Federal Register Notice: 89 FR 51554, Federal Register :: Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development, June 18, 2024.

Request for Information on the National Digital Twins R&D Strategic Plan

General Atomics

National Strategy for Scientific Innovation and Leadership Through Digital Twins and the Fusion Energy Exemplar

R. Nazikian, D.P. Schissel, *General Atomics*
T. Gibbs, J. Wells, *Nvidia*
B. Spears, T. Bremer, *Lawrence Livermore National Laboratory*
F. Wuerthwein, A. Majumdar, *San Diego Supercomputer Center, University of California, San Diego*

Key topics: Artificial Intelligence, Business Case, Ecosystem, International, Long Term, VVUQ, Workforce

1. EXECUTIVE SUMMARY

"Intelligent" digital twins (iDTs) [1], integrating artificial intelligence (AI) and machine learning (ML) into traditional digital twin frameworks offer unprecedented potential to accelerate scientific discovery and technological innovation. By creating high-fidelity virtual replicas of complex systems, iDTs can enable rapid iteration, optimization, and testing of ideas in a low-risk virtual environment, allowing researchers to explore vast multi-dimensional parameter spaces, identify complex interdependencies, and make data-driven decisions at a pace and scale not possible with traditional methods alone. As the United States seeks to maintain its leadership in science and technology, the development of a national strategy for AI-enhanced digital twins is crucial.

To capitalize on the rapid pace of AI innovation and evolving iDT capabilities, we recommend the immediate initiation of domain-specific iDT pilot programs with clear objective outcomes. These pilots should adopt a "Fail Fast" approach, prioritizing swift deployment, iterative refinement, and continuous learning. Critically, the funding and management framework for the pilots must seamlessly integrate the three core pillars of advanced scientific computing—models, data, and compute—with the domain-specific expertise found at, for example, Department of Energy (DOE) user facilities. This integration is essential given the intrinsically interconnected nature of models, data, and compute in iDTs, where the computing spans from edge systems at the user facility to leadership-class computing centers nationwide, and where latency requirements are crucial for linking virtual and physical systems. This aligns with the recent DOE Advanced Scientific Computing Research (ASCR) Advisory Committee recommendation [2] advocating for innovative governance and funding models to prepare for a future where DOE user facilities are intimately linked with national computing and data facilities, enabling real-time decision making in control rooms and accelerating scientific breakthroughs.

Recent advancements in AI for both magnetic and inertial fusion energy (MFE/IFE), coupled with successful proof of concept (POC) demonstrations, underscore the immense potential of iDTs in fusion research. Given fusion energy's recent technical breakthroughs, its substantial economic and societal promise, the success of early POCs, and the White House Bold Decadal Vision [3] for fusion energy development, fusion stands out as an ideal candidate for a national iDT pilot program. By enabling rapid, cost-effective iteration, risk mitigation, and AI-driven discovery, iDTs could significantly accelerate timelines, reduce costs, and increase the probability of successfully achieving commercial fusion energy. The proposed pilot would aim to develop a high-

performing iDT, tightly integrated with facility operations, combine high-fidelity AI models with live data streams, enabling bi-directional, low-latency interactions between the digital and physical realms. The iDT would continuously update its models based on real-time inputs, allowing for timely analysis, prediction, and optimization of fusion facility operations. To realize the transformative potential of iDTs in science and technology, we recommend the following actions:

- **Implement a domain-specific iDT pilot program** prioritizing rapid prototyping, iterative development, and continuous learning. Pilots must adopt a "Fail Fast" approach, have clearly defined objectives and integrated management for models, compute, and data.
- **Launch a national flagship initiative to develop a comprehensive fusion iDT,** building upon recent POCs in MFE and IFE, and integrated with experimental facilities to support the White House's Bold Decadal Vision for fusion energy.
- **Fund multi-disciplinary teams comprising domain science, computational, and data science experts.** These integrated teams will ensure the agile development of iDTs.

