Fundamental materials or fabrication R&D should continue to be considered "green," and when advances in this area lead to a "yellow"- or "red"-level technology (be it for computing, sensing, clocks, metrology, or communications), it should be considered for additional protection. But those metrics that determine categorization should stem directly from technology performance, and any advancement or breakthrough stemming from or largely due to materials and fabrication should be protected as such.

Removal of Technologies that are Export Controlled from the 2025 S&T Risk Matrix

New export controls for quantum technologies were introduced in September 2024 (https://www.federalregister.gov/d/2024-19633/p-47) that attempt to address the national security aspects of quantum technology. These include, but are not limited to:

- Quantum computing systems with 34 or more physical qubits with high enough two-qubit gate fidelity on a sliding scale based on system size,
- Quantum computing systems with 2,000 or more fully controlled, connected and working physical qubits,

- Materials for certain types of qubits, such as epitaxial materials with enhanced percentages of certain silicon and germanium isotopes,
- Cryogenic CMOS controls,
- Parametric amplifiers designed to operate between 2 GHz and 15 GHz, and
- Cryogenic cooling systems with more than 600 μW of cooling power.

The previous version of the QIST section of the S&T Risk Matrix included qubit-count metrics for systems size and controls on large cryogenic cooling systems, as well as some controls on amplifiers and electronics. The 2025 version removes redundant controls and attempts to align boundaries between sensitivity categories, where feasible, with the export control guidance.

Suggestions for Sensitivities for Topics Relevant to QIST

Here we divide technology areas relevant to QIST into "green", "yellow," and "red" categories. In general, for a technology to show sufficient potential to label it red requires a demonstration of this potential, rather than simply designing experiments or postulating their possibilities. Research and development efforts oriented towards incremental improvement for technology presently in the green category would typically be green also, but R&D efforts promising transformational, order-of-magnitude improvement in key metrics for some presently green technologies might be considered yellow; the same might hold for the relationship between the yellow and red categories. Research in QIST technology areas should be monitored at regular intervals to track potential movement across the green to yellow category.

An important metric for monitoring progress in quantum computing is the Total Logical Error (TLE), which is defined as the error of a complete quantum computation, i.e., a measure of accuracy for the outcome of a quantum computer computation. Assuming independent multiplicative errors, define $\text{TLE} = 1 - (1-\text{LER})^{\text{NLG}}$ and, conversely, the required logical error rate as $\text{LER} = 1 - (1 - \text{TLE})^{(1/\text{NLG})}$, where NLG is the number of logical gates.

Activities for NISQ and early QEC scientific applications are anticipated to require a TLE ranging from 10^{-1} to possibly 10^{-6}. As an example, chemical simulations with an accuracy of 10^{-6} are sufficient for scientific applications. Considering $\text{NLG} = 100^3$ and a TLE of 10^{-6}, the threshold sets a requirement for the LER as 10^{-12} for each logical 2-qubit operation. By contrast, highly accurate applications, including factoring integers, are expected to require a TLE of 10^{-12} to 10^{-16} [Gidney and Ekerå, Quantum 5, 433 (2021)], and a corresponding LER of 10^{-21} or lower. The significant difference in required LER and TLE creates a technological divide between scientific applications in the early QEC regime and large-scale fully fault-tolerant applications of quantum computing.

When considering the G/Y/R classifications for quantum computing, we must consider that current technologies have not demonstrated either fidelity or scaling that would yield these limits. Thus, the field of quantum computing remains wide open to scientific exploration, except in a few cases. Today's NISQ computers currently exhibit all the characteristics of a hard physics experiment. Nevertheless, over time, as continuing research propels this area forward, we could see technical advances that will push 'green' items to 'yellow' and 'yellow' items to 'red'.

High Performance Computing

International Landscape in HPC

Many nations have active research and development programs in high performance computing. At the highest system size and in breadth of use, the U.S. has long been the unquestioned leader in the field. This leadership has come from the research, development, and deployment activities at the DOE National Labs, along with the dominant role played by American technology firms including IBM, Intel, Nvidia, AMD, HPE, and their predecessors. Western Europe has been a major contributor and partner to the U.S. in research and development. In the 1990s, Japan made large investments in HPC and developed and deployed several of the fastest computers in the world. Japan continues its investments, as demonstrated by the strong performance of the Fugaku system encoded in the Top500 results from 2020 through today (2025), including having had the #1 system on the TOP500 list from June 2020 through June 2022. China has demonstrated a strategic commitment to HPC and has developed an indigenous technology base that resulted in machines that were, at times, faster than any in the U.S. The June 2025 Top500 list contains several strong entries from China, even while the U.S. dominates the very top of the list. At the top of the list are systems reserved for some of the largest calculations, representing distinctive national capabilities, and a realm where DOE leadership is apparent. Among the top ten systems on the June 2025 list are five U.S. systems, and the top three are at DOE Laboratories.

The 2025 HPC focus is on exascale systems and their immediate successors – machines capable of computing 10^{18} double-precision numerical operations per second. The U.S. was the first to reach this milestone with the Oak Ridge National Lab (ORNL) Frontier system operating at 1.353 Exaflops/s. In addition to efforts to increase the speed and scale of these traditional approaches to computing, the increasing focus on the utility of a converged set of solutions incorporating HPC, quantum computing, and AI is producing a level of complexity that has not been present in past. Motivations and program details vary, but all of these efforts foresee value in basic science, industrial competitiveness and national security. As the Council on Competitiveness has argued, to out-compute is to out-compete.

A study by the scientists in DOE and the NSA examined the potential consequences to the U.S. of falling behind a strategic rival in advanced computing. They identified several concerns including the following.

- HPC is a tool for deep insight. If an adversary fields a weapons system with new capabilities based on superior HPC, and the U.S. cannot accurately estimate its true capabilities, there is a serious possibility of over- or under-estimating the threat. The adversary might also be able to gain insight that we lack into our own military systems.
- The U.S. enjoys relatively cost-effective HPC due to our indigenous ecosystem and trusted access. If we were instead to rely on systems designed and built by an adversary, our cyber and supply-chain risks would be unacceptable. If instead, the U.S. Government were to require indigenous procurement without a healthy domestic ecosystem, the costs could be very high.

