

**Fusion Simulation Program
Planning Documents**

The attached document is the summary of the Fusion Simulation Program Execution Plan that was prepared for the U.S. Department of Energy Office of Science by the following institutions under these respective contracts:

Institution	DOE Contract or Cooperative Agreement
Princeton Plasma Physics Laboratory (PPPL)	DE-AC02-09CH11466
Massachusetts Institute of Technology (MIT)	DE-SC000-1682
Tech-X, Inc.	DE-SC0001290
Los Alamos National Laboratory (LANL)	DE-AC52-06NA2-5396
Lehigh University	DE-SC0001268
General Atomics (GA)	DE-SC0001266
Oak Ridge National Laboratory (ORNL)	DE-AC05-00OR22725
Lawrence Livermore National Laboratory (LLNL)	DE-AC52-07NA27344
Argonne National Laboratory (ANL)	DE-AC02-06CH11357
Lawrence Berkeley National Laboratory (LBNL)	DE-AC02-05CH11231
New York University	DE-FG02-06ER 54845
University of California San Diego	DE-FG02-04ER 54738
Purdue University	DE-FG02-08ER54990
University of Texas	DE-FG02-04ER-54742

Work was also supported under a number of PPPL subcontracts:

Contracting Institution	PPPL Subcontract number
LLNL	ICP010306-R
ANL	ICP010379-G
ORNL	ICP010500-G
Tech-X	S010433-F
General Atomics	S010467-G
MIT	S010550-G
Lehigh University	S010741-R

Many others in the fusion energy and advanced scientific computing communities participated in the development of this plan. The core planning team is grateful for their important contributions.

This summary is meant as a quick overview the FSP's purpose and intentions. There are several additional documents referenced within this one and all are supplemental or flow down from this Program Plan. They are not included herein because of their purpose, format, or size. All documentation can be found at <http://www.pppl.gov/fsp/Overview.html> :

- The Detailed execution plan - the governing document that establishes the means to execute, monitor, and control the proposed Fusion Simulation Program (Approximately 205 pages).
- WBS Cost and Schedule Worksheet - a multipage spreadsheet containing the proposed FSP work breakdown structure (WBS) and its dictionary, work schedule, required resources, and funding timelines (MS Excel workbook)
- Detailed Science Driver Planning Reports - a report containing bottoms-up definitions of plans for addressing critical programmatic issues that require integrated simulations. It was prepared by small groups of technical experts from the fusion energy and advanced scientific computing communities at the request of the FSP planning team (MS Word, approximately 182 pages)

FSP Execution Plan Summary

Table of Contents

A. SCIENCE OPPORTUNITIES	1
B. HOW THE FSP IS STRUCTURED TO ADDRESS SCIENCE PROBLEMS.....	7
C. INTERACTIONS WITH PARTNER PROGRAMS	16
D. EXECUTION PLAN.....	19
E. SUMMARY OF RESPONSIVENESS TO RFP	33

ACRONYMS

AM/SC	Applied Math and/or Computer Science
ASC	NNSA's Advanced Simulation & Computing Program
ASCR	Advanced Scientific Computing Research
DEMO	Demonstration Magnetic Fusion Reactor
DOE-SC	Department of Energy - Science (offices)
ELM	Edge Localized Mode
ET	Enabling Technologies
FES	DOE's Office of Fusion Energy Sciences
FSP	Fusion Simulation Program
HPC	High Performance Computing (or Computers)
INCITE	ASCR's Innovative and Novel Computational Impact on Theory and Experiment Program
ISA	An Integrated Science Application Team within FSP
ITER	International Thermonuclear Experimental Reactor being built in France
LCF	Leadership Computing Facility
MFE	Magnetic Fusion Energy
MHD	Magneto-Hydro-Dynamics
NBI	Neutral Beam Injection
NNSA	DOE's National Nuclear Security Administration
PAC	Program Advisory Committee
PHYS	Advanced Physics Components development team within FSP
PMI	Plasma Material Interface or Plasma Material Interaction
PMO	Project Management Office
PROD	Production Support Team within FSP
QoI	Quantity of Interest
R&D	Research and Development
RF	Radio Frequency
SIS	Software Integration Support Team within FSP
SOL	Scrape-Off-Layer
SQ	Software Quality or the Software Quality Team within FSP
UQ	Uncertainty Quantification
V&V or VV	Verification and Validation
VVUQ	Verification, Validation and Uncertainty Quantification
WDM	Whole Device Model or Whole Device Modeling

A. SCIENCE OPPORTUNITIES

A.1 What Science Problems will the FSP Address?

Focus for the First Three Years – The Plasma Edge And Whole Device Modeling. The overall science goal of the DOE Office of Fusion Energy Sciences' (FES) Fusion Simulation Program (FSP) is to develop predictive simulation capability for magnetically confined fusion plasmas at an unprecedented level of integration and fidelity. This will directly support and enable effective U.S. participation in International Thermonuclear Experimental Reactor (ITER) research and the overall mission of delivering practical fusion energy. The FSP will address a rich set of scientific issues together with experimental programs, producing validated integrated physics results. This is very well aligned with the mission of the ITER Organization to coordinate with its members the integrated modeling and control of fusion plasmas, including benchmarking and validation activities [1]. Initial FSP research will focus on two critical Integrated Science Application (ISA) areas: ISA1, the plasma edge; and ISA2, whole device modeling (WDM) including disruption avoidance. The first of these problems involves the narrow plasma boundary layer and its complex interactions with the plasma core and the surrounding material wall. The second requires development of a computationally tractable, but comprehensive model that describes all equilibrium and dynamic processes at a sufficient level of detail to provide useful prediction of the temporal evolution of fusion plasma experiments. The initial driver for the whole device model will be prediction and avoidance of discharge-terminating disruptions, especially at high performance, which are a critical impediment to successful operation of machines like ITER. If disruptions prove unable to be avoided, their associated dynamics and effects will be addressed in the next phase of the FSP.

The Pedestal-Boundary model developed in ISA1 will include boundary magnetic topology, cross-field transport of multi-species plasmas, parallel plasma transport, neutral transport, atomic physics and interactions with the plasma wall. It will address the origins and structure of the plasma electric field, rotation, the L-H transition, and the wide variety of pedestal relaxation mechanisms. This ISA is a key partial integration project – along with future ISA's – that will ultimately be integrated into the FSP WDM that will predict the entire discharge evolution given external actuators (i.e., magnets, power supplies, heating, current drive and fueling systems) and control strategies. Based on advanced components with improved physics fidelity operating within an appropriate integration framework, ISA2 will focus on modeling the plasma equilibrium, plasma sources, profile evolution, linear stability and nonlinear evolution toward a disruption (but not the full disruption dynamics).

To guide the planning, a set of criteria was used to prioritize the research targeted. These included: (i) a clear need for multi-scale, multi-physics integration; (ii) importance and urgency; (iii) readiness and tractability; and (iv) opportunity to open up new lines of research that produce new insights/potential breakthroughs inaccessible by other means. An overarching prioritization criterion is the "buy-in" from the "customer-base" for FSP products with respect to what software capabilities are in greatest demand and urgency from the user communities. The final choices made reflect a realistic level of "market analysis" with linkages to the FES 2009 community-wide Research Needs Workshop (ReNeW) document [7], to the priorities of the Fusion Facilities Coordinating Committee, to ITER, and also to other international facilities (e.g., in Asia with experimental capabilities not available in U.S. facilities).

Topics for the Out Years. The plan assumes that, as the FSP matures and demonstrates success, the program will evolve and grow, enabling additional science problems to be addressed. The next set of integration opportunities could include the following topics:

Disruption Mitigation: If disruptions – the large-scale macroscopic events leading to rapid termination of plasma discharges, including severe impulsive heat loads damaging material components – cannot be completely avoided, mitigating the associated dynamics is critical because ITER can sustain at most a very small number of such full-current events. The associated science goal is to minimize the impact of disruptions, including dealing with transient heat and mechanical loads and generation of run-away electrons. This will involve dealing with strongly nonlinear magnetohydrodynamics (MHD) phenomena in large Lundquist number plasmas, addressing

coupling to plasma pressure and current and also to atomic physics, neutral and impurity transport, radiation transport, and relativistic electron transport. It will also require assessment of the relationship to an electromagnetic model of the fusion device, including complex wall geometry, power supplies, coils, and control systems. If successful, the expected benefits include: (i) survivability of first wall tokamak components; and (ii) viable steady-state operation of a fusion device.

Core Profiles: The science goal here is improved predictive capability for the temperature, density, current, and rotation profiles in the plasma core, including the internal transport barrier region. This task includes dealing with 3D effects, mesoscale physics, and integration with the plasma edge dynamics. It involves producing self-consistent, global solutions of micro-and macro-nonlinear dynamics on transport time scales. Since mesoscale phenomena (between gyro-radius and device size), overlap with MHD scale, there is no justifiable strong scale separation that can be invoked to simplify this challenging problem. If successful, expected benefits include a predictive capability for plasma profiles that would enable providing profile information needed to determine operational limits (e.g., sustainable plasma pressure) and plasma performance (e.g., fusion yield, bootstrap current fraction), and also to provide confidence in extrapolating core confinement predictions to future devices.

Energetic Particles/Wave Physics: These are dynamical interactions between energetic particles and electromagnetic waves in a magnetic fusion energy (MFE) plasma that impact the efficacy of auxiliary heating and the fast-particle confinement of fusion products (alpha particles at 3.5 MeV) and of supra-thermal particles from radio-frequency (RF) and energetic neutral beam ion heating (NBI). Energetic particles represent potent sources of free energy available to drive instabilities, and their thermalization without loss is critically important. The associated science goals in this area include: (i) a self-consistent description of phase space distribution on long time scales (energy confinement or slowing-down) that are orders of magnitude longer than the Alfvénic time scales for underlying wave-particle interactions; (ii) dealing with strong nonlinearities and mutual coupling to transport through pressure, velocity and current profiles, and fluctuation spectra; and (iii) ultimately delivering reliable predictive capability for fast particle distributions that are self-consistent with fusion reactions, RF and NBI sources, MHD activity and short-wavelength turbulence. If successful, improved predictive capability of fusion yield and key aspects of steady-state operation would be achievable, thereby enabling information essential for ensuring steady-state (long-pulse) performance in burning plasmas such as ITER.

Basic research in these additional ISA areas would also be carried out by the five existing projects in DOE's Scientific Discovery through Advanced Computing Program (SciDAC) in MHD, micro-turbulence, wave-particles, and RF, all of which were recently renewed for five more years. During this interval, they would carry the principle load in building the foundation for integration activities. The choice for the next ISA would be made during year 3 of the FSP and would depend on progress in these areas.

The identification of the need for the FSP has been presented in a number of prominent past studies and reports over the past decade [2-6], and the importance of validated predictive simulation capability affirmed prominently in the 2009 ReNeW workshop. Most recently, the Program Advisory Committee (FSP) for the FSP¹ has strongly endorsed both the concept and potential of the FSP. In the Executive Summary of their report of May 8, 2011, these distinguished scientists have stated that after closely following the development of the FSP Plan over the preceding 18 months, they believe that the FSP will:

- enable significant advances in fusion science,
- substantially increase the value of ITER to the U.S.,
- make major contributions to build the knowledge base required for DEMO, a demonstration magnetic fusion reactor, and

¹ Douglass Post, *Chair (DoD)*, Allen Boozer (Columbia U), Leslie Greengard (NYU), Brian Gross (GFDL), Greg Hammett (PPPL), Wayne Houlberg (ITER), Earl Marmor (MIT), Dan Meiron (CalTech), Jon Menard (PPPL), Mike Norman (UCSD), Rick Stevens (ANL), Carl Sovinec (U Wisc), Tony Taylor (GA), Jim Van Dam (U Texas)

- provide one of the few opportunities available for the U.S. to provide recognized leadership in the international fusion science community.

They conclude with the statement: *“A Fusion Simulation Program of the type proposed provides the most credible path forward for the integrated whole device model that will be highly important for the realization of fusion energy.” [8]*

A.2 Why Are These Initial Topical Areas Important for Fusion Energy?

The initial topical areas of focus address important scientific questions with urgent practical impact. Burning plasmas in ITER and future fusion reactors represent a significant extrapolation in energy exhausted to the wall and a tremendous extrapolation in pulse length relative to current experiments. While ITER has been designed based on simplified models, actual operation will demand careful planning based on more accurate modeling of the integrated physics of the plasma. Validated comprehensive models of the plasma boundary and the interactions of the plasma with the wall will be required for predicting key physical processes including:

- heat and particle loads to the first wall along with the accompanying erosion and changes in surface morphology or chemistry impacting the selection and predicted lifetime of plasma-facing material components;
- levels of tritium co-deposition in re-deposited material and tritium trapping in bulk surface material, which is critical for the fusion fuel cycle;
- core plasma contamination by surface emitted material;
- safety issues associated with accumulation of dust that could be dispersed during an unintended vent;
- the plasma density limit; and
- pedestal relaxation phenomena with their potential for delivering large transient heat loads to the first wall.

