





*Exceptional service in the national interest*



$$\frac{\partial}{\partial a} \ln J_{a, \sigma^2}(\xi_1) = \frac{(\xi_1 - a)}{\sigma^2} f_{a, \sigma^2}(\xi_1)$$

$$\int_{\mathcal{R}_a} T(x) \cdot \frac{\partial}{\partial \theta} f(x, \theta) dx = M \left( T(\xi) \cdot \frac{\partial}{\partial \theta} \ln U(\theta) \right)$$



## Kokkos Sustaining Exascale Codes for the Long Haul

Unclassified Unlimited Release

**Christian R. Trott**, - Center for Computing Research  
Sandia National Laboratories/NM



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.




How does Kokkos help applications with sustainability concerns?




### Applications

### Libraries

### Frameworks




**SNL NALU**  
Wind Turbine CFD




**SNL LAMMPS**  
Molecular Dynamics

- O
- Al
- Co
- Ni
- Fe

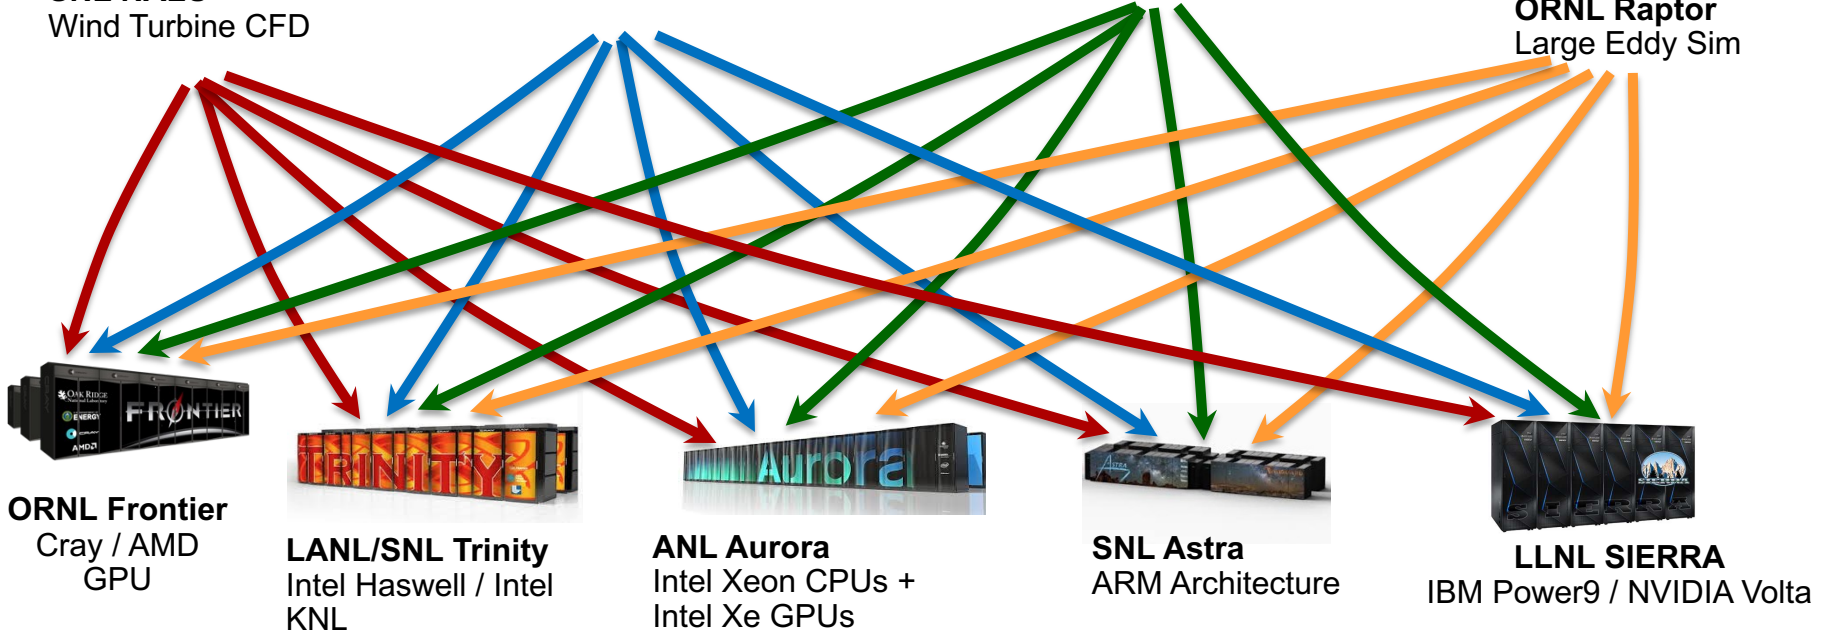


**UT Uintah**  
Combustion



**ORNL Raptor**  
Large Eddy Sim

$r: d = 1 \text{ mm}$   
 $d \times 2d \times 1d$   
 $T = 450 \text{ K (Jet)}$   
 $T = 1000 \text{ K (Cross-Flow)}$   
 $50 \text{ (Jet)}$   
 $50 \text{ (Cross-Flow)}$



