Pacific Northwest NATIONAL LABORATORY

Data-Model Convergence (DMC) Initiative

Multicore World 2019
Wellington, New Zealand
February 11-14, 2019

James A. Ang, Initiative Lead

Andrew Lumsdaine, Draguna Vrabie, Robert Rallo, Mark Raugas, and Kevin Barker
Initiative Leadership Team

U.S. DEPARTMENT OF ENERGY BATTELLE

PNNL is operated by Battelle for the U.S. Department of Energy

PNNL-SA-140780

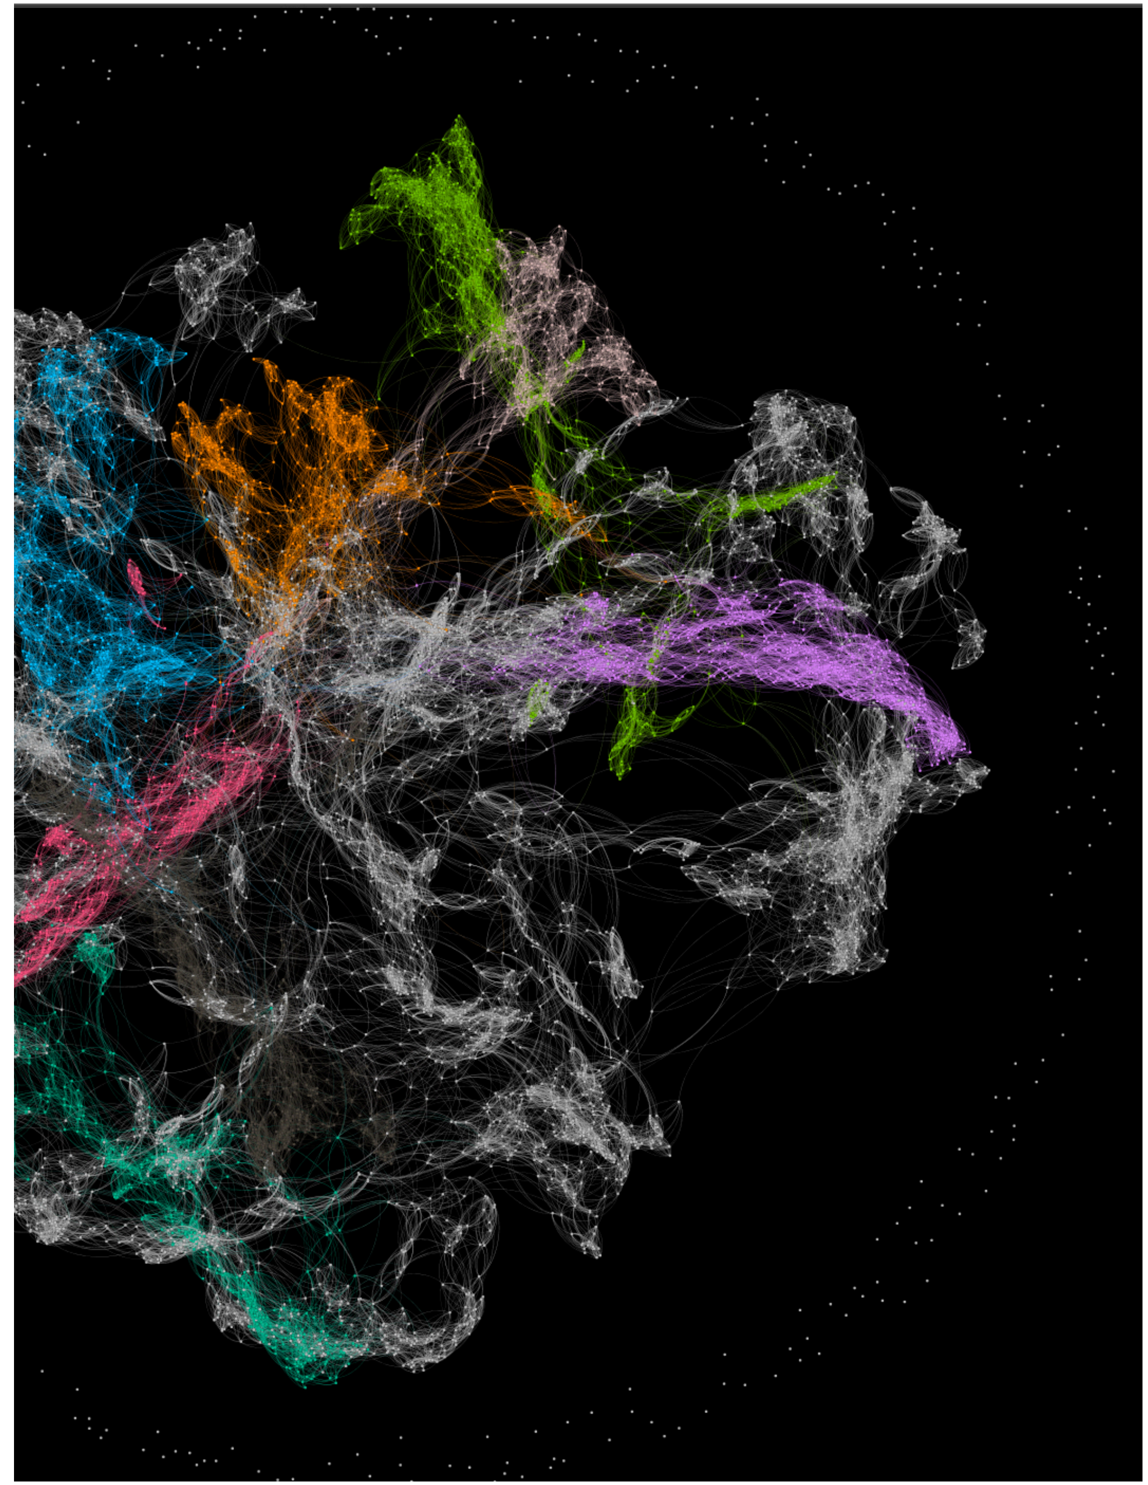

Outline

- Leadership Team
- Definition of DMC
- Approach
- DMC Scope
- PNNL Strengths
- Risks and Mitigations
- Exit Strategy
- PNNL Initiative Goals

DMC Initiative Leadership Team

Role	Name	Organization
Initiative Lead	Jim Ang	PCSD
Chief Scientist	Andrew Lumsdaine	PCSD
Domain Applications Thrust Lead	Draguna Vrabie	EED
Domain-Aware ML Thrust Lead	Robert Rallo	PCSD
Software Stack Thrust Lead	Mark Raugas	NSD
Hardware Architecture Thrust Lead	Kevin Barker	PCSD
Communications	Karyn Hede	CITD
Project Coordination	Becky Wattenburger	PCSD