Similarly, whole device models provide an essential tool for the analysis, planning, and optimization of discharge scenarios for ITER experiments, for current and planned experiments, and for design of future devices. Since the nonlinear interaction of the plasma with control actuators are too strong to investigate by empirical methods alone, modeling to guide exploration has proven to be a necessary and successful strategy. In particular, for ITER and reactor-scale devices, the consequences of disruptions will be severe. Thus reliable prediction of high-performance plasma evolution and the margin for avoiding disruptions will be an indispensable part of discharge planning. Validated whole device models of high-performance (H-mode) plasmas that include improved capability to avoid major disruptions will be required for delivering predictive tools needed to address:

- onset conditions for major disruptions in high-performance tokamak discharges, including ITER burning plasmas;
- disruption avoidance margins to support the design and implementation of diagnostics for the detection and feedback control systems needed to mitigate disruptions;
- estimates for amount of auxiliary heating power from radio-frequency and/or neutral-beam-injection needed to be deposited in the core plasma to produce conditions relevant to the evolution of a burning plasma;
- dependencies/sensitivities of overall fusion performance on properties of advanced plasma physics components in a properly integrated model;
- the validated simulation database needed to enable reliable extrapolation of key information from existing confinement experiments worldwide to ITER burning plasmas; and

- between-shots analysis capability for long-pulse ITER discharges needed to cost-effectively harvest the burning plasma science from this major investment; and
- foundations for further improved tools needed to design a demonstration magnetic fusion reactor.

A.3 Why Do These Topics Require Integrated Multi-Physics Simulation?

Models for simulating magnetically confined fusion plasmas must address a wide range of temporal and spatial scales, nonlinear interactions between the plasma and electromagnetic fields, strong anisotropies, and, particularly near the plasma edge, strong coupling between the plasma and (i) non-plasma physics phenomena; and (ii) significant non-thermal particle populations. Historically, the computational approach taken has been to divide the problem into separate domains, each with a limited range of scales. Examples include RF models that are valid on time scales comparable to the inverse cyclotron frequency, and turbulence codes that are applicable on the inverse diamagnetic frequency scale. While substantial progress has resulted, this approach is fundamentally inadequate for many realistic problems. Clean scale separation is only an ideal; in reality, strong ordering is often not justified. Further, additional physics enters in important ways – nuclear reactions, atomic physics, neutral transport, radiation transport, and plasma-material interactions – that cannot be ignored or treated as small perturbations. Thus, studies of isolated physical effects, while essential, must be integrated and embodied in more comprehensive codes. Predicting the dynamics of magnetically confined fusion plasmas is a true scientific grand challenge.

A.4 What Are the Main Physics and Computational Challenges?

The FSP will be addressing multi-scale, multi-physics problems in complex geometry. Integrated models must deal with the lack of spatial and temporal scale separation or strong ordering on which most current codes are based. Particularly in the edge, profile gradients, the ion gyro-radius, neutral and photon mean-free-path overlap, and key dimensionless parameters vary by orders of magnitude over very short distances. For realistic problems, the magnetic topology is complicated – including open and closed field lines, 3D structures, and stochastic regions. While the plasma is highly anisotropic, neutrals are not, and plasma wall interactions need to account for 3D material interfaces. Additional multi-physics phenomena need to be modeled including atomic physics for hydrogenic and impurity species and a wealth of materials issues concerning plasma-wall interactions and the evolution of surface chemistry and morphology. Whole device models that include disruption avoidance must couple all relevant phenomena at sufficiently high levels of fidelity. This will require a library of capable physics components, adapted for use in a reliable and flexible framework. The framework needs to be flexible enough to allow the coupling with 1d, 2d, 3d or higher components and to support explicit and implicit coupling (stable for $>10^5$ time steps) and dynamic parallelism. High fidelity, whole-device modeling capabilities in Fusion Energy Sciences will demand computing resources in the petascale range and beyond to address ITER burning plasma issues. Sustained petascale and even exascale (10^{18} floating point operations per second) platforms (expected to appear circa 2020) will be needed to meet the future challenges of designing the DEMO demonstration power plant. Effective use of computational resources on this scale will require significant coordinated research involving collaboration between fusion scientists and specialists in applied math and computer science [9].

A.5 What Are the Research Opportunities in Applied Mathematics and Computer Sciences?

Success of the FSP depends critically on leveraging expertise from within the broader advanced simulation community. Building on the foundations established at the FSP community planning workshop in San Diego (Feb. 2011), the FSP Plan features research contributions that can both accelerate progress in FES as well as in the research mission of DOE's Office of Advanced Scientific Computing Research (ASCR). Research from within the ASCR community will be required in seven general areas:

- (i) **Scalable Solvers** – solver techniques, especially for highly parallel or multithreaded hardware;

- (ii) **Time Integration** – improved time integration techniques, especially for coupled partial differential equations;
- (iii) **Formulation** – Innovative formulation of continuous and discrete models
- (iv) **Multi-scale/physics** – advanced methods for multi-scale and multi-physics coupling;
- (v) **Data/Meshing** – advanced methods for more efficient data management/analysis, including visualization and meshing;
- (vi) **Frameworks** – framework design, including the software challenges of componentization and coupling on high performance computing (HPC) systems (i.e., systems with fast interconnects such as the Leadership Computing Facilities (LCF)); and
- (vii) **Verification and Uncertainty Quantification (UQ)** – application and further development of methods in verification and uncertainty quantification.

Although specific to the targeting of the FES application domain in the context of the FSP, these categories in general span a wide range of current ASCR activities.

The WDM/Disruptions-avoidance ISA task clearly requires research that is mutually relevant to the achievement of scientific mission goals of both FES and ASCR. Fundamental to WDM is the coupling of multi-physics and multi-code models that span a variety of different physical domains and associated governing equations. In the near-to-intermediate term, the convergence and robustness of existing WDM methodologies can be expected to benefit from the application of: (i) more modern nonlinear solver techniques such as accelerated fixed-point or Jacobian-free Newton-Krylov with physics-based preconditioning; (ii) stronger coupling through implicit time integration techniques; and (iii) sensitivity analysis of critical off-diagonal block coupling terms. The robustness of computing magnetic equilibria could be improved through alternative nonlinear solvers such as Jacobian-free Newton-Krylov with physics-based preconditioning; (ii) stronger coupling through implicit time integration techniques; and globalization/backtracking methods, but further analysis is required to determine if convergence difficulties stem from the linear or nonlinear solvers (or both). Significant nearer-term improvement in predicting the onset of disruptions requires: (i) improved quantitative understanding of the impact of uncertainties and of errors in the input data and magnetic reconstructions on the stability and bifurcation diagrams and (ii) an efficient and robust data management system that helps to synthesize experimental and computational data. In the longer term, the challenge of coupling extended MHD models to kinetic models will require: (i) advances in fluid closures that are self-consistent with kinetic fluxes; (ii) strong coupling of overlapping-multiple-time-scale multi-physics, most likely through semi- or fully-implicit time integration with effective preconditioning; and (iii) self-consistent treatment of overlapping spatial scales.

While these ASCR research problems have been identified in the WDM/Disruptions-avoidance area, they are not an exhaustive list. Similar challenges have been identified for FSP tasks in the Edge/Pedestal ISA and for Advanced Components development, such as the need for developing innovative methods for kinetic/fluid coupling. Indeed, the general problem of coupling models for dynamical processes occurring at different scales and governed by different physics is a wide-ranging area of research with deep implications for the simulation community as a whole. If successful, advances in multi-scale coupling will not only enable accelerated progress for the FSP, but they will also influence advanced simulation as a whole with positive impact on the ASCR mission. Other examples of mutually beneficial, mission-critical research are outlined in the full report.

Given the tight integration of this type of research, the FSP plan does not isolate ASCR research within a single management construct. Rather, to ensure maximum integration with the motivating application, the FSP plan distributes these tasks naturally throughout the organizational structure. Specifically, the ISA and Advanced Component teams will be comprised of physicists, applied mathematicians, and computer scientists. The FES and ASCR research communities have already had significant experience in mutually beneficial partnerships with notable success over the years in the SciDAC fusion applications and Proto-FSP programs, which have been similarly structured in collaborative, interdisciplinary teams. By continuing to build on these established collaborations, a significant and growing portion of the Applied Math and Computer Science (AM/CS) community will develop an even stronger familiarity with the more challenging and critically important fusion

simulations. Looking forward, new ASCR programs, such as the SciDAC-3 Institutes and the Exascale Co-design Centers, will undoubtedly provide exciting collaborative opportunities in the near future.

A.6 Expected Outcomes from FSP

The principal aim of the FSP is to foster scientific discovery through development and application of experimentally validated integrated simulations. A successful FSP will enhance the return on investments in ITER and fusion experimental facilities in general and help ensure their success of by supporting experimental planning and enabling the harvesting of scientific insights. The two areas chosen for initial focus, the edge plasma and whole device modeling/disruption prediction, pose important scientific questions whose solution would have a particularly strong practical impact. Overall, the FSP will embody our state of knowledge in a suite of advanced codes under a unified framework and made widely available to the FES community. Benefits to the fusion program include:

- Developing predictive models which improve our capabilities for reliable scenario modeling, especially for ITER, and for the design and optimization of future machines such as DEMO;
- Addressing multi-physics and multi-scale problems that are now treated in isolation, leading to scientific discovery of new phenomena that emerge only with integration;
- Incorporating powerful HPC capabilities to help accelerate scientific understanding and modern software engineering approaches to ensure the reliability, robustness, and ease-of-use of the new tools that are developed;
- Carrying out a rigorous and systematic validation program in collaboration with experiments to put models on the firmest and most realistic possible foundation; and
- Addressing critical science problems through a set of ISAs that will be aided by computing at the extreme scale in collaboration with ASCR over the next decade.

The FSP research represents a unique opportunity for the U.S. to maintain and expand its world leadership in this critical scientific domain by building on the strong base program in theory and computation; collaborating with researchers in applied math and computer sciences; leveraging DOE investment in advanced computational platforms; and exploiting well-diagnosed experimental facilities for code validation. At the same time, the FSP will create a software infrastructure and a set of tools in which component codes interact efficiently and through which larger software collaborations can be created. The collaborations forged with experimental facilities for code validation along with the dedicated production computing support will put the most advanced simulation capabilities into the hands of end users. Based on current experience, we expect that graduate students and young scientists will particularly benefit from this approach. Education and training of the next generation of fusion computational scientists and code users is expected to be an important outcome of the FSP, enabling the U.S. to sustain its current leading position.

The FSP aims to strengthen the collaborations with the AM/CS communities. Capitalizing on experience from the SciDAC program, the FSP is organized to bring researchers with diverse specializations together into closely-coupled interdisciplinary teams focused on high-priority science applications. The benefits will be mutual. Fusion-specific numerical challenges will yield AM/CS research problems of broader relevance to the ASCR mission, and AM/CS expertise will directly support mission-critical fusion energy applications. This collaboration will also help prepare the fusion program to take advantage of new computational architectures, especially at the leadership class facilities.

A.7 FSP Mission and Vision Statements

Mission Statement: The mission of the FSP will be to enable scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration. This requires developing a predictive integrated simulation capability for magnetically-confined fusion plasmas that are properly validated against experiments in regimes relevant for producing practical fusion energy.

Vision Statement: The FSP will:

- Provide the capability to predict confidently toroidal magnetic confinement fusion device behavior with comprehensive and targeted science-based simulations of nonlinearly-coupled phenomena in the core plasma, edge plasma, and wall region on time and space scales required for fusion energy production.
- Integrate the knowledge from key multi-scale physical processes to improve fidelity continually for extending whole-device modeling capabilities beyond current applicability domains.
- Produce a framework in which physics component-codes interact efficiently to enable unprecedented capabilities to compute experimental observables, interpret experimental data, and explore the consequences of theoretical models.
- Incorporate modern software engineering and software quality assurance to ensure the reliability, robustness, and ease-of-use of the tools that are developed.
- Create the most advanced suite of predictive codes under a unified framework and distribute it to and provide support for the fusion community to maximize U.S. investments in experimental facilities (especially, ITER) and in HPC resources (especially, the Leadership Class Facilities) to produce the scientific basis for an economically and environmentally attractive source of energy.

B. HOW THE FSP IS STRUCTURED TO ADDRESS SCIENCE PROBLEMS²

The primary task of the FSP Plan is to develop an integrated computational toolkit for the broad user community that is capable both of carrying out innovative simulations for selected science applications and of providing the foundation for extension to a much broader range of problems. This section outlines the technical organization for the FSP, specifically, how the science application efforts interact with component development, code infrastructure, quality assurance, production computing, and AM/CS crosscutting efforts.

B.1 The Matrixed Approach

The mission of the FSP is not limited to a single science problem but rather aims to develop and integrate broader solutions that are applicable to a grand challenge class of science questions and will enable a more comprehensive whole device model. In doing so, the plan adopted is structured to achieve economy of scale as it grows in an accountable way to a program capable of producing increasingly greater scientific productivity. The associated goal is to deliver, for the first time, a suite of true community codes for production computing with user support that is coordinated by means of a national program dedicated to: software infrastructure and developer support – instead of the current practice of diffuse development and support of individual applications;

- enabling modern AM/CS technologies be deployed for rapid sharing of new tools and approaches between applications;
- efficient integration of best physics components with common interfaces and data structures guided by appropriate set of standards;
- producing FSP-standardized, well-documented tool sets for data preparation, code input validation, data analysis, and visualization FSP standards – instead of the current inefficient customization approach; and

² This summary and the full FSP Execution Plan upon which it is derived are based on assumptions and opinions of the FSP planning team concerning necessary minimal funding amounts and the most appropriate funding mechanisms – information not currently available to the team. Once these issues are clarified by the relevant DOE program offices, the team expects these planning documents to be revised, possibly substantially, to accommodate any new information or understanding.

