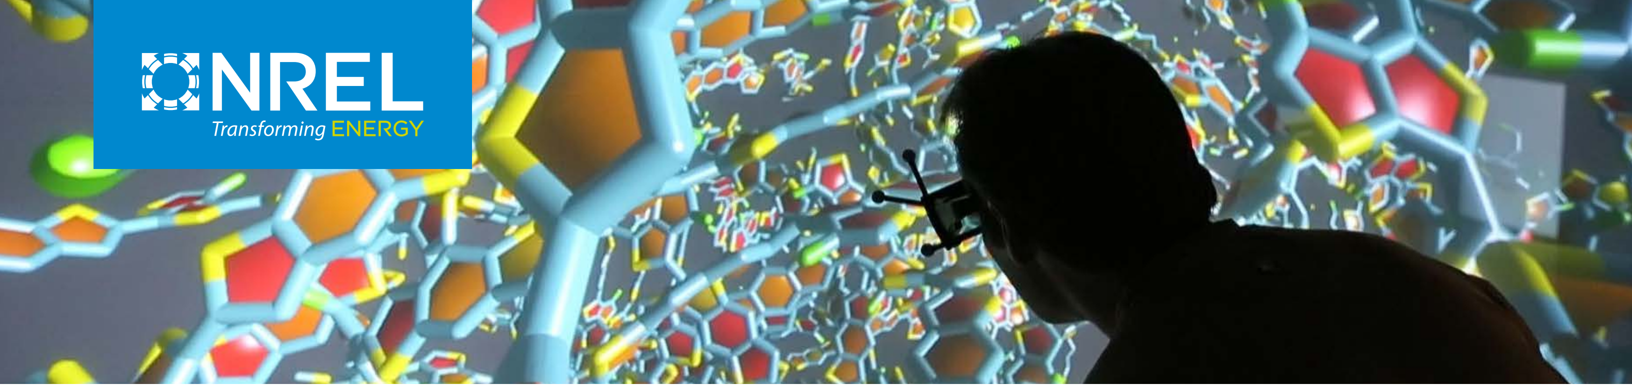

Transforming ENERGY through Computational Excellence

Computing for Clean Energy

Achieving a carbon-free power sector by 2035 as a step toward a decarbonized U.S. energy economy in 2050[1] will require major changes in power generation along with advances in autonomous energy systems, transportation, and buildings/communities.

Development of integrated modeling approaches for complex energy systems will be essential for deployment. Success requires developments in optimization and control theory, complemented by machine learning (ML) and artificial intelligence (AI), all of which in turn need targeted investments in breadth and scale of computing.

This document defines an opportunity space where:

- Embracing computing can link established research and development (R&D) to a decarbonization agenda
- Pursuing emerging approaches can accelerate the pace of technology advancement across the portfolio
- Leading by example could reduce the carbon footprint of computing worldwide.

Enabling Decarbonization through R&D

The National Renewable Energy Laboratory's (NREL's) advanced computing influence spans several common themes across the U.S. Department of Energy's Energy Efficiency and Renewable Energy (EERE) offices, including materials discovery, process modeling, fluid dynamics, resource mapping, and analysis of large-scale systems with real-time optimization.

NREL's strategic vision for the next decade is largely codified in terms of our "critical objectives": Integrated Energy Pathways, Electrons to Molecules, and Circular Economy for Energy Materials. To tie the aggressive agenda represented in part by these objectives to deep decarbonization, three major challenges must be confronted.

First, **increasing the scale of existing practice**, both by moving analysis from cities to national scale and by exploring complete design spaces (e.g., the search for new materials).

Second, **increasing the depth and comprehensiveness of current practice**, from improving the range of microscale phenomena included in battery models, to addressing the interaction between climate, energy, and mobility.

Third, **using new data sources to increase the pace of innovation and discovery**, from the ubiquitous sensors in smart buildings to live data streams from advanced experimental equipment.

Addressing these challenges while adapting to uncertainty is too complex, too large, and the stakes are too high to proceed through trial and error alone. Computing and advanced visualization (see photo on page 3) can close the gap between data and knowledge.

New Applications, Emerging Approaches

Rapidly advancing new solutions to achieve decarbonization requires computational capability to enable new activities:

- Scenario evaluation—scalable temporally, spatially, and computationally—to inform decisions within an uncertain and evolving state of the world
- Fusing simulations and sensor data sources in a complementary use of physics and data-driven approaches; thereby, closely coupling experimental observations and theory.