Arrayed against these dire concerns are the costs of an isolated national academic and industrial base in HPC. The HPC community has benefited enormously from the efficiencies of open-source software for systems, libraries and scientific applications, and integration of commodity computing components into HPC systems at a broad range of scales from workstations to departmental servers, to leadership class systems. As commercial leaders in the field, our computer vendors have profited greatly from an open international marketplace.

Approach to Managing HPC Technologies

Existing laws and regulations on export control and classified work already offer some protection for HPC technologies, and the guidelines developed here should not overlap with those regulations or policies. In the discussion below, considerations that may govern the dissemination of technology will be noted, although these may lie beyond the scope of this matrix. The Labs have well-established processes for managing access to export-controlled data and have successfully used those for partnerships with HPC vendors in programs like PathForward. In this case, the technology originates with a vendor, is protected by a non-disclosure agreement (NDA) that covers one or more Labs or (when Export Controlled information may be shared) protected by a more Restricted Secret Non-Disclosure Agreement (RSNDA) that covers named individuals at the Labs. Classification guidelines may also be used for HPC methods and software within clearly defined programs that have strict access controls. But much of the work at the Labs is done with the intent to publish with unfettered international collaborations and is often performed under a Fundamental Research Exemption to Export Control. The purpose of this document is to propose a management framework for HPC technology that has the potential to significantly impact economic and international competitiveness but is not covered by current Export Control or classification rules.

Facilities and Systems

The DOE Labs operate several large computing facilities, ranging from commodity turn-key systems to leadership class systems that push the boundaries of what is possible in terms of system size, complexity and capability. Systems on the high end of the TOP500 list are designed to run applications at scale and offer capabilities not available through other sources. Historically these have been built out of commodity components although features have been highly tailored to the needs of large-scale systems. The individual components will be covered in the next section; here we are concerned with the specific aspects that become relevant when procuring and operating systems at scale. The difficulties of assembling an HPC system include the scalability of interconnect hardware and communication software, node stability, failure management, and system administration on many nodes. There is a distinction between systems to support commercial cloud computing, and very large, high-performance systems that enable tightly coupled applications to run across the entire system.

The operation of computing and networking facilities is managed by an Authority to Operate (ATO) approval from DOE. This takes into consideration the types of data/software on the systems, whether it is classified, export controlled, proprietary, private (PII or PHI), or entirely open. As with Export Control and Classification, it is important that new guidance to manage HPC technologies does not redefine or contradict this established and tailored process. The summary below describes some general principles for access to and operation of HPC facilities, with the intent to leave specific details to the ATO approval process, which is done on a per Lab basis, often with a separate network enclave as well as a high-performance wide area network.

User Access

Large compute time allocations on some of the largest HPC systems can be used for competitive advantage, even if the computations are not classified or sensitive and do not reveal details of the underlying system that would allow the system to be replicated. Systems near the top of the TOP500 list (at least the top 25) tend to be configured to run tightly integrated jobs across a significant fraction of the system, so cyber security and access plans for such facilities should receive special attention. While the users are neither physical visitors nor employees, an access management plan should involve multi-factor authentication and some level of background check/screening for foreign nationals. All DOE HPC systems currently operate under such a model.

Systems Administration

HPC Facilities store application software from users in source and binary form, data produced by their applications, operational data on system usage, network accesses, and more. The sensitivity of this data depends on the mission and policies of the facilities, i.e., whether data may be classified, export controlled, or otherwise restricted. Some facility personnel require administrator access to the systems, and thus access to all the data, while others require no more access than a facility user. Facilities may therefore restrict systems administration personnel further to match the sensitivity of data on the system.

Cybersecurity Operations

Cybersecurity positions within these facilities are another special case, since cyber security data and plans may involve sensitive information about cyberattacks or vulnerabilities. Each facility implements cyber security based on a combination of best practices, site specific requirements, and higher-level policies. Cyber security controls are responsible for ensuring protection against unauthorized or hostile individuals and software. Cyber security is a combination of tools/applications that monitor usage or block access, and processes driven by internal and external policies, possibly in response to cyber events. Cyber security data and plans may involve sensitive information about cyberattacks or vulnerabilities and should therefore be considered sensitive.

Procurement activities and vendor partnerships

Design and development of the largest HPC systems begin years before delivery, and vendors therefore share their long-term roadmap, which can contain NDA and export-controlled information. In addition, the procurements often stretch the capabilities of vendors, and it is therefore common for procurement contracts to include allowances for non-recurring engineering (NRE) activities that the vendor might not otherwise pursue. The results of such NRE are IP specifically necessary to enable the largest scale HPC but also produce the ‘next smaller’ scale machine in a commercially viable way. Because procurement activities and contract management positions can access sensitive data, they are therefore subject to other controls (e.g. proprietary or export-controlled information).

A developing frontier includes the procurement of hybrid HPC-quantum systems, either with explicit coordination between the individual vendors or without. Because so many quantum hardware vendors are in the beginning stages of bringing their products to a wider market, the controls they institute around NDA and other IP-protecting procedures can be expected to lag those implemented by HPC vendors.

Components

Computer Company-Developed Technology R&D

Most HPC Component Technologies are based on a technology foundation of vendor-led hardware architecture R&D. In recent years DOE and the National Labs have influenced vendor roadmaps through co-design collaborations that have been facilitated by several rounds of "*Forward" computer vendor architecture R&D programs, and in some cases, follow-on leadership-class platform procurement contracts with their closely associated NRE technology development contracts. To facilitate these co-design collaborations DOE and several of its National Labs have established multi-party NDAs with each of the R&D computer vendors (as such, this information is not repeated on the risk matrix). The information associated with each of these vendor contracts is proprietary and often subject to additional export control restrictions, so some type of restricted "RSNDA" is often used. In some cases, information surrounding the topics and partners associated with these vendor R&D projects are further protected as "business" sensitive, and while not proprietary or export controlled, is still protected as For Official Use Only by DOE and its National Labs. The areas for HPC Component Technology architecture R&D investment have included:

- Advanced Processor Designs
- Memory Interfaces and processing in/near memory
- Interconnection Network Designs

In general, the level of protection for these HPC hardware component R&D projects start as *Red* and eventually transition to *Green*. In particular, the performance specifications are closely controlled and competition-sensitive before the computing components are generally available as a product. Because these components are targeting commercialization and may have national competitiveness implications, participation in co-design activities or technical oversight of R&D subcontracts is *Red*. After product release, performance specifications and the open interfaces (instruction set architectures, programming interfaces, etc.) are *Green,* although detailed chip designs or internal interfaces may remain *Red* if they continue to be covered by vendor NDAs. For DOE, it is important that performance characteristics and interfaces of HPC Component Technologies to be openly available, as is critical to the successful commercialization and associated market demand for new generations of commodity components.