- achieving a world standard for FES with application of modern verification, validation, and uncertainty quantification (VVUQ) methods to yield more rigor and efficiency via strong coordination with experimental facilities (national and international).
- Thus, it is appropriate to initiate the FSP by beginning with two ISAs under development at the same time. However, this will not by itself lead to an integrated project capable of satisfying the full FSP mission. There would be a strong natural tendency for each ISA team to work independently, resulting in, for example, multiple similar codes for different application areas and inevitably to wasteful duplication of effort. A collection of largely independent research activities has no real hope of cost-effectively establishing the necessary coordinated infrastructure essential to deliver the software products that would enable ITER to harvest, in a timely way, the targeted scientific knowledge base needed for its mission goals, i.e., laying the foundations for moving on to DEMO. The FSP will accordingly employ a matrixed approach, drawing staff from both ISA Teams and cross-cutting Enabling Technology (ET) Teams.

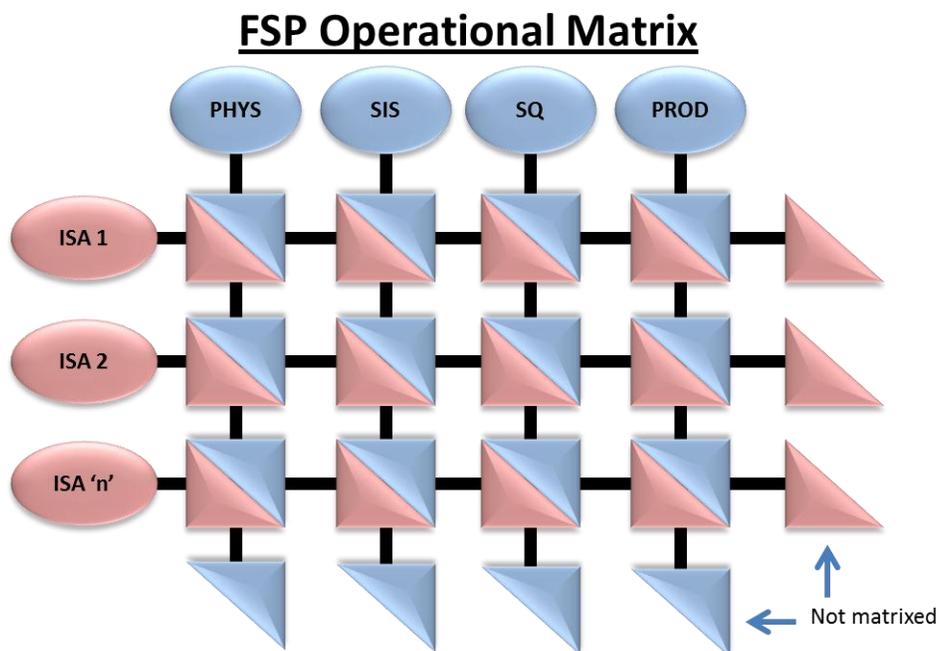


Figure 1: FSP Operational Matrix

Figure 1 graphically represents the FSP operational matrix and the ways in which contributors to the FSP mission relate to one another. It is not an organization chart describing a management structure (See Figure 3 below). In Figure 1, each horizontal ISA row of colored triangles represents a composition of members from both the ISA and ET teams who will work together to produce a particular simulation code release. The oval shapes are simply row or column identifiers and do not represent people. A vertical column represents the coordinated effort to leverage, as much as possible, common code solutions, infrastructure, quality control processes, etc. The unpaired triangles represent un-matrixed staff within a respective ISA or ET team. The lone blue triangles represent ET staff who might, for example, be assigned to work on a physics component that is deemed to be important for future, but not current, ISA efforts. A lone red triangle at the right end of an ISA might, for example, represent an analyst who is working with users to advance scientific discovery using codes developed in that ISA.

I. Integrated Science Applications (ISA) Teams

The FSP will be primarily organized around multiple integrated science application teams each with approximately seven to ten members. Each ISA team is multi-disciplinary and charged with executing the

simulation development plan for a particular science problem; e.g., those summarized within the science driver reports. Each ISA team will have a single technical lead for day-to-day direction. All ISA teams will be overseen by the Deputy Director for Science, who will drive each team toward their technical goals, manage resource distribution, ensure proper interaction with the broader scientific community, and recommend re-prioritization of allocations based on success/failure, new ideas from the community, and new directions from the Program Office.

Through its research and development, the ISA teams will define strategies, address critical problems, exercise the required range of code capabilities, and provide useful tools for the broader fusion community. Working with the Software Integration and Support (SIS) team, they will help define ISA requirements and architecture for common components, infrastructure, and any other enabling science or technology needed to implement the physics integration scheme. For physics components specific to the ISA, the team will work with the component team to adapt or build, ensuring that development work conforms to FSP standards for coding, data exchange and documentation. The ISA teams will ensure regular testing of software components and integrated codes and will document and repair anomalies found in testing using provided infrastructure. Working with the software quality group, the ISAs will carry out a program of verification and uncertainty quantification. They will coordinate partnerships with experimental facilities to plan, execute and analyze validation experiments. They will ensure proper documentation of V&V studies using FSP standard methodologies. At some predetermined interval or when work has progressed to a satisfactory point, the ISA team will authorize the official “release” of FSP code. Finally, the ISA will work with the production support group to prepare code documentation and user’s guides and to provide ISA-specific expertise for user support.

II. Enabling Technology (ET) Teams

The division into ISA teams will not by itself produce an integrated project. In the FSP, the application teams are matrixed with personnel from various ET teams that focus on the more global aspects of integrated code development. Each ET team will have a single technical lead and will be overseen by one of the two Deputy Directors.

Physics Component Team (PHYS)

The role of the component team is defined in relation to other parts of FSP, particularly the integrated science application effort. The component team reports to the Deputy Director for Science and is responsible for well-defined, reusable physics modules that service more than one application. Each identified physics component will have at least one embedded component team member. The role of the PHYS technical lead is to ensure that the group members working within an application area are developing from a common code base, that code improvements targeting a given application driver are being built into that code base, and that the methods used in the physics components are verified. There are three specific roles for the component team as a whole.

First, the component team is a capability organization. It holds the technical capability in developing advanced physics components to be integrated into one or multiple science applications to address one or more science drivers. The component team has both regular members, who are appointed for the entire FSP execution phase, and collaborative members, who participate on the basis of individual component projects.

Second, the component team provides stewardship of the FSP component library by continuously standardizing and maintaining the suite of physics component codes for FSP science applications. It is through the stewardship of the component library that a common set of standards and best practices are introduced and applied to the FSP physics component development. The three primary activities are (i) publishing the component standard on data interfaces and documentation for component developers; (ii) performing acceptance test and review of a newly developed component; (iii) carrying out further improvement, maintenance, and regular testing of the component in the component library.

Third, the component team plans and executes, or manages the execution of, new component adaptation/development projects and the related enabling exploratory research and prototyping.

As an illustrative example, the FSP component team recruits leading subject experts in computational MHD and maintains a suite of standardized MHD component codes. It initiates and carried out development projects of adapting existing fusion MHD code into FSP and of prototyping new physics models and/or numerical algorithms in response to the evolving science driver needs.

Software Integration and Support Team (SIS)

SIS has the mission of providing the composition software for integrated computation, providing and/or supporting the software for job setup, data analytics, visualization and data management, as well as providing support for software development throughout the project. Software integration has been separated into two areas: 1) *On-HPC* integration, which at present is nearly exclusively the integration of physics components to run together on a High-Performance Computer; and 2) Task Composition (*Off-HPC* integration), which is concerned with all other integration and development and/or support of the associated modules that are needed to go from initial concept to final research result. For example, included in task composition is the development of specialized fusion plug-ins for reading data into visualization tools. Developer support includes providing and maintaining the collaboration systems, such as software repositories and communication lists, and implementing or assisting with software engineering issues, such as build and test systems and performance measurement.

Software Quality Team (SQ)

The Software Quality team has responsibility for ensuring the reliability of the FSP software targeted for release to the community. The technical team leader receives oversight from the Deputy Director for Code Architecture and chairs the Software Quality Board, which is composed of members designated from the other areas. The SQ team leader coordinates the teams tasked with software quality management, which involves software quality assessment, testing, verification, validation, and UQ. The implementation of these activities is undertaken by all members of the team. Associated tasks of the SQ team include: (i) developing standards for software development and testing; (ii) reviewing plans and progress on software quality activities across the entire FSP program; and (iii) organizing software reviews prior to release. The SQ technical staff provides crosscutting tools and technologies such as testing systems. In addition, research into new techniques for verification and UQ falls under the auspices of the SQ team.

Each ISA team will have an identified contact from Software Quality. That individual will be responsible for coordinating the relevant activities involving software QA, testing, verification and validation. This coordinator will also serve on a Software Quality Review Board, chaired by a FSP-wide software quality manager. The SQ team reports to the Deputy Director for Code Architecture.

Production Support Team (PROD)

The Production Support Team also reports to the Deputy Director for Code Architecture. PROD will support production versions of FSP software for research applications by end users. Such end users can come from within the FSP project; e.g., FSP software quality experts performing uncertainty quantification or FSP analysts performing physics validation studies, or from outside the FSP; e.g., experimental facilities or theory and computation base programs. End users will be expected to understand physics issues involved in running the FSP software but will not be required to have FSP code developer skills.

The Production Support Team will deploy production versions of FSP software, as identified by the code development teams, on specific computational platforms supported by the SIS team. The team will make available the means for end users to prepare input, submit runs, monitor runs, and examine output. The team will make sure that run output are transferred to FSP data management facilities provided by SIS, with a record of the production code version and copies of all input data preserved.

With the active assistance of FSP developers, the Production Support Team will provide user documentation for FSP code, provide user support, and, with backing of the FSP code development teams, provide trouble-shooting for failed runs or for problems in the supporting software for data preparation, job submission, monitoring, etc.

The Production Support Team will work with users as a key part of the process for continuous improvement of FSP software.

B.2 Management Approach

The FSP management approach enables an achievable, requirements-driven plan for a focused R&D effort that is an order of magnitude larger and more challenging than any similar projects attempted in magnetic fusion research. It will follow the best practices of successful large scientific software programs/projects in other fields – such as DOE National Nuclear Security Administration’s Advance Simulation and Computing (ASC) Program and Climate Modeling. Important roles of the FSP management approach will be to manage risks proactively, utilize formal project management processes to plan, monitor, and manage the majority of the effort of the program, and to integrate quality control fully throughout the program. The FSP will, as it has done in its planning phase, carry out its outreach responsibilities by conducting a regular annual cycle of community workshops. The FSP will also continue to form individual ISA planning committees to ensure broad community engagement and to assure that the ongoing FSP understanding of science drivers and their requirements is current and responsive to evolving fusion energy research community needs.

The following two sections identify (i) the relationships that FSP has with external organizations or other entities from which the FSP receives direct support, direction, and/or guidance and (ii) the organizational components within the FSP.

External management relationships

FSP External Management Relationships

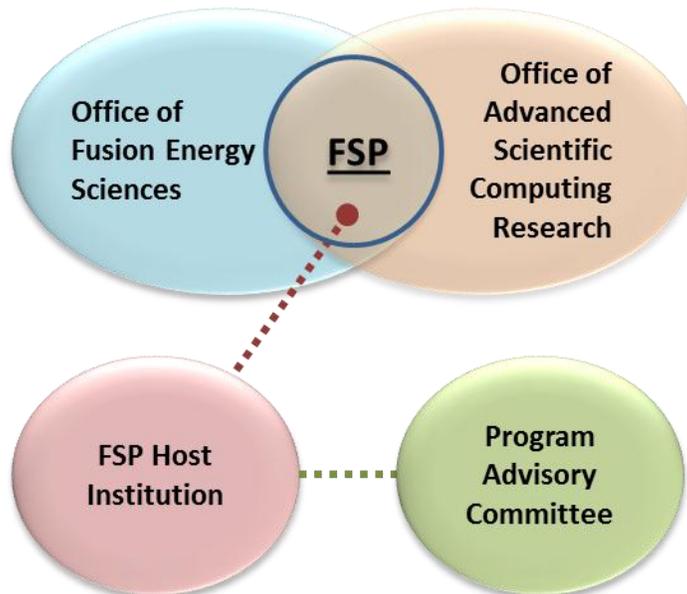


Figure 2: FSP External Management Relationships

FES

The FSP directly supports the efforts of the FES to develop integrated simulation capabilities to further its mission and support its strategic goal to “advance the fundamental science of magnetically confined plasmas to develop the predictive capability needed for a sustainable fusion energy source.”

ASCR

ASCR recognizes the synergistic advantages available to it in achieving its stated goals by joining FES in supporting the efforts of the FSP. These advantages are described further in Section C.3.

The work of the FSP does not solely reside within the purview of either program office, but at the intersection of interests of both organizations.

The Host Institution

The host institution provides the management processes and infrastructure together with practical financial control mechanisms and other institutional support (office space, communications, etc.) as well as the necessary day-to-day oversight required to help ensure a successful program. The FSP Director will be appointed by the host institution with the concurrence of DOE Office of Science (DOE-SC).

Program Advisory Committee (PAC)

This is an external group of experts, reporting to the Director of the lead institution and providing advice on a broad range of technical and management issues. The PAC will meet at least once per year and address a charge formulated by the Director of the host institution and the FSP Director.

Internal management roles and responsibilities

Figure 3 is an organization chart describing the internal reporting lines within the FSP. The entities represent one or more people within the organization who are ultimately responsible to the FSP Director. The chart entities do not necessarily represent unique individuals, especially in the initial FSP ramp-up periods. One person could hold several organizational positions.