NREL has seen great success conducting jurisdictional planning studies—the City of Los Angeles (LA) and Dallas/Fort Worth (DFW) regional airport are notable examples—that wrestle with the challenges of melding the plethora of clean energy technologies now available with aspirations to reduce their carbon footprints.

---

[1] The Biden Plan for a Clean Energy Revolution and Environmental Justice (2020)
https://joebiden.com/climate-plan/

LA- and DFW-type analyses aimed at clean electricity for most of the population would require scaling to state, regional, and national levels. Getting to a net zero economy will additionally require more robust analysis with rigorous treatment of cyber, climate resilience, environmental justice, and economy-wide modeling. This increases the data and computational intensity considerably.

Scale Scenarios

Because the timeframe for deployment planning is long (i.e., decades) relative to the timescale of potentially disruptive technology and behavior changes (i.e., years), the planning process would benefit from evolution beyond deterministic, scenario-based optimization to being able to incorporate a robust understanding of uncertainty. Bringing the mathematics of stochastic optimization enables making decisions based on available information and refining as new information arrives. Not only could we identify optimal combinations of clean energy sources, but we could shape the direction of R&D activity critical to achieving a net zero energy economy by midcentury.

Couple Experimental Observations and Theory

A new way of working—where sensors, simulation, and experiments are used together—presents an opportunity to transform the pace of research progress. Two key paradigms exemplify this space:

- When the system model is composed of both physical and computational models, a physical representation can be used as a component of a larger system alongside a virtual representation of components as part of a physical apparatus. For example, a digital twin of a regional or national electric grid could be coupled to a microgrid representing a particular distribution feeder.
- When simulation drives experimentation, experimental conditions can be chosen to quickly estimate model parameters, fundamental constants, or optimize control policies. Reinforcement learning and Bayesian inference are two extremely powerful tools for building data-driven models that enable operational decisions in real time; they require large data sets and the ability to evaluate the probability of success for a given decision. When simulations are paired with instruments in near real time, theory, experiment, and computation can be combined with learning to accelerate research progress.

Carbon-Free Computing

Data centers globally account for 200 terawatt hours annually—less than 1% of global electricity demand[2] worldwide. Limiting the growth of the sector's energy usage in the context of growing demand requires continued innovation. By extending the living lab approach piloted with the Energy Systems Integration Facility (ESIF) data center, the industry can move beyond power usage effectiveness and focus on the production side of the equation. Additionally, the built environment for data centers is poised to take advantage of renewable energy at scale. A high-profile demonstration—in partnership with the private sector—could catalyze the uptake of carbon-neutral data centers for computing across the industry, from cloud services to the backbone of the internet.

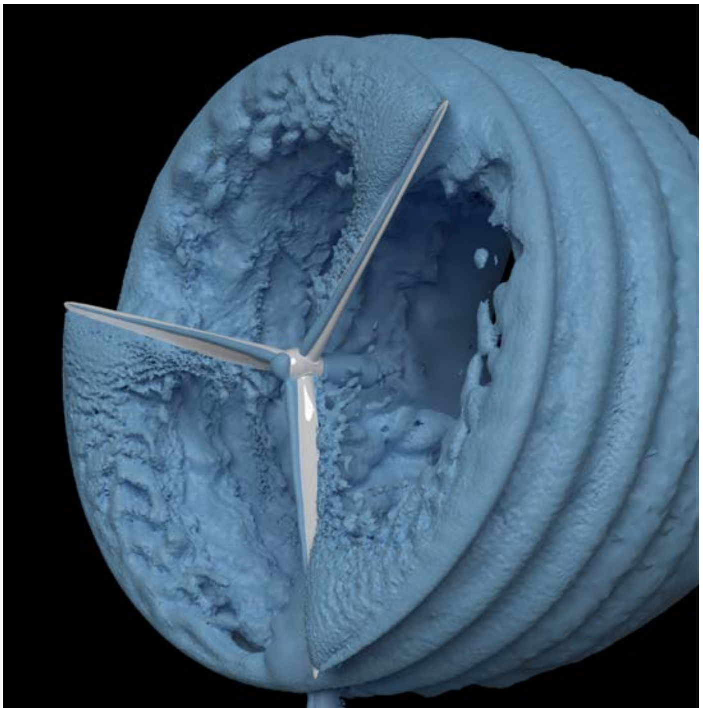

A 3D rendering of a theoretical wind turbine reveals some small-scale physics—previously beyond our ability to resolve— that will help advance our fundamental understanding of the flow dynamics that govern whole wind plant performance. *Image by Nicholas Brunhart-Lupo / NREL*

Power needs for on-premise computing approaching 20 megawatts (MW) on the horizon will soon exceed the capacity of the ESIF data center, providing a timely opportunity for a high-impact demonstration. A modular data center designed for advanced computing—integrating cloud, data, and simulation engines—could provide a rich research environment to demonstrate the possibility and practicality of a net-zero data center at scale.