- Physical and Computational Sciences Directorate
- Energy and Environment Directorate
- National Security Directorate
- Communications and Information Technology Directorate

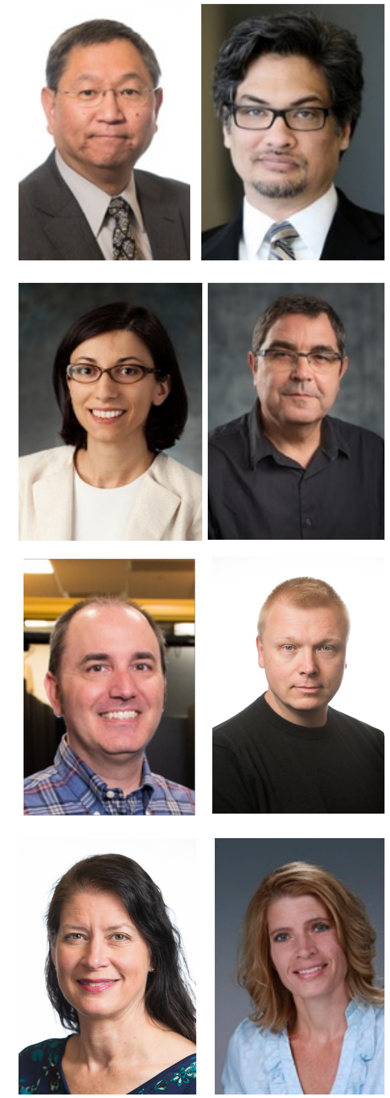

Definition of DMC

- Many contemporary science and engineering problems facing PNNL and DOE—such as grid optimization or materials discovery—are best solved by integration of:
  - High performance computing
  - Large-scale data analytics
  - Machine learning methods
- We call this integration "Data-Model Convergence (DMC)"
- Supports directly PNNL Lab Objective:

Accelerating Scientific Discovery through
Extreme-Scale Data Analytics and Simulation