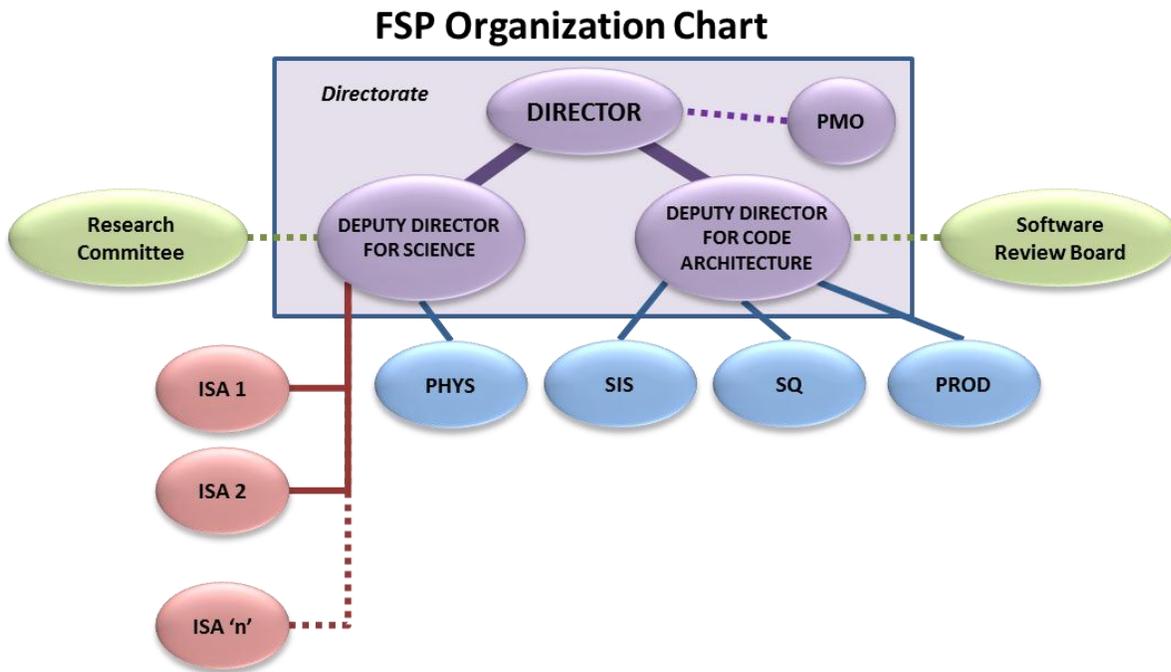


Figure 3: FSP Organization Chart

The FSP is managed through a Directorate comprised of a Director, two Deputy Directors, and a Project Management Office (PMO).

- Director
 - The FSP Director has overall responsibility to ensure that the scientific and software development goals are properly executed on time and within budget. The Director is the principal contact with DOE-SC and with the senior management of the FSP home institution. The Director has oversight of the core FSP management team. With guidance from DOE-SC, final decisions for project prioritization, resource allocations, and personnel are made by the Director with guidance provided by the Deputy Directors.
- Deputy Directors

The FSP Directorate also includes two Deputy Directors, one for Science and one for Code Architecture. These positions provide the top level communication channel to ensure the cross-cutting functional groups work together in a seamless way.

Deputy Director for Science has oversight of the Integrated Science Applications (ISA) and Physics Components (PHYS) teams. The Deputy Director for Science drives the scientific goals of the FSP, ensuring application projects are well-balanced and making appropriate progress. He/she also recommends funding allocations (by area) based on the need to balance short/long-term progress, to address the priorities of the DOE-SC Program Office, and to respond to community feedback.

Deputy Director for Code Architecture has oversight of the Software Integration (SIS), Software Quality (SQ), and Production Computing (PROD) teams. The Deputy Director for Code Architecture is responsible for the management of the overall FSP code repository as integrated software. Principal roles also include: (i) driving both the applied math/computer science research and applied project software goals of the FSP and (ii) ensuring that an integrated “community code” (suite of tools) flows from the ISA projects and lives within a proper software development lifecycle, including documentation, testing, versioning, and repository management.

- **The Project Management Office**

The PMO is responsible for establishing, in consultation with the Director and the host institution, the project management standards, policies, and procedures to be followed for all FSP activities established as projects, which is expected to be the majority of work performed by FSP. This will include such things as the level of project planning detail required, review processes, how cost and progress is to be reported, and how change and risk management activities will be monitored, controlled, and reported. The PMO is also responsible for coordinating the tracking of all project costs and progress and developing periodic reports for the Directorate and the host institution.

Research Committee

The Research Committee is composed of the FSP technical leadership associated with the ISA, PHYS, SIS, SQ, and PROD teams and the FSP directorate. It also includes representatives of major collaborating groups (in particular those providing experimental facilities for code validation). The Research Committee provides advice on a broad range of research planning activities including assessment of priorities for R&D, preparing work proposals, and organizing publications and presentations. This committee will provide well-documented findings and recommendations to the FSP Director.

Software Quality Review Board

The Software Quality Review Board is chaired by the technical lead of the SQ Team and is composed of members of the ISA, PHYS, SIS, and PROD teams. This board provides standards for software development and testing, reviews plans and progress on software quality activities across the entire FSP program, and facilitates software reviews prior to release.

B.3 Approach for Dealing with Distributed Project Team

A fundamental challenge of the FSP involves successfully addressing the multi-institutional, multi-disciplinary nature of the program. It is vital to ensure a concentrated level of effort with genuine commitment from the partner institutions. Associated plan requirements include:

- Some level of co-location to facilitate institutional commitments and efficient team-building. This can involve co-location at the team level; e.g., an institution could have overall responsibility for executing a specific mission/task with support from other partner institutions. The majority of the team of a particular task could also reside in one location if it holds the expertise to best accelerate development of the necessary tools.
- Making use of established relationships with a “critical mass” of talent. This requires sufficient engagement of performers; i.e., a low percentage engagement of performers should be avoided.

- Effective utilization of modern collaboration tools (video-conferences, teleconferences, wikis, etc.).
- All FSP codes and tools will be made available via a centralized repository located at an appropriate DOE-SC computational resource center. FSP data and documentation will be shared by all members of the program. A common approach for data management will be taken FSP-wide.
- Specific detailed contracts for performance ascribed to by all paid contributors to the FSP mission that assure that FSP standards and procedures will be followed. Otherwise, the FSP will lose management control of outcomes for which it will be held responsible.

Since the FSP is distributed with respect to people, computational facilities, and experimental facilities, there is obviously a great need for effective collaborations that require:

- A detailed understanding of FSP leadership roles and associated responsibilities beginning with DOE-SC;
- Frequent communication between the nationally distributed teams – flowing from DOE-SC to the front-lines where the FSP Plan will be executed; and
- Location of FSP analysts at key experimental facilities.

B.4 How Will FSP Stand Up and Operate

The FSP will stand up and operate by following an integrated project schedule/Gantt chart with measurable milestones and deliverables. This is presented in more detail in Section D below where the main FSP execution plan features are specified. The logic here is that the deliverables provide the means for measuring progress towards the goals and objectives of the FSP effort consistent with the associated project schedule. The different threads of the activities undertaken should be viewed as being connected to reflect the integrated aspect of schedule. This requires: (i) a clear understanding of roles/responsibilities; (ii) appropriate scheduling of meetings to help monitor feedback of results/progress; and (iii) implementing an effective change control strategy for re-evaluation/re-validation of the FSP requirements targeted.

FSP operations will follow a *work plan* with a fully budgeted and scheduled work breakdown structure (WBS) that provides specific information on who is to do what and when they are expected to do it. Interdependencies between WBS elements will be identified and tracked. Given the research nature of the FSP and the concomitant uncertainty of progress, the WBS will be updated at least annually after input from the community to accommodate any new understandings of requirements. However, between updates, cost and schedule baselines will be closely monitored to assure that effective progress is being made. Periodic review of performance will be conducted and actions taken if required.

The FSP plan addresses product delivery and the responsiveness of the FSP to key stakeholders, including, for example, the ITER directorate (including the U.S. ITER head), the FES facility directors, the ASC associate directors at the Leadership Class Facilities, etc. An associated integration strategy within the FSP Plan describes how the FSP will integrate with, coordinate with, and leverage other U. S. programs with shared interests and imperatives – such as the FES theory and experimental programs, the ASCR applied math and computer science programs, the HPC resource center programs at ORNL, ANL, and NERSC, and the FES/ASCR SciDAC programs. With respect to possible international collaborations, it will be necessary to develop an effective approach to interaction and coordination with integrated modeling efforts abroad (e.g., for ITER partners such as Europe and Japan) as well as with international facilities, such as the long-pulse tokamak facilities in China and Korea that have capabilities not available in U.S. facilities.

In standing up and operating the FSP, special attention will be paid to the integration of the program elements, including synergism with discovery science efforts in the FES and SciDAC programs. For example, advanced physics components/modules development can span the entire field of fusion plasma physics and the interdisciplinary area of plasma/material interactions, while making use of a wide variety of computer science

and applied mathematics expertise and resources. The first task of the FSP components development activity in this area will be to proceed with acting on the advanced physics components gaps identified by the science drivers/ISA teams that are needed to capture the understanding absent from existing FES simulation capabilities. Once the new scientific discovery capabilities are developed and demonstrated via verification and validation to be desirable new tools, a plan for community-wide assimilation of FSP component capabilities and user support for the use of these capabilities in scientific discovery must be developed. The FSP component and framework teams will also collaborate on the planning of possible reduced models that can be readily adapted into the integrated physics framework. These tasks will include the development of the principles and concepts needed to address issues such as flexibility in resource allocation to maintain balance among key FSP tasks areas.

B.5 Risk Management

Overall, risk management involves a continuous, forward-looking process that identifies and mitigates or otherwise manages issues that could endanger achievement of critical FSP objectives. Early, frequent, and aggressive detection of risk is important because it is typically easier, less costly, and less disruptive to make changes and correct work efforts during the earlier, rather than the later, phases of any activity. As illustrated in the following examples, a FSP risk is defined as a future event, action, or condition that might prevent, delay, or increase the cost of the successful execution of FSP activities. The FSP Risk Register in the FSP Plan, Appendix F, identifies and summarizes the most prominent set of FSP risks as well as provides an approach for managing these risks with proper triggers and mitigation actions identified. The list of FSP risks will be periodically updated, evaluated, and reported throughout the development of the FSP. All reasonable risks will be considered with respect to technical, programmatic, and budgetary aspects.

The FSP approach to risk management identifies potentially critical "bottlenecks" and provides alternatives in the event the proposed methods prove to be unsatisfactory. Key technical risks with possible mitigation approaches that could be deployed if needed are summarized in the Risk Register Table within the FSP Plan. In addition, other kinds of risks associated with funding variances, efficiency of execution, personnel issues, etc. are also articulated. The risk register will be periodically updated, evaluated, and reported throughout the development of the FSP. All reasonable risks will be considered with respect to technical, programmatic, and budgetary aspects.

Examples of risks from some (but not all) of the FSP work-scopes along with possible mitigation approaches include:

- A general ISA risk from an organizational perspective includes difficulty in recruiting teams with sufficient skills and ability to make required time commitment. This impacts schedule with respect to delay in starting the FSP activities as well as slower progress throughout. Possible mitigation includes more active talent recruitment from outside the FES program together with dedicated resources and time for training.
- A technical risk for ISA1 is that the gyrokinetic formulation for the edge/pedestal region may not be sufficiently complete and/or computationally tractable. Schedule delays could result while theoretical work is under way to remediate such issues. A mitigation approach would be to engage the theory community proactively with the support of FES as well as providing contingency FSP funds to address such problems, should they arise. Another ISA1 risk is in dealing with plasma-wall models that might be insufficiently complete and could reduce the fidelity of edge plasma predictive simulations. Mitigation could include shifting a modest level of FSP resources to wall modeling – while being cognizant of delays in lower priority/less time-urgent FSP tasks. However, if this proves inadequate and resources and expertise beyond levels appropriate for the FSP are required, it will then be necessary to work with the DOE-SC leadership to plan an effective partnership between the FES and Basic Energy Sciences (BES) offices to establish dedicated funded partnerships between these science communities.
- Representative technical risks for ISA2 include timely delivery of: (i) an appropriate FSP framework into which advanced physics modules can be integrated; and (ii) a reliable free-boundary Grad-Shafranov

solver for plasma equilibria. If the new FSP framework is not available in time, mitigation alternatives could include proceeding with WDM activities using a framework that is available in an existing 1.5D integrated modeling code. Another significant risk is the potential delay in the development of the FSP free-boundary equilibrium solver including structures and PF coils. The mitigation approach will be to use a fixed boundary equilibrium solver such as TEQ or VMEC until a suitable free boundary equilibrium is delivered.

- A key risk associated with experimental validation is insufficient experimental run-time on U.S. facilities. This would impact schedule and deliverables and could lead to reduced fidelity of FSP codes associated with a reduction of scope in the validation campaign. An associated mitigation approach would be to work closely with domestic experimental facilities as true partners in planning and also to seek allocations proactively of experimental run-time on international facilities that would benefit from collaborations with the FSP. Another significant risk is associated with the adequacy of diagnostic capabilities as validation tasks demand finer granularity. Without remediation this could also limit the validation scope and accordingly reduce an otherwise higher achievable level of fidelity. Mitigation approaches would require early planning with experimental facilities, working closely with FES experimental scientists/diagnosticians, and pro-actively engaging with international partners.

Any realistic risk management strategy must include an associated “change management” strategy/plan to deal with the inevitable changes required in multiple aspects of an ambitious effort such as the FSP that must deal with major R&D challenges. The FSP begins with the recognition that built-in redundancy of skill sets is not the most cost-effective/realistic path to dealing with redirection of efforts when needed. Moreover, much more than “technical coordination” alone will be demanded to ensure success. Change management is fundamentally needed for maintaining a coherent view (understandable by all) of a large distributed project such as the FSP. This includes articulation of a clear *decision authority* for such changes – an essential aspect of any distributed, multi-institutional project like the FSP. The plan for change process management deals with issues from an annual as well as a 5-year perspective. This includes identification of the inter-dependencies between the different functional and/or geographically distributed teams. Project Status information will be prominently displayed on the national FSB web-site, including issues, status, and final resolutions.