This request for information (RFI) response describes the rationale for a national strategy that focuses on domain-specific multi-disciplinary iDT pilots for the realization of fusion energy.

2. THE FUSION ENERGY CHALLENGE AND DIGITAL TWIN OPPORTUNITY

Fusion energy stands at the forefront of humanity's quest for sustainable, clean power, offering a potential solution to the pressing challenges of climate change and global energy demand. Recent breakthrough achievements in both inertial and magnetic confinement fusion have reignited enthusiasm for this transformative technology. The National Ignition Facility's (NIF’s) demonstration of fusion ignition [4], and the sustained high-performance conditions in magnetic confinement systems [5], mark significant milestones in the field. These advancements have catalyzed unprecedented growth in the private fusion industry [6], with $6B invested in numerous startups now racing to commercialize fusion energy. Recognizing the momentum, the United States government has put forth a bold decadal vision for fusion energy development, aiming to demonstrate commercial viability within the next decade. This ambitious goal underscores the critical role of the government in fostering innovation, supporting basic and translational research [7], and creating a regulatory framework for this emerging industry. As we stand on the cusp of a new era in energy production, the development of comprehensive iDTs for fusion systems emerges as a key enabler, promising to accelerate progress in this complex, multidisciplinary endeavor.

The National Academies of Science, Engineering, and Medicine (National Academies) study on burning plasma science [8] highlights several critical challenges in achieving commercially viable fusion energy. These include maintaining plasma confinement and stability, managing heat and particle exhaust, developing fusion-compatible materials, and integrating complex systems for continuous operation. Importantly, the report emphasizes the need for accelerated learning cycles and improved predictive capabilities. iDTs integrated with a performant data platform (see high-performance data facility progress [9]) and advanced computing capability (see integrated research infrastructure development [10]) offer a powerful approach to address these challenges and accelerate progress towards fusion energy.

The iDTs have the potential to dramatically accelerate the development of commercial fusion energy by revolutionizing the design, testing, and optimization processes. They enable rapid, cost-effective iteration of reactor designs in a low-risk environment, while integrating operations, maintenance, and control into the design process, potentially reducing cost and risk in developing and operating physical systems. These virtual environments allow for the exploration of vast parameter spaces and parallel testing of multiple design variants, potentially accelerating development cycles and increasing the number of design iterations. By enabling early identification of design and potential operational issues, iDTs mitigate risks and reduce costs associated with developing and operating physical systems. On existing facilities, they open up possibilities for AI-driven discoveries of novel operating conditions and fusion configurations, which can translate to improved reactor concepts. Collectively, these advantages can compress decades of traditional fusion development into years, significantly reducing overall costs and increasing the probability of success. This acceleration can be highly enabling for the commercial development of fusion energy and the realization of the ambitious timeline set forth in the United States government's decadal vision.

3. FUSION PILOT PROGRAM

The recent convergence of multiple technological advancements presents a unique opportunity for rapidly advancing fusion iDT pilots. These advancements include the successful integration of AI into real-time control systems, the seamless incorporation of leadership-class computing capabilities into facility operations, and the development of community-driven data platforms that facilitate data curation and efficient AI and ML model creation. The synergy of these cutting-edge technologies provides the essential foundation for the rapid development of transformative fusion iDTs.

MFE: MFE is fundamentally a continuous process of energy production, in contrast to IFE, which is fundamentally a discrete process of implosions occurring multiple times per second. To fully grasp the role and impact of MFE iDTs, it is essential to consider the key stages of the MFE lifecycle and the corresponding latency requirements for information flow between the virtual and physical assets.