Performance characterization of production HPC systems and components is an important research topic, whether there are open published performance information, and should remain *Green*.

National Lab-Developed Technology R&D

Success of the open-source model is not limited to software alone. There are several examples of open-source hardware that support a general Open Innovation framework, this includes Facebook's Open Compute Project, RISC-V (Reduced Instruction Set Computer, 5th generation), and Advanced RISC Machine (ARM) technologies. At the hardware architecture level, these Open Instruction Set Architectures like RISC-V are described as Free and Open Instruction Sets; whereas Open Hardware eco-system technologies like ARM are not free, but Open and licensable.

There are Lab-led architecture R&D projects that are leveraging either ARM or RISC-V Open Hardware eco-systems. These projects are targeting the opportunities for HPC-centric architectural specialization where co-design principles are applied to develop conceptual designs for specific computing components, e.g., processors, memory interfaces, interconnects, or storage interfaces. While these projects leverage Open Innovation frameworks, the conceptual designs may fall into one of two categories: exploratory high-level designs that are incomplete or have insufficient detail to be patented, and more concrete designs that, if proven to have sufficient value, may transition to chip level designs. As a result, hardware designs for HPC components or systems will be treated at *Yellow*, transitioning to *Red* if the design has reached a stage of maturity and promises that it will be held as DOE or Laboratory proprietary. For some cases the assessment of risk is straightforward, whereas in other cases advances multiple technologies must be simultaneously considered. For ultra-low-power computing, improvements in transistor efficiencies in the coming years should yield 2x improvements in efficiency that will become available on the open market. Power for Flash solid-state devices is increasing to enable higher bandwidths, somewhat countervailing this improvement. Any approach that yields an overall 5x improvement in efficiency will likely be the product of Laboratory R&D and should be protected. Ultra-high-density storage-class memories (i.e., non-volatile, byte-addressable memories such as Memristor, Resistive RAM (ReRAM), Phase Change Memory, Spin Torque Transfer RAM (STT-RAM)) and solid-state, block-addressable storage (e.g., NAND flash) are likely necessary for future beyond-Exascale HPC systems. Increases in density over 5x current technologies should be protected, as such improvements will likely be driven primarily by this use case, particular to the National Laboratories. The primary focus of research around enabling ultra-fast communications has primarily been on future high-speed optical hardware and associated software and standards. Any research that seeks to increase bandwidth or reduce latency by more than 5x currently available technologies would need to be protected. Should the design transition to a technology transfer and commercialization phase, the conceptual design will remain controlled under existing NDAs until the product development and NRE is completed, and the component design is successfully commercialized. At that point the HPC component use and interface is *Green*, while underlying details may remain controlled.

In addition to these specific estimates of evolving HPC ecosystem technologies, the development of Lab R&D on technology with the potential for orders of magnitude improvements in performance or energy efficiency, as measured on one or more applications of DOE mission or commercial interest, should also be labeled as *Red*. This may include Lab-led designs of non-traditional computing technologies, e.g., neuromorphic, analog, or approximate computing architectures that show sufficient promise that the designs should be protected. Evaluations of such technology, assuming the technology is open, will continue to be *Green*.

Platform Software

The landscape of HPC system software is complex and multidimensional, and includes Operating Systems and container software, job scheduling and systems monitoring tools, programming systems (compilers, runtime systems, communication software and tools for performance and correctness), input/output libraries, check point and restart software, data reduction, and visualization. In addition, workflow tools and distributed resource managers may operate across HPC systems, local and wide area networks, as well as other computing devices, systems, and instruments. Much of this HPC software ecosystem is based on open source and leverages the broader computing community—beyond DOE and HPC— especially for programming systems and node level software.

As a rule of thumb, it is unlikely that platform software developed or modified by the Labs would necessitate a *Red* designation as a standalone artifact. Indeed, the leverage enjoyed by HPC from software investments from other communities, as well as collaborations outside the Labs is essential to the health and ongoing progress in HPC software, so it is important that such software be designated *Green*. While some system software has direct access to low level hardware (e.g., the BIOS) or other essential functions of a system, these should be addressed as part of a cybersecurity plan rather than technology sensitivity.

Most of the system software running in production on high end systems is vendor proprietary, or updated and optimized software whose origin can be tracked to open-source packages developed in academia, national laboratories or by individual developers. Such proprietary vendor software could be in some cases export controlled. The source code for systems software may also reveal underlying hardware details, so in conjunction with systems containing proprietary or export-controlled technology the hardware-software stack would remain controlled in its entirety.

Application-Level Libraries and Frameworks

This section defines the grading of application-level libraries and frameworks but defers the grading of applications themselves to subject matter experts in the application domain. For example, the Accelerator Modeling technology plan would define how codes to design future particle accelerators will be graded or the AI/ML topic area would define how applications of those technologies will be graded.

Application-level libraries and frameworks are common software used for building multiple applications. Above the programming systems level (covered above), there are numerical libraries (e.g., solvers for linear equations, solvers for ordinary and partial differential equations, implicit and explicit time marching methods), other application-level frameworks (e.g., adaptive mesh generation, particle handling) and algorithms (e.g., fluid-particle calculations on hybrid CPU/GPU systems, load balancing).

DOE has a leadership role in application-level libraries and frameworks but also benefits enormously from collaborations with universities and other laboratories in the U.S. and abroad. The libraries and frameworks are largely open source, and on balance, the freedom of publication, exchange of ideas, and recruiting from a large pool of talent benefits the DOE mission and broader scientific goals. These open-source packages, which include hypre, PETSc, SuperLU, Trilinos, ScaLAPACK and LAPACK, AMReX, Chombo, and SAMRAI, are widely and openly available and will be graded *Green*. Lab employees that lead many of these efforts vet contributions to repositories, no matter the source of the contribution, a practice that is widely used by open-source communities and has proven very effective.