C. INTERACTIONS WITH PARTNER PROGRAMS

C.1 Interaction with Experiments

The success of FSP validation hinges on building an effective partnership between FSP and the experimental facilities. Both have to recognize that there are significant benefits towards advancing the fusion energy goals by sharing resources and making experimental validation a high programmatic priority. A document (“Principles for Collaboration on Major Experiments” in https://ice.txcorp.com/trac/2011_FspDefinitionWorkshop) describing how the partnership will be implemented has been prepared based on ideas drawn from existing collaboration agreements used by the three major fusion experimental facilities and their governing and planning processes. The managements of the major experimental facilities have reviewed and approved this document, providing constructive feedback. The document outlines general principles for collaboration and intellectual property sharing with major facilities:

- Use of collaboration agreements to support the partnerships and to ensure proper use and proper credit for shared data, codes and other intellectual property;
- Interactions with facilities on planning and how smooth running collaborations require interactions at both working and management levels and timely input into the planning processes (both run time and long-term) of each program or project;
- Well-defined roles for the FSP and the experimental teams wherein:

- FSP provides:
 - Codes suites and requests for associated allocation of computer time ;
 - Help in understanding code capabilities and limitations;
 - Dedicated analysts;
 - Consideration of code developments, based on needs of the experiments.
- Experiments provide:
 - Run time, subject to local planning processes;
 - Access to data;
 - Support for diagnostic data analysis; and
 - Considers and requests appropriate upgrades, including for diagnostics, based on the needs of the simulation program.
- The two collaborate on:
 - Setting priorities;
 - Run planning;
 - Experiment analysis and interpretation;
 - Development of synthetic diagnostics;
 - Physics interpretation;
 - Preparation, presentation and publication of results; and
 - Lessons learned from experimental facilities for FSP in terms of organizing its own research efforts.

Over many years the major facilities, which are large, heterogeneous organizations, have developed processes and policies that strengthen their research programs. The FSP will continue to learn from their experiences and adopt approaches that could benefit its own efforts.

C.2 Interactions with Fusion Theory/Computation/SciDAC

Base theory/computation program: The FSP component development will rely on the base theory program for developing the theoretical basis of the physics models that will be implemented in the component codes. These models include, but are not limited to, improved closures for the fluid moment equations, a more complete set of gyrokinetic equations that apply to the tokamak edge, and a mathematically rigorous formulation that couples neoclassical and turbulent transport to the quasi-static evolution of the three dimensional magnetic field. The research needs identified in the FSP program will feed into the base program, motivating new solicitations to address them. The base theory program also provides the exploratory research that, upon a proof-of-principle demonstration, can lead to a more complete component development under FSP. In the near term, the base theory/computation program will provide most of the code candidates to be adapted into the FSP for the two ISA's in the first two years. This is particularly true for almost all components for profile evolution and for the entire suite of stability codes.

SciDAC program: The FSP will rely heavily on the fusion SciDAC program for component candidates to be adapted into the component library for use in multiple science driver applications. The candidate components tend to be high fidelity, advanced physics codes that target complex plasma dynamics with first-principles models. Very often, such codes use large scale computing effectively to tackle the most difficult fusion science problems. The initial value MHD and gyrokinetic codes from the SciDAC programs are such examples. In areas where substantial gaps exist between the existing SciDAC capabilities and the FSP ISA requirements, the FSP will coordinate with the SciDAC program to make complementary investment on new capability development.

C.3 Interactions with AM/CS

The FSP will rely on ASCR-funded base program activities and will work with them broadly. Through programs like SciDAC, the FSP will look to FES and ASCR to continue to fund research into methods and models for treating physical mechanisms not included in the early FSP efforts (e.g., integration challenges associated with wave-particle interactions). At a later stage, the FSP will assimilate such capabilities into future coupled models. Similarly, the FSP will make use of leadership-class computing capabilities, and every effort will be made to

develop frameworks and algorithms that anticipate the next generation of computing architectures. It should also be kept in mind that the FSP will need to interact with a community where much of the scientific research will continue to utilize desktops, clusters, and capability-class machines. Nevertheless, since higher physics fidelity in the FSP suite of codes will always be in demand, the mission goals of this program will clearly require existing as well as evolving ASCR expertise in high-performance computing. While the dedicated mission and associated funding to address the exascale challenge lie outside the immediate scope of the FSP, the prominent ASCR investments in algorithms for future massively parallel architectures will be leveraged by the FSP in order for fusion simulation to make efficient use of the ASCR LCF capabilities at the petascale on the path toward the exascale and beyond.

C.4 Need and Plans for HPC Resources

The FSP planning activity has engaged in active discussions with NERSC, both with one-on-one discussions with the NERSC leadership at LBNL and also as part of a major NERSC Resource Requirements Workshop held in Washington, DC in 2010. It can be envisioned that, if fully funded in a sustained manner, the FSP could double the scale and scope of the current MFE computation program. Thus, a rough estimate can be made by extrapolation from related computational programs in MFE, especially the proto-FSP projects that have been operational from 2006 to the present. Over a five year horizon, such extrapolated estimates for FSP HPC resource requirements can be summarized as follows:

- Tens of thousands of small runs using thousands of cores;
- Hundreds of medium-scale runs using tens of thousands of cores;
- Dozens of large jobs using in aggregate >1,000 core-hours on one million cores; and
- Memory requirements from 0.1 GB/core for the largest jobs to 2 GB/core for small and medium runs.

The majority of HPC runs will take place on the *capacity computing* resources at NERSC. In addition to the above resource requirements, computing environment and policy needs were also identified:

- Support for production computing including a “Simulation as a Service” model.
- Integrated data management, long term storage, and advanced cataloging of modeling and experimental data.
- Support for off-line analysis of large data sets on systems that facilitate efficient data access from storage.
- The availability of required libraries and other supporting software.
- The ability to set job priorities within the FSP domain; e.g., fast turn-around is often required for smaller jobs, especially in support of code development, verification and validation.
- Adequate CPU hours for software development (advanced components and frameworks), for V&V and UQ testing, and for production services.
- Tools to enable the tuning of systems for job mix, i.e. helping users identify the most cost-effective platform for each job with flexibility.

With regard to *capability computing* requirements, associated resources reside primarily at the LCFs at ORNL and ANL. The FSP will pursue access to these resources following the guidelines from prominent ASCR programs such as ASCR's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program and other possible arrangements made available by ASCR.

With respect to FSP data storage needs, the goal is that all FSP runs across all platforms, regardless of physical location, will be catalogued. A Universal Access Layer is planned for location-independent data access. Only rough notional estimates of data storage needs can be made, and these indicate that (i) aggregate archival storage is likely to be in multi-petabyte range in thousands to tens of thousands of files per year and (ii) temporary storage needed by jobs during runs is also predicted to go into the petabyte range.

D. EXECUTION PLAN

D.1 Plan for Initiating the Program

In order to provide an initial schedule, the FSP Execution Plan was developed under the following assumptions:

- That the process for the eventual awarding of a DOE contract to execute the FSP will include decisions on which institution will host the FSP and who the Director and Deputy Directors will be. Therefore, the Directorate will be in place at Program start.
- That the funding for the Directorate will be in place at Program start and that management and oversight mechanisms for the entire FSP effort will be in place and fully functional within two months of Program start.
- That the initial ISA teams will be selected by some, as yet unspecified, process that will identify partner institutions and the ISA technical team lead within three months of start and that the ISA teams will be fully staffed within six months of start.
- That the ET teams (lead and other staff) will be selected by the FSP Directorate and be fully in place within three months of Program start.
- That the PMO will be staffed and fully functional within one month of Program start.
- That the effort levels start at \$12M for the initial year and increase by \$3M per year until \$24M when the FSP is maximally configured and operational.
- The milestones listed below are based on bottoms-up estimates of resource requirements, however they are deliberately aggressive as the planning team felt it was important that the program have a significant impact as soon as practical. In the assessment of program risks, most risks had schedule delay as a consequence. Risk mitigation strategies have been defined, but cannot eliminate the possibilities. The expectation is that progress will be reviewed regularly, especially in the early stages of program execution. Based on these assessments, the out-year schedule could be modified.

The table below provides a quick reference to the major milestones expected for the first five years.

Deliverable / Milestone	Applied Math / Computer Science Research	Expected Date from Award
1.4: FSP operational – management and technical teams sufficiently staffed, funding mechanisms in place, FSP policies and procedures in effect.		6 months
5.1.2.4: Common I/O capability with consistent metadata available to fusion community		21 months
5.4.3: Availability of all FSP software on line and continuously updated		21 months
3.2.8: First release of FSP 1.5 WDM code	<ul style="list-style-type: none"> ● Scalable Solvers ● Time Integration 	22 months
4.3.1-4: Release of library of adapted components including Grad-Shafranov solver and embedded turbulence model	<ul style="list-style-type: none"> ● Multi-scale/physics ● Frameworks 	23 months
2.2.4: First release of static model within FSP framework	<ul style="list-style-type: none"> ● Scalable Solvers 	32 months
5.1.5: Reference implementation of On-HPC integration software, Release 1 (concurrent components, low-dimensional couplings)	<ul style="list-style-type: none"> ● Data/Meshing ● Frameworks 	32 months

3.5.5 Release code with gyrokinetic turbulent simulations included	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Multi-scale/physics • VUQ 	41 months
2.3.3: First release of coupled Scrape-Off Layer (SOL) model	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Multi-scale/physics • Data/Meshing • VUQ 	53 months
4.4.1: Complete new component development of profile evolution with 3D equilibrium	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing 	59 months
2.5.3: First release of coupled kinetic SOL model	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing • VUQ 	68 months
3.4.8: Release WDM code with 3D core and pedestal models	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Multi-scale/physics • Data/Meshing • VUQ 	73 months
3.6.5: Release code with combined ISA1 and ISA2 components	<ul style="list-style-type: none"> • Multi-scale/physics • Data/Meshing 	83 months
2.6.2: First release of dynamic pedestal model	<ul style="list-style-type: none"> • Scalable Solvers • Time Integration • Formulation • Multi-scale/physics • Data/Meshing • VUQ 	88 months

Table 1: FSP Deliverables and Milestones (Quick Reference)*

**Relevant WBS reference numbers included. Definitions of Applied Mathematics / Computer Science terminology can be found in Section A.5.*

D.2 Plan for ISA's

Overview

The ISA teams have the end-to-end responsibility to develop, test and deliver integrated code suites. In carrying out this task, they will help set requirements for common tools, components, and infrastructure shared across the entire FSP program. As noted, based on programmatic urgency, current readiness, and applicability for integration, the initial focus will be in two application areas: pedestal-boundary and whole device modeling. With significant community involvement, detailed plans for each of these areas have been developed. Goals were defined along with requirements for physics components, framework and infrastructure. Critical physics requiring experimental validation was identified along with the corresponding measurement requirements.

Work was broken into a small set of discrete tasks and milestones defined with estimations of the required manpower resources.

What follows is a brief summary of these plans. Full details can be found in Section 3.2 and in Science Drivers Report in Appendix B of the full program plan.

Roadmaps for Edge Physics and WDM/Disruption Avoidance ISA's

The overall strategy is to build the models in stages with each stage providing new simulation capabilities in “releasable” code and an extensive program of code verification and experimental validation integral to each. Beginning with existing physics code components, each stage builds on the previous, adding new, more coupled physics models that should result in predictions with higher physics fidelity. In this way, the FSP will produce useful research tools and provide new scientific insights early and throughout its life.

Development of the Edge Physics ISA is divided into six task areas accomplished in successive steps.

- Refine a static linear model for pedestal structure. This task builds on existing work, first componentizing and verifying EPED, the linear MHD model that calculates the pedestal pressure profile by finding a solution that simultaneously satisfies global peeling-ballooning and local kinetic-ballooning stability criteria. This calculation would then be extended using quasi-linear gyrokinetics and including the effects of ExB stabilization.
- Model the region from the bottom of the pedestal to the wall by coupling a 3D fluid turbulence code to a 2D transport code including calculations for neutral and atomic physics and a highly simplified plasma-material model. This boundary model would be validated against experimental measurements of scrape-off layer (SOL) plasmas profiles and flows and various threshold behaviors such as detachment or the density limit.
- Carry out developments required in preparation for kinetic edge modeling, including coupling of fluid and kinetic models for plasma and neutrals; computation of prompt ion drift-orbit losses; and enhanced models for plasma-wall dynamics that would include more self-consistent calculation of recycling, impurity sources, re-deposition and fuel retention.
- Further develop the boundary model by coupling the new physics capabilities developed into a 4D kinetic transport model for ions and electrons that is coupled to a 3D turbulence model for transport-time simulations. This code would also include kinetic neutrals, either directly coupled to the plasma model or parameterized, and a dynamic wall model that includes impurities.
- Take on the physics of edge localized mode (ELM) dynamics and control with a model for the dynamic evolution of pedestal profiles. This code would model ELM's and other pedestal relaxation mechanisms that limit or control edge profiles using extended MHD, 2-fluid and/or kinetic-fluid codes. Starting with existing core gyrokinetic and extended MHD models, the code would be enhanced with a more accurate free-boundary equilibrium solver required for the pedestal and SOL and with the inclusion of magnetic perturbations into the gyrokinetic simulations.
- The final step would be direct, multi-scale simulation of the edge. Such a code would require a solution to nonlinear electromagnetic gyrokinetics appropriate for the pedestal and SOL plasma with a realistic collision operator coupled to neoclassical transport, sources and self-consistent wall interactions. The physics of this model would go beyond the current micro/macro paradigm. Approaches to be explored might include gyrokinetic treatments without the high-n approximation, hybrid kinetic-fluid methods, or a full 6D Vlasov treatment.