Computing Needs

There are two specific needs: expansion of readily available capabilities and embracing emerging technologies.

Significant Increase in Capability

The current role of modeling and simulation (see page 4, "Current Computing Capabilities and Practice") is exploding, with new generations of machines over-subscribed at deployment. Very high-fidelity simulations are voracious consumers of computing, such as the ExaWind project's blade-resolved wind turbine simulations (see image on page 2). EERE and NREL expect a 6-fold increase in computing needs over the lifetime of the successor to current machines from

---

[2] IEA (2020), Data Centres and Data Transmission Networks, IEA, Paris https://www.iea.org/reports/data-centres-and-data-transmission-networks

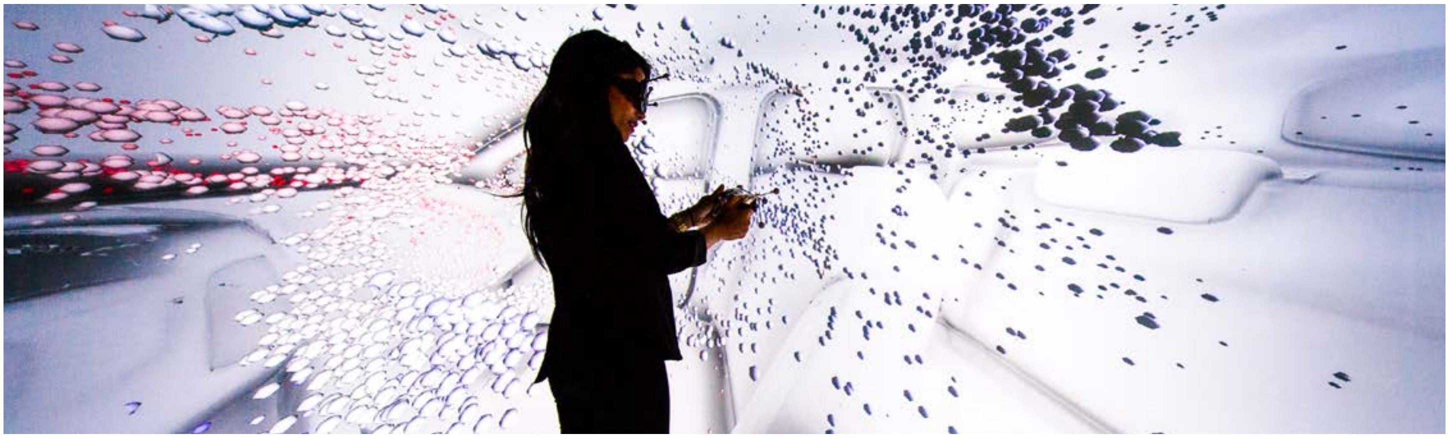

Members of the Department of Homeland Security sustainability team interact with an airflow visualization of dynamic interactive probing of vector fields, in the 3D Visualization Lab in the ESIF at NREL. *Dennis Schroeder / NREL 34548*

these needs alone. Additionally, new deployment-centered analysis activities—such as scaling to national levels—would require dedicating an equivalent of the current EERE flagship facility to the problem for an entire year. The addition of parametric uncertainty and continual refinement of the models as technologies, behavior, and deployment evolve could easily saturate such a machine between now and 2035. Of the big questions critical to the decarbonization agenda, it would likely be possible to address any one with a moderate increase in computing capability, but to solve them all at the same time requires a significant increase.

Emerging Computing Technologies

New workflows for interactive and real-time access to computing and streaming data hardware can enable ubiquitous tight coupling between simulations, experiments, other simulations, and real-world measurements. The resulting models utilize information available through data-driven methods along with traditional modeling and simulation. From a hardware view, such data-centric computing is anchored by sensors, data acquisition, and new technology platforms in the data center, in the cloud, and at the edge. Together, these buffer high-bandwidth data streams and extract information that can be exchanged with simulations at appropriate rates to fuse experimentation and theory.

Innovation in computing—from both the public and private sectors—continues at a lightning pace. Its key to success here is not just the hardware but the ability to leverage the rapid advances in AI (including ML and deep learning), data analytics, data streaming, and cloud computing that come along with intensive research activity in the underlying computational methods.

Investment Options

1. Increase baseline computing to 400 million allocation units/year (7x Eagle, see page 4).
2. Upsize ESIF data center power and cooling to 8 MW.
3. Advance methodologies for analytics workflows, R&M for increased system size.
4. Develop real-time computing for close sync with Integrated Energy Systems at Scale (IESS) and enable Advanced Research on Integrated Energy Systems (ARIES).
5. Construct combined Computational Science Center and Information Technology Services data center at Flat Irons to meet campus needs through 2035.
6. Increase baseline computing to 12x Eagle in 2024.
7. Build smart lab instrumentation and associated modeling, simulation (smart instrument capex, edge computing capex).
8. Develop instrument-to-edge computing for close tie between characterization and theory.

Investment Outcomes: Options 1-4 support approximately 10 LA, DFW, and regional mobility projects; adding options 5-7 support approximately 50 LA, DFW, and regional mobility projects as well as ARIES, catalyze change in research methodology towards 'lab of the future,' and anchor facility for computing for clean energy on a 15-year time horizon.

Investments Supporting Transformation

The analytic and computational capabilities discussed here would anchor an implementation blueprint for moving to a clean-energy grid, and eventually, pathways to a decarbonized economy.

Several specific investments could be made now to support the transformation over the coming decades (see "Investment Options" box). These range from natural follow-on from current practice to an aggressive push toward new concepts. Also included is major infrastructure estimated to be necessary to sustain computing additions in the clean-energy space through 2035.

Grappling with the climate/energy nexus is a grand challenge of our time, where computing plays a critical role. Getting enough computing—and the right type in the hands of our researchers so that they can focus their ingenuity and creativity on addressing the challenge—is of paramount importance.

Current Computing Capabilities & Practice

NREL's high-performance computing (HPC) data center, the Eagle supercomputing system, and visualization capabilities in the Insight Center, all located in the **ESIF**, propel technology innovation by providing NREL and industry partners the ability to tackle energy challenges that cannot be addressed through traditional experimentation alone.

The **Eagle** supercomputing system provides 8 petaflops of computing power and consists of two 114 compute nodes and 296 terabytes of total memory. A high-speed network and 14 petabytes of data storage enable the computing workflows that help advance renewable energy and energy-efficient technologies.

The **Insight Center** combines state-of-the-art visualization and collaboration tools to promote knowledge discovery for energy systems. A large-format, rear-projected display wall hosts the analysis of large-scale simulations, ensembles of simulations, and highly detailed visual analytics. A 3D, stereoscopic, and immersive visualization environment that illuminates both wall and floor surfaces provides an interactive experience with highly complex, large-scale data in many dimensions, from the atomistic scale of new materials to the atmospheric boundary scales used in multiturbine array simulations.

Computational materials and continuum mechanics are the mainstay of traditional simulation done with the aid of HPC, while forecasting and hybrid data-driven models are combined with digital twins and physical models.

Computational materials and chemistry form a workhorse capability for materials discovery that is expected to account for 10%–15% of traditional computing needs soon. Usage models range from ensembles of many small calculations to huge petascale computations using density functional theory, molecular dynamics, computational fluid dynamics (CFD), ML, and composite methods as integral parts of the design/test/build cycle for materials discovery and design. Emerging focus on plastics and polymer upcycling and recycling, biosourced material and fuel design, separations technologies, and new polymers with end-of-life and circularity designed into their molecular and morphological structure will continue to drive needs in this area.

Continuum mechanics, typified by CFD and structural mechanics, are the mainstays of modeling and simulation to understand and control energy carriers such as wind, water, and fuels. CFD is estimated to account for up to one-third of future computing needs. Ground truth calculations to understand driving processes require access to very large-scale computing leveraging the state of the art in scalable CFD, whereas, system optimization (e.g., siting, layout of offshore wind power plants, optimization of reactor or engine geometry) uses ensembles of smaller calculations that can be performed on a midlevel computing resource.

Evaluating future scenarios (supply, demand, congestion, system operations and optimization) through **forecasting** adds data-centric aspects to traditional modeling and simulation, from acquisition of sensor data streams (e.g., video, image or advanced metering data), ensembles of simulations (e.g., weather or digital twins), and advanced analytics (e.g., disaggregation, state estimation, uncertainty quantification, and stochastic optimization). These techniques, important for both planning and operation, are tightly coupled to markets and behavioral factors and are estimated to drive upward of one-fourth of future needs.

New approaches using cutting-edge AI techniques to interpolate regional data to the fine scale needed to evaluate resource potential is facilitated by farms of GPU-accelerated systems to train ML models.

Digital twins are virtual replications of the real world based on observations that allow researchers and policymakers to conduct what-if'explorations. This offers unprecedented opportunities to gain insights and inform science, infrastructure, planning, and operational decisions. Enabled by HPC, data-centric workflows integrate experimental data from actual hardware devices as hardware-in-the-loop with computational models, resulting in hybrid emulation/digital models. These facilitate extrapolating experimental data to real-world conditions, validating at-scale evaluation of disruptions and other emergent behaviors—phenomena that may only be seen at the full size of the system—in complex hybrid energy systems and provide virtual test benches for new control algorithms. This capability—effectively combining multiple techniques so that models become increasingly accurate and insightful—provides new opportunities to accelerate innovation.

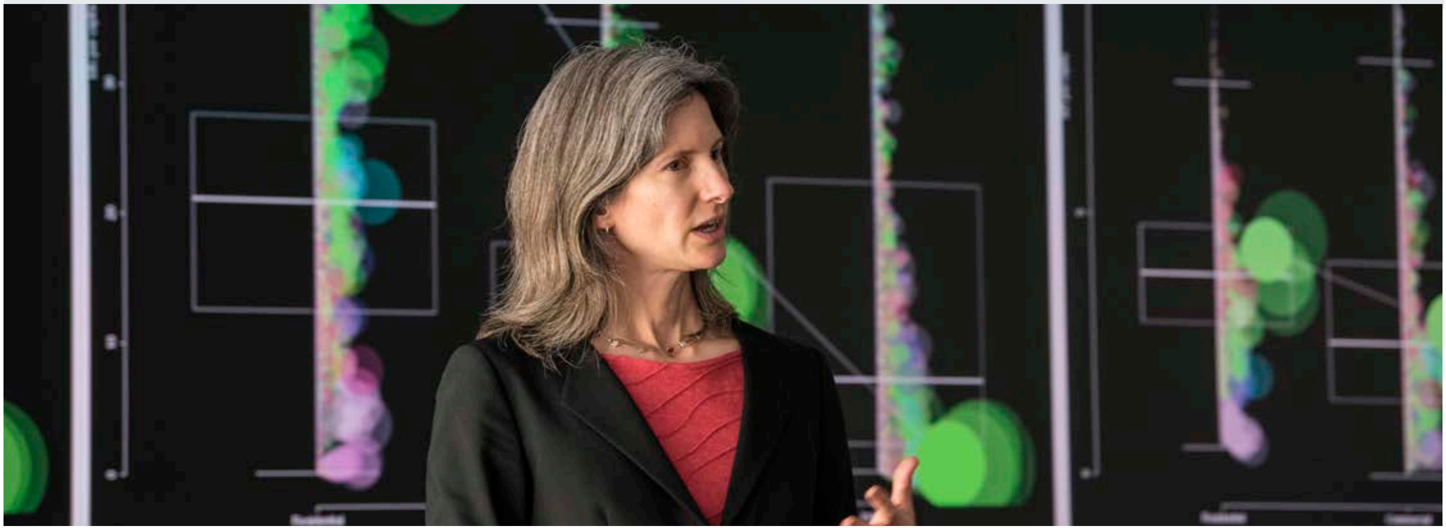

NREL researcher Megan Day presents 3D and 2D data from the Cities-LEAP research during a tour for the energy company representatives who were attending a Grid Modernization Workshop in Denver. *Dennis Schroeder / NREL 54137*

Cover Image: An NREL researcher explores photovoltaic molecular dynamics at work in the 3D Visualization Lab in the ESIF at NREL.
Photo by Malone Media

NREL
Transforming ENERGY

National Renewable Energy Laboratory
15013 Denver West Parkway, Golden, CO 80401
303-275-3000 • www.nrel.gov

NREL prints on paper that contains recycled content.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

NREL/FS-2C00-79695 • July 2021