**ORNL Frontier**  
Cray / AMD  
GPU



**LANL/SNL Trinity**  
Intel Haswell / Intel  
KNL



**ANL Aurora**  
Intel Xeon CPUs +  
Intel Xe GPUs




**SNL Astra**  
ARM Architecture




**LLNL SIERRA**  
IBM Power9 / NVIDIA Volta



**Applications**




**Computational  
Fluid Dynamics**




**Molecular Dynamics**

**Libraries**



**Electro Magnetics**

**Frameworks**



**Climate Simulation**

**Kokkos**



**ORNL Frontier  
AMD GPUs**



**LANL/SNL Trinity  
Intel CPUs**



**ANL Aurora  
Intel GPUs**



**Riken Fugaku  
ARM CPUs**



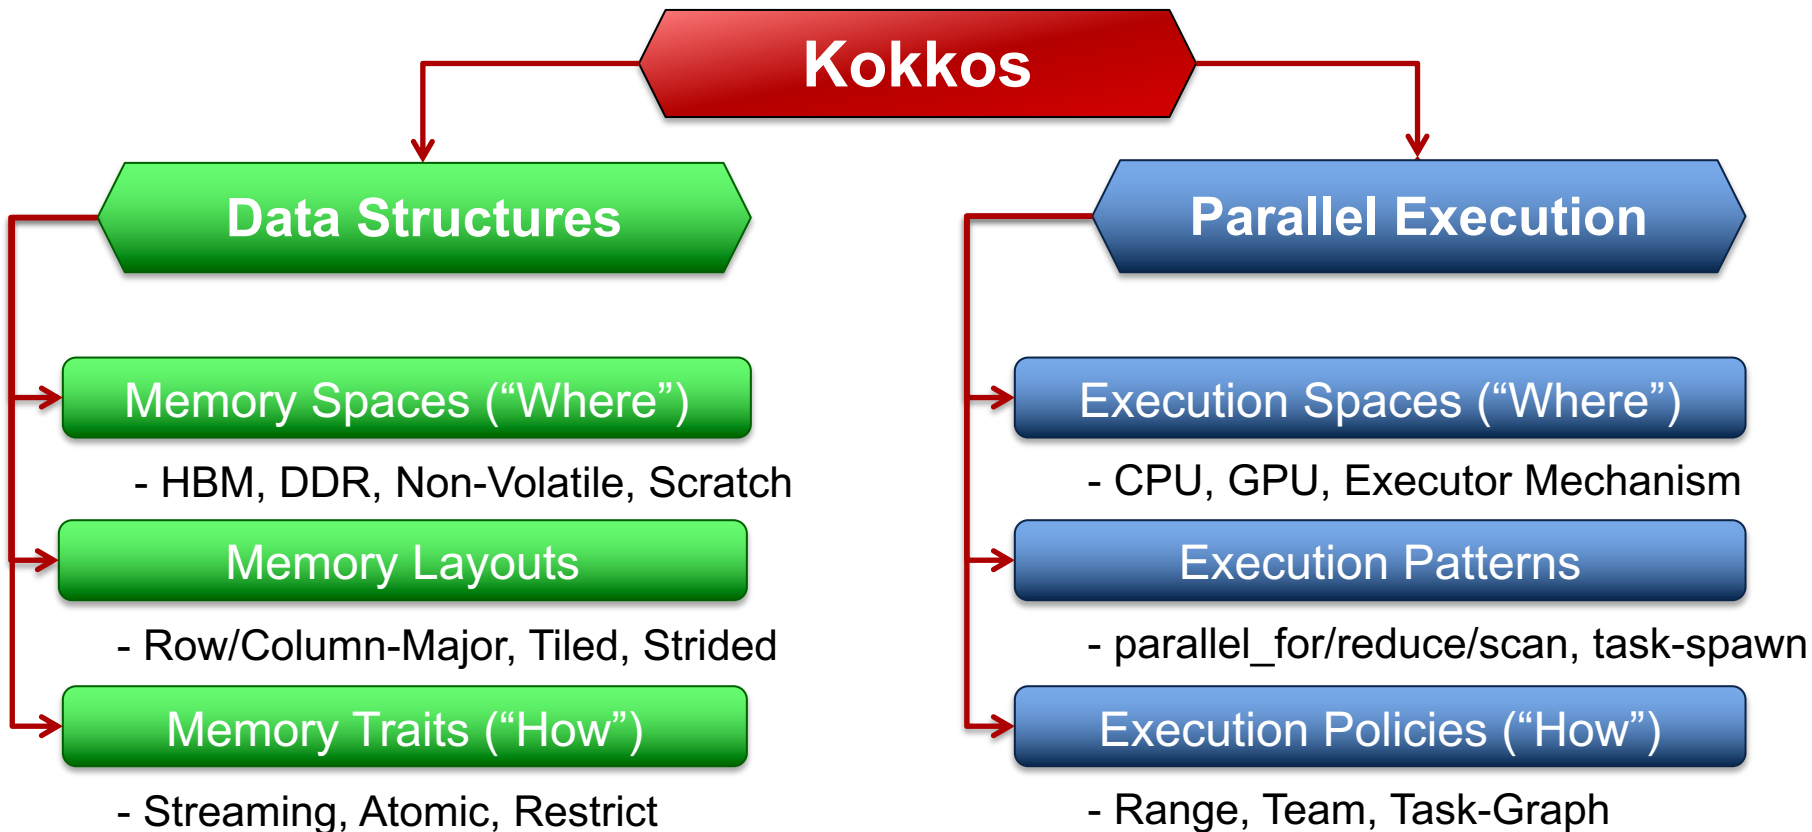
**NERSC Permuter  
NVIDIA GPUS**

# What is Kokkos?

- A C++ Programming Model for Performance Portability
  - Implemented as a template library on top of CUDA, OpenMP, HPX, ...
  - Aims to be descriptive not prescriptive
  - Aligns with developments in the C++ standard
- Expanding solution for common needs of modern science/engineering codes
  - Math libraries based on Kokkos
  - Tools which enable insight into Kokkos
- It is Open Source
  - Maintained and developed at <https://github.com/kokkos>
- It has many users at wide range of institutions.