The WDM/Disruption Avoidance ISA would target dynamical modeling across all regions of the plasma for all discharge phases from startup to shut down. The approach will be to build a modular code that allows use of alternate models with differing degrees of physics fidelity. Development will begin by adapting existing physics components for 2D MHD equilibrium, heating, particle, momentum, and current sources and for quasi-linear profile evolution to a FSP framework. This code will be extensively benchmarked against current codes with

similar capabilities. At the next stage, components for MHD linear stability will be added along with calculations of nonlinear evolution sufficient to begin studies of disruption prediction validated against experimental databases. The next major development will be creation of a 2.5D model based on 3D equilibrium and 1D transport components. This model will allow for interactions of finite-amplitude MHD instabilities and turbulence driven profile evolution. The next step will be the addition of new or adapted physics components, at successive levels of fidelity, for fast particle sources and evolution that takes into account fusion, beam, and RF sources, including RF coupling and self-consistent interaction with macroscopic instabilities. During this stage, algorithms capable of computing turbulent transport on transport times scales will be adapted for use in the WDM. Gyro-kinetic models capable of taking into account the self-consistent effects of energetic particles will be incorporated as they become available. Finally, as the pedestal-boundary ISA matures, models of core-edge coupling will be included in the WDM. Each WDM component will be verified and, as far as possible, validated as a stand-alone component and as a part of the whole device model throughout the development process.

Summary of Tasks and Milestones

Years 1-2

- Develop, test and release static linear model for pedestal structure based on linear MHD and gyrokinetic calculations.
- Develop and begin V&V for coupled fluid 3D turbulence to 2D transport code with models for neutral, atomic and simplified plasma-wall physics.
- Develop first 1.5D Whole Device Model based on existing physics components within FSP framework; benchmark against current codes. Establish WDM prototypes with parallel architectures. Develop production capabilities and plans for user support.
- Implement linear macroscopic stability calculations within WDM and validate against experimental database in support of improving prediction and avoidance of disruptions.

Years 3-5

- Develop coupled kinetic-transport/fluid-turbulence code for boundary plasma with improved wall/sheath models.
- Develop dynamic model for evolution of pedestal profiles.
- Develop 2.5D WDM by integrating 3D equilibrium solvers with 1D transport equations.
- Implement and validate nonlinear MHD calculation to determine whether discharge evolves to a disruption, integrating 3D equilibrium effects.
- Test algorithms for coupling turbulence into transport code (e.g., Profiles via Trinity/TGyro/FACETS approach, para-real algorithms for time integration). Incorporate core-edge coupling with linear model from Edge ISA.

Years 5-10

- Develop direct nonlinear simulation capabilities for pedestal boundary.
- Assess the role of 3D equilibrium on disruption onset and evolution including the role of impurities in nonlinear plasma evolution.
- Enhance WDM through adoption of nonlinear models from the Edge Physics ISA and higher fidelity calculations for fast-particle physics and turbulence.

Approach for adding additional ISAs

Additional Science Driver areas were identified while planning the FSP. As described in some detail in Section A.1, these included Disruption Mitigation, Core Profiles, and Energetic Particles/Wave Physics. Component teams would be formed to adapt and validate models in these areas, and basic research in would continue to be carried out by the five existing SciDAC projects in MHD, micro-turbulence, wave-particles and RF – all of which

were recently renewed for five more years. The choice for the next ISA would likely be made during the third year of the FSP and would depend on progress in these areas. Possible specific steps in carrying out associated ISA activities in each of the candidate topics include:

Disruption Mitigation

(i) begin with existing extended MHD codes and free-boundary disruption models associated with ISA2; (ii) add Fokker-Plank modeling of runaway electron generation and transport in stochastic, time-varying fields; (iii) include radiation and impurity dynamics with impurity sources and transport in a disrupting plasma; (iv) incorporate advances from ISA1 for plasma-material boundary interactions including sheath model, evolution of fuel, impurities and dust; (v) analysis of 3D mechanical interactions via inductive and conductive (halo) currents; (vi) improved modeling of gas jet and pellets for disruption mitigation; (vii) improved models for electron and ion (thermal and super-thermal) transport in stochastic field; and (viii) development of appropriate Kinetic-MHD hybrid models.

Core Profiles

(i) begin with a detailed comparison of current first-principles and reduced transport models through experimental validation; (ii) address discrepancies in models of electron transport; (iii) evaluate current approaches to global (full radius) and long-time simulations via local models, including improving parallelization and usability of 5D gyrokinetic codes; (iv) include appropriate gyrokinetic formulations for calculating momentum transport and radial electric field dynamics; (v) develop approach for efficiently extracting information from global simulations to allow prediction of profile evolution for temperature, density and momentum; (vi) begin investigation of mesoscale phenomena on transport time scales, including examining interactions between neoclassical tearing modes (NTM) and other MHD and micro-turbulence dynamics; (vii) characterize nature of boundary condition between the core and pedestal including fluctuations and flows

Energetic Particles/Wave Physics

(i) begin with 2-fluid and gyrokinetic models for finite Larmor radii effects on Alfvén eigenmodes; (ii) develop “edge to core” wave coupling and propagation; (iii) calculate nonlinear evolution and transport of fast ions in field of Alfvénic instabilities calculated via linear eigenmodes and extend to mode-saturation time scale; (iv) incorporate finite Larmor radii effects into ion cyclotron radio frequency energetic particle interactions; (v) address effects of edge instabilities on coupling and propagation of short wavelength modes; (vi) incorporate reduced models of energetic particle modes and transport into RF models; (vii) develop kinetic closure for MHD hierarchy to describe stabilization of NTMs and sawteeth; and (viii) more complete, self-consistent calculation of nonlinear evolution of fast ion distribution on slowing-down or transport time scales.

D.3 Plan for Physics Components

The integrated FSP component development plan takes advantage of substantial overlaps between the component requirements arising from different integrated science applications. To realize the economies of scale, this activity is planned and carried out across the entire range of FSP science drivers. For the first three years, the component team will focus on standing up the basic functionalities of ISAs for pedestal/boundary plasmas and whole device modeling/disruption prediction and avoidance. The associated projects will be organized into adaptation and development groups. The first group will adapt established codes from the base Theory/SciDAC program, transforming them into FSP components for the two ISAs. These adaptation projects will deliver adapted FSP component codes within the first two years of the FSP execution phase. The task of the second group consists of developmental projects with deliverables within the first five years of the FSP. Although some level of capability or pilot projects in the base Theory/SciDAC program, substantial gaps remain that must be addressed by sustained FSP investment in development over the first five years of the FSP.

The component adaptation projects for the first two years of FSP will target three classes of capabilities:

(1) Profile evolution. The target list of components in terms of their physics functionality include a free boundary Grad-Shafranov equilibrium solver (A-1) to compute the axisymmetric magnetic field configuration for a given

profile of the plasma as a function of the magnetic flux function and of the externally applied coil currents. The time update of the core plasma profiles (for density, temperature, and plasma current) will then be provided by 1D neoclassical and reduced turbulent transport models (A-2) or by a local turbulence-based transport model (A-3). The pedestal and scrape-off layer (SOL) require equilibrium models for temperature and density that are distinct from that of the core plasma. The pedestal model (A-4) will aim to set the pedestal width and height, while the SOL model will solve for the variation of the density and temperature both across the magnetic field lines and along the magnetic field. As a first step, an SOL transport model (A-5) based on the drift kinetic equation will be adapted for the pedestal/boundary ISA. The sources and sinks in tokamak plasma discharges, which are essential for modeling plasma profile evolution, will also be adapted. These include an RF heating and current drive model (A-6), a neutral beam injection model (A-7), gas puffing (A-8), and pellet injection (A-9) for axisymmetric plasmas. The boundary condition for the SOL transport will be provided by a 1D plasma/materials interaction (PMI) model (A-10) that supplies the particle recycling flux and sputtering yield of impurities. The neutral flux produced at the wall by recycling and sputtering will be followed into the plasma by a neutral particle transport model (A-11), which will also treat the ionization and, hence, will produce a volumetric source of impurity ions. This impurity source will be handed off to an impurity ion transport model (A-12).

(2) Instability Detection. For disruption onset prediction and avoidance, a suite of ideal and resistive magnetohydrodynamic stability codes (B-1) will be adapted into the FSP physics component library. These codes will evaluate the threshold for the onset of vertical displacement instability, ideal kink instability with no wall, a perfectly conducting wall, and a resistive wall, and low mode number (n,m) tearing instability. A family of pedestal stability codes (B-2) that focus on medium- n modes such as peeling-ballooning modes and kinetic ballooning modes will also be adapted to support the prediction of edge localized mode (ELM) onset and the pedestal equilibrium model. For completeness, the component team will also adapt a family of energetic particle stability codes (B-3) into the FSP component library. This capability will be augmented by year three with a quasi-linear model for energetic particle transport calculation.

(3) Nonlinear Evolution of Tokamak Plasmas after Onset of Macro- or Micro-instabilities. An initial value MHD code (C-1) will be adapted for core MHD and ELM dynamics, which have been identified in the WDM/disruption ISA and pedestal/boundary ISA, respectively. Both the pedestal MHD and the SOL turbulence can be evolved by an initial value two-fluid code with Braginskii closures (C-2). The coupling of C-2 to neutral production (A-10) and transport (A-11) introduce critical multi-physics to model the tokamak boundary plasma. There is also a plan to adapt an existing global gyrokinetic nonlinear solver (C-3) for the FSP component library. This kinetic model will provide a cross-check for A-2 and A-3 and information on core momentum transport and possibly energetic particle transport.

The initial set of component adaptation projects will build upon existing and relatively mature codes in the base Theory/SciDAC program. In combination, these projects will meet the initial needs for the FSP integrated physics modeling capabilities of the two ISAs. The performance target for these projects will be to release the first component library by the end of the second year of the FSP. A preliminary release for testing and review is scheduled for the end of the first year.

In addition to the adaptation projects with deliverables in the first two years, a number of new component development projects will be launched with target delivery dates before the end of first five years of the FSP. These new component development projects will cover the component needs for not only the two early ISAs but also those ISAs to be launched after a full FSP ramp-up. There will be two families of development projects, with one focusing on the plasma dynamics and the other focusing on plasma-materials interactions (PMI).

Plasma component projects will tackle the interaction between non-axisymmetric magnetic fields and plasma transport. The first task targets WDM with profile evolution with 3D equilibria (D-1); i.e. 2.5D transport modeling. This initial goal will be a low-dimensional transport model coupled to a general 3D equilibrium solver for the magnetic fields that can support islands and stochastic regions. The second project will address quasi-static evolution of coupled 3D fields and turbulence transport (D-2). It also aims at long-term profile evolution with 3D magnetic fields. The third project will be to develop an initial value gyrokinetic solver in 3D fields (D-3). It could be used for embedded calculations that update the plasma profiles for subsequent 3D equilibrium

calculations in D-2. The D-4 project will develop an initial value kinetic MHD solver to provide the effect of long mean-free-path transport on macro-stability in a tokamak plasma. Finally, the D-5 project will blur the line between conventional transport models and MHD macro-stability models, solving the initial value electromagnetic gyrokinetic problem in the regime of plasmas typical of the tokamak pedestal and SOL. This capability is specifically called out under the edge/pedestal ISA.

The second family of component development projects will address the PMI problem in a tokamak reactor. Since a comprehensive model is beyond the scope of the FSP, a more focused objective will be pursued in the first five years of the FSP: to construct a first-principles PMI code that resolves the plasma recycling physics at the wall (D-6). Such a code will require a kinetic plasma model for the boundary plasma (namely sheath/pre-sheath), a molecular dynamics-based material model for wall response to ion irradiation, and neutral transport and atomic physics models for neutral ionization and radiation.

D.4 Plan for Software Infrastructure and Services

In addition to providing support for software development throughout the project, Software Integration and Support (SIS) will need to develop the on-HPC composition software, the task composition (off-HPC) software, and the customizations of components needed by the integrated fusion modeling efforts. Furthermore, the SIS team must support the needs of data management, including data down-selection and archiving, developing the data catalog, and developing the software for these tasks. In this section, we summarize the SIS tasks to be undertaken in the Fusion Simulation Program.

Summary of tasks (from Science Drivers and other inputs)

The science driver analyses established a broad set of tasks for SIS. Immediate on-HPC integration tasks involve handling the amount and types of data needed for integration, the speed at which data must be exchanged, the development of algorithms (including implicit coupling), and the data layout and metadata. Longer-term efforts will include enhancing coupling by advancing the ability to smooth and refine data during the transfer. For Task Composition (off-HPC integration), data layout and metadata are also important. In addition, there are tasks associated with input file validation; software to manage sequential, interdependent execution of separate applications with appropriate information about data location; data reading plug-ins for general visualization tools; and the development of standard analyses and standard visualizations within the general visualization tools. The Task Composition effort must also provide the tools needed for data management, for which the requirements include having a searchable data server and a universal view to the data wherever it might exist. FSP will need to adopt rigorous software practices, such as revision control and unit and regression testing in order to provide developer support. Further, the FSP will need to adopt modern engineering tools, such as build and package management, that allow maintenance of the FSP software suite across the range of platforms that exist today as well those that will come on line in the future.