Current Approach for DA, HPC, ML Computing Paradigms

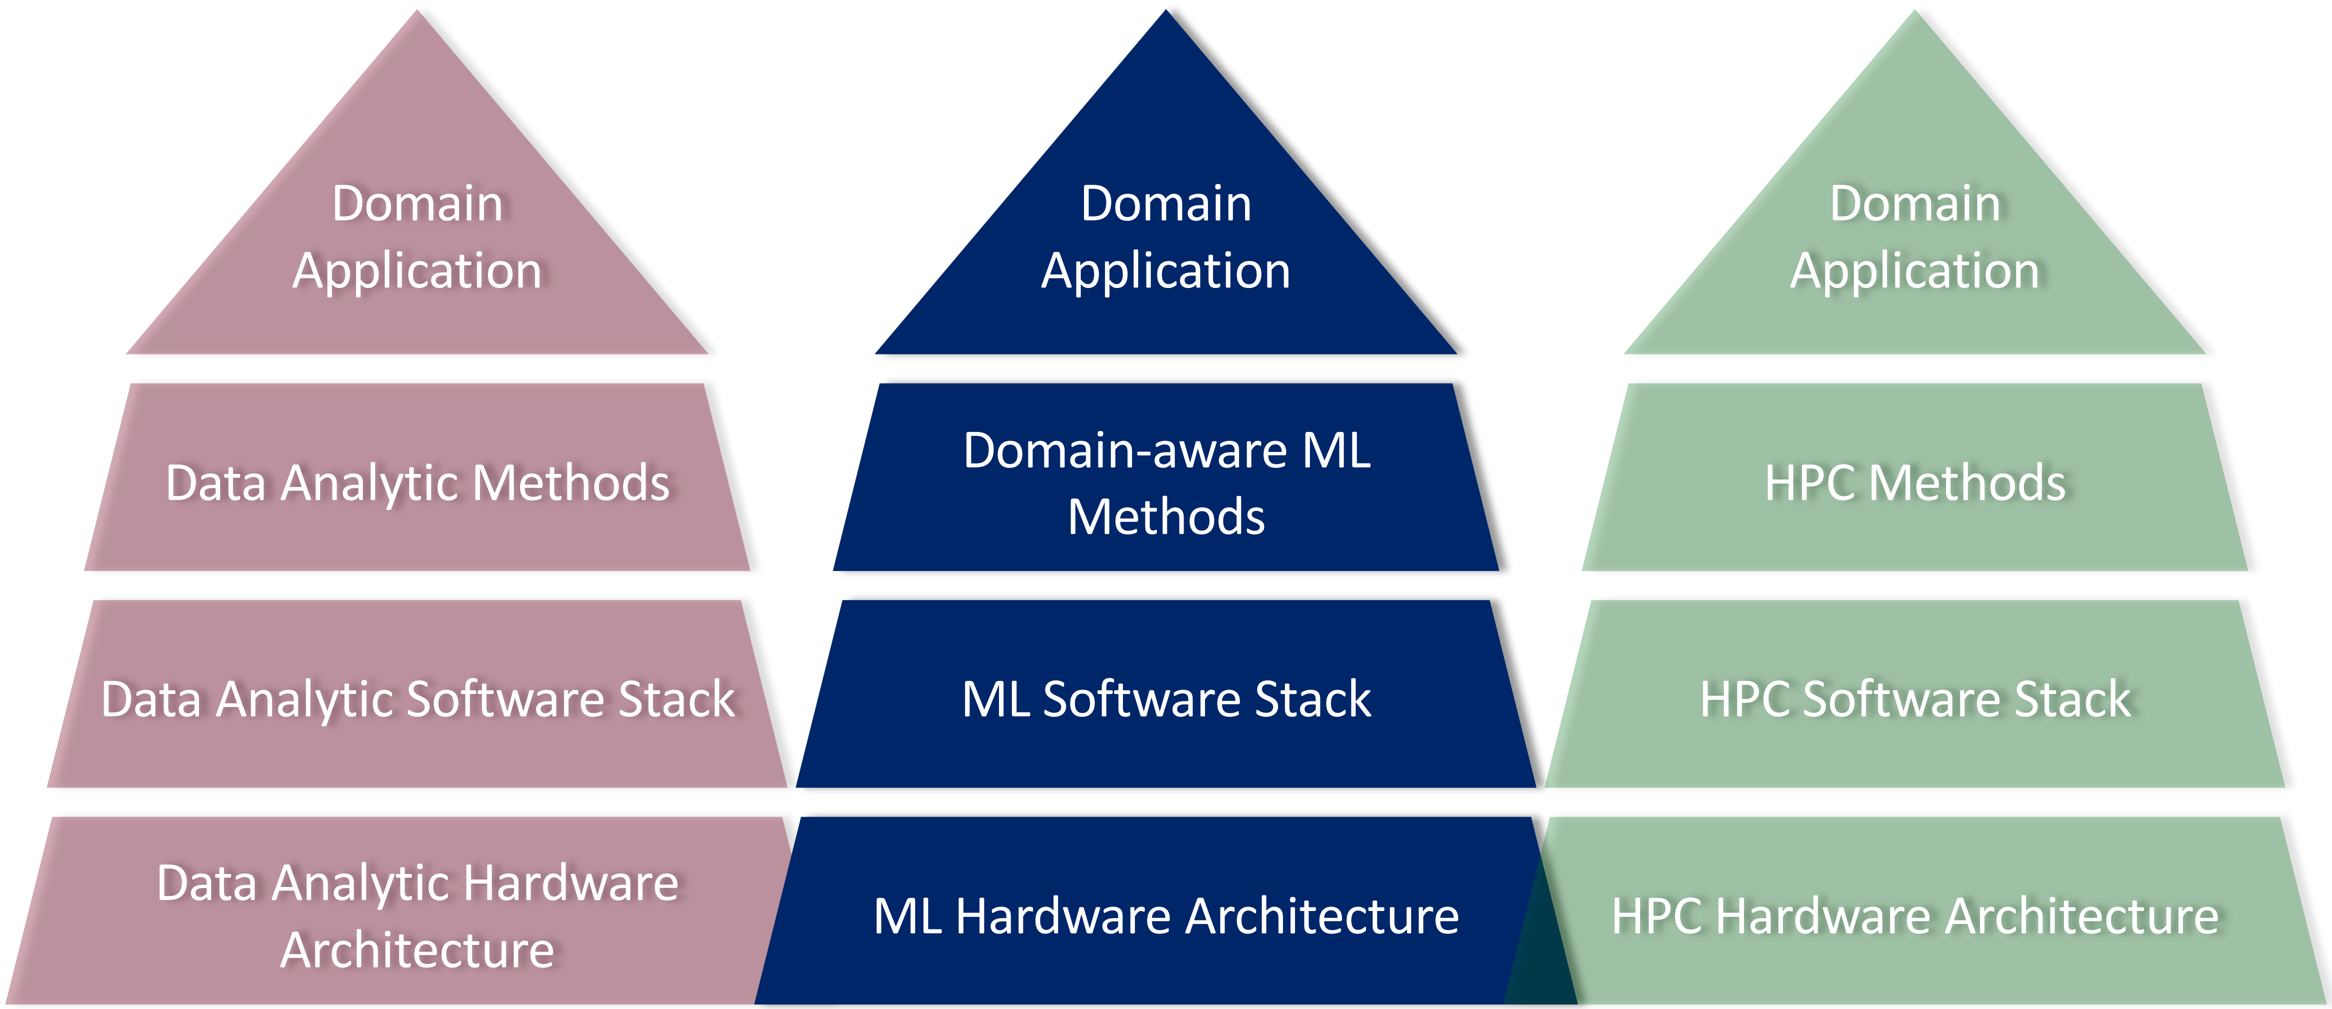

DMC Approach for Converged Computing Paradigms

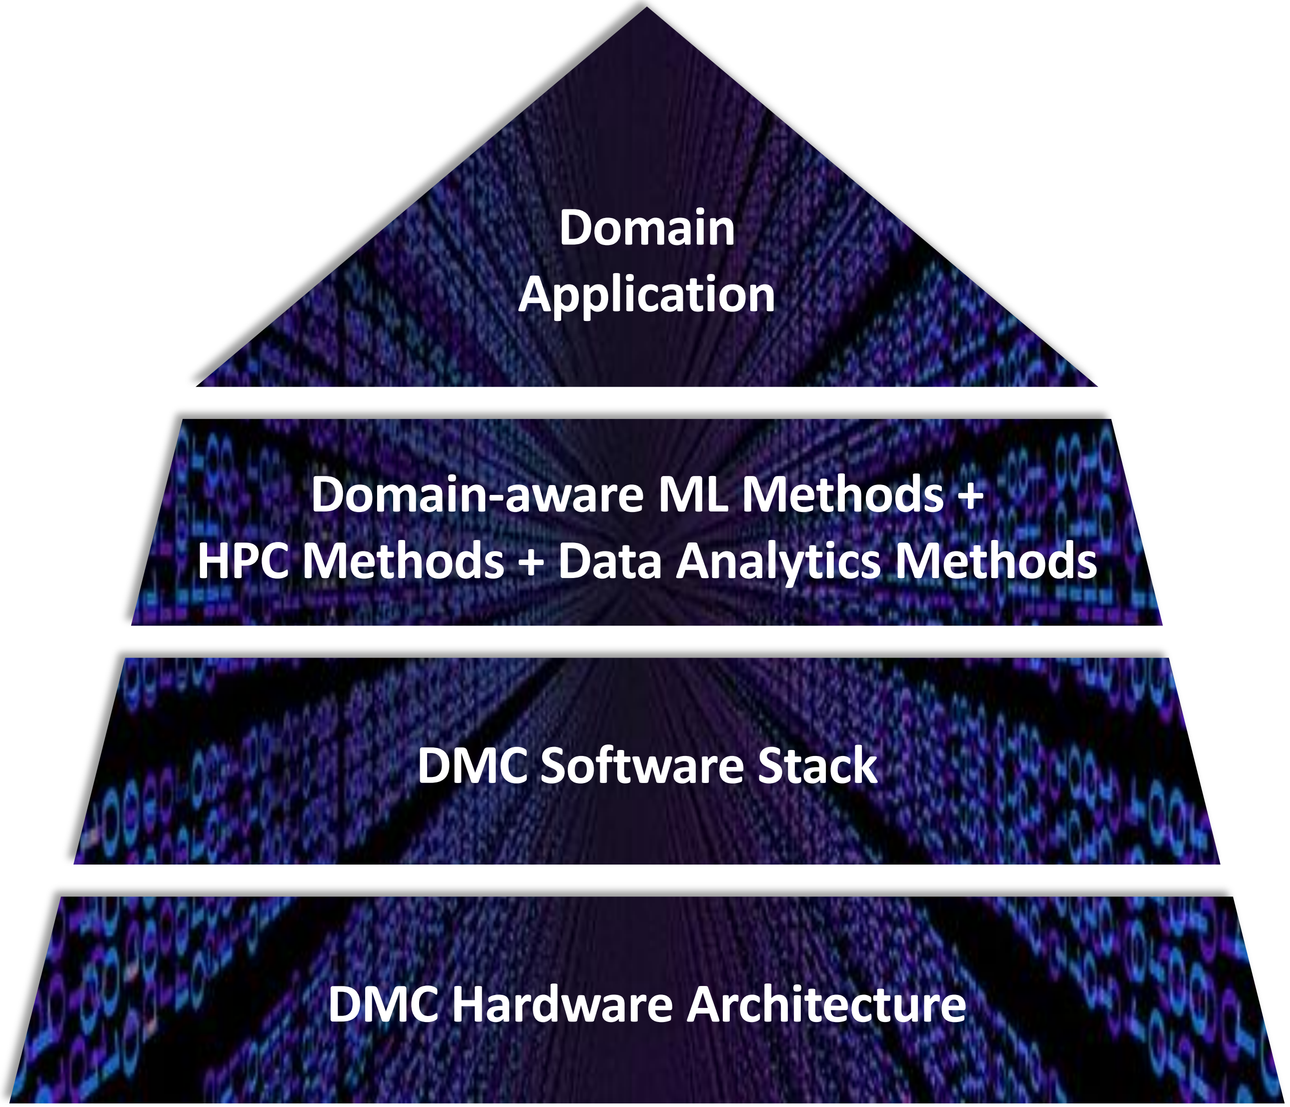

Approach: The Lead User

- Lead User is a key concept from *Democratizing Innovation* by Eric Von Hippel
- Lead Users are "a source of novel product concepts"
- With Open Source, innovation does not only come from manufacturers
- "Users are firms or individual consumers that expect to benefit from using a product or service. In contrast, manufacturers expect to benefit from selling a product or service."