# Kokkos Core Abstractions




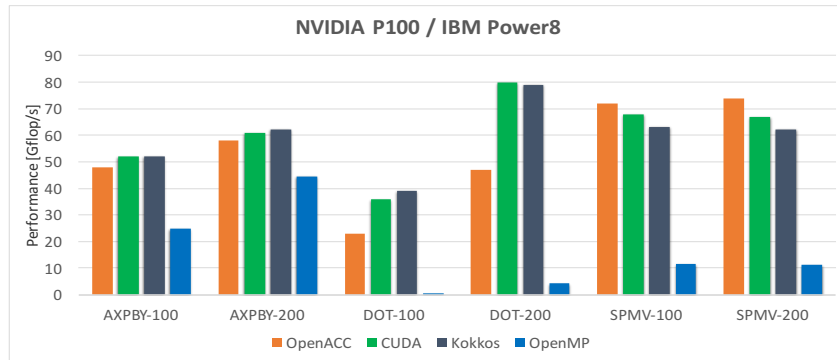


# Kokkos Core Capabilities

Concept	Example
Parallel Loops	<code>parallel_for( N, KOKKOS_LAMBDA (int i) { ...BODY... } );</code>
Parallel Reduction	<code>parallel_reduce( RangePolicy&lt;ExecSpace&gt;(0,N), KOKKOS_LAMBDA (int i, double&amp; upd) { ...BODY... upd += ... }, Sum&lt;&gt;(result));</code>
Tightly Nested Loops	<code>parallel_for(MDRangePolicy&lt;Rank&lt;3&gt; &gt; ({0,0,0},{N1,N2,N3},{T1,T2,T3}, KOKKOS_LAMBDA (int i, int j, int k) {...BODY...});</code>
Non-Tightly Nested Loops	<code>parallel_for( TeamPolicy&lt;Schedule&lt;Dynamic&gt;&gt;( N, TS ), KOKKOS_LAMBDA (Team team) { ... COMMON CODE 1 ... parallel_for(TeamThreadRange( team, M(N)), [&amp;] (int j) { ... INNER BODY... } ); ... COMMON CODE 2 ... });</code>
Task Dag	<code>task_spawn( TaskTeam( scheduler , priority), KOKKOS_LAMBDA (Team team) { ... BODY } );</code>
Data Allocation	<code>View&lt;double**, Layout, MemSpace&gt; a("A",N,M);</code>
Data Transfer	<code>deep_copy(a,b);</code>
Atomics	<code>atomic_add(&amp;a[i],5.0); View&lt;double*,MemoryTraits&lt;AtomicAccess&gt;&gt; a(); a(i)+=5.0;</code>
Exec Spaces	Serial, Threads, OpenMP, Cuda, HPX (experimental), HIP (experimental), OpenMPTarget (experimental)

# CG Solve: Performance 2016

- Comparison with other Programming Models
- Straight forward implementation of kernels
- OpenMP 4.5 was immature at that point
- Two problem sizes: 100x100x100 and 200x200x200 elements








# CG Solve Performance 2022

- CG-Solve as discussed above
- Also try replacing SPMV with TPL
- Running 100x100x100 heat conduction problem
  - "MiniFE" Proxyapp setup
- Measure effective Bandwidth
  - Algorithmical memory ops per time
- Why is this beating vendor libs?
  - Its complicated, but a real effect

## CG-Solve Effective Bandwidth





Developing a sustainable project.