The MFE iDT must operate across multiple timescales, each with distinct latency requirements (Fig. 1). For real-time control applications, the fastest timescale, latencies must be extremely low, ranging from milliseconds to seconds. This enables rapid adjustments to plasma conditions and machine parameters through edge computing, coupling sensor data with real-time inference engines to predict future states and command actuator responses. On a longer timescale of minutes to hours, the iDTs provide interpretive and predictive analysis to inform operations. This intermediate timescale allows for more complex calculations and data processing at leadership-class computing centers, guiding operational decision-making with advanced insights. For the DIII-D National User Facility [11], this step is enabled through DOE’s Integrated Research Infrastructure (IRI) operated by ASCR. Finally, for informing the design of future experiments or facility upgrades, the iDTs operate on the longest timescale, with latencies of hours to days. This extended timeframe enables comprehensive simulations and analyses to support long-term planning and design optimization. It also allows for the development of advanced inference engines, surrogate and “foundation models” that require significant time to optimize on high-performance GPUs, as well as data platforms capable of processing information at scale. For the DIII-D facility, this is enabled through the IRI and a fusion data platform (FDP) that we plan to integrate with DOE's High Performance Data Facility (HPDF). This multi-tiered approach ensures that the MFE iDT can address both immediate operational needs and long-term strategic goals in fusion energy research and development.

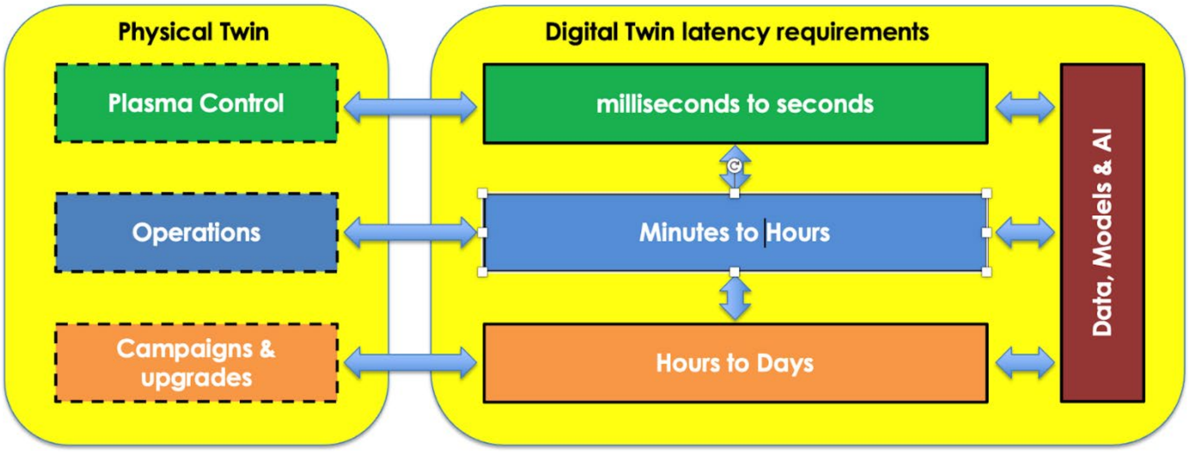

Figure 1: Latency requirements for a fusion iDT

In the following, we illustrate key advances in each of these different time scales that form the building blocks for integration into a performant MFE iDT.

Surrogate models for real-time control: A recent breakthrough in AI-driven real-time control for fusion plasmas has been demonstrated on the KSTAR tokamak [12]. Researchers from Princeton Plasma Physics Laboratory and collaborating institutions have developed an innovative approach that combines ML with adaptive control techniques to optimize the performance of tokamak fusion reactors. The study showcases a fully automated 3D-field optimization system that successfully suppresses edge magnetic instabilities while significantly enhancing fusion performance. By leveraging a ML surrogate model to rapidly calculate optimal magnetic field configurations, the system achieves real-time adaptability at millisecond timescales (Fig. 2). This approach led to remarkable improvements, including up to a 90% increase in fusion performance metrics and near-complete elimination of magnetic instabilities. The research represents a significant step towards realizing stable, high-performance fusion plasmas and demonstrates the potential of AI-driven control systems surrogate models derived from high-fidelity simulations to overcome long standing challenges in fusion energy development.