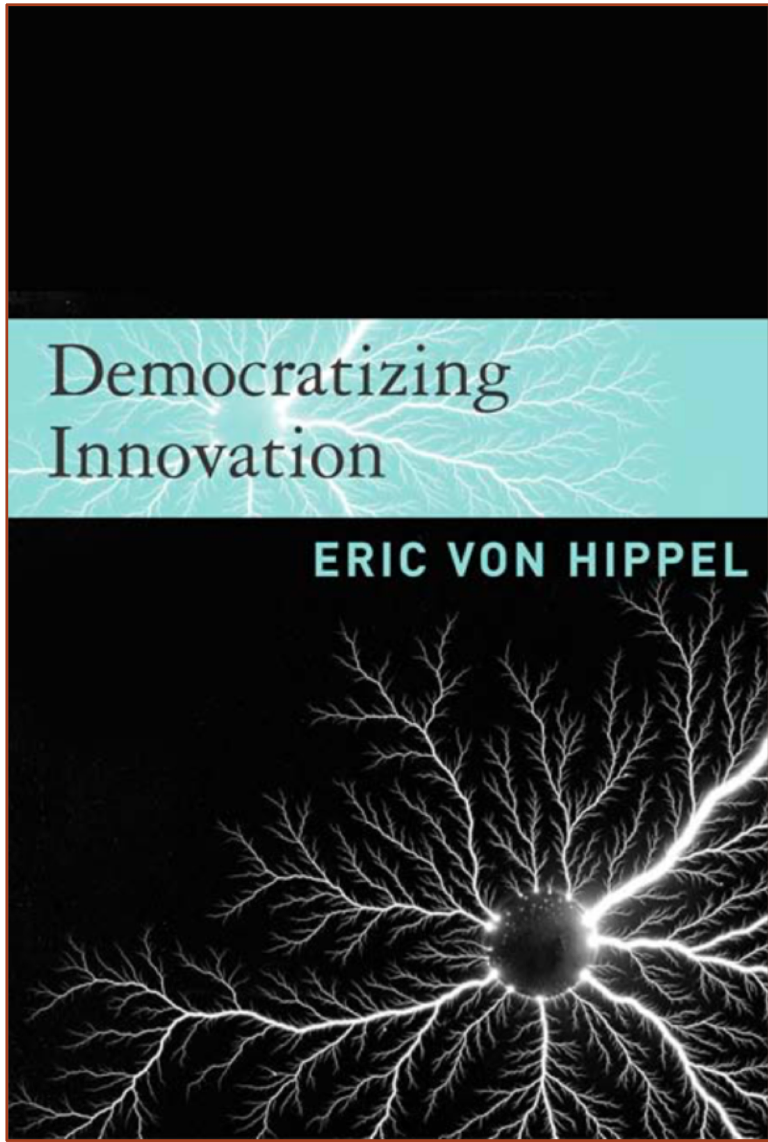

Approach: Holistic Co-design

- We can be Lead Users
  - Our goal is to co-design and develop DMC computer architectures and software stack for DMC workloads
- Purpose-designed Hardware Specialization
  - Processor and Memory manufacturers don’t have sufficient insight into our applications to know the most effective architectural innovations
  - As Lead Users, the DMC Initiative can establish multi-disciplinary collaborations to develop advanced design concepts