Software developed under non-disclosure agreements and domain-specific modules from classified subject matter areas will be governed by those legal or statutory restrictions. At the NNSA Labs, this is already the working model. For example, Trilinos is used with many application-specific components including domain-specific preconditioners, spatial discretization, and software for uncertainty quantification that contain parameters from export controlled or classified applications. These components never leave a secure network and are integrated at compile-time and runtime with the rest of the open source Trilinos APIs and components.

Edge and Hybrid HPC

The demands of connected scientific instruments, observational and sensor systems, and other technologies that produce and require the storage of substantial amounts of data at very high rates have pushed HPC concerns out of traditional data-center settings and into remote and non-traditional sites. DOE continues to be at the forefront of the development of these technologies, especially in the form of custom microelectronics designed to effectively process, filter, and transmit the resultant data, as well as the production and dissemination of reasonable communication and other standards. Where these activities intersect other technologies covered in various sections of the Risk Matrix (see, for example, the earlier mention of quantum science technology and the application of artificial intelligence in many of these solutions), care must be taken to identify places where HPC-related hardware and software development enables these other technologies to flourish in the context of a fully connected computational ecosystem.

Artificial Intelligence and Machine Learning

In the context of fundamental AI advances, the guiding principle is that steady and expected progression of techniques is better performed in the public domain to allow faster overall progress of the entire ecosystem. Controls should be considered in cases where significant, non-incremental changes are anticipated in topical areas in the matrix resulting from the availability of significantly larger data sets, faster training times, or other indicators that may portend an enduring techno-economic advantage. Similarly, capabilities should be monitored (and controls considered) when the integration of otherwise open information leads to the demonstration of capabilities controlled in these subject area domains. By design, the matrix does not consider security concerns, such as dual-use technologies or the impacts of incorrect or deliberately biased models.

The importance of judgement within specific situations is a tenet in the use of this reference document for AI and ML and the reviewers should take the following guiding principles and definitions into account:

- **Fundamental AI/ML**: The term fundamental AI/ML refers to all areas of relevant research including but not limited to new architectures, training algorithms, adversarial attack or defense, explainability and interpretability, safety, agentic systems, etc. Furthermore, given the speed of current advancements it is likely that new areas will become relevant in the future, such as in human machine alignment or interfaces, trust, or new applied mathematics concepts.

- **Order of magnitude**: In several instances the matrix refers to an "order of magnitude" improvement which is deliberately left to the reviewers to interpret. Depending upon the specific question or parameter, a factor of two improvements might be considered ground-breaking and enable unprecedented capabilities, while in other areas a factor of ten might be required for equivalent impact. Alternatively, the phrases "with significant economic interest" and "with acknowledged commercial interest" are also used to indicate important increases in performance. The reviewers are asked to use their best judgement to distinguish between a large but ultimately incremental advancement in any area and a potentially disruptive and economically significant improvement in capabilities.
- **Emergent capabilities**: Similar to the order of magnitude discussion above, considerations of emergent capabilities and the potential of models or algorithms to invoke knowledge and understanding well beyond the inputs must be a part of ongoing awareness efforts within the DOE research community. Attention, on a case-by-case basis, must be paid to models and/or agents that could significantly elevate the data they started with or were trained on and hence need a change in appropriate controls. For example, any surrogate model will provide the ability for some interpolation and extrapolation. Depending upon the quality of the results, this may represent a useful and expected improvement in the given application area or, possibly, a ground-breaking advancement. Two examples on the opposite end of the spectrum might be a materials model developed with significant input from subject matter experts, experimental campaigns, and years of work that results in a handful of viable candidate materials versus a system that instantaneously produces equation of state data for previously unknown or highly valuable materials. The former should be considered GREEN (unless the materials themselves are considered YELLOW/RED), while demonstration of the latter outcome may warrant a YELLOW designation, depending upon the application and market.
- **Full stack development**: Especially in security sensitive domains there exists a distinction between knowing or advancing individual aspects of a technology and comprehensive knowledge of a deployed system. In airport security, for example, research on X-ray computed tomography hardware or segmentation algorithms is generally open knowledge. However, the exact specifications and thresholds used at checkpoints are closely guarded Intellectual Property and security secrets. Reviewers are asked to make similar distinctions for AI/ML systems especially in the areas of adversarial attack/defense, cybersecurity, or data poisoning. Research on specific algorithms or approaches is most likely to be considered GREEN as it has limited immediate impact. On the contrary, the combination of even existing algorithms into a security framework potentially being deployed or sold commercially should be considered YELLOW or RED, to the extent it reveals capabilities of such systems.

Recommended actions regarding applications of these AI tools to topics in the other five sections of the Risk Matrix are also described in this portion of the Matrix. Both the information in this section and in the application-specific section must be consulted to inform judgements on how research information should be treated. Additional considerations on economic value must consider the importance of tacit knowledge gained through this research in areas such as workflow construction and veracity of collected data. Also related to data, it is important to note that if the underlying data are restricted, then ML/AI research that invokes these data as training sets is also restricted.

For all topics described in the Table below, some related information that is already protected by either classification or export control (or, in some cases, by industry vendor-specific non-disclosure agreements) is not reflected here, but their scope should be evaluated relative to these controls on a case-by-case basis prior to release or dissemination.

Further, these controls as outlined in this document may be affected by significant advances in algorithms and models. Some have speculated that advanced models may enable deeper understanding and/or enhanced capabilities via emergent phenomena. In addition, techniques to curate, combine, or augment data may be created that can enhance information content or reveal new relationships. Continued awareness, consultation, and coordination among DOE Laboratory research performers are necessary to monitor the onset of potential disruptive advances with economic or technological leadership value. Users of this risk matrix must monitor the ongoing progress of such work for these postulated outcomes.