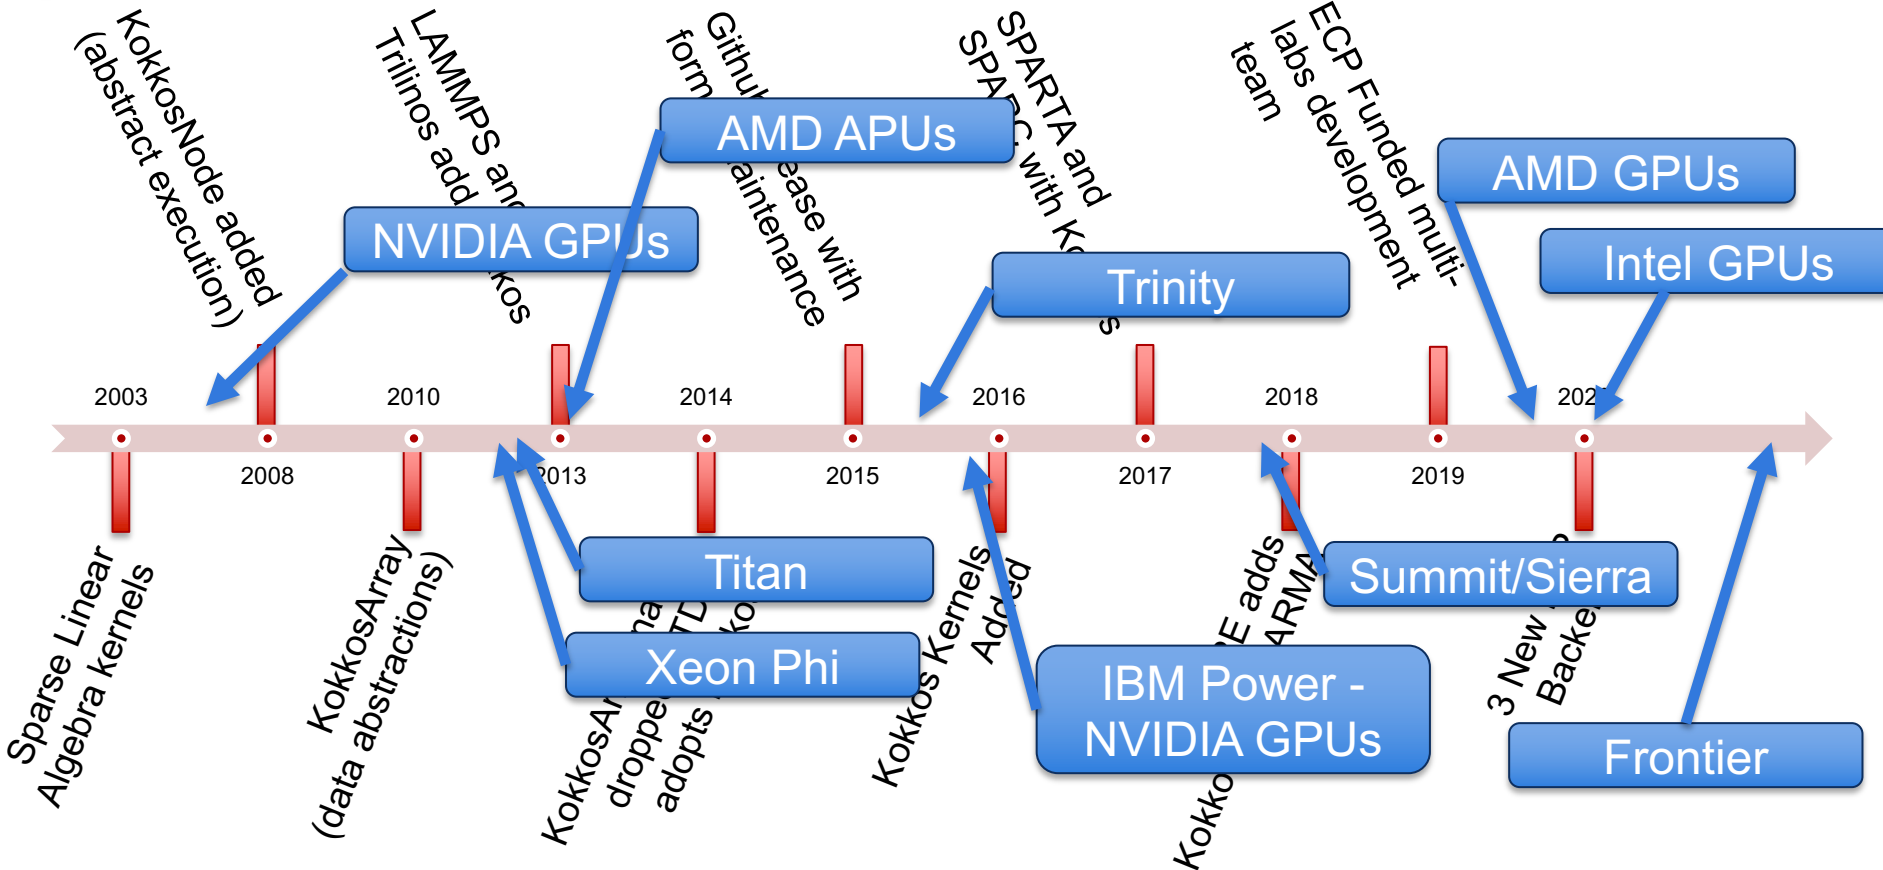


# Stepping Stones To Success

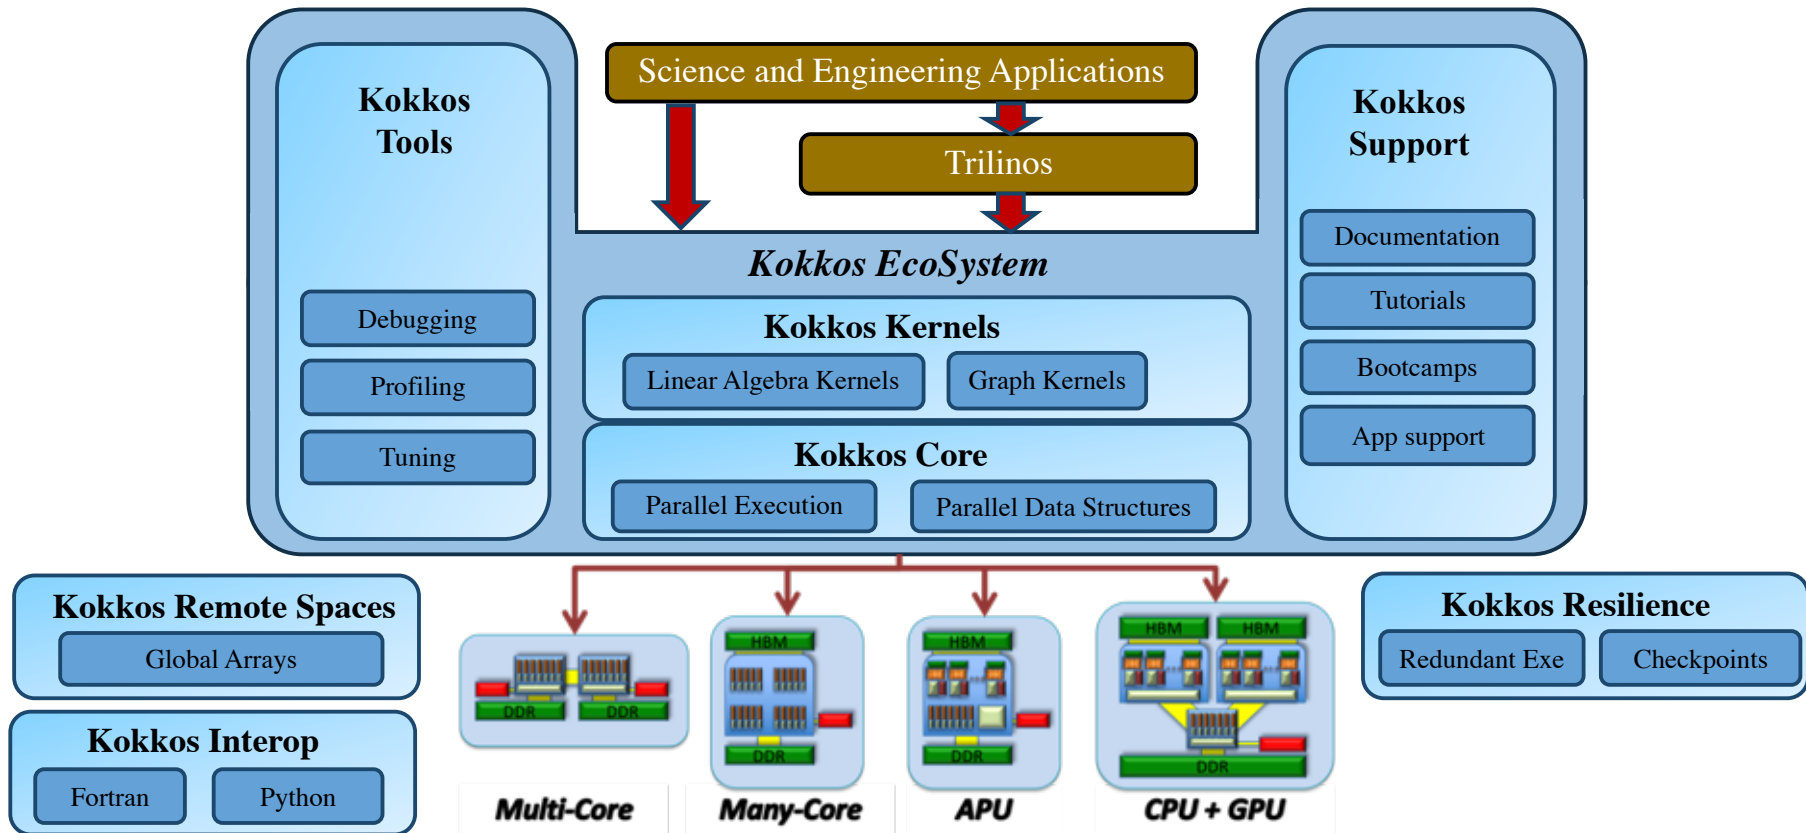
- Making the transition from research to actually deployed production software required years with a number of key aspects being critical:
  - **Address a Critical Need:** More platforms were coming and apps simply can't rewrite for every single one
  - **Demonstrate Viability:** The Kokkos team ported chunks of LAMMPS and Trilinos to demonstrate that it works and performs; We also demonstrated performance parity to native models in miniApps
  - **User outreach:** The Kokkos team spend large amounts of its time in outreach, training and support to build the initial community
  - **Long term stability:** Having the support from Sandia's HPC management ensured stable long term funding



# Kokkos Timeline at Sandia



# Kokkos EcoSystem



# Testing and Verification

- Github based developments with Pull Requests
- Pull Request Testing
  - 30 configurations
  - Includes: CUDA, HIP, OpenMPTarget, SYCL, Windows, Linux, OSX
- “Rule of three”: two additional developers need to approve a pull request
  - Approving means: “I feel I could maintain this if the author left”
- Strong backwards compatibility policy
  - Can’t break backwards compatibility too much, and even minor ones only every couple years or so
- Releases require integration testing with major customers
  - Includes multi-platform test of Trinos and LAMMPS among others



# Building a Community



# Kokkos Uptake




## Community Project



- 25 Developers
  - 15 for Core
- Regular contributions from HPC vendors

## Kokkos Slack Channel

- >1200 registered users
  - 150 Institutions
  - Every continent
    - (-Antarctica)



## Applications and Libraries

- Estimated 200-400 HPC projects using Kokkos
- On the order of two-dozen apps run science and engineering production runs with Kokkos
- Similar distribution as the Slack User





# Training and Documentation

- The Kokkos Lectures
  - 8 lectures covering most aspects of Kokkos
  - 15 hours of recordings
  - > 500 slides
  - >20 exercises
- Extensive Wiki
  - API Reference
  - Programming Guide
- Slack as primary direct support

<https://kokkos.github.io>

- Module 1: Introduction
  - Introduction, Basic Parallelism, Build System
- Module 2: Views and Spaces
  - Execution and Memory Spaces, Data Layout
- Module 3: Data Structures and MDRangePolicy
  - Tightly Nested Loops, Subviews, ScatterView,...
- Module 4: Hierarchical Parallelism
  - Nested Parallelism, Scratch Pads, Unique Token
- Module 5: Advanced Optimizations
  - Streams, Tasking and SIMD
- Module 6: Language Interoperability
  - Fortran, Python, MPI and PGAS
- Module 7: Tools
  - Profiling, Tuning , Debugging, Static Analysis
- Module 8: Kokkos Kernels
  - Dense LA, Sparse LA, Solvers, Graph Kernels



# Recently Added Capabilities



# New Capabilities: ISO C++ Standard Algorithms

- Kokkos provides vast majority of all ISO C++ standard algorithms since v. 3.6
- Take iterators or Kokkos::View
- Available at top level and at team level in hierarchical parallelism (new in 4.2)
  - Top level take optional “Label” like other top level Kokkos parallel patterns
  - Always take either Execution Space instance or Team Handle

```
auto iter = find("Label", DefaultExecutionSpace(),
                cbegin(view), cend(view), find_value);
auto iter = find("Label", DefaultExecutionSpace(),
                view, find_value);
auto iter = find(DefaultExecutionSpace(), cbegin(view), cend(view), find_value);
auto iter = find(DefaultExecutionSpace(), view, find_value);
auto iter = find(team_handle, cbegin(view), cend(view), find_value);
auto iter = find(team_handle, view, find_value);
```



# New Capabilities: ISO C++ SIMD Types

- - Previously available as separate repository ([github.com/kokkos/simd-math](https://github.com/kokkos/simd-math))
- - Now part of Kokkos Core
- - Implements proposed ISO C++ SIMD types (for C++26)
- - Work across all backends, may just fall back to scalar operations

```
using simd_t = Kokkos::Experimental::simd<double>, Kokkos::Experimental::simd_abi::native<double>;

int N = N_in / simd_t::size();
Kokkos::View<simd_t **> data("D", N, M);
Kokkos::View<simd_t *> results("R", N);

double a = 3.0;
Kokkos::parallel_for( "Combine", data.extent(0), KOKKOS_LAMBDA(const int i) {
    simd_t tmp = 0.0;
    double b = a;
    for (int j = 0; j < data.extent(1); j++) {
        tmp += b * data(i, j);
        b += a + 1.0 * (j + 1);
    }
    results(i) = tmp;
});
```



# New Capabilities: MDRangePolicy for Teams



- Support MDRangePolicy at nested hierarchical parallelism level
- Parallelizes with both vector and thread level if available
- Three policies in analogy to existing nested policies:
  - TeamThreadMDRange
  - TeamVectorMDRange
  - ThreadVectorMDRange
- Support all typical things: parallel\_for, parallel\_reduce, Rank argument etc.

```
Kokkos::parallel_for(Kokkos::TeamPolicy<ExecSpace>(leagueSize, Kokkos::AUTO),
  KOKKOS_LAMBDA(const TeamType& team) {
    int leagueRank = team.league_rank();

    auto teamRange =
      Kokkos::TeamThreadMDRange<Kokkos::Rank<2, Direction>, TeamType>(team, n0, n1);

    Kokkos::parallel_for(teamRange, [=](int i, int j) {
      v(leagueRank, i, j) += fillFlattenedIndex(leagueRank, i, j);
    });
  });
```



**Sandia  
National  
Laboratories**