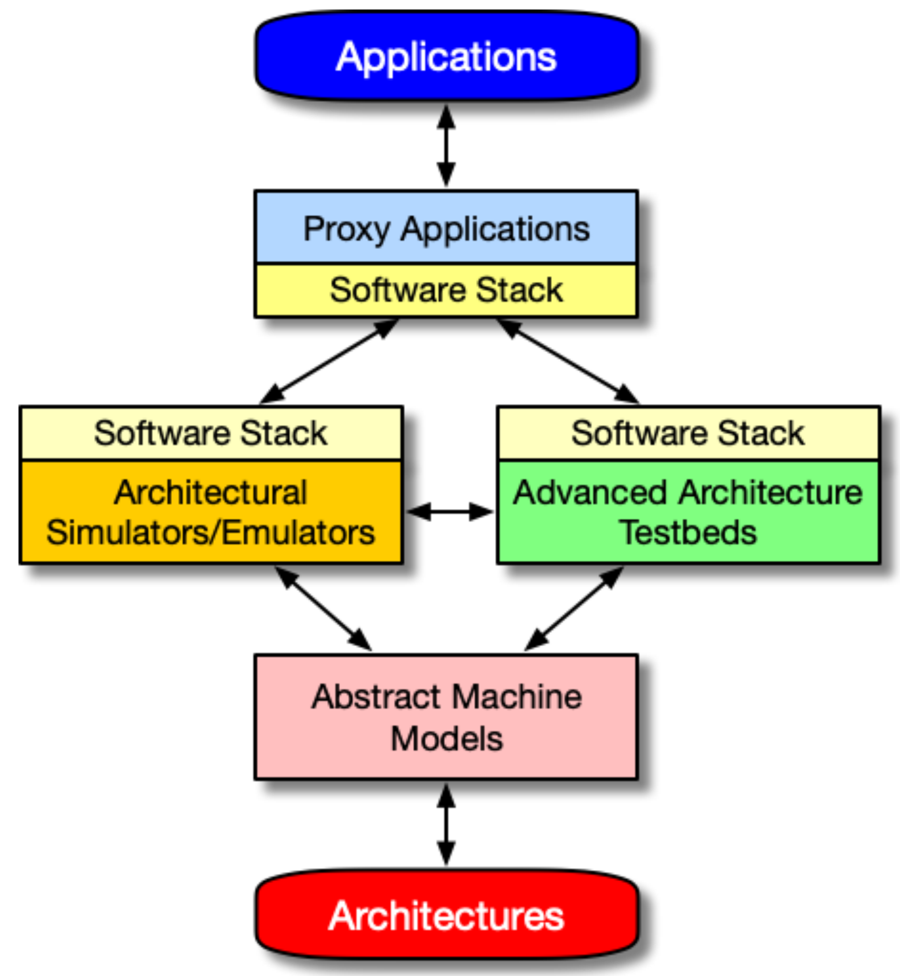

Approach: DMC Initiative Thrusts

- Domain Applications
  - Define DMC challenge problems
  - Deliver DMC capabilities to domain scientists
- Domain-aware Machine Learning
  - Address gap between domain models and scalable approaches to scientific ML
  - Deliver new theory and tools
- Software Stack
  - Address need for single programming framework for developing combined data, HPC, and ML applications
  - Deliver scalable software framework to support next-generation applications on heterogeneous hardware
- Hardware Architecture
  - Address need for novel architectures to improve system efficiency and performance in DMC applications
  - Deliver tools, methods, and prototype designs for next-generation architectures

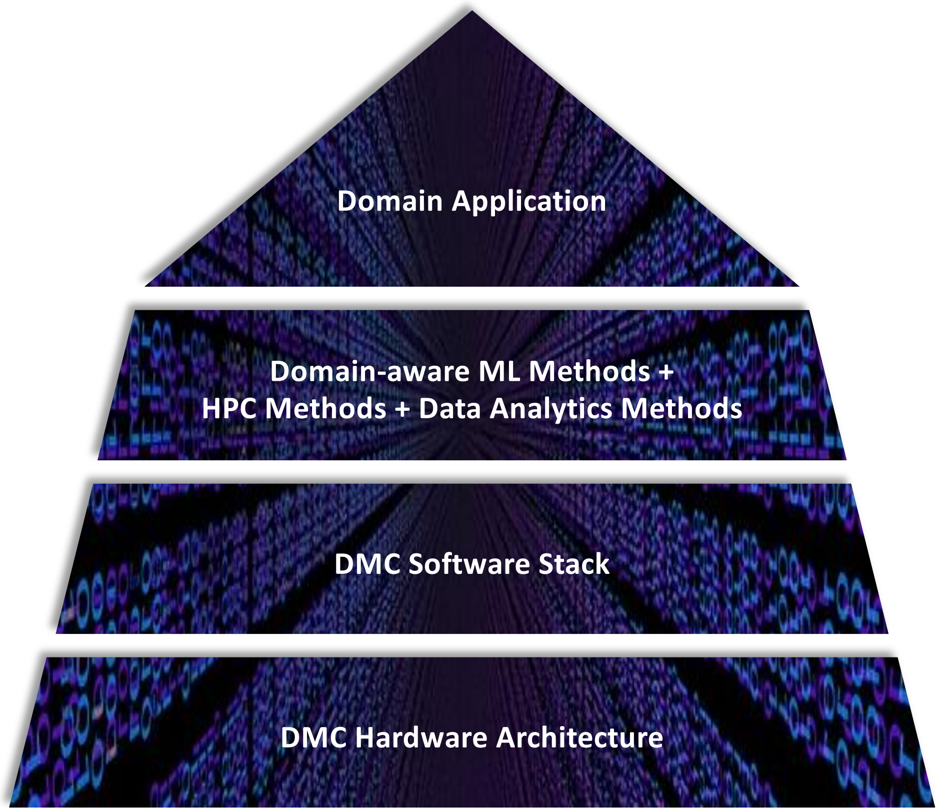

Domain Applications: Research Areas of Interest

- Address decision making and discovery under uncertainty at the convergence of data analytics, high performance computing and machine learning
  - Accelerate solutions to mathematical equations with hybrid dynamics**.** Integrate purpose designed hardware specialization, advanced numerical solvers, and data-informed algorithms for real-time monitoring, performance forecasting and robust decisions
  - Develop multi-agent decision systems. Integrate data-informed optimization, dynamic system decomposition, hierarchical machine learning and leverage infrastructure graph architecture knowledge, and data-driven control
  - Co-optimize design and control decisions. Integrate combinatorial optimization methods with data-driven modeling for decision steering

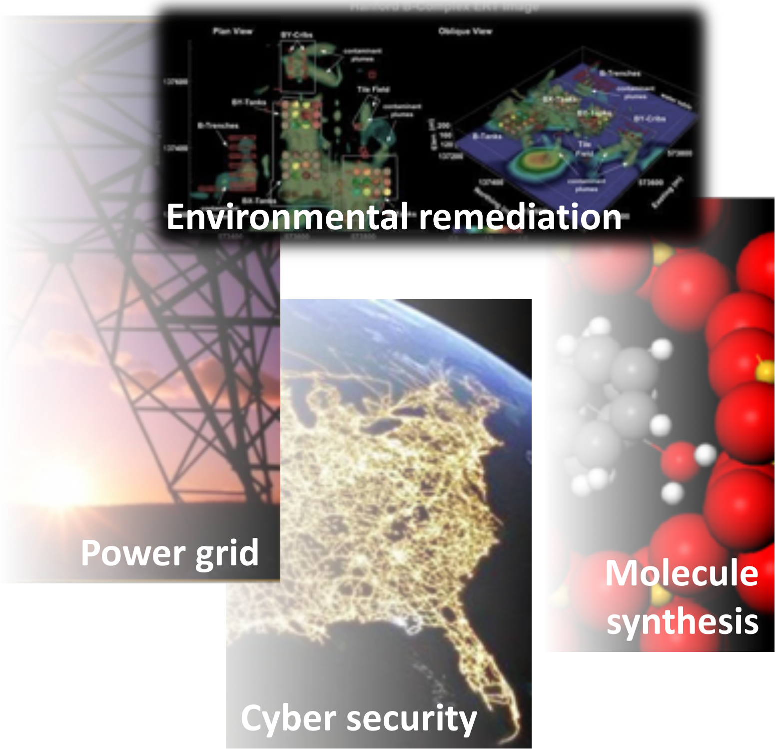

Domain Application: Power Grid

Goal

- Realizing a Secure, Flexible, and Resilient Electric Power System
  - DMC will enable real-time control of the future grid

Challenges

- Grid dynamics Control and Analysis
  - **Data-driven short timescale response** (continuous to milliseconds)
  - **Operator-driven moderate timescale response** (seconds to minutes to 10s of minutes)
  - **Model-driven long timescale optimization** (months to decades)
- Data volumes
  - Grid edge devices and smart meters amass data surpassing **1 petabyte/year**
- Machine Learning
  - Training to Improve Grid Resilience
  - Inference for distributed embedded controls

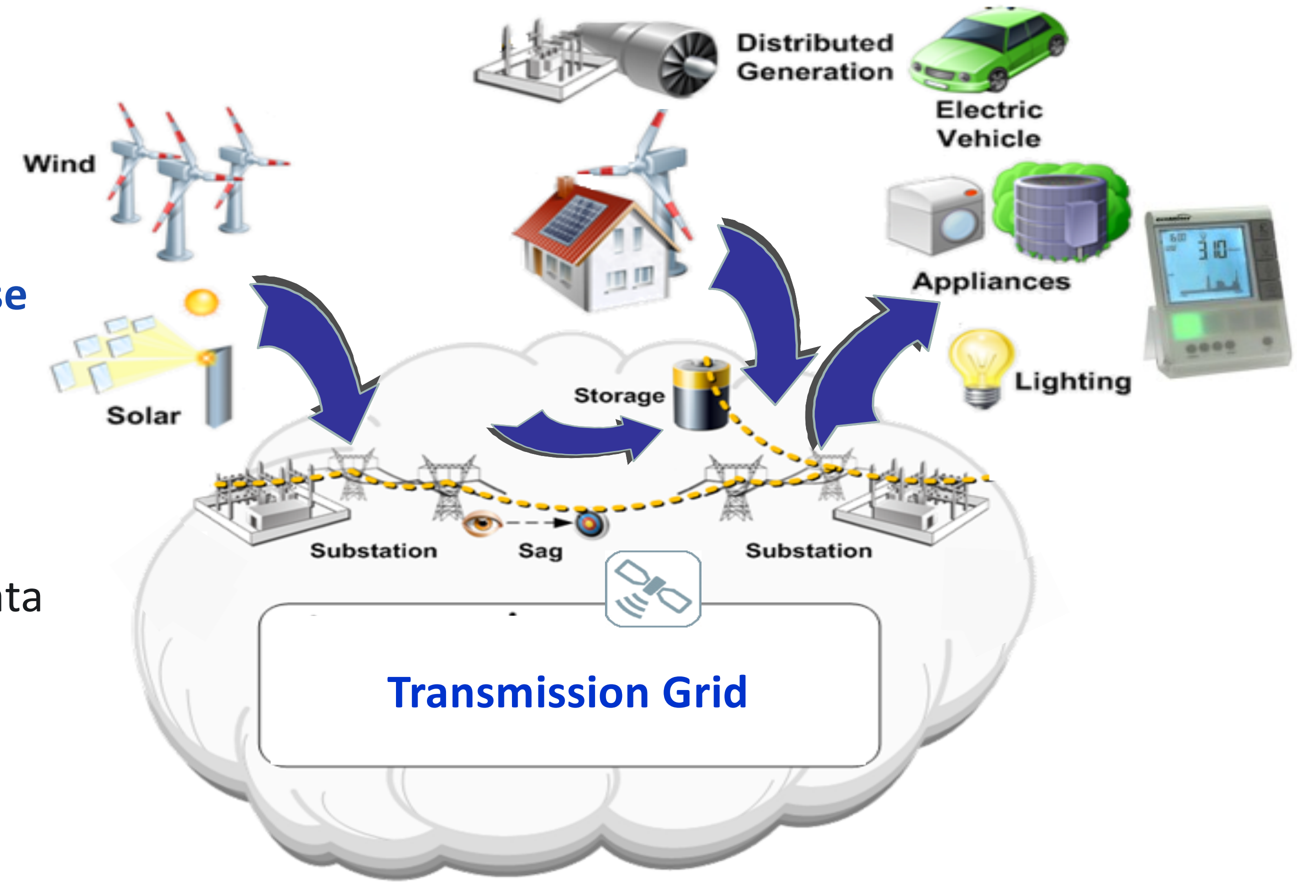

Domain-aware Machine Learning: Research Areas of Interest

- Embedding of Domain Knowledge
  - Knowledge representations for ML
  - Integration of domain knowledge
  - Data-driven scientific discovery
- Accelerated Learning
  - Methods to accelerate the convergence and stability of ML algorithms when (labeled) data are limited
  - Tools to speed-up the tuning and optimization of domain aware ML models

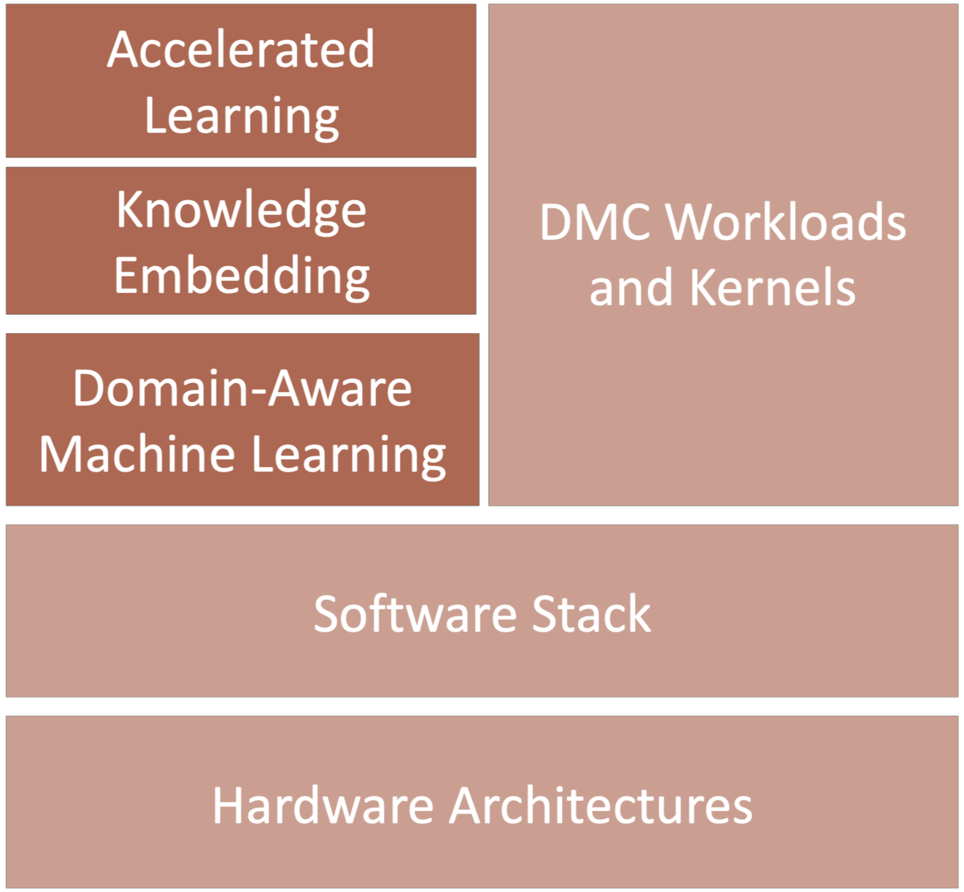

Software Stack: Research Areas of Interest

- DMC Software Runtimes
  - Asynchronous task-based, dataflow, functional, distributed actor-based methods
  - Support for processing-in-network, moving of compute to data (non von-Neumann)
  - Security and privacy mechanisms to support compliance
- Accelerator-based programming methodologies
  - Synthesis of custom hardware accelerators based on high-level algorithm specifications
  - Distributed coordination of workloads across heterogeneous, possibly non von-Neumann architectures
  - Provide efficient data movement (streaming, batch, parallel) to processing units
- Compiler or runtime synthesis and/or optimization of algorithms
  - Directive-based programming for parallel and streaming algorithms
  - Methods of discovering, exploiting, and maintaining parallelism in scientific codes

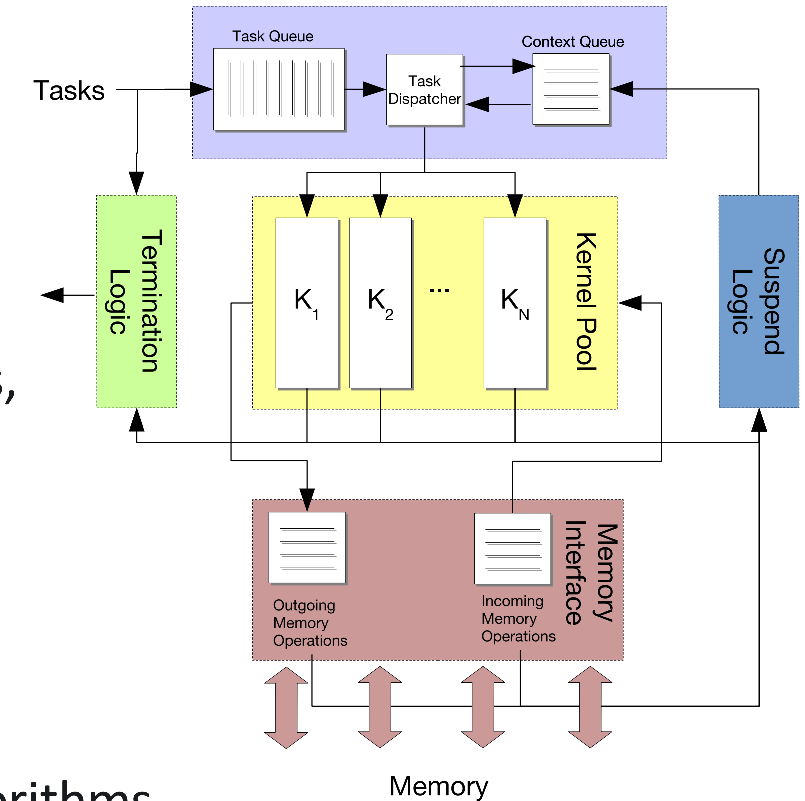

Hardware Architecture: Research Areas of Interest

- Tools
  - *Rapid* high-level design and synthesis
  - Enable design-space exploration
- Architectures
  - Targeting performance-critical DMC kernels
  - New computing paradigms (e.g., Non-von Neumann)
- Testbeds
  - Supporting DMC thrust areas
  - Enable measurement and assessment of performance impact

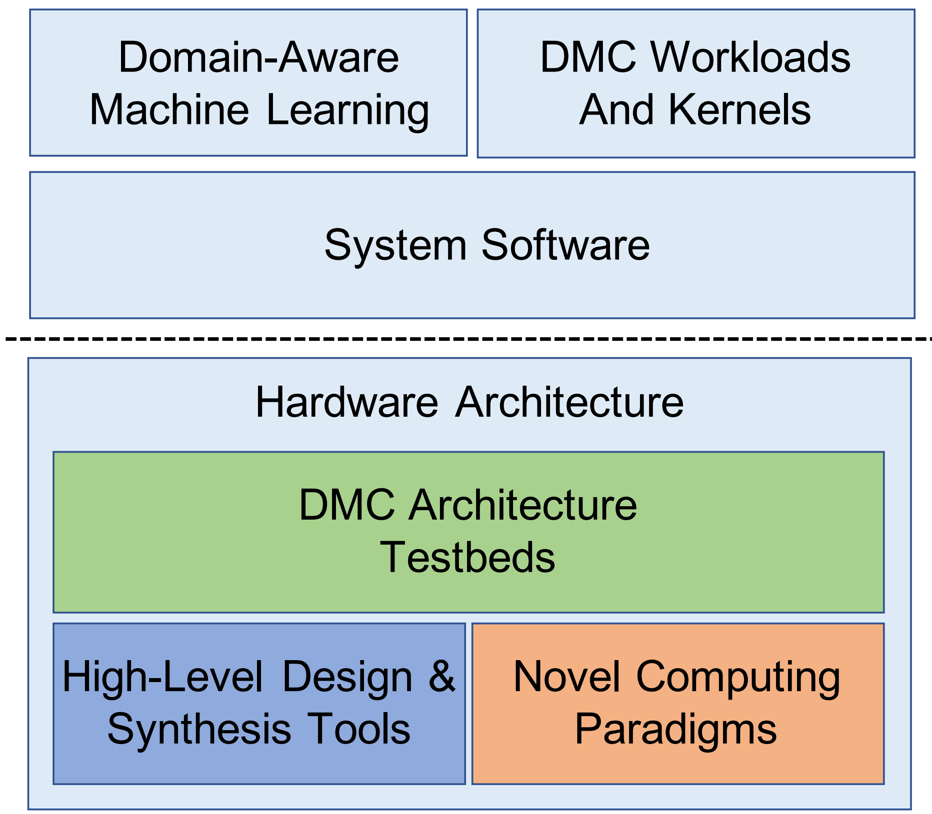

PNNL Strengths

Thrust Area	Leaders	Artifacts/Projects	Customers
Domain Applications	Carl Imhoff	Grid Modernization Lab Consortium	DOE/OE
	Henry Huang	ExaSGD, MMICCs: MACSER	ECP, ASCR, OE, CESER
	Jeff Dagle, et al	Dynamic Contingency Analysis Tool - DCAT	2018 R&D100 Awardee
Domain-aware Machine Learning	Nathan Baker	SciML BRN	ASCR
	Alex Tartakovsky	MMICCs: PhILMs	ASCR
	Court Corley, Nathan Hodas	Deep Science Agile Investment	PNNL LDRD
	Sutanay Choudhury, et al	StreamWorks	2018 R&D100 Awardee
	Hodas, Choudhury, Ang	ExaLearn	ECP, ASCR
Software Stack	Andrew Lumsdaine	MPI, PBGL, C++	ASCR, DoD, NSF
	Mark Raugas, John Feo	SHAD, CASS-MT	DoD/HPDA, DARPA
Hardware Architecture	Jim Ang	μelectronics BRN, ECP/HT Director	DOE-SC, ECP
	Ang, Kevin Barker, Feo	P38, AGILE	DoD, IARPA
	Barker, Tallent, Marquez	CENATE	ASCR
	Antonino Tumeo	DAEDALuS	DoD/HPDA

DOE-SC Basic Research Needs for Microelectronics

DOE's Long-term R&D Strategy for *Beyond Moore's Law*

Priority Research Directions

- Flip the current paradigm: Define innovative material, device, and architecture requirements driven by applications, algorithms, and software
- Revolutionize memory and data storage
- Reimagine information flow unconstrained by interconnects
- Redefine computing by leveraging unexploited physical phenomena
- Reinvent the electricity grid through new materials, devices, and architectures

https://science.energy.gov/ascr/community-resources/program-documents/

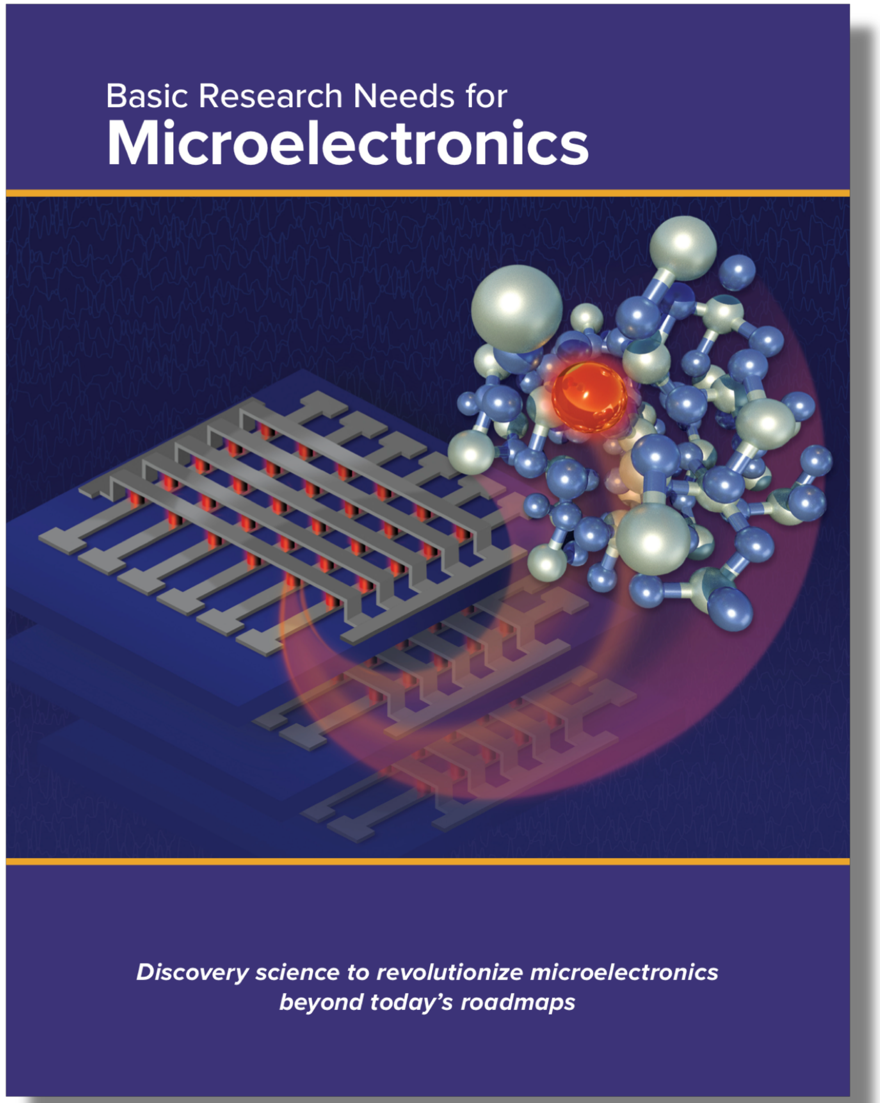

DMC Risks and Mitigations

- Technical Risks are expected / accepted to develop new game-changing capabilities
  - Concurrent development of Hardware Architecture and Software Stack creates moving targets
  - Holistic Co-design requires ongoing collaboration, not monthly or quarterly synchronization
  - DMC does not have direct control of open source projects that it may rely on
- The Leadership team has the responsibility to manage technical risk of the Overall DMC Initiative, through the selection and pruning of projects
  - Semi-annual project reviews and evaluation
  - Maintain queue of backup projects and tasks within projects
  - Agile Manifesto R&D approach leads to rapid test and evaluation of prototype capabilities by domain application users to inform project and initiative leaders

DMC Exit Strategy

- Targets
  - DMC Computing R&D shifts to ASCR, DoD, DARPA, CS R&D programs
  - DMC Domain Application support shifts to OE, CESER, BER, EERE, DoD, DHS programs
- Approach
  - "Off-ramps" integrated into the DMC structure to ensure success and mitigate risks
  - Agile development with regular delivery of high-quality functional products to actively engage customers/collaborators
- Areas of Emphasis
  - Initial application focus: power grid, building on existing electric power system relationships and capabilities
  - PNNL responsibility for deployment of DMC Architecture can cultivate skills as a future ASCR system prototype facility with strong vendor collaborations

PNNL Initiative Goals

- Development of enduring and game-changing capabilities
- Attract technical leadership to PNNL
- Pursue new opportunities
  - Program development
  - Academic partnerships
  - Strategic industry collaborations

Thank You

DMC will position PNNL to help define the Beyond Moore's Law computing world