Battery Science & Technology

State of the Art (SOTA) Metrics:

1. **High energy transportation/military/space**: Gr|liquid| NMC955: 320 Wh/kg_{cell}, 720 Wh/l_{cell} measured at C/3 rate, 1000 cycles to <20 percent capacity fade, 15 year calendar life, \$86/$\text{kWh}_{\text{cell,rated}}$, \$103/$\text{kWh}_{\text{pack, rated}}$, EUCAR safety
2. **Low cost**: Gr|liquid|LFP: 232 Wh/kg_{cell}, 477 Wh/l_{cell} measured at C/3 rate, 1000 cycles to <20 percent capacity fade, 15 year calendar life, \$82/$\text{kWh}_{\text{cell,rated}}$, \$101.7/$\text{kWh}_{\text{pack, rated}}$, \$300/$\text{kWh}_{\text{system}}$ (grid), EUCAR safety
3. **Grid storage**: 150-250 Wh/kg_{cell}, 3500 cycles to <20 percent capacity fade, 365 cycles per year, 1C max rate, \$130/$\text{kWh}_{\text{pack/rack}}$, \$75/$\text{kW}_{\text{power electronics}}$, \$300/$\text{kWh}_{\text{system}}$
4. **Cost of materials**: NMC955: \$22.5/kg; LFP: \$8.5/kg; Graphite: \$9/kg; Electrolyte: \$14/l; Separator: \$0.9/m^2, Copper: \$1.2/m^2, Aluminum: \$0.4/m^2

Bioscience & Biotechnology

Synthetic Biology

Synthetic biology is the application of genetic engineering tools to develop new biological parts, processes, devices, systems or organisms for a useful purpose. Genetic engineering has been used since the 1970s as a tool to investigate biological systems, but its utility has been dramatically expanded due to advances in DNA sequencing, omics, DNA synthesis, genome editing tools and bioinformatics. As a result, metabolic engineering of microbes has become commonplace and is used to conduct fundamental research as well as for biomanufacturing of fuels and value-added chemicals. Synthetic biology is also being applied to address human health challenges.

Key drivers for the advancement of synthetic biology tools and methods with the potential for significant economic and international competitiveness impact are:

- Significant advances in genome editing methods, tools, technologies that improve their utility, accuracy, cost, etc.
- Advances in computational and bioinformatic methods for biological design, data analysis, management and biological systems modeling. Publicly accessible data repositories, such as those hosted by JGI, EMSL, the Broad Institute and iGEM, are primary mechanisms through which scientists and facilities share information relating to the synthetic biology constructs and strains that have been generated in the past for a user community or published in the scientific literature. Such public databases are usually less sensitive than private repositories that collect information for commercial development.
- Facile synthesis and assembly of long sequences of nucleic acids encoding genes, gene clusters, or entire genomes. Booting, the steps required to go from DNA to make a new, viable organism, is a critical bottleneck in synthesizing organisms. Viral rescue from transfected cells can produce active virus from engineered nucleic acid.
- Adaptation of gene editing tools developed in model organisms (e.g., E. coli K-12, Arabidopsis thaliana, or Saccharomyces cerevisiae) to engineer hosts organisms or chassis that are well-suited for biomanufacturing.
- Utilization of advanced synthetic biology methods or automated interpretation of analytical results for chassis organism engineering. This could enable both dual use gain of function in potentially harmful organisms as well as substantial economic benefit for biomanufacturing.

Some strains and genetic elements may be subject to commerce restrictions (e.g., export control regulations and the Commerce Control List, Category 1 – Special materials and related equipment, chemicals, "microorganisms" and toxins). Information about biological threats that could impact humans, plants, animals and the environment could be controlled by national security restrictions. Other biological materials may be protected by non-disclosure agreements or licenses.

Omics and Automation Technologies

"Omics" technologies are a set of techniques used to measure the presence and, in some cases, the abundance and distribution of, specific molecular species or interactions within a biological organism or system. They include genomics (DNA), transcriptomics (RNA), proteomics (protein measurements), lipidomics (lipids), glycomics (oligosaccharides) and metabolomics (metabolites). Advances in omics technologies have been rapid and are contributing to dramatic increases in our understanding of how biological systems interact and function. Basic research funded by DOE plays a key role in the development and utilization of omics and automation technologies, and performers include national laboratories and academia. Research that is low on the technical readiness level (TRL) scale, including proof of concept of new omics and automation technologies, should remain open and accessible to the scientific community. However, new applications and portability also underscore the importance of protecting potentially high-value intellectual property. Higher TRL efforts (e.g., prototype systems in relevant environments), as described in the U.S. DOE Technology Readiness Assessment Guide and the NNSA Technology Readiness and Maturation Guide, may be protected if the promise of economic and/or disruptive technology application is clear, or if technologies are intended for battlefield, clandestine, or defense applications.

Data & Advanced Computational Biology

There is growing recognition that the nation needs to better understand how to protect biological data. Loosely defined, data security can be seen as:

- Protecting biological data, such as genomics, multi-omics, microbiome, clinical trial data, health trends and outcomes, process data, and product safety results.
- Enhancing bioinformatics analysis integrity.
- Enhancing process security and assurance.
- Building trust and validation into processes and data.
- Protecting against cyber and cyber-physical intrusions, manipulations, and theft of information technology infrastructure, control systems, biological data, equipment, systems, and medical devices.

Biomanufacturing and Biomaterials

Biomanufacturing uses biological organisms or their enzymes to transform heterogeneous materials, such as lignocellulosic biomass, algae, gaseous feedstocks or plastics, into value-added products such as biofuels, biomaterials, and bioplastics. Total direct biomanufacturing, including biomass production, contributes $439 billion annually to the U.S. Gross Domestic Product. In this context, biomaterials are naturally produced substances that can be used to create bio-derived materials.

Research, development, scaleup, and commercialization activities can be categorized in stages of Biomanufacturing Readiness Levels (BioMRLs) analogous to Technology Readiness Levels. The Risk Matrix for this subsection is structured with foundational research at low BioMRL stages (typically Green). Pre-competitive data are often published and publicly accessible; the Matrix usually assigns this information to the Green category. This includes publicly available data, genetic sequences, protein sequences, and related bioinformatic data. Development, scale-up, and deployment activities that add significant economic value and de-risk commercialization of bioproducts or biomaterials may be assigned as Yellow or Red in the Matrix to protect proprietary data or national security. This data might also be protected by classification, export control, licensing or nondisclosure agreements, particularly in collaborations with commercial entities. The Risk Matrix provides graded examples of data sensitivity for biomanufacturing and biomaterials production to assist Risk Matrix review determinations.