On-HPC software integration

On-HPC software integration must support two primary methodologies for coupling: 1) direct method invocations with use of MPI for communication and 2) file-based data sharing. For file-based data sharing, SIS will develop standards for metadata that can be applied to any of the implementations (HDF5, NetCDF, and ADIOS-BP). Such standards should be layered, with a minimal set being to specify the metadata needed for visualization and provenance. For FSP, a semantic level, where physical quantities such as electron temperature can be found in looked-up tables, will be needed. In addition, the on-HPC effort will need to define the basic APIs for component invocation. Ultimately, a reference implementation of the on-HPC coupling framework will be developed and in place for the development of new ISAs that are launched in the out years. In this context, addressing software complexity resulting from integrating many components from large teams of researchers could benefit from new approaches developed in the enterprise community.

Task composition (off-HPC integration or workflow)

Task composition must include: 1) the development of methodologies for describing and executing task composition; 2) the development of applications used in task composition; and 3) the development of plug-ins or enhancements to existing applications (like visualization tools) that are used in task composition. With regard to describing and executing task composition, the early work will concentrate on studying use cases since widely used task composition tools do not exist. Currently, most computational physicists write scripts in either bash or python to automate their work. Some applications, such as job preparation and submission, can be started on immediately, as the tasks are well defined, but others, such as tools for data analytics specific to FSP, need more definition, which will come from a study of the user processes.

Integrated data management

For integrated data management, the first step is the definition of protocols for storing and retrieving data in the system. The system will likely not store all of the data, as simulations tend to produce excessive amounts. Therefore, criteria and methods for down-selection will need to be developed prior to storage. An implementation, which may be distributed, will then need to be developed. A key element of the FSP plan is the deployment of a program-wide metadata catalog. Such an implementation would ideally allow anyone with a browser to search the FSP-generated data via keywords or other characteristics to help select the simulations to study, and the data could then be viewed both textually and graphically.

Software engineering practices

For a project of this complexity and scope, it will no longer be sufficient to follow the common historical practices for computational application development. The FSP will therefore mandate software engineering practices to ensure the quality of the software. Those practices will include the use of revision control systems (either svn or git); modern build systems that allow maintenance over the many platforms that FSP must support; and test systems, ideally with test-driven development, so that the capabilities of accumulated code are known at all times.

Developer support

Developer support will be needed to facilitate teams working efficiently. Collaboration tools (wikis, mailman lists, desktop sharing) will be set up. Furthermore, the ISA and ET development teams will need assistance with build and package management systems and test systems for their projects.

Release strategy

The goal for FSP software is to release the primary on-HPC computational applications yearly, but the release process requires further study. For existing, complex multi-physics projects, release processes are still evolving. Given the decoupling of the task composition software from the on-HPC composition software the two can be released separately. In addition, there will be a yearly release of the tools developed for job preparation, data analysis, visualization, etc. However, the development of data plug-ins for external visualization packages, e.g., VisIt, will be released on the schedules of those packages.

Research required

There are a number of areas in which research will be required to ensure a knowledge base for the out years of the program. In particular, applied math research is needed to understand items such as time-step limits for coupled components and the accuracy obtained with different coupling schemes. Computer science research is needed to understand how best to communicate or share the large volumes of data for 3D coupled simulations, including investigation of I/O staging and/or extended data structures. Computer science research is also needed regarding effective means to construct and maintain such complicated systems for data handling.

Summary of tasks and milestones

The detailed tasks and milestones are listed in the full Program Plan. Rather than repeat them here, we summarize the overall strategy. Infrastructure (e.g., for collaboration) must be set up immediately at the start of

the program. Next, one must settle on standards, for example, for metadata, down-selection, or APIs. Finally, one provides implementations. In a spiral model of software development, this occurs repeatedly through the life of the project to deliver ever greater capability.

D.5 Plan for Software Quality

As noted previously, in addition to traditional software quality assessment activities, the FSP considers verification, uncertainty quantification, and a robust experimental validation effort to be integral parts of software quality for fusion simulation. Both the physics and the code quality share equally high priority. Of course, a certain fraction of the initial effort in Software Quality will be to establish testing standards; to select and provide testing software; and to establish software release processes. The more substantive work will be in the instantiation of the verification, UQ, and validation efforts, which are summarized here.

D.5.1 Plan for Verification and Uncertainty Quantification

The complexity of the fusion simulation problem and the limited available resources necessitate V&UQ activities that are limited in scope and directed towards key questions. In support of experimental validation, the ultimate goal is to develop uncertainty estimates for computed quantities that accurately reflect the numerical errors and sensitivity to uncertainties from inputs and parameters. The minimal requirement for comparison with experiment is that systematic *code* verification is done on all components. Confidence can be increased with reliable *calculation* verification capabilities, as these can provide discretization error bounds on quantities of interest (QoIs). Ultimately, the ability to include the uncertainty of inputs and parameters in computed QoIs not only improves the conclusions of validation activities, but also provides the means to direct efforts to reduce uncertainties and improve models. Thus, V&UQ in support of validation need not be an all-or-nothing proposition, so long as the limitations of the conclusions are well understood and explicitly documented.

The V&UQ execution plan therefore initially emphasizes code, and to a lesser extent, calculation verification. An uncertainty analysis effort will commence with the start of the FSP, but its activities will be necessarily focused and exploratory in nature. Effective UQ methodologies for a limited set of specific questions will be developed. Over time, the UQ scope is expected to grow to encompass a broader array of applications and will more directly impact experimental validation efforts.

Plan for ISA verification activities

The activities for both ISA's and any additional component development are quite similar, so we do not distinguish. The verification activities will be led and primarily conducted within each ISA and component project, but a small number of numerical analysts from the V&UQ team will actively assist in study definition, analysis of results, review of documentation, and development and implementation of new *a posteriori* error estimation techniques.

Each component, whether part of an ISA activity or component development project, will have to undergo a formal verification process. The basic process, to be refined by the V&UQ team and adapted over time, will be (1) develop pre-requisite documentation, including the mathematical models, the discretizations, the parameters, *a priori* expected rates of convergence, test problems with and without solutions, physically important quantities of interest for the test problems, and prior verification and benchmarking activities; (2) conduct code verification on the identified problems using grid convergence or other *a posteriori* error estimation techniques; and (3) investigate calculation verification techniques, that is studies of code convergence and error estimation for problems with no known solutions. A similar process will be applied to integrated applications, although the results are expected to become ambiguous as the problem complexity increases. Further details are provided in Section 3.5.1 of the full plan.

Plan for ISA uncertainty quantification activities

Based on the Science Driver reports and the subsequent planning workshop, there appear to be no substantial existing UQ efforts in the fusion simulation community. Hence, the expectation is that the FSP UQ effort will start out small and highly focused, but grow in scope and program relevance over time. Uncertainty analysis will

not initially be attempted on large, integrated application codes. Instead, hierarchical sequences of increasing complexity will be investigated to develop the knowledge and the methodologies necessary to attempt more complete UQ analyses. The definition and execution of uncertainty analysis will require close collaboration of physicists and numerical analysts within each ISA with UQ experts from the V&UQ team.

The Boundary/Pedestal development plan (full Plan, Section 3.1.1) is already organized into a hierarchical sequence of models of increasing complexity suitable for a staged UQ analysis. Two potential areas for initial UQ investigation are (1) the determination of the uncertainty in ELM instability threshold and critical profile (height and width) with geometry and pedestal profiles and (2) the determination of uncertainty in profiles and fluctuations in the plasma boundary layer from input fluxes and geometry. Additional details can be found in Section 3.5.2.1.

In a manner similar to that for the Edge Physics ISA, a sequence of hierarchical models from the Whole Device Modeling/Disruption Avoidance ISA can be identified for a staged UQ analysis. Three potential problems to investigate initially are (1) the determination of uncertainty in the plasma with geometry and pedestal profiles as inputs using a hierarchy of core plasma transport models; (2) the determination of uncertainty in equilibrium reconstruction and instability thresholds with geometry and pedestal profiles; and (3) the determination of uncertainty in the onset of disruptions. The first study in particular would allow for comparison of uncertainties between reduced models and more complete models, and the third provides an opportunity to consider a UQ target of high value to ITER. Additional details can be found in the full Plan, Section 3.5.2.2.

Research opportunities

There are many open issues in the application of V&UQ techniques to large-scale simulation codes. These include robust *a posteriori* error estimators; propagation of error and uncertainties through coupled components; determination of errors due to coupling; error estimation in the presence of models for unresolved physics; error estimation and UQ for multi-scale problems; error estimation for solution-driven model changes; methods to reduce the work of sampling in high-dimensional parameter space; and data assimilation and the construction of parameter uncertainty distributions. Resources and therefore method research will be limited within the FSP, so the FSP will need to engage external research and development efforts (ASRC, ASC, NSF, etc.) for new tools and techniques.

Summary of Tasks and Milestones

In the first year, the code verification process will be begun on each code that is componentized or developed. One or two problems will be selected to be addressed hierarchically with UQ analysis; this will include the identification of desired QoIs, input parameters and data; the quantification of uncertainties in input parameters and data; and the exploration of parameter space. In the second year, components that make up the targeted Year 2 milestones will have completed and documented code verification results. Calculation verification investigations will be started, including work on advanced error estimation techniques. For UQ, exploration of parameter space will continue with the application of a variety of techniques that seek to reduce the dimensionality of parameter space. In years three through five, consolidated tools for verification will be made available and work on advanced error estimation techniques will continue. Full sensitivity and uncertainty analyses will be applied to the initial levels of the problem hierarchies defined above.

D.5.2 Plan for Experimental Validation

General guidance for planning validation

The FSP validation procedure will follow the Phenomena Identification and Ranking Table approach for planning. Five crucial steps will have to be completed before embarking on the experimental execution:

- Identify the hierarchical key physics quantities and processes relevant for the specified application;
- Determine whether the conceptual model(s) to be tested is(are) sufficient to describe the key physics;

- Conduct verification and uncertainty quantification (UQ) to ensure the accuracy and adequacy of the code solution(s);
- Lay out critical experiments and associated diagnostic capabilities (including synthetic diagnostics) that will provide the key data for use in testing the model; and
- Develop quantitative metrics for use in assessing the fidelity of the model based on results from the validation experiments.

Discussions of the critical physics, code readiness and V&UQ appear in Sections D.2, D.3 and D.4. Here we provide a template for laying out critical experiments that follows the hierarchical strategy (with timeline indicated):

1. Test processes that have predicted implications for one or more profiles (1-3 years).
2. Simultaneously, make measurements of phenomena that should appear when predicted limits are reached (e.g., rise in fluctuations with expected characteristics) (1-3 years).
3. Perform steps 1 and 2 over a wide range of plasma conditions chosen to stress the important parameters of the processes (1-3 years).
4. Longer term, make quantitative tests of the relevant phenomena (e.g., fluctuation amplitudes) (3-5 years).
5. For processes that survive steps 1-4, develop integrated models (transport models or frameworks that incorporate important processes) to test coupled physics (3-5 years or longer).

In order for this template to provide useful results, quantitative validation metrics will need to be developed for each model hierarchy of interest. These metrics are needed to both establish the fidelity of current models (and thus the confidence that should be assigned to their predictions) and to track improvements in model fidelity as they (and available computing resources) improve. The metrics should incorporate an assessment of the numerical error in the model results, as well as both model and experimental uncertainties, and reflect the inherent key sensitivities of the models being considered. In general, a suite of “simple” metrics which assess model fidelity for a single physical parameter will be needed, with these simple metrics combined into composite metrics to provide more holistic assessments of model performance.

Validation plan for Edge Physics ISA

Because of the uncertainties of the physics in the tokamak boundary, it is likely any modeling of this ISA in the next five years will be made up of a combination of first-principles and reduced (even empirical) models. The boundary/pedestal model should cover phenomena over a wide range of timescales from the steady-state (time-averaged) heat and particle fluxes to larger transient fluxes induced by off-normal events such as disruptions and ELM. A heavy burden will have to be put on experimental validation to quantify the fidelity of each component as well as the integrated model of this region. Fortunately a wide range of existing devices, with pulse length ranging from a few seconds to hundreds of seconds and operating with very different boundary conditions, are available for this purpose. Special attention should be paid to the following critical issues that can impact the heat and particle loads as well the edge transport barrier and the maximum plasma pressure at the top of the pedestal:

- L-mode, H-mode, L-H transition;
- Pedestal structure;
- ELM avoidance and mitigation;
- First wall and divertor PMI, loads on high heat flux poloidal field coils;
- Evolution of first wall and divertor poloidal field coils (material migration, mixed and re-deposited materials, etc.);
- RF antenna/SOL interactions;
- Impurity generation and transport; and

- Steady-state operations with self-consistent plasma and wall modeling.

The readiness of the code capability to model and the experimental capability to validate these critical physics is summarized in Section 3.5 of the full Plan. The critical physics can be grouped into validation tasks that range from being near-term ready to having major gaps requiring significant development before validation:

- Static models of pedestal/edge (validation underway);
- Micro-meso-scale dynamics (turbulence -> ELM's) and impact on plasma profiles across pedestal/edge/SOL (tools emerging);
- Slowly varying plasma profiles for density, momentum, heat and neutral transport (tools emerging);
- Resulting impacts on wall and divertor (major gap); and
- Back-reaction of wall/divertor on pedestal/edge/SOL (major gap).