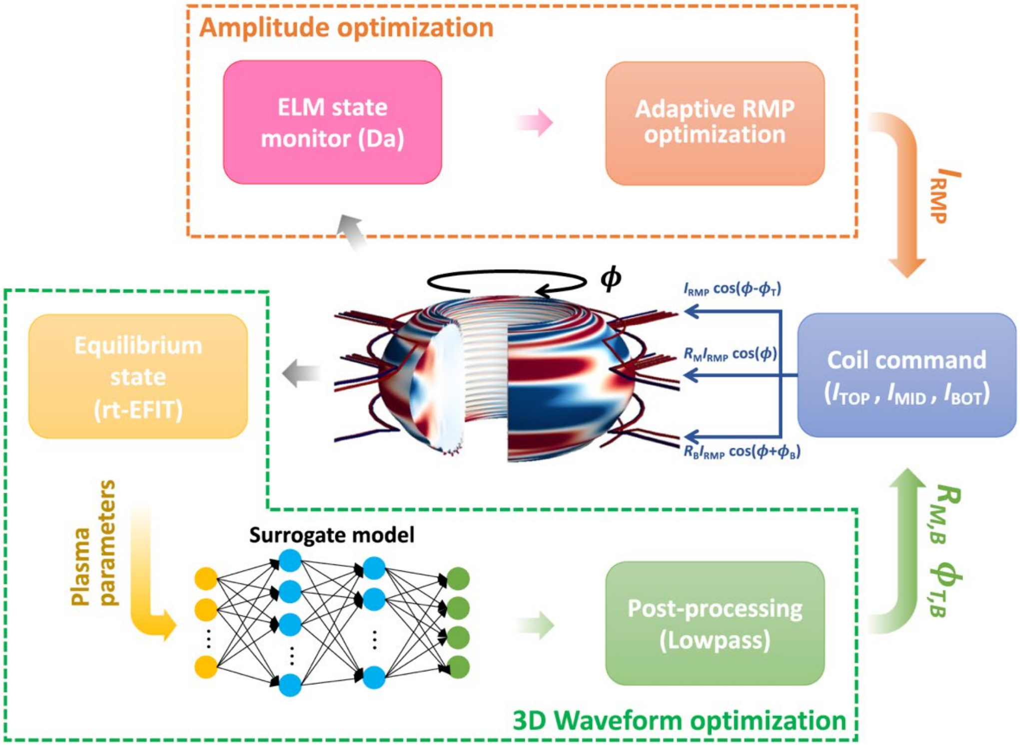

Figure 2: ML based real-time optimization using 3D magnetic field surrogate model in the KSTAR tokamak. [12]

Informing machine operations: In the minutes-to-hours timescale, the MFE iDT leverages DOE's IRI to rapidly perform complex calculations, such as neutral beam injection, ionization, and their impact on material surfaces, as shown in Fig. 3 [13]. By utilizing the IRI, DIII-D—an early IRI pathfinder—can seamlessly dispatch computations to high-performance computing centers like the National Energy Research Scientific Computing Center, the Argonne Leadership Computing Facility, and soon on the Oak Ridge Leadership Computing Facility, swiftly retrieve results to support decision making in the control room. This advanced tool can be integrated with other models, including surrogate and foundation models, to inform control room decisions on operational issues by comparing simulations with real-time machine data. Foundation models could provide a robust, pre-trained base for developing more specialized models for specific fusion-related tasks, potentially accelerating the development of accurate predictive capabilities. The combination of these capabilities with automated interpretive analysis between plasma pulses [14] enables informed decision-making on experiments and optimization of operational plans. This synergy of high-performance computing, advanced visualization, and AI-driven analysis exemplifies how the iDT can revolutionize the efficiency and effectiveness of fusion experimentation, allowing researchers to make timely data-driven decisions.

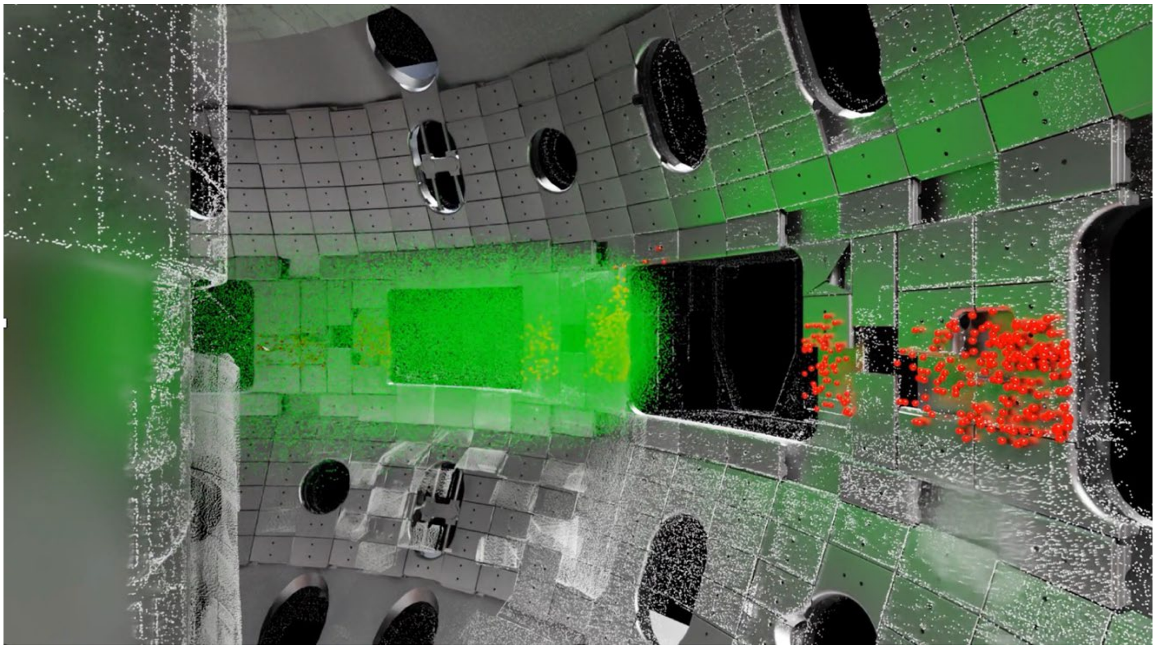

Figure 3: Calculations of neutral beam injection (right to left), beam ionization profile (green) and beam ion loss on tiles (red) inside the DIII-D vacuum vessel [13]

Long-term planning for experimental campaigns, upgrades, and new facilities: At the core of a fusion iDT's long-term planning capability is the development of a robust data platform that can analyze experimental and simulation data at an unprecedented scale. An FDP developed for the fusion research community is capable of handling several petabytes of data. This will expand significantly when integrated with the DOE HPDF and coupled to leadership class computing facilities through IRI. The platform's architecture enables the creation and refinement of AI/ML models derived from vast simulation and experimental datasets, which can then be deployed in real-time or between plasma pulses during operations to inform control room decisions. This could include the development and fine-tuning of foundation models specifically for fusion applications, leveraging the power of large-scale pre-training to enhance predictive capabilities across various aspects of fusion research and operations. Leveraging the open science data federation for efficient data distribution via caching, the platform maintains rigorous data provenance for all models, data, and artifacts through the open-source Common Metadata Framework [14]. This comprehensive approach to data management is absolutely essential for the iDT concept, as it provides a scalable solution for managing the enormous volume of models and data inherent to iDT. These various capabilities will form the backbone of a fusion iDT capacity for informed decision-making and long-term strategic planning in fusion energy development.

MFE pilot: In this RFI response, we present a few of the recent advancements in fusion research that lay a solid foundation for proceeding with confidence towards a fusion iDT pilot. Several critical components have been successfully demonstrated: AI-driven real-time control systems employing surrogate models of high-fidelity simulations used to suppress magnetic instabilities and optimize plasma performance, advanced data platforms capable of managing petabyte-scale experimental and simulation datasets, and the integration of national leadership class computing facilities for rapid between-pulse and long-term planning. These achievements, coupled with progress in simulation capabilities and visualization tools, provide a strong basis for embarking on a comprehensive iDT pilot. The fusion iDT pilot will build upon these proven capabilities, aiming to create a scalable, integrated framework that enhances our understanding, control and optimization of fusion systems from millisecond-scale control to campaign-level planning. This ambitious initiative represents a crucial next step in leveraging advanced computing and AI technologies to accelerate progress towards practical fusion energy.

IFE/IFT pilot: IFE provides a valuable complement to MFE for iDT development. IFE is fundamentally a batch-wise operation – making use of discrete implosions at multiple times per second. This contrasts with MFE, where operations are continuous. The complementary challenges introduced by this mode of operation provide a fertile ground for developing additional dimensions of iDTs and for further refining those dimensions that are shared with MFE.

IFE iDTs will make meaningful impact in three critical and interrelated stages of the IFE lifecycle. First, an iDT will provide critical information for target manufacturing operations. iDT of advanced and precision manufacturing processes will serve as a critical technology to reach production rates commensurate with the operation of an IFE reactor, or about 10 hertz in most applications. The iDT will ingest real-world metrology for targets in-situ, that is during the production process. This will produce a living in-silico copy of the target as it is being created. The iDT will present the in-silico target to a predictive model that can evaluate defects in real-time. That predictive capability can then assess the impact of observed defects while designing the mitigating “anti-defect” in real-time to be applied to the production job, such as in the next pass of a direct ink write printer. The two-way communication enabled by the iDT in this production example will enable high-throughput and high-quality operations required for future IFE power production. It will also deliver extensively metrologized targets and their iDT analogs for custom operation facilities to optimally match fusion driver execution with the particulars of the target being imploded at the current moment, thereby adding additional robustness to the operational cycle. Projects on iDT for in-situ metrology and real-time decision making are currently underway, paving the way for rapid expansion of the capability in the immediate future.

Second, the IFE iDT will require a strong and reliable predictive capability. It must be able to confidently predict what will happen to the physical twin of the system (the operating reactor and target system) based on the iDT (the simulated operations). Models have already been developed for inertial confinement fusion efforts at NIF. NIF researchers built AI-driven models of implosions using the combination of high-fidelity physics models (simulation codes) and experimental observations (actual implosion measurements). The resultant model successfully predicted ignition (the production of more fusion energy than used in the driving laser) for the first time and for the half-dozen subsequent successful events. Such a model is a critical component for ensuring two-way coupling of an iDT between the digital and physical world. It allows for real-time updating of model expectations that are both consistent with physical theory, but also accurate about observations in complex, evolving systems. This predictive capability will allow the iDT to “understand” the impacts of operational changes in the real-world system.

Finally, the ultimate goal of an iDT is to allow for optimal reasoning about a physical system based on the safer, cheaper, and faster digital model. For IFE, the iDT provides a route to full closed-loop or “self-driving” reactor operations. As presented here, the coupling of knowledge from an advanced target manufacturing system with a real-world-aware predictive model allows for optimal operation of facilities. For example, an iDT could perform rapid optimization of the fusion driver, such as a shaped laser pulse, to accommodate the particular target (and all of its individual defects and measurements) being fired in the reactor at the current moment.

4. EMERGING PUBLIC PRIVATE ECOSYSTEM FOR FUSION ENERGY

Scientific user facilities are the flagship of the DOE Office of Science ecosystem, advancing the frontiers of science, training the future high-tech workforce, and enhancing United States economic competitiveness. The rapid rise of the fusion industry, with recent investments exceeding $6 billion in MFE and IFE, and its growing engagement with national facilities like the DIII-D tokamak [11], underscores the critical role of public-private partnerships in developing a robust fusion ecosystem.

The creation of an iDT to support science and operations on DOE fusion facilities will provide a crucial link in achieving the bold decadal vision for fusion energy. Established world-leading AI companies, like Google DeepMind, are already engaged in AI research for fusion energy, along with many early-stage companies developing the underlying technologies needed for realizing commercial viability. Initiating a fusion iDT pilot program will further strengthen the synergy between national research activities and industry, both in the fusion technology and IT industries.

Recent DOE initiatives, such as the FIRE collaboratives for translational research, public-private partnerships, and the emerging public private consortia, combined with a national priority for a fusion iDT pilot and the DOE user facilities will accelerate fusion energy development, strengthen United States leadership in the field, and foster innovation and technology transfer. By bringing together industry, universities, and national laboratories, this pilot program will be instrumental in realizing the potential of commercial fusion energy.

5. DIGITAL TWIN PILOT PROGRAM ELEMENTS

To fully realize the potential of iDTs, the Office of Science and Technology Policy should consider the following key points in formulating a national policy for iDT development.

- Establish a cross-agency initiative for rapidly identifying high-impact iDT pilots, with fusion energy as a flagship project within the DOE. Such a cross-agency initiative will drive innovation across multiple fields while addressing critical national priorities, such as energy security and climate change mitigation.
- Implement an integrated funding and management model that breaks down traditional boundaries between domain sciences, computational sciences, and data sciences. This model should ensure that iDT pilots work efficiently towards clearly defined objective outcomes, avoid duplication of efforts and efficiently share progress and insights.
- Establish strategic partnerships among industry, national laboratories, and universities, on the regional and national level, integrating sector and regional strengths to accelerate iDT innovation, and enhance United States technological leadership and economic competitiveness. Within the DOE Office of Fusion Energy Science, these partnerships can

be enabled through the recently initiated FIRE collaborative program [15] for translational research and the proposed public private consortium framework [16].

- Prioritize the development of a flexible, adaptable computing and data infrastructure to support diverse iDT needs across domains. This infrastructure should be capable of supporting the development of large-scale foundation models tailored for iDT applications. DOE's IRI initiative aims to integrate edge computing at experimental facilities with leadership-class computing centers, addressing iDT latency demands. NSF's National AI Research Resource pilot uses cloud and on-premise infrastructure for large-scale AI. These programs should be informed by requirements to support the development of an intelligent twin for fusion.
- Invest in workforce development programs that cultivate the interdisciplinary skills required for effective iDT development and utilization. The traditional silos of domain science, computational science, and data science need to be bridged by an emerging workforce that comfortably spans these disciplines. University partnerships in these pilot programs can play an important role in cultivating a robust workforce pipeline.
- Establish guidelines for data sharing and interoperability standards to maximize the value of iDTs across different science domains and sectors. Within the DOE, the establishment of the HPDF will help to enable the development of such standards.

To accelerate the development of a functional iDT, the pilot program should adopt a "Fail Fast" management model. This approach emphasizes rapid prototyping of small-scale components within each functional area required for a complete DT. A crucial element of this strategy is the early deployment of a minimum viable product (MVP). By prioritizing the release of an MVP early in the development cycle, the project team can gather vital feedback from users at major scientific facilities. This real-world user experience provides invaluable insights, guiding subsequent development and enabling the team to swiftly iterate and refine the iDT capabilities. Such an approach ensures that the project's trajectory and outcomes are continuously improved from the earliest stages, maximizing the likelihood of success and accelerating the path to fully functional, high-impact iDTs.

6. CONCLUSION

The development of iDTs represents a critical opportunity for the United States to maintain its global leadership in science, technology, and innovation. As demonstrated by recent advancements in fusion energy research, iDTs have the potential to dramatically accelerate progress in complex, multidisciplinary fields of national importance.

To fully realize the potential of iDTs, we recommend the establishment of a cross-agency initiative for identifying high-impact iDT pilots, with fusion energy as a flagship project. This initiative should implement an integrated funding and management model that breaks down traditional boundaries between domain, computational, and data sciences, ensuring efficient progress towards clearly defined outcomes. The policy should invest in workforce development programs to cultivate interdisciplinary skills, foster public-private partnerships to accelerate the transition from research to commercial applications, and establish guidelines for data sharing and interoperability standards.

By implementing a coherent well-coordinated cross-agency initiative, the United States can create a robust ecosystem for iDT development that will drive scientific discovery, technological innovation, and economic growth. This forward-looking approach will position the nation at the forefront of the next wave of scientific and technological advancements, ensuring continued leadership in an increasingly competitive global digital landscape.

7. BIBLIOGRAPHY

[1] Crespi, N., Drobot, A. T., & Minerva, R. (Eds.). (2023). *The Digital Twin*. Springer International Publishing.
[2] Advanced Scientific Computing Advisory Committee, "Facilities Subcommittee Recommendations," 2024; https://science.osti.gov/-/media/ascr/ascac/pdf/reports/2024/52224Draft-ASCAC-Facilities-Report-003.pdf
[3] Developing a Bold Vision for Commercial Fusion Energy; https://www.whitehouse.gov/ostp/news-updates/2022/03/15/fact-sheet-developing-a-bold-vision-for-commercial-fusion-energy/
[4] A.B. Zylstra, et al., "Burning Plasma Achieved in Inertial Fusion," Nature **601**, 542-548 (2022); https://doi.org/10.1038/s41586-021-04281-w
[5] Sustained High Performance in MFE - what reference?
[6] United States Government Accountability Office, "Technology Assessment Fusion Energy," (2022); https://www.gao.gov/assets/gao-23-105813.pdf
[7] United States Department of Energy, "Fusion Energy Strategy 2024,"; https://www.energy.gov/sites/default/files/2024-06/fusion-energy-strategy-2024.pdf
[8] Committee on a Strategic Plan for U.S. Burning Plasma Research, "Final Report," (2019); https://nap.nationalacademies.org/download/25331
[9] United States Department of Energy, "U.S. Department of Energy Selects the High Performance Data Facility Lead," (2023); https://www.energy.gov/science/articles/us-department-energy-selects-high-performance-data-facility-lead
[10] W. Miller, et al., "Integrated Research Infrastructure Architecture Blueprint Activity Final Report," (2023); https://www.osti.gov/servlets/purl/1984466
[11] J.L. Luxon, "A Brief Introduction to DIII-D Tokamak," Fusion Sci. Technol., **48** 828 (2005)
[12] S.K. Kim, et al., "Highest Fusion Performance without Harmful Edge Energy Bursts in Tokamak," Nature Communications **15**, 3990 (2024); https://doi.org/10.1038/s41467-024-48415-w
[13] M. Kostuk, et al., "Digital Twin development for Thermal Event Detection and Avoidance in the DIII-D Tokamak," manuscript in preparation.
[14] S. Smith, et al., "Expediting Higher Fidelity Equilibrium Analysis for the DIII-D National Fusion Facility Using HPC Resources," to be submitted to the 6th Annual Workshop on Extreme-Scale Experiment-in-the-Loop Computing (XLOOP 2024).
[15] Common Metadata Framework; https://github.com/HewlettPackard/cmf
[16] DOE Announces New Decadal Fusion Energy Strategy (2024): https://www.energy.gov/articles/doe-announces-new-decadal-fusion-energy-strategy

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in the National Digital Twins R&D Strategic Plan and associated documents without attribution.