Biomedical Research & Technologies

Biomedical research aims to understand the biological processes that govern health and disease states, to measure markers of these states, or to alter these states in a beneficial way. Basic biomedical research is typically unclassified and publishable, with broad potential benefits to humanity. Applied biomedical research may have significant implications for economic competitiveness and national security; because of this, higher TRL efforts (e.g., prototype systems in relevant environments), as described in the U.S. Department of Energy Technology Readiness Assessment Guide and the NNSA Technology Readiness and Maturation Guide, may be protected if the promise of economic and/or disruptive technology application is clear, or if technologies are intended for battlefield, clandestine, or defense applications.

Agricultural & Environmental Technologies

Agriculture is a multi-billion-dollar enterprise, the tangible output of which is embodied in elite germplasm for crop species as seeds, rootstocks, rhizomes and other propagules. These propagules represent the culmination of years, or even decades, of intense study and breeding by agricultural companies. Crop Research Technology Levels (analogous to TRLs) have been proposed to characterize stages of agricultural technology development. Preliminary technology solutions, proof of concept experiments, and some field trials at low TRL usually describe fundamental research in the Green category of this Risk Matrix. By their nature, pre-release elite commercial lines are self-contained carriers of recombinant technology that are ready to deploy at scale. Theft of propagules, or their digital genomes, by a foreign entity for crop production allows foreign players to capitalize on critical U.S. technology development. Work at National Laboratories based on these pre-release elite commercial lines for pre-commercial assessment should therefore be carefully controlled and based on their potential for immediate deployment at the highest TRL levels, should fall into the Red category.

Pre-commercial germplasm testing in field trials, and the scale-up of processing technologies for harvested lignocellulosic biomass are categorized as Yellow, to protect commercial potential. In addition, harnessing the potential of plant crop species is currently bottlenecked by limited innovation in introducing and propagating recombinant DNA in plants. Technological breakthroughs in this area will substantially accelerate the field of plant synthetic biology and the derived economic benefits when implemented at pre-commercial scale. By contrast, stock centers and other resources of publicly available germplasm enable techniques for Genome-Wide Association Studies to identify phenotypes that contribute to specific traits. These populations and experiments use well established and publicly documented lines for fundamental discovery science. The identification of genotype-phenotype relationships underlying specific traits is fundamental science should therefore be classified as Green. Positive association of a gene with a phenotype is the start of a multi-dimensional process in which the gene must be introgressed with elite germplasm, often along with other specific genes, under the control of a number of promoters, promising lines must be tested at laboratory scale for efficacy before scale-up to field testing at many geographical locations, and under a diversity of environmental conditions to determine consistency of the desired phenotype and identify any changes in overall yield before commercialization can occur.

This process can be likened to drug development in which the linkage of gene to phenotype is analogous to the identification of novel drugs or therapeutics, that must pass screens for effective concentration, distribution within the body, toxicity and phase I, II and III testing before widespread deployment.

Bioscience & Biotechnology Applications of Artificial Intelligence

Artificial intelligence (AI) is increasingly an integral part of modern bioscience and biotechnology. Deep-learning models such as AlphaFold and a growing ecosystem of generative design tools cut the time and cost of discovering enzymes, vaccines, small-molecule leads, biomaterials, and climate-resilient crop traits; when paired with high-throughput robotics, they could enable “self-driving” laboratories that iterate thousands of biological designs per week and feed new data back into ever-improving models. These capabilities promise major gains in health, energy, agriculture, and environmental remediation, and are becoming core to U.S.

economic competitiveness in synthetic biology, pharmaceuticals, and biomanufacturing. These same AI technologies, however, can also reduce the skill, time, and tacit knowledge needed to create or optimize novel pathogens, toxins, or delivery systems. Systematic research into autonomy risks and AI safeguards is essential to anticipating threats, shaping effective policies, and ensuring safe and beneficial integration of AI in bioscience and biotechnology.

Additional guidance is provided for some, but not all, sections of the S&T risk matrix in the narrative.

Accelerator Science and Technology

In the national security arena, accelerator technology is evaluated against sponsor-based criteria using lab-specific policy guidelines. Sponsor criteria provide for coherence across the DOE Lab complex, and lab-centric policy allows for tailored adaptation to specific capability areas that may be unique. A similar approach is proposed to establish guidelines to inform and guide researchers in the field regarding communicating their work. The approach is based on establishing a position that general research in advancing the science and technology of accelerators be designated open communication areas, then identifying specific exclusions to the general position that are so designated to:

- Preserve a competitive market or intellectual property advantage
- Limit the advancement of an application or technology area by a foreign country leveraging the U.S. investment in restricted areas
- Maintain U.S. technical leadership and prominence in the area
- Allow time for establishing a market base to develop

In the following sections, the status of the following areas of active accelerator R&D in the U.S. is discussed:

- Superconducting Radio Frequency Technology
- Laser and Plasma Wakefield Acceleration
- Superconducting Magnets
- Cryogenic Plant Design & Operation
- Advanced Light Source Technologies
- Very High Average Current Beam Technologies: Electrons
- Very High Average Current Beam Technologies: Hadrons
- Accelerator-based Nuclear Systems
- Accelerator-based Isotope Production

Superconducting Radio Frequency Technology

Superconducting radio frequency (SRF) is an important and rapidly evolving technology for frontier science with tremendous, untapped potential for spin-off applications in industry, medicine, security and environmental applications. Basic R&D and development have been open and global through collaborations such as the TESLA Technology Collaboration, the U.S.-Japan Science and Technology Cooperation Program in High Energy Physics, the International Linear Collider (ILC) Global Design Effort process, collaboration with CERN and at labs and universities worldwide. Many current and future accelerator projects that utilize the SRF technology are of international character and require sharing relevant technologies with partner institutions and industry, e.g., PIP-II, Linac Coherent Light Source (LCLS)-II/LCLS-II-HE, and Electron-Ion Collider in the US, HL-LHC at CERN, European x-ray free electron laser (XFEL) in Germany, SHINE in China, future ILC in Japan.

The state of the art of bulk niobium technology is well established in all regions of the world at >30 MV/m (up to ~50 MV/m in some cases) and $Q > 10^{10}$ using 2 K liquid helium refrigeration. SRF cavities and cryomodule components are supplied by vendors around the globe. Recent large-scale projects such as European XFEL and LCLS-II have enhanced supply chain capabilities by pushing processes upstream to vendors so that they can supply fully test-ready assemblies.

Worldwide R&D

There are active R&D efforts worldwide to move beyond bulk niobium by exploring surface modifications (e.g., "doping"), new materials (e.g., Nb3Sn), thin film and multi-layer technologies, and innovative manufacturing methods. These are presently early stage and open but may prove transformational in terms of achievable gradient (~100 MV/m), operating temperature (4 K or higher), or cost with dramatic implications for both scientific and commercial applications. Staying engaged in open collaborations is very important on early-stage basic research as the U.S. is both a contributor and beneficiary of shared knowledge. For example, the U.S. is leading the field with the vapor diffusion process on depositing Nb3Sn on niobium. However, other countries are investigating different processes, which are not being studied in the U.S. currently. If they are successful, maintaining open collaborations would allow U.S. researchers to have access to those processes.

While research results are openly shared and published, the labs are keenly aware that details of processes developed because of years of effort and investment may have significant commercial value. This IP is protected via invention disclosures and patents. Joint developments with industry and other parties with potential commercial value are regulated by CRADA's in which the IP terms are defined.

SRF-based integrated systems

While the building blocks of SRF-based accelerators are well known globally, the integration of these elements into highly capable systems for specific purposes is an important capability. Such integrated designs can be transformative in areas outside of science such as industry, medicine, security and environment. Examples include compact light sources for lithography, e- beam sources for industrial or environmental applications, accelerators for isotope production, among others. Such developments are typically tightly controlled projects, and the resulting designs are protected by patents, copyrights and licensing agreements. Drawings, designs and process details are at a minimum "business sensitive" and depending on capabilities may be export controlled or even International Traffic in Arms Regulations.

Workforce considerations

The global SRF community is small, maybe ~1,000 specialists worldwide. It is imperative that the U.S. remain engaged in open collaborations on early-stage basic research as both a contributor and beneficiary of shared knowledge. It is also vital to retain access to the best and brightest students and researchers globally to maintain the pace of innovation required by the DOE missions. There are currently not enough SRF experts educated and trained in the U.S. A significant fraction of the workforce is recruited from abroad: non-U.S. citizens are vital to support SRF research. Training future generations of scientists and engineers should be a high priority.

Recommendation: The US should maintain its world leadership in advanced SRF technology utilizing advanced materials, coatings, and/or treatments beyond bulk niobium. International collaborations are beneficial at early R&D stages. However, IP protection via invention disclosures and patents is important at this stage. As technology areas mature, careful considerations should be given to identifying specific advanced techniques as sensitive technologies to be controlled. There is a very delicate balance between protecting such technologies and hampering further progress.

One area where SRF based technology can be transformative for security and defense applications is in the production of very high-current or high-beam-power systems (also covered below in the section on high current beams). While the building blocks of high current accelerators have been well established worldwide over many decades, the integration of these technologies into very high-power systems can deliver advanced capabilities.

Laser and Plasma Wakefield Acceleration

Basic research on the techniques of laser- and beam-driven plasma wakefield acceleration have been conducted since the 1990's. While extremely high accelerator gradients have been demonstrated, the realization of a practical accelerator has not been accomplished to date. The successful development of this technology for practical applications would open many new opportunities by providing much smaller accelerators and perhaps provide a method to achieve higher energy beams than conventional accelerators and at lower cost. The U.S. program has been world-leading but recently is being strongly outspent by Asian and European efforts. While in the near term the U.S. remains in position for breakthrough advances, long term progress will necessitate either increased work with the larger international programs or, preferably, a substantially increased domestic program. Restriction without investment would cause the U.S. to be left behind.

Evaluation of the underlying laser technology that enables laser-driven acceleration of particles is outside the scope of this section.

***Recommendation**:* The US should continue to develop and maintain a forefront R&D program in wakefield acceleration. While most of the research is basic physics and appropriate to keep open, specific topics that are close to application should be treated as sensitive technology to maximize our competitive advantage. Proactive use of the patent system to restrict access to substantial new developments will be important in this area.

7 Reference: Advanced Accelerator Development Strategy Report, USDOE Office of Science, 2016. DOI: 10.2172/1358081

Superconducting Magnets

Superconducting (SC) magnets were developed in the 1970's to enable higher magnetic fields than achievable with conventional copper conductors. Superconducting magnets using the now- standard Niobium-Titanium conductor are nowadays commercially built in large numbers for MRI machines and other applications. Industrial-scale production of even higher field magnets requires new technology currently under development. There is a worldwide effort to develop higher field accelerator magnets based on novel materials with higher transition temperatures and higher critical field, such as Nb3Sn and high-temperature superconductors (HTS), including a very active program in the U.S.

Worldwide R&D

Several university and laboratory groups around the world are engaged in this research. For now, the primary potential application of the Nb3Sn accelerator magnet technology is within the international HEP community. High temperature superconductors and the associated magnet technology is an important nascent research area that is developing rapidly in large part due to a vibrant international research community. Industrial partners are typically involved in low-temperature superconductor and HTS conductor R&D to fabricate both short samples and production runs of conductors for test magnets as well as to develop industrial applications, particularly in the energy sector (e.g., energy storage, wind generator applications, and fusion confinement). While the U.S. has a slight advantage in high field conductor and magnet development now, European institutions are catching up rapidly. Any restriction of information exchange would likely delay the progress to upgrade or possibly build new particle accelerators and would undermine U.S. leadership in the SC accelerator magnet field; therefore, U.S. roles in HEP experiments abroad would be negatively impacted.

Workforce considerations

The global SC magnet community, including industry, is relatively large, but only a small portion of it is involved in advanced R&D. It is imperative that the U.S. remain engaged in open collaborations on early-stage basic research as both a contributor and beneficiary of shared knowledge. It is also vital to retain access to the best and brightest students and researchers globally to maintain the pace of innovation required by the DOE missions. Non-U.S. citizens represent a relatively large and active portion of the workforce. They have contributed and are contributing critically to U.S. leadership in the field. Continuing to train future generations of scientists and engineers in the U.S. should be a high priority.

Recommendation: The U.S. should maintain a competitive position in advanced superconducting magnet technology using novel materials to achieve high magnetic fields. Existing international collaborations (e.g., with CERN) are vital for further progress.

Cryogenic Plant Design & Operation

The development and deployment of large-scale refrigerators, greater than 4kW @ 2K, for cooling superconducting RF accelerators first started in the late 1980's in the U.S. Since that time the technology has been utilized for accelerators in Europe and the U.S. with extensive publication in professional literature of hardware designs and operating parameters for the refrigerators, compressor systems, cold boxes, and distribution systems. The refrigerator systems are substantially built from components (turbo machinery, warm helium compressors, transfer lines and 4K refrigeration cycles) provided by industry and produced in the U.S., Europe and Asia. The key to successful stable and reliable operation of superconducting cavities is the precise control of the sub-atmospheric helium bath pressure/temperature in a high dynamic heat load environment. In the U.S. this is generally done with a completely cold vacuum pumping system coupled with dynamic heat load management programs to rapidly stabilize heat input and resulting helium bath pressure fluctuations. These stabilization algorithms and hardware must evolve with the improvements in superconducting radio frequency technology. Recent projects in the U.S. continue to push development of these control algorithms to support more extreme heat management challenges in particle accelerators.

Advanced light source technologies

Accelerator research is the cornerstone for the development of new technologies that facilitate the construction of next-generation x-ray sources. Accelerator research discussed in this section focuses on creating, manipulating, transporting, and diagnosing ultra-low emittance electron beams relevant to high-brightness storage ring and FEL user facilities. The brightness of storage ring-based synchrotron radiation sources has increased dramatically over the last several decades and will continue to increase into the future based on advanced accelerator lattice designs, notably multibend achromat lattices paired with progress on undulator technologies. These continue to support new and improved techniques to probe material, chemical, biological, and environmental systems in ever-increasing detail. Similarly, the development of FELs over the last decade that produce intense and highly coherent ultrafast x-ray pulses enables many new imaging and scattering modalities to probe nanoscale dynamics at the attosecond timescale. Advanced FEL seeding techniques, undulators, and superconducting RF technologies continue to advance the capabilities of FELs at a rapid pace.

***Recommendation**:* The U.S. should maintain a forefront R&D program in storage rings and FELs for future light sources. While much of the research is basic physics and appropriate to keep open, specific topics that are close to application should be treated as restricted to maximize our competitive advantage.

Very High Average Current Beam Technologies: Electrons and Hadrons

Electron beams – Very high current electron beams, especially with high beam quality, remain an area of expertise for the U.S., though strong efforts exist outside this country. The U.S. has produced the highest current from a photocathode source and the highest beam power in an ERL. The technologies involved in circulating high currents will not only be useful in the future electron-ion collider designs but may also be useful for producing new high-power light sources, either for industrial defense/security applications.

Research into energy recovery of low current beams (<250 mA) and of high brightness beams for short wavelength light sources should remain completely open. High brightness sources of polarized or magnetized beams should also, in general, be unrestricted.

Hadron beams – High power hadron beams is an area of expertise in the U.S. The highest power hadron linac presently operating in the world is the ORNL Spallation Neutron Source: the 1.9 MW, CW proton beam is directed on a mercury target to generate neutrons for a wide range of scientific research and applications. High-power hadron facilities are operating in

Europe and Japan, new high-power hadron facilities are being commissioned in China and South Korea and the European Spallation neutron Source, a facility that will be ultimately capable of operating at five MW, is being built in Sweden as a European collaboration. It is imperative for the U.S. to remain competitive in high-power hadron beams and to push the envelope into the 10-20 MW regime.

A series of strategic applications can be enabled by such high-power hadrons beams: generation of neutrons via spallation, nuclear technology applications, and accelerator driven systems for nuclear energy production, irradiation facilities, isotope production, high-intensity muon sources. Accelerator-based nuclear systems have applications in tritium production, production of fissile materials, processing nuclear waste, and advanced nuclear reactor designs. The systems often involve high power (>2MW) proton beams, which are discussed in the previous section on very high current beam technologies. Additional aspects are associated with coupling of accelerators to reactors or target systems. As these technologies are relevant to issues of national security, many restrictions and export controls are already in place.

Accelerator-based Isotope Production and other Industrial and Medical Accelerators

Isotope production via reactors and accelerators has applications for medicine, scientific research, national security, and industry. Accelerator-based production methods using electron, proton, and ion beams, typically with energies in the 10s to 100s of MeV, provide customized production paths for a wide range of isotopes. Secondary particle production, i.e., spallation neutrons and photons, produced from the interaction of these beams with a target, further expands the production paths available. Conventional RF LINACs and cyclotrons are the most utilized accelerators for isotope production.

Based on recommendations from the 2015 Long Range Plan for the DOE-NP Isotope Program8, one focus of current isotope program research is the "development of advanced, cost-effective and efficient technologies for the production, processing, recycling and distribution of isotopes in short supply9." Active basic R&D efforts include research into more energy efficient and higher current proton/ion acceleration techniques suitable for effective isotope production, simulation and design of targets to provide optimized isotope production, exploration of new isotope production pathways, and improved processing and purification of the resulting isotopes. It should be noted that the currently available production facilities also provide unique opportunities for irradiation studies of materials of relevance to high energy physics as well as fission and fusion reactor applications. While basic accelerator R&D that can synergistically benefit the isotope program should remain open, isotope production technology is generally considered business sensitive and controls are in place to limit access to and distribution of information.

Recommendation: The U.S. should maintain its leading role in developing efficient and cost- effective routes to fulfill society's needs for a reliable supply of isotopes. Secondary utilization of national laboratory production facilities for development and testing of materials relevant to other SC thrusts should also be preserved. Additionally, the U.S. should continue to maintain a forefront R&D program in compact accelerators for industrial and medical applications, and the DOE Laboratories need to continue to be aware of when technical advances lead to realized commercial advantages or potential dual-use national security applications.

8 Meeting Isotope Needs and Capturing Opportunities for the Future: The 2015 Long Range Plan for the DOE-NP Isotope Program, pp. 79-89. 9 FY 2019 DOE Office of Nuclear Physics Funding Opportunity Announcement DE-FOA-0001968.