Validation plan for WDM/Disruption Avoidance ISA

The goal for this ISA is to validate the growing capabilities for WDM, beginning with existing framework approaches and including components for profile evolution, stability assessment and nonlinear evolution (disruption prediction) including active control. As an abstraction for the plasma control system, WDM will integrate all of the necessary physics to simulate the plasma response to external influences. Magnetic field coils, heating and current drive sources, and plasma transport properties determine equilibrium shape and profiles. Pedestal/ELM's, fueling, and impurities strongly influence fusion performance. Heating, current drive, fueling, and 3D field actuators strongly influence plasma MHD stability and thus disruption avoidance. Experimental validation will have to be planned to test the fidelity of each physics element, as well as binary and multiply coupled physics.

A possible flow-chart for WDM-based stability forecasting is shown below. Each box in this figure represents an extensive validation campaign. Within each campaign, the critical physics can be broken down into finer granularity. For example, validating the topic of near-by equilibrium states can start with axisymmetric equilibrium and its sensitivities to profile measurements as well as accuracy in the equilibrium reconstruction. The next hierarchy up in validation will include error fields, toroidal field ripples, resonant magnetic perturbation coils and magnetic islands. Further considerations will include the impacts of 3D fields on transport and equilibrium profile modifications. Other campaigns can be similarly constructed.

The readiness of the code capability to model and the experimental capability to validate the critical physics in

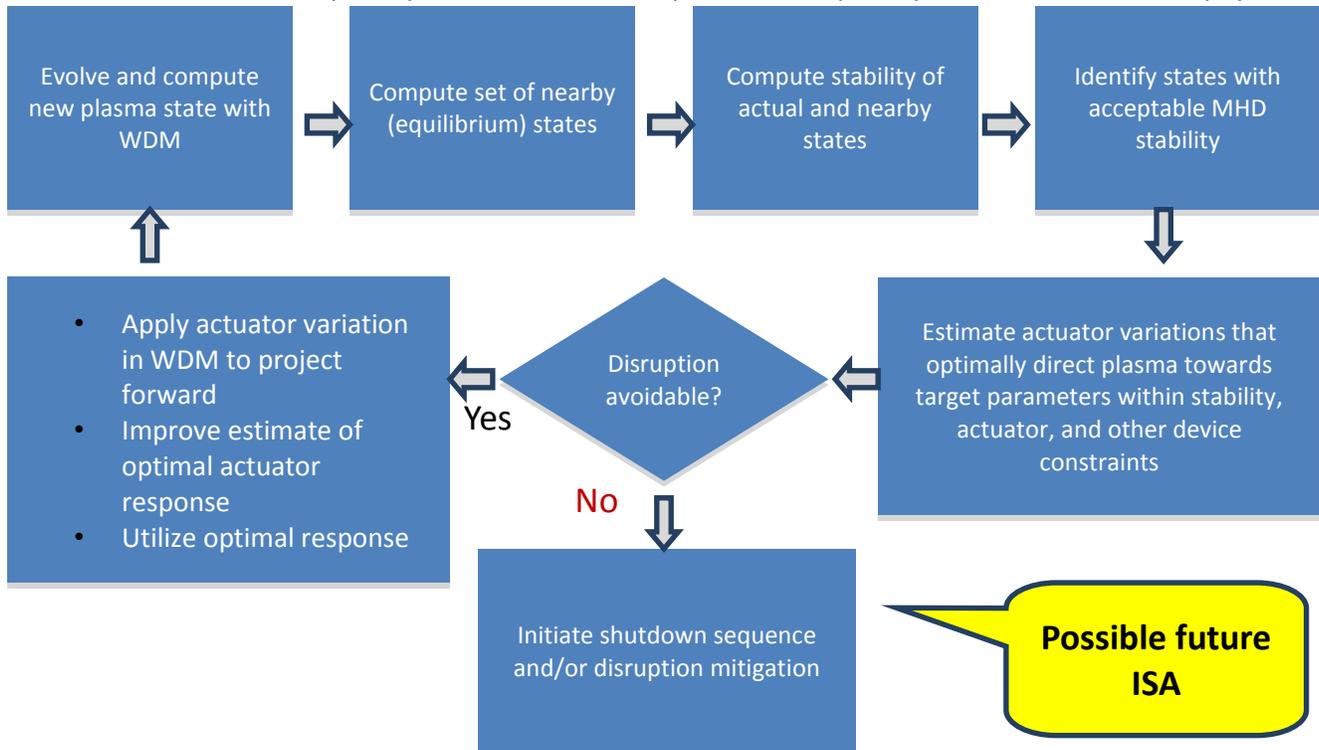


Figure 4: Flow-chart for WDM-based stability forecasting. (from Figure 4.1.4-2 of the full FSP Plan)

these boxes is summarized in Section 3.5 of the full Plan.

D.6 Plan for Production Support

This section summarizes the plan for FSP production support assuming that the FSP will ultimately have hundreds of users and production numbering 10,000s of runs per year.

Description of expected user base

FSP production uses include all scientific application of FSP software outside of direct code development. Therefore, elements of the user base come from both inside and outside the FSP program. From inside, there are the FSP supported analysts who develop applications and validation studies in collaboration with the base program, as well as FSP software quality team members performing such activities as UQ. From outside of the FSP come researchers in the experiment and computation/theory base programs with scientific applications for WDM and/or Pedestal Boundary modeling. Such applications can involve planning studies for ITER and current tokamaks, as well as studies for detailed validation against experimental data of computation/theory models that have matured to the extent that they are available as components in FSP production software.

Summary of requirements

The Production Support Team will supply the practical tools to enable knowledgeable researchers to take advantage of FSP software modeling capabilities. In the main, the physicist user should be able to operate with physics knowledge and a clear set of instructions for use, without having to become involved with code development issues such as build systems, regression testing, and the like. The production Support Team will support users in the ability to prepare input data and submit runs; to monitor progress of runs with possibility for run interruption at any time; to retrieve, examine, and analyze output of runs; and to restore and repeat an existing run “as is” or with modifications. The production team will assure that metadata is stored with run results sufficient to identify code versions, and copies of all input data will be stored so as to allow rerun of simulations at any point in the future.

Release Support

The Production Support team will work with Software Infrastructure and code development teams to design a means of detailed identification of versions of the FSP framework and component software with automated means for extraction of identifying data for inclusion with archived run output. Production releases will not be kept operational indefinitely, but each production platform will support access to a “beta” release (with new features) and at least one well-tested stable release of the FSP software.

Plans for User Support and Documentation

The Production Support team – in collaboration with developers – provides the detailed instructions for actual execution of simulations and access to the output data for interpretation of results. This includes clear, concise documentation of procedures for use of the FSP software, covering all phases: input preparation, job submission, monitoring, output visualization and analysis. When problems arise, the Production Support team will provide a first response, seeking to distinguish, e.g., between user input errors, system errors, and actual FSP code bugs. Where necessary, Production Support will involve FSP code development team members in such trouble-shooting efforts. In general, due to labor/cost considerations, such support will only be available on specifically designated production platforms. The basic responsibility of the operations teams is to provide user-level documentation for each step needed for production use of FSP software.

Plans for Education, Outreach, and Training

Since the FSP is expected to extend beyond ten years, education is an imperative associated with the need for the infusion of young talent. This requires connections to university programs that produce the best young scientists in theoretical, experimental, and computational plasma physics – as well as graduate students in AM/CS with interests in the plasma science applications domain. Establishing FSP post-doctoral positions will be especially important as a foundational component in developing “analysts” whose multi-disciplinary skills are key to a vibrant and productive FSP. Outreach will be carried out by the FSP at venues such as major conferences and in visits to prominent institutions. Individual scientists wishing to learn more and to acquire skills as users will be encouraged. The training of the requisite talent base for FSP activities will span a multi-disciplinary set of topics. For example, the FSP production computing support teams can be expected to help provide some of the key training assistance for broader FSP project efforts. Within this context, operations teams will encourage and assist the formation and training of user groups with experience often demonstrating that the most efficient education and training normally happens when users help each other. Knowledge of FSP capabilities and plans will be part of the education of a new generation of Fusion Energy scientists. Individual scientists wishing to learn more and to acquire skills as users will be encouraged. The Production Support team will also make the necessary user documentation available and will identify existing users with similar interests with whom the new user might productively collaborate.

Summary of Tasks and Milestones

In the first year of the FSP, the production support team will establish contact with user groups and help set up an ongoing means of communication in order to facilitate user input into FSP program priorities. A production system will be established using a mixture of legacy and FSP tools on a midrange cluster system – or on a set of closely match systems – at one or more major laboratories. An initial set of documentation and user tools will be provided, as well as a procedure for trouble-shooting of runs. In the second year of the FSP, an HPC production queue will be established at NERSC. FSP components will be integrated into production-ready legacy frameworks until such time as FSP native frameworks are ready for production. By the third year, FSP-provided production facilities should be in wide use, with strong governance of research user groups setting priorities for further development of production capabilities.

E. SUMMARY OF RESPONSIVENESS TO RFP

In the current full FSP Plan, the following topics from the original Request for Proposals (RFP) have been addressed:

FSP Deliverables – With the co-leads (Kritz and Keyes) of the 2007 FSP workshop report as part of the current FSP Planning Team, the list of prioritized deliverables outlined in that document have been critically evaluated and modified as appropriate in articulating the science opportunities and goals. The associated roadmaps were guided by both the near-term and longer-term priorities of FES stakeholders with respect to national as well as international needs, including ITER. The planning study has included a systematic assessment of the resources (in terms of Full Time Equivalents [FTE]) and the mix of expertise (plasma physics, material science, applied math, and computer science) necessary to successfully accomplish the ISA goals. The study led to detailed descriptions of the approach that will be followed for determining the required resources and reassessing the list of deliverables for the FSP, as well as for developing clear and compelling Work Breakdown Structures. More specifically, the planning effort has involved:

- Comprehensive assessment of the present computational capabilities of the fusion community in terms of major simulation codes, numerical algorithms, computational science tools (data management, visualization, code performance tools, etc.), computational frameworks, interface standards, code scalability, and other related issues. Detailed information of this kind can be found in the full length versions of Science Drivers reports in the Appendix B of the Full FSP Plan. Identification of major gaps and weaknesses and suggestions for the path forward are addressed – with respect to scientific opportunities contained in the Science Drivers discussions within Section 2 of the full Plan – as well as in the targeted goals of the ISA’s.
- Integration and coordination of the FSP with the projects in the FES SciDAC portfolio, including the process for incorporating results from the FES SciDAC Centers into the FSP, have been addressed by a detailed assessment of the SciDAC proto-FSP projects and by the articulation of targeted relationships/collaborations in the FSP Program Execution (Section 4) of the Plan. This has also encompassed:
 - Integration and coordination of the FSP with other SciDAC (non-FES) Centers, in particular with SciDAC Institutes and Centers for Enabling Technologies, as well as with efforts supported by the ASCR Applied Mathematics program, as described in Section 4.8 of the Plan;
 - Integration and coordination with the FES analytic theory and modeling program, including the process for incorporating improved theoretical models into the FSP simulation codes and engaging the help of the FES theory community to address gaps in the physics models implemented in the FSP codes. Associated examples are provided in the Full FSP Plan in Section 4.3 as well as in the Science Drivers reports in Appendix B.
 - Integration and coordination with the materials community for the purpose of addressing the plasma-materials interaction challenges – especially with respect to the ISA on Edge Physics as found in Sections 2.2 and within the Science Drivers reports in the Full FSP Plan, Appendix B.
- Details of the FSP vision and approach for developing a successful and credible Verification and Validation plan, including interaction and coordination with the FES experimental and diagnostic communities is addressed in detail in both Sections 3.5 of the full Plan.
- Interaction and coordination with international integrated modeling efforts, in particular those undertaken by our ITER partners in support of the needs of the international ITER Organization, are specified in various parts of Sections 3 of the Plan. This has been informed during the planning process by productive interactions in workshops and meetings involving; e.g., the European integrated modeling activities.
- High Performance Computing Resource Requirements. As a major computational activity, the success of the FSP will critically depend on the availability of HPC resources. In Section 4.7 of the full Plan, the current approach is described for determining the required HPC resources for carrying out the various

FSP tasks, including the appropriate mix of capacity and capability resources. Resources to be considered include the current and projected capabilities at the DOE-SC leadership computing facilities, as well as other resources (national or local) that can be reasonably expected to be available to the FSP researchers.

Referenced reports used in the FSP planning activity include:

- (1) "The ITER Integrated Modelling Programme," W. Houlberg, (private communication) May 3, 2011.
- (2) J. Dahlburg, et al., FESAC Report, [http://www.isofs.info/FSP_Final_Report.pdf (2002)]
- (3) D. Post, et al., FES Report leading to "Proto-FSP Projects [Journal of Fusion Energy 23, 1 (2004)].
- (4) A. Kritz, D. Keyes, et al., FSP Workshop Report (2007),
http://science.energy.gov/~media/fes/pdf/workshop-reports/Fsp_workshop_report_may_2007.pdf
- (5) W. Tang, et al., FESAC FSP Report (2007)
http://science.energy.gov/~media/fes/fesac/pdf/2007/Fesac_fsp_report.pdf
- (6) F. R. Bailey, et al., ASCAC FSP Report (2008)
http://science.energy.gov/~media/ascr/ascac/pdf/reports/Ascac_fsp_report_final.pdf
- (7) R. Hazeltine, D. Hill, et al., Report of the Research Needs Workshop (ReNeW) (2009)
<http://burningplasma.org/web/ReNeW/ReNeW.report.web2.pdf>
- (8) FSP Program Advisory Committee Final Report (May, 2011)
<http://www.pppl.gov/fsp/documents/FSPAC%20reportMAY2011.pdf>
- (9) "Scientific Grand Challenges: Fusion Energy Sciences and the Role of Computing at the ExtremeScale," PNNL-19404, 212pp(2010).
<http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf>